INNOVATE

AI/ML EDITION

24 February 2022

Build machine learning models with Amazon SageMaker optimal for your use case

Romina Sharifpour
AI/ML Specialist Solutions Architect, AWS

Agenda

- Considerations when building machine learning models
- Various ways we can build machine learning models
- Building ML models on Amazon SageMaker
- Demo
- Take away and best practices
- Resources

Build a model that classifies the blogs users post on a website into a number of classes

How much data I have?

How am I going to label my data?

Where am I going to build it?

How much time do I have?

aws

Considerations when building ML models

Data

- Size
- Type
- Data sources/ pipelines
- Amount of labeled data

Choice of algorithm

- Problem to be solved
- DL vs ML
- Complexity ; Implementation, Number of parameters
- Production and Management

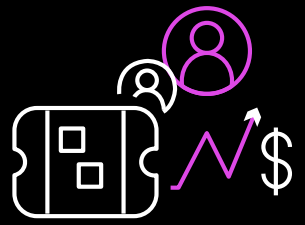

Resources

- Machine Learning expertise
- Infrastructure (i.e. Memory, CPU, GPU, managed or self managed?)
- Cost ($$$)

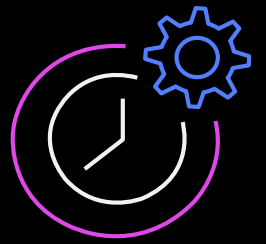

Time to production

- Time to experiment
- Time to market

Build ML models : full to no customization

CUSTOM BUILD MODELS

- Your own machine learning code
- Opensource/ proprietary implementations or frameworks

- More time to build
- More costly
- High level of ML expertise
- Most flexible

aws

Build ML models : full to no customization

PRE-TRAINED MODELS: WHY TRAIN WHEN YOU CAN FINE TUNE?

- Pre-trained model: a model that was trained on a large benchmark dataset to solve a problem similar to the one you want to solve.
  - Use the deployed version via an API call (managed)
  - Deploy directly (self-managed)
  - Fine tune (Managed or self managed)

- Saves time
- Saves experimentation and cost
- ML expertise
- None to some flexibility

