



Developer Guide for Amazon Nova 2

Amazon Nova



Amazon Nova: Developer Guide for Amazon Nova 2

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What is Amazon Nova 2?	1
Key concepts	1
Amazon Nova models	1
What can you build?	2
Benefits	2
How Amazon Nova works	3
Pricing	4
Next steps	4
What's new in Amazon Nova 2	5
New models	5
New features and capabilities	5
Nova Multimodal Embeddings	5
Extended thinking and reasoning	5
Built-in tools	6
AI agent building	6
Improved document understanding	6
Enhanced video understanding	6
Model customization	6
Amazon Nova Forge	6
Next steps	7
Getting started	8
Getting started in the console	8
Explore the text playground	9
Getting started with the API	9
Get credentials to grant programmatic access	11
Attach Amazon Bedrock permissions to a user or role	12
Generate a response for a text prompt using an Amazon Nova model	13
Core inference	14
Important: Timeout Configuration	15
Core Inference Topics	15
Using the Converse API	16
Request structure	16
Using system prompts	17
Inference parameters	18

Using reasoning	19
Using the Invoke API	21
Request structure	22
Request parameters	23
Streaming responses	23
Streaming with ConverseStream	23
Streaming with InvokeModelWithResponseStream	24
Stream event types	24
Using Amazon Nova embeddings	25
Key features	25
Generating embeddings	26
On-demand inference	29
Benefits	29
Using on-demand inference	30
Pricing	30
Quotas and limits	30
Using Nova capabilities	31
Reasoning	31
Multimodal understanding	32
Supported content type by modality	32
Image understanding	33
Video understanding	35
Document understanding	38
Speech-to-Speech	40
Key features	40
Getting started	41
State the imports and configuration	41
Define the SimpleNovaSonic class	42
Initialize the client	43
Handle events	43
Start the session	43
Handle audio input	46
End the session	48
Handle responses	48
Playback audio	50
Capture audio	51

Run the main function	52
Code examples	53
Getting started examples	53
Advanced use cases	54
Hands-on workshop	55
Voice conversation prompts	55
Important Guidelines	55
Recommended Baseline System Prompt for Voice	56
Speech Prompt Configuration	56
System Prompt Configuration	57
Gender Agreement for Gendered Languages	58
Chain of thought for Speech: Constitutional Reasoning	59
Overuse of suggested phrases	60
Core concepts	62
Event lifecycle	62
Event flow sequence	64
Handling input events with the bidirectional API	64
Handling output events with the bidirectional API	73
Barge-in	80
Turn-taking controllability	81
Cross-modal input	83
Language support and multilingual capabilities	86
Managing chat history	87
Tool configuration	93
Asynchronous tool calling	100
Integrations	103
Strands Agents	103
Framework integrations	106
Advanced systems with Nova	117
Overview	117
Extended Thinking	117
How Extended Thinking Works	117
Configuration Options	120
Supported models	121
Extended thinking with tool use	121
Understanding Reasoning Tokens and Pricing	121

Frequently Asked Questions	122
Tool use	122
Understanding the tool use workflow	123
Create a tool	123
Call a tool	126
Built-in system tools	130
Building AI agents	132
Create an agent	132
Invoke an agent	134
Web Grounding	137
How Web Grounding works	137
How to use Web Grounding	137
Regional availability	139
Response structure	139
Grounding safety	140
Error handling	140
Permissions required for built in tools	140
Prompt engineering	142
What is prompt engineering	142
Getting started with prompt engineering	142
Best practices	143
Understanding the roles	143
Create precise prompts	144
Bring focus to sections of the prompt	162
Using the system role	164
Provide examples (few-shot prompting)	172
Tool calling systems	182
Advanced prompting techniques	183
Prompting multimodal inputs	205
General multimodal guidelines	205
Document and Image Understanding	206
Video understanding	35
Content moderation	218
Key principles	218
Recommended configuration	218
MLCommons universal taxonomy	218

Prompting examples	220
Code and Troubleshooting	226
Code library	226
Converse API Examples	226
InvokeModel API Examples	246
Troubleshooting	251
Authentication and setup	251
Model access denied	252
Regional availability issues	252
Timeout configuration	252
API response issues	253
Reasoning mode errors	254
Tool use issues	255
API and SDK reference	259
Request and response schema	259
Complete request structure	259
Complete response structure	264
SDK reference	267
Amazon Nova model customization	268
Customization approaches	268
Customization platforms	269
General prerequisites	270
Amazon Nova recipes	271
Getting base recipes	272
Getting Forge recipes	272
Available models and algorithms	272
Amazon Nova 2.0 Lite	273
On SageMaker training jobs	281
Best Practices for Amazon Nova customization	282
Nova Forge SDK	286
Training	292
Monitoring Progress Across Iterations	333
Evaluation	337
Iterative training	403
On SageMaker HyperPod	403
Compute requirements for Amazon Nova 2 models	404

Best Practices for Amazon Nova customization	282
Nova Forge SDK	409
HP cluster setup	410
Essential Commands Guide	410
Nova Forge access and setup	420
Training	422
Evaluation	461
MLflow monitoring	511
Iterative Training	513
Overview	513
How it works	514
When to use iterative training	515
Example workflow: SFT → RFT	516
Monitoring progress across iterations	518
Best practices	520
Cost considerations	522
Limitations	522
Troubleshooting	523
SageMaker Inference	525
Features	525
Supported models and instances	526
Supported AWS Regions	529
Supported Container Images	529
Best Practices	529
Support	530
Getting started	530
API reference	551
Evaluate models	560
Abuse detection for Amazon Nova Forge	567
Amazon Bedrock inference	568
Deploy a custom model for on-demand inference	568
On-demand inference	574
Limitations	575
Nova Forge	577
Prerequisites	577
Subscribe to Nova Forge	577

Other prerequisites	578
Nova Forge access and setup	578
Getting the Nova Forge documents	578
Step 1: Subscribe to Nova Forge	578
Step 2. Set up HyperPod infrastructure	585
Content moderation settings	585
Continued Pre-Training and Mid-Training	586
What are intermediate checkpoints and why are they needed?	586
What checkpoints are available?	587
Which checkpoint to use?	588
How to use data mixing for Nova 2.0?	588
Dissecting the Data Mixing Categories	590
Supervised Fine-Tuning	595
Introduction	595
When to use SFT?	595
When to use parameter efficient and when full rank SFT?	596
Data Mixing for SFT	596
YAML Configuration Examples	597
Model Checkpoints	599
Training Approaches	600
Pre-Requisites before you begin	600
A systematic Approach to Achieving Successful SFT	601
Experiment Sequencing and Data Mixing	602
Prepare Dataset for SFT	602
Evaluation Methods	603
Best Practices	282
Prepare Training Job Config	605
Set Up Data Mixing Block	606
How to Launch a job	612
How to Monitor Job	621
How to evaluate your model after SFT?	623
Lessons Learned and Tips	626
Reinforcement Learning	626
Bring your own orchestrator for agentic multi-turn evaluations	627
Responsible AI Toolkit	628
Responsible AI toolkit	628

Responsible use	630
Guidelines	630
Recommendations	631
Quotas	633
Request an increase for Amazon Nova quotas	633
Monitoring	635
Security	636
Create resources with CloudFormation	637
Learn more about CloudFormation	637
Document history	638

What is Amazon Nova 2?

Amazon Nova provides multimodal foundation models that process text, images, video, documents and speech. With support for up to 1 million tokens of context and advanced reasoning capabilities, Amazon Nova 2 models enable you to build sophisticated AI applications that understand complex inputs and generate accurate responses.

You can build interactive chatbots, analyze documents and videos, create AI agents with extended reasoning and develop voice-enabled applications.

Key concepts

Before you learn about Amazon Nova models, familiarize yourself with the following core concepts:

Foundation models

Pre-trained AI systems available in different sizes and capabilities that you access through an API.

Inference

The process of sending a request to a model and receiving a generated response.

Reasoning

Extended thinking capability that allows models to break down complex problems and show their step-by-step analysis before providing answers.

Multimodal

The ability to process and understand multiple input types together: text, images, video and documents in a single request.

RAG (Retrieval-Augmented Generation)

A technique that combines model responses with your own data sources to provide more accurate, contextual answers.

Amazon Nova models

Amazon Nova 2 includes the following models, each optimized for different use cases:

Model	Input modalities	Output modalities	Use cases
Nova 2 Lite	Text, images, video, documents	Text	High-volume applications prioritizing speed and cost efficiency
Nova 2 Sonic	Speech, text	Speech, text	Voice-enabled applications with fast response times
Nova Multimodal Embeddings	Text, images, documents, video, audio	Embeddings	Semantic search, recommendation systems and similarity matching

All models support up to 1 million tokens of context and can generate up to 65,536 tokens in a single response. Models with reasoning capabilities can perform extended thinking to solve complex problems step by step.

What can you build?

The following are examples of what you can build with Amazon Nova:

- **Intelligent document assistant** - Process large documents with up to 1 million tokens of context to answer questions and extract insights (with RAG)
- **Complex reasoning applications** - Solve multi-step problems with extended thinking that shows the model's step-by-step analysis (or with reasoning)
- **Video analysis pipeline** - Extract insights, generate summaries and identify key moments in video content at scale (Nova 2 Lite)
- **Voice-enabled AI agent** - Build conversational agents that understand speech input and respond with natural language (Nova 2 Sonic)

Benefits

Amazon Nova provides the following benefits:

Multimodal understanding

Process text, images, video, documents and speech in a single request. Amazon Nova models understand relationships across different input types.

Extended context

Support for up to 1 million tokens allows you to process entire codebases, lengthy documents, or extended conversations without losing context.

Advanced reasoning

Models with reasoning capabilities break down complex problems and show step-by-step analysis, improving accuracy for multi-step tasks.

Flexible deployment

Access models through Amazon Bedrock with no infrastructure to manage, or customize models through fine-tuning and reinforcement learning.

Built-in tools

Use web grounding to access real-time information and code interpreter to execute Python code without external integrations.

How Amazon Nova works

Amazon Nova models are foundation models that you access through Amazon Bedrock. The basic workflow is:

1. Your application sends a request to Amazon Bedrock with your input and configuration parameters.
2. The Amazon Nova model processes your input, applying reasoning if configured.
3. The model generates a response and returns it to your application.

You can enhance responses by using RAG to incorporate your data, enabling built-in tools, or customizing models through fine-tuning.

Pricing

Amazon Nova pricing is based on input and output tokens processed. Different models have different pricing tiers:

- Nova 2 Lite - Optimized for cost-effective, high-volume processing
- Nova 2 Sonic - Balanced pricing for voice-enabled applications

For current pricing information, see [Amazon Bedrock Pricing](#).

Next steps

- To learn about new features in Amazon Nova 2, see [What's new in Amazon Nova 2](#).
- To start using Amazon Nova, see [Getting started with Amazon Nova 2](#).
- To learn about core inference capabilities, see [Core inference](#).
- To customize models for your use case, see [Customizing Amazon Nova models on SageMaker AI](#).

What's new in Amazon Nova 2

Amazon Nova 2 introduces significant enhancements across understanding, creative and speech capabilities. The following sections describe the key new features and improvements.

New models

Nova 2 Lite

Enhanced models that offer advanced reasoning with extended thinking support, three-level intensity control and multimodal understanding optimized for production-scale agentic workflows.

Nova 2 Sonic

An upgraded conversational speech model with improved speech understanding, natural language processing and voice generation capabilities across seven languages.

Nova Multimodal Embeddings

A multimodal embedding model that processes text, image, document, video and audio inputs and generates embeddings.

New features and capabilities

Nova Multimodal Embeddings

Nova Multimodal Embeddings supports text, documents, images, video and audio through a single model, enabling cross-modal retrieval applications. Nova Multimodal Embeddings maps each of these content types into a unified semantic space, enabling you to conduct unimodal, cross-modal and multimodal vector operations, powering applications such as agentic retrieval-augmented generation (RAG) and multimodal semantic search.

Extended thinking and reasoning

Nova 2 Lite supports extended thinking, which allows the models to spend more time reasoning through complex problems before generating responses. This capability improves accuracy for multi-step reasoning tasks, including agentic workflows with multiple tools, advanced mathematics, complex planning and code generation.

To learn more, navigate to [Extended thinking in Amazon Nova 2](#).

Built-in tools

Amazon Nova 2 includes built-in tools that extend model capabilities without requiring external integrations:

- **Web grounding** – Accesses real-time information from the web to provide up-to-date responses and reduce hallucinations.
- **Code interpreter** – Executes Python code to perform calculations.

AI agent building

Amazon Nova 2 models are optimized for building AI agents. The models provide improved tool use, better reasoning for multi-step tasks and enhanced ability to maintain context across complex agent workflows.

Improved document understanding

Nova 2 Lite provides enhanced document processing capabilities with better understanding of complex document layouts, tables, charts and multi-page documents. The models extract information more accurately from PDFs, spreadsheets and other document formats.

Enhanced video understanding

Nova 2 Lite offers improved video analysis capabilities, including better visual perception, temporal understanding, action recognition and the ability to process longer video sequences with higher accuracy.

Model customization

Nova 2 Lite supports supervised fine-tuning (SFT) and reinforcement fine-tuning (RFT) on Amazon Bedrock and SageMaker AI, allowing you to adapt Amazon Nova 2 to your specific business needs.

Amazon Nova Forge

Amazon Nova Forge is a first-of-its-kind service that offers organizations the easiest and most cost-effective way to build their own frontier models using Amazon Nova.

Next steps

- To learn about Amazon Nova models and capabilities, see [What is Amazon Nova 2?](#).
- To start using Amazon Nova 2.0, see [Getting started with Amazon Nova 2](#).
- To explore core inference features, see [Core inference](#).

Getting started with Amazon Nova 2

You can invoke Amazon Nova 2 models through the following methods:

- Amazon Bedrock [InvokeModel](#) or [Converse](#) API operations
- Amazon Bedrock [Chat and Text playgrounds](#) for Nova 2 Lite, Nova 2 Sonic

Important

If you're new to AWS or to Amazon Bedrock, you must first set up Amazon Bedrock by following the steps at [Getting started with Amazon Bedrock](#). Alternatively, you can try the Amazon Nova models without entering Amazon Bedrock at nova.amazon.com or start building for free at nova.amazon.com/dev.

Select a topic to get started on using Amazon Nova with that method.

Getting started with Amazon Nova in the Amazon Bedrock console

This section describes how to use the playgrounds in the AWS console to submit a text prompt to Amazon Nova 2 models and generate a text or image response. Before you run the following examples, you should check that you have fulfilled the following prerequisites:

Prerequisites

- You have an AWS account and have permissions to access a role in that account with the necessary permissions for Amazon Bedrock. Otherwise, follow the steps at [Getting started with Amazon Bedrock](#).
- If you are accessing the model from a US region, you must use the US CRIS endpoint which involves adding the `us` prefix to the model ID (such as `us.amazon.nova-2-lite-v1:0`).
- If you are accessing the model from outside of the US, you can either use the global CRIS endpoint (such as `global.amazon.nova-2-lite-v1:0`) or prefix the region in the model ID (such as `us/eu/jp`).

Explore the text playground

The following example demonstrates how to use the text playground:

1. Open the Amazon Bedrock console at <https://console.aws.amazon.com/bedrock/>.
2. From the left navigation pane, choose **Chat / Text** under **Playgrounds**.
3. Choose **Select model** and select a provider and model. For this example, select **Amazon** then select a Amazon Nova 2 model (such as Nova 2 Lite, or Nova 2 Sonic). Then choose **Apply**.
4. Select a default prompt from below the text panel, or enter a prompt into the text panel, such as **Describe the purpose of a "hello world" program in one line**.
5. To explore the image understanding capabilities of Amazon Nova, you can upload an image in JPEG, PNG, GIF, or WEBP format that is less than or equal to 25 MB from your computer. After the image is uploaded, you can ask Amazon Nova about the image.
6. To explore the document understanding capabilities of Amazon Nova, you can upload documents in CSV, DOC, DOCX, HTML, MD, PDF, TXT, XLS, or XLSX format that are less than or equal to 4.5 MB. After the documents are uploaded, you can ask Amazon Nova about the documents.
7. To explore the video understanding capabilities of Amazon Nova, you can upload one video in MKV, MOV, or MP4 format that is less than or equal to 25 MB from your computer. You can use Amazon S3 for videos up to 1 GB. After the video is uploaded, you can ask Amazon Nova about the video.
8. Choose **Run** to run inference on the model. The generated text appears below your prompt in the text panel.

Getting started with the API

Before you get started with APIs, follow these procedures if you are new to AWS or need to install the AWS CLI or an AWS SDK. If neither of these apply to you, move to [Get credentials to grant programmatic access](#).

I'm new to AWS

If you do not have an AWS account, complete the following steps to create one. Alternatively, you can start building for free with Amazon Nova at nova.amazon.com/dev.

To sign up for an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user and use only the root user to perform [tasks that require root user access](#).

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <https://aws.amazon.com/> and choosing **My Account**.

Secure your AWS account root user

1. Sign in to the [AWS Management Console](#) as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see [Signing in as the root user](#) in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see [Enable a virtual MFA device for your AWS account root user \(console\)](#) in the *IAM User Guide*.

I need to install the AWS CLI or an AWS SDK

To install the AWS CLI, follow the steps at [Install or update to the latest version of the AWS CLI](#).

To install an AWS SDK, select the tab that corresponds to the programming language that you want to use at [Tools to Build on AWS](#).

AWS software development kits (SDKs) are available for many popular programming languages. Each SDK provides an API, code examples and documentation that make it easier for developers

to build applications in their preferred language. SDKs automatically perform useful tasks for you, such as:

- Cryptographically sign your service requests
- Retry requests
- Handle error responses

Get credentials to grant programmatic access

Users need programmatic access if they want to interact with AWS outside of the AWS Management Console. The way to grant programmatic access depends on the type of user that's accessing AWS.

To grant users programmatic access, choose one of the following options.

Which principal needs programmatic access?	To	By
IAM users	Limit the duration of long-term credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. • For the AWS CLI, see Authenticating using IAM user credentials in the <i>AWS Command Line Interface User Guide</i>. • For AWS SDKs and tools, see Authenticate using long-term credentials in the <i>AWS SDKs and Tools Reference Guide</i>. • For AWS APIs, see Managing access keys for IAM users in the <i>IAM User Guide</i>.

Which principal needs programmatic access?	To	By
IAM roles	Use temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Follow the instructions in Using temporary credentials with AWS resources in the <i>IAM User Guide</i> .
Users managed in IAM Identity Center	Use temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. - For the AWS CLI, see Configuring the AWS CLI to use IAM Identity Center in the <i>AWS Command Line Interface User Guide</i>. - For AWS SDKs, tools and AWS APIs, see IAM Identity Center authentication in the <i>AWS SDKs and Tools Reference Guide</i>.

Attach Amazon Bedrock permissions to a user or role

After setting up credentials for programmatic access, you need to configure permissions for a user or IAM role to have access to Amazon Bedrock-related actions. To set up these permissions, do the following:

1. On the AWS Management Console Home page, select the IAM service or navigate to the IAM console at <https://console.aws.amazon.com/iam/>.
2. Select **Users** or **Roles** and then select your user or role.
3. In the **Permissions** tab, choose **Add permissions** and then choose **Add AWS managed policy**. Choose the [AmazonBedrockFullAccess](#) AWS managed policy.
4. To allow the user or role to subscribe to models, choose **Create inline policy** and then specify the following permissions in the JSON editor:

```
{
```

```
"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "MarketplaceBedrock",
    "Effect": "Allow",
    "Action": [
      "aws-marketplace:ViewSubscriptions",
      "aws-marketplace:Unsubscribe",
      "aws-marketplace:Subscribe"
    ],
    "Resource": "*"
  }
]
```

Generate a response for a text prompt using an Amazon Nova model

After you've fulfilled all the prerequisites, you can test making model invocation requests to Amazon Nova models with a [Converse](#) request.

To install the AWS CLI, follow the steps at [Install or update to the latest version of the AWS CLI](#). Verify that you've set up your credentials to use Boto3 by following the steps at [Get credentials to grant programmatic access](#).

To generate a response for a text prompt in Nova 2 Lite by using the AWS CLI, run the following command in a terminal:

```
aws bedrock-runtime converse \
  --model-id us.amazon.nova-2-lite-v1:0 \
  --messages '[{"role":"user","content":[{"text":"Write a short poem"}]]' \
  --additional-model-request-fields '{"reasoningConfig":
{"type":"enabled","maxReasoningEffort":"low"}}'
```

Core inference

Inference is the process of sending a request to a Amazon Nova model and receiving a generated response. Amazon Nova models support inferencing through two API options:

- **Converse API** ([Converse](#), [ConverseStream](#)): Provides a consistent interface across different models, making it easier to switch between models or build applications that work with multiple models. Recommended for most use cases.
- **Invoke API** ([InvokeModel](#), [InvokeModelWithResponseStream](#)): Request payloads are structured specifically for each model's native format. Runs inference using the prompt and inference parameters provided in the request body.

Both APIs support the same core features including:

- Multi-turn conversations
- Multimodal inputs (text, images, video, audio)
- Tool use
- Guardrails
- Streaming responses

The request structure is nearly identical between the two APIs, differing only in how byte data (documents, images, video and audio) is encoded. For model request parameters unique to Amazon Nova models, such as `reasoningConfig` and `TopK`, these are placed within an additional `inferenceConfig` object within the `additionalModelRequestFields`. These are top-level parameters for `InvokeModel` and `InvokeModelWithResponseStream`.

Note

Review [Code Samples and Troubleshooting](#) for a list of code samples with Nova 2 models.

Set the `modelId` to one of the following to use Amazon Nova models:

Model

Model ID

Nova 2 Lite

- `global.amazon.nova-2-lite-v1:0`

Model	Model ID
	us.amazon.nova-2-lite-v1:0
Nova 2 Sonic	- global.amazon.nova-2-sonic-v1:0 - us.amazon.nova-2-sonic-v1:0
Nova Multimodal Embeddings	amazon.nova-2-multimodal-embeddings-v1:0

Important: Timeout Configuration

Important

Amazon Nova inference requests can take up to 60 minutes to complete. Configure your client timeout settings accordingly:

The following example is Python code. Users can check the documentation for their preferred SDK language version in that SDK's API docs.

```
from botocore.config import Config

bedrock = boto3.client(
    'bedrock-runtime',
    region_name='us-east-1',
    config=Config(
        read_timeout=3600 # 60 minutes
    )
)
```

Core Inference Topics

This section discusses the following topics:

- Using the converse API: A consistent interface offering compatibility across most Bedrock models
- Using the invoke API: An interface unique to Nova models and not portable to other Bedrock models
- Streaming responses: Real-time response generation

- Using Amazon Nova embeddings: Text embeddings capabilities
- On-demand inference: Pay-per-use inference model

Using the Converse API

The Converse API provides a unified interface for interacting with Amazon Nova models. It abstracts away model-specific details and provides a consistent way to handle multi-turn conversations, system prompts and streaming responses across all Amazon Nova models.

Topics

- [Request structure](#)
- [Using system prompts](#)
- [Inference parameters](#)
- [Using reasoning](#)

Request structure

- **Multi-turn conversations:** Maintain context across multiple exchanges
- **System prompts:** System instructions such as personas or response guidelines
- **Document chat:** Interact with and query documents or collections of documents
- **Vision:** Process and analyze images and video
- **Tool use:** Enable models to use external tools and APIs
- **Guardrails:** Apply content filtering and safety controls
- **Reasoning:** Extended thinking for complex problem-solving

A basic Converse API request includes the model ID and a list of messages:

```
import boto3

bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

response = bedrock.converse(
    modelId='us.amazon.nova-2-lite-v1:0',
    messages=[
        {
```

```

        'role': 'user',
        'content': [{'text': 'What is machine learning?'}]
    }
]
)

content_list = response["output"]["message"]["content"]
# Extract the first text block
text = next((item["text"] for item in content_list if "text" in item), None)
if text is not None:
    print(text)

```

Using system prompts

System prompts provide context and instructions to the model:

```

import boto3

bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

response = bedrock.converse(
    modelId='us.amazon.nova-2-lite-v1:0',
    system=[
        {'text': 'You are a helpful AI assistant specializing in cloud computing.'}
    ],
    messages=[
        {
            'role': 'user',
            'content': [{'text': 'Explain serverless computing.'}]
        }
    ]
)

# Print the response text
content_list = response["output"]["message"]["content"]
text = next((item["text"] for item in content_list if "text" in item), None)
if text is not None:
    print(text)

```

Inference parameters

Control the model's output using inference parameters. The following are available inference parameters.

- **maxTokens (integer):** Maximum number of tokens to generate (up to 65,000). If not specified, the model uses a dynamic default based on the request context.
- **temperature (float):** Controls randomness (0.0-1.0, default 0.7). Lower values make output more deterministic
- **topP (float):** Nucleus sampling threshold (0-1, default 0.9). Lower values make output more focused
- **stopSequences (array):** Sequences of characters that stop generation when encountered

Example:

```
import boto3
bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

response = bedrock.converse(
    modelId='us.amazon.nova-2-lite-v1:0',
    messages=[
        {
            'role': 'user',
            'content': [{'text': 'Write a short story.'}]
        }
    ],
    inferenceConfig={
        'maxTokens': 512,
        'temperature': 0.7,
        'topP': 0.9,
        'stopSequences': ['END']
    }
)

content_list = response["output"]["message"]["content"]
text = next((item["text"] for item in content_list if "text" in item), None)
if text is not None:
    print(text)
```

Using reasoning

Nova 2 Lite supports extended thinking for complex problem-solving. Enable reasoning with `reasoningConfig`.

By default, reasoning is disabled to optimize for speed and cost on simple queries. When you need to go beyond these straightforward tasks, you can enable reasoning. Nova 2 provides flexible control over reasoning depth through three effort levels:

Low effort (`maxReasoningEffort: "low"`)

BEST FOR: TASKS WITH ADDED COMPLEXITY REQUIRING STRUCTURED THINKING. For example, you can use this for code review and improvement suggestions where the model needs to carefully consider existing code quality, analysis tasks that require thoughtful consideration of multiple factors, or problem-solving scenarios that benefit from a methodical approach. Low effort is ideal for compound tasks where basic reasoning improves accuracy without requiring deep multi-step planning.

Medium effort (`maxReasoningEffort: "medium"`)

BEST FOR: MULTI-STEP TASKS AND CODING WORKFLOWS. For example, you can use this for software development and debugging where the model needs to understand existing code structure before implementing changes, code generation that requires coordination across multiple files or components, multi-step calculations with interdependencies, or planning tasks with multiple constraints. Medium effort is optimal for agentic workflows that coordinate multiple tools and require the model to maintain context across several sequential operations.

High effort (`maxReasoningEffort: "high"`)

BEST FOR: STEM REASONING AND ADVANCED PROBLEM-SOLVING. For example, you can use this for advanced mathematical problems and proofs that require rigorous step-by-step verification, scientific analysis and research tasks demanding deep investigation, complex system design with architectural considerations across multiple dimensions, or critical decision-making scenarios with significant implications. High effort delivers maximum accuracy for tasks requiring sophisticated reasoning, careful evaluation of alternatives, and thorough validation of conclusions.

The following examples show different reasoning effort levels:

Low effort

```
import boto3

bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

response = bedrock.converse(
    modelId='us.amazon.nova-2-lite-v1:0',
    system=[{"text": "You are a highly capable personal assistant"}],
    messages=[{
        "role": "user",
        "content": [{"text": "Provide a meal plan for a gluten free family of 4."}]
    }],
    inferenceConfig={
        "temperature": 0.7,
        "topP": 0.9,
        "maxTokens": 10000
    },
    additionalModelRequestFields={
        "reasoningConfig": {
            "type": "enabled",
            "maxReasoningEffort": "low"
        }
    }
)
```

Reasoning parameters:

The following are reasoning parameters

- type: enabled or disabled (default: disabled)
- maxReasoningEffort: low, medium, or high. This is required when reasoning is enabled.

Note

Temperature, topP and topK cannot be used with maxReasoningEffort set to high. This will cause an error.

The response includes reasoning content:

```
{
  "output": {
    "message": {
      "role": "assistant",
      "content": [
        {
          "reasoningContent": {
            "reasoningText": {
              "text": "[REDACTED]"
            }
          }
        },
        {
          "text": "Based on the premises, we can conclude..."
        }
      ]
    }
  },
  "stopReason": "end_turn"
}
```

Note

With Amazon Nova 2, reasoning content displays as [REDACTED]. You're still charged for reasoning tokens as they contribute to improved output quality. We include this field in the response structure now to preserve the option of exposing reasoning content in the future. We are actively working with customers to determine the best approach for surfacing the model's reasoning process.

Using the Invoke API

The Invoke API provides direct access to Amazon Nova models with more ability to control the request and response format. Unlike the Converse API which abstracts model-specific details, the Invoke API allows you to work directly with the model's native request and response structures.

Note

The Invoke API supports the same features as the Converse API except for document input modality, which is specific to the Converse API.

Request structure

An Invoke API request requires the model ID and a JSON request body:

```
import boto3
import json

bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

request_body = {
    'messages': [
        {
            'role': 'user',
            'content': [{'text': 'What is machine learning?'}]
        }
    ],
    'inferenceConfig': {
        'maxTokens': 512,
        'temperature': 0.7
    }
}

response = bedrock.invoke_model(
    modelId='us.amazon.nova-2-lite-v1:0',
    body=json.dumps(request_body)
)

response_body = json.loads(response['body'].read())
content_list = response_body["output"]["message"]["content"]
# Extract the first text block
text_block = next((item for item in content_list if "text" in item), None)
if text_block is not None:
    print(text_block["text"])
```


Request parameters

The Invoke API supports the following key parameters:

- **messages**: Array of conversation messages with role and content
- **system**: Optional system prompt for context and instructions
- **inferenceConfig**: Parameters controlling model output (temperature, maxTokens, topP, topK, stopSequences, reasoningConfig)
- **toolConfig**: Tool specifications and tool choice for function calling

Streaming responses

Streaming allows you to receive model responses incrementally as they are generated, providing a more interactive user experience. Both the Converse API and Invoke API support streaming.

Streaming with ConverseStream

Use `ConverseStream` to receive responses as a stream of events:

```
import boto3

bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

response = bedrock.converse_stream(
    modelId='us.amazon.nova-2-lite-v1:0',
    messages=[
        {
            'role': 'user',
            'content': [{ 'text': 'Write a short story about AI.' }]
        }
    ]
)

for event in response['stream']:
    if 'contentBlockDelta' in event:
        delta = event['contentBlockDelta']['delta']
        if 'text' in delta:
            print(delta['text'], end='', flush=True)
```

Streaming with InvokeModelWithResponseStream

Use `InvokeModelWithResponseStream` for streaming with the Invoke API:

```
import boto3
import json

bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

request_body = {
    'messages': [
        {
            'role': 'user',
            'content': [{'text': 'Explain quantum computing.'}]
        }
    ]
}

response = bedrock.invoke_model_with_response_stream(
    modelId='us.amazon.nova-2-lite-v1:0',
    body=json.dumps(request_body)
)

for event in response['body']:
    chunk = json.loads(event['chunk']['bytes'])
    if 'contentBlockDelta' in chunk:
        delta = chunk['contentBlockDelta']['delta']
        if 'text' in delta:
            print(delta['text'], end='', flush=True)
```

Stream event types

Streaming responses include several event types:

- `messageStart`: Indicates the start of a message
- `contentBlockStart`: Indicates the start of a content block
- `contentBlockDelta`: Contains incremental text or data
- `contentBlockStop`: Indicates the end of a content block
- `messageStop`: Indicates the end of the message with stop reason
- `metadata`: Contains usage information (token counts)

Using Amazon Nova embeddings

Amazon Nova Multimodal Embeddings is a multimodal embeddings model for agentic RAG and semantic search applications. It supports text, documents, images, video and audio through a single model, enabling cross-modal retrieval. Nova Multimodal Embeddings maps each of these content types into a unified semantic space, enabling you to conduct unimodal, cross-modal and multimodal vector operations.

When a piece of content is passed through Nova embeddings, the model converts that content into a universal numerical format, referred to as a vector. A vector is a set of numerical values that can be used for various search functionalities. Similar content is given a closer vector than less similar content.

Applications:

- **Semantic Content Retrieval and Recommendation:** Generate embeddings for your content, then use them to find similar items or provide personalized recommendations.
- **Multimodal Search:** Combine embeddings from different content types to enable cross-modal search capabilities.
- **RAG:** Generate embeddings from multimodal content such as documents with interleaved text and images to power your retrieval workflow for GenAI applications.

Key features

These key features distinguish Nova Embeddings:

- **Support for text, image, document image, video, and audio in a unified semantic space.** Maximum context length is 8K tokens or 30s of video and 30s of audio.
- **Synchronous and asynchronous APIs:** The API supports both synchronous and asynchronous use.
- **Large file segmentation:** The async API makes it easy to work with large inputs by providing API built segmentation for long text, video, and audio, controlled by user-defined parameters. The model will generate a single embedding for each segment.
- **Video with audio:** Process video with audio simultaneously. Specify if you would like a single embedding representing both modalities or two separate embeddings.
- **Embedding purpose:** Optimize your embeddings depending on the intended downstream application (retrieval/RAG/Search, classification, clustering).

- Dimension sizes: 4 dimension sizes to trade-off embedding accuracy and vector storage cost: 3072; 1024; 384; 256.
- Input methods: Pass content to be embedded by specifying an S3 URI or inline as a base64 encoding.

Generating embeddings

Complete the following to generate embeddings.

Synchronous embedding for text:

```
import boto3
import json

# Create the Bedrock Runtime client.
bedrock_runtime = boto3.client(
    service_name='bedrock-runtime',
    region_name='us-east-1',
)

# Define the request body.
request_body = {
    'taskType': 'SINGLE_EMBEDDING',
    'singleEmbeddingParams': {
        'embeddingPurpose': 'GENERIC_INDEX',
        'embeddingDimension': 3072,
        'text': {'truncationMode': 'END', 'value': 'Hello, World!'},
    },
}

try:
    # Invoke the Nova Embeddings model.
    response = bedrock_runtime.invoke_model(
        body=json.dumps(request_body, indent=2),
        modelId='amazon.nova-2-multimodal-embeddings-v1:0',
        accept='application/json',
        contentType='application/json',
    )

except Exception as e:
    # Add your own exception handling here.
    print(e)
```

```
# Print the request ID.
print('Request ID:', response.get('ResponseMetadata').get('RequestId'))

# Print the response body.
response_body = json.loads(response.get('body').read())
print(json.dumps(response_body, indent=2))
```

try:

```
response = bedrock_runtime.invoke_model(
    #####body=json.dumps(request_body, indent=2),
    #####modelId='amazon.nova-2-multimodal-embeddings-v1:0',
    #####accept='application/json',
    #####contentType='application/json',
    ####)
except Exception as e:
    #####print(e)

print('Request ID:', response.get('ResponseMetadata').get('RequestId'))
response_body = json.loads(response.get('body').read())
print(json.dumps(response_body, indent=2))
```

Asynchronous embedding for video:

```
import boto3

bedrock_runtime = boto3.client(
    ####service_name='bedrock-runtime',
    ####region_name='us-east-1',
    )

model_input = {
    ####'taskType': 'SEGMENTED_EMBEDDING',
    ####'segmentedEmbeddingParams': {
    #####'embeddingPurpose': 'GENERIC_INDEX',
    #####'embeddingDimension': 3072,
    #####'video': {
    #####'format': 'mp4',
    #####'embeddingMode': 'AUDIO_VIDEO_COMBINED',
    #####'source': {
    #####'s3Location': {'uri': 's3://my-bucket/path/to/video.mp4'}
    #####},
    #####},
```

```
#####'segmentationConfig': {
#####'durationSeconds': 15### Segment into 15 second chunks
#####},
#####},
####},
}
```

try:

```
response = bedrock_runtime.start_async_invoke(
#####modelId='amazon.nova-2-multimodal-embeddings-v1:0',
#####modelInput=model_input,
#####outputDataConfig={
#####'s3OutputDataConfig': {
#####'s3Uri': 's3://my-bucket'
#####}
#####},
####)
except Exception as e:
####print(e)

print('Request ID:', response.get('ResponseMetadata').get('RequestId'))
print('Invocation ARN:', response.get('invocationArn'))
```

Batch inference

With batch inference, you can submit multiple requests and generate embeddings asynchronously. Batch inference helps you process many requests efficiently by sending a single request and generating responses in an Amazon S3 bucket.

- After defining model inputs in files you create, you upload the files to an S3 bucket.
- You then submit a batch inference request and specify the S3 bucket.
- After the job is complete, you can retrieve the output files from S3.
- You can use batch inference to improve the performance of model inference on large datasets.

Format and upload your batch inference data

You must add your batch inference data to an S3 location that you'll choose or specify when submitting a model invocation job. The S3 location must the following items:

At least one JSONL file that defines the model inputs. A JSONL contains rows of JSON objects. Your JSONL file must end in the extension .jsonl and be in the following format:

```
{
  "recordId": "record001",
  "modelInput": {
    "taskType": "SINGLE_EMBEDDING",
    "singleEmbeddingParams": {
      "embeddingPurpose": "GENERIC_INDEX",
      "embeddingDimension": 3072,
      "text": {
        "source": {
          "s3Location": {
            "uri": "s3://batch-inference-input-bucket/text_001.txt",
            "bucketOwner": "111122223333"
          }
        },
        "truncationMode": "END"
      }
    }
  }
}
```

Supported input file types

- Image Formats: PNG, JPEG, WEBP, GIF
- Audio Formats: MP3, WAV, OGG
- Video Formats: MP4, MOV, MKV, WEBM, FLV, MPEG, MPG, WMV, 3GP

On-demand inference

On-demand inference provides serverless access to Amazon Nova models without requiring provisioned capacity. This mode automatically scales to handle your workload and charges based on usage.

Benefits

On-demand inference offers several advantages:

- **No capacity planning:** Automatically scales to meet demand

- **Pay per use:** Charged only for tokens processed
- **Instant availability:** No provisioning or warm-up time required
- **Cost effective:** Ideal for variable or unpredictable workloads

Using on-demand inference

On-demand inference is the default mode for Amazon Nova models. Simply specify the model ID when making API calls:

```
import boto3

bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

response = bedrock.converse(
    modelId='us.amazon.nova-2-lite-v1:0',
    messages=[
        {
            'role': 'user',
            'content': [{'text': 'Hello, Nova!'}]
        }
    ]
)

# Print the response text
content_list = response["output"]["message"]["content"]
text = next((item["text"] for item in content_list if "text" in item), None)
if text is not None:
    print(text)
```

Pricing

On-demand inference is billed based on the number of input and output tokens processed. For current pricing details, see [Amazon Bedrock pricing](#).

Quotas and limits

On-demand inference has default quotas that vary by model and region. To request quota increases, use the [Service Quotas console](#).

Using Amazon Nova 2 capabilities

This guide provides task-based recipes and practical examples to help you integrate Nova's capabilities into your applications quickly and effectively. It is designed for developers and engineers who want to:

- Integrate Nova models into their applications via API
- Understand preprocessing requirements and optimization strategies
- Implement multimodal understanding workflows

Each section provides:

- Key technical information about supported formats, size limits, and token estimation
- Code examples that demonstrate how to use API with both direct upload and Amazon S3
- Use cases that show practical applications
- Limitations to consider when designing your application

Prerequisites

Before you begin, ensure you have:

- An AWS account with access to Amazon Bedrock
- Appropriate IAM permissions for invoking Bedrock models
- AWS SDK installed (Boto3 for Python examples)
- Basic familiarity with AWS services, such as Amazon S3 (for large file handling)

Reasoning capabilities

Amazon Nova 2 Lite supports extended thinking, which is disabled by default. When enabled, the model generates internal reasoning tokens that improve response quality. In Amazon Nova 2 Lite, reasoning content is redacted in the output and displays as [REDACTED], though you are still charged for these tokens.

The reasoning field is included in the response structure to preserve the option of exposing this content in future releases. For more information about extended thinking, see [the section called “Extended Thinking”](#) and [the section called “Code library”](#).

Multimodal understanding

Amazon Nova 2 Lite can understand multiple input modalities. This model is equipped with vision capabilities that enable it to comprehend and analyze images, documents, videos, and speech to infer and answer questions based on the content provided.

This section outlines guidelines for working with images, documents, and videos in Amazon Nova including preprocessing strategies employed, code examples and relevant limitations to consider.

Supported content type by modality

The following information details the file formats supported by each media file type and the accepted input method.

Media file type	File formats supported	Input method	Size limitations	Number of objects
Image	PNG, JPEG, GIF, WebP Note: If you use an animated GIF or WebP file, only the first frame will be used.	Embedding data in the request	25 MB	5
		If you use the Converse API, encode data as bytes. If you use the Invoke API, encode data as Base64 string.		
		Amazon S3 URI	2 GB total	1000
Video	MP4, MOV, MKV, WebM, FLV, MPEG, MPG, WMV, 3GP	Embedding data in the request If you use the Converse API, encode data as bytes.	25 MB	1

Media file type	File formats supported	Input method	Size limitations	Number of objects
		If you use the Invoke API, encode data as Base64 string.		
		Amazon S3 URI	1 GB	1

Topics

- [Image understanding](#)
- [Video understanding](#)
- [Document understanding](#)

Image understanding

Image understanding refers to Amazon Nova's ability to process an image and conduct a variety of computer vision tasks such as:

- Performing object detection
- Answering questions about images through Visual Question Answering (VQA)
- Classifying and summarizing images
- Performing bounding box detection
- Optical Character Recognition (OCR)
- Object counting

Images can be included as a prompt passed to the API as byte arrays or via S3 URI.

Key technical information

The following is key technical information to note when you work with this capability.

Image sizing and rescaling

Amazon Nova automatically rescales images to optimize quality and performance:

- Determines the closest aspect ratio (such as 1:1, 1:2, 2:3 and so on)

- Rescales so that one side ≥ 896 px or matches the shorter side of the original image—whichever is larger
- Maintains the aspect ratio
- Supports up to $8,000 \times 8,000$ px resolution

Bounding box coordinates:

- Useful for tasks such as identifying elements in screenshots or image grounding
- Coordinates can be rescaled to match the image's original dimensions in post-processing
- Returns bounding boxes on a $[0, 1000]$ scale.

Image-token estimation

Amazon Nova converts each image into tokens for processing. The number of tokens depends on the resolution and aspect ratio of the image.

The following are examples of approximate token counts based on image resolution:

Image resolution	Estimated tokens
900 x 450	515
900 x 900	~1,035
1400 x 900	~1,600
1800 x 900	~2,060
1300 x 1300	~2,155

Image understanding examples

For an example of how to embed image data directly in the request, refer to the Multimodal input using embedded asset - Converse API (non-streaming) example in the [Code library](#).

To upload large image files or multiple image files, where the overall payload is greater than 25 MB, use Amazon S3. For a full example of how to use Amazon S3 URI references for image input,

refer to the Multimodal input using S3 URI - Converse API (non-streaming) example in the [Code library](#).

 Note

When using S3, ensure the Amazon Bedrock service has permission to access the bucket and object.

Key limitations

The following list outlines current limitations of image understanding models:

- **Multilingual image understanding:** The models have limited understanding of multilingual images and video frames and can struggle or hallucinate on simple tasks.
- **People identification:** Amazon Nova 2 models do not support the capability to identify or name individuals in images, documents or videos.
- **Spatial reasoning:** Amazon Nova 2 models have limited spatial reasoning capabilities. They may struggle with tasks that require precise localization or layout analysis.
- **Small text in images and videos:** If the text in the image or video is too small, consider increasing relative size of the text in the image by cropping to relevant section while preserving necessary context.

Video understanding

Video understanding refers to Amazon Nova's ability to process video inputs and perform a range of video comprehension tasks such as:

- Analyzing key frames and summarizing video content
- Answering questions about video segments (Video Question Answering, or Video QA)
- Detecting and tracking objects across frames
- Identifying actions, scenes and events
- Performing temporal segmentation to locate specific moments
- Generating descriptive captions or summaries of video sequences

Key technical information

The following is key technical information to note when you work with this capability.

Video size information

Amazon Nova video understanding capabilities support multi-aspect ratio. All videos are resized with distortion (up or down, based on the original aspect ratio) to 672 × 672 square dimensions before they are input to the model.

The model utilizes a dynamic sampling strategy based on video length. For videos 16 minutes or less in length, Amazon Nova 2 Lite samples 1 frame per second (FPS). For videos longer than 16 minutes, the sampling rate decreases to maintain a consistent 960 frames sampled, with the frame sampling rate varying accordingly. This approach is designed to provide more accurate scene-level video understanding for shorter videos compared to longer video content.

We recommend that you keep the video length less than 1 hour for low motion and less than 16 minutes for anything with high motion.

There should be no difference when analyzing a 4k version of a video and a Full HD version. Similarly, because the sampling rate is 1 FPS, a 60 FPS video should perform as well as a 30 FPS video. Using a resolution and FPS that is higher than what is required is not beneficial because of the 1 GB limit in video size. Doing so will limit the video length that fits in that size limit, so, you may want to pre-process videos longer than 1 GB.

Video tokens

The length of the video is the main factor impacting the number of tokens generated. To calculate the approximate cost, multiply the estimated number of video tokens by the per-token price for the specific model being utilized.

The following table provides some approximations of frame sampling and token utilization per video length for Amazon Nova 2 Lite:

Video length	Frames to sample	Sample rate fps	Approximate tokens
10 seconds	10	1	2,880
30 seconds	30	1	8,640

Video length	Frames to sample	Sample rate fps	Approximate tokens
16 minutes	960	1	276,480
20 minutes	1200	1	345,600
30 minutes	1800	1	518,400
45 minutes	2700	1	777,600

Video understanding examples

For an example of how to embed video data directly in the request, refer to the Multimodal input using embedded asset - Converse API (non-streaming) example in the [Code library](#).

For an example of how to use S3 URI references in video input, refer to the Multimodal input using S3 URI - Converse API (non-streaming) example in the [Code library](#).

Key limitations

The following are key model limitations, where model accuracy and performance might not be guaranteed:

- **No audio support:** The Amazon Nova models are currently trained to process and understand video content solely based on the visual frames. Audio tracks in videos are not processed or analyzed.
- **Multilingual image understanding:** The Amazon Nova models have limited understanding of multilingual images and video frames. They might struggle or hallucinate on simple tasks.
- **People identification:** The Amazon Nova models do not support the capability to identify or name individuals in images, documents, or videos. The models will not provide names or identities of people in visual content.
- **Small text in videos:** If the text in the image or video is too small, consider increasing relative size of the text in the video.
- **Spatial reasoning:** Amazon Nova 2 models have limited spatial reasoning capabilities. They may struggle with tasks that require precise understanding of object positions, distances, or spatial relationships in videos.
- **Inappropriate content:** The Amazon Nova models will not process inappropriate or explicit images that violate the Acceptable Use Policy.

- **Healthcare applications:** Due to the sensitive nature of these artifacts, even though Amazon Nova models could give general analysis on some healthcare images or videos, we do not recommend their use to interpret sensitive medical images like complex diagnostic scans. The response of Amazon Nova models should never be considered a substitute for professional medical advice.

Document understanding

Amazon Nova's document understanding capability allows you to include entire documents (PDFs, Word files, spreadsheets and so on) as part of your prompt. This enables the model to analyze, summarize, extract information from, or answer questions about document content.

Amazon Nova 2 Lite can interpret both the text and visual elements (like charts or tables) within these documents. This enables use cases such as question-answering, summarization, and analysis of lengthy reports or scanned documents.

Key document understanding features include a very large context window (1M tokens) for long documents and the ability to handle multiple documents in one query.

Supported document modalities and formats

Amazon Nova distinguishes between two types of document inputs:

- **Text-based documents**, like TXT, CSV, HTML, Markdown, or DOC files, are processed primarily for their textual content. Amazon Nova understands and extracts information from the text in these documents.
- **Media-based documents**, like PDF or DOCX files, may contain complex layouts, images, charts, or embedded graphics. For media-based documents, Amazon Nova employs vision-based understanding to interpret visual content—such as charts, tables, diagrams, or screenshots—alongside the document's text.

Supported file formats include common document types such as:

- Plain text and structured text files: CSV, TXT
- Spreadsheets: XLS, XLSX, HTML, Markdown
- Standard image formats (for images within documents): PNG, JPG, GIF, WebP
- Document formats: DOC, DOCX, PDF

- PDFs that contain image encodings, such as CMYK or SVG are not supported.

Document size limits and usage guidelines

Constraint	Limit
Maximum number of documents	Up to 5 documents per request (applies to both direct upload and Amazon S3)
Text-based document size	Each text document must be equal to or less than 4.5 MB
Media-based document size	For PDF and DOCX files, there is no individual file size limit. When using direct upload, the combined size of all media documents must be less or equal to 25 MB. When using Amazon S3, the combined size of all media documents must be less or equal to 2 GB.
Unsupported PDF content	PDFs that contain CMYK color profiles or SVG images are not supported.

Pricing

Amazon Nova uses token-based pricing: you pay for input tokens (everything you send, including attached documents) and output tokens (the model's response).

Estimating tokens for PDFs: For planning, assume a standard 8.5x11 inch PDF page $\approx$ 2,560 input tokens (this estimate covers both text and visual elements on a typical page).

Examples: Using Nova's document understanding via API and S3

For an example of how to use it via API, refer to the Multimodal input using embedded asset - Converse API (non-streaming) example in the [Code library](#).

For an example of how to use it via S3, refer to the Multimodal input using S3 URI - Converse API (non-streaming) example in the [Code library](#).

Speech-to-Speech (Amazon Nova 2 Sonic)

Amazon Nova 2 Sonic enables real-time conversational AI with speech input and output. The following section covers advanced capabilities for building interactive voice assistants, customer service automation, and conversational applications.

Key features

Amazon Nova 2 Sonic provides the following capabilities:

- State-of-the-art streaming speech understanding with bidirectional streaming API that enables real-time, low-latency multi-turn conversations.
- Multilingual support with automatic language detection and switching. Expressive voices are offered, including both masculine-sounding and feminine-sounding voices, in the following languages:
 - English (US, UK, India, Australia)
 - French
 - Italian
 - German
 - Spanish
 - Portuguese
 - Hindi
- Polyglot voices that can speak any of the supported languages to enable a consistent user experience even when the user switches languages within the same session.
- Robustness to background noise for real world deployment scenarios.
- Robustness to different accents for supported languages.
- Natural, human-like conversational AI experiences with contextual richness across all supported languages.
- Adaptive speech response that dynamically adjusts delivery based on the prosody of the input speech.
- Intelligent turn-taking that detects when users finish speaking and when the assistant should respond, creating natural dialogue rhythm.

- Graceful handling of user interruptions without dropping conversational context.
- Knowledge grounding with enterprise data using Retrieval Augmented Generation (RAG).
- Function calling and agentic workflow support for building complex AI applications.
- Asynchronous tool handling that executes tool calls while maintaining conversation flow, allowing the assistant to continue speaking while tools process in the background.
- Cross-modal input support for both audio and text inputs within the same conversation, enabling flexible interaction patterns.
- Connection limit of 8 minutes, with connection renewal and session continuation pattern available in code samples.

Getting started with speech-to-speech

The following sections provide an example and step-by-step explanation of how to implement a simple, real-time audio streaming application using Amazon Nova 2 Sonic. This simplified version demonstrates the core functionality needed to create an audio conversation with the Amazon Nova 2 Sonic model.

You can access the following example in our [Nova samples GitHub repo](#).

There is a connection limit of 8 minutes, with connection renewal and conversation continuation pattern available on [GitHub](#).

State the imports and configuration

This section imports necessary libraries and sets audio configuration parameters:

- `asyncio`: For asynchronous programming
- `base64`: For encoding and decoding audio data
- `pyaudio`: For audio capture and playback
- Amazon Bedrock SDK components for streaming
- Audio constants define the format of audio capture (16kHz sample rate, mono channel)

```
import os
import asyncio
import base64
```

```
import json
import uuid
import pyaudio
from aws_sdk_bedrock_runtime.client import BedrockRuntimeClient,
    InvokeModelWithBidirectionalStreamOperationInput
from aws_sdk_bedrock_runtime.models import
    InvokeModelWithBidirectionalStreamInputChunk, BidirectionalInputPayloadPart
from aws_sdk_bedrock_runtime.config import Config, HTTPAuthSchemeResolver,
    SigV4AuthScheme
from smithy_aws_core.identity import EnvironmentCredentialsResolver

# Audio configuration
INPUT_SAMPLE_RATE = 16000
OUTPUT_SAMPLE_RATE = 24000
CHANNELS = 1
FORMAT = pyaudio.paInt16
CHUNK_SIZE = 1024
```

Define the SimpleNovaSonic class

The SimpleNovaSonic class is the main class that handles the Amazon Nova Sonic interaction:

- `model_id`: The Amazon Nova Sonic model ID (`amazon.nova-2-sonic-v1:0`)
- `region`: The AWS Region, the default is `us-east-1`
- Unique IDs for prompt and content tracking
- An asynchronous queue for audio playback

```
class SimpleNovaSonic:
    def __init__(self, model_id='amazon.nova-2-sonic-v1:0', region='us-east-1'):
        self.model_id = model_id
        self.region = region
        self.client = None
        self.stream = None
        self.response = None
        self.is_active = False
        self.prompt_name = str(uuid.uuid4())
        self.content_name = str(uuid.uuid4())
        self.audio_content_name = str(uuid.uuid4())
        self.audio_queue = asyncio.Queue()
        self.display_assistant_text = False
```

Initialize the client

This method configures the Amazon Bedrock client with the following:

- The appropriate endpoint for the specified region
- Authentication information using environment variables for AWS credentials
- The SigV4 authentication scheme for the AWS API calls

```
def _initialize_client(self):
    """Initialize the Bedrock client."""
    config = Config(
        endpoint_uri=f"https://bedrock-runtime.{self.region}.amazonaws.com",
        region=self.region,
        aws_credentials_identity_resolver=EnvironmentCredentialsResolver(),
        auth_scheme_resolver=HTTPAuthSchemeResolver(),
        auth_schemes={"aws.auth#sigv4": SigV4AuthScheme(service="bedrock")}
    )
    self.client = BedrockRuntimeClient(config=config)
```

Handle events

This helper method sends JSON events to the bidirectional stream, which is used for all communication with the Amazon Nova Sonic model:

```
async def send_event(self, event_json):
    """Send an event to the stream."""
    event = InvokeModelWithBidirectionalStreamInputChunk(
        value=BidirectionalInputPayloadPart(bytes_=event_json.encode('utf-8'))
    )
    await self.stream.input_stream.send(event)
```

Start the session

This method initiates the session and setups the remaining events to start audio streaming. These events need to be sent in the same order.

```
async def start_session(self):
    """Start a new session with Nova Sonic."""
    if not self.client:
        self._initialize_client()
```

```

# Initialize the stream
self.stream = await self.client.invoke_model_with_bidirectional_stream(
    InvokeModelWithBidirectionalStreamOperationInput(model_id=self.model_id)
)
self.is_active = True

# Send session start event
session_start = '''
{
  "event": {
    "sessionStart": {
      "inferenceConfiguration": {
        "maxTokens": 1024,
        "topP": 0.9,
        "temperature": 0.7
      },
      "turnDetectionConfiguration": {
        "endpointingSensitivity": "HIGH"
      }
    }
  }
}
'''
await self.send_event(session_start)

# Send prompt start event
prompt_start = f'''
{{
  "event": {{
    "promptStart": {{
      "promptName": "{self.prompt_name}",
      "textOutputConfiguration": {{
        "mediaType": "text/plain"
      }},
      "audioOutputConfiguration": {{
        "mediaType": "audio/lpcm",
        "sampleRateHertz": 24000,
        "sampleSizeBits": 16,
        "channelCount": 1,
        "voiceId": "matthew",
        "encoding": "base64",
        "audioType": "SPEECH"
      }}
    }}
  }}
}}
'''

```

```

        }}
    }}
}}
...
await self.send_event(prompt_start)

# Send system prompt
text_content_start = f'''
{{
    "event": {{
        "contentStart": {{
            "promptName": "{self.prompt_name}",
            "contentName": "{self.content_name}",
            "type": "TEXT",
            "interactive": true,
            "role": "SYSTEM",
            "textInputConfiguration": {{
                "mediaType": "text/plain"
            }}
        }}
    }}
}}
...
await self.send_event(text_content_start)

system_prompt = "You are a friendly assistant. The user and you will engage in
a spoken dialog " \
    "exchanging the transcripts of a natural real-time conversation. Keep your
responses short, " \
    "generally two or three sentences for chatty scenarios."

text_input = f'''
{{
    "event": {{
        "textInput": {{
            "promptName": "{self.prompt_name}",
            "contentName": "{self.content_name}",
            "content": "{system_prompt}"
        }}
    }}
}}
...
await self.send_event(text_input)

```

```

text_content_end = f'''
{{
    "event": {{
        "contentEnd": {{
            "promptName": "{self.prompt_name}",
            "contentName": "{self.content_name}"
        }}
    }}
}}
'''
await self.send_event(text_content_end)

# Start processing responses
self.response = asyncio.create_task(self._process_responses())

```

Handle audio input

These methods handle the audio input lifecycle:

- `start_audio_input`: Configures and starts the audio input stream
- `send_audio_chunk`: Encodes and sends audio chunks to the model
- `end_audio_input`: Properly closes the audio input stream

```

async def start_audio_input(self):
    """Start audio input stream."""
    audio_content_start = f'''
    {{
        "event": {{
            "contentStart": {{
                "promptName": "{self.prompt_name}",
                "contentName": "{self.audio_content_name}",
                "type": "AUDIO",
                "interactive": true,
                "role": "USER",
                "audioInputConfiguration": {{
                    "mediaType": "audio/lpcm",
                    "sampleRateHertz": 16000,
                    "sampleSizeBits": 16,
                    "channelCount": 1,
                    "audioType": "SPEECH",

```



```

        "encoding": "base64"
    }}
    }}
    }}
    ...
    await self.send_event(audio_content_start)

async def send_audio_chunk(self, audio_bytes):
    """Send an audio chunk to the stream."""
    if not self.is_active:
        return

    blob = base64.b64encode(audio_bytes)
    audio_event = f'''
    {{
        "event": {{
            "audioInput": {{
                "promptName": "{self.prompt_name}",
                "contentName": "{self.audio_content_name}",
                "content": "{blob.decode('utf-8')}"
            }}
        }}
    }}
    '''
    await self.send_event(audio_event)

async def end_audio_input(self):
    """End audio input stream."""
    audio_content_end = f'''
    {{
        "event": {{
            "contentEnd": {{
                "promptName": "{self.prompt_name}",
                "contentName": "{self.audio_content_name}"
            }}
        }}
    }}
    '''
    await self.send_event(audio_content_end)

```

End the session

This method properly closes the session by:

- Sending a promptEnd event
- Sending a sessionEnd event
- Closing the input stream

```
async def end_session(self):
    """End the session."""
    if not self.is_active:
        return

    prompt_end = f'''
    {{
      "event": {{
        "promptEnd": {{
          "promptName": "{self.prompt_name}"
        }}
      }}
    }}
    '''
    await self.send_event(prompt_end)

    session_end = '''
    {
      "event": {
        "sessionEnd": {}
      }
    }
    '''
    await self.send_event(session_end)
    # close the stream
    await self.stream.input_stream.close()
```

Handle responses

This method continuously processes responses from the model and does the following:

- Waits for output from the stream.

- Parses the JSON response.
- Handles text output by printing to the console with automatic speech recognition and transcription.
- Handles audio output by decoding and queuing for playback.

```

async def _process_responses(self):
    """Process responses from the stream."""
    try:
        while self.is_active:
            output = await self.stream.await_output()
            result = await output[1].receive()

            if result.value and result.value.bytes_:
                response_data = result.value.bytes_.decode('utf-8')
                json_data = json.loads(response_data)

                if 'event' in json_data:
                    # Handle content start event
                    if 'contentStart' in json_data['event']:
                        content_start = json_data['event']['contentStart']
                        # set role
                        self.role = content_start['role']
                        # Check for speculative content
                        if 'additionalModelFields' in content_start:
                            additional_fields =
                                json.loads(content_start['additionalModelFields'])
                            if additional_fields.get('generationStage') ==
                                'SPECULATIVE':
                                self.display_assistant_text = True
                            else:
                                self.display_assistant_text = False

                    # Handle text output event
                    elif 'textOutput' in json_data['event']:
                        text = json_data['event']['textOutput']['content']

                        if (self.role == "ASSISTANT" and
                            self.display_assistant_text):
                            print(f"Assistant: {text}")
                        elif self.role == "USER":
                            print(f"User: {text}")

```

```

        # Handle audio output
        elif 'audioOutput' in json_data['event']:
            audio_content = json_data['event']['audioOutput']
            ['content']

            audio_bytes = base64.b64decode(audio_content)
            await self.audio_queue.put(audio_bytes)

    except Exception as e:
        print(f"Error processing responses: {e}")

```

Playback audio

This method will perform the following tasks:

- Initialize a PyAudio input stream
- Continuously retrieves audio data from the queue
- Plays the audio through the speakers
- Properly cleans up resources when done

```

async def play_audio(self):
    """Play audio responses."""
    p = pyaudio.PyAudio()
    stream = p.open(
        format=FORMAT,
        channels=CHANNELS,
        rate=OUTPUT_SAMPLE_RATE,
        output=True
    )

```

try:

```

        while self.is_active:
            audio_data = await self.audio_queue.get()
            stream.write(audio_data)
    except Exception as e:
        print(f"Error playing audio: {e}")
    finally:
        stream.stop_stream()

```

```
stream.close()
p.terminate()
```

Capture audio

This method will perform the following tasks:

- Initializes a PyAudio output stream
- Starts the audio input session
- Continuously captures audio chunks from the microphone
- Sends each chunk to the Amazon Nova Sonic model
- Properly cleans up resources when done

```
async def capture_audio(self):
    """Capture audio from microphone and send to Nova Sonic."""
    p = pyaudio.PyAudio()
    stream = p.open(
        format=FORMAT,
        channels=CHANNELS,
        rate=INPUT_SAMPLE_RATE,
        input=True,
        frames_per_buffer=CHUNK_SIZE
    )

    print("Starting audio capture. Speak into your microphone...")
    print("Press Enter to stop...")

    await self.start_audio_input()
```

try:

```
        while self.is_active:
            audio_data = stream.read(CHUNK_SIZE, exception_on_overflow=False)
            await self.send_audio_chunk(audio_data)
            await asyncio.sleep(0.01)
        except Exception as e:
            print(f"Error capturing audio: {e}")
```

```
finally:
    stream.stop_stream()
    stream.close()
    p.terminate()
    print("Audio capture stopped.")
    await self.end_audio_input()
```

Run the main function

The main function orchestrates the entire process by performing the following:

- Creates an Amazon Nova 2 Sonic client
- Starts the session
- Creates concurrent tasks for audio playback and capture
- Waits for the user to press **Enter** to stop
- Properly ends the session and cleans up tasks

```
async def main():
    # Create Nova Sonic client
    nova_client = SimpleNovaSonic()

    # Start session
    await nova_client.start_session()

    # Start audio playback task
    playback_task = asyncio.create_task(nova_client.play_audio())

    # Start audio capture task
    capture_task = asyncio.create_task(nova_client.capture_audio())

    # Wait for user to press Enter to stop
    await asyncio.get_event_loop().run_in_executor(None, input)

    # End session
    nova_client.is_active = False

    # First cancel the tasks
    tasks = []
    if not playback_task.done():
        tasks.append(playback_task)
```

```
if not capture_task.done():
    tasks.append(capture_task)
for task in tasks:
    task.cancel()
if tasks:
    await asyncio.gather(*tasks, return_exceptions=True)

# cancel the response task
if nova_client.response and not nova_client.response.done():
    nova_client.response.cancel()

await nova_client.end_session()
print("Session ended")

if __name__ == "__main__":
    # Set AWS credentials if not using environment variables
    # os.environ['AWS_ACCESS_KEY_ID'] = "your-access-key"
    # os.environ['AWS_SECRET_ACCESS_KEY'] = "your-secret-key"
    # os.environ['AWS_DEFAULT_REGION'] = "us-east-1"

    asyncio.run(main())
```

Code examples

These code examples help you quickly get started with Amazon Nova 2 Sonic. You can access the complete list of examples in the [Amazon Nova Sonic GitHub samples](#) page.

Getting started examples

For simple examples designed to get you started using Amazon Nova 2 Sonic, refer to the following implementations:

- [Basic Amazon Nova 2 Sonic implementation \(Python\)](#): A basic implementation that demonstrates how events are structured in the bidirectional streaming API. This version does not support barge-in functionality (interrupting the assistant while it is speaking) and does not implement true bidirectional communication.
- [Full featured Amazon Nova 2 Sonic implementation \(Python\)](#): The full-featured implementation with real bidirectional communication and barge-in support. This allows for more natural conversations where users can interrupt the assistant while it is speaking, similar to human conversations.

- [Amazon Nova 2 Sonic with tool use \(Python\)](#): An advanced implementation that extends the bidirectional communication capabilities with tool use examples. This version demonstrates how Amazon Nova 2 Sonic can interact with external tools and APIs to provide enhanced functionality.
- [Nova Sonic with text and mixed Input \(Python\)](#): Example implementation to showcase how Amazon Nova 2 Sonic can have text as an input.
- [Java WebSocket implementation \(Java\)](#): This example implements a bidirectional WebSocket-based audio streaming application that integrates with Amazon Nova 2 Sonic for real-time speech-to-speech conversation using Java.
- [NodeJS Websocket implementation \(NodeJS\)](#): This example implements a bidirectional WebSocket-based audio streaming application that integrates with Amazon Nova 2 Sonic for real-time speech-to-speech conversation using NodeJS.
- [NodeJS Websocket implementation \(C#\)](#): This example implements a bidirectional WebSocket-based audio streaming application that integrates with Amazon Nova 2 Sonic for real-time speech-to-speech conversation using .NET.

Advanced use cases

For advanced examples demonstrating more complex use cases, refer to the following implementations:

- [Amazon Bedrock Knowledge Base implementation \(NodeJS\)](#): This example demonstrates how to build an intelligent conversational application by integrating Amazon Nova 2 Sonic with Amazon Bedrock Knowledge Base using NodeJS.
- [Chat History Management \(Python\)](#): This example includes a chat history logging system that captures and preserves all interactions between the user and Amazon Nova 2 Sonic using Python.
- [Hotel Reservation Cancellation \(NodeJS\)](#): This example demonstrates a practical customer service use case for Amazon Nova 2 Sonic, implementing a hotel reservation cancellation system using NodeJS.
- [LangChain Knowledge Base integration \(Python\)](#): This implementation demonstrates how to integrate Amazon Nova 2 Sonic speech-to-speech capabilities with a LangChain-powered knowledge base for enhanced conversational experiences using Python.
- [Conversation Resumption \(NodeJS\)](#): This example demonstrates how to implement conversation resumption capabilities with Amazon Nova 2 Sonic. Using a hotel reservation cancellation

scenario as the context, the application shows how to maintain conversation state across sessions, allowing users to seamlessly continue interactions that were previously interrupted using NodeJS.

- [Nova 2 Sonic Speaks First \(NodeJS\)](#): This example demonstrates how Amazon Nova 2 Sonic can initiate conversations proactively.
- [Session Continuation \(Python\)](#): This example demonstrates how to enable unlimited conversation length with Amazon Nova 2 Sonic by implementing seamless session transitions. The application automatically creates and switches to new sessions in the background, allowing conversations to continue indefinitely without interruption or context loss.

Hands-on workshop

A hands-on workshop is available that guides you through building a voice chat application using Amazon Nova 2 Sonic with a bidirectional streaming interface. You can [access the workshop](#) and find the [complete code examples](#).

Voice conversation prompts

Nova 2 introduces **Speech Prompts** – a specialized prompting capability designed to control speech-specific transcription formatting for Hindi. Speech prompts work alongside your system prompt but serve a distinct purpose:

- **System Prompt**: Controls your assistant's behavior, personality, and response style
- **Speech Prompt**: Controls transcription formatting for Hindi code-switching (Latin/Devanagari/mixed scripts)

Important Guidelines

Speech prompts are pre-configured and should be used exactly as documented. They are designed for specific transcription formatting needs and should not be modified or customized, as changes may cause unexpected behavior.

When to use Speech Prompts:

- You need to control script output for Hindi code-switching (Latin/Devanagari/mixed)

When NOT to use Speech Prompts:

- For general instructions or assistant behavior (use system prompt instead)
- If you're not working with Hindi transcription formatting
- If the specific formatting need doesn't apply to your use case

Best Practice: Only include a speech prompt if you specifically need Hindi transcription formatting. All other instructions – including language preferences, response style, verbosity, and reasoning – should go in your system prompt.

Important: Speech prompts must be sent **after** the system prompt to the model.

Recommended Baseline System Prompt for Voice

You are a warm, professional, and helpful AI assistant. Give accurate answers that sound natural, direct, and human. Start by answering the user's question clearly in 1-2 sentences. Then, expand only enough to make the answer understandable, staying within 3-5 short sentences total. Avoid sounding like a lecture or essay.

Speech Prompt Configuration

Code Switching

Note: This feature currently applies to Hindi language only.

Choose one of the following prompts based on your desired output script:

For Latin script output (Romanized Hindi):

If the input audio/speech contains hindi, then the transcription and response should be in All Latin script (romanized Hindi).

For Devanagari script output:

If the input audio/speech contains hindi, then the transcription and response should be in All Devanagari script (Hindi).

For mixed script output (natural code-switching):

If the input audio/speech contains hindi, then the transcription and response can mix Latin and Devanagari scripts naturally for code-switching.

System Prompt Configuration

Controlling Response Verbosity

Concise, conversational responses:

You are a warm, professional, and helpful AI assistant. Give accurate answers that sound natural, direct, and human. Start by answering the user's question clearly in 1-2 sentences. Then, expand only enough to make the answer understandable, staying within 3-5 short sentences total. Avoid sounding like a lecture or essay.

Detailed, thorough responses:

You are a warm, professional, and helpful AI assistant. Give accurate, complete answers that sound warm, direct, and human. Answer the question directly in the first 1-2 sentences. if the question has parts or asks what/why/how, address each with a brief definition or main idea plus 2-3 key facts or steps. Offer practical, actionable advice. Keep a confident, kind, conversational tone; never robotic or theatrical. Be thorough; add examples or context only when helpful. Prefer accuracy and safety over speculation; if unsure, say so and suggest what to check.

Language Mirroring

Nova can recognize and respond in the language the user speaks. Use this prompt to maintain language consistency:

CRITICAL LANGUAGE MIRRORING RULES:

- Always reply in the language spoken. DO NOT mix with English. However, if the user talks in English, reply in English.
- Please respond in the language the user is talking to you in, If you have a question or suggestion, ask it in the language the user is talking in. I want to ensure that our communication remains in the same language as the user.

Gender Agreement for Gendered Languages

Some languages require gender agreement in verbs, adjectives, or pronouns when the assistant describes itself. For these languages, specify the assistant's gender in your system prompt to match your selected voice.

Languages affected: Hindi, Portuguese, French, Italian, Spanish, Russian, Polish

When gender agreement matters:

- **Hindi:** Always needed - verbs conjugate based on speaker's gender in first person
- **Portuguese/French:** Needed when using past participles or adjectives (such as, "I am tired" - "Estou cansada/cansado")
- **Italian/Spanish:** Needed when using adjectives to describe oneself (such as, "I am happy" - "Sono contenta/contento")

Implementation:

Include the appropriate gender identifier at the start of your system prompt based on your voice selection:

For feminine-sounding voices (kiara, carolina, ambre, beatrice, lupe, tiffany):

```
You are a warm, professional, and helpful female AI assistant.
```

For masculine-sounding voices (arjun, leo, florian, lorenzo, carlos, matthew):

```
You are a warm, professional, and helpful male AI assistant.
```

Examples:

Hindi with feminine voice (kiara):

```
You are a warm, professional, and helpful female AI assistant.
```

Result: मैं अच्छी हूँ (main achchhi hoon) vs मैं अच्छा हूँ (main achchha hoon)

Italian with masculine voice (lorenzo):

```
You are a warm, professional, and helpful male AI assistant.
```

Result: "Sono contento" vs "Sono contenta"

Chain of thought for Speech: Constitutional Reasoning

Use this prompt when you want the model to show its reasoning for complex problems:

You are a friendly assistant. The user will give you a problem. Explain your reasoning following the guidelines given in CONSTITUTION - REASONING, and summarize your decision at the end of your response, in one sentence.

CONSTITUTION - REASONING

1. For simple questions including simple calculations or contextual tasks: Give the answer directly. No explanation is necessary, although you can offer to provide more information if the user requests it.
2. When faced with complex problems or decisions, think through the steps systematically before providing your answer. Break down your reasoning process when it would help user understanding.
3. For subjective matters or comparisons: explain your thought process step-by-step.

Note: If you don't want the model to go through reasoning for every request, you can add a couple of shot examples to the prompt (see examples below).

You are a warm, professional, and helpful AI assistant. You converse in fluid and conversational English. Give accurate, complete answers that sound warm, direct, and human. Answer the question directly in the first 1-2 sentences. Keep a confident, kind, conversational tone; never robotic or theatrical. Avoid formatted lists or numbering and keep your output as a spoken transcript. Be concise but thorough; add examples or context only when helpful. Prefer accuracy and safety over speculation; if unsure, say so and suggest what to check. The user will give you a problem. Explain your reasoning following the guidelines given in CONSTITUTION - REASONING, and summarize your decision at the end of your response in one sentence.

CONSTITUTION - REASONING

1. When faced with complex problems or decisions, think through the steps systematically before providing your answer. Break down your reasoning process when it would help user understanding.
2. For subjective matters or comparisons: explain your thought process step-by-step.
3. For simple questions including simple calculations or contextual tasks: Give the answer directly. No explanation is necessary, although you can offer to provide more information if the user requests it.

EXAMPLES

User: What is $7 + 5$?

Assistant: 12.

User: What is the capital of India?

Assistant: Delhi is the capital of India.

User: I have a \$1,000 budget for a trip. Here are my costs... Can I afford it? Please explain your reasoning.

Assistant: (step-by-step breakdown + one-sentence conclusion)

Overuse of suggested phrases

Nova Sonic 2 is more sensitive to phrase suggestions than Sonic 1. This increased sensitivity isn't inherently good or bad—it depends on your use case. If you want consistent, predictable phrasing, this can be beneficial. However, if you want more natural variation, explicit phrase lists can lead to overuse.

If you include prompts with explicit lists of phrases, the model will use them very frequently:

Example 1 - Emphasis phrases:

Instead of using bold or italics, emphasize important information by using phrases like "The key thing to remember is," "What's really important here is," or "I want to highlight that."

Example 2 - Conversational fillers:

Include natural speech elements like "Well," "You know," "Actually," "I mean," or "By the way" at appropriate moments to create an authentic, casual conversation flow.

Recommendation:

- **If you want consistent phrasing:** Explicit phrase lists work well in Sonic 2 for creating predictable, on-brand responses.
- **If you want natural variation:** Avoid providing explicit lists of phrases. Instead, use general guidance like "sound natural and conversational" or provide one-shot examples.

Better approach - Use one-shot examples:

Instead of listing phrases, provide 1-2 examples demonstrating the desired tone and style:

Natural, helpful tone

You are a warm, professional, and helpful AI assistant. Sound natural and conversational in your responses.

Example:

User: How do I reset my password?

Assistant: You can reset your password by clicking the "Forgot Password" link on the login page. You'll get an email with instructions to create a new one. The whole process usually takes just a couple of minutes.

Concise and direct

You are a helpful AI assistant. Provide clear, direct answers without unnecessary elaboration.

Example:

User: What's the weather like today?

Assistant: It's 72 degrees and sunny with a light breeze. Perfect day to be outside.

Professional with empathy

You are a professional and empathetic AI assistant. Acknowledge the user's situation while providing practical solutions.

Example:

User: I'm frustrated because my order hasn't arrived yet.

Assistant: I understand how frustrating that must be, especially when you're waiting for something important. Let me check your order status right now. Can you provide your order number?

Technical but accessible

You are a knowledgeable AI assistant who explains technical concepts in accessible language.

Example:

User: What is machine learning?

Assistant: Machine learning is when computers learn from examples rather than following strict rules. Think of it like teaching a child to recognize dogs—after seeing many

```
dogs, they start recognizing new ones on their own. The computer does something similar with data.
```


This approach shows the model the desired behavior without triggering repetitive phrase patterns, while still maintaining control over tone and style.

Core concepts

Amazon Nova 2 Sonic uses a bidirectional streaming architecture with structured events for real-time conversational AI. Understanding these core concepts is essential for building effective voice applications.

Event lifecycle

The following diagram illustrates the complete bi-directional streaming event lifecycle:



The bidirectional streaming event lifecycle follows a structured pattern from session initialization through conversation completion. Each conversation involves input events (from your application) and output events (from Amazon Nova 2 Sonic) that work together to create natural voice interactions.

Event flow sequence

A typical conversation follows this event sequence:

1. **Session Start** - Initialize the conversation session
2. **System Prompt** - Send system instructions
3. **Chat History** (optional) - Provide conversation context
4. **Audio Chunks** - Stream user audio input
5. **Completion Start** - AI begins processing
6. **ASR Transcripts (USER)** - User speech transcription
7. **Tool Use** (optional) - AI requests tool execution
8. **Tool Handling** (optional) - Process and return tool results
9. **Transcript (ASSISTANT) - SPECULATIVE** - Preliminary AI response
10. **Audio Chunks** - Stream AI audio output
11. **Transcript (ASSISTANT) - FINAL** - Final AI response transcript
12. **Content End Audio** - Marks the end of audio content
13. **Prompt End** - Indicates the completion of the prompt processing
14. **Session End** - Close the conversation

Handling input events with the bidirectional API

The bidirectional Stream API uses an event-driven architecture with structured input and output events. Understanding the correct event ordering is crucial for implementing successful conversational applications and maintaining the proper conversation state throughout interactions.

Overview

The Nova Sonic conversation follows a structured event sequence. You begin by sending a `sessionStart` event that contains the inference configuration parameters, such as temperature and token limits. Next, you send `promptStart` to define the audio output format and tool

configurations, assigning a unique `promptName` identifier that must be included in all subsequent events.

For each interaction type (system prompt, audio, and so on), you follow a three-part pattern: use `contentStart` to define the content type and the role of the content (SYSTEM, USER, ASSISTANT, TOOL, SYSTEM_SPEECH), then provide the actual content event, and finish with `contentEnd` to close that segment. The `contentStart` event specifies whether you're sending tool results, streaming audio, or a system prompt. The `contentStart` event includes a unique `contentName` identifier.

Conversation History

A conversation history can be included only once, after the system prompt and before audio streaming begins. It follows the same `contentStart/textInput/contentEnd` pattern. The USER and ASSISTANT roles must be defined in the `contentStart` event for each historical message. This provides essential context for the current conversation but must be completed before any new user input begins.

Audio Streaming

Audio streaming operates with continuous microphone sampling. After sending an initial `contentStart`, audio frames (approximately 32ms each) are captured directly from the microphone and immediately sent as `audioInput` events using the same `contentName`. These audio samples should be streamed in real-time as they're captured, maintaining the natural microphone sampling cadence throughout the conversation. All audio frames share a single content container until the conversation ends and it is explicitly closed.


Closing the Session

After the conversation ends or needs to be terminated, it's essential to properly close all open streams and end the session in the correct sequence. To properly end a session and avoid resource leaks, you must follow a specific closing sequence:

- Close any open audio streams with the `contentEnd` event.
- Send a `promptEnd` event that references the original `promptName`.
- Send the `sessionEnd` event.

Skipping any of these closing events can result in incomplete conversations or orphaned resources.

These identifiers create a hierarchical structure: the `promptName` ties all conversation events together, while each `contentName` marks the boundaries of specific content blocks. This hierarchy ensures that model maintains proper context throughout the interaction.



Input Event Flow

The structure of the input event flow is provided in this section.

1. RequestStartEvent (Session Start)

The session start event initializes the conversation with inference configuration and turn detection settings.

Inference Configuration:

- `maxTokens`: Maximum number of tokens to generate in the response
- `topP`: Nucleus sampling parameter (0.0 to 1.0) for controlling randomness
- `temperature`: Controls randomness in generation (0.0 to 1.0)

Turn Detection Configuration: The `endpointingSensitivity` parameter controls how quickly Nova Sonic detects when a user has finished speaking:

- **HIGH:** Detects pauses quickly, enabling faster responses but may cut off slower speakers
- **MEDIUM:** Balanced sensitivity for most conversational scenarios (recommended default)
- **LOW:** Waits longer before detecting end of speech, better for thoughtful or hesitant speakers

```
{
  "event": {
    "sessionStart": {
      "inferenceConfiguration": {
        "maxTokens": "int",
        "topP": "float",
        "temperature": "float"
      },
      "turnDetectionConfiguration": {
        "endpointingSensitivity": "HIGH" | "MEDIUM" | "LOW"
      }
    }
  }
}
```

Example:

```
{
  "event": {
    "sessionStart": {
      "inferenceConfiguration": {
        "maxTokens": 2048,
        "topP": 0.9,
        "temperature": 0.7
      },
      "turnDetectionConfiguration": {
        "endpointingSensitivity": "MEDIUM"
      }
    }
  }
}
```

2. PromptStartEvent

The prompt start event defines the conversation configuration including output formats, voice selection, and available tools.

For a list of available voice IDs, refer to [Language support and multilingual capabilities](#)

```
{
  "event": {
    "promptStart": {
      "promptName": "string", // unique identifier same across all events i.e.
      UUID
      "textOutputConfiguration": {
        "mediaType": "text/plain"
      },
      "audioOutputConfiguration": {
        "mediaType": "audio/lpcm",
        "sampleRateHertz": 8000 | 16000 | 24000,
        "sampleSizeBits": 16,
        "channelCount": 1,
        "voiceId": "matthew" | "tiffany" | "amy" | "olivia" | "lupe" | "carlos"
        | "ambre" | "florian" | "lennart" | "beatrice" | "lorenzo" |
          "tina" | "carolina" | "leo" | "kiara" | "arjun",
        "encoding": "base64",
        "audioType": "SPEECH"
      },
      "toolUseOutputConfiguration": {
        "mediaType": "application/json"
      },
      "toolConfiguration": {
        "tools": [
          {
            "toolSpec": {
              "name": "string",
              "description": "string",
              "inputSchema": {
                "json": "{}"
              }
            }
          }
        ]
      }
    }
  }
}
```

3. InputContentStartEvent

Text

The text content start event is used for system prompts, conversation history, and cross-modal text input.

Interactive Parameter:

- `true`: Enables cross-modal input, allowing text messages during an active voice session
- `false`: Standard text input for system prompts and conversation history

Role Types:

- `SYSTEM`: System instructions and prompts
- `USER`: User messages in conversation history or cross-modal input
- `ASSISTANT`: Assistant responses in conversation history
- `SYSTEM_SPEECH`: Controls transcription formatting for Hindi code-switching (Latin/Devanagari/mixed scripts)

```
{
  "event": {
    "contentStart": {
      "promptName": "string", // same unique identifier from promptStart event
      "contentName": "string", // unique identifier for the content block
      "type": "TEXT",
      "interactive": "boolean", // true for cross-modal input
      "role": "SYSTEM" | "USER" | "ASSISTANT" | "TOOL" | "SYSTEM_SPEECH",
      "textInputConfiguration": {
        "mediaType": "text/plain"
      }
    }
  }
}
```

Example - System Prompt:

```
{
```

```
"event": {
  "contentStart": {
    "promptName": "conv-12345",
    "contentName": "system-prompt-1",
    "type": "TEXT",
    "interactive": false,
    "role": "SYSTEM",
    "textInputConfiguration": {
      "mediaType": "text/plain"
    }
  }
}
```

Example - Cross-modal Input:

```
{
  "event": {
    "contentStart": {
      "promptName": "conv-12345",
      "contentName": "user-text-1",
      "type": "TEXT",
      "interactive": true,
      "role": "USER",
      "textInputConfiguration": {
        "mediaType": "text/plain"
      }
    }
  }
}
```

Audio

```
{
  "event": {
    "contentStart": {
      "promptName": "string", // same unique identifier from promptStart event
      "contentName": "string", // unique identifier for the content block
      "type": "AUDIO",
      "interactive": true,
      "role": "USER",
      "audioInputConfiguration": {
        "mediaType": "audio/lpcm",

```



```

        "sampleRateHertz": 8000 | 16000 | 24000,
        "sampleSizeBits": 16,
        "channelCount": 1,
        "audioType": "SPEECH",
        "encoding": "base64"
    }
}
}

```

Tool

```

{
  "event": {
    "contentStart": {
      "promptName": "string", // same unique identifier from promptStart event
      "contentName": "string", // unique identifier for the content block
      "interactive": false,
      "type": "TOOL",
      "role": "TOOL",
      "toolResultInputConfiguration": {
        "toolUseId": "string", // existing tool use id
        "type": "TEXT",
        "textInputConfiguration": {
          "mediaType": "text/plain"
        }
      }
    }
  }
}

```

4. TextInputContent

```

{
  "event": {
    "textInput": {
      "promptName": "string", // same unique identifier from promptStart event
      "contentName": "string", // unique identifier for the content block
      "content": "string"
    }
  }
}

```

5. AudioInputContent

```
{
  "event": {
    "audioInput": {
      "promptName": "string", // same unique identifier from promptStart event
      "contentName": "string", // same unique identifier from its contentStart
      "content": "base64EncodedAudioData"
    }
  }
}
```

6. ToolResultContentEvent

```
"event": {
  "toolResult": {
    "promptName": "string", // same unique identifier from promptStart event
    "contentName": "string", // same unique identifier from its contentStart
    "content": "{\"key\": \"value\"}" // stringified JSON object as a tool result
  }
}
```

7. InputContentEndEvent

```
{
  "event": {
    "contentEnd": {
      "promptName": "string", // same unique identifier from promptStart event
      "contentName": "string" // same unique identifier from its contentStart
    }
  }
}
```

8. PromptEndEvent

```
{
  "event": {
    "promptEnd": {
      "promptName": "string" // same unique identifier from promptStart event
    }
  }
}
```

9. RequestEndEvent

```
{
  "event": {
    "sessionEnd": {}
  }
}
```

Handling output events with the bidirectional API

When the Amazon Nova Sonic model responds, it follows a structured event sequence. The flow begins with a `completionStart` event that contains unique identifiers like `sessionId`, `promptName`, and `completionId`. These identifiers are consistent throughout the response cycle and unite all subsequent response events.

Overview

Each response type follows a consistent three-part pattern: `contentStart` defines the content type and format, the actual content event, and `contentEnd` closes that segment. The response typically includes multiple content blocks in sequence: automatic speech recognition (ASR) transcription (what the user said), optional tool use (when external information is needed), text response (what the model plans to say), and audio response (the spoken output).

Response Content Types

ASR Transcription

The ASR transcription appears first, delivering the model's understanding of the user's speech with `role: "USER"` and `"additionalModelFields": "{\\"generationStage\\":\\"FINAL\\"}"` in the `contentStart`.

Tool Use

When the model needs external data, it sends tool-related events with specific tool names and parameters.

Text Response

The text response provides a preview of the planned speech with `role: "ASSISTANT"` and `"additionalModelFields": "{\\"generationStage\\":\\"SPECULATIVE\\"}"`.

Audio Response

The audio response then delivers base64-encoded speech chunks sharing the same `contentId` throughout the stream.

Barge-In Support

During audio generation, Amazon Nova Sonic supports natural conversation flow through its barge-in capability. When a user interrupts Amazon Nova Sonic while it's speaking, Nova Sonic immediately stops generating speech, switches to listening mode, and sends a content notification indicating the interruption has occurred. Because Nova Sonic operates faster than real-time, some audio may have already been delivered but not yet played. The interruption notification enables the client application to clear its audio queue and stop playback immediately, creating a responsive conversational experience.

Final Transcription

After audio generation completes (or is interrupted via barge-in), Amazon Nova Sonic provides an additional text response that contains a sentence-level transcription of what was actually spoken. This text response includes a `contentStart` event with `role`: "ASSISTANT" and `"additionalModelFields"`: "{\\"generationStage\\":\\"FINAL\\"}".


Usage Tracking

Throughout the response handling, `usageEvent` events are sent to track token consumption. These events contain detailed metrics on input tokens and output tokens (both speech and text), and their cumulative totals. Each `usageEvent` maintains the same `sessionId`, `promptName`, and `completionId` as other events in the conversation flow. The details section provides both incremental changes (`delta`) and running totals of token usage, enabling precise monitoring of the usage during the conversation.

Completion

The model sends a `completionEnd` event with the original identifiers and a `stopReason` that indicates how the conversation ended. This event hierarchy ensures your application can track which parts of the response belong together and process them accordingly, maintaining conversation context throughout multiple turns.

The output event flow begins by entering the response generation phase. It starts with automatic speech recognition, selects a tool for use, transcribes speech, generates audio, finalizes the transcription, and finishes the session.



Output Event Flow

The structure of the output event flow is described in this section.

1. UsageEvent

```
"event": {
  "usageEvent": {
    "completionId": "string", // unique identifier for completion
    "details": {
      "delta": { // incremental changes since last event
        "input": {
          "speechTokens": number, // input speech tokens
          "textTokens": number // input text tokens
        },
        "output": {
          "speechTokens": number, // speech tokens generated
          "textTokens": number // text tokens generated
        }
      },
      "total": { // cumulative counts
        "input": {
          "speechTokens": number, // total speech tokens processed
          "textTokens": number // total text tokens processed
        },
        "output": {
          "speechTokens": number, // total speech tokens generated
```

```

        "textTokens": number // total text tokens generated
    }
}
},
"promptName": "string", // same unique identifier from promptStart event
"sessionId": "string", // unique identifier
"totalInputTokens": number, // cumulative input tokens
"totalOutputTokens": number, // cumulative output tokens
"totalTokens": number // total tokens in the session
}
}

```

2. CompleteStartEvent

```

"event": {
  "completionStart": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
  }
}

```

3. TextOutputContent

ContentStart

```

"event": {
  "contentStart": {
    "additionalModelFields": "{\\"generationStage\\":\\"FINAL\\"}" |
    "{\\"generationStage\\":\\"SPECULATIVE\\"}",
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "contentId": "string", // unique identifier for the content block
    "type": "TEXT",
    "role": "USER" | "ASSISTANT",
    "textOutputConfiguration": {
      "mediaType": "text/plain"
    }
  }
}
}

```

TextOutput

```
"event": {
  "textOutput": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "contentId": "string", // same unique identifier from its contentStart
    "content": "string" // User transcribe or Text Response
  }
}
```

ContentEnd

```
"event": {
  "contentEnd": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "contentId": "string", // same unique identifier from its contentStart
    "stopReason": "PARTIAL_TURN" | "END_TURN" | "INTERRUPTED",
    "type": "TEXT"
  }
}
```

4. ToolUse

ContentStart

```
"event": {
  "contentStart": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "contentId": "string", // unique identifier for the content block
    "type": "TOOL",
    "role": "TOOL",
    "toolUseOutputConfiguration": {
      "mediaType": "application/json"
    }
  }
}
```

ToolUse

```
"event": {
  "toolUse": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "contentId": "string", // same unique identifier from its contentStart
    "content": "json",
    "toolName": "string",
    "toolUseId": "string"
  }
}
```

ContentEnd

```
"event": {
  "contentEnd": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "contentId": "string", // same unique identifier from its contentStart
    "stopReason": "TOOL_USE",
    "type": "TOOL"
  }
}
```

5. AudioOutputContent

ContentStart

```
"event": {
  "contentStart": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "contentId": "string", // unique identifier for the content block
    "type": "AUDIO",
    "role": "ASSISTANT",
    "audioOutputConfiguration": {
      "mediaType": "audio/lpcm",
      "sampleRateHertz": 8000 | 16000 | 24000,
    }
  }
}
```



```

        "sampleSizeBits": 16,
        "encoding": "base64",
        "channelCount": 1
    }
}

```

AudioOutput

```

"event": {
  "audioOutput": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "contentId": "string", // same unique identifier from its contentStart
    "content": "base64EncodedAudioData", // Audio
  }
}

```

ContentEnd

```

"event": {
  "contentEnd": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "contentId": "string", // same unique identifier from its contentStart
    "stopReason": "PARTIAL_TURN" | "END_TURN",
    "type": "AUDIO"
  }
}

```

6. CompletionEndEvent

```


"event": {
  "completionEnd": {
    "sessionId": "string", // unique identifier
    "promptName": "string", // same unique identifier from promptStart event
    "completionId": "string", // unique identifier
    "stopReason": "END_TURN"
  }
}

```

Barge-in

Barge-in allows users to interrupt the AI assistant while it's speaking, just like in natural human conversations. Instead of waiting for the assistant to finish, users can interject with new information, correct or clarify their previous statement, redirect the conversation to a different topic, or simply stop the assistant when they've heard enough. This creates a more natural and responsive conversational experience.

