

AWS Announces Nine New Amazon SageMaker Capabilities

Amazon SageMaker Data Wrangler provides the fastest and easiest way for developers to prepare data for machine learning

Amazon SageMaker Feature Store delivers a purpose-built data store for storing, updating, retrieving, and sharing machine learning features

Amazon SageMaker Pipelines gives developers the first purpose-built, easy-to-use continuous integration and continuous delivery (CI/CD) service for machine learning

Amazon SageMaker Clarify provides developers with greater visibility into their training data so they can limit bias in machine learning models and explain predictions

Deep profiling for Amazon SageMaker Debugger monitors machine learning training performance to help developers train models faster

Distributed Training on Amazon SageMaker delivers new capabilities that can train large models up to two times faster than would otherwise be possible with today's machine learning processors

Amazon SageMaker Edge Manager delivers machine learning model monitoring and management for edge devices to ensure that models deployed in production are operating correctly

Amazon SageMaker JumpStart provides a developer portal for pre-trained models and pre-built workflows

SEATTLE, WA – **December 8, 2020** – Today at AWS re:Invent, Amazon Web Services, Inc. (AWS), an Amazon.com, Inc. company (NASDAQ: AMZN), announced nine new capabilities for its industry-leading machine learning service, Amazon SageMaker, making it even easier for developers to automate and scale all steps of the end-to-end machine learning workflow. Today's announcements bring together powerful new capabilities like faster data preparation, a purpose-built repository for prepared data, workflow automation, greater transparency into training data to mitigate bias and explain predictions, distributed training capabilities to train large models up to two times faster, and model monitoring on edge devices. To get started with Amazon SageMaker, visit: <https://aws.amazon.com/sagemaker>

Machine learning is becoming more mainstream, but it is still evolving at a rapid clip. With all the attention machine learning has received, it seems like it should be simple to create machine learning models, but it isn't. In order to create a model, developers need to start with the highly manual process of preparing the data. Then they need to visualize it in notebooks, pick the right algorithm, set up the framework, train the model, tune millions of possible parameters, deploy the model, and monitor its performance. This process needs to be continuously repeated to ensure that the model is performing as expected over time. In the past, this process put machine learning out of the reach of all but the most skilled developers. However, Amazon SageMaker has changed that. Amazon SageMaker is a fully managed service that removes challenges from each stage of the machine learning process, making it radically easier and faster for everyday developers and data scientists to build, train, and deploy machine learning models. Tens of thousands of customers utilize Amazon SageMaker to help accelerate their machine learning deployments, including 3M, ADP, AstraZeneca, Avis, Bayer, Bundesliga, Capital One, Cerner, Chick-fil-A, Convoy, Domino's Pizza, Fidelity Investments, GE Healthcare, Georgia-Pacific,

Hearst, iFood, iHeartMedia, JPMorgan Chase, Intuit, Lenovo, Lyft, National Football League, Nerdwallet, T-Mobile, Thomson Reuters, and Vanguard.

Today's announcements build on the more than 50 new Amazon SageMaker capabilities that AWS has delivered in the past year to make it even easier for developers and data scientists to prepare, build, train, deploy, and manage machine learning models, including:

- **Amazon SageMaker Data Wrangler automated data preparation:** Amazon SageMaker Data Wrangler provides the fastest and easiest way to prepare data for machine learning. Data preparation for machine learning is a difficult process. This difficulty arises from the fact that data attributes (known as features) used to train a machine learning model often come from different sources and exist in various formats. This means that developers must spend considerable time extracting and normalizing this data so it's consistently easy to use with machine learning. Customers might also want to combine features into composite features to give the machine learning model more helpful inputs. For example, a customer might want to create a feature that describes a group of customers that are prolific spenders so they can be offered loyalty program rewards by combining features for items previously purchased, amount spent, and frequency of purchases. The work associated with transforming data into features is called feature engineering, and it consumes a lot of time for developers when they're building machine learning models. Amazon SageMaker Data Wrangler radically simplifies the process of data preparation and feature engineering. With Amazon SageMaker Data Wrangler, customers can choose the data they want from their various data stores and import it with a single click. Amazon SageMaker Data Wrangler contains over 300 built-in data transformers that can help customers normalize, transform, and combine features without having to write any code, while managing all of the processing infrastructure under the hood. Customers can quickly preview and inspect that these transformations are what was intended by viewing them in SageMaker Studio (the first end-to-end Integrated Development Environment for machine learning). Once the features have been engineered, Amazon SageMaker Data Wrangler will save them for reuse in the Amazon SageMaker Feature Store.
- **Amazon SageMaker Feature Store feature storage and management:** Amazon SageMaker Feature Store provides a new repository that makes it easy to store, update, retrieve, and share machine learning features for training and inference. Today, customers can save their features to Amazon Simple Storage Service (S3). This works well for a simple set of features that are mapped to a single model, but most features are not mapped to only one model. Most features are used repeatedly by multiple models and multiple developers and data scientists, and as new features are created, developers also want to be able to reuse them repeatedly. This leads to multiple S3 objects to manage, which can quickly become difficult to manage. Developers and data scientists try to solve this by using spreadsheets, paper notes, and emails. Sometimes they even try to build a custom application to keep track of the features, but this is a lot of work and error-prone. Further, developers and data scientists need the same features not only to train multiple models with all of the data available and where this activity can happen over hours, but also to use during inference when the predictions need to be returned in milliseconds and often use just a subset of the data in relevant features. For example, a developer might want to create a model that predicts the next best song in a playlist. To do this, developers would train the model on thousands of songs and then provide the model the last three songs played during inference to predict the next song. Training and inference are very different uses cases. During training, the models can access the features offline and in batch, but for inference, the model needs only a subset of the features in near real-time. Since machine learning models have a

single source of features that need to be consistent, these different access patterns make it challenging to keep the features consistent and up to date. Amazon SageMaker Feature Store solves this problem by providing a purpose-built feature store where developers can access and share features that make it much easier to name, organize, find, and share sets of features among teams of developers and data scientists. Since Amazon SageMaker Feature Store resides in Amazon SageMaker Studio—close to where machine learning models are run—it provides single-digit millisecond latency for inference. Amazon SageMaker Feature Store makes it simple and easy to organize and update large batches of features for training and smaller instantiations of them for inference. That way, there's one consistent view of features for machine learning models to use and it becomes significantly easier to generate models that produce highly accurate predictions.

