6 steps to machine learning success

The path toward leveraging the full power of machine learning

INTRODUCTION

Forging ahead

When deployed with the right strategy, artificial intelligence (AI) and machine learning (ML) can increase agility, streamline processes, boost revenue by creating new products and improving existing ones, and enable better, faster decision-making.

There's no doubt that AI and ML can help companies achieve more—by 2025, global spending on AI will reach $204 billion.[1] It's also clear that adopters continue to have confidence in AI's ability to drive value and competitive advantage.

While ML has been around for decades, its accessibility as a tool to transform businesses is relatively new. Additionally, the lack of a singular proven path to ML success is keeping some businesses waiting on the sidelines, unsure of how to take the next (or even the first) step on the journey.

It's time for organizations to overcome ML barriers, stop playing catch-up, and forge ahead with confidence. This eBook outlines a proven path—from the first step to measuring results—with insights from Amazon ML best practices and its experience helping thousands of customers realize their own initiatives.

75%

of enterprises will start operationalizing AI by the end of 2024[2], and 66% view AI as critical to success[3].

[1] "Investment in Artificial Intelligence Solutions Will Accelerate as Businesses Seek Insights, Efficiency, and Innovation, According to a New IDC Spending Guide," IDC, 2021
[2] "Gartner Identifies Top 10 Data and Analytics Technology Trends for 2020," Gartner, 2020
[3] "Becoming an AI-fueled Organization," Deloitte, 2021

WHY ML

What are artificial intelligence and machine learning?

You've probably heard AI and ML described in a number of ways, so let's take a step back and review what each term means. AI is a way to describe any system that can replicate tasks that previously required human intelligence.

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, which 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.

Why machine learning?

Before diving into the steps of the ML journey, let's explore why businesses should go on that journey in the first place. After all, even with the guidance in this eBook, completing the steps outlined here will require continued investment and dedication. A business will need to regularly remind itself what it's working toward—staying focused on the precise business benefits that can be unlocked by fully leveraging ML technology.

Businesses are already realizing the impact of:

1

Optimizing business with new efficiencies

ML can be used to create greater efficiency through sophisticated demand planning and forecasting models. While this is true in almost every industry, manufacturing provides some specific evidence. During the COVID-19 pandemic, **Foxconn**, the world's largest electronics manufacturer and technology solutions provider, faced unprecedented volatility in customer demand, supplies, and capacity. The company collaborated with the **Amazon Machine Learning Solutions Lab** to develop a demand forecasting model using **Amazon Forecast** for its factory in Mexico to increase the accuracy of its net order forecasts by eight percent. According to a 2022 report from McKinsey, companies in telecommunications, electric power, natural gas, and healthcare have found that AI forecasting engines can automate up to 50 percent of workforce management tasks, leading to cost reductions of 10–15 percent.[4]

2

Making smarter, faster decisions

Informed by data and analytics sources that grow smarter through ML, a business and its workforce can make more informed, faster decisions that allow it to act on opportunities sooner and achieve better results. The ML team at **Zendesk** is responsible for enhancing customer experience teams to better serve customers. By combining the power of data and people to automate manual work, Zendesk delivers intelligent products that make its customers more productive.[5]

[4] Amar, J., et al., "AI-Driven Operations Forecasting in Data-Light Environments," McKinsey & Company, 2022
[5] Jaffry, S., et al., "How to Scale Machine Learning Inference for Multi-Tenant SaaS Use Cases," AWS Machine Learning Blog, 2022

3

Increasing customer engagement by providing personalized experiences

Warner Bros. Discovery, Inc. has a wide range of content that appeals to a broad base of customer audiences.[6] The company wanted to help customers find curated content that matched their specific interests and turned to Amazon Web Services (AWS) and **Amazon Personalize** to enable tailored content suggestions for its discovery+ streaming platform.[7] Instead of offering viewers generic content, Warner Bros. Discovery, Inc. now uses the ML-powered recommendations of Amazon Personalize to customize the viewer experience and improve the overall customer journey.