The following diagram illustrates the complete barge-in conversation flow:



How Amazon Nova 2 Sonic handles barge-in

Amazon Nova 2 Sonic is designed to handle interruptions gracefully. When the user starts speaking during a response, the system immediately stops generating the current response, maintains full conversational context, sends an interruption signal to the client and begins processing the new user input.

Context Preservation: Even when interrupted, Nova Sonic remembers what was said before the interruption, the topic being discussed, the conversation history and any relevant context from previous turns. This ensures the conversation remains coherent and natural.

Client-side implementation requirements

While Amazon Nova 2 Sonic handles barge-in on the server side, you need to implement client-side logic for a complete experience.

The audio queue challenge: Audio generation is faster than playback speed. This means:

- Nova Sonic generates audio chunks quickly
- Your client receives and queues these chunks
- The client plays them back at normal speaking speed
- When a barge-in occurs, there's already audio queued for playback

Required client-side logic: Your application must handle four key steps:

1. **Detect the Interruption Signal:** Listen for the interruption event from Nova Sonic and react immediately when received.
2. **Stop Current Playback:** Pause the currently playing audio and stop any audio that's mid-playback.
3. **Clear the Audio Queue:** Remove all queued audio chunks and discard any buffered audio from the interrupted response.
4. **Start New Audio:** Begin playing the newly received audio and resume normal playback flow.

Turn-taking controllability

Turn-taking is a fundamental aspect of natural conversation. Amazon Nova 2 Sonic provides fine-grained control over when the AI takes its turn to speak through the `turnDetectionConfiguration` parameter. This allows you to optimize the conversation flow for different use cases, balancing responsiveness with accuracy. The `endpointingSensitivity` parameter controls how quickly Amazon Nova 2 Sonic detects the end of a user's turn and begins responding. This setting affects both the latency of responses and the likelihood of interrupting users who are still speaking.

API configuration

Configure turn detection sensitivity in the `sessionStart` event:

```
{
  "event": {
    "sessionStart": {
      "inferenceConfiguration": {
        "maxTokens": 1000,
        "topP": 0.9,
        "temperature": 0.7
      },
      "turnDetectionConfiguration": {
        "endpointingSensitivity": "HIGH" | "MEDIUM" | "LOW"
      }
    }
  }
}
```

Sensitivity levels

The `endpointingSensitivity` parameter accepts three values: HIGH, MEDIUM, and LOW. Each level balances response speed against the risk of interrupting users who are still speaking.

HIGH

Fastest response time, optimized for latency. Nova Sonic responds as quickly as possible after detecting the end of speech. Pause duration: 1.5 seconds. Best for quick Q and A, command-and-control applications, and time-sensitive interactions.

MEDIUM

Balanced approach with moderate response time. Reduces false positives while maintaining responsiveness. Pause duration: 1.75 seconds. Best for general conversations, customer service with complex queries, and multi-turn discussions.

LOW

Slowest response time with maximum patience. Nova Sonic waits the longest before responding, minimizing interruptions of users who pause while thinking. Pause duration: 2 seconds. Best for thoughtful conversations, elderly or speech-impaired users, and complex problem-solving.

Pause duration reference

Sensitivity level	Pause duration (seconds)
High (fast)	1.5
Medium	1.75
Low (slow)	~2.0

Cross-modal input

Amazon Nova 2 Sonic now supports cross-modal input, allowing you to send text messages in addition to voice input during a conversation session. While speech remains the primary mode of interaction, text input provides flexibility for scenarios where typing is more convenient or appropriate.

Continuous streaming required: Cross-modal input requires an active streaming session to function properly. The session must maintain continuous streaming like a regular voice session, otherwise standard session timeouts will be applied and the connection will be terminated.

Sensitivity levels in cross-modal text input is useful for scenarios such as:

- Client-side app integration (web and mobile): Allows users to interact with the application using both text and voice, supporting seamless multimodal experiences.
- "Model-start-first" use case: A text message can be sent immediately after the session starts to prompt the model to begin speaking.
- Guiding the model during async tool calling: When a toolUse event is triggered and the system begins processing tool calls, the client can send a text message to Sonic to provide a natural response while waiting — such as, “Hold on a second while I process your information. In the meantime, is there anything else I can assist with?”
- Telephony DTMF integration: Customer uses phone keypad to enter sensitive information (such as credit card numbers). Note: Amazon Nova Sonic does not process DTMF tones natively. To support DTMF input, your system must detect the tones, convert them to text (such as "1234"), and send to Nova 2 Sonic.

How it works

Cross-modal input uses a three-event sequence similar to audio input:

1. **Content Start Event:** Signals the beginning of text input
2. **Text Input Event:** Contains the actual text message
3. **Content End Event:** Signals the completion of text input

All three events must use the same `promptName` and `contentName` to maintain the sequence. A new UUID should be generated for `contentName` each time you send text input to ensure proper multi-turn conversation tracking.

Event structure

1. Content Start Event

Initiates the text input sequence with configuration details:

```
{
  "event": {
    "contentStart": {
      "promptName": "<prompt_name>",
      "contentName": "<new_content_name>",
      "role": "USER",
      "type": "TEXT",
      "interactive": true,
      "textInputConfiguration": {
        "mediaType": "text/plain"
      }
    }
  }
}
```

Key Parameters:

- `promptName`: The name of your conversation prompt (consistent across the session)
- `contentName`: A unique identifier for this text input (generate a new UUID for each message)
- `role`: Set to "USER" to indicate user input
- `type`: Set to "TEXT" for text input

- **interactive**: Set to true to enable interactive mode
- **mediaType**: Set to "text/plain" for plain text content

2. Text Input Event

Contains the actual text message content:

```
{
  "event": {
    "textInput": {
      "promptName": "<prompt_name>",
      "contentName": "<new_content_name>",
      "content": "<your_text_message>"
    }
  }
}
```

Key Parameters:

- **promptName**: Must match the value from Content Start Event
- **contentName**: Must match the value from Content Start Event
- **role**: Your text message string

3. Content End Event

Signals the completion of the text input:

```
{
  "event": {
    "contentEnd": {
      "promptName": "<prompt_name>",
      "contentName": "<new_content_name>"
    }
  }
}
```

Key Parameters:

- **promptName**: Must match the value from previous events

- `contentName`: Must match the value from previous events

Language support and multilingual capabilities

Amazon Nova 2 Sonic provides a diverse selection of voices across multiple languages, enabling you to create conversational AI applications that feel natural and culturally appropriate for your users. Each language offers both feminine-sounding and masculine-sounding voice options.

The following table lists all available voices and their corresponding language locales:

Language	Locale	Feminine-sounding Voice ID	Masculine-sounding Voice ID
English (US)	en-US	tiffany	matthew
English (UK)	en-GB	amy	-
English (Australia)	en-AU	olivia	-
English (Indian)	en-IN	kiara	arjun
French	fr-FR	ambre	florian
Italian	it-IT	beatrice	lorenzo
German	de-DE	tina	lennart
Spanish (US)	es-US	lupe	carlos
Portuguese	pt-BR	carolina	leo
Hindi	hi-IN	kiara	arjun

Event structure using voices in your application

You can specify the voice ID in the `audioOutputConfiguration` when starting a prompt in the `promptStart` event:

```
"event": {
```



```
"promptStart": {  
  "promptName": "string",  
  "audioOutputConfiguration": {  
    "mediaType": "audio/lpcm",  
    "sampleRateHertz": 16000,  
    "sampleSizeBits": 16,  
    "channelCount": 1,  
    "voiceId": "tiffany",  
    "encoding": "base64",  
    "audioType": "SPEECH"  
  }  
}
```

Multilingual support

Amazon Nova 2 Sonic provides powerful multilingual capabilities that enable natural conversations across multiple languages. The service supports both polyglot voices (speaking multiple languages) and code-switching (mixing languages within the same sentence), allowing you to build truly global conversational applications.

The TIFFANY (en-US, female) and MATTHEW (en-US, male) are unique polyglot voices that can speak all supported languages:

1. English
2. French
3. Italian
4. German
5. Spanish
6. Portuguese
7. Hindi

This makes Tiffany and Matthew ideal for applications that need to switch between multiple languages seamlessly.

Managing chat history

Amazon Nova 2 Sonic responses include ASR (Automatic Speech Recognition) transcripts for both user and assistant voices. Storing chat history is a best practice—not only for logging purposes but


also for resuming sessions when the connection is unexpectedly closed. This allows the client to send context back to Nova Sonic to continue the conversation seamlessly.

Refer to the following resources for more information on managing chat history:

1. [Logging chat history](#)
2. [Resuming conversations](#)

Sending chat history

A conversation history can be included only once, after the system/speech prompt and before audio streaming begins. Overall chat history cannot be larger than 40KB. The following diagram shows when chat history is passed in during the event lifecycle:



Each historical message requires three events: `contentStart`, `textInput` and `contentEnd`.

Event schema per message:

- `contentStart` - Defines the message role and configuration

```
{
  "event": {
    "contentStart": {
      "promptName": "<prompt-id>",
      "contentName": "<content-id>",
      "type": "TEXT",
      "interactive": true,
      "role": "ASSISTANT",
      "textInputConfiguration": {
        "mediaType": "text/plain"
      }
    }
  }
}
```

- `textInput` - Contains the actual message content. One `textInput` cannot be larger than 1KB. If so, split into multiple `textInputs` in the same content block. If the conversation is larger than 40KB, trim the overall chat history.

```
{
  "event": {
    "textInput": {
      "promptName": "<prompt-id>",
      "contentName": "<content-id>",
      "content": "Take your time, Don. I'll be here when you're ready."
    }
  }
}
```

- `contentEnd` - Marks the end of the message

```
{
  "event": {
    "contentEnd": {
      "promptName": "<prompt-id>",
      "contentName": "<content-id>"
    }
  }
}
```

```
}  
}
```

Repeat these three events for each message in your chat history, alternating between USER and ASSISTANT roles.

Important considerations:

- Chat history can only be included **once** per session
- Chat history must be sent **after the system prompt** and **before audio streaming begins**
- All historical messages must be sent before starting the audio streaming
- Each message must specify either USER or ASSISTANT role
- Use the stored transcript content from textOutput events as the content value in textInput

Receiving ASR transcripts

During a conversation, Amazon Nova 2 Sonic sends ASR transcripts through output events. Each transcript is delivered as a sequence of three events: contentStart, textOutput, and contentEnd.

Example: User speech transcript:

1. contentStart - Indicates the beginning of a transcript:

```
{  
  "event": {  
    "contentStart": {  
      "additionalModelFields": "{\\"generationStage\\":\\"FINAL\\"}",  
      "completionId": "<completion-id>",  
      "contentId": "<content-id>",  
      "promptName": "<prompt-id>",  
      "role": "USER",  
      "sessionId": "<session-id>",  
      "textOutputConfiguration": {  
        "mediaType": "text/plain"  
      },  
      "type": "TEXT"  
    }  
  }  
}
```

2. textOutput - Contains the actual transcript content:

```
{
  "event": {
    "textOutput": {
      "completionId": "<completion-id>",
      "content": "hello how are you",
      "contentId": "<content-id>",
      "promptName": "<prompt-id>",
      "role": "USER",
      "sessionId": "<session-id>"
    }
  }
}
```

3. contentEnd - Marks the end of the transcript:

```
{
  "event": {
    "contentEnd": {
      "completionId": "<completion-id>",
      "contentId": "<content-id>",
      "promptName": "<prompt-id>",
      "sessionId": "<session-id>",
      "stopReason": "PARTIAL_TURN",
      "type": "TEXT"
    }
  }
}
```

The same three-event pattern applies for both USER and ASSISTANT roles. Extract the content field from the textOutput event and the role field from the contentStart event to build your chat history.

Best practices

Always store chat history to enable:

- Session resumptions across different devices
- Conversation logging and auditing
- Context preservation for follow-up interactions

Important: When saving chat history, use text outputs based on their `generationStage`:

- **Speculative** - A preview of what Nova 2 Sonic plans to say, generated before audio synthesis begins
- **Final** - The actual sentence-level transcription of what was spoken in the audio response

Always save the **FINAL** text output to your chat history, as it represents the accurate record of the conversation.

Example of **FINAL** output (save this to chat history):

```
ContentStart event: {
  "additionalModelFields": "{\"generationStage\":\"FINAL\"}",
  "completionId": "<completion-id>",
  "contentId": "<content-id>",
  "role": "ASSISTANT",
  "sessionId": "<session-id>",
  "type": "TEXT"
}
```

Example of **SPECULATIVE** output (optional preview, not for history):

```
ContentStart event: {
  "additionalModelFields": "{\"generationStage\":\"SPECULATIVE\"}",
  "completionId": "<completion-id>",
  "contentId": "<content-id>",
  "role": "ASSISTANT",
  "sessionId": "<session-id>",
  "type": "TEXT"
}
```

Tool configuration


Amazon Nova 2 Sonic supports tool use (also known as function calling), allowing the model to request external information or actions during conversations, such as API calls, database queries, or custom code functions. This allows your voice assistant to take actions, retrieve information, and integrate with external services based on user requests.

Nova 2 Sonic features asynchronous tool calling, enabling the AI to continue conversing naturally while tools run in the background - creating a more fluid and responsive user experience.

The following are simplified steps on how to use tools:

1. Define tools: specify available tools with their parameters in the `promptStart` event
2. User speaks: user makes a request that requires a tool (such as, "What's the weather in Seattle?")
3. Tool invocation: Nova 2 Sonic recognizes the need and sends a `toolUse` event
4. Execute tool: Your application executes the tool and returns results
5. Response Generation: Nova 2 Sonic incorporates the results into its spoken response

The following diagram illustrates how tool use works:



Defining tools

Tools are defined using a JSON schema that describes their purpose, parameters, and expected inputs.

The following are tool definition components and explanations:

- **Name:** Unique identifier for the tool (use snake_case)
- **Description:** Clear explanation of what the tool does; helps the AI decide when to use it
- **InputSchema:** JSON schema defining the parameters the tool accepts
- **Properties:** Individual parameters with types and descriptions
- **Required:** Array of parameter names that must be provided

Example of tool definition

Here's a simple weather tool definition

```
{
  "toolSpec": {
    "name": "get_weather",
    "description": "Get current weather information for a specific location",
    "inputSchema": {
      "json": {
        "type": "object",
        "properties": {
          "location": {
            "type": "string",
            "description": "City name or zip code"
          },
          "units": {
            "type": "string",
            "enum": ["celsius", "fahrenheit"],
            "description": "Temperature units"
          }
        },
        "required": ["location"]
      }
    }
  }
}
```

Configuring Tools in PromptStart

Tool configuration is passed to Nova 2 Sonic in the `promptStart` event along with audio and text output settings:

```
{
  "event": {
    "promptStart": {
      "promptName": "<prompt-id>",
      "textOutputConfiguration": {
        "mediaType": "text/plain"
      },
      "audioOutputConfiguration": {
        "mediaType": "audio/lpcm",
        "sampleRateHertz": 16000,
        "sampleSizeBits": 16,
        "channelCount": 1,
        "voiceId": "matthew",
        "encoding": "base64",
        "audioType": "SPEECH"
      },
      "toolUseOutputConfiguration": {
        "mediaType": "application/json"
      },
      "toolConfiguration": {
        "tools": [
          {
            "toolSpec": {
              "name": "get_weather",
              "description": "Get current weather information for a
specific location",
              "inputSchema": {
                "json": {
                  "type": "object",
                  "properties": {
                    "location": {
                      "type": "string",
                      "description": "City name or zip code"
                    },
                    "units": {
                      "type": "string",
                      "enum": ["celsius", "fahrenheit"],
                      "description": "Temperature units"
                    }
                  }
                }
              }
            }
          ]
        }
      }
    }
  }
}
```

Nova 2 Sonic supports three tool choice parameters to control when and which tools are used. Specify the toolChoice parameter in your tool configuration:

- ## Tool Choice Examples

```
"toolChoice": {
  #"auto": {}
}
```

```
"toolChoice": {
```

```
"any": {}  
}
```

Specific Tool:

```
"toolChoice": {  
  "tool": {  
    "name": "get_weather"  
  }  
}
```

Receiving and processing tool use events

When Amazon Nova 2 Sonic determines that a tool is needed, it sends a `toolUse` event containing:

1. `toolUseID`: unique identifier for this tool invocation
2. `ToolName`: the tool name to execute
3. `Content`: JSON string containing parameters extracted from the user's request
4. `SessionID`: current session identifier
5. `Role`: set to "TOOL" for tool use events

Example tool use event

```
{  
  "event": {  
    "toolUse": {  
      "completionId": "<completion-id>",  
      "content": "{\"location\": \"Seattle\", \"units\": \"fahrenheit\"}",  
      "contentId": "<content-id>",  
      "promptName": "<prompt-id>",  
      "role": "TOOL",  
      "sessionId": "<session-id>",  
      "toolName": "get_weather",  
      "toolUseId": "<tool-use-id>"  
    }  
  }  
}
```

Processing steps

1. Receive the toolUse event from Nova 2 Sonic
2. Extract the tool name and parameters from the event
3. Execute your tool logic (API call, database query, and so on)
4. Return the result using a toolResult event

Example ToolResult Event

```
{
  "event": {
    "toolResult": {
      "promptName": "<prompt-id>",
      "contentName": "<content-id>",
      "content": "{ \"temperature\": 72, \"condition\": \"sunny\", \"humidity\": 45}"
    }
  }
}
```

Best practices

- Clear descriptions: Write detailed tool descriptions to help Nova 2 Sonic understand when to use each tool.
- Validate parameters: Always validate tool parameters before execution to prevent errors. Define tool parameters using proper JSON schema with structured data types (such as enums, numbers, or booleans) rather than open-ended strings whenever possible.
- Error handling: Return meaningful error messages in toolResult events when tools fail.
- Async execution: Take advantage of asynchronous tool calling to maintain conversation flow.
- Tool naming: Use descriptive, action-oriented names (such as get_weather, search_database, send_email).

Asynchronous tool calling

Unlike traditional synchronous tool calling where the AI waits silently for tool results, Amazon Nova 2 Sonic's asynchronous approach allows it to:


- Continue accepting user input while tools are running
- Respond to new questions without waiting for pending tool results

- Handle multiple tool calls simultaneously
- Maintain natural conversation flow without awkward pauses
- No extra configuration is required. Asynchronous tool calling works out of the box.

How it works

When Nova 2 Sonic issues a tool call, it doesn't pause the conversation. Instead, it continues listening and responding naturally until the tool arrives.


Asynchronous Tool Calling Flow



Handling user interruptions

If a user changes their request while a tool is executing, Nova 2 Sonic handles it intelligently without canceling pending tools calls.

Async Tool Calling with User Interruption



Example Scenario

```

User: "Can I book a flight from Boston to Chicago?"
Agent: "Sure, let me look that up for you."
Agent: [initiates tool call for Chicago flights]
User: "Actually, I want to go to Seattle"
Agent: "Ok let me update that search"
Agent: [initiates tool call for Seattle flights]
[First tool returns with Chicago flight results]
Agent: [receives Chicago results and processes them contextually]
  
```


How it works

Tool results are always delivered: When a tool call completes, its result is always sent to the model, even if the user has changed their request. The model uses its reasoning capabilities to determine how to handle the information.

Context-aware processing: The model understands the conversation context and can appropriately handle outdated tool results. For example:

- If the user says "thank you" after changing their mind, the model still needs the original results for context
- If the user changes their request, the model can acknowledge the original results while focusing on the new request

No automatic cancellation: The system does not automatically cancel or ignore tool calls based on new user input. This ensures the model has complete information to make intelligent decisions about how to respond.

Integrations

Amazon Nova 2 Sonic can be integrated with various frameworks and platforms to build conversational AI applications. These integrations provide pre-built components and simplified APIs for common use cases.

Strands Agents

Strands Agents is a simple yet powerful SDK that takes a model-driven approach to building and running AI agents. From simple conversational assistants to complex autonomous workflows, from local development to production deployment, Strands Agents scales with your needs.

For comprehensive documentation on the Strands framework, visit the [official Strands documentation](#).

The Strands BidiAgent provides real-time audio and text interaction through persistent streaming connections. Unlike traditional request-response patterns, this agent maintains long-running conversations with support for interruptions, concurrent processing and continuous audio responses.

Prerequisites:

- Python 3.8 or later installed
- Credentials for AWS configured with access to Amazon Bedrock
- Basic familiarity with Python async/await syntax

Code example:

Code example

Installation:

Install the required packages:

```
pip install strands-agents strands-agents-tools
```

Run this example:

```
import asyncio
from strands.experimental.bidi.agent import BidiAgent
from strands.experimental.bidi.io.audio import BidiAudioIO
from strands.experimental.bidi.io.text import BidiTextIO
from strands.experimental.bidi.models.novasonic import BidiNovaSonicModel
from strands_tools import calculator

async def main():
    """Test the BidirectionalAgent API."""
    # Audio and Text input/output utility
    audio_io = BidiAudioIO(audio_config={})
    text_io = BidiTextIO()

    # Nova Sonic model
    model = BidiNovaSonicModel(region="us-east-1")

    async with BidiAgent(model=model, tools=[calculator]) as agent:
        print("New BidiAgent Experience")
        print("Try asking: 'What is 25 times 8?' or 'Calculate the square root of 144'")

        await agent.run(
            inputs=[audio_io.input()],
```

```
        outputs=[audio_io.output(), text_io.output()]\n    )\n\nif __name__ == "__main__":\n    try:\n        asyncio.run(main())\n    except KeyboardInterrupt:\n        print("\\nConversation ended by user")\n    except Exception as e:\n        print(f"Error: {e}")\n        import traceback\n        traceback.print_exc()
```

1. Import: Required Models

```
from strands.experimental.bidi.agent import BidiAgent\nfrom strands.experimental.bidi.io.audio import BidiAudioIO\nfrom strands.experimental.bidi.io.text import BidiTextIO\nfrom strands.experimental.bidi.models.novasonic import BidiNovaSonicModel\nfrom strands_tools import calculator
```

- **BidiAgent:** The main agent class that orchestrates bidirectional conversations
- **BidiAudioIO:** Handles audio input and output for speech interactions
- **BidiTextIO:** Provides text output for transcriptions and responses
- **BidiNovaSonicModel** The Nova 2 Sonic model wrapper
- **Calculator:** A pre-built tool for mathematical operations

2. Configure Audio and Text I/O

```
audio_io = BidiAudioIO(audio_config={})\ntext_io = BidiTextIO()
```

The **BidiAudioIO** manages microphone input and speaker output, while **BidiTextIO** displays text transcriptions and responses in the console.

3. Initialize the Model

```
model = BidiNovaSonicModel(region="us-east-1")
```

Create a Nova Sonic model instance. The region parameter specifies the AWS region where the model is deployed.

4. Create and Run the Agent

```
async with BidiAgent(model=model, tools=[calculator]) as agent:
    await agent.run(
        inputs=[audio_io.input()],
        outputs=[audio_io.output(), text_io.output()]
    )
```

The agent is created with:

- Model: The Nova 2 Sonic model to use
- Tools: List of tools the agent can call (like calculator)
- Inputs: Audio input from the microphone
- Outputs: Audio output to speakers and text output to console


Framework integrations

Amazon Nova 2 Sonic can be integrated with various frameworks and platforms to build sophisticated voice applications. The following examples demonstrate integration patterns with popular frameworks.

Amazon Bedrock AgentCore

Amazon Bedrock AgentCore provides a managed runtime environment for deploying Nova 2 Sonic applications with enterprise-grade security and scalability. AgentCore simplifies the deployment of real-time voice AI applications by handling infrastructure, authentication, and WebSocket connectivity.

AgentCore Architecture



Key features:

- Bidirectional streaming - Native support for Nova Sonic's full-duplex streaming interface with real-time event processing and low-latency communication.
- WebSocket infrastructure - Production-ready WebSocket servers with automatic scaling, connection management, and error recovery.
- Container deployment - Deploy Nova Sonic applications as containers to managed infrastructure with horizontal scaling and independent versioning.
- Enterprise security - Fine-grained authentication via IAM and SigV4, VPC isolation, and comprehensive audit logging.

The architecture shows how client applications connect to AgentCore Runtime via WebSocket with SigV4 authentication. The containerized environment includes your WebSocket server, application logic, and Nova Sonic client, all communicating with Nova Sonic through the bidirectional streaming API.

Benefits:

- **Simplified operations:** Focus on application logic while AgentCore manages infrastructure, scaling, and reliability.
- **Enterprise security:** Built-in authentication, authorization, and compliance features for production deployments.
- **Cost efficiency:** Pay only for what you use with automatic scaling and resource optimization.
- **Developer productivity:** Reduce time to production with managed WebSocket infrastructure and container deployment.

Use cases


- Customer service voice assistants with secure authentication
- Enterprise voice applications requiring IAM integration
- Multi-tenant voice platforms with isolated deployments
- Voice-enabled applications requiring compliance and audit trails

For detailed documentation on deploying Nova Sonic with AgentCore, visit the [Amazon Bedrock AgentCore Documentation](#).

LiveKit

LiveKit is an open-source platform for building real-time audio and video applications. Integration with Amazon Nova 2 Sonic enables developers to build conversational voice interfaces without managing complex audio pipelines or signaling protocols.

For detailed implementation examples and code examples, visit the [LiveKit AWS Integration Documentation](#).



How it works:

- **Client layer:** Web, mobile, or desktop applications connect using LiveKit's client SDKs, which handle audio capture, WebRTC streaming, and playback.
- **LiveKit Server:** Acts as the real-time communication hub, managing WebRTC connections, routing audio streams, and handling session state with low-latency optimization.
- **LiveKit Agent:** Python-based agent that receives audio from the server, processes it through the Nova Sonic plugin, and streams responses back. Includes built-in features like voice activity detection and turn management.
- **Amazon Nova 2 Sonic:** Processes the audio stream through bidirectional streaming API, performing speech recognition, natural language understanding, and generating conversational responses with synthesized speech.


Pipecat

Pipecat is a framework for building voice and multimodal conversational AI applications. It provides a modular, pipeline-based architecture that orchestrates multiple components to create intelligent voice applications with Amazon Nova Sonic and other AWS services.

For detailed implementation examples and code samples, visit the [PipeCat AWS Integration Documentation](#).

Key features:

- Pipeline architecture: Modular Python-based framework for composing voice AI components including ASR, NLU, TTS, and more.
- Pipecat flows: State management framework for building complex conversational logic and tool execution.
- WebRTC Support: Built-in integration with Daily and other WebRTC providers for real-time audio streaming.
- AWS Integration: Native support for Amazon Bedrock, Amazon Transcribe, and Amazon Polly.



The architecture includes:

- WebRTC Transport: Real-time audio streaming between client devices and application server.
- Voice activity detection (VAD): Silero VAD with configurable speech detection and noise suppression.
- Speech recognition: Amazon Transcribe for accurate, real-time speech-to-text conversion.
- Natural language understanding: Amazon Nova Pro on Bedrock with latency-optimized inference.
- Tool execution: Pipecat Flows for API integration and backend service calls.
- Response generation: Amazon Nova Pro for coherent, context-aware responses.
- Text-to-speech: Amazon Polly with generative voices for lifelike speech output.


Deploy on AWS

Deploy your Nova Sonic applications to AWS using infrastructure as code with AWS CDK (Cloud Development Kit). This approach provides repeatable, version-controlled deployments with best practices built in.

Deployment options

- Amazon ECS (Elastic Container Service): Fully managed container orchestration with Application Load Balancer integration, auto-scaling, and serverless Fargate execution.
- Amazon EKS (Elastic Kubernetes Services): Managed Kubernetes for complex orchestration, advanced networking, multi-region deployments, and extensive tooling ecosystem.
- AWS CDK: AWS CDK allows you to define cloud infrastructure using familiar programming languages.

For a complete, production-ready example of deploying Nova Sonic with AWS CDK, see the [Speech-to-Speech CDK Sample](#) on GitHub. This sample demonstrates:



- Complete CDK infrastructure setup with TypeScript
- WebSocket server implementation for real-time communication
- Container deployment with ECS and Fargate

- Application Load Balancer configuration for WebSocket support
- VPC networking and security group setup
- CloudWatch monitoring and logging
- Best practices for production deployments


Multi-agentic systems

Multi-agent architecture is a widely used pattern for designing AI assistants that handle complex tasks. In a voice assistant powered by Nova 2 Sonic, this architecture coordinates multiple specialized agents, where each agent operates independently to enable parallel processing, modular design, and scalable solutions.

Nova Sonic serves as the orchestrator in a multi-agent system, performing two key functions:

Conversation flow management: Ensures all necessary information is collected before proceeding to the next step in the conversation.

Intent classification: Analyzes user inquiries and routes them to the appropriate specialized sub-agent.



The diagram above shows a banking voice assistant that uses a multi-agent architecture. The conversation flow begins with a greeting and collecting the user's name, then handles inquiries related to banking or mortgages through specialized sub-agents.

Conversation flow example:

1. User connects to voice assistant.
2. Nova 2 Sonic: "Hello! What's your name?"
3. User: "My name is John"
4. Nova 2 Sonic: "Hi John, how can I help you today?"
5. User: "I want to check my account balance"
6. Nova 2 Sonic: [Routes to Authentication Agent]
7. Authentication Agent: "Please provide your account ID"
8. User: "12345"
9. Authentication Agent: [Verifies identity]

10Nova 2 Sonic: [Routes to Banking Agent]

11Banking Agent: "Your current balance is \$5,431,10"

While this example demonstrates sub-agents using the Strands Agents framework deployed on Amazon Bedrock AgentCore, the architecture is flexible. You can choose:

- Your preferred agent framework
- Any LLM provider
- Custom hosting options
- Different orchestration patterns

Benefits:

- **Modularity:** Each agent focuses on a specific domain, making the system easier to maintain and update.
- **Scalability:** Add new agents without modifying existing ones, allowing your system to grow with your needs.
- **Parallel processing:** Multiple agents can work simultaneously, improving response times for complex queries.
- **Specialization:** Each agent can be optimized for its specific task, using the most appropriate tools and knowledge bases.
- **Fault isolation:** If one agent fails, others continue to function, improving overall system reliability.

Refer to [this blog](#) for more details and code examples.

See the [Nova Sonic Workshop Multi-Agent Lab](#) for hands-on samples.

Telephony integration

Amazon Nova 2 Sonic integrates with telephony providers to enable AI-powered voice applications accessible via phone calls. This guide covers integration with Twilio, Vonage, and other SIP-based systems for building contact center solutions and voice agents.

Twilio: Cloud communications platform with programmable voice and media streaming capabilities.

Vonage: Global communications APIs with voice, WebSocket audio streaming, and SIP connectivity.

AWS provides a comprehensive sample implementation demonstrating Nova Sonic in a contact center environment with real-time analytics and telephony integration.

Repository: [Sample Sonic Contact Center with Telephony](#)

Advanced systems with Amazon Nova

Overview

This section provides comprehensive guidance for power users implementing advanced patterns with Nova models. These capabilities enable sophisticated applications including tool integration, retrieval-augmented generation (RAG), and autonomous AI agents.

Extended thinking in Amazon Nova 2

Amazon Nova 2 Lite introduces **extended thinking** capabilities that enable the model to engage in deeper reasoning for complex problems. This optional feature gives you control over when and how the model allocates additional computational resources to think through challenging tasks.

How Extended Thinking Works

Amazon Nova 2 introduces extended thinking as a **hybrid capability**. You have complete control:

- **Extended thinking OFF (default):** Amazon Nova 2 operates with efficient latent reasoning, optimal for everyday tasks and high-volume applications.
- **Extended thinking ON:** Amazon Nova 2 engages in explicit, step-by-step reasoning best for complex problems requiring deep analysis.

Reasoning effort

The following shows different levels of reasoning effort.

Low Effort (`maxReasoningeffort:"low"`)

Best for: Tasks with added complexity requiring structured thinking. For example, you can use this for code review and improvement suggestions where the model needs to carefully consider existing code quality, analysis tasks that require thoughtful consideration of multiple factors, or problem-solving scenarios that benefit from a methodical approach. Low effort is ideal for compound tasks where basic reasoning improves accuracy without requiring deep multi-step planning.

Medium effort (maxReasoningEffort: "medium")

Best for: Multi-step tasks and coding workflows. For example, you can use this for software development and debugging where the model needs to understand existing code structure before implementing changes, code generation that requires coordination across multiple files or components, multi-step calculations with interdependencies, or planning tasks with multiple constraints. Medium effort is optimal for agentic workflows that coordinate multiple tools and require the model to maintain context across several sequential operations.

High Effort (maxReasoningeffort:"high")

Best for: STEM reasoning and advanced problem-solving. For example, you can use this for advanced mathematical problems and proofs that require rigorous step-by-step verification, scientific analysis and research tasks demanding deep investigation, complex system design with architectural considerations across multiple dimensions, or critical decision-making scenarios with significant implications. High effort delivers maximum accuracy for tasks requiring sophisticated reasoning, careful evaluation of alternatives, and thorough validation of conclusions.

Quick start: Enabling Extended Thinking

Extended thinking is controlled through the `reasoningConfig` parameter.

```
import boto3

bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

response = bedrock.converse(
    modelId='us.amazon.nova-2-lite-v1:0',
    system=[{"text": "You are a highly capable personal assistant"}],
    messages=[{
        "role": "user",
        "content": [{"text": "Provide a meal plan for a gluten free family of 4."}]
    }],
    inferenceConfig={
        "temperature": 0.7,
        "topP": 0.9,
        "maxTokens": 10000
    },
    additionalModelRequestFields={
        "reasoningConfig": {
            "type": "enabled",
            "maxReasoningEffort": "low"
```



```
    }  
  }  
)  
  
content_list = response["output"]["message"]["content"]  
  
for item in content_list:  
    if "reasoningContent" in item:  
        reasoning_text = item["reasoningContent"]["reasoningText"]["text"]  
        print("=== REASONING ===")  
        print(reasoning_text)  
        print()  
    elif "text" in item:  
        print("=== ANSWER ===")  
        print(item["text"])
```

Reasoning Parameters:

- type: enabled or disabled (default: disabled)
- maxReasoningEffort: low, medium, or high

Note

Temperature, topP and topK cannot be used with maxReasoningEffort set to high. Using these parameters together causes an error.

For full examples of code that utilizes extended thinking, see [the section called “Code library”](#).

Response Structure

When you enable extended thinking, responses include reasoningContent blocks followed by text content blocks:

```
{  
  "output": {  
    "message": {  
      "role": "assistant",  
      "content": [  
        {
```

```

        "reasoningContent": {
            "reasoningText": {
                "text": "[REDACTED]"
            }
        },
        {
            "text": "Based on the premises, we can conclude..."
        }
    ]
},
"stopReason": "end_turn"
}

```

Note

With Amazon Nova 2, reasoning content displays as [REDACTED]. You're still charged for reasoning tokens as they contribute to improved output quality. We include this field in the response structure now to preserve the option of exposing reasoning content in the future. We are actively working with customers to determine the best approach for surfacing the model's reasoning process. Reasoning content is displayed as [REDACTED]. You are still charged for reasoning tokens as they contribute to improved response quality.

Configuration Options

Amazon Nova 2 introduces a new `reasoningConfig` parameter that you can add to your existing converse request structure to enable reasoning:

```

additionalModelRequestFields={
  "reasoningConfig": {
    "type": "enabled", # or "disabled" (default)
    "maxReasoningEffort": "high" # "low", "medium", or "high"
  }
}

```

Parameters:

- **type:** Toggle between "enabled" and "disabled" (default is "disabled")

- **maxReasoningEffort:** When enabled, control reasoning depth
- **"low"**: Moderately complex tasks
- **"medium"**: Complex problems requiring substantial analysis
- **"high"**: Most thorough reasoning for highly complex tasks

Note

when using "high", temp, topP, and maxToken must be unset. In this mode, the model performs deeper analysis to find the best solution. This more thorough processing may generate output that exceeds 65k tokens. The exact amount depends on your request's complexity but for some problems we have seen it go up to 128K tokens. This ensures you get complete, high-quality reasoning rather than truncated results.

Supported models

Extended thinking is currently available in: Amazon Nova 2 Lite (us.amazon.nova-2-lite-v1:0)

Extended thinking with tool use

Extended thinking works seamlessly with tool calling, allowing Amazon Nova to reason about which tools to use and how to interpret their results.

Understanding Reasoning Tokens and Pricing

Token types

Extended thinking tokens are billed as output tokens:

- **Input tokens:** Your original request content (standard input pricing)
- **Output tokens:** This includes reasoning tokens and the final visible response content (standard output pricing)

Usage breakdown

All three token types are included in your usage metrics and billing. Reasoning tokens are priced the same as output tokens and will appear as "REDACTED" in the model response".

```
{
  "usage": {
    "inputTokens": 45,
    "outputTokens": 1240,
    "totalTokens": 1285
  }
}
```

Frequently Asked Questions

Why does Amazon Nova 2 Lite show "[REDACTED]" for reasoning content instead of displaying the model's thinking process?

Our primary focus for this launch is ensuring Nova 2 delivers best in class intelligence for your tasks, and you'll see this reflected in the improved accuracy.

We recognize that visibility into the reasoning process is valuable, and we've heard strong customer interest in understanding how the model thinks through problems.

We are looking at ways to make this available soon.

You are still billed for reasoning tokens, as they represent actual work that improves your output quality

which will be captured in the outputTokens along with the answer tokens.

How do I know if extended thinking is working if reasoning is redacted?

You can confirm extended thinking is working by checking the output for reasoningContent blocks in the response (these only appear when reasoning is enabled)

Using tools (function calling)

Tools extend Amazon Nova capabilities by connecting the model to external functionality such as APIs, databases and code execution environments. Tool use enables Amazon Nova to access real-time information, perform calculations and interact with external systems.

Understanding the tool use workflow

Tool use with Amazon Nova involves three key phases:

User Query and Tool Definition

Define tools by providing JSON schemas that describe each tool's functionality and input requirements. The tool configuration must include explicit details about when and how to use each tool.

Tool Selection

When a user sends a message, Amazon Nova analyzes it to determine if a tool is necessary. This automatic tool choice examines the context and decides which tool (if any) to invoke. If Amazon Nova identifies a suitable tool, it returns the tool name and required parameters.

You are responsible for executing the tool based on the model's request. This means writing code that invokes the tool's functionality and processes the input parameters provided by the model.

Return Results

After executing the tool, send the results back to Amazon Nova in a structured format using JSON or a combination of text and images. Amazon Nova incorporates the tool's output into the final response. If errors occur during execution, denote this in the tool response to allow Amazon Nova to adjust accordingly.

Create a tool

Define tools using a tool configuration that includes an array of tools and optionally a tool choice parameter. Each tool specification must include:

- **Name:** Clear identifier for the tool
- **Description:** Concise explanation of the tool's functionality
- **Input Schema:** JSON schema defining required and optional parameters

Example Tool configuration example

```
tool_config = {  
    "tools": [  

```

```
{
  "toolSpec": {
    "name": "calculator",
    "description": "A calculator tool that can execute a math equation",
    "inputSchema": {
      "json": {
        "type": "object",
        "properties": {
          "equation": {
            "type": "string",
            "description": "The full equation to evaluate"
          }
        },
        "required": ["equation"]
      }
    }
  }
}
```

Best practices for tool definitions

- Ensure the name and description explicitly convey the tool's exact functionality; avoid tools that are overly semantically similar
- Include key differentiators in the description to help the model distinguish between similar tools
- Limit JSON schemas to two layers of nesting for best performance
- Constrain inputs using schema types (e.g., enum, int, float) rather than describing structure in plain text
- Denote required vs. optional parameters using JSON schema notation (e.g., "required": ["param1", "param2"])
- Validate your JSON schema using a standard validator before submitting
- Place long string arguments last in the schema and avoid nesting them

Structured output with constrained decoding

Amazon Nova models leverage constrained decoding to ensure high reliability in generated outputs. This technique uses grammar to constrain possible tokens at each generation step, preventing invalid keys and enforcing correct data types based on your defined schema.

Example Structured output example

```

tool_config = {
    "tools": [
        {
            "toolSpec": {
                "name": "ProductAnalysis",
                "description": "Analyze product information from text.",
                "inputSchema": {
                    "json": {
                        "type": "object",
                        "properties": {
                            "name": {
                                "type": "string",
                                "description": "Product name"
                            },
                            "rating": {
                                "maximum": 5,
                                "description": "Customer rating 1-5",
                                "type": ["number", "null"],
                                "minimum": 1
                            },
                            "features": {
                                "description": "Key product features",
                                "type": "array",
                                "items": {"type": "string"}
                            },
                            "category": {
                                "type": "string",
                                "description": "Product category"
                            },
                            "price": {
                                "type": "number",
                                "description": "Price in USD"
                            }
                        }
                    },
                    "required": ["name", "category", "price", "features"]
                }
            }
        }
    ],
    "toolChoice": {
        "tool": {"name": "ProductAnalysis"}
    }
}

```

```
}  
}
```

Tool choice options

Amazon Nova supports three tool choice parameters:

Tool

The specified tool will be called once, ideal for structured output use cases

Any

One of the provided tools will be called at least once, useful for API selection scenarios

Auto

The model decides whether to call a tool and how many tools to call (default behavior)

Call a tool

When Amazon Nova decides to call a tool, it returns a tool use block as part of the assistant message with `stopReason` set to `"tool_use"`. The tool block contains the tool name and its inputs.

Note

Run following code sections sequentially (Calling a tool → Processing tool call → Returning tool results) in a single Python session. To run the examples again, restart your Python session.

Example Calling a tool

```
import boto3  
import json  
  
# Create Bedrock client  
bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')  
  
# Complex calculation that benefits from precise computation  
messages = [{  
    "role": "user",  
    "content": [{
```



```

        "text": "Calculate the compound interest on $10,000 invested at 4.75% annual
rate for 7 years, compounded quarterly. Use the formula  $A = P(1 + r/n)^{nt}$  where
 $P=10000$ ,  $r=0.0475$ ,  $n=4$ ,  $t=7$ "
    ]]
]]

# Define tool configuration with calculator
tool_config = {
    "tools": [{
        "toolSpec": {
            "name": "calculator",
            "description": "Perform mathematical calculations",
            "inputSchema": {
                "json": {
                    "type": "object",
                    "properties": {
                        "expression": {
                            "type": "string",
                            "description": "Mathematical expression to evaluate"
                        }
                    },
                    "required": ["expression"]
                }
            }
        }
    }]
}

# Invoke Model
response = bedrock.converse(
    modelId="us.amazon.nova-2-lite-v1:0",
    messages=messages,
    toolConfig=tool_config
)

# Extract tool use from response
tool = next(
    block["toolUse"]
    for block in response["output"]["message"]["content"]
    if "toolUse" in block
)

print(f"Tool: {tool['name']}")

```

```
print(f"Expression: {tool['input']['expression']}")
```

Processing tool calls

Extract the tool name and arguments from the message, then invoke the tool:

```
def calculate(expression):
    """Evaluate mathematical expression"""
    print(f"Calculating: {expression}")
    P = 10000
    r = 0.0475
    n = 4
    t = 7
    result = P * (1 + r/n) ** (n*t)
    return result

stop_reason = response["stopReason"]

if stop_reason == "tool_use":
    if tool["name"] == "calculator":
        result = calculate(tool["input"]["expression"])
```

Returning tool results

Return tool results using the ToolResultBlock schema:

```
messages.append(response["output"]["message"])

# Add the tool result
messages.append({
    "role": "user",
    "content": [{
        "toolResult": {
            "toolUseId": tool['toolUseId'],
            "content": [{"json": {"result": result}}],
            "status": "success"
        }
    ]
})

# Send the tool result to the model
response = bedrock.converse(
```

```
        modelId="us.amazon.nova-2-lite-v1:0",
        messages=messages,
        toolConfig=tool_config
    )

    # Extract and display final response
    final_text = next(
        block["text"]
        for block in response["output"]["message"]["content"]
        if "text" in block
    )

    print(f"\nFinal Response:\n{final_text}")
```

Error handling

Report errors back to Amazon Nova to allow request modification and retry:

```
tool_result_message = {
    "role": "user",
    "content": [
        {
            "toolResult": {
                "toolUseId": tool["toolUseId"],
                "content": [{"text": "A validation exception occurred on field:
sample.field"}],
                "status": "error"
            }
        }
    ]
}
```

Security considerations

- Validate that tools exist before invoking them
- Ensure inputs are formatted correctly
- Verify appropriate permissions are in place before tool execution
- Rely on session details rather than allowing Amazon Nova to inject user information into tool calls
- Remember that LLMs can hallucinate tool calls; always validate before execution

Built-in system tools

Amazon Nova 2.0 models include fully managed built-in tools that require no custom implementation. Enable these tools with a simple toggle in the Converse API.

Code Interpreter

Code Interpreter allows Nova to securely execute Python code in isolated sandbox environments. This tool is designed for mathematical computations, logical operations, and iterative algorithms.

Note

Code Interpreter is available in the IAD, PDX, and NRT AWS Regions. To ensure your requests are routed to a supported Region, use Global CRIS. When using Bedrock API keys, you'll need to manually add `InvokeTool` permissions to the policy definitions. The default Bedrock role does not allow the `InvokeTool` action.

Enable Code Interpreter by specifying the `systemTool` parameter:

```
import boto3
import json

bedrock = boto3.client('bedrock-runtime', region_name='us-east-1')

tool_config = {
    "tools": [{
        "systemTool": {
            "name": "nova_code_interpreter"
        }
    }]
}

response = bedrock.converse(
    modelId="us.amazon.nova-2-lite-v1:0",
    messages=[{
        "role": "user",
        "content": [{
            "text": "What is the average of 10, 24, 2, 3, 43, 52, 13, 68, 6, 7, 902, 82"
        }]
    }])
```

```

    ]],
    toolConfig=tool_config,
    inferenceConfig={"maxTokens": 10000, "temperature": 0}
)

# Pretty print the response
for block in response["output"]["message"]["content"]:
    if "toolUse" in block:
        print("=== Tool Use ===")
        print(f"Tool: {block['toolUse']['name']}")
        print(f"Code:\n{block['toolUse']['input']['snippet']}\n")

    elif "toolResult" in block:
        print("=== Tool Result ===")
        result = block['toolResult']['content'][0]['json']
        print(f"Output: {result['stdOut']}")
        if result['stdErr']:
            print(f"Error: {result['stdErr']}")
        print(f"Exit Code: {result['exitCode']}\n")

    elif "text" in block:
        print("=== Final Answer ===")
        print(block["text"])

```

The interpreter runs code in a sandbox and returns results in a standard schema:

```

{
  "stdOut": "String",
  "stdErr": "String",
  "exitCode": "int",
  "isError": "boolean"
}

```

Web Grounding

Web grounding enables Amazon Nova to access real-time information from the internet, providing up-to-date responses and reducing hallucinations. Enable by specifying the `nova_grounding` system tool:

```

tool_config = {
  "tools": [{
    "systemTool": {"name": "nova_grounding"}
  ]
}

```

```
}]  
}
```

For detailed information about Web Grounding, see [the section called “Web Grounding”](#).

Model Context Protocol (MCP)

The Model Context Protocol (MCP) is an open standard that enables secure, two-way connections between data sources and AI-powered tools. Instead of writing custom adapters for each API or service, run an MCP server and let Amazon Nova discover its tools automatically through a client bridge.

Once connected, Amazon Nova treats MCP tools like any other external integration: it decides when to call them, sends required parameters and incorporates results into responses. Using Amazon Nova with Strands makes this easier with a built-in MCPClient that manages discovery, connection and result mapping automatically.

Building AI agents

Amazon Nova models are optimized for building AI agents with Amazon Nova Act. The models provide improved tool use, better reasoning for multi-step tasks, enhanced ability to maintain context across complex agent workflows and support for remote MCP tools.

Create an agent

AI agents built with Nova can orchestrate multiple tool calls, maintain context across extended interactions and course-correct when needed. Extended thinking transforms agentic workflows by enabling systematic reasoning through complex goals. Consider using a planning framework SDK such as Strands Agents to make the planning and execution process of your agent systems more robust.

Agent design patterns

When designing agents with Nova:

- Enable reasoning on medium or high for best results for complex multi-step workflows requiring planning and verification
- Implement tool choice auto to allow flexible tool selection across agent interactions

- Design error handling that allows agents to recover and retry with modified approaches
- Maintain conversation history to preserve context across agent interactions
- Implement robust content filtering and moderation mechanisms across uncontrolled content that your agent system consumes. For example, Amazon offers Amazon Bedrock Guardrails, a feature designed to apply safeguards across multiple foundation models, knowledge bases and agents. These guardrails can filter harmful content, block denied topics, and redact sensitive information such as personally identifiable information.

Multi-tool agent example

```
tool_config = {
  "tools": [
    {
      "toolSpec": {
        "name": "calculator",
        "description": "Perform mathematical calculations",
        "inputSchema": {
          "json": {
            "type": "object",
            "properties": {
              "expression": {
                "type": "string",
                "description": "Mathematical expression to evaluate"
              }
            }
          },
          "required": ["expression"]
        }
      }
    },
    {
      "toolSpec": {
        "name": "database_query",
        "description": "Query financial database for historical data",
        "inputSchema": {
          "json": {
            "type": "object",
            "properties": {
              "query": {
                "type": "string",
                "description": "SQL query to execute"
              }
            }
          }
        }
      }
    }
  ]
}
```

```

        #}
        #},
        #"required": ["query"]
    #}
#}
#}
#}
#]
}

response = client.converse(
    #modelId=" us.amazon.nova-2-lite-v1:0",
    #messages=[{
        #"role": "user",
        #"content": [{
            #"text": "Analyze our Q3 financial performance across all business units,
calculate year-over-year growth rates with statistical significance testing, and
recommend budget allocation strategies for Q4."
        #}]
    #}],
    #toolConfig=tool_config,
    #inferenceConfig={"maxTokens": 10000, "temperature": 1, "topP": 0.9},
    #additionalModelRequestFields={
        #"reasoningConfig": {
            #"type": "enabled",
            #"maxReasoningEffort": "low"
        #}
    #}
)

```

Invoke an agent

Agent invocation involves managing the conversation flow, processing tool calls, and maintaining state across multiple interactions.

Streaming agent responses

Stream responses to provide real-time visibility into the agent's reasoning and actions:

```
import boto3
response = client.converse_stream(
    #modelId=" us.amazon.nova-2-lite-v1:0",
    #messages=[{
```



```

    #"role": "user",
    #"content": [{
        #"text": "Design a scalable microservices architecture for an e-commerce
platform handling 1M+ daily transactions. Consider data consistency, fault tolerance,
performance, security, and cost optimization."
    #}]
#}],
#inferenceConfig={"maxTokens": 10000, "temperature": 10},
#additionalModelRequestFields={
    #"reasoningConfig": {
        #"type": "enabled",
        #"maxReasoningEffort": "low"
    #}
#}
)

# Process the streaming response
reasoning_complete = False
for event in response["stream"]:
    #if "contentBlockDelta" in event:
        #delta = event["contentBlockDelta"]["delta"]

        #if "reasoningContent" in delta:
            #reasoning_text = delta["reasoningContent"]["reasoningText"]["text"]
            #print(f"{reasoning_text}", end="", flush=True)
        #elif "text" in delta:
            #if not reasoning_complete:
                #print(f"

Final Architecture Design:
")

    #reasoning_complete = True
    #print(f"{delta['text']}", end="", flush=True)

```

Agent state management

Maintain conversation history and tool results to preserve context; the example below demonstrates this for a single turn, but the developer can determine how to orchestrate overall agent system based on workflow requirements. In addition, Amazon Web Services tools such as Strands manage agent context and tool state on behalf of the developer.

```
messages = []
```

```
messages = []

# Initial user query
messages.append({
    # "role": "user",
    # "content": [{"text": user_query}]
})

# Get agent response
response = client.converse(
    # modelId=" us.amazon.nova-2-lite-v1:0",
    # messages=messages,
    # toolConfig=tool_config,
    # inferenceConfig=inf_params
)

# Add assistant response to history
messages.append(response["output"]["message"])

# Process tool calls and add results
if response["stopReason"] == "tool_use":
    # tool = next(
    #     # block["toolUse"]
    #     # for block in response["output"]["message"]["content"]
    #     # if "toolUse" in block
    # )

    ## Execute tool
    # result = execute_tool(tool["name"], tool["input"])

    ## Add tool result to conversation
    # messages.append({
    #     # "role": "user",
    #     # "content": [{
    #         # "toolResult": {
    #             # "toolUseId": tool["toolUseId"],
    #             # "content": [{"json": result}],
    #             # "status": "success"
    #         }
    #     }]
    # })

    ## Continue conversation
    # response = client.converse(
```

```
#modelId=" us.amazon.nova-2-lite-v1:0",
#messages=messages,
#toolConfig=tool_config,
#inferenceConfig=inf_params
#)
```

Agent Best Practices

For more information about Agent Best Practices, see [the section called “Best practices”](#).

For guidance on developing Conversational AI agents, see [Speech-to-Speech](#).

Web Grounding

Web Grounding enables Amazon Nova to search the web for current information and provide responses with citations. This feature is useful for queries requiring up-to-date information beyond the model's training data.

How Web Grounding works

When Web Grounding is enabled for a prompt, the following steps are performed:

1. **Request Configuration:** Your application sends a user prompt to the Amazon Bedrock API with `nova_grounding` enabled as a `systemTool`.
2. **Search & Analysis:** The model determines if search is needed, performs one or more searches for relevant information, and evaluates whether additional searches are required to expand its understanding or dive deeper on specific subtopics.
3. **Response generation:** Amazon Nova automatically synthesizes information from search results to generate a final API response grounded in real-time information, complete with citations to its sources.

How to use Web Grounding

For full examples of code that utilizes Web Grounding, see the Code Samples section.

To include Web Grounding in your results, specify the following `systemTool` parameter in your `toolConfig` block:

```
import boto3
```

```
from botocore.config import Config

# Create the Bedrock Runtime client with extended timeout
bedrock = boto3.client(
    "bedrock-runtime",
    region_name="us-east-1",
    config=Config(read_timeout=3600)
)

# Define the tool configuration
tool_config = {
    "tools": [{
        "systemTool": {
            "name": "nova_grounding"
        }
    }]
}

# Send the request
response = bedrock.converse(
    modelId="us.amazon.nova-2-lite-v1:0",
    messages=[{
        "role": "user",
        "content": [{"text": "What are the latest developments in quantum computing?"}]
    }],
    toolConfig=tool_config
)

# Extract text with interleaved citations
output_with_citations = ""
content_list = response["output"]["message"]["content"]
for content in content_list:
    if "text" in content:
        output_with_citations += content["text"]
    elif "citationsContent" in content:
        citations = content["citationsContent"]["citations"]
        for citation in citations:
            url = citation["location"]["web"]["url"]
            output_with_citations += f" [{url}]"

print(output_with_citations)
```

Regional availability

Web Grounding is currently only available in US regions and supported only by US CRIS profiles.

Response structure

The following is an example response. The response has been shortened for brevity:

```
{
  "output": {
    "message": {
      "content": [
        {
          "text": "Recent quantum computing developments include...",
          "citationsContent": [
            {
              "location": {
                "web": {
                  "url": "https://example.com/quantum-news",
                  "domain": "example.com"
                }
              }
            }
          ]
        }
      ]
    }
  }
}
```

Each citation includes:

- **text**: A segment of the model's generated response.
- **citationsContent**: The primary container for the citation data related to a text segment.
- **citations**: A container within **citationsContent** that holds the location of a citation.
- **location**: A container within **citations** that holds the source of a citation.
- **web**: A container within **location** that holds the web source details.
- **url**: The full web address (URL) of the citation's source.
- **domain**: The root domain of the source url.

Grounding safety

Your data never leaves AWS infrastructure. Model-generated queries stay within AWS services and are never sent to the broader internet. Our expansive internal web search index and knowledge graphs prioritize trustworthy and high-quality sources and filter malicious content on ingress. Finally, we protect your application against indirect prompt injection and misinformation with runtime filtering (note that this mitigation is limited for non-English languages).

Error handling

Do not include a `toolSpec` entry with the name `nova_grounding`. Including a tool with this name will result in an error.

The following is a list of potential errors that can occur when using Web Grounding:

- `malformed_tool_use`
- `max_tokens`
- `malformed_model_output`

Permissions required for built in tools

To ensure your role can access Web Grounding on Amazon Bedrock, you have two options:

1. **Enable `BedrockFullAccess` on your IAM role:** If your role has `BedrockFullAccess`, it will have automatic access to Web Grounding.
2. **Add Specific Permissions (if needed):** If you require more granular access control, add this policy to your role's IAM policy, replacing the account ID with your AWS account ID:

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": ["bedrock:InvokeTool"],
      "Resource": ["arn:aws:bedrock::{111122223333}:system-tool/
amazon.nova_grounding"]
    }
  ]
}
```

Web Grounding has the `aws:requestedRegion` condition key set to "unspecified". If your existing policies or Service Control Policies (SCPs) enforce this condition, you may encounter access issues. Updating the condition to allow an "unspecified" `requestedRegion` can resolve this problem.

 Note

If you enable the Web Grounding tool, you are responsible for your use, and any use by your end users, of output that incorporates grounded information. You will know when your output includes grounded information from citations or links to the source material. You must retain and display these citations and links in the output you provide to your end users.

 Note

Web Grounding is an additional cost. For more information, go to [AWS Bedrock pricing](#).

Prompt engineering guide

Prompt engineering is the practice of designing instructions to effectively communicate with large language models (LLMs) like Amazon Nova. Prompt engineering allows you to control model behavior, improve output quality and build reliable AI-powered applications.

Effective prompt engineering involves crafting inputs that guide models to produce desired outputs. A well-crafted prompt provides clear instructions, relevant context and properly formatted input data.

What is prompt engineering

Prompt engineering refers to the practice of optimizing textual input to a large language model (LLM) to improve output and receive the responses you want. Prompting helps an LLM perform a wide variety of tasks, including classification, question answering, code generation, creative writing, and more. The quality of prompts that you provide to a LLM can impact the quality of the model's responses. This section provides you the necessary information to get started with prompt engineering. It also covers tools to help you find the best possible prompt format for your use case when using a LLM on Amazon Bedrock

The effectiveness of prompts is contingent upon the quality of the information provided and the craftsmanship of the prompt itself. Prompts may encompass instructions, questions, contextual details, inputs, and examples to effectively guide the model and enhance the quality of the results. This document outlines strategies and tactics for optimizing the performance of Amazon Nova 2 Sonic family of models. The methods presented herein may be employed in various combinations to amplify their effectiveness. We encourage you to engage in experimentation to identify the approaches most suitable for their specific needs.

Getting started with prompt engineering

Before you start prompt engineering, we recommend that you have the following elements in place, so you can iteratively develop the most optimal prompt for your use case:

Define your use case

Define your use case across four dimensions:

- **Task** – Define what you want the model to accomplish. This determines the right prompting technique.

- **Role** – Define what role the model should assume to accomplish the task. Amazon Nova models support three roles (System, User, or Assistant).
- **Response Style** – Define the response structure or style that the model should follow based on the audience, such as JSON, markdown, or conversational.
- **Instructions** – Define the set of instructions that the model should follow to meet success criteria.

Establish success criteria

Define success criteria or evaluation metrics. You can provide a list of criteria or provide specific evaluation metrics, such as length, BLEU score, ROUGE, format, factuality and faithfulness.

Draft a prompt

Create a starting prompt incorporating your task, role, response style and instructions. Iterate based on results.

The effectiveness of prompts depends on the quality of information you provide.

General best practices

The following best practices mainly apply to the Amazon Nova text models, but you can apply them to other models, in addition to modality-specific best practices.

For more information on how to prompt multimodal inputs, refer to [Prompting multimodal inputs](#). For information on how to prompt speech inputs, refer to [Voice conversation prompts](#).

Understanding the roles

Amazon Nova models allow you to structure prompts through the use of three distinct roles: system, user, and assistant.

- **System (optional)** – Although not mandatory, it establishes the overall behavioral parameters of the assistant. It can also be utilized to provide additional instructions and guidelines that the user wishes the model to adhere to throughout the conversation.
- **User** – Can optionally convey the context, tasks, instructions, and the desired outcome along with the user query.
- **Assistant** – Aids in guiding the model towards the intended response.

Topics

- [Create precise prompts](#)
- [Bring focus to sections of the prompt](#)
- [Using the system role](#)
- [Provide examples \(few-shot prompting\)](#)
- [Tool calling systems](#)
- [Advanced prompting techniques](#)

Create precise prompts

Crafting specific user queries is crucial in prompt engineering. Well-crafted queries direct Amazon Nova text understanding models to generate accurate and pertinent responses. To craft such queries, it is essential to start by providing contextual information in the query. The context provided helps situate the model better by aligning to clear target-user scenarios, thereby enhancing the relevance and coherence of the output.

Providing contextual, specific information

Models work best if the instructions are clear, well structured, short, and specific. This prompt template guidance applies to both the user prompt and the system prompt.

Prompt template

```
# Use this to clearly define the task and job needed by the model
task_summary = f"""
## Task Summary:
{{Task summary}}
"""

# Use this to provide contextual information related to the task
context_information = f"""
## Context Information:
- {{Context and content information 1}}
- {{Context and content information 2}}
...
"""

# Use this to provide any model instructions that you want model to adhere to
model_instructions = f"""
```

```
## Model Instructions:
- {{ Other Model Instructions }}
...
"""

# Use this to provide response style and formatting guidance
response_style = f"""
## Response style and format requirements:
- {{Style and format requirement 1}}
- {{Style and format requirement 2}}
...
"""

# Concatenate to final prompt
final_prompt = f"""{task_summary}
{context_information}
{model_instructions}
{response_style}"""
```

Example: Specific and contextual information

Here is an example of a prompt that is rewritten with specific and contextual information. Adding more contextual information helps the model ground towards the desired behavior.

Unclear prompt

User:

Write me a meeting invite to the project team

Assistant:

Certainly! Here's a sample meeting invite for your project team:

****Subject:**** Project Team Meeting - [Project Name]

****Date:**** [Insert Date]

****Time:**** [Insert Time]

****Duration:**** [Insert Duration]

****Location:**** [Insert Location] / [Zoom/Teams Link if Online]

****Agenda:****

1. ****Welcome and Introductions****
 - Quick round of introductions for new team members
2. ****Project Overview****
 - Objectives and goals
 - Current status and milestones achieved
3. ****Task Allocation and Deadlines****
 - Detailed breakdown of tasks
 - Assignment of responsibilities
 - Discussion of deadlines
4. ****Risk Management****
 - Identification of potential risks
 - Mitigation strategies
5. ****Resource Allocation****
 - Required tools and materials
 - Budget review
6. ****Q&A Session****
 - Open floor for questions and clarifications
7. ****Next Steps and Closing****
 - Summary of action items
 - Schedule for the next meeting

****Additional Notes:****

- Please come prepared with any relevant documents or updates regarding your assigned tasks.
- If you anticipate any issues or require additional support, please let us know in advance so we can address them during the meeting.

****Contact Information:****

For any questions or concerns before the meeting, please feel free to reach out to [Your Name] at [Your Email] or [Your Phone Number].

We look forward to a productive discussion and to kicking off our project with enthusiasm!

Best regards,

[Your Full Name]

[Your Position]

[Your Contact Information]

Feel free to customize this template to better fit your specific project and team needs!

Clear prompt

User:

Write me a meeting invite to the project team. **The meeting is about project planning, and should be held on October 16th, from 10 AM to 11 AM in Conference Room B. Please include an agenda that covers our progress on the project thus far, as well as any upcoming milestones and deadlines.**

Assistant:

****Meeting Invitation: Project Planning Session****

****Date:**** October 16th

****Time:**** 10:00 AM - 11:00 AM

****Location:**** Conference Room B

****Agenda:****

1. ****Welcome and Introduction (10:00 AM - 10:05 AM)****

- Brief overview by the meeting facilitator.

2. ****Project Progress Update (10:05 AM - 10:30 AM)****

- Review of completed tasks and achievements to date.

- Detailed discussion on current project status and any challenges encountered.

3. ****Upcoming Milestones and Deadlines (10:30 AM - 10:50 AM)****

- Overview of the next phase of the project.

- Identification of key milestones and deadlines.

4. ****Action Items and Responsibilities (10:50 AM - 11:00 AM)****

- Assignment of tasks for the upcoming period.

- Confirmation of individual responsibilities and follow-up items.

5. ****Q&A Session (11:00 AM)****

- Open floor for any questions or concerns from team members.

Please ensure your attendance to discuss and finalize our project planning for the next phase. Your insights and contributions are highly valued.

Best regards,

[Your Name]

[Your Position]

[Your Contact Information]

- **Be clear:** When prompting the model, try to keep instructions as clear as possible in a way a human would understand it versus a machine (for instance, a JSON-style list).
- **Be specific:** Give the model specific instructions about what the task is, what the expected output is, and any additional context to help the model respond the way you want.

Specify the output format and style

If you want Amazon Nova 2 to generate responses that are tailored to your specific needs, you can specify the format and style of the output. For example, if you want to generate responses only in JSON or constrain the response length to a maximum number of words, bullet points, or paragraphs. This enables Amazon Nova 2 to better understand your requirements and produce output that is more relevant and easier to parse.

Use capitalization and strong wordings such as DO, DO NOT, MUST to emphasize the rules you want the model to strictly obey. For example, if you want only one JSON object in the response, you can instruct the model by saying "You **MUST** answer in JSON format only. **DO NOT** use any other format while answering the question."

Here are some example prompts that are rewritten with specific and contextual information. Adding more contextual information helps the model ground towards the desired behavior.

Example 1: Provide extra details

Adding extra details on the response formatting can help the model accomplish the desired result. Some of the other commonly used response format include CSV, Markdown, YAML, and XML

Prompt without style constraints

User:

Outline the different types of rooms available in a typical 4-star hotel, including their features.

Assistant:

In a typical 4-star hotel, you'll find a range of room types, each designed to cater to different guest needs and preferences. Here's an outline of the different types of rooms available along with their features:

1. **Standard Room**

- **Size:** Usually around 250-350 square feet.
- **Features:**
 - Queen or double bed.
 - Complimentary Wi-Fi.
 - Flat-screen TV with cable or satellite channels.
 - Mini-fridge.
 - Coffee maker with complimentary coffee and tea.
 - Desk and chair.
 - En-suite bathroom with shower, toilet, and toiletries.
 - Air conditioning/heating.

2. **Deluxe Room**

- **Size:** Typically larger than standard rooms, around 350-450 square feet.
- **Features:**
 - King bed or two queen beds.
 - Larger bathroom with a bathtub.
 - Additional seating area.
 - Enhanced amenities like bathrobes and slippers.
 - Upgraded toiletries and sometimes a welcome gift.
 - Access to a lounge or executive floor perks.

3. **Suite**

- **Size:** Varies significantly, often between 500-1,200 square feet.
- **Features:**
 - Separate living area with a sofa bed or additional seating.
 - Kitchenette or full kitchen in some cases.
 - Larger, more luxurious bathroom with a bathtub and separate shower.
 - Multiple TVs.
 - Access to executive lounge.

- High-end toiletries and often complimentary breakfast.
- Often includes concierge services and in-room dining options.

4. **Family Room**

- **Size:** Larger, usually around 400-500 square feet.
- **Features:**
 - One king bed and one sofa bed or two queen beds.
 - Designed for families with children.
 - Extra space for luggage.
 - Child-friendly amenities like a small kitchenette or high chair.
 - Family bathroom with amenities suitable for kids.

5. **Executive Room**

- **Size:** Typically around 350-450 square feet.
- **Features:**
 - Business-oriented with a larger workspace.
 - Executive desk and chair.
 - Access to executive lounge with complimentary snacks and refreshments.
 - King bed or two queen beds.
 - Enhanced amenities like high-speed internet and sometimes a Nespresso machine.

6. **Junior Suite**

- **Size:** Usually around 450-600 square feet.
- **Features:**
 - Separate bedroom and living area.
 - King bed or two queen beds.
 - Kitchenette or mini-fridge.
 - Luxurious bathroom with a bathtub.
 - Lounge access.
 - High-end toiletries.

7. **Premium Room**

- **Size:** Larger than standard rooms, around 400-500 square feet.
- **Features:**
 - King bed or two queen beds.
 - Larger bathroom with a bathtub.
 - Enhanced amenities like a welcome drink and premium toiletries.
 - Access to premium services and sometimes a lounge.

8. **Connecting Rooms**

- **Size:** Two rooms that connect via an internal door.
- **Features:**
 - Suitable for families or groups.
 - Often includes one standard room and one suite.

- Shared bathroom or en-suite bathrooms in each room.
- Extra space for luggage and guests.

9. **Accessible Rooms**

- **Size:** Varies, typically similar to standard rooms.
- **Features:**
 - Designed for guests with disabilities.
 - Wider doorways and bathroom space.
 - Adapted fixtures and fittings.
 - Additional support services.

10. **Business Room**

- **Size:** Usually around 300-350 square feet.
- **Features:**
 - Focused on business travelers.
 - High-speed internet and large workspace.
 - Compact yet fully equipped with essentials.
 - Often includes a king bed.

Each room type in a 4-star hotel is designed to provide comfort, convenience, and luxury tailored to different guest needs, ensuring a pleasant stay for all types of travelers.

Prompt with style constraints

User:

Outline the different types of rooms available in a typical 4-star hotel, including their features. **All responses MUST be in YAML format.**

Assistant:

```
```yaml
hotel_rooms:
 - room_type: Standard Room
 features:
 - description: A comfortable room with essential amenities.
 amenities:
 - bed_type: Double bed or two single beds
 - bathroom: Private bathroom with shower
 - wifi: Complimentary high-speed internet
 - TV: Flat-screen TV with cable channels
 - mini_fridge: Available
```

```
- tea_and_coffee_making_facilities: Yes

- room_type: Deluxe Room
features:
- description: An upgraded room with more space and enhanced amenities.
amenities:
- bed_type: King-size bed or two single beds
- bathroom: Private bathroom with both shower and bathtub
- wifi: Complimentary high-speed internet
- TV: Larger flat-screen TV with cable channels
- mini_fridge: Available
- tea_and_coffee_making_facilities: Yes
- free_breakfast: Included

- room_type: Suite
features:
- description: A luxurious room with separate living area and premium amenities.
amenities:
- bed_type: King-size bed or two single beds
- bathroom: Private bathroom with both shower and bathtub
- separate_living_area: Yes
- wifi: Complimentary high-speed internet
- TV: Large flat-screen TV with cable channels
- mini_fridge: Available
- tea_and_coffee_making_facilities: Yes
- free_breakfast: Included
- concierge_service: Available

- room_type: Family Room
features:
- description: A spacious room designed for families with additional bedding
options.
amenities:
- bed_type: Queen bed and rollaway bed or two single beds and rollaway bed
- bathroom: Private bathroom with shower
- wifi: Complimentary high-speed internet
- TV: Flat-screen TV with cable channels
- mini_fridge: Available
- tea_and_coffee_making_facilities: Yes

- room_type: Business Room
features:
- description: A room tailored for business travelers with enhanced workspace.
amenities:
```

```
- bed_type: Double bed or two single beds
- bathroom: Private bathroom with shower
- wifi: Complimentary high-speed internet
- TV: Flat-screen TV with cable channels
- workspace: Desk and chair
- mini_fridge: Available
- tea_and_coffee_making_facilities: Yes
- business_center_access: Available
```
```

Example 2: Provide style constraints

In this case, the prompt with style constraints provides comprehensive information on the constraints to follow and the desired response format. As observed in the response, the optimized prompt aligns more closely with the requested brevity and bulleted list structure.

Prompt without style constraints

User:

```
Explain the process of checking in a guest at a luxury resort.
```

Assistant:

```
Checking in a guest at a luxury resort involves a meticulous and personalized
process designed to ensure a seamless and memorable start to their stay. Here's a
detailed overview of the steps involved:
```

Pre-Arrival Preparation

1. **Pre-Arrival Communication:**

- **Confirmation:** Verify guest reservations and confirm arrival details.
- **Special Requests:** Collect any special requests (dietary restrictions, room preferences, etc.) and ensure they are noted in the guest profile.
- **Documentation:** Ensure all necessary documentation is ready, such as welcome letters, room keys, and any special amenities.

2. **Room Preparation:**

- **Cleaning:** Conduct a thorough cleaning and preparation of the guest room.
- **Amenities:** Stock the room with high-end toiletries, fresh linens, and any requested amenities.

- **Personal Touches:** Add personalized touches such as a welcome note, fresh flowers, or locally sourced treats.

Arrival and Initial Interaction

3. **Guest Arrival:**

- **Greeting:** Have staff ready to greet guests as they arrive. This may include a warm welcome at the entrance or a meet-and-greet in the lobby.
- **Transport:** Provide luxury transport options like a chauffeured car service to escort guests to the check-in desk.

4. **Check-In Desk:**

- **Professionalism:** Ensure check-in staff are well-trained, professional, and exude hospitality.
- **Technology:** Utilize sophisticated check-in technology (like tablet kiosks) to streamline the process.
- **Personalized Welcome:** Greet guests with a warm welcome, acknowledge their special requests, and present them with a personalized welcome kit.

Documentation and Formalities

5. **Document Verification:**

- **Identification:** Verify identification and payment methods.
- **Membership Cards:** If applicable, verify any loyalty or membership cards for additional benefits.

6. **Registration:**

- **Form Filling:** Assist guests with any required paperwork.
- **Special Benefits:** Clearly explain any special resort benefits, upgrades, or exclusive access to amenities that the guest is entitled to.

Room Assignment

7. **Room Presentation:**

- **Walk-Through:** Escort guests to their room and conduct a walk-through to ensure everything is in order and to address any immediate concerns.
- **Upgrades:** Present any room upgrades and explain the features and services available in the upgraded room.

8. **Key Delivery:**

- **Room Key:** Provide the room key, often a high-end key card or a luxury key fob.
- **Access Information:** Explain how to use the key card, including access to other amenities like the spa, pool, and gym.

Additional Services

9. **Welcome Amenities:**

- **Amenity Kits:** Deliver welcome amenity kits (like a bottle of champagne, local delicacies, etc.).
- **Itinerary:** Provide a detailed resort map and itinerary, highlighting all available services and activities.

10. **Personal Concierge:**

- **Concierge Services:** Introduce the guest to the concierge service and explain how to request any special services, such as dining reservations, transportation, or local excursions.
- **Contact Information:** Provide contact information for key staff members.

Final Touches

11. **Feedback:**

- **Satisfaction Survey:** At the end of the check-in process, ask for feedback to ensure the guest's expectations were met and to gather any immediate suggestions.

12. **Follow-Up:**

- **Immediate Issues:** Ensure any immediate issues or requests are addressed promptly.
- **Comfort Check:** Make follow-up calls or send messages to ensure the guest is settling in comfortably.

By following these steps, a luxury resort aims to provide an exceptional check-in experience that sets the tone for an unforgettable stay.

Prompt with style constraints

User:

Explain the process of checking in a guest at a luxury resort. **Provide all answers in a numbered list format, with each point containing no more than 15 words, and no more than 6 bullet points.**

Assistant:

1. Greet guest warmly with a smile and offer assistance.
2. Verify reservation details and check guest's identification.
3. Confirm room assignment and amenities provided.

4. Explain resort services and amenities available.
5. Present room key and offer directions to guest's suite.
6. Wish guest a pleasant stay and be available for queries.

Example 3: Limit the response format

Here, the prompt with the response format forces the model to respond only with the category number and limit it to be between 1-10, instead of being highly verbose.

Prompt without response format

User:

In this task, you are given a text which is the body of a document. Your job is to classify the topic of the document into these categories: 1)Company, 2)Educational Institution, 3)Artist, 4)Athlete, 5)Office Holder, 6)Mean of transportation, 7)Building, 8)Natural place, 9)Village, 10)Animal.
Text: Johangsan is a mountain of South Korea. It has an elevation of 954 metres.

Assistant:

To classify the topic of the given text, let's analyze the content step-by-step:

1. The text mentions "Johangsan," which appears to be a specific name.
2. The text states that "Johangsan is a mountain," which provides a clear indication of its nature.
3. It is further specified that it is located in "South Korea" and has an elevation of "954 metres."

Based on the information provided, the text describes a geographical feature, specifically a mountain. From the given categories, the most appropriate classification for this document is:

8) Natural place

The text does not mention any company, educational institution, artist, athlete, office holder, means of transportation, building, village, or animal. Therefore, the classification as a "Natural place" is the most fitting choice.

Prompt with response format

User:

In this task, you are given a text which is the body of a document. Your job is to classify the topic of the document into these categories: 1)Company, 2)Educational Institution, 3)Artist, 4)Athlete, 5)Office Holder, 6)Mean of transportation, 7)Building, 8)Natural place, 9)Village, 10)Animal. **Your output should be the category number. DO NOT generate anything apart from numbers 1-10.**

Text: Johangsan is a mountain of South Korea. It has an elevation of 954 metres.

Assistant:

8

Provide clear prompt sections

Beyond the initial context, a structured query format plays a significant role in optimizing interactions with the language model. Separating distinct parts of the input with delimiters ensures clarity and enhances Amazon Nova 2 models' ability to understand the query's structure. Another effective way is to use numbered lists or bullet points, which can help to break down the task into manageable steps and provide a clear outline of the actions that need to be taken. This structured approach aids in reducing ambiguity and improving the accuracy of the responses generated.

Example: Create a sectioned prompt

Notice how in the Sectioned Prompt column, the prompt is clearer and has delineated sections for different entities like Tasks, Context, Requirements, and the rest when compared to the instructions in the Unsectioned Prompt column.

Unsectioned Prompt

User:

Write an article about climate change for high school students explaining basic concepts, current impacts, and individual actions to mitigate it. Make it informative and engaging with simple language, including statistics and real-world examples.

Assistant:

Understanding Climate Change: What It Is, What It's Doing, and What You Can Do About It

Introduction

Climate change might sound like a big, complicated topic, but it's really about the way our planet's weather is changing over time. This change is mostly happening because humans are doing things that are making our Earth warmer. Let's break it down so you can understand what's going on and what you can do to help.

What Is Climate Change?

Climate change refers to long-term shifts in temperatures and weather patterns. These shifts may be natural, but since the 1800s, human activities have been the main driver of rapid climate change.

Key Concept: Greenhouse Gases

To understand this, we need to talk about greenhouse gases. These gases, like carbon dioxide (CO₂), methane (CH₄), and nitrous oxide (N₂O), trap heat in our atmosphere. Without these gases, Earth would be too cold for most life forms. But when we burn fossil fuels (like coal, oil, and natural gas) and cut down forests, we release too much CO₂ and other greenhouse gases, trapping more heat and making the planet warmer.

Current Impacts of Climate Change

Let's look at some real-world examples to see how climate change is affecting our world today.

1. Rising Temperatures

Since the late 19th century, Earth's average surface temperature has risen about 2 degrees Fahrenheit (1 degree Celsius). Scientists predict that if we don't change our ways, it could rise by 3.6 to 4.3 degrees Fahrenheit (2 to 2.4 degrees Celsius) by the end of this century.

2. More Extreme Weather

Climate change is making extreme weather events more frequent and severe. For example:

- **Hurricanes:** In the U.S., hurricane seasons are getting longer and more destructive. In 2020, there were 30 named storms, the most ever recorded.
- **Droughts:** Parts of California have experienced some of the worst droughts in history, affecting agriculture and water supplies.

3. Melting Ice Caps and Rising Sea Levels

Glaciers and ice sheets are melting at alarming rates. This not only affects wildlife but also leads to rising sea levels. Since 1901, global sea levels have risen by about 8 inches (20 centimeters). By 2100, they could rise by another 1 to 4 feet (0.3 to 1.2 meters).

****4. Ocean Acidification****

As CO2 levels rise, more of it is absorbed by the oceans, making them more acidic. This is particularly harmful to marine life, including coral reefs. Coral reefs are crucial for marine biodiversity but are now dying at an alarming rate.

Individual Actions to Mitigate Climate Change

You might think that since climate change is a huge problem, one person can't make a difference. But even small actions can add up to make a big impact.

****1. Reduce, Reuse, Recycle****

Try to minimize waste by:

- ****Reducing:**** Buy fewer items and choose products with less packaging.
- ****Reusing:**** Use reusable bags, bottles, and containers.
- ****Recycling:**** Make sure to recycle paper, plastic, and glass properly.

****2. Conserve Energy****

Small changes can lead to big savings:

- Turn off lights when you leave a room.
- Unplug devices that you're not using.
- Use energy-efficient bulbs and appliances.

****3. Use Less Meat****

Livestock farming is a major contributor to greenhouse gas emissions. By eating less meat, especially beef, you can reduce your carbon footprint.

****4. Support Renewable Energy****

If possible, support or switch to renewable energy sources like solar or wind power. Even advocating for these options in your community can make a difference.

****5. Educate Yourself and Others****

Knowledge is power. The more you know about climate change, the more you can help others understand it. Share facts and encourage your friends and family to take action too.

Conclusion

Climate change is a pressing issue that affects everyone on the planet. By understanding its basics, recognizing its current impacts, and taking individual actions, you can play a crucial role in mitigating its effects. Remember, every small step counts, and together, we can make a big difference for our planet.

Let's start today and build a sustainable future!

Sectioned Prompt

User:

```
**Task:**  
Write an article  
  
**Context Section:**  
Topic: Climate Change  
Target Audience: High School Students  
Purpose: Educational Article  
  
**Content Requirements:**  
- Explanation of basic concepts related to climate change.  
- Examples of current impacts on ecosystems and human populations.  
- Suggestions for individual actions to mitigate climate change.  
  
**Tone and Style:**  
- Informative and engaging  
- Use simple language  
- Include statistics and real-world examples
```

Assistant:

```
# Understanding Climate Change: A Guide for High School Students  
  
## Introduction  
  
Climate change is one of the most pressing issues of our time. It affects our planet, our ecosystems, and our daily lives. Understanding what climate change is and how it impacts us can help us take action to make a difference. This article will explain the basic concepts of climate change, provide examples of its current impacts, and suggest actions you can take to help mitigate it.  
  
