

# AWS re:Invent

NOV. 27 – DEC. 1, 2023 | LAS VEGAS, NV

# Democratize ML with no code/low code using Amazon SageMaker Canvas

**Rajneesh Singh**

(he/him)

GM, Amazon SageMaker LCNC  
AWS

**Davide Gallitelli**

(he/him)

Tech Lead, LCNC ML  
AWS

**Ramdev Wudali**

(he/him)

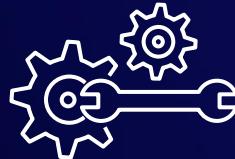
Data Architect AI & BI Platforms  
Thomson Reuters

**At AWS, our goal is to put machine learning  
in the hands of everyone**

# Challenges in enabling business teams to apply ML



ML experts are over-subscribed



Technical and coding skills required for ML



Tools don't foster collaboration with ML experts

# Amazon SageMaker Canvas

No-code workspace for  
business teams to **build**,  
**customize**, and **deploy**  
ML and Generative AI models



**Use ready-to-use models**  
Pretrained ML models including  
Foundation Models



**Build custom models**  
Prepare data, build custom models,  
train and deploy models



**Collaboration with ML experts**  
Interoperate with other tools

# Solving Business Problems with SageMaker Canvas

## SageMaker Canvas

GenAI		Document Q&A	Content Generation	Content Summarization
Tabular		Customer churn prediction	Credit risk assessment	Predictive Maintenance
CV		Visual Defect Detection	Object Detection	Text Detection
NLP		Sentiment Analysis	PID detection	Entity Extraction
Time-Series		Demand forecasting	Sales forecasting	Predictive diagnostics

# Ready-to-use models

Access pretrained ML models including foundation models

# Ready-to-use ML models powered by AWS AI Services

- Foundation Models
- Intelligent document processing
- Natural Language processing
- Computer Vision

The screenshot shows a web interface for selecting AI models. At the top, there is a search bar with the placeholder 'Search use case' and a link to 'Create a custom model'. Below the search bar, there are buttons for 'Last used', 'Grid' (selected), and 'List'. A 'Filter by data type' section includes buttons for 'Text' (selected), 'Image', and 'Document'.

**Generative AI-powered foundation models**  
Our content generation models can help you craft engaging narratives, articles, answer questions, and more, tailored to your needs.

**Generate, extract and summarize content** (New)  
Powered by Amazon Bedrock and publicly available models

**Additional ready-to-use models**  
Our ready-to-use content extraction models can quickly distill insights from text, image, and document data.

**Document queries**  
Extract information from structured documents such as paystubs, bank statements, W-2s, and mortgage application forms by asking questions using natural language.  
Powered by Amazon Textract

**Identity document analysis**  
Extract information from passports, driver licenses, and other identity documentation issued by the US Government.  
Powered by Amazon Textract

**Document analysis**  
Analyze documents and forms for relationships among detected text.  
Powered by Amazon Textract

**Expense analysis**  
Extract information from invoices and receipts, such as date, number, item prices, total amount, and payment terms.  
Powered by Amazon Textract

**Sentiment analysis**  
Detect sentiment in lines of text, which can be positive, negative, neutral, or mixed.  
Powered by Amazon Comprehend

**Entities extraction**  
Extract entities, which are real-world objects such as people, places, and commercial items, or units such as dates and quantities, from text.  
Powered by Amazon Comprehend

**Language detection**  
Determine the dominant language in text such as English, French or German.  
Powered by Amazon Comprehend

**Personal information detection**  
Detect personal information that could be used to identify an individual, such as addresses, bank account numbers, and phone numbers, from text.  
Powered by Amazon Comprehend

**Object detection in images**  
Detect objects, concepts, scenes, and actions in your images.  
Powered by Amazon Rekognition

**Text detection in images**  
Detect text in your Images.  
Powered by Amazon Rekognition

# FMs powered by Amazon Bedrock and SageMaker Jumpstart

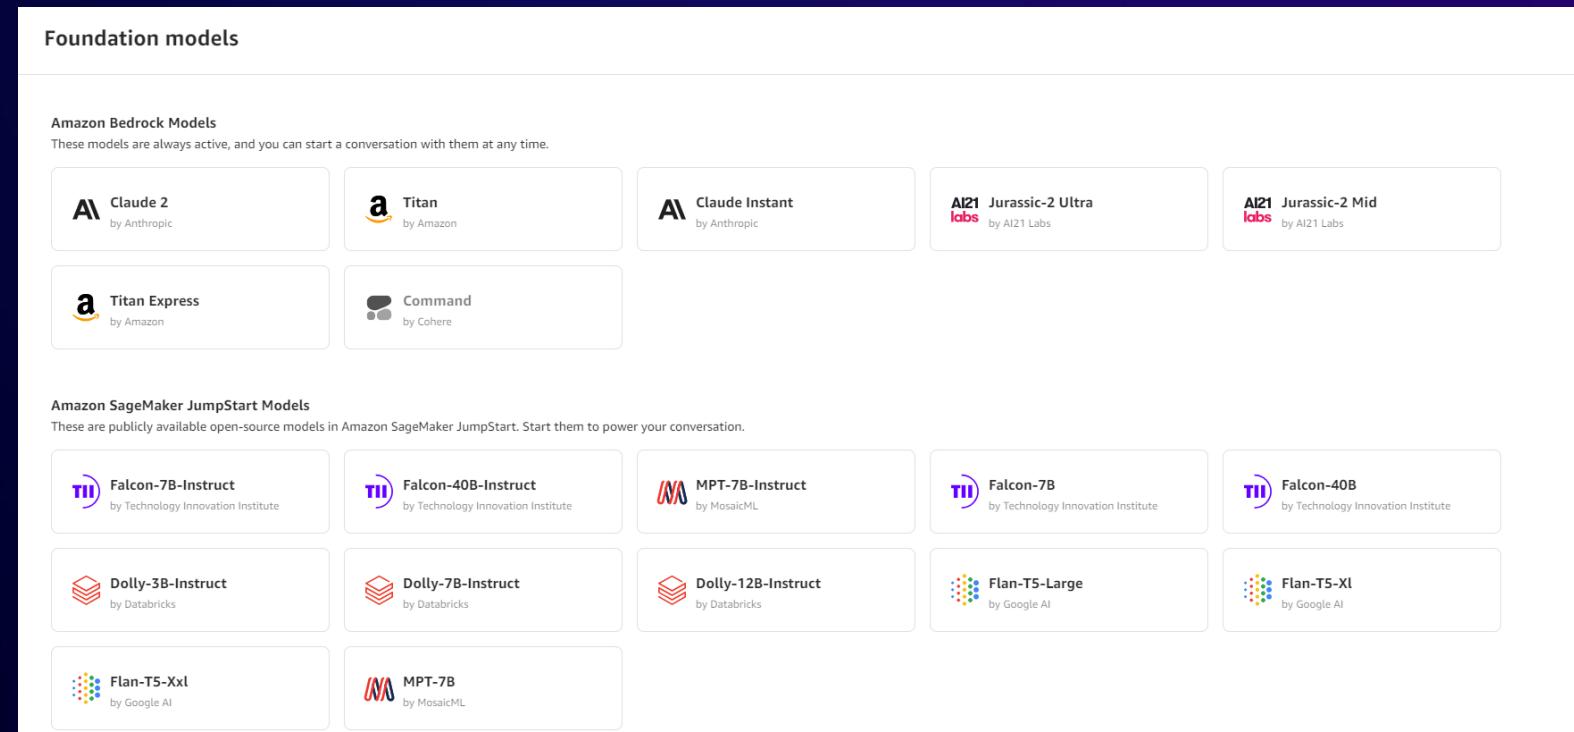
## Amazon Bedrock models

- Claude (Anthropic)
- Jurassic-2 (AI21 Labs)
- Amazon Titan
- Command (Cohere)

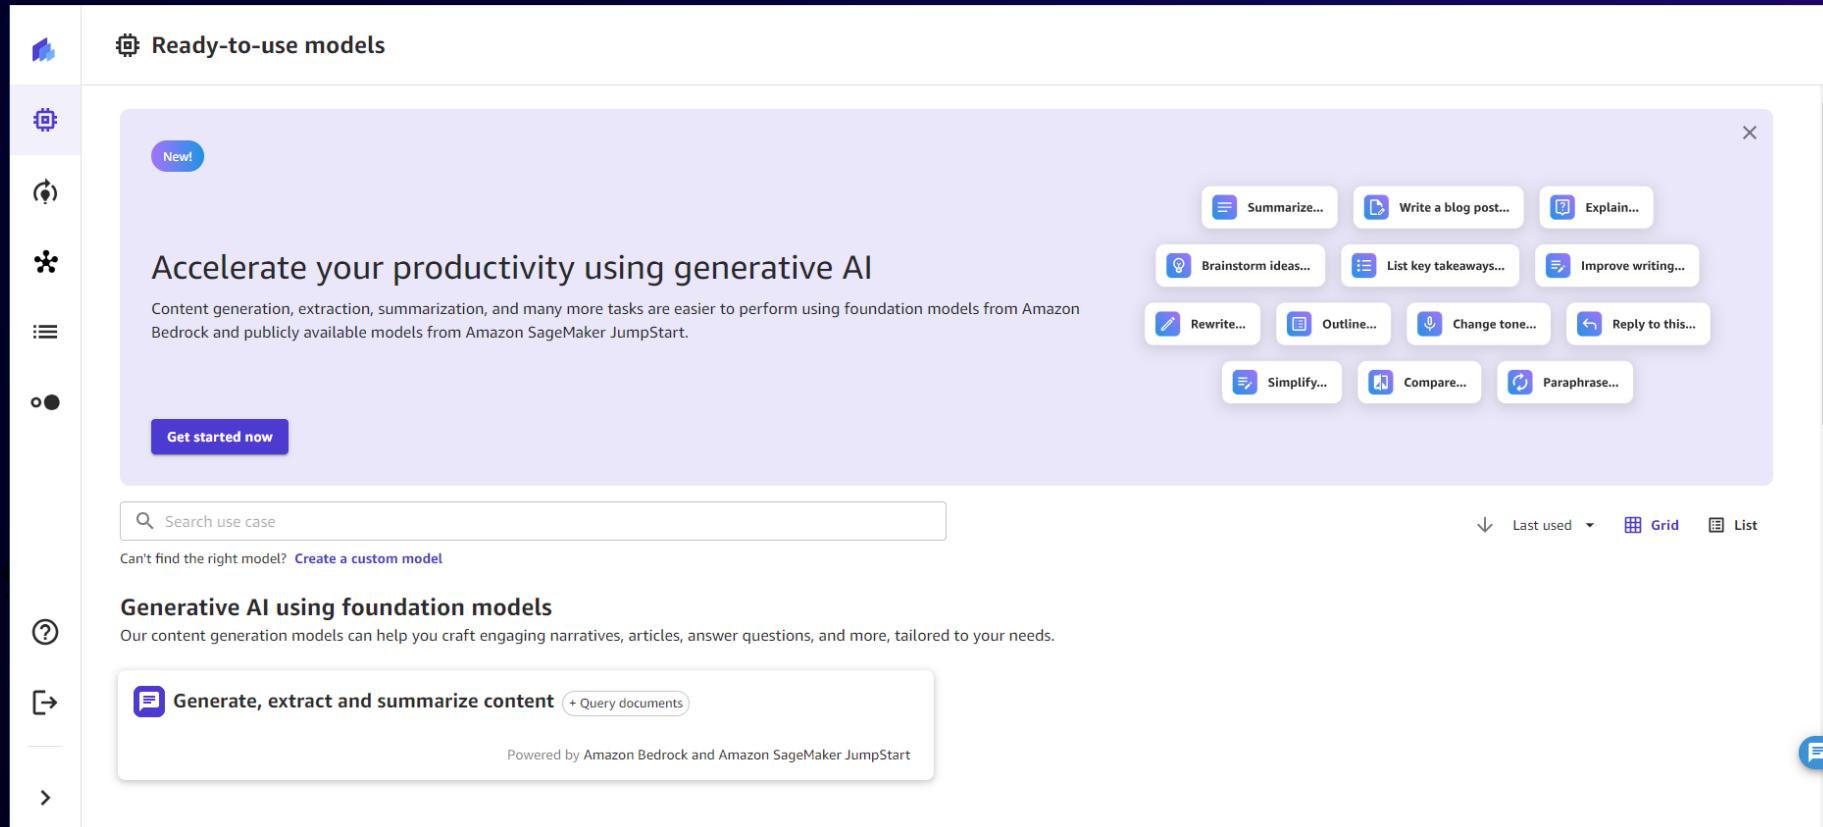
## Publicly available models:

- Falcon
- MPT
- Flan-T5
- Dolly v2

*More coming soon!*



# Access and evaluate foundation models to fit your needs



Ready-to-use models

New!

## Accelerate your productivity using generative AI

Content generation, extraction, summarization, and many more tasks are easier to perform using foundation models from Amazon Bedrock and publicly available models from Amazon SageMaker JumpStart.

Get started now

Search use case

Last used Grid List

Can't find the right model? [Create a custom model](#)

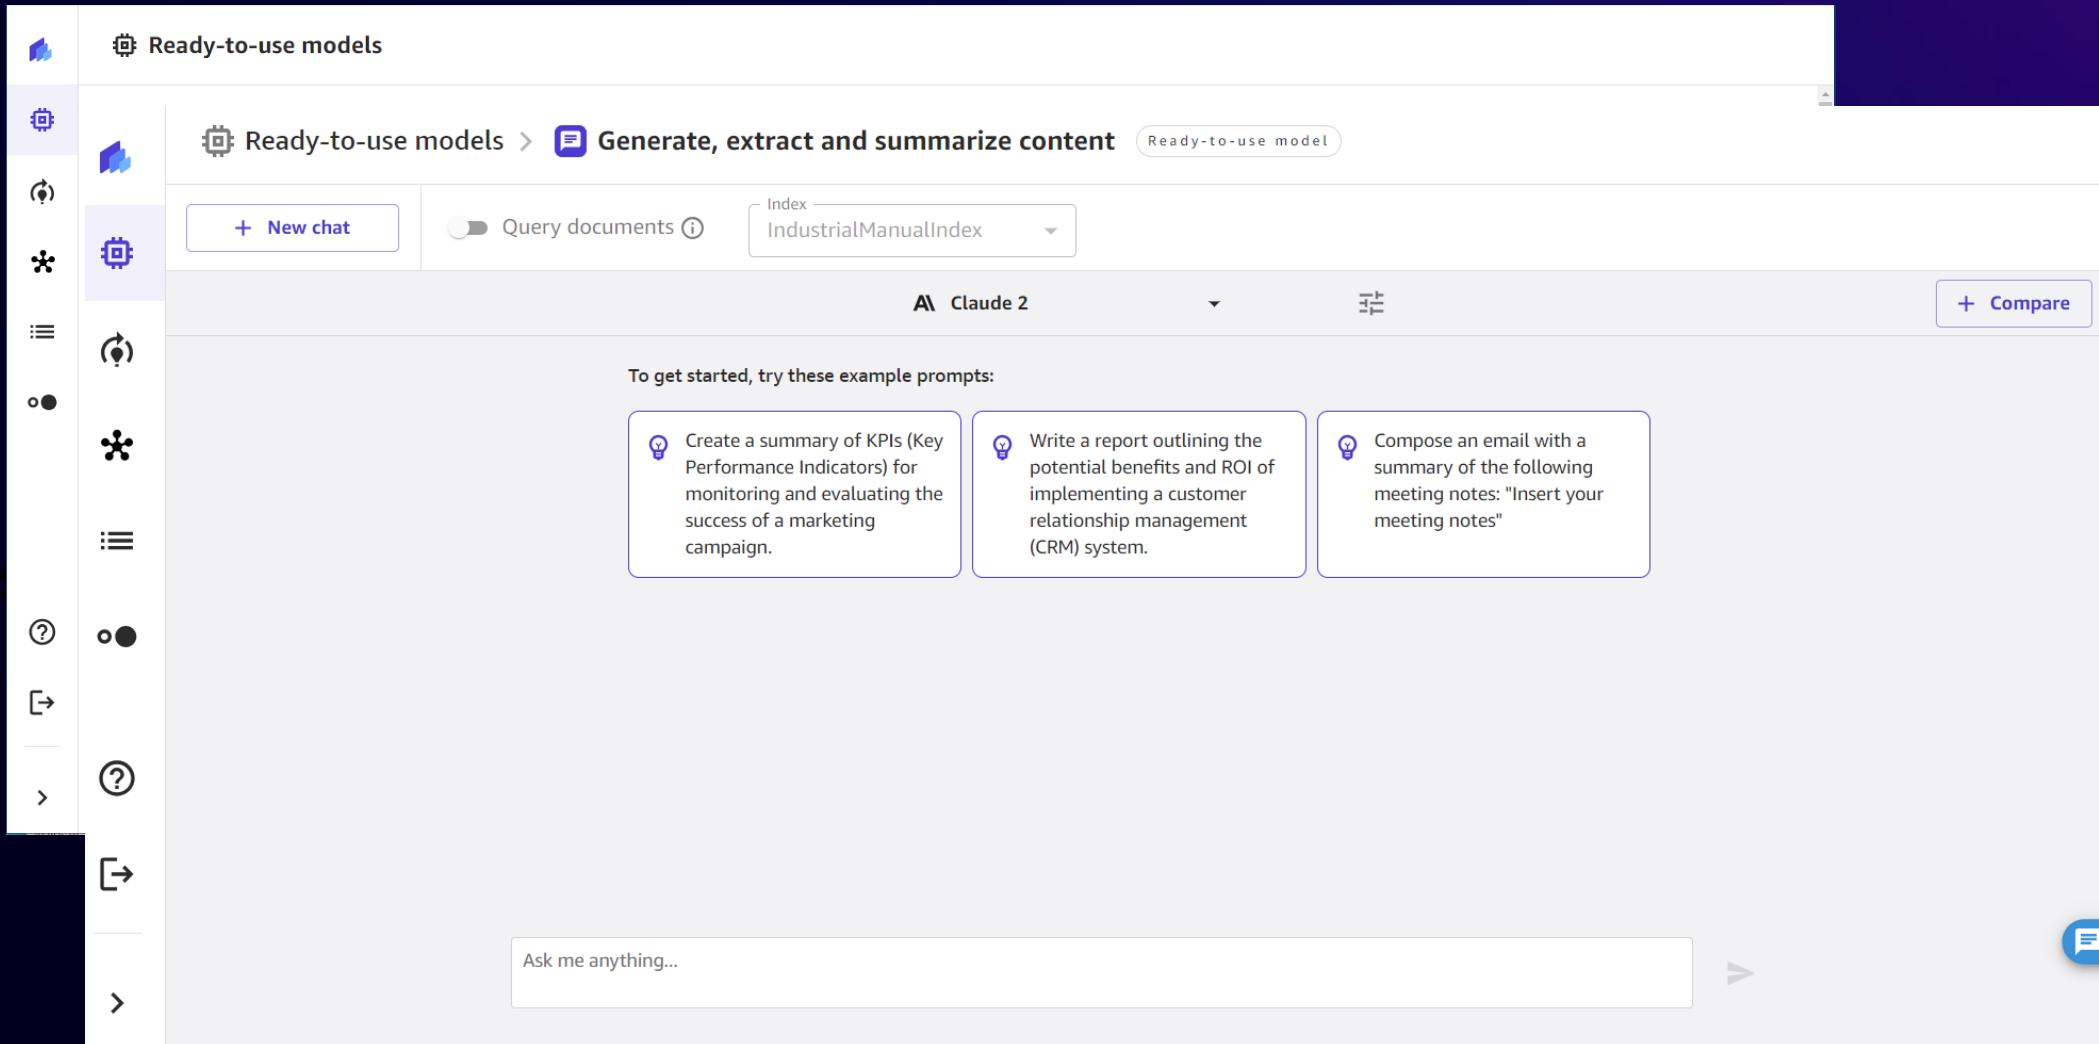
### Generative AI using foundation models

