

Revolutionizing Fintech with AWS SageMaker AND Go in ML

By: Likhit Mada



The Fintech Revolution: Market Overview

\$1009.10B

Projected Market

The global fintech market value by 2033, representing massive growth opportunities for institutions embracing technological innovation.

40%

Development Boost

Average reduction in model development time experienced by financial institutions implementing AWS SageMaker for their ML workflows.

Accuracy Increase

Improvement in fraud detection accuracy when implementing sophisticated machine learning models through AWS SageMaker.

Financial institutions face increasing pressure to innovate while maintaining strict compliance standards. The adoption of advanced machine learning solutions like AWS SageMaker represents a critical competitive advantage in this rapidly evolving marketplace.

65%

Challenges in Fintech ML Implementation



Financial institutions encounter formidable challenges when implementing machine learning solutions. Protecting sensitive customer financial data remains the highest priority, while adhering to stringent regulatory frameworks like GDPR and PSD2 introduces additional layers of complexity. Organizations frequently struggle to acquire and retain the specialized technical talent needed to design, train, and operationalize sophisticated ML models at enterprise scale. AWS SageMaker directly addresses these pain points by offering a comprehensive, integrated platform that dramatically simplifies infrastructure management while maintaining rigorous enterprise-grade security protocols and compliance standards essential for financial services.

Introducing AWS SageMaker

Integrated ML Workflow

SageMaker provides a complete environment for building, training, and deploying machine learning models with minimal infrastructure management.

Framework Flexibility

Support for popular ML frameworks including TensorFlow, PyTorch, and MXNet, allowing teams to work with familiar tools while leveraging AWS infrastructure.

Pre-built Algorithms

Access production-ready algorithms for common fintech use cases including fraud detection (XGBoost), credit scoring (Linear Learner), and customer segmentation (K-Means).

Automated Infrastructure

Eliminates the complexity of provisioning and managing specialized ML infrastructure while ensuring optimal performance and cost efficiency.

AWS SageMaker democratizes machine learning by removing infrastructure barriers and providing tooling that supports the entire ML lifecycle. This empowers fintech organizations to focus on building sophisticated models rather than managing complex infrastructure.

AWS Sagemaker Studio and Unified Studio



Platform

A fully managed service that enables data scientists and developers to seamlessly build, train, and deploy machine learning models.With unified Studio advanced project management and collaboration tools are also provided



IDE

Unified web-based interface that provides integrated development environment (IDE) through which you can perform all machine learning development steps and deploy at scale too. Including JupyterLab, Code Editor based on Code – OSS (Visual Studio Code Open Source), and RStudio



Kernels and Infra

Development environments can be spun up quickly with just a few clicks, and they come pre-configured with the necessary machine learning libraries and kernels. Users can also select specific instance types to optimize performance and cost

The platform also offers capabilities for model tuning, monitoring, debugging, experimenting and updating. Check all the training jobs, deployed endpoints, build and maintain any pipelines unifying the tools needed for machine learning development into one interface, making it easier for data scientists and developers to collaborate and manage their machine learning (ML) workflows.

End 2 End FLow

AWS Sagemaker Studio Setup

Amazon SageMaker × Getting started Studio Studio Lab [2] Canvas RStudio TensorBoard ▼ Admin configurations Domains Role manager	Amazon SageMaker SageMak The first developn (IDE) for r	ter Studio fully integrate nent environm machine learn	d eet Star Select user default-20 nent open s ing.	ted profile 0240530t070037 wudio	
Images Lifecycle configurations	How it works				
SageMaker Studio > Jupyterlab > Ju	pyterLab			🗭 Provide feedback	()
Image: Applications (5) ^ Image: Applications (5) ^	JupyterLab JupyterLab • 5 GB • ml.t3.lar C Run space Status Status Status Status	Private ge Instance ml.t3.medium	Image 🕢 SageMaker Distribution 1.7	•	
 A Home 	Space Settings (New) A space is a named, self-contained	, durable storage container (like a fi	lesystem), to which an app can be att	GLearn about Spac	res
로 Data ✓ 옪 Auto ML 표 Experiments	Storage (GB) 5 Enter a value from 5 to 100 GB. Pleas administrator for larger storage volue	Lifecycle Config No Script e contact your ne.	juration	Attach custom EFS filesystem - optional None	
Name	Application	Status	Туре	Last modified	Action
JupyterLab	C JupyterLab	🕢 Running	Private	50 seconds ago	Stop Open 🖉
1 results Results are c	ached C Refresh			Go to page 1	Page 1 of 1 < >

6

AWS Sagemaker Studio Setup





JumpStart Deploy, fine-tune, and evaluate pre-trained models from the most popular model hubs.