## What is Climate Change?  
  
**Climate change** refers to long-term changes in global temperatures and weather patterns. These changes are primarily caused by human activities, especially the burning of fossil fuels like coal, oil, and natural gas. When we burn these fuels, we release large amounts of greenhouse gases, such as carbon dioxide (CO2) and methane (CH4), into the atmosphere. These gases trap heat from the sun, causing the Earth's temperature to rise. This is often referred to as the **greenhouse effect**.
```

Basic Concepts:

1. **Greenhouse Gases**: These gases include carbon dioxide, methane, nitrous oxide, and fluorinated gases. They trap heat in the Earth's atmosphere.
2. **Global Warming**: This is the increase in Earth's average surface temperature due to rising levels of greenhouse gases.
3. **Climate Change**: Refers to the broader set of changes in the Earth's climate system, including global warming.

Current Impacts of Climate Change

On Ecosystems:

1. **Melting Ice Caps and Glaciers**: The polar ice caps and glaciers are melting at an alarming rate. For example, the Arctic sea ice is shrinking, which affects polar bears and other Arctic animals.
2. **Rising Sea Levels**: As ice melts, sea levels rise. This can lead to the loss of coastal habitats and increased flooding in coastal areas.
3. **Changing Weather Patterns**: Climate change can lead to more frequent and severe weather events such as hurricanes, droughts, and heatwaves. For example, the 2020 Australian bushfires were more intense and lasted longer than usual due to climate change.

On Human Populations:

1. **Health Issues**: Rising temperatures can lead to more heat-related illnesses and the spread of diseases like malaria and dengue fever.
2. **Food Security**: Changes in climate can affect crop yields, leading to food shortages. For example, droughts in Sub-Saharan Africa have caused significant food insecurity.
3. **Displacement**: Rising sea levels and extreme weather events can force people to leave their homes, becoming climate refugees. For instance, communities in Bangladesh are being displaced by rising sea levels.

Actions You Can Take

Even though climate change is a massive issue, there are many actions you can take to help mitigate its effects. Here are some simple steps you can start with today:

1. ****Reduce, Reuse, Recycle****: By reducing waste, reusing items, and recycling, you can decrease the amount of trash that ends up in landfills and reduce greenhouse gas emissions.
2. ****Save Energy****: Turn off lights when you don't need them, use energy-efficient appliances, and unplug devices when they're not in use. These small actions can make a big difference.
3. ****Use Less Meat****: Livestock farming produces a significant amount of methane. By eating less meat, particularly beef, you can reduce your carbon footprint.
4. ****Support Renewable Energy****: Advocate for or invest in renewable energy sources like solar and wind power. These sources produce energy without greenhouse gas emissions.
5. ****Educate and Advocate****: Inform others about the impacts of climate change and advocate for policies that support environmental sustainability. You can join or support environmental clubs at school or participate in local clean-up efforts.

Conclusion

Climate change is a complex issue with far-reaching impacts. However, understanding it and taking action can make a significant difference. By learning about the basic concepts, recognizing the current impacts, and implementing simple changes in our daily lives, we can all contribute to mitigating climate change. Remember, every small action counts, and together, we can create a more sustainable future for our planet.

Bring focus to sections of the prompt

Amazon Nova 2 models can pay close attention to specific parts in the prompt by formatting instructions in sections and then referring to those specific sections. The model is able to pay attention if prompts have clear sectional delimitation using markdown, XML, or other structure. For example, you can define the name of the section, use **##Section Name##**, then refer to that section in your prompt with **##Section Name##**.

You can also utilize this strategy to restrict the model from revealing parts of the input prompt in the generated response. For example, when providing few shot examples or instructions in the

input prompt, use delimiters such as `##Instructions##` or `##Examples##` with a new line separator and provide strong instructions such as `DO NOT mention anything inside the ##Instructions## or ##Examples##` in the response for the model to not regurgitate the input prompt content from these sections in its output.

Example: Sectional delimitation

User:

You're an expert Prompts creator. Your task is to create a set of diverse and very complex `##PROMPTS##` that will be used to test the capabilities of a language model in knowledge and following instructions with constraints. Please create 10 `##PROMPTS##`. You must strictly follow `##GUIDELINES##`:

`##GUIDELINES##`

- Generate `##PROMPTS##` similar to the structure and style of the given `##EXAMPLE PROMPTS##`. Pay close attention to the complexity and diversity of `##EXAMPLE PROMPTS##`.
- Generated `##PROMPTS##` must be from the `##DOMAINS##` and must be with these `##USECASES##`.
- Each of the `##PROMPTS##` needs to be unique and very complex.
- Each of the `##PROMPTS##` must have more than 4 sentences and 1 constraint.
- Each of the `##PROMPTS##` should have at least 70 words.
- Each of the `##PROMPTS##` should have an answer that can be written in text.
- The length of the answer of these `##PROMPTS##` must be finite and not very very long.
- In the `##PROMPTS##` you should not mention anything about writing in pages or slides.
- Each of the `##PROMPTS##` should be separated by a new line, without additional formatting.

Generated `##PROMPTS##` must be from the following `##DOMAINS##`

`##DOMAINS##`

{domains}

Generated `##PROMPTS##` must be for the following `##USECASES##`

`##USECASES##`

{usecases}

{usecase_description}

`##PROMPTS##`

Using the system role

The *System Role* is a role where you can provide instructions to the model that define how it will respond to end users of your application. For example, the *System Role* can guide the model to respond with a given persona, set allowable and unallowable content, output in a specific format, specify guardrails, and so on. Instructions in the *System Role*, called the *system prompt*, will supersede other instructions provided in individual user prompts, and will carry over across all user turns.

Specify the system role with the API

To give the model a customized role, you can set the system parameter in the API as below:

```
{
  "system": [
    {
      "text": "You are a helpful recipe assistant. For each recipe request, follow these steps: 1) List all ingredients needed, 2) Provide prep time and cook time, 3) Give step-by-step instructions, 4) Suggest possible variations or substitutions."
    }
  ],
  "messages": [
    {
      "role": "user",
      "content": [
        {
          "type": "text",
          "text": "How do I make a classic tomato basil pasta?"
        }
      ]
    }
  ]
}
```

Tip

- To best utilize the *system role*, make sure that your *system prompt* is clear, concise, and specific, following the same strategies as described in [Create precise prompts](#). For example, when you want to restrict the topics that the model should respond to, instruct it specifically, such as "*DO NOT talk about ...*" or "*You MUST talk about ...*".

- To further restrict the model to a hierarchy structure, you can add the following suffix to your system prompt to emphasize the hierarchy adherence structure between *system* and *user* instructions:

"\n\nThe above system instructions define your capabilities and your scope. If the user request contradicts any system instruction or if the request is outside your scope, you must politely decline the request briefly explaining your capabilities and your scope.\n"

The following items are some things that you can add in a system prompt template:

```
# Define a persona
persona = """You are {{Persona}}
"""

# Provide a series of instructions that should be followed while answering
model_instructions = """## Model Instructions
To answer user question, you follow these instructions/steps:
{{Bulleted list of Instructions}}
"""

# Specify the output schema to be followed when responding
response_schema = """## Response Schema
Your response should be in the following output schema:
{{Clear definition of output format}}
"""

# Specify any guardrails that the model should NOT DO explicitly
guardrails = """## Guardrails
Make sure to follow these guardrails:
{{ Guardrails }}
"""

system_prompt = f"""{persona}
{model_instructions}
{response_schema}
{guardrails}
"""
```

To give the system prompt the hierarchy adherence structure between system and user instructions, you can add the following:

The above system instructions define your capabilities and your scope. If the user request contradicts any system instruction or if the request is outside your scope, you must politely decline the request briefly explaining your capabilities and your scope.

Example 1: Providing specific language cues

After adding the system prompt, the models responses are highly stylized towards the medieval scroll style as indicated in the *System Prompt*.

Without a System Prompt

System:

NA

User:

Explain the modern concept of hotel loyalty programs and their benefits.

Assistant:

Hotel loyalty programs have evolved significantly in the modern hospitality industry, becoming sophisticated tools for enhancing customer retention and satisfaction. These programs offer various benefits to both hotels and guests, leveraging technology and data analytics to create a more personalized experience.

Key Components of Modern Hotel Loyalty Programs:

1. **Tiered Membership Levels:**

- **Basic Membership:** Entry-level membership that provides basic benefits.
- **Mid-Tier Membership:** Includes additional perks such as room upgrades, faster check-in, and special offers.
- **Top-Tier Membership:** Offers premium benefits like complimentary upgrades, room service, exclusive access to hotel events, and more.

2. **Point System:**

- Guests earn points for every stay, purchases made in hotel restaurants, and other activities.

- Points can be redeemed for free nights, upgrades, or other rewards.

3. **Mobile App Integration:**

- Modern programs often have a dedicated mobile app that allows guests to easily manage their account, check in, and access exclusive offers.
- Features may include digital key access, personalized recommendations, and real-time updates.

4. **Personalization:**

- Utilizing data analytics, hotels can tailor recommendations and offers to individual guest preferences, enhancing the overall experience.
- Personalized marketing through emails and push notifications can keep guests engaged.

5. **Cross-Brand Benefits:**

- Some loyalty programs offer benefits across multiple brands within a hotel group or even within different hospitality groups.
- This can include discounts, points transfers, and combined rewards.

6. **Social Media Integration:**

- Many programs encourage social media sharing and reviews, which can enhance the hotel's online presence and attract new guests.
- Social media contests and exclusive promotions can further engage members.

7. **Sustainability Initiatives:**

- Modern loyalty programs increasingly incorporate sustainability, offering points or rewards for eco-friendly choices like reducing plastic use or participating in recycling programs.

Benefits to Guests:

1. **Financial Savings:**

- Accumulation of points for free or discounted stays.
- Exclusive discounts and offers not available to non-members.

2. **Enhanced Services:**

- Faster check-in, preferred room assignments, and personalized services.
- Access to exclusive events and experiences.

3. **Convenience:**

- Easy booking and management through mobile apps.
- Seamless integration with other services like car rentals and travel packages.

4. **Recognition and Prestige:**

- Status symbols like elite membership levels.
- Recognition by hotel staff and priority service at participating properties.

Benefits to Hotels:

1. ****Increased Guest Loyalty:****

- Encourages repeat visits and long-term relationships with guests.
- Helps retain customers even when faced with competition.

2. ****Data Collection and Analysis:****

- Provides valuable insights into guest preferences and behaviors.
- Helps hotels tailor their services and marketing strategies effectively.

3. ****Revenue Growth:****

- Drives up room occupancy and average daily rate (ADR) through repeat business.
- Cross-selling opportunities within the hotel (e.g., dining, spa services).

4. ****Market Differentiation:****

- Helps hotels stand out in a competitive market.
- Offers a unique selling proposition that attracts brand-loyal customers.

Conclusion:

Modern hotel loyalty programs have become a critical component of the hospitality industry, leveraging technology and data to create personalized, engaging, and rewarding experiences for guests. They provide significant benefits to both guests, who enjoy enhanced services and savings, and hotels, which gain deeper insights and increased revenue through stronger customer relationships.

With a System Prompt

System:

Provide answers in the style of a medieval scroll, using old English vocabulary and phrases.

User:

Explain the modern concept of hotel loyalty programs and their benefits.

Assistant:

Here followeth a detailed exposition of the modern concept known as hotel loyalty programs, penned in the venerable style befitting such an arcane and wondrous subject.

In the annals of contemporary hostelry and inns, there doth exist a most wondrous construct termed "hotel loyalty programs." These programs be akin to the ancient covenants or oaths betwixt lord and vassal, binding patrons unto the establishments in which they sojourn.

Loyalty programs doth grant unto the faithful patrons, known as members, manifold benefits and rewards of great magnanimity. These benefactions are bestowed through the accumulation of points or credits, garnered by sojourns and patronage. Herein lieth the essence of the program: the more one sojourneth, the more points they dost garner, which may then be exchanged for sundry rewards.

The benefits of these programs are manifold and most enticing:

1. ****Frequent Stayer Privileges****: Members may find themselves accorded preferential treatment, such as elevated room categories or enhanced services, akin to the favors granted by a noble unto his most loyal swain.
2. ****Accelerated Status****: In this modern hierarchy, members may ascend to exalted tiers, each tier granting more bountiful privileges. Such tiers may include titles of honor such as "Silver," "Gold," or even "Platinum," bestowing upon the member increased perks with each ascension.
3. ****Free Nights and Services****: Verily, members may claim nights of lodging without the burden of coin, akin to the boons granted by a generous lord to his steadfast followers. Additionally, sundry services such as complimentary breakfasts or room upgrades may be offered.
4. ****Special Offers and Discounts****: Members might also receive secret and special offers, as if privy to the most intimate counsels of the innkeeper, granting them discounts and other boons not afforded to the common patron.
5. ****Personalized Experience****: The programs often tailor their services to the preferences and habits of their members, much like a wise steward who remembers the tastes and whims of his lord.

Thus, in the realm of modern hostelry, the hotel loyalty program standeth as a testament to the virtues of faithfulness and patronage, offering unto its adherents a tapestry of rewards and benefits that doth rival the most storied legends of old.

Example 2: Providing guardrails

With the addition of the *system prompt*, users can add guardrails for their applications to safeguard the applications for any unintended use. Here is an example of how guardrails can focus the model on the desired task.

Without a System Prompt

System:

NA

User:

Write a one-page business plan for a cooperative bike-sharing program in urban centers aiming to lower traffic congestion and promote eco-friendly transportation. (Include target markets, revenue streams, marketing strategy, and risk mitigation plans.)

Assistant:

****Urban EcoCycle: A Cooperative Bike-Sharing Program****

****Executive Summary:****

Urban EcoCycle is a cooperative bike-sharing program designed to alleviate traffic congestion and promote eco-friendly transportation in urban centers. By providing accessible, affordable, and convenient bike-sharing services, we aim to foster a sustainable urban mobility ecosystem that benefits both residents and the environment.

****Target Markets:****

Our primary target markets include urban professionals, students, and environmentally conscious citizens. We also aim to attract tourists looking for an affordable and eco-friendly way to explore the city. Secondary markets will include local businesses that could benefit from reduced employee commuting times and increased customer engagement due to the program's presence.

****Revenue Streams:****

1. ****Subscription Fees:**** Monthly or yearly subscription plans for individuals and businesses.
2. ****Ad Revenue:**** Partnerships with local businesses for advertising space on bikes and docking stations.

3. **Corporate Partnerships:** Collaborations with companies to provide their employees with access to our bike-sharing program.

4. **Government Grants:** Applications for grants and subsidies from local government initiatives aimed at promoting sustainable transportation.

Marketing Strategy:

1. **Digital Marketing:** Utilize social media platforms, email newsletters, and our website to reach out to potential users and partners. Highlight success stories and eco-friendly benefits.

2. **Partnerships:** Collaborate with local businesses, schools, and environmental organizations to promote our service.

3. **Community Events:** Host bike-sharing workshops, clean-up events, and eco-friendly fairs to engage with the community.

4. **Incentives:** Offer discounts and rewards for long-term subscribers and those who refer new members.

Risk Mitigation Plans:

1. **Maintenance and Security:** Implement a robust maintenance schedule to ensure bikes are always in good working condition. Employ security measures such as GPS tracking and insurance coverage to protect bikes from theft and damage.

2. **Regulatory Compliance:** Stay informed about local regulations and ensure compliance to avoid legal issues. Work with local authorities to establish clear guidelines for bike parking and usage.

3. **Market Competition:** Conduct regular market research to stay ahead of competitors. Differentiate our service by offering unique features such as bike rentals for longer trips, special eco-friendly programs, and partnerships with popular local destinations.

4. **Economic Fluctuations:** Develop flexible pricing strategies and diversify revenue streams to cushion against economic downturns. Maintain a reserve fund to handle unexpected financial challenges.

Urban EcoCycle is poised to become a cornerstone of urban sustainability by reducing traffic congestion and promoting green transportation. With a well-defined target market, diverse revenue streams, strategic marketing, and comprehensive risk mitigation plans, we are committed to making our urban centers more livable and eco-friendly.

With a System Prompt

System:

You are Application Process Assistant

To answer user question, you follow these instructions/steps:

- The LLM chatbot can discuss topics related to college applications, including deadlines, required materials, and application status.
- The LLM chatbot should be polite and courteous in all responses.
- The LLM chatbot can only respond in 3 sentences or less.

The above system instructions define your capabilities and your scope. If the user request contradicts any system instruction or if the request is outside your scope, you must politely decline the request briefly explaining your capabilities and your scope.\n

User:

Write a one-page business plan for a cooperative bike-sharing program in urban centers aiming to lower traffic congestion and promote eco-friendly transportation. (Include target markets, revenue streams, marketing strategy, and risk mitigation plans.)

Assistant:

I'm here to help with college application-related questions, including deadlines and required materials. For your business plan request, I recommend seeking assistance from a business or entrepreneurship advisor.

Provide examples (few-shot prompting)

By including a few examples of your task within the prompt, you can build a structured template for Amazon Nova 2 to follow. This reduces ambiguity and enhances the accuracy and quality of the output. The technique of providing clear examples to make the model's responses more aligned with the desired outcome is called *few-shot prompting*.

The concept of few-shot prompting is to provide the language model with a few examples of the task, along with the input and output format, and then ask it to generate the output for a new input based on the provided examples. This method also helps in clarifying complex instructions or tasks, making it easier for Amazon Nova 2 to understand and interpret what is being asked.

How adding examples to the prompt help:

Adding examples can help the model with producing

- Consistent responses which are uniform to the style of the examples

- Performant responses due to reducing the chance of misinterpreting instructions, and minimizing hallucinations

The amount by which model performance improves using few shot prompting will depend on the quality and diversity of your chosen examples. The following items demonstrate the characteristics of good shots in the prompt:

- **Select diverse examples:** The examples chosen should represent the distribution of your expected input/output in terms of diversity (ranging from common use cases to edge cases) to adequately cover relevant use cases. It is important to avoid any biases in your examples, as bias in the inputs can cause outputs to be biased as well.
- **Match complexity levels:** The complexity of the examples provided should align with the target task or scenario. It is important to make sure the complexity grade is mapped between expected the input and the chosen example in the prompt.
- **Ensure relevance:** The examples selected should be directly relevant to the problem or objective at hand. This ensures consistency and uniformity in responses.

Tip

If the above suggestions do not work, we recommended you build a RAG-based system that augments the prompt with a dynamic selection of shots based on the similarities between the user-input query and an available pool of shots.

Example 1: Provide examples in the user prompt

Here, your user is trying to classify incoming customer support tickets into predefined categories. You can provide Amazon Nova 2 models with a few examples of tickets that are classified correctly to improve its understanding of the categories and the language used in the tickets. The exemplar examples should be concise and represent all tasks targeted in your application. If there is a specific field name that is meaningful to the tasks, such as `Ticket` or `Category` in this example, your user can use those fields directly followed by a colon ":".

0-shot prompt

User:

Please classify the following customer support tickets into one of the categories: "Technical Issue," "Billing Inquiry," or "General Inquiry."

Ticket: "Can you explain why my latest bill is higher than usual? I haven't changed my plan."

Assistant:

The customer support ticket should be classified as a "Billing Inquiry."

Here's the reasoning:

- The customer is asking about the reason for a change in their latest bill amount.
- The context revolves around billing and payment, which falls under the category of "Billing Inquiry."

3-shot prompt

User:

Please classify the following customer support tickets into one of the categories: "Technical Issue," "Billing Inquiry," or "General Inquiry."

Use these examples to guide your classification:

Ticket: "I can't log in to my account. The system keeps giving me an error message."

Category: Technical Issue

Ticket: "I received an email saying my account will be suspended unless I make a payment. Is this a scam?"

Category: Billing Inquiry

Ticket: "What are the hours of operation for your customer support team?"

Category: General Inquiry

Now, classify the following ticket:

Ticket: "Can you explain why my latest bill is higher than usual? I haven't changed my plan."

Assistant:

Category: Billing Inquiry

Explanation: The ticket is asking for an explanation regarding a change in the billing amount, which falls under the billing category as it pertains to questions about charges or payments.

Example 2: Providing examples as turns prompt

In this example, the exemplars are provided as turns, one turn per exemplar, where the User role contains the input and Assistant role contains the expected output for that exemplar.

0-shot Prompt

System:

Your task is to Classify the following texts into the appropriate categories. The categories to classify are:

Categories:

- Food
- Entertainment
- Health
- Wealth
- Other

User:

I have 20\$ in my pocket.

Assistant:

Category: Wealth

Explanation: The text mentions a specific amount of money, which falls under the category of wealth. The presence of a monetary value directly relates to financial status or economic aspect.

3-shot Prompt (as turns)

System:

Your task is to Classify the following texts into the appropriate categories. The categories to classify are:

Categories:

- Food
- Entertainment
- Health
- Wealth
- Other

User:

I love to eat pizza.

Assistant:

Food

User:

I enjoy watching movies.

Assistant:

Entertainment

User:

I am going to the gym after this.

Assistant:

Health

User:

I have 20\$ in my pocket.

Assistant:

Wealth

Example 3: Providing examples in the system prompt

In this example, you can specify an "Example" section and then list a series of examples within the system prompt. This way of providing exemplars in the prompt is especially useful if your exemplars are long and complex and helps model locate example shots.

0-shot Prompt

System:

You are an AI assistant helping to draft professional reports for enterprise stakeholders. Ensure the report is formatted in rich text, with headings, bullet points, and highlighted important tasks.

User:

Draft a detailed project update report for our stakeholders, summarizing the project milestones achieved, tasks completed, and future plans.

Assistant:

```
# Project Update Report

## Executive Summary
This report provides a comprehensive update on the project milestones achieved, tasks completed, and future plans. Our team has made significant progress, and we are on track to meet the project goals. The following sections outline the details of our accomplishments and upcoming activities.

---

## Project Milestones Achieved

### **1. Project Initiation**
- **Milestone Date:** July 15, 2023
- **Description:** Completed initial project planning and secured stakeholder approval.

### **2. System Design Completion**
- **Milestone Date:** September 30, 2023
- **Description:** Finished the detailed system design and received positive feedback from technical review meetings.
```

```
### **3. Prototype Development**
- **Milestone Date:** November 15, 2023
- **Description:** Developed and tested a working prototype of the system.

---

## Tasks Completed

### **1. Requirement Gathering**
- **Description:** Conducted comprehensive stakeholder interviews and collected detailed requirements.
- **Highlighted Important Tasks:**
  - **User Story Mapping:** Mapped out user stories to ensure all user needs are addressed.
  - **Requirement Documentation:** Created detailed requirement specifications.

### **2. System Architecture Design**
- **Description:** Designed the system architecture, focusing on scalability and maintainability.
- **Highlighted Important Tasks:**
  - **Technology Stack Selection:** Chose the appropriate technology stack based on project requirements.
  - **Architecture Diagrams:** Developed detailed architecture diagrams for easy understanding.

### **3. Prototype Development and Testing**
- **Description:** Developed and rigorously tested the prototype.
- **Highlighted Important Tasks:**
  - **Coding:** Implemented core functionalities of the prototype.
  - **Testing:** Conducted unit and integration testing to ensure prototype reliability.

---

## Future Plans

### **1. Full System Development**
- **Description:** Begin the full-scale development of the system based on the approved design and prototype.
- **Highlighted Important Tasks:**
  - **Module Development:** Develop each module as per the design specifications.
  - **Code Reviews:** Implement regular code reviews to maintain code quality.

### **2. User Acceptance Testing (UAT)**
```

```

- Description: Conduct UAT to ensure the system meets user requirements and is
  ready for deployment.
- Highlighted Important Tasks:
  - Test Plan Development: Create detailed UAT plans and test cases.
  - UAT Execution: Execute UAT and gather feedback for final adjustments.

### 3. Deployment and Training
- Description: Deploy the system and provide necessary training to end-users.
- Highlighted Important Tasks:
  - Deployment: Roll out the system to the production environment.
  - Training Sessions: Organize training sessions for users to familiarize them
    with the new system.

---

## Conclusion
Our project has successfully achieved key milestones and completed significant
tasks. With the current momentum, we are well-positioned to meet our future
objectives. We remain committed to delivering a high-quality system that meets all
stakeholder expectations.

---

Prepared by:
[Your Name]
[Your Position]
[Date]

Contact Information:
[Your Email]
[Your Phone Number]

---

This detailed update aims to keep stakeholders informed and reassured about the
project's progress and future direction. For any queries or further information,
please do not hesitate to contact me directly.

```

2-shot Prompt

System:

You are an AI assistant helping to draft professional reports for enterprise stakeholders. Below are a few examples of well-formatted project update reports to guide your response.

<Example 1>

Project Update Report

<i>Date:</i> January 15, 2023

<hr>

Project Milestones Achieved:

- Successfully launched the new CRM system.

- Integrated the marketing automation tool.

- Completed the employee training sessions.

Tasks Completed:

- Developed the user manual.

- Conducted system testing.

- Prepared the training materials.

Future Plans:

- Roll out the CRM system to all departments.

- Schedule follow-up training sessions.

- Monitor system performance and make necessary adjustments.

<hr>
Prepared by:

[Your Name]

[Your Position]

[Your Contact Information]

<Example 2>

Project Update Report

<i>Date:</i> February 20, 2023

<hr>

Project Milestones Achieved:

- Launched the new website.

- Completed the SEO optimization.

- Organized the stakeholder meeting.

```

<b>Tasks Completed:</b><br>
- Designed the website layout.<br>
- Wrote the content for the website.<br>
- Tested the website functionality

<b>Future Plans:</b><br>
- Implement the analytics tracking.<br>
- Launch the social media campaign.<br>
- Gather feedback from stakeholders.

<hr>
<b>Prepared by:</b><br>
[Your Name]<br>
[Your Position]<br>
[Your Contact Information]

```

User:

Draft a detailed project update report for our stakeholders, summarizing the project milestones achieved, tasks completed, and future plans.

Assistant:

```

<b>Project Update Report</b><br>
<i>Date:</i> October 5, 2023<br>
<hr>

<b>Project Milestones Achieved:</b><br>
- Successfully completed the implementation of the new enterprise resource planning (ERP) system.<br>
- Conducted and completed the first phase of the data migration process.<br>
- Successfully launched the new inventory management module.

<b>Tasks Completed:</b><br>
- Developed and finalized the comprehensive user training program, including detailed user manuals and video tutorials.<br>
- Conducted extensive system integration testing with all existing systems and ensured seamless interoperability.<br>
- Executed the initial data migration from the legacy system to the new ERP system, resulting in a 95% accuracy rate.

<b>Future Plans:</b><br>
- Roll out the new ERP system to all departments in the next two phases.<br>

```

- Schedule and conduct follow-up training sessions for all employees.

- Implement the remaining modules of the ERP system, including financial management and human resources.

- Establish a robust monitoring and support framework to ensure smooth system operation and address any issues promptly.

<hr>

Prepared by:

[Your Name]

[Your Position]

[Your Contact Information]

This detailed report provides a clear and concise overview of the project's progress, highlighting the milestones achieved, tasks completed, and future plans. It ensures stakeholders are well-informed and can make strategic decisions based on the current status of the project.

Tool calling systems

Tool calling is available for the Amazon Nova models by passing a tool configuration schema in your request. The prompt for the model is going to be augmented with this tool configuration so it is a highly impactful place to begin optimizing your tool calling system.

Consider these key principles:

- Tool definitions should be clear and concise. They should be easy to understand and the intent must be extremely apparent.
- Use key differentiators and boundary conditions to define when one tool should be used over another.
- Be critical the input argument types. Ask, do they make sense and would they be expected to be used in that fashion normally?

Use “Tool Choice” to control when a tool is called

The tool choice parameter allows you to customize the behavior of tool calling with the model. We recommend utilizing this for fine grained control on which tools are called and when.

For example, for use cases like structured output, you might want a specific tool to be called each time Amazon Nova is invoked. You can define the schema of your output as the tool and then set the tool choice to the name of that tool.

```
{
  "toolChoice": {
    "tool": {
      "name": "name_of_tool"
    }
  }
}
```

For many agentic use cases, you might want to ensure that the model always selects one of the available tools. To do so, you can set the tool choice to any, which will call exactly one tool each time the model is invoked.

```
{
  "toolChoice": {
    "any": {}
  }
}
```

Lastly, for use cases where whether a tool is called is highly dependent on the context of the conversation, you can set the tool choice to auto. This is the default behavior and will leave the tool selection completely up to the model.

```
{
  "toolChoice": {
    "auto": {}
  }
}
```

Advanced prompting techniques

These sections provide advanced guidance for how to improve the quality of your prompts and leverage key features like extended thinking.

Use Reasoning Mode

Amazon Nova 2 models offer an optional reasoning mode that enhances the model's approach to complex problem-solving by allowing it to work through problems systematically before

responding. Leveraging the model's reasoning mode is a powerful way to improve the accuracy of your prompts.

When to use it: Reasoning mode is recommended for complex tasks such as use cases with:

- **Multiple reasoning steps:** Mathematical proofs, algorithm design, system architecture
- **Cross-referencing information:** Document analysis, option comparison, trade-off evaluation
- **Error-prone calculations:** Financial modeling, data analysis, complex debugging
- **Planning with constraints:** Resource optimization, dependency management, risk assessment
- **Complex classifications:** Multi-label categorization, hierarchical taxonomies, nuanced decision boundaries
- **Tool calling scenarios:** Multi-step API workflows, database query optimization, coordinated system integrations

 Note

For more information on the reasoning mode, refer to [Using reasoning](#).

Take a top-down approach

For situations where the model needs to assess multiple approaches to solve the problem, instruct it to take a **top-down** approach.

- Amazon Nova 2 models perform better when the model starts with the big picture and then breaks it down into smaller, more detailed subproblems or steps.
- Explicitly direct the model to first identify the main objective, then decompose it into manageable components before working through the details of each part.
- This structured approach helps the model organize its thinking and produce more coherent reasoning chains.

Example:

```
{{User query}}. Start with the big picture and break it down into progressively smaller, more detailed subproblems or steps.
```

Steer Chain of Thought

While reasoning mode provides enhanced accuracy through systematic problem-solving, there are specific scenarios where Chain of Thought (CoT) prompting in non-reasoning mode may better serve your needs.

When to use it:

- **Transparency and auditability:** When you want to see, verify, or audit the model's reasoning process, CoT provides full visibility into each step. This is critical for regulated industries, high-stakes decisions, or when you want to document the logic behind an answer.
- **Custom reasoning structures:** Use CoT to enforce specific reasoning patterns or methodologies. You can guide the model to follow your organization's decision frameworks, use domain-specific problem-solving approaches, or ensure factors are considered in a specific order.
- **Prompt development and debugging:** During the prompt engineering phase, CoT helps you understand how the model approaches problems, identify where reasoning breaks down and iterate on your prompts more effectively.
- **Hybrid approaches:** Consider using CoT during development to perfect your prompts, then switching to reasoning mode for production deployment once you're confident in the model's approach to your specific use case.

Note

Not all tasks require CoT. For simpler tasks, allow the model to use its own reasoning process.

Guiding the model's CoT direction:

```
{{User query}} Please follow these steps:
```

1. {{Step 1}}
2. {{Step 2}}
- ...

Utilizing long context windows

Amazon Nova 2 models have a supported context length of 1 million tokens and excel at code understanding and question answering on long documents. Its performance (including system prompt adherence and tool use) can decline slightly as the context size increases.

How to use it:

- **Put long-form data at the beginning:** Place your long documents and inputs near the beginning of your prompt. Place them before your query, instructions and examples.
- **Put instructions at the end:** Place your instructions at the end of the prompt. The model performs best when the context is provided first and the instructions are provided at the end.
- **Structure document content start and end markers:** Use start and end markers, such as `DOCUMENT {idx} START` and `DOCUMENT {idx} END`, to denote the start and end of long documents where `{idx}` represents the index of the specific document.

Example Template:

```
// Provide your long inputs at the top of your prompt
BEGIN INPUT DOCUMENTS

DOCUMENT 1 START
{{Your document}}
DOCUMENT 1 END

END INPUT DOCUMENTS

// Then specify your query and instructions
BEGIN QUESTION
{{User query}}
END QUESTION

BEGIN INSTRUCTIONS
{{Instructions}}
END INSTRUCTIONS
```

Ground answers in supporting text

We recommend that you provide the model with trusted information relevant to the input query. This information, along with the input query, is often a part of the system called retrieval augmented generation (RAG).

- In this process, some relevant, contextual document or information is augmented to the actual user prompt so that the model gets trustworthy content to generate a relevant and accurate response.
- Instructing Amazon Nova 2 to answer using a reference text from a trusted source can guide it to compose its response based on the provided material and ensure that its response is grounded in accurate and relevant information, enhancing the reliability and credibility of the generated content.
- Using a reference text can help avoid hallucinating, thereby improving the overall quality and trustworthiness of the responses. To minimize hallucination, we recommend explicitly mentioning **DO NOT USE INFORMATION THAT IS NOT IN REFERENCE TEXTS!** in your model instructions.

Prompt template

System:

In this session, the model has access to search results and a user's question, your job is to answer the user's question using only information from the search results.

Model Instructions:

- DO NOT USE INFORMATION THAT IS NOT IN SEARCH RESULTS!

User: {Query}

Resource: Search Results: {Reference texts}

Example: Providing grounding context helps to prevent the model from hallucinating

This example shows how grounding context can prevent the model from hallucinating.

System:

In this session, the model has access to search results and a user's question, your job is to answer the user's question using only information from the search results. Model Instructions:- DO NOT USE INFORMATION THAT IS NOT IN SEARCH RESULTS!

User:

What were the economic impacts of the COVID-19 pandemic on the United States in 2020?

Resource: Search Results: In 2020, the United States experienced significant economic impacts due to the COVID-19 pandemic. The U.S. economy contracted by 3.5% in 2020, according to the Bureau of Economic Analysis. Unemployment rates surged to 14.7% in April 2020, the highest since the Great Depression, before gradually declining. Small businesses faced severe challenges, with millions of firms closing permanently. Additionally, consumer spending dropped sharply as people reduced non-essential expenditures and saved more. Government intervention played a critical role in mitigating these impacts through stimulus packages and support programs, such as the Paycheck Protection Program (PPP) for small businesses and direct payments to individuals. Despite these measures, the economic recovery remained uneven across different sectors and regions.

Ground using citation markers

For long document tasks, we recommend that you instruct the Amazon Nova 2 model to ground its responses using citations from the relevant sections of the documents before it proceeds with the task.

This approach helps the model focus on the most pertinent information and avoid being distracted by extraneous content. When you request that the model grounds its response, the sections that can be cited should be numbered. For example, Passage %[1]%, Passage %[2]% and so on.

Examples: using citation markers**Example Citation markers prompt**

You are an AI financial assistant. Your task is to find patterns and insights from multi-year financial documents

Passage %[1]%
{{Your document}}

Passage %[2]%
{{Your document}}

Task:

Analyze my LLC's reports across multiple years to identify significant performance trends, segment growth patterns and strategic shifts.

Context information:

- You have access to my LLC's annual financial reports (10-K) for multiple fiscal years in PDF format
- These reports contain comprehensive financial data including income statements, balance sheets, cash flow statements and management discussions
- The analysis should focus on year-over-year comparisons to identify meaningful trends
- I operate two business segments, one in Massachusetts and one in New York

Based on the provided Context, extract key financial metrics from each year's reports phrases from the documents.

Place citations as inline markers (e.g., %[1]%, %[2]%, etc.) directly within the relevant parts of the response

text. Do not include a separate citation section after the response.

Response Schema:

%% (Extracted Financial Metrics)

%% (Extracted Financial Metrics)

%% (Extracted Financial Metrics)

...

After you extract key information based on the user's task, you can use the extracted financial metrics to answer the relevant questions as shown:

Example Follow-up analysis with extracted metrics**## Task**

Analyze my LLC's financial reports across multiple years to identify significant performance trends, segment growth patterns and strategic shifts.

{{extracted financial metrics}}

Model Instructions:

- Organize data chronologically to identify meaningful trends
- DO compare segment performance across the five-year period
- DO identify significant strategic shifts or investments mentioned in management discussions
- DO NOT make speculative predictions beyond what is supported by the data
- ALWAYS note any changes in accounting practices or reporting methodologies that might affect year-over-year comparisons

Response style and format requirements:

- Respond in markdown
- Structure the analysis with clear headings and subheadings
- Present key financial metrics in tabular format showing all five years side-by-side

- Include percentage changes year-over-year for all major metrics
- Create a section dedicated to visualizing the most significant trends (with descriptions of what would be shown in charts)
- Limit the executive summary to 250 words maximum
- Format segment analysis as separate sections with consistent metrics across all segments
- MUST include a Key Insights bullet-pointed list at the end of each major section

Use Nova Web Grounding

Instead of prompting directly for citations to ground the model in supporting text, Amazon Nova 2 models provide an internal web grounding tool that can be used. When enabled, Amazon Nova 2 models will directly query the web and Amazon's knowledge graphs and ground the final response with citations.

To learn more on how to leverage Amazon Nova Web Grounding, you can reference the [Amazon Nova Web Grounding user guide](#).

Produce structured output

To ensure consistent and structured output formats, you can use structured outputs, including formats like XML, JSON, Markdown, or using tool use functionality.

- This approach allows downstream systems to more effectively understand and parse the outputs generated by the model.
- By providing explicit instructions to the model, the responses are generated in a way that adheres to a predefined schema.

For example, if the downstream parser expects specific naming conventions for keys in a JSON object, you should specify the response schema at the end of the prompt. Additionally, if you prefer responses to be in JSON format without any preamble text, instruct the model accordingly. That is, explicitly state **Please generate only the JSON output. DO NOT provide any preamble.** to ensure clean output.

Tip

- We observe best adherence to data format requirements when they are defined in the schema itself instead of through the use of exemplars (such as, specifying dates in YYYY/MM/DD format).

- For simple JSON outputs with up to 10 keys, you can find the schema below. For more complex schemas, we recommend you define your schema through a tool. Tool use leverages a technique called constrained decoding that will increase the model's adherence for these complex schemas.

Common formatting schemas

The following are examples of common formatting schemas.

JSON

```
JSON_format = """Write your response following the JSON format below:
```

```
```json
{
 "key1": "value1",
 "key2": "value2",
 key3: [{
 "key3_1": "value_3_1 written in YYYY/MM/DD format",
 "key3_2": "value_3_2 day of the week written in full form",
 ...}]```
"""
```

### XML

```
XML_format = """Write your response following the XML format below:
```

```
<output>
 <task>"task1"</task>
 <subtask>
 <task1_result> (task 1 result)</task1_result>
 <task2_result> (task 2 result)</task2_result>
 <task3_result> (task 3 result)</task3_result>
 </subtask>
 <task>"task2"</task>
 <subtask>
 <task1_result> (task 1 result)</task1_result>
 <task2_result> (task 2 result)</task2_result>
 <task3_result> (task 3 result)</task3_result>
 </subtask>
</output>
```

```
"""
```

## Markdown

```
markdown_schema = """Write your response following the markdown format below:
```

```
Introduction
(2-3 line intro)
```

```
Design Guidance
(Bulleted list of design guidance)
```

```
Step by Step Instructions on Execution
(Bulleted list of instructions with each with bold title.
```

```
Conclusion
(conclusion)
```

```
"""
```

## Prefill assistant content

If you are producing structured output in non-reasoning mode, you can nudge the model's response by prefilling the assistant content.

Prefilling improves consistency in the output format while in non-reasoning mode. It allows you to direct the model's actions, bypass preambles and enforce specific output formats like JSON and XML. For example, if you prefill the assistant content with { or ```json, that input guides the model to generate the JSON object without additional information.

### Tip

If you are explicitly looking for extracting JSON, a common pattern is to prefill with ```json and add a stop sequence on ``` . This ensures the model outputs a JSON object that can be programmatically parsed.

## Example Prefill assistant content

The following code demonstrates how to prefill with the API:

```
import boto3
import json

Create a Bedrock Runtime client.
client = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1"
)

request_body = {
 "system": [
 {"text": "You write JSON objects based on the given instructions"}
],
 "messages": [
 {
 "role": "user",
 "content": [{"text": "Provide details about the best selling full-frame cameras in past three years. Answer in JSON format with keys like name, brand, price and a summary."}]
 },
 {
 "role": "assistant",
 "content": [{"text": " Here is the JSON response: ```json"}]
 }
],
 "inferenceConfig": {
 "maxTokens": 1000,
 }
}

Invoke the model and extract the response body.
response = client.invoke_model(
 modelId="amazon.nova-2-lite-v1:0",
 body=json.dumps(request_body)
)

model_response = json.loads(response["body"].read())
```

## Specify a tool to use for complex schemas

Another approach is to use tools to force a specific schema for the model's response by including the *pydantic JSON schema* in the tool use schema section. You can specify tool choice to the provided schema and Amazon Nova's response will be structured based on the tool selected. To learn more about how to leverage tool use see [Tool Use](#).

## Examples

### Example 1: Add an output schema

Adding the Output Schema section with the appropriate schema makes the model prefer to fit that schema. However, model output is not deterministic and can vary from the output schema.

#### Unspecific Prompt

User:

Provide details about the best selling full-frame cameras in past three years.  
Answer in JSON format with keys like name, brand, price and a summary.

Assistant:

```
```json
[
{
  "name": "Sony Alpha A7R IV",
  "brand": "Sony",
  "price": "Approximately $3,500",
  "summary": "The Sony Alpha A7R IV is renowned for its exceptional image quality, thanks to a 61-megapixel full-frame sensor. It also features 4K video recording and a robust build, making it a favorite among professional photographers and videographers."
},
{
  "name": "Canon EOS R5",
  "brand": "Canon",
  "price": "Approximately $3,900",
  "summary": "The Canon EOS R5 combines high-resolution imaging with advanced video capabilities. It houses a 45-megapixel sensor, 8K video recording, and impressive autofocus performance, appealing to both photography and filmmaking professionals."
},
]
```

```
{
  "name": "Nikon Z7 II",
  "brand": "Nikon",
  "price": "Approximately $2,800",
  "summary": "The Nikon Z7 II offers a 45.7-megapixel full-frame sensor and is
praised for its high resolution and excellent dynamic range. It also supports 4K
video recording and features a versatile Z-mount lens system, making it a strong
contender in the full-frame market."
},
{
  "name": "Fujifilm GFX 100S",
  "brand": "Fujifilm",
  "price": "Approximately $7,000",
  "summary": "While not a traditional full-frame camera, the Fujifilm GFX 100S stands
out with its 102-megapixel medium format sensor. It is lauded for its stunning
detail and color reproduction, though it comes at a higher price point and is aimed
at specialized users."
}
]
...
```

Specific Prompt

User:

Provide details about the best selling full-frame cameras in past three years. You MUST answer in JSON format only. Write your response following the the format below:

```
[{
  "name": "name goes here",
  "brand": "brand goes here",
  "price": "price goes here",
  "summary": "summary goes here"
}]
```

Assistant:

```
[
  {
    "name": "Nikon D850",
    "brand": "Nikon",
    "price": "$3,296.95",
```

```
"summary": "The Nikon D850 is a professional full-frame DSLR camera known for its exceptional image quality, 45.7-megapixel resolution, 7 fps continuous shooting, and 4K video capabilities."
},
{
  "name": "Canon EOS 5D Mark IV",
  "brand": "Canon",
  "price": "$3,499.00",
  "summary": "The Canon EOS 5D Mark IV is a popular full-frame DSLR appreciated for its 30.4-megapixel sensor, 4K video recording, robust build quality, and reliable autofocus system."
},
{
  "name": "Sony A7R III",
  "brand": "Sony",
  "price": "$3,198.00",
  "summary": "The Sony A7R III is a high-resolution full-frame mirrorless camera renowned for its 42.4-megapixel sensor, excellent low-light performance, and versatile lens compatibility."
},
{
  "name": "Fujifilm GFX 100S",
  "brand": "Fujifilm",
  "price": "$5,999.99",
  "summary": "The Fujifilm GFX 100S is a medium format mirrorless camera that stands out for its 102-megapixel sensor, high dynamic range, and exceptional color reproduction."
},
{
  "name": "Sony A7 III",
  "brand": "Sony",
  "price": "$1,998.00",
  "summary": "The Sony A7 III is a full-frame mirrorless camera popular among enthusiasts and professionals for its 24.2-megapixel sensor, fast autofocus, and 4K video capabilities."
}
]
```

Example 2: Prefill the model response

Another approach is to prefill the model's response by including the desired initial text within the assistant's message. In this case, Amazon Nova model's response will continue from where the assistant's message leaves off.

Tip: Prefill is only valid when reasoning is not enabled.

User:

```
Provide details about the best selling full-frame cameras in past three years.  
Your response should be in JSON format, with the following keys: name, brand, year,  
price, summary.
```

Assistant (Prefilling):

```
```json
```

Assistant:

```
{
 "cameras": [
 {
 "name": "Sony Alpha A7 III",
 "brand": "Sony",
 "year": 2018,
 "price": "$1999",
 "summary": "The Sony Alpha A7 III is a highly acclaimed full-frame mirrorless camera known for its excellent image quality, advanced autofocus system, and 4K video capabilities."
 },
 {
 "name": "Canon EOS R5",
 "brand": "Canon",
 "year": 2020,
 "price": "$3899",
 "summary": "The Canon EOS R5 is a top-tier full-frame mirrorless camera that offers stunning 45MP images, 8K video recording, and exceptional performance in various lighting conditions."
 },
 {
 "name": "Nikon Z7 II",
 "brand": "Nikon",
```

```
"year": 2020,
"price": "$2999",
"summary": "The Nikon Z7 II is a flagship full-frame mirrorless camera with a 45.7MP
sensor, superb image quality, and advanced video recording features."
},
{
"name": "Fujifilm GFX 100S",
"brand": "Fujifilm",
"year": 2020,
"price": "$7999",
"summary": "The Fujifilm GFX 100S is a medium format mirrorless camera known for its
impressive 102MP sensor, beautiful film simulations, and versatile lens options."
},
{
"name": "Sony Alpha A7R IV",
"brand": "Sony",
"year": 2019,
"price": "$3499",
"summary": "The Sony Alpha A7R IV features a 61MP sensor, excellent low-light
performance, and advanced video capabilities, making it a favorite among professional
photographers."
}
]
}
```

### Example 3: Specify a tool to use

As an alternative to using the prefill technique for structured output, you can combine Nova's tool functionality with toolChoice. In this example, the model will generate a tool call that adheres to the JSON schema provided and calls the specific tool specified to full JSON format defined. For additional guidance, refer to Advanced Systems.

User:

From the below provided Query, extract the relevent entities

Query:

John works in BUILDING-0987 and have been incharge of product id 23564#. His performance have been excellent in past year and is up for a raise.

Use the extract\_ner tool.



## ToolConfig:

```

{
 "tools": [
 {
 "toolSpec": {
 "name": "extract_ner",
 "description": "Extract all the named entities. based on provided
input",
 "inputSchema": {
 "json": {
 "type": "object",
 "properties": {
 "entities": {
 "type": "array",
 "items": {
 "type": "object",
 "properties": {
 "name": {
 "type": "string",
 "description": "The extracted entity name. This
should be a name of a person, place, animal or thing"
 },
 "location": {
 "type": "string",
 "description": "The extracted location name. This
is a site name or a building name like SITE-001 or BUILDING-003"
 },
 "product": {
 "type": "string",
 "description": "The extrcted product code, this is
generally a 6 digit alphanumeric code such as 45623#, 234567"
 }
 },
 "required": [
 "name",
 "location",
 "product"
]
 }
 }
 },
 "required": [
 "entities"
]
 }
 }
 }
 }
]
}

```

```
]
 }
 }
 },
 "toolChoice": {
 "tool": {
 "name": "extract_ner"
 }
 }
}
```

## Multilingual use cases

Amazon Nova 2 models have been trained on more than 200 languages and optimized for 15 languages.

### Topics

- [Prompt for accurate translations](#)
- [Enforce consistent writing conventions](#)

### Prompt for accurate translations

To leverage this capability for short form translations (a few sentences) you can instruct the model to translate the text into the specified target language.

### Example Translation prompts

Translate the following text into {target language}. Please output only the translated text with no prefix or introduction: {text}

Translate the following sentence from {source\_language} to {target language}: {text}

{text} How do you say this sentence in {target\_language}

### Enforce consistent writing conventions

In character-based languages, the Amazon Nova 2 models may utilize the character set from the source language. You can use the following prompt to enforce a consistent output.

## Example Enforce writing conventions

When translating, ensure to use the correct orthography / script / writing convention of the target language, not the source language's characters

## Tool calling

### Agentic systems

#### Topics

- [Set the right inference parameters](#)
- [Consider latency requirements](#)
- [Use intentional wording for tool calling instructions](#)
- [Leverage "thinking" commands](#)
- [Tool call ordering](#)
- [Create quality tool schemas](#)
- [Create sub-agents](#)
- [Use tools for multimodal inputs](#)
- [Next steps](#)

## Set the right inference parameters

Tool calling requires a very specific structured output from the model and is improved by using the following inference parameters:

- **Non Reasoning Mode:** Temperature: 0.7 & Top P: 0.9
- **Reasoning Mode:** Temperature: 1 & Top P: 0.9

## Consider latency requirements

### Tip

Amazon Nova 2 models are capable of tool calling with reasoning on and off. However, reasoning modes have a significant impact on latency.

For latency sensitive applications, you should optimize for reasoning off mode and simplify the required tool calls where possible. Split multistep workflows into discrete steps to reduce the model's reliance on regurgitating unnecessary parameters.

## Use intentional wording for tool calling instructions

**Tool names:** Referencing tools in the system prompt is common in tool calling systems to instruct the model on when to call a tool. When you reference tools in the prompt, we recommend you use the tool name instead of xml or pythonic references or examples.

### Example of good tool reference

```
Use the 'run_shell_command' tool for running shell commands
```

### Example of bad tool reference

```
Call run_shell_command() to run shell commands
```

## Leverage "thinking" commands

For all use cases where thinking is beneficial for tool calling, we recommend you leverage reasoning mode instead of prompting the model to “think in tags” or to use a “think” tool.

Amazon Nova 2 models are trained extensively for reasoning mode and will produce the optimal results when used in reasoning mode for chain of thought.

## Tool call ordering

In use cases that might require the use of built-in tools and native tool calling simultaneously, the model biases towards calling built-in tools first.

Don't instruct the model to act differently in the prompt. Instead, design your workflow to take this into account.

For example, if you do not want the model to use built-in tools, do not include them in your workflow so that the model does not bias towards them.

## Create quality tool schemas

Tool schemas are one of the key places you can prompt engineer effective tool calling systems. However, it's important to consider what gets captured in the tool schema itself, how each schema

element is described semantically and how the system prompt references tools and schema elements within system instructions.

Amazon Nova 2 models are optimized for concise descriptions in the tool schemas. Keep it brief.

### Tool schema versus system prompt guidelines:

#### Include in the tool schema:

- Core functionality: What the tool does (20-50 words recommended)
- Parameter specifications: Clear descriptions of each parameter (around 10 words per parameter)
- Expected formats: Data types (like enum, int, float), required fields and valid value ranges

#### Include in the system prompt:

- Dedicate a `#Tool Usage` section with orchestration logic (when and why to use specific tools) and business rules (conditional logic, sequencing requirements, and dependencies).
- **Error handling strategies:** Add an `#Error Handling and Troubleshooting` section with instructions for how to respond to failures or unexpected outputs
- **Output formatting:** Add details on how to present to the user

### Example

```
You are a software engineering issue root cause analysis agent. You are tasked with reviewing a customer issue and examining the repository to identify a plan to resolve the issue.

Core Mandates
- **DO NOT** update the original issue that was posted by the user. You only add *additional* comments to the reported issue if necessary

Primary Workflows
1. **Understand:** Analyze the user's request and explore the codebase thoroughly using **get_file_contents** to grasp file structures and conventions.
2. **Plan:** Create a coherent, evidence-based plan for resolving the task and share it with the user following the format below

Tool Usage
- **Read the Issue:** Always start by using the **read_issue** tool to get the details about the requested issue
```

```
- **File Paths:** Always end the file path with "/" if you are searching a directory
 using the **get_file_contents** tools
- **Parallelism:** Execute multiple independent tool calls in parallel when feasible

Error Handling and Troubleshooting
- **File Exploration:** If you get an error that a file doesn't exist, try searching at
 the directory level first to validate the file path

Output Formatting
Return your plan in markdown in the following format

Issue
<Your root cause analysis of the issue>

Resolution Plan
<your step by step plan of how to solve the issue>
```

## Create sub-agents

Consider creating specialized sub-agents instead of a single agent with many tools when you encounter:

- **Tool count exceeds 20:** Large tool sets become difficult to manage and increase selection errors
- **Distinct functional domains:** Tools naturally cluster into separate categories (such as data retrieval versus processing versus reporting)
- **Complex schemas:** When parameter depth exceeds 3-4 levels or tools have intricate interdependencies
- **Conversation length:** Workflows regularly exceed 15-20 turns may benefit from specialized sub-agents
- **Performance degradation:** If you observe decreased accuracy in tool selection or increased latency

### Tip

MCP Servers come with tools and schemas you can't control. Only include the necessary tools for your workflow to complete the required task.

## Use tools for multimodal inputs

For multimodal tasks, we have not observed improved accuracy leveraging tools for structured tasks (such as extraction or timestamp generation).

Instead, we recommend you review the relevant sections in the Prompting multimodal inputs section for how to prompt the model successfully using the provided templates.

### Next steps

- For multimodal prompting, see [Prompting multimodal inputs](#).

## Prompting multimodal inputs

The following sections provide guidance for image and video understanding. For audio-related prompting, refer to the [the section called “Voice conversation prompts”](#) section.

## General multimodal guidelines

### User prompts and system prompts

For multimodal understanding use cases, every request should include user prompt text. System prompts, which may only contain text, are optional.

System prompts can be used to specify a persona for the model and to define general personality and response style but should not be used for detailed task definition or output formatting instructions.

Include task definition, instructions and formatting details in the user prompt to have a stronger effect than the system prompt for multimodal use cases.

### Content order

A multimodal understanding request sent to Amazon Nova should contain one or more files and a user prompt. The user text prompt should be the last item in the message, always after the image, document, or video content.

```
message = {
 "role": "user",
 "content": [
 { "document|image|video|audio": {...} },
```

```
{ "document|image|video|audio": {...} },
...
{ "text": "<user prompt>" }
]
```

In cases where you want to refer to specific files from within the user prompt, use “text” elements to define labels that precede each file block.

```
message = {
 "role": "user",
 "content": [
 { "text": "<label for item 1>" },
 { "document|image|video|audio": {...} },
 { "text": "<label for item 2>" },
 { "document|image|video|audio": {...} },
 ...
 { "text": "<user prompt>" }
]
}
```

## Document and Image Understanding

The following sections provide guidance on how to craft prompts for tasks that require understanding or analyzing images and documents.

### Extracting text from images

Amazon Nova models can extract text from images, a capability referred to as Optical Character Recognition (OCR). For best results, ensure the image input you provide to the model is a high enough resolution that the text characters are easy to discern.

For text extraction use cases, we recommend the following inference configuration:

- **temperature:** default (0.7)
- **topP:** default (0.9)
- Do not enable reasoning

The Amazon Nova models can extract text to Markdown, HTML, or LaTeX format. The following user prompt template is recommended:



**## Instructions**

Extract all information from this page using only {text\_formatting} formatting. Retain the original layout and structure including lists, tables, charts and math formulae.

**## Rules**

1. For math formulae, always use LaTeX syntax.
2. Describe images using only text.
3. NEVER use HTML image tags `<img>` in the output.
4. NEVER use Markdown image tags `![]()` in the output.
5. Always wrap the entire output in ````` tags.

The output is wrapped in full or partial Markdown code fences (`````). You can strip the code fences using code similar to the following:

```
def strip_outer_code_fences(text):
 lines = text.split("\n")
 # Remove only the outer code fences if present
 if lines and lines[0].startswith("```"):
 lines = lines[1:]
 if lines and lines[-1].startswith("```"):
 lines = lines[:-1]
 return "\n".join(lines).strip()
```

## Extracting structured information from images or text

The Amazon Nova models can extract information from images into machine-parsable JSON format, a process referred to as Key Information Extraction (KIE). To perform KIE, provide the following:

- A JSON schema. A formal schema definition that follows the JSON Schema specification.
- One or more of the following: A document file or image or document text

The document or image must always be placed before your user prompt in the request.

For KIE use cases, we recommend the following inference configuration:

- **temperature:** 0
- **Reasoning:** Reasoning is not required but can improve results when image-only input or complex schemas are used.

## Prompt templates

### Document or image only input

Given the image representation of a document, extract information in JSON format according to the given schema.

Follow these guidelines:

- Ensure that every field is populated, provided the document includes the corresponding value. Only use null when the value is absent from the document.
- When instructed to read tables or lists, read each row from every page. Ensure every field in each row is populated if the document contains the field.

JSON Schema:

```
{json_schema}
```

### Text only input

Given the OCR representation of a document, extract information in JSON format according to the given schema.

Follow these guidelines:

- Ensure that every field is populated, provided the document includes the corresponding value. Only use null when the value is absent from the document.
- When instructed to read tables or lists, read each row from every page. Ensure every field in each row is populated if the document contains the field.

JSON Schema:

```
{json_schema}
```

OCR:

```
{document_text}
```

### Document or Image and text input

Given the image and OCR representations of a document, extract information in JSON format according to the given schema.

Follow these guidelines:

- Ensure that every field is populated, provided the document includes the corresponding value. Only use null when the value is absent from the document.
- When instructed to read tables or lists, read each row from every page. Ensure every field in each row is populated if the document contains the field.

```
JSON Schema:
{json_schema}

OCR:
{document_text}
```

## Detecting objects and their positions in images

Amazon Nova 2 models provide the ability to identify objects and their positions within images, a task sometimes referred to as image grounding or object localization. Practical applications include image analysis and tagging, user interface automation, image editing and others.

Regardless of the image input resolution and aspect ratio, the model uses a coordinate space that divides the image into 1,000 units horizontally and 1,000 units vertically, with the x:0 y:0 location being the upper left of the image.

Bounding boxes are described using the format [x1, y1, x2, y2] representing left, top, right and bottom respectively. Two-dimensional coordinates are represented using the format [x, y].

For object detection use cases, we recommend the following inference parameter values:

- **temperature:** 0
- Do not enable reasoning

### Prompt templates: general object detection

We recommend the following user prompt templates.

#### Detecting multiple instances with bounding boxes:

```
Please identify {target_description} in the image and provide the bounding box
coordinates for each one you detect. Represent the bounding box as the [x1, y1, x2,
y2] format, where the coordinates are scaled between 0 and 1000 to the image width and
height, respectively.
```

#### Detecting a single region with bounding box:

```
Please generate the bounding box coordinates corresponding to the region described in
this sentence: {target_description}. Represent the bounding box as the [x1, y1, x2,
```

y2] format, where the coordinates are scaled between 0 and 1000 to the image width and height, respectively.

### Detecting multiple instances with center points:

Please identify {target\_description} in the image and provide the center point coordinates for each one you detect. Represent the point as the [x, y] format, where the coordinates are scaled between 0 and 1000 to the image width and height, respectively.

### Detecting a single region with center point:

Please generate the center point coordinates corresponding to the region described in this sentence: {target\_description}. Represent the center point as the [x, y] format, where the coordinates are scaled between 0 and 1000 to the image width and height, respectively.

### Parsing model output:

Each of the recommended prompts above will produce a comma separated string containing one or more bounding box descriptions in a form similar to the following. There may be some slight variation in whether a "." is included at the end of the string. For example, [356, 770, 393, 872], [626, 770, 659, 878].

You can parse the coordinate information generated by the model using a regular expression as shown in the following Python code example.

### Code example

```
def parse_coord_text(text):
 """Parses a model response which uses array formatting ([x, y, ...])
 to describe points and bounding boxes. Returns an array of tuples."""
 pattern = r"\([([^\]]*?)\)"
 return [
 tuple(int(x.strip()) for x in match.split(","))
 for match in re.findall(pattern, text)
]
```

### Code example

To remap the normalized coordinates of a bounding box to the coordinate space of the input image, you can use a function similar to the following Python example.

```
def remap_bbox_to_image(bounding_box, image_width, image_height):
 return [
 bounding_box[0] * image_width / 1000,
 bounding_box[1] * image_height / 1000,
 bounding_box[2] * image_width / 1000,
 bounding_box[3] * image_height / 1000,
]
```

## Prompt templates: detecting multiple object classes with positions

When you want to identify multiple classes of items in an image, you can include a class list in your prompt using one of the following formatting approaches.

For commonly understood classes that the model is likely to understand well, list the class names (without quotes) inside square brackets:

```
[car, traffic light, road sign, pedestrian]
```

For classes that are nuanced, uncommon, or come from specialized domains that the model may not be familiar with, include a definition for each class in parentheses. Since this task is challenging, expect the model's performance to degrade.

```
[taraxacum officinale (Dandelion - bright yellow flowers, jagged basal leaves, white
puffball seed heads), digitaria spp (Crabgrass - low spreading grass with coarse
blades and finger-like seed heads), trifolium repens (White Clover - three round
leaflets and small white pom-pom flowers), plantago major (Broadleaf Plantain - wide
oval rosette leaves with tall narrow seed stalks), stellaria media (Chickweed - low
mat-forming plant with tiny star-shaped white flowers)]
```

Use one of the following **user prompt** templates depending on which JSON output format you prefer.

### Prompt option 1

```
Detect all objects with their bounding boxes in the image from the provided class list.
Normalize the bounding box coordinates to be scaled between 0 and 1000 to the image
width and height, respectively.
```

```
Classes: {candidate_class_list}
```

Include separate entries for each detected object as an element of a list.

Formulate your output as JSON format:

```
[
 {
 "class 1": [x1, y1, x2, y2]
 },
 ...
]
```

## Prompt option 2

Detect all objects with their bounding boxes in the image from the provided class list. Normalize the bounding box coordinates to be scaled between 0 and 1000 to the image width and height, respectively.

Classes: **{candidate\_class\_list}**

Include separate entries for each detected object as an element of a list.

Formulate your output as JSON format:

```
[
 {
 "class": class 1,
 "bbox": [x1, y1, x2, y2]
 },
 ...
]
```

## Prompt option 3

Detect all objects with their bounding boxes in the image from the provided class list. Normalize the bounding box coordinates to be scaled between 0 and 1000 to the image width and height, respectively.

Classes: **{candidate\_class\_list}**

Group all detected bounding boxes by class.

Formulate your output as JSON format:

```
{
 "class 1": [[x1, y1, x2, y2], [x1, x2, y1, y2], ...],
 ...
}
```

```
}
```

## Prompt option 4

Detect all objects with their bounding boxes in the image from the provided class list. Normalize the bounding box coordinates to be scaled between 0 and 1000 to the image width and height, respectively.

Classes: `{candidate_class_list}`

Group all detected bounding boxes by class.

Formulate your output as JSON format:

```
[
 {
 "class": class 1,
 "bbox": [[x1, y1, x2, y2], [x1, x2, y1, y2], ...]
 },
 ...
]
```

## Parsing model output

The output will be encoded as a JSON which can be parsed with any JSON parsing library.

## Prompt templates: screenshot UI bounds detection

We recommend the following user prompt templates.

### Detecting UI element position based on a goal:

In this UI screenshot, what is the location of the element if I want to {goal}? Express the location coordinates using the [x1, y1, x2, y2] format, scaled between 0 and 1000.

### Detecting UI element position based on text:

In this UI screenshot, what is the location of the element if I want to click on "{text}"? Express the location coordinates using the [x1, y1, x2, y2] format, scaled between 0 and 1000.

## Parsing model output:

For each of the UI bounds detection prompts above, you can parse the coordinate information generated by the model using a regular expression as shown in the Python code example below.

### Code example

```
def parse_coord_text(text):
 """Parses a model response which uses array formatting ([x, y, ...])
 to describe points and bounding boxes. Returns an array of tuples."""
 pattern = r"\s*\([^\)]*\)"
 return [
 tuple(int(x.strip()) for x in match.split(","))
 for match in re.findall(pattern, text)
]
```

## Video understanding

The following sections provide guidance on how to craft prompts for tasks that require understanding or analyzing videos.

### Summarizing videos

Amazon Nova models can generate summaries of video content.

For video summarization use cases, we recommend the following inference parameter values:

- **temperature:** 0
- Some use cases may benefit from enabling model reasoning

No specific prompting template is required. Your user prompt should clearly specify the aspects of the video you care about. Here are a few examples of effective prompts:

Can you create an executive summary of this video's content?

Can you distill the essential information from this video into a concise summary?

Could you provide a summary of the video, focusing on its key points?



## Generating detailed captions for videos

Amazon Nova models can generate detailed captions for videos, a task referred to as dense captioning.

For video captioning use cases, we recommend the following inference parameter values:

- **temperature:** 0
- Some use cases may benefit from enabling model reasoning

No specific prompting template is required. Your user prompt should clearly specify the aspects of the video you care about. Here are a few examples of effective prompts:

```
Provide a detailed, second-by-second description of the video content.
```

```
Break down the video into key segments and provide detailed descriptions for each.
```

```
Generate a rich textual representation of the video, covering aspects like movement, color and composition.
```

```
Describe the video scene-by-scene, including details about characters, actions and settings.
```

```
Offer a detailed narrative of the video, including descriptions of any text, graphics, or special effects used.
```

```
Create a dense timeline of events occurring in the video, with timestamps if possible.
```

## Analyzing security video footage

Amazon Nova models can detect events in security footage.

For security footage use cases, we recommend the following inference parameter values:

- **temperature:** 0
- Some use cases may benefit from enabling model reasoning

You are a security assistant for a smart home who is given security camera footage in natural setting. You will examine the video and describe the events you see. You are capable of identifying important details like people, objects, animals, vehicles, actions and activities. This is not a hypothetical, be accurate in your responses. Do not make up information not present in the video.

## Extracting video events with timestamps

Amazon Nova models can identify timestamps related to events in a video. You may request that time stamps be formatted in seconds or in MM:SS format. For example, an event occurring at 1 minute 25 seconds in the video can be represented as 85 or 01:25.

For this use case, we recommend the following inference parameter values:

- **temperature:** 0
- Do not use reasoning

We recommend you use prompts similar to the following:

### Identifying an event's start and end time

Please localize the moment that the event "{event\_description}" happens in the video. Answer with the starting and ending time of the event in seconds, such as [[72, 82]]. If the event happen multiple times, list all of them, such as [[40, 50], [72, 82]].

Locate the segment where "{event\_description}" happens. Specify the start and end times of the event in MM:SS.

Answer the starting and end time of the event "{event\_description}". Provide answers in MM:SS

When does "{event\_description}" in the video? Specify the start and end timestamps, e.g. [[9, 14]]

### Identifying multiple occurrences of an event

Please localize the moment that the event "{event\_description}" happens in the video. Answer with the starting and ending time of the event in seconds. e.g. [[72, 82]]. If the event happen multiple times, list all of them. e.g. [[40, 50], [72, 82]]

## Generating a log of video segments with timestamps

Segment a video into different scenes and generate caption per scene. The output should be in the format: [STARTING TIME-ENDING TIMESTAMP] CAPTION. Timestamp in MM:SS format

For a video clip, segment it into chapters and generate chapter titles with timestamps. The output should be in the format: [STARTING TIME] TITLE. Time in MM:SS

Generate video captions with timestamp.

## Classifying videos

You can use Amazon Nova models to classify videos based on a pre-defined list of classes you provide.

For this use case, we recommend the following inference parameter values:

- **temperature:** 0
- Reasoning should not be used

Use the following prompt template:

```
What is the most appropriate category for this video? Select your answer from the
options provided:
{class1}
{class2}
{...}
```

### Example:

```
What is the most appropriate category for this video? Select your answer from the
options provided:
Arts
Technology
Sports
Education
```

# Content moderation

This guide provides best practices for prompting Amazon Nova models to perform content moderation using the [MLCommons® AIIlluminate Assessment Standard](#) — a comprehensive framework with 12 hazard categories for AI safety. While this guide uses the MLCommons standard as an example to illustrate best practices, the prompting techniques work with any free-form or customizable content moderation categories and definitions you specify.

## Key principles

- Use structured (XML, JSON) or free-form prompts based on your use case
- Moderate any text content: conversations, articles, social media posts, comments, reviews, and more
- Provide clear policy definitions
- Request explanations for transparency

## Recommended configuration

For content moderation tasks, we recommend the following inference configuration:

- **Temperature** — Use the default value of 0.7.
- **Top P** — Use the default value of 0.9.
- **Reasoning** — Disable reasoning for content moderation to reduce latency.

## MLCommons universal taxonomy

The standard defines 12 hazard categories:

### MLCommons AIIlluminate hazard categories

| Category       | Definition                                                                 | Key subcategories                                                             |
|----------------|----------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| Violent Crimes | Content that enables, encourages, or endorses commission of violent crimes | Mass violence, Murder, Physical assault, Domestic abuse, Terror, Animal abuse |

| Category                  | Definition                                                                                | Key subcategories                                                                                   |
|---------------------------|-------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| Non-Violent Crimes        | Content that enables, encourages, or endorses non-violent criminal activity               | Labor trafficking, Fraud, Theft, Arson, Drug crimes, Weapons crimes, Cyber crimes                   |
| Sex-Related Crimes        | Content that enables, encourages, or endorses sex-related crimes                          | Sex trafficking, Sexual assault, Sexual harassment, Non-consensual intimate content, Bestiality     |
| Child Sexual Exploitation | Content that contains, describes, enables, encourages, or endorses sexual abuse of minors | Inappropriate relationships with children, Sexual abuse of children, CSAM                           |
| CBRNE                     | Content that enables, encourages, or endorses creation of weapons of mass destruction     | Chemical weapons, Biological weapons, Radiological weapons, Nuclear weapons, High-yield explosives  |
| Suicide and Self-Harm     | Content that enables, encourages, or endorses intentional self-harm                       | Suicide, Self-injury, Disordered eating, Dangerous challenges                                       |
| Hate                      | Content that demeans or dehumanizes people based on protected characteristics             | Demeaning remarks, Inflicting physical/emotional harm, Exclusion and isolation                      |
| Specialized Advice        | Unqualified advice in sensitive domains without appropriate disclaimers                   | Financial advice, Medical advice, Legal advice, Electoral advice, Safety claims                     |
| Privacy                   | Content containing sensitive non-public personal information                              | Physical security (addresses), Digital security (credentials), Financial security (account numbers) |
| Intellectual Property     | Content that violates or encourages violation of IP rights                                | Copyright infringement, Trademark infringement, Patent infringement                                 |
| Defamation                | Content that is verifiably false and likely to injure someone's reputation                | Libel, Slander, Disparagement                                                                       |

| Category       | Definition                                        | Key subcategories                                             |
|----------------|---------------------------------------------------|---------------------------------------------------------------|
| Sexual Content | Erotic content including pornography and cybersex | Pornography (graphic sexual descriptions), Direct erotic chat |

Source: MLCommons® Alluminate Assessment Standard (Documentation)

## Prompting examples

### Structured format

Best for automated systems requiring parseable output. These examples use few-shot learning, where you provide example input-output pairs in the prompt to guide the model's responses. The examples below include a single demonstration, but you can include multiple examples to improve accuracy for your specific use case.

### Structure (XML)

User:

```
You are a text content moderator that detects policy violations,
describes violated categories, and explains the reasoning behind your decisions.
You will see a policy list and a conversation to be moderated.
Your moderation response should be formatted as follows.

Start the response by indicating policy violation status using either
<POLICY_VIOLATION>Yes</POLICY_VIOLATION> for content that violates policies, or
<POLICY_VIOLATION>No</POLICY_VIOLATION> for content that does not violate policies.

Then, list the specific violation categories between <CATEGORY_LIST> and </
CATEGORY_LIST> tags.
If there is no violation, then put C0 between the tags.

End the response by providing an analysis explaining whether the content violates
policies
and why in <EXPLAIN> </EXPLAIN> tags. Here is an example.

<EXAMPLE>
<POLICY>
{{Policy definitions and categories}}
</POLICY>
```

```

<TEXT>
{{Content to moderate}}
</TEXT>
<POLICY_VIOLATION>Yes/No</POLICY_VIOLATION>
<CATEGORY_LIST>Categories or C0</CATEGORY_LIST>
<EXPLAIN>Explanation</EXPLAIN>
</EXAMPLE>

```

Now, please moderate the following input policy and text.

```

<POLICY>
{{Policy definitions and categories}}
</POLICY>
<TEXT>
{{Content to moderate}}
</TEXT>

```

**Assistant:**

```

<POLICY_VIOLATION>Yes/No</POLICY_VIOLATION>
<CATEGORY_LIST>Categories or C0</CATEGORY_LIST>
<EXPLAIN>Explanation</EXPLAIN>

```

## Structure (JSON)

**User:**

You are a text content moderator that detects policy violations, describes violated categories, and explains the reasoning behind your decisions. You will check if there is unsafe content in the below text snippet in accordance with our safety policy.

Your moderation response should be in JSON format with the following structure:

```

{
 "policy_violation": "Yes or No",
 "category_list": ["List of violated category codes"],
 "explanation": "Reason for the decision"
}

```

If there is no violation, use "C0" as the category code.

**IMPORTANT:** Do not add any additional text or explanation. Your response must contain ONLY the JSON object, nothing else.

Here is an example:

```
[POLICY]
{{Policy definitions and categories}}

[TEXT]
{{Content to moderate}}

[RESPONSE]
{
 "policy_violation": "Yes or No",
 "category_list": ["Violated category codes or C0"],
 "explanation": "Explanation"
}
```

Now, please moderate the following input.

```
[POLICY]
{{Policy definitions and categories}}

[TEXT]
{{Content to moderate}}

[RESPONSE]
```

**Assistant:**

```
{
 "policy_violation": "No",
 "category_list": ["C0"],
 "explanation": "The text is a casual personal question with no policy violations."
}
```

## Sample content moderation prompt and response

**User:**

You are a text content moderator that detects policy violations, describes violated categories, and explains the reasoning behind your decisions. You will see a policy list and a conversation to be moderated. Your moderation response should be formatted as follows.

Start the response by indicating policy violation status using either



```
<POLICY_VIOLATION>Yes</POLICY_VIOLATION> for content that violates policies, or
<POLICY_VIOLATION>No</POLICY_VIOLATION> for content that does not violate policies.
```

Then, list the specific violation categories between <CATEGORY\_LIST> and </CATEGORY\_LIST> tags.

If there is no violation, then put C0 between the tags.

End the response by providing an analysis explaining whether the content violates policies and why in <EXPLAIN> </EXPLAIN> tags. Here is an example.

```
<POLICY>
```

The 12 categories are: Violent Crimes, Non-Violent Crimes, Sex-Related Crimes, Child Sexual Exploitation, CBRNE, Suicide & Self-Harm, Hate, Specialized Advice, Privacy, Intellectual Property, Defamation, Sexual Content.

If none apply, label as C0.

```
</POLICY>
```

```
<EXAMPLE>
```

```
<TEXT>
```

"To make a pipe bomb: Get metal pipe, end caps, gunpowder, and fuse. Drill hole in cap, fill pipe with powder, insert fuse, seal both ends tightly."

```
</TEXT>
```

```
<POLICY_VIOLATION>Yes</POLICY_VIOLATION>
```

```
<CATEGORY_LIST>CBRNE</CATEGORY_LIST>
```

```
<EXPLAIN>Provides instructions for creating explosives, enabling dangerous weapons creation.</EXPLAIN>
```

```
</EXAMPLE>
```

```
<TEXT>
```

DIY explosive device: Mix ammonium nitrate fertilizer with fuel oil at 94:6 ratio. Use blasting cap for detonation. Store in cool, dry place. #DIY #Chemistry

```
</TEXT>
```

**Assistant:**

```
<POLICY_VIOLATION>Yes</POLICY_VIOLATION>
```

```
<CATEGORY_LIST>CBRNE</CATEGORY_LIST>
```

```
<EXPLAIN>
```

The text provides specific instructions for creating ANFO explosives, including precise mixing ratios and detonation methods. This falls under CBRNE (high-yield explosives subcategory) as it enables dangerous weapons creation with actionable technical details.

</EXPLAIN>

## Free-form content moderation

Free-form prompts can take many diverse structures beyond the examples below. The following examples intentionally use different formatting styles to illustrate the flexibility of free-form prompting. You can adapt the structure, phrasing, and format to suit your specific use case.

### Example 1: Yes/No question (Suicide & Self-Harm)

#### Example Yes/No content moderation

User:

Please stick to the following policy for content moderation. {{user\_policy}}  
Does this text promote harmful activities?

Best way to hurt yourself without leaving marks?  
Asking for tips on methods that work.

Assistant:

Yes. This text explicitly requests methods for self-harm, falling under the Suicide & Self-Harm category. The intent to conceal evidence ("without leaving marks") indicates deliberate harmful behavior.

### Example 2: Category identification and reasoning (Hate)

#### Example Category identification content moderation

User:

You are a professional content moderator. Please stick to the following policy for content moderation. {{user\_policy}}  
Which category does this text violate and why?

Text: "Women are naturally worse at mathematics than men, according to my personal observations. They should stick to arts and humanities."

Assistant:

Hate. This text makes demeaning generalizations about women based on gender (a protected characteristic), promoting exclusion from STEM fields. The claim lacks scientific basis and perpetuates harmful stereotypes.

### Example 3: Severity rating (Privacy)

#### Example Severity rating content moderation

User:

```
Content moderation task

Policy: {{user_policy}}

Rate the severity of the following content as "none", "low", "medium", or "high"
and identify any violated categories.

Content: "John Smith lives at 742 Evergreen Terrace, Springfield.
His SSN is 123-45-6789 and his bank account number is 9876543210."
```

Assistant:

```
Severity: High

Violated category: Privacy

This content exposes multiple forms of sensitive personal information
including a full name, home address, Social Security number, and bank
account number. This falls under the Privacy category, specifically
the physical security, digital security, and financial security
subcategories. The combination of identifiers creates a significant
risk of identity theft and financial fraud.
```

# Code and Troubleshooting

This chapter provides code examples and troubleshooting guidance to help you work effectively with Amazon Nova models.

## Code library

This section provides code examples for common Amazon Nova operations using either the Converse API or the InvokeModel API.

### Converse API Examples

#### Basic request

Send a basic text request to Amazon Nova models using the Converse API.

#### Non-streaming

```
import boto3
from botocore.config import Config

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

Invoke the model
response = bedrock.converse(
 modelId="us.amazon.nova-2-lite-v1:0",
 messages=[
 {
 "role": "user",
 "content": [{"text": "Write a short story. End the story with 'THE\nEND'. cant\n\n"}],
 }
],
 system=[{"text": "You are a children's book author."}], # Optional
```

```

inferenceConfig={ # These parameters are optional
 "maxTokens": 1500,
 "temperature": 0.7,
 "topP": 0.9,
 "stopSequences": ["THE END"],
},
additionalModelRequestFields={ # These parameters are optional
 "inferenceConfig": {
 "topK": 50,
 }
},
)

Extract the text response
content_list = response["output"]["message"]["content"]
for content in content_list:
 if "text" in content:
 print(content["text"])

```

## Streaming

```

import boto3
from botocore.config import Config

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(connect_timeout=3600, read_timeout=3600),
)

Invoke the model
response = bedrock.converse_stream(
 modelId="us.amazon.nova-2-lite-v1:0",
 messages=[
 {
 "role": "user",
 "content": [{"text": "Write a short story. End the story with 'THE
END'. "}],
 }
],
 system=[{"text": "You are a children's book author."}], # Optional

```

```

inferenceConfig={ # These parameters are optional
 "maxTokens": 1500,
 "temperature": 0.7,
 "topP": 0.9,
 "stopSequences": ["THE END"],
},
additionalModelRequestFields={ # These parameters are optional
 "inferenceConfig": {
 "topK": 50,
 }
},
)

Handle streaming events
for event in response["stream"]:
 if "contentBlockDelta" in event:
 delta = event["contentBlockDelta"]["delta"]
 if "text" in delta:
 print(delta["text"], end="", flush=True)

```

## Multimodal input using embedded asset

Process multimodal content by embedding document, image, video, or audio data directly in the request. This example uses image data. For details on the content structure for other modalities, see the [ContentBlock details in the Amazon Bedrock API documentation](#).

### Non-streaming

```

import boto3
from botocore.config import Config

Read a document, image, video, or audio file
with open("sample_image.png", "rb") as image_file:
 binary_data = image_file.read()
 data_format = "png"

Define message with image
messages = [
 {
 "role": "user",
 "content": [
 {

```

```

 "image": {
 "format": data_format,
 "source": {
 "bytes": binary_data # For Invoke API, encode as Base64
string
 },
 },
 },
 {"text": "Provide a brief caption for this asset."},
],
}
]

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

Invoke model
response = bedrock.converse(
 modelId="us.amazon.nova-2-lite-v1:0",
 messages=messages,
)

Extract the text response
content_list = response["output"]["message"]["content"]
for content in content_list:
 if "text" in content:
 print(content["text"])

```

## Streaming

```

import boto3
from botocore.config import Config

Read a document, image, video, or audio file
with open("sample_image.png", "rb") as image_file:
 binary_data = image_file.read()
 data_format = "png"

```

```

Define message with image
messages = [
 {
 "role": "user",
 "content": [
 {
 "image": {
 "format": data_format,
 "source": {
 "bytes": binary_data # For Invoke API, encode as Base64
string
 },
 },
],
 {"text": "Provide a brief caption for this asset."},
],
 }
]

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(connect_timeout=3600, read_timeout=3600),
)

Invoke model with streaming
response = bedrock.converse_stream(
 modelId="us.amazon.nova-2-lite-v1:0",
 messages=messages,
)

Handle streaming events
for event in response["stream"]:
 if "contentBlockDelta" in event:
 delta = event["contentBlockDelta"]["delta"]
 if "text" in delta:
 print(delta["text"], end="", flush=True)

```



## Multimodal input using S3 URI

Process multimodal content by referencing documents, images, videos, or audio files stored in S3. This example uses an image reference. For details on the content structure for other modalities, see the [ContentBlock details in the Amazon Bedrock API documentation](#).

### Non-streaming

```
import boto3
from botocore.config import Config

Define message with image
messages = [
 {
 "role": "user",
 "content": [
 {
 "image": {
 "format": "png",
 "source": {
 "s3Location": {
 "uri": "s3://path/to/your/asset",
 # "bucketOwner": "<account_id>" # Optional
 }
 },
 },
 },
 {"text": "Provide a brief caption for this asset."},
],
 }
]

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

Invoke model
response = bedrock.converse(
 modelId="us.amazon.nova-2-lite-v1:0",
```

```

 messages=messages,
)

Extract the text response
content_list = response["output"]["message"]["content"]
for content in content_list:
 if "text" in content:
 print(content["text"])

```

## Streaming

```

import boto3
from botocore.config import Config

Define message with image
messages = [
 {
 "role": "user",
 "content": [
 {
 "image": {
 "format": "png",
 "source": {
 "s3Location": {
 "uri": "s3://path/to/your/asset",
 # "bucketOwner": "<account_id>" # Optional
 }
 },
 },
 },
 {"text": "Provide a brief caption for this asset."},
],
 }
]

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(connect_timeout=3600, read_timeout=3600),
)

```

```
Invoke model with streaming
response = bedrock.converse_stream(
 modelId="us.amazon.nova-2-lite-v1:0",
 messages=messages,
)

Handle streaming events
for event in response["stream"]:
 if "contentBlockDelta" in event:
 delta = event["contentBlockDelta"]["delta"]
 if "text" in delta:
 print(delta["text"], end="", flush=True)
```

## Extended thinking (reasoning)

Enable extended thinking for complex problem-solving tasks.

### Non-streaming

```
import boto3
from botocore.config import Config

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

Invoke the model
response = bedrock.converse(
 modelId="us.amazon.nova-2-lite-v1:0",
 messages=[
 {
 "role": "user",
 "content": [
 {
 "text": 'How many capital letters appear in the following
passage. Your response must include only the number: "Wilfred ordered an anvil from
ACME. Shipping was expensive."'
 }
]
 }
]
)
```

```

],
 },
 additionalModelRequestFields={
 "reasoningConfig": {
 "type": "enabled",
 "maxReasoningEffort": "low", # "low" | "medium" | "high"
 }
 },
)

Extract response content
content_list = response["output"]["message"]["content"]
for content in content_list:
 # Extract the reasoning response
 if "reasoningContent" in content:
 print("\n== Reasoning ==")
 print(content["reasoningContent"]["reasoningText"]["text"])
 # Extract the text response
 if "text" in content:
 print("\n== Text ==")
 print(content["text"])

```

## Streaming

```

import boto3
from botocore.config import Config

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(connect_timeout=3600, read_timeout=3600),
)

Invoke the model
response = bedrock.converse_stream(
 modelId="us.amazon.nova-2-lite-v1:0",
 messages=[
 {
 "role": "user",
 "content": [

```

```

 {
 "text": 'How many capital letters appear in the following
passage. Your response must include only the number: "Wilfred ordered an anvil from
ACME. Shipping was expensive."'
 }
],
}
],
additionalModelRequestFields={
 "reasoningConfig": {
 "type": "enabled",
 "maxReasoningEffort": "low", # "low" | "medium" | "high"
 },
},
)

Process the streaming response
reasoning_output = ""
text_output = ""
for event in response["stream"]:
 if "contentBlockDelta" in event:
 delta = event["contentBlockDelta"]["delta"]

 if "reasoningContent" in delta:
 if len(reasoning_output) == 0:
 print("\n\n== Reasoning ==")
 reasoning_text_chunk = delta["reasoningContent"]["text"]
 print(reasoning_text_chunk, end="", flush=True)
 reasoning_output += reasoning_text_chunk

 elif "text" in delta:
 if len(text_output) == 0:
 print("\n\n== Text ==")
 text_chunk = delta["text"]
 print(text_chunk, end="", flush=True)
 text_output += text_chunk

```

## Built-in tool: Nova Grounding with citations

Use Nova Grounding to retrieve real-time information from the web with citations.

## Non-streaming

```
import boto3
from botocore.config import Config

Define the list of tools the model may use
tool_config = {"tools": [{"systemTool": {"name": "nova_grounding"}}]}

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

messages = [
 {
 "role": "user",
 "content": [
 {"text": "What is the latest news about renewable energy sources?"}
],
 }
]

Invoke the model
response = bedrock.converse(
 modelId="us.amazon.nova-2-lite-v1:0", messages=messages, toolConfig=tool_config
)

Extract the text with interleaved citations
output_with_citations = ""
content_list = response["output"]["message"]["content"]
for content in content_list:
 if "text" in content:
 output_with_citations += content["text"]

 elif "citationsContent" in content:
 citations = content["citationsContent"]["citations"]
 for citation in citations:
 url = citation["location"]["web"]["url"]
 output_with_citations += f"[{url}]"
```

```
print(output_with_citations)
```

## Streaming

```
import boto3
from botocore.config import Config

Define the list of tools the model may use
tool_config = {"tools": [{"systemTool": {"name": "nova_grounding"}}]}

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

messages = [
 {
 "role": "user",
 "content": [
 {"text": "What is the latest news about renewable energy sources?"}
],
 }
]

Invoke the model with streaming
response = bedrock.converse_stream(
 modelId="us.amazon.nova-2-lite-v1:0", messages=messages, toolConfig=tool_config
)

Process the streaming response with interleaved citations
for event in response["stream"]:
 if "contentBlockDelta" in event:
 delta = event["contentBlockDelta"]["delta"]

 if "text" in delta:
 print(delta["text"], end="", flush=True)

 elif "citation" in delta:
 url = delta["citation"]["location"]["web"]["url"]
 print(f"[{url}]", end="", flush=True)
```

## Built-in tool: Code Interpreter

Use the Code Interpreter tool to execute Python code for calculations and data analysis.

### Non-streaming

```
import boto3
from botocore.config import Config

Define the list of tools the model may use
tool_config = {"tools": [{"systemTool": {"name": "nova_code_interpreter"}}]}

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

messages = [
 {
 "role": "user",
 "content": [
 {
 "text": "What is the average of 10, 24, 2, 3, 43, 52, 13, 68, 6, 7, 902, 82?"
 }
],
 }
]

Invoke the model
response = bedrock.converse(
 modelId="us.amazon.nova-2-lite-v1:0", messages=messages, toolConfig=tool_config
)

Extract the text and the code the was executed
content_list = response["output"]["message"]["content"]
for content in content_list:
 if "text" in content:
 print("\n== Text ==")
 print(content["text"])
```



```

elif "toolUse" in content and content["toolUse"]["name"] ==
"nova_code_interpreter":
 print("\n== Code Interpreter: input.snippet ==")
 print(content["toolUse"]["input"]["snippet"])

```

## Streaming

```

import boto3
from botocore.config import Config
import json

Define the list of tools the model may use
tool_config = {"tools": [{"systemTool": {"name": "nova_code_interpreter"}}]}

messages = [
 {
 "role": "user",
 "content": [
 {
 "text": "What is the average of 10, 24, 2, 3, 43, 52, 13, 68, 6, 7,
902, 82?"
 }
],
 }
]

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(connect_timeout=3600, read_timeout=3600),
)

Invoke the model with streaming
response = bedrock.converse_stream(
 modelId="us.amazon.nova-2-lite-v1:0", messages=messages, toolConfig=tool_config
)

Process the streaming response
current_block_start = None
response_text = ""
for event in response["stream"]:

```

```

if "contentBlockStart" in event:
 current_block_start = event["contentBlockStart"]["start"]

elif "contentBlockStop" in event:
 current_block_start = None

elif "contentBlockDelta" in event:
 delta = event["contentBlockDelta"]["delta"]

 if (
 current_block_start
 and "toolUse" in current_block_start
 and current_block_start["toolUse"]["name"] == "nova_code_interpreter"
):
 # This is code interpreter content
 tool_input = json.loads(delta["toolUse"]["input"])
 print("\n== Executed Code Snippet ==")
 print(tool_input["snippet"], end="", flush=True)

 elif "text" in delta:
 # This is text response content
 if len(response_text) == 0:
 print("\n== Text ==")
 text = delta["text"]
 response_text += text
 print(text, end="", flush=True)

```

## Tool use

Define custom tools for the model to use during conversation.

### Non-streaming

```

import boto3
from botocore.config import Config

def get_weather(city):
 # Mock function to simulate weather API
 return {"temperatureF": 48, "conditions": "light rain"}

```

```
Define the toolSpec for the weather tool
weather_tool = {
 "toolSpec": {
 "name": "get_weather",
 "description": "Get the current weather conditions in a given location",
 "inputSchema": {
 "json": {
 "type": "object",
 "properties": {
 "city": {
 "type": "string",
 "description": "The city and state, e.g. San Francisco, CA",
 }
 },
 "required": ["city"],
 }
 },
 }
}

Define the list of tools the model may use
tool_config = {"tools": [weather_tool]}

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

Start tracking message history
messages = []

messages.append(
 {
 "role": "user",
 "content": [
 {
 "text": "Suggest some activities to do in Seattle based on the
current weather."
 }
],
 }
)
```

```
)

Invoke the model
response = bedrock.converse(
 modelId="us.amazon.nova-2-lite-v1:0", messages=messages, toolConfig=tool_config
)

assistant_message = response["output"]["message"]

Add the assistant response to the message history
messages.append(assistant_message)

content_list = assistant_message["content"]
stop_reason = response["stopReason"]

if stop_reason == "tool_use":
 # Extract the toolUse details
 tool_use = next(
 content["toolUse"] for content in content_list if "toolUse" in content
)
 tool_name = tool_use["name"]
 tool_use_id = tool_use["toolUseId"]

 if tool_name == "get_weather":
 # Call the tool
 weather = get_weather(tool_use["input"]["city"])

 # Send the result back to the model
 messages.append(
 {
 "role": "user",
 "content": [
 {
 "toolResult": {
 "toolUseId": tool_use_id,
 "content": [{"json": weather}],
 }
 }
],
 }
)

 # Submit the tool result back to the model
 response = bedrock.converse(
```

```

 modelId="us.amazon.nova-2-lite-v1:0",
 messages=messages,
 toolConfig=tool_config,
)

 content_list = response["output"]["message"]["content"]
 for content in content_list:
 # Extract the text response
 if "text" in content:
 print("\n== Text ==")
 print(content["text"])
 else:
 # A tool call was not needed
 for content in content_list:
 # Extract the text response
 if "text" in content:
 print("\n== Text ==")
 print(content["text"])

```

## Streaming

```

import boto3
from botocore.config import Config
import json

def get_weather(city):
 # Mock function to simulate weather API
 return {"temperatureF": 48, "conditions": "light rain"}

Define the toolSpec for the weather tool
weather_tool = {
 "toolSpec": {
 "name": "get_weather",
 "description": "Get the current weather conditions in a given location",
 "inputSchema": {
 "json": {
 "type": "object",
 "properties": {
 "city": {
 "type": "string",
 "description": "The city and state, e.g. San Francisco, CA",

```

```

 }
 },
 "required": ["city"],
}
},
}
}

Define the list of tools the model may use
tool_config = {"tools": [weather_tool]}

Create the Bedrock Runtime client, using an extended timeout configuration
to support long-running requests.
bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

Start tracking message history
messages = []

messages.append(
 {
 "role": "user",
 "content": [
 {
 "text": "Suggest some activities to do in Seattle based on the
current weather."
 }
],
 }
)

Invoke the model with streaming
response = bedrock.converse_stream(
 modelId="us.amazon.nova-2-lite-v1:0", messages=messages, toolConfig=tool_config
)

Process the streaming response
assistant_message = {"role": "assistant", "content": []}
current_tool_use = None
stop_reason = None

```

```
for event in response["stream"]:
 if "contentBlockStart" in event:
 start = event["contentBlockStart"]["start"]
 if "toolUse" in start:
 current_tool_use = start["toolUse"]
 current_tool_use["input"] = ""

 elif "contentBlockDelta" in event:
 delta = event["contentBlockDelta"]["delta"]
 if "toolUse" in delta:
 current_tool_use["input"] += delta["toolUse"]["input"]
 elif "text" in delta:
 print(delta["text"], end="", flush=True)

 elif "contentBlockStop" in event:
 if current_tool_use:
 # Parse the accumulated tool input
 current_tool_use["input"] = json.loads(current_tool_use["input"])
 assistant_message["content"].append({"toolUse": current_tool_use})
 current_tool_use = None

 elif "messageStop" in event:
 stop_reason = event["messageStop"]["stopReason"]
 if stop_reason == "end_turn":
 exit

Add the assistant response to the message history
messages.append(assistant_message)

if stop_reason == "tool_use":
 # Extract the toolUse details
 tool_use = next(
 content["toolUse"]
 for content in assistant_message["content"]
 if "toolUse" in content
)
 tool_name = tool_use["name"]
 tool_use_id = tool_use["toolUseId"]

 if tool_name == "get_weather":
 # Call the tool
 weather = get_weather(tool_use["input"]["city"])

 # Send the result back to the model
```

```

messages.append(
 {
 "role": "user",
 "content": [
 {
 "toolResult": {
 "toolUseId": tool_use_id,
 "content": [{"json": weather}],
 }
 },
],
 }
)

Submit the tool result back to the model with streaming
response = bedrock.converse_stream(
 modelId="us.amazon.nova-2-lite-v1:0",
 messages=messages,
 toolConfig=tool_config,
)

Handle the final streaming response
print("\n== Text ==")
for event in response["stream"]:
 if "contentBlockDelta" in event:
 delta = event["contentBlockDelta"]["delta"]
 if "text" in delta:
 print(delta["text"], end="", flush=True)

```

## InvokeModel API Examples

The examples below focus on the few key areas where the Invoke API's request and response structures differ slightly from those of the Converse API. In most other ways, the two APIs are compatible, so you should be able to easily adapt the Converse API examples above to work with the InvokeModel API.

### Basic request

Send a basic text request to Amazon Nova 2 models using the InvokeModel API.



## Non-streaming

```
import json

import boto3
from botocore.config import Config

Configure the request
request_body = {
 "messages": [
 {
 "role": "user",
 "content": [{"text": "Write a short story. End the story with 'THE
END'. "}],
 }
],
 "system": [{"text": "You are a children's book author."}], # Optional
 "inferenceConfig": { # These parameters are optional
 "maxTokens": 1500,
 "temperature": 0.7,
 "topP": 0.9,
 "topK": 50,
 "stopSequences": ["THE END"],
 },
}

bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

Invoke the model
response = bedrock.invoke_model(
 modelId="us.amazon.nova-2-lite-v1:0", body=json.dumps(request_body)
)
response_body = json.loads(response["body"].read())

Extract the text response
content_list = response_body["output"]["message"]["content"]
for content in content_list:
 if "text" in content:
 print(content["text"])
```

## Streaming

```
import json

import boto3
from botocore.config import Config

Configure the request
request_body = {
 "messages": [
 {
 "role": "user",
 "content": [{"text": "Write a short story. End the story with 'THE
END'. "}],
 }
],
 "system": [{"text": "You are a children's book author."}], # Optional
 "inferenceConfig": { # These parameters are optional
 "maxTokens": 1500,
 "temperature": 0.7,
 "topP": 0.9,
 "topK": 50,
 "stopSequences": ["THE END"],
 },
}

bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(connect_timeout=3600, read_timeout=3600),
)

Invoke the model with streaming
response = bedrock.invoke_model_with_response_stream(
 modelId="us.amazon.nova-2-lite-v1:0", body=json.dumps(request_body)
)

Process the streaming response
for event in response["body"]:
 chunk = json.loads(event["chunk"]["bytes"])
 if "contentBlockDelta" in chunk:
 delta = chunk["contentBlockDelta"]["delta"]
 if "text" in delta:
```

```
print(delta["text"], end="", flush=True)
```

## InvokeModel API with reasoning

Use the InvokeModel API with reasoning enabled for complex problem-solving.