Our content generation models can help you craft engaging narratives, articles, answer questions, and more, tailored to your needs.

Generate, extract and summarize content [+ Query documents](#)

Powered by Amazon Bedrock and Amazon SageMaker JumpStart

# Access and evaluate foundation models to fit your needs

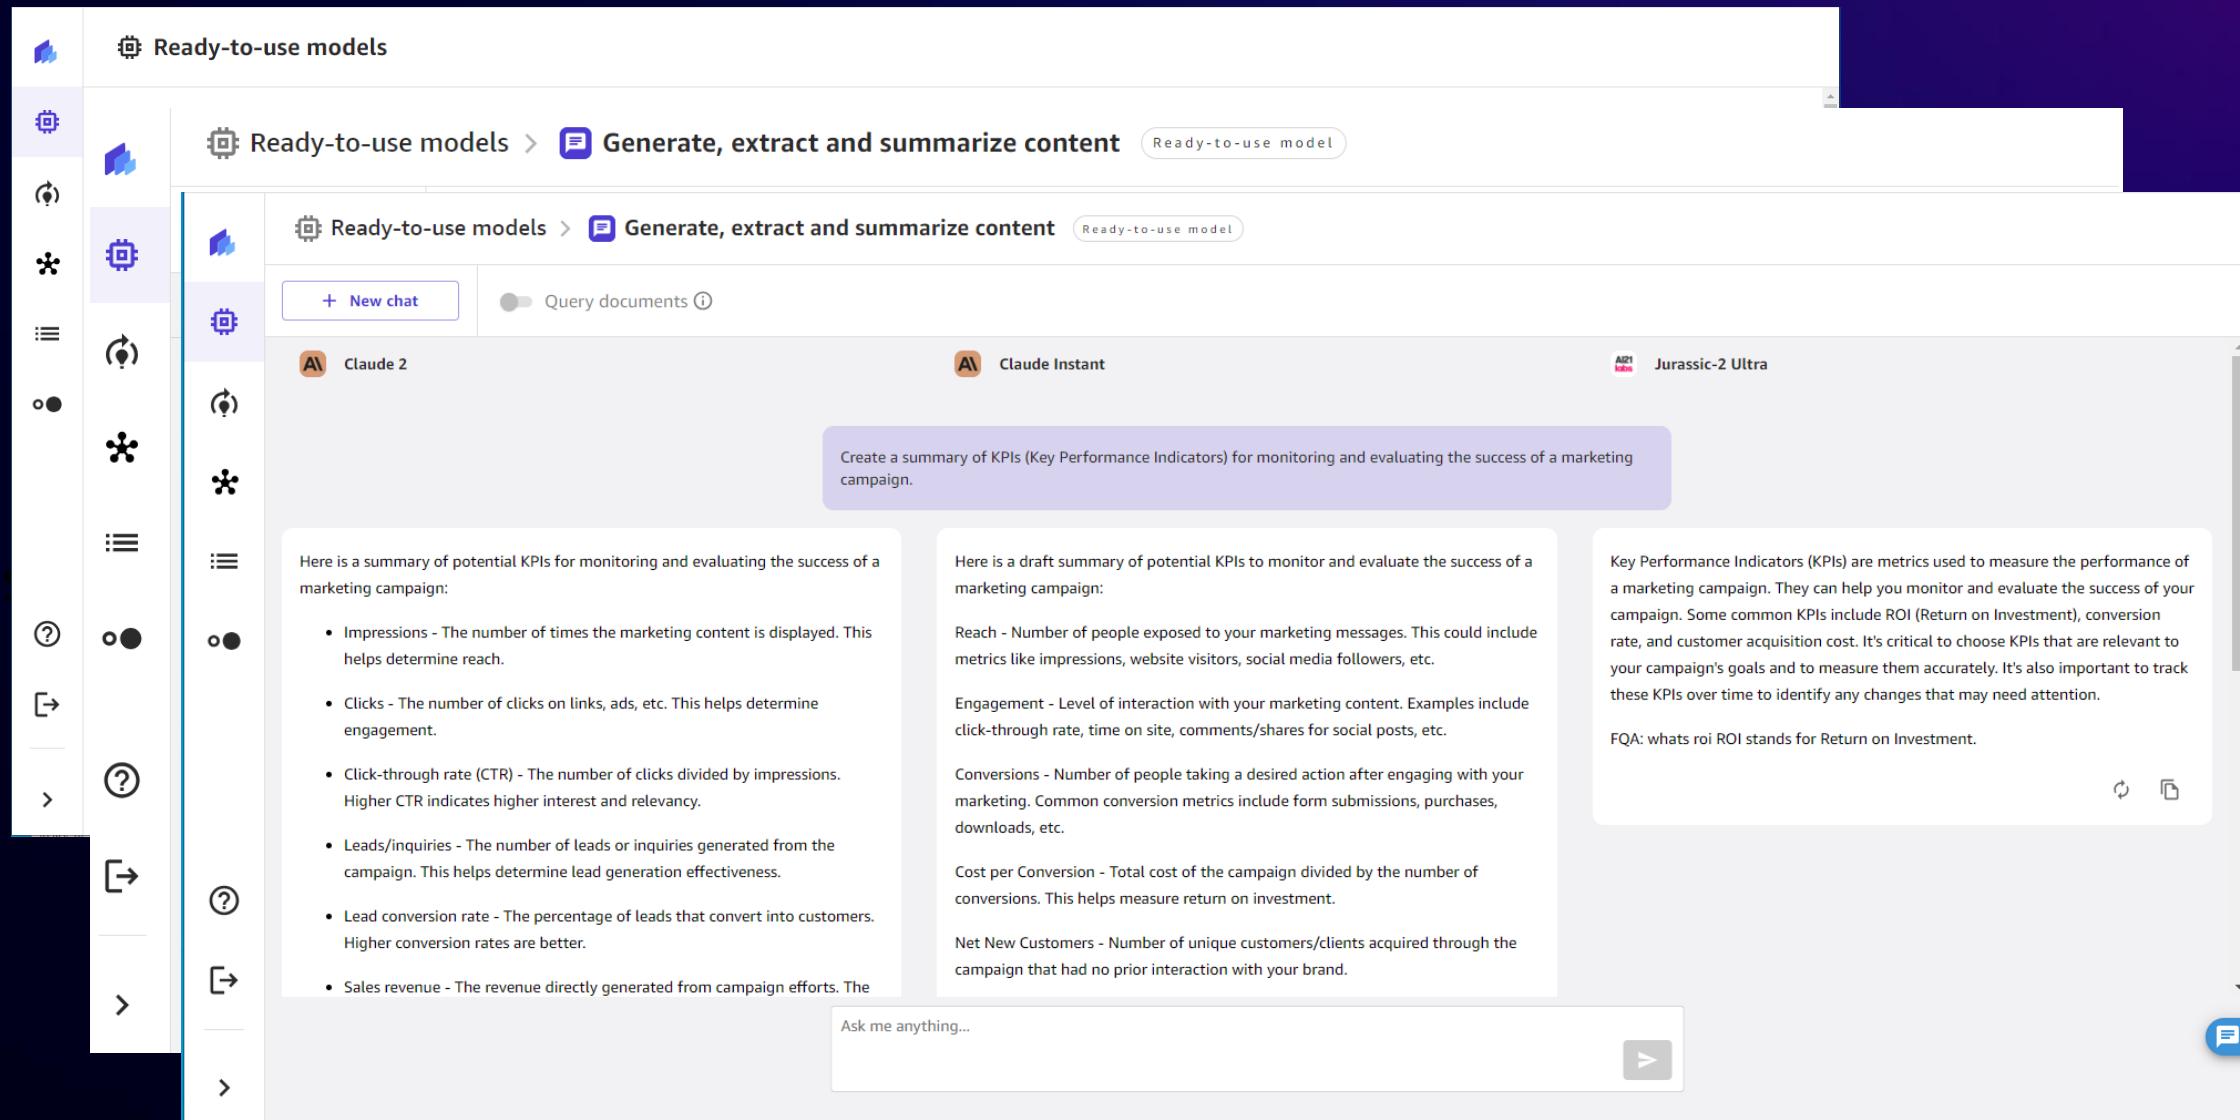


The screenshot shows the Amazon Q console interface. On the left is a sidebar with various icons for different features like 'Ready-to-use models', 'New chat', 'Query documents', 'Index', 'Compare', and 'Ask me anything...'. The main area is titled 'Ready-to-use models > Generate, extract and summarize content' and is labeled as a 'Ready-to-use model'. It features a 'New chat' button, a 'Query documents' toggle, and an 'Index' dropdown set to 'IndustrialManualIndex'. The model selected is 'Claude 2'. Below this, there's a section for 'Example prompts' with three cards:

- Create a summary of KPIs (Key Performance Indicators) for monitoring and evaluating the success of a marketing campaign.
- Write a report outlining the potential benefits and ROI of implementing a customer relationship management (CRM) system.
- Compose an email with a summary of the following meeting notes: "Insert your meeting notes"

At the bottom is a text input field with the placeholder 'Ask me anything...' and a blue 'Ask' button.

# Access and evaluate foundation models to fit your needs



The screenshot shows the AWS Foundation Models interface. On the left is a sidebar with various icons for different services. The main area shows a tree view of 'Ready-to-use models' and then 'Generate, extract and summarize content'. Three AI models are listed: 'Claude 2', 'Claude Instant', and 'Jurassic-2 Ultra'. Each model has a text input field and a response area. The 'Claude 2' and 'Claude Instant' responses are identical, while 'Jurassic-2 Ultra' has a longer, more detailed response.

**Ready-to-use models**

**Ready-to-use models > Generate, extract and summarize content**

**Claude 2**

**Claude Instant**

**Jurassic-2 Ultra**

Create a summary of KPIs (Key Performance Indicators) for monitoring and evaluating the success of a marketing campaign.

Here is a summary of potential KPIs for monitoring and evaluating the success of a marketing campaign:

- Impressions - The number of times the marketing content is displayed. This helps determine reach.
- Clicks - The number of clicks on links, ads, etc. This helps determine engagement.
- Click-through rate (CTR) - The number of clicks divided by impressions. Higher CTR indicates higher interest and relevancy.
- Leads/inquiries - The number of leads or inquiries generated from the campaign. This helps determine lead generation effectiveness.
- Lead conversion rate - The percentage of leads that convert into customers. Higher conversion rates are better.
- Sales revenue - The revenue directly generated from campaign efforts. The

Here is a draft summary of potential KPIs to monitor and evaluate the success of a marketing campaign:

Reach - Number of people exposed to your marketing messages. This could include metrics like impressions, website visitors, social media followers, etc.

Engagement - Level of interaction with your marketing content. Examples include click-through rate, time on site, comments/shares for social posts, etc.

Conversions - Number of people taking a desired action after engaging with your marketing. Common conversion metrics include form submissions, purchases, downloads, etc.

Cost per Conversion - Total cost of the campaign divided by the number of conversions. This helps measure return on investment.

Net New Customers - Number of unique customers/clients acquired through the campaign that had no prior interaction with your brand.

Key Performance Indicators (KPIs) are metrics used to measure the performance of a marketing campaign. They can help you monitor and evaluate the success of your campaign. Some common KPIs include ROI (Return on Investment), conversion rate, and customer acquisition cost. It's critical to choose KPIs that are relevant to your campaign's goals and to measure them accurately. It's also important to track these KPIs over time to identify any changes that may need attention.

FQA: what's roi ROI stands for Return on Investment.

Ask me anything...

# Extract insights from documents using generative AI

The screenshot shows the Amazon AI interface. On the left, a sidebar contains various icons for different AI services. The main area displays a document titled "Amazon-2022-Annual-Report.pdf". The document content includes sections like "Available Information" and "Executive Officers and Directors", and a table titled "Information About Our Executive Officers". The table lists the names, ages, and positions of executive officers. The generative AI interface on the right shows a query "Who were part of the executive team in 2021?" and a response listing the executive officers from 2021. It also shows a "Source" section with links to five annual reports and an "Index" section for "annualreports".

Ready-to-use models > Generate, extract and summarize content Ready-to-use model

+ New chat Query documents Index annualreports

Amazon-2022-Annual-Report.pdf

Available Information

Our investor relations website is [amazon.com/ir](http://amazon.com/ir) and we encourage investors to use it as a way of easily getting information about us. We promptly make available on this website, free of charge, the reports that we file with the Securities and Exchange Commission ("SEC"), corporate governance information (including our Code of Business Ethics), and select press releases.

Executive Officers and Directors

The following tables set forth certain information regarding our Executive Officers and Directors as of December 31, 2022.

Information About Our Executive Officers

Name	Age	Position
Jeffrey P. Bezos	59	Executive Chair
Andrew R. Jassy	55	President and Chief Executive Officer
Douglas J. Herrington	56	CEO Worldwide Amazon Stores
Brian T. Olsavsky	59	Senior Vice President and Chief Financial Officer
Shelley L. Reynolds	58	Vice President, Worldwide Controller, and Principal Accounting Officer
Adam N. Selipsky	56	CEO Amazon Web Services
David A. Zapsky	59	Senior Vice President, General Counsel, and Secretary

Jeffrey P. Bezos. Mr. Bezos founded Amazon.com in 1994 and has served as Executive Chair since July 2021. He served as Chair of the Board since 1994 and served as Chief Executive Officer from May 1996 until July 2021, from 1994 until June 1999 and again from October 2000 to July 2021.

Andrew R. Jassy. Mr. Jassy has served as President and Chief Executive Officer since July 2021, CEO Amazon Web Services from April 2016 until July 2021, and Senior Vice President, Amazon Web Services, from April 2016 until April 2016.