Amazon Comprehend real time text analysis

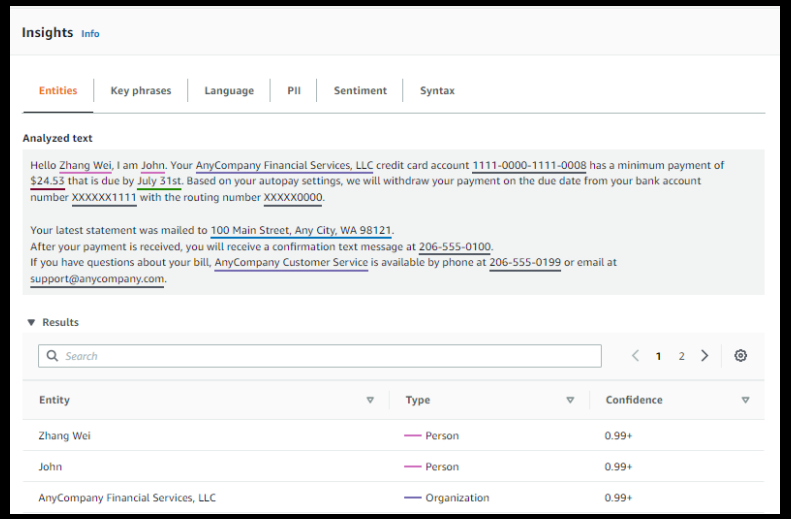

Amazon Rekognition label detection

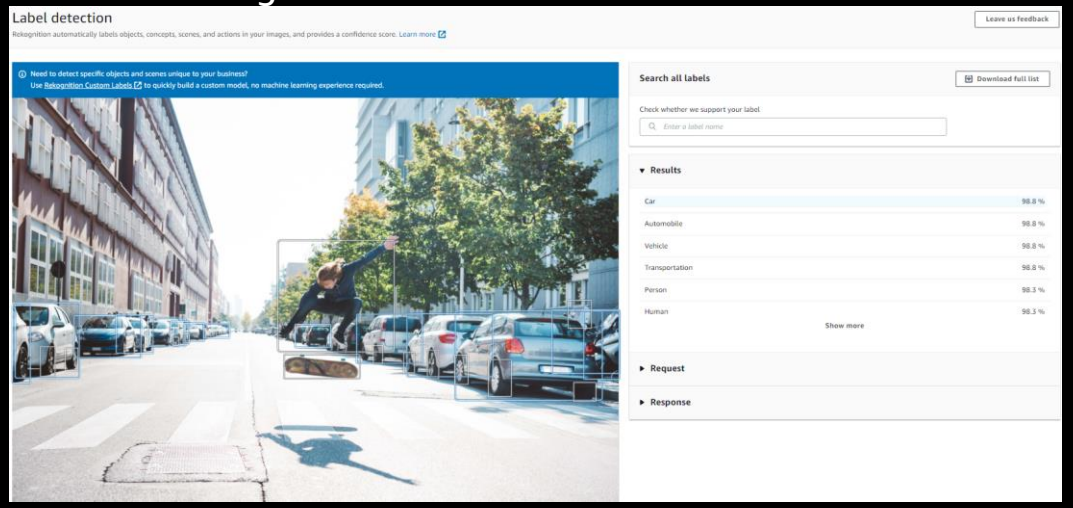

Build ML models : full to no customization

AUTOML

- AutoML is automating the end-to-end cycle of applying Machine Learning to a problem.
- Whether you are new to ML or an experienced practitioner, AutoML will simplify your workflow.
- Wide variety of options available; proprietary or open source with varied level of automation

AutoGluon

AutoKeras

- Save time
- Save experimentation
- Cost can vary
- Low code to no code
- Little to no ML expertise
- None to little flexibility

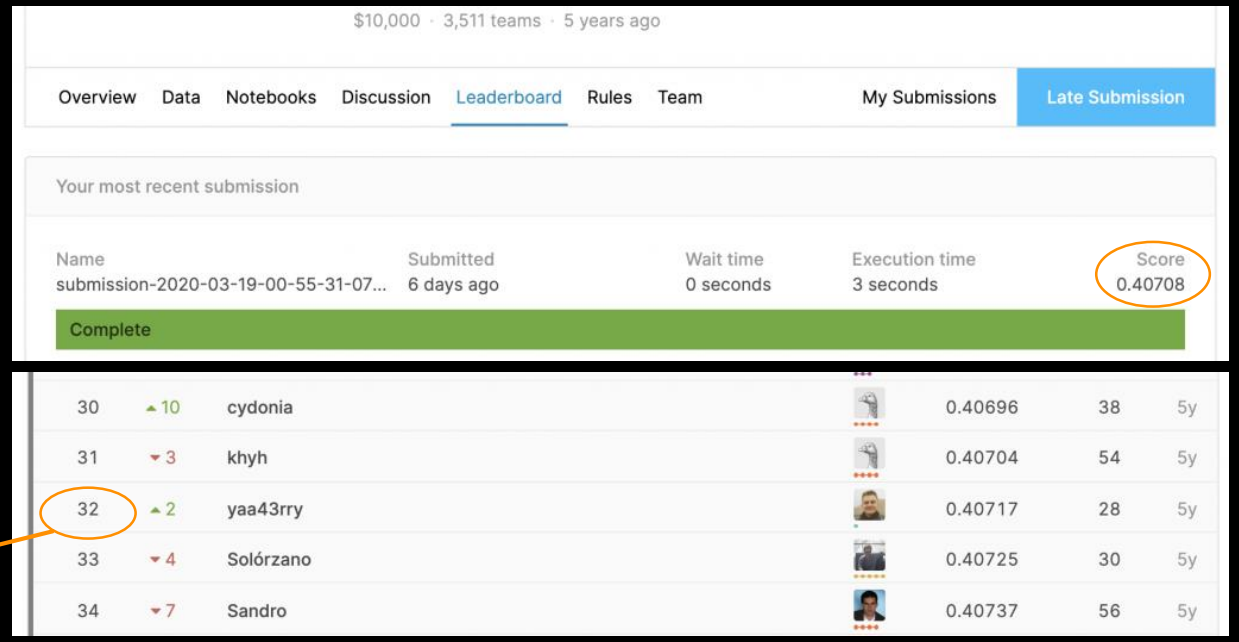

32nd position out of 3,511 submissions (under top 1% of all submissions)

Source: https://aws.amazon.com/blogs/opensource/machine-learning-with-autogluon-an-open-source-automl-library/