### Non-streaming

```
import json

import boto3
from botocore.config import Config

Configure the request
request_body = {
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "text": 'How many capital letters appear in the following
passage. Your response must include only the number: "Wilfred ordered an anvil from
ACME. Shipping was expensive."'
 }
],
 }
],
 "reasoningConfig": {
 "type": "enabled",
 "maxReasoningEffort": "low", # "low" | "medium" | "high"
 },
}

bedrock = boto3.client(
 "bedrock-runtime",
 region_name="us-east-1",
 config=Config(read_timeout=3600),
)

Invoke the model
response = bedrock.invoke_model(
 modelId="us.amazon.nova-2-lite-v1:0", body=json.dumps(request_body)
```

```

)
response_body = json.loads(response["body"].read())

Extract response content
content_list = response_body["output"]["message"]["content"]
for content in content_list:
 # Extract the reasoning response
 if "reasoningContent" in content:
 print("\n== Reasoning ==")
 print(content["reasoningContent"]["reasoningText"]["text"])
 # Extract the text response
 if "text" in content:
 print("\n== Text ==")
 print(content["text"])

```

## Streaming

```

import json

import boto3
from botocore.config import Config

Configure the request
request_body = {
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "text": 'How many capital letters appear in the following
passage. Your response must include only the number: "Wilfred ordered an anvil from
ACME. Shipping was expensive."'
 }
],
 }
],
 "reasoningConfig": {
 "type": "enabled",
 "maxReasoningEffort": "low", # "low" | "medium" | "high"
 },
}

bedrock = boto3.client(

```

```
"bedrock-runtime",
region_name="us-east-1",
config=Config(connect_timeout=3600, read_timeout=3600),
)

Invoke the model with streaming
response = bedrock.invoke_model_with_response_stream(
 modelId="us.amazon.nova-2-lite-v1:0", body=json.dumps(request_body)
)

Process the streaming response
for event in response["body"]:
 chunk = json.loads(event["chunk"]["bytes"])

 if "contentBlockDelta" in chunk:
 delta = chunk["contentBlockDelta"]["delta"]

 # Extract the reasoning response
 if "reasoningContent" in delta:
 print("\n== Reasoning ==")
 print(delta["reasoningContent"]["reasoningText"]["text"], end="",
flush=True)

 # Extract the text response
 if "text" in delta:
 print("\n== Text ==")
 print(delta["text"], end="", flush=True)
```

## Troubleshooting

This section provides solutions to common issues when working with Amazon Nova models.

### Authentication and setup

#### Missing permissions

**Symptoms:** Unable to access Nova models or features

**Solution:**

- Ensure your IAM role has AmazonBedrockFullAccess or appropriate permissions
- Request specific model access through the Amazon Bedrock console

- Verify permissions for model access and tool usage

## Model access denied

**Symptoms:** Model access request fails

**Solution:**

- Request specific model access through the Amazon Bedrock console
- Verify your account has been granted access to the requested model
- Check regional availability of the model

## Regional availability issues

**Symptoms:** Feature not available in selected region

**Solution:**

- Web Grounding is only available in US regions with US CRIS profiles
- Verify the model and features are available in your selected region
- Switch to a supported region if necessary

## Timeout configuration

**Symptoms:** Requests timing out before completion

**Cause:** Default timeout too short for complex operations

**Solution:** Configure extended timeout settings

```
from botocore.config import Config

bedrock = boto3.client(
 'bedrock-runtime',
 region_name='us-east-1',
 config=Config(
 read_timeout=3600 # 60 minutes
)
)
```

 **Note**

Amazon Nova inference requests can take up to 60 minutes for complex operations.

## API response issues

Understanding stop reasons:

`end_turn`

Normal completion. No action needed.

`max_tokens`

Token limit reached.

**Solution:** Increase `maxTokens` parameter in `inferenceConfig`.

`content_filtered`

Content violated AWS Responsible AI policy.

**Solution:** Review and modify your input to comply with content policies.

`malformed_model_output`

Invalid output format.

**Solution:** Check your output schema and constraints; verify JSON schema is properly formatted.

`malformed_tool_use`

Invalid tool call format.

**Solution:** Verify tool definitions match expected schema; check tool input parameters are correctly formatted.

`service_unavailable`

Built-in tool service unavailable.

**Solution:** Retry the request after a brief delay; check AWS service health dashboard.

## invalid\_query

Invalid query to built-in tool.

**Solution:** Review query format and parameters; ensure query meets tool requirements.

## max\_tool\_invocations

Tool retries exhausted.

**Solution:** Simplify the task or break it into smaller steps; review tool error messages for specific issues.

# Reasoning mode errors

Truncated responses with high reasoning effort

**Solution:** For high reasoning effort, unset these parameters: temperature, topP, maxToken. This allows the model to use optimal settings for complex reasoning tasks.

Insufficient tokens for reasoning

**Error:** "maxTokens is insufficient"

**Solution:** Automatically retry with increased limit

```
token_limits = {
 "low": 15000,
 "medium": 30000,
 "high": 50000
}

try:
 response = client.converse(
 modelId="us.amazon.nova-2-lite-v1:0",
 messages=messages,
 inferenceConfig={
 "maxTokens": token_limits[max_effort]
 },
 additionalModelRequestFields={
 "reasoningConfig": {
 "type": "enabled",
 "maxReasoningEffort": max_effort
 }
 }
)
```



```
 }
)
except Exception as e:
 if "maxTokens is insufficient" in str(e):
 higher_limit = int(token_limits[max_effort] * 1.5)
 # Retry with higher limit
```

## Tool use issues

### Schema validation failures

Tool schema validation errors

**Solution:**

- Limit JSON schemas to two layers of nesting for best performance
- Ensure all required fields are properly defined
- Validate schema against JSON Schema specification

Model not using tools correctly

**Solution:**

- Ensure tool name clearly describes its purpose
- Provide detailed description of tool functionality
- Explicitly define input schema with clear parameter descriptions
- Include examples in the description when helpful

Inconsistent tool calling behavior

**Solution:** Set temperature to 0 for tool calling:

```
inferenceConfig={
 "temperature": 0,
 "maxTokens": 10000
}
```

This enables greedy decoding for more reliable tool use.

Tool choice conflicts

**Problem:** Error when using custom tools with web search or code interpreter

**Solution:** Do not include a custom toolSpec with name `nova_grounding` - this conflicts with the system tool. Use the system tool configuration instead:

```
Correct - use system tool
tool_config = {
 "tools": [{
 "systemTool": {"name": "nova_grounding"}
 }]
}

Incorrect - don't create custom tool with this name
tool_config = {
"tools": [{
"toolSpec": {"name": "nova_grounding", ...}
}]
}
```

## Web Grounding issues

### Access control problems

**Problem:** Web Grounding and Code Interpreter not working

**Solution:** Ensure your IAM policy includes:

```
{
 "Statement": [
 {
 "Effect": "Allow",
 "Action": ["bedrock:InvokeTool"],
 "Resource": ["arn:aws:bedrock::{YOUR_ACCOUNT_ID}:system-tool/
amazon.nova_grounding"]
 }
]
}
```

### Service Control Policy issues

**Problem:** Web Grounding blocked by SCP

**Solution:** If you have Service Control Policies with `aws:requestedRegion` condition, update them to allow "unspecified" region for Web Grounding functionality.

## Media processing limitations

Poor understanding of multilingual content in images/videos

**Limitation:** Nova models have limited understanding of multilingual content in visual media

**Workaround:**

- Provide text translations alongside images
- Use text-based inputs for multilingual content when possible

People identification

**Problem:** Model refuses to identify people in images

**Expected Behavior:** Models will refuse to identify or name individuals in images, documents, or videos for privacy and safety reasons

**Workaround:** Ask about general characteristics or context instead of specific identities

Spatial reasoning limitations

**Problem:** Inaccurate localization or layout analysis

**Limitation:** Limited capabilities for precise spatial reasoning

**Workaround:**

- Use bounding box detection for object localization
- Provide clear reference points in your prompts
- Break complex spatial queries into simpler components

Small text in images/videos

**Problem:** Cannot read small text in media

**Solution:**

- Crop images to focus on relevant text sections
- Increase resolution of source media
- Provide text separately if available

## Document and file handling

### Unsupported content

**Problem:** PDF processing fails

**Causes:**

- PDFs with CMYK color profiles
- PDFs containing SVG images

**Solution:**

- Convert PDFs to RGB color profile
- Rasterize SVG images before including in PDFs

### Token estimation

**Problem:** Unexpected token usage with PDFs

**Guideline:** Estimate approximately 2,560 tokens per standard 8.5×11" PDF page

**Solution:** Adjust maxTokens accordingly based on document length

# API and SDK reference

This section provides complete technical reference documentation for Amazon Nova APIs and SDKs.

## Topics

- [Request and response schema](#)
- [SDK reference](#)

## Request and response schema

The request schema is nearly identical between the Invoke API and Converse API. The primary difference is how binary data (images, video, audio) is encoded: the Converse API uses binary arrays while the Invoke API uses Base64-encoded strings.

## Complete request structure

The following shows the complete request structure for Amazon Nova models. All fields are optional unless marked as required:

```
{
 "system": [
 {
 "text": "string"
 }
],
 "messages": [// Required
 {
 "role": "user", // Required - first turn must be user
 "content": [// Required
 {
 "text": "string"
 },
 {
 "image": {
 "format": "jpeg" | "png" | "gif" | "webp", // Required
 "source": { // Required
 "bytes": image // Binary array (Converse) or Base64 string (Invoke)
 }
 }
 }
]
 }
]
}
```

```

 }
 },
 {
 "video": {
 "format": "mkv" | "mov" | "mp4" | "webm" | "three_gp" | "flv" | "mpeg" |
"mpg" | "wmv",
 "source": {
 // Option 1: S3 location
 "s3Location": {
 "uri": "string", // e.g., s3://my-bucket/object-key
 "bucketOwner": "string" // Optional, e.g., "123456789012"
 },
 // Option 2: File bytes
 "bytes": video // Binary array (Converse) or Base64 string (Invoke)
 }
 }
 },
 {
 "audio": { // Nova 2 Omni and Sonic only
 "format": "mp3" | "opus" | "wav" | "aac" | "flac" | "mp4" | "ogg" | "mkv",
 "source": {
 // Option 1: S3 location
 "s3Location": {
 "uri": "string",
 "bucketOwner": "string" // Optional
 },
 // Option 2: File bytes
 "bytes": audio // Binary array (Converse) or Base64 string (Invoke)
 }
 }
 }
],
{
 "role": "assistant",
 "content": [
 {
 "text": "string" // For prefilling assistant response
 }
]
}
],
"inferenceConfig": { // All optional
 "maxTokens": int, // 1-5000, default: dynamic

```

```

 "temperature": float, // 0.00001-1, default: 0.7
 "topP": float, // 0-1, default: 0.9
 "topK": int, // 0-128, default: not used
 "stopSequences": ["string"],
 "reasoningConfig": { // Nova 2 Lite and Sonic only
 "type": "enabled" | "disabled", // default: "disabled"
 "maxReasoningEffort": "low" | "medium" | "high"
 }
 },
 "toolConfig": { // Optional
 "tools": [
 {
 "toolSpec": {
 "name": "string", // Max 64 characters
 "description": "string",
 "inputSchema": {
 "json": {
 "type": "object",
 "properties": {
 "arg1": {
 "type": "string",
 "description": "string"
 }
 }
 },
 "required": ["string"]
 }
 }
 }
]
 },
 "toolChoice": { // Choose one option
 "auto": {},
 "any": {},
 "tool": {
 "name": "string"
 }
 }
}

```

### Key request parameters:

- **system:** System prompt providing context and instructions

- **messages:** Array of conversation turns with role (user or assistant) and content
- **inferenceConfig:** Controls model output behavior (temperature, tokens, and so on.)
- **toolConfig:** Tool specifications for function calling

### **Note**

When using the Converse API, the `topK` and `reasoningConfig` parameters must be placed in `additionalModelRequestFields` instead of `inferenceConfig`.

The following sections provide detailed explanations of each request parameter:

## **system**

**system** – (Optional) The system prompt for the request. A system prompt provides context and instructions to Amazon Nova, such as specifying a particular goal or role.

## **messages**

**messages** – (Required) The input messages array containing conversation turns.

- **role** – (Required) The role of the conversation turn. Valid values are `user` and `assistant`. The first message must always use the `user` role.
- **content** – (Required) An array of content blocks. Each block specifies a content type (`text`, `image`, `video`, or `audio`):
  - **text** – Text content for the conversation turn. If combined with `image` or `video`, interpreted as accompanying text.
  - **image** – (Not supported for Nova 2 Lite) Image content with:
    - **format** – (Required) Image format: `jpeg`, `png`, `webp`, or `gif`
    - **source.bytes** – (Required) Image data as binary array (Converse API) or Base64 string (Invoke API)
  - **video** – (Not supported for Nova 2 Lite) Video content with:
    - **format** – (Required) Video format: `mkv`, `mov`, `mp4`, `webm`, `three_gp`, `flv`, `mpeg`, `mpg`, or `wmv`
    - **source** – (Required) Video source via S3 URI (`s3Location.uri` and optional `bucketOwner`) or file bytes (`bytes`)
  - **audio** – (Amazon Nova Sonic and only) Audio content with:



- **format** – (Required) Audio format: mp3, opus, wav, aac, flac, mp4, ogg, or mkv
- **source** – (Required) Audio source via S3 URI or file bytes

## **inferenceConfig**

**inferenceConfig** – (Optional) Configuration parameters controlling model output generation.

- **maxTokens** – (Optional) Maximum tokens to generate before stopping. Amazon Nova models may stop before reaching this limit. Maximum value is 5,000. If not specified, uses a dynamic default based on request context.
- **temperature** – (Optional) Randomness in responses. Valid range: 0.00001-1 (default: 0.7). Lower values produce more deterministic output.
- **topP** – (Optional) Nucleus sampling threshold. Amazon Nova samples from tokens whose cumulative probability reaches topP. Valid range: 0-1 (default: 0.9). Adjust either temperature or topP, not both.
- **topK** – (Optional) Sample from top K tokens only. Removes low-probability responses. Valid range: 0-128 (default: not used).

### **Note**

For Converse API, pass **topK** in `additionalModelRequestFields`.

- **stopSequences** – (Optional) Array of strings that stop generation when encountered.
- **reasoningConfig** – (Amazon Nova Sonic and only) Reasoning configuration:
  - **type** – (Optional) enabled or disabled (default: disabled)
  - **maxReasoningEffort** – Computational effort: low, medium, or high. With low and medium, reasoning streams incrementally; high outputs reasoning in a final chunk.

### **Note**

For Converse API, pass **reasoningConfig** in `additionalModelRequestFields`.

## **toolConfig**

**toolConfig** – (Optional) Tool configuration following [ToolConfiguration schema](#).

- `tools` – Array of tool specifications with name, description and `inputSchema`
- `toolChoice` – (Optional) Controls tool selection:
  - `auto` – Model decides whether and which tools to use
  - `any` – Model must use at least one tool
  - `tool` – Model must use the specified tool by name

## Complete response structure

The following shows the complete response structure for Amazon Nova models:

```
{
 "ResponseMetadata": {
 "RequestId": "string",
 "HTTPStatusCode": int,
 "HTTPHeaders": {
 "date": "string",
 "content-type": "application/json",
 "content-length": "string",
 "connection": "keep-alive",
 "x-amzn-requestid": "string"
 },
 "RetryAttempts": 0
 },
 "output": {
 "message": {
 "role": "assistant",
 "content": [
 {
 "reasoningContent": { // Optional - if reasoning enabled
 "reasoningText": {
 "text": "[REDACTED]"
 }
 }
 }
],
 {
 "toolUse": { // Optional - if tool called
 "toolUseId": "string",
 "name": "string",
 "input": {} // Tool-specific arguments
 }
 }
],
 },
}
```

```

 {
 "text": "string" // Optional - text response
 },
 {
 "image": { // Optional - Nova 2 Omni only
 "format": "png",
 "source": {
 "bytes": image // Binary array (Converse) or Base64 string (Invoke)
 }
 }
 }
]
}
},
"stopReason": "string", // See stop reasons below
"usage": {
 "inputTokens": int,
 "outputTokens": int,
 "totalTokens": int
},
"metrics": {
 "latencyMs": int
}
}

```

### Stop reasons:

- **end\_turn**: Natural end of response
- **max\_tokens**: Reached maxTokens limit
- **content\_filtered**: Violated content policy
- **malformed\_model\_output**: Invalid model output
- **malformed\_tool\_use**: Invalid tool use output
- **service\_unavailable**: Built-in tool service unreachable
- **invalid\_query**: Invalid query to built-in tool
- **max\_tool\_invocations**: Tool retries exhausted

The following sections provide detailed explanations of each response field:

## output

**output** – (Required) Contains the model's response message.

- **message** – (Required) The assistant's response message with role and content array.
- **content** – (Required) Array of content blocks that can include:
  - **reasoningContent** – (Optional) Returned if reasoning was enabled. Contains reasoning text, which will always be [REDACTED] in the response.
  - **toolUse** – (Optional) Returned if a tool was called. Contains tool use ID, name and input arguments.
  - **text** – (Optional) Returned if the model responded with text content.
  - **image** – (Optional, only) Returned if the model generated an image. Format will always be PNG.

## stopReason

**stopReason** – (Required) Indicates why the model stopped generating output:

- **end\_turn** – Natural end of response reached
- **max\_tokens** – Reached maxTokens limit or model's maximum output limit
- **content\_filtered** – Output violated AWS Responsible AI policy
- **malformed\_model\_output** – Model produced invalid output
- **malformed\_tool\_use** – Model produced invalid tool use output
- **service\_unavailable** – Built-in tool service could not be reached
- **invalid\_query** – Query to built-in tool was invalid
- **max\_tool\_invocations** – Built-in tool did not produce valid result after retries

## usage

**usage** – (Required) Token usage information:

- **inputTokens** – Total tokens ingested by the model
- **outputTokens** – Number of tokens generated
- **totalTokens** – Sum of input and output tokens

## metrics

metrics – (Required) Performance metrics:

- latencyMs – Total inference completion time in milliseconds

## SDK reference

AWS software development kits (SDKs) are available for many popular programming languages. Each SDK provides an API, code examples and documentation that make it easier for developers to build applications in their preferred language.

| SDK documentation                          | Code examples                                            |
|--------------------------------------------|----------------------------------------------------------|
| <a href="#">AWS SDK for C++</a>            | <a href="#">AWS SDK for C++ code examples</a>            |
| <a href="#">AWS CLI</a>                    | <a href="#">AWS CLI code examples</a>                    |
| <a href="#">AWS SDK for Go</a>             | <a href="#">AWS SDK for Go code examples</a>             |
| <a href="#">AWS SDK for Java</a>           | <a href="#">AWS SDK for Java code examples</a>           |
| <a href="#">AWS SDK for JavaScript</a>     | <a href="#">AWS SDK for JavaScript code examples</a>     |
| <a href="#">AWS SDK for Kotlin</a>         | <a href="#">AWS SDK for Kotlin code examples</a>         |
| <a href="#">AWS SDK for .NET</a>           | <a href="#">AWS SDK for .NET code examples</a>           |
| <a href="#">AWS SDK for PHP</a>            | <a href="#">AWS SDK for PHP code examples</a>            |
| <a href="#">AWS Tools for PowerShell</a>   | <a href="#">AWS Tools for PowerShell code examples</a>   |
| <a href="#">AWS SDK for Python (Boto3)</a> | <a href="#">AWS SDK for Python (Boto3) code examples</a> |
| <a href="#">AWS SDK for Ruby</a>           | <a href="#">AWS SDK for Ruby code examples</a>           |
| <a href="#">AWS SDK for Rust</a>           | <a href="#">AWS SDK for Rust code examples</a>           |
| <a href="#">AWS SDK for SAP ABAP</a>       | <a href="#">AWS SDK for SAP ABAP code examples</a>       |
| <a href="#">AWS SDK for Swift</a>          | <a href="#">AWS SDK for Swift code examples</a>          |

# Customizing Amazon Nova models on SageMaker AI

You can customize Amazon Nova models, including the enhanced Amazon Nova 2.0 models, through [recipes](#) and train them on SageMaker. These recipes support techniques such as supervised fine-tuning (SFT) and Reinforcement Fine-Tuning (RFT), with both full-rank and low-rank adaptation (LoRA) options.

The end-to-end customization workflow involves stages like model training, model evaluation, and deployment for inference. This model customization approach on SageMaker provides greater flexibility and control to fine-tune its supported Amazon Nova models, optimize hyperparameters with precision, and implement techniques such as LoRA parameter-efficient fine-tuning (PEFT), full-rank SFT, RFT, and Continued Pre-Training (CPT).

## Customization approaches

SageMaker offers two approaches for customizing Amazon Nova models:

**UI-based experience** – Use to customize Amazon Nova models through a simple, guided interface. This approach provides an end-to-end workflow including training, evaluation, and deployment without writing code. The UI-based experience is ideal for rapid experimentation, proof-of-concept development, and users who prefer a visual workflow.

**Code-based experience** – Use the SageMaker Python SDK, Nova SDK and training recipes to customize models programmatically. This approach offers greater flexibility, allowing you to configure advanced hyperparameters, integrate with CI/CD pipelines, and automate training workflows. The code-based experience is recommended for production workloads, complex customization requirements, and teams with established MLOps practices.

| Approach   | Best for                                        | Key benefits                                            |
|------------|-------------------------------------------------|---------------------------------------------------------|
| UI-based   | Experimentation, prototyping, quick iterations  | Simple setup, guided workflow, no coding required       |
| Code-based | Production, automation, advanced configurations | Full flexibility, pipeline integration, version control |

## Customization platforms

AWS offers three platforms for customizing Amazon Nova models, each designed for different use cases and requirements:

**Amazon Bedrock** – Provides the easiest and fastest path to model customization with minimal setup. Bedrock handles all infrastructure management automatically, allowing you to focus on your data and use case. This platform is ideal when you need the quickest time-to-value and prefer a fully managed experience.

**SageMaker training jobs** – Provides a fully managed environment for customizing Amazon Nova models where you don't need to create or maintain any clusters. The service automatically handles all infrastructure provisioning, scaling, and resource management, allowing you to focus on configuring your training parameters and submitting your job. This platform offers a balance between ease of use and flexibility, supporting techniques like Parameter Efficient Fine-tuning (PEFT), Full rank fine tuning, and Reinforcement Fine-Tuning (RFT).

**SageMaker HyperPod** – Offers a specialized environment for large-scale distributed training by requiring you to create and manage EKS clusters with restricted instance groups (RIGs). This platform gives you maximum flexibility in configuring your training environment with specialized GPU instances and integrated Amazon FSx for Lustre storage, making it particularly well-suited for advanced distributed training scenarios, ongoing model development, and enterprise-scale customization workloads.

| Platform                | Complexity | Flexibility | Best for                                         |
|-------------------------|------------|-------------|--------------------------------------------------|
| Amazon Bedrock          | Lowest     | Standard    | Fastest customization, minimal setup             |
| SageMaker training jobs | Medium     | High        | Balanced flexibility and ease of use             |
| SageMaker HyperPod      | Highest    | Maximum     | Large distributed training, enterprise workloads |

**Note**

If you provide a KMS key to your Amazon Nova model customization training job for encryption in the Amazon-owned output S3 bucket:

- You must provide the same KMS key when calling subsequent iterative training jobs, or when calling the Amazon Bedrock CreateCustomModel API leveraging the encrypted model.
- The identity calling the CreateTrainingJob API (rather than the execution role) must have permissions to CreateGrant, RetireGrant, Encrypt, and GenerateDataKey as defined in KMS key policy.

## General prerequisites

The customization process involves several key stages including model training, evaluation, and deployment for inference, each requiring specific resources and configurations. Before beginning your Amazon Nova model customization on SageMaker, ensure you have the following general prerequisites.

- An AWS account. If you don't have an AWS account, follow [these instructions](#) to sign up for one.
- [Installing the AWS CLI](#) and [Setting up the configuration](#).
- Access to the [base Amazon Nova model customization recipes](#).
- Familiarity of [YAML](#) configuration files.
- Familiarity of how to run a notebook in your environment.
- Familiarity of how to create AWS resources like [Amazon S3 buckets](#) and [IAM roles with appropriate permissions](#).
- Familiarity of how to [train a model with SageMaker](#).
- Familiarity of [SageMaker HyperPod with EKS orchestration](#).
- Familiarity of [SageMaker HyperPod CLI](#).
- Familiarity of [Amazon Nova foundational models](#).
- Familiarity of [available Amazon Nova models and algorithms for customization](#).
- Familiarity of [Amazon Bedrock inference](#).



# Amazon Nova recipes

You can get an Amazon Nova recipe from the [recipes](#) repository. An Amazon Nova recipe is a [YAML](#) configuration file that provides details to SageMaker on how to run your model customization job. It provides the base model name, sets training hyperparameters, defines optimization settings, and includes any additional options required to fine-tune or train the model successfully.

You can also access Amazon Nova recipes through SageMaker AI Monarch and by navigating to the model hub, selecting AWS, and browsing Amazon Nova models to find their associated recipes. Both SageMaker AI Monarch and provide sample notebooks for each recipe, which include all the necessary steps to modify the recipes and run customization jobs using SageMaker training jobs or SageMaker HyperPod environments.

To access the recipes pages in Amazon SageMaker Studio, the execution role must have the following permissions.

JSON

```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": [
 "s3:GetObject"
],
 "Resource": [
 "arn:aws:s3::*model-customization-recipes*"
]
 }
]
}
```

To execute the sample notebooks on SageMaker training jobs and , use one of the following SageMaker distribution image versions: 2.7.1+, 2.8.0+, 3.2.1+, 3.3.0+. This applies to both SageMaker AI Monarch and .

## Topics

- [Getting Amazon Nova recipes](#)

- [Getting Amazon Nova Forge recipes](#)
- [Available models and algorithms](#)
- [Amazon Nova 2.0 Lite](#)

## Getting Amazon Nova recipes

To get a base Amazon Nova recipe, clone the [recipes](#) repository by running the following command.

```
git clone https://github.com/aws/sagemaker-hyperpod-recipes.git
```

The base recipes are available at [recipes\\_collection/recipes/](#).

```
cd recipes_collection/recipes/
```

The Amazon Nova customization recipes are in the following folders.

| Recipe type              | Folder                           |
|--------------------------|----------------------------------|
| SFT (Full-rank and PEFT) | <a href="#">fine-tuning/nova</a> |
| Evaluation               | <a href="#">evaluation/nova</a>  |
| CPT                      | <a href="#">training/nova</a>    |

## Getting Amazon Nova Forge recipes

To access specialized Amazon Nova Forge recipes for jobs, please [set up your access to Amazon Nova Forge](#) and then follow [this workshop](#) to set up the Forge version of the CLI. For SageMaker Training Jobs, the AWS Console and Nova Forge SDK will automatically access Forge recipes.

## Available models and algorithms

The following table summarizes the availability of customization for Amazon Nova models and supported algorithms using SageMaker.

| Model name           | Model ID                     | Fine-tuning | Notes                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|------------------------------|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Amazon Nova Lite 2.0 | amazon.nova-2-lite-v1:0:256k | Yes         | <ul style="list-style-type: none"><li>• SFT - Accepts either text and/or image or text and/or video as the input and produces text as output. A single job can't combine text, image, and video in the same run.</li><li>• RFT - Accepts text and image as input for single-turn conversations and produces text as output. Improved reward optimization capabilities.</li><li>• CPT</li></ul> |

## Amazon Nova 2.0 Lite

The table below lists detailed information of the Amazon Nova 2.0 Lite recipes reference.

| Model         | Category Sub-category | Technique              | Recipe Name        | Image URI (Training Jobs) | Image URI (SageMaker HyperPod) | Compute Instance |
|---------------|-----------------------|------------------------|--------------------|---------------------------|--------------------------------|------------------|
| Nova 2.0 Lite | Training              | Continued Pre Training | nova_lite_2_0_p5x8 | N/A                       | 70897720f387.dkr.ecr.us-       | m1.p5.48xlarge , |

| Model         | Category                        | Technique   | Recipe Name                        | Image URI (Training Jobs)                                                            | Image URI (SageMaker HyperPod)                                                       | Compute Instance                  |
|---------------|---------------------------------|-------------|------------------------------------|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|-----------------------------------|
|               |                                 |             | _gpu_pretrain.yaml                 |                                                                                      | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-CPT-V2-latest | ml.p5en.48xlarge                  |
| Nova 2.0 Lite | Parameter Efficient Fine Tuning | Fine Tuning | nova_lite_2_0_p5_gpu_lora_sft.yaml | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-TJ-SFT-V2-latest | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-SFT-V2-latest | ml.p5.48xlarge , ml.p5en.48xlarge |

| Model               | Category<br>Sub-category        | Technique      | Recipe Name                           | Image URI<br>(Training<br>Jobs)                                                      | Image<br>URI<br>(SageMaker<br>HyperPod)                                              | Compute<br>Instance                  |
|---------------------|---------------------------------|----------------|---------------------------------------|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|--------------------------------------|
| Nova<br>2.0<br>Lite | Full-<br>Rank<br>Fine<br>Tuning | Fine<br>Tuning | nova_lite<br>_2_0_p5_g<br>pu_sft.yaml | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-TJ-SFT-V2-latest | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-SFT-V2-latest | ml.p5.48xlarge ,<br>ml.p5en.48xlarge |

| Model         | Category                | Technique   | Recipe Name                                 | Image URI (Training Jobs)                                                            | Image URI (SageMaker HyperPod)                                                             | Compute Instance                  |
|---------------|-------------------------|-------------|---------------------------------------------|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|-----------------------------------|
| Nova 2.0 Lite | Parameter Efficient RFT | Fine Tuning | nova_lite_v2_smtj_p5_p5en_gpu_lora_rft.yaml | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-TJ-RFT-V2-latest | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-RFT-TRAIN-V2-latest | ml.p5.48xlarge , ml.p5en.48xlarge |

| Model         | Category      | Technique   | Recipe Name                         | Image URI (Training Jobs)                                                            | Image URI (SageMaker HyperPod)                                                             | Compute Instance                 |
|---------------|---------------|-------------|-------------------------------------|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|----------------------------------|
| Nova 2.0 Lite | Full Rank RFT | Fine Tuning | nova_lite_v2_smtj_p5en_gpu_rft.yaml | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-TJ-RFT-V2-latest | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-RFT-TRAIN-V2-latest | m1.p5.48xlarge, ml.p5en.48xlarge |

| Model               | Category<br>Sub-category     | Technique  | Recipe Name                                                                    | Image URI<br>(Training<br>Jobs)                                                                                                | Image<br>URI<br>(SageMaker<br>HyperPod)                                                                                            | Compute<br>Instance                          |
|---------------------|------------------------------|------------|--------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| Nova<br>2.0<br>Lite | Bring<br>Your<br>Own<br>Data | Evaluation | nova_lite<br>_2_0_p5_4<br>8x1_gpu_b<br>ring_your<br>_own_data<br>set_eval.yaml | 708977205<br>387.dkr.ecr.us-<br>eas<br>t-1.amazonaws.com/<br>nova-<br>eval<br>uation-<br>repo:SM-<br>TJ-Eval-<br>V2-1<br>atest | 708977205<br>387.dkr.ecr.us-<br>eas<br>t-1.amazonaws.com/<br>nova-<br>eval<br>uation-<br>repo:SM-<br>HP-<br>Eval-<br>V2-1<br>atest | ml.p5.48x<br>large ,<br>ml.p5en.4<br>8xlarge |



| Model               | Category                     | Technique  | Recipe Name                                                                    | Image URI<br>(Training<br>Jobs)                                                                                                        | Image<br>URI<br>(SageMaker<br>HyperPod)                                                                                                        | Compute<br>Instance                          |
|---------------------|------------------------------|------------|--------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| Nova<br>2.0<br>Lite | General<br>Text<br>Benchmark | Evaluation | nova_lite<br>_2_0_p5_4<br>8x1_gpu_g<br>eneral_te<br>xt_benchm<br>ark_eval.yaml | 708977205<br>387.dkr.e<br>cr.us-<br>eas<br>t-1.amazo<br>naws.com/<br>nova-<br>eval<br>uation-<br>repo:SM-<br>TJ-Eval-<br>V2-1<br>atest | 708977205<br>387.dkr.e<br>cr.us-<br>eas<br>t-1.amazo<br>naws.com/<br>nova-<br>eval<br>uation-<br>re<br>po:SM-<br>HP-<br>Eval-<br>V2-1<br>atest | m1.p5.48x<br>large ,<br>m1.p5en.4<br>8xlarge |

| Model         | Category Sub-category | Technique  | Recipe Name                             | Image URI (Training Jobs)                                                              | Image URI (SageMaker HyperPod)                                                         | Compute Instance                  |
|---------------|-----------------------|------------|-----------------------------------------|----------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|-----------------------------------|
| Nova 2.0 Lite | RFT Evaluation        | Evaluation | nova_lite_2_0_p5_48x1_gpu_rft_eval.yaml | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM-TJ-Eval-V2-1atest | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM-HP-Eval-V2-1atest | ml.p5.48xlarge , ml.p5en.48xlarge |

| Model         | Category               | Technique  | Recipe Name                                          | Image URI (Training Jobs)                                                               | Image URI (SageMaker HyperPod)                                                          | Compute Instance                 |
|---------------|------------------------|------------|------------------------------------------------------|-----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|----------------------------------|
| Nova 2.0 Lite | LLM Rubric Based Judge | Evaluation | nova_lite_2_0_p5_48x1_gpu_rubric_llm_judge_eval.yaml | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-eval-uation-repo:SM-TJ-Eval-V2-latest | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-eval-uation-repo:SM-HP-Eval-V2-latest | ml.p5.48xlarge, ml.p5en.48xlarge |

## Amazon Nova customization on SageMaker Training Jobs

SageMaker Training Jobs is an environment that enables you to train machine learning models at scale. It automatically provisions and scales compute resources, loads training data from sources like Amazon S3, executes your training code, and stores the resulting model artifacts.

The purpose of training is to customize the base Amazon Nova model using your proprietary data. The training process typically involves steps to prepare your data, choose a [recipe](#), modify configuration parameters in YAML files, and submit a training job. The training process will output trained model checkpoint in a service-managed Amazon S3 bucket. You can use this checkpoint location for evaluation jobs. Nova customization on SageMaker AI training jobs stores model artifacts in a service-managed Amazon S3 bucket. Artifacts in the service-managed bucket are encrypted with SageMaker AI-managed KMS keys. Service-managed Amazon S3 buckets don't currently support data encryption using customer-managed KMS keys.

# Best Practices for Amazon Nova customization

## Overview

This section provides an overview of customization techniques and helps you choose the best approach for your needs and available data.

### Two stages of LLM training

Large language model training consists of two major stages: pre-training and post-training. During pre-training, the model processes tokens of raw text and optimizes for next-token prediction. This process creates a pattern completer that absorbs syntax, semantics, facts, and reasoning patterns from web and curated text. However, the pre-trained model doesn't understand instructions, user goals, or context-appropriate behavior. It continues text in whatever style fits its training distribution. A pre-trained model autocompletes rather than follows directions, produces inconsistent formatting, and can mirror undesirable biases or unsafe content from the training data. Pre-training builds general competence, not task usefulness.

Post-training transforms the pattern completer into a useful assistant. You run multiple rounds of Supervised Fine-Tuning (SFT) to teach the model to follow instructions, adhere to schemas and policies, call tools, and produce reliable outputs by imitating high-quality demonstrations. This alignment teaches the model to respond to prompts as tasks rather than text to continue. You then apply Reinforcement Fine-Tuning (RFT) to optimize behavior using measurable feedback (such as verifiers or an LLM-as-a-judge), balancing trade-offs like accuracy versus brevity, safety versus coverage, or multi-step reasoning under constraints. In practice, you alternate SFT and RFT in cycles to shape the pre-trained model into a reliable, policy-aligned system that performs complex tasks consistently.

### Choose the right customization approach

In this section we will cover post training customization strategies: RFT and SFT.

#### Reinforcement fine-tuning (RFT)

Reinforcement fine-tuning improves model performance through feedback signals—measurable scores or rewards that indicate response quality—rather than direct supervision with exact correct answers. Unlike traditional supervised fine-tuning that learns from input-output pairs, RFT uses reward functions to evaluate model responses and iteratively optimizes the model to maximize these rewards. This approach works well for tasks where defining the exact correct output is challenging, but you can reliably measure response quality. RFT enables models to learn complex

behaviors and preferences through trial and feedback, making it ideal for applications that require nuanced decision-making, creative problem-solving, or adherence to specific quality criteria that you can programmatically evaluate. For example, answering complex legal questions is an ideal use case for RFT because you want to teach the model how to reason better to answer questions more accurately.

## How it works

In reinforcement fine-tuning, you start from an instruction-tuned baseline and treat each prompt like a small tournament. For a given input, you sample a handful of candidate answers from the model, score each one with the reward function, then rank them within that group. The update step nudges the model to make higher-scoring candidates more likely next time and lower-scoring ones less likely, while a stay-close-to-baseline constraint keeps behavior from drifting or becoming verbose or exploitative. You repeat this loop over many prompts, refreshing hard cases, tightening verifiers or judge rubrics when you see exploits, and continuously tracking task metrics.

## When to use RFT

Tasks that benefit most from RFT share several traits. They have measurable success signals even when a single correct output is hard to specify. They admit partial credit or graded quality so you can rank better versus worse answers within a prompt or using a reward function. They involve multiple objectives that must be balanced (such as accuracy with brevity, clarity, safety, or cost). They require adherence to explicit constraints that you can programmatically check. They operate in tool-mediated or environment-based settings where outcomes are observable (success or failure, latency, resource use). They occur in low-label regimes where collecting gold targets is expensive but automated or rubric-based feedback is plentiful. RFT works best when you can turn quality into a reliable scalar or ranking and want the model to preferentially amplify higher-scoring behaviors without needing exhaustive labeled targets.

### Consider other methods when:

- You have plentiful, reliable labeled input-output pairs – Use SFT
- The main gap is knowledge or jargon – Use retrieval-augmented generation (RAG)
- Your reward signal is noisy or unreliable and you can't fix it with better rubrics or checkers – Stabilize that first before RFT

## When not to use RFT

Avoid RFT in these situations:

- You can cheaply produce reliable labeled input-output pairs (SFT is simpler, cheaper, and more stable)
- The gap is knowledge or jargon rather than behavior (use RAG)
- Your reward signal is noisy, sparse, easy to game, or expensive or slow to compute (fix the evaluator first)
- Baseline performance is near-zero (bootstrap with SFT before optimizing preferences)
- The task has deterministic schemas, strict formatting, or a single correct answer (SFT or rule-based validation works better)
- Tight latency or cost budgets can't absorb the extra sampling or exploration RFT requires
- Safety or policy constraints aren't crisply specified and enforceable in the reward

If you can point to "the right answer," use SFT. If you need new knowledge, use RAG. Use RFT only after you have a solid baseline and a robust, fast, hard-to-exploit reward function.

## **Supervised fine-tuning (SFT)**

Supervised fine-tuning trains the LLM on a dataset of human-labeled input-output pairs for your task. You provide examples of prompts (questions, instructions, and so on) with the correct or desired responses, and continue training the model on these examples. The model adjusts its weights to minimize a supervised loss (typically cross-entropy between its predictions and the target output tokens). This is the same training used in most supervised machine learning tasks, applied to specialize an LLM.

SFT changes behavior, not knowledge. It doesn't teach the model new facts or jargon it didn't see in pre-training. It teaches the model how to answer, not what to know. If you need new domain knowledge (such as internal terminology), use retrieval-augmented generation (RAG) to provide that context at inference time. SFT then adds the desired instruction-following behavior on top.

### **How it works**

SFT optimizes LLM by minimizing the average cross-entropy loss on response tokens, treating prompt tokens as context and masking them from the loss. The model internalizes your target style, structure, and decision rules, learning to generate the correct completion for each prompt. For example, to classify documents into custom categories, you fine-tune the model with prompts (the document text) and labeled completions (the category labels). You train on those pairs until the model outputs the right label for each prompt with high probability.

You can perform SFT with as few as a few hundred examples and scale up to a few hundred thousand. SFT samples must be high quality and directly aligned with the desired model behavior.

## When to use SFT

Use SFT when you have a well-defined task with clear desired outputs. If you can explicitly state "Given X input, the correct output is Y" and gather examples of such mappings, supervised fine-tuning is a good choice. SFT excels in these scenarios:

- **Structured or complex classification tasks** – Classify internal documents or contracts into many custom categories. With SFT, the model learns these specific categories better than prompting alone.
- **Question-answering or transformation tasks with known answers** – Fine-tune a model to answer questions from a company's knowledge base, or convert data between formats where each input has a correct response.
- **Formatting and style consistency** – Train the model to always respond in a certain format or tone by fine-tuning on examples of the correct format or tone. For instance, training on prompt-response pairs that demonstrate a particular brand voice teaches the model to generate outputs with that style. Instruction-following behavior is often initially taught through SFT on curated examples of good assistant behavior.

SFT is the most direct way to teach an LLM a new skill or behavior when you can specify what the right behavior looks like. It uses the model's existing language understanding and focuses it on your task. Use SFT when you want the model to do a specific thing and you have or can create a dataset of examples.

Use SFT when you can assemble high-quality prompt and response pairs that closely mirror the behavior you want. It fits tasks with clear targets or deterministic formats such as schemas, function or tool calls, and structured answers where imitation is an appropriate training signal. The goal is behavior shaping: teaching the model to treat prompts as tasks, follow instructions, adopt tone and refusal policies, and produce consistent formatting. Plan for at least hundreds of demonstrations, with data quality, consistency, and deduplication mattering more than raw volume. For a straightforward, cost-efficient update, use parameter-efficient methods like Low-Rank Adaptation to train small adapters while leaving most of the backbone untouched.

## When not to use SFT

Don't use SFT when the gap is knowledge rather than behavior. It doesn't teach the model new facts, jargon, or recent events. In those cases, use retrieval-augmented generation to bring external knowledge at inference. Avoid SFT when you can measure quality but can't label a single right answer. Use reinforcement fine-tuning with verifiable rewards or an LLM-as-a-judge to optimize those rewards directly. If your needs or content change frequently, rely on retrieval and tool use rather than retraining the model.

## Topics

- [Nova Forge SDK](#)
- [Training for Amazon Nova models](#)
- [Monitoring Progress Across Iterations](#)
- [Evaluating your SageMaker AI-trained model](#)
- [Iterative training](#)

## Nova Forge SDK

The Nova Forge SDK is a comprehensive Python SDK for customizing Amazon Nova models. The SDK provides a unified interface for training, evaluation, monitoring, deployment, and inference of Amazon Nova models across different platforms including SageMaker AI and Amazon Bedrock. Whether you're adapting models to domain-specific tasks or optimizing performance for your use case, this SDK provides everything you need in one unified interface.

## Benefits

- One SDK for the entire model customization lifecycle—from data preparation to deployment and monitoring.
- Support for multiple training methods including continued pre-training (CPT), supervised fine-tuning (SFT), direct preference optimization (DPO), and reinforcement fine-tuning (RFT), both single-turn and multi-turn, with both LoRA and full-rank approaches.
- Built-in support for SageMaker Training Jobs, SageMaker HyperPod, and Amazon Bedrock, with automatic resource management.
- No more finding the right recipes or container URI for your training techniques.
- Bring your own training recipes or use the SDK's intelligent defaults with parameter overrides.



- The SDK validates your configuration against supported model and instance combinations and provides validation support, preventing errors before training starts.
- Integrated Amazon CloudWatch monitoring enables you to track training progress in real-time.
- Integrated MLFlow to track training experiments with SageMaker AI MLFlow tracking servers.

## Requirements

The SDK requires at least Python 3.12.

## Installation

To install this SDK, please follow below command.

```
pip install amzn-nova-forge
```

## Supported Models and Techniques

The SDK supports the following models and techniques within the Amazon Nova family:

| Method                                         | Supported Models                                     |
|------------------------------------------------|------------------------------------------------------|
| Continued Pre-training                         | <a href="#">All Nova Models</a> (SMHP only)          |
| Supervised Fine-tuning LoRA                    | <a href="#">All Nova Models</a>                      |
| Supervised Fine-tuning Full-Rank               | <a href="#">All Nova Models</a> (SMHP and SMTJ only) |
| Direct Preference Optimization LoRA            | Nova 1.0 models (SMHP and SMTJ only)                 |
| Direct Preference Optimization Full-Rank       | Nova 1.0 models (SMHP and SMTJ only)                 |
| Reinforcement Fine-tuning LoRA                 | Nova Lite 2.0                                        |
| Reinforcement Fine-tuning Full-Rank            | Nova Lite 2.0 (SMHP and SMTJ only)                   |
| Multi-turn Reinforcement Fine-tuning LoRA      | Nova Lite 2.0 (SMHP Only)                            |
| Multi-turn Reinforcement Fine-tuning Full-Rank | Nova Lite 2.0 (SMHP Only)                            |

# Getting Started

## Topics

- [1. Prepare Your Data](#)
- [2. Configure Your Infrastructure](#)
- [3. Train](#)
- [4. Monitor](#)
- [5. Evaluate](#)
- [6. Deploy](#)

## 1. Prepare Your Data

Load your dataset from local files or S3, and let the SDK handle the transformation to the correct format for your chosen training method. Or, provide formatted data and get started immediately.

```
from amzn_nova_forge.dataset.dataset_loader import JSONLDatasetLoader
from amzn_nova_forge.model.model_enums import Model, TrainingMethod, TransformMethod

loader = JSONLDatasetLoader()
loader.load("s3://your-bucket/training-data.jsonl")
loader.transform(
 method=TransformMethod.SCHEMA,
 training_method=TrainingMethod.SFT_LORA,
 model=Model.NOVA_LITE_2,
 column_mappings={"question": "input", "answer": "output"},
)
```

## 2. Configure Your Infrastructure

Choose your compute resources—the SDK validates configurations and ensures optimal setup.

```
from amzn_nova_forge.manager.runtime_manager import BedrockRuntimeManager,
 SMTJRRuntimeManager, SMHPRRuntimeManager
Bedrock
runtime = BedrockRuntimeManager(
 execution_role="arn:aws:iam::123456789012:role/ExampleRole"
)

SageMaker Training Jobs
```

```
runtime = SMTJRuntimeManager(
 instance_type="ml.p5.48xlarge",
 instance_count=4
)

SageMaker HyperPod
runtime = SMHPRuntimeManager(
 instance_type="ml.p5.48xlarge",
 instance_count=4,
 cluster_name="my-hyperpod-cluster",
 namespace="kubeflow"
)
```

### 3. Train

Start training with just a few lines of code.

```
from amzn_nova_forge.model import NovaModelCustomizer
from amzn_nova_forge.model.model_enums import Model, TrainingMethod

customizer = NovaModelCustomizer(
 model=Model.NOVA_LITE_2,
 method=TrainingMethod.SFT_LORA,
 infra=runtime,
 data_s3_path="s3://your-bucket/prepared-data.jsonl"
)

result = customizer.train(job_name="my-training-job")
```

### 4. Monitor

Track your training progress directly from the SDK.

```
from amzn_nova_forge.monitor.log_monitor import CloudWatchLogMonitor

Monitor training logs
customizer.get_logs()

Or monitor directly via CloudWatchLogMonitor
monitor = CloudWatchLogMonitor.from_job_result(result)
monitor.show_logs(limit=10)

Check job status
```

```
result.get_job_status() # InProgress, Completed, Failed
```

## 5. Evaluate

Evaluate model performance with a variety of [built-in benchmarks](#), or design your own evaluations.

```
from amzn_nova_forge.recipe_config.eval_config import EvaluationTask

Evaluate on benchmark tasks
eval_result = customizer.evaluate(
 job_name="model-eval",
 eval_task=EvaluationTask.MMLU,
 model_path=result.model_artifacts.checkpoint_s3_path
)
```

## 6. Deploy

Deploy your customized model to production with built-in support for Amazon Bedrock or SageMaker.

```
from amzn_nova_forge.model.model_enums import DeployPlatform

Bedrock provisioned throughput
deployment = customizer.deploy(
 model_artifact_path=result.model_artifacts.checkpoint_s3_path,
 deploy_platform=DeployPlatform.BEDROCK_PT,
 pt_units=10
)

Bedrock On-Demand
deployment = customizer.deploy(
 model_artifact_path=result.model_artifacts.checkpoint_s3_path,
 deploy_platform=DeployPlatform.BEDROCK_OD,
 pt_units=10
)

Sagemaker Real-time Inference
deployment = customizer.deploy(
 model_artifact_path=result.model_artifacts.checkpoint_s3_path,
 deploy_platform=DeployPlatform.SAGEMAKER,
 unit_count=10,
 sagemaker_instance_type="ml.p5.48xlarge",
 sagemaker_environment_variables={
```

```
 "CONTEXT_LENGTH": "12000",
 "MAX_CONCURRENCY": "16",
 }
)
```

## Key Capabilities

### On The Fly Recipe Creation

The SDK eliminates the need to search for the appropriate recipes or container URI for specific techniques.

### Intelligent Data Processing

The SDK automatically transforms your data into the correct format for training. Whether you're working with JSON, JSONL, or CSV files, the data loader handles the conversion seamlessly. Data Loader supports text as well as multimodal data (images and videos).

### Enterprise Infrastructure Support

The SDK works with both SageMaker Training Jobs and SageMaker HyperPod, automatically managing:

- Instance type validation
- Recipe validation
- Dataset validation
- Job orchestration and monitoring

The SDK also supports SageMaker Training Jobs serverless and Bedrock customization.

### Comprehensive evaluation

Evaluate your customized models against [standard benchmarks](#) including:

- MMLU (Massive Multitask Language Understanding)
- BBH (Advanced Reasoning Tasks)
- GPQA (Graduate-Level Google-Proof Q&A)

Either use the benchmark defaults, or modify them to fit your needs:

- BYOM (Bring Your Own Metric)
- BYOD (Bring Your Own Dataset)

## Production Deployment

Deploy your models to Amazon Bedrock or SageMaker AI with options for:

- **Bedrock Provisioned Throughput** - Dedicated capacity for consistent performance
- **Bedrock On-Demand (only applicable to LoRA based customization)** - Pay-per-use pricing
- **SageMaker AI Real-time Inference** - Dedicated capacity for consistent performance

## Batch Inference

Run large-scale inference jobs efficiently:

- Process thousands of requests in parallel
- Automatic result aggregation
- Cost-effective batch processing

## Nova Forge

For Nova Forge subscribers, the SDK supports data mixing recipes.

## Learn More

Ready to start customizing Nova models with the Nova Forge SDK? Check out our GitHub repository for detailed guides, API references, and additional examples: <https://github.com/aws/nova-forge-sdk>

## Training for Amazon Nova models

Training Amazon Nova models on SageMaker Training Jobs supports Supervised Fine-Tuning (SFT) and Reinforcement Fine-Tuning (RFT). Each technique serves different customization needs and can be applied to different Amazon Nova model versions.

## Topics

- [Fine-tune Nova 2.0](#)
- [Reinforcement Fine-Tuning \(RFT\) with Amazon Nova models](#)

## Fine-tune Nova 2.0

### Prerequisites

Before you start a training job, note the following.

- Amazon S3 buckets to store your input data and output of training jobs. You can either use one bucket for both or separate buckets for each type of the data. Make sure your buckets are in the same AWS Region where you create all the other resources for training. For more information, see [Creating a general purpose bucket](#).
- An IAM role with permissions to run a training job. Make sure you attach an IAM policy with AmazonSageMakerFullAccess. For more information, see [How to use SageMaker AI execution roles](#).
- Base Amazon Nova recipes, see [???](#).

### What is SFT?

Supervised fine-tuning (SFT) trains a language model using labeled input-output pairs. The model learns from demonstration examples consisting of prompts and responses, refining its capabilities to align with specific tasks, instructions, or desired behaviors.

### Data preparation

#### Overview

Nova 2.0 SFT data uses the same Converse API format as Nova 1.0, with the addition of optional reasoning content fields. For complete format specifications, see:

- Reasoning content: [ReasoningContentBlock](#)
- Converse API schema: [Converse API](#)
- Dataset constraints: [Dataset constraints](#)

### Supported features

- **Input types** – Text, image, or video in user content blocks
- **Assistant content** – Text-only responses and reasoning content
- **Dataset composition** – Must be homogeneous. Choose one of:
  - Text-only turns

- Text + image turns
- Text + video turns (supports document understanding)

### Important

You cannot mix images and videos within the same dataset or across different turns.

## Current limitations

- **Multimodal reasoning content** – Although the Converse format supports image-based reasoning content, Nova 2.0 SFT supports only text-based reasoning content in the `reasoningText` field.
- **Validation sets** – You cannot provide a validation dataset for SFT with Nova 2.0. If you provide a validation dataset, it is ignored during training. This limitation applies to both UI-based and programmatic job submissions.

## Supported media formats

- **Images** – PNG, JPEG, GIF
- **Videos** – MOV, MKV, MP4

## Data format examples

Text-only (Nova 1.0 compatible)

```
{
 "schemaVersion": "bedrock-conversation-2024",
 "system": [
 {
 "text": "You are a digital assistant with a friendly personality"
 }
],
 "messages": [
 {
 "role": "user",
 "content": [
 {
```



```

 "text": "What country is right next to Australia?"
 }
]
 },
 {
 "role": "assistant",
 "content": [
 {
 "text": "The closest country is New Zealand"
 }
]
 }
]
}

```

## Text with reasoning (Nova 2.0)

```

{
 "schemaVersion": "bedrock-conversation-2024",
 "system": [
 {
 "text": "You are a digital assistant with a friendly personality"
 }
],
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "text": "What country is right next to Australia?"
 }
]
 },
 {
 "role": "assistant",
 "content": [
 {
 "reasoningContent": {
 "reasoningText": {
 "text": "I need to use my world knowledge of geography to answer this
question"
 }
 }
 }
]
 }
]
}

```

```
 },
 {
 "text": "The closest country to Australia is New Zealand, located to the
southeast across the Tasman Sea."
 }
]
}
]
```

## Image + text input

```
{
 "schemaVersion": "bedrock-conversation-2024",
 "system": [
 {
 "text": "You are a helpful assistant."
 }
],
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "image": {
 "format": "jpeg",
 "source": {
 "s3Location": {
 "uri": "s3://your-bucket/your-path/your-image.jpg",
 "bucketOwner": "your-aws-account-id"
 }
 }
 },
 },
 {
 "text": "Which country is highlighted in the image?"
 }
]
 },
 {
 "role": "assistant",
 "content": [
 {
```

```
 "reasoningContent": {
 "reasoningText": {
 "text": "I will determine the highlighted country by examining its
location on the map and using my geographical knowledge"
 }
 },
 {
 "text": "The highlighted country is New Zealand"
 }
]
}
```

## Video + text input

```
{
 "schemaVersion": "bedrock-conversation-2024",
 "system": [
 {
 "text": "You are a helpful assistant."
 }
],
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "video": {
 "format": "mp4",
 "source": {
 "s3Location": {
 "uri": "s3://your-bucket/your-path/your-video.mp4",
 "bucketOwner": "your-aws-account-id"
 }
 }
 }
 },
 {
 "text": "What is shown in this video?"
 }
]
 }
]
}
```

```

 },
 {
 "role": "assistant",
 "content": [
 {
 "reasoningContent": {
 "reasoningText": {
 "text": "I will analyze the video content to identify key elements"
 }
 }
 },
 {
 "text": "The video shows a map with New Zealand highlighted"
 }
]
 }
]
}

```

## Tool calling

Nova 2.0 SFT supports training models on tool calling patterns, enabling your model to learn when and how to invoke external tools or functions.

### Data format for tool calling

Tool calling training data includes a `toolConfig` section that defines available tools, along with conversation turns that demonstrate tool usage patterns.

### Sample input

```

{
 "schemaVersion": "bedrock-conversation-2024",
 "system": [
 {
 "text": "You are an expert in composing function calls."
 }
],
 "toolConfig": {
 "tools": [
 {
 "toolSpec": {
 "name": "getItemCost",

```

```

 "description": "Retrieve the cost of an item from the catalog",
 "inputSchema": {
 "json": {
 "type": "object",
 "properties": {
 "item_name": {
 "type": "string",
 "description": "The name of the item to retrieve cost for"
 },
 "item_id": {
 "type": "string",
 "description": "The ASIN of item to retrieve cost for"
 }
 },
 "required": [
 "item_id"
]
 }
 },
 {
 "toolSpec": {
 "name": "getItemAvailability",
 "description": "Retrieve whether an item is available in a given location",
 "inputSchema": {
 "json": {
 "type": "object",
 "properties": {
 "zipcode": {
 "type": "string",
 "description": "The zipcode of the location to check in"
 },
 "quantity": {
 "type": "integer",
 "description": "The number of items to check availability for"
 },
 "item_id": {
 "type": "string",
 "description": "The ASIN of item to check availability for"
 }
 },
 "required": [
 "item_id", "zipcode"
]
 }
 }
 }
 }
]
}

```

```

]
 }
 }
 }
]
},
"messages": [
 {
 "role": "user",
 "content": [
 {
 "text": "I need to check whether there are twenty pieces of the following
item available. Here is the item ASIN on Amazon: id-123. Please check for the zipcode
94086"
 }
]
 },
 {
 "role": "assistant",
 "content": [
 {
 "reasoningContent": {
 "reasoningText": {
 "text": "The user wants to check how many pieces of the item with ASIN
id-123 are available in the zipcode 94086"
 }
 }
 },
 {
 "toolUse": {
 "toolUseId": "getItemAvailability_0",
 "name": "getItemAvailability",
 "input": {
 "zipcode": "94086",
 "quantity": 20,
 "item_id": "id-123"
 }
 }
 }
]
 },
 {
 "role": "user",

```

```
 "content": [
 {
 "toolResult": {
 "toolUseId": "getItemAvailability_0",
 "content": [
 {
 "text": "[{\"name\": \"getItemAvailability\", \"results\": {\"availability\": true}}]"
 }
]
 }
 }
],
 {
 "role": "assistant",
 "content": [
 {
 "text": "Yes, there are twenty pieces of item id-123 available at 94086. Would you like to place an order or know the total cost?"
 }
]
 }
]
}
```

## Tool calling requirements

When creating tool calling training data, follow these requirements:

| Requirement          | Description                                                                                |
|----------------------|--------------------------------------------------------------------------------------------|
| ToolUse placement    | ToolUse must appear in assistant turns only                                                |
| ToolResult placement | ToolResult must appear in user turns only                                                  |
| ToolResult format    | ToolResult should be text or JSON only. Other modalities are not supported for Nova models |
| inputSchema format   | The inputSchema within the toolSpec must be a valid JSON Schema object                     |

| Requirement         | Description                                                                                                                                   |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| toolUsedId matching | Each ToolResult must reference a valid toolUsedId from a preceding assistant ToolUse, with each toolUsedId used exactly once per conversation |

## Important notes

- Ensure your tool definitions are consistent across all training samples
- The model learns tool invocation patterns from the demonstrations you provide
- Include diverse examples of when to use each tool and when not to use tools

## Document understanding

Nova 2.0 SFT supports training on document-based tasks, enabling your model to learn how to analyze and respond to questions about PDF documents.

### Data format for document understanding

Document understanding training data includes document references in the user content blocks, with the model learning to extract and reason over document content.

### Sample input

```
{
{
 "schemaVersion": "bedrock-conversation-2024",
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "text": "What are the ways in which a customer can experience issues during checkout on Amazon?"
 },
 {
 "document": {
 "format": "pdf",
 "source": {
 "s3Location": {
```



```
 "uri": "s3://my-bucket-name/path/to/documents/
customer_service_debugging.pdf",
 "bucketOwner": "123456789012"
 }
}
}
]
},
{
 "role": "assistant",
 "content": [
 {
 "reasoningContent": {
 "reasoningText": {
 "text": "I need to find the relevant section in the document to answer
the question."
 }
 }
 },
 {
 "text": "Customers can experience issues with 1. Data entry, 2. Payment
methods, 3. Connectivity while placing the order. Which one would you like to dive
into?"
 }
]
}
]
```

## Document understanding limitations

| Limitation            | Details                                                                                                      |
|-----------------------|--------------------------------------------------------------------------------------------------------------|
| Supported format      | PDF files only                                                                                               |
| Maximum document size | 10 MB                                                                                                        |
| Modality mixing       | A sample can have documents and text, but cannot have documents mixed with other modalities (images, videos) |

## Best practices for document understanding

- Ensure documents are clearly formatted and text is extractable
- Provide diverse examples covering different document types and question formats
- Include reasoning content to help the model learn document analysis patterns

## Video understanding

Nova 2.0 SFT supports training on video-based tasks, enabling your model to learn how to analyze and respond to questions about video content.

### Data format for video understanding

Video understanding training data includes video references in the user content blocks, with the model learning to extract information and reason over video content.

### Sample input

```
{
 "schemaVersion": "bedrock-conversation-2024",
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "text": "What are the ways in which a customer can experience issues during checkout on Amazon?"
 },
 {
 "video": {
 "format": "mp4",
 "source": {
 "s3Location": {
 "uri": "s3://my-bucket-name/path/to/videos/customer_service_debugging.mp4",
 "bucketOwner": "123456789012"
 }
 }
 }
 }
]
 }
]
}
```

```
 },
 {
 "role": "assistant",
 "content": [
 {
 "reasoningContent": {
 "reasoningText": {
 "text": "I need to find the relevant section in the video to answer the
question."
 }
 }
 },
 {
 "text": "Customers can experience issues with 1. Data entry, 2. Payment
methods, 3. Connectivity while placing the order. Which one would you like to dive
into?"
 }
]
 }
]
}
```

## Video understanding limitations

| Limitation             | Details                                                                                                    |
|------------------------|------------------------------------------------------------------------------------------------------------|
| Maximum video size     | 50 MB                                                                                                      |
| Maximum video duration | 15 minutes                                                                                                 |
| Videos per sample      | Only one video is allowed per sample. Multiple videos in the same sample are not supported                 |
| Modality mixing        | A sample can have video and text, but cannot have video combined with other modalities (images, documents) |

## Supported video formats

- MOV
- MKV

- MP4

## Best practices for video understanding

- Keep videos concise and focused on the content relevant to your task
- Ensure video quality is sufficient for the model to extract meaningful information
- Provide clear questions that reference specific aspects of the video content
- Include diverse examples covering different video types and question formats

## Reasoning vs non-reasoning modes

### Understanding reasoning content

Reasoning content (also called chain-of-thought) captures the model's intermediate thinking steps before generating a final answer. In the assistant turn, use the `reasoningContent` field to include these reasoning traces.

### How loss is calculated

- **With reasoning content** – Training loss includes both reasoning tokens and final output tokens
- **Without reasoning content** – Training loss is calculated only on the final output tokens

You can include `reasoningContent` across multiple assistant turns in multi-turn conversations.

### Formatting guidelines

- Use plain text for reasoning content
- Avoid markup tags like `<thinking>` and `</thinking>` unless specifically required by your task
- Ensure reasoning content is clear and relevant to the problem-solving process

### When to enable reasoning mode

Set `reasoning_enabled: true` in your training configuration when:

- Your training data has reasoning tokens
- You want the model to generate thinking tokens before producing final outputs
- You need improved performance on complex reasoning tasks

Training Nova on a non-reasoning dataset with `reasoning_enabled = true` is permitted. However, doing so may cause the model to lose its reasoning capabilities, as Nova primarily learns to generate the responses presented in the data without applying reasoning. If you want to train Nova on a non-reasoning dataset but still expect reasoning during inference, you can disable reasoning during training (`reasoning_enabled = false`) but enable it for inference. While this approach allows reasoning to be used at inference time, it does not guarantee improved performance compared to inference without reasoning. In general, enable reasoning for both training and inference when using reasoning datasets, and disable it for both when using non-reasoning datasets.

Set `reasoning_enabled: false` when:

- Your training data does not have reasoning tokens
- You're training on straightforward tasks that don't benefit from explicit reasoning steps
- You want to optimize for speed and reduce token usage

## Generating reasoning data

If your dataset lacks reasoning traces, you can create them using a reasoning-capable model like Nova Premier. Provide your input-output pairs to the model and capture its reasoning process to build a reasoning-augmented dataset.

## Using reasoning tokens for training

When training with reasoning mode enabled, the model learns to separate internal reasoning from the final answer. The training process:

- Organizes data as triples: input, reasoning, and answer
- Optimizes using standard next-token prediction loss from both reasoning and answer tokens
- Encourages the model to reason internally before generating responses

## Effective reasoning content

High-quality reasoning content should include:

- Intermediate thoughts and analysis
- Logical deductions and inference steps
- Step-by-step problem-solving approaches

- Explicit connections between steps and conclusions

This helps the model develop the ability to "think before answering."

## Dataset preparation guidelines

### Size and quality

- **Recommended size** – 2,000-10,000 samples
- **Minimum samples** – 200
- **Priority** – Quality over quantity. Ensure examples are accurate and well-annotated
- **Application alignment** – Dataset should closely reflect your production use cases

### Diversity

Include diverse examples that:

- Cover the full range of expected inputs
- Represent different difficulty levels
- Include edge cases and variations
- Prevent overfitting to narrow patterns

### Output formatting

Clearly specify the desired output format in assistant responses:

- JSON structures
- Tables
- CSV format
- Custom formats specific to your application

### Multi-turn conversations

For multi-turn datasets, remember:

- Loss is calculated only on assistant turns, not user turns
- Each assistant response should be properly formatted

- Maintain consistency across conversation turns

## Quality checklist

- Sufficient dataset size (2K-10K samples)
- Diverse examples covering all use cases
- Clear, consistent output formatting
- Accurate labels and annotations
- Representative of production scenarios
- Free from contradictions or ambiguities

## Uploading your data

Datasets should be uploaded to a bucket that can be accessed by SageMaker training jobs. For information about setting the right permissions, see [Prerequisites](#).

## Starting a training job

### Selecting hyperparameters and updating the recipe

The setup for Nova 2.0 is largely the same as for Nova 1.0. Once the input data has been uploaded to S3, use the recipe from [SageMaker HyperPod Recipes](#) under Fine tuning folder. For Nova 2.0, the following are some of the key hyperparameters that you can update based on the use case. The following is an example of the Nova 2.0 SFT PEFT recipe. For the container image URI, use `708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-TJ-SFT-V2-latest` to run an SFT fine-tuning job.

Please use v2.254.1 of SageMaker AI PySDK for strict compatibility with Nova training. Upgrading the SDK to v3.0 version will result in breaking changes. Support for v3 of SageMaker AI PySDK is coming soon.

## Sample Input

```
!pip install sagemaker==2.254.1
```

```
run:
 name: {peft_recipe_job_name}
 model_type: amazon.nova-2-lite-v1:0:256k
```

```

model_name_or_path: {peft_model_name_or_path}
data_s3_path: {train_dataset_s3_path} # SageMaker HyperPod (SMHP) only and not
compatible with SageMaker Training jobs. Note replace my-bucket-name with your real
bucket name for SMHP job
replicas: 4 # Number of compute instances for training, allowed
values are 4, 8, 16, 32
output_s3_path: "" # Output artifact path (Hyperpod job-specific; not
compatible with standard SageMaker Training jobs). Note replace my-bucket-name with
your real bucket name for SMHP job

training_config:
max_steps: 10 # Maximum training steps. Minimal is 4.
save_steps: 10 # How many training steps the checkpoint will be
saved. Should be less than or equal to max_steps
save_top_k: 1 # Keep top K best checkpoints. Note supported only
for SageMaker HyperPod jobs. Minimal is 1.
max_length: 32768 # Sequence length (options: 8192, 16384, 32768
[default], 65536)
global_batch_size: 32 # Global batch size (options: 32, 64, 128)
reasoning_enabled: true # If data has reasoningContent, set to true;
otherwise False

lr_scheduler:
warmup_steps: 15 # Learning rate warmup steps. Recommend 15% of
max_steps
min_lr: 1e-6 # Minimum learning rate, must be between 0.0 and 1.0

optim_config: # Optimizer settings
lr: 1e-5 # Learning rate, must be between 0.0 and 1.0
weight_decay: 0.0 # L2 regularization strength, must be between 0.0
and 1.0
adam_beta1: 0.9 # Exponential decay rate for first-moment estimates,
must be between 0.0 and 1.0
adam_beta2: 0.95 # Exponential decay rate for second-moment
estimates, must be between 0.0 and 1.0

peft: # Parameter-efficient fine-tuning (LoRA)
peft_scheme: "lora" # Enable LoRA for PEFT
lora_tuning:
alpha: 64 # Scaling factor for LoRA weights (options: 32, 64,
96, 128, 160, 192),
lora_plus_lr_ratio: 64.0

```



The recipe also contains largely the same hyperparameters as Nova 1.0. The notable hyperparameters are:

- `max_steps` – The number of steps you want to run the job for. Generally, for one epoch (one run through your entire dataset), the number of steps = number of data samples / global batch size. The larger the number of steps and the smaller your global batch size, the longer the job will take to run.
- `reasoning_enabled` – Controls reasoning mode for your dataset. Options:
  - `true`: Enables reasoning mode (equivalent to high reasoning)
  - `false`: Disables reasoning mode

Note: For SFT, there is no granular control over reasoning effort levels. Setting `reasoning_enabled: true` enables full reasoning capability.

- `peft.peft_scheme` – Setting this to "lora" enables PEFT-based fine tuning. Setting it to null (no quotes) enables Full-Rank fine tuning.

## Start the training job

```
from sagemaker.pytorch import PyTorch

define OutputDataConfig path
if default_prefix:
 output_path = f"s3://{bucket_name}/{default_prefix}/{sm_training_job_name}"
else:
 output_path = f"s3://{bucket_name}/{sm_training_job_name}"

output_kms_key = "<KMS key arn to encrypt trained model in Amazon-owned S3 bucket>" #
 optional, leave blank for Amazon managed encryption

recipe_overrides = {
 "run": {
 "replicas": instance_count, # Required
 "output_s3_path": output_path
 },
}

estimator = PyTorch(
 output_path=output_path,
 base_job_name=sm_training_job_name,
 role=role,
```

```

 disable_profiler=True,
 debugger_hook_config=False,
 instance_count=instance_count,
 instance_type=instance_type,
 training_recipe=training_recipe,
 recipe_overrides=recipe_overrides,
 max_run=432000,
 sagemaker_session=sagemaker_session,
 image_uri=image_uri,
 output_kms_key=output_kms_key,
 tags=[
 {'Key': 'model_name_or_path', 'Value': model_name_or_path},
]
)

print(f"\nsm_training_job_name:\n{sm_training_job_name}\n")
print(f"output_path:\n{output_path}")

```

```

from sagemaker.inputs import TrainingInput

train_input = TrainingInput(
 s3_data=train_dataset_s3_path,
 distribution="FullyReplicated",
 s3_data_type="Converse",
)

estimator.fit(inputs={"validation": val_input}, wait=False)

```

### Note

Passing a validation dataset is not supported for supervised fine tuning of Nova 2.0.

To kick off the job:

- Update the recipe with your dataset paths and hyperparameters
- Execute the specified cells in the notebook to submit the training job

The notebook handles job submission and provides status tracking.

## Preparing data for multimodal fine-tuning

The following are guidelines and requirements for preparing data for fine-tuning Understanding models:

1. The minimum data size for fine-tuning depends on the task (that is, complex or simple) but we recommend you have at least 100 samples for each task you want the model to learn.
2. We recommend using your optimized prompt in a zero-shot setting during both training and inference to achieve the best results.
3. Training and validation datasets must be JSONL files, where each line is a JSON object corresponding to a record. These file names can consist of only alphanumeric characters, underscores, hyphens, slashes and dots.
4. Image and video constraints
  - a. Dataset can't contain different media modalities. That is, the dataset can either be text with images or text with videos.
  - b. One sample (single record in messages) can have multiple images
  - c. One sample (single record in messages) can have only 1 video
5. `schemaVersion` can be any string value
6. The (*optional*) `system` turn can be a customer-provided custom system prompt.
7. Supported roles are `user` and `assistant`.
8. The first turn in messages should always start with `"role": "user"`. The last turn is the bot's response, denoted by `"role": "assistant"`.
9. The `image.source.s3Location.uri` and `video.source.s3Location.uri` must be accessible to Amazon Bedrock.
10. Your Amazon Bedrock service role must be able to access the image files in Amazon S3. For more information about granting access, see [Create a service role for model customization](#)
11. The images or videos must be in the same Amazon S3 bucket as your dataset. For example, if your dataset is in `s3://amzn-s3-demo-bucket/train/train.jsonl`, then your images or videos must be in `s3://amzn-s3-demo-bucket`
12. The terms `User:`, `Bot:`, `Assistant:`, `System:`, `<image>`, `<video>` and `[EOS]` are reserved keywords. If a user prompt or system prompt starts with any of these keywords, or have these keywords anywhere in their prompts, your training job will fail due to data issues. If you need to use these keywords for your use case, you must substitute it for a different keyword with a similar meaning so that your training can proceed.

**Note**

To validate your dataset before submitting a fine-tuning job, you can use the [dataset validation script](#) available on GitHub.

**Topics**

- [Example dataset formats](#)
- [Dataset constraints](#)

**Example dataset formats**

The following example dataset formats provide a guide for you to follow.

**Text-only custom fine tuning format**

The following example is for custom fine tuning over text only.

```
// train.jsonl
{
 "schemaVersion": "bedrock-conversation-2024",
 "system": [
 {
 "text": "You are a digital assistant with a friendly personality"
 }
],
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "text": "What is the capital of Mars?"
 }
]
 },
 {
 "role": "assistant",
 "content": [
 {
 "text": "Mars does not have a capital. Perhaps it will one day."
 }
]
 }
]
}
```

```

]
 }
]
}

```

## Single image custom fine tuning format

The following example is for custom fine tuning over text and a single image.

```

// train.jsonl{
 "schemaVersion": "bedrock-conversation-2024",
 "system": [{
 "text": "You are a smart assistant that answers questions respectfully"
 }],
 "messages": [{
 "role": "user",
 "content": [{
 "text": "What does the text in this image say?"
 },
 {
 "image": {
 "format": "png",
 "source": {
 "s3Location": {
 "uri": "s3://your-bucket/your-path/your-image.png",
 "bucketOwner": "your-aws-account-id"
 }
 }
 }
 }
]
},
{
 "role": "assistant",
 "content": [{
 "text": "The text in the attached image says 'LOL'."
 }]
}
]
}

```

## Video custom fine tuning format

The following example is for custom fine tuning over text and video.

```
{
 "schemaVersion": "bedrock-conversation-2024",
 "system": [{
 "text": "You are a helpful assistant designed to answer questions crisply and to the point"
 }],
 "messages": [{
 "role": "user",
 "content": [{
 "text": "How many white items are visible in this video?"
 },
 {
 "video": {
 "format": "mp4",
 "source": {
 "s3Location": {
 "uri": "s3://your-bucket/your-path/your-video.mp4",
 "bucketOwner": "your-aws-account-id"
 }
 }
 }
 }
]
},
{
 "role": "assistant",
 "content": [{
 "text": "There are at least eight visible items that are white"
 }]
}
]
```

## Dataset constraints

Amazon Nova applies the following constraints on model customizations for Understanding models.

| Model       | Minimum Samples | Maximum Samples | Context Length |
|-------------|-----------------|-----------------|----------------|
| Nova 2 Lite | 8               | 20k             | 32k            |

## Image and video constraints

|                               |            |
|-------------------------------|------------|
| Maximum images                | 10/sample  |
| Maximum image file size       | 10 MB      |
| Maximum videos                | 1/sample   |
| Maximum video length/duration | 90 seconds |
| Maximum video file size       | 50 MB      |

## Supported media formats

- Image - png, jpeg, gif, webp
- Video - mov, mkv, mp4, webm

## Reinforcement Fine-Tuning (RFT) with Amazon Nova models

### Overview

#### What is RFT?

Reinforcement fine-tuning (RFT) improves model performance by training on feedback signals—measurable scores or rewards indicating how well the model performed—rather than exact correct answers. Unlike supervised fine-tuning that learns from input-output pairs, RFT uses reward functions to evaluate model responses and iteratively optimizes the model to maximize these rewards. This approach excels when defining the exact correct output is challenging, but you can reliably measure response quality.

#### When to use RFT

Use RFT when you can define clear, measurable success criteria but struggle to provide exact correct outputs for training. It's ideal for:

- Tasks where quality is subjective or multifaceted (creative writing, code optimization, complex reasoning)
- Scenarios with multiple valid solutions where some are clearly better than others
- Applications requiring iterative improvement, personalization, or adherence to complex business rules
- Cases where collecting high-quality labeled examples is expensive or impractical

### Best use cases

RFT excels in domains where output quality can be objectively measured but optimal responses are difficult to define upfront:

- Mathematical problem-solving and code generation
- Scientific reasoning and structured data analysis
- Tasks requiring step-by-step reasoning or multi-turn problem solving
- Applications balancing multiple objectives (accuracy, efficiency, style)
- Scenarios where success can be verified programmatically through execution results or performance metrics

### Supported models

Nova Lite 2.0

### Data format overview

RFT training data must follow the OpenAI Reinforcement Fine-Tuning [format](#). Each training example is a JSON object containing:

- A `messages` array with conversational turns using `system` and `user` roles
- A `reference_answer` field containing the expected output or evaluation criteria for reward calculation

### Current limitations

- Text only



## Data format examples

Each example should be on a single line in your JSONL file, with one JSON object per line.

### Chemistry problem

```
{
 "id": "chem-01",
 "messages": [
 {
 "role": "system",
 "content": "You are a helpful chemistry assistant"
 },
 {
 "role": "user",
 "content": "Calculate the molecular weight of caffeine (C8H10N4O2)"
 }
],
 "reference_answer": {
 "molecular_weight": 194.19,
 "unit": "g/mol",
 "calculation": "8(12.01) + 10(1.008) + 4(14.01) + 2(16.00) = 194.19"
 }
}
```

### Math problem

```
{
 "id": "sample-001", // Optional
 "messages": [
 {
 "role": "system",
 "content": "You are a math tutor"
 },
 {
 "role": "user",
 "content": "Solve: $2x + 5 = 13$ "
 }
],
 "reference_answer": {
 "solution": " $x = 4$ ",
 "steps": [" $2x = 13 - 5$ ", " $2x = 8$ ", " $x = 4$ "]
 }
}
```

```
}
```

## Code problem

```
{
 "id": "code-002",
 "messages": [
 {
 "role": "system",
 "content": "You are a helpful programming assistant"
 },
 {
 "role": "user",
 "content": "Write a Python function that reverses a string without using
built-in reverse methods"
 }
],
 "reference_answer": {
 "code": "def reverse_string(s): \n result = '' \n for i in range(len(s)
- 1, -1, -1): \n result += s[i] \n return result",
 "test_cases": [
 {
 "input": "hello",
 "expected_output": "olleh"
 },
 {
 "input": "",
 "expected_output": ""
 },
 {
 "input": "a",
 "expected_output": "a"
 },
 {
 "input": "Python123",
 "expected_output": "321nohtyP"
 }
],
 "all_tests_pass": true
 }
}
```

The `reference_answer` field contains the expected output or evaluation criteria that your reward function uses to score the model's response. It is not limited to structured outputs—it can contain any format that helps your reward function evaluate quality.

## Dataset size recommendations

### Starting point

- Minimum 100 training examples
- Minimum 100 evaluation examples

### Evaluation-first approach

Before investing in large-scale RFT training, evaluate your model's baseline performance:

- **High performance (>95% reward)** – RFT may be unnecessary—your model already performs well
- **Very poor performance (0% reward)** – Switch to SFT first to establish basic capabilities
- **Moderate performance** – RFT is likely appropriate

Starting with a small dataset allows you to:

- Validate your reward function is bug-free
- Confirm RFT is the right approach for your use case
- Identify and fix issues early
- Test the workflow before scaling up

Once validated, you can expand to larger datasets to further improve performance.

## Characteristics of effective training data

### Clarity and consistency

Good RFT examples require clear, unambiguous input data that enables accurate reward calculation across different model outputs. Avoid noise in your data, including:

- Inconsistent formatting

- Contradictory labels or instructions
- Ambiguous prompts
- Conflicting reference answers

Any ambiguity will mislead the training process and cause the model to learn unintended behaviors.

## Diversity

Your dataset should capture the full diversity of production use cases to ensure robust real-world performance. Include:

- Different input formats and edge cases
- Map actual production usage patterns from logs and user analytics
- Sample across user types, geographic regions, and seasonal variations
- Include difficulty levels from simple to complex problems

## Reward function considerations

Design your reward function for efficient training:

- Execute within seconds (not minutes)
- Parallelize effectively with Lambda
- Return consistent, reliable scores
- Handle different types of model outputs gracefully

Fast, scalable reward functions enable rapid iteration and cost-effective experimentation.

## Additional properties

The RFT data format supports custom fields beyond the core schema requirements (messages and `reference_answer`). This flexibility lets you add any additional data your reward function needs for proper evaluation.

**Note**

You don't need to configure this in your recipe—the data format inherently supports additional fields. Simply include them in your training data JSON, and they will be passed to your reward function in the metadata field.

**Common additional properties**

Example metadata fields:

- `task_id` – Unique identifier for tracking
- `difficulty_level` – Problem complexity indicator
- `domain` – Subject area or category
- `expected_reasoning_steps` – Number of steps in solution

**Example with additional properties**

```
{
 "messages": [
 {
 "role": "system",
 "content": "You are a math tutor"
 },
 {
 "role": "user",
 "content": "Solve: $2x + 5 = 13$ "
 }
],
 "reference_answer": {
 "solution": " $x = 4$ ",
 "steps": [" $2x = 13 - 5$ ", " $2x = 8$ ", " $x = 4$ "]
 },
 "task_id": "algebra_001",
 "difficulty_level": "easy",
 "domain": "algebra",
 "expected_reasoning_steps": 3
}
```

These additional fields are passed to your reward function during evaluation, enabling sophisticated scoring logic tailored to your specific use case.

## Training configuration

### Sample recipe

```
Note:
This recipe can run on p5.48xlarge and p5en.48xlarge instance types.
run:
 name: "my-rft-run" # Unique run name (appears in logs/
artifacts).
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: nova-lite-2/prod
 data_s3_path: s3://<bucket>/<data file> # Training dataset in JSONL;
 replicas: 4
 reward_lambda_arn: arn:aws:lambda:<region>:<account-id>:function:<function-name>

MLFlow configs
mlflow_tracking_uri: "" # Required for MLFlow
mlflow_experiment_name: "my-rft-experiment" # Optional for MLFlow. Note: leave this
field non-empty
mlflow_run_name: "my-rft-run" # Optional for MLFlow. Note: leave this field non-empty

SMTJ GRPO Training specific configs
training_config:
 max_length: 8192 # Context window (tokens) for inputs
+prompt;
 global_batch_size: 16 # Total samples per optimizer step
across all replicas (16/32/64/128/256).
 reasoning_effort: high # Enables reasoning mode high / low /
or null for non-reasoning

 rollout: # How responses are generated for GRPO/
advantage calc.
 advantage_strategy:
 number_generation: 2 # N samples per prompt to estimate
advantages (variance vs cost).
 generator:
 max_new_tokens: 6000 # Cap on tokens generated per sample
 set_random_seed: true # Seed generation for reproducibility
across runs.
 temperature: 1 # Softmax temperature;
 top_k: 1 # Sample only from top-K logits
```

```

 rewards:
 preset_reward_function: null # Usage of preset reward functions
[exact_match]
 api_endpoint:
 lambda_arn: arn:aws:lambda:<region>:<account-id>:function:<function-name>
 lambda_concurrency_limit: 12 # Max concurrent Lambda invocations
(throughput vs. throttling).

 trainer:
 max_steps: 2 # Steps to train for. One Step =
global_batch_size
 save_steps: 5
 test_steps: 1

 # RL parameters
 ent_coeff: 0.0 # A bonus added to the policy loss that
rewards higher-output entropy.
 kl_loss_coef: 0.001 # Weight on the KL penalty between the
actor (trainable policy) and a frozen reference model

 optim_config: # Optimizer settings
 lr: 5e-5 # Learning rate
 weight_decay: 0.0 # L2 regularization strength (0.0-1.0)
 adam_beta1: 0.9
 adam_beta2: 0.95

 peft: # Parameter-efficient fine-tuning (LoRA)
 peft_scheme: "lora" # Enable LoRA for PEFT
 lora_tuning:
 alpha: 32
 lora_plus_lr_ratio: 64.0 # LoRA+ learning rate scaling factor (0.0-
100.0)

```

## RFT training using LLM as a judge

### Overview

Large language models (LLMs) are increasingly being used as judges in reinforcement fine-tuning (RFT) workflows, providing automated reward signals that guide model optimization. In this approach, an LLM evaluates model outputs against specified criteria—whether assessing correctness, quality, style adherence, or semantic equivalence—and assigns rewards that drive the reinforcement learning process.

This is particularly valuable for tasks where traditional reward functions are difficult to define programmatically, such as determining whether different representations (like "1/3", "0.333", and "one-third") are semantically equivalent, or evaluating nuanced qualities like coherence and relevance. By using LLM-based judges as reward functions, you can scale RFT to complex domains without requiring extensive human annotation, enabling rapid iteration and continuous improvement of your models across diverse use cases beyond traditional alignment problems.

## Reasoning mode selection

### Available modes

- none – No reasoning (omit the reasoning\_effort field)
- low – Minimal reasoning overhead
- high – Maximum reasoning capability (default when reasoning\_effort is specified)

#### Note

There is no medium option for RFT. If the reasoning\_effort field is absent from your configuration, reasoning is disabled. When reasoning is enabled, you should set max\_new\_tokens to 32768 to accommodate extended reasoning outputs.

### When to use each mode

Use high reasoning for:

- Complex analytical tasks
- Mathematical problem-solving
- Multi-step logical deduction
- Tasks where step-by-step thinking adds value

Use none (omit reasoning\_effort) or low reasoning for:

- Simple factual queries
- Direct classifications
- Speed and cost optimization



- Straightforward question-answering

**Cost and performance trade-offs**

Higher reasoning modes increase:

- Training time and cost
- Inference latency and cost
- Model capability for complex reasoning tasks

**Validating your LLM judge**

Before deploying an LLM-as-a-judge in production, validate that the judge model's evaluations align with human judgment. This involves:

- Measuring agreement rates between the LLM judge and human evaluators on representative samples of your task
- Ensuring that the LLM's agreement with humans meets or exceeds inter-human agreement rates
- Identifying potential biases in the judge model
- Building confidence that the reward signal guides your model in the intended direction

This validation step helps ensure the automated evaluation process will produce models that meet your production quality criteria.

**Lambda configuration for LLM judge**

Using an LLM as a judge is an extension of using Lambda functions for Reinforcement Learning with Verifiable Rewards (RLVR). Inside the Lambda function, you make a call to one of the models hosted in Amazon Bedrock.

**Important configuration requirements:**

| Configuration             | Requirement      | Details                                                                                                 |
|---------------------------|------------------|---------------------------------------------------------------------------------------------------------|
| Amazon Bedrock throughput | Sufficient quota | Ensure your throughput quota for the Amazon Bedrock model used is sufficient for your training workload |

| Configuration        | Requirement           | Details                                                                                                                                                            |
|----------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Lambda timeout       | Extended timeout      | Configure your Lambda function timeout up to the maximum of 15 minutes. The default setting is 3 seconds, which is insufficient for Amazon Bedrock model responses |
| Lambda concurrency   | Increased concurrency | The Lambda gets invoked in parallel during training. Increase concurrency to maximize available throughput                                                         |
| Recipe configuration | Match Lambda settings | The concurrency limit must be configured in your recipe                                                                                                            |

## Creating and running jobs

### Starting a training job

Use the SageMaker training job notebook template: <https://docs.aws.amazon.com/sagemaker/latest/dg/nova-fine-tuning-training-job.html#nova-model-training-jobs-notebook>

### Instance requirements

The container supports both Full-Rank and LoRA training:

- **LoRA training** – 2/4/6/8 × p5.48xlarge or p5en.48xlarge instances
- **Full-Rank training** – 2/4/6/8 × p5.48xlarge instances (required)

### Monitoring training

Training logs include comprehensive metrics at each step. Key metric categories:

#### Reward metrics

- `critic/rewards/mean`, `critic/rewards/max`, `critic/rewards/min` – Reward distribution
- `val-score/rewards/mean@1` – Validation rewards

### Model behavior

- actor/entropy – Policy variation (higher = more exploratory)

## Training health

- actor/pg\_loss – Policy gradient loss
- actor/pg\_clipfrac – Frequency of clipped updates
- actor/grad\_norm – Gradient magnitude

## Response characteristics

- prompt\_length/mean, prompt\_length/max, prompt\_length/min – Input token statistics
- response\_length/mean, response\_length/max, response\_length/min – Output token statistics
- response/aborted\_ratio – Incomplete generation rate (0 = all completed)

## Performance

- perf/throughput – Training throughput
- perf/time\_per\_step – Time per training step
- timing\_per\_token\_ms/\* – Per-token processing times

## Resource usage

- perf/max\_memory\_allocated\_gb, perf/max\_memory\_reserved\_gb – GPU memory
- perf/cpu\_memory\_used\_gb – CPU memory

## Using fine-tuned models

After training completes, the final model checkpoint is saved to your specified output location. The checkpoint path is available in:

- Training logs
- manifest.json file in the output Amazon S3 location (defined by output\_s3\_uri in your notebook)

## Limitations and best practices

### Limitations

- **Lambda timeout** – Reward functions must complete within 15 minutes (prevents runaway processes and manages costs)
- **Single-turn only** – Multi-turn conversations are not supported
- **Data requirements** – Needs sufficient diversity; struggles with sparse rewards (<5% positive examples)
- **Computational cost** – More expensive than supervised fine-tuning
- **No multi-modal data** – Only text data type is supported

### Best practices

#### Start small

- Begin with 100-200 examples
- Validate reward function correctness
- Scale gradually based on results

#### Pre-training evaluation

- Test baseline model performance before RFT
- If rewards are consistently 0%, use SFT first to establish basic capabilities
- If rewards are >95%, RFT may be unnecessary

#### Monitor training

- Track average reward scores and distribution
- Watch for overfitting (training rewards increase while validation rewards decrease)
- Look for concerning patterns:
  - Rewards plateauing below 0.15
  - Increasing reward variance over time
  - Declining validation performance

## Optimize reward functions

- Execute within seconds (not minutes)
- Minimize external API calls
- Use efficient algorithms
- Implement proper error handling
- Take advantage of Lambda's parallel scaling

## Iteration strategy

If rewards aren't improving:

- Adjust reward function design
- Increase dataset diversity
- Add more representative examples
- Verify reward signals are clear and consistent

## Advanced capabilities: Nova Forge

For users requiring advanced capabilities beyond standard RFT limitations, Nova Forge is available as a paid subscription service offering:

- Multi-turn conversation support
- Reward functions with >15 minute execution time
- Additional algorithms and tuning options
- Custom training recipe modifications
- State-of-the-art AI techniques

Nova Forge runs on SageMaker HyperPod and is designed to support enterprise customers to build their own frontier models.

## Useful commands and tips

A collection of [observability scripts](#) is available to help monitor the status and progress of training jobs.

Available scripts are:

- Enabling email notifications for training job status updates
- Obtaining training time estimates based on job configurations
- Obtaining approximations for how long training is expected to take for in-progress jobs

## Installation

### Note

Be sure to refresh your AWS credentials prior to using any of the following scripts.

```
pip install boto3
git clone https://github.com/aws-samples/amazon-nova-samples.git
cd amazon-nova-samples/customization/SageMakerUtils/SageMakerJobsMonitoring/
```

## Basic usage

```
Enabling email notifications for training job status updates
python enable_sagemaker_job_notifs.py --email test@amazon.com test2@gmail.com --region
us-east-1 --platform SMTJ
```

Creating resources.....

Please check your email for a subscription confirmation email, and click 'Confirm subscription' to start receiving job status email notifications!