4

Adding new capabilities to existing products

ML can enrich existing products, improving customer engagement and attracting new users through deeper experiences. For example, **Livongo** is a platform and mobile app that works with smart devices such as connected blood glucose monitors to help people manage medical issues. It uses ML to translate data from blood glucose readings, physical activity, and meal logs, in addition to smartphone data and other important data, into timely and actionable "health nudges." These personalized messages around diet, exercise, medications, and more—delivered in real time to members on their connected devices—help them avoid complications that could land them in the hospital and allow the system and its users to save money.[8]

5

Inventing net-new products

In its drive to make healthcare technology better serve individuals, **Cambia** used AWS—including ML technology—to develop Journi, a digital all-in-one health solution guided by data-driven intelligence and human expertise. Journi helps health plan members and their families make the most of their health benefits.[9]

6

Creating a competitive advantage

ML isn't a niche choice anymore—it's a necessity to stay relevant and competitive, regardless of an organization's field and domain. **OnPoint,** a subsidiary of Koch Engineered Solutions, recently collaborated with AWS to co-develop an intelligent process optimization platform (IPOP). Designed to enable OnPoint's solutions, the platform utilizes the cloud and edge digital infrastructure to provide advanced ML and industrial computing capabilities for OnPoint's connected ecosystem. This enables industrial operators and their preferred partners to integrate and act upon the best available knowledge—anywhere, anytime, at scale.[10]

Now that we've outlined the *why* of ML, it's time to explore the *how*. The next sections will demonstrate the steps of the ML journey using AWS best practices and those of customers to exemplify the necessary changes that must take place in order to successfully implement, deploy, and scale ML solutions.

[6] "Warner Bros. Discovery, Inc. Is Enhancing the Customer Experience with Amazon Personalize," Amazon Web Services, 2021, YouTube video
[7] "Discovery Enhances Customer Experience Using Amazon Personalize," AWS Case Study, 2021
[8] "Data Empowers Patients and Providers," Wall Street Journal Article Commissioned by AWS, 2020
[9] "Cambia Health Solutions Builds Integrated, Human-Centered Healthcare Solution on AWS," AWS Case Study, 2021
[10] Robinson, M., et al., "Acoustic Anomaly Detection Using Amazon Lookout for Equipment," AWS Machine Learning Blog, 2021

The machine learning journey

The ML journey is not always a straightforward path. Achieving success with ML requires more than just great technology—it also means ensuring the organization is aligned to the right goals. Identifying and reaching those goals will necessitate broad changes in processes, management, and culture. The next sections will explore how organizations can overcome common challenges that often impede progress and take the right steps to implement ML in efficient, sustainable ways.

STEP 1

Championing a machine learning culture

Unlocking the full business potential of ML requires cultural changes in team organization, objectives, and outlook.

For ML to proliferate through an organization, both business and technical teams must work together and share the same priorities. To achieve this at the outset, the ML effort must be supported from the highest levels, with goals set by executive champions and an investment in the technology and processes that enable success. This includes, among others, a commitment to build more responsible AI from the beginning—working to identify and mitigate bias, improve explainability, and help keep data private and secure. By taking a people-centric approach, organizations can work to educate their workforce on responsible AI and build more diverse teams to bring more perspectives to the table to improve fairness.

It's important for management to take a wide-scale view while fostering ML initiatives. Executives must be firm in their goals but flexible in how the organization reaches them. Mistakes are sure to be made. But by staying focused on the long-term outlook and not allowing discouragement, organizations can glean wisdom from every error and apply those learnings to champion ML culture throughout the business.

Perhaps the largest cultural change organizations must undergo is utilizing the opportunity inherent in mistakes. ML is an iterative process that can only succeed through constant experimentation. Often, these experiments will result in failure. Only by learning from mistakes—and refusing to grind progress to a halt in the name of determining "what went wrong?"—can organizations consistently reach the breakthrough successes waiting on the other side.

How Amazon did it

Amazon has been using ML for over 20 years. But after about a decade of leveraging the technology, our leadership team asked every business leader in the organization—irrespective of whether they ran a research team, a fulfillment center, or an HR organization—to answer the question of how they planned to use ML in their businesses.