Build Machine Learning models on Amazon SageMaker

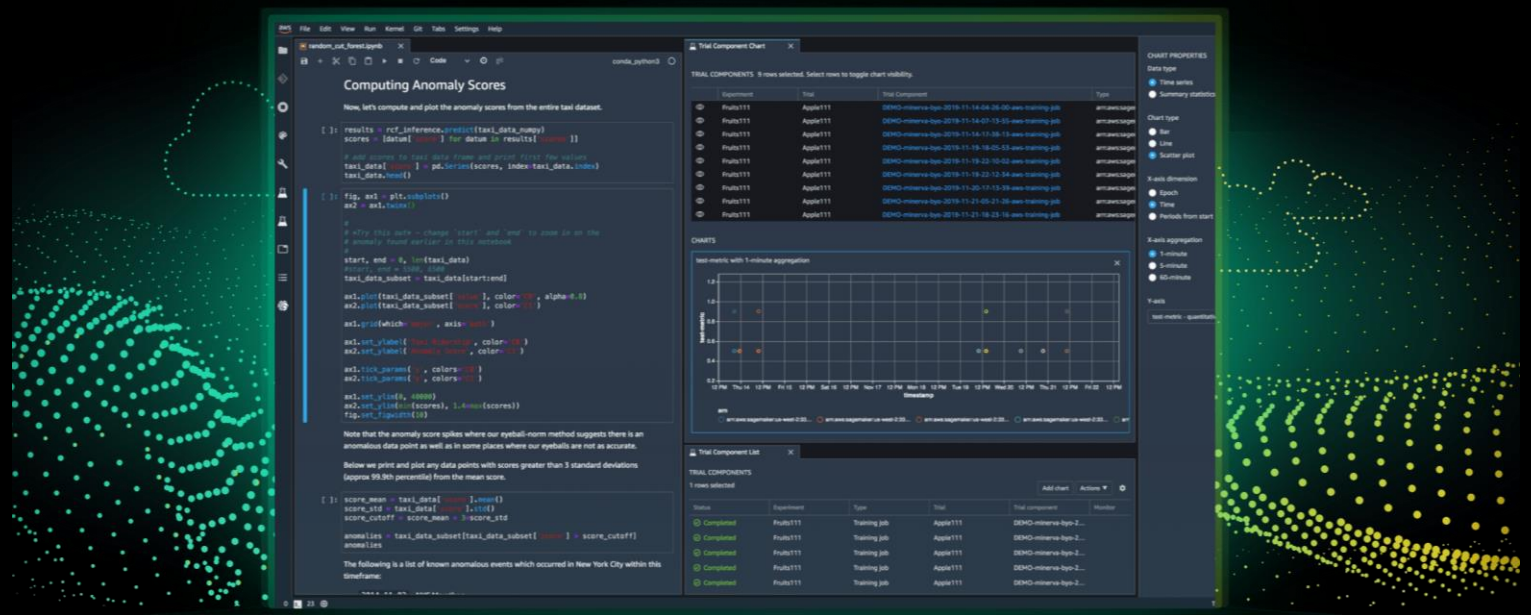

Amazon SageMaker overview

Amazon SageMaker

PREPARE

SageMaker Ground Truth
Label training data for machine learning

SageMaker Data Wrangler NEW
Aggregate and prepare data for machine learning

SageMaker Processing
Built-in Python, BYO R/Spark

SageMaker Feature Store
Store, update, retrieve, and share features

SageMaker Clarify
Detect bias and understand model predictions

BUILD

SageMaker Studio Notebooks
Jupyter notebooks with elastic compute and sharing

Built-in and Bring your-own Algorithms
Dozens of optimized algorithms or bring your own

Local Mode
Test and prototype on your local machine

SageMaker Autopilot
Automatically create machine learning models with full visibility

SageMaker JumpStart
Pre-built solutions for common use cases

TRAIN & TUNE

Managed Training
Distributed infrastructure management

SageMaker Experiments
Capture, organize, and compare every step

Automatic Model Tuning
Hyperparameter optimization

Distributed Training Libraries NEW
Training for large datasets and models

SageMaker Debugger NEW
Debug and profile training runs

Managed Spot Training
Reduce training cost by 90%

DEPLOY & MANAGE

Managed Deployment
Fully managed, ultra low latency, high throughput

Kubernetes & Kubeflow Integration
Simplify Kubernetes-based machine learning

Multi-Model Endpoints
Reduce cost by hosting multiple models per instance

SageMaker Model Monitor
Maintain accuracy of deployed models

SageMaker Edge Manager NEW
Manage and monitor models on edge devices

SageMaker Pipelines NEW
Workflow orchestration and automation

SageMaker Studio
Integrated development environment (IDE) for ML

Amazon SageMaker Studio Notebook

Perform data engineering, analytics, and ML workflows in one notebook

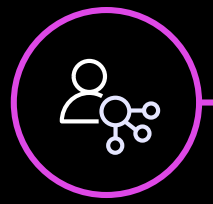

Connect with Amazon EMR, Amazon S3, and more

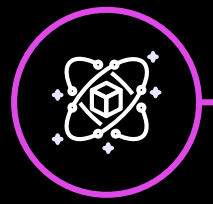

Interactively access, transform, and analyze a wide range of data

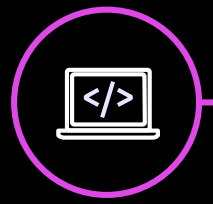

Build, train, and deploy models using your preferred framework

Model options

AMAZON SAGEMAKER

Training code

No code required

- **XGBoost**
- **Matrix Factorization**
- **Regression**
- **Principal Component Analysis**
- **K-Means Clustering**
- **And More!**

Built-in Algorithms (17)
No ML coding required

Bring your Own Script
Amazon SageMaker build the container
Open source containers

