



# Azure Machine Learning Competitive Analysis

Spring 2020

# Contents

Introduction [3](#)

Our Product [11](#)

Direct Competitors [24](#)

Indirect Competitor [46](#)

Influencers [58](#)

Summary [64](#)

# Introduction

# Executive Summary

This analysis compares our product against 5 competitors and influencers:

- Direct Competitors AWS Sagemaker and GCP AI Platform
- Indirect Competitor IBM Watson Machine Learning
- Influencers SAP Artificial Intelligence and Salesforce Einstein

It shows detailed comparison across 5 areas:

- Overview of services
- Pricing
- Documentation
- Testimonials
- Getting started

These are the top 5 recommendations:

- Provide a visual of the basic architecture of our product on home page
- Validate that feature areas highlighted on home page are the most relevant ones
- Showcase our commitment to cutting-edge new ML features in our product
- Introduce more links to product detail pages to allow further reading
- Ensure that users new to AML but not to Azure are suitably onboarded

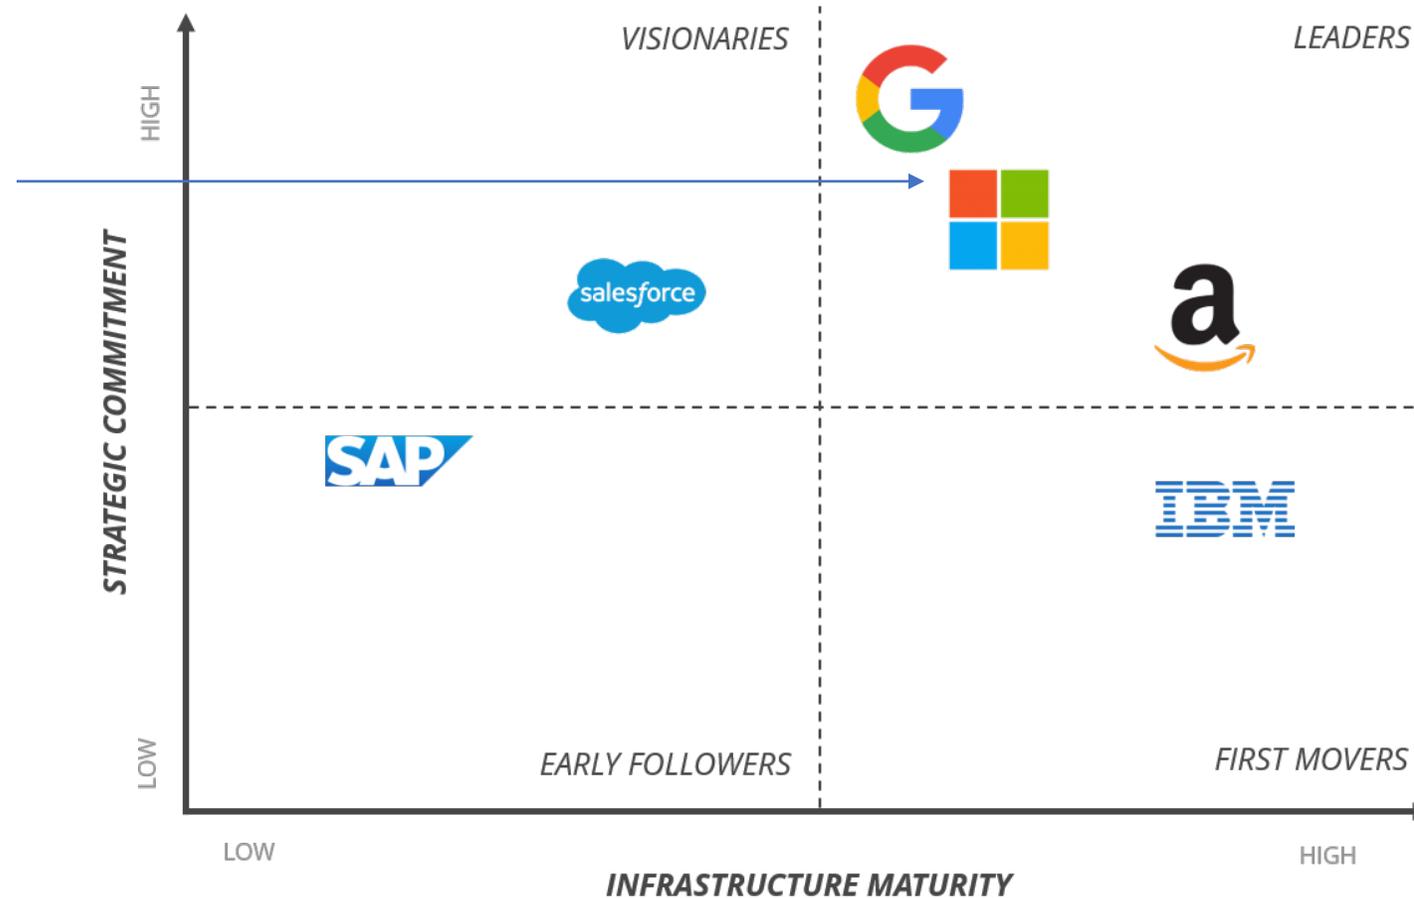
# Business Overview

- **Our business area has as its goal** to democratize artificial intelligence (AI), **lowering the barrier of entry for bringing machine learning (ML) capabilities to enterprise applications** without requiring advanced skills in AI or data science (DS)
- Cloud AI services are not strictly speaking required, as ML can be done with open source frameworks, but will likely run into issues scaling workloads because training real-world models typically requires large compute clusters
- AI services span a spectrum of general-purpose services with high flexibility on one end and special-purpose services with high ease-of-use on the other
- On the first end, and **our main focus, is what we're calling MLaaS (machine learning as a service), which are machine learning services for custom predictive analytics**. These are cloud-based platforms that cover most infrastructure issues such as data pre-processing, model training, and model evaluation, with further prediction. They allow for fast model training and deployment and handle the lifecycle of a machine learning model. Users don't have to provision and manage compute, storage, and networking environments to run jobs, or set up and configure DS environments for training, tuning, and hosting the model. The services do **expect data scientists to bring their own dataset and code so they can train a model against custom data**. They **require some ML experience**. Our product in this area is Azure Machine Learning
- On the other end and related to MLaaS, there are also machine learning APIs for services with pre-trained models that are plug and play, requiring users to just feed in their data. Examples include computer vision, natural language processing, and speech services. These don't require ML experience. Our products in this area are under Cognitive Services

MLaaS services are strongly linked to their associated cloud providers. Here is the overall landscape

We mitigated not being first to the cloud market by making strong commitments to develop and innovate

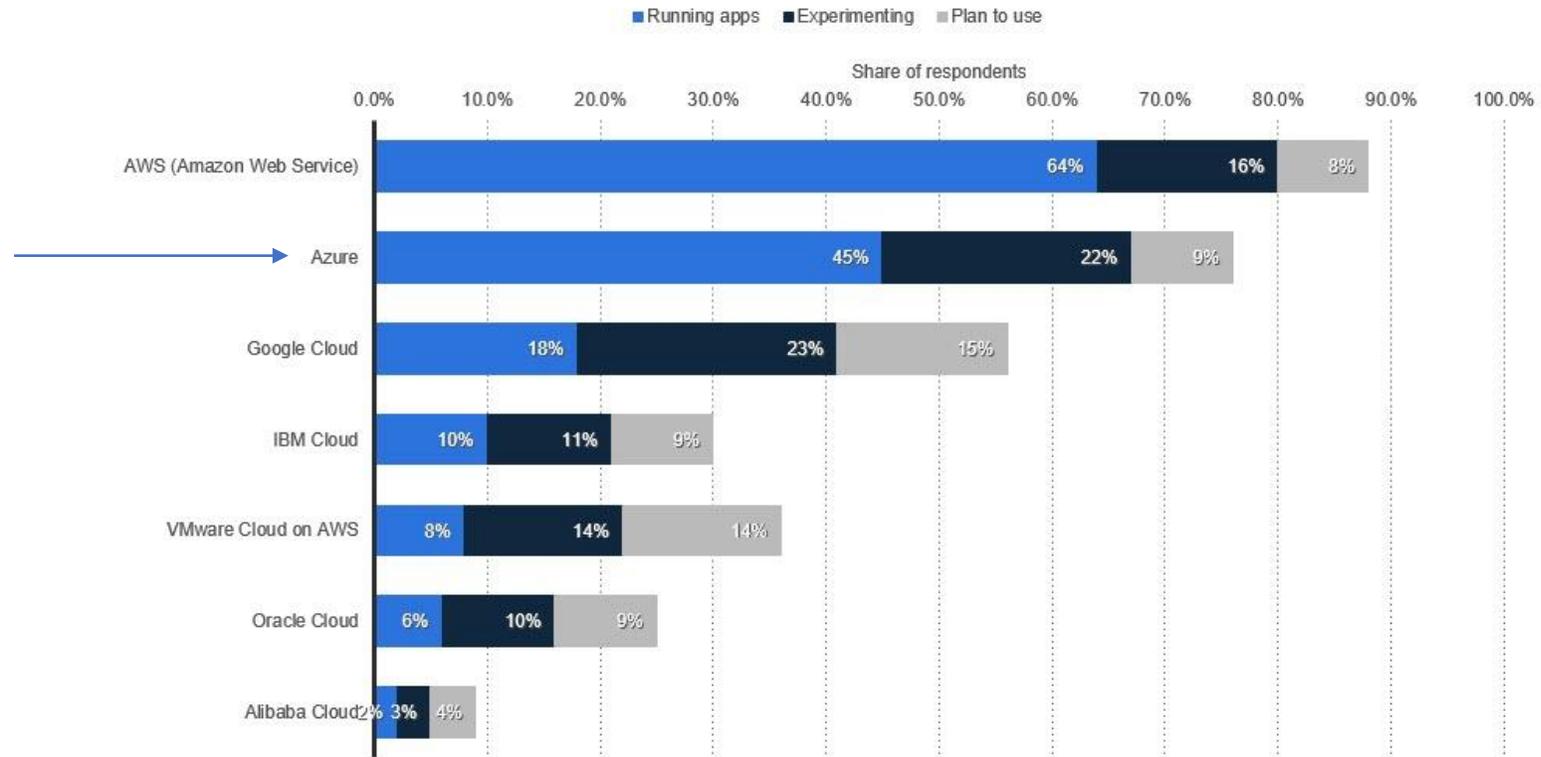
## AI LANDSCAPE: CLOUD PLATFORM VENDORS



Public cloud platform usage worldwide 2018

## Current and planned usage of public cloud platform services running applications worldwide in 2018

Though we're currently #2 in the cloud market, we grow significantly year on year



Note: Worldwide; January 2018; 997 Respondents; Technical executives, managers, and practitioners of cloud technologies

Further information regarding this statistic can be found on .

Source: RightScale;



# Method Introduction and Goals

**Competitive analysis is a UX research method that seeks to learn from competitors**, examining their best practices and bad user experiences, what types of customer segments use their products, and how current digital solutions address or fail to meet the needs of target customers, **with the end goal of devising a solution that creates a competitive advantage**

**The goal is to benchmark how competitors present and differentiate themselves in this market and use the insights to improve our own positioning**

# Competitors and Influencers

**Direct Competitors** offer a very similar value proposition to a similar customer base, current or future

- **AWS Sagemaker**
- **GCP AI Platform**

**Indirect Competitors** offer a related but different value proposition to a similar customer base

- **IBM Watson Machine Learning** offers a select bundle of ML services on various locales; IBM's MLaaS is split between this service and Watson Studio

**Influencers** offer a related value proposition to a related customer base

- **SAP Artificial Intelligence** offers a limited number of pre-made AI services
- **Salesforce Einstein** does the same; these are more like our Cognitive Services than AML

# Scope of Study

- **Comparison across 5 areas of public, pre-paywall content** on each company's website (except influencers, who are examined for select features). Focuses on opening matter, not detail pages
  - **Overview of services**
  - **Pricing**
  - **Documentation**
  - **Testimonials**
  - **Getting started**
- Not focused on feature-level offerings, compatibilities, or product functionalities
- **Comparison focuses on UX interest areas** around information architecture, content strategy, visual design, and interaction design
- **Assumes user is not committed to a solution or feature**, but is shopping for their company, consulting client, or personal project
- Focuses on full-service, flexible MLaaS whenever possible over single-serve, pre-made AI services

# Our Product

# Microsoft's Azure Machine Learning



Overview: Empower developers and data scientists with a wide range of productive experiences for building, training, and deploying machine learning models faster

<b>Founded:</b> 2014	<b>Market position:</b> #2
<b>Pricing:</b> pay as you go (2 tiers, basic with less functionality and enterprise with full functionality)	<b>Main segments:</b> autoML, pipelines, integrated notebooks, compute instance, SDK support, compute, data for ML, MLOps, labeling

# Overview of Services

Tabbed navigation shows contents of home page

+ Advance on scroll instead of click flows nicely

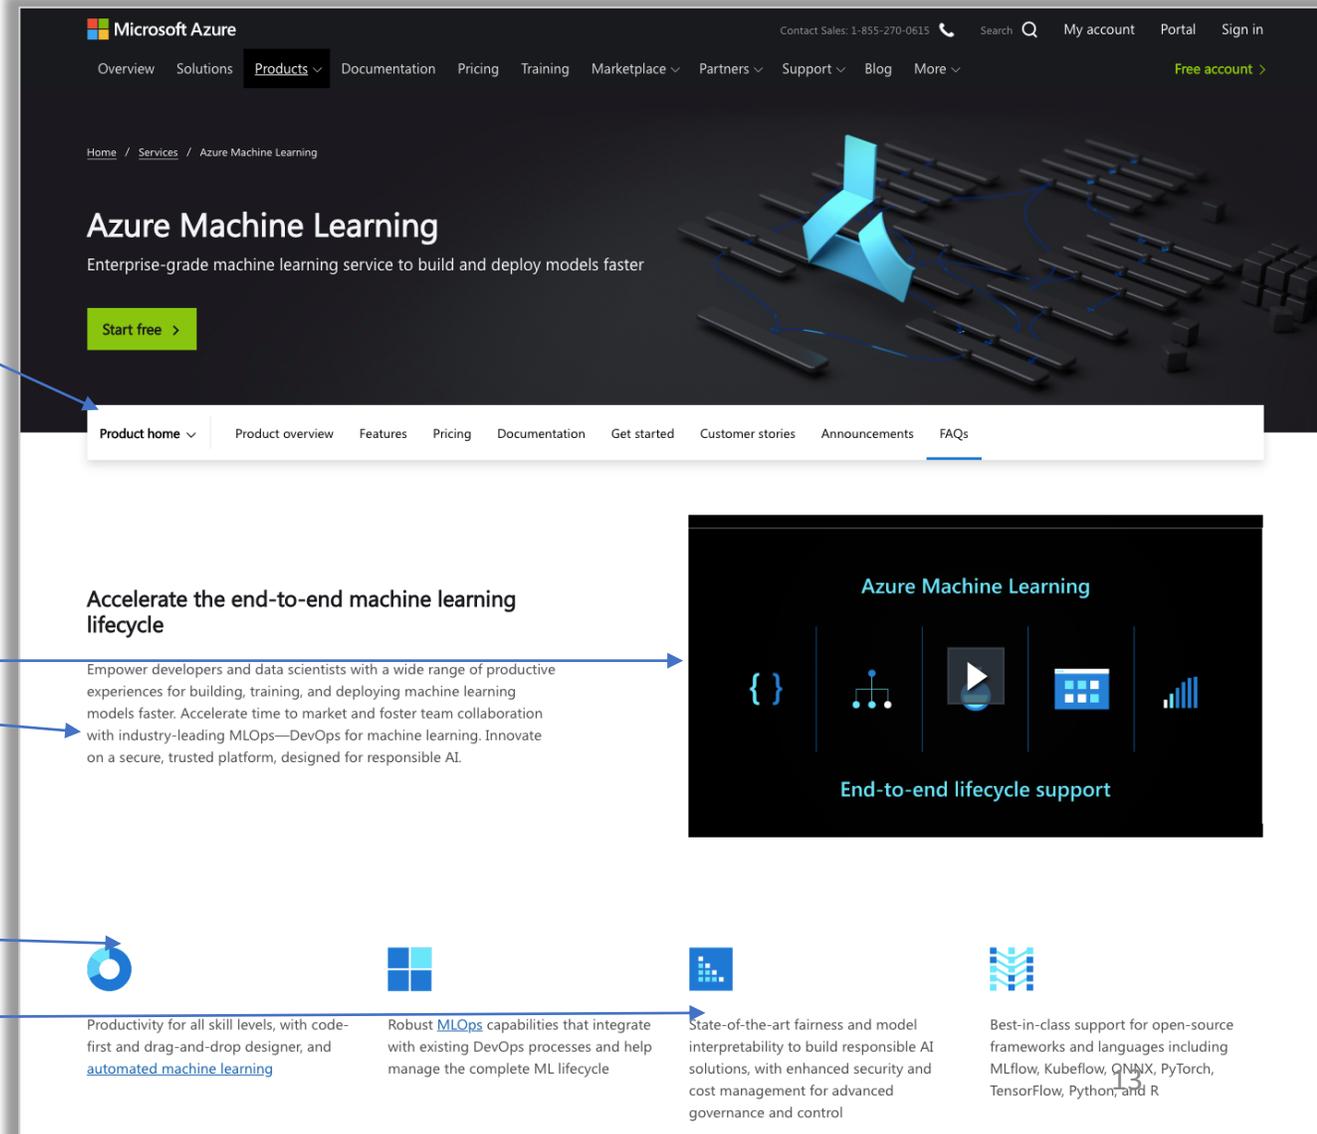
- It also prevents from surfacing much information