In most cases, "We don't plan to" wasn't an acceptable answer. This forced the leadership, domain experts, and technical experts to collaborate on ML initiatives and let nothing halt their progress—even in instances where tangible benefits were still years down the road.

In addition to hiring external data scientists, Amazon also created the Machine Learning University (MLU), which trained many of its developers to use ML more effectively. The company also built tools like **Amazon SageMaker**, which simplifies the process of creating models and lowers the barrier to entry, enabling ML technologies and initiatives to scale more effectively. Additionally, Amazon created a set of pre-built **AI services** that provides ready-made intelligence to address common business use cases for customers without them having to build their own models.

Today, there isn't a single department at Amazon that hasn't been touched by ML. Personalization technology from Amazon, which provides unique customer recommendations based on user preferences and habits, has significantly improved since its first model 20 years earlier and has been applied to other areas of the business.

Let's take a look at some examples of how Amazon is leveraging ML. The company uses ML throughout its fulfillment process and leverages a forecast system that can predict demand for nearly every product in its enormous inventory. These prediction models allow Amazon to better deliver on customer expectations of convenience, cost, and delivery speed.

"We forecast millions of products every single day across all of our Amazon sites worldwide," said Jenny Freshwater, director of forecasting at Amazon. "And without machine learning, we would not be able to produce those forecasts."

The examples go on and on. Amazon created Alexa, which provides customers with an entirely new way to interact with technology. The company developed groundbreaking technology with autonomous flight via Amazon Prime Air drones. And it uses robotics in its fulfillment centers to get packages to customers faster.

Achieving these successes required great investments in technology, research, and talent. But those investments would've gone to waste without the cultural changes that pushed them forward through many failures and unexpected challenges. Every organization must foster this same fault-tolerant culture of experimentation and innovation before the ML journey can truly begin.

Amazon is using ML to minimize packaging waste, reducing outbound packaging weight by **33%** and eliminating **915k tons** of packaging material worldwide.[11]

[11] "How Amazon Is Using Machine Learning to Eliminate 915,000 Tons of Packaging," VentureBeat, 2020

STEP 2

Reinventing data strategy

Success with ML is highly dependent on quality data. Without a proper data strategy in place, progress will slow to a crawl and hamper the effectiveness of the final model. Worse yet, if the model is informed by bad data, the results it generates may be misleading—or even flat-out wrong.

"[Machine learning models are] highly sensitive to data quality," Freshwater said. "So we learned—in many cases the hard way—that the time spent on getting data of high quality on the way in paid dividends in our models on the way out."

The right data strategy for ML should aim to break down silos, enabling IT teams to easily, quickly, and securely access and collect the data they need.

While modern data strategies take many forms, data lakes are becoming an increasingly popular core component of the most efficient models. Data lakes offer more agility and flexibility than traditional data management systems, allowing companies to manage multiple data types from a wide variety of sources and store that data—structured or unstructured—in a centralized repository.

Once stored, data can be analyzed by many types of analytics and ML services faster and more efficiently than with traditional, siloed approaches. Data lake architectures also enable multiple groups within the organization to benefit from analyzing a consistent pool of data that spans the entire business.

For help developing a more holistic data strategy that includes data lakes, visit **AWS for Data**.

How Georgia-Pacific did it

Hundreds of paper and tissue parent rolls are produced every day at **Georgia-Pacific** manufacturing facilities across North America. If tears or breaks occur frequently, it leads to paper machine and converting-line downtime that can cost Georgia-Pacific millions of dollars per year, per line.

Georgia-Pacific started by migrating 50 TB of structured and unstructured production data from a legacy database infrastructure to a cloud-based data lake. By layering AWS databases and analytics tools on all of that data, Georgia-Pacific was able to optimize key manufacturing processes to predict equipment failure 60–90 days in advance. By reducing paper tears and unplanned downtime, the company increased profits by millions of dollars for one production line.

Read the full story ›

"We are using AWS data-analysis technologies to predict…precisely how fast converting lines should run to avoid tearing. By reducing paper tears, we have increased profits by millions of dollars for one production line."

