



Tackling our world's hardest problems with machine learning



Table of contents

- Tackling our world’s hardest problems with machine learning . . . 4
- Building a sustainable future 5
- Creating financial opportunities for more 9
- Improving health outcomes 12
- Changing lives for the better 16





A letter from Swami Sivasubramanian

Machine learning (ML) has graduated from the realm of science fiction to become a core, transformative technology for organizations across industries and categories. The unique potential and power of ML are sparking significant innovation, powering the ideas that are improving lives and protecting our planet right now. With ML, organizations are making inroads toward protecting and supporting our veterans, finding homes for homeless people, understanding climate change, improving health outcomes, and more. But this is just the beginning.

The technology is ripe, and it now has the ability to provide new and significant solutions for some of the world's biggest challenges.

More than a hundred thousand companies and organizations worldwide have turned to Amazon Web Services (AWS) for ML—to help track disease outbreaks worldwide, find new ways to treat cancer, and more.

However, access to ML, a new technology to so many of these organizations, can often come with a skills and technology deficit. That's where AWS steps in, partnering with innovators to bridge the gap and bring pioneering solutions that can help tackle our most urgent and important challenges.

Swami Sivasubramanian, VP, Database, Analytics & ML, AWS

Tackling our world's hardest problems with machine learning

AWS puts ML in the hands of every business analyst, data engineer, developer, data scientist, and expert practitioner. And every day, these individuals leverage and apply ML in new ways for the benefit of society. From decreasing carbon emissions to improving safety in the workplace, the methods by which today's leaders apply ML reveal a tremendous opportunity to improve lives and the world in which we live.

What are artificial intelligence and machine learning?

You've probably heard artificial intelligence (AI) and ML described in several ways, so let's take a step back and review their exact definitions:

AI is a way to describe any system that can replicate tasks that previously required human intelligence. Almost always, this is related to some complex decision-making task where human judgment would normally be required. Most use cases for AI are looking for a probabilistic outcome—making predictions, classifications, or decisions with a high degree of certainty and in ways that are similar to human judgment.

Almost all AI systems today are created using ML. ML uses large amounts of data to create and validate decision logic. This is known as a model. The AI system feeds input data into that model, and then the model outputs human-like predictions or classifications. Essentially, ML is the underlying technology that powers intelligent systems.



Building a sustainable future

As companies and organizations across the globe make their commitments to sustainability, ML will play its part in helping to ensure success. Amazon is committed to building a sustainable future for our customers and the planet, from creating sustainability in the cloud to the ambitious goals of The Climate Pledge.



The University of Oxford analyzes cloud patterns to better understand climate change

As a recipient of the Amazon Research Award, the **University of Oxford** and its climate scientists are working to unearth new ways to combat climate change. The award program offers unrestricted funds and AWS Promotional Credits to support research at academic institutions and nonprofit organizations in areas that align with our mission to advance the science of customer obsession.

ML is an essential tool for climate change research because climate science is such a data-intensive field. Climate models are enormous, requiring supercomputers to run them, and analysis requires a huge amount of Earth observation data. As data continues to grow, along with its complexity, it becomes impossible to explore all avenues of research manually. This is one of many ways the University of Oxford and AWS continue to work together, including a collaboration to fund a test bed of new research in AI and data science across the university.

The Climate Processes Research Group in the Department of Physics at Oxford hopes that by studying the effects of aerosol pollution on clouds, they'll be able to break new ground in global warming research, leveraging tools such as AWS Deep Learning AMIs running on Amazon Elastic Compute Cloud (Amazon EC2). Clouds reflect sunlight back to space, acting like an umbrella that cools the Earth. So, even small changes in clouds in response to global warming or air pollution could have a big impact on environmental health and serve to accelerate or dampen the greenhouse effect. ML models can track these changes to understand why clouds change, which could be the key to addressing global warming.

Now, Oxford scientists can analyze satellite data covering the entire Earth multiple times a day, providing countless images of aerosol-impacted clouds, which they're able to process in the AWS Cloud, thanks to the AWS IMAGINE Grant program. Such scalable ML techniques allow experts to make rapid progress in an area where researchers previously spent months of their time manually identifying features in fairly limited datasets.

Amazon eliminates 915,000 tons of packaging with machine learning

Amazon sells hundreds of millions of different products and ships billions of items a year. To ship with minimal packaging at maximum speed and ensure the customer's order arrives at their doorstep undamaged, the company must innovate at a large scale. The goal is to scale decision-making across the hundreds of millions of products that are shipped—not to automatically default to boxes but instead to identify items that can be shipped in flexible packaging options such as mailers (padded paper envelopes) or bags, which are more sustainable.