+ Succinct description of product benefits

+ Video introduction alternative

+ 4 key areas identified as selling points with short descriptions

- No links provided for latter 2 key areas



+ Videos are often positively perceived as interactive

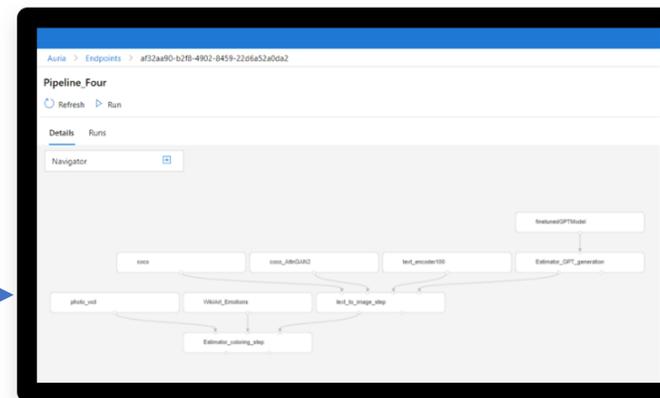
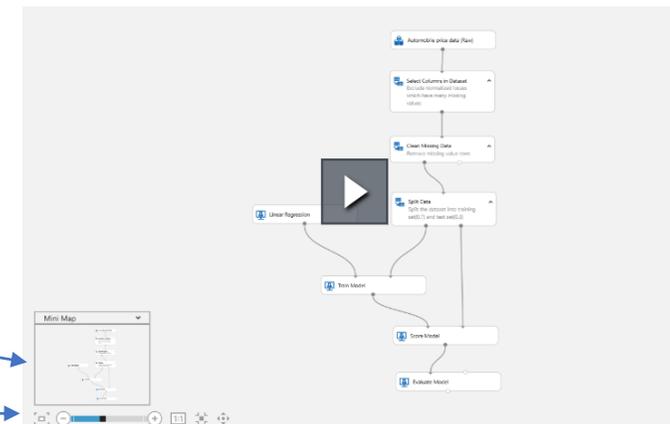
- Unclear why this is not presented in a frame like the mock-ups

- Unclear why the other sections are mock-ups rather than videos

- No link for more info

### Boost productivity and access ML for all skills

Rapidly build and deploy machine learning models using tools that meet your needs regardless of skill level. Use the no-code designer to get started, or use built-in Jupyter notebooks for a code-first experience. Accelerate model creation with the [automated machine learning](#) UI, and access built-in feature engineering, algorithm selection, and hyperparameter sweeping to develop highly accurate models.

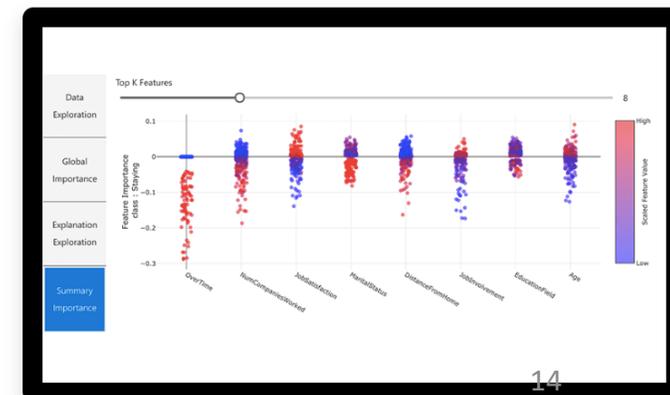


### Operationalize at scale with robust MLOps

[MLOps](#), or DevOps for machine learning, streamlines the machine learning lifecycle, from building models to deployment and management. Use ML pipelines to build repeatable workflows, and use a rich model registry to track your assets. Manage production workflows at scale using advanced alerts and machine learning automation capabilities. Profile, validate, and deploy machine learning models anywhere, from the cloud to the edge, to manage production ML workflows at scale in an enterprise-ready fashion.

### Build responsible AI solutions

Access state-of-the-art technology for fairness and machine learning model transparency. Use model interpretability for explanations about predictions to better understand model behavior. Reduce model bias by applying common fairness metrics, automatically making comparisons and using recommended mitigations.



+ Pointing out open-source compatibility is nice

- Uncertain why this is presented with a different, bulleted layout with check marks as opposed to paragraph form like above

- No link for more info

- This wasn't shown in the curation of 4 main areas above

- Unclear how these 5 areas were chosen

Some of these are for different audiences. 1, 3, 4 more data scientists, 2 and 5 more for ops/IT

The screenshot displays the Microsoft Azure Security Center interface. At the top, there are three columns: 'Development tools' (Visual Studio, Jupyter, VS Code, PyCharm), 'Languages' (Python, R), and 'Frameworks' (ONNX, Chainer, TensorFlow, PyTorch, Keras, mxnet, dmlc, XGBoost). Below this is a section titled 'Advanced security, governance, and control' which shows a list of recommendations with progress bars and details. The interface includes a search bar, navigation tabs for Overview, VMs and Computers, Cloud services, and App services, and a sidebar with various Azure services.

### Innovate on an open and flexible platform

Get built-in support for [open-source tools and frameworks](#) for machine learning model training and inferencing. Use familiar frameworks like PyTorch, TensorFlow, and scikit-learn, or the open and interoperable ONNX format. Choose the development tools that best meet your needs, including popular IDEs, Jupyter notebooks, and CLIs—or languages such as Python and R. Use ONNX Runtime to optimize and accelerate inferencing across cloud and edge devices.

### Advanced security, governance, and control

- ✓ Build machine learning models using the enterprise-grade security, compliance, and virtual network support of Azure.
- ✓ Protect your assets using built-in controls for identity, data, and network access, including custom roles.
- ✓ Restrict access to only your corporate network or apply Azure security policies.
- ✓ Manage governance and controls with audit trail, quota and cost management, and a comprehensive compliance portfolio.

# Pricing

This is all that's visible of pricing on the home page



Pay only for what you need, with no upfront cost  
For details, go to the [Azure Machine Learning pricing page](#).

After clicking on the link from home, users are brought to this pricing page

Visible CTA but not large

Blue branding is different from the black of the home page. Unclear why

Unclear why this is repeated from before. Redundant

+ Succinct summary of pricing benefits

- Long paragraph that hides the fact that Enterprise is in preview

The screenshot shows the Azure Machine Learning pricing page. The header is dark blue with the Microsoft Azure logo and navigation links. The main content area is also dark blue, featuring the title 'Azure Machine Learning pricing' and a sub-headline 'Enterprise-grade machine learning service to build and deploy models faster'. Below this are three green checkmarks with text: 'No upfront cost', 'No termination fees', and 'Pay only for what you use'. A prominent green button labeled 'Try for free >' is positioned below these benefits. The page also includes a 'Pricing Tiers' section with a long paragraph of text. Annotations with arrows point to various elements: the 'Pricing' navigation link, the 'Try for free' button, the three benefit checkmarks, the 'Pricing Tiers' section, and the long paragraph in the 'Pricing Tiers' section.

+ Nice table for comparing options



- Pricing seems like a place where users might want help, especially if they're interested in the enterprise preview. None is available

- Not sure if intent is to eventually upsell customers but there's no incitement to upgrade

### Azure Machine Learning editions

FEATURES	BASIC	ENTERPRISE
	For open source development at cloud scale with a code-first experience.	Basic + UI capabilities + secure and comprehensive machine learning lifecycle management for all skill levels.
<b>Automated machine learning</b>		
Create and run experiments in notebooks	✓	✓
Create and run experiments in studio web experience		✓
Industry leading forecasting capabilities		✓
Support for deep learning and other advanced learners		✓
Large data support (up to 100GB)		✓
Interpretability in UI		✓
<b>Machine Learning Pipelines</b>		
Create, run, and publish pipelines using the Azure ML SDK	✓	✓
Create pipeline endpoints using the Azure ML SDK	✓	✓
Create, edit, and delete scheduled runs of pipelines using the Azure ML SDK	✓	✓
Create and publish custom modules using the Azure ML SDK	✓	✓
View pipeline run details in studio	✓	✓
Create, run, visualize, and publish pipelines in Azure ML designer		✓
Create pipeline endpoints in Azure ML designer		✓
Create, edit, and delete scheduled runs of pipelines in Azure ML designer		✓
Create and publish custom modules in Azure ML designer		✓
<b>Integrated notebooks</b>		
Workspace notebook and file sharing	✓	✓
R and Python support	✓	✓
Notebook collaboration	✓	✓

— This is where the architecture of the product/feature list is visible for the first time. Odd that it's in pricing rather than the main page

— This breakdown makes me think we're in danger of shipping the org chart--presenting parts of the product in the way we internally organize, rather than how the user breaks it down

Compute instance		
Managed compute Instances for integrated Notebooks	✓	✓
Sharing of compute instances	✓	✓
Collaborative debugging of models	✓	✓
Jupyter, JupyterLab, Visual Studio Code	✓	✓
Virtual Network (VNet) support for deployment	✓	✓
SDK Support		
R and Python SDK support	✓	✓
Security		
Role Based Access Control (RBAC) support	✓	✓
Virtual Network (VNet) support for training	✓	✓
Virtual Network (VNet) support for inference	✓	✓
Scoring endpoint authentication	✓	✓
Compute		
Cross workspace capacity sharing and quotas		✓
Data for machine learning		
Create, view or edit datasets and datastores from the SDK	✓	✓
Create, view or edit datasets and datastores from the UI	✓	✓
View, edit, or delete dataset drift monitors from the SDK	✓	✓
View, edit, or delete dataset drift monitors from the UI		✓
MLOps		
Create ML pipelines in SDK	✓	✓
Batch inferencing	✓	✓
Model profiling	✓	✓
Interpretability in UI		✓
Labeling		
Labeling Project Management Portal	✓	✓
Labeler Portal	✓	✓
Labeling using private workforce	✓	✓

+ Nice intro to pricing tables

+ Nice filter option to cut down on results

+ Breakdown chart leads directly into VM pricing, which is a lot of text but easily scannable (the page continues, showing the huge range of VMs)

Could potentially incorporate another CTA in-between sections to break it up and remind users to start a trial

## Pricing details

The table below explains the ML surcharge for a broad category of VM's. For details please select the region and other information below to see all available VM's and associated pricing.

EDITION	CPU (GENERAL PURPOSE, COMPUTE OPTIMIZED, MEMORY OPTIMIZED, STORAGE OPTIMIZED)	GPU
Basic	No ML surcharge, VM pricing only for training and inferencing	No ML surcharge, VM pricing only for training and inferencing
Enterprise	\$0 per core hour ML surcharge	From \$0+ onwards per core ML surcharge

Edition:  Region:  Currency:  Display pricing by:

Category: [All](#) [General purpose](#) [Compute optimized](#) [Memory optimized](#) [Storage optimized](#) [GPU](#) [High performance compute](#) [FPGA](#)

### General purpose

For websites, small-to-medium databases, and other everyday applications

### Bs-series

INSTANCE	VCPU(S)	RAM	LINUX VM PRICE	MACHINE LEARNING SERVICE SURCHARGE	PAY AS YOU GO TOTAL PRICE	ONE YEAR RESERVED (% SAVINGS) TOTAL PRICE	THREE YEAR RESERVED (% SAVINGS) TOTAL PRICE
B1S	1	1 GiB	\$0.011/hour	\$0/hour	\$0.011/hour	\$0.007/hour (~42%)	\$0.004/hour (~62%)
B2S	2	4 GiB	\$0.042/hour	\$0/hour	\$0.042/hour	\$0.025/hour (~42%)	\$0.016/hour (~62%)
B1LS	1	0.5 GiB	\$0.006/hour	\$0/hour	\$0.006/hour	\$0.004/hour (~41%)	\$0.002/hour (~62%)
B1MS	1	2 GiB	\$0.021/hour	\$0/hour	\$0.021/hour	\$0.014/hour (~35%)	\$0.009/hour (~58%)
B2MS	2	8 GiB	\$0.084/hour	\$0/hour	\$0.084/hour	\$0.049/hour (~41%)	\$0.032/hour (~62%)
B4MS	4	16 GiB	\$0.166/hour	\$0/hour	\$0.166/hour	\$0.098/hour (~41%)	\$0.063/hour (~62%)
B8MS	8	32 GiB	\$0.333/hour	\$0/hour	\$0.333/hour	\$0.195/hour (~42%)	\$0.126/hour (~62%)
B12MS	12	48 GiB	\$0.499/hour	\$0/hour	\$0.499/hour	\$0.293/hour (~41%)	\$0.188/hour (~62%)
B16MS	16	64 GiB	\$0.666/hour	\$0/hour	\$0.666/hour	\$0.390/hour (~42%)	\$0.251/hour (~62%)
B20MS	20	80 GiB	\$0.832/hour	\$0/hour	\$0.832/hour	\$0.487/hour (~41%)	\$0.314/hour (~62%)

# Documentation

- 3-panel design. Not immediately obvious that each part is clickable, and no auto-advance. Could be more effective as a carousel

- Shift in tone here from imperative “do this...” to “you can...”