Providers 14			
Q Search providers or models			
•	Meta	Al21 labs	stability ai
HuggingFace	Meta	AI21	Stability Al
Explore hundreds of popular and trending models from HuggingFace.	Explore popular and trending models from Meta including Llama, Code Llama, and more.	Explore popular and trending models from AI21 Labs including Jurassic and more.	Explore popular and trending models from Stability.ai including Stable Diffusion and more.
View 334 models >	View 36 models >	View 6 models >	View 11 models >
Cohere	•	^{••} PyTorch	🤨
Cohere	TensorFlow	PyTorch	Upstage
Explore popular and trending models from Cohere including Command, Rerank, and more.	Explore popular and trending models from TensorFlow for computer vision and NLP tasks.	Explore popular and trending models from PyTorch for computer vision and NLP tasks.	Explore popular and trending models from Upstage including Solar mini chat model and more.
View 10 models >	View 319 models >	View 34 models >	View 2 models >



	Mistral 7B
About	Notebooks
Task: F	Ine-Tune: Instruction Tuning
() This	is a read-only preview of the sample notebook.
	SageMaker JumpStart - text generation instruction tuning
	This notebook demonstrates how to use the SageMaker Python SDK to fine-tune a JumpStart text generation r
[]:	<pre>import json from sagemaker.jumpstart.estimator import JumpStartEstimator from sagemaker.jumpstart.utils import get_jumpstart_content_bucket</pre>
	Select your desired model ID. You can search for available models in the Built-in Algorithms with pre-trained M
11	model_id = "huggingface-llm-mistral-7b"
	If your selected model is gated, you will need to set accept_eula to True to accept the model end-user licens
[]:	accept_oula = False
	Prepare dataset
	The following cell identifies the S3 location of a dataset that can be used to fine-tune this text generation mod
	contains questions posed by human annotators on a set of Wikipedia articles. In addition to questions with ans
	questions. Such questions are plausible, but cannot be directly answered from articles' content. The curated of
	section A.1. Dataset format.
[]:	train_data_bucket = get_jumpstart_content_bucket()
	train data location _ fraining-datasets/genuq/seatt/"

	Train	Deploy	Evaluate	
				_
		C Open in	lupyterLab	
odel using an instruction t	uning datas			
ouel using an instruction o	uning uata:			
del Table.				
agreement (EUI A)				
agreement (corx).				
l. The dataset is a subset o	f the SQUA	D2.0 datase	that	
vers, SQuAD2.0 contains a	bout 50k u	nanswerable		
taset only uses unanswer	ble questio	ons for demo	instration	
For details on the instruct	tion tuning	dataset form	iat, see	

Sagemaker and AWS Integration



AWS SageMaker integrates seamlessly with various other AWS services, simplifying the ML development, training, and deployment workflows.

Data preparation



JupyterLab

- Interactive exploration and analysis
- Flexible, user-friendly interface
- Ideal for individual exploration and experimentation
- You have experience with coding and prefer a flexible approach.
- The data size is manageable for individual processing.



Data Wrangler

- Visual interface for data preparation
- Pre-built functions for common tasks
- User-friendly for non-coders
- You prefer a visual interface for data preparation.
- You have limited coding experience or want to avoid writing complex code.
- The data processing tasks are relatively straightforward.

\sim	7		7
W	S EMR:	AW	S Glu
•	Large-scale data processing	•	Aut pre
•	Handles diverse data	•	' Ser
	formats	•	Inte
•	Advanced functionalities		ser
	for complex	•	Υοι
	transformations		pre
•	You have large-scale data processing needs.	•	Ma effi
•	You need to handle	•	Ser
	complex data formats and		anc
	transformations.		are
•	Parallel processing is		
	crucial for faster data		

preparation.

Glue:

utomates data reparation tasks erverless infrastructure tegrates with other AWS ervices ou need to automate data reparation tasks. lanaging large-scale data fficiently is a priority. erverless infrastructure nd pay-per-use pricing re desired.