In practice, this means creating ML algorithms built on terabytes of product data, from product descriptions to customer feedback. Working closely with AWS Professional Services, these terabytes of data are cleaned, cataloged, and ready for mining. The ML algorithms then ingest that data to identify the best packaging with the least waste. Using Amazon SageMaker, the packaging team can analyze hundreds of millions of products, billions of customer shipments, and multiple channels of customer feedback, providing actionable insights in real time.

Some of the most impactful ML models identify products that don't need any packaging at all—such as diapers. Others are designed to determine product categories such as toys and differentiate collectibles, where the type of the original packaging is important to the customer. These products are shipped with the protection of an Amazon shipping box, whereas, for example, a stuffed animal can be shipped safely in a bag.

ML has accelerated changes in the packaging mix significantly, reducing the use of boxes from 69 percent to 42 percent, and it's still Day 1!

ML is an essential tool to drive innovation to support sustainability. It empowers the potential for real business change and innovation, including efficiency improvements, renewable energy, materials reductions, and other carbon emission elimination strategies.

Flexible packaging is 75% lighter than a similarly sized box and will conform around a product, taking up to 40% less space than a box during shipping—which means a lot fewer trucks on the road.¹

This use case is critical, considering that leakage accounts for an estimated 1.7 trillion gallons of lost water—almost 30% of all treated water.

Mueller Water improves leak detection by 40%

Utilities have not traditionally been early adopters of digital technologies, but the climate crisis and the COVID-19 pandemic have set the stage for a new operating environment. **Mueller Water Products, Inc.**, a global maker of water infrastructure and solutions, leverages AWS AI and ML to provide utility customers with critical insights into various aspects of water, including leak and anomaly detection.

In the US alone, leakage accounts for an estimated 1.7 trillion gallons of water—almost 30 percent of all treated water. It's a major problem for both the environment and the funding and sustainability of utility companies, which see no revenue from lost water. Leakage translates to a high carbon footprint from the water industry overall.

Mueller answered the challenge by consolidating its cloud and on-premises systems into a single comprehensive cloud system and building a data lake on Amazon Simple Storage Service (Amazon S3). Using SageMaker, the company was able to increase anomaly detection, which allowed them to build, train, and deploy ML models quickly. As a result, Mueller can identify acoustic signatures that predict the probability of a leak with 90 percent accuracy—a roughly 40 percent improvement.

Mueller now has a committed software-as-a-service (SaaS) strategy in place to provide customers with critical insights into water quality evaluation, pressure monitoring, and other aspects of water distribution. By empowering customers through increased problem-solving potential, Mueller Water Products and AWS are helping lead the water market toward a future of innovation, growth, and sustainability.

Creating financial opportunities for more

ML is a valuable tool that can unlock new business potential for organizations around the world, but it also has the power to serve and reach individuals and provide them with new financial opportunities. From small-business owners to underserved consumers looking for access to credit-based economies, ML creates new opportunities.



Tala helps millions unlock financial potential with machine learning

About 2.5 billion people around the world are underserved by traditional financial institutions. According to a One World Bank estimate, approximately 68 percent of adults have no credit data and, therefore, no credit score. And in most of these areas, national IDs do not exist or are inconsistently enforced throughout the country. As a result, about 85 percent of the world's transactions are cash-based, limiting or preventing their access to the global economy and creating a frustrating, vicious cycle.

Without financial data, it's hard for financial services companies to identify, reach, and transact with this population. But financial services company **Tala** recognized that they could find this population of unbanked consumers and serve them with a modern credit infrastructure built from scratch using AWS-powered ML. In effect, Tala would be the first company to offer loans to the population of previously unbanked consumers, unlocking their financial potential in the marketplace and giving them the ability to participate as both buyers and business owners and contribute to their community.

ML solutions often hinge on human understanding. In the case of Tala, sitting down with real customers provided the foundational insights, and ML provided the ability to scale. With SageMaker, Tala has adopted an end-to-end automated process, leveraging ML at each step of the decision process to approve or decline a loan, determine loan amounts, or make decisions concerning collections and acquisitions. The entire flow uses ML to automate as much as possible while reducing costs, enabling Tala to serve more customers at a better price.

To date, Tala's innovative ML solution has offered more than 5 million customers over \$2 billion in loans.

ADP fosters a more diverse, inclusive workforce with machine learning and analytics

Although organizations express a commitment to closing the gender pay gap, many struggle to make significant improvements. By analyzing data in areas such as recruiting, payroll, promotions, and salary, organizations can increase transparency and accountability around pay equity. But sifting through large volumes of data in multiple formats to find insights is time-consuming and error-prone.