- This paragraph in particular is less coherent and polished than others. Not all elements are strongly related

- Plain-looking documentation link section

+ Nice breakdown of beginner vs. advanced

+ Nice that there are videos

- Unclear what “featured” means

## How to use Azure Machine Learning

Go to your studio web experience      Build and train      Deploy and manage

**Welcome!**

Create new ▾

- Automated ML**  
Automatically train and tune a model using a target metric.  
[Start now](#)
- Designer**  
Drag-and-drop interface from prepping data to deploying models.  
[Start now](#)
- Notebooks**  
Code with Python SDK and run sample experiments.  
[Start now](#)

**My recent resources**

Run Number	Experiment	Status Updated Time	Status
1	Sample_1_-_Regression...	9/27/2019, 1:38:37 PM	Completed
1474	category-based-prope...	9/18/2019, 4:37:10 PM	Completed
1475	category-based-prope...	9/18/2019, 3:49:21 PM	Completed

**Resources**

- Beginner tutorials**
  - Try a Jupyter notebook with Python
  - Drag-and-drop experiments
  - Use the automated ML UI
  - Configure your dev environment
- Advanced tutorials**
  - Predict taxi fares through automated ML
  - Classify images with scikit-learn
  - Batch score with Azure ML pipelines
- Featured videos**
  - Get started with Azure Machine Learning
  - Use automated machine learning to build models
  - Build zero-code models with Azure Machine Learning designer
  - MLOps for managing the end-to-end lifecycle
  - Incorporating ONNX Runtime into your models
  - Model interpretability and transparency
  - Building models with R

You can author new models and store your compute targets, models, deployments, metrics, and run histories in the cloud.

# Testimonials

+ This element is a carousel, which works nicely as a showcase

+ Link to read more does a good job of highlighting key content and making the rest available elsewhere

+ Good use of quote  
- This one needs a proofread

Customers using Azure Machine Learning

Walgreens Boots Alliance Schneider Electric bp OSOS Fermilab Borrowell

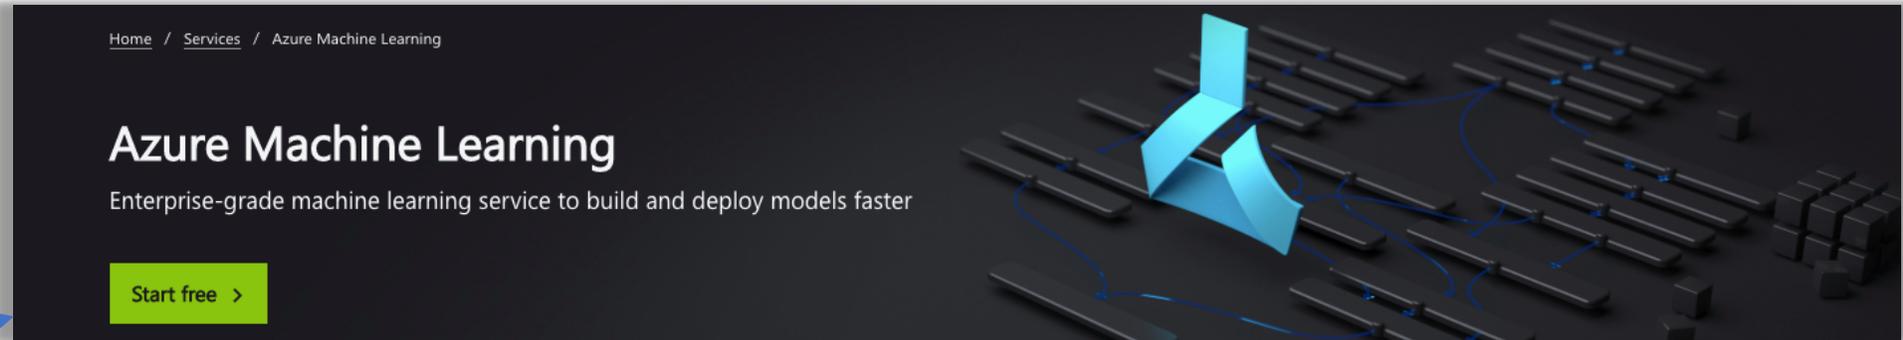
"If I have 200 models to train—I can just do this all at once. It can be farmed out to a huge compute cluster, and it can be done in minutes. So I'm not waiting for days."

Dean Riddlesden, Senior Data Scientist, Global Analytics, Walgreens Boots Alliance

[Read the story >](#)

# Getting Started

- + On home, 3 different reminders to get started, which is a good reinforcement of the message



Top banner of page

### Start using Azure Machine Learning today

- 

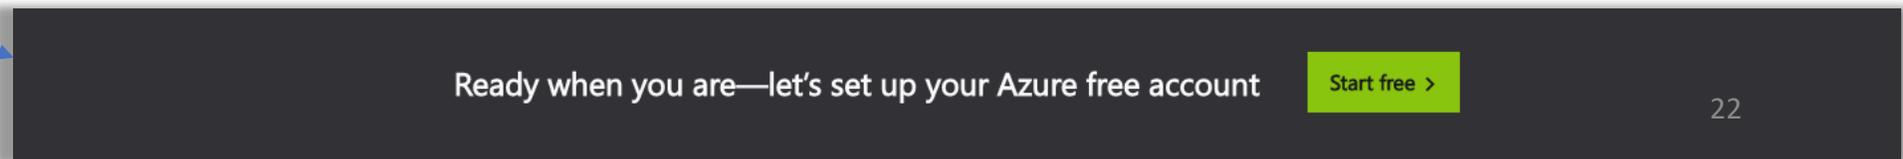
Get instant access and a \$200 credit by signing up for an [Azure free account](#).
- 

Sign in to the [Azure portal](#).
- 

Explore the [documentation and tutorials](#).

Middle of page

Bottom of page



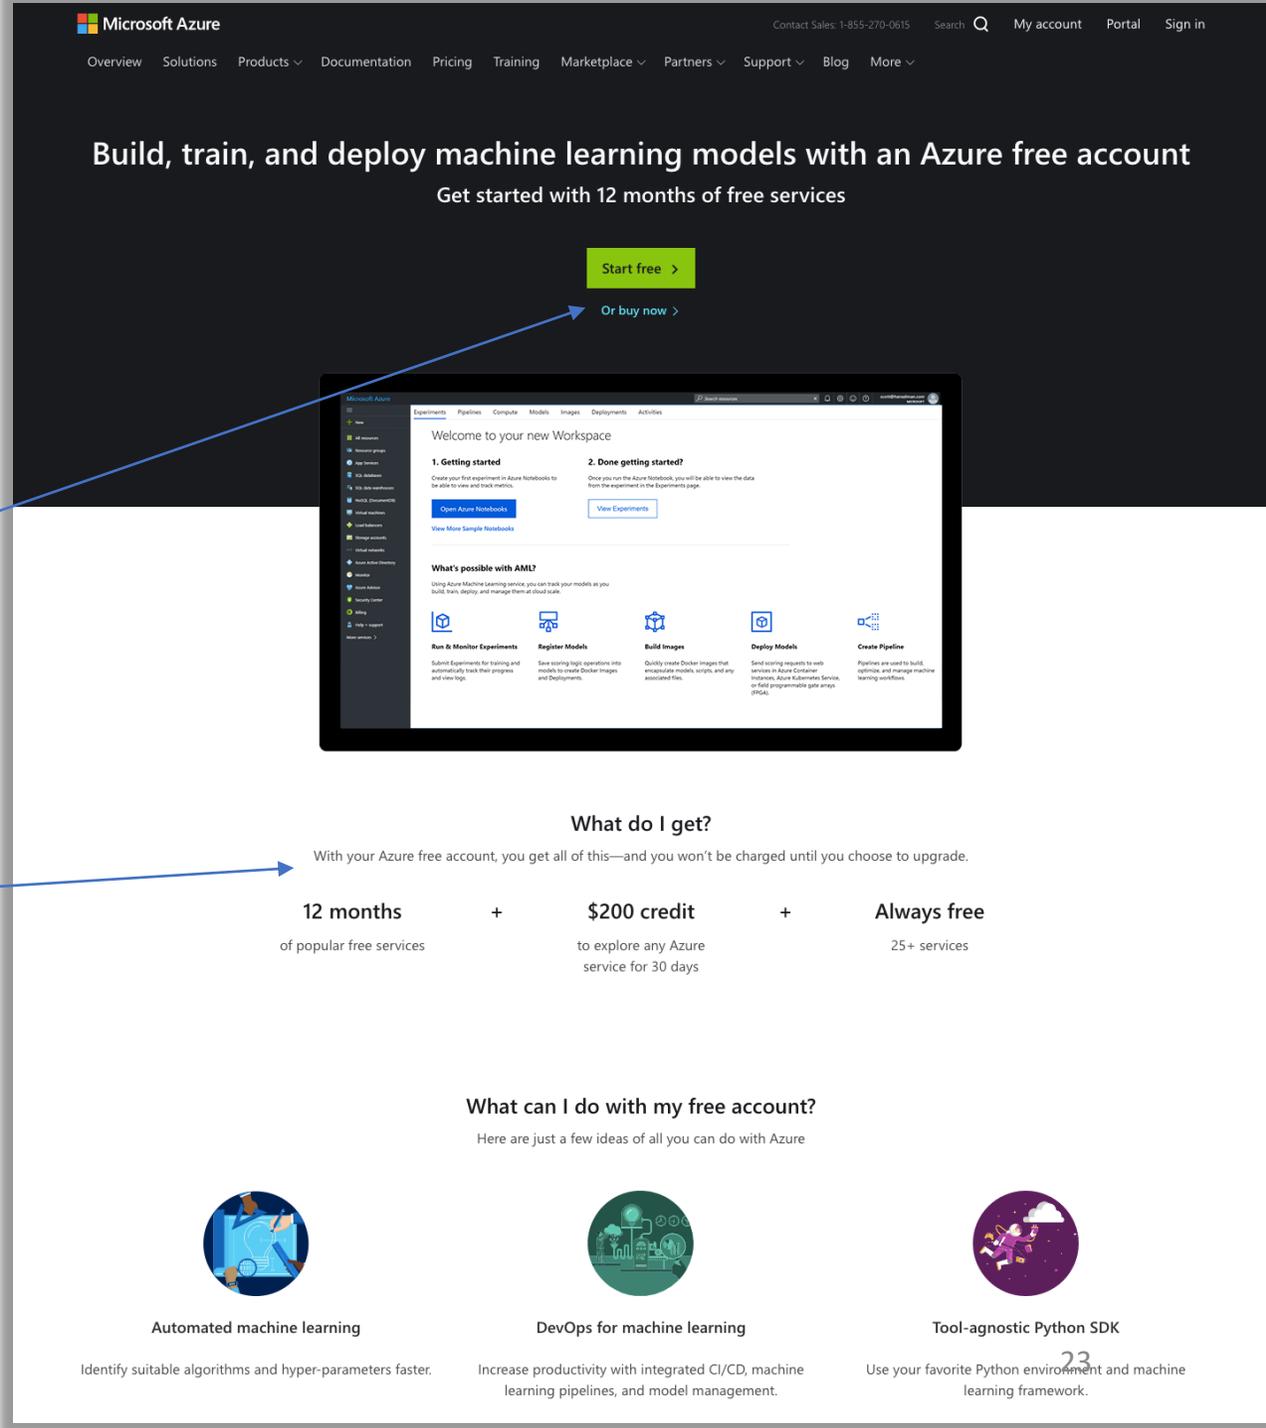
Clicking “start free” goes to this page

Users might expect that clicking from the previous page takes them directly to sign-up. This is an extra click

Option to buy now was not surfaced on the pricing page (or anywhere else)

Good info and structure on this page

Not an optimal flow for users of Azure who are new to AML. Hopefully the page would dynamically adapt



# Direct Competitors

# Amazon's AWS Sagemaker



Amazon SageMaker

Overview: Fully managed service that provides every developer and data scientist with the ability to build, train, and deploy machine learning models quickly

<b>Founded:</b> 2015	<b>Market position:</b> #1
<b>Pricing:</b> pay as you go	<b>Main segments:</b> label (1 service), build (2 services), train & tune (3 services), deploy & manage (3 services), 1 service across build--train & tune, 1 service across build--train & tune--deploy & manage

# Overview of Services

Tabbed navigation on click offers a lot of information hidden away

+ Appealing language of democratization

+ Events are intriguing

Nice description but users may not care for the longer read

+ Pitting SageMaker against traditional ML is a smart move

+ Graphic of services with descriptions broken down by stages is highly informative. Impressive sense-making and presentation

The screenshot shows the Amazon SageMaker website. At the top, there is a navigation bar with the AWS logo, a search bar, and links for 'Contact Sales', 'Support', 'English', 'My Account', and 'Create an AWS Account'. Below this is a secondary navigation bar with 'Amazon SageMaker' and 'Overview' (highlighted), along with dropdown menus for 'Features', 'Pricing', 'FAQs', 'Developer Resources', and 'Customers'. The main header features the 'Amazon SageMaker' logo and the tagline 'Machine learning for every developer and data scientist', with a 'Create an AWS Account' button. A 'FEATURED EVENT' section highlights the 'AWS JPL Open Source Rover Challenge' with a 'Register now' button. The main content area contains a paragraph describing SageMaker as a fully managed service that simplifies the machine learning process. Below this is a graphic titled 'Amazon SageMaker Studio' showing a workflow from 'Label' to 'Build' to 'Train & Tune' to 'Deploy & Manage'. The 'Label' stage includes 'Amazon SageMaker Ground Truth'. The 'Build' stage includes 'Amazon SageMaker Notebooks' and 'AWS Marketplace'. The 'Train & Tune' stage includes 'Amazon SageMaker Studio', 'Amazon SageMaker Autopilot', 'Amazon SageMaker Experiments', 'Amazon SageMaker Debugger', and 'Automatic Model Tuning'. The 'Deploy & Manage' stage includes 'Amazon SageMaker Model Monitor', 'Amazon SageMaker Neo', and 'Amazon Augmented AI'. A page number '26' is visible in the bottom right corner.

## Build machine learning models

ONLY ON AMAZON SAGEMAKER

### Improve productivity using Amazon SageMaker Studio, the first fully integrated development environment (IDE) for machine learning

Amazon SageMaker Studio provides a single, web-based visual interface where you can perform all ML development steps. SageMaker Studio gives you complete access, control, and visibility into each step required to build, train, and deploy models. You can quickly upload data, create new notebooks, train and tune models, move back and forth between steps to adjust experiments, compare results, and deploy models to production all in one place, making you much more productive. All ML development activities including notebooks, experiment management, automatic model creation, debugging, and model drift detection can be performed within the unified SageMaker Studio visual interface.



Use an IDE for ML development. For example, make updates to models inside a notebook and see how changes impact model quality using a side-by-side view of your notebook and training experiments.

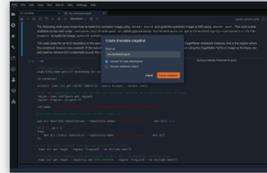
[Click to enlarge](#)

ONLY ON AMAZON SAGEMAKER

### Build and collaborate faster using Amazon SageMaker Notebooks

Managing compute instances to view, run, or share a notebook is tedious. Now available in preview, Amazon SageMaker Notebooks provide one-click Jupyter notebooks that you can start working with in seconds. The underlying compute resources are fully elastic, so you can easily dial up or down the available resources and the changes take place automatically in the background without interrupting your work. SageMaker also enables one-click sharing of notebooks. All code dependencies are automatically captured, so you can easily collaborate with others. They'll get the exact same notebook, saved in the same place.

You can choose from dozens of pre-built notebooks within SageMaker for different use cases. You can also get hundreds of algorithms and pre-trained models available in [AWS Marketplace](#) making it easy to get started quickly.



Generate a shareable link without manually tracking dependencies, to reproduce the notebook code.

[Click to enlarge](#)

ONLY ON AMAZON SAGEMAKER

### Automatically build, train, and tune models with full visibility and control, using Amazon SageMaker Autopilot

Amazon SageMaker Autopilot is the industry's first automated machine learning capability that gives you complete control and visibility into your ML models. Typical approaches to automated machine learning do not give you the insights into the data used in creating the model or the logic that went into creating the model. As a result, even if the model is mediocre, there is no way to evolve it. Also, you don't have the flexibility to make trade-offs such as sacrificing some accuracy for lower latency predictions since typical automated ML solutions provide only one model to choose from.