Model Development and training and Deployment



No Code/ Low Code - Built in Algorithms

- Require no/less coding the only inputs you need to provide are the data, hyperparameters, and compute resources. This allows you to run experiments more quickly, with less overhead for tracking results and code changes.
- Library of foundational models which are based on Gen Al models like Llama, GPT-3,4, Stable diffusion and also many other ML Prebuilt algorithms like XGBoost, CatBoost, and LightGBM(Linear) esNet-18, EfficientNet (Deep learning models) k-means(Clustering algorithms) are available



Script Mode

- If the algorithm required for your model isn't available as a built-in option and you're adept at coding your own solution, you might want to use an SageMaker supported framework, often called "script mode." In this mode, you create your custom code or script in a code file using the provided IDEs.
- It provides popular machine learning frameworks such as TensorFlow, PyTorch, and Apache MXNet, making it easier for users to develop and deploy models.



5

 \bullet

External Model

- The built-in algorithms and supported frameworks should cover most use cases, but there are times when you may need to use an algorithm from a package not included in any of the supported frameworks
- In such cases you can write your own model externally and create a custom docker image with model code and necessary dependencies.
- Once your model is containerized it can be pushed to an online repository and used sagemaker for further steps

Training

Once your model is ready you can train your model against training data configuring the training instance, instance count and tuning parameters. After the model has been trained, SageMaker will automatically save the model artifacts to the specified S3 bucket. The output includes the model file (e.g., a .tar.gz file containing the trained model parameters) and other output files specified in the training script.

Built in model

)		
Edit View	Insert Cell Kernel Widgets Help Not Trusted	Python 3 O
+ % 🛙 🕻	h \bigstar ψ H Run \blacksquare C ψ Code \checkmark	
Out[39]:	'us-east-1'	*
In [40]: 谢	<pre>1 dockercontainer=sagemaker.amazon.amazon_estimator.get_image_uri(sagemakerSess.boto_region_name,'linear-learner',' 4</pre>	'latest')
	'get_image_uri' method will be deprecated in favor of 'ImageURIProvider' class in SageMaker Python SDK v2.	
In [41]: 🛛	<pre>LogisticModel=sagemaker.estimator.Estimator(image_name=ECRdockercontainer,</pre>	
	WARNING:root:Parameter image_name will be renamed to image_uri in SageMaker Python SDK v2.	
In [42]: 🕨	<pre>1 LogisticModel.set_hyperparameters(predictor_type='binary_classifier',mini_batch_size=100)</pre>	
In [43]: 🕨	1 LogisticModel.hyperparameters()	
Out[43]:	{'predictor_type': 'binary_classifier', 'mini_batch_size': 100}	
In [44]: 🕨	<pre>1 trainConfig=sagemaker.session.s3_input(s3_data=s3Train,content_type='text/csv')</pre>	
	WARNING:sagemaker:'s3_input' class will be renamed to 'TrainingInput' in SageMaker Python SDK v2.	
Tn [45] · N	1 LogisticModel_fit({'train': trainConfig})	

from sagemaker import image_uris container = image_uris.retrieve('xgboost', region='us-east-1', version='latest')

sess = sagemaker.Session() xgb = sagemaker.estimator.Estimator(container, role, instance_count=1, instance_type='ml.m4.xlarge', sagemaker_session=sess) xgb.set_hyperparameters(max_depth=5, eta=0.2, gamma=4, min_child_weight=6, subsample=0.8, silent=0, objective='binary:logistic', num_round=100)

kgb.fit({'train': s3_input_train, 'validation': s3_input_validation})

Ð

```
output_path='s3://{}/{}/output'.format(bucket, prefix),
```

Script Mode - Using Tensor Flow



38 shuffle=True)() 39

```
output path=model artifact location,
code location=custom code upload location,
train instance type='ml.c4.xlarge',
```

```
INFO: sagemaker: Creating training-job with name: sagemaker-tensorflow-2018-06-02-19-50-01
```