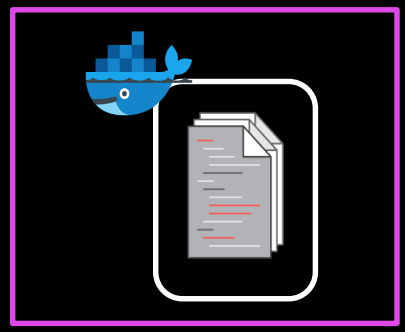

Bring your Own Container
Full control, you build the container
R, C++, etc

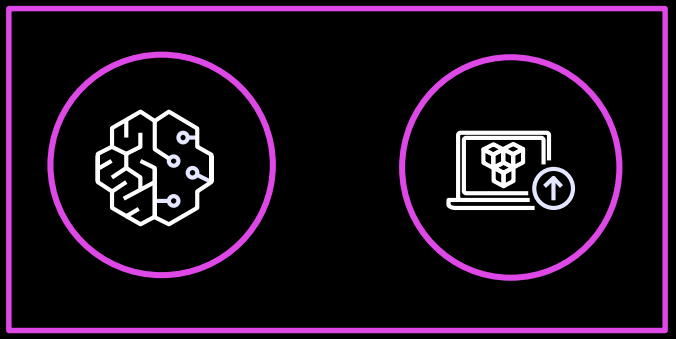

Amazon SageMaker Autopilot

Amazon SageMaker Jumpstart

Fully Managed, Distributed, Auto-Scaled and Secure

Amazon SageMaker in-built algorithms

Amazon SageMaker
has built-in algorithms
or bring your own

Computer vision
Image classification | Object detection | Semantic segmentation

Topic modeling
LDA | NTM

Classification
Linear Learner | XGBoost | KNN

Recommendation
Factorization machines

Forecasting
DeepAR

Working with text
BlazingText | Supervised | Unsupervised

Regression
Linear Learner | XGBoost | KNN

Clustering
KMeans

Sequence translation
Seq2Seq

Anomaly detection
Random cut forests | IP Insights

Feature reduction
PCA

Amazon SageMaker in-built algorithms

SAMPLE CODE

```
region=boto3.Session().region_name
container = sagemaker.image_uris.retrieve(region=region, framework='xgboost', version='latest')
print( "Using SageMaker XGBoost container: {}, ({})".format (container, region))
```

```
using SageMaker XGBoost container: 544295431143.dkr.ecr.ap-southeast-2.amazonaws.com/xgboost:latest, (ap-southeast-2)
```

```
sess = sagemaker.Session()

xgb = sagemaker.estimator.Estimator(container,
                                    role,
                                    instance_count=1,
                                    instance_type='ml.m4.xlarge',
                                    output_path='s3://{}/{}/output'.format(bucket, prefix),
                                    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)

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

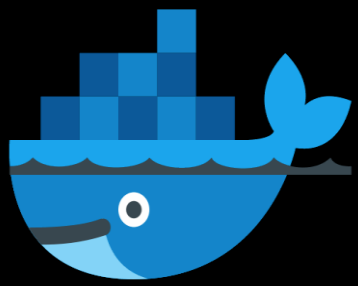

Amazon Elastic Container Registry (Amazon ECR)

aws

Bring your own script ('script mode')

HIGH LEVEL WORKFLOW

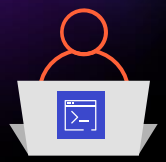

Local laptop or desktop with Amazon SageMaker SDK

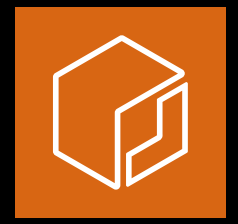

Amazon Elastic Container Registry (Amazon ECR)

Amazon Simple Storage Service (Amazon S3)

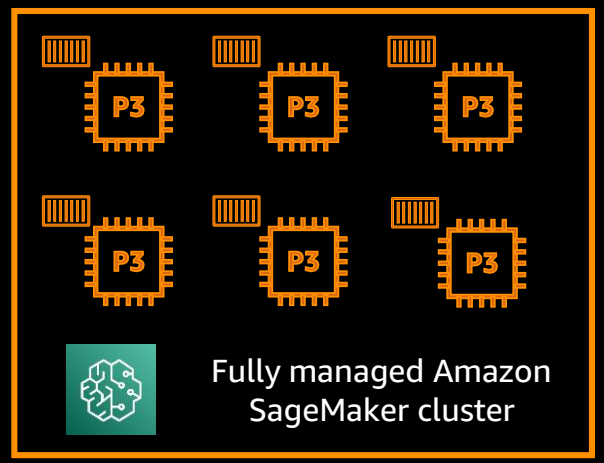

Available frameworks

Run your own script on supported frameworks

SAMPLE CODE