SageMaker Autopilot automatically inspects raw data, applies feature processors, picks the best set of algorithms, trains and tunes multiple models, tracks their performance, and then ranks the models based on performance, all with just a few clicks. The result is the best performing model that you can deploy at a fraction of the time normally required to train the model. You get full visibility into how the model was created and what's in it and SageMaker Autopilot integrates with Amazon SageMaker Studio. You can explore up to 50 different models generated by SageMaker Autopilot inside SageMaker Studio so its easy to pick the best model for your use case. SageMaker Autopilot can be used by people without machine learning experience to easily produce a model or it can be used by experienced developers to quickly develop a baseline model on which teams can further iterate.

[Learn more](#)



Automatically create machine learning models and pick the one that best suits your use case. For example, review the leaderboard to see how each option performs and pick the model that meets your model accuracy and latency requirements.

[Click to enlarge](#)

ONLY ON AMAZON SAGEMAKER

### Reduce data labeling costs by up to 70% using Amazon SageMaker Ground Truth

Successful machine learning models are built on the shoulders of large volumes of high-quality training data. But, the process to create the training data necessary to build these models is often expensive, complicated, and time-consuming. Amazon SageMaker Ground Truth helps you build and manage highly accurate training datasets quickly. Ground Truth offers easy access to labelers through Amazon Mechanical Turk and provides them with pre-built workflows and interfaces for common labeling tasks. You can also use your own labelers or use vendors recommended by Amazon through AWS Marketplace. Additionally, Ground Truth continuously learns from labels done by humans to make high quality, automatic annotations to significantly lower labeling costs.

[Learn more](#)

### Amazon SageMaker supports the leading deep learning frameworks

Supported frameworks include TensorFlow, PyTorch, Apache MXNet, Chainer, Keras, Gluon, Horovod, Scikit-learn, and Deep Graph Library.



## Train machine learning models

ONLY ON AMAZON SAGEMAKER

### Organize, track, and evaluate training runs using Amazon SageMaker Experiments

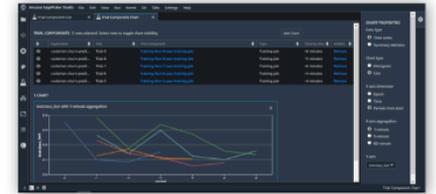
Amazon SageMaker Experiments helps you organize and track iterations to machine learning models. Training an ML model typically entails many iterations to isolate and measure the impact of changing data sets, algorithm versions, and model parameters. You produce hundreds of artifacts such as models, training data, platform configurations, parameter settings, and training metrics during these iterations. Often cumbersome mechanisms like spreadsheets are used to track these experiments.

SageMaker Experiments helps you manage iterations by automatically capturing the input parameters, configurations, and results, and storing them as 'experiments'. You can work within the visual interface of SageMaker Studio, where you can browse active experiments, search for previous experiments by their characteristics, review previous experiments with their results, and compare experiment results visually.

ONLY ON AMAZON SAGEMAKER

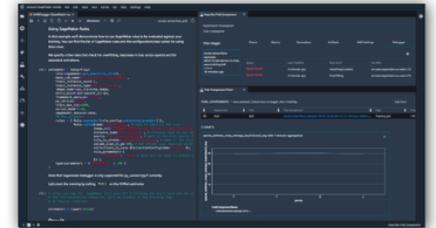
### Analyze, detect, and alert problems for machine learning using Amazon SageMaker Debugger

The ML training process is largely opaque and the time it takes to train a model can be long and difficult to optimize. As a result, it is often difficult to interpret and explain models. Amazon SageMaker Debugger makes the training process more transparent by automatically capturing real-time metrics during training such as training and validation, confusion



Track thousands of training experiments to understand the accuracy of your model. For example, view in a chart of how different time series datasets impact model accuracy.

[Click to enlarge](#)



Analyze and debug anomalies. For example, training a neural network will cease if gradients are determined to be vanishing. SageMaker Debugger identifies vanishing gradients so you can remediate before training is impacted.

[Click to enlarge](#)

Clever repackaging of controversial services (MTurk as backbone of SageMaker Ground Truth labeling)



SageMaker Studio for easy understanding. SageMaker can detect common training problems. With SageMaker Debugger, vanishing gradients are detected. With SageMaker Experiments, representing an early step towards model explainability.

### AWS is the best place to run TensorFlow

AWS' TensorFlow optimizations provide near-linear scaling efficiency across hundreds of GPUs to operate at cloud scale without a lot of processing overhead to train more accurate, more sophisticated models in much less time.

90%

SCALING EFFICIENCY WITH 256 GPUs

### Lower training costs by 90%

Amazon SageMaker provides Managed Spot Training to help you to reduce training costs by up to 90%. This capability uses Amazon EC2 Spot instances, which is spare AWS compute capacity. Training jobs are automatically run when compute capacity becomes available and are made resilient to interruptions caused by changes in capacity, allowing you to save cost when you have flexibility with when to run training jobs.

90%

COST REDUCTION WITH MANAGED SPOT TRAINING

+ Clear breakdown into 3 main steps (see previous page)

+ Change in background color across 3 stages helps break up page

Gives a huge amount of detail on features

Still easily scannable due to titles

+ Dark theme evokes developer-friendly associations

+ Exclusivity of “only on SageMaker” is an effective marker of uniqueness

+ Nice ability to magnify

Audience is focused on data scientists

+ Effective callout of what value add is

## Deploy machine learning models

### One-click deployment

Amazon SageMaker makes it easy to deploy your trained model into production with a single click so that you can start generating predictions for real-time or batch data. You can one-click deploy your model onto auto-scaling Amazon ML instances across multiple availability zones for high redundancy. Just specify the type of instance, and the maximum and minimum number desired, and SageMaker takes care of the rest. SageMaker will launch the instances, deploy your model, and set up the secure HTTPS endpoint for your application. Your application simply needs to include an API call to this endpoint to achieve low latency, high throughput inference. This architecture allows you to integrate your new models into your application in minutes because model changes no longer require application code changes.

ONLY ON AMAZON SAGEMAKER

### Keep models accurate over time using Amazon SageMaker Model Monitor

Amazon SageMaker Model Monitor allows developers to detect and remediate concept drift. Today, one of the big factors that can affect the accuracy of deployed models is if the data being used to generate predictions differs from data used to train the model. For example, changing economic conditions could drive new interest rates affecting home purchasing predictions. This is called concept drift, whereby the patterns the model uses to make predictions no longer apply. SageMaker Model Monitor automatically detects concept drift in deployed models and provides detailed alerts that help identify the source of the problem. All models trained in SageMaker automatically emit key metrics that can be collected and viewed in SageMaker Studio. From inside SageMaker Studio you can configure data to be collected, how to view it, and when to receive alerts.



Monitor models in production. For example, view charts with important model features and summary statistics, watch them over time and compare with the features used in training. Some features drift when the model is running in production, which can indicate the need to retrain your model.

Click to enlarge

ONLY ON AMAZON SAGEMAKER

### Validate predictions through human review

Many machine learning applications require humans to review low confidence predictions to ensure the results are correct. But, building human review into the workflow can be time consuming and expensive involving complex processes. Amazon Augmented AI is a service that makes it easy to build the workflows required for human review of ML predictions. Augmented AI provides built-in human review workflows for common machine learning use cases. You can also create your own workflows for models built on Amazon SageMaker. With Augmented AI, you can allow human reviewers to step in when a model is unable to make high confidence predictions.

[Learn more »](#)

### Lower machine learning inference costs by up to 75% using Amazon Elastic Inference

In most deep learning applications, making predictions using a trained model - a process called inference - can be a major factor in the compute costs of the application. A full GPU instance may be over-sized for model inference. In addition, it can be difficult to optimize the GPU, CPU, and memory needs of your deep-learning application. Amazon Elastic Inference solves these problems by allowing you to attach just the right amount of GPU-powered inference acceleration to any Amazon EC2 or Amazon SageMaker instance type or Amazon ECS task with no code changes. With Elastic Inference, you can choose the instance type that is best suited to the overall CPU and memory needs of your application, and then separately configure the amount of inference acceleration that you need to use resources efficiently and to reduce the cost of running inference.

**75%**  
LOWER INFERENCE COSTS

### Integrate with Kubernetes for orchestration and management

Kubernetes is an open source system used to automate the deployment, scaling, and management of containerized applications. Many customers want to use the fully managed capabilities of Amazon SageMaker for machine learning, but also want platform and infrastructure teams to continue using Kubernetes for orchestration and managing pipelines. SageMaker lets users train and deploy models in SageMaker using Kubernetes operators and pipelines. Kubernetes users can access all of SageMaker's capabilities natively from KubeFlow.

# Pricing

- Very text-heavy
- No CTA to incite purchasing or trying
- Overall fairly opaque explanation of pricing
- Unclear what happens after 2-month free trial, and what the tiers are
- No intro to VM pricing tables. A new user might need guidance

aws

Contact Sales Support English My Account Create an AWS Account

Products Solutions Pricing Documentation Learn Partner Network AWS Marketplace Customer Enablement Events Explore More Q

Amazon SageMaker Overview Features Pricing FAQs Developer Resources Customers

Products / Machine Learning / Amazon SageMaker / ...

## Amazon SageMaker Pricing

With Amazon SageMaker, you pay only for what you use. Building, training, and deploying ML models is billed by the second, with no minimum fees and no upfront commitments. Pricing within Amazon SageMaker is broken down by on-demand ML instances, ML storage, and fees for data processing in hosting instances.

### Try Amazon SageMaker for two months, free!

As part of the [AWS Free Tier](#), you can get started with Amazon SageMaker for free. If you have never used Amazon SageMaker before, for the first two months, you are offered a monthly free tier of 250 hours of t2.medium or t3.medium notebook usage for building your models, plus 50 hours of m4.xlarge or m5.xlarge for training, plus 125 hours of m4.xlarge or m5.xlarge for deploying your machine learning models for real-time inferencing and batch transform with Amazon SageMaker. Your free tier starts from the first month when you create your first SageMaker resource.

### Included with Amazon SageMaker Training and Hosting

When you train your models in Amazon SageMaker and enable Amazon SageMaker Debugger, you can use built-in rules for debugging or write your own custom rules, or both. SageMaker Debugger provides a fully managed experience for running both built-in and custom rules as Amazon SageMaker Processing jobs. **For built-in rules, there is no charge and Amazon SageMaker Debugger will automatically select an instance type.** For custom rules, you will need to choose an instance (e.g. ml.m5.xlarge) and you will be charged for the duration for which the instance is in use for the Amazon SageMaker Processing job.

When you deploy your models as Amazon SageMaker endpoints for real-time inference and enable Amazon SageMaker Model Monitor, you can use built-in rules to monitor your models or write your own custom rules, or both. Model Monitor provides a fully managed experience for running both built-in and custom rules as Amazon SageMaker Processing jobs. **For built-in rules with ml.m5.xlarge instance, you get up to 30 hours of monitoring aggregated across all endpoints each month, at no charge.** Additional usage beyond 30 hours or usage for other ML instance types will be charged for the duration for which the instance is in use at the Amazon SageMaker Processing on demand rate.

### Choice of Amazon EC2 On-demand and Spot Instances

With Amazon SageMaker you have the choice of choosing from Amazon EC2 On-Demand instances or Amazon EC2 Spot instances. For building, training, and deploying your models on Amazon SageMaker, on-demand ML instances let you pay for machine learning compute capacity by the second, with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware, and transforms what are commonly large fixed costs into much smaller variable costs. Pricing is per instance-hour consumed for each instance, from the time an instance is available for use until it is terminated or stopped. Each partial instance-hour consumed will be billed per-second.

For training your ML models, you have the choice of using Amazon EC2 Spot instances with Managed Spot Training. This option can help reduce the cost of training your machine learning models by up to 90%. Once a Managed Spot Training job completes, you can calculate the cost savings as the percentage difference between the duration for which the training job ran and the duration for which you were billed. The cost savings is also visible on the AWS management console.

### Inference Acceleration

When building and deploying your model, you have the option to attach fractional GPU compute capacity to your Amazon SageMaker endpoint using Amazon Elastic Inference in select AWS Regions. If you choose to add an Amazon Elastic Inference accelerator, you will be billed for the accelerator hours. For more details on inference acceleration with Amazon SageMaker, see the [Amazon Elastic Inference website](#). For regional availability of Amazon Elastic Inference, see the [Region table](#).

### ML General Purpose storage

For model training, Amazon SageMaker provides you with the ability to select up to 6 TB of associated General Purpose (SSD) storage capacity for your training data. For notebook, model training, and model hosting, General Purpose (SSD) storage capacity is also added for temporary data storage. With General Purpose (SSD), you will be charged for this storage. However, you will not be charged for the I/Os consumed.

### Data processed

When hosting your models, data processed by Amazon SageMaker is pulled into and out of model hosting instances.

PAGE CONTENT

- US East (N. Virginia)
- US East (Ohio)
- US West (Oregon)
- US West (N. California)
- AWS GovCloud (US)
- Canada Central (Montreal)
- EU (Ireland)
- EU (Frankfurt)

US East (N. Virginia)

Building	+
Processing	+
Model Training	+
Model Deployment	+

# Documentation

+ The benefit of showing this on a separate page is clear, as the full depth of resourcing is on display

+ Good breakdown that reflects categories previously established

+ Nice intro video

+ Visual design is different on this page, which calls attention to its different functionality, but background of video ties back to previous design elements

aws

Contact Sales Support English My Account Create an AWS Account

Products Solutions Pricing Documentation Learn Partner Network AWS Marketplace Customer Enablement Events Explore More

Amazon SageMaker Overview Features Pricing FAQs Developer Resources Customers

Products / Machine Learning / Amazon SageMaker / ...

## Amazon SageMaker developer resources

PAGE CONTENT

- Getting started
- Build ML models
- Train and tune ML models
- Deploy ML models
- Additional resources

### Getting started

Amazon SageMaker is a fully managed, modular service that helps developers and data scientists to build, train, and deploy machine learning models at scale. Get started with these developer resources, so you can quickly move from concept to production.

#### Introduction to Amazon SageMaker

Learn about the build, train, and deploy modules of Amazon SageMaker.

Amazon SageMaker (1:03)

DEVELOPER GUIDE

[Getting started guide](#)

Follow this step-by-step guide to start using Amazon SageMaker quickly.

TUTORIAL

[Amazon SageMaker tutorial](#)

Learn how to get started with Amazon SageMaker in 10 minutes.

TRAINING COURSE

[Amazon SageMaker: Simplifying Machine Learning Application Development](#)

In this training course, learn how to use Amazon SageMaker to simplify the integration of machine learning into your applications. Key topics include: an overview of machine learning and problems it can help solve, using a Jupyter Notebook to train a model based on Amazon SageMaker's built-in algorithms and, using Amazon SageMaker to publish the validated model. You will finish the course by building a serverless application that integrates with the Amazon SageMaker published endpoint.

TRAINING COURSE

[Build an Object Detection model using images labeled with Amazon SageMaker Ground Truth](#)

In this training course, you will learn how to implement a machine learning pipeline using Amazon SageMaker and Amazon SageMaker Ground Truth. First you will create a labeled dataset, then you'll create a training job to train your object detection model, and finally you will use Amazon SageMaker to create and update your model.

30

## Build machine learning models

Use the Build module of SageMaker to collect and prepare training data, access pre-built notebooks, and leverage the built-in, high performance algorithms.

### DEVELOPER GUIDE

#### A step-by-step guide to building ML models

Learn to build an ML model with the steps and resources outlined in this guide.

### VIDEO

#### Fully-managed notebook instances

In this video, learn all about the fully-managed notebook instances with Amazon SageMaker.