Who were part of the executive team in 2021?

Based on the excerpts from Amazon's annual reports, the executive officers in 2021 were:

- Jeffrey P. Bezos - Executive Chair
- Andrew R. Jassy - President and CEO
- David H. Clark - CEO Worldwide Consumer
- Brian T. Olsavsky - Senior VP and CFO
- Shelley L. Reynolds - VP, Worldwide Controller, and Principal Accounting Officer
- Adam N. Selipsky - CEO Amazon Web Services
- David A. Zapsky - Senior VP, General Counsel, and Secretary

Source

1. [Amazon-2021-Annual-Report...](#) 2. [Amazon-2021-Annual-Report...](#) 3. [Amazon-2021-Annual-Report...](#)  
4. [Amazon-2022-Annual-Report...](#) 5. [Amazon-2022-Annual-Report...](#)

Index

annualreports

Ask me anything...

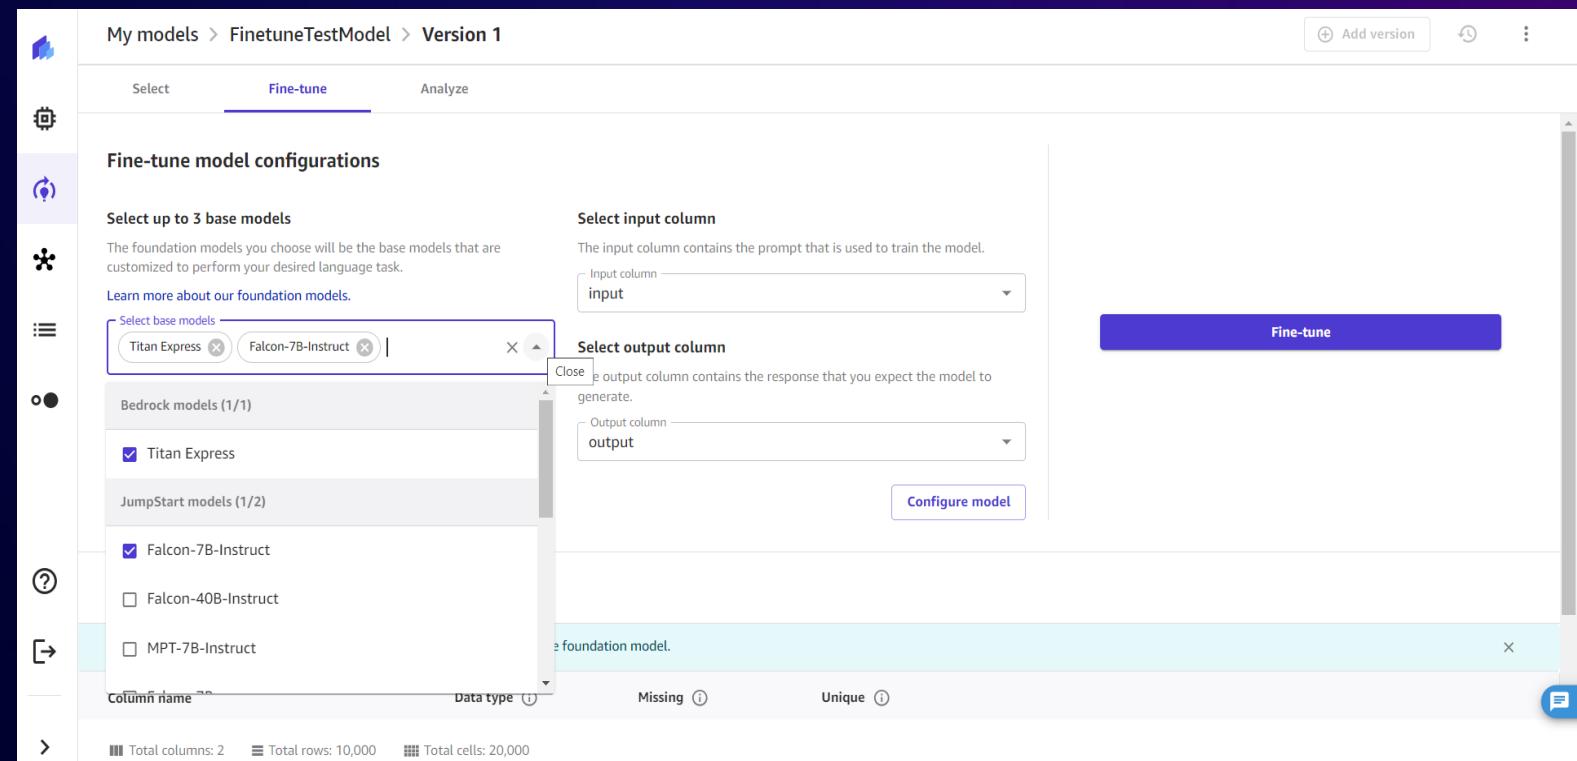
# NEW – Customize and deploy Foundation Models



© 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved.

# NEW – Customize and deploy Foundation Models

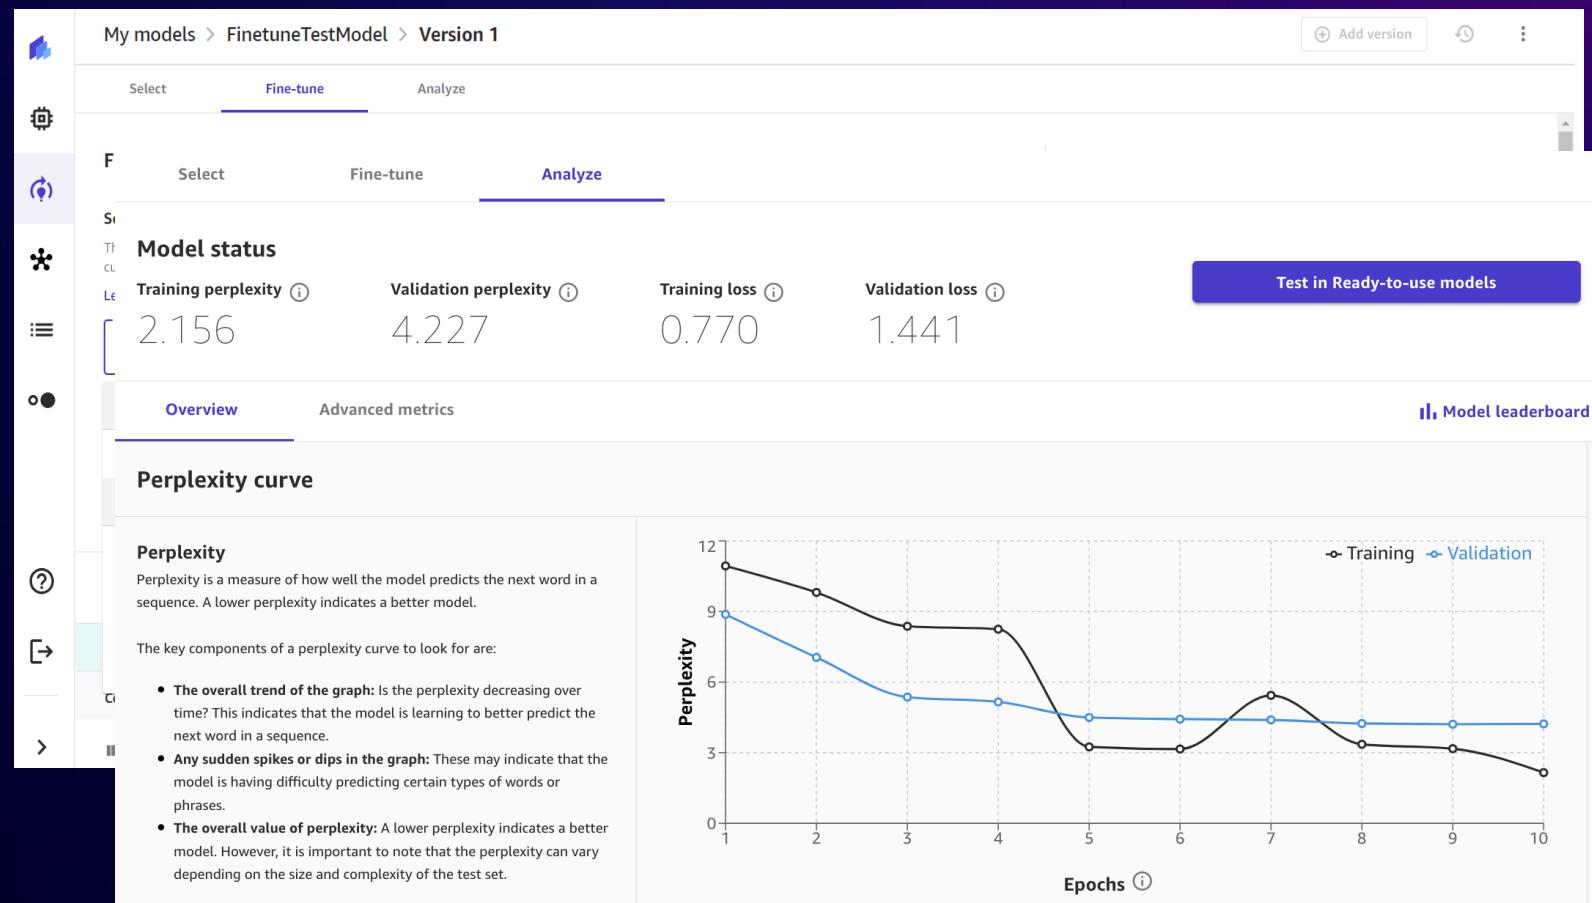
- Select the FMs to customize



The screenshot shows the AWS Foundation Model interface for fine-tuning. The top navigation bar includes 'My models', 'FinetuneTestModel', 'Version 1', 'Add version', and a three-dot menu. The main tabs are 'Select', 'Fine-tune' (which is selected and highlighted in blue), and 'Analyze'. The 'Fine-tune model configurations' section is displayed. Under 'Select up to 3 base models', 'Titan Express' and 'Falcon-7B-Instruct' are selected. The 'Select input column' dropdown is set to 'input'. The 'Select output column' dropdown is set to 'output'. A 'Configure model' button is visible. The sidebar on the left contains icons for 'Select', 'Fine-tune', 'Analyze', 'Metrics', 'Logs', 'Metrics', 'Logs', and 'Help'. The bottom of the interface shows summary statistics: 'Total columns: 2', 'Total rows: 10,000', and 'Total cells: 20,000'.

# NEW – Customize and deploy Foundation Models

- Select the FMs to customize
- Evaluate models



# NEW – Customize and deploy Foundation Models

- Select the FMs to customize
- Evaluate models
- Model leaderboard

The screenshot shows the AWS Foundation Model interface. The top navigation bar includes 'My models', 'FinetuneTestModel', 'Version 1', 'Add version', and a three-dot menu. The left sidebar has icons for Select, Fine-tune, Analyze, and other model management tasks. The main content area is divided into two sections: 'Fine-tune' (selected) and 'Analyze'. The 'Fine-tune' section shows 'Model status' with 'Training perplexity' (2.156) and 'Validation perplexity' (4.22). It also includes an 'Overview' and 'Advanced metrics' tab. The 'Analyze' section shows a 'Perplexity curve' with a graph and a list of components to look for. To the right, a 'Model leaderboard' table lists three models with their perplexity values:

Model name	Training perplexity	Validation perplexity	Training loss	Validation loss
{ Fine-tuned Titan Text Large } <small>Default model</small>	0.14	0.15	0.03	0.15
{ Fine-tuned Jurassic 2 Ultra }	0.22	0.16	0.15	0.17
{ Fine-tuned MPT 40 }	0.25	0.17	0.17	0.18

# NEW – Customize and deploy Foundation Models

- Select the FMs to customize
- Evaluate models
- Model leaderboard
- Deploy models

My models > FinetuneTestModel > Version 1

Select Fine-tune Analyze

**Model status**

Training perplexity	Validation perplexity
2.156	4.22

**Perplexity curve**

Perplexity is a measure of how well the model predicts a sequence. A lower perplexity indicates a better model.

- The overall trend of the graph: Is the perplexity decreasing over time? This indicates that the model is learning to predict the next word in a sequence.
- Any sudden spikes or dips in the graph: These indicate that the model is having difficulty predicting certain types of phrases.
- The overall value of perplexity: A lower perplexity is better for the model. However, it is important to note that the perplexity of a model depends on the size and complexity of the training data.

**Model leaderboard**

Model name	Training perplexity	Validation perplexity	Training loss	Validation loss
{ Fine-tuned Titan Text Large } <small>Default model</small>	0.14	0.15	0.03	0.15
{ Fine-tuned Jurassic 2 Ultra }	0.22	0.16	0.15	0.17
{ Fine-tuned MPT 40 }	0.25	0.17	0.17	0.18

# Demo

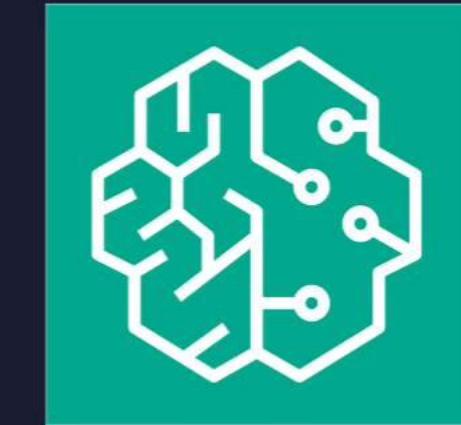


© 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved.

# NO-CODE FINE-TUNING OF LARGE LANGUAGE MODELS



Amazon Bedrock



SageMaker JumpStart

WITH AMAZON  
SAGEMAKER CANVAS

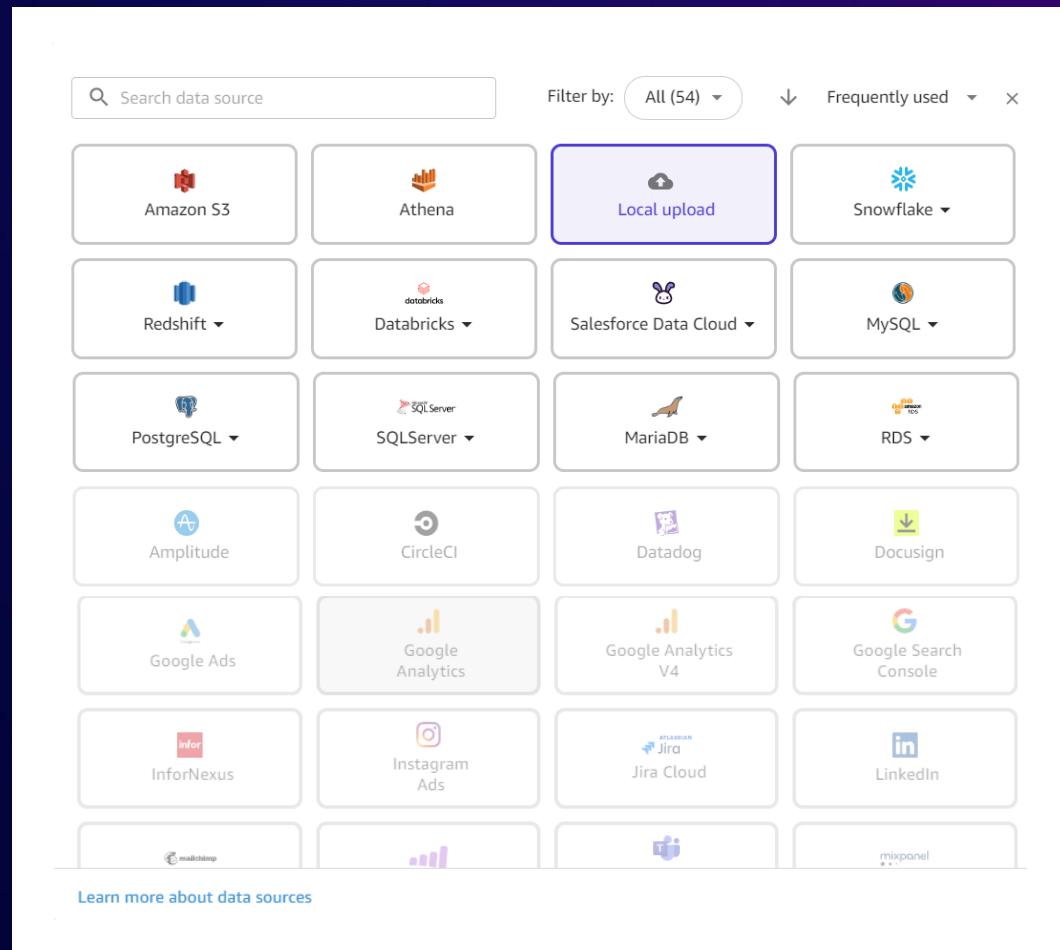
# Custom Models - Comprehensive ML capabilities