Own model - Custom Image			Custom container		
Decision_Tree Dockerfile			import tensorflow a ing import ing import fri impo fri import os fri from fri from fri from tensor	orflow as tf arse flow import keras	
<pre>**sh # The name of our algorithm algorithm_name=custom-algorithm-sklearn cd Algo_Container chmod +x Decision_Tree/train chmod +x Decision_Tree/serve fullname="\${algorithm_name}:latest" sm-docker buildrole Sagemaker build rolerepository \${fullname} ``</pre>		@ ↑ ↓ 古 무 ■	frei from frei from frei from HEI HEIG DEF WIDT NUR DEPT NUR DEPTH = 32 DEPTH = 3 NUR_CLASSES	flow.keras.layers import Inpu flow.keras.models import Mode flow.keras.utils import multi flow.keras.optimizers import a de files	
bucket sagemaker-tutorials-mlhub		Amazon SageMaker > Training jobs			
<pre>[9]: param_dict = { "max_leaf_nodes":3, "random_state": 0, "criterion":"gini" }</pre>		Training jobs Info Q. Search training jobs Name	Creation time	C Actions Create training job < 1 > Job status ▼ Warm pool status Time left	
<pre>[*]: estimator = sage.estimator.Estimator(image, role, 1, "ml.c4.2xlarge", output_path="s3://{}/output".format(sess.default_bucket()), sagemaker_session=sess, hyperparameters=param_dict, train_use_spot_instances=True, I# Specify the use of spot instances train_max_run=3600, # Maximum time for training job in seconds (optional) train_max_wait=7200 }</pre>		Custom-algorithm-sklearn-2023-09-16-06-32- 24-237 ▲ S3 model artifact s3://sagemaker-ap-south-1 sklearn-2023-09-16-06-32-	9/16/2023, 12:02:24 РМ -179822996285/output/ 24-237/output/model.ta	⊘InProgress	
<pre>26]: estimator.fit(data_location, logs=True)</pre>	◎ ↑ ↓ 봅 두 章			1	