Dive deep into fully-managed notebook instances (16:44)

### BLOG

#### Use Amazon SageMaker notebook instances on Github for a wide range of use cases and machine learning workflows

Read this blog and learn to use common workflows using Amazon SageMaker notebook instances.

### HANDS-ON LAB

#### Amazon SageMaker sample notebooks

Access a rich repository of SageMaker notebooks, on GitHub.

### HANDS-ON LAB

#### Leverage the built-in, high-performance algorithms on Amazon SageMaker

Utilize algorithms built into Amazon SageMaker that are faster and cheaper than popular alternatives.

### VIDEO

#### Address a wide range of use cases with multiple built-in ML algorithms

In this video, learn about the high-performance algorithms, built-in with Amazon SageMaker.



Leverage high-performance built-in machine learning algorithms (15:37)

## Train and tune machine learning models

Use the Train module to set up training environments with one click and optimize your model using automatic module tuning

### DEVELOPER GUIDE

#### Train a machine learning model with Amazon SageMaker

Read an overview of how to train machine learning models using Amazon SageMaker.

### VIDEO

#### Train and tune ML models

In this video, learn how to train and tune your machine learning models to the highest accuracy with Amazon SageMaker.



Train and tune ML models with Amazon SageMaker (18:29)

### BLOG

#### Track training runs of your ML models with Amazon SageMaker

Identify your best ML models for your use case and get to production faster. Track, search, filter and sort your machine learning training runs using the steps outlined in this blog. You can now get to the best ML model across all your experiments using key model attributes, such as hyperparameter values and accuracy metrics, with Amazon SageMaker.

### HANDS-ON LAB

#### Hyperparameter tuning with Amazon SageMaker

Try these examples of using hyperparameter tuning across different algorithms and deep learning frameworks.

### BLOG

#### Amazon SageMaker automatic model tuning produces better models, faster

Learn how to automatically tune the hyperparameter values of the algorithm in your machine learning model to obtain the most accurate predictions.

### WEBINAR

#### Train machine learning models using TensorFlow on Amazon SageMaker

In this on-demand tech talk, learn to train TensorFlow-based machine learning models. Understand the unique combination of TensorFlow and Amazon SageMaker to accelerate training of your machine learning models and bring them to production.



Appealing card layout with descriptions of what each guide or course accomplishes



## Deploy machine learning models

Use the Deploy module to deploy your machine learning models to production with a single click.

### DEVELOPER GUIDE

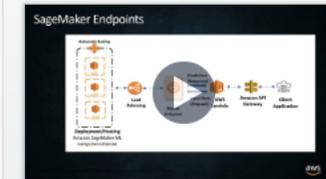
#### Deploy machine learning models with Amazon SageMaker

Follow the step-by-step guide to deploy machine learning models on the highest performing infrastructure.

### VIDEO

#### Move your ML models from experimentation to production

In this video, learn how to deploy your ML models to production on the most scalable infrastructure.



Deploy ML models from experimentation to production (7:52)

### HANDS-ON LAB

#### Automate Amazon SageMaker custom models

Follow the examples on GitHub to use Amazon SageMaker and AWS Step Functions to automate the building, training, and deploying of custom machine learning models.

### BLOG

#### Deploy trained Keras or TensorFlow models using Amazon SageMaker

Learn to use the deployment capabilities of SageMaker including A/B testing and Auto Scaling, delivering high performance and high availability for your machine learning models.

### WEBINAR

#### Integrate Amazon SageMaker into your enterprise

In this on-demand tech talk, learn about the machine learning life cycle, best practices for using Amazon SageMaker in your enterprise, and how to integrate Amazon SageMaker with other AWS services.

### BLOG

#### Build, train, and deploy fast.ai models with Amazon SageMaker

In this blog, learn how you can build, train, and deploy fast.ai models into Amazon SageMaker training and hosting by using the Amazon SageMaker Python SDK and a PyTorch base image. You can avoid the extra steps of building your own container.

# Testimonials

- Too minimal. Seeing only the companies (and forcing users to click in) doesn't reveal anything about how they used SageMaker or why a new user should care



## Case studies



Keep farmland healthy and optimize crop yield



Accelerate drug development



Identify fraud and improve the security of financial transactions



Increase efficiency of manufacturing operations and optimize the supply chain



Keep truckers happy and reduce carbon emissions



Add intelligence to medical devices

# Getting Started

— Not an optimal flow for users of AWS who are new to SageMaker. Hopefully the page would dynamically adapt

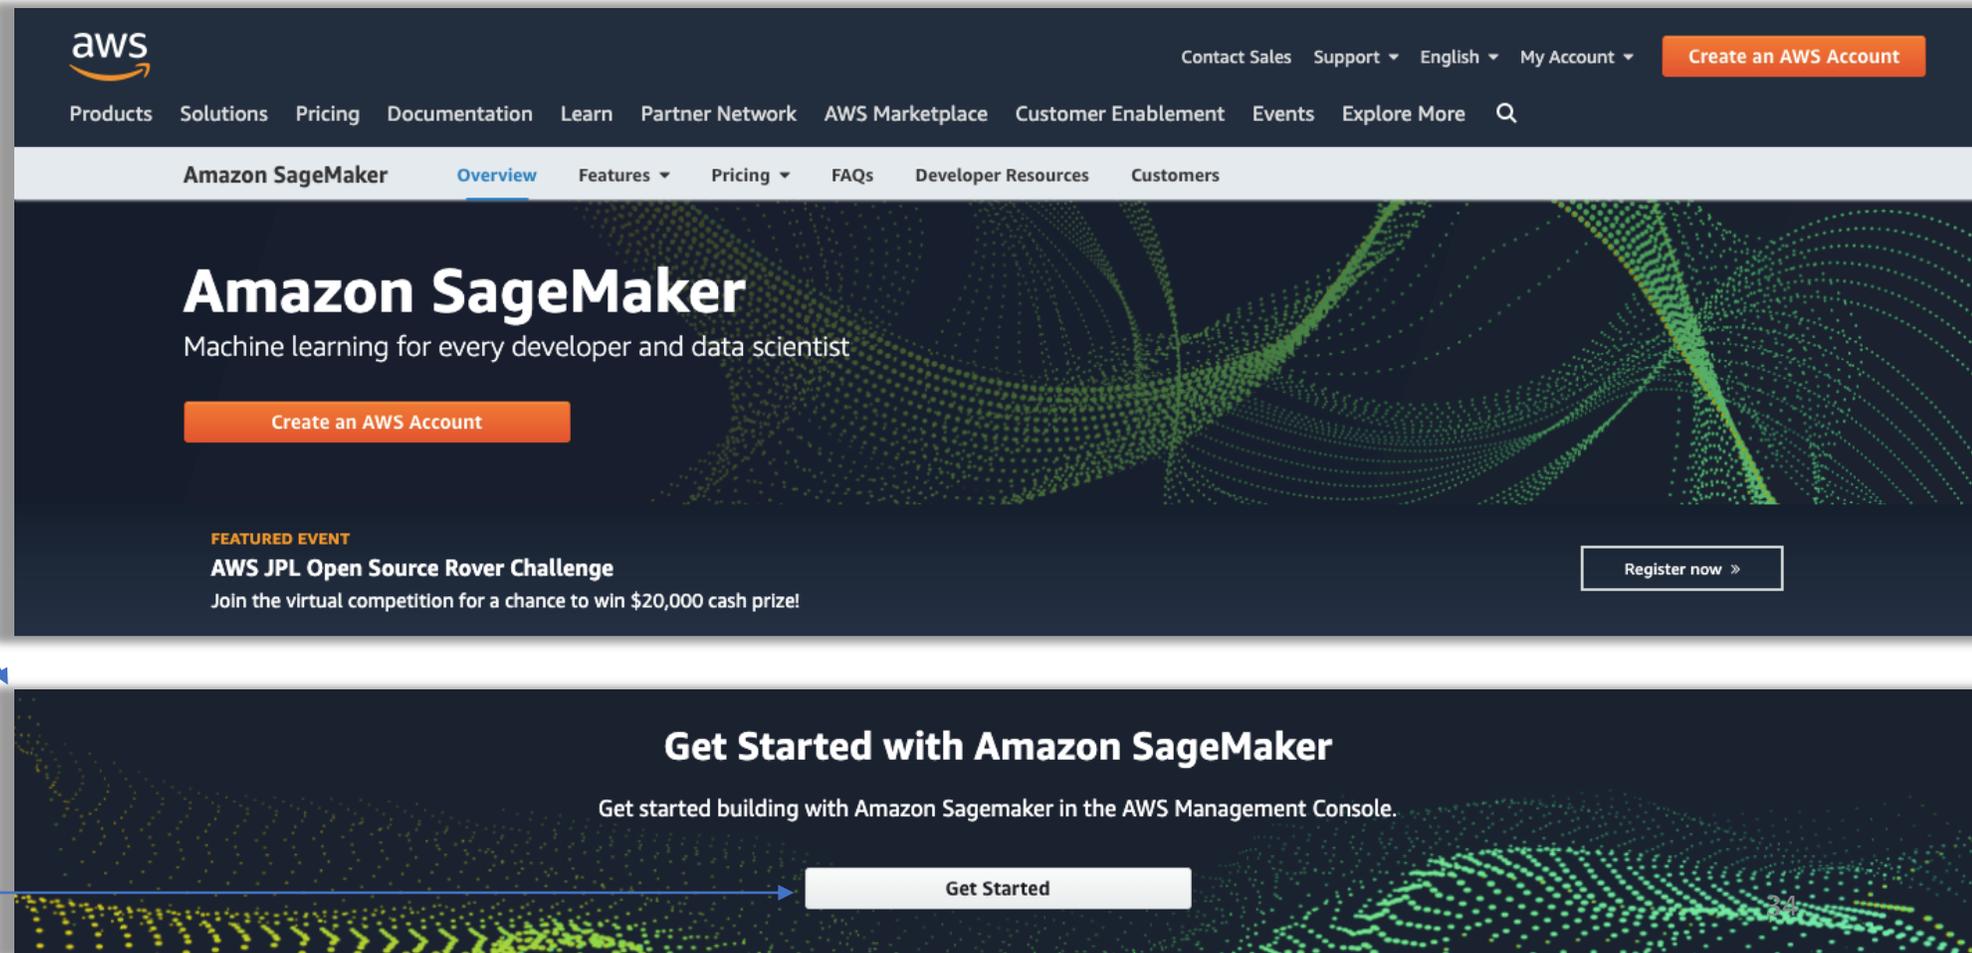
2 incitements to sign up is likely enough to hook interested users

Getting started banner at top of page →

+ Nice cohesive design across the two CTAs

Second one on bottom of page

— Unclear why bottom CTA isn't orange to call attention to itself and match AWS branding →



+ Clicking from the previous page brings the user to sign-up. Nice that it's direct

## AWS Accounts Include 12 Months of Free Tier Access

Including use of Amazon EC2, Amazon S3, and Amazon DynamoDB

Visit [aws.amazon.com/free](https://aws.amazon.com/free) for full offer terms

+ Nice that a link for more info is offered

- Concerning that this isn't covered on the pricing page

+ Option to sign in is good for users not new to AWS

### Create an AWS account

Email address

Password

Confirm password

AWS account name ⓘ

Continue

[Sign in to an existing AWS account](#)

© 2020 Amazon Web Services, Inc. or its affiliates.  
All rights reserved.

[Privacy Policy](#) | [Terms of Use](#)

# Google's GCP AI Platform



Overview: Code-based data science development environment empowers machine learning developers, data scientists, and data engineers to take their projects from ideation to deployment, quickly and cost-effectively

<b>Founded:</b> 2016	<b>Market position:</b> #3
<b>Pricing:</b> n/a (in beta)	<b>Main segments:</b> prepare (2 services), build & run (5 services), manage (3 services), share (2 services)

# Overview of Services

Layout is a single page of info. Tab navigation is for GCP in general. Extremely streamlined presentation

+ Pointing out on-prem compatibility is nice

+ Nice minimalist design and graphic

Leaning on other Google technology like TensorFlow is good for building credibility but might call into question if non-Google tech is supported

Google Cloud Why Google Solutions **Products** Pricing Getting Started

AI & Machine Learning Products [Contact Sales](#)

## AI Platform

Create your AI applications once, then run them easily on both GCP and on-premises.

[Contact sales](#) [Try it free](#)

### Take your machine learning projects to production

AI Platform makes it easy for machine learning developers, data scientists, and data engineers to take their ML projects from ideation to production and deployment, quickly and cost-effectively. From data engineering to “no lock-in” flexibility, AI Platform’s integrated tool chain helps you build and run your own machine learning applications.

AI Platform supports Kubeflow, Google’s open-source platform, which lets you build portable ML pipelines that you can run on-premises or on Google Cloud without significant code changes. And you’ll have access to cutting-edge Google AI technology like TensorFlow, TPUs, and TFX tools as you deploy your AI applications to production.



Effective breakdown of 4 main categories with products



“Related” products sounds more ancillary than it probably is



Writing here falls flat and should be more dynamic. It doesn't sell the product. More telling than showing

Also mostly targeted for data scientists/ML engineers

The screenshot shows the Google Cloud AI Platform documentation page, divided into four main sections: Prepare, Build and run, Manage, and Share. Each section includes a brief description and a list of related products and services. A blue arrow points from the 'Related products and services' list in the 'Build and run' section to the 'Related products and services' list in the 'Manage' section.

**Prepare**

You can use Cloud Storage or [BigQuery](#) to store your data. Then use the built-in data labeling service to label your training data by applying classification, object detection, and entity extraction, etc., for images, videos, audio, and text. You can also import the labeled data to AutoML and train a model directly.

Related products and services:  
[BigQuery](#)  
[Data Labeling Service](#)

**Build and run**

You can build your ML applications on GCP with a managed Jupyter Notebook service that provides fully configured environments for different ML frameworks using [Deep Learning VM Image](#). Then you can use AI Platform Training and Prediction services to train your models and deploy them to production on GCP in a serverless environment, or do so on-premises using the training and prediction microservices provided by [Kubeflow](#).

Related products and services:  
[Deep Learning VM Image](#)  
[AI Platform Notebooks](#)  
[AI Platform Training](#)  
[AI Platform Prediction](#)  
[Kubeflow](#)

**Manage**

You can manage your models, experiments, and end-to-end workflows using the AI Platform interface within the GCP console, or do so on-premises using [Kubeflow Pipelines](#). AI Platform offers advanced tooling to help you understand your model results and explain them to business users.