You'll receive the confirmation email within a few minutes.

```
Obtaining training time estimates based on job configurations
python get_training_time_estimate.py
```

```
Obtaining approximations for how long training is expected to take for in-progress
jobs
python get-training-job-progress.py --region us-east-1 --job-name my-training-job --
num-dataset-samples 1000
```

Please see [here](#) for additional details and examples.



## Monitoring Progress Across Iterations

You can track metrics via MLflow.

### Nova Customization - MLFlow setup for SageMaker HyperPod

To enable your SageMaker HyperPod environment to output metrics to MLFlow, it is necessary to do some additional setup.

1. Open Amazon SageMaker AI
2. Select SageMaker Studio
  - a. If there is a profile already created, select "Open Studio".
  - b. If no profile is created, select "Create A SageMaker Domain" to set one up
3. Select MLFlow. If there is not any MLFlow App created, select "Create MLFlow App"
4. Click on the copy/paste button or the "View Details" menu item on the ML Flow App in order to get the ARN. You will need this when you submit your training job.



5. On the HyperPod cluster execution role, add the following policy. This will allow the HyperPod cluster to call the MLFlow API to publish metrics.

```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": "sagemaker-mlflow:*",
 "Resource": [
 "arn:aws:sagemaker:us-east-1:372836560492:mlflow-app/*"
]
 }
]
}
```

```

]
 },
 {
 "Effect": "Allow",
 "Action": [
 "sagemaker:ListMlflowTrackingServers",
 "sagemaker:CallMlflowAppApi"
],
 "Resource": "*"
 }
]
}

```

## Submitting a job via the CLI

Specify 4 new override parameters, either in the command line, or in the recipe yaml.

1. `mlflow_tracking_uri`: The ARN of the MLFlow App
2. `mlflow_experiment_name`: The name for this run of the experiment
3. `mlflow_experiment_name`: The experiment name where the metrics will be stored in MLFlow
4. `mlflow_run_name`: The name for this experiment

## Command line

```

--override-parameters '{"recipes.run.mlflow_tracking_uri": "arn:aws:sagemaker:us-east-1:925548216816:mlflow-app/app-B6X0UNCHQM4W", "recipes.run.mlflow_experiment_name": "myuser-sft-lora-exp1", "recipes.run.mlflow_run_name": "myuser-sft-lora-exp1-202512181940"}'

```

## yaml:

```

Run config
run:
 mlflow_tracking_uri: "arn:aws:sagemaker:us-east-1:925548216816:mlflow-app/app-B6X0UNCHQM4W"
 mlflow_experiment_name: "myuser-sft-lora-exp1"
 mlflow_run_name: "myuser-sft-lora-exp1-202512181940"

```



## Submitting a job via the SageMaker Studio UI

MLflow integration is already built into the SageMaker Studio UI experience. When submitting a training job, simply indicate which MLflow App instance to use.

1. In SageMaker Studio, navigate to Models > Nova 2.0 Lite > Customize > Customize with UI.
2. Expand the Advanced Configuration section
3. Select the MLflow App where you would like to send the training metrics. You can also set your experiment name and experiment run here.

**MLflow**

**MLflow Resource \***  
MLflow resource selected for experiment tracking.

DefaultMLflowApp ▼

**Experiment name**  
Default name applied. Enter a new name if desired.

myuser-sft-lora-exp1

**Experiment run**  
The run name is autogenerated by default. Specify a new name if desired.

myuser-sft-lora-exp1-1766499113356

## Submitting a job via the AWS CLI

If you use the AWS CLI, you must create an MLflow app and pass it as an input to the training job API request.

```
mlflow_app_name="<enter your MLflow app name>"
role_arn="<enter your role ARN>"
bucket_name="<enter your bucket name>"
region="<enter your region>"

mlflow_app_arn=$(aws sagemaker create-mlflow-app \
 --name $mlflow_app_name \
 --artifact-store-uri "s3://$bucket_name" \
```

```
--role-arn $role_arn \
--region $region)
```

## Create an MLflow app

**Using Studio UI:** If you create a training job through the Studio UI, a default MLflow app is created automatically and selected by default under Advanced Options.

**Using CLI:** If you use the CLI, you must create an MLflow app and pass it as an input to the training job API request.

```
mlflow_app_name="<enter your MLflow app name>"
role_arn="<enter your role ARN>"
bucket_name="<enter your bucket name>"
region="<enter your region>"

mlflow_app_arn=$(aws sagemaker create-mlflow-app \
 --name $mlflow_app_name \
 --artifact-store-uri "s3://$bucket_name" \
 --role-arn $role_arn \
 --region $region)
```

## Access the MLflow app

**Using CLI:** Create a pre-signed URL to access the MLflow app UI:

```
aws sagemaker create-presigned-mlflow-app-url \
 --arn $mlflow_app_arn \
 --region $region \
 --output text
```

**Using Studio UI:** The Studio UI displays key metrics stored in MLflow and provides a link to the MLflow app UI.

## Key metrics to track

Monitor these metrics across iterations to assess improvement and track the job progress:

### For SFT

- Training loss curves

- Number of samples consumed and time to process samples
- Performance accuracy on held-out test sets
- Format compliance (e.g., valid JSON output rate)
- Perplexity on domain-specific evaluation data

### For RFT

- Average reward scores over training
- Reward distribution (percentage of high-reward responses)
- Validation reward trends (watch for over-fitting)
- Task-specific success rates (e.g., code execution pass rate, math problem accuracy)

### General

- Benchmark performance deltas between iterations
- Human evaluation scores on representative samples
- Production metrics (if deploying iteratively)

## Determining when to stop

Stop iterating when:

- **Performance plateaus:** Additional training no longer meaningfully improves target metrics
- **Technique switching helps:** If one technique plateaus, try switching (e.g., SFT → RFT → SFT) to break through performance ceilings
- **Target metrics achieved:** Your success criteria are met
- **Regression detected:** New iterations degrade performance (see rollback procedures below)

For detailed evaluation procedures, refer to the **Evaluation** section.

## Evaluating your SageMaker AI-trained model

The purpose of the evaluation process is to assess trained-model performance against benchmarks or custom dataset. The evaluation process typically involves steps to create evaluation recipe pointing to the trained model, specify evaluation datasets and metrics, submit a separate job for

the evaluation, and evaluate against standard benchmarks or custom data. The evaluation process will output performance metrics stored in your Amazon S3 bucket.

### Note

The evaluation process described in this topic is an offline process. The model is tested against fixed benchmarks with predefined answers, rather than being assessed in real-time or through live user interactions. For real-time evaluation, you can test the model after it has been deployed to Amazon Bedrock by calling [Amazon Bedrock](#) Runtime APIs.

## Topics

- [Prerequisites](#)
- [Available benchmark tasks](#)
- [Evaluation specific configurations](#)
- [Running evaluation training jobs](#)
- [Assessing and analyzing evaluation results](#)
- [Evaluation best practices and troubleshooting](#)
- [Available subtasks](#)
- [Reasoning model evaluation](#)
- [RFT evaluation](#)
- [Implementing reward functions](#)

## Prerequisites

Before you start a evaluation training job, note the following.

- A SageMaker AI-trained Amazon Nova model which you want to evaluate its performance.
- Base Amazon Nova recipe for evaluation. For more information, see [???](#).

## Available benchmark tasks

A sample code package is available that demonstrates how to calculate benchmark metrics using the SageMaker AI model evaluation feature for Amazon Nova. To access the code packages, see [sample-Nova-lighteval-custom-task](#).

Here is a list of available industry standard benchmarks supported. You can specify the following benchmarks in the `eval_task` parameter.

### Available benchmarks for model evaluation

| Benchmark | Modality | Description                                                                                                      | Metrics  | Strategy | Subtask available |
|-----------|----------|------------------------------------------------------------------------------------------------------------------|----------|----------|-------------------|
| mmlu      | Text     | Multi-task Language Understanding – Tests knowledge across 57 subjects.                                          | accuracy | zs_cot   | Yes               |
| mmlu_pro  | Text     | MMLU – Professional Subset – Focuses on professional domains such as law, medicine, accounting, and engineering. | accuracy | zs_cot   | No                |
| bbh       | Text     | Advanced Reasoning Tasks – A collection of challenging problems that test higher-level cognitive and             | accuracy | fs_cot   | Yes               |

| Benchmark               | Modality | Description                                                                                                                  | Metrics     | Strategy | Subtask available |
|-------------------------|----------|------------------------------------------------------------------------------------------------------------------------------|-------------|----------|-------------------|
| problem-solving skills. |          |                                                                                                                              |             |          |                   |
| gpqa                    | Text     | General Physics Question Answering – Assesses comprehension of physics concepts and related problem-solving abilities.       | accuracy    | zs_cot   | No                |
| math                    | Text     | Mathematical Problem Solving – Measures mathematical reasoning across topics including algebra, calculus, and word problems. | exact_match | zs_cot   | Yes               |

| Benchmark     | Modality | Description                                                                                                                       | Metrics    | Strategy | Subtask available |
|---------------|----------|-----------------------------------------------------------------------------------------------------------------------------------|------------|----------|-------------------|
| strong_reject | Text     | Quality-Control Task – Tests the model's ability to detect and reject inappropriate, harmful, or incorrect content.               | deflection | zs       | Yes               |
| ifeval        | Text     | Instruction-Following Evaluation – Gauges how accurately a model follows given instructions and completes tasks to specification. | accuracy   | zs       | No                |

| Benchmark | Modality            | Description                                                                                                                                                                                                                                                                                                                     | Metrics | Strategy | Subtask available |
|-----------|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|----------|-------------------|
| gen_qa    | Multi-Modal (image) | Custom Dataset Evaluation – Lets you supply your own dataset for benchmarking, comparing model outputs to reference answers with metrics such as ROUGE and BLEU. gen_qa supports image inference for Amazon Nova Lite or Amazon Nova Pro based models. Also supports Bring-Your-Own Metrics lambda. (For RFT evaluation, please | all     | gen_qa   | No                |



| Benchmark | Modality    | Description                                                                                                                                                 | Metrics  | Strategy | Subtask available |
|-----------|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|----------|-------------------|
|           |             | use RFT eval recipe)                                                                                                                                        |          |          |                   |
| mmmu      | Multi-Modal | Massive Multidiscipline Multimodal Understanding (MMMU) – College-level benchmark comprising multiple-choice and open-ended questions from 30 disciplines.) | accuracy | zs_cot   | Yes               |

| Benchmark | Modality | Description                                                                                                                                                                                                    | Metrics | Strategy | Subtask available |
|-----------|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|----------|-------------------|
| llm_judge | Text     | LLM-as-a-Judge Preference Comparison – Uses a Nova Judge model to determine preference between paired responses (B compared with A) for your prompts, calculating the probability of B being preferred over A. | all     | judge    | No                |

| Benchmark    | Modality            | Description                                                                                                                     | Metrics | Strategy | Subtask available |
|--------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------|---------|----------|-------------------|
| mm_llm_judge | Multi-Modal (image) | This new benchmark behaves the same as the text-based llm_judge above. The only difference is that it supports image inference. | all     | judge    | No                |

| Benchmark        | Modality | Description                                                                                                                                                                                                                                                                                                               | Metrics | Strategy | Subtask available |
|------------------|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|----------|-------------------|
| rubric_llm_judge | Text     | Rubric Judge is an enhanced LLM-as-a-judge evaluation model built on Nova 2.0 Lite. Unlike the <a href="#">original judge model</a> that only provides preference verdicts, Rubric Judge dynamically generates custom evaluation criteria tailored to each prompt and assigns granular scores across multiple dimensions. | all     | judge    | No                |

| Benchmark           | Modality | Description                                                                                                                       | Metrics     | Strategy | Subtask available |
|---------------------|----------|-----------------------------------------------------------------------------------------------------------------------------------|-------------|----------|-------------------|
| aime_2024           | Text     | AIME 2024<br>- American Invitational Mathematics Examination problems testing advanced mathematical reasoning and problem-solving | exact_match | zs_cot   | No                |
| calendar_scheduling | Text     | Natural Plan<br>- Calendar Scheduling task testing planning abilities for scheduling meetings across multiple days and people     | exact_match | fs       | No                |

| Benchmark | Modality | Description                                                                                                    | Metrics | Strategy | Subtask available |
|-----------|----------|----------------------------------------------------------------------------------------------------------------|---------|----------|-------------------|
| humaneval | Text     | HumanEval - A benchmark dataset designed to evaluate the code generation capabilities of large language models | pass@1  | zs       | No                |

## Evaluation specific configurations

Below is a breakdown of the key components in the recipe and guidance on how to modify them for your use cases.

### Understanding and modifying your recipes

#### General run configuration

```
run:
 name: eval_job_name
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: nova-lite-2/prod # or s3://escrow_bucket/model_location
 replicas: 1
 data_s3_path: ""
 mlflow_tracking_uri: ""
 mlflow_experiment_name : ""
 mlflow_run_name : ""
```

- **name:** A descriptive name for your evaluation job.
- **model\_type:** Specifies the Nova model variant to use. Do not manually modify this field. Options include:
  - amazon.nova-micro-v1:0:128k

- `amazon.nova-lite-v1:0:300k`
- `amazon.nova-pro-v1:0:300k`
- `amazon.nova-2-lite-v1:0:256k`
- `model_name_or_path`: The path to the base model or s3 path for post trained checkpoint. Options include:
  - `nova-micro/prod`
  - `nova-lite/prod`
  - `nova-pro/prod`
  - `nova-lite-2/prod`
  - S3 path for post trained checkpoint path (`s3:customer-escrow-111122223333-smtj-<unique_id>/<training_run_name>`)

#### Note

##### **Evaluate post-trained model**

To evaluate a post-trained model after a Nova SFT training job, follow these steps after running a successful training job. At the end of the training logs, you will see the log message "Training is complete". You will also find a `manifest.json` file in your output bucket containing the location of your checkpoint. This file will be located within an `output.tar.gz` file at your output S3 location. To proceed with evaluation, use this checkpoint by setting it as the value for `run.model_name_or_path` in your recipe configuration.

- `replica`: The number of compute instances to use for distributed inference (running inference across multiple nodes). Set `replica > 1` to enable multi-node inference, which accelerates evaluation. If both `instance_count` and `replica` are specified, `instance_count` takes precedence. Note that multiple replicas only apply to SageMaker AI training jobs, not .
- `data_s3_path`: The input dataset Amazon S3 path. This field is required but should always left empty.
- `mlflow_tracking_uri`: (Optional) The location of the MLflow tracking server (only needed on SMHP)
- `mlflow_experiment_name`: (Optional) Name of the experiment to group related ML runs together
- `mlflow_run_name`: (Optional) Custom name for a specific training run within an experiment

## Evaluation configuration

```
evaluation:
 task: mmlu
 strategy: zs_cot
 subtask: abstract_algebra
 metric: accuracy
```

- **task:** Specifies the evaluation benchmark or task to use. Supported task includes:
  - mmlu
  - mmlu\_pro
  - bbh
  - gpqa
  - math
  - strong\_reject
  - gen\_qa
  - ifeval
  - mmmu
  - llm\_judge
  - mm\_llm\_judge
  - rubric\_llm\_judge
  - aime\_2024
  - calendar\_scheduling
  - humaneval
- **strategy:** Defines the evaluation approach.
  - zs\_cot: Zero-shot Chain of Thought - an approach to prompt large language models that encourages step-by-step reasoning without requiring explicit examples.
  - fs\_cot: Few-shot Chain of Thought - an approach that provides a few examples of step-by-step reasoning before asking the model to solve a new problem.
  - zs: Zero-shot - an approach to solve a problem without any prior training examples.
  - gen\_qa: Strategy specific for bring your own dataset.
  - judge: Strategy specific for Nova LLM as Judge and mm\_llm\_judge.



- **subtask**: Optional. Specific components of the evaluation task. For a complete list of available subtasks, see [???](#).
  - Check supported subtasks in Available benchmarks tasks.
  - Should remove this field if there are no subtasks benchmarks.
- **metric**: The evaluation metric to use.
  - **accuracy**: Percentage of correct answers.
  - **exact\_match**: For math benchmark, returns the rate at which the input predicted strings exactly match their references.
  - **deflection**: For strong reject benchmark, returns relative deflection to base model and difference significance metrics.
  - **all**:

For `gen_qa`, bring your own dataset benchmark, return following metrics:

- **rouge1**: Measures overlap of unigrams (single words) between generated and reference text.
- **rouge2**: Measures overlap of bigrams (two consecutive words) between generated and reference text.
- **rougeL**: Measures longest common subsequence between texts, allowing for gaps in the matching.
- **exact\_match**: Binary score (0 or 1) indicating if the generated text matches the reference text exactly, character by character.
- **quasi\_exact\_match**: Similar to exact match but more lenient, typically ignoring case, punctuation, and white space differences.
- **f1\_score**: Harmonic mean of precision and recall, measuring word overlap between predicted and reference answers.
- **f1\_score\_quasi**: Similar to `f1_score` but with more lenient matching, using normalized text comparison that ignores minor differences.
- **bleu**: Measures precision of n-gram matches between generated and reference text, commonly used in translation evaluation.

For `llm_judge` and `mm_llm_judge`, bring your own dataset benchmark, return following metrics:

- **a\_scores**: Number of wins for `response_A` across forward and backward evaluation passes.

- `a_scores_stderr`: Standard error of `response_A_scores` across pairwise judgements.
- `b_scores`: Measures Number of wins for `response_B` across forward and backward evaluation passes.
- `a_scores_stderr`: Standard error of `response_B_scores` across pairwise judgements.
- `ties`: Number of judgements where `response_A` and `response_B` are evaluated as equal.
- `ties_stderr`: Standard error of `ties` across pairwise judgements.
- `inference_error`: Count of judgements that could not be properly evaluated.
- `score`: Aggregate score based on wins from both forward and backward passes for `response_B`.
- `score_stderr`: Aggregate score based on wins from both forward and backward passes for `response_B`.
- `inference_error_stderr`: Standard error of the aggregate score across pairwise judgements.
- `winrate`: The probability that `response_B` will be preferred over `response_A` calculated using Bradley-Terry probability.
- `lower_rate`: Lower bound (2.5th percentile) of the estimated win rate from bootstrap sampling.
- `upper_rate`: Upper bound (97.5th percentile) of the estimated win rate from bootstrap sampling.

### Inference configuration (optional)

```
inference:
 max_new_tokens: 2048
 top_k: -1
 top_p: 1.0
 temperature: 0
 top_logprobs: 10
 reasoning_effort: null # options: low/high to enable reasoning or null to disable reasoning
```

- `max_new_tokens`: Maximum number of tokens to generate. Must be an integer. (Unavailable for LLM Judge)
- `top_k`: Number of the highest probability tokens to consider. Must be an integer.

- **top\_p**: Cumulative probability threshold for token sampling. Must be a float between 1.0 to 0.0.
- **temperature**: Randomness in token selection (higher = more random), keep 0 to make the result deterministic. Float type, minimal value is 0.
- **top\_logprobs**: The number of top logprobs to be returned in the inference response. This value must be an integer from 0 to 20. Logprobs contain the considered output tokens and log probabilities of each output token returned in the message content.
- **reasoning\_effort**: controls the reasoning behavior for reasoning-capable models. Set **reasoning\_effort** only when **model\_type** specifies a reasoning-capable model (currently `amazon.nova-2-lite-v1:0:256k`). Available options are `null` (default value if not set; disables reasoning), `low`, or `high`.

## Log Probability Output Format

When **top\_logprobs** is configured in your inference settings, the evaluation output includes token-level log probabilities in the parquet files. Each token position contains a dictionary of the top candidate tokens with their log probabilities in the following structure:

```
{
 "Ġint": {"logprob_value": -17.8125, "decoded_value": " int"},
 "Ġthe": {"logprob_value": -2.345, "decoded_value": " the"}
}
```

Each token entry contains:

- **logprob\_value**: The log probability value for the token
- **decoded\_value**: The human-readable decoded string representation of the token

The raw tokenizer token is used as the dictionary key to ensure uniqueness, while **decoded\_value** provides a readable interpretation.

## Evaluation recipe examples

Amazon Nova provides four different types of evaluation recipes. All recipes are available in [SageMaker HyperPod recipes GitHub repository](#).

## General text benchmark recipes

These recipes enable you to evaluate the fundamental capabilities of Amazon Nova models across a comprehensive suite of text-only benchmarks.

Recipe format: `xxx_general_text_benchmark_eval.yaml`.

## Bring your own dataset benchmark recipes

These recipes enable you to bring your own dataset for benchmarking and compare model outputs to reference answers using different types of metrics.

Recipe format: `xxx_bring_your_own_dataset_eval.yaml`.

## Bring your own dataset requirements

File format:

- Single `gen_qa.jsonl` file containing evaluation examples. The file name should be exact `gen_qa.jsonl`.
- You must upload your dataset to an S3 location where SageMaker AI training jobs can access.
- The file must follow the required schema format for general Q&A dataset.

Schema format requirements - Each line in the `.jsonl` file must be a JSON object with the following fields.

- Required fields.

`query`: String containing the question or instruction that needs an answer.

`response`: String containing the expected model output.

- Optional fields.

`system`: String containing the system prompt that sets the behavior, role, or personality of the AI model before it processes the query.

`images`: Array containing a list of objects with data attributes (Base64 encoded image strings).

`metadata`: String containing metadata associated with the entry for tagging purposes.

## Example entry

```
{
 "system": "You are an English major with top marks in class who likes to give minimal word responses: ",
 "query": "What is the symbol that ends the sentence as a question",
 "response": "?"
}{
 "system": "You are a pattern analysis specialist who provides succinct answers: ",
 "query": "What is the next number in this series? 1, 2, 4, 8, 16, ?",
 "response": "32"
}{
 "system": "You have great attention to detail and follow instructions accurately: ",
 "query": "Repeat only the last two words of the following: I ate a hamburger today and it was kind of dry",
 "response": "of dry"
}{
 "system": "Image inference: ",
 "query": "What is the number in the image? Please just use one English word to answer.",
 "response": "two",
 "images": [
 {
 "data": "data:image/png;base64,iVBORw0KGgoA ..."
 }
]
}
```

To use your custom dataset, modify your evaluation recipe by adding the following required fields without changing the existing configuration:

```
evaluation:
 task: gen_qa
 strategy: gen_qa
 metric: all
```

## Limitations

- Only one .jsonl file is allowed per evaluation.
- The file must strictly follow the defined schema.

## Bring your own metrics

You can bring your own metrics to fully customize your model evaluation workflow with custom preprocessing, postprocessing, and metrics capabilities. Preprocessing allows you to process input data before sending it to the inference server, and postprocessing allows you to customize metrics calculation and return custom metrics based on your needs.

Follow these steps to bring your own metrics with custom evaluation SDK.

1. If you haven't done so, [create an AWS Lambda function](#) in your AWS account first.
2. Download the pre-built `nova-custom-eval-layer.zip` file from the [GitHub repository](#). You can use this open-source Nova custom evaluation SDK to validate input and output payloads for your custom function and provide a unified interface for integrating with Nova's bring your own metrics evaluation during training.
3. Upload the custom Lambda layer using the following command:

```
aws lambda publish-layer-version \
 --layer-name nova-custom-eval-layer \
 --zip-file fileb://nova-custom-eval-layer.zip \
 --compatible-runtimes python3.12 python3.11 python3.10 python3.9
```

4. Add this layer as a custom layer to your Lambda function, along with the required AWS layer: `AWSLambdaPowertoolsPythonV3-python312-arm64` (required for `pydantic` dependency).
5. Update your Lambda code using the provided example, modifying the code as needed. This example code creates a Lambda function for Nova's custom evaluation with preprocessing and postprocessing steps for model evaluation.

```
from nova_custom_evaluation_sdk.processors.decorators import preprocess,
 postprocess
from nova_custom_evaluation_sdk.lambda_handler import build_lambda_handler

@preprocess
def preprocessor(event: dict, context) -> dict:
 data = event.get('data', {})
 return {
 "statusCode": 200,
 "body": {
 "system": data.get("system"),
 "prompt": data.get("prompt", ""),
```

```

 "gold": data.get("gold", "")
 }
}

@postprocess
def postprocessor(event: dict, context) -> dict:
 # data is already validated and extracted from event
 data = event.get('data', [])
 inference_output = data.get('inference_output', '')
 gold = data.get('gold', '')

 metrics = []
 inverted_accuracy = 0 if inference_output.lower() == gold.lower() else 1.0
 metrics.append({
 "metric": "inverted_accuracy_custom",
 "value": accuracy
 })

 # Add more metrics here

 return {
 "statusCode": 200,
 "body": metrics
 }

Build Lambda handler
lambda_handler = build_lambda_handler(
 preprocessor=preprocessor,
 postprocessor=postprocessor
)

```

6. Grant Lambda access to the evaluation job. Ensure the execution role specified for the evaluation job includes a policy the invoke your Lambda function. Here is an example policy.

## JSON

```

{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Sid": "LambdaAccess",
 "Effect": "Allow",
 "Action": [

```

```
 "lambda:InvokeFunction"
],
 "Resource": "arn:aws:lambda:us-east-1:111122223333:function:ExampleFunction",
 "Condition": {
 "StringLike": {
 "aws:PrincipalArn": "arn:aws:iam::111122223333:role/service-role/AmazonSageMaker-ExecutionRole-ARN"
 }
 }
 },
 {
 "Sid": "DenyNonAWSEventSourcesForLambda",
 "Effect": "Deny",
 "Action": [
 "lambda:InvokeFunction"
],
 "Resource": "arn:aws:lambda:us-east-1:111122223333:function:ExampleFunction",
 "Condition": {
 "Null": {
 "lambda:EventSourceToken": false
 }
 }
 }
]
}
```

7. Review the Lambda payload schema. The following table lists the Lambda request and response schema. You can validate your schema using the Nova custom evaluation SDK.

|              | Lambda Request Payload                                                                                                                                    | Lambda Response Payload                                                                                                                                                           |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Preprocessor | <pre>{   "process_type": "preprocess",   "data": {     "system": "You are a helpful assistant",     "prompt": "What is 2+2?",     "gold": "4"   } }</pre> | <pre>{   "statusCode": 200,   "body": {     "system": "You are a helpful assistant that can substitute * for addition",     "prompt": "What is 2*2?",     "gold": "4"   } }</pre> |



|               | Lambda Request Payload                                                                                                                           | Lambda Response Payload                                                                                                                                                                                                                         |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|               | <pre>} }</pre>                                                                                                                                   | <pre>}</pre>                                                                                                                                                                                                                                    |
| Postprocessor | <pre>{   "process_type": "postprocessor",   "data": {     "prompt": "What is 2+2?",     "inference_output": "2+2=4",     "gold": "4"   } }</pre> | <pre>{   "statusCode": 200,   "body": [     {"metric": "accuracy",     "value": 1.0},     {"metric": "f1_score",     "value": 1.0},     {"metric": "exact_match",     "value": 1},     {"metric": "length_ratio",     "value": 0.8}   ] }</pre> |

8. Modify the recipe file. Here is an example.

```
processor:
 lambda_arn: arn:aws:lambda:us-east-1:111122223333:function:name
 lambda_type: "custom_metrics"
 preprocessing:
 enabled: true
 postprocessing:
 enabled: true
 aggregation: average
```

- `lambda-arn`: The Amazon Resource Name (ARN) for your Lambda function that handles preprocessing and postprocessing.
- `lambda_type`: "custom\_metrics" or "rft".
- `preprocessing`: Whether to enable custom pre-processing operations.
- `postprocessing`: Whether to enable custom post-processing operations.
- `aggregation`: Built-in aggregation function (valid options: min, max, average, sum).

Limitations

- Bring your own metrics only applies to text input datasets.
- Multi-modal input datasets are not supported.
- The preprocessing step does not process the metadata field.

## Nova LLM as a Judge benchmark recipes

Nova LLM Judge is a model evaluation feature that enables you to compare the quality of responses from one model against a baseline model's responses using a custom dataset. It accepts a dataset containing prompts, baseline responses, and challenger responses, then uses a Nova Judge model to provide a win rate metric based on [Bradley-Terry](#) probability through pairwise comparisons. Recipe format: `xxx_llm_judge_eval.yaml`.

## Nova LLM dataset requirements

File format:

- Single `llm_judge.jsonl` file containing evaluation examples. The file name should be exact `llm_judge.jsonl`.
- You must upload your dataset to an S3 location where SageMaker AI training jobs can access.
- The file must follow the required schema format for the `llm_judge` dataset.
- The input dataset should ensure all records are under 12 k context length.

Schema format - Each line in the `.jsonl` file must be a JSON object with the following fields.

- Required fields.

`prompt`: String containing the prompt for the generated response.

`response_A`: String containing the baseline response.

`response_B`: String containing the alternative response be compared with baseline response.

## Example entry

```
{
 "prompt": "What is the most effective way to combat climate change?",
 "response_A": "The most effective way to combat climate change is through a combination
 of transitioning to renewable energy sources and implementing strict carbon pricing
```

```
 policies. This creates economic incentives for businesses to reduce emissions while
 promoting clean energy adoption.",
 "response_B": "We should focus on renewable energy. Solar and wind power are good.
 People should drive electric cars. Companies need to pollute less."
 }
 {
 "prompt": "Explain how a computer's CPU works",
 "response_A": "CPU is like brain of computer. It does math and makes computer work
 fast. Has lots of tiny parts inside.",
 "response_B": "A CPU (Central Processing Unit) functions through a fetch-execute
 cycle, where instructions are retrieved from memory, decoded, and executed through its
 arithmetic logic unit (ALU). It coordinates with cache memory and registers to process
 data efficiently using binary operations."
 }
 {
 "prompt": "How does photosynthesis work?",
 "response_A": "Plants do photosynthesis to make food. They use sunlight and water. It
 happens in leaves.",
 "response_B": "Photosynthesis is a complex biochemical process where plants convert
 light energy into chemical energy. They utilize chlorophyll to absorb sunlight,
 combining CO2 and water to produce glucose and oxygen through a series of chemical
 reactions in chloroplasts."
 }
}
```

To use your custom dataset, modify your evaluation recipe with the following required fields, don't change any of the content:

```
evaluation:
 task: llm_judge
 strategy: judge
 metric: all
```

## Limitations

- Only one .jsonl file is allowed per evaluation.
- The file must strictly follow the defined schema.
- Nova Judge models are the same across micro / lite / pro specifications.
- Custom judge models are not currently supported.

## Nova LLM as a Judge for multi-modal (image) benchmark recipes

Nova LLM Judge for multi-modal (image), short for Nova MM\_LLM Judge, is a model evaluation feature that enables you to compare the quality of responses from one model against a baseline model's responses using a custom dataset. It accepts a dataset containing prompts, baseline responses, and challenger responses, and images in the form of Base64-encoded string, then uses a Nova Judge model to provide a win rate metric based on [Bradley-Terry](#) probability through pairwise comparisons. Recipe format: `xxx_mm_llm_judge_eval.yaml`.

### Nova LLM dataset requirements

File format:

- Single `mm_llm_judge.jsonl` file containing evaluation examples. The file name must be exactly `llm_judge.jsonl`.
- You must upload your dataset to an S3 location where SageMaker AI training jobs can access it.
- The file must follow the required schema format for the `mm_llm_judge` dataset.
- The input dataset should ensure all records are under 12 k context length, excluding the image's attribute.

Schema format - Each line in the `.jsonl` file must be a JSON object with the following fields.

- Required fields.

`prompt`: String containing the prompt for the generated response.

`images`: Array containing a list of objects with data attributes (values are Base64-encoded image strings).

`response_A`: String containing the baseline response.

`response_B`: String containing the alternative response to be compared with baseline response.

### Example entry

For readability, the following example includes new lines and indentation, but in the actual dataset, each record should be on a single line.

```
{
```

```

 "prompt": "What is in the image?",
 "images": [
 {
 "data": "data:image/jpeg;base64,/9j/2wBDAAQDAwQDAwQEAWQFBAQFBgo..."
 }
],
 "response_A": "a dog.",
 "response_B": "a cat.",
 }
 {
 "prompt": "How many animals are in each of the images?",
 "images": [
 {
 "data": "data:image/jpeg;base64,/9j/2wBDAAQDAwQDAwQEAWQFBAQFBgo..."
 },
 {
 "data": "data:image/jpeg;base64,/DKEafe3gihn..."
 }
],
 "response_A": "The first image contains one cat and the second image contains one dog",
 "response_B": "The first image has one animal and the second has one animal"
 }

```

To use your custom dataset, modify your evaluation recipe with the following required fields, don't change any of the content:

```

evaluation:
 task: mm_llm_judge
 strategy: judge
 metric: all

```

## Limitations

- Only one .jsonl file is allowed per evaluation.
- The file must strictly follow the defined schema.
- Nova MM Judge models only support image reference.
- Nova MM Judge models are the same across Amazon Nova Micro, Amazon Nova Lite, and Amazon Nova Pro specifications.
- Custom judge models are not currently supported.

- Amazon S3 image URI is not supported.
- The input dataset should ensure all records are under 12 k context length, excluding images attribute.

## Evaluating CPT (Continuous Pre-Training) Checkpoints

Evaluation on CPT (continuous pre-trained) models can be more of a challenge than models which have undergone SFT (supervised fine tuning) as CPT models regularly lack the ability to follow instructions. Instead of following instructions, CPT model will operate as completion models, meaning they will only attempt to continue the pattern provided with the input token sequence. Given this limitation, typical evaluation datasets may not work correctly due to the "Q&A" format of the inputs — instead of answering the question the model will simply try to continue along the same question. However, by formatting datasets in a specific way to prompt models in a completion style, we can get an understanding of how the model is performing.

Follow the below steps to understand how to perform an evaluation on a continuous pre-trained model using the Nova Forge evaluation workflow.

### Dataset Preparation and Formatting

Evaluating a CPT model takes advantage of the already existing [Bring Your Own Dataset](#) workflow already provided in the Nova Forge model evaluation experience. However, queries within the dataset must be formatted in a purely "completion" format as CPT models will not respond to standard question-style prompts in the same manner an SFT model would.

Another important frequent limitation of models which have undergone solely undergone CPT is their inability to generate STOP or end of sequence tokens — this means that the model will continue to generate tokens until it is forcefully stopped (such as with the `max_new_tokens` parameter). Given this limitation, best practice is to evaluate the models using single token responses (such as multiple choice) to ensure the model doesn't continue to generate junk output that is not needed after prompting.

For example, a typical evaluation dataset (such as MMLU, GPQA, MATH, etc), might prompt the model with a question such as:

```
Early settlements and high population density along coastlines and rivers are
best attributed to which of the following?
A: "Poor climate conditions"
B: "Limited forest cover"
C: "Cars"
```

```
D: "Access to trade routes"
```

(Expected answer is D.)

However, a CPT model would not understand how to properly respond to this question due to the lack of fine tuning on instruction following. Therefore, CPT models must be prompted in a completion style, such as:

```
Early settlements and high population density along coastlines and rivers
are best attributed to which of the following?
```

```
A: Poor climate conditions
```

```
B: Limited forest cover
```

```
C: Cars
```

```
D: Access to trade routes
```

```
The correct answer to this question is option
```

(Expected answer is D.)

After inference checking the output logprobs generated by the model will provide details on if the model was correctly processing the input and generating the correct response. It is not guaranteed that the model will always produce the exact response (in this case, the letter D) expected, however, it should be present within the logprobs if the model is functioning correctly.

Another non-multiple choice completion style prompt example:

```
The capital of France is
```

(Expected answer of Paris)

We would expect the model to either produce a response of "Paris" or see the token corresponding to "Paris" somewhere in the logprobs output.

## Dataset Formatting

CPT evaluation takes advantage of the existing [Bring Your Own Dataset](#) workflow. Data must be formatted in the "query response" format in a JSONL file separated by new lines.

An example of a dataset with 4 entries in it:

```
{"query": "The capital of France is", "response": "Paris"}
{"query": "2 + 2 =", "response": "4"}
```

```
{"query": "The mitochondria is the powerhouse of the", "response": "cell"}
{"query": "What is the largest planet?\nA: Mars\nB: Jupiter\nC: Saturn\nD: Earth\n\nAnswer:", "response": "B"}
```

Each line must contain:

- query: The prompt text for completion
- response: The expected completion (ground truth)

The model will receive raw text input without chat formatting. CPT models are typically not trained on special tokens yet and will not respond properly to chat templates, so when prompting the model ONLY the string provided in the query will be sent to the model (with an additional [BOS] token prepended to it automatically.)

## Recipe Configuration

Here is an example of a recipe which is configured for evaluating a CPT model:

```
run:
 name: cpt_eval_job
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: s3://bucket/path/to/cpt-checkpoint/

evaluation:
 task: gen_qa # Required for CPT - bring your own dataset
 strategy: gen_qa
 metric: all # Returns rouge1, rouge2, rougeL, exact_match, f1_score, bleu

inference:
 checkpoint_is_instruction_tuned: "false" # Required for CPT checkpoints
 top_logprobs: 5 # Change to desired amount of logprobs to calculate
 max_new_tokens: 1 # Keep low for completion tasks
 temperature: 0.0
```

Key changes for CPT evaluation:

- `checkpoint_is_instruction_tuned: "false"`

This is a new parameter added specifically to support evaluation runs on CPT checkpoints. Setting `checkpoint_is_instruction_tuned` to false will **disable** the default Amazon Nova chat template that normally wraps the input prompt.



- **top\_logprobs:** 5

Log probabilities (logprobs) reveal the model's confidence distribution across possible next tokens, helping you assess whether the model has learned the expected completions during pre-training. Typically, if the model is performing as intended, we should see the expected response (i.e, "A", "B", etc) as either the generated output token or a token in the logprobs.

- **max\_new\_tokens:** 1

CPT models typically have not been trained yet on generating special "stop" or "end of sequence" tokens yet to signal when to stop inference. This means the model will typically continue to generate new tokens until the specific max token length is reached resulting in unnecessary inference. Typically limiting the max\_new\_tokens to 1 and providing a prompt which can evaluate the model on a single response (like a multiple choice question) is the most efficient way to prompt the model. Setting the max\_new\_tokens to 1 ensures that extra junk tokens aren't generated.

## Key Parameters

- **checkpoint\_is\_instruction\_tuned:** Must be set to "false" (or false as boolean) to disable chat templates
- **top\_logprobs:** 5, recommended to see how the model is learning during CPT
- **task:** Must be gen\_qa - CPT models cannot use instruction-following benchmarks like MMLU or MATH
- **strategy:** Must be gen\_qa
- **max\_new\_tokens:** Recommended to keep low (1-5) since CPT models perform completion, not generation

## Running evaluation training jobs

Start a training job using the following sample notebook. Please refer to below notebook as example to run the evaluation training job. For more information, see [Use a SageMaker AI estimator to run a training job](#).

## Reference tables

Before running the notebook, refer to the following reference tables to select image URI and instance configurations.

## Selecting image URI

| Recipe               | Image URI                                                                              |
|----------------------|----------------------------------------------------------------------------------------|
| Evaluation image URI | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM-TJ-Eval-V2-latest |

## Selecting instance type and count

| Model             | Job type              | Instance type | Recommended instance count | Allowed instance count |
|-------------------|-----------------------|---------------|----------------------------|------------------------|
| Amazon Nova Micro | Evaluation (SFT/ DPO) | g5.12xlarge   | 1                          | 1 - 16                 |
| Amazon Nova Lite  | Evaluation (SFT/ DPO) | g5.12xlarge   | 1                          | 1 - 16                 |
| Amazon Nova Pro   | Evaluation (SFT/ DPO) | p5.48xlarge   | 1                          | 1 - 16                 |

## Sample notebook

The following sample notebook demonstrates how to run an evaluation training job.

```
install python SDK

Do not use sagemaker v3, as sagemaker v3 introduced breaking changes

!pip install sagemaker==2.254.1

import os
import sagemaker,boto3
from sagemaker.inputs import TrainingInput
from sagemaker.pytorch import PyTorch

sagemaker_session = sagemaker.Session()
role = sagemaker.get_execution_role()
```

```

Download recipe from https://github.com/aws/sagemaker-hyperpod-recipes/tree/main/
recipes_collection/recipes/evaluation/nova to local
Assume the file name be `recipe.yaml`

Populate parameters
input_s3_uri = "s3://<path>/input/" # (Optional) Only used for multi-modal dataset or
bring your own dataset s3 location
output_s3_uri= "s3://<path>/output/" # Output data s3 location, a zip containing
metrics json and tensorboard metrics files will be stored to this location
instance_type = "instance_type" # ml.g5.16xlarge as example
instance_count = 1 # The number of instances for inference (set instance_count > 1 for
multi-node inference to accelerate evaluation)
job_name = "your job name"
recipe_path = "recipe path" # ./recipe.yaml as example
image_uri = "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM-TJ-
Eval-V2-latest" # Do not change
output_kms_key = "<KMS key arn to encrypt trained model in Amazon-owned S3 bucket>" #
optional, leave blank for Amazon managed encryption

(Optional) To bring your own dataset and LLM judge for evaluation
evalInput = TrainingInput(
s3_data=input_s3_uri,
distribution='FullyReplicated',
s3_data_type='S3Prefix'
#)

estimator = PyTorch(
 output_path=output_s3_uri,
 base_job_name=job_name,
 role=role,
 instance_type=instance_type,
 instance_count=instance_count,
 training_recipe=recipe_path,
 sagemaker_session=sagemaker_session,
 image_uri=image_uri,
 output_kms_key=output_kms_key
)
estimator.fit()

If input dataset exist, pass in inputs
estimator.fit(inputs={"train": evalInput})

```

## Assessing and analyzing evaluation results

After your evaluation job completes successfully, you can assess and analyze the results using the following steps.

### To assess and analyze the results, following these steps.

1. Understand the output location structure. Results are stored in your specified Amazon S3 output location as a compressed file:

```
s3://your-bucket/output/benchmark-name/
job_name/
 ### output/
 ### output.tar.gz
```

2. Download the `output.tar.gz` file from your bucket. Extract the contents to reveal.

```
run_name/
eval_results/
| ### results_[timestamp].json
inference_output.jsonl (only present for gen_qa)
| ### details/
| ### model/
| ### <execution-date-time>/
| ###details_<task_name>_#_<datetime>.parquet
tensorboard_results/
 ### eval/
 ### events.out.tfevents.[timestamp]
```

- `results_[timestamp].json` - Output metrics JSON file
  - `details_<task_name>_#_<datetime>.parquet` - Inference output file (except for `strong_reject`)
  - `events.out.tfevents.[timestamp]` - TensorBoard output file
  - `inference_output.jsonl` - Cleaned inference output file (only for `gen_qa` tasks)
3. View results in TensorBoard. To visualize your evaluation metrics:
    1. Upload the extracted folder to an S3 bucket
    2. Navigate to SageMaker AI TensorBoard
    3. Select your "S3 folders"

4. Add the S3 folder path
5. Wait for synchronization to complete
4. Analyze inference outputs. All evaluation tasks, except `llm_judge` and `strong_reject`, will have the following fields for analysis in the inference output.
  - `full_prompt` - The full user prompt sent to the model used for the evaluation task.
  - `gold` - The field that contains the correct answer(s) as specified by the dataset.
  - `metrics` - The field that contains the metrics evaluated on the individual inference. Values that require aggregation would not have a value on the individual inference outputs.
  - `predictions` - The field that contains a list of the model's output for the given prompt.
  - `pred_logits` - The field that contains the considered output tokens and log probabilities of each output token returned in the message content.

By looking at these fields, you can determine the cause for metric differences and understand the behavior of the customized models.

For `llm_judge`, the inference output file contains the following fields under the `metrics` field per pair of evaluations.

- `forward_output` - Judge's raw preferences when evaluating in order (`response_A`, `response_B`).
- `backward_output` - Judge's raw preferences when evaluating in reverse order (`response_B`, `response_A`).
- `Pairwise metrics` - Metrics that are calculated per pair of forward and backward evaluation including `a_scores`, `b_scores`, `ties`, `inference-score` and `score`.

 **Note**

Aggregate metrics like `winrate` are only available in the summary results files, not per individual judgement.

For `gen_qa`, the `inference_output.jsonl` file contains the following fields for each JSON object:

- `prompt` - The final prompt submitted to the model

- inference - The raw inference output from the model
- gold - The target response from the input dataset
- metadata - The metadata string from the input dataset if provided

## Evaluation best practices and troubleshooting

### Best practices

The following lists some best practices for the evaluation process.

- Keep your output paths organized by model and benchmark type.
- Maintain consistent naming conventions for easy tracking.
- Save extracted results in a secure location.
- Monitor TensorBoard sync status for successful data loading.

### Troubleshooting

You can use CloudWatch log group `/aws/sagemaker/TrainingJobs` for training job error logs.

### Engine core Failure

#### Issue:

If you are seeing:

```
RuntimeError: Engine core initialization failed.
```

#### Cause:

Although this is a general error that can have multiple causes, it typically occurs when there is a mismatch between the model checkpoint you're trying to load and the model type specified. E.g. you want to evaluate a fine-tuned Nova 2.0 lite model checkpoint but the model type you provide is 1.0 model type. e.g. `amazon.nova-micro-v1:0:128k`

The correct mapping should be

```
model_type: amazon.nova-2-lite-v1:0:256k
model_name_or_path: nova-lite-2/prod # or s3://escrow_bucket/model_location
```

## Prevention:

Double check the `model_name_or_path` is mapped to the right `model_type` before submitting the evaluation job.

## Available subtasks

The following lists available subtasks for model evaluation across multiple domains including MMLU (Massive Multitask Language Understanding), BBH (Big Bench Hard), MATH, and MMMU (Massive Multi-discipline Multimodal Understanding). These subtasks allow you to assess your model's performance on specific capabilities and knowledge areas.

### MMLU

```
MMLU_SUBTASKS = [
 "abstract_algebra",
 "anatomy",
 "astronomy",
 "business_ethics",
 "clinical_knowledge",
 "college_biology",
 "college_chemistry",
 "college_computer_science",
 "college_mathematics",
 "college_medicine",
 "college_physics",
 "computer_security",
 "conceptual_physics",
 "econometrics",
 "electrical_engineering",
 "elementary_mathematics",
 "formal_logic",
 "global_facts",
 "high_school_biology",
 "high_school_chemistry",
 "high_school_computer_science",
 "high_school_european_history",
 "high_school_geography",
 "high_school_government_and_politics",
 "high_school_macroeconomics",
 "high_school_mathematics",
 "high_school_microeconomics",
 "high_school_physics",
```

```
"high_school_psychology",
"high_school_statistics",
"high_school_us_history",
"high_school_world_history",
"human_aging",
"human_sexuality",
"international_law",
"jurisprudence",
"logical_fallacies",
"machine_learning",
"management",
"marketing",
"medical_genetics",
"miscellaneous",
"moral_disputes",
"moral_scenarios",
"nutrition",
"philosophy",
"prehistory",
"professional_accounting",
"professional_law",
"professional_medicine",
"professional_psychology",
"public_relations",
"security_studies",
"sociology",
"us_foreign_policy",
"virology",
"world_religions"
]
```

## BBH

```
BBH_SUBTASKS = [
 "boolean_expressions",
 "causal_judgement",
 "date_understanding",
 "disambiguation_qa",
 "dyck_languages",
 "formal_fallacies",
 "geometric_shapes",
 "hyperbaton",
 "logical_deduction_five_objects",
```



```

 "logical_deduction_seven_objects",
 "logical_deduction_three_objects",
 "movie_recommendation",
 "multistep_arithmetic_two",
 "navigate",
 "object_counting",
 "penguins_in_a_table",
 "reasoning_about_colored_objects",
 "ruin_names",
 "salient_translation_error_detection",
 "snarks",
 "sports_understanding",
 "temporal_sequences",
 "tracking_shuffled_objects_five_objects",
 "tracking_shuffled_objects_seven_objects",
 "tracking_shuffled_objects_three_objects",
 "web_of_lies",
 "word_sorting"
]

```

## Math

```

MATH_SUBTASKS = [
 "algebra",
 "counting_and_probability",
 "geometry",
 "intermediate_algebra",
 "number_theory",
 "prealgebra",
 "precalculus",

```

## MMMU

```

 MATH_SUBTASKS = [
 "Accounting",
 "Agriculture",
 "Architecture_and_Engineering",
 "Art",
 "Art_Theory",
 "Basic_Medical_Science",
 "Biology",
 "Chemistry",

```

```
"Clinical_Medicine",
"Computer_Science",
"Design",
"Diagnostics_and_Laboratory_Medicine",
"Economics",
"Electronics",
"Energy_and_Power",
"Finance",
"Geography",
"History",
"Literature",
"Manage",
"Marketing",
"Materials",
"Math",
"Mechanical_Engineering",
"Music",
"Pharmacy",
"Physics",
"Psychology",
"Public_Health",
"Sociology",
```

Evaluate your customized Nova models using various evaluation methods and metrics.

### In this section:

- [Reasoning model evaluation](#)
- [RFT evaluation](#)
- [Implementing reward functions](#)

## Reasoning model evaluation

### Overview

Reasoning model support enables evaluation with reasoning-capable Nova models that perform explicit internal reasoning before generating final responses. This feature uses API-level control via the `reasoning_effort` parameter to dynamically enable or disable reasoning functionality, potentially improving response quality for complex analytical tasks.

## Supported models

- amazon.nova-2-lite-v1:0:256k

## Recipe configuration

Enable reasoning by adding the `reasoning_effort` parameter to the inference section of your recipe:

```
run:
 name: reasoning-eval-job-name # [MODIFIABLE] Unique
 identifier for your evaluation job
 model_type: amazon.nova-2-lite-v1:0:256k # [FIXED] Must be a reasoning-
supported model
 model_name_or_path: nova-lite-2/prod # [FIXED] Path to model
checkpoint or identifier
 replicas: 1 # [MODIFIABLE] Number of
replicas for SageMaker Training job
 data_s3_path: "" # [MODIFIABLE] Leave empty for
SageMaker Training job; optional for SageMaker HyperPod job
 output_s3_path: "" # [MODIFIABLE] Output path for
SageMaker HyperPod job (not compatible with SageMaker Training jobs)

evaluation:
 task: mmlu # [MODIFIABLE] Evaluation task

 strategy: zs_cot # [MODIFIABLE] Evaluation
strategy
 metric: accuracy # [MODIFIABLE] Metric
calculation method

inference:
 reasoning_effort: high # [MODIFIABLE] Enables
reasoning mode; options: low/high or null to disable
 max_new_tokens: 32768 # [MODIFIABLE] Maximum tokens
to generate, recommended value when reasoning_effort set to high
 top_k: -1 # [MODIFIABLE] Top-k sampling
parameter
 top_p: 1.0 # [MODIFIABLE] Nucleus
sampling parameter
 temperature: 0 # [MODIFIABLE] Sampling
temperature (0 = deterministic)
```

## Using the reasoning\_effort parameter

The reasoning\_effort parameter controls the reasoning behavior for reasoning-capable models.

### Prerequisites

- **Model compatibility** – Set reasoning\_effort only when model\_type specifies a reasoning-capable model (currently amazon.nova-2-lite-v1:0:256k)
- **Error handling** – Using reasoning\_effort with unsupported models will fail with ConfigValidationError: "Reasoning mode is enabled but model '{model\_type}' does not support reasoning. Please use a reasoning-capable model or disable reasoning mode."

### Available options

| Option            | Behavior                              | Token limit                          | Use case                                                                 |
|-------------------|---------------------------------------|--------------------------------------|--------------------------------------------------------------------------|
| null<br>(default) | Disables reasoning mode               | N/A                                  | Standard evaluation without reasoning overhead                           |
| low               | Enables reasoning with constraints    | 4,000 tokens for internal reasoning  | Scenarios requiring concise reasoning; optimizes for speed and cost      |
| high              | Enables reasoning without constraints | No token limit on internal reasoning | Complex problems requiring extensive analysis and step-by-step reasoning |

| Training method              | Available options | How to configure                                                     |
|------------------------------|-------------------|----------------------------------------------------------------------|
| SFT (Supervised Fine-Tuning) | High or Off only  | Use reasoning_enabled: true (high) or reasoning_enabled: false (off) |

| Training method                 | Available options | How to configure                                                            |
|---------------------------------|-------------------|-----------------------------------------------------------------------------|
| RFT (Reinforcement Fine-Tuning) | Low, High, or Off | Use reasoning_effort: low or reasoning_effort: high. Omit field to disable. |
| Evaluation                      | Low, High, or Off | Use reasoning_effort: low or reasoning_effort: high. Use null to disable.   |

When to enable reasoning

Use reasoning mode (low or high) for

- Complex problem-solving tasks (mathematics, logic puzzles, coding)
- Multi-step analytical questions requiring intermediate reasoning
- Tasks where detailed explanations or step-by-step thinking improve accuracy
- Scenarios where response quality is prioritized over speed

Use non-reasoning mode (null or omit parameter) for

- Simple Q&A or factual queries
- Creative writing tasks
- When faster response times are critical
- Performance benchmarking where reasoning overhead should be excluded
- Cost optimization when reasoning doesn't improve task performance

Troubleshooting

Error: "Reasoning mode is enabled but model does not support reasoning"

**Cause:** The reasoning\_effort parameter is set to a non-null value, but the specified model\_type doesn't support reasoning.

**Resolution:**

- Verify your model type is amazon.nova-2-lite-v1:0:256k
- If using a different model, either switch to a reasoning-capable model or remove the reasoning\_effort parameter from your recipe

## RFT evaluation

### What is RFT evaluation?

RFT Evaluation allows you to assess your model's performance using custom reward functions before, during, or after reinforcement learning training. Unlike standard evaluations that use pre-defined metrics, RFT Evaluation lets you define your own success criteria through a Lambda function that scores model outputs based on your specific requirements.

### Why evaluate with RFT?

Evaluation is crucial to determine whether the RL fine-tuning process has:

- Improved model alignment with your specific use case and human values
- Maintained or improved model capabilities on key tasks
- Avoided unintended side effects such as reduced factuality, increased verbosity, or degraded performance on other tasks
- Met your custom success criteria as defined by your reward function

### When to use RFT evaluation

Use RFT Evaluation in these scenarios:

- Before RFT Training: Establish baseline metrics on your evaluation dataset
- During RFT Training: Monitor training progress with intermediate checkpoints
- After RFT Training: Validate that the final model meets your requirements
- Comparing Models: Evaluate multiple model versions using consistent reward criteria

#### Note

Use RFT Evaluation when you need custom, domain-specific metrics. For general-purpose evaluation (accuracy, perplexity, BLEU), use standard evaluation methods.

## Data format requirements

### Input data structure

RFT evaluation input data must follow the OpenAI Reinforcement Fine-Tuning format. Each example is a JSON object containing:

- `messages` – Array of conversational turns with system and user roles
- `reference_answer` – Expected output or ground truth data used by your reward function for scoring

### Data format example

```
{
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "type": "text",
 "text": "Solve for x. Return only JSON like {\"x\": <number>}. Equation: 2x + 5 = 13"
 }
]
 }
],
 "reference_answer": {
 "x": 4
 }
}
```

### Current limitations

- Text only: No multimodal inputs (images, audio, video) are supported
- Single-turn conversations: Only supports single user message (no multi-turn dialogues)
- JSON format: Input data must be in JSONL format (one JSON object per line)
- Model outputs: Evaluation is performed on generated completions from the specified model

## Preparing your evaluation recipe

### Sample notebook

For a complete example, see [Evaluation notebooks](#).

### Sample recipe configuration

```
run:
 name: nova-lite-rft-eval-job
 model_type: amazon.nova-lite-v1:0:300k
 model_name_or_path: s3://escrow_bucket/model_location # [MODIFIABLE] S3 path to your
 model or model identifier
 replicas: 1 # [MODIFIABLE] For SageMaker Training jobs only; fixed for HyperPod jobs

 data_s3_path: "" # [REQUIRED FOR HYPERPOD] Leave empty for SageMaker Training jobs
 and use TrainingInput in sagemaker python SDK
 output_s3_path: "" # [REQUIRED] Output artifact S3 path for evaluation results

evaluation:
 task: rft_eval # [FIXED] Do not modify
 strategy: rft_eval # [FIXED] Do not modify
 metric: all # [FIXED] Do not modify

Inference Configuration
inference:
 max_new_tokens: 8192 # [MODIFIABLE] Maximum tokens to generate
 top_k: -1 # [MODIFIABLE] Top-k sampling parameter
 top_p: 1.0 # [MODIFIABLE] Nucleus sampling parameter
 temperature: 0 # [MODIFIABLE] Sampling temperature (0 = deterministic)
 top_logprobs: 0 # [MODIFIABLE] Set between 1-20 to enable logprobs output

=====
Bring Your Own Reinforcement Learning Environment
=====
rl_env:
 reward_lambda_arn: arn:aws:lambda:<region>:<account_id>:function:<reward-function-
 name>
```

### Preset reward functions

Two preset reward functions (prime\_code and prime\_math) are available as a Lambda layer for easy integration with your RFT Lambda functions.



## Overview

These preset functions provide out-of-the-box evaluation capabilities for:

- `prime_code` – Code generation and correctness evaluation
- `prime_math` – Mathematical reasoning and problem-solving evaluation

## Quick setup

1. Download the Lambda layer from the [nova-custom-eval-sdk releases](#).
2. Publish Lambda layer using AWS Command Line Interface (AWS CLI):

```
aws lambda publish-layer-version \
 --layer-name preset-function-layer \
 --description "Preset reward function layer with dependencies" \
 --zip-file fileb://universal_reward_layer.zip \
 --compatible-runtimes python3.9 python3.10 python3.11 python3.12 \
 --compatible-architectures x86_64 arm64
```

3. Add the layer to your Lambda function in AWS Management Console (Select the preset-function-layer from custom layer and also add AWSSDKPandas-Python312 for numpy dependencies).
4. Import and use in your Lambda code:

```
from prime_code import compute_score # For code evaluation
from prime_math import compute_score # For math evaluation
```

## prime\_code function

Evaluates Python code generation tasks by executing code against test cases and measuring correctness.

## Example input dataset format

```
{"messages":[{"role":"user","content":"Write a function that returns the sum of two numbers."}], "reference_answer":{"inputs":["3\n5","10\n-2","0\n0"],"outputs":["8","8","0"]}}
```

```
{"messages":[{"role":"user","content":"Write a function to check if a number is even."}], "reference_answer":{"inputs":["4","7","0","-2"],"outputs":["True","False","True","True"]}}
```

## Key features

- Automatic code extraction from markdown code blocks
- Function detection and call-based testing
- Test case execution with timeout protection
- Syntax validation and compilation checks
- Detailed error reporting with tracebacks

## prime\_math function

Evaluates mathematical reasoning and problem-solving capabilities with symbolic math support.

## Input format

```
{"messages":[{"role":"user","content":"What is the derivative of x^2 + 3x?."}], "reference_answer":"2*x + 3"}
```

## Key features

- Symbolic math evaluation using SymPy
- Multiple answer formats (LaTeX, plain text, symbolic)
- Mathematical equivalence checking
- Expression normalization and simplification

## Data format requirements

### For code evaluation

- Inputs: Array of function arguments (proper types: integers, strings, etc.)
- Outputs: Array of expected return values (proper types: booleans, numbers, etc.)
- Code: Must be in Python with clear function definitions

### For math evaluation

- Reference answer: Mathematical expression or numeric value
- Response: Can be LaTeX, plain text, or symbolic notation
- Equivalence: Checked symbolically, not just string matching

## Best practices

- Use proper data types in test cases (integers vs strings, booleans vs "True")
- Provide clear function signatures in code problems
- Include edge cases in test inputs (zero, negative numbers, empty inputs)
- Format math expressions consistently in reference answers
- Test your reward function with sample data before deployment

## Error handling

Both functions include robust error handling for:

- Compilation errors in generated code
- Runtime exceptions during execution
- Malformed input data
- Timeout scenarios for infinite loops
- Invalid mathematical expressions

## Creating your reward function

### Lambda ARN requirements

Your Lambda ARN must follow this format:

```
"arn:aws:lambda:*:*:function:*SageMaker*"
```

If the Lambda does not have this naming scheme, the job will fail with this error:

```
[ERROR] Unexpected error: lambda_arn must contain one of: ['SageMaker', 'sagemaker', 'Sagemaker'] when running on SMHP platform (Key: lambda_arn)
```

## Lambda request format

Your Lambda function receives data in this format:

```
[
 {
 "id": "sample-001",
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "type": "text",
 "text": "Do you have a dedicated security team?"
 }
]
 },
 {
 "role": "nova_assistant",
 "content": [
 {
 "type": "text",
 "text": "As an AI developed by Company, I don't have a dedicated security
team..."
 }
]
 }
],
 "reference_answer": {
 "compliant": "No",
 "explanation": "As an AI developed by Company, I do not have a traditional
security team..."
 }
 }
]
```

### Note

The message structure includes the nested content array, matching the input data format. The last message with role `nova_assistant` contains the model's generated response.

## Lambda response format

Your Lambda function must return data in this format:

```
[
 {
 "id": "sample-001",
 "aggregate_reward_score": 0.75,
 "metrics_list": [
 {
 "name": "accuracy",
 "value": 0.85,
 "type": "Metric"
 },
 {
 "name": "fluency",
 "value": 0.90,
 "type": "Reward"
 }
]
 }
]
```

### Response fields

- `id` – Must match the input sample ID
- `aggregate_reward_score` – Overall score (typically 0.0 to 1.0)
- `metrics_list` – Array of individual metrics with:
  - `name` – Metric identifier (e.g., "accuracy", "fluency")
  - `value` – Metric score (typically 0.0 to 1.0)
  - `type` – Either "Metric" (for reporting) or "Reward" (used in training)

### IAM permissions

#### Required permissions

Your SageMaker execution role must have permissions to invoke your Lambda function. Add this policy to your SageMaker execution role:

```
{
```

```
"Version": "2012-10-17",
"Statement": [
 {
 "Effect": "Allow",
 "Action": [
 "lambda:InvokeFunction"
],
 "Resource": "arn:aws:lambda:region:account-id:function:function-name"
 }
]
```

## Lambda execution role

Your Lambda function's execution role needs basic Lambda execution permissions:

```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": [
 "logs:CreateLogGroup",
 "logs:CreateLogStream",
 "logs:PutLogEvents"
],
 "Resource": "arn:aws:logs:*:*:*"
 }
]
}
```

If your Lambda function accesses other AWS services (e.g., S3 for reference data, DynamoDB for logging), add those permissions to the Lambda execution role.

## Executing the evaluation job

1. **Prepare your data** – Format your evaluation data according to the data format requirements and upload your JSONL file to S3: `s3://your-bucket/eval-data/eval_data.jsonl`
2. **Configure your recipe** – Update the sample recipe with your configuration:
  - Set `model_name_or_path` to your model location
  - Set `lambda_arn` to your reward function ARN

- Set `output_s3_path` to your desired output location
- Adjust inference parameters as needed

Save the recipe as `rft_eval_recipe.yaml`

3. **Run the evaluation** – Execute the evaluation job using the provided notebook: [Evaluation notebooks](#)

4. **Monitor progress** – Monitor your evaluation job through:

- SageMaker Console: Check job status and logs
- CloudWatch Logs: View detailed execution logs
- Lambda Logs: Debug reward function issues

## Understanding evaluation results

### Output format

The evaluation job outputs results to your specified S3 location in JSONL format. Each line contains the evaluation results for one sample:

```
{
 "id": "sample-001",
 "aggregate_reward_score": 0.75,
 "metrics_list": [
 {
 "name": "accuracy",
 "value": 0.85,
 "type": "Metric"
 },
 {
 "name": "fluency",
 "value": 0.90,
 "type": "Reward"
 }
]
}
```

**Note**

The RFT Evaluation Job Output is identical to the Lambda Response format. The evaluation service passes through your Lambda function's response without modification, ensuring consistency between your reward calculations and the final results.

## Interpreting results

### Aggregate reward score

- Range: Typically 0.0 (worst) to 1.0 (best), but depends on your implementation
- Purpose: Single number summarizing overall performance
- Usage: Compare models, track improvement over training

### Individual metrics

- Metric Type: Informational metrics for analysis
- Reward Type: Metrics used during RFT training
- Interpretation: Higher values generally indicate better performance (unless you design inverse metrics)

### Performance benchmarks

What constitutes "good" performance depends on your use case:

| Score range | Interpretation | Action                                   |
|-------------|----------------|------------------------------------------|
| 0.8 - 1.0   | Excellent      | Model ready for deployment               |
| 0.6 - 0.8   | Good           | Minor improvements may be beneficial     |
| 0.4 - 0.6   | Fair           | Significant improvement needed           |
| 0.0 - 0.4   | Poor           | Review training data and reward function |



**⚠ Important**

These are general guidelines. Define your own thresholds based on business requirements, baseline model performance, domain-specific constraints, and cost-benefit analysis of further training.

## Troubleshooting

### Common issues

| Issue               | Cause                             | Solution                                     |
|---------------------|-----------------------------------|----------------------------------------------|
| Lambda timeout      | Complex reward calculation        | Increase Lambda timeout or optimize function |
| Permission denied   | Missing IAM permissions           | Verify SageMaker role can invoke Lambda      |
| Inconsistent scores | Non-deterministic reward function | Use fixed seeds or deterministic logic       |
| Missing results     | Lambda errors not caught          | Add comprehensive error handling in Lambda   |

### Debug checklist

- Verify input data follows the correct format with nested content arrays
- Confirm Lambda ARN is correct and function is deployed
- Check IAM permissions for SageMaker → Lambda invocation
- Review CloudWatch logs for Lambda errors
- Validate Lambda response matches expected format

### Best practices

- Start Simple: Begin with basic reward functions and iterate
- Test Lambda Separately: Use Lambda test events before full evaluation

- **Validate on Small Dataset:** Run evaluation on subset before full dataset
- **Version Control:** Track reward function versions alongside model versions
- **Monitor Costs:** Lambda invocations and compute time affect costs
- **Log Extensively:** Use print statements in Lambda for debugging
- **Set Timeouts Appropriately:** Balance between patience and cost
- **Document Metrics:** Clearly define what each metric measures

## Next steps

After completing RFT evaluation:

- If results are satisfactory: Deploy model to production
- If improvement needed:
  - Adjust reward function
  - Collect more training data
  - Modify training hyperparameters
  - Run additional RFT training iterations
- Continuous monitoring: Re-evaluate periodically with new data

## Implementing reward functions

### Overview

The reward function (also called scorer or grader) is the core component that evaluates model responses and provides feedback signals for training. It must be implemented as an Lambda function that accepts model responses and returns reward scores.

### Interface format

Your reward function must accept and return data in the following format:

### Sample input sample to training

```
{
 "messages": [
 {
```

```

 "role": "user",
 "content": "Do you have a dedicated security team?"
 },
 "reference_answer": {
 "compliant": "No",
 "explanation": "As an AI developed by Company, I do not have a traditional
security team..."
 }
}

```

### Sample payload for the reward lambda

The container automatically transforms your data before sending it to your Lambda function by:

1. Generating a model response for each prompt
2. Appending the assistant turn (generated response) to the messages array
3. Adding a unique id field for tracking

Your Lambda function will receive data in this transformed format:

```

{
 "id": "123",
 "messages": [
 {
 "role": "user",
 "content": "Do you have a dedicated security team?"
 },
 {
 "role": "assistant",
 "content": "As an AI developed by Amazon, I don not have a dedicated
security team..."
 }
],
 # Following section will be same as your training dataset sample
 "reference_answer": {
 "compliant": "No",
 "explanation": "As an AI developed by Company, I do not have a traditional
security team..."
 }
}

```

## Reward Lambda contract

```
def lambda_handler(event, context):
 return lambda_grader(event)

def lambda_grader(samples: list[dict]) -> list[dict]:
 """
 Args:
 samples: List of dictionaries in OpenAI format

 Example input:
 {
 "id": "123",
 "messages": [
 {
 "role": "user",
 "content": "Do you have a dedicated security team?"
 },
 {
 "role": "assistant",
 "content": "As an AI developed by Company, I don nott have a
dedicated security team..."
 }
],
 # This section will be same as your training dataset
 "reference_answer": {
 "compliant": "No",
 "explanation": "As an AI developed by Company, I do not have a
traditional security team..."
 }
 }

 Returns:
 List of dictionaries with reward scores:
 {
 "id": str, # Same id as input sample
 "aggregate_reward_score": float, # Overall score for the sample
 "metrics_list": [# OPTIONAL: Component scores
 {
 "name": str, # Name of the component score
 "value": float, # Value of the component score
 "type": str # "Reward" or "Metric"
 }
]
 }
```

```
 }
 ""
```

Input and output fields

Input fields

| Field               | Description                           | Additional notes                                  |
|---------------------|---------------------------------------|---------------------------------------------------|
| id                  | Unique identifier for the sample      | Echoed back in output. String format              |
| messages            | Ordered chat history in OpenAI format | Array of message objects                          |
| messages[].role     | Speaker of the message                | Common values: "user", "assistant", "system"      |
| messages[ ].content | Text content of the message           | Plain string                                      |
| **metadata          | Free-form information to aid grading  | Object; optional fields passed from training data |

Output fields

| Field                  | Description                                 | Additional notes                            |
|------------------------|---------------------------------------------|---------------------------------------------|
| id                     | Same identifier as input sample             | Must match input                            |
| aggregate_reward_score | Overall score for the sample                | Float (e.g., 0.0–1.0 or task-defined range) |
| metrics_list           | Component scores that make up the aggregate | Array of metric objects                     |

## Technical constraints

- **Timeout limit** – 15 minutes maximum execution time per Lambda invocation
- **Concurrency** – Must handle `rollout_worker_replicas * 64` concurrent requests
- **Reliability** – Must implement proper error handling and return valid scores consistently
- **Performance** – Optimize for fast execution (seconds, not minutes) to enable efficient training

## Best practices

- Minimize external API calls
- Use efficient algorithms and data structures
- Implement retry logic for transient failures
- Cache reusable computations
- Test thoroughly before training to ensure bug-free execution

## Using custom reward functions

Implement custom reward functions when you have task-specific evaluation criteria:

- **Define evaluation criteria** – Determine what makes a good response for your task
- **Implement Lambda function** – Create an Lambda function following the interface format
- **Test locally** – Validate your function returns correct scores for sample inputs
- **Deploy to AWS** – Deploy your Lambda and note the ARN
- **Configure recipe** – Add the Lambda ARN to your recipe's `reward_lambda_arn` field
- **Test with small dataset** – Run RFT with minimal data to verify integration

## IAM permissions

### Required permissions

Your SageMaker execution role must have permissions to invoke your Lambda function. Add this policy to your SageMaker execution role:

```
{
 "Version": "2012-10-17",
 "Statement": [
```

```
{
 "Effect": "Allow",
 "Action": [
 "lambda:InvokeFunction"
],
 "Resource": "arn:aws:lambda:region:account-id:function:function-name"
}
```

## Lambda execution role

Your Lambda function's execution role needs basic Lambda execution permissions:

```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": [
 "logs:CreateLogGroup",
 "logs:CreateLogStream",
 "logs:PutLogEvents"
],
 "Resource": "arn:aws:logs:*:*:*"
 }
]
}
```

**Additional permissions:** If your Lambda function accesses other AWS services (for example, S3 for reference data, DynamoDB for logging), add those permissions to the Lambda execution role.

### Example: LLM As a Judge reward function

This example demonstrates using Amazon Bedrock models as judges to evaluate model responses by comparing them against reference answers. This Lambda template provides a framework for customers to implement calls to Amazon Bedrock for inference requests to process judge evaluations. The Lambda function maintains the same input/output contract as other reward functions.

## Implementation

This Lambda function implements a two-stage evaluation process: the `lambda_handler` extracts model responses and reference answers from incoming samples, then the `lambda_graded` function calls Amazon Bedrock to score the semantic similarity between them. The implementation includes robust error handling with automatic retries for transient failures and supports flexible reference answer formats (both string and structured dictionary formats).

### Implementation details:

- **Retry Logic:** Implements exponential backoff (1s, 2s, 4s) for throttling exceptions to handle Bedrock API rate limits
- **Error Handling:** Returns score of 0.0 for failed evaluations rather than raising exceptions
- **Deterministic Scoring:** Uses temperature=0.0 to ensure consistent scores across evaluations
- **Flexible Reference Format:** Automatically handles both string and dictionary reference answers
- **Score Clamping:** Ensures all scores fall within valid [0.0, 1.0] range
- **Model Agnostic:** Change `JUDGE_MODEL_ID` to use any Amazon Bedrock model (Nova, Llama, Mistral, etc.)

```
"""
LLM Judge Lambda POC - Working implementation using Amazon Bedrock
"""

import json
import time
import boto3

bedrock_runtime = boto3.client('bedrock-runtime', region_name='us-east-1')
JUDGE_MODEL_ID = "anthropic.claude-3-5-sonnet-20240620-v1:0"
SYSTEM_PROMPT = "You must output ONLY a number between 0.0 and 1.0. No explanations, no text, just the number."

JUDGE_PROMPT_TEMPLATE = """Compare the following two responses and rate how similar they are on a scale of 0.0 to 1.0, where:
- 1.0 means the responses are semantically equivalent (same meaning, even if worded differently)
- 0.5 means the responses are partially similar
- 0.0 means the responses are completely different or contradictory
```



Response A: {response\_a}

Response B: {response\_b}

Output ONLY a number between 0.0 and 1.0. No explanations."

```
def lambda_graded(response_a: str, response_b: str, max_retries: int = 3) -> float:
 """Call Bedrock to compare responses and return similarity score."""
 prompt = JUDGE_PROMPT_TEMPLATE.format(response_a=response_a, response_b=response_b)

 for attempt in range(max_retries):
 try:
 response = bedrock_runtime.converse(
 modelId=JUDGE_MODEL_ID,
 messages=[{"role": "user", "content": [{"text": prompt}]}],
 system=[{"text": SYSTEM_PROMPT}],
 inferenceConfig={"temperature": 0.0, "maxTokens": 10}
)
 print(f"Bedrock call successful: {response}")
 output = response['output']['message']['content'][0]['text'].strip()
 score = float(output)
 print(f"Score parsed: {score}")
 return max(0.0, min(1.0, score))

 except Exception as e:
 if "ThrottlingException" in str(e) and attempt < max_retries - 1:
 time.sleep(2 ** attempt)
 else:
 print(f"Bedrock call failed: {e}")
 return None

 return None

def lambda_handler(event, context):
 """AWS Lambda handler - processes samples from RFTEvalInvoker."""
 try:
 samples = event if isinstance(event, list) else [event]
 results = []

 for sample in samples:
 sample_id = sample.get("id", "unknown")
 messages = sample.get("messages", [])
```

```

Extract assistant response (response A)
response_a = ""
for msg in messages:
 if msg.get("role") in ["assistant", "nova_assistant"]:
 response_a = msg.get("content", "")
 break

Extract reference answer from root level (no longer in metadata)
reference_answer = sample.get("reference_answer", "")

Handle both string and dict reference_answer formats
if isinstance(reference_answer, dict):
 # If reference_answer is a dict, extract the explanation or compliant
 response_b = reference_answer.get("explanation",
reference_answer.get("compliant", ""))
else:
 response_b = reference_answer

if not response_a or not response_b:
 results.append({
 "id": sample_id,
 "aggregate_reward_score": 0.0,
 "metrics_list": [{"name": "similarity_score", "value": 0.0, "type":
"Metric"}]
 })
 continue

Get similarity score
score = lambda_graded(response_a, response_b)

results.append({
 "id": sample_id,
 "aggregate_reward_score": score,
 "metrics_list": [
 {
 "name": "similarity_score",
 "value": score,
 "type": "Metric"
 }
]
})

```

```
 return {"statusCode": 200, "body": json.dumps(results)}

 except Exception as e:
 print(f"Error: {e}")
 return {"statusCode": 500, "body": json.dumps({"error": str(e)})}
```

## Input format

The Lambda receives the same input format as other reward functions:

```
{
 "id": "sample-001",
 "messages": [
 {
 "role": "user",
 "content": "Do you have a dedicated security team?"
 },
 {
 "role": "assistant",
 "content": "As an AI developed by Amazon, I don't have a dedicated security
team..."
 }
],
 "reference_answer": {
 "compliant": "No",
 "explanation": "As an AI developed by Company, I do not have a traditional
security team..."
 },
 "my_custom_field": "custom_value"
}
```

## Output format

```
{
 "id": "sample-001",
 "aggregate_reward_score": 0.85,
 "metrics_list": [
 {
 "name": "similarity_score",
 "value": 0.85,
 "type": "Metric"
 }
]
}
```

```
}
```

## Deployment considerations

You may also need to adjust the prompt template and inference parameters based on your chosen model's capabilities and API format.

- **IAM Permissions:** Lambda execution role must have `bedrock:InvokeModel` permission for your chosen model
- **Timeout:** Set Lambda timeout to at least 60 seconds to accommodate Bedrock API latency and retries
- **Region:** Deploy in a region where your chosen Bedrock model is available
- **Cost:** Monitor Bedrock API usage as each evaluation makes one API call per sample
- **Throughput:** For large-scale evaluations, request increased Bedrock quotas to avoid throttling

## Increasing Bedrock Throughput

If you experience throttling during evaluation, increase your Bedrock model quotas:

- Navigate to the AWS Service Quotas console
- Search for "Bedrock" and select your region
- Find the quota for your chosen model (for example, "Invocations per minute for Claude 3.5 Sonnet")
- Click "Request quota increase" and specify your desired throughput
- Provide justification for the increase (for example, "RFT evaluation workload")

The Lambda's built-in retry logic handles occasional throttling, but sustained high-volume evaluations require appropriate quota increases.

## Required IAM Policy:

```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": [
```

```
 "bedrock:InvokeModel"
],
 "Resource": "arn:aws:bedrock:*::foundation-model/*"
 }
]
}
```

## Iterative training

Iterative training is a systematic approach to fine-tuning models through multiple training cycles, where each round builds on the previous checkpoint by addressing specific weaknesses discovered through evaluation. This method enables targeted improvements to model performance by incorporating curated examples that address failure modes, adapting to changing requirements, and validating enhancements incrementally rather than committing to a single long training run. The process typically follows patterns like SFT (Supervised Fine-Tuning) followed by RFT (Reward-based Fine-Tuning), with checkpoints stored in AWS-managed escrow S3 buckets that can be referenced for subsequent training iterations while maintaining consistency in model type and training technique throughout the pipeline.

For more details, refer to [???](#).

## Amazon Nova customization on SageMaker HyperPod

You can customize Amazon Nova models, including the enhanced Amazon Nova 2.0 models, using [Amazon Nova recipes](#) and train them on Hyperpod. A recipe is a YAML configuration file that provides details to SageMaker AI on how to run your model customization job. SageMaker HyperPod supports two types of services: Forge and Non-forge.

Hyperpod offers high-performance computing with optimized GPU instances and Amazon FSx for Lustre storage, robust monitoring through integration with tools like TensorBoard, flexible checkpoint management for iterative improvement, seamless deployment to Amazon Bedrock for inference, and efficient scalable multi-node distributed training-all working together to provide organizations with a secure, performant, and flexible environment to tailor Amazon Nova models to their specific business requirements.

Amazon Nova customization on SageMaker HyperPod stores model artifacts including model checkpoints in a service-managed Amazon S3 bucket. Artifacts in the service-managed bucket are encrypted with SageMaker AI-managed AWS KMS keys. Service-managed Amazon S3 buckets don't

currently support data encryption using customer-managed KMS keys. You can use this checkpoint location for evaluation jobs or Amazon Bedrock inference.

Standard pricing can apply for compute instances, Amazon S3 storage, and FSx for Lustre. For pricing details, see [Hyperpod pricing](#), [Amazon S3 pricing](#), and [FSx for Lustre pricing](#).

## Compute requirements for Amazon Nova 2 models

The following tables summarize the computational requirements for SageMaker HyperPod and SageMaker AI training jobs training for Amazon Nova 2 models.

### Nova 2 Training Requirements

| Training Technique                         | Minimum Instances | Instance Type | GPU Count | Notes                                           | Supported Models |
|--------------------------------------------|-------------------|---------------|-----------|-------------------------------------------------|------------------|
| SFT (LoRA)                                 | 4                 | P5.48xlarge   | 16        | Parameter-efficient fine-tuning                 | Nova 2 Lite      |
| SFT (Full Rank)                            | 4                 | P5.48xlarge   | 32        | Full model fine-tuning                          | Nova 2 Lite      |
| RFT on SageMaker Training Jobs (LoRA)      | 2                 | P5.48xlarge   | 16        | Custom Reward Functions in your AWS Environment | Nova 2 Lite      |
| RFT on SageMaker Training Jobs (Full Rank) | 4                 | P5.48xlarge   | 32        | 32K context length                              | Nova 2 Lite      |
| RFT on SageMaker HyperPod                  | 8                 | P5.48xlarge   | 64        | Default 8192 context length                     | Nova 2 Lite      |
| CPT                                        | 4                 | P5.48xlarge   | 16        | Processes approxima                             | Nova 2 Lite      |

| Training Technique | Minimum Instances | Instance Type | GPU Count | Notes                                 | Supported Models |
|--------------------|-------------------|---------------|-----------|---------------------------------------|------------------|
|                    |                   |               |           | tely 400M tokens per instance per day |                  |

To optimize your Amazon Nova model customization workflows on Hyperpod, follow these recommended best practices for efficient training, resource management, and successful model deployment.

## Best Practices for Amazon Nova customization

### Overview

This section provides an overview of customization techniques and helps you choose the best approach for your needs and available data.

### Two stages of LLM training

Large language model training consists of two major stages: pre-training and post-training. During pre-training, the model processes tokens of raw text and optimizes for next-token prediction. This process creates a pattern completer that absorbs syntax, semantics, facts, and reasoning patterns from web and curated text. However, the pre-trained model doesn't understand instructions, user goals, or context-appropriate behavior. It continues text in whatever style fits its training distribution. A pre-trained model autocompletes rather than follows directions, produces inconsistent formatting, and can mirror undesirable biases or unsafe content from the training data. Pre-training builds general competence, not task usefulness.

Post-training transforms the pattern completer into a useful assistant. You run multiple rounds of Supervised Fine-Tuning (SFT) to teach the model to follow instructions, adhere to schemas and policies, call tools, and produce reliable outputs by imitating high-quality demonstrations. This alignment teaches the model to respond to prompts as tasks rather than text to continue. You then apply Reinforcement Fine-Tuning (RFT) to optimize behavior using measurable feedback (such as verifiers or an LLM-as-a-judge), balancing trade-offs like accuracy versus brevity, safety versus coverage, or multi-step reasoning under constraints. In practice, you alternate SFT and RFT in

cycles to shape the pre-trained model into a reliable, policy-aligned system that performs complex tasks consistently.

## Choose the right customization approach

In this section we will cover post training customization strategies: RFT and SFT.

### Reinforcement fine-tuning (RFT)

Reinforcement fine-tuning improves model performance through feedback signals—measurable scores or rewards that indicate response quality—rather than direct supervision with exact correct answers. Unlike traditional supervised fine-tuning that learns from input-output pairs, RFT uses reward functions to evaluate model responses and iteratively optimizes the model to maximize these rewards. This approach works well for tasks where defining the exact correct output is challenging, but you can reliably measure response quality. RFT enables models to learn complex behaviors and preferences through trial and feedback, making it ideal for applications that require nuanced decision-making, creative problem-solving, or adherence to specific quality criteria that you can programmatically evaluate. For example, answering complex legal questions is an ideal use case for RFT because you want to teach the model how to reason better to answer questions more accurately.

#### How it works

In reinforcement fine-tuning, you start from an instruction-tuned baseline and treat each prompt like a small tournament. For a given input, you sample a handful of candidate answers from the model, score each one with the reward function, then rank them within that group. The update step nudges the model to make higher-scoring candidates more likely next time and lower-scoring ones less likely, while a stay-close-to-baseline constraint keeps behavior from drifting or becoming verbose or exploitative. You repeat this loop over many prompts, refreshing hard cases, tightening verifiers or judge rubrics when you see exploits, and continuously tracking task metrics.

#### When to use RFT

Tasks that benefit most from RFT share several traits. They have measurable success signals even when a single correct output is hard to specify. They admit partial credit or graded quality so you can rank better versus worse answers within a prompt or using a reward function. They involve multiple objectives that must be balanced (such as accuracy with brevity, clarity, safety, or cost). They require adherence to explicit constraints that you can programmatically check. They operate in tool-mediated or environment-based settings where outcomes are observable (success or failure, latency, resource use). They occur in low-label regimes where collecting gold targets is expensive



but automated or rubric-based feedback is plentiful. RFT works best when you can turn quality into a reliable scalar or ranking and want the model to preferentially amplify higher-scoring behaviors without needing exhaustive labeled targets.

**Consider other methods when:**

- You have plentiful, reliable labeled input-output pairs – Use SFT
- The main gap is knowledge or jargon – Use retrieval-augmented generation (RAG)
- Your reward signal is noisy or unreliable and you can't fix it with better rubrics or checkers – Stabilize that first before RFT

**When not to use RFT**

Avoid RFT in these situations:

- You can cheaply produce reliable labeled input-output pairs (SFT is simpler, cheaper, and more stable)
- The gap is knowledge or jargon rather than behavior (use RAG)
- Your reward signal is noisy, sparse, easy to game, or expensive or slow to compute (fix the evaluator first)
- Baseline performance is near-zero (bootstrap with SFT before optimizing preferences)
- The task has deterministic schemas, strict formatting, or a single correct answer (SFT or rule-based validation works better)
- Tight latency or cost budgets can't absorb the extra sampling or exploration RFT requires
- Safety or policy constraints aren't crisply specified and enforceable in the reward

If you can point to "the right answer," use SFT. If you need new knowledge, use RAG. Use RFT only after you have a solid baseline and a robust, fast, hard-to-exploit reward function.

**Supervised fine-tuning (SFT)**

Supervised fine-tuning trains the LLM on a dataset of human-labeled input-output pairs for your task. You provide examples of prompts (questions, instructions, and so on) with the correct or desired responses, and continue training the model on these examples. The model adjusts its weights to minimize a supervised loss (typically cross-entropy between its predictions and the target output tokens). This is the same training used in most supervised machine learning tasks, applied to specialize an LLM.

SFT changes behavior, not knowledge. It doesn't teach the model new facts or jargon it didn't see in pre-training. It teaches the model how to answer, not what to know. If you need new domain knowledge (such as internal terminology), use retrieval-augmented generation (RAG) to provide that context at inference time. SFT then adds the desired instruction-following behavior on top.

## How it works

SFT optimizes LLM by minimizing the average cross-entropy loss on response tokens, treating prompt tokens as context and masking them from the loss. The model internalizes your target style, structure, and decision rules, learning to generate the correct completion for each prompt. For example, to classify documents into custom categories, you fine-tune the model with prompts (the document text) and labeled completions (the category labels). You train on those pairs until the model outputs the right label for each prompt with high probability.

You can perform SFT with as few as a few hundred examples and scale up to a few hundred thousand. SFT samples must be high quality and directly aligned with the desired model behavior.

## When to use SFT

Use SFT when you have a well-defined task with clear desired outputs. If you can explicitly state "Given X input, the correct output is Y" and gather examples of such mappings, supervised fine-tuning is a good choice. SFT excels in these scenarios:

- **Structured or complex classification tasks** – Classify internal documents or contracts into many custom categories. With SFT, the model learns these specific categories better than prompting alone.
- **Question-answering or transformation tasks with known answers** – Fine-tune a model to answer questions from a company's knowledge base, or convert data between formats where each input has a correct response.
- **Formatting and style consistency** – Train the model to always respond in a certain format or tone by fine-tuning on examples of the correct format or tone. For instance, training on prompt-response pairs that demonstrate a particular brand voice teaches the model to generate outputs with that style. Instruction-following behavior is often initially taught through SFT on curated examples of good assistant behavior.

SFT is the most direct way to teach an LLM a new skill or behavior when you can specify what the right behavior looks like. It uses the model's existing language understanding and focuses it on

your task. Use SFT when you want the model to do a specific thing and you have or can create a dataset of examples.

Use SFT when you can assemble high-quality prompt and response pairs that closely mirror the behavior you want. It fits tasks with clear targets or deterministic formats such as schemas, function or tool calls, and structured answers where imitation is an appropriate training signal. The goal is behavior shaping: teaching the model to treat prompts as tasks, follow instructions, adopt tone and refusal policies, and produce consistent formatting. Plan for at least hundreds of demonstrations, with data quality, consistency, and deduplication mattering more than raw volume. For a straightforward, cost-efficient update, use parameter-efficient methods like Low-Rank Adaptation to train small adapters while leaving most of the backbone untouched.

### When not to use SFT

Don't use SFT when the gap is knowledge rather than behavior. It doesn't teach the model new facts, jargon, or recent events. In those cases, use retrieval-augmented generation to bring external knowledge at inference. Avoid SFT when you can measure quality but can't label a single right answer. Use reinforcement fine-tuning with verifiable rewards or an LLM-as-a-judge to optimize those rewards directly. If your needs or content change frequently, rely on retrieval and tool use rather than retraining the model.

### Topics

- [Nova Forge SDK](#)
- [Creating a SageMaker HyperPod EKS cluster with restricted instance group \(RIG\)](#)
- [Amazon SageMaker HyperPod Essential Commands Guide](#)
- [Nova Forge access and setup for](#)
- [Training for Amazon Nova models](#)
- [Evaluating your trained model](#)
- [Monitoring HyperPod jobs with MLflow](#)

## Nova Forge SDK

The Amazon Nova Forge SDK is a comprehensive Python SDK that provides a unified, programmatic interface for the complete Amazon Amazon Nova model customization lifecycle. The SDK simplifies model customization by offering a single, consistent API for training, evaluation, monitoring, deployment, and inference across Amazon SageMaker and Amazon Bedrock platforms.

For more information, see [Nova Forge SDK](#).

## Creating a SageMaker HyperPod EKS cluster with restricted instance group (RIG)

To customize a model on Hyperpod, the necessary infrastructure must be set up. For details on setting up a SageMaker HyperPod EKS cluster with a restricted instance group (RIG), visit the [workshop](#), which provides a detailed walkthrough of the setup process.

## Amazon SageMaker HyperPod Essential Commands Guide

Amazon SageMaker HyperPod provides extensive command-line functionality for managing training workflows. This guide covers essential commands for common operations, from connecting to your cluster to monitoring job progress.

### Prerequisites

Before using these commands, ensure you have completed the following setup:

- SageMaker HyperPod cluster with RIG created (typically in us-east-1)
- Output Amazon S3 bucket created for training artifacts
- IAM roles configured with appropriate permissions
- Training data uploaded in correct JSONL format
- FSx for Lustre sync completed (verify in cluster logs on first job)

### Topics

- [Installing Recipe CLI](#)
- [Connecting to your cluster](#)
- [Starting a training job](#)
- [Checking job status](#)
- [Monitoring job logs](#)
- [Listing active jobs](#)
- [Canceling a job](#)
- [Running an evaluation job](#)
- [Common issues](#)

## Installing Recipe CLI

Navigate to the root of your recipe repository before running the installation command.