Steve Bakalar, VP – IT: Digital Transformation, Georgia-Pacific

STEP 3

Finding the right business problem to address

One mistake organizations often make in their ML journeys is employing discrete data scientists who work in silos to build ML models as proofs of concepts rather than solve real business problems. With no specific business problem to solve, IT executives will find it increasingly difficult to demonstrate the value of ML projects to their business executive counterparts. This can stall or even stop progress on ML initiatives.

Here are some important questions organizations should ask before embarking on an ML journey:

1. Is the project important enough to get attention and adoption?
2. Does the project solve a real business problem?
3. Are there places where the organization already has a lot of untapped data?
4. Does the project require ML?
5. Can the project be done by a single business?
6. Can the project eventually be operationalized?

"A first step is to identify a problem that is rich in data, but [one that] you haven't been able to solve through traditional methods," Freshwater said.

The **AI Use Case Explorer** is a business-outcome-centric search and navigation site that enables users to find the right AI use cases, discover relevant customer success stories, and mobilize their teams toward deployment.

In a successful ML journey, organizations create ML teams to address specific business problems. This requires including both technical and domain experts within these teams. While the technical experts will take on the brunt of model creation, they need the field knowledge of domain experts to define precise business challenges and identify the data most important to finding a solution.

This approach is also critical to change management. When technical and domain experts collaborate to create ML models, employees will feel more confident in making decisions based on the algorithm's logic.

Together, these teams should also work through how to measure success. "Make sure you…have very crisp and clear metrics as you embark on the machine learning journey," Freshwater said. "Many times, your models are taking over for something existing and you want to make sure that they're actually better and that you can measure it."

For more on measuring the success of ML initiatives, refer to **Step 6** in this eBook.

Some organizations have the talent in-house to identify the problems that would be best addressed by ML and to implement the appropriate pilot programs. Organizations can leverage the **Amazon Machine Learning Solutions Lab** to work backward from business challenges and go step-by-step through the process of creating ML projects to solve them.

How the NFL did it

For decades, the **NFL** has worked to provide deeper insights into its players and teams to satisfy both the need for improved player safety and the insatiable fan appetites for data and statistics.

To address these needs, the NFL worked with AWS to create the ML-powered NFL Next Gen Stats (NGS) program. Since data science and football are wildly different disciplines, the NFL wisely included both technical and domain experts in the creation of NGS, ensuring both groups could work hand in hand to identify the right data and develop statistics.

Leveraging RFID tags to track player movement, NGS provides real-time location data, speed, and acceleration for every player during every play on every inch of the field. By simulating different situations within a game environment, the NFL aims to foster a better understanding of how to treat and rehabilitate injuries in the near term and eventually predict and intervene to prevent injuries in the future.

NGS also uses ML models to calculate more than 20 different advanced statistics that are compelling to fans. One example is the Expected Rushing Yards metric, which is designed to show how many rushing yards a ball carrier is expected to gain on a given carry based on the relative location, speed, and direction of blockers and defenders.

Insights like Expected Rushing Yards wouldn't exist without the partnership between technical experts who can build and train the models that crunch the necessary data and domain experts who know what data to measure to create the most exciting statistics.

This partnership also helps to build acceptance for NGS, as broadcasters are more likely to cite advanced statistics that football experts—and, in some cases, the broadcasters themselves—had a hand in creating.