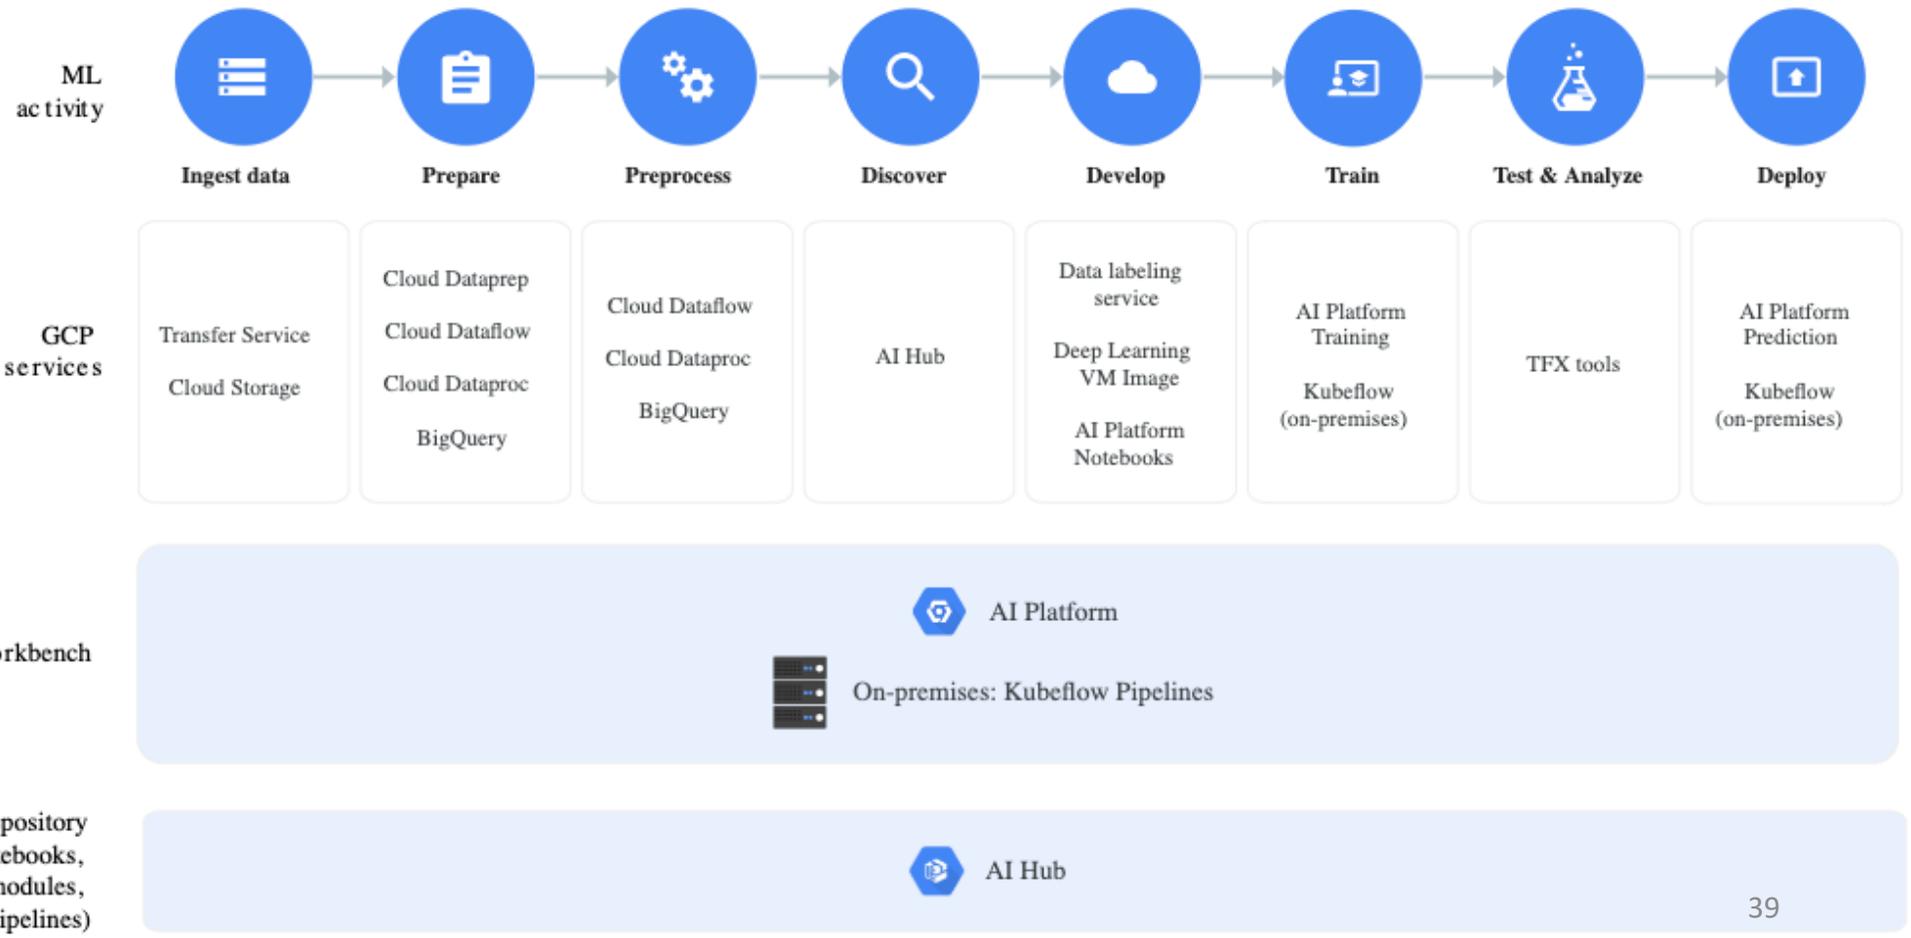
Related products and services:  
[AI Platform docs](#)  
[Kubeflow Pipelines](#)  
[Explainable AI](#)

**Share**

You can discover ML pipelines, notebooks, and other AI content via [AI Hub](#) and leverage [Kubeflow Pipelines](#) to build reusable end-to-end ML pipelines that you can share with other users and deploy on GCP or on-premises.

Related products and services:  
[AI Hub](#)  
[Kubeflow Pipelines](#)

# Machine learning development: the end-to-end cycle



+ Great graphic that ties user activities with GCP services. Recalls similar AWS graphic

- Visual presentation is not as minimal as other items on the page (and font appears different)

# Pricing

—  
No CTA to incite purchasing or trying

Very minimal pricing info on the home page

## Pricing

Kubeflow, AI Hub, and notebooks can be used for no charge. You can learn about the pricing of our managed services like AI Platform Training, AI Platform Predictions, Compute Engine, Google Kubernetes Engine, BigQuery, and Cloud Storage [here](#). You can also use our [pricing calculator](#) to estimate the costs of running your workloads.

— To the right is the pricing calculator, which is uncharacteristically maximalist for a Google design

+ It does offer help on hover

— Clicking on the first link leads to a loop back to the home page, possibly because the product is in beta with selected groups prior to GA

The screenshot shows the Google Cloud Platform Pricing Calculator interface. At the top, there are navigation links for Google Cloud, Why Google, Solutions, Products, Pricing, and Getting Started. Below this is the 'Google Cloud Platform Pricing Calculator' header. A row of icons represents various services: Compute Engine, App Engine, Kubernetes Engine, Cloud Run, Cloud Storage, Networking, BigQuery, and BigQuery ML. A search bar is provided for finding products. The main section is titled 'Instances' and contains several configuration options, each with a help icon (question mark):

- Number of instances \*
- What are these instances for?
- Operating System / Software: Free: Debian, CentOS, CoreOS, Ubuntu, or other User Provided OS
- Machine Class: Regular
- Machine Family: General purpose
- Series: N1
- Machine type: f1-micro (vCPUs: shared, RAM: 0.60 GB)
- Add GPUs. (checkbox)
- Local SSD: 0
- Datacenter location: Iowa (us-central1)
- Instances using ephemeral public IP
- Instances using static public IP
- Committed usage: None
- Average hours per day each server is running \*: 24 hours per day
- Average days per week each server is running \*: 7
- Include Always Free usage in my estimate. (checkbox)

At the bottom right, there is a link to 'Learn More about Always Free' and an 'ADD TO ESTIMATE' button.

# Documentation

Minimal display of resources that hide full-service documentation

+ Card-based visual design is clean

- Orange icons differ in branding from other items on the page, and it's unclear why

- Vague categories. What is "and more?" How is content from AI hub different from the other content?

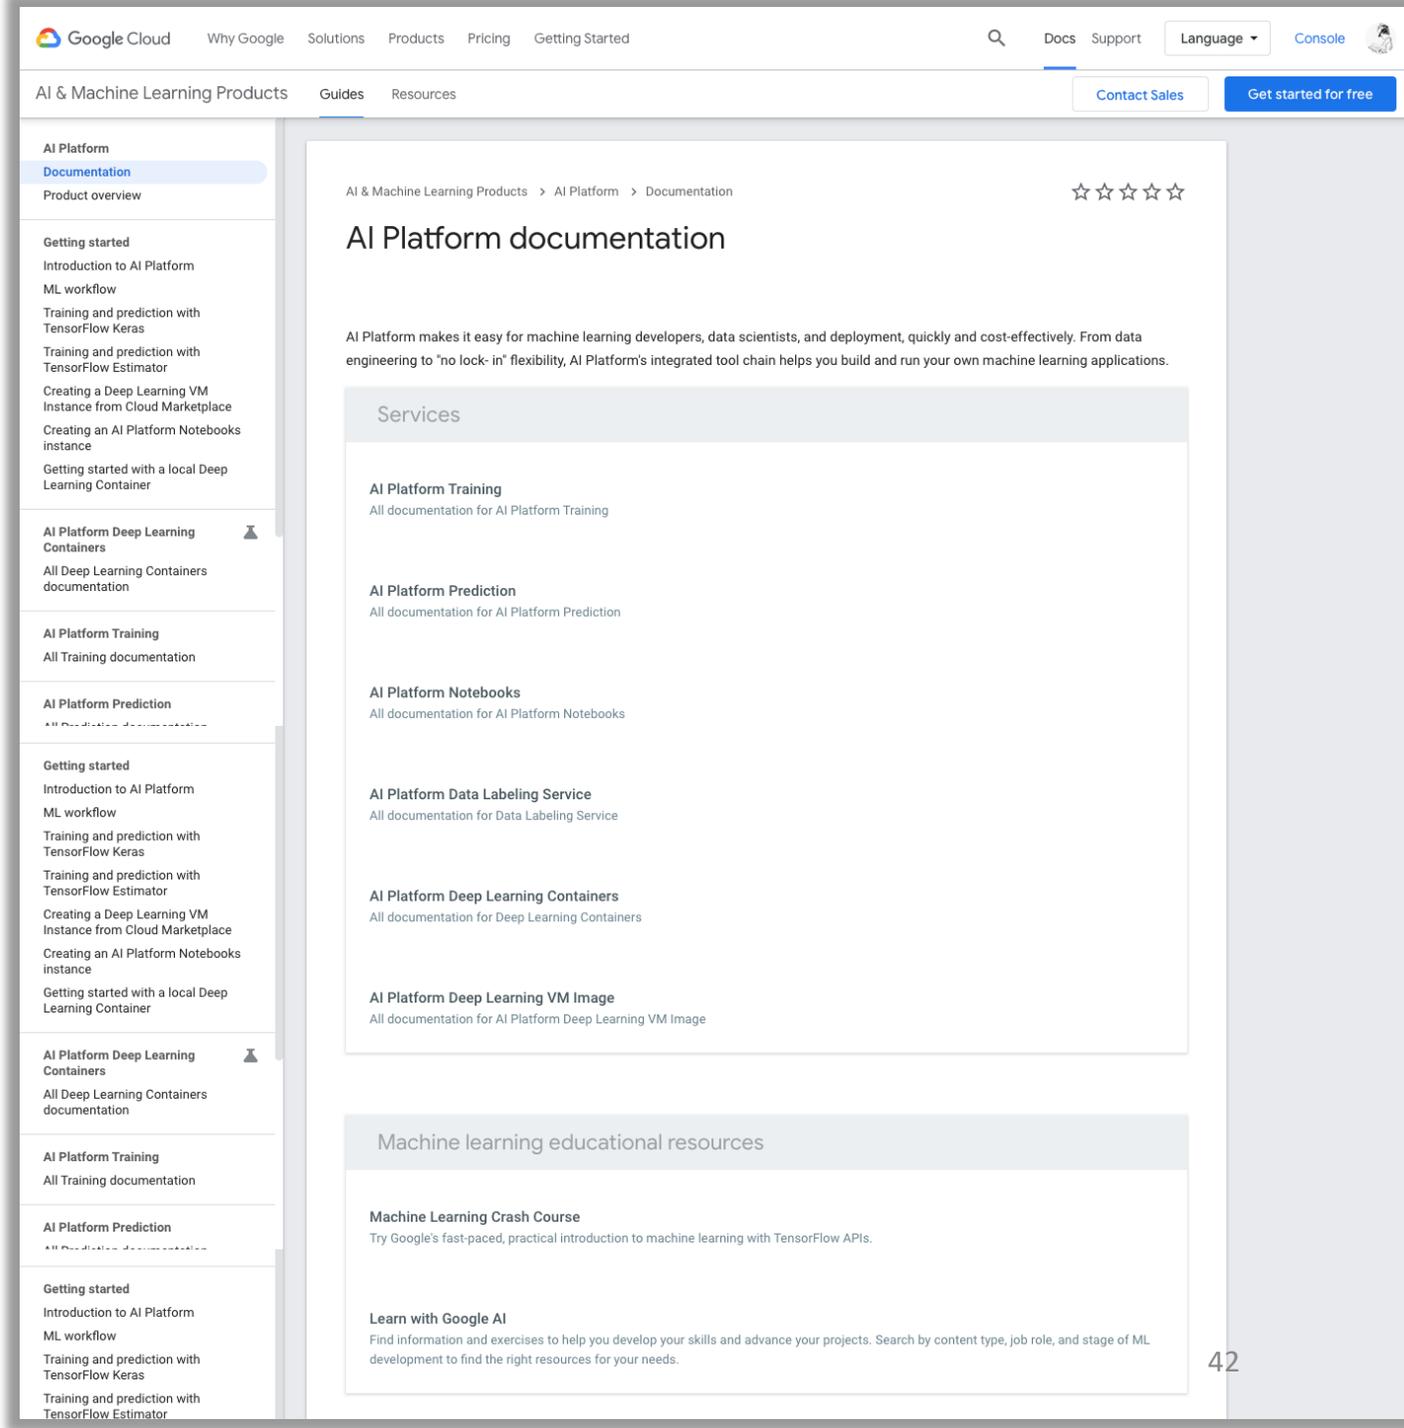
**Resources**

-  Discover content from AI Hub  
[Learn more →](#)
-  Start building with our fully configured notebooks  
[Get started →](#)
-  Try tutorials, quickstarts, and more.  
[Get started →](#)

Clicking on “tutorials, quickstarts, and more” goes to this page

Text-heavy and maximalist, with little hierarchy

Not immediately obvious where to start



# Testimonials

+ Nice minimal carousel of user quotes

- Doesn't seem to be more info about this client or a case study

- Unclear what partners means even after reading the description

+ The idea is intriguing

Not shown, but there's a see all partners that gives more info on a separate page

The screenshot shows a testimonial carousel with a light blue background. The main text is a quote from Lucas Ngoo, co-founder, CTO, of Carousell. The quote is enclosed in a light blue box with a white border. To the right of the quote is the Carousell logo, which consists of a red square with a white camera icon and the word "carousell" in lowercase. Below the quote is the name "Lucas Ngoo, co-founder, CTO, Carousell". At the bottom of the carousel is a "See all customers" link. Below the carousel is a section titled "Partners" with a subtitle: "Google Cloud Machine Learning Partners come with deep AI expertise and can help you incorporate ML for a wide range of use cases across every stage of model development and serving." At the bottom of the page are logos for Intel, Cisco, Pluto7, Atos, SpringML, and NVIDIA.