Own model - Custom Image

Job name vortube-1 Mainum of 63 alphanumeric duracters. Can include hyphens (-), but not spaces. Must be unique within your account is na MXS Region. Vortice Marcen SagetAaler reguites permissions to call other services on your behalt. Choose a role or let us create a role tot. Vortice Algorithm options Use annacon SagetAaler built-in algorithm, your own algorithm, or a third-party algorithm from AVXS Marketpalce. Vortown algorithm resource Output data configuration Vortown algorithm resource Vortown algorithm resource Vortown algorithm source Nanzon SagetAaler built-in algorithm Learn more (C) Anazon SagetAaler built-in algorithm team more (C) Anazon SagetAaler built-in algorithm source Vortown algorithm resource Provide container ECR path Container Region annozonous ECR, Learn more In registry path where the training intage is stored in Anazon ECR, Learn more In registry path where the training intage is stored in Anazon ECR, Learn more In registry path where the training intage is stored in Anazon ECR, Learn more In registry path where the training intage is stored in Anazon ECR, Learn more In registry path where the training intage is stored in Anazon ECR, Learn more In registry path where the training intage is stored in Anazon ECR, Learn more In registry path where the training intage is stored in Anazon ECR, Learn more In registry path where the training intage is stored in Anazon ECR, Learn more In registry path where the training intage is st	Job settings	Data source S3
Auto name youtube: Namium of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an WS Region. Namium of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an WS Region. Algorithm options Use an Anzeon SageMaker built-in algorithm, your own algorithm, or a third-party algorithm from AWS Marketplace. Algorithm source In anzeon SageMaker built-in algorithm cean more (2) Anazon SageMaker built-in algorithm cean more (2) Anagorithm source Provide container ECR path In registry path where the training image is stored in Anazon ECR. Lean more In registry path where the training image is stored in Anazon ECR. Lean more In registry path where the training image is stored in Anazon ECR. Lean more In container In registry path where the training image is stored in Anazon ECR. Lean more In container In registry path where the training image is stored in Anazon ECR. Lean more In container Add tag	Joh name	S3 data type S3 data distribution type
Voltage: Values:	Job name	S3Prefix FullyReplicated
Maintum of & Jaghanumer: characters. Can include hyphens [-], but not spaces. Must be unique within your account in an AWS Bighon man and WS Bighon man and	youtube-1	S3 location
IAM role Amazon SageMaker requires permissions to call other services on your behalf. Choose a role or let us create a role that have an SageMaker built-in algorithm. your own algorithm from AWS Marketplace. ✓ Algorithm options Use an Amazon SageMaker built-in algorithm, your own algorithm from AWS Marketplace. ✓ Algorithm source ③ Amazon SageMaker built-in algorithm Learn more ③ ④ Your own algorithm resource ④ Your own algorithm container in ECR Learn more ④ ④ An algorithm subscription from AWS Marketplace. ✓ Provide container ECR path Container The registry path where the training image is stored in Amazon ECR. Learn more ① Accountid.dkr.ecr. Region amazonews.com/repository(tog) or [@digest]	Maximum of 63 alphanumeric characters. Can include hyphens (-), but not spaces. Must be unique within your account in an AWS Region.	s3://bucket/path-to-your-data/
Algorithm options Use an Amazon SageMaker built-in algorithm, your own algorithm, or a third-party algorithm from AWS Marketplace. Algorithm source Algorithm source Your own algorithm resource Your own algorithm container in ECR Learn more Your own algorithm subscription from AWS Marketplace Provide container ECR path Container The registry path where the training image is stored in Amazon ECR. Learn more accountid dir. der. Region. annazonows.com/repository[tog] or [@digest]	IAM role Amazon SageMaker requires permissions to call other services on your behalf. Choose a role or let us create a role that has the AmazonSageMakerFullAccess IAM policy attached.	Add channel
Algorithm options Use an Amazon SageMaker built- in algorithm, your own algorithm from AWS Marketplace. Algorithm source Amazon SageMaker built- in algorithm Learn more Your own algorithm resource Your own algorithm container in ECR Learn more An algorithm subscription from AWS Marketplace Provide container ECR path Container The registry path where the training image is stored in Amazon ECR. Learn more accountid.dkr.ecr. Region. amazonaves.com/repository[stog] or (@digest]	•	
 Algorithm source Amazon SageMaker built-in algorithm Learn more C Your own algorithm resource Your own algorithm container in ECR Learn more C An algorithm subscription from AWS Marketplace Provide container ECR path Container The registry path where the training image is stored in Amazon ECR. Learn more c accountif.d.dkr.etr. Region amazonows.com/tepositoryf.tag] or [@digest] 	Algorithm options Use an Amazon SageMaker built-in algorithm, your own algorithm, or a third-party algorithm from AWS Marketplace.	Output data configuration
Amazon SageMaker built-in algorithm Learn more Your own algorithm resource Your own algorithm container in ECR Learn more An algorithm subscription from AWS Marketplace	▼ Algorithm source	S3 output path
 Amazon SageMaker built-in algorithm Learn more Your own algorithm resource Your own algorithm container in ECR Learn more An algorithm subscription from AWS Marketplace Provide container ECR path Container The registry path where the training image is stored in Amazon ECR, Learn more accountid.dkr.ecr.Region.amazonaws.com/repository[.tog] or [@digest] 		s3://sagemaker-us-east-2-468564051655/output
 Your own algorithm resource Your own algorithm container in ECR Learn more An algorithm subscription from AWS Marketplace Provide container ECR path Container Tags - optional Key Value Remove Add tag 	Amazon SageMaker built-in algorithm Learn more	Encryption key - optional
Your own algorithm container in ECR Learn more An algorithm subscription from AWS Marketplace Provide container ECR path Container The registry path where the training image is stored in Amazon ECR. Learn more <i>accountid dk: ecr. Region. amazonows.com/repository[:tag] or [@digest]</i> Add tag	Your own algorithm resource	instead of the default S3 service key, provide its ID or ARN.
An algorithm subscription from AWS Marketplace Provide container ECR path Container The registry path where the training image is stored in Amazon ECR. Learn more occountid.dkr.ecr. Region.amazanows.com/repository[:tag] or [@digest] Not based:	Your own algorithm container in ECR Learn more	
 Provide container ECR path Container Container The registry path where the training image is stored in Amazon ECR. Learn more accountid.dki.ecr.Region.amazanows.com/repository[:tag] or [@digest] Insutement 	An algorithm subscription from AWS Marketplace	
Container The registry path where the training image is stored in Amazon ECR: Learn more occountid dkr.ecr. Region amazonows.com/repository[:tag] or [@digest] New Value Add tag	Provide container ECR path	
Container The registry path where the training image is stored in Amazon ECR. Learn more occountid_dkr.ecr.Region_amazonows.com/repository[:tag] or [@digest] Key Value Add tag		▼ Tags - optional
The registry path where the training image is stored in Amazon ECR. Learn more occountid_dkr.ecr.Region.omozonows.com/repository[:tog] or [@digest] Remove Add tag	Container	Kau Value D
occountid_dkr.ecr.Region.amazanaws.com/repository[:tag] or [@digest] Add tag	The registry path where the training image is stored in Amazon ECR. Learn more	Remove
to a second s	occountid.dkr.ecr.Region.amazonows.com/repository[:tag] or [@digest]	Add tag
	teres teres de	