- **Amazon SageMaker Pipelines workflow management and automation:** Amazon SageMaker Pipelines is the first purpose-built, easy-to-use continuous integration and continuous delivery (CI/CD) service for machine learning. As customers can see with feature engineering, machine learning comprises multiple steps that can benefit from orchestration and automation. This is not dissimilar to traditional programming, where customers have tools like CI/CD to help them develop and deploy applications more quickly. However, with machine learning, CI/CD tools are rarely used because they don't exist or because they are hard to set up, configure, and manage. With Amazon SageMaker Pipelines, developers can define each step of an end-to-end machine learning workflow. These workflows include the data-load steps, transformations from Amazon SageMaker Data Wrangler, features stored in Amazon SageMaker Feature Store, training configuration and algorithm set up, debugging steps, and optimization steps. With Amazon SageMaker Pipelines, developers can easily re-run an end-to-end workflow from Amazon SageMaker Studio, using the same settings to get the exact same model every time, or they can re-run the workflow on a regular schedule with new data to update a model. Amazon SageMaker Pipelines logs each step in Amazon SageMaker Experiments (an Amazon SageMaker capability that organizes and tracks machine learning experiments and model versions) every time a workflow is run. This helps developers visualize and compare machine learning model iterations, training parameters, and outcomes. With Amazon SageMaker Pipelines, workflows can be shared and re-used between teams, either to recreate a model or to act as a starting point for making improvements through new features, algorithms, or optimizations.
- **Amazon SageMaker Clarify bias detection and explainability:** Amazon SageMaker Clarify provides bias detection across the machine learning workflow, enabling developers to build greater fairness and transparency into their machine learning models. Once developers have prepared data for training and inference, they need to try to ensure the data is free from statistical bias and that model predictions are transparent, so they can explain how the model features are contributing to predictions. Today, developers sometimes try to use open source tools to detect statistical bias in their data, but these tools require a lot of manual effort and coding and are typically error prone. With Amazon SageMaker Clarify, developers can now more easily detect statistical bias across the entire machine learning workflow and provide explanations for predictions their machine learning models are making. Amazon SageMaker Clarify integrates with Amazon SageMaker Data Wrangler where it runs a set of algorithms on features to identify bias during data preparation with visualizations that include a description of the sources and severity of possible bias. This way, developers can take steps for mitigation. Amazon SageMaker Clarify also integrates with Amazon SageMaker Experiments to make it easier to check trained models for statistical bias. It also details how each feature inputted into the model is affecting predictions. Finally, Amazon SageMaker Clarify integrates with Amazon SageMaker Model Monitor (an Amazon SageMaker capability that continuously monitors the

quality of machine learning models in production) to alert developers if the importance of model features shifts and causes model behavior to change.

- **Deep Profiling for Amazon SageMaker Debugger model training profiler:** Deep Profiling for Amazon SageMaker Debugger now enables developers to train their models faster by automatically monitoring system resource utilization and providing alerts for training bottlenecks. Today, developers don't have a standard way to monitor system utilization (e.g. GPU, CPU, network throughput, and memory I/O) to identify and troubleshoot bottlenecks in their training jobs. As a result, developers can't train models as quickly and cost effectively as possible. Amazon SageMaker Debugger solves this problem with Deep Profiling's newly announced capabilities, which provide developers the ability to visually profile and monitor system resource utilization in Amazon SageMaker Studio. This makes it easier to root cause issues and reduce the time and cost of training machine learning models. With these new capabilities, Amazon SageMaker Debugger expands its scope to monitor the utilization of system resources, send out alerts on problems during training in Amazon SageMaker Studio or via AWS CloudWatch, and correlate usage to different phases in the training job or a specific point in time during training (e.g. 28 minutes after the training job started). Amazon SageMaker Debugger can also trigger actions based on alerts (e.g. stop a training job when irregularities in GPU usage are detected). Amazon SageMaker Debugger's Deep Profiling works across frameworks (PyTorch, Apache MXNet, and TensorFlow) and collects necessary system and training metrics automatically without requiring any code changes in training scripts. This allows developers to visualize how their system resources were used during training in Amazon SageMaker Studio.
- **Distributed Training on Amazon SageMaker accelerates training times:** New Distributed Training on Amazon SageMaker makes it possible to train large, complex deep learning models up to two times faster than current approaches. Today, advanced machine learning use cases—such as natural language processing for intelligent assistants, object detection and classification for autonomous vehicles, and image classification for large-scale content moderation—demand increasingly large datasets and more graphics processing unit (GPU) memory for training. However, some of these models are too big to fit in the memory provided by a single GPU. Customers can attempt to split models across multiple GPUs, but finding the best way to split the model and adjusting training code can often take weeks of tedious experimentation. To overcome these challenges, Distributed Training on Amazon SageMaker offers two distributed training capabilities that enable developers to train large models up to two times faster at no additional cost. Distributed Training with Amazon SageMaker's Data Parallelism engine scales training jobs from one GPU to hundreds or thousands by automatically splitting data across multiple GPUs, improving training time by up to 40%. The reduction in training time is possible because Amazon SageMaker's Data Parallelism engine manages GPUs for optimal synchronization using algorithms that are purposefully built to fully utilize AWS infrastructure with near-linear scaling efficiency. Distributed Training with Amazon SageMaker's Model Parallelism engine can efficiently split large, complex models with billions of parameters across multiple GPUs by automatically profiling and identifying the best way to partition models. They do this by using graph partitioning algorithms to optimally balance computation and minimize communication between GPUs, resulting in minimal code changes and fewer errors caused by GPU memory constraints.
- **Amazon SageMaker Edge Manager model management for edge devices:** Amazon SageMaker Edge Manager allows developers to optimize, secure, monitor, and maintain machine learning models deployed on fleets of edge devices. Today, customers use Amazon SageMaker Neo to create optimized models for edge devices that run up to twice as fast, with less than a tenth of