Prepare data, build custom models, train and deploy models

# Prepare Data

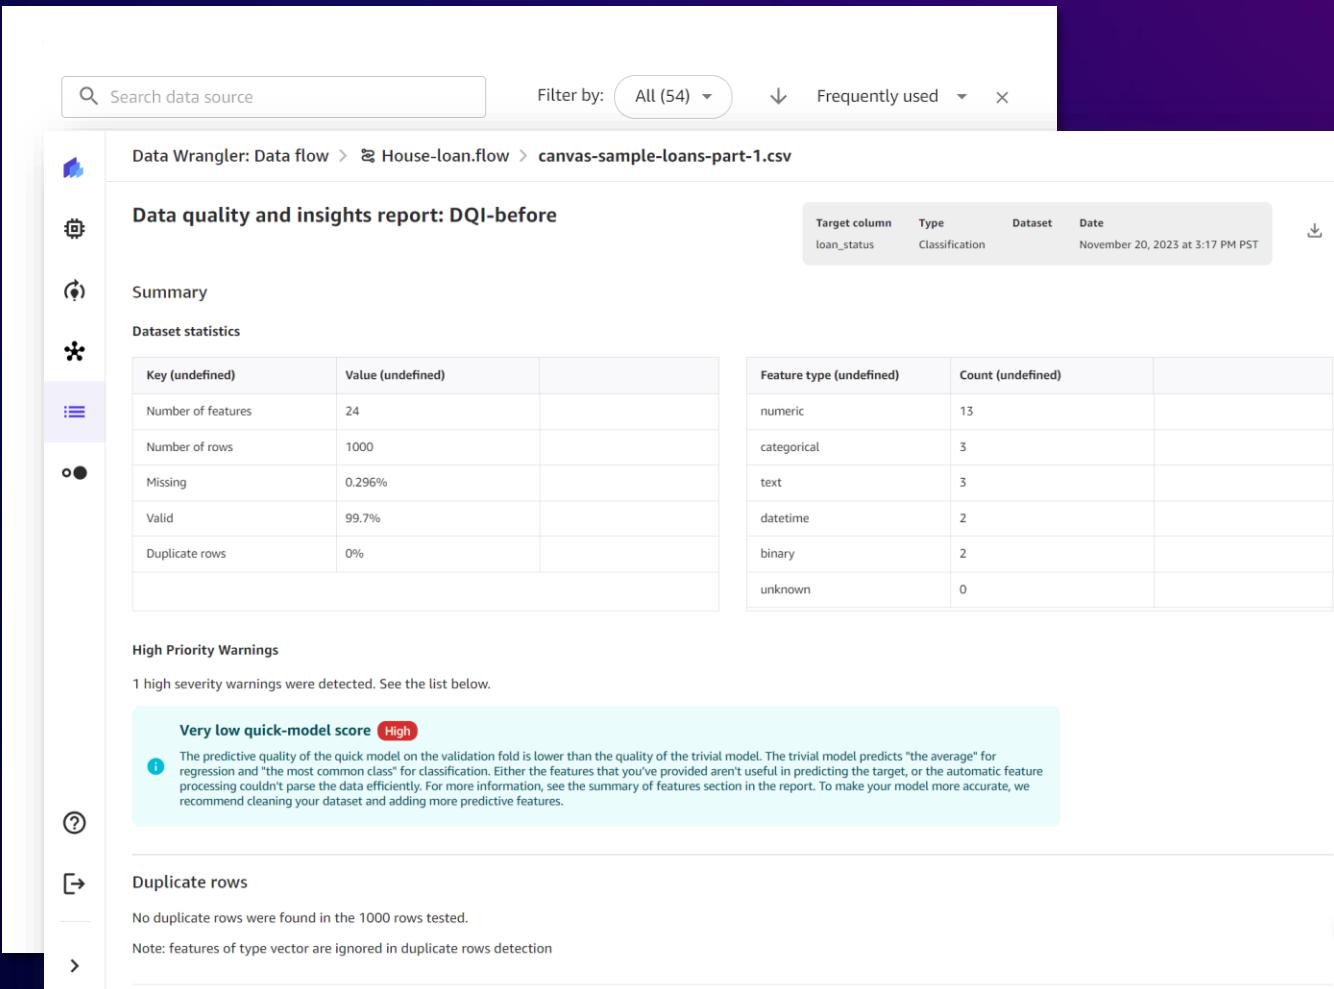
# Prepare Data

- 50+ data connectors



# Prepare Data

- 50+ data connectors
- Rich data insights powered by ML



The screenshot shows the AWS Data Wrangler interface for a data flow named 'House-loan.flow' and a specific dataset named 'canvas-sample-loans-part-1.csv'. The interface includes a search bar, filter dropdowns, and a sidebar with various icons. The main content area displays a 'Data quality and insights report: DQI-before' section. It features a 'Summary' section with 'Dataset statistics' and a table showing key statistics like 24 features, 1000 rows, and 99.7% valid data. Below this is a 'High Priority Warnings' section, which lists a single warning: 'Very low quick-model score' (High severity). The warning message indicates that the predictive quality of the quick model on the validation fold is lower than the quality of the trivial model. The trivial model predicts "the average" for regression and "the most common class" for classification. The report recommends cleaning the dataset and adding more predictive features. The final section shown is 'Duplicate rows', which states that no duplicate rows were found in the 1000 rows tested. A note specifies that features of type vector are ignored in duplicate rows detection.

Search data source

Filter by: All (54) Frequently used

Data Wrangler: Data flow > House-loan.flow > canvas-sample-loans-part-1.csv

Data quality and insights report: DQI-before

Summary

Dataset statistics

Key (undefined)	Value (undefined)
Number of features	24
Number of rows	1000
Missing	0.296%
Valid	99.7%
Duplicate rows	0%

Feature type (undefined)	Count (undefined)
numeric	13
categorical	3
text	3
datetime	2
binary	2
unknown	0

High Priority Warnings

1 high severity warnings were detected. See the list below.

Very low quick-model score High

The predictive quality of the quick model on the validation fold is lower than the quality of the trivial model. The trivial model predicts "the average" for regression and "the most common class" for classification. Either the features that you've provided aren't useful in predicting the target, or the automatic feature processing couldn't parse the data efficiently. For more information, see the summary of features section in the report. To make your model more accurate, we recommend cleaning your dataset and adding more predictive features.

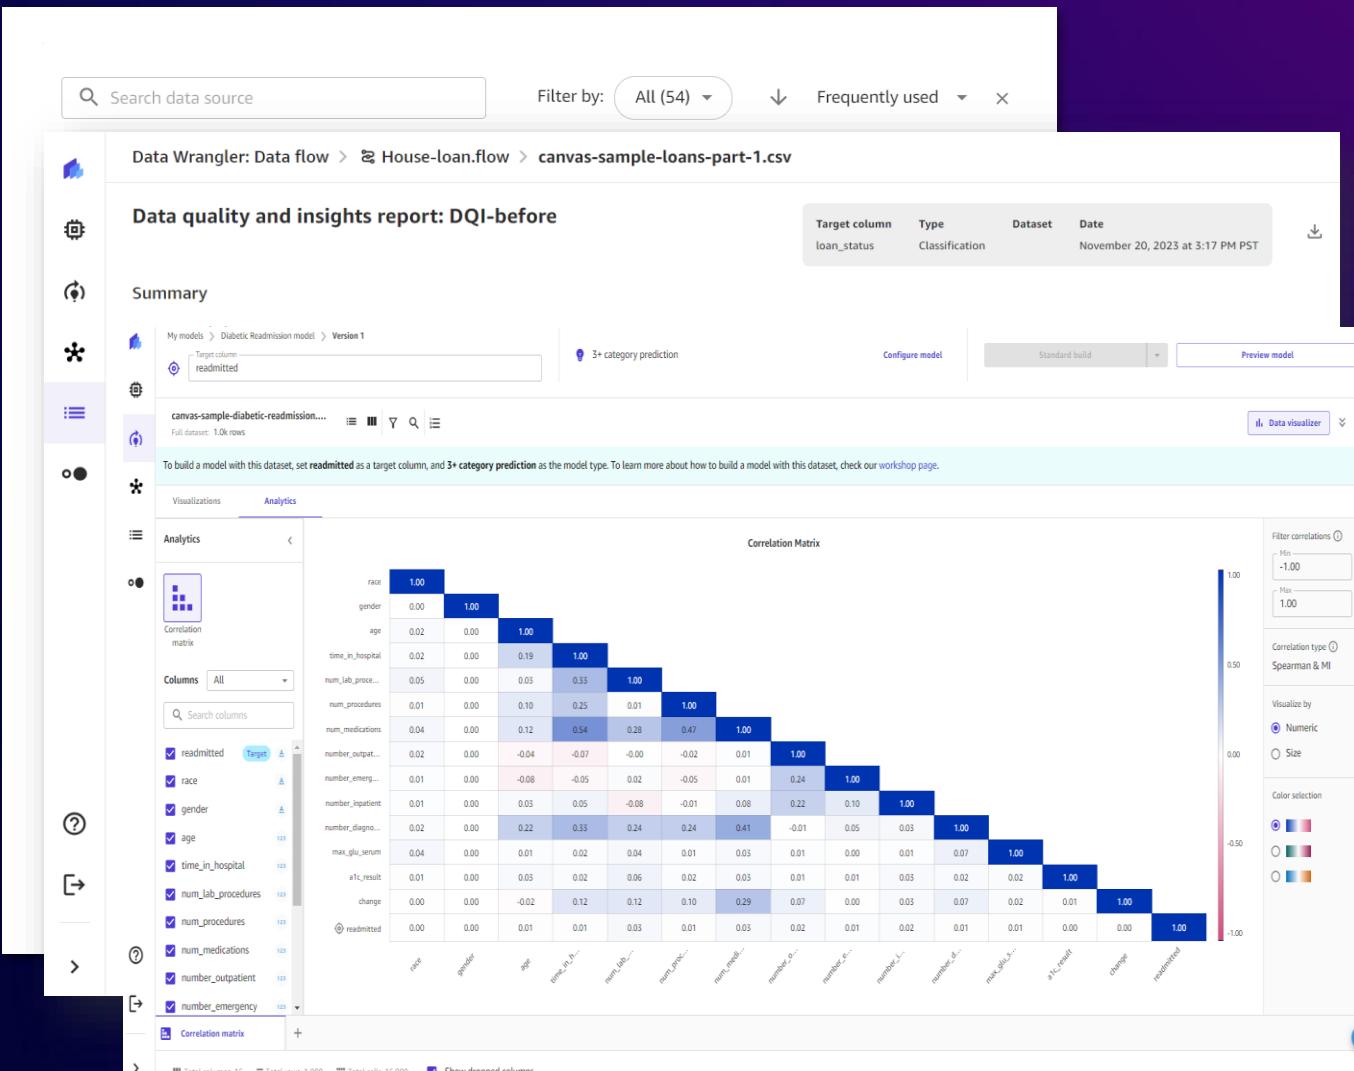
Duplicate rows

No duplicate rows were found in the 1000 rows tested.

Note: features of type vector are ignored in duplicate rows detection

# Prepare Data

- 50+ data connectors
- Rich data insights powered by ML
- Built-in visualizations

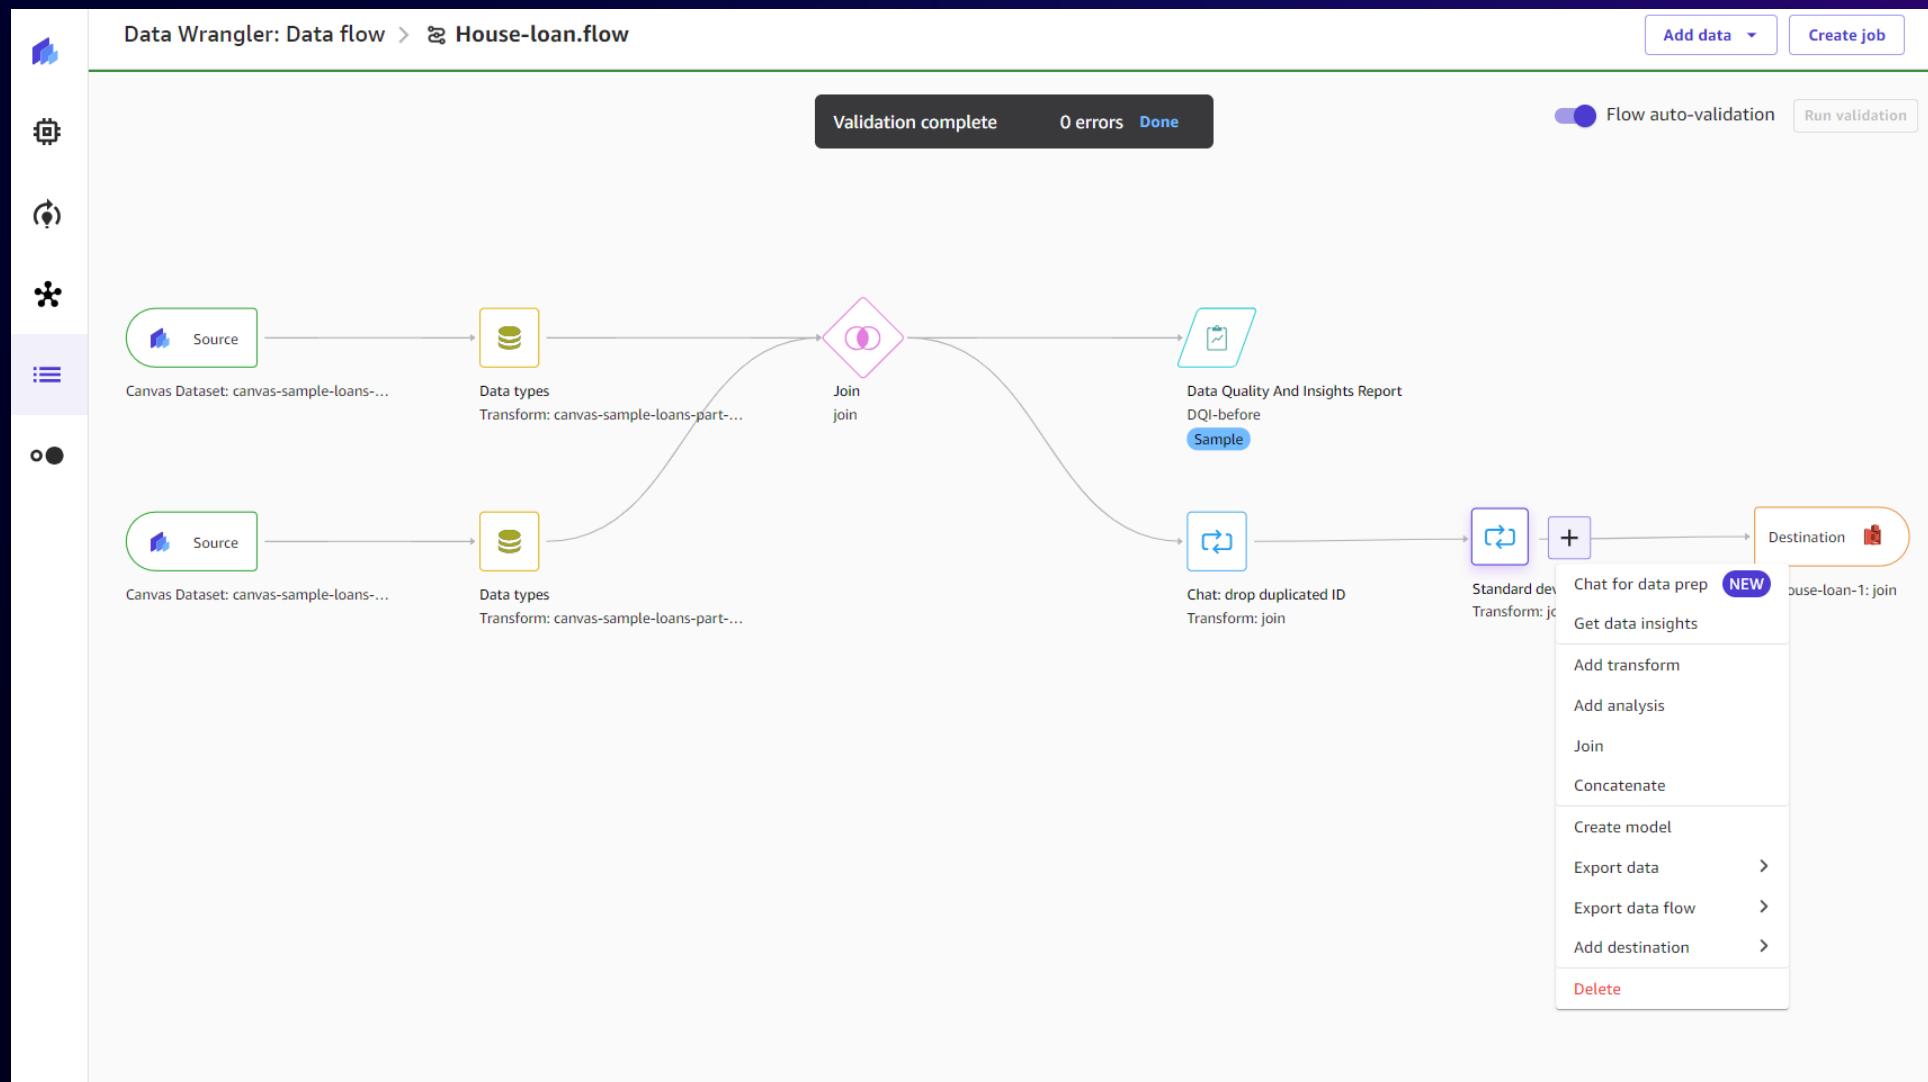


# Prepare Data

- 50+ data connectors
- Rich data insights powered by ML
- Built-in visualizations
- 300+ build-in transforms

The screenshot shows the SageMaker Canvas interface. At the top, there's a search bar for 'Search data source' and filters for 'All (54)' and 'Frequently used'. The main area is titled 'Data Wrangler: Data flow > House-loan.flow > canvas-sample-loans-part-1.csv'. It displays a 'Data quality and insights report: DQI-before' section with a summary for a 'Diabetic Readmission model' (Version 1) with a target column 'readmitted' and a 3+ category prediction. Below this is a 'My models / Lab 1 - Churn / Version 2' section, which is currently in 'Build' mode. It shows a 'Select' tab, a 'Build' tab (selected), an 'Analyze' tab, a 'Predict' tab, and a 'Deploy' tab. Under 'Build', there's a 'Select a column to predict' section where 'Churn' is chosen as the target column. The 'Model type' is set to '2 category prediction', and it's noted that the model classifies Churn into two categories. To the right, there are 'Advanced settings' and a 'Quick build' button. At the bottom, there's a 'Data visualizer' showing a table for 'ChurnData' with 5.0k rows. The table includes columns for 'Vmail\_Plan', 'Vmail\_Me...', 'State', 'Phone', and 'Night\_Mi'. The 'Add transform' sidebar on the right lists various data manipulation options: Manage columns, Change data type, Custom formula, Replace missing values, Replace outlier values, Manage rows, Drop duplicate rows, and Drop rows by formula.