“ In retail, it’s important to provide customers with easy access to alternative products or recommended add-ons. We train our own machine learning models with TensorFlow on AI Platform, and we automate the periodic retraining of these models with Kubeflow Pipelines. Together with AI Hub, useful for sharing models between data scientists, we can now iterate faster on our models, and automatically deploy them to staging and production. ”

Lucas Ngoo, co-founder, CTO, Carousell

See all customers

## Partners

Google Cloud Machine Learning Partners come with deep AI expertise and can help you incorporate ML for a wide range of use cases across every stage of model development and serving.

intel | CISCO | Pluto7 | Atos | SpringML | NVIDIA

# Getting Started

Banner at top of page

+ Option to contact sales is nice. Hopefully they also answer questions

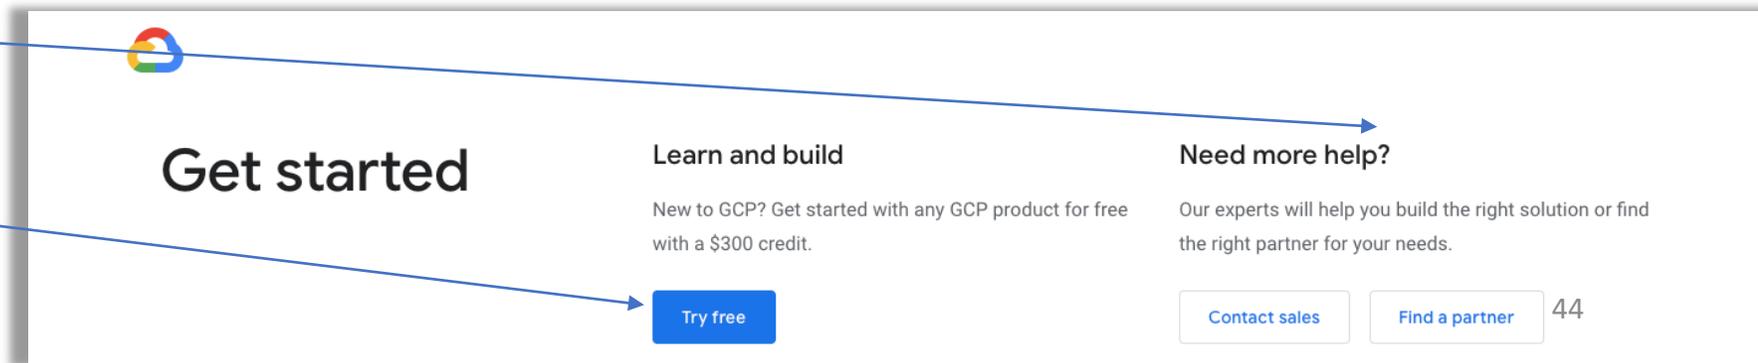
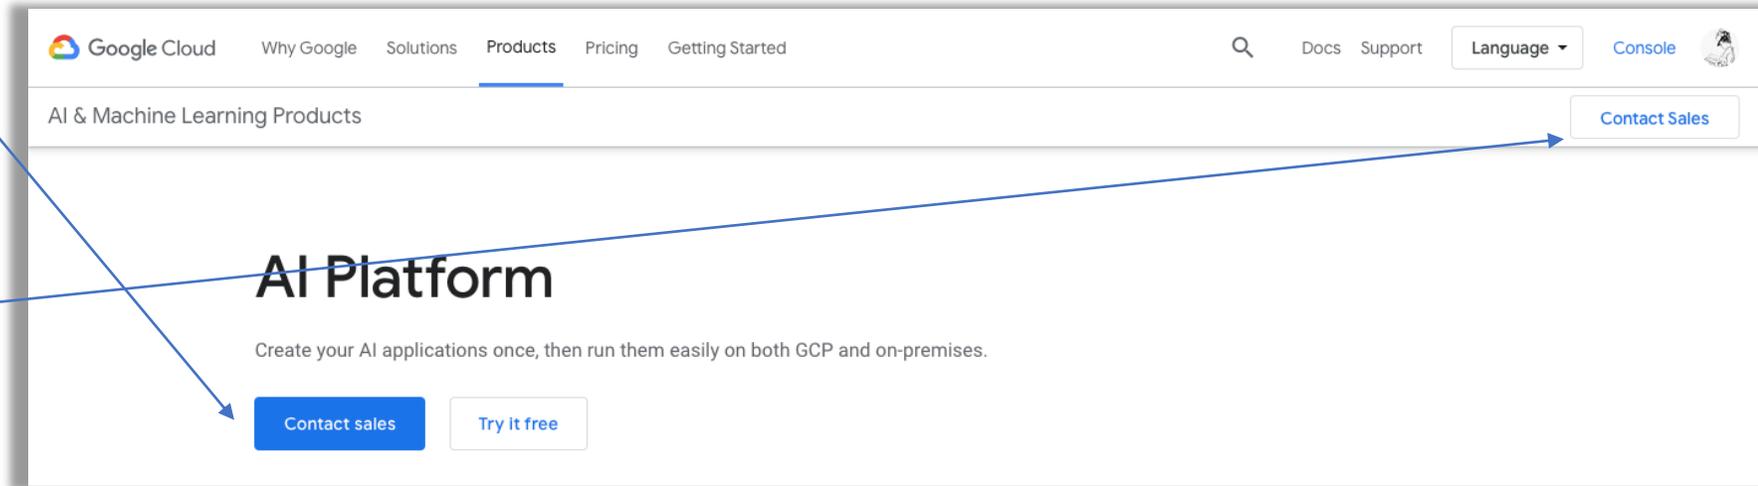
- Contact sales is potentially blunt as a CTA

- Redundant second contact sales button

Bottom of page banner

+ Help is more visible with GCP than Azure and AWS

- Unclear why bottom of page CTA is for the free trial while the top is for contact sales



Clicking “try free” goes to this page

— Users might expect that clicking from the previous page takes them directly to sign-up. This is an extra click

— Not an optimal flow for users of GCP who are new to the AI Platform. Hopefully the page would dynamically adapt

— This information feels relevant to the pricing page so unclear why it’s not there

+ Nice layout of benefits

Google Cloud

Language Console

# Google Cloud Platform Free Tier

Get started for free

**Always free products**  
You'll get free usage (up to monthly limits) of select products, including BigQuery and Compute Engine.

**\$300 free credit**  
New customers will also get \$300 to spend on Google Cloud Platform products during your first 12 months.

[View details →](#)

Always free products

There is no charge to use these products up to their specified free usage limit. The free usage limit does not expire, but is subject to change. Available for eligible customers.

[View details →](#)

COMPUTE	DATABASE	COMPUTE
<a href="#">Cloud Run</a>	<a href="#">Firestore</a>	<a href="#">Compute Engine</a>
<b>2 million</b>	<b>1 GB</b>	<b>1</b>
Requests per month	Storage	F1-micro instance per month
A fully managed environment to run stateless containers.	Scalable NoSQL document database.	Scalable, high-performance virtual machines.
<a href="#">View details →</a>	<a href="#">View details →</a>	<a href="#">View details →</a>

45

# Indirect Competitor

# IBM's Watson Machine Learning



Overview: Run machine-learning models anywhere, across any cloud. Bring your open source AI projects into production. For data scientists and developers

<b>Founded:</b> 2017	<b>Market position:</b> #4
<b>Pricing:</b> pay as you go or subscription	<b>Main segments:</b> AI lifecycle management, 1-click deployment, model operations, integrated UI end-to-end, deploy any model at scale, dynamic retraining



— Link to other Watson services feels like a distracting move away from this product

— Use of numbers is eye-catching but they don't feel consequential here. A user might wonder why should these stats matter to them

+ Nice breakdown of features

— Could use some visual interest to break up the page and make it scannable

<p>Watson Machine Learning simplifies and speeds deployment at scale</p> <p>→ <a href="#">Learn the economic impact of Watson Studio and Watson Knowledge Catalog</a></p>	<p>Moving beyond experimentation and maintaining production-level accuracy of AI models are still big challenges. Watson Machine Learning can accelerate the time to value of any model, with a projected ROI of 459 percent over three years, with payback in less than six months.<sup>1</sup></p>	<p>51%</p> <p>find optimizing, sustaining and expanding AI capabilities challenging<sup>2</sup></p>	<p>80%</p> <p>of AI pioneers are investing in the process to train algorithms<sup>3</sup></p>
<p>Watson Machine Learning features</p>	<p><b>AutoAI: AI lifecycle management</b></p> <p>Gain consistency and repeatability of end-to-end AI lifecycle management with AutoAI.</p> <p><a href="#">Learn more</a></p>	<p><b>One-click deployment</b></p> <p>When you are ready, save and select services to deploy. Get started with experimentation, evaluation and deployment.</p> <p><a href="#">Learn more</a></p>	<p><b>Model operations</b></p> <p>Use a machine learning framework for saving, deploying and versioning models, and for deployment of infrastructure management.</p> <p><a href="#">Learn more</a></p>
	<p><b>Integrated UI end-to-end</b></p> <p>Automate data prep, advanced data refinery and model explainability with Watson Studio and Watson Machine Learning.</p> <p><a href="#">Learn more</a></p>	<p><b>Deploy any model at scale</b></p> <p>Bring your favorite data science tools – open source, IBM SPSS® Modeler or Watson Studio.</p> <p><a href="#">Learn more</a></p>	<p><b>Dynamic retraining</b></p> <p>Enable continuous learning and automatically retrains your model to maintain model quality.</p> <p><a href="#">Learn more</a></p>

+ Effective comparison chart that makes the product look good, though it's easy to tout features above conventional ML tools

— What's concerning is that conventional ML tools aren't defined. There's a huge range and usage varies widely. It may not actually say much or be verifiably true

Compare features	
Watson Machine Learning	Conventional ML tools
Hybrid multicloud deployment	☑
Integrated with AutoAI	☑
Prescriptive analytics solvers	☑
Visual data science deployment	☑
Access to data preparation	☑
Advanced data refinery	☑
Dynamic continuous learning	☑
Neural network search	☑
1-click open-source deployment	☑
Automated AI lifecycle mgmt.	☑
Monitoring and debiasing	☑
Natural language integration	☑
Visual recognition service	☑

Strongest targeting to ML engineers

# Pricing

+ 2 CTAs for starting a trial offers good visibility

+ “Book a consultation” is friendlier language than Google’s “contact sales”

- Booking feels less immediate as an action

- Video is out of place on this page

IBM | Cloud | Why IBM | Cloud Paks | Products | Solutions | Pricing | Partners | Blog | Learn | Docs | More | Search | Cloud sign-up/log-in

Watson Machine Learning | Details | Pricing | Resources

## Watson Machine Learning: Pricing

Deploy self-learning models into production at scale.

[Start your free trial](#) [Book a consultation](#)

IBM Watson Studio | My Projects | Machine Learning | Bank marketing (sample)

Pipeline leaderboard

RANK	PIPELINE	RUC AUC	MODEL	ENHANCEMENTS	
1	Pipeline.6	0.895	XGB classifier	(HPS) (FE) (HPS+FE)	Save as model
2	Pipeline.2	0.895	XGB classifier	(HPS) (FE)	Save as model
3	Pipeline.2	0.895	XGB classifier	(HPS)	Save as model
4	Pipeline.5	0.892	XGB classifier	None	Save as model

Create a Watson Machine Learning model 05:03

### Pricing details

To suit all user needs, IBM Watson® Machine Learning has a variety of pricing plans, as well as cloud, server and local deployment options. Our plans give you the flexibility to deploy and run anywhere in a hybrid environment, helping you accelerate time to value for your data science projects. IBM Watson Machine Learning is also a key component for training and deploying AutoAI with IBM Watson Studio.

+ This is the best pricing page of the competitors. Clearly lays out options, features, and actual dollars. Easy to read, and strikes a balance between thorough and spare

+ Not shown, but in addition to 3 cloud plans, 1 server and 1 local plan is offered

The fastest way to move models from experimentations on your desktop or in your lab to deployments for production workloads. Runs in a managed service in a fully managed IBM Cloud™ environment.

**Features:**

- Pay-as-you-go consumption with as-a-service delivery, allowing you to be up and running in seconds
- Fully managed service, removing the need for infrastructure cost overhead
- Flexible interfaces available for managing models and configuring deployments
- End-to-end lifecycle support for machine learning, from training to deployment and evaluation
- First-class support for the most popular open source frameworks, as well as IBM SPSS®, IBM Decision Optimization for Watson Studio, and AutoAI with IBM Watson Studio

Plan	Features	Pricing
Watson Machine Learning Cloud Lite	Includes a maximum of five deployed models, 5,000 predictions per month, and 50 capacity unit hours per month for training or batch deployments.  During this period, models can be trained, evaluated, and deployed to return model predictions.	<a href="#">Free</a>
Watson Machine Learning Cloud Standard	Billed at a flat rate per thousand predictions and per capacity unit hour.  Predictions are consumed by making REST calls to a deployed online model; capacity unit hours are consumed by executing training and batch scoring jobs.	USD 0.50 /1,000 predictions  USD 0.50 /capacity unit hour  <a href="#">Buy now</a> <a href="#">Get a cost estimate</a>
Watson Machine Learning Cloud Professional	Billed at a fixed amount per month, including 2,000,000 predictions and 1,000 capacity unit hours. Overage is billed at a flat rate per thousand predictions and per capacity unit hour.  Predictions are consumed by making REST calls to a deployed online model; capacity unit hours are consumed by executing training and batch scoring jobs.	USD 1,000.00 /month  Overages: USD 0.40 /1,000 predictions  USD 0.40 /capacity unit hour  <a href="#">Buy now</a> <a href="#">Get a cost estimate</a>

# Documentation

Banner on home page

+ Nice modern visual design and content

- Doesn't match other branding

+ This is the most intriguing feature and one that no other competitor offers

Expert resources to help you succeed

**Community**  
Get technical tips and insights from other users in the Watson Studio community.

**Knowledge center**  
Find answers quickly in the IBM Watson Machine Learning documentation.

**Support**  
Learn more about technical support.

Additional resources page from tab navigation

This is one page where I'd advocate not having prominent trial and consultation CTAs

Again, video doesn't directly relate to the topic

Unclear what "featured" means. Promoted? If so, why?

Surprising that documentation isn't more prominently shown

Study seems to be at a lower level than some of these other higher-level items like documentation

+ Nice links to GitHub and Stack Overflow establish credibility by association

"Quick links" is also a puzzling header

FAQ not labeled

The screenshot shows the IBM Watson Machine Learning Resources page. The main heading is "Watson Machine Learning: Resources" with the subtext "Deploy self-learning models into production at scale". There are two prominent CTAs: "Start your free trial" and "Book a consultation". A video player is embedded, showing a "Pipeline leaderboard" with a table of metrics. Below the video is a grid of featured resources including "Data science community", "Github repository", "Stack Overflow", "Forrester study", "Documentation", and "Solution brief". A "Quick links" section follows, containing "Watson products and services". At the bottom is an FAQ section with questions like "What is Watson Machine Learning?", "Do the models need to be built using only IBM offerings?", and "Does Watson Machine Learning integrate with Decision Optimization?".

Rank	Model	Accuracy	Model Size	Model Type	Model Status	Model Link
1	Classification	0.895	100 MB	IBM Classifier	DEVELOPMENT	See as model
2	Classification	0.875	100 MB	IBM Classifier	DEVELOPMENT	See as model
3	Classification	0.875	100 MB	IBM Classifier	DEVELOPMENT	See as model
4	Classification	0.870	100 MB	IBM Classifier	DEVELOPMENT	See as model

# Testimonials

— Description text in the imperative is momentarily confusing. If this is what that client achieved, writing it that way would be more effective

— These look like generic stock photos, which isn't appealing

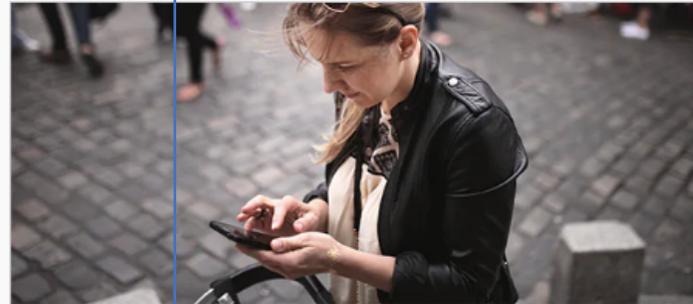
## Watson Machine Learning in action



### Anyline

Train and deploy deep learning models for offline OCR (optional character recognition) on mobile and embedded devices in record time.

→ [Watch the video \(01:30\)](#)



### Caixa Geral de Depósitos France

Accelerate customer service by delivering faster decisions using an innovative hybrid-cloud app enabled by machine learning.