With the AWS Cloud tools, organizations can gather and combine data from multiple sources and apply analytics and ML to uncover diversity, equity, and inclusion (DE&I) inequities, set targets for improvements, and track progress across teams. **ADP**, the world's largest provider of payroll and human capital management solutions, moved its data to the cloud in 2019. Using AWS ML and advanced analytics, ADP built a collection of tools, data, and best practices to help organizations foster a more diverse, equitable, and inclusive workplace.

One result of their cloud migration was ADP's Diversity, Equity, and Inclusion Dashboard. It enables companies to examine their workforce by organization, department, and job level along several demographics, ranging from gender, ethnicity, and race to veteran status, age, and disability. ADP customers can also compare their metrics against other businesses and get recommendations on talent investments to remain competitive.

ADP also used AWS solutions to build the ADP Pay Equity Storyboard, which combines analytics and benchmarking to help employers understand potential pay gaps. The Pay Equity Storyboard further provides action plans for organizations to close gaps and benchmark data, helping to ensure that pay is both equitable and market competitive.

Several months after launching these tools, about **two-thirds of ADP customers showed improvement in pay equity**, and over 210,000 people saw a substantial increase in their pay rate.

Using AWS solutions, **two-thirds** of ADP customers reported improvement in pay equity.

Improving health outcomes

The healthcare industry has long been an early adopter of advances in technology. Today, ML is a key tool in revolutionizing the expansive field of health and medicine. From developing new procedures to collecting and analyzing data, ML has the unique ability to power new insights and predictions to improve patient care and population health.

Inspire connects patients to caregivers and critical resources

Inspire, the leading healthcare social network, has a mission to connect patients to their caregivers and other patients with online tools, resources, and condition-specific support groups. It also connects pharmaceutical companies and other medical institutions conducting clinical-trial research to health-outcome studies on patients. Partnering with over 100 nonprofit patient-advocacy organizations, Inspire creates and manages support communities for more than one million patients and caregivers, representing more than 3,600 health conditions.

But building on its success and capitalizing on its growth potential required the company to overcome the scaling challenges posed by its legacy on-premises infrastructure. The challenge called for a scalable cloud infrastructure that could keep up with growth and an ML solution to improve its content recommendation engine for patients and caregivers.

Using managed solutions from AWS, Inspire gained numerous advantages in the cloud, including faster iteration, greater flexibility, and multi-region availability. It chose Amazon Aurora as part of its new infrastructure and SageMaker, a fully managed service that provides every developer and data scientist with the ability to build, train, and deploy ML models quickly.

Leveraging these tools, the company saw a substantial increase in user engagement across all channels and was enabled to accurately connect patients and caregivers with more personalized content and resources, including rare-disease information and treatment pathways. Inspire was also able to streamline the process by which pharmaceutical companies conducting clinical trials or medical research could connect with relevant patient data—a critical milestone in the development of lifesaving therapies.

“AWS has given the whole software side of the house the ability to do what the business side needs done quickly and simply. To bring patients who have rare diseases together in a single place where they can share their experiences, regardless of where they live or what language they speak—has a profound impact.”

Anthony Sheetz, VP of Information Security & Infrastructure, Inspire

Cortica revolutionizes healthcare for autistic children

Cortica is on a mission to revolutionize healthcare for children with autism and other neurodevelopmental differences and to fix the fragmented journey families typically navigate while seeking diagnoses and therapies for their children. With multiple vendors providing systems and data solutions, Cortica found itself in a situation all too common in the healthcare industry: massive amounts of data with multiple formats, making it a complex challenge to match patients from system to system. The company turned to AWS to simplify the process.

Cortica's team imported all its data into an Amazon S3 data lake. But that was only part of the solution. Because healthcare data is driven by medical terminologies, the Cortica team needed a solution that could help unify data from multiple healthcare fields to present a clear patient journey through the different specialties Cortica offers. To exploit this deeper value, Cortica chose Amazon HealthLake to standardize data and scale insights.

The team was able to quickly load a basic set of resources into HealthLake, enabling them to develop a proof of concept for understanding their patient's journey from the perspective of behavior therapy goals and medical comorbidities. They were able to quickly see the length of time it took for patients to attain their goals across their entire patient population. With HealthLake, the Cortica team spends more time analyzing and understanding data patterns rather than sourcing and formatting data and joining it into a usable state.