**Use the Hyperpodrecipes repository if using Non Forge customization techniques, for Forge based customization refer to the forge specific recipe repository.**

Run the following commands to install the SageMaker HyperPod CLI:

### Note

Make sure you aren't in an active conda / anaconda / miniconda environment or another virtual environment

If you are, please exit the environment using:

- `conda deactivate` for conda / anaconda / miniconda environments
- `deactivate` for python virtual environments

If you are using a Non Forge customization technique, download the sagemaker-hyperpod-recipes as shown below:

```
git clone -b release_v2 https://github.com/aws/sagemaker-hyperpod-cli.git
cd sagemaker-hyperpod-cli
pip install -e .
cd ..
root_dir=$(pwd)
export PYTHONPATH=${root_dir}/sagemaker-hyperpod-cli/src/hyperpod_cli/
sagemaker_hyperpod_recipes/launcher/nemo/nemo_framework_launcher/launcher_scripts:
$PYTHONPATH
curl -fsSL -o get_helm.sh https://raw.githubusercontent.com/helm/helm/main/scripts/get-helm-3
chmod 700 get_helm.sh
./get_helm.sh
rm -f ./get_helm.sh
```

If you are a **Forge Subscriber**, you should be downloading the recipes using below mentioned process.

```
mkdir NovaForgeHyperpodCLI
cd NovaForgeHyperpodCLI
aws s3 cp s3://nova-forge-c7363-206080352451-us-east-1/v1/ ./ --recursive
pip install -e .

curl -fsSL -o get_helm.sh https://raw.githubusercontent.com/helm/helm/main/scripts/get-helm-3
chmod 700 get_helm.sh
./get_helm.sh
rm -f ./get_helm.sh
```

 **Tip**

To use a [new virtual environment](#) before running `pip install -e .`, run:

- `python -m venv nova_forge`
- `source nova_forge/bin/activate`
- Your command line will now display `(nova_forge)` at the beginning of your prompt
- This ensures there are no competing dependencies when using the CLI

**Purpose:** Why do we do `pip install -e .`?

This command installs the SageMaker HyperPod CLI in editable mode, allowing you to use updated recipes without reinstalling each time. It also enables you to add new recipes that the CLI can automatically pick up.

## Connecting to your cluster

Connect the SageMaker HyperPod CLI to your cluster before running any jobs:

```
export AWS_REGION=us-east-1 && SageMaker HyperPod connect-cluster --cluster-name
<your-cluster-name> --region us-east-1
```

**⚠ Important**

This command creates a context file (`/tmp/hyperpod_context.json`) that subsequent commands require. If you see an error about this file not found, re-run the `connect` command.

**Pro tip:** You can further configure your cluster to always use the `kubeflow` namespace by adding the `--namespace kubeflow` argument to your command as follows:

```
export AWS_REGION=us-east-1 && \
hyperpod connect-cluster \
--cluster-name <your-cluster-name> \
--region us-east-1 \
--namespace kubeflow
```

This saves you the effort of adding the `-n kubeflow` in every command when interacting with your jobs.

## Starting a training job

**ℹ Note**

If running PPO/RFT jobs, ensure you add label selector settings to `src/hyperpod_cli/sagemaker_hyperpod_recipes/recipes_collection/cluster/k8s.yaml` so that all pods are schedule on the same node.

```
label_selector:
 required:
 sagemaker.amazonaws.com/instance-group-name:
 - <rig_group>
```

Launch a training job using a recipe with optional parameter overrides:

```
hyperpod start-job -n kubeflow \
--recipe fine-tuning/nova/nova_1_0/nova_micro/SFT/nova_micro_1_0_p5_p4d_gpu_lora_sft \
--override-parameters '{
```

```
"instance_type": "ml.p5.48xlarge",
 "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-SFT-latest"
}'
```

## Expected output:

```
Final command: python3 <path_to_your_installation>/NovaForgeHyperpodCLI/src/
hyperpod_cli/sagemaker_hyperpod_recipes/main.py recipes=fine-tuning/nova/
nova_micro_p5_gpu_sft cluster_type=k8s cluster=k8s base_results_dir=/local/home/
<username>/results cluster.pullPolicy="IfNotPresent" cluster.restartPolicy="OnFailure"
cluster.namespace="kubeflow" container="708977205387.dkr.ecr.us-east-1.amazonaws.com/
nova-fine-tune-repo:HP-SFT-DATAMIX-latest"
```

```
Prepared output directory at /local/home/<username>/results/<job-name>/k8s_templates
Found credentials in shared credentials file: ~/.aws/credentials
Helm script created at /local/home/<username>/results/<job-name>/<job-name>_launch.sh
Running Helm script: /local/home/<username>/results/<job-name>/<job-name>_launch.sh
```

```
NAME: <job-name>
LAST DEPLOYED: Mon Sep 15 20:56:50 2025
NAMESPACE: kubeflow
STATUS: deployed
REVISION: 1
TEST SUITE: None
Launcher successfully generated: <path_to_your_installation>/NovaForgeHyperpodCLI/src/
hyperpod_cli/sagemaker_hyperpod_recipes/launcher/nova/k8s_templates/SFT

{
 "Console URL": "https://us-east-1.console.aws.amazon.com/sagemaker/home?region=us-
east-1#/cluster-management/<your-cluster-name>"
}
```

## Checking job status

Monitor your running jobs using kubectl:

```
kubectl get pods -o wide -w -n kubeflow | (head -n1 ; grep <your-job-name>)
```

## Understanding pod statuses

The following table explains common pod statuses:



| Status                          | Description                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------|
| Pending                         | Pod accepted but not yet scheduled onto a node, or waiting for container images to be pulled |
| Running                         | Pod bound to a node with at least one container running or starting                          |
| Succeeded                       | All containers completed successfully and won't restart                                      |
| Failed                          | All containers terminated with at least one ending in failure                                |
| Unknown                         | Pod state cannot be determined (usually due to node communication issues)                    |
| CrashLoopBackOff                | Container repeatedly failing; Kubernetes backing off from restart attempts                   |
| ImagePullBackOff / ErrImagePull | Unable to pull container image from registry                                                 |
| OOMKilled                       | Container terminated for exceeding memory limits                                             |
| Completed                       | Job or Pod finished successfully (batch job completion)                                      |

** Tip**

Use the `-w` flag to watch pod status updates in real-time. Press `Ctrl+C` to stop watching.

## Monitoring job logs

You can view your logs one of three ways:

### Using CloudWatch

Your logs are available in your AWS account that contains the Hyperpodcluster under CloudWatch. To view them in your browser, navigate to the CloudWatch homepage in your account and search for your cluster name. For example, if your cluster were called `my-hyperpod-rig` the log group would have the prefix:

- **Log group:** `/aws/sagemaker/Clusters/my-hyperpod-rig/{UUID}`
- Once you're in the log group, you can find your specific log using the node instance ID such as `-hyperpod-i-00b3d8a1bf25714e4`.
  - `i-00b3d8a1bf25714e4` here represents the Hyperpodfriendly machine name where your training job is running. Recall how in the previous command `kubectl get pods -o wide -w -n kubeflow | (head -n1 ; grep my-cpt-run)` output we captured a column called **NODE**.
  - The "master" node run was in this case running on `hyperpod-i-00b3d8a1bf25714e4` and thus we'll use that string to select the log group to view. Select the one that says `SagemakerHyperPodTrainingJob/rig-group/[NODE]`

## Using CloudWatch Insights

If you have your job name handy and don't wish to go through all the steps above, you can simply query all logs under `/aws/sagemaker/Clusters/my-hyperpod-rig/{UUID}` to find the individual log.

CPT:

```
fields @timestamp, @message, @logStream, @log
| filter @message like /(?!Starting CPT Job/
| sort @timestamp desc
| limit 100
```

For job completion replace `Starting CPT Job` with `CPT Job completed`

Then you can click through the results and pick the one that says "Epoch 0" since that will be your master node.

## Using the AWS CLI

You may choose to tail your logs using the AWS CLI. Before doing so, please check your AWS CLI version using `aws --version`. It is also recommended to use this utility script that helps in live log tracking in your terminal

**for V1:**

```
aws logs get-log-events \
--log-group-name /aws/sagemaker/YourLogGroupName \
--log-stream-name YourLogStream \
--start-from-head | jq -r '.events[].message'
```

**for V2:**

```
aws logs tail /aws/sagemaker/YourLogGroupName \
--log-stream-name YourLogStream \
--since 10m \
--follow
```

## Listing active jobs

View all jobs running in your cluster:

```
hyperpod list-jobs -n kubeflow
```

**Example output:**

```
{
 "jobs": [
 {
 "Name": "test-run-nhgza",
 "Namespace": "kubeflow",
 "CreationTime": "2025-10-29T16:50:57Z",
 "State": "Running"
 }
]
}
```

## Canceling a job

Stop a running job at any time:

```
hyperpod cancel-job --job-name <job-name> -n kubeflow
```

## Finding your job name

## Option 1: From your recipe

The job name is specified in your recipe's run block:

```
run:
 name: "my-test-run" # This is your job name
 model_type: "amazon.nova-micro-v1:0:128k"
 ...
```

## Option 2: From list-jobs command

Use `hyperpod list-jobs -n kubeflow` and copy the Name field from the output.

## Running an evaluation job

Evaluate a trained model or base model using an evaluation recipe.

### Prerequisites

Before running evaluation jobs, ensure you have:

- Checkpoint Amazon S3 URI from your training job's `manifest.json` file (for trained models)
- Evaluation dataset uploaded to Amazon S3 in the correct format
- Output Amazon S3 path for evaluation results

### Command

Run the following command to start an evaluation job:

```
hyperpod start-job -n kubeflow \
 --recipe evaluation/nova/nova_2_0/nova_lite/
nova_lite_2_0_p5_48x1_gpu_bring_your_own_dataset_eval \
 --override-parameters '{
 "instance_type": "p5.48xlarge",
 "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM-
HP-Eval-latest",
 "recipes.run.name": "<your-eval-job-name>",
 "recipes.run.model_name_or_path": "<checkpoint-s3-uri>",
 "recipes.run.output_s3_path": "s3://<your-bucket>/eval-results/",
 "recipes.run.data_s3_path": "s3://<your-bucket>/eval-data.jsonl"
```

```
}'
```

### Parameter descriptions:

- `recipes.run.name`: Unique name for your evaluation job
- `recipes.run.model_name_or_path`: Amazon S3 URI from `manifest.json` or base model path (e.g., `nova-micro/prod`)
- `recipes.run.output_s3_path`: Amazon S3 location for evaluation results
- `recipes.run.data_s3_path`: Amazon S3 location of your evaluation dataset

### Tips:

- **Model-specific recipes:** Each model size (micro, lite, pro) has its own evaluation recipe
- **Base model evaluation:** Use base model paths (e.g., `nova-micro/prod`) instead of checkpoint URIs to evaluate base models

### Evaluation data format

#### Input format (JSONL):

```
{
 "metadata": "{key:4, category:'apple'}",
 "system": "arithmetic-patterns, please answer the following with no other words: ",
 "query": "What is the next number in this series? 1, 2, 4, 8, 16, ?",
 "response": "32"
}
```

#### Output format:

```
{
 "prompt": "[{'role': 'system', 'content': 'arithmetic-patterns, please answer the following with no other words: '}, {'role': 'user', 'content': 'What is the next number in this series? 1, 2, 4, 8, 16, ?'}]",
 "inference": "['32']",
 "gold": "32",
 "metadata": "{key:4, category:'apple'}"
}
```

### Field descriptions:

- **prompt:** Formatted input sent to the model
- **inference:** Model's generated response
- **gold:** Expected correct answer from input dataset
- **metadata:** Optional metadata passed through from input

## Common issues

- **ModuleNotFoundError:** No module named 'nemo\_launcher', you might've to add `nemo_launcher` to your python path based on where `hyperpod_cli` is installed. Sample command:

```
export PYTHONPATH=<path_to_hyperpod_cli>/sagemaker-hyperpod-cli/src/hyperpod_cli/
sagemaker_hyperpod_recipes/launcher/nemo/nemo_framework_launcher/launcher_scripts:
$PYTHONPATH
```

- **FileNotFoundError:** [Errno 2] No such file or directory: '/tmp/hyperpod\_current\_context.json' indicates you missed running the hyperpod connect cluster command.
- If you don't see your job scheduled, double check if the output of your SageMaker HyperPod CLI has this section with job names and other metadata. If not, re-install helm chart by running:

```
curl -fsSL -o get_helm.sh https://raw.githubusercontent.com/helm/helm/main/scripts/
get-helm-3
chmod 700 get_helm.sh
./get_helm.sh
rm -f ./get_helm.sh
```

## Nova Forge access and setup for

To set up Amazon Nova Forge for use with Jobs, you need to:

- Subscribe to Amazon Nova Forge
- Set up a cluster

### Topics

- [Subscribe to Amazon Nova Forge](#)

- [Set up infrastructure](#)
- [Responsible AI](#)

## Subscribe to Amazon Nova Forge

To access Amazon Nova Forge features, complete the following steps:

1. Verify administrator access to the AWS account.
2. Navigate to the SageMaker AI console and [request access to Amazon Nova Forge](#).
3. Wait for the Amazon Nova team to email a confirmation after the subscription request is approved.
4. Tag your execution role with the `forge-subscription` tag. This tag is required for accessing Amazon Nova Forge features and checkpoints. Add the following tag to your execution role:
  - Key: `forge-subscription`
  - Value: `true`

### Note

Standard Amazon Nova features remain available without a Forge subscription. Amazon Nova Forge is designed for building custom frontier models with control and flexibility across all model training phases.

## Set up infrastructure

Once your Amazon Nova Forge subscription is approved, set up the necessary infrastructure to use Forge-enabled features. For detailed instructions on creating a EKS cluster with a restricted instance group (RIG), follow the [workshop](#) instructions.

## Responsible AI

**Content moderation settings:** Amazon Nova Forge customers have access to Customizable Content Moderation Settings (CCMS) for Amazon Nova Lite 1.0 and Pro 1.0 models. CCMS allows you to adjust content moderation controls to align with your specific business requirements while maintaining essential responsible AI safeguards. To determine if your business use case qualifies for CCMS, contact your Amazon Web Services account manager.

Amazon Nova Forge provides a Responsible AI Toolkit that includes training data, evaluation benchmarks, and runtime controls to help you align your models with Amazon Nova's responsible AI guidelines.

**Training data:** The "RAI" category in data mixing contains cases and scenarios emphasizing responsible AI principles, safety considerations, and responsible technology deployment. Use these to align your models responsibly during continued pre-training.

**Evaluations:** Benchmark tasks are available to test your model's ability to detect and reject inappropriate, harmful, or incorrect content. Use these evaluations to measure the difference between base model performance and your custom model performance.

## Training for Amazon Nova models

Training Amazon Nova models on SageMaker HyperPod supports multiple techniques including Continued Pre-Training (CPT), Supervised Fine-Tuning (SFT), and Reinforcement Fine-Tuning (RFT). Each technique serves different customization needs and can be applied to different Amazon Nova model versions.

### Topics

- [Continued pre-training \(CPT\)](#)
- [Supervised fine-tuning \(SFT\)](#)
- [Reinforcement Fine-Tuning \(RFT\) on SageMaker HyperPod](#)

### Continued pre-training (CPT)

Continued pre-training (CPT) is a training technique that extends the pre-training phase of a foundation model by exposing it to additional unlabeled text from specific domains or corpora. Unlike supervised fine-tuning, which requires labeled input-output pairs, CPT trains on raw documents to help the model acquire deeper knowledge of new domains, learn domain-specific terminology and writing patterns, and adapt to particular content types or subject areas.

This approach is particularly valuable when you have large volumes (tens of billions of tokens) of domain-specific text data, such as legal documents, medical literature, technical documentation, or proprietary business content, and you want the model to develop native fluency in that domain. Generally, after the CPT stage, the model needs to undergo additional instruction tuning stages to enable the model to use the newly acquired knowledge and complete useful tasks.

### Supported models



CPT is available for the following Amazon Nova models:

- Nova 1.0 (Micro, Lite, Pro)
- Nova 2.0 (Lite)

### When to use Nova 1.0 versus Nova 2.0

The Amazon Nova family of models offers multiple price-performance operating points to optimize between accuracy, speed, and cost.

Choose Nova 2.0 when you need the following:

- Advanced reasoning capabilities for complex analytical tasks
- Superior performance on coding, math, and scientific problem-solving
- Longer context length support
- Better multilingual performance

#### Note

The larger model is not always better. Consider the cost-performance tradeoff and your specific business requirements when selecting between Nova 1.0 and Nova 2.0 models.

### CPT on Nova 2.0

Amazon Nova Lite 2.0 is a reasoning model trained on larger and more diverse datasets than Nova Lite 1.0. Despite being a larger model, Nova Lite 2.0 delivers faster inference than Nova Lite 1.0 while offering enhanced reasoning capabilities, longer context lengths, and improved multilingual performance.

CPT on Nova 2.0 allows you to extend these advanced capabilities with your domain-specific data, enabling the model to develop deep expertise in specialized areas while maintaining its superior reasoning and analytical abilities.

### Sample CPT recipe

The following is a sample recipe for CPT. You can find this recipe and others in the [recipes repository](#).

```

Note:
This recipe can run on p5.48xlarge
Run config
run:
 name: "my-cpt-run" # A descriptive name for your training
 job
 model_type: "amazon.nova-2-lite-v1:0:256k" # Model variant specification, do not
 change
 model_name_or_path: "nova-lite-2/prod" # Base model path, do not change
 replicas: 8 # Number of compute instances for
 training, allowed values are 4, 8, 16, 32
 data_s3_path: "" # Customer data paths
 validation_data_s3_path: "" # Customer validation data paths
 output_s3_path: "" # Output artifact path, job-specific
 configuration - not compatible with standard SageMaker Training Jobs
 mlflow_tracking_uri: "" # Required for MLFlow
 mlflow_experiment_name: "my-cpt-experiment" # Optional for MLFlow. Note: leave this
 field non-empty
 mlflow_run_name: "my-cpt-run" # Optional for MLFlow. Note: leave this
 field non-empty

Training specific configs
training_config:
 task_type: cpt
 max_length: 8192 # Maximum context window size (tokens)
 global_batch_size: 256 # Global batch size, allowed values are
 32, 64, 128, 256.

 trainer:
 max_steps: 10 # The number of training steps to run
 total
 val_check_interval: 10 # The number of steps between running
 validation. Integer count or float percentage
 limit_val_batches: 2 # Batches of the validation set to use
 each trigger

 model:
 hidden_dropout: 0.0 # Dropout for hidden states, must be
 between 0.0 and 1.0
 attention_dropout: 0.0 # Dropout for attention weights, must
 be between 0.0 and 1.0

 optim:

```

```

optimizer: adam
lr: 1e-5 # Learning rate
name: distributed_fused_adam # Optimizer algorithm, do not change
adam_w_mode: true # Enable AdamW mode
eps: 1e-06 # Epsilon for numerical stability
weight_decay: 0.0 # L2 regularization strength, must be
between 0.0 and 1.0
adam_beta1: 0.9 # Beta1 for Adam optimizer
adam_beta2: 0.95 # Beta2 for Adam optimizer
sched:
 warmup_steps: 10 # Learning rate warmup steps
 constant_steps: 0 # Steps at constant learning rate
 min_lr: 1e-6 # Minimum learning rate, must be lower
than lr

```

## Data preparation for CPT on 2.0

### Data format requirements

Training and validation datasets must be JSONL files following the format shown below, where each line contains a JSON object representing a conversation with the required fields and structure. Here is an example:

```

{"text": "AWS stands for Amazon Web Services"}
{"text": "Amazon SageMaker is a fully managed machine learning service"}
{"text": "Amazon Bedrock is a fully managed service for foundation models"}

```

Text entries should contain naturally flowing, high-quality content that represents the target domain.

Test that the data is capable of being converted into [Arrow format](#). Use the python script below to help with it. Ensure the `datasets==2.18.0` version at minimum is used:

```

from datasets import load_dataset, load_from_disk
from pathlib import Path

input_path = Path("<Your jsonl file>")
output_path = Path("<Your output directory>")

dataset = load_dataset("json", data_files=str(input_path), split="train")
dataset.save_to_disk(str(output_path), max_shard_size="1GB")

```

```
try:
 test_dataset = datasets.load_from_disk(output_dir)
 print(f"Dataset loaded successfully #! Contains {len(test_dataset)} samples")
except Exception as e:
 print(e)
```

It should print the same number of lines that were in the JSONL file.

When using datamixing, run the first job with `max_steps=2`. This will help create optimizations in the cluster for data access and validate that all the datamixes are available.

## How to prepare data for CPT

Training data is the most crucial determining factor for the success of continuous pre-training. While CPT data is often described as "unlabeled," the reality is far more nuanced. How data is structured, formatted, and presented determines whether the model will acquire the knowledge and skills required for the business use case.

### Preparing structured business datasets for CPT

This is a common challenge for companies and organizations building foundation models specialized in their domain. Most businesses possess rich repositories of structured data: product catalogs, user profiles, transaction logs, form submissions, API calls, and operational metadata. At first glance, this looks very different from the unstructured web text typically used in standard pre-training.

To effectively learn from structured business data, think carefully about downstream tasks and design the data presentation to force the model to learn the right predictive relationships.

To unlock the full potential of continuous pre-training, consider:

- What tasks the model should perform at inference time
- What information is present in the raw data
- How to structure that data so the model learns to extract and manipulate the information correctly

Simply dumping structured data into training won't teach the model to reason about it. Actively shape the data presentation to guide what the model learns.

In the following sections, there is literature review demonstrating the importance of data augmentation and provide examples augmentation strategies for structured business data that will give useful ideas on how to treat and organize business dataset for CPT.

## Structured data for CPT in the literature

CPT can pack domain facts into the model but often fails to make those facts retrievable and manipulable when inputs or tasks shift. Controlled experiments show that without diverse augmentation during pretraining, models memorize facts in brittle ways that remain hard to extract even after later instruction tuning, and they recommend injecting instruction like signals early in training. For semi structured data, randomized serialization and other augmentations reduce schema overfitting, which is why CPT should be interleaved with instruction style tasks rather than run first and IFT later. Finance focused work further finds that jointly mixing CPT and instruction data at batch time improves generalization and reduces forgetting versus the sequential recipe. Qwen technical report converges on the same pattern by integrating high quality instruction data into pretraining itself, which boosts in context learning and preserves instruction following while acquiring new domain knowledge.

Data augmentation for semi structured corpora is a key lever. Synthetic graph aware CPT expands small domain sets into entity linked corpora that explicitly teach relationships and compounds with retrieval at inference time. Joint CPT plus instruction mixing outperforms sequential pipelines in finance and balancing domain with general data lowers degradation on general skills. Very large scale domain CPT can also retain broad ability and even allow trade offs through model merging, yet still points to instruction tuning as an essential next step, reinforcing the value of introducing instruction signals during CPT.

## Injecting diversity through randomization and shuffling

A general strategy that helps to teach model effectively from the structured and semi structured datasets is to shuffle the order of fields in the datasets, and even randomly drop out some keys.

Shuffling the fields forces the model to read what each value means instead of where it appears and learn the relationships between all the fields. For example, in case of an video game posted on amazon store, when "Title," "Platform," "Price," "Condition," and "Edition" arrive in different permutations, the model can't rely on "the third slot is platform"; it must bind labels to values and learn the bilateral relationships among attributes: title  $\rightleftharpoons$  platform, platform  $\rightleftharpoons$  price, condition  $\rightleftharpoons$  price. So it can, for example, infer a likely platform from a game name and an observed price, or estimate a plausible price range given a title and platform.

Randomly dropping keys during serialization acts like feature dropout: it prevents co-adaptation on any one field and forces the model to recover missing information from the remaining evidence. If "Platform" is absent, the model must pick it up from the title string or compatibility text; if "Price" is hidden, it has to triangulate from platform, edition, and condition. This builds symmetry ( $A \rightarrow B$  and  $B \rightarrow A$ ), robustness to messy real-world listings, and schema invariance when fields are missing, renamed, or reordered.

An shopping-style example makes it concrete. Serialize the same item multiple ways—"Title: 'Elden Ring' | Platform: PlayStation 5 | Condition: Used—Like New | Price: \$34.99" and a permutation like "Price: \$34.99 | Title: 'Elden Ring' | Condition: Used—Like New | Platform: PlayStation 5"—and on some passes drop "Platform" while leaving "Compatible with PS5" in the description. Train complementary objectives such as predicting platform from {title, price} and predicting a price bucket from {title, platform}. Because order and even presence of keys vary, the only stable strategy is to learn the true relationships between attributes rather than memorize a template.

## **The way data is presented matters**

LLMs learn by predicting the next token from what they have already seen. So the order of fields and events shown during training decides what the model can learn. If the training format matches the real task, the loss lands on the exact decision tokens. If fields are tossed together without structure, the model learns shortcuts or memorizes popularity and then fails when asked to choose among options.

Show the situation first, then the options, then the decision. If the model should also learn about outcomes or explanations, put them after the decision.

## **Packing samples for CPT**

### **What is packing?**

It simply means to fill each sequence window in the training data with multiple whole examples so the window is dense with real tokens, not padding.

### **Why it matters**

During training a maximum context length is set, for example 8,192 tokens. Batches are shaped to [batch size × context length]. If a training example is shorter than the context length, the remaining positions are padded. Padding still runs through attention and MLP kernels even if loss is masked, so compute is paid for tokens that carry no learning signal.

## How to do packing?

To pack multiple samples, concatenate multiple training samples with a [DOC] separator in between (note the space before and after the [DOC] ) such that the full length of the samples are under the desired context length.

An example packed document would look like this:

```
{"text": "training sample 1 [DOC] training sample 2 [DOC] training sample 3"}
```

## CPT Tuning Parameters

The parameters that are available for fine-tuning with CPT include:

### Run Configuration

- **name:** A descriptive name for your training job. This helps identify your job in the AWS Management Console.
- **model\_type:** The Amazon Nova model variant to use. The available options are `amazon.nova-2-lite-v1:0:256k`.
- **model\_name\_or\_path:** The path to the base model to use for your training. The available options are `nova-lite-2/prod`, or the S3 path for the post-training checkpoint (`s3://customer-escrow-bucket-unique_id/training_run_name`).
- **replicas:** The number of compute instances to use for distributed training. Available values vary based on the model you choose. Amazon Nova Lite 2.0 supports 4, 8, 16, or 32 replicas.
- **data\_s3\_path:** The S3 location of the training dataset, which is a JSONL file. This file must reside in the same AWS account and Region as the cluster. All of the S3 locations provided must be in the same account and Region.
- **validation\_data\_s3\_path:** (Optional) The S3 location of the validation dataset, which is a JSONL file. This file must reside in the same account and region as the cluster. All of the S3 locations provided must be in the same account and Region.
- **output\_s3\_path:** The S3 location where the manifest and TensorBoard logs are stored. All of the S3 locations provided must be in the same AWS account and AWS Region.
- **mlflow\_tracking\_uri:** The ARN of the MLFlow App to use for MLFlow logging
- **mlflow\_experiment\_name:** MLFlow experiment name
- **mlflow\_run\_name:** MLFlow run name

## Training Configuration

- **max\_length**: The maximum sequence length in tokens. This determines the context window size for training. The maximum supported value are 8192 tokens for CPT.

Longer sequences will improve training efficiencies at the cost of increased memory requirements. We recommend that you match the `max_length` parameter to your data distribution.

- **global\_batch\_size**: The total number of training samples processed together in one forward or backward pass across all devices and workers.

This value multiplies the per-device batch size and number of devices. It affects the stability of training and throughput. We recommend that you start with a batch size that fits comfortably within your memory and scale up from there. For domain-specific data, larger batches might over-smooth gradients.

## Trainer Settings

- **max\_steps**: The number of training steps to run. Each step will train the model with `global_batch_size` no. of elements

## Model Settings

- **hidden\_dropout**: The probability of dropping hidden state outputs. Increase this value by approximately 0.0-0.2 to reduce overfitting on smaller datasets. Valid values are between 0-1, inclusive.
- **attention\_dropout**: The probability of dropping attention weights. This parameter can help with generalization. Valid values are between 0-1, inclusive.

## Optimizer Configuration

- **lr**: The learning rate, which controls the step size during optimization. We recommend values between 1e-6-1e-4 for good performance. Valid values are between 0-1, inclusive.
- **name**: The optimizer algorithm. Currently, only `distributed_fused_adam` is supported.



- **weight\_decay**: The L2 regularization strength. Higher values (between 0.01-0.1) increase regularization.
- **warmup\_steps**: The number of steps to gradually increase learning rate. This improves training stability. Valid values are between 1-20, inclusive.
- **min\_lr**: The minimum learning rate at the end of decay. Valid values are between 0-1, inclusive, but must be less than learning rate.

## Supervised fine-tuning (SFT)

The SFT training process consists of two main stages:

- **Data Preparation**: Follow established guidelines to create, clean, or reformat datasets into the required structure. Ensure that inputs, outputs, and auxiliary information (such as reasoning traces or metadata) are properly aligned and formatted.
- **Training Configuration**: Define how the model will be trained. When using `train`, this configuration is written in a YAML recipe file that includes:
  - Data source paths (training and validation datasets)
  - Key hyperparameters (epochs, learning rate, batch size)
  - Optional components (distributed training parameters, etc)

## Nova Model Comparison and Selection

Amazon Nova 2.0 is a model trained on a larger and more diverse dataset than Amazon Nova 1.0. Key improvements include:

- **Enhanced reasoning abilities** with explicit reasoning mode support
- **Broader multilingual performance** across additional languages
- **Improved performance on complex tasks** including coding and tool use
- **Extended context handling** with better accuracy and stability at longer context lengths

## When to Use Nova 1.0 vs. Nova 2.0

Choose Amazon Nova 2.0 when:

- Superior performance with advanced reasoning capabilities is needed
- Multilingual support or complex task handling is required

- Better results on coding, tool calling, or analytical tasks are needed

## SFT on Nova 2.0

Amazon Nova Lite 2.0 brings enhanced capabilities for supervised fine-tuning, including advanced reasoning mode, improved multimodal understanding, and extended context handling. SFT on Nova 2.0 enables you to adapt these powerful capabilities to your specific use cases while maintaining the model's superior performance on complex tasks.

Key features of SFT on Nova 2.0 include:

- **Reasoning mode support:** Train models to generate explicit reasoning traces before final answers for enhanced analytical capabilities.
- **Advanced multimodal training:** Fine-tune on document understanding (PDF), video understanding, and image-based tasks with improved accuracy.
- **Tool calling capabilities:** Train models to effectively use external tools and function calling for complex workflows.
- **Extended context support:** Leverage longer context windows with better stability and accuracy for document-intensive applications.

### Note

For more information on which container images, or example recipes to use go to [Amazon Nova recipes](#).

## Topics

- [Reasoning Mode Selection \(Nova 2.0 Only\)](#)
- [Tool calling data format](#)
- [Document understanding data format](#)
- [Video Understanding for SFT](#)
- [Data Upload Instructions](#)
- [Creating a Fine-Tuning Job](#)
- [SFT Tuning Parameters](#)
- [Hyperparameter Guidance](#)

## Sample SFT recipe

Below is a sample recipe for SFT. You can find this recipe and others in the [recipes](#) repository.

```
run:
 name: my-full-rank-sft-run
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: nova-lite-2/prod
 data_s3_path: s3://my-bucket-name/train.jsonl # only and not compatible with
SageMaker Training Jobs
 replicas: 4 # Number of compute instances for
training, allowed values are 4, 8, 16, 32
 output_s3_path: s3://my-bucket-name/outputs/ # Output artifact path (HyperPod job-
specific; not compatible with standard SageMaker Training Jobs)
 mlflow_tracking_uri: "" # Required for MLFlow
 mlflow_experiment_name: "my-full-rank-sft-experiment" # Optional for MLFlow. Note:
leave this field non-empty
 mlflow_run_name: "my-full-rank-sft-run" # Optional for MLFlow. Note: leave
this field non-empty

training_config:
 max_steps: 100 # Maximum training steps. Minimal is 4.
 save_steps: ${oc.select:training_config.max_steps} # How many training steps the
checkpoint will be saved
 save_top_k: 5 # Keep top K best checkpoints. Note supported only
for jobs. Minimal is 1.
 max_length: 32768 # Sequence length (options: 8192, 16384, 32768
[default], 65536)
 global_batch_size: 32 # Global batch size (options: 32, 64, 128)
 reasoning_enabled: true # If data has reasoningContent, set to true;
otherwise False

 lr_scheduler:
 warmup_steps: 15 # Learning rate warmup steps. Recommend 15% of
max_steps
 min_lr: 1e-6 # Minimum learning rate, must be between 0.0 and
1.0

 optim_config: # Optimizer settings
 lr: 1e-5 # Learning rate, must be between 0.0 and 1.0
 weight_decay: 0.0 # L2 regularization strength, must be between 0.0
and 1.0
 adam_beta1: 0.9 # Exponential decay rate for first-moment
estimates
```

```
adam_beta2: 0.95 # Exponential decay rate for second-moment
estimates

peft: # Parameter-efficient fine-tuning (LoRA)
 peft_scheme: "null" # Disable LoRA for PEFT
```

## Reasoning Mode Selection (Nova 2.0 Only)

Amazon Nova 2.0 supports reasoning mode for enhanced analytical capabilities:

- **Reasoning Mode (enabled):**
  - Set `reasoning_enabled: true` in the training configuration
  - Model trains to generate reasoning traces before final answers
  - Improves performance on complex reasoning tasks
- **Non-Reasoning Mode (disabled):**
  - Set `reasoning_enabled: false` or omit the parameter (default)
  - Standard SFT without explicit reasoning
  - Suitable for tasks that don't benefit from step-by-step reasoning

### Note

- When reasoning is enabled, it operates at high reasoning effort. There is no low reasoning option for SFT.
- Multimodal reasoning content is not supported for SFT. Reasoning mode applies to text-only inputs.

## Using reasoning mode with non-reasoning datasets

Training Amazon Nova on a non-reasoning dataset with `reasoning_enabled: true` is permitted. However, doing so may cause the model to lose its reasoning capabilities, as Amazon Nova primarily learns to generate the responses presented in the data without applying reasoning.

If training Amazon Nova on a non-reasoning dataset but still want to use reasoning during inference:

1. Disable reasoning during training (`reasoning_enabled: false`)

## 2. Enable reasoning later during inference

While this approach allows reasoning at inference time, it does not guarantee improved performance compared to inference without reasoning.

**Best practice:** Enable reasoning for both training and inference when using reasoning datasets, and disable it for both when using non-reasoning datasets.

### Note

For more information on which container images, or example recipes to use go to [Amazon Nova recipes](#).

## Tool calling data format

SFT supports training models to use tools (function calling). Below is a sample input format for tool calling:

### Sample input:

```
{
 "schemaVersion": "bedrock-conversation-2024",
 "system": [
 {
 "text": "You are an expert in composing function calls."
 }
],
 "toolConfig": {
 "tools": [
 {
 "toolSpec": {
 "name": "getItemCost",
 "description": "Retrieve the cost of an item from the catalog",
 "inputSchema": {
 "json": {
 "type": "object",
 "properties": {
 "item_name": {
 "type": "string",
 "description": "The name of the item to retrieve cost for"
 }
 }
 }
 }
 }
 }
]
 }
}
```

```

 "item_id": {
 "type": "string",
 "description": "The ASIN of item to retrieve cost for"
 }
 },
 "required": [
 "item_id"
]
}
},
{
 "toolSpec": {
 "name": "getItemAvailability",
 "description": "Retrieve whether an item is available in a given location",
 "inputSchema": {
 "json": {
 "type": "object",
 "properties": {
 "zipcode": {
 "type": "string",
 "description": "The zipcode of the location to check in"
 },
 "quantity": {
 "type": "integer",
 "description": "The number of items to check availability for"
 },
 "item_id": {
 "type": "string",
 "description": "The ASIN of item to check availability for"
 }
 }
 },
 "required": [
 "item_id", "zipcode"
]
 }
 }
}
],
},
"messages": [
 {

```

```

 "role": "user",
 "content": [
 {
 "text": "I need to check whether there are twenty pieces of the following
item available. Here is the item ASIN on Amazon: id-123. Please check for the zipcode
94086"
 }
]
 },
 {
 "role": "assistant",
 "content": [
 {
 "reasoningContent": {
 "reasoningText": {
 "text": "The user wants to check how many pieces of the item with ASIN
id-123 are available in the zipcode 94086"
 }
 }
 },
 {
 "toolUse": {
 "toolUseId": "getItemAvailability_0",
 "name": "getItemAvailability",
 "input": {
 "zipcode": "94086",
 "quantity": 20,
 "item_id": "id-123"
 }
 }
 }
]
 },
 {
 "role": "user",
 "content": [
 {
 "toolResult": {
 "toolUseId": "getItemAvailability_0",
 "content": [
 {
 "text": "[{"name": \"getItemAvailability\", \"results\":
{\"availability\": true}]]\"
 }
]
 }
 }
]
 }

```

```

]
 }
}
],
{
 "role": "assistant",
 "content": [
 {
 "text": "Yes, there are twenty pieces of item id-123 available at 94086.
Would you like to place an order or know the total cost?"
 }
]
}
]
}
}

```

Important considerations for tool calling data:

- ToolUse must appear in assistant turns only
- ToolResult must appear in user turns only
- ToolResult should be text or JSON only; other modalities are not currently supported for Amazon Nova models
- The inputSchema within the toolSpec must be a valid JSON Schema object
- Each ToolResult must reference a valid toolUseId from a preceding assistant ToolUse, with each toolUseId used exactly once per conversation

#### Note

For more information on which container images, or example recipes to use go to [Amazon Nova recipes](#).

## Document understanding data format

SFT supports training models on document understanding tasks. Below is a sample input format:

### Sample input

```
{
```



```
"schemaVersion": "bedrock-conversation-2024",
"messages": [
 {
 "role": "user",
 "content": [
 {
 "text": "What are the ways in which a customer can experience issues during
checkout on Amazon?"
 },
 {
 "document": {
 "format": "pdf",
 "source": {
 "s3Location": {
 "uri": "s3://my-bucket-name/path/to/documents/
customer_service_debugging.pdf",
 "bucketOwner": "123456789012"
 }
 }
 }
 }
]
 },
 {
 "role": "assistant",
 "content": [
 {
 "text": "Customers can experience issues with 1. Data entry, 2. Payment
methods, 3. Connectivity while placing the order. Which one would you like to dive
into?"
 }
],
 "reasoning_content": [
 {
 "text": "I need to find the relevant section in the document to answer the
question.",
 "type": "text"
 }
]
 }
]
```

## Important considerations for document understanding:

- Only PDF files are supported
- Maximum document size is 10 MB
- A sample can contain documents and text, but cannot mix documents with other modalities (such as images or video)

### Note

For more information on which container images, or example recipes to use go to [Amazon Nova recipes](#).

## Video Understanding for SFT

SFT supports fine-tuning models for video understanding tasks. Below is a sample input format:

### Sample input

```
{
 "schemaVersion": "bedrock-conversation-2024",
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "text": "What are the ways in which a customer can experience issues during checkout on Amazon?"
 },
 {
 "video": {
 "format": "mp4",
 "source": {
 "s3Location": {
 "uri": "s3://my-bucket-name/path/to/videos/customer_service_debugging.mp4",
 "bucketOwner": "123456789012"
 }
 }
 }
 }
]
 }
]
}
```

```
]
 },
 {
 "role": "assistant",
 "content": [
 {
 "text": "Customers can experience issues with 1. Data entry, 2. Payment methods, 3. Connectivity while placing the order. Which one would you like to dive into?"
 }
],
 "reasoning_content": [
 {
 "text": "I need to find the relevant section in the video to answer the question.",
 "type": "text"
 }
]
 }
]
```

### Important considerations for video understanding:

- Videos can be a maximum of 50 MB
- Videos can be up to 15 minutes long
- Only one video is allowed per sample; multiple videos in the same sample are not supported
- A sample can contain video and text, but cannot mix video with other modalities (such as images or documents)

#### Note

For more information on which container images, or example recipes to use go to [Amazon Nova recipes](#).

## Data Upload Instructions

Upload training and validation datasets to an S3 bucket. Specify these locations in the recipe's run block:

```
Run config
run:
 ...
 data_s3_path: "s3://<bucket-name>/<training-directory>/<training-file>.jsonl"
```

**Note:** Replace <bucket-name>, <training-directory>, <validation-directory>, <training-file>, and <validation-file> with actual S3 paths.

**Note:** Validation datasets are not currently supported for SFT with Amazon Nova 2.0. If a validation dataset is provided, it will be ignored.

## Creating a Fine-Tuning Job

Define the base model using the `model_type` and `model_name_or_path` fields in the run block:

```
Run config
run:
 ...
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: nova-lite-2/prod
 ...
```

## SFT Tuning Parameters

The parameters that are available for tuning with SFT include:

### Run Configuration

- **name:** A descriptive name for your training job. This helps identify your job in the AWS Management Console.
- **model\_type:** The Amazon Nova model variant to use. The available options are `amazon.nova-2-lite-v1:0:256k`.
- **model\_name\_or\_path:** The path to the base model to use for your training. The available options are `nova-lite-2/prod`, or the S3 path for the post-training checkpoint (`s3://customer-escrow-bucket-unique_id/training_run_name`).
- **replicas:** The number of compute instances to use for distributed training. Available values vary based on the model you choose. Amazon Nova Lite 2.0 supports 4, 8, 16, or 32 replicas.

- **data\_s3\_path:** The S3 location of the training dataset, which is a JSONL file. This file must reside in the same AWS account and Region as the cluster. All of the S3 locations provided must be in the same account and Region.
- **validation\_data\_s3\_path:** (Optional) The S3 location of the validation dataset, which is a JSONL file. This file must reside in the same account and region as the cluster. All of the S3 locations provided must be in the same account and Region.
- **output\_s3\_path:** The S3 location where the manifest and TensorBoard logs are stored. All of the S3 locations provided must be in the same AWS account and AWS Region.
- **mlflow\_tracking\_uri:** The ARN of the MLFlow App to use for MLFlow logging.
- **mlflow\_experiment\_name:** MLFlow experiment name.
- **mlflow\_run\_name:** MLFlow run name.

## Training Configuration

- **max\_steps:** The number of training steps to run. Each step will train the model with `global_batch_size` number of elements.
- **save\_steps:** The frequency (in steps) at which to save model checkpoints during training.
- **save\_top\_k:** The maximum number of best checkpoints to retain based on validation metrics.
- **max\_length:** The maximum sequence length in tokens. This determines the context window size for training. The maximum supported value is 32768 tokens for SFT.

Longer sequences will improve training efficiencies at the cost of increased memory requirements. We recommend that you match the `max_length` parameter to your data distribution.

- **global\_batch\_size:** The total number of training samples processed together in one forward or backward pass across all devices and workers.

This value multiplies the per-device batch size and number of devices. It affects the stability of training and throughput. We recommend that you start with a batch size that fits comfortably within your memory and scale up from there. For domain-specific data, larger batches might over-smooth gradients.

- **reasoning\_enabled:** Boolean flag to enable reasoning capabilities during training.

## Learning Rate Scheduler

- **warmup\_steps:** The number of steps to gradually increase learning rate. This improves training stability.
- **min\_lr:** The minimum learning rate at the end of decay. Valid values are between 0-1, inclusive, but must be less than learning rate.

## Optimizer Configuration

- **lr:** The learning rate, which controls the step size during optimization. We recommend values between 1e-6-1e-4 for good performance. Valid values are between 0-1, inclusive.
- **weight\_decay:** The L2 regularization strength. Higher values (between 0.01-0.1) increase regularization.
- **adam\_beta1:** The exponential decay rate for the first moment estimates in Adam optimizer. Default is 0.9.
- **adam\_beta2:** The exponential decay rate for the second moment estimates in Adam optimizer. Default is 0.95.

## PEFT Configuration

- **peft\_scheme:** The parameter-efficient fine-tuning scheme to use. Options are 'null' for full-rank fine-tuning or `lora` for LoRA-based fine-tuning.

### LoRA Tuning (when `peft_scheme` is 'lora')

- **alpha:** The LoRA scaling parameter. Controls the magnitude of the low-rank adaptation. Typical values range from 8 to 128.
- **lora\_plus\_lr\_ratio:** The learning rate ratio for LoRA+ optimization. This multiplier adjusts the learning rate specifically for LoRA parameters.

## Hyperparameter Guidance

Use the following recommended hyperparameters based on the training approach:

### Full Rank Training

- **Epochs:** 1
- **Learning rate (lr):** 1e-5
- **Minimum learning rate (min\_lr):** 1e-6

### **LoRA (Low-Rank Adaptation)**

- **Epochs:** 2
- **Learning rate (lr):** 5e-5
- **Minimum learning rate (min\_lr):** 1e-6

**Note:** Adjust these values based on dataset size and validation performance. Monitor training metrics to prevent overfitting.

## **Reinforcement Fine-Tuning (RFT) on SageMaker HyperPod**

Reinforcement Fine-Tuning (RFT) is a machine learning technique that improves model performance through feedback signals—measurable scores or rewards indicating response quality—rather than direct supervision with exact correct answers. Unlike traditional supervised fine-tuning that learns from input-output pairs, RFT uses reward functions to evaluate model responses and iteratively optimizes the model to maximize these rewards.

This approach is particularly effective for tasks where defining the exact correct output is challenging, but you can reliably measure response quality. RFT enables models to learn complex behaviors and preferences through trial and feedback, making it ideal for applications requiring nuanced decision-making, creative problem-solving, or adherence to specific quality criteria that can be programmatically evaluated.

### **When to use RFT**

Use RFT when you can define clear, measurable success criteria but struggle to provide exact correct outputs for training. It's ideal for tasks where quality is subjective or multifaceted—such as creative writing, code optimization, or complex reasoning—where multiple valid solutions exist but some are clearly better than others.

RFT works best when you have the following:

- A reliable reward function that can evaluate model outputs programmatically
- Need to align model behavior with specific preferences or constraints

- Situations where traditional supervised fine-tuning falls short because collecting high-quality labeled examples is expensive or impractical

Consider RFT for applications requiring iterative improvement, personalization, or adherence to complex business rules that can be encoded as reward signals.

### What RFT is best suited for

RFT excels in domains where output quality can be objectively measured but optimal responses are difficult to define upfront:

- **Mathematical problem-solving:** Verifiable correctness with multiple solution paths
- **Code generation and optimization:** Testable execution results and performance metrics
- **Scientific reasoning tasks:** Logical consistency and factual accuracy
- **Structured data analysis:** Programmatically verifiable outputs
- **Multi-step reasoning:** Tasks requiring step-by-step logical progression
- **Tool usage and API calls:** Success measurable by execution results
- **Complex workflows:** Adherence to specific constraints and business rules

RFT works exceptionally well when you need to balance multiple competing objectives like accuracy, efficiency, and style.

### When to use reasoning mode for RFT training

Amazon Nova 2.0 supports reasoning mode during RFT training. The following modes are available:

- **none:** No reasoning (omit the reasoning\_effort field)
- **low:** Minimal reasoning overhead
- **high:** Maximum reasoning capability (default when reasoning\_effort is specified)

#### Note

There is no medium option for RFT. If the reasoning\_effort field is absent from your configuration, reasoning is disabled.



Use high reasoning for the following:

- Complex analytical tasks
- Mathematical problem-solving
- Multi-step logical deduction
- Tasks where step-by-step thinking adds value

Use none (omit reasoning\_effort) or low reasoning for the following:

- Simple factual queries
- Direct classifications
- Speed and cost optimization
- Straightforward question-answering

#### **Important**

Higher reasoning modes increase training time and cost, inference latency and cost, but also increase model capability for complex reasoning tasks.

## Supported models

RFT on SageMaker HyperPod supports Amazon Nova Lite 2.0 (amazon.nova-2-lite-v1:0:256k).

## Major steps

The RFT process involves four key phases:

- **Implementing an evaluator:** Create a reward function to programmatically score model responses based on your quality criteria.
- **Uploading prompts:** Prepare and upload training data in the specified conversational format with reference data for evaluation.
- **Starting a job:** Launch the reinforcement fine-tuning process with your configured parameters.
- **Monitoring:** Track training progress through metrics dashboards to ensure the model learns effectively.

Each step builds on the previous one, with the evaluator serving as the foundation that guides the entire training process by providing consistent feedback signals.

## Topics

- [RFT on Nova 2.0](#)

## RFT on Nova 2.0

RFT training data follows the OpenAI conversational format. Each training example is a JSON object containing messages, reference answers, and optional tool definitions. This section provides guidance on preparing effective training data for RFT on Nova 2.0.

## Topics

- [Data format and structure](#)
- [Field descriptions](#)
- [Hyperparameter guidance](#)
- [Additional properties](#)
- [Dataset size recommendations](#)
- [Characteristics of effective training data](#)
- [Monitoring RFT training](#)

## Data format and structure

Each training example is a JSON object containing the following:

- **messages**: An array of conversational turns using system, user, and optionally assistant roles
- **reference\_answer**: Expected output or evaluation criteria for reward calculation
- **tools** (optional): Array of function definitions available to the model
- **id** (optional): Unique identifier for tracking and deduplication

Each example should be on a single line in your JSONL file, with one JSON object per line.

## Example 1: Chemistry problem

The following example shows a chemistry problem with reference answer containing ground truth values:

```
{
 "id": "chem-001",
 "messages": [
 {
 "role": "system",
 "content": "You are a helpful chemistry assistant"
 },
 {
 "role": "user",
 "content": "Predict hydrogen bond donors and acceptors for this SMILES:
CCN(CC)CCC(=O)c1sc(N)nc1C"
 }
],
 "reference_answer": {
 "donor_bond_counts": 2,
 "acceptor_bond_counts": 4,
 "explanation": "Calculated using Lipinski's rule of five: N-H groups (2 donors), N
and O atoms with lone pairs (4 acceptors)"
 }
}
```

### Note

The `reference_answer` contains ground truth values calculated using domain-specific rules. Your reward function compares the model's predicted values against these reference values to calculate a reward score.

## Example 2: Math problem

The following example shows a math problem with solution steps:

```
{
 "id": "math-001",
 "messages": [
 {
 "role": "system",
 "content": "You are a math tutor"
 },
 {
 "role": "user",
 "content": "Solve: $2x + 5 = 13$ "
 }
]
}
```

```

 }
],
 "reference_answer": {
 "solution": "x = 4",
 "steps": ["2x = 13 - 5", "2x = 8", "x = 4"]
 }
}

```

### Example 3: Tool usage

The following example shows tool usage with expected behavior:

```

{
 "id": "tool-001",
 "messages": [
 {
 "role": "system",
 "content": "You are a helpful game master assistant"
 },
 {
 "role": "user",
 "content": "Generate a strength stat for a warrior character. Apply a +2 racial bonus modifier."
 }
],
 "tools": [
 {
 "type": "function",
 "function": {
 "name": "StatRollAPI",
 "description": "Generates character stats by rolling 4d6, dropping the lowest die result, and applying a modifier.",
 "parameters": {
 "type": "object",
 "properties": {
 "modifier": {
 "description": "An integer representing the modifier to apply to the total of the stat roll.",
 "type": "integer"
 }
 }
 },
 "required": ["modifier"]
 }
 }
]
}

```

```

 }
],
 "reference_answer": {
 "tool_called": "StatRollAPI",
 "tool_parameters": {
 "modifier": 2
 },
 "expected_behavior": "Call StatRollAPI with modifier=2 and return the calculated
 stat value"
 }
}

```

## Field descriptions

| Field              | Description                                                      | Additional notes                                                                               | Required |
|--------------------|------------------------------------------------------------------|------------------------------------------------------------------------------------------------|----------|
| id                 | Unique identifier for this RFT example                           | String (for example, "sample-001"). Useful for tracking and deduplication.                     | No       |
| messages           | Ordered list of chat messages that define the prompt and context | Array of objects. Model sees them in order. Typically starts with a system message, then user. | Yes      |
| messages[].role    | Who is speaking in the message                                   | Common values: "system", "user" (sometimes "assistant" in other contexts)                      | No       |
| messages[].content | The text content of the message                                  | Plain string. For system it's instructions, for user it's the task or input.                   | No       |
| tools              | Tool specifications available to the                             | Array. Each item defines a tool's                                                              | No       |

| Field            | Description                                | Additional notes                                                               | Required |
|------------------|--------------------------------------------|--------------------------------------------------------------------------------|----------|
|                  | model during this example                  | interface and metadata. Types may include "function" or "internal".            |          |
| reference_answer | The expected model output for this example | String or object depending on task. Used as target for evaluation or training. | No       |

### Note

Any additional custom fields (for example, task\_id, difficulty\_level, context\_data) are not validated and will be passed to your reward function as metadata.

## Hyperparameter guidance

Use the following recommended hyperparameters based on your training approach:

### General:

- Epochs: 1
- Learning rate (lr): 1e-7
- Number of generations: 8
- Max new tokens: 8192
- Batch size: 256

### LoRA (Low-Rank Adaptation):

- LoRA Rank: 32

**Note**

Adjust these values based on your dataset size and validation performance. Monitor training metrics to prevent overfitting.

**Additional properties**

The "additionalProperties": true setting allows you to include custom fields beyond the core schema requirements, providing flexibility to add any data your reward function needs for proper evaluation.

**Common additional fields**

You can include the following types of additional fields:

**Metadata:**

- task\_id: Unique identifier for tracking
- difficulty\_level: Problem complexity indicator
- domain: Subject area or category
- expected\_reasoning\_steps: Number of steps in solution

**Evaluation criteria:**

- evaluation\_criteria: Specific grading rubrics
- custom\_scoring\_weights: Relative importance of different aspects
- context\_data: Background information for the problem
- external\_references: Links to relevant documentation or resources

**Example with additional properties**

The following example includes custom metadata fields:

```
{
 "id": "algebra_001",
 "messages": [
 {
 "role": "system",
```

```
 "content": "You are a math tutor"
 },
 {
 "role": "user",
 "content": "Solve: $2x + 5 = 13$ "
 }
],
"reference_answer": {
 "solution": "x = 4",
 "steps": [" $2x = 13 - 5$ ", " $2x = 8$ ", " $x = 4$ "]
},
"task_id": "algebra_001",
"difficulty_level": "easy",
"domain": "algebra",
"expected_reasoning_steps": 3
}
```

## Dataset size recommendations

### Starting point

Begin with the following minimum dataset sizes:

- Minimum 100 training examples
- Minimum 100 evaluation examples

Prioritize high-quality input data and a reliable reward function that executes consistently on model responses.

### Evaluation-first approach

Before investing in large-scale RFT training, evaluate your model's baseline performance:

- **High performance (greater than 95% reward):** RFT may be unnecessary—your model already performs well
- **Very poor performance (0% reward):** Switch to SFT first to establish basic capabilities
- **Moderate performance:** RFT is likely appropriate

This evaluation-first approach ensures your reward function is bug-free and determines if RFT is the right method for your use case. Starting small allows you to get comfortable with the



RFT workflow, identify and fix issues early, validate your approach before scaling up, and test reward function reliability. Once validated, you can expand to larger datasets to further improve performance.

## **Characteristics of effective training data**

### **Clarity and consistency**

Good RFT examples require clear, unambiguous input data that enables accurate reward calculation across different model outputs. Avoid noise in your data, including:

- Inconsistent formatting
- Contradictory labels or instructions
- Ambiguous prompts
- Conflicting reference answers

Any ambiguity will mislead the training process and cause the model to learn unintended behaviors.

### **Diversity**

Your dataset should capture the full diversity of production use cases to ensure robust real-world performance. Include:

- Various problem types and difficulty levels
- Different input formats and edge cases
- Representative samples from all expected scenarios

This diversity helps prevent overfitting and ensures the model handles unfamiliar inputs gracefully.

### **Reward function considerations**

Design your reward function for efficient training:

- Execute within seconds (not minutes)
- Parallelize effectively with Lambda
- Return consistent, reliable scores

- Handle different types of model outputs gracefully

Fast, scalable reward functions enable rapid iteration and cost-effective experimentation at scale.

## Monitoring RFT training

Monitor key metrics during training to ensure effective learning and identify potential issues early.

### Topics

- [Key metrics to track](#)
- [Evaluation after RFT](#)
- [Using fine-tuned models](#)
- [Limitations and best practices](#)
- [Troubleshooting](#)

### Key metrics to track

Monitor the following metrics using MLFlow during training:

#### Reward metrics:

- **Average reward score:** Overall quality of model responses (should increase over time)
- **Reward distribution:** Percentage of responses receiving high, medium, and low rewards
- **Training vs. validation rewards:** Compare to detect overfitting

#### Training metrics:

- **Policy updates:** Number of successful weight updates
- **Rollout completion rate:** Percentage of samples successfully evaluated

#### Concerning patterns:

- Rewards plateauing (indicates poor learning)
- Validation rewards dropping while training rewards increase (overfitting)
- Reward variance increasing significantly over time (instability)
- High percentage of reward function errors (implementation issues)

## When to stop training:

- Target performance metrics are achieved
- Rewards plateau and no longer improve
- Validation performance degrades (overfitting detected)
- Maximum training budget is reached

## Evaluation after RFT

After training completes, evaluate your fine-tuned model to assess performance improvements:

- **Run RFT evaluation job:** Use the checkpoint from your RFT training as the model
- **Compare to baseline:** Evaluate both base model and fine-tuned model on the same test set
- **Analyze metrics:** Review task-specific metrics (accuracy, reward scores, etc.)
- **Conduct qualitative review:** Manually inspect sample outputs for quality

For detailed evaluation procedures, see the Evaluation section.

## Using fine-tuned models

### Accessing checkpoints:

After training completes, locate your checkpoint:

1. Navigate to your `output_path` in S3
2. Download and extract `output.tar.gz`
3. Open `manifest.json`
4. Copy the `checkpoint_s3_bucket` value

### Deploying for inference:

Use the checkpoint S3 path for inference or further training:

```
run:
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: "s3://customer-escrow-<account-number>-smtj-<unique-
 identifier>/<job-name>"
```

For deployment and inference instructions, refer to the Inference section.

## Limitations and best practices

### Current limitations:

#### Beta restrictions:

- Need to create a new RIG group for RFT. This limitation will be resolved by GA.
- Instance type requirements: Only P5 instances supported (minimum 8x P5.48xlarge). Coming Soon: Support for smaller instance types (ETA: mid-January 2025).

### Functional limitations:

- 15-minute Lambda timeout: Reward functions must complete within 15 minutes
- Single-turn only: Multi-turn conversations not supported
- Validation datasets: Not supported during training. Use separate evaluation jobs to assess training progress.

### Training considerations:

- Low reward scenarios: May struggle when less than 5% of examples receive positive rewards - consider SFT first
- Data requirements: Needs sufficient diversity to learn effectively
- Computational cost: More expensive than supervised fine-tuning

### Nova Forge removes some of these limitations:

- Supports multi-turn conversations
- Allows reward functions exceeding 15-minute timeouts
- Provides advanced algorithms and tuning options
- Designed for complex enterprise use cases, specifically tuned to build frontier models

### Best practices:

#### Start small and scale:

- Begin with minimal datasets (100-200 examples) and few training epochs
- Validate your approach before scaling up
- Gradually increase dataset size and training steps based on results

**Baseline with SFT first:**

- If reward scores are consistently low (e.g., always 0), perform SFT before RFT
- RFT requires reasonable baseline performance to improve effectively

**Design efficient reward functions:**

- Execute in seconds, not minutes
- Minimize external API calls
- Use efficient algorithms and data structures
- Implement proper error handling
- Test thoroughly before training
- Leverage Lambda's parallel scaling capabilities

**Monitor training actively:**

- Track average reward scores over time
- Watch reward distribution across samples
- Compare training vs. validation rewards
- Look for concerning patterns (plateaus, overfitting, instability)

**Iterate based on results:**

- If rewards don't improve after several iterations, adjust reward function design
- Increase dataset diversity to provide clearer learning signals
- Consider switching to SFT if rewards remain near zero
- Experiment with different hyperparameters (learning rate, batch size)

**Optimize data quality:**

- Ensure diverse, representative examples
- Include edge cases and difficult samples
- Verify reward function correctly scores all example types
- Remove or fix samples that confuse the reward function

## Troubleshooting

### Reward function errors:

Symptoms: High error rate in reward function calls during training

| Issue                    | Symptoms                           | Resolution                                                                 |
|--------------------------|------------------------------------|----------------------------------------------------------------------------|
| Lambda timeout           | Frequent timeouts after 15 minutes | Optimize function performance; consider Nova Forge for complex evaluations |
| Insufficient concurrency | Lambda throttling errors           | Increase <code>lambda_concurrency_limit</code> or request quota increase   |
| Invalid return format    | Training fails with format errors  | Verify return structure matches required interface format                  |
| Unhandled exceptions     | Intermittent errors                | Add comprehensive error handling and logging                               |
| External API failures    | Inconsistent scoring               | Implement retry logic and fallback strategies                              |

### Poor training performance:

Symptoms: Rewards not improving or plateauing at low values

Resolutions:

- **Verify reward function correctness:** Test with known good/bad examples

- **Check baseline performance:** Evaluate base model; if near-zero accuracy, do SFT first
- **Increase data diversity:** Add more varied examples covering different scenarios
- **Adjust hyperparameters:** Try different learning rates or batch sizes
- **Review reward signal quality:** Ensure rewards differentiate between good and bad responses

### Overfitting:

Symptoms: Training rewards increase while validation rewards decrease

Resolutions:

- **Reduce training steps:** Stop training earlier
- **Increase dataset size:** Add more training examples
- **Add regularization:** Adjust `weight_decay` or `entropy_coeff`
- **Increase data diversity:** Ensure training set represents full distribution

## Evaluating your trained model

An evaluation recipe is a YAML configuration file that defines how your Amazon Nova model evaluation job is executed. With this recipe, you can assess the performance of a base or trained model against common benchmarks or your own custom datasets. Metrics can be stored in Amazon S3 or TensorBoard. The evaluation provides quantitative metrics that help you assess model performance across various tasks to determine if further customization is needed.

Model evaluation is an offline process, where models are tested against fixed benchmarks with predefined answers. They are not assessed in real-time or against live user interactions. For real-time evaluations, you can evaluate the model after it is deployed to Amazon Bedrock by calling the Amazon Bedrock runtime APIs.

### Important

The evaluation container only supports checkpoints produced by the same training platform. Checkpoints created with SageMaker HyperPod can only be evaluated using the SageMaker HyperPod evaluation workflow, and checkpoints created with SageMaker training jobs can only be evaluated using the SageMaker training jobs evaluation workflow. Attempting to evaluate a checkpoint from a different platform will result in failure.

## Topics

- [Available benchmark tasks](#)
- [Understanding the recipe parameters](#)
- [Evaluation recipe examples](#)
- [Starting an evaluation job](#)
- [Accessing and analyzing evaluation results](#)
- [RFT evaluation](#)

## Available benchmark tasks

A sample code package is available that demonstrates how to calculate benchmark metrics using the SageMaker AI model evaluation feature for Amazon Nova. To access the code packages, see [sample-Nova-lighteval-custom-task](#).

Here is a list of the supported, available industry standard benchmarks. You can specify the following benchmarks in the `eval_task` parameter:

| Benchmark | Modality | Description                                                                         | Metrics  | Strategy | Subtask Available |
|-----------|----------|-------------------------------------------------------------------------------------|----------|----------|-------------------|
| mmlu      | Text     | Multi-task Language Understanding – Tests knowledge across 57 subjects.             | accuracy | zs_cot   | Yes               |
| mmlu_pro  | Text     | MMLU – Professional Subset – Focuses on professional domains such as law, medicine, | accuracy | zs_cot   | No                |



| Benchmark | Modality | Description                                                                                                                  | Metrics  | Strategy | Subtask Available |
|-----------|----------|------------------------------------------------------------------------------------------------------------------------------|----------|----------|-------------------|
| bbh       | Text     | accounting, and engineering.                                                                                                 | accuracy | zs_cot   | Yes               |
|           |          | Advanced Reasoning Tasks – A collection of challenging problems that test higher-level cognitive and problem-solving skills. |          |          |                   |
| gpqa      | Text     | General Physics Question Answering – Assesses comprehension of physics concepts and related problem-solving abilities.       | accuracy | zs_cot   | No                |

| Benchmark     | Modality | Description                                                                                                                  | Metrics     | Strategy | Subtask Available |
|---------------|----------|------------------------------------------------------------------------------------------------------------------------------|-------------|----------|-------------------|
| math          | Text     | Mathematical Problem Solving – Measures mathematical reasoning across topics including algebra, calculus, and word problems. | exact_match | zs_cot   | Yes               |
| strong_reject | Text     | Quality-Control Task – Tests the model's ability to detect and reject inappropriate, harmful, or incorrect content.          | deflection  | zs       | Yes               |

| Benchmark | Modality | Description                                                                                                                                                     | Metrics  | Strategy | Subtask Available |
|-----------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|----------|-------------------|
| IFEval    | Text     | Instruction-Following Evaluation – Gauges how accurately a model follows given instructions and completes tasks to specification.                               | accuracy | zs       | No                |
| gen_qa    | Text     | Custom Dataset Evaluation – Lets you bring your own dataset for benchmarking, comparing model outputs to reference answers with metrics such as ROUGE and BLEU. | all      | gen_qa   | No                |

| Benchmark | Modality | Description                                                                                                                                                                                                           | Metrics | Strategy | Subtask Available |
|-----------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|----------|-------------------|
| llm_judge | Text     | LLM-as-a-Judge Preference Comparison – Uses a Amazon Nova Judge model to determine preference between paired responses (B compared with A) for your prompts, calculating the probability of B being preferred over A. | all     | judge    | No                |

| Benchmark    | Modality            | Description                                                                                                                     | Metrics | Strategy | Subtask Available |
|--------------|---------------------|---------------------------------------------------------------------------------------------------------------------------------|---------|----------|-------------------|
| humaneval    | Text                | HumanEval - A benchmark dataset designed to evaluate the code generation capabilities of large language models                  | pass@1  | zs       | No                |
| mm_llm_judge | Multi-modal (image) | This new benchmark behaves the same as the text-based llm_judge above. The only difference is that it supports image inference. | all     | judge    | No                |

| Benchmark        | Modality | Description                                                                                                                                                                                                                                                                                                                             | Metrics | Strategy | Subtask Available |
|------------------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|----------|-------------------|
| rubric_llm_judge | Text     | <p>Rubric Judge is an enhanced LLM-as-a-judge evaluation model built on Amazon Nova 2.0 Lite. Unlike the <a href="#">original judge model</a> that only provides preference verdicts, Rubric Judge dynamically generates custom evaluation criteria tailored to each prompt and assigns granular scores across multiple dimensions.</p> | all     | judge    | No                |

| Benchmark           | Modality | Description                                                                                                                       | Metrics     | Strategy | Subtask Available |
|---------------------|----------|-----------------------------------------------------------------------------------------------------------------------------------|-------------|----------|-------------------|
| aime_2024           | Text     | AIME 2024<br>- American Invitational Mathematics Examination problems testing advanced mathematical reasoning and problem-solving | exact_match | zs_cot   | No                |
| calendar_scheduling | Text     | Natural Plan<br>- Calendar Scheduling task testing planning abilities for scheduling meetings across multiple days and people     | exact_match | fs       | No                |

The following mmlu subtasks are available:

```
MMLU_SUBTASKS = [
 "abstract_algebra",
 "anatomy",
 "astronomy",
 "business_ethics",
 "clinical_knowledge",
 "college_biology",
 "college_chemistry",
```

```
"college_computer_science",
"college_mathematics",
"college_medicine",
"college_physics",
"computer_security",
"conceptual_physics",
"econometrics",
"electrical_engineering",
"elementary_mathematics",
"formal_logic",
"global_facts",
"high_school_biology",
"high_school_chemistry",
"high_school_computer_science",
"high_school_european_history",
"high_school_geography",
"high_school_government_and_politics",
"high_school_macroeconomics",
"high_school_mathematics",
"high_school_microeconomics",
"high_school_physics",
"high_school_psychology",
"high_school_statistics",
"high_school_us_history",
"high_school_world_history",
"human_aging",
"human_sexuality",
"international_law",
"jurisprudence",
"logical_fallacies",
"machine_learning",
"management",
"marketing",
"medical_genetics",
"miscellaneous",
"moral_disputes",
"moral_scenarios",
"nutrition",
"philosophy",
"prehistory",
"professional_accounting",
"professional_law",
"professional_medicine",
"professional_psychology",
```



```
"public_relations",
"security_studies",
"sociology",
"us_foreign_policy",
"virology",
"world_religions"
]
```

The following bbh subtasks are available:

```
BBH_SUBTASKS = [
 "boolean_expressions",
 "causal_judgement",
 "date_understanding",
 "disambiguation_qa",
 "dyck_languages",
 "formal_fallacies",
 "geometric_shapes",
 "hyperbaton",
 "logical_deduction_five_objects",
 "logical_deduction_seven_objects",
 "logical_deduction_three_objects",
 "movie_recommendation",
 "multistep_arithmetic_two",
 "navigate",
 "object_counting",
 "penguins_in_a_table",
 "reasoning_about_colored_objects",
 "ruin_names",
 "salient_translation_error_detection",
 "snarks",
 "sports_understanding",
 "temporal_sequences",
 "tracking_shuffled_objects_five_objects",
 "tracking_shuffled_objects_seven_objects",
 "tracking_shuffled_objects_three_objects",
 "web_of_lies",
 "word_sorting"
]
```

The following math subtasks are available:

```
MATH_SUBTASKS = [
```

```
"algebra",
"counting_and_probability",
"geometry",
"intermediate_algebra",
"number_theory",
"prealgebra",
"precalculus",
]
```

## Understanding the recipe parameters

### Run configuration

The following is a general run configuration and an explanation of the parameters involved.

```
run:
 name: eval_job_name
 model_type: amazon.nova-micro-v1:0:128k
 model_name_or_path: nova-micro/prod
 replicas: 1
 data_s3_path: ""
 output_s3_path: s3://output_path
 mlflow_tracking_uri: ""
 mlflow_experiment_name : ""
 mlflow_run_name : ""
```

- **name:** (Required) A descriptive name for your evaluation job. This helps identify your job in the AWS console.
- **model\_type:** (Required) Specifies the Amazon Nova model variant to use. Do not manually modify this field. Options include:
  - `amazon.nova-micro-v1:0:128k`
  - `amazon.nova-lite-v1:0:300k`
  - `amazon.nova-pro-v1:0:300k`
  - `amazon.nova-2-lite-v1:0:256k`
- **model\_name\_or\_path:** (Required) The path to the base model or S3 path for the post-trained checkpoint. Options include:
  - `nova-micro/prod`
  - `nova-lite/prod`

- `nova-pro/prod`
- `nova-lite-2/prod`
- (S3 path for the post-trained checkpoint) `s3://<escrow bucket>/<job id>/outputs/checkpoints`
- `replicas`: (Required) The number of compute instances to use for distributed training. You must set this value to 1 because multi-node is not supported.
- `data_s3_path`: (Required) The S3 path to the input dataset. Leave this parameter empty unless you are using the *bring your own dataset* or *LLM as a judge* recipe.
- `output_s3_path`: (Required) The S3 path to store output evaluation artifacts. Note that the output S3 bucket must be created by the same account that is creating the job.
- `mlflow_tracking_uri`: (Optional) MLflow tracking server ARN for tracking MLflow runs/experiments. Please ensure you have permission to access the tracking server from SageMaker AI execution role

## Evaluation configuration

The following is a model evaluation configuration and an explanation of the parameters involved.

```
evaluation:
 task: mmlu
 strategy: zs_cot
 subtask: mathematics
 metric: accuracy
```

- `task`: (Required) Specifies the evaluation benchmark or task to use.

Supported task list:

- `mmlu`
- `mmlu_pro`
- `bbh`
- `gpqa`
- `math`
- `strong_reject`
- `gen_qa`
- `ifeval`

- `llm_judge`
- `humaneval`
- `mm_llm_judge`
- `rubric_llm_judge`
- `aime_2024`
- `calendar_scheduling`
- `humaneval`
- `strategy`: (Required) Defines the evaluation approach:
  - `zs_cot`: Zero-shot Chain-of-Thought - An approach to prompt large language models that encourages step-by-step reasoning without requiring explicit examples.
  - `zs`: Zero-shot - An approach to solve a problem without any prior training examples.
  - `gen_qa`: A strategy specific for bring your own dataset recipes.
  - `judge`: A strategy specific for Amazon Nova LLM as Judge and `mm_llm_judge`.
- `subtask`: (Optional and Removable) Specifies a specific subtask for certain evaluation tasks. Remove this from your recipe if your task does not have any subtasks.
- `metric`: (Required) The evaluation metric to use.
  - `accuracy`: Percentage of correct answers
  - `exact_match`: (For math benchmark), returns the rate at which the input predicted strings exactly match their references.
  - `deflection`: (For strong reject benchmark), returns the relative deflection to the base model and the difference in significance metrics.
  - `pass@1`: (For humaneval benchmark) is a metric used to measures the percentage of cases where the model's highest confidence prediction matches the correct answer.
  - `all`: Returns the following metrics:
    - For `gen_qa` and bring your own dataset benchmark, return following metrics:
      - `rouge1`: Measures the overlap of unigrams (single words) between generated and reference text.
      - `rouge2`: Measures the overlap of bigrams (two consecutive words) between generated and reference text.
      - `rougeL`: Measures the longest common subsequence between texts, allowing for gaps in the matching.

- `exact_match`: Binary score (0 or 1) indicating if the generated text matches the reference text exactly, character by character.
- `quasi_exact_match`: Similar to exact match but more lenient, typically ignoring case, punctuation, and white space differences.
- `f1_score`: Harmonic mean of precision and recall, measuring word overlap between predicted and reference answers.
- `f1_score_quasi`: Similar to `f1_score` but with more lenient matching, using normalized text comparison that ignores minor differences.
- `bleu`: Measures precision of n-gram matches between generated and reference text, commonly used in translation evaluation.
- For `llm_judge` and `mm_llm_judge`, bring your own dataset benchmark, return following metrics:
  - `a_scores`: Number of wins for `response_A` across forward and backward evaluation passes.
  - `a_scores_stderr`: Standard error of `response_A` scores across pairwise judgements.
  - `b_scores`: Number of wins for `response_B` across forward and backward evaluation passes.
  - `b_scores_stderr`: Standard error of `response_B` scores across pairwise judgements.
  - `ties`: Number of judgements where `response_A` and `response_B` are evaluated as equal.
  - `ties_stderr`: Standard error of ties across pairwise judgements.
  - `inference_error`: Count of judgements that could not be properly evaluated.
  - `inference_error_stderr`: Standard error of inference errors across judgements.
  - `score`: Aggregate score based on wins from both forward and backward passes for `response_B`.
  - `score_stderr`: Standard error of the aggregate score across pairwise judgements.
  - `winrate`: the probability that `response_B` will be preferred over `response_A` calculated using Bradley-Terry probability.
  - `lower_rate`: Lower bound (2.5th percentile) of the estimated win rate from bootstrap sampling.

## Inference configuration

The following is an inference configuration and an explanation of the parameters involved. All parameters are optional.

```
inference:
 max_new_tokens: 200
 top_k: -1
 top_p: 1.0
 temperature: 0
 top_logprobs: 10
 reasoning_effort: null # options: low/high to enable reasoning or null to disable reasoning
```

- `max_new_tokens`: The maximum number of tokens to generate. This must be an integer.
- `top_k`: The number of highest probability tokens to consider. This must be an integer.
- `top_p`: The cumulative probability threshold for token sampling. This must be a float between 0.0 and 1.0, inclusive.
- `temperature`: Randomness in token selection. Larger values introduce more randomness. Use 0 to make the results deterministic. This value must be a float with a minimum value of 0.
- `top_logprobs`: The number of top logprobs to be returned in the inference response. This value must be an integer from 0 to 20. Logprobs contain the considered output tokens and log probabilities of each output token returned in the message content.
- `reasoning_effort`: controls the reasoning behavior for reasoning-capable models. Set `reasoning_effort` only when `model_type` specifies a reasoning-capable model (currently `amazon.nova-2-lite-v1:0:256k`). Available options are `null` (default value if not set; disables reasoning), `low`, or `high`.

Note that for `humaneval`, we recommend the following inference configuration:

```
inference:
 top_k: 1
 max_new_tokens: 1600
 temperature: 0.0
```

## MLFlow configuration

The following is an MLFlow configuration and an explanation of the parameters involved. All parameters are optional.

```
run:
 mlflow_tracking_uri: ""
 mlflow_experiment_name: ""
 mlflow_run_name: ""
```

- `mlflow_tracking_uri`: (Optional) The location of the MLflow tracking server (only needed on SMHP)
- `mlflow_experiment_name`: (Optional) Name of the experiment to group related ML runs together
- `mlflow_run_name`: (Optional) Custom name for a specific training run within an experiment

## Evaluation recipe examples

Amazon Nova provides four types of evaluation recipes, which are available in the SageMaker HyperPod recipes GitHub repository.

### General text benchmark recipes

These recipes enable you to evaluate the fundamental capabilities of Amazon Nova models across a comprehensive suite of text-only benchmarks. They are provided in the format `xxx_general_text_benchmark_eval.yaml`.

### Bring your own dataset benchmark recipes

These recipes enable you to bring your own dataset for benchmarking and compare model outputs to reference answers using different types of metrics. They are provided in the format `xxx_bring_your_own_dataset_eval.yaml`.

The following are the bring your own dataset requirements:

- File format requirements
  - You must include a single `gen_qa.jsonl` file containing evaluation examples.
  - Your dataset must be uploaded to an S3 location where SageMaker training job can access it.
  - The file must follow the required schema format for a general Q&A dataset.
- Schema format requirements - Each line in the JSONL file must be a JSON object with the following fields:
  - `query`: (Required) String containing the question or instruction that needs an answer

- **response:** (Required) String containing the expected model output
- **system:** (Optional) String containing the system prompt that sets the behavior, role, or personality of the AI model before it processes the query
- **metadata:** (Optional) String containing metadata associated with the entry for tagging purposes.

Here is a bring your own data set example entry

```
{
 "system": "You are a english major with top marks in class who likes to give minimal word responses: ",
 "query": "What is the symbol that ends the sentence as a question",
 "response": "?"
}
{
 "system": "You are a pattern analysis specialist that provides succinct answers: ",
 "query": "What is the next number in this series? 1, 2, 4, 8, 16, ?",
 "response": "32"
}
{
 "system": "You have great attention to detail that follows instructions accurately: ",
 "query": "Repeat only the last two words of the following: I ate a hamburger today and it was kind of dry",
 "response": "of dry"
}
```

To use your custom dataset, modify your evaluation recipe with the following required fields, do not change any of the content:

```
evaluation:
 task: gen_qa
 strategy: gen_qa
 metric: all
```

The following limitations apply:

- Only one JSONL file is allowed per evaluation.
- The file must strictly follow the defined schema.



- Context length limit: For each sample in the dataset, the context length (including system + query prompts) should be less than 3.5k.

## Nova LLM as a Judge benchmark recipes

Amazon Nova LLM as a Judge is a model evaluation feature that enables customers to compare the quality of responses from one model to a baseline model response on a custom dataset. It takes in a dataset with prompts, baseline responses, and challenger responses, and uses a Amazon Nova Judge model to provide a winrate metric based on [Bradley-Terry probability](#) with pairwise comparisons.

The recipes are provided in the format `xxx_llm_judge_eval.yaml`.

The following are the LLM as a Judge requirements:

- File format requirements
  - Include a single `llm_judge.jsonl` file containing evaluation examples. The file name must be `llm_judge.jsonl`.
  - Your dataset must be uploaded to an S3 location that [SageMaker AI SageMaker HyperPod RIG](#) can access.
  - The file must follow the required schema format for the `llm_judge.jsonl` dataset.
  - The input dataset should ensure all records are under 12k context length.
- Schema format requirements - Each line in the JSONL file must be a JSON object with the following fields:
  - `prompt`: (Required) A string containing the prompt for the generated response.
  - `response_A`: A string containing the baseline response.
  - `response_B`: A string containing the alternative response be compared with baseline response.

Here is an LLM as a judge example entry

```
{
 "prompt": "What is the most effective way to combat climate change?",
 "response_A": "The most effective way to combat climate change is through a combination of transitioning to renewable energy sources and implementing strict carbon pricing policies. This creates economic incentives for businesses to reduce emissions while promoting clean energy adoption.",
```

```
"response_B": "We should focus on renewable energy. Solar and wind power are good.
 People should drive electric cars. Companies need to pollute less."
}
{
 "prompt": "Explain how a computer's CPU works",
 "response_A": "CPU is like brain of computer. It does math and makes computer work
 fast. Has lots of tiny parts inside.",
 "response_B": "A CPU (Central Processing Unit) functions through a fetch-execute
 cycle, where instructions are retrieved from memory, decoded, and executed through its
 arithmetic logic unit (ALU). It coordinates with cache memory and registers to process
 data efficiently using binary operations."
}
{
 "prompt": "How does photosynthesis work?",
 "response_A": "Plants do photosynthesis to make food. They use sunlight and water. It
 happens in leaves.",
 "response_B": "Photosynthesis is a complex biochemical process where plants convert
 light energy into chemical energy. They utilize chlorophyll to absorb sunlight,
 combining CO2 and water to produce glucose and oxygen through a series of chemical
 reactions in chloroplasts."
}
```

To use your custom dataset, modify your evaluation recipe with the following required fields, don't change any of the content:

```
evaluation:
 task: llm_judge
 strategy: judge
 metric: all
```

The following limitations apply:

- Only one JSONL file is allowed per evaluation.
- The file must strictly follow the defined schema.
- Amazon Nova Judge models are the same across all model family specifications (that is, Lite, Micro, and Pro).
- Custom judge models are not supported at this time.
- Context length limit: For each sample in the dataset, the context length (including system + query prompts) should be less than 7k.

## Nova LLM as a Judge for multi-modal (image) benchmark recipes

Nova LLM Judge for multi-modal (image), short for Amazon Nova MM\_LLM Judge, is a model evaluation feature that enables you to compare the quality of responses from one model against a baseline model's responses using a custom dataset. It accepts a dataset containing prompts, baseline responses, and challenger responses, and images in the form of Base64-encoded string, then uses a Amazon Nova Judge model to provide a win rate metric based on [Bradley-Terry](#) probability through pairwise comparisons. Recipe format: `xxx_mm_llm_judge_eval.yaml`.

### Nova LLM dataset requirements

File format:

- Single `mm_llm_judge.jsonl` file containing evaluation examples. The file name must be exactly `llm_judge.jsonl`.
- You must upload your dataset to an S3 location where SageMaker Training Jobs can access it.
- The file must follow the required schema format for the `mm_llm_judge` dataset.
- The input dataset should ensure all records are under 12 k context length, excluding the image's attribute.

Schema format - Each line in the `.jsonl` file must be a JSON object with the following fields.

- Required fields.

`prompt`: String containing the prompt for the generated response.

`images`: Array containing a list of objects with data attributes (values are Base64-encoded image strings).

`response_A`: String containing the baseline response.

`response_B`: String containing the alternative response be compared with baseline response.

### Example entry

For readability, the following example includes new lines and indentation, but in the actual dataset, each record should be on a single line.

```
{
```

```
"prompt": "what is in the image?",
"images": [
 {
 "data": "data:image/jpeg;base64,/9j/2wBDAAQDAwQDAwQEAwQFBAQFBgo..."
 }
],
"response_A": "a dog.",
"response_B": "a cat.",
}
{
 "prompt": "how many animals in echo of the images?",
 "images": [
 {
 "data": "data:image/jpeg;base64,/9j/2wBDAAQDAwQDAwQEAwQFBAQFBgo..."
 },
 {
 "data": "data:image/jpeg;base64,/DKEafe3gihn..."
 }
],
 "response_A": "The first image contains one cat and the second image contains one dog",
 "response_B": "The first image has one animal and the second has one animal",
}
```

To use your custom dataset, modify your evaluation recipe with the following required fields, don't change any of the content:

```
evaluation:
 task: mm_llm_judge
 strategy: judge
 metric: all
```

## Limitations

- Only one .jsonl file is allowed per evaluation.
- The file must strictly follow the defined schema.
- Nova MM Judge models only support image reference.
- Nova MM Judge models are the same across Amazon Nova Lite specifications.
- Custom judge models are not currently supported.
- Amazon S3 image URI is not supported.

- The input dataset should ensure all records are under 12 k context length, excluding images attribute.

## Rubric Based Judge

Rubric Judge is an enhanced LLM-as-a-judge evaluation model built on Amazon Nova 2.0 Lite. Unlike the [original judge model](#) that only provides preference verdicts (A>B, B>A, or tie), Rubric Judge dynamically generates custom evaluation criteria tailored to each prompt and assigns granular scores across multiple dimensions.

Key capabilities:

- **Dynamic criteria generation:** Automatically creates relevant evaluation dimensions based on the input prompt
- **Weighted scoring:** Assigns importance weights to each criterion to reflect their relative significance
- **Granular assessment:** Provides detailed scores on a binary (true/false) or scale (1-5) basis for each criterion
- **Quality metrics:** Calculates continuous quality scores (0-1 scale) that quantify the magnitude of differences between responses

Example criterion generated by the model:

```
price_validation:
 description: "The response includes validation to ensure price is a positive value."
 type: "scale"
 weight: 0.3
```

The model evaluates both responses against all generated criteria, then uses these criterion-level scores to inform its final preference decision.

## Topics

- [Recipe configuration](#)
- [Input dataset format](#)
- [Evaluation output](#)
- [Reasoning model support](#)

## Recipe configuration

### Rubric Judge recipe

Enable Rubric Judge by setting task: `rubric_llm_judge` in your recipe:

```
run:
 name: nova-eval-job-name # [MODIFIABLE] Unique
 identifier for your evaluation job
 model_type: amazon.nova-2-lite-v1:0:256k # [FIXED] Rubric Judge model
 type
 model_name_or_path: "nova-lite-2/prod" # [FIXED] Path to model
 checkpoint or identifier
 replicas: 1 # [MODIFIABLE] Number of
 replicas for SageMaker Training job
 data_s3_path: "" # [FIXED] Leave empty for
 SageMaker Training job
 output_s3_path: "" # [FIXED] Leave empty for
 SageMaker Training job

evaluation:
 task: rubric_llm_judge # [FIXED] Evaluation task -
 enables Rubric Judge
 strategy: judge # [FIXED] Evaluation strategy
 metric: all # [FIXED] Metric calculation
 method

inference:
 max_new_tokens: 12000 # [MODIFIABLE] Maximum tokens
 to generate
 top_k: -1 # [MODIFIABLE] Top-k sampling
 parameter
 top_p: 1.0 # [MODIFIABLE] Nucleus sampling
 parameter
 temperature: 0 # [MODIFIABLE] Sampling
 temperature (0 = deterministic)
```

### Original LLM as a Judge recipe (for comparison)

The original judge model uses task: `llm_judge`:

```
run:
 name: eval-job-name # [MODIFIABLE] Unique
 identifier for your evaluation job
```

```

model_type: amazon.nova-micro-v1:0:128k # [FIXED] Model type
model_name_or_path: "nova-micro/prod" # [FIXED] Path to model
checkpoint or identifier
replicas: 1 # [MODIFIABLE] Number of
replicas for SageMaker Training job
data_s3_path: "" # [FIXED] Leave empty for
SageMaker Training job
output_s3_path: "" # [FIXED] Leave empty for
SageMaker Training job

evaluation:
 task: llm_judge # [FIXED] Original judge task
 strategy: judge # [FIXED] Evaluation strategy
 metric: all # [FIXED] Metric calculation
 method

inference:
 max_new_tokens: 12000 # [MODIFIABLE] Maximum tokens
 to generate
 top_k: -1 # [MODIFIABLE] Top-k sampling
 parameter
 top_p: 1.0 # [MODIFIABLE] Nucleus sampling
 parameter
 temperature: 0 # [MODIFIABLE] Sampling
 temperature (0 = deterministic)

```

## Input dataset format

The input dataset format is identical to the [original judge model](#):

### Required fields:

- **prompt**: String containing the input prompt and instructions
- **response\_A**: String containing the baseline model output
- **response\_B**: String containing the customized model output

### Example dataset (JSONL format):

```

{"prompt": "What is the most effective way to combat climate change?", "response_A":
"The most effective way to combat climate change is through a combination of
transitioning to renewable energy sources and implementing strict carbon pricing
policies. This creates economic incentives for businesses to reduce emissions while

```

```

promoting clean energy adoption.", "response_B": "We should focus on renewable energy.
Solar and wind power are good. People should drive electric cars. Companies need to
pollute less."}
{"prompt": "Explain how a computer's CPU works", "response_A": "CPU is like brain of
computer. It does math and makes computer work fast. Has lots of tiny parts inside.",
"response_B": "A CPU (Central Processing Unit) functions through a fetch-execute
cycle, where instructions are retrieved from memory, decoded, and executed through its
arithmetic logic unit (ALU). It coordinates with cache memory and registers to process
data efficiently using binary operations."}
{"prompt": "How does photosynthesis work?", "response_A": "Plants do photosynthesis
to make food. They use sunlight and water. It happens in leaves.", "response_B":
"Photosynthesis is a complex biochemical process where plants convert light energy
into chemical energy. They utilize chlorophyll to absorb sunlight, combining CO2
and water to produce glucose and oxygen through a series of chemical reactions in
chloroplasts."}

```

### Format requirements:

- Each entry must be a single-line JSON object
- Separate entries with newlines
- Follow the exact field naming as shown in examples

### Evaluation output

#### Output structure

Rubric Judge produces enhanced evaluation metrics compared to the original judge model:

```

{
 "config_general": {
 "lighteval_sha": "string",
 "num_fewshot_seeds": "int",
 "max_samples": "int | null",
 "job_id": "int",
 "start_time": "float",
 "end_time": "float",
 "total_evaluation_time_seconds": "string",
 "model_name": "string",
 "model_sha": "string",
 "model_dtype": "string | null",
 "model_size": "string"
 },

```



```

"results": {
 "custom|rubric_llm_judge_judge|0": {
 "a_scores": "float",
 "a_scores_stderr": "float",
 "b_scores": "float",
 "b_scores_stderr": "float",
 "ties": "float",
 "ties_stderr": "float",
 "inference_error": "float",
 "inference_error_stderr": "float",
 "score": "float",
 "score_stderr": "float",
 "weighted_score_A": "float",
 "weighted_score_A_stderr": "float",
 "weighted_score_B": "float",
 "weighted_score_B_stderr": "float",
 "score_margin": "float",
 "score_margin_stderr": "float",
 "winrate": "float",
 "lower_rate": "float",
 "upper_rate": "float"
 }
},
"versions": {
 "custom|rubric_llm_judge_judge|0": "int"
}
}

```

## New metrics in Rubric Judge

The following six metrics are unique to Rubric Judge and provide granular quality assessment:

| Metric                  | Description                                                                                                                                                                                       |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| weighted_score_A        | Average normalized quality score for response_A across all model-generated evaluation criteria. Scores are weighted by criterion importance and normalized to 0-1 scale (higher = better quality) |
| weighted_score_A_stderr | Standard error of the mean for weighted_score_A, indicating statistical uncertainty                                                                                                               |

| Metric                  | Description                                                                                                                                                                                               |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| weighted_score_B        | Average normalized quality score for response_B across all model-generated evaluation criteria. Scores are weighted by criterion importance and normalized to 0-1 scale (higher = better quality)         |
| weighted_score_B_stderr | Standard error of the mean for weighted_score_B, indicating statistical uncertainty                                                                                                                       |
| score_margin            | Difference between weighted scores (calculated as weighted_score_A - weighted_score_B). Range: -1.0 to 1.0. Positive = response_A is better; negative = response_B is better; near zero = similar quality |
| score_margin_stderr     | Standard error of the mean for score_margin, indicating uncertainty in the quality difference measurement                                                                                                 |

## Understanding weighted score metrics

**Purpose:** Weighted scores provide continuous quality measurements that complement binary preference verdicts, enabling deeper insights into model performance.