```
estimator = TensorFlow(entry_point='train.py',
                       train_instance_type='ml.g4dn.xlarge',
                       train_instance_count=1,
                       hyperparameters=hyperparameters,
                       framework_version='2.6.2',
                       py_version='py3',
                       script_mode=True)

estimator.fit({'training': inputs})
```

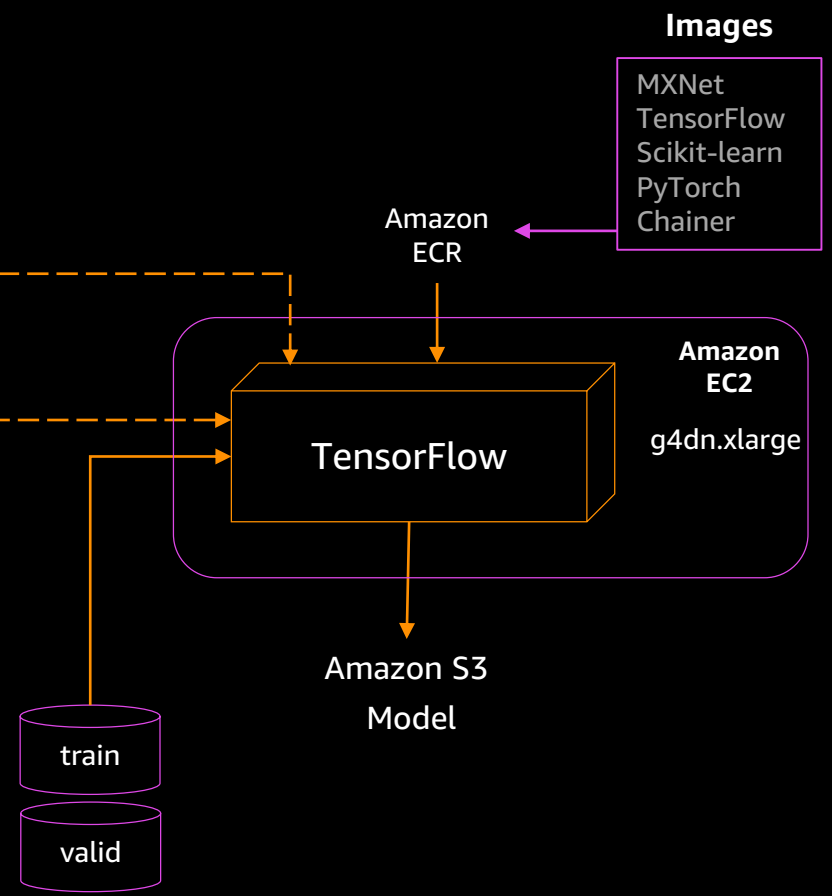

Bring your own container

HIGH LEVEL WORKFLOW

Local laptop or desktop with Amazon SageMaker SDK

Docker build

Custom container

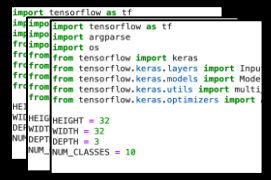

Code files

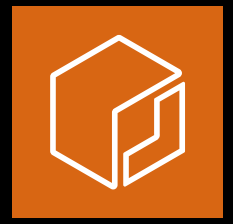

Amazon Elastic Container Registry (Amazon ECR)

Amazon Simple Storage Service (Amazon S3)

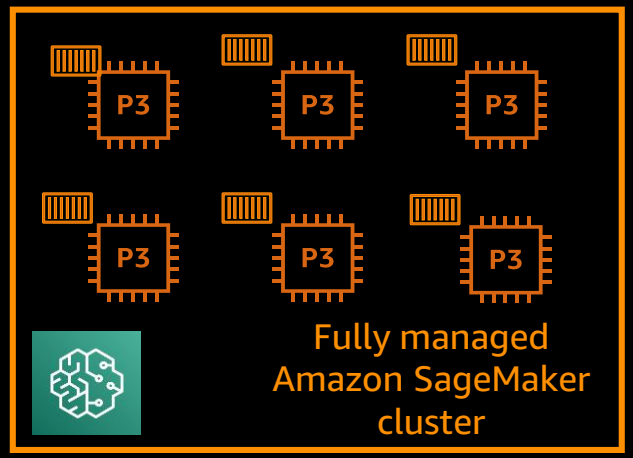

1. Create your custom container

2. Build & push docker Image to ECR

Bring your own container

SAMPLE CODE

```
estimator = Estimator(image='custom_container',
                      train_instance_type='ml.c4.xlarge',
                      train_instance_count=1,
                      hyperparameters=hyperparameters,
                      py_version='py3')