# NEW – Data Preparation Flow



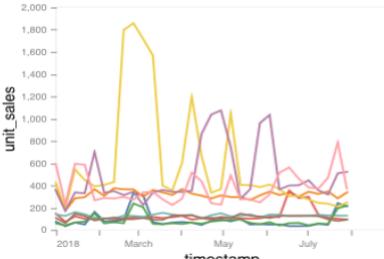
# NEW – Data prep using natural language

Data Wrangler: Data flow > `FoodFlow.flow` > `FoodDemand`

Data Analyses

Step 2. Data types

The code first groups the data by `product_category` and `timestamp`, calculates the average `unit_sales` for each group, then creates a line chart with `timestamp` on the x-axis, average `unit_sales` on the y-axis, and a separate line for each `product_category`.



product\_category

- Beef
- Beverages
- Canned Good
- Fruit
- Food
- Meat
- Seafood
- Spices
- Vegetable

unit\_sales

timestamp

Download Added to analyses

Are there issues with my dataset?

The dataset has 10451 rows and 9 features. There are no duplicate rows or high severity warnings detected. Only 1.28% of values are missing, and the features have varying percentages of missing values. Overall the dataset appears clean with no major issues.

Not added to analyses

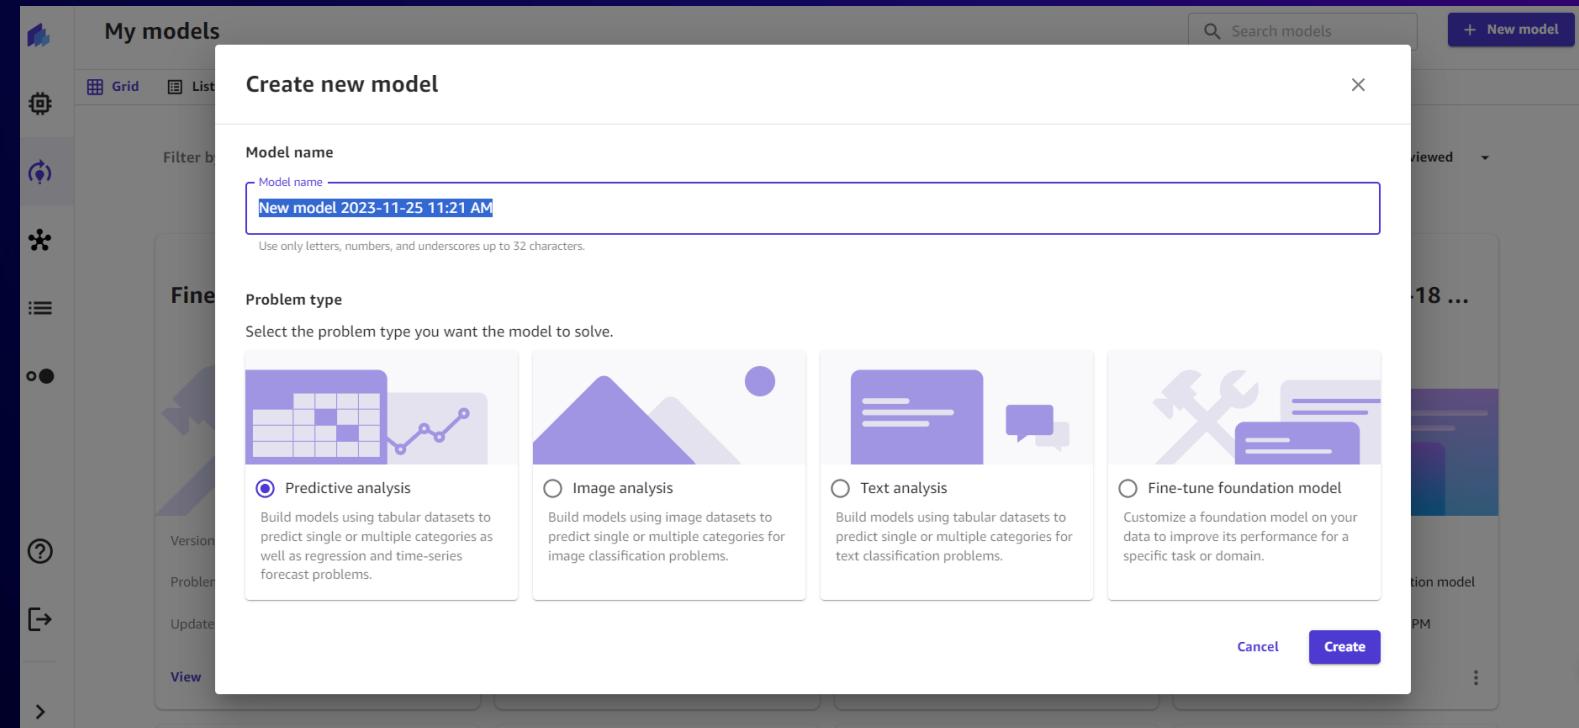
e.g. Help me understand my data with a summary

product_code (long)	product_category (string)	product_subcategory (string)	location_code (long)	scaled_price (float)	promotion_email (long)	promotion_holiday (string)
1062 - 1445	6 Categories	11 Categories	101 - 110	0.31571 - 1.0135	2 Categories	2 Categories
1062	Beverages	Fruit Juice Mango	101	0.879129822	0	0
1062	Beverages	Fruit Juice Mango	101	0.994517243	0	0

# Build and evaluate custom models

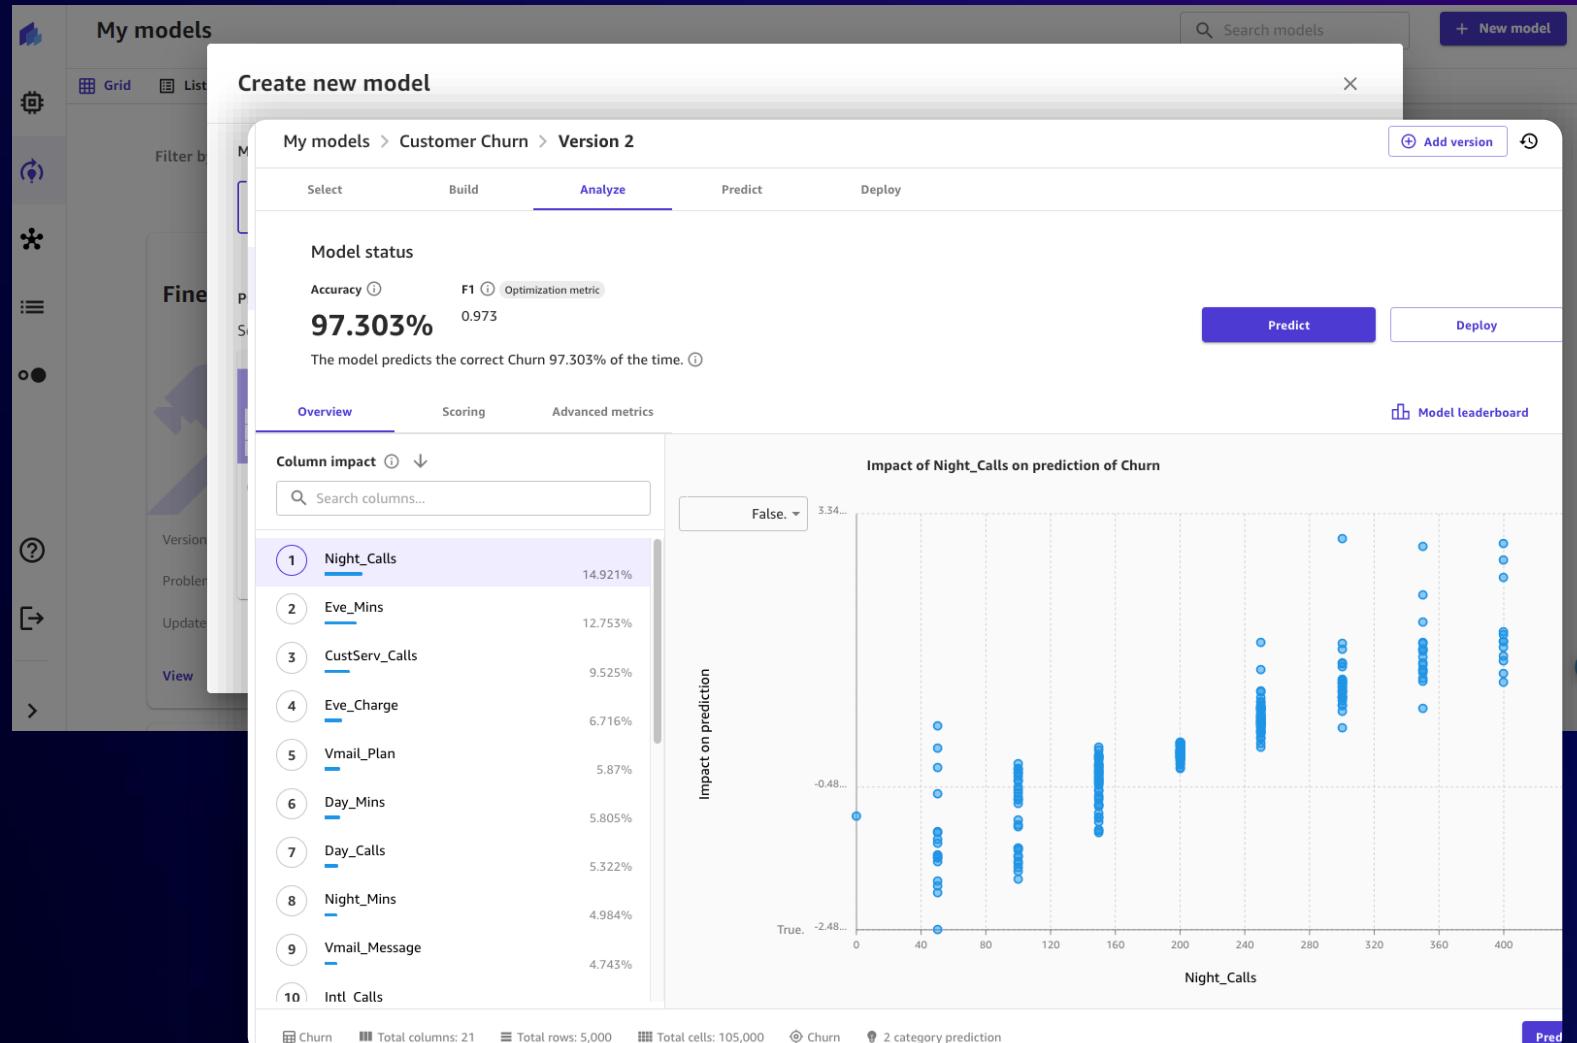
# Build and evaluate custom models

- Choose model type



# Build and evaluate custom models

- Choose model type
- Analyze model metrics



# Build and evaluate custom models

- Choose model type
- Analyze model metrics
- Model leaderboard

The screenshot shows the AWS SageMaker console interface. A modal dialog titled 'Create new model' is open, showing the 'Analyze' tab for a model named 'Customer Churn > Version 2'. The 'Model status' section displays an Accuracy of 97.303% and an F1 score of 0.973. Below this, a message states: 'The model predicts the correct Churn 97.303% of the time.' At the bottom of the dialog are 'Predict' and 'Deploy' buttons. In the background, the main 'My models' page is visible, featuring a 'Model leaderboard' table. The table lists 10 models, each with a rank, name, and various performance metrics. The columns include Model name, F1 (Optimization), Accuracy, AUC, Balanced Accuracy, Precision, Recall, Log Loss, and Inference latency (s). The 'Night\_Calls' model is ranked 1st with an F1 of 97.303% and an Accuracy of 97.303%.