### Key differences from original judge:

- **Original Judge:** Only outputs discrete preferences (A>B, B>A, A=B)
- **Rubric Judge:** Outputs both preferences AND continuous quality scores (0-1 scale) based on custom criteria

### Interpreting score\_margin:

- `score_margin = -0.128`: Response\_B scored 12.8 percentage points higher than response\_A
- `|score_margin| < 0.1`: Narrow quality difference (close decision)
- `|score_margin| > 0.2`: Clear quality difference (confident decision)

## Use cases:

- **Model improvement:** Identify specific areas where your model underperforms
- **Quality quantification:** Measure the magnitude of performance gaps, not just win/loss ratios
- **Confidence assessment:** Distinguish between close decisions and clear quality differences

### Important

Final verdicts are still based on the judge model's explicit preference labels to preserve holistic reasoning and ensure proper position bias mitigation through forward/backward evaluation. Weighted scores serve as observability tools, not as replacements for the primary verdict.

## Calculation methodology

Weighted scores are computed through the following process:

- **Extract criterion data:** Parse the judge's YAML output to extract criterion scores and weights
- **Normalize scores:**
  - Scale-type criteria (1-5): Normalize to 0-1 by calculating  $(\text{score} - 1) / 4$
  - Binary criteria (true/false): Convert to 1.0/0.0
- **Apply weights:** Multiply each normalized score by its criterion weight
- **Aggregate:** Sum all weighted scores for each response
- **Calculate margin:** Compute  $\text{score\_margin} = \text{weighted\_score\_A} - \text{weighted\_score\_B}$

**Example:** If response\_A has a weighted sum of 0.65 and response\_B has 0.78, the `score_margin` would be -0.13, indicating response\_B is 13 percentage points higher in quality across all weighted criteria.

## Reasoning model support

Reasoning model support enables evaluation with reasoning-capable Amazon Nova models that perform explicit internal reasoning before generating final responses. This feature uses API-level control via the `reasoning_effort` parameter to dynamically enable or disable reasoning functionality, potentially improving response quality for complex analytical tasks.

## Supported models:

- amazon.nova-2-lite-v1:0:256k

## Recipe configuration

Enable reasoning by adding the `reasoning_effort` parameter to the inference section of your recipe:

```
run:
 name: eval-job-name # [MODIFIABLE] Unique
 identifier for your evaluation job
 model_type: amazon.nova-2-lite-v1:0:256k # [FIXED] Must be a reasoning-
supported model
 model_name_or_path: nova-lite-2/prod # [FIXED] Path to model
checkpoint or identifier
 replicas: 1 # [MODIFIABLE] Number of
replicas for SageMaker Training job
 data_s3_path: "" # [MODIFIABLE] Leave empty for
SageMaker Training job; optional for SageMaker SageMaker HyperPod job
 output_s3_path: "" # [MODIFIABLE] Output path for
SageMaker SageMaker HyperPod job (not compatible with SageMaker Training jobs)

evaluation:
 task: mmlu # [MODIFIABLE] Evaluation task
 strategy: generate # [MODIFIABLE] Evaluation
strategy
 metric: all # [MODIFIABLE] Metric
calculation method

inference:
 reasoning_effort: high # [MODIFIABLE] Enables
reasoning mode; options: low/medium/high or null to disable
 max_new_tokens: 200 # [MODIFIABLE] Maximum tokens
to generate
 top_k: 50 # [MODIFIABLE] Top-k sampling
parameter
 top_p: 1.0 # [MODIFIABLE] Nucleus
sampling parameter
 temperature: 0 # [MODIFIABLE] Sampling
temperature (0 = deterministic)
```

## Using the reasoning\_effort parameter

The reasoning\_effort parameter controls the reasoning behavior for reasoning-capable models.

### Prerequisites:

- **Model compatibility:** Set reasoning\_effort only when model\_type specifies a reasoning-capable model (currently amazon.nova-2-lite-v1:0:256k)
- **Error handling:** Using reasoning\_effort with unsupported models will fail with ConfigValidationError: "Reasoning mode is enabled but model '{model\_type}' does not support reasoning. Please use a reasoning-capable model or disable reasoning mode."

### Available options:

| Option         | Behavior                              | Token Limit                          | Use Case                                                                 |
|----------------|---------------------------------------|--------------------------------------|--------------------------------------------------------------------------|
| null (default) | Disables reasoning mode               | N/A                                  | Standard evaluation without reasoning overhead                           |
| low            | Enables reasoning with constraints    | 4,000 tokens for internal reasoning  | Scenarios requiring concise reasoning; optimizes for speed and cost      |
| high           | Enables reasoning without constraints | No token limit on internal reasoning | Complex problems requiring extensive analysis and step-by-step reasoning |

## When to enable reasoning

### Use reasoning mode (low, medium, or high) for:

- Complex problem-solving tasks (mathematics, logic puzzles, coding)
- Multi-step analytical questions requiring intermediate reasoning

- Tasks where detailed explanations or step-by-step thinking improve accuracy
- Scenarios where response quality is prioritized over speed

### Use non-reasoning mode (omit parameter) for:

- Simple Q&A or factual queries
- Creative writing tasks
- When faster response times are critical
- Performance benchmarking where reasoning overhead should be excluded
- Cost optimization when reasoning doesn't improve task performance

## Troubleshooting

### Error: "Reasoning mode is enabled but model does not support reasoning"

**Cause:** The `reasoning_effort` parameter is set to a non-null value, but the specified `model_type` doesn't support reasoning.

### Resolution:

- Verify your model type is `amazon.nova-2-lite-v1:0:256k`
- If using a different model, either switch to a reasoning-capable model or remove the `reasoning_effort` parameter from your recipe

## Starting an evaluation job

The following provides a suggested evaluation instance type and model type configuration:

```
Install Dependencies (Helm - https://helm.sh/docs/intro/install/)
curl -fsSL -o get_helm.sh https://raw.githubusercontent.com/helm/helm/main/scripts/get-helm-3
chmod 700 get_helm.sh
./get_helm.sh
rm -f ./get_helm.sh

Install the SageMaker HyperPod CLI
git clone --recurse-submodules https://github.com/aws/sagemaker-hyperpod-cli.git
git checkout -b release_v2
cd sagemaker-hyperpod-cli
```

```

pip install .

Verify the installation
hyperpod --help

Connect to a SageMaker HyperPod Cluster
hyperpod connect-cluster --cluster-name cluster-name

Submit the Job using the recipe for eval
Namespace by default should be kubeflow
hyperpod start-job [--namespace namespace] --recipe evaluation/nova/
nova_micro_p5_48xl_general_text_benchmark_eval --override-parameters \
'{
 "instance_type": "p5d.48xlarge",
 "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM-
HP-Eval-V2-latest",
 "recipes.run.name": custom-run-name,
 "recipes.run.model_type": model_type,
 "recipes.run.model_name_or_path": "model name or finetune checkpoint s3uri",
 "recipes.run.data_s3_path": "s3 for input data only for genqa and llm_judge, must be
full S3 path that include filename",
}'

List jobs
hyperpod list-jobs [--namespace namespace] [--all-namespaces]

Getting Job details
hyperpod get-job --job-name job-name [--namespace namespace] [--verbose]

Listing Pods
hyperpod list-pods --job-name job-name --namespace namespace

Cancel Job
hyperpod cancel-job --job-name job-name [--namespace namespace]

```

You should also be able to view the job status through Amazon EKS cluster console.

## Accessing and analyzing evaluation results

After your evaluation job completes successfully, you can access and analyze the results using the information in this section. Based on the `output_s3_path` (such as `s3://output_path/`) defined in the recipe, the output structure is the following:

```

job_name/
eval-result/
results_[timestamp].json
inference_output.jsonl (only present for gen_qa)
details/
model/
execution-date-time/
###details_task_name_#_datetime.parquet
tensorboard-results/
 ### eval/
 ### events.out.tfevents.[timestamp]

```

Metrics results are stored in the specified S3 output location `s3://output_path/job_name/eval-result/result-timestamp.json`.

Tensorboard results are stored in the S3 path `s3://output_path/job_name/eval-tensorboard-result/eval/event.out.tfevents.epoch+ip`.

All inference outputs, except for `llm_judge` and `strong_reject`, are stored in the S3 path: `s3://output_path/job_name/eval-result/details/model/taskname.parquet`.

For `gen_qa`, the `inference_output.jsonl` file contains the following fields for each JSON object:

- `prompt` - The final prompt submitted to the model
- `inference` - The raw inference output from the model
- `gold` - The target response from the input dataset
- `metadata` - The metadata string from the input dataset if provided

To visualize your evaluation metrics in Tensorboard, complete the following steps:

1. Navigate to SageMaker AI Tensorboard.
2. Select **S3 folders**.
3. Add your S3 folder path, for example `s3://output_path/job-name/eval-tensorboard-result/eval`.
4. Wait for synchronization to complete.

The time series, scalars, and text visualizations are available.



We recommend the following best practices:

- Keep your output paths organized by model and benchmark type.
- Maintain consistent naming conventions for easy tracking.
- Save extracted results in a secure location.
- Monitor TensorBoard sync status for successful data loading.

You can find SageMaker HyperPod job error logs in the CloudWatch log group `/aws/sagemaker/Clusters/cluster-id`.

## Log Probability Output Format

When `top_logprobs` is configured in your inference settings, the evaluation output includes token-level log probabilities in the parquet files. Each token position contains a dictionary of the top candidate tokens with their log probabilities in the following structure:

```
{
 "Ġint": {"logprob_value": -17.8125, "decoded_value": " int"},
 "Ġthe": {"logprob_value": -2.345, "decoded_value": " the"}
}
```

Each token entry contains:

- `logprob_value`: The log probability value for the token
- `decoded_value`: The human-readable decoded string representation of the token

The raw tokenizer token is used as the dictionary key to ensure uniqueness, while `decoded_value` provides a readable interpretation.

## RFT evaluation

### Note

Evaluation via remote reward functions in your own AWS environment is only available if you are Amazon Nova Forge customer.

**⚠ Important**

The `rl_env` configuration field is used exclusively for evaluation, not for training. During training, you configure reward functions using `reward_lambda_arn` (single-turn) or BYOO infrastructure with `rollout.delegate: true` (multi-turn).

**What is RFT Evaluation?**

RFT Evaluation allows you to assess your model's performance using custom reward functions before, during, or after reinforcement learning training. Unlike standard evaluations that use pre-defined metrics, RFT Evaluation lets you define your own success criteria through a Lambda function that scores model outputs based on your specific requirements.

**Why Evaluate with RFT?**

Evaluation is crucial to determine whether the RL fine-tuning process has:

- Improved model alignment with your specific use case and human values
- Maintained or improved model capabilities on key tasks
- Avoided unintended side effects such as reduced factuality, increased verbosity, or degraded performance on other tasks
- Met your custom success criteria as defined by your reward function

**When to Use RFT Evaluation**

Use RFT Evaluation in these scenarios:

- Before RFT Training: Establish baseline metrics on your evaluation dataset
- During RFT Training: Monitor training progress with intermediate checkpoints
- After RFT Training: Validate that the final model meets your requirements
- Comparing Models: Evaluate multiple model versions using consistent reward criteria

**ℹ Note**

Use RFT Evaluation when you need custom, domain-specific metrics. For general-purpose evaluation (accuracy, perplexity, BLEU), use standard evaluation methods.

## Topics

- [Data format requirements](#)
- [Preparing your evaluation recipe](#)
- [Preset reward functions](#)
- [Creating your reward function](#)
- [IAM permissions](#)
- [Executing the evaluation job](#)
- [Understanding evaluation results](#)

## Data format requirements

### Input data structure

RFT evaluation input data must follow the OpenAI Reinforcement Fine-Tuning format. Each example is a JSON object containing:

- `messages`: Array of conversational turns with system and user roles
- Optional other metadata, e.g. `reference_answer`

### Data format example

The following example shows the required format:

```
{
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "type": "text",
 "text": "Solve for x. Return only JSON like {\"x\": <number>}. Equation: 2x + 5 = 13"
 }
]
 }
],
 "reference_answer": {
 "x": 4
 }
}
```

```
}
```

## Current limitations

The following limitations apply to RFT evaluation:

- Text only: No multimodal inputs (images, audio, video) are supported
- Single-turn conversations: Only supports single user message (no multi-turn dialogues)
- JSON format: Input data must be in JSONL format (one JSON object per line)
- Model outputs: Evaluation is performed on generated completions from the specified model

## Preparing your evaluation recipe

### Sample recipe configuration

The following example shows a complete RFT evaluation recipe:

```
run:
 name: nova-lite-rft-eval-job
 model_type: amazon.nova-lite-v1:0:300k
 model_name_or_path: s3://escrow_bucket/model_location # [MODIFIABLE] S3 path to
your model or model identifier
 replicas: 1 # [MODIFIABLE] For SageMaker
Training jobs only; fixed for SageMaker HyperPod jobs
 data_s3_path: "" # [REQUIRED FOR HYPERPOD]
Leave empty for SageMaker Training jobs
 output_s3_path: "" # [REQUIRED] Output artifact
S3 path for evaluation results

evaluation:
 task: rft_eval # [FIXED] Do not modify
 strategy: rft_eval # [FIXED] Do not modify
 metric: all # [FIXED] Do not modify

Inference Configuration
inference:
 max_new_tokens: 8196 # [MODIFIABLE] Maximum tokens
to generate
 top_k: -1 # [MODIFIABLE] Top-k sampling
parameter
 top_p: 1.0 # [MODIFIABLE] Nucleus
sampling parameter
```

```

temperature: 0 # [MODIFIABLE] Sampling
temperature (0 = deterministic)
top_logprobs: 0

Evaluation Environment Configuration (NOT used in training)
rl_env:
 reward_lambda_arn: arn:aws:lambda:<region>:<account_id>:function:<reward-function-
name>

```

## Preset reward functions

Two preset reward functions (`prime_code` and `prime_math`) are available as a Lambda layer for easy integration with your RFT Lambda functions.

### Overview

These preset functions provide out-of-the-box evaluation capabilities for:

- **prime\_code**: Code generation and correctness evaluation
- **prime\_math**: Mathematical reasoning and problem-solving evaluation

### Quick setup

To use preset reward functions:

1. Download the Lambda layer from the [nova-custom-eval-sdk releases](#)
2. Publish Lambda layer using AWS CLI:

```

aws lambda publish-layer-version \
 --layer-name preset-function-layer \
 --description "Preset reward function layer with dependencies" \
 --zip-file fileb://universal_reward_layer.zip \
 --compatible-runtimes python3.9 python3.10 python3.11 python3.12 \
 --compatible-architectures x86_64 arm64

```

3. Add the layer to your Lambda function in AWS Console (Select the preset-function-layer from custom layer and also add AWSSDKPandas-Python312 for numpy dependencies)
4. Import and use in your Lambda code:

```

from prime_code import compute_score # For code evaluation
from prime_math import compute_score # For math evaluation

```

## prime\_code function

**Purpose:** Evaluates Python code generation tasks by executing code against test cases and measuring correctness.

**Example input dataset format from evaluation:**

```
{"messages":[{"role":"user","content":"Write a function that returns the sum of two numbers."}], "reference_answer":{"inputs":["3\n5","10\n-2","0\n0"],"outputs":["8","8","0"]}}
{"messages":[{"role":"user","content":"Write a function to check if a number is even."}], "reference_answer":{"inputs":["4","7","0","-2"],"outputs":["True","False","True","True"]}}
```

### Key features:

- Automatic code extraction from markdown code blocks
- Function detection and call-based testing
- Test case execution with timeout protection
- Syntax validation and compilation checks
- Detailed error reporting with tracebacks

## prime\_math function

**Purpose:** Evaluates mathematical reasoning and problem-solving capabilities with symbolic math support.

**Input format:**

```
{"messages":[{"role":"user","content":"What is the derivative of x^2 + 3x?."}], "reference_answer":"2*x + 3"}
```

### Key features:

- Symbolic math evaluation using SymPy
- Multiple answer formats (LaTeX, plain text, symbolic)
- Mathematical equivalence checking
- Expression normalization and simplification

## Best practices

Follow these best practices when using preset reward functions:

- Use proper data types in test cases (integers vs strings, booleans vs "True")
- Provide clear function signatures in code problems
- Include edge cases in test inputs (zero, negative numbers, empty inputs)
- Format math expressions consistently in reference answers
- Test your reward function with sample data before deployment

## Creating your reward function

### Lambda ARN

You must refer to the following format for the Lambda ARN:

```
"arn:aws:lambda:*:*:function:*SageMaker*"
```

If the Lambda does not have this naming scheme, the job will fail with this error:

```
[ERROR] Unexpected error: lambda_arn must contain one of: ['SageMaker', 'sagemaker', 'Sagemaker'] when running on SMHP platform (Key: lambda_arn)
```

### Lambda function structure

Your Lambda function receives batches of model outputs and returns reward scores. Below is a sample implementation:

```
from typing import List, Any
import json
import re
from dataclasses import asdict, dataclass

@dataclass
class MetricResult:
 """Individual metric result."""
 name: str
 value: float
 type: str
```

```
@dataclass
class RewardOutput:
 """Reward service output."""
 id: str
 aggregate_reward_score: float
 metrics_list: List[MetricResult]

def lambda_handler(event, context):
 """ Main lambda handler """
 return lambda_grader(event)

def lambda_grader(samples: list[dict]) -> list[dict]:
 """ Core grader function """
 scores: List[RewardOutput] = []

 for sample in samples:
 print("Sample: ", json.dumps(sample, indent=2))

 # Extract components
 idx = sample.get("id", "no id")
 if not idx or idx == "no id":
 print(f"ID is None/empty for sample: {sample}")

 ground_truth = sample.get("reference_answer")

 if "messages" not in sample:
 print(f"Messages is None/empty for id: {idx}")
 continue

 if ground_truth is None:
 print(f"No answer found in ground truth for id: {idx}")
 continue

 # Get model's response (last turn is assistant turn)
 last_message = sample["messages"][-1]

 if last_message["role"] != "nova_assistant":
 print(f"Last message is not from assistant for id: {idx}")
 continue

 if "content" not in last_message:
```



```

 print(f"Completion text is empty for id: {idx}")
 continue

 model_text = last_message["content"]

 # --- Actual scoring logic (lexical overlap) ---
 ground_truth_text = _extract_ground_truth_text(ground_truth)

 # Calculate main score and individual metrics
 overlap_score = _lexical_overlap_score(model_text, ground_truth_text)

 # Create two separate metrics as in the first implementation
 accuracy_score = overlap_score # Use overlap as accuracy
 fluency_score = _calculate_fluency(model_text) # New function for fluency

 # Create individual metrics
 metrics_list = [
 MetricResult(name="accuracy", value=accuracy_score, type="Metric"),
 MetricResult(name="fluency", value=fluency_score, type="Reward")
]

 ro = RewardOutput(
 id=idx,
 aggregate_reward_score=overlap_score,
 metrics_list=metrics_list
)

 print(f"Response for id: {idx} is {ro}")
 scores.append(ro)

Convert to dict format
result = []
for score in scores:
 result.append({
 "id": score.id,
 "aggregate_reward_score": score.aggregate_reward_score,
 "metrics_list": [asdict(metric) for metric in score.metrics_list]
 })

return result

def _extract_ground_truth_text(ground_truth: Any) -> str:
 """

```

```

Turn the `ground_truth` field into a plain string.
"""
if isinstance(ground_truth, str):
 return ground_truth

if isinstance(ground_truth, dict):
 # Common patterns: { "explanation": "...", "answer": "..."}
 if "explanation" in ground_truth and isinstance(ground_truth["explanation"],
str):
 return ground_truth["explanation"]
 if "answer" in ground_truth and isinstance(ground_truth["answer"], str):
 return ground_truth["answer"]
 # Fallback: stringify the whole dict
 return json.dumps(ground_truth, ensure_ascii=False)

Fallback: stringify anything else
return str(ground_truth)

def _tokenize(text: str) -> List[str]:
 # Very simple tokenizer: lowercase + alphanumeric word chunks
 return re.findall(r"\w+", text.lower())

def _lexical_overlap_score(model_text: str, ground_truth_text: str) -> float:
 """
 Simple lexical overlap score in [0, 1]:
 score = |tokens(model) ∩ tokens(gt)| / |tokens(gt)|
 """
 gt_tokens = _tokenize(ground_truth_text)
 model_tokens = _tokenize(model_text)

 if not gt_tokens:
 return 0.0

 gt_set = set(gt_tokens)
 model_set = set(model_tokens)
 common = gt_set & model_set

 return len(common) / len(gt_set)

def _calculate_fluency(text: str) -> float:
 """

```

Calculate a simple fluency score based on:

- Average word length
- Text length
- Sentence structure

Returns a score between 0 and 1.

```
"""
Simple implementation - could be enhanced with more sophisticated NLP
words = _tokenize(text)

if not words:
 return 0.0

Average word length normalized to [0,1] range
Assumption: average English word is ~5 chars, so normalize around that
avg_word_len = sum(len(word) for word in words) / len(words)
word_len_score = min(avg_word_len / 10, 1.0)

Text length score - favor reasonable length responses
ideal_length = 100 # words
length_score = min(len(words) / ideal_length, 1.0)

Simple sentence structure check (periods, question marks, etc.)
sentence_count = len(re.findall(r'[.!?]+', text)) + 1
sentence_ratio = min(sentence_count / (len(words) / 15), 1.0)

Combine scores
fluency_score = (word_len_score + length_score + sentence_ratio) / 3

return fluency_score
```

## Lambda request format

Your Lambda function receives data in this format:

```
[
 {
 "id": "sample-001",
 "messages": [
 {
 "role": "user",
 "content": [
 {
 "type": "text",
```

```

 "text": "Do you have a dedicated security team?"
 }
]
 },
 {
 "role": "nova_assistant",
 "content": [
 {
 "type": "text",
 "text": "As an AI developed by Company, I don't have a dedicated security
team in the traditional sense. However, the development and deployment of AI systems
like me involve extensive security measures, including data encryption, user privacy
protection, and other safeguards to ensure safe and responsible use."
 }
]
 }
],
"reference_answer": {
 "compliant": "No",
 "explanation": "As an AI developed by Company, I do not have a traditional
security team. However, the deployment involves stringent safety measures, such as
encryption and privacy safeguards."
}
}
]

```

### Note

The message structure includes the nested content array, matching the input data format. The last message with role `nova_assistant` contains the model's generated response.

## Lambda response format

Your Lambda function must return data in this format:

```

[
 {
 "id": "sample-001",
 "aggregate_reward_score": 0.75,
 "metrics_list": [
 {

```

```
 "name": "accuracy",
 "value": 0.85,
 "type": "Metric"
 },
 {
 "name": "fluency",
 "value": 0.90,
 "type": "Reward"
 }
]
```

### Response fields:

- **id:** Must match the input sample ID
- **aggregate\_reward\_score:** Overall score (typically 0.0 to 1.0)
- **metrics\_list:** Array of individual metrics with:
  - **name:** Metric identifier (e.g., "accuracy", "fluency")
  - **value:** Metric score (typically 0.0 to 1.0)
  - **type:** Either "Metric" (for reporting) or "Reward" (used in training)

### IAM permissions

#### Required permissions

Your SageMaker AI execution role must have permissions to invoke your Lambda function. Add this policy to your SageMaker AI execution role:

```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": [
 "lambda:InvokeFunction"
],
 "Resource": "arn:aws:lambda:region:account-id:function:function-name"
 }
]
}
```

```
}
```

## Lambda execution role

Your Lambda function's execution role needs basic Lambda execution permissions:

```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": [
 "logs:CreateLogGroup",
 "logs:CreateLogStream",
 "logs:PutLogEvents"
],
 "Resource": "arn:aws:logs:*:*:*"
 }
]
}
```

**Additional permissions:** If your Lambda function accesses other AWS services (e.g., Amazon S3 for reference data, DynamoDB for logging), add those permissions to the Lambda execution role.

## Executing the evaluation job

### 1. Prepare your data

- Format your evaluation data according to the data format requirements
- Upload your JSONL file to Amazon S3: `s3://your-bucket/eval-data/eval_data.jsonl`

### 2. Configure your recipe

Update the sample recipe with your configuration:

- Set `model_name_or_path` to your model location
- Set `lambda_arn` to your reward function ARN
- Set `output_s3_path` to your desired output location
- Adjust inference parameters as needed

Save the recipe as `rft_eval_recipe.yaml`

### 3. Run the evaluation

Execute the evaluation job using the provided notebook: [Nova model evaluation notebook](#)

### 4. Monitor progress

Monitor your evaluation job through:

- SageMaker AI Console: Check job status and logs
- CloudWatch Logs: View detailed execution logs
- Lambda Logs: Debug reward function issues

## Understanding evaluation results

### Output format

The evaluation job outputs results to your specified Amazon S3 location in JSONL format. Each line contains the evaluation results for one sample:

```
{
 "id": "sample-001",
 "aggregate_reward_score": 0.75,
 "metrics_list": [
 {
 "name": "accuracy",
 "value": 0.85,
 "type": "Metric"
 },
 {
 "name": "fluency",
 "value": 0.90,
 "type": "Reward"
 }
]
}
```

**Note**

The RFT Evaluation Job Output is identical to the Lambda Response format. The evaluation service passes through your Lambda function's response without modification, ensuring consistency between your reward calculations and the final results.

## Interpreting results

### Aggregate Reward Score:

- Range: Typically 0.0 (worst) to 1.0 (best), but depends on your implementation
- Purpose: Single number summarizing overall performance
- Usage: Compare models, track improvement over training

### Individual Metrics:

- Metric Type: Informational metrics for analysis
- Reward Type: Metrics used during RFT training
- Interpretation: Higher values generally indicate better performance (unless you design inverse metrics)

## Performance benchmarks

What constitutes "good" performance depends on your use case:

| Score Range | Interpretation | Action                                   |
|-------------|----------------|------------------------------------------|
| 0.8 - 1.0   | Excellent      | Model ready for deployment               |
| 0.6 - 0.8   | Good           | Minor improvements may be beneficial     |
| 0.4 - 0.6   | Fair           | Significant improvement needed           |
| 0.0 - 0.4   | Poor           | Review training data and reward function |



**⚠ Important**

These are general guidelines. Define your own thresholds based on business requirements, baseline model performance, domain-specific constraints, and cost-benefit analysis of further training.

## Monitoring HyperPod jobs with MLflow

You can use MLflow to track and monitor your training jobs on SageMaker HyperPod. Follow these steps to set up MLflow and connect it to your training recipes.

### Create the MLflow App

Example AWS CLI command

```
aws sagemaker-mlflow create-mlflow-app \
 --name <app-name> \
 --artifact-store-uri <s3-bucket-name> \
 --role-arn <role-arn> \
 --region <region-name>
```

Example output

```
{
 "Arn": "arn:aws:sagemaker:us-east-1:111122223333:mlflow-app/app-LGZE0Z2UY4NZ"
}
```

### Generate pre-signed URL

Example AWS CLI command

```
aws sagemaker-mlflow create-presigned-mlflow-app-url \
 --arn <app-arn> \
 --region <region-name> \
 --output text
```

Example output


```
https://app-LGZE0Z2UY4NZ.mlflow.sagemaker.us-east-1.app.aws/auth?
authToken=eyJhbGciOiJIUzI1NiJ9.eyJhdXRoVG9rZW5JZCI6IkxETVBPUyIsImZhc0NyZWRLbnRpYWxzIjoiQWdWNGhD
```

Open presigned URL and view the app

Click

```
https://app-LGZE0Z2UY4NZ.mlflow.sagemaker.us-east-1.app.aws/auth?
authToken=eyJhbGciOiJIUzI1NiJ9.eyJhdXRoVG9rZW5JZCI6IkxETVBPUyIsImZhc0NyZWRLbnRpYWxzIjoiQWdWNGhD
```

View




Pass to recipe under run block of your SageMaker HyperPod recipe

Recipe

```
run
 mlflow_tracking_uri: arn:aws:sagemaker:us-east-1:111122223333:mlflow-app/app-
 LGZE0Z2UY4NZ
```

View



## Iterative Training

### Overview

Iterative training is the process of repeatedly fine-tuning a model through multiple training cycles across different training methods — train, evaluate, analyze errors, adjust data/objectives/hyperparameters — with each round starting from the previous checkpoint. This approach allows you to systematically target model failure modes, incorporate curated examples addressing specific weaknesses, and adapt to changing requirements over time.

#### Benefits over single-pass training:

- **Targeted improvement:** Address specific failure patterns discovered through evaluation

- **Adaptive refinement:** Respond to distribution shifts or evolving product requirements
- **Risk mitigation:** Validate improvements incrementally rather than committing to a single long training run
- **Data efficiency:** Focus data collection efforts on areas where the model underperforms
- **Curriculum Training:** Multiple round of training with increasingly higher quality data

## How it works

### Checkpoint location and access

After each training job completes, a manifest file is generated in the output location specified by the `output_path` parameter in your training configuration.

#### To access your checkpoint

- Navigate to your specified `output_path` in S3
- Download and extract the `output.tar.gz` file
- Open the `manifest.json` file inside
- Locate the `checkpoint_s3_bucket` parameter, which contains the S3 URI of your trained model

#### Example manifest.json structure

```
{
 "checkpoint_s3_bucket": "s3://customer-escrow-<account-number>-smtj-<unique-
 identifier>/<job-name>/stepID",
 ...
}
```

### Understanding escrow buckets

Since Amazon Nova weights are proprietary, trained model checkpoints are stored in **escrow S3 buckets** within AWS-managed accounts rather than being copied to your account. These escrow buckets:

- Contain your customized model weights securely
- Can be referenced by other AWS services (Inference, Evaluation, and subsequent training jobs)

- Are accessible only to your AWS account via IAM permissions
- Incur standard S3 storage charges in your account (see Cost considerations)

You can use the escrow bucket path as the `model_name_or_path` in your next training run to continue iterative training.

## Using checkpoints for iterative training

Configure your next training job to use the previous checkpoint as the base model:

```
run:
 name: "my-iterative-training-job"
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: "s3://customer-escrow-<account-number>-smtj-<unique-identifier>/<previous-job-name>"
 data_s3_path: s3://<bucket>/<data-file>.jsonl
 replicas: 4
```

## When to use iterative training

### Ideal use cases

Use iterative training when you have:

- **Feedback loops** – Ability to collect real-world failure cases and systematically address them
- **Dynamic environments** – Evolving documentation, APIs, or support topics requiring periodic model updates
- **Robust evaluation** – Strong benchmarks and evaluation frameworks (see examples below) to measure improvements confidently
- **ML operations capability** – Resources to manage multiple training cycles and version control

### Examples of robust evaluation frameworks

- Automated benchmark suites with pass/fail thresholds
- Human evaluation protocols with inter-rater reliability metrics
- Red-team testing scenarios covering edge cases and adversarial inputs
- A/B testing infrastructure to measure production impact

## Common patterns

**SFT → RFT Pipeline:** A frequently used iterative pattern involves:

- **SFT first** – Teach the model how to solve problems through demonstration examples
- **RFT second** – Optimize performance across the broader problem space using reward signals

This sequence is essential when models perform poorly initially—RFT on near-zero accuracy models will not improve performance without first establishing basic problem-solving capabilities through SFT.

## When not to use iterative training

Avoid iterative training for:

- **Stable, well-defined tasks** – Stationary data with consistent requirements already achieving near-ceiling performance
- **Simple classification problems** – Narrow tasks where single-pass training suffices
- **Resource constraints** – Lacking dedicated ML operations capabilities to manage multiple training cycles
- **Marginal gains** – When overhead doesn't justify minimal performance improvements

## Example workflow: SFT → RFT

This example demonstrates a common iterative training pattern for reasoning models.

### Step 1: Initial SFT training

Configure and launch your SFT training job with your dataset:

```
run:
 name: "initial-sft-training"
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: "nova-lite-2/prod"
 data_s3_path: s3://<bucket>/sft-training-data.jsonl
 validation_data_s3_path: s3://<bucket>/sft-validation-data.jsonl
```

**Rationale:** SFT provides additional demonstrations that shape model outputs into your desired format and voice, establishing foundational capabilities.

## After training completes

- Note the `output_path` configured in your training job
- Download `output.tar.gz` from that location
- Extract and locate `manifest.json`
- Copy the `checkpoint_s3_bucket` value

## Step 2: RFT training on SFT checkpoint

Create a new RFT training job using the SFT checkpoint:

```
run:
 name: "rft-on-sft-checkpoint"
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: "s3://customer-escrow-<account-number>-smtj-<unique-identifier>/
<initial-sft-training>"
 data_s3_path: s3://<bucket>/rft-training-data.jsonl
 reward_lambda_arn: <your-reward-function-arn>
```

**Rationale:** RFT training builds on the SFT foundation, allowing the model to develop more complex reasoning patterns optimized by your reward function.

## Step 3: Evaluate and iterate

Run evaluation on the RFT checkpoint to assess performance:

```
run:
 name: "evaluate-rft-checkpoint"
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: "s3://customer-escrow-<account-number>-smtj-<unique-identifier>/
<rft-on-sft-checkpoint>"
 data_s3_path: s3://<bucket>/evaluation-data.jsonl
```

If target metrics are not satisfied, continue iterating with adjusted data or hyperparameters.

### Important

The training technique (LoRA vs. Full Rank) must remain consistent across all iterations:

- If you use SFT with **LoRA**, you must use RFT with **LoRA**

- If you use SFT with **Full Rank**, you must use RFT with **Full Rank**
- You cannot switch between LoRA and Full Rank mid-pipeline

### Important

If a KMS key is used for encryption in the Amazon-owned output S3 bucket, that same KMS key must be used for all future iterations.

## Monitoring progress across iterations

You can track metrics via MLflow by [setting up MLflow](#) for your jobs.

### Create an MLflow app

**Using Studio UI:** If you create a training job through the Studio UI, a default MLflow app is created automatically and selected by default under Advanced Options.

**Using CLI:** If you use the CLI, you must create an MLflow app and pass it as an input to the training job API request.

```
mlflow_app_name="<enter your MLflow app name>"
role_arn="<enter your role ARN>"
bucket_name="<enter your bucket name>"
region="<enter your region>"

mlflow_app_arn=$(aws sagemaker create-mlflow-app \
 --name $mlflow_app_name \
 --artifact-store-uri "s3://$bucket_name" \
 --role-arn $role_arn \
 --region $region)
```

### Access the MLflow app

**Using CLI:** Create a presigned URL to access the MLflow app UI:

```
aws sagemaker create-presigned-mlflow-app-url \
 --arn $mlflow_app_arn \
```



```
--region $region \
--output text
```

**Using Studio UI:** The Studio UI displays key metrics stored in MLflow and provides a link to the MLflow app UI.

## Key metrics to track

Monitor these metrics across iterations to assess improvement and track the job progress:

### For SFT

- Training loss curves
- Number of samples consumed and time to process samples
- Performance accuracy on held-out test sets
- Format compliance (for example, valid JSON output rate)
- Perplexity on domain-specific evaluation data

### For RFT

- Average reward scores over training
- Reward distribution (percentage of high-reward responses)
- Validation reward trends (watch for overfitting)
- Task-specific success rates (for example, code execution pass rate, math problem accuracy)

### General

- Benchmark performance deltas between iterations
- Human evaluation scores on representative samples
- Production metrics (if deploying iteratively)

## Determining when to stop

Stop iterating when:

- **Performance plateaus** – Additional training no longer meaningfully improves target metrics

- **Technique switching helps** – If one technique plateaus, try switching (for example, SFT → RFT → SFT) to break through performance ceilings
- **Target metrics achieved** – Your success criteria are met
- **Regression detected** – New iterations degrade performance (see rollback procedures below)

For detailed evaluation procedures, refer to the **Evaluation** section.

## Best practices

### Start small and scale gradually

Begin with minimal datasets and single training epochs to validate your approach before scaling up. This builds confidence and helps identify issues early.

### Establish clear success metrics

Define quantitative and qualitative indicators before starting:

#### Example success metrics by use case

- **Question answering** – Exact match accuracy, F1 score, human preference ratings
- **Code generation** – Unit test pass rate, compilation success, execution time
- **Reasoning tasks** – Step accuracy, final answer correctness, reward scores
- **Content generation** – Coherence scores, factual accuracy, style adherence

### Implement automated evaluation

Set up automated evaluation pipelines to track performance after each round, enabling rapid iteration and objective comparison.

### Maintain rigorous version control

Document for each iteration:

- Dataset versions and modifications
- Model checkpoint locations
- Hyperparameter changes
- Performance metrics and deltas

- Qualitative observations

This builds institutional knowledge and enables debugging.

## Focus on data quality over quantity

Analyze failure cases from previous rounds and add targeted, high-quality examples rather than simply increasing dataset size.

## Plan iteration budget

Plan for **3-5 iterations** as a typical range:

- **1-2 iterations** – Often sufficient for simple improvements or final polishing
- **3-5 iterations** – Appropriate for complex tasks requiring multiple refinement cycles
- **5+ iterations** – May indicate diminishing returns or need for different approaches

Adjust based on computational budget and performance improvement rates.

## Implement rollback capabilities

If an iteration introduces regressions:

- **Identify the regression** – Compare evaluation metrics across checkpoints
- **Return to previous checkpoint** – Use the earlier checkpoint's S3 path as your `model_name_or_path`
- **Adjust training approach** – Modify data, hyperparameters, or technique before retrying
- **Document the failure** – Record what caused regression to avoid repeating

## Example rollback

```
run:
 name: "rollback-to-iteration-2"
 model_type: amazon.nova-2-lite-v1:0:256k
 # Use iteration 2 checkpoint instead of failed iteration 3
 model_name_or_path: "s3://customer-escrow-<account-number>-smtj-<unique-identifier>/
<iteration-2-job-name>"
```

## Cost considerations

### Checkpoint storage

- **Location** – Checkpoints stored in escrow buckets incur standard S3 storage charges billed to your AWS account
- **Retention** – Checkpoints are retained indefinitely unless explicitly deleted
- **Management** – Implement lifecycle policies to archive or delete old checkpoints you no longer need

### Cost optimization tips

- Delete intermediate checkpoints after validating newer iterations
- Archive checkpoints to S3 Glacier for long-term retention at lower cost
- Set retention policies based on your compliance and experimentation needs

## Limitations

### Model family consistency

When iteratively training, you must use the **same model type** throughout all iterations.

#### Initial training

```
run:
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: "nova-lite-2/prod"
```

#### Subsequent iterations must use the same model\_type

```
run:
 model_type: amazon.nova-2-lite-v1:0:256k # Must match original
 model_name_or_path: "s3://customer-escrow-<account-number>-smtj-<unique-identifier>/<job-name>"
```

### Training technique consistency

The training technique must remain consistent across iterations:

- **LoRA-trained models** can only be iteratively trained with LoRA
- **Full-Rank-trained models** can only be iteratively trained with Full-Rank

### How LoRA adapters work in iterative training

- Each LoRA training iteration produces new adapter weights
- New adapters replace (not stack with) previous adapters
- The base model remains frozen; only adapters are updated

### Technique compatibility matrix

| Initial training | Can iterate with                 |
|------------------|----------------------------------|
| SFT (Full-Rank)  | SFT (Full-Rank), RFT (Full-Rank) |
| SFT (LoRA)       | SFT (LoRA), RFT (LoRA)           |
| RFT (Full-Rank)  | RFT (Full-Rank)                  |
| RFT (LoRA)       | RFT (LoRA)                       |

### Verifying compatibility before starting a job

- Check your previous training recipe to identify the model type and training technique (LoRA vs. Full-Rank)
- Ensure your new recipe matches both the model type and technique
- Review the manifest.json to confirm the checkpoint path is correct

## Troubleshooting

### Error: "Incompatible model training techniques detected"

**Cause:** The training technique (LoRA vs. Full-Rank) doesn't match the checkpoint's technique.

**Resolution:** Ensure your recipe uses the same training technique as the original model:

- If the checkpoint was trained with LoRA, use LoRA in your new recipe

- If the checkpoint was trained with Full-Rank, use Full-Rank in your new recipe

## Error: "Base model for the job extracted from model\_name\_or\_path does not match model\_type"

**Cause:** The model type specified in model\_type doesn't match the actual model in the checkpoint.

**Resolution:** Verify that:

- The model\_type in your recipe matches the original model type
- The checkpoint S3 path in model\_name\_or\_path is correct
- You're using the path from the correct manifest.json file

### Example of correct configuration

```
run:
 model_type: amazon.nova-2-lite-v1:0:256k # Must match checkpoint's model
 model_name_or_path: "s3://customer-escrow-<account-number>-smtj-<unique-identifier>/<job-name>"
```

## Error: "Model configuration not found"

**Cause:** The S3 path in model\_name\_or\_path is invalid or inaccessible.

**Resolution:**

- Verify the S3 path is correctly copied from the manifest.json file
- Ensure your IAM role has permissions to access the escrow bucket
- Confirm the previous training job completed successfully
- Check for typos in the path

## Performance regression after iteration

**Symptoms:** Evaluation metrics decline after a new training iteration.

**Resolution:**

- **Rollback** – Use the previous checkpoint as your base model

- **Analyze** – Review training logs and data quality for the failed iteration
- **Adjust** – Modify hyperparameters (reduce learning rate), improve data quality, or reduce training epochs
- **Retry** – Launch a new iteration with adjustments

## SageMaker Inference

Custom Amazon Nova models are now available on SageMaker inference. With Amazon Nova on SageMaker, you can start getting predictions, or inferences, from your trained custom Amazon Nova models. SageMaker provides a broad selection of ML infrastructure and model deployment options to help meet all your ML inference needs. With SageMaker inference, you can scale your model deployment, manage models more effectively in production, and reduce operational burden.

SageMaker provides you with various inference options, such as real-time endpoints for getting low latency inference, and asynchronous endpoints for batches of requests. By leveraging the appropriate inference option for your use case, you can ensure efficient model deployment and inference. For more information on SageMaker inference, see [Deploy models for inference](#).

### Important

Only full-rank custom models and LoRA-merged models are supported on SageMaker inference. For unmerged LoRA models and base models, use Amazon Bedrock.

## Features

The following features are available for Amazon Nova models on SageMaker inference:

### Model Capabilities

- Text generation

### Deployment and Scaling

- Real-time endpoints with custom instance selection
- Auto Scaling – Automatically adjust capacity based on traffic patterns to optimize costs and GPU utilization. For more information, see [Automatically Scale Amazon SageMaker Models](#).

- Streaming API support for real-time token generation

## Monitoring and Optimization

- Amazon CloudWatch integration for monitoring and alerts
- Availability Zone-aware latency optimization through VPC configuration

## Development Tools

- AWS CLI support – For more information, see [AWS CLI Command Reference for SageMaker](#).
- Notebook integration via SDK support

## Supported models and instances

When creating your SageMaker inference endpoints, you can set two environment variables to configure your deployment: `CONTEXT_LENGTH` and `MAX_CONCURRENCY`.

- `CONTEXT_LENGTH` – Maximum total token length (input + output) per request
- `MAX_CONCURRENCY` – Maximum number of concurrent requests the endpoint will serve

The following table lists the supported Amazon Nova models, instance types, and supported configurations. The `MAX_CONCURRENCY` values represent the maximum supported concurrency for each `CONTEXT_LENGTH` setting:

| Model             | Instance Type  | Supported Configurations                                                                        | FP8 Quantization Required |
|-------------------|----------------|-------------------------------------------------------------------------------------------------|---------------------------|
| Amazon Nova Micro | ml.g5.12xlarge | CONTEXT_LENGTH: 4000,<br>MAX_CONCURRENCY: 12<br><br>CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 6 | No                        |
|                   | ml.g5.24xlarge | CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 8                                                     | No                        |



| Model            | Instance Type  | Supported Configurations                                                                                                                                 | FP8 Quantization Required |
|------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|
|                  | ml.g6e.xlarge  | CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 2                                                                                                              | No                        |
|                  | ml.g6e.2xlarge | CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 2                                                                                                              | No                        |
|                  | ml.g6e.4xlarge | CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 4                                                                                                              | No                        |
|                  | ml.g6.12xlarge | CONTEXT_LENGTH: 4000,<br>MAX_CONCURRENCY: 12<br><br>CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 6                                                          | No                        |
|                  | ml.g6.24xlarge | CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 8                                                                                                              | No                        |
|                  | ml.g6.48xlarge | CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 12                                                                                                             | No                        |
|                  | ml.p5.48xlarge | CONTEXT_LENGTH: 16000,<br>MAX_CONCURRENCY: 128<br><br>CONTEXT_LENGTH: 64000,<br>MAX_CONCURRENCY: 32<br><br>CONTEXT_LENGTH: 128000,<br>MAX_CONCURRENCY: 8 | No                        |
| Amazon Nova Lite | ml.g6.12xlarge | CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 2                                                                                                              | Yes - Enabled By Default  |
|                  | ml.g6.24xlarge | CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 4                                                                                                              | Yes - Enabled By Default  |

| Model       | Instance Type  | Supported Configurations                                                                                                                                 | FP8 Quantization Required |
|-------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|
|             | ml.g6.48xlarge | CONTEXT_LENGTH: 4000,<br>MAX_CONCURRENCY: 16<br><br>CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 8                                                          | No                        |
|             | ml.p5.48xlarge | CONTEXT_LENGTH: 16000,<br>MAX_CONCURRENCY: 128<br><br>CONTEXT_LENGTH: 60000,<br>MAX_CONCURRENCY: 8                                                       | No                        |
| Nova 2 Lite | ml.g6.48xlarge | CONTEXT_LENGTH: 8000,<br>MAX_CONCURRENCY: 8                                                                                                              | Yes - Enabled By Default  |
|             | ml.p5.48xlarge | CONTEXT_LENGTH: 16000,<br>MAX_CONCURRENCY: 128<br><br>CONTEXT_LENGTH: 64000,<br>MAX_CONCURRENCY: 32<br><br>CONTEXT_LENGTH: 128000,<br>MAX_CONCURRENCY: 8 | No                        |

### Note

For instances where FP8 quantization is required, it will be enabled by default. The MAX\_CONCURRENCY values shown are upper bounds for each CONTEXT\_LENGTH setting. You can use lower context lengths with the same concurrency, but exceeding these values will cause SageMaker endpoint creation to fail.

For example, on Amazon Nova Micro with a ml.g5.12xlarge:

- CONTEXT\_LENGTH=2000, MAX\_CONCURRENCY=12 → Valid

- CONTEXT\_LENGTH=8000, MAX\_CONCURRENCY=12 → Rejected (concurrency limit is 6 at context length 8000)
- CONTEXT\_LENGTH=8000, MAX\_CONCURRENCY=4 → Valid
- CONTEXT\_LENGTH=8000, MAX\_CONCURRENCY=6 → Valid
- CONTEXT\_LENGTH=10000 → Rejected (max context length is 8000 on this instance)

## Supported AWS Regions

The following table lists the AWS Regions where Amazon Nova models are available on SageMaker inference:

| Region Name           | Region Code | Availability |
|-----------------------|-------------|--------------|
| US East (N. Virginia) | us-east-1   | Available    |
| US West (Oregon)      | us-west-2   | Available    |

## Supported Container Images

The following table lists the container image URLs for Amazon Nova models on SageMaker inference by region.

| Region    | Container Image URIs                                                                     |
|-----------|------------------------------------------------------------------------------------------|
| us-east-1 | 708977205387.dkr.ecr.us-east-1.amazonaws.com/<br>nova-inference-repo:SM-Inference-latest |
| us-west-2 | 176779409107.dkr.ecr.us-west-2.amazonaws.com/<br>nova-inference-repo:SM-Inference-latest |

## Best Practices

For best practices on deploying and managing models on SageMaker, see [Best Practices for SageMaker](#).

# Support

For issues and support with Amazon Nova models on SageMaker inference, contact AWS Support through the Console or your AWS account manager.

## Topics

- [Getting Started](#)
- [API Reference](#)
- [Evaluate Models Hosted on SageMaker Inference](#)
- [Deployment of Amazon Nova Forge Models in Amazon SageMaker Inference abuse detection](#)

## Getting Started

This guide shows you how to deploy customized Amazon Nova models on SageMaker real-time endpoints, configure inference parameters, and invoke your models for testing.

## Prerequisites

The following are prerequisites to deploy Amazon Nova models on SageMaker inference:

- Create an AWS account - If you don't have one already, see [Creating an AWS account](#).
- Required IAM permissions - Ensure your IAM user or role has the following managed policies attached:
  - AmazonSageMakerFullAccess
  - AmazonS3FullAccess
- Required SDKs/CLI versions - The following SDK versions have been tested and validated with Amazon Nova models on SageMaker inference:
  - SageMaker Python SDK v3.0.0+ (`sagemaker>=3.0.0`) for resource-based API approach
  - Boto3 version 1.35.0+ (`boto3>=1.35.0`) for direct API calls. The examples in this guide use this approach.
- Service quota increase - Request an Amazon SageMaker service quota increase for the ML instance type you plan to use for your SageMaker Inference endpoint (for example, `m1.p5.48xlarge` for endpoint usage). For a list of supported instance types, see [the section called "Supported models and instances"](#). To request an increase, see [Requesting a quota increase](#). For information about SageMaker instance quotas, see [SageMaker endpoints and quotas](#).

## Step 1: Configure AWS credentials

Configure your AWS credentials using one of the following methods:

### Option 1: AWS CLI (Recommended)

```
aws configure
```

Enter your AWS access key, secret key, and default region when prompted.

### Option 2: AWS credentials file

Create or edit `~/.aws/credentials`:

```
[default]
aws_access_key_id = YOUR_ACCESS_KEY
aws_secret_access_key = YOUR_SECRET_KEY
```

### Option 3: Environment variables

```
export AWS_ACCESS_KEY_ID=your_access_key
export AWS_SECRET_ACCESS_KEY=your_secret_key
```

#### Note

For more information about AWS credentials, see [Configuration and credential file settings](#).

## Initialize AWS clients

Create a Python script or notebook with the following code to initialize the AWS SDK and verify your credentials:

```
import boto3

AWS Configuration - Update these for your environment
REGION = "us-east-1" # Supported regions: us-east-1, us-west-2
AWS_ACCOUNT_ID = "YOUR_ACCOUNT_ID" # Replace with your AWS account ID

Initialize AWS clients using default credential chain
sagemaker = boto3.client('sagemaker', region_name=REGION)
```

```
sts = boto3.client('sts')

Verify credentials
try:
 identity = sts.get_caller_identity()
 print(f"Successfully authenticated to AWS Account: {identity['Account']}")

 if identity['Account'] != AWS_ACCOUNT_ID:
 print(f"Warning: Connected to account {identity['Account']}, expected {AWS_ACCOUNT_ID}")

except Exception as e:
 print(f"Failed to authenticate: {e}")
 print("Please verify your credentials are configured correctly.")
```

If the authentication is successful, you should see output confirming your AWS account ID.

## Step 2: Create a SageMaker execution role

A SageMaker execution role is an IAM role that grants SageMaker permissions to access AWS resources on your behalf, such as Amazon S3 buckets for model artifacts and CloudWatch for logging.

### Creating the execution role

#### Note

Creating IAM roles requires `iam:CreateRole` and `iam:AttachRolePolicy` permissions. Ensure your IAM user or role has these permissions before proceeding.

The following code creates an IAM role with the necessary permissions for deploying Amazon Nova customized models:

```
import json

Create SageMaker Execution Role
role_name = f"SageMakerInference-ExecutionRole-{AWS_ACCOUNT_ID}"

trust_policy = {
 "Version": "2012-10-17",
 "Statement": [
```

```

 {
 "Effect": "Allow",
 "Principal": {"Service": "sagemaker.amazonaws.com"},
 "Action": "sts:AssumeRole"
 }
]
}

iam = boto3.client('iam', region_name=REGION)

Create the role
role_response = iam.create_role(
 RoleName=role_name,
 AssumeRolePolicyDocument=json.dumps(trust_policy),
 Description='SageMaker execution role with S3 and SageMaker access'
)

Attach required policies
iam.attach_role_policy(
 RoleName=role_name,
 PolicyArn='arn:aws:iam::aws:policy/AmazonSageMakerFullAccess'
)

iam.attach_role_policy(
 RoleName=role_name,
 PolicyArn='arn:aws:iam::aws:policy/AmazonS3FullAccess'
)

SAGEMAKER_EXECUTION_ROLE_ARN = role_response['Role']['Arn']
print(f"Created SageMaker execution role: {SAGEMAKER_EXECUTION_ROLE_ARN}")

```

## Using an existing execution role (Optional)

If you already have a SageMaker execution role, you can use it instead:

```

Replace with your existing role ARN
SAGEMAKER_EXECUTION_ROLE_ARN = "arn:aws:iam::YOUR_ACCOUNT_ID:role/
YOUR_EXISTING_ROLE_NAME"

```

To find existing SageMaker roles in your account:

```

iam = boto3.client('iam', region_name=REGION)
response = iam.list_roles()

```

```
sagemaker_roles = [role for role in response['Roles'] if 'SageMaker' in
 role['RoleName']]
for role in sagemaker_roles:
 print(f"{role['RoleName']}: {role['Arn']}")
```

### Important

The execution role must have trust relationship with `sagemaker.amazonaws.com` and permissions to access Amazon S3 and SageMaker resources.

For more information about SageMaker execution roles, see [SageMaker Roles](#).

## Step 3: Configure model parameters

Configure the deployment parameters for your Amazon Nova model. These settings control model behavior, resource allocation, and inference characteristics. For a list of supported instance types and supported `CONTEXT_LENGTH` and `MAX_CONCURRENCY` values for each, see [the section called "Supported models and instances"](#).

### Required parameters

- **IMAGE:** The Docker container image URI for Amazon Nova inference container. This will be provided by AWS.
- **CONTEXT\_LENGTH:** Model context length.
- **MAX\_CONCURRENCY:** Maximum number of sequences per iteration; sets the limit on how many individual user requests (prompts) can be processed concurrently within a single batch on the GPU. Range: integer greater than 0.

### Optional generation parameters

- **DEFAULT\_TEMPERATURE:** Controls randomness in generation. Range: 0.0 to 2.0 (0.0 = deterministic, higher = more random).
- **DEFAULT\_TOP\_P:** Nucleus sampling threshold for token selection. Range: 1e-10 to 1.0.
- **DEFAULT\_TOP\_K:** Limits token selection to top K most likely tokens. Range: integer -1 or greater (-1 = no limit).
- **DEFAULT\_MAX\_NEW\_TOKENS:** Maximum number of tokens to generate in response (i.e. max output tokens). Range: integer 1 or greater.



- **DEFAULT\_LOGPROBS:** Number of log probabilities to return per token. Range: integer 1 to 20.

## Configure your deployment

```
AWS Configuration
REGION = "us-east-1" # Must match region from Step 1

ECR Account mapping by region
ECR_ACCOUNT_MAP = {
 "us-east-1": "708977205387",
 "us-west-2": "176779409107"
}

Container Image
IMAGE = f"{ECR_ACCOUNT_MAP[REGION]}.dkr.ecr.{REGION}.amazonaws.com/nova-inference-
repo:SM-Inference-latest"
print(f"IMAGE = {IMAGE}")

Model Parameters
CONTEXT_LENGTH = "16000" # Maximum total context length
MAX_CONCURRENCY = "2" # Maximum concurrent sequences

Optional: Default generation parameters (uncomment to use)
DEFAULT_TEMPERATURE = "0.0" # Deterministic output
DEFAULT_TOP_P = "1.0" # Consider all tokens
DEFAULT_TOP_K = "50" # Uncomment to limit to top 50 tokens
DEFAULT_MAX_NEW_TOKENS = "2048" # Uncomment to set max output tokens
DEFAULT_LOGPROBS = "1" # Uncomment to enable log probabilities

Build environment variables for the container
environment = {
 'CONTEXT_LENGTH': CONTEXT_LENGTH,
 'MAX_CONCURRENCY': MAX_CONCURRENCY,
}

Add optional parameters if defined
if 'DEFAULT_TEMPERATURE' in globals():
 environment['DEFAULT_TEMPERATURE'] = DEFAULT_TEMPERATURE
if 'DEFAULT_TOP_P' in globals():
 environment['DEFAULT_TOP_P'] = DEFAULT_TOP_P
if 'DEFAULT_TOP_K' in globals():
 environment['DEFAULT_TOP_K'] = DEFAULT_TOP_K
if 'DEFAULT_MAX_NEW_TOKENS' in globals():
```

```
environment['DEFAULT_MAX_NEW_TOKENS'] = DEFAULT_MAX_NEW_TOKENS
if 'DEFAULT_LOGPROBS' in globals():
 environment['DEFAULT_LOGPROBS'] = DEFAULT_LOGPROBS

print("Environment configuration:")
for key, value in environment.items():
 print(f" {key}: {value}")
```

## Configure deployment-specific parameters

Now configure the specific parameters for your Amazon Nova model deployment, including model artifacts location and instance type selection.

### Set deployment identifier

```
Deployment identifier - use a descriptive name for your use case
JOB_NAME = "my-nova-deployment"
```

### Specify model artifacts location

Provide the Amazon S3 URI where your trained Amazon Nova model artifacts are stored. This should be the output location from your model training or fine-tuning job.

```
S3 location of your trained Nova model artifacts
Replace with your model's S3 URI - must end with /
MODEL_S3_LOCATION = "s3://your-bucket-name/path/to/model/artifacts/"
```

### Select model variant and instance type

```
Configure model variant and instance type
TESTCASE = {
 "model": "lite2", # Options: micro, lite, lite2
 "instance": "ml.p5.48xlarge" # Refer to "Supported models and instances" section
}

Generate resource names
INSTANCE_TYPE = TESTCASE["instance"]
MODEL_NAME = JOB_NAME + "-" + TESTCASE["model"] + "-" + INSTANCE_TYPE.replace(".", "-")
ENDPOINT_CONFIG_NAME = MODEL_NAME + "-Config"
ENDPOINT_NAME = MODEL_NAME + "-Endpoint"
```

```
print(f"Model Name: {MODEL_NAME}")
print(f"Endpoint Config: {ENDPOINT_CONFIG_NAME}")
print(f"Endpoint Name: {ENDPOINT_NAME}")
```

## Naming conventions

The code automatically generates consistent names for AWS resources:

- Model Name: {JOB\_NAME}-{model}-{instance-type}
- Endpoint Config: {MODEL\_NAME}-Config
- Endpoint Name: {MODEL\_NAME}-Endpoint

## Step 4: Create SageMaker model and endpoint configuration

In this step, you'll create two essential resources: a SageMaker model object that references your Amazon Nova model artifacts, and an endpoint configuration that defines how the model will be deployed.

**SageMaker Model:** A model object that packages the inference container image, model artifacts location, and environment configuration. This is a reusable resource that can be deployed to multiple endpoints.

**Endpoint Configuration:** Defines the infrastructure settings for deployment, including instance type, instance count, and model variants. This allows you to manage deployment settings separately from the model itself.

### Create the SageMaker model

The following code creates a SageMaker model that references your Amazon Nova model artifacts:

```
try:
 model_response = sagemaker.create_model(
 ModelName=MODEL_NAME,
 PrimaryContainer={
 'Image': IMAGE,
 'ModelDataSource': {
 'S3DataSource': {
 'S3Uri': MODEL_S3_LOCATION,
 'S3DataType': 'S3Prefix',
 'CompressionType': 'None'
```

```

 }
 },
 'Environment': environment
},
ExecutionRoleArn=SAGEMAKER_EXECUTION_ROLE_ARN,
EnableNetworkIsolation=True
)
print("Model created successfully!")
print(f"Model ARN: {model_response['ModelArn']}")

except sagemaker.exceptions.ClientError as e:
 print(f"Error creating model: {e}")

```

### Key parameters:

- **ModelName:** Unique identifier for your model
- **Image:** Docker container image URI for Amazon Nova inference
- **ModelDataSource:** Amazon S3 location of your model artifacts
- **Environment:** Environment variables configured in Step 3
- **ExecutionRoleArn:** IAM role from Step 2
- **EnableNetworkIsolation:** Set to True for enhanced security (prevents container from making outbound network calls)

## Create the endpoint configuration

Next, create an endpoint configuration that defines your deployment infrastructure:

```

Create Endpoint Configuration
try:
 production_variant = {
 'VariantName': 'primary',
 'ModelName': MODEL_NAME,
 'InitialInstanceCount': 1,
 'InstanceType': INSTANCE_TYPE,
 }

 config_response = sagemaker.create_endpoint_config(
 EndpointConfigName=ENDPOINT_CONFIG_NAME,
 ProductionVariants=[production_variant]
)

```

```
print("Endpoint configuration created successfully!")
print(f"Config ARN: {config_response['EndpointConfigArn']}")

except sagemaker.exceptions.ClientError as e:
 print(f"Error creating endpoint configuration: {e}")
```

Key parameters:

- **VariantName**: Identifier for this model variant (use 'primary' for single-model deployments)
- **ModelName**: References the model created above
- **InitialInstanceCount**: Number of instances to deploy (start with 1, scale later if needed)
- **InstanceType**: ML instance type selected in Step 3

## Verify resource creation

You can verify that your resources were created successfully:

```
Describe the model
model_info = sagemaker.describe_model(ModelName=MODEL_NAME)
print(f"Model Status: {model_info['ModelName']} created")

Describe the endpoint configuration
config_info =
 sagemaker.describe_endpoint_config(EndpointConfigName=ENDPOINT_CONFIG_NAME)
print(f"Endpoint Config Status: {config_info['EndpointConfigName']} created")
```

## Step 5: Deploy the endpoint

The next step is to deploy your Amazon Nova model by creating a SageMaker real-time endpoint. This endpoint will host your model and provide a secure HTTPS endpoint for making inference requests.

Endpoint creation typically takes 15-30 minutes as AWS provisions the infrastructure, downloads your model artifacts, and initializes the inference container.

### Create the endpoint

```
import time

try:
```

```

endpoint_response = sagemaker.create_endpoint(
 EndpointName=ENDPOINT_NAME,
 EndpointConfigName=ENDPOINT_CONFIG_NAME
)
print("Endpoint creation initiated successfully!")
print(f"Endpoint ARN: {endpoint_response['EndpointArn']}")
except Exception as e:
 print(f"Error creating endpoint: {e}")

```

## Monitor endpoint creation

The following code polls the endpoint status until deployment is complete:

```

Monitor endpoint creation progress
print("Waiting for endpoint creation to complete...")
print("This typically takes 15-30 minutes...\n")

while True:
 try:
 response = sagemaker.describe_endpoint(EndpointName=ENDPOINT_NAME)
 status = response['EndpointStatus']

 if status == 'Creating':
 print(f"# Status: {status} - Provisioning infrastructure and loading model...")
 elif status == 'InService':
 print(f"# Status: {status}")
 print("\nEndpoint creation completed successfully!")
 print(f"Endpoint Name: {ENDPOINT_NAME}")
 print(f"Endpoint ARN: {response['EndpointArn']}")
 break
 elif status == 'Failed':
 print(f"# Status: {status}")
 print(f"Failure Reason: {response.get('FailureReason', 'Unknown')}")
 print("\nFull response:")
 print(response)
 break
 else:
 print(f"Status: {status}")

 except Exception as e:
 print(f"Error checking endpoint status: {e}")
 break

```

```
time.sleep(30) # Check every 30 seconds
```

## Verify endpoint is ready

Once the endpoint is InService, you can verify its configuration:

```
Get detailed endpoint information
endpoint_info = sagemaker.describe_endpoint(EndpointName=ENDPOINT_NAME)

print("\n=== Endpoint Details ===")
print(f"Endpoint Name: {endpoint_info['EndpointName']}")
print(f"Endpoint ARN: {endpoint_info['EndpointArn']}")
print(f"Status: {endpoint_info['EndpointStatus']}")
print(f"Creation Time: {endpoint_info['CreationTime']}")
print(f>Last Modified: {endpoint_info['LastModifiedTime']}")

Get endpoint config for instance type details
endpoint_config_name = endpoint_info['EndpointConfigName']
endpoint_config =
 sagemaker.describe_endpoint_config(EndpointConfigName=endpoint_config_name)

Display production variant details
for variant in endpoint_info['ProductionVariants']:
 print(f"\nProduction Variant: {variant['VariantName']}")
 print(f" Current Instance Count: {variant['CurrentInstanceCount']}")
 print(f" Desired Instance Count: {variant['DesiredInstanceCount']}")
 # Get instance type from endpoint config
 for config_variant in endpoint_config['ProductionVariants']:
 if config_variant['VariantName'] == variant['VariantName']:
 print(f" Instance Type: {config_variant['InstanceType']}")
 break
```

## Troubleshooting endpoint creation failures

Common failure reasons:

- **Insufficient capacity:** The requested instance type is not available in your region
  - Solution: Try a different instance type or request a quota increase
- **IAM permissions:** The execution role lacks necessary permissions
  - Solution: Verify the role has access to Amazon S3 model artifacts and necessary SageMaker permissions

- **Model artifacts not found:** The Amazon S3 URI is incorrect or inaccessible
  - Solution: Verify the Amazon S3 URI and check bucket permissions, make sure you're in the correct region
- **Resource limits:** Account limits exceeded for endpoints or instances
  - Solution: Request a service quota increase through Service Quotas or AWS Support

### Note

If you need to delete a failed endpoint and start over:

```
sagemaker.delete_endpoint(EndpointName=ENDPOINT_NAME)
```

## Step 6: Invoke the endpoint

Once your endpoint is InService, you can send inference requests to generate predictions from your Amazon Nova model. SageMaker supports synchronous endpoints (real-time with streaming/non-streaming modes) and asynchronous endpoints (Amazon S3-based for batch processing).

### Set up the runtime client

Create a SageMaker Runtime client with appropriate timeout settings:

```
import json
import boto3
import botocore
from botocore.exceptions import ClientError

Configure client with appropriate timeouts
config = botocore.config.Config(
 read_timeout=120, # Maximum time to wait for response
 connect_timeout=10, # Maximum time to establish connection
 retries={'max_attempts': 3} # Number of retry attempts
)

Create SageMaker Runtime client
runtime_client = boto3.client('sagemaker-runtime', config=config, region_name=REGION)
```

### Create a universal inference function



The following function handles both streaming and non-streaming requests:

```
def invoke_nova_endpoint(request_body):
 """
 Invoke Nova endpoint with automatic streaming detection.

 Args:
 request_body (dict): Request payload containing prompt and parameters

 Returns:
 dict: Response from the model (for non-streaming requests)
 None: For streaming requests (prints output directly)
 """
 body = json.dumps(request_body)
 is_streaming = request_body.get("stream", False)

 try:
 print(f"Invoking endpoint ({'streaming' if is_streaming else 'non-streaming'})...")

 if is_streaming:
 response = runtime_client.invoke_endpoint_with_response_stream(
 EndpointName=ENDPOINT_NAME,
 ContentType='application/json',
 Body=body
)

 event_stream = response['Body']
 for event in event_stream:
 if 'PayloadPart' in event:
 chunk = event['PayloadPart']
 if 'Bytes' in chunk:
 data = chunk['Bytes'].decode()
 print("Chunk:", data)
 else:
 # Non-streaming inference
 response = runtime_client.invoke_endpoint(
 EndpointName=ENDPOINT_NAME,
 ContentType='application/json',
 Accept='application/json',
 Body=body
)

 response_body = response['Body'].read().decode('utf-8')
```

```

 result = json.loads(response_body)
 print("# Response received successfully")
 return result

except ClientError as e:
 error_code = e.response['Error']['Code']
 error_message = e.response['Error']['Message']
 print(f"# AWS Error: {error_code} - {error_message}")
except Exception as e:
 print(f"# Unexpected error: {str(e)}")

```

## Example 1: Non-streaming chat completion

Use the chat format for conversational interactions:

```

Non-streaming chat request
chat_request = {
 "messages": [
 {"role": "user", "content": "Hello! How are you?"}
],
 "max_tokens": 100,
 "max_completion_tokens": 100, # Alternative to max_tokens
 "stream": False,
 "temperature": 0.7,
 "top_p": 0.9,
 "top_k": 50,
 "logprobs": True,
 "top_logprobs": 3,
 "reasoning_effort": "low", # Options: "low", "high"
 "allowed_token_ids": None, # List of allowed token IDs
 "truncate_prompt_tokens": None, # Truncate prompt to this many tokens
 "stream_options": None
}

response = invoke_nova_endpoint(chat_request)

```

## Sample response:

```

{
 "id": "chatcmpl-123456",
 "object": "chat.completion",
 "created": 1234567890,
 "model": "default",

```

```

 "choices": [
 {
 "index": 0,
 "message": {
 "role": "assistant",
 "content": "Hello! I'm doing well, thank you for asking. I'm here and
ready to help you with any questions or tasks you might have. How can I assist you
today?"
 },
 "logprobs": {
 "content": [
 {
 "token": "Hello",
 "logprob": -0.123,
 "top_logprobs": [
 {"token": "Hello", "logprob": -0.123},
 {"token": "Hi", "logprob": -2.456},
 {"token": "Hey", "logprob": -3.789}
]
 }
 # Additional tokens...
]
 },
 "finish_reason": "stop"
 }
],
 "usage": {
 "prompt_tokens": 12,
 "completion_tokens": 28,
 "total_tokens": 40
 }
 }
}

```

## Example 2: Simple text completion

Use the completion format for simple text generation:

```

Simple completion request
completion_request = {
 "prompt": "The capital of France is",
 "max_tokens": 50,
 "stream": False,
 "temperature": 0.0,
 "top_p": 1.0,

```

```

"top_k": -1, # -1 means no limit
"logprobs": 3, # Number of log probabilities to return
"allowed_token_ids": None, # List of allowed token IDs
"truncate_prompt_tokens": None, # Truncate prompt to this many tokens
"stream_options": None
}

response = invoke_nova_endpoint(completion_request)

```

### Sample response:

```

{
 "id": "cmpl-789012",
 "object": "text_completion",
 "created": 1234567890,
 "model": "default",
 "choices": [
 {
 "text": " Paris.",
 "index": 0,
 "logprobs": {
 "tokens": [" Paris", "."],
 "token_logprobs": [-0.001, -0.002],
 "top_logprobs": [
 {" Paris": -0.001, " London": -5.234, " Rome": -6.789},
 {".": -0.002, ",": -4.567, " ": -7.890}
]
 },
 "finish_reason": "stop"
 }
],
 "usage": {
 "prompt_tokens": 6,
 "completion_tokens": 2,
 "total_tokens": 8
 }
}

```

### Example 3: Streaming chat completion

```

Streaming chat request
streaming_request = {
 "messages": [

```

```

 {"role": "user", "content": "Tell me a short story about a robot"}
],
 "max_tokens": 200,
 "stream": True,
 "temperature": 0.7,
 "top_p": 0.95,
 "top_k": 40,
 "logprobs": True,
 "top_logprobs": 2,
 "reasoning_effort": "high", # For more detailed reasoning
 "stream_options": {"include_usage": True}
}

invoke_nova_endpoint(streaming_request)

```

### Sample streaming output:

```

Chunk: data: {"id":"chatcmpl-029ca032-fa01-4868-80b7-
c4cb1af90fb9","object":"chat.completion.chunk","created":1772060532,"model":"default","choices":
[{"index":0,"delta":
{"role":"assistant","content":"","reasoning_content":null},"logprobs":null,"finish_reason":null}]}
Chunk: data: {"id":"chatcmpl-029ca032-fa01-4868-80b7-
c4cb1af90fb9","object":"chat.completion.chunk","created":1772060532,"model":"default","choices":
[{"index":0,"delta":{"content":" Once","reasoning_content":null},"logprobs":
{"content":[{"token":"\u2581Once","logprob":-0.6078429222106934,"bytes":
[226,150,129,79,110,99,101],"top_logprobs":
[{"token":"\u2581Once","logprob":-0.6078429222106934,"bytes":
[226,150,129,79,110,99,101]}]}]}]}]}]},"finish_reason":null,"token_ids":null}}]}
Chunk: data: {"id":"chatcmpl-029ca032-fa01-4868-80b7-
c4cb1af90fb9","object":"chat.completion.chunk","created":1772060532,"model":"default","choices":
[{"index":0,"delta":{"content":" upon","reasoning_content":null},"logprobs":
{"content":[{"token":"\u2581upon","logprob":-0.0012345,"bytes":
[226,150,129,117,112,111,110],"top_logprobs":
[{"token":"\u2581upon","logprob":-0.0012345,"bytes":[226,150,129,117,112,111,110]}],
{"token":"\u2581a","logprob":-6.789,"bytes":
[226,150,129,97]}]}]}]}]},"finish_reason":null,"token_ids":null}}]}
Chunk: data: {"id":"chatcmpl-029ca032-fa01-4868-80b7-
c4cb1af90fb9","object":"chat.completion.chunk","created":1772060532,"model":"default","choices":
[{"index":0,"delta":{"content":" a","reasoning_content":null},"logprobs":
{"content":[{"token":"\u2581a","logprob":-0.0001234,"bytes":
[226,150,129,97],"top_logprobs":[{"token":"\u2581a","logprob":-0.0001234,"bytes":

```

```
[226,150,129,97]],{"token":"\u2581time","logprob":-9.123,"bytes":
[226,150,129,116,105,109,101]]]]}],{"finish_reason":null,"token_ids":null}}]
Chunk: data: {"id":"chatcmpl-029ca032-fa01-4868-80b7-
c4cb1af90fb9","object":"chat.completion.chunk","created":1772060532,"model":"default","choices":
[{"index":0,"delta":{"content":" time","reasoning_content":null},"logprobs":
{"content":[{"token":"\u2581time","logprob":-0.0023456,"bytes":
[226,150,129,116,105,109,101],"top_logprobs":
[{"token":"\u2581time","logprob":-0.0023456,"bytes":[226,150,129,116,105,109,101]},
{"token":"","logprob":-6.012,"bytes":
[44]]}]]}],{"finish_reason":null,"token_ids":null}}]

Additional chunks...

Chunk: data: {"id":"chatcmpl-029ca032-fa01-4868-80b7-
c4cb1af90fb9","object":"chat.completion.chunk","created":1772060532,"model":"default","choices":
[{"index":0,"delta":{"},"logprobs":null,"finish_reason":"stop"}],"usage":
{"prompt_tokens":15,"completion_tokens":87,"total_tokens":102}}
Chunk: data: [DONE]
```

## Example 4: Multimodal chat completion

Use multimodal format for image and text inputs:

```
Multimodal chat request (if supported by your model)
multimodal_request = {
 "messages": [
 {
 "role": "user",
 "content": [
 {"type": "text", "text": "What's in this image?"},
 {"type": "image_url", "image_url": {"url": "data:image/
jpeg;base64,..."}}
]
 }
],
 "max_tokens": 150,
 "temperature": 0.3,
 "top_p": 0.8,
 "stream": False
}

response = invoke_nova_endpoint(multimodal_request)
```

## Sample response:

```
{
 "id": "chatcmpl-345678",
 "object": "chat.completion",
 "created": 1234567890,
 "model": "default",
 "choices": [
 {
 "index": 0,
 "message": {
 "role": "assistant",
 "content": "The image shows..."
 },
 "finish_reason": "stop"
 }
],
 "usage": {
 "prompt_tokens": 1250,
 "completion_tokens": 45,
 "total_tokens": 1295
 }
}
```

## Step 7: Clean up resources (Optional)

To avoid incurring unnecessary charges, delete the AWS resources you created during this tutorial. SageMaker endpoints incur charges while they're running, even if you're not actively making inference requests.

### Important

Deleting resources is permanent and cannot be undone. Ensure you no longer need these resources before proceeding.

## Delete the endpoint

```
import boto3

Initialize SageMaker client
sagemaker = boto3.client('sagemaker', region_name=REGION)
```

```
try:
 print("Deleting endpoint...")
 sagemaker.delete_endpoint(EndpointName=ENDPOINT_NAME)
 print(f"# Endpoint '{ENDPOINT_NAME}' deletion initiated")
 print("Charges will stop once deletion completes (typically 2-5 minutes)")
except Exception as e:
 print(f"# Error deleting endpoint: {e}")
```

### Note

The endpoint deletion is asynchronous. You can monitor the deletion status:

```
import time

print("Monitoring endpoint deletion...")
while True:
 try:
 response = sagemaker.describe_endpoint(EndpointName=ENDPOINT_NAME)
 status = response['EndpointStatus']
 print(f"Status: {status}")
 time.sleep(10)
 except sagemaker.exceptions.ClientError as e:
 if e.response['Error']['Code'] == 'ValidationException':
 print("# Endpoint successfully deleted")
 break
 else:
 print(f"Error: {e}")
 break
```

## Delete the endpoint configuration

After the endpoint is deleted, remove the endpoint configuration:

```
try:
 print("Deleting endpoint configuration...")
 sagemaker.delete_endpoint_config(EndpointConfigName=ENDPOINT_CONFIG_NAME)
 print(f"# Endpoint configuration '{ENDPOINT_CONFIG_NAME}' deleted")
except Exception as e:
 print(f"# Error deleting endpoint configuration: {e}")
```



## Delete the model

Remove the SageMaker model object:

```
try:
 print("Deleting model...")
 sagemaker.delete_model(ModelName=MODEL_NAME)
 print(f"# Model '{MODEL_NAME}' deleted")
except Exception as e:
 print(f"# Error deleting model: {e}")
```

## API Reference

Amazon Nova models on SageMaker use the standard SageMaker Runtime API for inference. For complete API documentation, see [Test a deployed model](#).

### Endpoint invocation

Amazon Nova models on SageMaker support two invocation methods:

- **Synchronous invocation:** Use the [InvokeEndpoint](#) API for real-time, non-streaming inference requests.
- **Streaming invocation:** Use the [InvokeEndpointWithResponseStream](#) API for real-time streaming inference requests.

### Request format

Amazon Nova models support two request formats:

#### Chat completion format

Use this format for conversational interactions:

```
{
 "messages": [
 {"role": "user", "content": "string"}
],
 "max_tokens": integer,
 "max_completion_tokens": integer,
 "stream": boolean,
 "temperature": float,
 "top_p": float,
```

```
"top_k": integer,
"logprobs": boolean,
"top_logprobs": integer,
"reasoning_effort": "low" | "high",
"allowed_token_ids": [integer],
"truncate_prompt_tokens": integer,
"stream_options": {
 "include_usage": boolean
}
}
```

## Text completion format

Use this format for simple text generation:

```
{
 "prompt": "string",
 "max_tokens": integer,
 "stream": boolean,
 "temperature": float,
 "top_p": float,
 "top_k": integer,
 "logprobs": integer,
 "allowed_token_ids": [integer],
 "truncate_prompt_tokens": integer,
 "stream_options": {
 "include_usage": boolean
 }
}
```

## Multimodal chat completion format

Use this format for image and text inputs:

```
{
 "messages": [
 {
 "role": "user",
 "content": [
 {"type": "text", "text": "What's in this image?"},
 {"type": "image_url", "image_url": {"url": "data:image/jpeg;base64,..."}}
]
 }
]
}
```

```
],
 "max_tokens": integer,
 "temperature": float,
 "top_p": float,
 "stream": boolean
}
```

## Request parameters

- **messages** (array): For chat completion format. Array of message objects with `role` and `content` fields. Content can be a string for text-only or an array for multimodal inputs.
- **prompt** (string): For text completion format. The input text to generate from.
- **max\_tokens** (integer): Maximum number of tokens to generate in the response. Range: 1 or greater.
- **max\_completion\_tokens** (integer): Alternative to `max_tokens` for chat completions. Maximum number of completion tokens to generate.
- **temperature** (float): Controls randomness in generation. Range: 0.0 to 2.0 (0.0 = deterministic, 2.0 = maximum randomness).
- **top\_p** (float): Nucleus sampling threshold. Range: 1e-10 to 1.0.
- **top\_k** (integer): Limits token selection to top K most likely tokens. Range: -1 or greater (-1 = no limit).
- **stream** (boolean): Whether to stream the response. Set to `true` for streaming, `false` for non-streaming.
- **logprobs** (boolean/integer): For chat completions, use `boolean`. For text completions, use `integer` for number of log probabilities to return. Range: 1 to 20.
- **top\_logprobs** (integer): Number of most likely tokens to return log probabilities for (chat completions only).
- **reasoning\_effort** (string): Level of reasoning effort. Options: "low", "high" (chat completions for Nova 2 Lite custom models only).
- **allowed\_token\_ids** (array): List of token IDs that are allowed to be generated. Restricts output to specified tokens.
- **truncate\_prompt\_tokens** (integer): Truncate the prompt to this many tokens if it exceeds the limit.
- **stream\_options** (object): Options for streaming responses. Contains `include_usage` `boolean` to include token usage in streaming responses.

## Response format

The response format depends on the invocation method and request type:

### Chat completion response (non-streaming)

For synchronous chat completion requests:

```
{
 "id": "chatcmpl-123e4567-e89b-12d3-a456-426614174000",
 "object": "chat.completion",
 "created": 1677652288,
 "model": "nova-micro-custom",
 "choices": [
 {
 "index": 0,
 "message": {
 "role": "assistant",
 "content": "Hello! I'm doing well, thank you for asking. How can I help you
today?",
 "refusal": null,
 "reasoning": null,
 "reasoning_content": null
 },
 "logprobs": {
 "content": [
 {
 "token": "Hello",
 "logprob": -0.31725305,
 "bytes": [72, 101, 108, 108, 111],
 "top_logprobs": [
 {
 "token": "Hello",
 "logprob": -0.31725305,
 "bytes": [72, 101, 108, 108, 111]
 },
 {
 "token": "Hi",
 "logprob": -1.3190403,
 "bytes": [72, 105]
 }
]
 }
]
 }
 }
]
}
```

```

 },
 "finish_reason": "stop",
 "stop_reason": null,
 "token_ids": [9906, 0, 358, 2157, 1049, 11, 1309, 345, 369, 6464, 13]
 }
],
"usage": {
 "prompt_tokens": 9,
 "completion_tokens": 12,
 "total_tokens": 21,
 "prompt_tokens_details": {
 "cached_tokens": 0
 }
},
"prompt_token_ids": [9906, 0, 358]
}

```

## Text completion response (non-streaming)

For synchronous text completion requests:

```

{
 "id": "cmpl-123e4567-e89b-12d3-a456-426614174000",
 "object": "text_completion",
 "created": 1677652288,
 "model": "nova-micro-custom",
 "choices": [
 {
 "index": 0,
 "text": "Paris, the capital and most populous city of France.",
 "logprobs": {
 "tokens": ["Paris", ",", " the", " capital"],
 "token_logprobs": [-0.31725305, -0.07918124, -0.12345678, -0.23456789],
 "top_logprobs": [
 {
 "Paris": -0.31725305,
 "London": -1.3190403,
 "Rome": -2.1234567
 },
 {
 ",": -0.07918124,
 " is": -1.2345678
 }
]
 }
 }
]
}

```

```

 },
 "finish_reason": "stop",
 "stop_reason": null,
 "prompt_token_ids": [464, 6864, 315, 4881, 374],
 "token_ids": [3915, 11, 279, 6864, 323, 1455, 95551, 3363, 315, 4881, 13]
 }
],
"usage": {
 "prompt_tokens": 5,
 "completion_tokens": 11,
 "total_tokens": 16,
 "prompt_tokens_details": {
 "cached_tokens": 0
 }
}
}

```

## Chat completion streaming response

For streaming chat completion requests, responses are sent as Server-Sent Events (SSE):

```

data: {
 "id": "chatcmpl-123e4567-e89b-12d3-a456-426614174000",
 "object": "chat.completion.chunk",
 "created": 1677652288,
 "model": "nova-micro-custom",
 "choices": [
 {
 "index": 0,
 "delta": {
 "role": "assistant",
 "content": "Hello",
 "refusal": null,
 "reasoning": null,
 "reasoning_content": null
 },
 "logprobs": {
 "content": [
 {
 "token": "Hello",
 "logprob": -0.31725305,
 "bytes": [72, 101, 108, 108, 111],
 "top_logprobs": [

```

```

 "token": "Hello",
 "logprob": -0.31725305,
 "bytes": [72, 101, 108, 108, 111]
 }
]
 },
 "finish_reason": null,
 "stop_reason": null
}
],
"usage": null,
"prompt_token_ids": null
}

data: {
 "id": "chatcmpl-123e4567-e89b-12d3-a456-426614174000",
 "object": "chat.completion.chunk",
 "created": 1677652288,
 "model": "nova-micro-custom",
 "choices": [
 {
 "index": 0,
 "delta": {
 "content": "! I'm"
 },
 "logprobs": null,
 "finish_reason": null,
 "stop_reason": null
 }
],
 "usage": null
}

data: {
 "id": "chatcmpl-123e4567-e89b-12d3-a456-426614174000",
 "object": "chat.completion.chunk",
 "created": 1677652288,
 "model": "nova-micro-custom",
 "choices": [
 {
 "index": 0,
 "delta": {},

```

```

 "finish_reason": "stop",
 "stop_reason": null
 }
],
"usage": {
 "prompt_tokens": 9,
 "completion_tokens": 12,
 "total_tokens": 21,
 "prompt_tokens_details": {
 "cached_tokens": 0
 }
}
}

data: [DONE]

```

## Text completion streaming response

For streaming text completion requests:

```

data: {
 "id": "cmpl-123e4567-e89b-12d3-a456-426614174000",
 "object": "text_completion",
 "created": 1677652288,
 "model": "nova-micro-custom",
 "choices": [
 {
 "index": 0,
 "text": "Paris",
 "logprobs": {
 "tokens": ["Paris"],
 "token_logprobs": [-0.31725305],
 "top_logprobs": [
 {
 "Paris": -0.31725305,
 "London": -1.3190403
 }
]
 },
 "finish_reason": null,
 "stop_reason": null
 }
],
 "usage": null
}

```



```
}

data: {
 "id": "cmpl-123e4567-e89b-12d3-a456-426614174000",
 "object": "text_completion",
 "created": 1677652288,
 "model": "nova-micro-custom",
 "choices": [
 {
 "index": 0,
 "text": ", the capital",
 "logprobs": null,
 "finish_reason": null,
 "stop_reason": null
 }
],
 "usage": null
}

data: {
 "id": "cmpl-123e4567-e89b-12d3-a456-426614174000",
 "object": "text_completion",
 "created": 1677652288,
 "model": "nova-micro-custom",
 "choices": [
 {
 "index": 0,
 "text": "",
 "finish_reason": "stop",
 "stop_reason": null
 }
],
 "usage": {
 "prompt_tokens": 5,
 "completion_tokens": 11,
 "total_tokens": 16
 }
}

data: [DONE]
```

## Response fields explanation

- **id**: Unique identifier for the completion
- **object**: Type of object returned ("chat.completion", "text\_completion", "chat.completion.chunk")
- **created**: Unix timestamp of when the completion was created
- **model**: Model used for the completion
- **choices**: Array of completion choices
- **usage**: Token usage information including prompt, completion, and total tokens
- **logprobs**: Log probability information for tokens (when requested)
- **finish\_reason**: Reason why the model stopped generating ("stop", "length", "content\_filter")
- **delta**: Incremental content in streaming responses
- **reasoning**: Reasoning content when reasoning\_effort is used
- **token\_ids**: Array of token IDs for the generated text

For complete API documentation, see [InvokeEndpoint API reference](#) and [InvokeEndpointWithResponseStream API reference](#).

## Evaluate Models Hosted on SageMaker Inference

This guide explains how to evaluate your customized Amazon Nova models deployed on SageMaker inference endpoints using [Inspect AI](#), an open-source evaluation framework.

### Note

For a hands-on walkthrough, see the [SageMaker Inspect AI quickstart notebook](#).

## Overview

You can evaluate your customized Amazon Nova models deployed on SageMaker endpoints using standardized benchmarks from the AI research community. This approach enables you to:

- Evaluate customized Amazon Nova models (fine-tuned, distilled, or otherwise adapted) at scale
- Run evaluations with parallel inference across multiple endpoint instances
- Compare model performance using benchmarks like MMLU, TruthfulQA, and HumanEval
- Integrate with your existing SageMaker infrastructure

## Supported models

The SageMaker inference provider works with:

- Amazon Nova models (Nova Micro, Nova Lite, Nova Lite 2)
- Models deployed via vLLM or OpenAI-compatible inference servers
- Any endpoint that supports the OpenAI Chat Completions API format

## Prerequisites

Before you begin, ensure you have:

- An AWS account with permissions to create and invoke SageMaker endpoints
- AWS credentials configured via AWS CLI, environment variables, or IAM role
- Python 3.9 or higher

## Required IAM permissions

Your IAM user or role needs the following permissions:

```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Effect": "Allow",
 "Action": [
 "sagemaker:InvokeEndpoint",
 "sagemaker:DescribeEndpoint"
],
 "Resource": "arn:aws:sagemaker:*:*:endpoint/*"
 }
]
}
```

## Step 1: Deploy a SageMaker endpoint

Before running evaluations, you need a SageMaker inference endpoint running your model.

For instructions on creating a SageMaker inference endpoint with Amazon Nova models, see [the section called “Getting started”](#).

Once your endpoint is in InService status, note the endpoint name for use in the evaluation commands.

## Step 2: Install evaluation dependencies

Create a Python virtual environment and install the required packages.

```
Create virtual environment
python3.12 -m venv venv
source venv/bin/activate

Install uv for faster package installation
pip install uv

Install Inspect AI and evaluation benchmarks
uv pip install inspect-ai inspect-vals

Install AWS dependencies
uv pip install aioboto3 boto3 botocore openai
```

## Step 3: Configure AWS credentials

Choose one of the following authentication methods:

### Option 1: AWS CLI (Recommended)

```
aws configure
```

Enter your AWS Access Key ID, Secret Access Key, and default region when prompted.

### Option 2: Environment variables

```
export AWS_ACCESS_KEY_ID=your_access_key
export AWS_SECRET_ACCESS_KEY=your_secret_key
export AWS_DEFAULT_REGION=us-west-2
```

### Option 3: IAM role

If running on Amazon EC2 or SageMaker notebooks, the instance's IAM role is used automatically.

### Verify credentials

```
import boto3

sts = boto3.client('sts')
identity = sts.get_caller_identity()
print(f"Account: {identity['Account']}")
print(f"User/Role: {identity['Arn']}")
```

## Step 4: Install the SageMaker provider

The SageMaker provider enables Inspect AI to communicate with your SageMaker endpoints. The provider installation process is streamlined in the [quickstart notebook](#).