the memory footprint and no loss in accuracy. However, after deployment on edge devices, customers still need to manage and monitor the models to ensure they continue to perform with high accuracy. Amazon SageMaker Edge Manager optimizes models to run faster on target devices and provides model management for edge devices, so customers can prepare, run, monitor, and update deployed machine learning models across fleets of devices at the edge. Amazon SageMaker Edge Manager gives customers the ability to cryptographically sign their models, upload prediction data from their devices to Amazon SageMaker for monitoring and analysis, and view a dashboard that tracks and visually reports on the operation of the deployed models within the Amazon SageMaker console. Amazon SageMaker Edge Manager extends capabilities that were previously only available in the cloud by sampling data from edge devices and sending it to Amazon SageMaker Model Monitor for analysis, so developers can continuously improve model quality by retraining them when their accuracy declines over time.

- **Amazon SageMaker JumpStart enables the machine learning journey:** Amazon SageMaker JumpStart provides developers an easy-to-use, searchable interface to find best-in-class solutions, algorithms, and sample notebooks. Today, some customers that lack experience with machine learning have difficulty getting started with machine learning deployments, while more advanced developers find it difficult to adopt machine learning for all of their use cases. With today's launch of Amazon SageMaker JumpStart, customers can now quickly find relevant information specific to their machine learning use cases. Developers new to machine learning will be able to select from several complete end-to-end machine learning solutions (e.g. fraud detection, customer churn prediction, or forecasting) and deploy them directly in their Amazon SageMaker Studio environments. And, experienced users will be able to choose from more than a hundred machine learning models to quickly get started on building and training models.

"Hundreds of thousands of everyday developers and data scientists have used our industry-leading machine learning service, Amazon SageMaker, to remove barriers to building, training, and deploying custom machine learning models. One of the best parts about having such a widely-adopted service like SageMaker is that we get lots of customer suggestions which fuel our next set of deliverables," said Swami Sivasubramanian, Vice President, Amazon Machine Learning, Amazon Web Services, Inc. "Today, we are announcing a set of tools for Amazon SageMaker that makes it much easier for developers to build end-to-end machine learning pipelines to prepare, build, train, explain, inspect, monitor, debug, and run custom machine learning models with greater visibility, explainability, and automation at scale."

With corporate operations in 70 countries and sales in 200, 3M is creating the technology and products that advance every company, enhance every home, and improve everyday life. "3M's success is grounded in our entrepreneurial researchers and our constant focus on science. One way we have advanced the science of our products is the adoption of machine learning on AWS," said David Frazee, Technical Director at 3M Corporate Systems Research Lab. "Using machine learning, 3M is improving tried-and-tested products, like sandpaper, and driving innovation in several other spaces, including healthcare. As we plan to scale machine learning to more areas of 3M, we see the amount of data and models growing rapidly – doubling every year. We are enthusiastic about the new Amazon SageMaker features because they will help us scale. Amazon SageMaker Data Wrangler makes it much easier to prepare data for model training, and Amazon SageMaker Feature Store will eliminate the need to create the same model features over and over. Finally, Amazon SageMaker Pipelines will help us automate data prep, model building, and model deployment into an end-to-end workflow so we can speed time to market for our models. Our researchers are looking forward to taking advantage of the new speed of science at 3M."