→ [Read the case study](#)



### KIST Europe

Explore how machine learning can automate quality management processes, reduce defect rates and save costs.

→ [Read the case study](#)

# Getting Started

Top of page banner

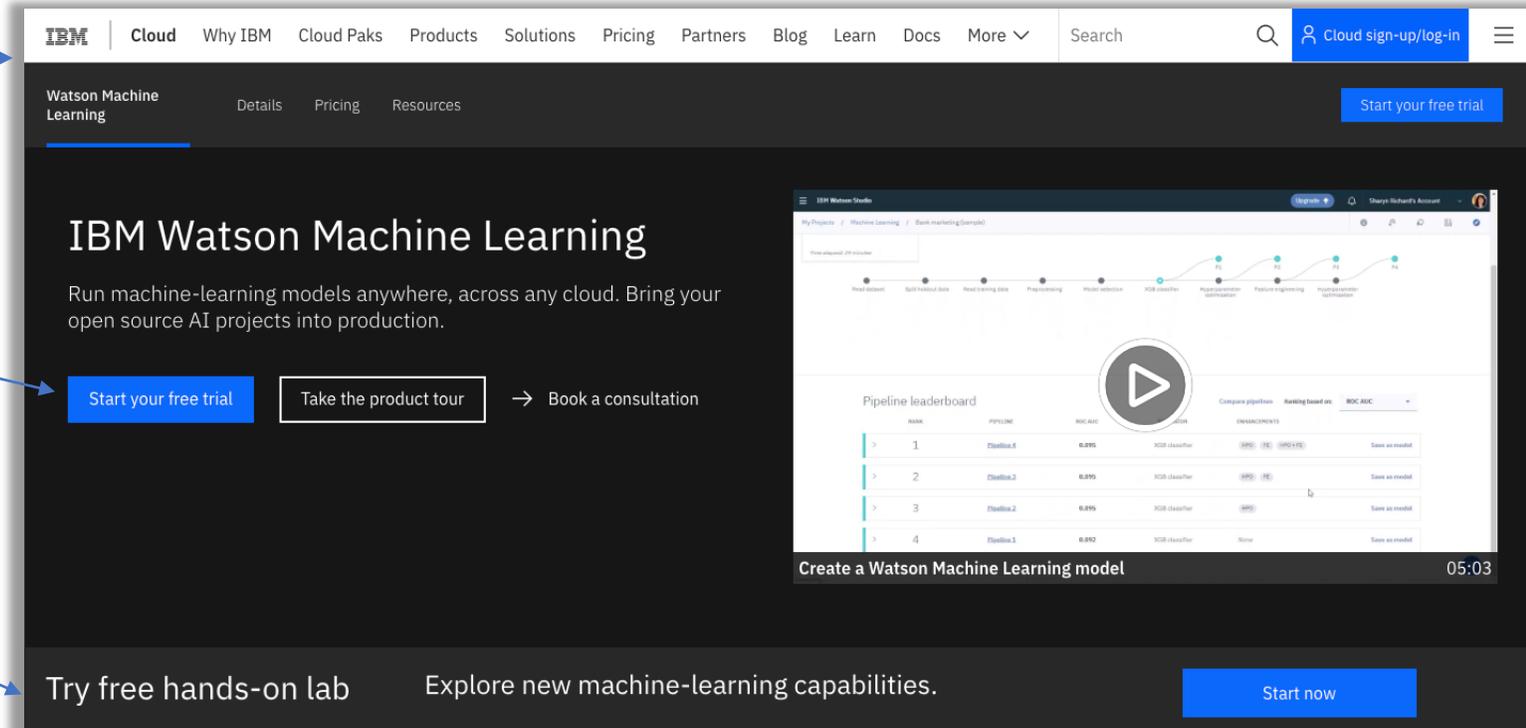
4 CTAs is likely too much

Vague description but intriguing lab

Bottom of page

Different and less visually salient look, which may go unnoticed

Seeing news doesn't seem to belong with getting started



+ As with AWS, clicking from the previous page brings the user directly to sign-up

- Different visual branding again. Site may have been pieced together by disparate design teams

+ Nice that no credit card is required

+ Good reminder of benefits

+ Nice ability to learn more

- Nothing has been done yet but the screen is already showing an error message

+ Accounts for users already on their cloud who are trying the ML service

Already have an IBM Cloud account? [Log in](#)

### Create a free account

Join us in the cloud and start building today.

Email

Enter an email address.

First Name

Last Name

Country or Region

United States

Password

57

# Influencers (Select Features Only)

# SAP's Artificial Intelligence

- Overview: Make confident decisions, automate repetitive tasks, and deliver human experiences with embedded artificial intelligence and machine learning
- Main segments: intelligent robotic process automation, conversational AI, data intelligence, service ticket intelligence, cash application
- Low market share

# See how customers are succeeding with SAP

89%

Of crises diverted, on average

Find out how HarrisLogic is keeping potential offenders off the street by gathering data from seven different data sources.

[Watch the video >](#)

15%

Decrease in inventory levels

Take a look at how Proximus Group opens up a world of digital opportunities to help people live better and work smarter.

[Watch the video >](#)

+ Very effective use of stats in selling the product

[View customer success stories >](#)

+ Nice to have endorsements from experts. Not seen elsewhere

+ Nice link to additional content with an intriguing title

- On second glance it appears to be in-house content, which is contrary to my expectation of articles from academics and field experts

## What are leading analysts saying about SAP?

FORRESTER®

Demand for machine learning continues to grow

Examine the role of predictive analytics and machine learning models in becoming intelligence-driven, and get practical tips for integrating them business-wide.

[Read the study >](#)

IDC | ANALYZE THE FUTURE

SAP named a major player in AI-enabled marketing

Review the IDC MarketScape report detailing how SAP Customer Experience solutions help customers apply AI and machine learning to diverse and complex use cases.

[Read the report excerpt >](#)

## Get the latest news and trends from experts



Susan Galer  
Marketing Strategy and Thought Leadership,  
SAP

Building intelligent bots in three minutes

Find out how developers are creating intelligent bots to automate workflows and streamline processes using SAP Intelligent Robotic Process Automation.

[Read the blog >](#)



Markus Noga  
VP Machine Learning,  
SAP

Bringing transparency into AI

Learn about the need for technical transparency, as companies are increasingly using machine learning models to make decisions that affect people's lives.

[Read the blog >](#)



Maricel Cabahug  
Chief Design Officer,  
SAP

Designing AI for people

Consider possible guidelines to ensure that AI systems are designed to maximize human potential, while still protecting our best interests.

[Read the blog >](#)

+ Of all competitors, this is the most user-friendly help section

+ Nice that other country numbers are offered

- Call and chat options should be clickable

- Unclear what "offline" means

# Questions? Get in touch!

Call us at

United States

**+1-800-872-1727**

Or see our complete list of [local country numbers](#)

---

-  **Call Offline**  
SAP can call you to discuss any questions you have.
-  **Chat Offline**  
Get live help and chat with an SAP representative.
-  **Contact Us**  
E-mail us with comments, questions or feedback.

# Salesforce's Einstein

- Overview: Get more done with Einstein, your smart CRM assistant. Make decisions faster, make employees more productive, and make customers happier using AI across the Salesforce Customer 360 Platform
- Main segments: machine learning (discovery, prediction builder, next best action), natural language processing, computer vision, automatic speech recognition
- Low market share

MACHINE LEARNING

NATURAL LANGUAGE PROCESSING

COMPUTER VISION

AUTOMATIC SPEECH RECOGNITION

Nice tabular breakdown of product areas with advance on scroll



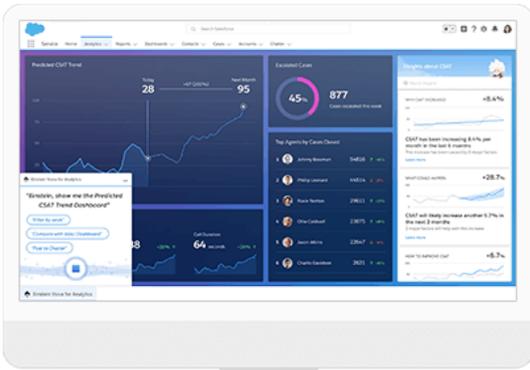
MACHINE LEARNING

Get more predictive about your business and customers.

Machine learning uses past data to predict what will happen in the future with minimal programming.

Language indicates these services are more for AI beginners

It's likely strategic that both small competitors (here marked influencers) prominently offer help. It's an area where they can offer value to make up for a less mature product



Einstein Discovery

Boost productivity and discover relevant patterns in all your data, whether it lives in Salesforce or outside. Find simple AI insights and recommendations to tough problems. Then, take action on your findings without ever leaving Salesforce.

WATCH DEMO >

Also a nice help section

Different color background works here to draw attention to the offer

Questions? We'll put you on the right path.

Ask about Salesforce products, pricing, implementation, or anything else. Our highly trained reps are standing by, ready to help.

CONTACT US

OR CALL 1-800-667-6389



# Summary

# Comparison Chart

These are grades aggregated for each focus area by provider:

Scale:		Overview	Pricing	Documen- tation	Testimonials	Getting Started
Excellent						
Good	<b>Azure</b>	Average	Average	Poor	Excellent	Average
Average	<b>AWS</b>	Excellent	Very Poor	Excellent	Very Poor	Average
Poor	<b>GCP</b>	Good	Very Poor	Very Poor	Average	Average
Very Poor	<b>IBM</b>	Poor	Excellent	Poor	Poor	Average

# Findings Summary

- As expected, **direct competitors have the most similarity with us in product structure and functionality, as well as presentation.** Direct competitors structure their products as full-service, end-to-end MLaaS tied to their respective clouds
- **As competitors get smaller they partition their product more, selling it as parts, and tie it less to the cloud.** Our indirect competitor shares moderate similarity with us, as they break their service into 2 parts and offer non-cloud options
- Influencers have the least similarity with us. These **small competitors** are on the level of selling one-off AI services rather than a cohesive solution, **resembling our Cognitive Services rather than AML.** Their main businesses are in tangential areas, and their AI offerings are basic
- **Big players are in the market for enterprise customers,** and may also pick up consultancies and startups. **Small players are after low-hanging fruit from small clients**
- **No single competitor does the best across the 5 interest areas.** Each has strengths and weaknesses we can learn from. Nobody does a great job with onboarding, which is a big opportunity
- There are similarities in messaging and tone throughout, with nuances in which are more effective. However, the generalizations point to the fact that in some or even many cases, the **buyer is not the end user, so being careful of language that is inclusive of many parties is critical**

- In fitting with their branding, **Amazon has the most info-rich, maximalist presentation, while Google has the most edited, minimalist presentation** (though clicking into some pages brings the user to a few extra-dense pages). **We are somewhere in the middle**, and I'd argue, should stay that way, but with improvements inspired by learnings from this study
- Competitive advantages
  - **AWS has the most critical mass behind their cloud, and likely their MLaaS product.** They're likely able to funnel considerable cloud profits into developing Sagemaker and convert existing customers. They have strong credibility with startups and the best reputation
  - **GCP has strong investments in open source** like TensorFlow, which give them considerable influence outside their cloud, and other popular investments like Colab and Kaggle, and are likely to pick up many more acquisitions through their parent Alphabet
  - **We have tech-forward and interesting features coming out of MSR** like AI fairness, and have strong hook-ups with related areas like MLOps, data to AI, and notebooks/IDE options
- See the next slide for product feature differentiation

# Overview Page Comparison

These are feature novelties presented on the home page by each provider:

Azure Machine Learning	AWS SageMaker	GCP AI Platform	IBM Watson ML
<ul style="list-style-type: none"><li>• MLOps</li><li>• Responsible AI</li><li>• Security</li></ul>	<ul style="list-style-type: none"><li>• Data labeling by humans</li><li>• Prediction validation by humans</li><li>• IDE for ML</li><li>• Debugger</li><li>• Elastic inference</li></ul>	<ul style="list-style-type: none"><li>• Kubeflow and other in-house open source tools</li></ul>	<ul style="list-style-type: none"><li>• API generation</li><li>• Lifecycle management</li><li>• Dynamic retraining</li></ul>

# Recommendations: Overview of Services

Recommendation	Inspiration	Priority
<b>Introduce more links to product detail pages throughout to offset the streamlined quantity of content the scroll tab navigation can accommodate</b>	n/a. our product weakness	high
Unify the information structure and visual presentation across main product demo areas	n/a	medium
<b>Ensure these highlighted areas are the most relevant ones for users to see. Ours is more feature-related while AWS and GCP are more ML stage-related (Salesforce too, but they're selling individual services)</b>	AWS, GCP, Salesforce	high
<b>Provide indication of the basic architecture of the product on home page. This could take the form of a graphic like with AWS and GCP</b>	AWS, GCP	high
Explore linking other engagement forms like events or conferences (maybe AML presence at upcoming Build)	AWS	low
Show the improvement compared with traditional ML tools/methods	AWS	low
Showcase the exclusive benefits of our product that competitors don't have	AWS	medium
Use statistics to show concrete value add with our product	AWS, SAP	medium
Potentially address on-prem compatibility (GCP) or compatibility across other clouds (IBM), but be careful, as we still want to convert users to our product and cloud	GCP, IBM	low

# Recommendations: Pricing

Recommendation	Inspiration	Priority
Offer pricing help (maybe following “consultation” language like IBM)	IBM	medium
Give info on how to sign up for enterprise beta (GCP also has beta but currently also doesn't give sign-up info)	GCP	medium
Offer buy now info in addition to free trial	n/a	low
Redesign the page to be more scannable with bulleted or chunked info, tables, highlights	IBM	medium

# Recommendations: Documentation

Recommendation	Inspiration	Priority
Better integrate the two parts of the page dedicated to how to use the product (AWS has a more cohesive design)	AWS	medium
Build out a more prominent documentation section, potentially with the use of cards (AWS) or other visual cues	AWS	medium
Explore whether there's any appetite for a community feature like IBM's	IBM	low

# Recommendations: Testimonials

Recommendation	Inspiration	Priority
Explore additional endorsements (like GCP's partners) or experts (SAP)	GCP, SAP	low

# Recommendations: Getting Started

Recommendation	Inspiration	Priority
Either remove the extra click to sign up (like AWS and IBM), or rephrase the CTA to make clear that there's an extra step (GCP also has extra step but fails to warn users)	AWS, IBM, GCP	medium
<b>Ensure that users new to AML but not to Azure don't create another account. Ideally they're welcomed as a current user</b>	<b>AWS, IBM</b>	<b>high</b>
Offer a way to reach sales	GCP	medium
Offer help in general	GCP, SAP, Salesforce	medium

# Conclusion

To improve our pre-paywall experience, it's recommended to focus on the **overview experience first**, as it outlines our offerings and value proposition to potential users, and the **getting started experience second**, because it's the locus of conversion from potential to actual users. **Documentation is third**, as it details product functionality and is our weakest area

**Top 5 recommendations** include:

- Providing a visual of the basic architecture of our product on home page
- Validating that feature areas highlighted on home page are the most relevant ones
- Showcasing our commitment to cutting-edge new ML features in our product
- Introducing more links to product detail pages to allow further reading
- Ensuring that users new to AML but not to Azure are suitably onboarded

# Next Steps/Further Research

- If desired, add analysis of low-priority indirect competitor IBM Watson Studio and influencer Oracle
- **Make recommended changes; test usability if possible; record telemetry and analyze for effects**
- Competitive usability research on free trial experience for new users or ML service addition for existing cloud users
- Competitive usability research on a selected number of core tasks (i.e. building a model, training a model, deploying a model) with standardized materials
- Competitive analysis across machine learning feature specifics, APIs, or specialized tools as part of MLaaS services

# Thank you

Contact me for more information or with questions