Read the full story ›

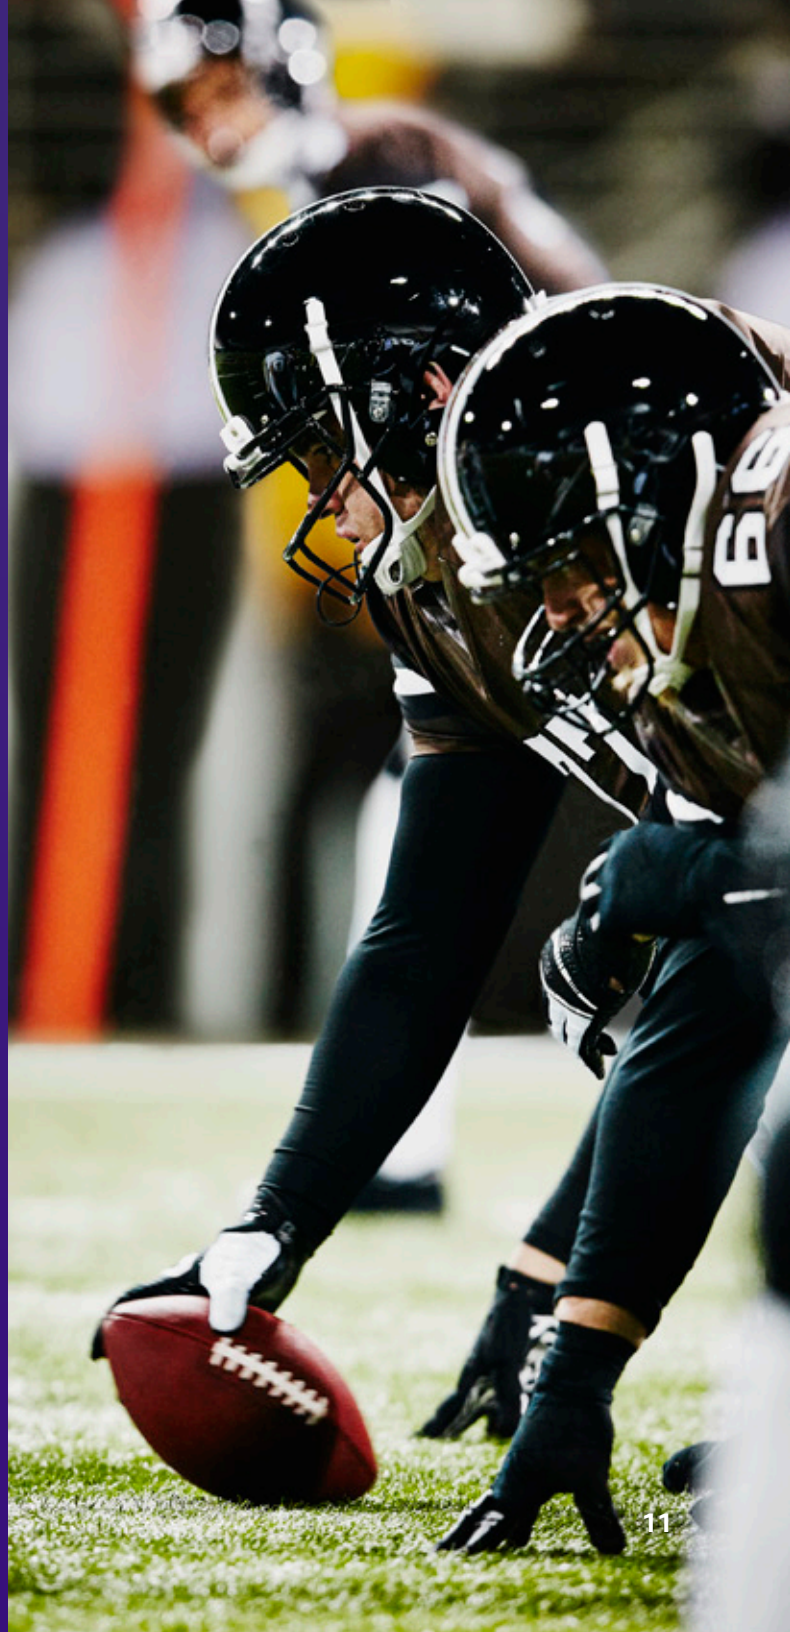

STEP 4

Upskilling teams

In parallel with creating a data strategy, an organization must focus on arming its teams with the right skills.

Organizations are growing increasingly aware of the ML skills gap—the expanding separation between technologies and the ability of internal IT specialists to take full advantage of them. A 2021 O'Reilly report, which surveyed more than 3,500 business leaders, found that a lack of skilled people and hiring difficulties topped the list of challenges in AI, with 19 percent of respondents citing it as a "significant" barrier.[12] Closing this gap for ML will require a combination of training and recruiting. The reality is that there aren't enough data scientists to lead the ML transformation that's coming. This requires organizations that want to leverage ML to invest in talent development and the right software tools for success.