Data analysis that took months to complete can now be completed in days and even hours. With AWS tools, Cortica can enhance analysis by adding layers of richness to the data and providing different views of the same analysis. Furthermore, analysis that required chart abstraction can now be done through automated data pipelines, allowing Cortica to process hundreds or thousands of documents and derive important insights from notes that were previously available to only a few clinicians.

HealthLake allows Cortica to process large amounts of data and derive insights in hours instead of days.



Cancer Commons treats cancer with the power of the cloud

Cancer Commons helps patients with advanced cancer access personalized treatment options and seeks to overcome a core problem in cancer research: The best cancer treatments involve combinations of therapies, with far too many plausible regimens to test in clinical trials.

Facing heavy computing and data requirements, Cancer Commons moved from multiple cloud service providers to AWS compute infrastructure and services, including Amazon Comprehend Medical, to enable physicians and patients to leverage the collective knowledge of the world's top institutions.

"Now, almost every part of our process is powered by AWS in some way," said Erika Vial-Monteverdi, executive director at Cancer Commons.

When necessary, Cancer Commons convenes virtual tumor boards (VTBs), developed on AWS HIPAA-compliant architecture, where experts help patients refine their options based on their individual medical history and preferences. The discussion that happens in the VTBs is mined using natural language parsing and other AI techniques, and a recommended treatment plan is created from that information.

"Whether it's raw computing, specialty services such as NLP, specialty computing such as GPU, or optimization services such as Amazon SageMaker, AWS has the tools and services we need," Erika said. "Thanks to AWS, we are helping accelerate cancer research and save lives."

Changing lives for the better

Some of the toughest challenges societies face are managed by nonprofits, government programs, and community-focused organizations. Technology has enabled these organizations to build networks and collaborate on modern solutions to today's biggest challenges. From helping veterans in need to addressing homelessness, ML has immense power to improve the lives of individuals and our communities in need.



PATH addresses homelessness by providing help faster

PATH, a Los Angeles-based organization founded to address the ever-increasing issue of homelessness, applied for an AWS IMAGINE Grant to develop a way to shorten the time it takes to match homeless individuals and families with homes of their own. With the grant and support from the AWS team, the organization developed LeaseUp! to connect clients with the best possible housing for their needs.

Amazon Personalize captures relevant information about available units of housing so case managers can recommend the best housing option to their clients in real time. By integrating this technology, the organization has been able to match over 600 individuals experiencing homelessness with housing—and reduce the time it takes to do so. Timing in these situations is often critical; a person who is ready to come in and get help one day may not return the next.

LeaseUp! aims to add 2,000 new units to its database over the next year to help even more people make it home. Bringing more existing apartments onto the platform and working more seamlessly with the landlords to list rental units are important steps in not just addressing homelessness but also ending it.

Talkspace increases access to mental healthcare with machine learning

Approximately 20 percent of Americans have a mental health condition, and over half of these individuals do not receive treatment for their conditions. **Talkspace**, an online therapy company, offers therapy sessions to users with the touch of a button on their devices. But matching patients with the exact care they need requires another level of sophistication.

Talkspace uses SageMaker and Amazon SageMaker Ground Truth to help improve the therapy it promises to provide. SageMaker helps match patients to therapists across multiple dimensions. And SageMaker Ground Truth, a fully managed data-labeling service, makes it easy to build highly accurate training datasets for ML, helping Talkspace prevent early attrition through specialized models that continually improve the matching affinity score as clients engage with their providers.

With AWS ML, Talkspace can develop a diagnostic profile of each patient, identify behavioral patterns and potential harm risks, and provide helpful insights. It can also send push notifications to therapists in real time if elevated risk is detected.

Additional ML-assisted features provide therapists with recommended tips and actions that help them retain patients and evaluate treatment courses. Talkspace also implemented autoprogess notes to enable therapists to quickly create session reports and easily visualize a patient's progress over a specific period. These time-saving benefits led to a 20 percent increase in the number of notes written by therapists.

"Pairing each client with a mental health professional who meets their needs is important for the success of therapy," says Gil Margolin, chief technology officer at Talkspace. "By standardizing our machine learning [ML] workloads on AWS, we're able to understand our clients better, increase our level of service, and provide time-saving tools for therapists to improve client outcomes."

"AWS helped us focus on what really counts: delivering quality therapy at scale to people who deserve but are unable to receive that level of care."

Gil Margolin, CTO, Talkspace

Getting started with machine learning

Organizations are constantly working on innovative techniques to solve the most important issues the world faces today, making a profound and significant impact. We've seen what's possible, and AWS is committed to helping our customers bring their ML solutions to life.

[Learn more >](#)