Deployment

After training your machine learning model, you can deploy it with Amazon Sagemaker to make predictions. Amazon SageMaker provides several deployment options to suit different needs:



Real-time Hosting Services

For continuous, real-time inferences against the model. Best in the cases where there is continuous traffics The payload is of smaller size. Deployed preferred instance type

Serverless Inference

Suitable for workloads that experience periods of inactivity interspersed with spikes in traffic and can handle occasional startup delays, opt for Serverless Inference.

For scenarios requiring the
processing of large payloads up to
1GB, extended processing times,
and near real-time response
speeds, consider using Amazon

Asynchronous Inference

2	 5	

SageMaker Asynchronous

create Hugging Face Model Class huggingface model = HuggingFaceModel(env=hub, # configuration for loading model from Hub role=role, # iam role with permissions to create an Endpoint transformers_version="4.6", # transformers version used pytorch version="1.7", # pytorch version used py_version='py36', # python version used # create Transformer to run our batch job batch_job = huggingface_model.transformer(instance count=1, instance type='ml.p3.2xlarge', output path=upload path, # we are using the same s3 path to save the output with the input strategy='SingleRecord') # starts batch transform job and uses s3 data as input batch job.transform(data=s3 file uri, content type='application/json', split_type='Line')



Batch Transform

If you need to make predictions on an entire dataset, SageMaker Al's batch transform feature is the appropriate choice. See Batch transform for inference with Amazon SageMaker AI for further

SageMaker ML Lifecycle Management



AWS SageMaker delivers specialized tools for each critical phase of the machine learning lifecycle. Financial institutions benefit from streamlined data labeling with SageMaker Ground Truth, comprehensive exploratory analysis through SageMaker Studio notebooks, and sophisticated hyperparameter optimization that eliminates guesswork in model refinement.

In the fintech sector, where model precision directly translates to financial outcomes and risk management, SageMaker's advanced monitoring capabilities provide essential safeguards. The platform employs real-time performance tracking in production environments, automatically identifying statistical deviations that could compromise accuracy in mission-critical operations such as fraud prevention, credit risk assessment, and algorithmic trading decisions.

Design algorithms for precise financial predictions

Maximize accuracy with automated optimizations

Seamlessly transition models to production

Additional Features



Sagemaker Experiments

Building an ML model requires numerous iterations of training to fine-tune the algorithm, model architecture, and parameters for optimal prediction accuracy.

With Amazon SageMaker Experiments, you can monitor the inputs and outputs throughout these training iterations, enhancing the repeatability of trials and facilitating team collaboration.

Additionally, you can track parameters, metrics, datasets, and various other artifacts associated with your model training jobs.

It provides a consolidated interface that allows you to view your ongoing training jobs, collaborate with your team on experiments, and directly deploy models from an experiment.



Sagemaker Pipelines

Amazon SageMaker Pipelines is a fully managed CI/CD service for machine learning workflows, enabling *automated, scalable, and reproducible ML model development. It allows data scientists and engineers to define end-to-end ML workflows—from data preprocessing and training to model evaluation and deployment—using a **directed acyclic graph (DAG) * structure. Pipelines integrate seamlessly with SageMaker features like Processing Jobs, Training Jobs, and Model Registry, while supporting *parameterization, conditional

execution, and reuse of pipeline steps*.

adoption.



Canal	Makar	Madal	Manitar
baye	viakei	model	WOIIILOI

5	
5	
5	
	-
	•
	•
l	
1	
l	
l	
1	
•	
,	

Amazon SageMaker Model Monitor	Ac
oversees the quality of AI machine	са
learning models in production on	Ur
Amazon SageMaker.	CU
It enables continuous monitoring	ар
through real-time endpoints, regularly	Th
scheduled batch transform jobs, or	wi
on-schedule monitoring for	m
asynchronous batch transform jobs.	ma
	W

Model Monitor also comes with alert functionalities to notify users of any deviations in model quality, facilitating early detection and enabling proactive corrective measures

Bedrock and Unified studio

- ccess Amazon Bedrock's
- pabilities through SageMaker
- nified Studio to quickly build and
- stomize your generative Al
- pplications.
- his intuitive interface lets you work
- th high-performing foundation
- odels (FMs)Customize FMs to
- atch your requirements, data,
- orkflows, and responsible Al
- standards.
- Access a wide range of
- high-performing FMs from leading Al
- companies through the generative AI
- playground. You can compare
- different models and configurations
- to evaluate their performance easily.