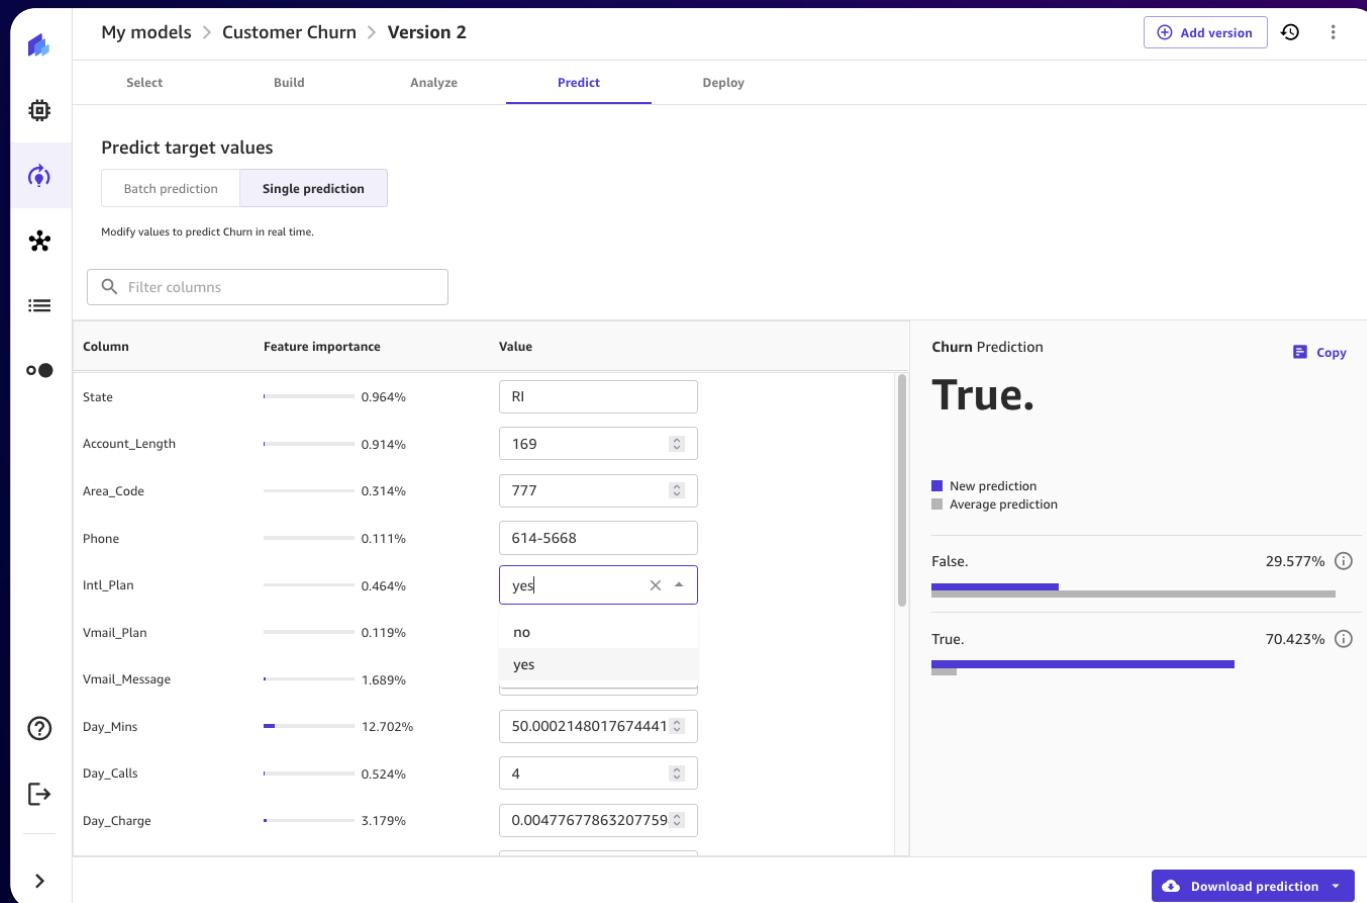
Rank	Model name	F1 (Optimization)	Accuracy	AUC	Balanced Accuracy	Precision	Recall	Log Loss	Inference latency (s)
1	Night_Calls	97.303%	97.303%	0.991	97.305%	97.206%	97.400%	0.200	0.514
2	Eve_Mins	97.303%	97.303%	0.991	97.305%	97.206%	97.400%	0.200	0.514
3	CustServ_L2	91.929%	91.808%	0.970	91.810%	90.504%	93.400%	0.251	0.222
4	Eve_Charge	97.114%	97.103%	0.993	97.103%	96.634%	97.600%	0.102	0.119
5	Vmail_Plans	94.466%	94.406%	0.986	94.407%	93.359%	95.600%	0.155	0.111
6	Day_Mins	94.466%	94.406%	0.986	94.407%	93.359%	95.600%	0.155	0.107
7	Day_Calls	94.466%	94.406%	0.986	94.407%	93.359%	95.600%	0.155	0.108
8	Night_Mins	97.114%	97.103%	0.993	97.103%	96.634%	97.600%	0.102	0.121
9	Vmail_Mess	97.114%	97.103%	0.993	97.103%	96.634%	97.600%	0.102	0.121
10	Intl_Calls	94.466%	94.406%	0.986	94.407%	93.359%	95.600%	0.155	0.107
	Churn	97.114%	97.103%	0.993	97.103%	96.634%	97.600%	0.102	0.118

# Generate highly accurate predictions



# Generate highly accurate predictions

- In-app predictions & what-if analysis



# Generate highly accurate predictions

- In-app predictions & what-if analysis
- Automate predictions

The screenshot displays the AWS Machine Learning console interface. At the top, a navigation bar shows 'My models > Customer Churn > Version 2' with tabs for 'Select', 'Build', 'Analyze', 'Predict' (which is selected), and 'Deploy'. A sub-menu on the left includes icons for 'Predict target values', 'Batch prediction', 'Single prediction', and 'Modify values to predict Churn in real time'. A modal window titled 'Automate batch prediction' is open, explaining that it allows generating predictions for a selected dataset whenever it is updated. It shows 'Prediction frequency' set to 'Immediate' (every time the dataset is updated) and a 'Select a dataset to generate predictions on' section with a 'Select dataset' button. Below the modal, a preview of a 'Churn Prediction' card is shown, stating 'True.' with a 70.423% probability. A legend indicates 'New prediction' (blue bar) and 'Average prediction' (grey bar). A 'Download prediction' button is at the bottom right of the card.

# Generate highly accurate predictions

- In-app predictions & what-if analysis
- Automate predictions
- One-click model deployment

The screenshot displays the AWS SageMaker console interface for generating highly accurate predictions. It includes three main sections: a navigation bar, a prediction interface, and a deployment details view.

**Navigation Bar:** My models > Customer Churn > Version 2. Includes tabs for Select, Build, Analyze, Predict (selected), and Deploy, along with an Add version button and a settings icon.

**Prediction Interface:** Predict target values. Options: Batch prediction (selected) and Single prediction. Subtext: Modify values to predict Churn in real time.

**Automate batch prediction dialog:** Select a dataset to generate predictions for. Every time the selected dataset is updated (either manually or automatically), a new automatic batch prediction job runs.

**Prediction frequency:** A dropdown menu showing the frequency of predictions.

**Operations: Deployment / canvas-customer-churn-prediction-model:** Deployment details.

- Deployment name:** canvas-customer-churn-prediction-model
- Status:** In service (highlighted with a red box)
- Created:** 10/21/23 09:10 AM
- Average predictions per day:** --
- Deployment type:** Real-time
- Model:** Churn Prediction
- Instance type:** ml.m5.xlarge
- Instance count:** 1 (highlighted with a red box)
- Last prediction:** --
- Inference response content:** predicted\_label, probability, probabilities, labels
- Input format:** text/csv
- Deployment URL:** <https://runtime.sagemaker.us-east-1.amazonaws.com/endpoints/canvas-customer-churn-prediction-model/invocations>
- View sample code:**

**Churn Prediction:** A chart showing the distribution of predictions. The chart has two bars: "False." at 29.577% and "True." at 70.423%.

**Download prediction:** A button to download the prediction results.

# Generate highly accurate predictions

- In-app predictions & what-if analysis
- Automate predictions
- One-click model deployment
- Share predictions to Amazon QuickSight

The screenshot illustrates the AWS SageMaker Canvas interface, highlighting its capabilities for generating highly accurate predictions.

**Top Navigation:** My models > Customer Churn > Version 2. The Predict tab is selected. A sidebar on the left includes icons for Select, Build, Analyze, Predict, Deploy, and a gear.

**Predict Target Values:** Options for Batch prediction and Single prediction are shown. A note: "Modify values to predict Churn in real time."

**Automate batch prediction:** A modal window titled "Automate batch prediction" explains that it allows generating predictions for a selected dataset, which triggers automatic batch prediction jobs when updated.

**Prediction frequency:** A section for configuring prediction frequency.

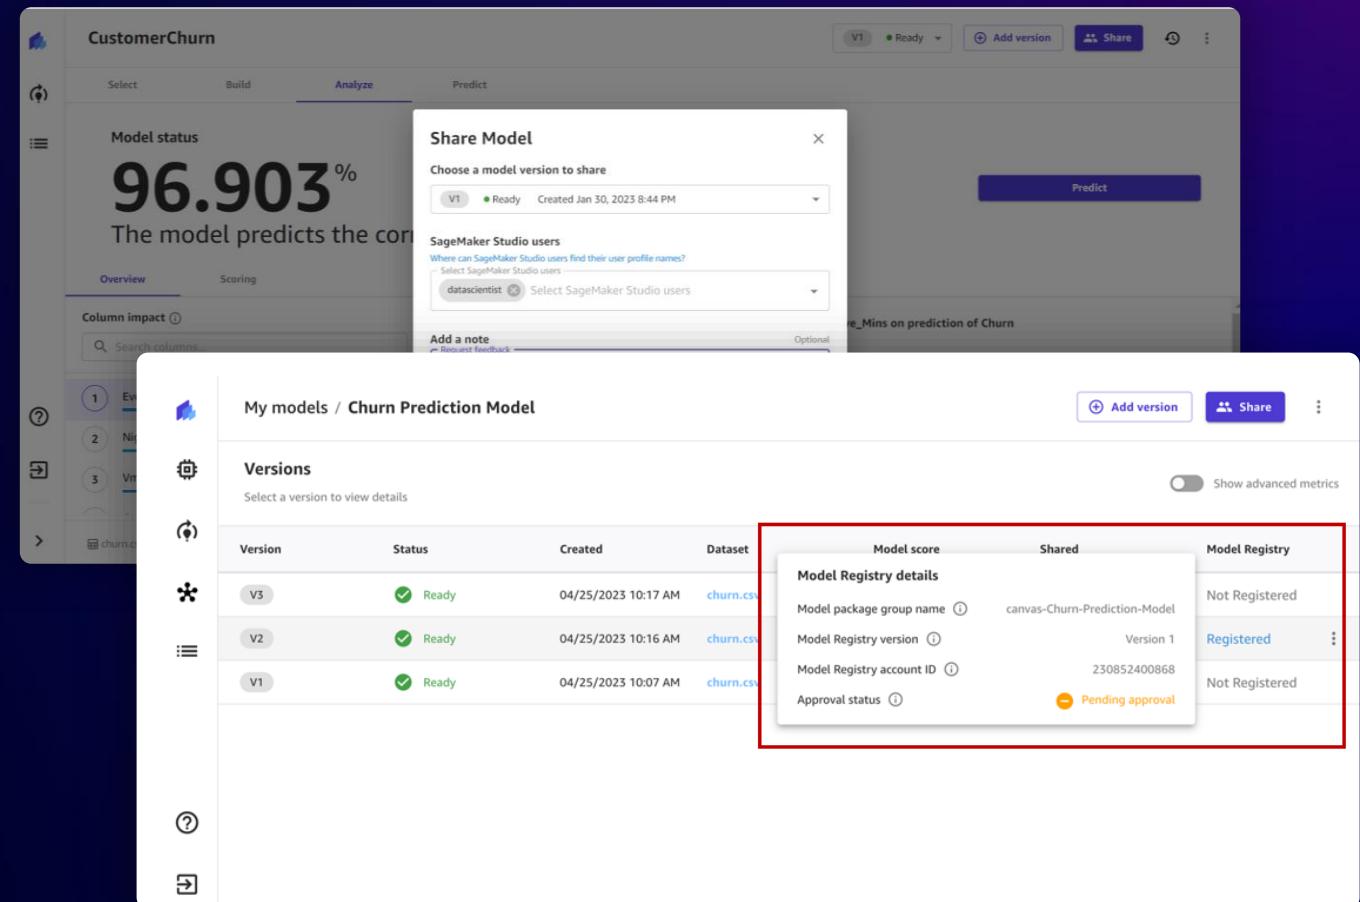
**Deployment Details:** Operations: Deployment / canvas-customer-churn-prediction-model. Deployment name: canvas-customer-churn-prediction-model. Status: In service (highlighted with a red box). Deployment type: Real-time. Model: Churn Prediction. Created: 10/21/23 09:10 AM. Average predictions per day: --. Last prediction: --. Instance type: ml.m5.xlarge. Instance count: 1 (highlighted with a red box). Deployment URL: <https://runtime.sagemaker.us-east-1.amazonaws.com/endpoints/canvas-customer-churn-prediction-model/invocations>. Inference response content: predicted\_label, probability, probabilities, labels. Input format: text/csv. A "View sample code" link is also present.

**Churn Prediction:** A bar chart showing the distribution of predictions. The legend indicates "New prediction" (blue) and "Average prediction" (grey). The chart shows "True." at 70.423% and "False." at 29.577%.

**Actions:** "Copy" (for the prediction chart), "Download prediction" (for deployment details), and a "Download" button (for the deployment URL).

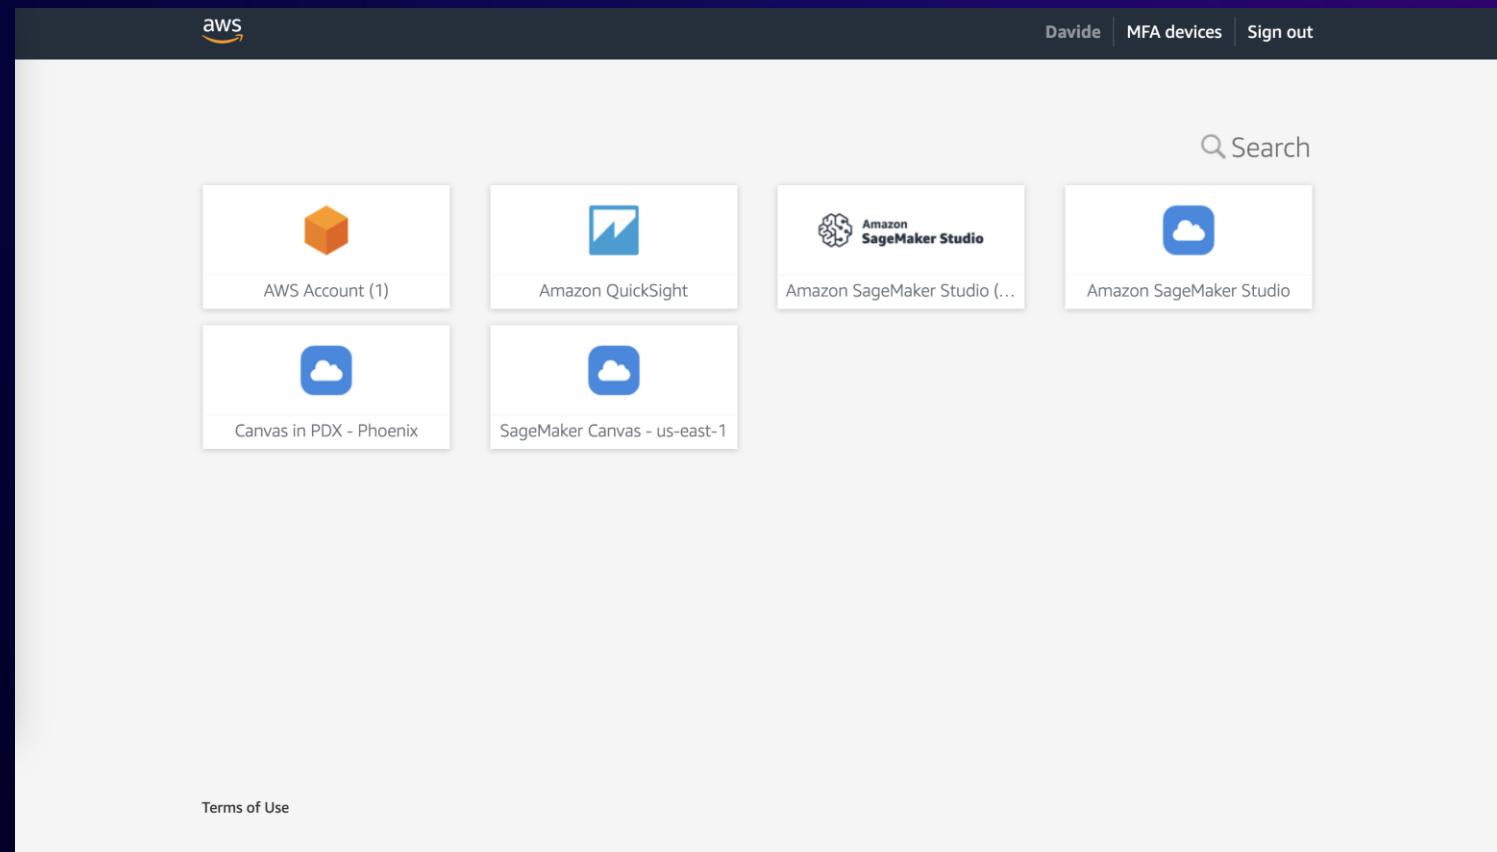
# Collaborate with pro-code ML users

- Share model artifacts to SageMaker Studio
- Register model to SageMaker Model Registry