estimator.fit({'training': inputs})
```

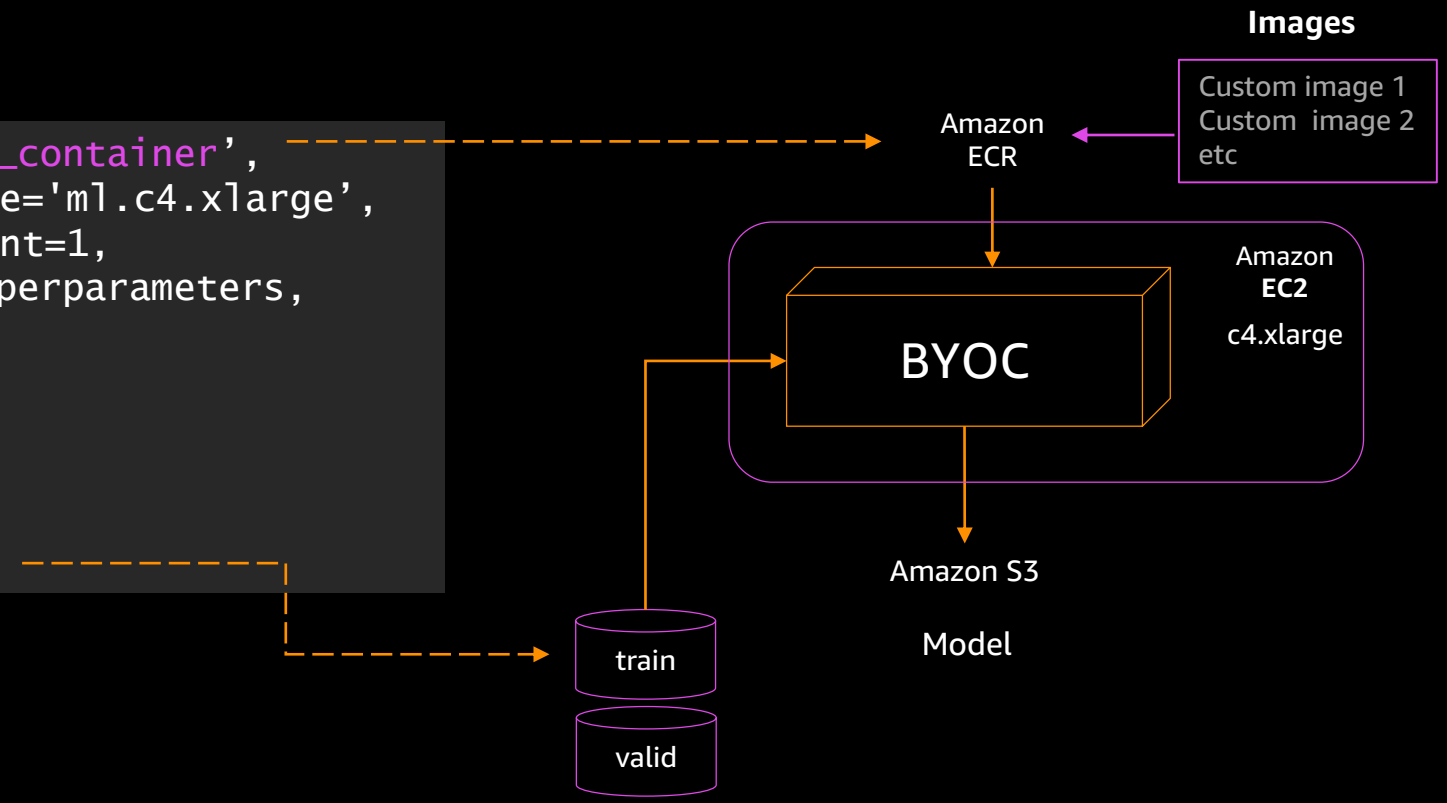

Tips : Amazon SageMaker local mode

DEVELOP CODE LOCALLY

- Amazon SageMaker Python SDK supports local mode; create estimators and deploy them to your local environment.
- Great way to test your deep learning scripts before running them in Amazon SageMaker's managed training or hosting environments.
- Local Mode is supported for frameworks images (TensorFlow, MXNet, Chainer, PyTorch, and Scikit-Learn) and images you supply yourself.
- You can choose to use an Amazon SageMaker notebook instance as your local environment.

```
estimator = PyTorch (py_version='py3',
entry_point='source/cifar10.py',
role=role,
framework_version='1.7.1',
instance_count=1,
instance_type='local')