Deloitte is helping transform organizations around the globe. The organization continuously evolves how it works and how it looks at marketplace challenges so it can continue to deliver measurable, sustainable results for its clients and communities. “Amazon SageMaker Data Wrangler enables us to hit the ground running to address our data preparation needs with a rich collection of transformation tools that accelerate the process of machine learning data preparation needed to take new products to market,” said Frank Farrall, Principal, AI Ecosystems and Platforms Leader at Deloitte. “In turn, our clients benefit from the rate at which we scale deployments, enabling us to deliver measurable, sustainable results that meet the needs of our clients in a matter of days rather than months.”

A subsidiary of Koch Industries since 2004, INVISTA brings to market the proprietary ingredients for nylon 6,6 and recognized brands including STAINMASTER, CORDURA, and ANTRON. It is one of the world’s largest integrated producers of chemical intermediates, polymers, and fibers. “At INVISTA, we are driven by transformation and look to develop products and technologies that benefit customers around the globe,” said Caleb Wilkinson, Lead Data Scientist at INVISTA. “We see machine learning as a way to improve the customer experience, but with datasets that span hundreds of millions of rows, we needed a solution to help us prepare data, and develop, deploy, and manage machine learning models at scale. To speed these processes, we worked with the AWS team on several new features. With Amazon SageMaker Data Wrangler, we can now interactively select, clean, explore, and understand our data effectively, empowering our data science team to create feature engineering pipelines that can scale effortlessly to datasets that span hundreds of millions of rows. We can also easily automate and manage machine learning workflows at scale with Amazon SageMaker Pipelines, so we can easily stitch together individual steps of the machine learning workflow. Together with Amazon SageMaker Data Wrangler and Amazon SageMaker Pipelines, we can operationalize our machine learning workflows faster.”

Snowflake Data Cloud shatters the barriers that have prevented organizations of all sizes from unleashing the true value from their data. “One of the biggest challenges our enterprise customers face is preparing data for machine learning projects,” said Christian Kleinerman, SVP of Product at Snowflake. “We’re excited about Amazon SageMaker Data Wrangler, which makes it easier for organizations to aggregate and prepare data for machine learning. With the addition of Snowflake as a data source in Amazon SageMaker Data Wrangler, joint customers will soon be able to leverage the integrated platform capabilities of Snowflake, together with the interactive data preparation and machine learning capabilities of Amazon SageMaker. Customers will have the ability to go from raw data to machine learning models and insights faster than previously possible.”

Founded in 2013 by the original creators of Apache Spark™, Delta Lake and MLflow, Databricks brings together data engineering, science, and analytics on an open, unified platform so data teams can collaborate and innovate faster. “At Databricks, we are committed to bringing together data engineering and science and analytics so data teams can collaborate and innovate faster,” said Adam Conway, SVP of Products at Databricks. “We are looking forward to continuing our partnership with AWS in 2021, especially with the seamless integration our customers can experience with AWS on Amazon SageMaker Data Wrangler. With this partnership, our customers can leverage Delta Lake with Amazon SageMaker to prepare robust training data so they can create the most accurate machine learning models.”

MongoDB Atlas is the fully managed service for MongoDB, the popular database designed to help teams build, scale, and iterate quickly. “Our mission at MongoDB is to free the genius within everyone by making data stunningly easy to work with. MongoDB Atlas runs more than 1.5 million database clusters,

AMAZON CONFIDENTIAL

powering critical applications for our customers; we want to make it easy to build, train, and deploy machine learning models based on the data those applications generate,” said Mark Porter, CTO at MongoDB. “We are excited that our customers now have a faster, visual way to aggregate and prepare data for machine learning using Amazon SageMaker Data Wrangler. Coming in 2021, our customers will soon be able to query and analyze data across Amazon S3 and MongoDB Atlas within Amazon SageMaker Data Wrangler, enabling them to get more value from their data faster.”