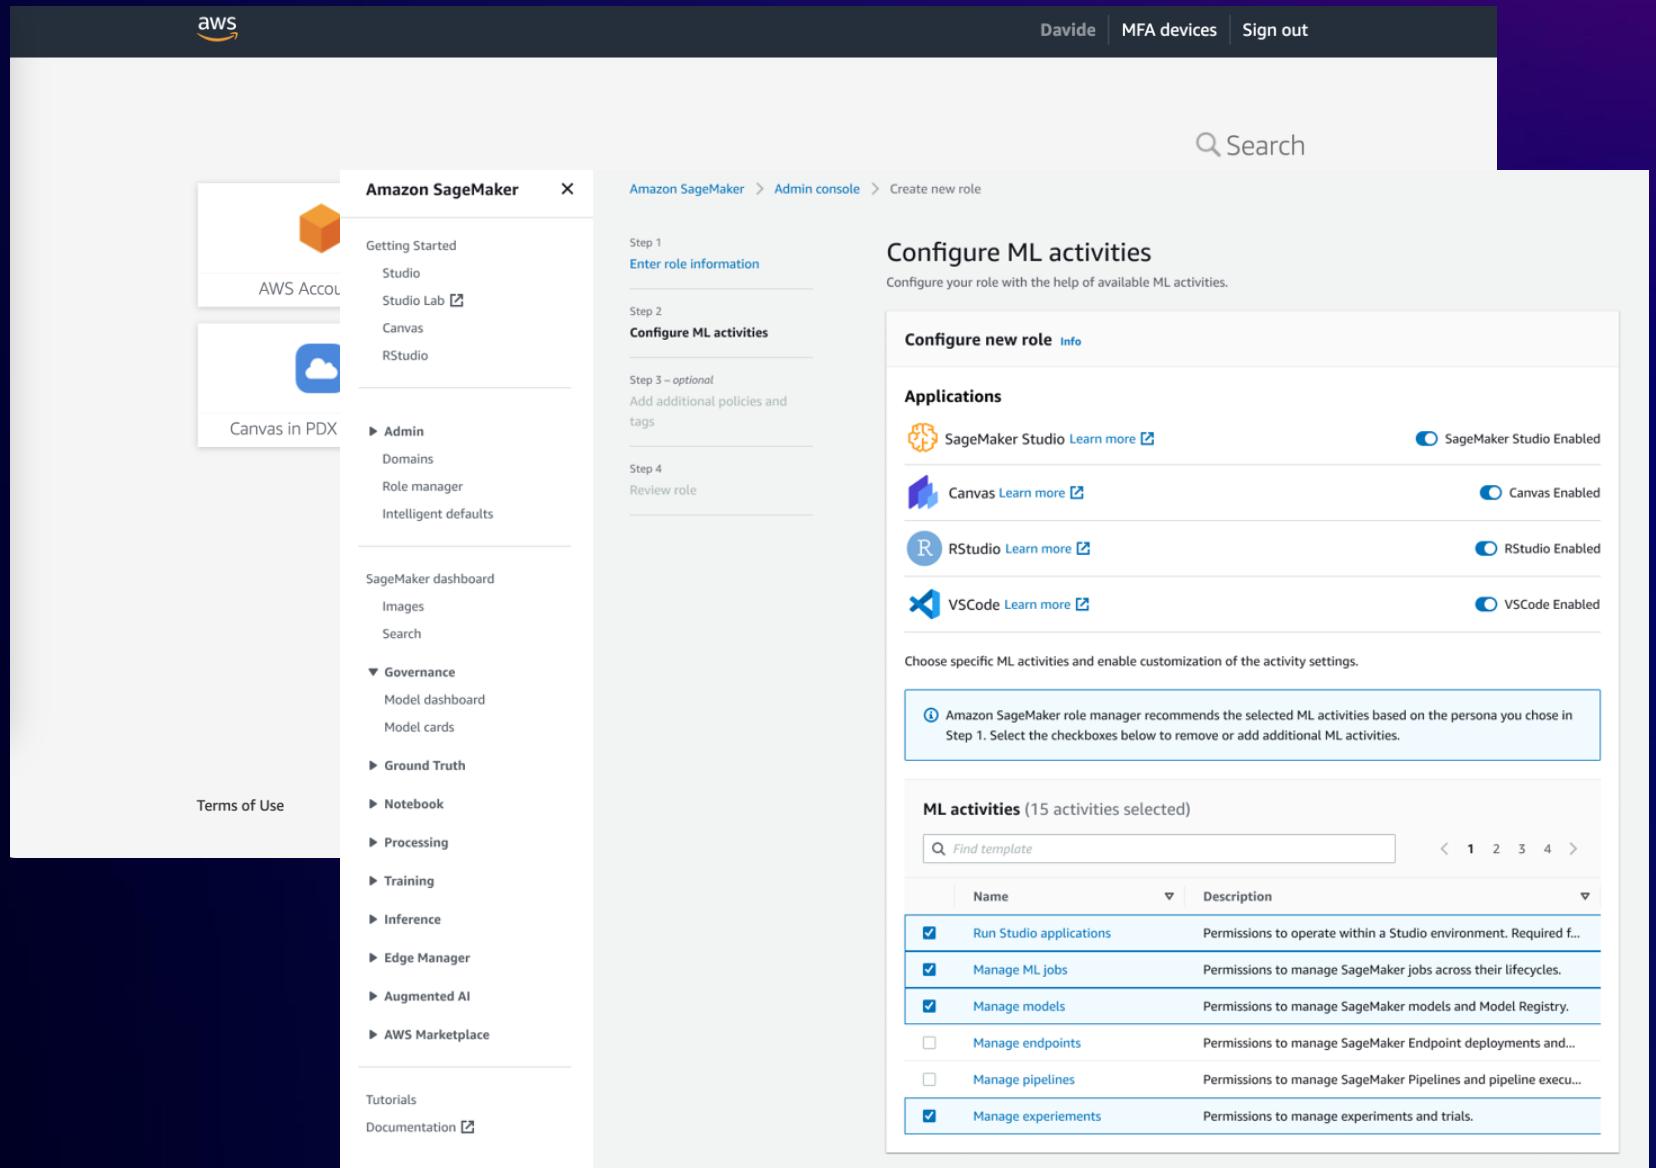
# Management controls

- Single sign-on setup



# Management controls

- Single sign-on setup
- Setup permissions



The screenshot shows the Amazon SageMaker Admin console with the 'Create new role' wizard open. The left sidebar includes links for Getting Started, Studio, Studio Lab, Canvas, RStudio, Admin (selected), Domains, Role manager, Intelligent defaults, SageMaker dashboard, Images, Search, Governance, Model dashboard, Model cards, Ground Truth, Notebook, Processing, Training, Inference, Edge Manager, Augmented AI, AWS Marketplace, Tutorials, and Documentation. The main content area shows the 'Create new role' process in four steps: Step 1 (Enter role information), Step 2 (Configure ML activities), Step 3 (optional: Add additional policies and tags), and Step 4 (Review role). Step 2 is currently active, titled 'Configure ML activities' with the sub-section 'Configure new role'. It lists four applications: SageMaker Studio (Enabled), Canvas (Enabled), RStudio (Enabled), and VSCode (Enabled). Below this, a note says: 'Amazon SageMaker role manager recommends the selected ML activities based on the persona you chose in Step 1. Select the checkboxes below to remove or add additional ML activities.' A table titled 'ML activities (15 activities selected)' lists 15 items, each with a checkbox, name, and description. The items are: Run Studio applications, Manage ML jobs, Manage models, Manage endpoints, Manage pipelines, and Manage experiments.

Amazon SageMaker

Getting Started

Studio

Studio Lab

Canvas

RStudio

Admin

Domains

Role manager

Intelligent defaults

SageMaker dashboard

Images

Search

Governance

Model dashboard

Model cards

Ground Truth

Notebook

Processing

Training

Inference

Edge Manager

Augmented AI

AWS Marketplace

Tutorials

Documentation

Step 1  
Enter role information

Step 2  
Configure ML activities

Step 3 – optional  
Add additional policies and tags

Step 4  
Review role

Configure ML activities

Configure new role

Applications

SageMaker Studio Learn more

Canvas Learn more

RStudio Learn more

VSCode Learn more

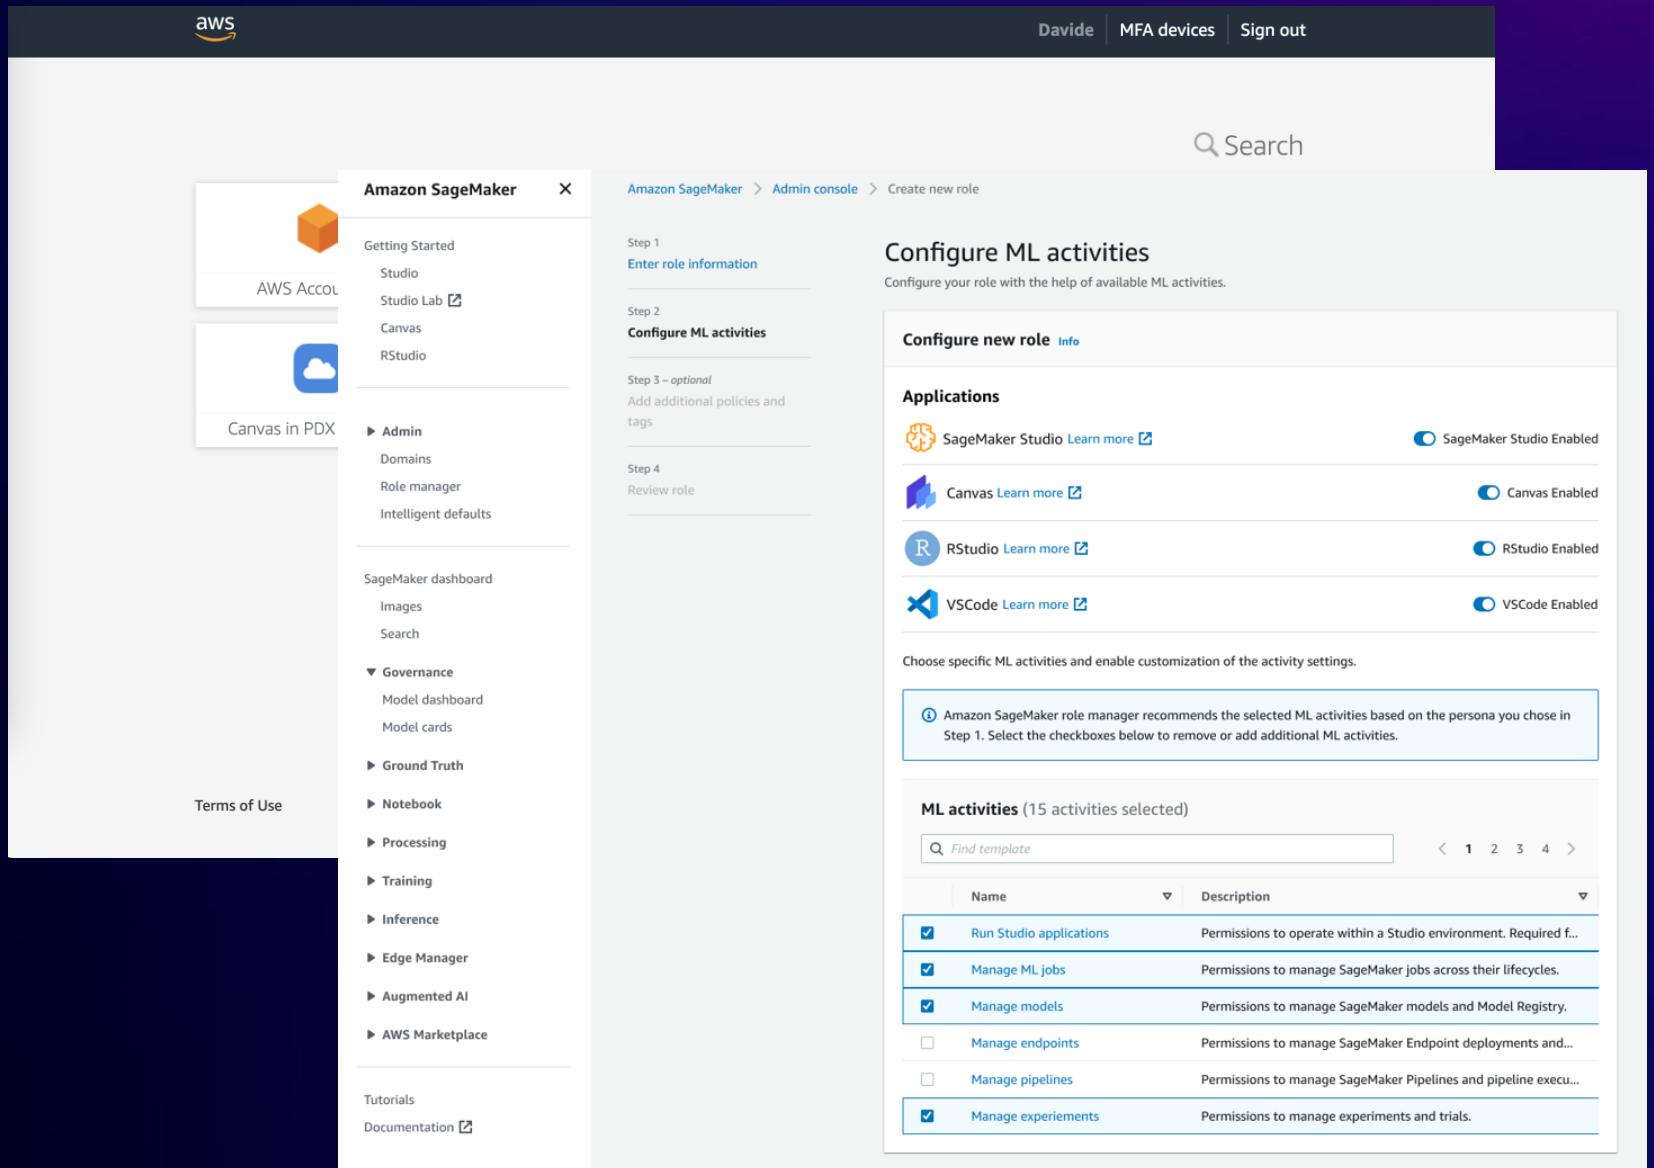
Choose specific ML activities and enable customization of the activity settings.

Amazon SageMaker role manager recommends the selected ML activities based on the persona you chose in Step 1. Select the checkboxes below to remove or add additional ML activities.

Name	Description
<input checked="" type="checkbox"/> Run Studio applications	Permissions to operate within a Studio environment. Required for...
<input checked="" type="checkbox"/> Manage ML jobs	Permissions to manage SageMaker jobs across their lifecycles.
<input checked="" type="checkbox"/> Manage models	Permissions to manage SageMaker models and Model Registry.
<input type="checkbox"/> Manage endpoints	Permissions to manage SageMaker Endpoint deployments and...
<input type="checkbox"/> Manage pipelines	Permissions to manage SageMaker Pipelines and pipeline execu...
<input checked="" type="checkbox"/> Manage experiments	Permissions to manage experiments and trials.

# Management controls

- Single sign-on setup
- Setup permissions
- Automate shutdown



The screenshot shows the Amazon SageMaker Admin console with the 'Create new role' wizard open. The left sidebar includes links for Getting Started, Studio, Studio Lab, Canvas, RStudio, Admin (Domains, Role manager, Intelligent defaults), SageMaker dashboard (Images, Search), Governance (Model dashboard, Model cards), Ground Truth, Notebook, Processing, Training, Inference, Edge Manager, Augmented AI, AWS Marketplace, Tutorials, and Documentation. The main content area shows the 'Create new role' wizard with four steps: Step 1 (Enter role information), Step 2 (Configure ML activities), Step 3 (optional: Add additional policies and tags), and Step 4 (Review role). Step 2 is currently active, titled 'Configure ML activities'. It includes a section for 'Configure new role' with checkboxes for SageMaker Studio (Enabled), Canvas (Enabled), RStudio (Enabled), and VSCode (Enabled). A note states: 'Amazon SageMaker role manager recommends the selected ML activities based on the persona you chose in Step 1. Select the checkboxes below to remove or add additional ML activities.' Below this is a table titled 'ML activities (15 activities selected)' with columns for Name and Description, showing items like Run Studio applications, Manage ML jobs, Manage models, Manage endpoints, Manage pipelines, and Manage experiments.

Name	Description
<input checked="" type="checkbox"/> Run Studio applications	Permissions to operate within a Studio environment. Required for...
<input checked="" type="checkbox"/> Manage ML jobs	Permissions to manage SageMaker jobs across their lifecycles.
<input checked="" type="checkbox"/> Manage models	Permissions to manage SageMaker models and Model Registry.
<input type="checkbox"/> Manage endpoints	Permissions to manage SageMaker Endpoint deployments and...
<input type="checkbox"/> Manage pipelines	Permissions to manage SageMaker Pipelines and pipeline execu...
<input checked="" type="checkbox"/> Manage experiments	Permissions to manage experiments and trials.

# Demo



© 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Import data



Prepare data



Scale data operations



Build models

To get started...

Import data ▾

Select existing dataset ▾



# Customer Story: Thomson Reuters

**Ramdev Wudali**

(he/him)

Data Architect AI & BI Platforms  
Thomson Reuters



© 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved.