estimator.fit(inputs)
```

aws

Amazon SageMaker JumpStart

EASILY AND QUICKLY BRING MACHINE LEARNING APPLICATIONS TO MARKET

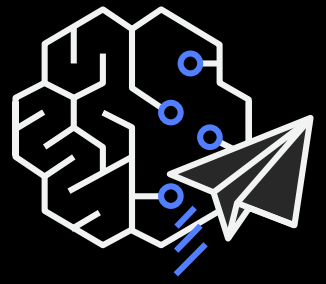

Open source models

150+ pre-trained models
from PyTorch Hub
& TensorFlow Hub
with 1-click deploy

Solutions

Solutions for common use cases
so you can move quickly from
concept to production
with 1-click deploy

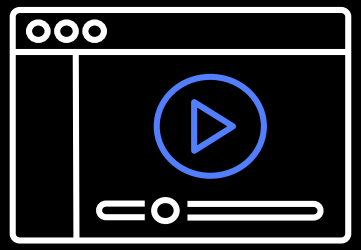

Getting started content

Examples and tutorials built to
help you get started with
machine learning faster
with 1-click deploy

aws

Open source models

150+ PRE-TRAINED OPEN SOURCE MODELS FROM PYTORCH HUB & TENSORFLOW HUB

	TASKS	MODELS
 VISION	Image Classification Object Detection	ResNet, MobileNet, SSD & More
 TEXT	Sentence Classification Text Classification Question Answering	BERT, RoBERTa, & DistilBERT

Solutions

SAMPLE ARCHITECTURE: DETECT MALICIOUS USERS AND TRANSACTIONS

- Over 15 business use cases.
- Solutions cover a wide variety of AWS services and architectures to get into production quickly.
- Can be deployed with just a few clicks.
- Customizable so you can easily modify to fit your specific use case and dataset.

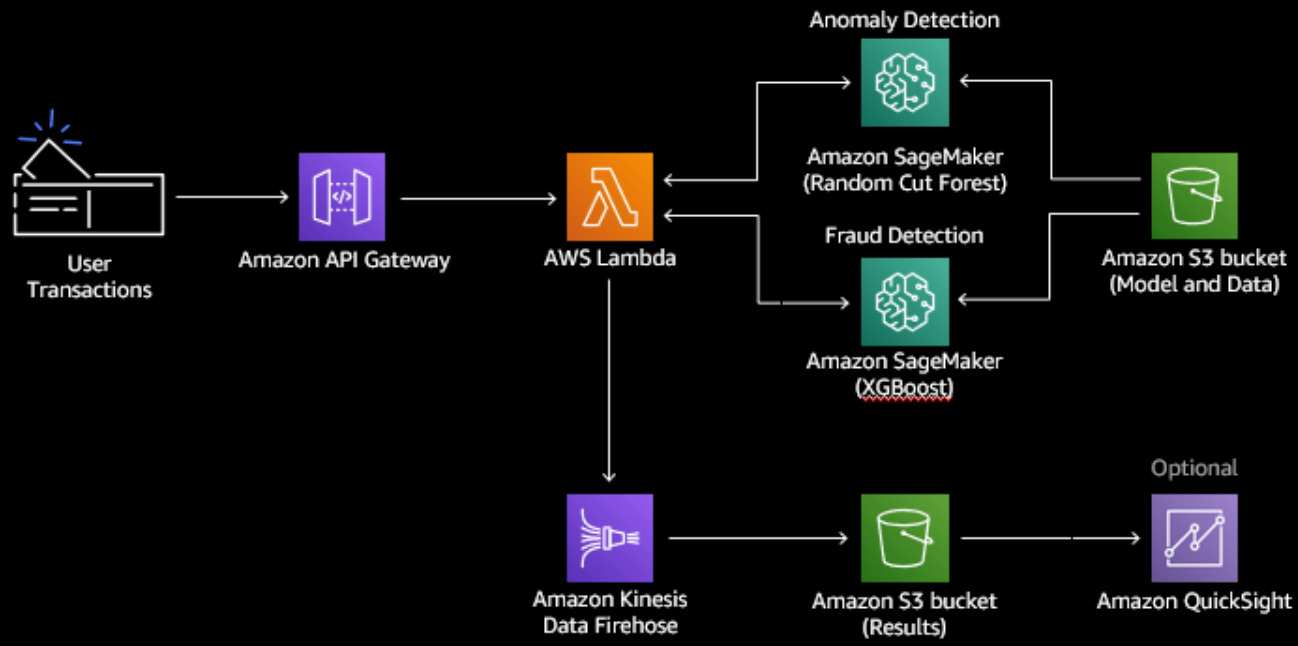

Building models with Amazon SageMaker Autopilot

Amazon SageMaker Autopilot is natively integrated with SageMaker and is the automated machine learning capability for tabular data that gives you complete visibility into your ML models

How Amazon SageMaker Autopilot works

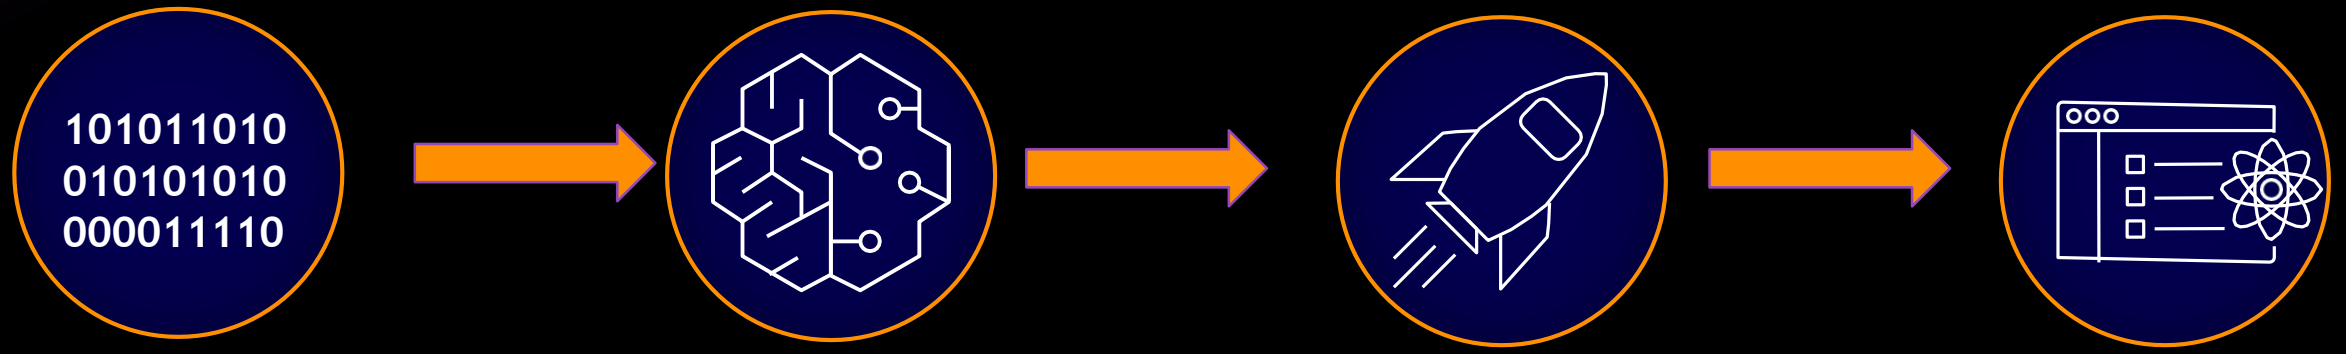

1. Bring data
2. Run Autopilot
3. Deploy model
4. Make predictions

Demo

Take-aways and best practices

- ML code is a small part of a machine learning product
- Explore and evaluate all your options properly before picking up the tools
- Explore the solutions available out there, that solve similar problems with the one you have (e.g. Amazon AI services) ; why re-invent the wheel when speed matters?
- Utilize prebuilt Amazon SageMaker open source framework containers; use requirements.txt or extend the containers to meet your requirements, before trying to build custom containers from the scratch
- Use local mode (on your local machine or an Amazon SageMaker notebook) to develop, test and debug your code, scale on more costly instances
- Other available model build environments such as fully managed Rstudio on Amazon SageMaker and Amazon SageMaker Canvas

Resources

1. https://aws.amazon.com/sagemaker/build/