Using Go for Machine Learning Model Development

While Go might not be the first language that comes to mind for machine learning, it offers several advantages for specific tasks and situations:

Performance and Efficiency:

- Go is a compiled language, resulting in fast execution times, critical for real-time applications and large datasets.
- Go's garbage collection is efficient, minimizing pauses and ensuring consistent performance.
- Go is lightweight and has a small memory footprint, making it suitable for deployment on resource-constrained environments like servers or edge devices.

Concurrency and Parallelism:

- Go's built-in concurrency features (goroutines and channels) make it easy to parallelize tasks, significantly speeding up training and inference.
- Go's concurrency model is lightweight and efficient, avoiding the complexities of thread management in other languages.

Integration and Interoperability:

- Go integrates seamlessly with C code, allowing access to existing libraries and optimized implementations.
- Go bindings are available for many popular ML libraries like TensorFlow, PyTorch, and scikit-learn. This enables leveraging existing code and algorithms within Go projects.
- Go's strong network capabilities facilitate communication with other services and APIs, crucial for distributed training and cloud deployments.

Go Lang Libraries for ML:

- **Gonum** Numerical computing library (like NumPy)
- Gorgonia- is a modern, high-performance deep learning library specifically designed for the Go programming language, It provides a clean and expressive API for building and training neural networks, making it a suitable choice for various ML tasks.
- **GoLearn** is a versatile and comprehensive machine learning library built for the Go programming language. It provides a wide range of algorithms and tools for various tasks, making it a valuable resource for ML practitioners.
- **GoCV** GoCV is a powerful and versatile computer vision library specifically designed for the Go programming language. It provides a comprehensive set of tools and algorithms for image and video processing, analysis, and manipulation.

Using Go with AWS Sagemaker

Here are a few scenarios on how Go can be used with AWS SageMaker:

- Data Handling and Pre/Post Processing: Although the actual machine learning models might be implemented in other languages better suited for data analysis (like Python), Go can be used for the tasks that involve data gathering, pre-processing, data storage, or post-processing results. Go's efficiency and speed make it a great choice for tasks that require high performance such as processing large volumes of data or streaming data.
- Invoking Endpoint Models: Once you have your model trained and deployed on AWS SageMaker as an endpoint, you can make predictions by invoking 2. this endpoint. You can write a Go application that sends requests to the SageMaker endpoint using the AWS SDK for Go or sagemaker-go. This is useful for integrating model predictions into Go-based applications, such as web services or backend systems. Use the AWS SDK for Go to interact with SageMaker resources programmatically. This SDK provides APIs allowing you to manage SageMaker resources, automate training job deployments, setup endpoints, and automate other workflow steps. For example, you can list training jobs, describe specific jobs, or manage lifecycle configurations directly from your Go code
- 3. AWS Lambda: Deploy a Go-based Lambda function to build event driven ML lifecycle.
- **Microservices Architecture**: If your infrastructure involves a microservices architecture, you can use Go to create lightweight, efficient microservices that 4. interact with AWS SageMaker. For instance, separate microservices could handle different aspects of the model training pipeline or could act as intermediaries for pre-processing data before it is sent to SageMaker for training or inferences
- Custom Algorithms in Container: If you need more control and want to use Go directly for developing algorithms, Use Go ML libraries like 5. [gonum](https://gonum.org/), [gorgonia](https://gorgonia.org/), or call Python models via os/exec. You can build ML models you can package your Go code as a Docker container with Go runtime and dependencies. Since AWS SageMaker supports Docker containers, you can train and deploy your ML model written in Go.

Key Fintech Use Cases

Fraud Detection

Real-time transaction monitoring using XGBoost algorithms to identify suspicious patterns and prevent fraudulent activity before losses occur. SageMaker's ability to process streaming data enables immediate response to potential fraud.

Credit Risk Assessment Advanced credit scoring models using Linear Learners and Neural Networks to evaluate borrower risk with greater accuracy. These models incorporate traditional and alternative data sources for comprehensive risk profiles.