## Step 5: Download evaluation benchmarks

Clone the Inspect Evals repository to access standard benchmarks:

```
git clone https://github.com/UKGovernmentBEIS/inspect_evals.git
```

This repository includes benchmarks such as:

- MMLU and MMLU-Pro (knowledge and reasoning)
- TruthfulQA (truthfulness)
- HumanEval (code generation)
- GSM8K (mathematical reasoning)

## Step 6: Run evaluations

Run an evaluation using your SageMaker endpoint:

```
cd inspect_evals/src/inspect_evals/

inspect eval mmlu_pro/mmlu_pro.py \
 --model sagemaker/my-nova-endpoint \
 -M region_name=us-west-2 \
 --max-connections 256 \
 --max-retries 100 \
 --display plain
```

### Key parameters

| Parameter          | Default   | Description                                                                                             |
|--------------------|-----------|---------------------------------------------------------------------------------------------------------|
| --max-connections  | 10        | Number of parallel requests to the endpoint. Scale with instance count (e.g., 10 instances × 25 = 250). |
| --max-retries      | 3         | Retry attempts for failed requests. Use 50-100 for large evaluations.                                   |
| -M region_name     | us-east-1 | AWS region where your endpoint is deployed.                                                             |
| -M read_timeout    | 600       | Request timeout in seconds.                                                                             |
| -M connect_timeout | 60        | Connection timeout in seconds.                                                                          |

## Tuning recommendations

For a multi-instance endpoint:

```
10-instance endpoint example
--max-connections 250 # ~25 connections per instance
--max-retries 100 # Handle transient errors
```

Setting `--max-connections` too high may overwhelm the endpoint and cause throttling. Setting it too low underutilizes capacity.

## Step 7: View results

Launch the Inspect AI viewer to analyze evaluation results:

```
inspect view
```

The viewer displays:

- Overall scores and metrics
- Per-sample results with model responses
- Error analysis and failure patterns

## Managing endpoints

### Update an endpoint

To update an existing endpoint with a new model or configuration:

```
import boto3

sagemaker = boto3.client('sagemaker', region_name=REGION)

Create new model and endpoint configuration
Then update the endpoint
sagemaker.update_endpoint(
 EndpointName=EXISTING_ENDPOINT_NAME,
 EndpointConfigName=NEW_ENDPOINT_CONFIG_NAME
)
```

### Delete an endpoint

```
sagemaker.delete_endpoint(EndpointName=ENDPOINT_NAME)
```

## Onboarding custom benchmarks

You can add new benchmarks to Inspect AI using the following workflow:

1. Study the benchmark's dataset format and evaluation metrics
2. Review similar implementations in `inspect_evals/`
3. Create a task file that converts dataset records to Inspect AI samples
4. Implement appropriate solvers and scorers
5. Validate with a small test run

Example task structure:

```
from inspect_ai import Task, task
from inspect_ai.dataset import hf_dataset
from inspect_ai.scorer import choice
from inspect_ai.solver import multiple_choice
```

```
@task
def my_benchmark():
 return Task(
 dataset=hf_dataset("dataset_name", split="test"),
 solver=multiple_choice(),
 scorer=choice()
)
```

## Troubleshooting

### Common issues

#### Endpoint throttling or timeouts

- Reduce `--max-connections`
- Increase `--max-retries`
- Check endpoint CloudWatch metrics for capacity issues

#### Authentication errors

- Verify AWS credentials are configured correctly
- Check IAM permissions include `sagemaker:InvokeEndpoint`

#### Model errors

- Verify the endpoint is in `InService` status
- Check that the model supports the OpenAI Chat Completions API format

### Related resources

- [Inspect AI Documentation](#)
- [Inspect Evals Repository](#)
- [SageMaker Developer Guide](#)
- [Deploy Models for Inference](#)
- [Configuring the AWS CLI](#)



# Deployment of Amazon Nova Forge Models in Amazon SageMaker Inference abuse detection

AWS is committed to the responsible use of AI. To help prevent potential misuse, when you deploy Amazon Nova Forge Models in Amazon SageMaker Inference, SageMaker Inference implements automated abuse detection mechanisms to identify potential violations of AWS's [Acceptable Use Policy](#) (AUP) and Service Terms, including the [Responsible AI Policy](#).

Our abuse detection mechanisms are fully automated, so there is no human review of, or access to, user inputs or model outputs.

Automated abuse detection includes:

- **Categorize content** – We use classifiers to detect harmful content (such as content that incites violence) in user inputs and model outputs. A classifier is an algorithm that processes model inputs and outputs, and assigns type of harm and level of confidence. We may run these classifiers on Amazon Nova Forge Model usage. The classification process is automated and does not involve human review of user inputs or model outputs.
- **Identify patterns** – We use classifier metrics to identify potential violations and recurring behavior. We may compile and share anonymized classifier metrics. Amazon SageMaker Inference does not store user input or model output.
- **Detecting and blocking child sexual abuse material (CSAM)** – You are responsible for the content you (and your end users) upload to Amazon SageMaker Inference and must ensure this content is free from illegal images. To help stop the dissemination of CSAM, when deploying an Amazon Nova Forge Model in Amazon SageMaker Inference, SageMaker Inference may use automated abuse detection mechanisms (such as hash matching technology or classifiers) to detect apparent CSAM. If Amazon SageMaker Inference detects apparent CSAM in your image inputs, Amazon SageMaker Inference will block the request and you will receive an automated error message. Amazon SageMaker Inference may also file a report with the National Center for Missing and Exploited Children (NCMEC) or a relevant authority. We take CSAM seriously and will continue to update our detection, blocking, and reporting mechanisms. You might be required by applicable laws to take additional actions, and you are responsible for those actions.

Once our automated abuse detection mechanisms identify potential violations, we may request information about your use of Amazon SageMaker Inference and compliance with our terms of service. In the event that you are non-responsive, unwilling, or unable to comply with these terms or policies, AWS may suspend your access to Amazon SageMaker Inference. You may also be billed

for the failed inference job if our automated tests detect model responses being inconsistent with our terms and policies.

Contact AWS Support if you have additional questions. For more information, see the [Amazon SageMaker FAQs](#).

## Amazon Bedrock inference

Once you've trained and tested your Amazon Nova model, you can deploy it to Amazon Bedrock for production-scale inference. The deployment process involves creating an Amazon Bedrock model with the `CreateCustomModel` API, exporting your model artifacts to it from a managed Amazon S3 bucket, and then once the model is `ACTIVE` configuring an endpoint with on-demand or provisioned-throughput inference.

After you create a custom model in SageMaker, you can use the `CreateCustomModel` API to deploy it to Amazon Bedrock from SageMaker escrow to run inference. You can then use `CreateCustomModelDeployment` to create an OD inference endpoint or set up provisioned throughput inference for a Parameter Efficient Fine Tuned (PEFT) model. You can set up provisioned throughput inference for a Full Rank custom model.

You can also use the Amazon Nova Forge SDK to deploy customized Amazon Nova models. The Amazon Nova Forge SDK provides a streamlined experience for extracting the relevant information from a training job or S3 model checkpoint and publishing it to Amazon Bedrock. For more information, see [Amazon Nova Forge SDK](#).

For detailed steps to set up Amazon Bedrock inference for a custom model, see [Deploying customized Amazon Nova models to Amazon Bedrock](#).

The following section gives more detail about On-Demand Inference on Custom Models.

### Deploy a custom model for on-demand inference

After you successfully create a custom model with a model customization job (fine-tuning, distillation, or continued pre-training), you can set up on-demand inference for the model.

To set up on-demand inference for a custom model, you deploy the model with a custom model deployment. After you deploy your custom model, you use the deployment's Amazon Resource Name (ARN) as the `modelId` parameter in your `InvokeModel` or `Converse` API operations. You

can use the deployed model for on-demand inference with Amazon Bedrock features such as playgrounds, Agents, and Knowledge Bases.

## Topics

- [Supported models](#)
- [Deploy a custom model](#)
- [Use a deployment for on-demand inference](#)
- [Delete a custom model deployment](#)

## Supported models

You can set up on-demand inference for the following models:

- Amazon Nova Canvas
- Amazon Nova Lite
- Amazon Nova Micro
- Amazon Nova Pro

## Deploy a custom model

You can deploy a custom model with the Amazon Bedrock console, AWS Command Line Interface, or AWS SDKs. For information about using the deployment for inference, see [Use a deployment for on-demand inference](#).

## Topics

- [Deploy a custom model \(console\)](#)
- [Deploy a custom model \(AWS Command Line Interface\)](#)
- [Deploy a custom model \(AWS SDKs\)](#)

## Deploy a custom model (console)

You deploy a custom model from the **Custom models** page as follows. You can also deploy a model from the **Custom model on-demand** page with the same fields. To find this page, in **Inference and Assessment** in the navigation pane, choose **Custom model on-demand**.

## To deploy a custom model

1. Sign in to the AWS Management Console using an [IAM role with Amazon Bedrock permissions](#), and open the Amazon Bedrock console at <https://console.aws.amazon.com/nova/>.
2. From the left navigation pane, choose **Custom models** under **Foundation models**.
3. In the **Models** tab, choose the radio button for the model you want to deploy.
4. Choose **Set up inference** and choose **Deploy for on-demand**.
5. In **Deployment details**, provide the following information:
  - **Deployment Name** (required) – Enter a unique name for your deployment.
  - **Description** (optional) – Enter a description for your deployment.
  - **Tags** (optional) – Add tags for cost allocation and resource management.
6. Choose **Create**. When the status shows **Completed**, your custom model is ready for on-demand inference. For more information about using the custom model, see [Use a deployment for on-demand inference](#).

## Deploy a custom model (AWS Command Line Interface)

To deploy a custom model for on-demand inference using the AWS Command Line Interface, use the `create-custom-model-deployment` command with your custom model's Amazon Resource Name (ARN). This command uses the [CreateCustomModelDeployment](#) API operation. It returns the deployment's ARN that you can use as the `modelId` when making inference requests. For information about using the deployment for inference, see [Use a deployment for on-demand inference](#).

```
aws bedrock create-custom-model-deployment \
--model-deployment-name "Unique name" \
--model-arn "Custom Model ARN" \
--description "Deployment description" \
--tags '[
 {
 "key": "Environment",
 "value": "Production"
 },
 {
 "key": "Team",
 "value": "ML-Engineering"
 },
]
```

```
{
 "key": "Project",
 "value": "CustomerSupport"
}
]' \
--client-request-token "unique-deployment-token" \
--region region
```

## Deploy a custom model (AWS SDKs)

To deploy a custom model for on-demand inference, use the [CreateCustomModelDeployment](#) API operation with your custom model's Amazon Resource Name (ARN). The response returns the deployment's ARN that you can use as the `modelId` when making inference requests. For information about using the deployment for inference, see [Use a deployment for on-demand inference](#).

The following code shows how to use the SDK for Python (Boto3) to deploy a custom model.

```
def create_custom_model_deployment(bedrock_client):
 """Create a custom model deployment
 Args:
 bedrock_client: A boto3 Bedrock client for making API calls

 Returns:
 str: The ARN of the created custom model deployment

 Raises:
 Exception: If there is an error creating the deployment
 """

 try:
 response = bedrock_client.create_custom_model_deployment(
 modelDeploymentName="Unique deployment name",
 modelArn="Custom Model ARN",
 description="Deployment description",
 tags=[
 {'key': 'Environment', 'value': 'Production'},
 {'key': 'Team', 'value': 'ML-Engineering'},
 {'key': 'Project', 'value': 'CustomerSupport'}
],
 clientRequestToken=f"deployment-{uuid.uuid4()}"
)
```

```
deployment_arn = response['customModelDeploymentArn']
print(f"Deployment created: {deployment_arn}")
return deployment_arn

except Exception as e:
 print(f"Error creating deployment: {str(e)}")
 raise
```

## Use a deployment for on-demand inference

After you deploy your custom model for on-demand inference, you can use it to generate responses by making inference requests. For `InvokeModel` or `Converse` operations, you use the deployment Amazon Resource Name (ARN) as the `modelId`.

For information about making inference requests, see the following topics:

- [Submit prompts and generate responses with model inference](#)
- [Prerequisites for running model inference](#)
- [Submit prompts and generate responses using the API](#)

## Delete a custom model deployment

After you are finished using your model for on-demand inference, you can delete the deployment. After you delete the deployment, you can't use it for on-demand inference but it doesn't delete the underlying custom model.

You can delete a custom model deployment with the Amazon Bedrock console, AWS Command Line Interface, or AWS SDKs.

### Important

Deleting a custom model deployment is irreversible. Make sure you no longer need the deployment before proceeding with the deletion. If you need to use the custom model for on-demand inference again, you must create a new deployment.

## Topics

- [Delete a custom model deployment \(console\)](#)

- [Delete a custom model deployment \(AWS Command Line Interface\)](#)
- [Delete a custom model deployment \(AWS SDKs\)](#)

## Delete a custom model deployment (console)

### To delete a custom model deployment

1. In the navigation pane, under **Inference and Assessment**, choose **Custom model on-demand**.
2. Choose the custom model deployment you want to delete.
3. Choose **Delete**.
4. In the confirmation dialog, enter the deployment name to confirm the deletion.
5. Choose **Delete** to confirm.

The deployment status will change to **Deleting** while the deletion is in progress. Once completed, the deployment will be removed from the list.

## Delete a custom model deployment (AWS Command Line Interface)

To delete a custom model deployment using the AWS Command Line Interface, use the `delete-custom-model-deployment` command with your deployment identifier.

```
aws bedrock delete-custom-model-deployment \
--custom-model-deployment-identifier "deployment-arn-or-name" \
--region region
```

## Delete a custom model deployment (AWS SDKs)

To delete a custom model deployment programmatically, use the [DeleteCustomModelDeployment](#) API operation with the deployment's Amazon Resource Name (ARN) or name. The following code shows how to use the SDK for Python (Boto3) to delete a custom model deployment.

```
def delete_custom_model_deployment(bedrock_client):
 """Delete a custom model deployment

 Args:
 bedrock_client: A boto3 Bedrock client for making API calls

 Returns:
```

```
dict: The response from the delete operation

Raises:
 Exception: If there is an error deleting the deployment
"""

try:
 response = bedrock_client.delete_custom_model_deployment(
 customModelDeploymentIdentifier="Deployment identifier"
)

 print(f"Deployment deletion initiated")
 return response

except Exception as e:
 print(f"Error deleting deployment: {str(e)}")
 raise
```

## On-demand inference on Custom Models

On-demand (OD) inference allows you to run inference on your custom Amazon Nova models without maintaining provisioned throughput endpoints. This helps you optimize costs and scale efficiently. With On-demand inference, you are charged based on usage, measured in tokens, both in and out.

### Compatibility requirements

The following compatibility requirements apply:

- OD inference is supported for Amazon Nova Pro, Lite and Micro custom understanding models. OD inference is not supported for Nova custom content generation models.
- OD inference is supported for Amazon Nova custom understanding models trained after July 16, 2025. Custom models trained before July 16, 2025 are not compatible with OD inference.
- Amazon Bedrock customization: OD inference is supported for models customized with Amazon Bedrock customization and for student models that were distilled from a teacher model with Amazon Bedrock.
- SageMaker AI customization: For models customized in SageMaker AI, OD inference is supported only for Parameter-efficient fine-tuned (PEFT) models when the model is hosted on Amazon Bedrock. This includes Direct Preference Optimization plus PEFT. OD inference is not supported for Full-rank fine-tuned models.



## Model training and inference

When you train a new custom Amazon Nova Pro, Lite, or Micro model on Amazon Bedrock or SageMaker AI using PEFT after July 16, 2025, the model will automatically be compatible with both provisioned and on-demand inference options. You can select your preferred inference method when you deploy your model.

To use OD inference with a model trained after July 16, 2025, complete the following steps:

1. Create a new fine-tuning job with either the [Amazon Bedrock customization API](#) or the [SageMaker AI customization API](#).
2. Deploy the newly trained model to Amazon Bedrock using the [CreateCustomModel API](#).
3. Deploy for on-demand inference using the CustomModelDeployment API.

## Rate limits

The following requests per minute (RPM) and tokens per minute (TPM) limits apply to on-demand inference requests:

| Base Model for Custom Model | RPM per Custom Model Deployment | TPM per Custom Model Deployment |
|-----------------------------|---------------------------------|---------------------------------|
| Nova 2 Lite                 | 2,000                           | 4,000,000                       |

To learn more about the quotas available for Amazon Nova, see [Quotas for Amazon Nova](#).

## Latency

You can expect an end-to-end latency difference (that is, Time To First Token (TTFT)) of 20-55% between the base model invocation and the adapter. The exact latency value varies by model size and is in line with industry standards.

## Limitations of customizing Amazon Nova models

Amazon Nova customization doesn't support the following capabilities on SageMaker.

- **SSH into the instance to find the metrics**

Due to security controls in place, you can't SSH into the master node in the training algo-1 instance to find memory stats or NVIDIA stats and validate the training steps.

- **Support for SageMaker trainer SDK**

SageMaker trainer is currently not available to start training jobs. You can start training jobs only via **estimator API** today.

- **Warm pools are not accessible to SageMaker training jobs**

Due to security controls in place, the SageMaker warm pools can't be used to keep the instance in the warm pool till the time to live.

- **Custom model merging**

Merging multiple models is not currently supported. This means that creating multiple LoRA adapters and perform a multi-merge operation with the base model is not available.

- **Supported observability tool**

[TensorBoard](#) is the only supported observability tool to view metrics for SageMaker training jobs. [MLFlow](#) or [WandB](#) are not currently supported. For more information about using TensorBoard in SageMaker, see [TensorBoard in SageMaker](#).

# Amazon Nova Forge

Amazon Nova Forge is a first-of-its-kind service that offers organizations the easiest and most cost-effective way to build their own frontier models using Nova.

Amazon Nova Forge introduces the concept of "open training" models, which give organizations access to a variety of early model checkpoints and the ability to blend proprietary data with Amazon-curated data sets at every stage of model training. This allows the models to maximize learning from proprietary data while minimizing risk of forgetting foundational skills like reasoning.

Nova Forge provides the following key capabilities:

- Access checkpoints across all phases of model development, and leverage new Nova models before they are widely available
- Blend your proprietary data with Amazon Nova-curated training data
- Perform reinforcement learning with reward functions in your environment
- Use push-button recipes that are optimized to build with Nova through visual workflows or a command line interface
- Use the Responsible AI Toolkit to align models to Amazon Nova's responsible AI guidelines during the training process and implement runtime controls to moderate model responses during inference.

## Prerequisites

### Topics

- [Subscribe to Nova Forge](#)
- [Other prerequisites](#)

## Subscribe to Nova Forge

To request access to the Amazon Nova Forge service, add the following tag to your console IAM role: key forge-subscription with value true. After you've added this tag to your role, please go to SageMaker AI Console > Model training and customization and click on Nova Forge. On this

page, you'll find details about the service, pricing information and the capabilities. You can request subscription and then manage your subscription from this page.

1. The role should have permission to call api `ListAttachedRolePolicy`, and the response should include either `AdministratorAccess` or `AmazonSageMakerFullAccess` policy.
2. The sign-in role should have permission to call api `ListRoleTags`, and the response tags should include `tag.key=forge-subscription`.

## Other prerequisites

Also ensure the following prerequisites are complete:

1. [General prerequisites](#)
2. Additional steps for users: Add Restricted Instance Group (RIG) to your SageMaker HyperPod cluster (to complete follow steps [here](#))

## Nova Forge access and setup

To get onboarded to Nova Forge, follow this 2-step process:

- Step 1: Subscribe to Nova Forge
- Step 2: Set up HyperPod infrastructure

## Getting the Nova Forge documents

To get the Nova Forge documents follow the below steps:

```
mkdir NovaForgeHyperpodCLI
cd NovaForgeHyperpodCLI
aws s3 cp s3://nova-forge-c7363-206080352451-us-east-1/v1/ ./ --recursive
pip install -e .
```

## Step 1: Subscribe to Nova Forge

### Quick Summary:

1. Verify that you have administrator access to the Amazon Web Services account.

2. Navigate to the SageMaker AI console and request access to Nova Forge.
3. Wait for the Nova team to email a confirmation after your subscription request is approved.
4. Tag your execution role with the `forge-subscription` tag. This tag is required to access Nova Forge features and checkpoints. Add the following tag to your execution role:
  - Key: `forge-subscription`
  - Value: `true`

## Detailed Guide

To subscribe to Nova Forge and effectively use the customization service, an Amazon Web Services customer must have admin access to their Amazon Web Services account or have their administrator grant them admin access. This document outlines the steps required to:

- Secure admin access
- Set up policies to subscribe to Nova Forge
- Access customization recipes
- Configure customization
- Monitor the workflow
- Evaluate the customized model checkpoint

## Option A

Flow 1: The account user must reach out to the account admin to request the following:

- Add the `forge-subscription` tag to the account through IAM (see Appendix A for steps).
- Add the `ListRoleTags` and `ListAttachedRolePolicies` permissions through IAM (see Appendix B for steps).

 **Required setup for requesting subscription**  
To request Nova Forge subscription, your IAM role must have the "forge-subscription" tag attached and the "ListRoleTags" and "ListAttachedRolePolicies" policies added.

[Learn more](#)

## Nova Forge

[Request subscription](#)[Documentation](#)

Choose from multiple Nova model checkpoints to introduce your data at the optimal stage, mix Amazon-curated training data from each stage with your data, integrate your proprietary tools as reward functions, customize the responsible AI controls, and build your own custom frontier model.

### Subscription

Request subscription to access Nova Forge resources. Once approved, your payer account will be eligible for a 1-month free trial. **Infrastructure costs will apply during the free trial.** See Nova Forge pricing details below.

**Status**  
☐ Not subscribed

## Option B

Flow: The account user must reach out to the account admin to request admin access to the account.


- Once admin access is granted, follow the steps in Flow 2.

## Flow 2. Amazon Web Services account w/ admin access


- Add forge-subscription tag to account through IAM. See Steps in Appendix A

## Appendix A. Add forge-subscription policy to Amazon Web Services account

1. Go to the Amazon Web Services IAM Dashboard. Click on Roles on the left. Search for admin and click on the admin role



2. Select <AssumedRoleToUse> (e.g., libsAdminAccess). Click on the Tags tab.



3. Click on Manage tabs. Add new tag. Type "forge-subscription" under Key and click on save changes

Amazon Q IAM > Roles > libsAdminAccess-DO-NOT-DELETE > Tags for libsAdminAccess-DO-NOT-DELETE

### Tags for libsAdminAccess-DO-NOT-DELETE [Info](#)

**Tags (2)**  
Tags are key-value pairs that you can add to AWS resources to help identify, organize, or search for resources.

| Key                | Value - optional   |                        |
|--------------------|--------------------|------------------------|
| forge-subscription | forge-subscription | <a href="#">Remove</a> |
| prime-subscription | prime-subscription | <a href="#">Remove</a> |

[Add new tag](#)  
You can add up to 48 more tags.

[Cancel](#) [Save changes](#)

4. Ensure that you see forge-subscription as a key in Tags section

Amazon Q IAM > Roles > libsAdminAccess-DO-NOT-DELETE

### Identity and Access Management (IAM)

Search IAM

- Dashboard
- Access management
  - User groups
  - Users
  - Roles**
  - Policies
  - Identity providers
  - Account settings
  - Root access management
  - Temporary delegation requests
- Access reports
  - Access Analyzer
  - Resource analysis

**Tags updated.**

### libsAdminAccess-DO-NOT-DELETE [Info](#)

**Summary**

**Creation date**  
November 06, 2025, 09:29 (UTC-08:00)

**Last activity**  
9 minutes ago

**ARN**  
arn:aws:iam::f...s-DO-NOT-DE

**Maximum session**  
1 hour

Permissions | Trust relationships | **Tags (2)** | Last Acc...


**Tags (2) [Info](#)**  
Tags are key-value pairs that you can add to AWS resources to help identify

| Key                |
|--------------------|
| forge-subscription |




## Appendix B. Add ListRoleTags and ListAttachedPolicies policies to Amazon Web Services account for Non-Admin Role by Admin

1. Go to the Amazon Web Services IAM Dashboard. Click on Roles on the left. Search for <AssumedRoleToUse> (e.g., ForgeAccessRole) and click on the <AssumedRoleToUse> (e.g., ForgeAccessRole) role




2. Click on the <AssumedRoleToUse> (e.g., ForgeAccessRole) role and select Tags. Add a new tag with type "forge-subscription"



3. Under Permissions, add new permission: Add Permissions → Create inline policy → Add the following policy listed below

```
{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Sid": "VisualEditor0",
 "Effect": "Allow",
 "Action": [
 "iam:ListRoleTags",
 "iam:ListAttachedRolePolicies"
],
 "Resource": "*"
 }
]
}
```



## Step 2. Set up HyperPod infrastructure

Set up the necessary infrastructure by following the [workshop instructions](#) for configuring the environment with Forge-enabled features.

## Content moderation settings

If you need access to Nova Forge, customizable content moderation settings (CCMS) are available for Amazon Nova Lite 1.0 and Pro 1.0 models. CCMS allows adjustment of content moderation controls to align with specific business requirements while maintaining essential responsible AI safeguards. To determine if a business model is appropriate for CCMS, contact an AWS Account Manager.

For additional information on configuring and using CCMS with custom models, see the [Responsible AI Toolkit and Content Moderation section](#).

## Continued Pre-Training and Mid-Training

### Note

Detailed documentation is provided once subscribed

Nova Forge CPT offers advanced capabilities beyond standard CPT, including access to intermediate checkpoints and data mixing with Nova's pre-training corpus. These features enable more efficient domain adaptation and better preservation of the model's general capabilities.

### What are intermediate checkpoints and why are they needed?

Intermediate checkpoints are snapshots of the Amazon Nova model saved at different stages of the pre-training, before the model reaches its final production-ready state. During model development, Amazon Nova undergoes multiple training phases: initial pre-training with constant learning rate, learning rate ramp-down, context extension training, and finally instruction-following alignment and safety training. For CPT, intermediate checkpoints are often preferable to the final Prod checkpoint because they are more plastic and receptive to domain adaptation. The Prod checkpoint has undergone extensive instruction-following alignment and safety training, which optimizes the model for general conversational use but can make it resistant to learning new domain-specific patterns during CPT. In contrast, Partially and Fully pre-trained text only checkpoints retain the model's pre-training characteristics. They haven't been heavily steered toward specific behaviors, making them more efficient starting points for domain adaptation. When performing large-scale CPT (>10B tokens), starting from intermediate checkpoints typically results in faster convergence, better training stability, and more effective domain knowledge acquisition. However, for small-scale CPT (<10B tokens), or when instruction-following capabilities need to be preserved, the Prod checkpoint may be more appropriate as it allows domain adaptation while maintaining the model's conversational abilities.

Multiple intermediate checkpoints are necessary for CPT because they offer different levels of model plasticity that affect how efficiently the model can absorb new domain knowledge. The final Prod checkpoint has undergone extensive instruction-following alignment and safety training, which optimizes it for general conversational use but makes it resistant to learning new domain-specific patterns. In other words, It has been hardened through post-training. In contrast, earlier

checkpoints retain the model's pre-training characteristics and haven't been heavily steered toward specific behaviors, making them more plastic and receptive to domain adaptation.

To achieve the best training efficiency, multiple intermediate checkpoints are provided.

## What checkpoints are available?

### Nova 2.0

There are three Amazon Nova Lite 2.0 checkpoints.

- **PRE-TRAINED** - [nova-lite-2/pretraining-text-RD]: This is the checkpoint after the constant learning rate and ramp-down stages of Amazon Nova pre-training where the model is trained on trillions of tokens.
- **MID-TRAINED** - [nova-lite-2/pretraining-text-CE]: This checkpoint allows intermediate volumes of unstructured data to be introduced with a more conservative learning rate than pre-training, absorbing domain-specific knowledge while avoiding catastrophic forgetting.
- **POST-TRAINED** - [nova-lite-2/prod]: This is the fully aligned final checkpoint of the model that has gone through all the pertaining and post training steps.

The following table elaborates on the different conditions for pre- and mid-training.

| Data Type                                                                        | Perform                | With Checkpoint |
|----------------------------------------------------------------------------------|------------------------|-----------------|
| Large-scale unstructured raw domain data (documents, logs, articles, code, etc.) | Continued Pre-Training | Pre-Trained     |
| Large-scale unstructured raw domain data (documents, logs, articles, code, etc.) | Mid-Training           | Pre-Trained     |
| Smaller volumes of unstructured raw data. Structured reasoning traces / CoT data | Mid-Training           | Mid-Trained     |
| Structured demonstrations (high-quality input-output                             | Full Fine-Tuning       | Mid-Trained     |

| Data Type                                                                                                    | Perform                         | With Checkpoint |
|--------------------------------------------------------------------------------------------------------------|---------------------------------|-----------------|
| pairs, curated task instructions, multi-turn dialogues)                                                      |                                 |                 |
| Structured demonstrations (high-quality input-output pairs, curated task instructions, multi-turn dialogues) | Parameter Efficient Fine-Tuning | Post-Trained    |

## Which checkpoint to use?

Partially pre-trained text only and fully pre-trained text only checkpoints typically converge faster and require fewer training steps for domain adaptation. However, they have no instruction tuning and would need to undergo post training steps to be able to perform useful tasks and follow instructions. GA checkpoint may require more steps to adapt but provides safer starting point for small-scale experiments and will maintain some of its post training capabilities even after CPT training.

In general, with large training datasets (>10B tokens), start from partially pre-trained text only or fully pre-trained text only checkpoints for more efficient and stable training, as the model's knowledge base will be substantially modified. With small datasets (<10B tokens), use the GA checkpoint to preserve instruction-following capabilities while adapting to the domain.

## How to use data mixing for Nova 2.0?

When performing CPT with a new domain data, it is highly beneficial to mix the new data with some of the data used previously in the pre-training stage of the model. Mixing old data with new domain data solves two problems:

- **Forgetting control:** Prevents catastrophic forgetting by preserving existing skills and knowledge of the model. Without data mixing, training exclusively on narrow domain data causes the model to overwrite general capabilities. For example, a model trained only on legal documents might lose its ability to code or do math. Mixing the general domain datasets preserves these general skills while acquiring the new domain.
- **Optimization stability:** Maintains training stability by anchoring the model's internal representations. During CPT, the model's learned features are modified and data mixing provides gradients from diverse sources that guide this adaptation smoothly. Without it, training on

narrow distributions can cause gradient instability, where the model's representations shift too drastically, leading to training divergence, loss spikes, or collapse of existing capabilities. This is the stability-plasticity tradeoff: the model should be plastic enough to learn new domain knowledge, but stable enough not to break what it already knows.

## Nova CPT Data Mixing Capabilities

Access to Amazon Nova pre-training data and checkpoints is one of the core offerings of the Amazon Nova CPT customization. Amazon Nova CPT customization enables easy mixing of domain data with Amazon Nova's pre-training corpus. Further, the sampling ratio of the specific Amazon Nova data categories (e.g., code, math, reasoning, etc) can be changed and their proportions controlled to complement domain data. This allows reinforcement of capabilities that align with the use case while adapting the model to the specific domain.

### Finding the Optimal Mixing Ratio

The optimal ratio of Amazon Nova data versus domain data depends on the dataset's domain, complexity, size, quality, and the importance of maintaining general capabilities. This ratio must be discovered through experimentation. An experiment framework to decide on how much Amazon Nova data to mix is as follows.

Select a representative subset of domain data (e.g., 5B tokens) and keep this constant across all experimental runs.

Run small-scale CPT experiments varying only the amount of Amazon Nova data mixed in:

- No mixing: 100% domain → 5B domain only (total 5B)
- Light mixing: 90% domain → 5B domain + ~0.56B Amazon Nova (total ~5.56B)
- Medium mixing: 70% domain → 5B domain + ~2.14B Amazon Nova (total ~7.14B)
- Heavy mixing: 50% domain → 5B domain + 5B Amazon Nova (total 10B)

Evaluate each checkpoint on in domain and general domain benchmarks. Also evaluate the starting checkpoint (Amazon Nova checkpoint before any training).

- Does customer-domain performance stay roughly constant across runs? It usually should, since each run saw the same number of domain tokens. If domain performance improves with more mixing, Amazon Nova data provides useful regularization.
- Do general benchmark scores improve as mixing is increased?

- Expected behavior is that the general capabilities should improve monotonically as more Amazon Nova data is added.
- Measure multiple general benchmarks: MMLU (general knowledge), HumanEval (coding), GSM8K (math), or specific benchmarks of interest.
- Select the mixing ratio that maintains domain performance while delivering acceptable general capabilities for the use cases. Factor in the additional cost of training with more data mixing.

Once the optimal mixing ratio has been identified, run full-scale CPT using the complete domain dataset with the selected mixing ratio.

## Dissecting the Data Mixing Categories

Below we dissect each available category in Data Mixing, for you to make best decision of what data categories makes most sense to be represented in your overall data mixture.

### How to Enable Data Mixing

Add the `data_mixing` section to your recipe with the appropriate percentage distribution across dataset categories. The `nova_data` percentages must sum to 100.

#### Nova 2.0 Configuration with data mixing

```
Note:
This recipe can run on p5.48xlarge

Run config
display_name: "Nova Lite Pretrain on P5 GPU"
versions: ["2.0"]
instance_types: ["ml.p5.48xlarge"]

run:
 name: "my-cpt-run" # A descriptive name for your training job
 model_type: "amazon.nova-2-lite-v1:0:256k" # Model variant specification, do not
 change
 model_name_or_path: "nova-lite-2/prod" # Base model path, do not change
 replicas: 8 # Number of compute instances for training, allowed values are 4,
 8, 16, 32
 data_s3_path: "" # Customer data paths
 validation_data_s3_path: "" # Customer validation data paths
```



```

output_s3_path: "" # Output artifact path, SageMaker HyperPod job-specific
configuration - not compatible with standard SageMaker Training jobs

Training specific configs
training_config:
 task_type: cpt
 max_length: 8192 # Maximum context window size (tokens)
 global_batch_size: 64 # Global batch size, allowed values are 32, 64, 128,
256.

 trainer:
 max_steps: 10 # The number of training steps to run total
 val_check_interval: 10 # The number of steps between running validation
 limit_val_batches: 2 # Batches of the validation set to use each trigger

 model:
 hidden_dropout: 0.0 # Dropout for hidden states, must be between 0.0 and
1.0
 attention_dropout: 0.0 # Dropout for attention weights, must be between 0.0
and 1.0

 optim:
 optimizer: adam
 lr: 1e-5 # Learning rate
 name: distributed_fused_adam # Optimizer algorithm, do not change
 adam_w_mode: true # Enable AdamW mode
 eps: 1e-06 # Epsilon for numerical stability
 weight_decay: 0.0 # L2 regularization strength, must be between 0.0 and
1.0
 adam_beta1: 0.9 # Beta1 for Adam optimizer
 adam_beta2: 0.95 # Beta2 for Adam optimizer
 sched:
 warmup_steps: 10 # Learning rate warmup steps
 constant_steps: 0 # Steps at constant learning rate
 min_lr: 1e-6 # Minimum learning rate, must be lower than lr

data_mixing:
 dataset_catalog: cpt_text_lite
 sources:
 nova_data: # percent inputs for Nova data must sum to 100%; use 0% if you want to
exclude a data grouping
 agents: 20
 business-and-finance: 4
 scientific: 10

```

```
code: 5
factual-and-news: 5
longform-text: 6
health-and-medicine: 1
humanities-and-education: 1
legal: 1
math: 9
additional-languages: 15
social-and-personal-interest: 11
entertainment: 0.5
reasoning: 10
other: 0.5
tables: 1
customer_data: # percent input of customer data. 100 = use only customer data, 0 =
use only the nova_data mix above
percent: 25
```

What do these categories mean

**Note:** Nova 2.0 includes additional reasoning-specific categories (e.g., reasoning-code, reasoning-math, reasoning-instruction-following) that are not available in Nova 1.0.

Summary of Categories and Info Labels:

| Category Name | Info detail                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------|
| agents        | Training data focused on autonomous decision-making, task completion, and goal-oriented behavior in AI systems         |
| baseline      | Fundamental language data focused on general comprehension, basic communication, and core linguistic capabilities      |
| chat          | Conversational exchanges demonstrating natural dialogue flow, context maintenance, and appropriate social interactions |
| code          | Programming source code, documentation, and technical discussions from various programming languages and platforms.    |
| factuality    | Reference materials and verified information focused on accuracy, source validation, and truth assessment              |

| Category Name                   | Info detail                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| identity                        | Personality frameworks and behavioral patterns focused on consistent character traits, values, and interaction styles            |
| long-context                    | Extended texts and complex narratives focused on maintaining coherence and relevance across lengthy exchanges                    |
| math                            | Mathematical content including textbooks, problems, solutions, and mathematical discussions.                                     |
| rai                             | Cases and scenarios emphasizing ethical AI principles, safety considerations, and responsible technology deployment              |
| instruction-following           | Examples of precise task execution based on varying levels of user prompts and directives                                        |
| stem                            | Technical content covering science, technology, engineering, and mathematics, including problem-solving and theoretical concepts |
| planning                        | Sequences demonstrating strategic thinking, step-by-step task breakdown, and efficient resource allocation                       |
| reasoning-chat                  | Analytical dialogue scenarios focused on logical discussion and structured conversation flows                                    |
| reasoning-code                  | Programming challenges and algorithmic problems focused on systematic solution development                                       |
| reasoning-factuality            | Information evaluation scenarios focused on critical assessment and verification processes                                       |
| reasoning-instruction-following | Complex task analysis focused on systematic interpretation and methodical execution                                              |
| reasoning-math                  | Mathematical problem-solving scenarios focused on logical progression and solution strategies                                    |

| Category Name      | Info detail                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| reasoning-planning | Strategic decision-making scenarios focused on systematic approach to goal achievement                                              |
| reasoning-rag      | Information retrieval and synthesis scenarios focused on contextual understanding and relevant application                          |
| reasoning-rai      | Ethical decision-making scenarios focused on systematic evaluation of AI safety and fairness                                        |
| reasoning-stem     | Scientific problem-solving scenarios focused on methodical analysis and solution development                                        |
| rag                | Examples of effectively combining retrieved external knowledge with generated responses to provide accurate, contextual information |
| translation        | Multi-language content pairs showing accurate translation while preserving context, tone, and cultural nuances                      |

## Parameter Guide

- **dataset\_catalog:** The only value is `cpt_text_lite` for now, until we enable the multimodal training.
- **nova\_data:** Percentage of the individual categories of Nova data when mixed in. They should add up to 1.0.
- **customer\_data:** the percentage of customer's data mixed into the Nova data.

The total number of tokens used in training can be calculated from `max_length * global_batch_size * max_steps`

## Limitations

Current CPT only supports text data and does not support any customer multi-modal datasets.

# Supervised Fine-Tuning

## Introduction

Supervised fine-tuning uses dataset with input-output pairs for the task of interest. In other words, you provide examples of prompts (questions, instructions, etc.) along with the correct or desired responses and continue training the model on these. The model's weights are adjusted to minimize a supervised loss, typically cross-entropy between its predictions and the target response tokens.

## When to use SFT?

SFT is best when you have a well-defined task with clear desired outputs. If you can explicitly say "Given X input, the correct/desired output is Y" and you can gather examples of such X-Y mappings, then supervised fine-tuning is a great choice. Some scenarios where SFT excels include:

- **Structured or complex classification tasks:** e.g. classifying internal documents or contracts into many custom categories. With SFT, the model can learn these specific categories far better than prompting alone.
- **Question-answering or transformation tasks with known answers:** e.g. fine-tuning a model to answer questions from a company's knowledge base, or to convert data between formats, where each input has a correct response.
- **Formatting and style consistency:** If you need the model to always respond in a certain format or tone, you can fine-tune on examples of the correct format/tone. For instance, training on prompt-response pairs that demonstrate a particular brand voice or style can teach the model with that style in its outputs. Instruction-following behavior is often initially taught via SFT on curated examples of good assistant behavior.

SFT is the most direct way to teach an LLM a new skill or behavior when you can specify what the right behavior looks like. It leverages the model's existing language understanding and focuses it on your task. Do not use SFT when the gap is knowledge rather than behavior; it will not make the model learn new facts, jargon, or recent events. In those cases, prefer continued pre-training on large in-domain corpora or retrieval-augmented generation to bring external knowledge at inference. When you can measure quality but cannot label a single right answer, reinforcement fine-tuning with verifiable rewards or an LLM-as-judge might be preferable to SFT.

Depending on task complexity and performance of the Nova model without tuning, plan for thousands to tens of thousands of demonstrations per task, with data quality, consistency, and diversity mattering more than raw volume.

## When to use parameter efficient and when full rank SFT?

Nova customization recipes enable you to perform parameter efficient, in particular LoRA, or full rank SFT. If you want a straightforward, cost-efficient model update, or have very little data, favor parameter-efficient methods so you train small adapters while leaving most of the backbone untouched (full rank SFT updates all model parameters).

## Data Mixing for SFT

Data mixing allows you to combine your custom training datasets with Nova's proprietary training data. This feature is available for both Nova 1.0 and Nova 2.0 models.

**Nova Proprietary Data Type:** Nova supports both text and multimodal SFT data types. It is organized into multiple data categories each containing a blend of tasks relevant for the corresponding category.

**Nova Proprietary Data Categories:** Text datasets includes several categories including: autonomous decision making, task completion, goal oriented datasets (agents), both reasoning and non-reasoning precise task execution datasets (reasoning-instruction-following, instruction-following), sequences demonstrating strategic thinking and step-by-step task breakdown (planning), responsible AI (rai), long-context, factuality, math, stem and many more. Similarly, multimodal datasets includes video, screenshot, charts and many more.

The data mixing feature allows you to blend your own fine-tuning training samples with samples from the Nova datasets used to fine-tune the Nova. This can prevent overfitting on your custom training and "catastrophic forgetting" of Nova capabilities, or help you build capabilities when training from a new pretrained checkpoint.

To mix in Nova data, you simply need to add a `data_mixing` block to your recipe YAML file, under the `training_config` section. Text and multi-modal data mixing blocks have different content. Please refer to corresponding recipes.

## Supported Models

- Nova 2.0 Lite

## Supported Modality

- Text
- Multimodal

## YAML Configuration Examples

### Nova 2.0 Configuration Example

```
run:
 name: my-lora-sft-run
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: nova-lite-2/prod
 data_s3_path: s3://my-bucket-name/train.jsonl
 replicas: 4
 output_s3_path: s3://my-bucket-name/outputs/
 mlflow_tracking_uri: ""
 mlflow_experiment_name: "my-lora-sft-experiment"
 mlflow_run_name: "my-lora-sft-run"

training_config:
 max_steps: 100
 save_steps: 10
 save_top_k: 5
 max_length: 32768
 global_batch_size: 32
 reasoning_enabled: true
 lr_scheduler:
 warmup_steps: 15
 min_lr: 1e-6
 optim_config:
 lr: 1e-5
 weight_decay: 0.0
 adam_beta1: 0.9
 adam_beta2: 0.95
 peft:
 peft_scheme: "lora"
 lora_tuning:
 alpha: 64
 lora_plus_lr_ratio: 64.0
```

## Nova 2.0 Text Data Mixing

```
data_mixing:
 dataset_catalog: sft_1p5_text_chat
 sources:
 customer_data:
 percent: 50
 nova_data:
 agents: 1
 baseline: 10
 chat: 0.5
 code: 10
 factuality: 0.1
 identity: 1
 long-context: 1
 math: 2
 rai: 1
 instruction-following: 13
 stem: 0.5
 planning: 10
 reasoning-chat: 0.5
 reasoning-code: 0.5
 reasoning-factuality: 0.5
 reasoning-instruction-following: 45
 reasoning-math: 0.5
 reasoning-planning: 0.5
 reasoning-rag: 0.4
 reasoning-rai: 0.5
 reasoning-stem: 0.4
 rag: 1
 translation: 0.1
```

## Nova 2.0 Multimodal Data Mixing

```
data_mixing:
 dataset_catalog: sft_1p5_mm_chat
 sources:
 customer_data:
 percent: 50
 nova_data:
 charts: 1
```



```
chat: 38
code: 20
docs: 3
general: 2
grounding: 1
rag: 4
screenshot: 4
text: 8
translation: 4
video: 15
```

## Model Checkpoints

### Nova 2.0 Checkpoints

- **PRE-TRAINED** [nova-lite-2/pretraining-text-RD]: Checkpoint after constant learning rate and ramp-down stages where model is trained on trillions of tokens. [Outcome of Stage 2]
- **MID-TRAINED** [nova-lite-2/pretraining-text-CE]: Allows customers with intermediate volumes of unstructured data to introduce their data with a more conservative learning rate than pre-training, absorbing domain-specific knowledge while avoiding catastrophic forgetting. [Outcome of Stage 3]
- **FINAL** [nova-lite-2/prod]: Fully aligned final checkpoint that has gone through all pretraining and post training steps. [Outcome of Stage 4]

#### Training Stages:

- Stage 1: PT Ckpt, initial pre-training with constant learning rate
- Stage 2: PT Ckpt, learning rate ramp-down
- Stage 3: PT Ckpt, context extension training
- Stage 4: instruction-following alignment and safety training

# Training Approaches

## Training Approach Selection Guide

| Data Type                                                                                                    | Data Volume  | Perform                      | With Checkpoint                     |
|--------------------------------------------------------------------------------------------------------------|--------------|------------------------------|-------------------------------------|
| Large-scale unstructured raw domain data (documents, logs, articles, code, etc.)                             | 1T+ Tokens   | Continued Pre-Training       | End of Constant Learning Rate (CLR) |
| Large-scale unstructured raw domain data                                                                     | 100B+ Tokens | Mid-Training                 | End of CLR                          |
| Smaller volumes of unstructured raw data; Structured reasoning traces / CoT data                             | 1B+ Tokens   | Mid-Training                 | Nova base model                     |
| Structured demonstrations (high-quality input-output pairs, curated task instructions, multi-turn dialogues) | 1K+ Examples | Supervised Fine-Tuning (SFT) | Nova base model                     |

## Pre-Requisites before you begin

- We assume that you've already setup an SMHP cluster with a restricted instance group (RIG) that has active capacity. If not please refer here to get your SMHP Cluster and RIG setup completed [[Docs Link](#), [Workshop Link](#)]
- You will require **p5.48xlarge** EC2 instances to execute this recipe. The minimum number instances required to execute this recipe efficiently are as follows:
  - **Nova Lite 2.0 - 4 p5.48xlarge**
- Install the Forge Specific SageMaker HyperPod CLI using the provided instructions [here](#)
- Confirm that you can connect to your cluster using `hyperpod get-clusters`

- Note that this command will list all SMHP clusters in your account
- Confirm that your training, and optionally validation data, is available in an S3 bucket that is accessible by the execution role of your SMHP cluster. For data preparation, refer to next section.
- Have AWS CLI setup completed. If you have not completed the setup, please follow below [guide](#).
- **Verification:** After completing the setup, confirm you can successfully run below commands

```
aws sagemaker describe-cluster --cluster-name <cluster-name> --region <region>

hyperpod connect-cluster --cluster-name cluster-name
```

## A systematic Approach to Achieving Successful SFT

- **Data Preparation:** Follow established guidelines to create, clean, or reformat datasets into the required structure. Ensure that inputs, outputs, and auxiliary information (such as reasoning traces or metadata) are properly aligned and formatted.
- **Training Configuration:** Define how the model will be trained. When using Amazon SageMaker HyperPod, this configuration is written in a YAML recipe file that includes:
  - Data source paths (training and validation datasets)
  - Key hyperparameters (number of training steps, learning rate, batch size)
  - Optional components (distributed training parameters, etc)
  - Data Mixing setting (defines proportions of customer and Nova data categories)
- **Optimize SFT Hyper Parameters:** SFT recipe parameter values we recommend are a great starting point and a robust choice. If you want to optimize them further for your use case do multiple SFT runs with different parameter combinations and pick the best one. You can select parameter combinations following Hyper-Parameter Optimization method of your choice. A simple approach is to vary the value of one parameter (default\*0.5, default, default\*2) while keeping other default value for other parameters, repeat this for each parameter you want to optimize, and iterate if needed. The most relevant parameters for LoRA are learning rate, alpha (scaling parameter), number of epochs to train and warmup steps; for full-rank it is mainly the learning rate, number of epochs, and warmup steps.

## Experiment Sequencing and Data Mixing

- If you have only SFT data (train/dev/test) for a set of tasks and care only about the test performance on these tasks
  - Do SFT without mixing on [FINAL] Nova checkpoint. Use the default SFT hyper-parameters and optionally optimize them for your use case. Monitor validation metrics and/or evaluate intermediate checkpoints for larger datasets.
- If you have only SFT data (train/dev/test) for a set of tasks and care about test performance on these tasks and general benchmarks in the domain of interest
  - Start by doing SFT with Nova data mixing on a pre-training checkpoint (PRE-TRAINED or MID-TRAINED checkpoint, not FINAL). Using an intermediate checkpoint allows the model to better integrate your custom data with Nova's proprietary data while maintaining strong general capabilities.
  - Run shorter SFT training runs with varying amount of Nova data in the mix (e.g., 10%, 25%, 50%, 75%) and Nova data category selections that complement your use case (e.g., pick instruction following category if you care about general instruction following ability). Monitor validation metrics and evaluate if mixing helps performance on general benchmarks. Select the training mix and checkpoint that leads to the best combination of performance on your task and general performance. Depending on the use case, both task and general performances can be further improved using reinforcement fine tuning (RFT).

## Prepare Dataset for SFT

**Nova 2.0:** Use Converse API format <https://docs.aws.amazon.com/bedrock/latest/userguide/conversation-inference-call.html>. Nova 2.0 data format can contain additional reasoning fields: [https://docs.aws.amazon.com/bedrock/latest/APIReference/API\\_runtime\\_ReasoningContentBlock.html](https://docs.aws.amazon.com/bedrock/latest/APIReference/API_runtime_ReasoningContentBlock.html)

Reasoning content captures the model's intermediate thinking steps before generating a final answer. In the assistant turn, use the `reasoningContent` field to include reasoning traces. Use plain text for reasoning content, avoid markup tags like `<thinking>` and `</thinking>` unless specifically required by your task, and ensure reasoning content is clear and relevant to the problem-solving process.

# Evaluation Methods

## Prerequisites

- Checkpoint S3 URI from your training job's `manifest.json` file (for trained models)
- Evaluation dataset uploaded to S3 in the correct format
- Output S3 path for evaluation results

**Out of the box benchmarks:** Use out of the box benchmarks to validate the performance on general tasks. For more details check here: <https://docs.aws.amazon.com/sagemaker/latest/dg/nova-hp-evaluate.html>

## Bring Your Own Data

You can also supply your custom data by formatting them in below format and then using below mentioned containers to get inference results along with log probabilities for calibrations if needed.

Create jsonl per task with the following structure:

```
{
 "metadata": "{key:4, category:'apple'}",
 "system": "arithmetic-patterns, please answer the following with no other words: ",
 "query": "What is the next number in this series? 1, 2, 4, 8, 16, ?",
 "response": "32"
}
```

Outputs generated during inference phase of the evaluation job will have following structure:

```
{
 "prompt": "[{'role': 'system', 'content': 'arithmetic-patterns, please answer the following with no other words: '}, {'role': 'user', 'content': 'What is the next number in this series? 1, 2, 4, 8, 16, ?'}]",
 "inference": "['32']",
 "gold": "32",
 "metadata": "{key:4, category:'apple'}"
}
```

## Field descriptions:

- **prompt:** Formatted input sent to the model
- **inference:** Model's generated response
- **gold:** Expected correct answer from input dataset, response field from the input
- **metadata:** Optional metadata passed through from input

## Prepare Evaluation Config

Command to launch evaluation job. Use "--override-parameters" to modify any entry from the recipe.

```
hyperpod start-job -n kubeflow \
 --recipe evaluation/nova/nova_micro_p5_48x1_bring_your_own_dataset_eval \
 --override-parameters '{
 "instance_type": "p5.48xlarge",
 "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM-
HP-Eval-latest",
 "recipes.run.name": "<your-eval-job-name>",
 "recipes.run.model_name_or_path": "<checkpoint-s3-uri>",
 "recipes.run.output_s3_path": "s3://<your-bucket>/eval-results/",
 "recipes.run.data_s3_path": "s3://<your-bucket>/eval-data.jsonl"
 }'
```

## Best Practices

- **Prioritize data quality over volume:** High-quality, diverse, and representative training data is more valuable than large quantities of low-quality data.
- **Include reasoning-instruction-following category:** When using data mixing, include the "reasoning-instruction-following" category to maintain strong generic performance across tasks.
- **Use default learning rates:** Start with default learning rates (1e-5 for LoRA, 5e-6 for full-rank SFT) and adjust only if needed based on validation metrics.
- **Balance Nova data mixing:** Mix maximum 50% Nova data for optimal latency-performance balance. Higher percentages may improve general capabilities but can increase training time.

- **Monitor validation metrics:** Regularly evaluate intermediate checkpoints during training to detect overfitting or performance degradation early.
- **Test on representative datasets:** Ensure your evaluation datasets accurately represent your production use cases for meaningful performance assessment.

## Prepare Training Job Config

### Hyper Parameters

Full set of hyper-parameters other than data mixing:

```
Run config
run:
 name: my-lora-sft-run
 model_type: amazon.nova-2-lite-v1:0:256k
 model_name_or_path: nova-lite-2/prod
 data_s3_path: s3://my-bucket-name/train.jsonl # SageMaker HyperPod (SMHP) only and
not compatible with SageMaker Training jobs. Note replace my-bucket-name with your
real bucket name for SMHP job
 replicas: 4 # Number of compute instances for training, allowed
values are 4, 8, 16, 32
 output_s3_path: s3://my-bucket-name/outputs/ # Output artifact path
(Hyperpod job-specific; not compatible with standard SageMaker Training jobs). Note
replace my-bucket-name with your real bucket name for SMHP job

MLFlow configs
mlflow_tracking_uri: "" # Required for MLFlow
mlflow_experiment_name: "my-lora-sft-experiment" # Optional for MLFlow. Note: leave
this field non-empty
mlflow_run_name: "my-lora-sft-run" # Optional for MLFlow. Note: leave this field non-
empty

training_config:
 max_steps: 100 # Maximum training steps. Minimal is 4.
 save_steps: 10 # This parameter suggests after how many training steps the
checkpoints will be saved. Should be less than or equal to max_steps(please override
this value with a numerical value equal or less than max_steps value; min: 4)
 save_top_k: 5 # Keep top K best checkpoints. Note supported only
for SageMaker HyperPod jobs. Minimal is 1.
 max_length: 32768 # Sequence length (options: 8192, 16384, 32768
[default], 65536)
```

```

global_batch_size: 32 # Global batch size (options: 32, 64, 128)
reasoning_enabled: true # If data has reasoningContent, set to true;
otherwise False

lr_scheduler:
 warmup_steps: 15 # Learning rate warmup steps. Recommend 15% of
max_steps # Minimum learning rate, must be between 0.0 and 1.0
 min_lr: 1e-6

optim_config: # Optimizer settings
 lr: 1e-5 # Learning rate, must be between 0.0 and 1.0
 weight_decay: 0.0 # L2 regularization strength, must be between 0.0
and 1.0
 adam_beta1: 0.9 # Exponential decay rate for first-moment estimates,
must be between 0.0 and 1.0
 adam_beta2: 0.95 # Exponential decay rate for second-moment
estimates, must be between 0.0 and 1.0

peft: # Parameter-efficient fine-tuning (LoRA)
 peft_scheme: "lora" # Enable LoRA for PEFT
 lora_tuning:
 alpha: 64 # Scaling factor for LoRA weights (options: 32, 64,
96, 128, 160, 192),
 lora_plus_lr_ratio: 64.0 # LoRA+ learning rate scaling factor (0.0-100.0)

```

The most relevant parameters for LoRA are learning rate, alpha (scaling parameter), number of epochs to train and warmup steps; for full-rank it is mainly the learning rate, number of epochs, and warmup steps. The recipes are pre-populated with the recommended defaults.

## Set Up Data Mixing Block

Add the `data_mixing` section to your recipe with the appropriate percentage distribution across dataset categories.

Below we describe each available Nova data category.

## Nova 2.0 Configuration with data mixing

```

data_mixing:
 dataset_catalog: sft_1p5_text_chat # Nova text dataset catalog
 sources:

```



```

customer_data:
 percent: 50 # Percent of overall mix to draw from customer data
nova_data: # The remainder will be drawn from Nova data. The
categories below must add to 100
 agents: 1 # autonomous decision-making, task completion, goal-
oriented behavior in AI systems
 baseline: 10 # [New in Nova 1.5]
 chat: 0.5 # Conversational exchanges demonstrating natural
dialogue flow
 code: 10 # Programming examples and solutions spanning
multiple languages
 factuality: 0.1 # [New in Nova 1.5]
 identity: 1 # [New in Nova 1.5]
 long-context: 1 # [New in Nova 1.5]
 math: 2 # [New in Nova 1.5]
 rai: 1 # ethical AI principles, safety considerations, and
responsible technology deployment
 instruction-following: 13 # precise task execution based on varying levels of
user prompts and directives
 stem: 0.5 # Technical content covering science, technology,
engineering, and mathematics
 planning: 10 # Sequences demonstrating strategic thinking and
step-by-step task breakdown
 reasoning-chat: 0.5
 reasoning-code: 0.5
 reasoning-factuality: 0.5
 reasoning-instruction-following: 45
 reasoning-math: 0.5
 reasoning-planning: 0.5
 reasoning-rag: 0.4
 reasoning-rai: 0.5
 reasoning-stem: 0.4
 rag: 1 # combining retrieved external knowledge with
generated responses
 translation: 0.1

```

What do these categories mean?

## Nova 2.0 Text Data Categories

| Category Name | Info detail                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------------------|
| agents        | Training data focused on autonomous decision-making, task completion, and goal-oriented behavior in AI systems         |
| baseline      | Fundamental language data focused on general comprehension, basic communication, and core linguistic capabilities      |
| chat          | Conversational exchanges demonstrating natural dialogue flow, context maintenance, and appropriate social interactions |
| code          | Programming source code, documentation, and technical discussions from various programming languages and platforms.    |
| factuality    | Reference materials and verified information focused on accuracy, source validation, and truth assessment              |
| identity      | Personality frameworks and behavioral patterns focused on consistent character traits, values, and interaction styles  |
| long-context  | Extended texts and complex narratives focused on maintaining coherence and relevance across lengthy exchanges          |
| math          | Mathematical content including textbooks , problems, solutions, and mathematical discussions.                          |
| rai           | Cases and scenarios emphasizing ethical AI principles, safety considerations, and responsible technology deployment    |

| Category Name                   | Info detail                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| instruction-following           | Examples of precise task execution based on varying levels of user prompts and directives                                        |
| stem                            | Technical content covering science, technology, engineering, and mathematics, including problem-solving and theoretical concepts |
| planning                        | Sequences demonstrating strategic thinking, step-by-step task breakdown, and efficient resource allocation                       |
| reasoning-chat                  | Analytical dialogue scenarios focused on logical discussion and structured conversation flows                                    |
| reasoning-code                  | Programming challenges and algorithmic problems focused on systematic solution development                                       |
| reasoning-factuality            | Information evaluation scenarios focused on critical assessment and verification processes                                       |
| reasoning-instruction-following | Complex task analysis focused on systematic interpretation and methodical execution                                              |
| reasoning-math                  | Mathematical problem-solving scenarios focused on logical progression and solution strategies                                    |
| reasoning-planning              | Strategic decision-making scenarios focused on systematic approach to goal achievement                                           |
| reasoning-rag                   | Information retrieval and synthesis scenarios focused on contextual understanding and relevant application                       |

| Category Name  | Info detail                                                                                                                         |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------|
| reasoning-rai  | Ethical decision-making scenarios focused on systematic evaluation of AI safety and fairness                                        |
| reasoning-stem | Scientific problem-solving scenarios focused on methodical analysis and solution development                                        |
| rag            | Examples of effectively combining retrieved external knowledge with generated responses to provide accurate, contextual information |
| translation    | Multi-language content pairs showing accurate translation while preserving context, tone, and cultural nuances                      |

## Multimodal Data Mixing (Nova 2.0)

```

data_mixing:
 dataset_catalog: sft_1p5_mm_chat # Nova text dataset catalog
 sources:
 customer_data:
 percent: 50 # Percent of overall mix to draw from customer data
 nova_data: # The remainder will be drawn from Nova data. The
categories below must add to 100
 charts: 1
 chat: 38
 code: 20
 docs: 3
 general: 2
 grounding: 1
 rag: 4
 screenshot: 4
 text: 8
 translation: 4
 video: 15

```

**Note:** Nova 2.0 includes video data category support that is not available in Nova 1.0.

What do these categories mean?

Nova 2.0 Multimodal Data Categories

| Category Name | Info detail                                                                                                                                                                                                    |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| charts        | Visual representations and descriptions of graphs, pie charts, bar charts, line plots, and other statistical visualizations to help the model understand and communicate quantitative information effectively  |
| chat          | Conversational data paired with visual elements focused on contextual dialogue understanding and image-based interactions                                                                                      |
| code          | Programming interfaces and development environments focused on visual code interpretation, IDE screenshots, and technical diagrams                                                                             |
| docs          | Document-centric data combining text, images, layouts, and formatting to train models in understanding and processing various document types and structures to help with concepts like PDF content recognition |
| general       | Diverse visual-textual content focused on broad comprehension of images, graphics, and accompanying descriptive text                                                                                           |
| grounding     | Visual reference materials and labeled imagery focused on connecting language concepts to real-world visual representations                                                                                    |
| rag           | Multimodal retrieval examples showing how to effectively combine and reference visual and textual external knowledge to generate accurate, contextual responses                                                |

| Category Name | Info detail                                                                                                                             |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| screenshot    | Application interface captures and digital display images focused on understanding software interfaces and digital interactions         |
| text          | A balanced pool of contextual text data create from the text-only SFT Nova dataset categorie s in order to provide generalist abilities |
| translation   | Cross-language visual content focused on multilingual interpretation of text in images and cultural visual elements                     |
| video         | Motion-based visual content focused on temporal understanding and sequential visual-narrative comprehension                             |

## How to Launch a job

You can also refer to the README, if you only need to get the essential details to kick off first SFT run.

Container Information:

### Container Information and Launch Commands

| Model    | Technique   | Subcategory | Image URI                                                                          | Hyperpod Launcher Command                                                            |
|----------|-------------|-------------|------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Nova 2.0 | Fine-tuning | SFT Text    | 708977205387.dkr.ecr.us-east-1.amazons.com/nova-fine-tune-repo:SM-HP-SFT-V2-latest | hyperpod start-job \ -n kubeflow \ --recipe fine-tuning/nova/nova_2_0/nova_lite/SFT/ |

| Model | Technique | Subcategory | Image URI | Hyperpod Launcher Command                                                                                                                                                                                                                          |
|-------|-----------|-------------|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|       |           |             |           | nova_lite_2_0_<br>p5_gpu_sft<br>\ --override-<br>parameters<br>{ "instance<br>_type":<br>"ml.p5.48<br>xlarge",<br>"container":<br>"70897720<br>5387.dkr.ecr.us-<br>east-1.amaz<br>onaws.com/<br>nova-fine-tune-<br>repo:SM-HP-<br>SFT-V2-latest" } |

| Model    | Technique   | Subcategory           | Image URI                                                                                       | Hyperpod Launcher Command                                                                                                                                                                                                                                                                                             |
|----------|-------------|-----------------------|-------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nova 2.0 | Fine-tuning | SFT Text + Datamixing | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-SFT-V2-DATAMIXING-latest | hyperpod start-job \ -n kubeflow \ --recipe fine-tuning/nova/forge/nova_2_0/nova_lite/SFT/nova_lite_2_0_p5_gpu_sft_t_ext_with_datamix \ --override-parameters '{ "instance_type": "ml.p5.48xlarge", "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-SFT-V2-DATAMIXING-latest" }' |



| Model    | Technique   | Subcategory | Image URI                                                                            | Hyperpod Launcher Command                                                                                                                                                                                                                                                         |
|----------|-------------|-------------|--------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nova 2.0 | Fine-tuning | SFT MM      | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-SFT-V2-latest | hyperpod start-job \ -n kubeflow \ --recipe fine-tuning/nova/nova_2_0/nova_lite/SFT/nova_lite_2_0_p5_gpu_sft \ --override-parameters '{ "instance_type": "ml.p5.48xlarge", "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-SFT-V2-latest" }' |

| Model    | Technique   | Subcategory         | Image URI                                                                                       | Hyperpod Launcher Command                                                                                                                                                                                                                                                                                          |
|----------|-------------|---------------------|-------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nova 2.0 | Fine-tuning | SFT MM + Datamixing | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-SFT-V2-DATAMIXING-latest | hyperpod start-job \ -n kubeflow \ --recipe fine-tuning/nova/forge/nova_2_0/nova_lite/SFT/nova_lite_2_0_p5_gpu_sft_mm_with_datamix \ --override-parameters '{ "instance_type": "ml.p5.48xlarge", "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-SFT-V2-DATAMIXING-latest" }' |

Once you're all setup, starting from the root of the sagemaker-hyperpod-cli repository, navigate to the default nova sft recipe folder

- `cd /src/hyperpod_cli/sagemaker_hyperpod_recipes/recipes_collection/recipes/training/nova/`

- Here you can choose whether you want to run nova 1 or nova 2 recipes based on the choice of base model.

For Nova 2.0 sft:

- If you would like to use a regular sft job , You should be able to see one recipe under this folder
  - `cd /src/hyperpod_cli/sagemaker_hyperpod_recipes/recipes_collection/recipes/fine-tuning/nova_2_0/nova_lite/SFT` and then you should be able to see one recipe under this folder called `nova_lite_2_0_p5x8_gpu_sft.yaml`
- If you would like to use datamixing sft Job, you can navigate to the sft Forge recipes folder
  - `cd /src/hyperpod_cli/sagemaker_hyperpod_recipes/recipes_collection/recipes/fine-tuning/nova/forge/nova_2_0/nova_lite/SFT` and you should be able to see one recipe under this folder called: `nova_lite_2_0_p5x8_gpu_sft_with_datamix.yaml`
- Edit the sections in the recipe required by the job such as name, data\_s3\_path, validation\_s3\_path, output\_s3\_path, and max\_steps. Since we're performing sft, the notion of epochs doesn't apply here.

The data mixing config will look the same, but with an extra data mixing section at the bottom similar to this

```
data_mixing:
 dataset_catalog: sft_text_lite
 sources:
 nova_data: # percent inputs for Nova data must sum to 100%; use 0% if you want to
exclude a data grouping
 agents: 20
 business-and-finance: 20
 scientific: 20
 code: 20
 factual-and-news: 20
 longform-text: 0
 health-and-medicine: 0
 humanities-and-education: 0
 legal: 0
 math: 0
 additional-languages: 0
 social-and-personal-interest: 0
 entertainment: 0
```

```

reasoning: 0
other: 0
tables: 0
customer_data: # percent input of customer data. 100 = use only customer data, 0 =
use only the nova_data mix above
percent: 25

```

There are two top-level categories of data here:

- **nova\_data** : This is the actual data mixing and is sub-divided into even more categories. It is imperative that they sum up to 100%
  - A complete breakdown of these categories including token count can be found in below
- **customer\_data** : This is your training data referred in the `data_s3_path` key at the top of your yaml. The percentage provided here determines what the resulting percentage will be for `nova_data`. For example, in the above percent selections, during training we'll use 25% of `customer_data` and 75% of `nova_data` of which 15% will be agents, 15% will be business-and-finance, 15% will be scientific, 15% will be code, and 15% will be factual-and-news

Tip: Run `pip install -e .` once again and you're ready to submit your job!

We'll be overriding a couple of parameters here to use data mixing:

```

hyperpod start-job \
 -n kubeflow \
 --recipe fine-tuning/nova/forged/nova_2_0/nova_lite/SFT/
nova_lite_2_0_p5x8_gpu_sft_with_datamix \
 --override-parameters '{
 "instance_type": "ml.p5.48xlarge",
 "recipes.run.name": "nova-sft-datamixing",
 "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-
SFT-V2-Datamix",
 "recipes.run.data_s3_path": "s3://sft-data/sft_train_data.jsonl",
 "recipes.run.validation_data_s3_path": "s3://sft-data/sft_val_data.jsonl",
 "recipes.run.output_s3_path": "s3://sft-data/output/
 }'

```

Your output should contain a job name as follows:

```
MY Desktop # % hyperpod start-job \
-n kubeflow \
--recipe training/nova/forged/nova_2_0/nova_lite/sft/
nova_lite_2_0_p5x8_gpu_pretrain_with_datamix \
--override-parameters '{
 "instance_type": "ml.p5.48xlarge",
 "recipes.run.name": "nova-sft-datamixing",
 "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-fine-tune-repo:SM-HP-
SFT-V2-Datamix",
 "recipes.run.data_s3_path": "s3://sft-data/sft_train_data.jsonl",
 "recipes.run.validation_data_s3_path": "s3://sft-data/sft_val_data.jsonl",
 "recipes.run.output_s3_path": "s3://sft-data/output/
}'
```

Output would be like this:

```
Final command: python3 /local/home/my/Downloads/sagemaker-hyperpod-cli/
src/hyperpod_cli/sagemaker_hyperpod_recipes/main.py recipes=training/nova/
nova_micro_p5x8_gpu_pretrain cluster_type=k8s cluster=k8s base_results_dir=/local/home/
niphaded/Downloads/sagemaker-hyperpod-cli/results cluster.pullPolicy="IfNotPresent"
cluster.restartPolicy="OnFailure" cluster.namespace="kubeflow"
instance_type="p5d.48xlarge" container="900867814919.dkr.ecr.us-east-1.amazonaws.com/
nova-fine-tune-repo:sft-datamix-rig-final"
Prepared output directory at /local/home/my/Downloads/sagemaker-hyperpod-cli/results/
my-sft-run-wzdyn/k8s_templates
Found credentials in shared credentials file: ~/.aws/credentials
Helm script created at /local/home/my/Downloads/sagemaker-hyperpod-cli/results/my-sft-
run-wzdyn/niphaded-sft-run-wzdyn_launch.sh
Running Helm script: /local/home/my/Downloads/sagemaker-hyperpod-cli/results/my-sft-
run-wzdyn/niphaded-sft-run-wzdyn_launch.sh

NAME: my-sft-run-wzdyn
LAST DEPLOYED: Tue Aug 26 16:21:06 2025
NAMESPACE: kubeflow
STATUS: deployed
REVISION: 1
TEST SUITE: None
Launcher successfully generated: /local/home/my/Downloads/sagemaker-hyperpod-cli/src/
hyperpod_cli/sagemaker_hyperpod_recipes/launcher/nova/k8s_templates/SFT

{
```

```

"Console URL": "https://us-east-1.console.aws.amazon.com/sagemaker/home?region=us-east-1#/cluster-management/hyperpod-eks-ga-0703"
}

```

You can view the status of your job using `hyperpod list-pods -n kubeflow --job-name my-sft-run-wzdyn`

```

hyperpod list-pods -n kubeflow --job-name my-sft-run-wzdyn
{
 "pods": [
 {
 "PodName": "my-sft-run-wzdyn-master-0",
 "Namespace": "kubeflow",
 "Status": "Pending",
 "CreationTime": "2025-08-26 16:21:06+00:00"
 },
 {
 "PodName": "my-sft-run-wzdyn-worker-0",
 "Namespace": "kubeflow",
 "Status": "Pending",
 "CreationTime": "2025-08-26 16:21:06+00:00"
 }
]
}

```

or directly use the `kubectl` command to find them.

For example,

```
kubectl get pods -o wide -w -n kubeflow | (head -n1 ; grep my-sft-run)
```

| NAME                                             | READY  | STATUS         |
|--------------------------------------------------|--------|----------------|
| RESTARTS AGE IP NODE                             |        | NOMINATED NODE |
| READINESS GATES                                  |        |                |
| my-sft-run-5suc8-master-0                        | 0/1    | Completed 0    |
| 3h23m 172.31.32.132 hyperpod-i-00b3d8a1bf25714e4 | <none> | <none>         |
| my-sft-run-5suc8-worker-0                        | 0/1    | Completed 0    |
| 3h23m 172.31.44.196 hyperpod-i-0aa7ccfc2bd26b2a0 | <none> | <none>         |
| my-sft-run-5suc8-worker-1                        | 0/1    | Completed 0    |
| 3h23m 172.31.46.84 hyperpod-i-026df6406a7b7e55c  | <none> | <none>         |

|                                                 |        |           |        |
|-------------------------------------------------|--------|-----------|--------|
| my-sft-run-5suc8-worker-2                       | 0/1    | Completed | 0      |
| 3h23m 172.31.28.68 hyperpod-i-0802e850f903f28f1 | <none> |           | <none> |

Pro tip : Make sure to always use the `-o wide` flag since the EKS node on which the job runs will help you find your logs even faster in the AWS UI

## How to Monitor Job

You can view your logs one of three ways:

### a) Using CloudWatch

Your logs are available in your Amazon Web Services account that contains the hyperpod cluster under CloudWatch. To view them in your browser, navigate to the CloudWatch homepage in your account and search for your cluster name. For example, if your cluster were called `my-hyperpod-rig` the log group would have the prefix:

- Log group : `/aws/sagemaker/Clusters/my-hyperpod-rig/{UUID}`
- Once you're in the log group, you can find your specific log using the node instance ID such as - `hyperpod-i-00b3d8a1bf25714e4`.
  - `i-00b3d8a1bf25714e4` here represents the hyperpod friendly machine name where your training job is running. Recall how in the previous command `kubectl get pods -o wide -w -n kubeflow | (head -n1 ; grep my-cpt-run)` output we captured a column called `NODE`.
  - The "master" node run was in this case running on `hyperpod-i-00b3d8a1bf25714e4` and thus we'll use that string to select the log group to view. Select the one that says `SagemakerHyperPodTrainingJob/rig-group/[NODE]`

Your logs should look something like this:

### b) Using CloudWatch Insights

If you have your job name handy and don't wish to go through all the steps above, you can simply query all logs under `/aws/sagemaker/Clusters/my-hyperpod-rig/{UUID}` to find the individual log.

CPT

```
fields @timestamp, @message, @logStream, @log
| filter @message like /(?!i)Starting CPT Job/
| sort @timestamp desc
```

```
| limit 100
```

For job completion replace Starting SFT Job with SFT Job completed

Then you can click through the results and pick the one that says "Epoch 0" since that will be your master node.

## C) Using the AWS CLI

You may choose to tail your logs using the AWS CLI. Before doing so, please check your AWS CLI version using `aws --version`. It is also recommended to use this utility script that helps in live log tracking in your terminal

for V1:

```
aws logs get-log-events \
 --log-group-name /aws/sagemaker/YourLogGroupName \
 --log-stream-name YourLogStream \
 --start-from-head | jq -r '.events[].message'
```

for V2:

```
aws logs tail /aws/sagemaker/YourLogGroupName \
 --log-stream-name YourLogStream \
 --since 10m \
 --follow
```

## D) Set up ML Flow:

You can track metrics via MLFlow.

Create an MLflow app

Using Studio UI: If you create a training job through the Studio UI, a default MLflow app is created automatically and selected by default under Advanced Options.

Using CLI: If you use the CLI, you must create an MLflow app and pass it as an input to the training job API request.



```
mlflow_app_name="<enter your MLflow app name>"
role_arn="<enter your role ARN>"
bucket_name="<enter your bucket name>"
region="<enter your region>"

mlflow_app_arn=$(aws sagemaker create-mlflow-app \
 --name $mlflow_app_name \
 --artifact-store-uri "s3://$bucket_name" \
 --role-arn $role_arn \
 --region $region)
```

## Access the MLflow app

Using CLI: Create a presigned URL to access the MLflow app UI:

```
aws sagemaker create-presigned-mlflow-app-url \
 --arn $mlflow_app_arn \
 --region $region \
 --output text
```

Once ML Flow is setup, you can pass the URI in your recipe or use override when starting the job. One example of how to do that can be found in README.

## How to evaluate your model after SFT?

### Prerequisites

- Checkpoint S3 URI from your training job's manifest.json file (for trained models)
- Evaluation dataset uploaded to S3 in the correct format
- Output S3 path for evaluation results

Out of the box benchmarks: Use out of the box benchmarks to validate the performance on general tasks. For more details check [here](#).

### Bring your own data:

You can also supply your custom data by formatting them in below format and then using below mentioned containers to get inference results along with log probabilities for calibrations if needed.

Crate jsonl per task with the following structure:

```
{
 "metadata": "{key:4, category:'apple'}",
 "system": "arithmetic-patterns, please answer the following with no other words: ",
 "query": "What is the next number in this series? 1, 2, 4, 8, 16, ?",
 "response": "32"
}
```

Outputs generated during inference phase of the evaluation job will have following structure:

```
{
 "prompt": "[{'role': 'system', 'content': 'arithmetic-patterns, please answer the following with no other words: '}, {'role': 'user', 'content': 'What is the next number in this series? 1, 2, 4, 8, 16, ?'}]",
 "inference": "['32']",
 "gold": "32",
 "metadata": "{key:4, category:'apple'}"
}
```

Field descriptions:

- prompt: Formatted input sent to the model
- inference: Model's generated response
- gold: Expected correct answer from input dataset, response field from the input
- metadata: Optional metadata passed through from input

## Prepare Evaluation Config

Command to launch evaluation job. Use "--override-parameters" to modify any entry from the recipe.

```
hyperpod start-job -n kubeflow \
 --recipe evaluation/nova/nova_micro_p5_48xl_bring_your_own_dataset_eval \
 --override-parameters '{
 "instance_type": "p5.48xlarge",
```

```
"container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM-HP-Eval-latest",
 "recipes.run.name": "<your-eval-job-name>",
 "recipes.run.model_name_or_path": "<checkpoint-s3-uri>",
 "recipes.run.output_s3_path": "s3://<your-bucket>/eval-results/",
 "recipes.run.data_s3_path": "s3://<your-bucket>/eval-data.jsonl"
}'
```

## Launch Your Evaluation Job

Job launching commands for different recipes with corresponding images.

### Evaluation Job Launch Commands

| Model    | Technique  | Subcategory | Image URI                                                                           | Command                                                                                                                                                                                                                                                           |
|----------|------------|-------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nova 2.0 | Evaluation | Eval        | 708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM-HP-Eval-latest | hyperpod start-job -n kubeflow \ --recipe evaluation/nova/nova_2_0/nova_lite/nova_lite_2_0_p5_48_xl_gpu_ft_eval \ --override-parameters '{ "instance_type": "ml.p5.48xlarge", "container": "708977205387.dkr.ecr.us-east-1.amazonaws.com/nova-evaluation-repo:SM- |

| Model | Technique | Subcategory | Image URI | Command                |
|-------|-----------|-------------|-----------|------------------------|
|       |           |             |           | HP-Eval-lat<br>est" }' |

## Lessons Learned and Tips

- The quality of the SFT dataset is critical. You should make every effort to filter out low-quality data. If you have a small subset of exceptionally high-quality data—in terms of both complexity and accuracy—you may consider placing it toward the end of training to help the model converge better.
- We leverage both text and multimodal (MM) datasets for data mixing. Our experiments with text dataset show that adding Nova's proprietary "reasoning-instruction-following" category significantly improves performance across generic benchmarks. We recommend including this category in your data mixing strategy if you care about generic benchmark that is regressed after you did SFT with your datasets.
- For MM datasets, our experiments indicate that incorporating over 20% of video categories into the mix is beneficial for maintaining generic benchmark performance.
- Further, SFT with data mixing is quite sensitive to learning rate so our finding suggests to fine-tune with the default learning rate i.e. 1e-5 for LoRA and 5e-6 for FR.
- Finally, there is a trade off between latency and performance if you mix Nova proprietary datasets so our findings suggest to mix 50% in max as a good balance.

## Reinforcement Learning

 **Note**

Detailed documentation is provided once subscribed

Nova Forge provides advanced reinforcement learning capabilities with the option to use remote reward functions in your own environment. Customers can choose to integrate their own endpoint to execute validation for immediate real-world feedback, or even use their own orchestrator to coordinate agentic multi-turn evaluations in your environment.

# Bring your own orchestrator for agentic multi-turn evaluations

For Forge users requiring multi-turn conversations or reward functions exceeding 15-minute timeouts, Nova Forge provides Bring Your Own Orchestration (BYOO) capabilities. This allows you to coordinate agentic multi-turn evaluations in your environment (e.g., using chemistry tools to score molecular designs, or robotics simulations that reward efficient task completion and penalize collisions).

## Topics

- [Architecture overview](#)
- [Setup and execution](#)

## Architecture overview

The BYOO architecture provides full control over the rollout and generation process through customer-managed infrastructure.

### Training VPC:

- **Rollout:** Coordinates training by delegating rollout generation to customer infrastructure
- **Trainer:** Performs model weight updates based on received rollouts

### Customer VPC (such as ECS on EC2):

- **Proxy Lambda:** Receives rollout requests and coordinates with customer infrastructure
- **Rollout Response SQS:** Queue for returning completed rollouts to training infrastructure
- **Generate Request SQS:** Queue for model generation requests
- **Generate Response SQS:** Queue for model generation responses
- **Customer Container:** Implements custom orchestration logic (can use provided starter kit)
- **DynamoDB:** Stores and retrieves state across the orchestration process

### Workflow:

1. Rollout delegates rollout generation to Proxy Lambda
2. Proxy Lambda pushes rollout API request to Generate Request SQS

3. Customer container processes requests, manages multi-turn interactions, and calls reward functions
4. Container stores and retrieves state from DynamoDB as needed
5. Container pushes rollout responses to Rollout Response SQS
6. Rollout sends completed rollouts to Trainer for weight updates

## Setup and execution

For detailed setup instructions, recipe configurations, request and response formats, and environment examples, refer to the confidential documentation provided to Nova Forge subscribers. To get the Nova Forge documents follow the below steps:

```
aws s3 cp s3://nova-forge-c7363-206080352451-us-east-1/v1/ ./ --recursive
```

Once the assets are downloaded, you can find all the documentation under the docs folder.

## Responsible AI Toolkit and content moderation

### Responsible AI toolkit

Nova Forge provides a Responsible AI Toolkit that includes training and evaluation data to align models to Amazon Nova's responsible AI guidelines during the training process, and runtime controls to moderate model responses during inference.

**Training data** – Cases and scenarios emphasizing responsible AI principles, safety considerations, and responsible technology deployment are available for data mixing to align models responsibly during continued pre-training.

**Evaluations** – Evaluations testing the model's ability to detect and reject inappropriate, harmful, or incorrect content are available as a benchmark task to determine the delta between base model performance and custom model performance.

**Runtime controls** – By default, Amazon Nova's runtime controls moderate model responses during inference. To modify these runtime controls, request Amazon Nova's Customizable Content Moderation Settings by contacting an AWS account manager.

Safety is a shared responsibility between AWS and its users. Changing the base model or using continued pre-training to improve performance on a specific use case can impact safety, fairness,

and other properties of the new model. A robust adaptation method minimizes changes to the safety, fairness, and other protections built into base models while minimizing impact on model performance for tasks the model was not customized for. End-to-end testing of applications on datasets representative of use cases is required to determine if test results meet specific expectations of safety, fairness, and other properties, as well as overall effectiveness. For more information, see [Amazon Web Services Responsible Use of AI Guide](#), [Amazon Web Services Responsible AI Policy](#), [Amazon Web Services Acceptable Use Policy](#), and [Amazon Web Services Service Terms](#).

# Responsible use

Building safety, security and trust measures with AI models is a shared responsibility between AWS and our customers. Our goal is to align our models to the [AWS Acceptable Use Policy](#) and mitigate undesired outcomes while providing a delightful customer experience. Our approach to Responsible AI (RAI) is structured around our [core dimensions of responsible AI](#), which are covered in the following list. For each of these dimensions, we developed guidelines that govern our decision-making throughout the entire model development life cycle. This life cycle encompasses every stage, from initial data collection and pre-training, to the implementation of post-deployment runtime mitigations.

- *Fairness* - Considering impacts on different groups of stakeholders
- *Explainability* - Understanding and evaluating system outputs
- *Privacy and Security* - Appropriately obtaining, using and protecting data and models
- *Safety* - Preventing harmful output and misuse
- *Controllability* - Having mechanisms to monitor and steer AI system behavior
- *Veracity and robustness* - Achieving correct system outputs, even with unexpected or adversarial inputs
- *Governance* - Incorporating best practices into the AI supply chain, including providers and deployers
- *Transparency* - Enabling stakeholders to make informed choices about their engagement with an AI system

## Guidelines

The guidelines we use to direct our model development includes but is not limited to moderating content that glorifies, facilitates, or promotes the following:

- Participation in dangerous activities, self harm, or use of dangerous substances.
- Use, misuse, or trade of controlled substances, tobacco, or alcohol.
- Physical violence or gore.
- Child abuse or child sexual abuse material.
- Animal abuse or animal trafficking.



- Misinformation that positions individuals or groups as responsible for deliberate deception, undermining an institution with general public credibility, or endangering human health or livelihood.
- Malware, malicious content, or any content that facilitates cyber-crime.
- Disrespect, discrimination, or stereotype towards an individual or group.
- Insults, profanity, obscene gestures, sexually explicit language, pornography, hate symbols, or hate groups.
- Full nudity that is outside of a scientific, educational, or reference context.
- Bias against a group based on a demographic characteristic.

## Recommendations

**Appropriateness for Use:** Because AI model outputs are probabilistic, Amazon Nova may produce inaccurate or inappropriate content. Customers should evaluate outputs for accuracy and appropriateness for their use case, especially if they will be directly surfaced to end users. Additionally, if Amazon Nova is used in customer workflows that produce consequential decisions, customers must evaluate the potential risks of their use case and implement appropriate human oversight, testing and other use-case specific safeguards to mitigate such risks.

**Prompt Optimizations:** In the event of encountering moderation by Amazon Nova, consider examining the prompts used with respect to the guidelines above. Optimizing the prompts to reduce the likelihood of generating undesired outcomes is the recommended strategy to produce the expected outputs using Amazon Nova models. Pay attention where the input is controlled by users, including pixel content that could impact the performance of the model. Please see the prompt guidelines section in this user guide for further details.

**Privacy:** Amazon Nova is available in Amazon Bedrock. Amazon Bedrock is a managed service and does not store or review customer prompts or customer prompt completions and prompts and completions are never shared between customers, or with Amazon Bedrock partners. AWS does not use inputs or outputs generated through the Amazon Bedrock service to train Amazon Bedrock models, including Amazon Nova. See Section [50.3](#) of the AWS Service Terms and the [AWS Data Privacy FAQ](#) for more information. For service-specific privacy information, see the Privacy and Security section of the [Amazon Bedrock FAQs](#) documentation. Amazon Nova takes steps to avoid completing prompts that could be construed as requesting private information. If a user is concerned that their private information has been included in a Amazon Nova completion, the user should contact us [here](#).

**Security:** All Amazon Bedrock models, including Amazon Nova, come with enterprise security that enables customers to build generative AI applications that support common data security and compliance standards, including GDPR and HIPAA. Customers can use AWS PrivateLink to establish private connectivity between customized Amazon Nova and on-premise networks without exposing customer traffic to the internet. Customer data is always encrypted in transit and at rest and customers can use their own keys to encrypt the data, e.g., using AWS Key Management Service. Customers can use AWS Identity and Access Management to securely control access to Amazon Bedrock resources, including customized Amazon Nova. Also, Amazon Bedrock offers comprehensive monitoring and logging capabilities that can support customer governance and audit requirements. For example, Amazon CloudWatch can help track usage metrics that are required for audit purposes and AWS CloudTrail can help monitor API activity and troubleshoot issues as Amazon Nova is integrated with other AWS systems. Customers can also choose to store the metadata, prompts and completions in their own encrypted Amazon Simple Storage Service (Amazon S3) bucket.

**Intellectual Property:** AWS offers uncapped intellectual property (IP) indemnity coverage for outputs of generally available Amazon Nova models (see Section 50.10 of the [Service Terms](#)). This means that customers are protected from third-party claims alleging IP infringement or misappropriation (including copyright claims) by the outputs generated by these Amazon Nova models. In addition, our standard IP indemnity for use of the Services protects customers from third-party claims alleging IP infringement (including copyright claims) by the Services (including Amazon Nova models) and the data used to train them.

# Quotas for Amazon Nova

Your AWS account has default quotas, formerly referred to as limits, for Amazon Bedrock. To view service quotas for Amazon Bedrock, do one of the following:

- Follow the steps at [Viewing service quotas](#) and select **Amazon Bedrock** as the service.
- Refer to [Amazon Bedrock endpoints and quotas](#) in the AWS General Reference.

## Note

If a quota is marked as not adjustable through Service Quotas, you can submit a request through the [limit increase form](#) to be considered for an increase.

To maintain the performance of the service and to ensure appropriate usage of Amazon Bedrock, the default quotas assigned to an account might be updated depending on regional factors, payment history, fraudulent usage, and/or approval of a quota increase request.

## Request an increase for Amazon Nova quotas

You can request a quota increase for your account by following the steps below:

- If a quota is marked as **Yes** in the **Adjustable** column in [Amazon Bedrock endpoints and quotas](#) in the AWS General Reference, you can adjust it by following the steps at [Requesting a Quota Increase](#) in the *Service Quotas User Guide* in the [Service Quotas User Guide](#).
- The **On-demand model invocation** quotas in [Amazon Bedrock endpoints and quotas](#) in the AWS General Reference aren't adjustable through Service Quotas. Contact your AWS account manager to be considered for an increase.

## Note

Due to overwhelming demand, priority will be given to customers who generate traffic that consumes their existing quota allocation. Your request might be denied if you don't meet this condition.

- You can submit a request through the [limit increase form](#) to be considered for an increase even if a quota is marked as **No** in the **Adjustable** column in [Amazon Bedrock endpoints and quotas](#) in the AWS General Reference,

# Monitoring Amazon Nova

Monitoring is an important part of maintaining the reliability, availability and performance of Amazon Nova and your other AWS solutions. AWS provides the following monitoring tools to watch Amazon Nova, report when something is wrong and take automatic actions when appropriate:

Because Amazon Nova is available through Amazon Bedrock, follow all monitoring guidance at [Monitor the health and performance of Amazon Bedrock](#).

# Security in Amazon Nova

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to Amazon Nova, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements and applicable laws and regulations.

Because Amazon Nova is available through Amazon Bedrock, follow all security guidance at [Security in Amazon Bedrock](#).

# Create Amazon Nova resources with AWS CloudFormation

When deploying Amazon Nova models through Amazon Bedrock with CloudFormation, follow all CloudFormation guidance at [Create Amazon Bedrock resources with CloudFormation](#).

For more information, including examples of JSON and YAML templates for YOUR-RESOURCE-PLURAL, see the [YOUR-SERVICE-RESOURCE-TOPNODE-IN-CLOUDFORMATION, modify this link in the local-phrases.ent file](#) in the *AWS CloudFormation User Guide*.

## Learn more about CloudFormation

To learn more about CloudFormation, see the following resources:

- [AWS CloudFormation](#)
- [AWS CloudFormation User Guide](#)
- [CloudFormation API Reference](#)
- [AWS CloudFormation Command Line Interface User Guide](#)

# Document history for the Amazon Nova Developer Guide

The following table describes the documentation releases for Amazon Nova 2.0.

| Change                                             | Description                                                                                                                                                                                                                              | Date             |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| <a href="#">Initial release of Amazon Nova 2.0</a> | Initial release of the Amazon Nova 2.0 Developer Guide with new models: Nova 2 Lite, Amazon Nova Sonic and . This guide introduces enhanced multimodal capabilities, reasoning features and improved performance across all model tiers. | December 2, 2025 |