2. https://aws.amazon.com/getting-started/hands-on/build-train-deploy-machine-learning-model-sagemaker/
3. https://aws.amazon.com/sagemaker/autopilot/
4. https://aws.amazon.com/blogs/machine-learning/bring-your-own-model-with-amazon-sagemaker-script-mode/
5. https://aws.amazon.com/blogs/machine-learning/bringing-your-own-custom-container-image-to-amazon-sagemaker-studio-notebooks/
6. https://docs.aws.amazon.com/sagemaker/latest/dg/studio-jumpstart.html
7. https://docs.aws.amazon.com/sagemaker/latest/dg/docker-containers.html
8. https://github.com/aws/deep-learning-containers
9. https://github.com/aws/deep-learning-containers/blob/master/available_images.md

Visit the AI & Machine Learning resource hub for more resources

Dive deeper into these resources, get inspired and learn how you can use AI and machine learning to accelerate your business outcomes.

- The machine learning journey e-book
- 7 leading machine learning use cases e-book
- A strategic playbook for data, analytics, and machine learning e-book
  Accelerate machine learning innovation with the right cloud services & infrastructure e-book
- Choosing the right compute infrastructure for machine learning e-book
- Improving service and reducing costs in contact centers e-book
- Why ML is essential in your fight against online fraud e-book
- … and more!

https://bit.ly/3mwi59V

Visit resource hub

AWS Machine Learning (ML) Training and Certification

AWS is how you build machine learning skills

Courses built on the curriculum leveraged by Amazon's own teams. Learn from the experts at AWS.

aws.training/machinelearning

Flexibility to learn your way

Learn online with on-demand digital courses or live with virtual instructor-led training, plus hands-on labs and opportunities for practical application.

explore.skillbuilder.aws/learn

Validate your expertise

Demonstrate expertise in building, training, tuning, and deploying machine learning models with an industry-recognized credential.

aws.amazon.com/certification

aws

Thank you for attending AWS Innovate – AI/ML Edition

We hope you found it interesting! A kind reminder to **complete the survey.**
Let us know what you thought of today's event and how we can improve the event experience for you in the future.


aws-apj-marketing@amazon.com


twitter.com/AWSCloud


facebook.com/AmazonWebServices


youtube.com/user/AmazonWebServices


slideshare.net/AmazonWebServices


twitch.tv/aws

Thank you!

Romina Sharifpour