Customer Partechigenice engines using K-Means clustering and Deep Learning to segment customers and tailor financial products. These models analyze transaction history and behavioral patterns to predict customer needs.



Portfolio Management

- Algorithmic trading
- strategies and portfolio
- optimization using
- reinforcement learning and
- time series forecasting.
- SageMaker's distributed
- training capabilities handle
- the computational demands
- of these complex models.

Security & Compliance Capabilities

Data Encryption

SageMaker automatically encrypts data at rest using AWS KMS and in transit using TLS, ensuring financial information remains protected throughout the ML lifecycle. This meets stringent requirements for PCI-DSS and other financial data protection standards.

VPC Integration

Training and inference containers run within your Virtual Private Cloud, ensuring models never process sensitive financial data on the public internet. This provides network isolation for all ML workloads handling customer financial information.

IAM Integration

AWS Identity and Access models and data. Role-based requirements for data access controls.

SageMaker's comprehensive security features address the unique compliance requirements of the financial industry, including GDPR, PSD2, and internal risk management frameworks. The platform maintains detailed audit logs of all ML activities, supporting the documentation requirements for model governance and regulatory examinations.

- Fine-grained access controls through
- Management ensure only authorized
- personnel can access specific
- permissions align with regulatory

Case Study: Major Bank Fraud Detection

A Fortune 500 bank implemented AWS SageMaker to overhaul their fraud detection system, which was suffering from high false positive rates and missed fraud attempts. By leveraging SageMaker's XGBoost implementation and real-time inference capabilities, they achieved dramatic improvements across all key performance indicators.



False Positives Reduced by **69**%

False positive rates dropped dramatically from 12.4% to just 3.8%, meaning fewer legitimate customer transactions were incorrectly flagged as suspicious.





60% Faster Model
Development
Model development time was
slashed from 45 days to just 18
days, allowing for more frequent
updates and refinements to the
fraud detection models.

Speed

This performance enables the bank to block fraudulent transactions before they complete, significantly reducing fraud losses while improving customer experience by reducing legitimate transaction declines.

Detection Rate Improved to 94%

The bank increased fraud detection success from 76% to 94%, capturing substantially more fraudulent transactions before completion.



5x Faster Processing

- Alert processing time was reduced
- from 1,200ms to just 250ms,
- enabling the system to process
- over 10,000 transactions per
- second during peak periods.

Best Practices for SageMaker Implementation

Start With Clear Business Metrics

Define specific KPIs for your ML initiative, such as fraud detection rate or credit approval accuracy. Establish baseline measurements before implementation to properly evaluate success.

Begin With Pre-built Algorithms

Leverage SageMaker's built-in algorithms for common fintech use cases before developing custom models. Use foundational models and develop your own models for better fine tuning and control

Implement CI/CD for Model Development Use AWS Lambda, SageMaker Pipelines to create automated workflows for model training, evaluation, and deployment. This ensures consistent quality and enables rapid iteration on financial models.

Mon Perfe Imple Monit produ perfor applic quick result comp

Organizations should approach SageMaker implementation with a cross-functional team that includes data scientists, compliance specialists, and business stakeholders. This ensures models meet both technical and regulatory requirements while driving measurable business results.

Monitor Model

Performance Vigilantly

- Implement SageMaker Model
- Monitor to continuously track
- production model
- performance. In financial
- applications, model drift can
- quickly impact business
- results and potentially create
- compliance issues.

Next Steps and

Resources

Assessment Phase

2

3

4

Evaluate your organization's ML maturity and identify high-value use cases that align with business objectives. Look for opportunities where existing processes have clear metrics that can be improved through predictive analytics.

Production Implementation

Scale successful models to production using SageMaker's deployment capabilities. Ensure proper monitoring and governance procedures are in place to meet compliance requirements.

To get started with AWS SageMaker for fintech applications, explore these resources:

- AWS Financial Services Industry Lens AWS Well-Architected Framework •
- SageMaker Example Notebooks for financial use cases ۲
- AWS Training and Certification courses for ML in Finance ۲
- Explore Golang ML libraries and implement models. ۲

Proof of Concept Implement a focused pilot project using AWS SageMaker to

demonstrate value. Fraud detection and credit scoring typically provide the fastest ROI for financial institutions new to machine learning.

Organizational Transformation

Develop internal ML capabilities through training and hiring. Create a center of excellence to share best practices and accelerate adoption across the organization.

Thank you