Intuit is a mission-driven, global financial platform company and proud maker of TurboTax, QuickBooks, and Mint. “We chose to build Intuit’s new machine learning platform on AWS in 2017, combining Amazon SageMaker’s powerful capabilities for model development, training, and hosting with Intuit’s own capabilities in orchestration and feature engineering,” said Mammad Zadeh, Intuit Vice President of Engineering, Data Platform. “As a result, we cut our model development lifecycle dramatically. What used to take six full months now takes less than a week, making it possible for us to push AI capabilities into our TurboTax, QuickBooks, and Mint products at a greatly accelerated rate. We have worked closely with AWS in the lead up to the release of Amazon SageMaker Feature Store, and we are excited by the prospect of a fully managed feature store so that we no longer have to maintain multiple feature repositories across our organization. Our data scientists will be able to use existing features from a central store and drive both standardization and reuse of features across teams and models.”

The Climate Corporation (Climate) is a subsidiary of Bayer and the industry leader in bringing digital innovation to farmers around the world by increasing their productivity using digital tools. Climate is focused on helping farmers understand their fields in ways that have never been possible before and derive impactful recommendations from agricultural data. “At Climate, we believe in providing the world’s farmers with accurate information to make data-driven decisions and maximize their return on every acre,” said Daniel McCaffrey, Vice President, Data and Analytics at Climate. “To achieve this, we have invested in technologies such as machine learning tools to build models using measurable entities known as features, such as yield for a grower’s field. With Amazon SageMaker Feature Store, we can accelerate the development of machine learning models with a central feature store to access and reuse features across multiple teams easily. Amazon SageMaker Feature Store makes it easy to access features in real-time using the online store or run features on a schedule using the offline store for different use cases. With Amazon SageMaker Feature Store, we can develop machine learning models faster.”

DeNA is a leading provider of mobile and online services, including games, e-commerce, and entertainment content distribution in Japan. “At DeNA, our mission is to deliver impact and delight customers using artificial intelligence and machine learning. Providing value-based services is our primary goal, and we want to ensure our businesses and services are ready to achieve that goal,” said Kenshin Yamada, General Manager, AI Systems at DeNA. “One of our key initiatives is to expand our capabilities in artificial intelligence and machine learning. Amazon SageMaker has helped us in our path to implement machine learning in many of our businesses by providing extensive capabilities to train and deploy accurate models. One of the areas we want to focus is on data preparation and make it easy for our engineering teams. With Amazon SageMaker Data Wrangler, we believe we can hit the ground running with a rich collection of transformation tools without the need to write additional code. As we become more efficient with data preparation, we also want to ensure our teams across our diverse businesses do not repeat or duplicate work in building similar features for our applications. We would like to discover and reuse features across the organization, and Amazon SageMaker Feature Store helps us with an easy and efficient way to reuse features for different applications. Amazon SageMaker Feature Store also helps us in maintaining standard feature definitions and helps us with a consistent methodology as we train models and deploy them to production. With these new capabilities of Amazon

AMAZON CONFIDENTIAL

SageMaker, we can train and deploy machine learning models faster, keeping us on our path to delight our customers with the best services.”

iFood is an online food delivery portal and one of the largest food delivery companies in Latin America offering quality services to consumers. “At iFood, we strive to delight our customers through our services using technology such as machine learning,” said Sandor Caetano, Chief Data Scientist at iFood. “We have been using Amazon SageMaker for our machine learning models to build high-quality applications throughout our business. Building a complete and seamless workflow to develop, train, and deploy models has been a critical part of our journey to scale machine learning. Amazon SageMaker Pipelines helps us to quickly build multiple scalable automated machine learning workflows, and makes it easy to deploy and manage our models effectively. Amazon SageMaker Pipelines enables us to be more efficient with our development cycle. We continue to emphasize our leadership in using artificial intelligence and machine learning to deliver superior customer service and efficiency with all these new capabilities of Amazon SageMaker.”