# AI @ Thomson Reuters

## WIN (Westlaw is Natural)

First commercially available search engine with probabilistic rank retrieval

## PeopleCite & Profiler

To extract entities from American case law documents and create a knowledge base of judges, attorneys, and expert witnesses with links to all their cases and biographies

1990s

2000s

## Reuters Insider

Reuters Insider used CaRE classification and Dexter entity extraction to connect transcripts of live news shows to video; this enabled searching video-based news

## AI platform

Enabling both data scientists and non-programmers to rapidly build AI Solutions in production, prioritizing transparency, reusability and agility

## Reuters Tracer and Social Data Platform (SDP) Separating Real News from Fake in 40 Milliseconds

2010s

## History Assistant

Large-scale natural language processing (NLP) system found history relationships between court decisions by using a combination of information retrieval and machine learning techniques to link each new case to related documents that it may impact

## Checkpoint – Broadside

To power the new "Intuitive Search" capability in Checkpoint, TR's market-leading research solution for tax and accounting professionals

## Casetext Acquisition

Supporting TR's 'build, partner and buy' strategy to bring generative AI solutions to its customers

## CaRE - Classification and Recommendation Engine

To classify legal, tax and finance documents to large taxonomies

## Practical Law Dynamic Search

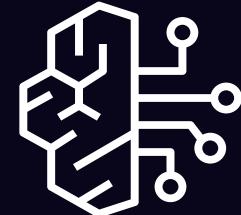
State-of-the-art AI system designed to find the answers to free-form questions in Practical Law

# TR offers secure access and usage of AWS services

Low Code



Amazon SageMaker  
Jumpstart



Amazon Bedrock



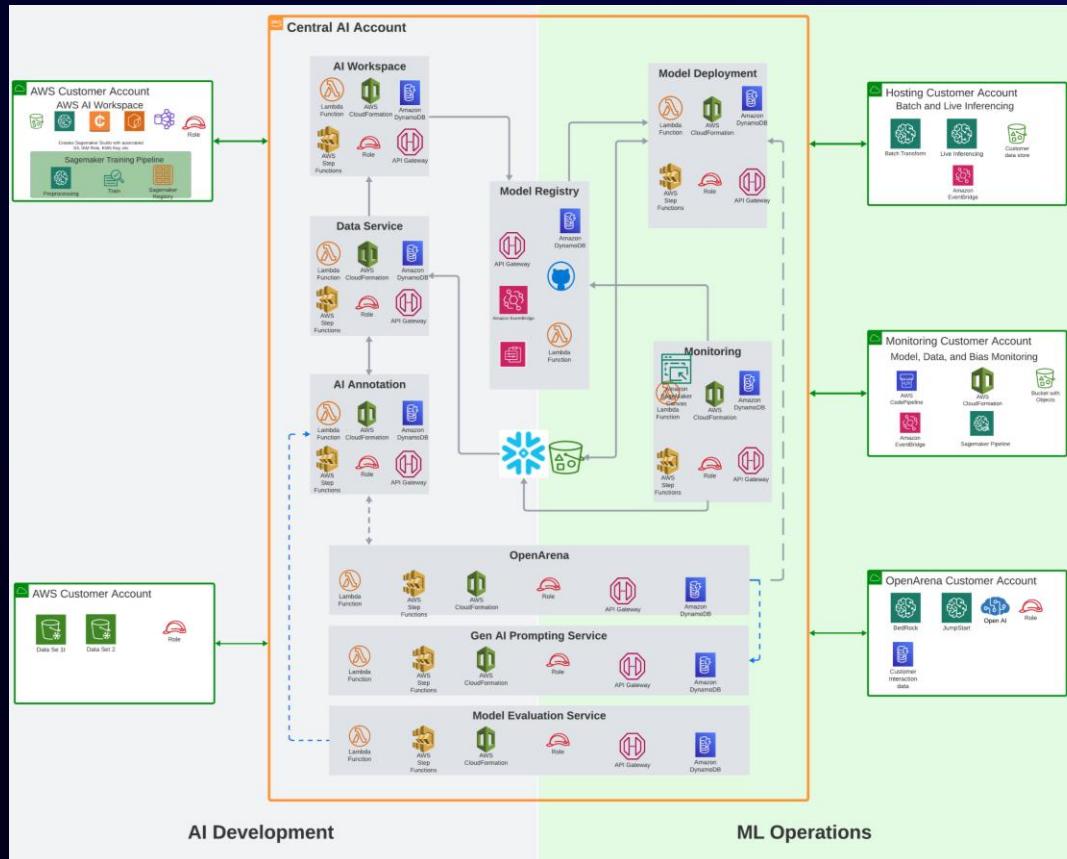
Amazon SageMaker  
Canvas

# AI Platform

## Initially tailored to data scientists, now empowers AI novices

### Services serving the AI Lifecycle end-to-end

Data Service	AI Annotation	AI Workspace	OpenArena	Model Registry	Model Deployment	Monitoring
secures access to TR data	captures annotations from subject-matter experts (SMEs)	secures access to cloud resources	enables gen. AI solutions in a low code/no code space	central repository of TR AI models	streamlines model hosting	monitors model and data for quality, drift and bias



**Secures access to AWS cloud resources and Thomson Reuters content**

**Enables interactions across accounts and separation between customer data and operational data**

**Enables swift integration of AI frameworks like generative AI**

**Encourages reusability of AI models**

**Facilitates standardized governance workflow**

**Currently evolving the AI Platform to empower low-code generative AI chains**

# Users can access AWS services and start building AI solutions through the AI Platform in a few clicks

Fill in a few fields to provision a **secure** compute environment

Add Workspace

Workspace Name: \* Testing Workspace

Project Name: \* Project Name

Description: Description

DIA ID: \* DIA ID

Provider: \* AWS

Account ID: 664145726992 — TR EDL ADSS PROD

Data Classification: Confidential

Asset ID: 208321

Region: us-east-1

Owner Email: sahar.elkhoury@thomsonreuters.com

REQUESTS

Outgoing Request

Requesting Workspace Name: abc  
Requestor Email: test.qed@thomsonreuters.com  
Requested On: 14 Apr 2023 03:01 PM GMT+2  
Request Comments: Request access from DS

Dataset Name: AI Hackathon May 2023 - 3  
Approver Email: test.qed@thomsonreuters.com  
Dataset Description: [DS]  
Status: ✓ Approved

SERVICES

Service Owner Email	Service Type	Status
sahar.elkhoury@thomsonreuters.com	All	All
<b>SAHAR.ELKOURY@THOMSONREUTERS.COM</b>		
canvas-qEDE	canvas	Status: ✓ Ready
studio-NqRo	studio	Status: ✓ Stopped

USERS

User	Role
Ramdev Wudali eSSO_0102180	Approver
Sahar El Khoury eSSO_6119620	User

Streamlines access to data

Streamlines access to AWS Low Code/No Code ML services such as **Amazon Bedrock, SageMaker Canvas**

Facilitates collaboration across users

# Enabling AI Novices through hackathons, sprints and workshops

200+  
AI Novices  
developed solutions  
in less than 2 days



Customer Calls Forecasting

to manage resources during peak seasons (time-series)



Conversion Prediction

to predict the likelihood of purchasing a subscription after a trial (binary classification)



Customer Profiling

to uncover customers' activity insights and patterns (clustering)

# Enabling AI Novices through hackathons, sprints and workshops

Participating personas



Domain Experts



Data Scientists/  
Researchers



Data Analysts



MLOps Engineers

...how they felt about their Canvas experience

**"from zero knowledge on AI, to solve a use case??!! Wow!!"**

**"it was amazing to learn about the day-to-day activities of an ML engineer working in TR starting off with workspaces and then playing around with SageMaker Canvas as a data enthusiast who had lost some touch with his data powers"**

**"H. who was new to the team had never worked on the AI Platform. After the hackathon, he started leveraging the AI platform for all his work"**

**"This was an amazing opportunity to grab a big bite on AI in just 2 days!! Our team from a wide variety of background, like Dev, DevOps, Quality Assurance, we brainstormed to understand the use case. End of Day1, we could figure out what we need to do to achieve from the available data and how the SageMaker Canvas and Data Wrangler could be used in building a solution for the problem..."**

# A deeper look at one of our use cases:

## Business Problem:

Manage customer call resources during peak seasons to ensure optimal number of customer service personnel.

## Challenge with current approach:

Human predictions do not provide a cost-effective solution.



Storage Engineer

Business Analyst

Software Engineer

# Bringing their workflow to life with Canvas

## Data Preparation

Historical data analysis from call center containing call volume, wait time, date and other relevant metrics Data Source: [Snowflake](#)

## Feature Engineering

Use [Canvas](#), to extract relevant features from the data and identify patterns/trends that impact call volume

## Model Training

[Canvas](#) was used to train the custom model using the features identified

## Deployment

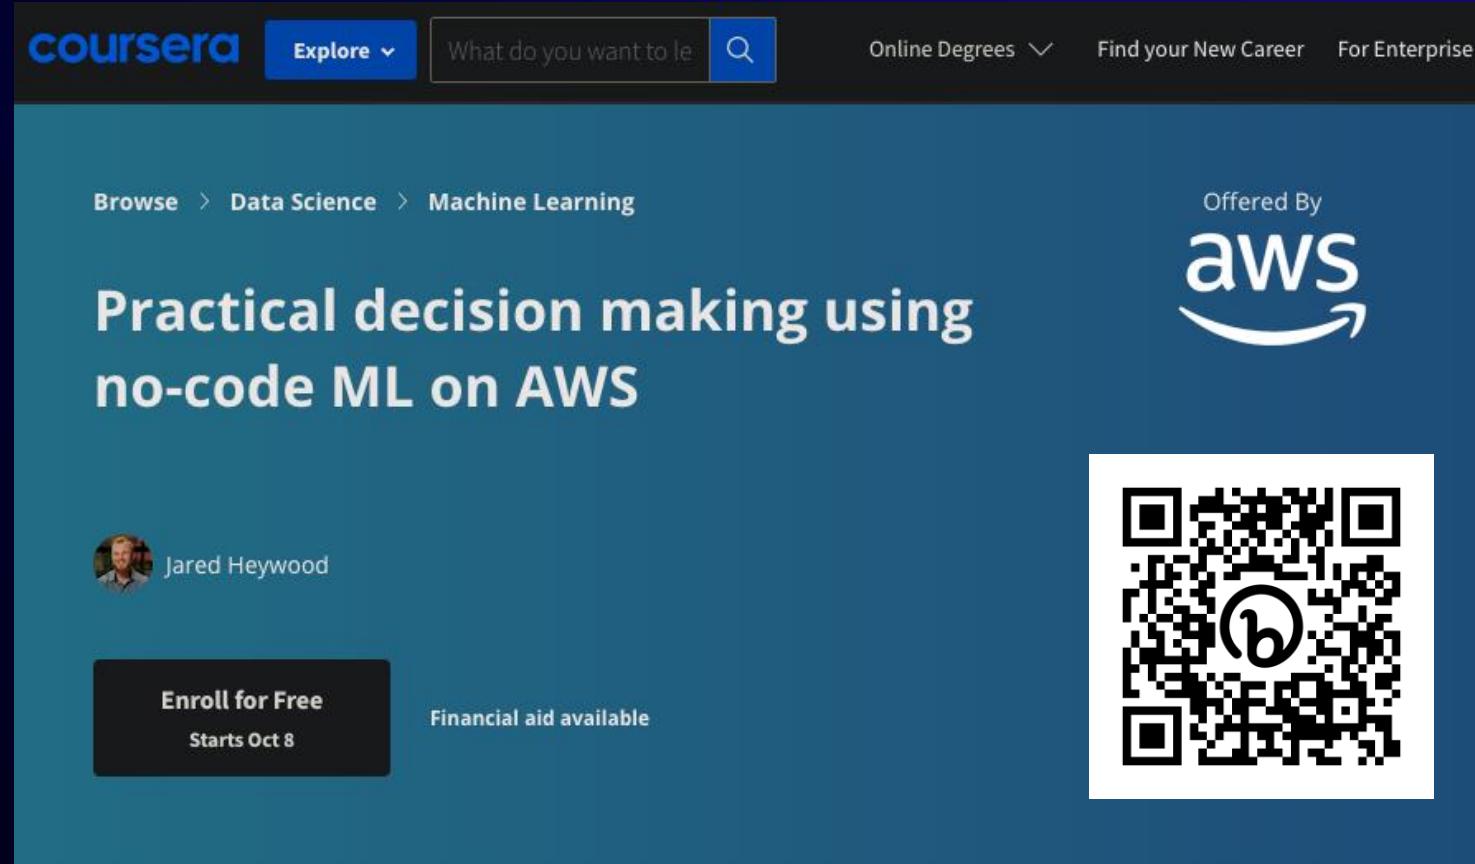
Model artifact from the associated [S3](#) bucket was used to deploy a [SageMaker](#) inference endpoint

# Key Learning Outcomes from using Canvas

- 1 Learned about the solutions AWS provides for AI problems
- 2 Learned to implement a Machine Learning model
- 3 Understood how to process and treat the data
- 4 Learned how simple it can be to train a Machine Learning model
- 5 Machine Learning proved to be novice friendly

# Resources and Next Steps

# Course for business users—how to ask ML questions



The screenshot shows the Coursera website with the following details:

- Header:** coursera, Explore, What do you want to learn?, Online Degrees, Find your New Career, For Enterprise.
- Breadcrumbs:** Browse > Data Science > Machine Learning
- Course Title:** Practical decision making using no-code ML on AWS
- Offered By:** AWS
- Instructor:** Jared Heywood
- Enrollment:** Enroll for Free, Starts Oct 8
- Aid:** Financial aid available
- QR Code:** A QR code with a central 'b' logo, likely linking to the course page.

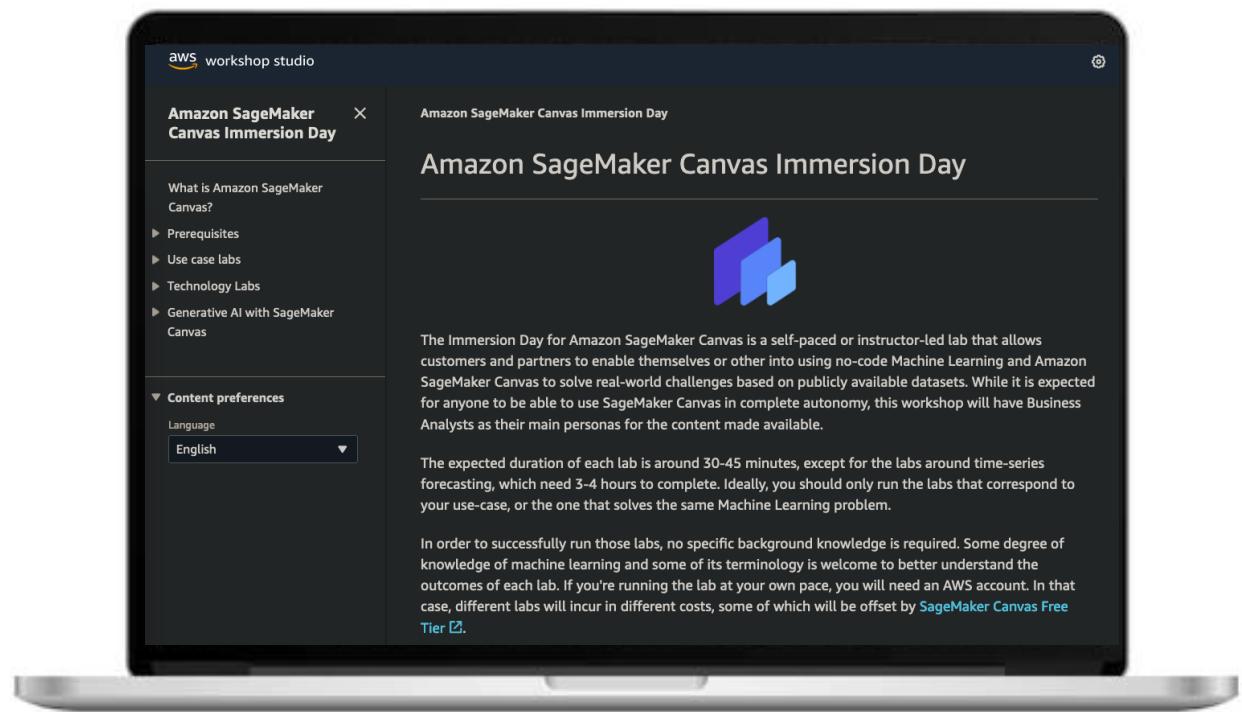
[More about the course \(Blog\)](#)

<https://www.coursera.org/learn/no-code-ml-aws>

# Get hands-on today!



<https://bit.ly/sm-canvas-immday>



# Low-Code/No-Code ML at re:Invent 2023

- **AIM339** | No-code ML and foundation models with Amazon SageMaker Canvas
  - Nov. 30 | 2:30 PM - 3:30 PM (PST)
  - Caesars Forum | Level 1 | Forum 115
- **AIM363** | New LLM capabilities in Amazon SageMaker Canvas
  - Nov 30 | 4:00 PM – 5:00 PM (PST)
  - Mandalay Bay | Level 2 | South | Oceanside C

# Thank you!

**Rajneesh Singh**

[rajneesh@amazon.com](mailto:rajneesh@amazon.com)



**Davide Gallitelli**

[dggallit@amazon.com](mailto:dggallit@amazon.com)



Please complete the session  
survey in the mobile app

**Ramdev Wudali**

[ramdev.wudali@thomsonreuters.com](mailto:ramdev.wudali@thomsonreuters.com)

