

SageMaker JumpStart

João Moura
AI/ML Specialist SA



SageMaker JumpStart

**Easily and quickly
bring machine
learning
applications to
market**



Accelerate time to deploy over 300 open source models

Use one-click deployable ML models and algorithms from popular model zoos



18 pre-built solutions for common ML use cases

Leverage solutions out-of-the-box or customize for a specific business problem






Get started with just a few clicks

Easily bring ML applications to market using pre-built solutions, ML models and algorithms from popular model zoos, and getting started content

SageMaker JumpStart open source models

300+ pre-trained, state-of-the-art, open source models from PyTorch Hub, TensorFlow Hub, Hugging Face ...

	TASKS	MODELS
 VISION	Image Classification Image Embedding Object Detection Semantic Segmentation	ResNet, Inception, MobileNet, SSD, Faster RCNN, YOLO, & More
 TEXT	Sentence Classification Text Classification Question Answering Summarization Text Generation, Translation, Named Entity Recognition	BERT, RoBERTa, DistilBERT, Distillbart xsum, GPT2, ELECTRA, & More
 TABULAR	Classification, Regression	LightGBM, CatBoost, XGBoost, AutoGluon TabTransformer

SageMaker JumpStart Solutions

Predictive Maintenance

*Detecting potential
equipment failure for
manufacturing*

*Predictive maintenance
for vehicle fleets*

Demand Forecasting

*Deep Demand
Forecasting*

Fraud Detection

*Detect Malicious Users
and Transactions*

*Fraud Detection in
Financial Transactions*

Credit Risk Prediction

*Corporate Credit
Rating Prediction*

*Explain Credit
Decisions*

Extract & Analyze Data from Documents

*Document
Understanding*

*Handwriting
Recognition*

*Intelligently fill in
missing form data*

Privacy-based NLP

Computer Vision

*Product Defect
Detection*

Autonomous Driving

*Visual Perception with
Active Learning*

*Reinforcement
Learning for
Battlesnake AI*

Personalized Recommendations

*Entity resolution for
smart advertising*

Purchase Modeling

Churn Prediction

Churn Prediction

Easily deploy or fine-tune models

Deploy or fine-tune pre-trained models with a single click

Easily manage assets from Amazon SageMaker JumpStart

Open pre-populated notebooks to perform inference on deployed models

The screenshot displays the Amazon SageMaker Studio interface. On the left, the 'SAGEMAKER JUMPSTART LAUNCHED ASSETS' sidebar lists several pre-trained models, all marked as 'In Service' and updated '4 days ago'. The main workspace shows a notebook titled 'ResNet 18' with a 'Deploy Model' section containing a 'Deploy' button and a 'Fine-tune Model' section containing a 'Train' button. The interface includes a top menu bar with options like File, Edit, View, Run, Kernel, Git, Tabs, Settings, and Help, and a bottom status bar showing 'Git: idle'.

Easily launch pre-built solutions

Launch solutions through cloud formation with a single click

Easily manage assets from Amazon SageMaker JumpStart

Open pre-populated notebooks for solutions to solve the business problems end to end

The screenshot displays the Amazon SageMaker Studio interface. On the left, a sidebar titled 'SAGEMAKER JUMPSTART LAUNCHED ASSETS' lists various solutions with their status and last updated time. The main panel shows the 'Demand Forecasting with Deep Learning' solution page, which includes a 'Launch Solution' button and a detailed overview of the project, including a table of normalized data.

time	ts_0	ts_1	ts_2	ts_3	ts_4	ts_5	ts_6	ts_7	ts_8	...
0 2012-01-01 00:00:00	0.102190	0.233108	0.389351	0.354701	0.393053	0.532258	0.176829	0.416254	0.376667	...
1 2012-01-01 01:00:00	0.131387	0.310811	0.519135	0.475214	0.533821	0.686996	0.176829	0.546085	0.451667	...
2 2012-01-01 02:00:00	0.153285	0.324324	0.519135	0.478632	0.497258	0.625000	0.176829	0.507929	0.450000	...

Use JumpStart models through public APIs

- Value to customers
 - Use all 300+ models for 15 ML tasks in SageMaker notebooks through SageMaker Python SDK
 - Automate training and inference in production pipelines
- SDK documentation:
 - <https://sagemaker.readthedocs.io/en/stable/overview.html#use-prebuilt-models-with-sagemaker-jumpstart>
- Launch blog:
 - <https://aws.amazon.com/blogs/machine-learning/amazon-sagemaker-jumpstart-models-and-algorithms-now-available-via-api/>

Example code:

```
# Create SageMaker Python SDK Estimator instance
od_estimator = Estimator(
    role = aws_role,
    image_uri = train_image_uri,
    source_dir = train_source_uri,
    model_uri = train_model_uri,
    entry_point = "transfer_learning.py",
    instance_count = 1,
    instance_type = training_instance_type,
    max_run = 360000,
    hyperparameters = hyperparameters,
    output_path = s3_output_location,
)

od_estimator.fit( {"training": training_dataset_s3_path} )

finetuned_predictor = od_estimator.deploy(
    initial_instance_count=1,
    instance_type=inference_instance_type,
    predictor_cls=Predictor, endpoint_name=endpoint_name
)
```

Problem types available in JumpStart

Task	Inference with pre-trained models	Training on custom dataset	Frameworks supported	Example Notebooks
Image Classification	yes	yes	PyTorch, TensorFlow	Image Classification
Object Detection	yes	yes	PyTorch, TensorFlow, MXNet	Object Detection
Semantic Segmentation	yes	yes	MXNet	Semantic Segmentation
Instance Segmentation	yes	no	MXNet	Instance Segmentation
Image Embedding	yes	no	TensorFlow, MXNet	Image Embedding
Text Classification	yes	yes	TensorFlow	Text Classification
Sentence Pair Classification	yes	yes	TensorFlow, Hugging Face	Sentence Pair Classification
Question Answering	yes	yes	PyTorch	Question Answering
Named Entity Recognition	yes	no	Hugging Face	Named Entity Recognition
Text Summarization	yes	no	Hugging Face	Text Summarization
Text Generation	yes	no	Hugging Face	Text Generation
Machine Translation	yes	no	Hugging Face	Machine Translation
Text Embedding	yes	no	TensorFlow, MXNet	Text Embedding
Tabular Classification	yes	yes	LightGBM, CatBoost, XGBoost, Linear Learner, AutoGluon, TabTransformer	Tabular Classification – LightGBM , CatBoost , XGBoost , Linear Learner , AutoGluon , TabTransformer
Tabular Regression	yes	yes	LightGBM, CatBoost, XGBoost, Linear Learner, AutoGluon, TabTransformer	Tabular Regression – LightGBM , CatBoost , XGBoost , Linear Learner , AutoGluon , TabTransformer

Thank you!