Since naming AWS as its official technology provider in January 2020, the DFL Deutsche Fußball Liga – organizer and marketer of Germany’s top soccer leagues Bundesliga and Bundesliga 2 – and AWS have embarked on a journey together to bring advanced sports analytics, by way of Bundesliga Match Facts powered by AWS, to life for fans and TV broadcasters around the globe. “Amazon SageMaker Clarify seamlessly integrates with the rest of the Bundesliga Match Facts digital platform and is a key part of our long-term strategy of standardizing our machine learning workflows on Amazon SageMaker,” said Andreas Heyden, Executive Vice President of Digital Innovations for the DFL Group. “By using AWS’s innovative technologies, such as machine learning, to deliver more in-depth insights and provide fans a better understanding of the split-second decisions made on the pitch, Bundesliga Match Facts enables viewers to gain deeper insights into the key decisions in each match.”

CS DISCO is a SaaS provider that offers solutions to automate and simplify a variety of legal tasks, including discovery. “At CS DISCO we have revolutionized the way legal evidence is reviewed with our DISCO AI platform for ediscovery,” said Alan Lockett, Principal Data Scientist at CS DISCO. “We are always looking to improve how quickly our advanced deep learning models train. Our team has worked with the Amazon SageMaker team at AWS and believes that technologies such as distributed training and others can help accelerate our AI use cases.”

Turbine is a simulation-driven drug discovery company delivering targeted cancer therapies to patients. “We use machine learning to train our in silico human cell model, called Simulated Cell™, based on a proprietary network architecture. By accurately predicting various interventions on the molecular level, Simulated Cell™ helps us to discover new cancer drugs and find combination partners for existing therapies,” said Kristóf Szalay, CTO at Turbine. “Training of our simulation is something we continuously iterate on, but on a single machine each training takes days, hindering our ability to iterate on new ideas quickly. We are very excited about Distributed Training on Amazon SageMaker, which we are expecting to decrease our training times by 90% and to help us focus on our main task: to write a best-of-the-breed codebase for the cell model training. Amazon SageMaker ultimately allows us to become more effective in our primary mission: to identify and develop novel cancer drugs for patients.”

Latent Space is a startup focused on building the world's first fully AI-rendered 3D game engine. “At Latent Space, we're building a neural rendered game engine where anyone can create at the speed of thought. Driven by advances in language modelling, we're working to incorporate semantic understanding of both text and images to determine what to generate,” said Sara Jane, Co-founder and

AMAZON CONFIDENTIAL

Chief Science Officer at Latent Space. “Our current focus is on utilizing information retrieval to augment large-scale model training, for which we have sophisticated machine learning pipelines. This setup presents a challenge on top of distributed training since there are multiple data sources and models being trained at the same time. As such, we're leveraging the new distributed training capabilities in Amazon SageMaker to efficiently scale training for large generative models.”

Lenovo is the world's largest maker of personal computers. Lenovo designs and manufactures devices such as laptops, tablets, smartphones and a variety of smart IoT devices. “At Lenovo, we're more than a hardware provider and are committed to being a trusted partner in transforming customers' device experience and delivering on their business goals. Lenovo Device Intelligence is a great example of how we're doing this with the power of machine learning, enhanced by Amazon SageMaker,” said Igor Bergman, Lenovo Vice President, Cloud & Software of PCs & Smart Devices. “With Lenovo Device Intelligence, IT administrators can proactively diagnose PC issues and help predict potential system failures before they occur, helping to decrease downtime and increase employee productivity. By incorporating Amazon SageMaker Neo, we've already seen a substantial improvement in the execution of our on-device predictive models – an encouraging sign for the new Amazon SageMaker Edge Manager that will be added in the coming weeks. The new Amazon SageMaker Edge Manager will help eliminate the manual effort required to optimize, monitor, and continuously improve the models after deployment. With it, we expect our models will run faster and consume less memory than with other comparable machine-learning platforms. As we extend AI to new applications across the Lenovo services portfolio, we will continue to require a high-performance pipeline that is flexible and scalable both in the cloud and on millions of edge devices. That's why we selected the Amazon SageMaker platform. With its rich edge-to-cloud and CI/CD workflow capabilities, we can effectively bring our machine learning models to any device workflow for much higher productivity.”

Basler AG is a leading manufacturer of high-quality digital cameras and accessories for industry, medicine, transportation and a variety of other markets. “Basler AG delivers intelligent computer vision solutions in a variety of industries, including manufacturing, medical, and retail applications. We are excited to extend our software offering with new features made possible by Amazon SageMaker Edge Manager,” said Mark Hebbel, Head of Software Solutions at Basler. “To ensure our machine learning solutions are performant and reliable, we need a scalable edge to cloud MLOps tool that allows us to continuously monitor, maintain, and improve machine learning models on edge devices. Amazon SageMaker Edge Manager allows us to automatically sample data at the edge, send it securely to the cloud, and monitor the quality of each model on each device continuously after deployment. This enables us to remotely monitor, improve, and update the models on our edge devices around the world and at the same time saves us and our customers' time and costs.”

Mission Automate handcrafts software solutions on behalf of their global customers. “We constantly look for new solutions that can provide the best quality software to our customers, but as a small organization, we don't have the same ability to specialize in silos like other organizations,” said Alex Panait, CEO at Mission Automate. “Amazon SageMaker JumpStart now provides us a way to get started with machine learning faster, including new techniques that we can use in our own workflows to increase our service offerings and reduce costs. The option to select machine learning models and algorithms from popular model zoos allows us to quickly train customized machine learning models, which helps our customers get to market faster. Thanks to Amazon SageMaker JumpStart, we are able to launch machine learning solutions within days to fulfill machine learning prediction needs faster and more reliably.”

MyCase offers a powerful legal practice management software that helps law firms run efficiently from anywhere, provide an exceptional client experience, and easily track firm performance. “We have several business and product elements that can be improved with machine learning,” said Gus Nguyen, Software Engineer at MyCase. “Amazon SageMaker JumpStart allows us to launch end-to-end solutions with one click and access a collection of notebooks to help us more deeply understand customers and use predictions to better serve their needs. Thanks to Amazon SageMaker JumpStart, we can have better starting points which makes it so that we can deploy a machine learning solution for our own use cases in four to six weeks instead of three to four months.”

About Amazon Web Services

For 14 years, Amazon Web Services has been the world’s most comprehensive and broadly adopted cloud platform. AWS offers over 175 fully featured services for compute, storage, databases, networking, analytics, robotics, machine learning and artificial intelligence (AI), Internet of Things (IoT), mobile, security, hybrid, virtual and augmented reality (VR and AR), media, and application development, deployment, and management from 77 Availability Zones (AZs) within 24 geographic regions, with announced plans for 18 more Availability Zones and six more AWS Regions in Australia, India, Indonesia, Japan, Spain, and Switzerland. Millions of customers—including the fastest-growing startups, largest enterprises, and leading government agencies—trust AWS to power their infrastructure, become more agile, and lower costs. To learn more about AWS, visit aws.amazon.com.

About Amazon

Amazon is guided by four principles: customer obsession rather than competitor focus, passion for invention, commitment to operational excellence, and long-term thinking. Customer reviews, 1-Click shopping, personalized recommendations, Prime, Fulfillment by Amazon, AWS, Kindle Direct Publishing, Kindle, Fire tablets, Fire TV, Amazon Echo, and Alexa are some of the products and services pioneered by Amazon. For more information, visit amazon.com/about and follow [@AmazonNews](https://twitter.com/AmazonNews).