While there's no one-size-fits-all solution to the ML skills gap, there are proven methods that can maximize the abilities of existing staff, reducing the need to make large investments in buying or borrowing pretrained expert talent.

These methods include:

Defining the skills gap: Before closing the skills gap, an organization must identify the precise differences between what it needs or wants its employees to do and what its employees currently can do.

Understanding how skills are mapped: Since ML initiatives are interdisciplinary efforts, an organization should map the skills needed across data scientists, data engineers, business analysts, application developers, statisticians, and other subject matter experts in the business.

Customizing training for specific needs: If an organization has existing training curriculums that could be useful, it should work to tailor those materials to the business' specific ML needs. Leaders should also investigate pretrained AI services that provide ready-made intelligence for business applications and workflows.

Evaluating no-code options: Many ML use cases are within the grasp of existing business analyst teams through the use of business intelligence tools that make predictions without the need to write code or have ML experience. These might be ideal for use cases like forecasting and churn prediction.

[12] Loukides, M., "AI Adoption in the Enterprise 2021," O'Reilly, 2021

In addition to training, organizations need to align teams to successfully tackle ML problems. They can achieve this by:

- **Promoting a culture of empowered teams:** ML project teams must be cross-functional, with the authority to execute individual objectives and the freedom to organically cross-pollinate with other teams as demands dictate and opportunities arise. To make this kind of teamwork possible, management will need to embrace new structures—letting go of the strictly hierarchical and departmentally siloed organizational models of the past.
- **Starting with a pilot team:** The business should establish a pilot team of engineers, IT and ML practitioners, and line-of-business leaders and task it with an ML project. "I'd recommend putting a couple of really smart people on trying to figure out what metrics you want to optimize for or predict…just start really small," Freshwater said.
- **Enabling organic transformation:** Once the pilot project is complete, the business can split up the team, add people to create new teams, and task them with new projects. This process continues, allowing knowledge to organically spread from veteran team members to new recruits and pollinate between teams.

By following this guidance, many organizations are finding that the people they currently have are actually the people they need to close their ML skills gaps. While some recruiting may still be required, investment in the right tools, processes, and management changes can do much of the work to upskill talent for ML success.

It's also important to ensure that business leaders are trained to understand ML, including what makes a good use case and how to speak the ML language. Many courses are available for business leaders, including **Machine Learning Essentials for Business and Technical Decision Makers**. In addition, line-of-business analysts supporting finance, marketing, operations, and sales can get started simply by uploading data to the cloud and using **Amazon SageMaker Canvas** to generate predictions. SageMaker Canvas is a visual, no-code service to generate ML predictions with just a few clicks.

How Morningstar did it

Investment research firm **Morningstar** uses ML to automate data collection processes and expand the number of funds it covers. The company does this by leveraging predictions from an ML model trained to emulate the fund evaluation process for its analysts.

To train its employees and accelerate ML application, Morningstar uses **AWS DeepRacer**—a tool that facilitates hands-on ML training through a fully autonomous 1/18th scale race car driven by reinforcement learning, a 3D racing simulator, and a global racing league. More than 445 Morningstar employees from multiple functions and eight countries—including 35 percent of its technology function—have been engaged in the DeepRacer League.[13]

Morningstar had dozens of ML projects in the pipeline for 2021. These included a reinforcement learning program that searches for patterns in regulatory filings and an algorithm that identifies and fixes broken links to the websites of financial institutions.

[13] "In the News: Morningstar Launches Global AWS DeepRacer Corporate Competition to Accelerate Application of Machine Learning," AWS for Industries, 2021

"Our DeepRacer challenge harnesses our employees' enthusiasm for machine learning and artificial intelligence. It provides hands-on training across the company and accelerates Morningstar's practical application of machine learning across our investing products, services, and processes. The response from our teams has gone well beyond my expectations, and it has been a fun way to unite our global teams, whether in technology or other functions."

James Rhodes, CTO, Morningstar

Read the full story ›

STEP 5

Scaling beyond pilot projects

After the first few successful pilots, organizations must take the next step on the journey: sustainably scaling ML across the business. This is both a technical and cultural challenge.

Achieving scalability requires organizations to make it easier for their developers to use ML. Building ML models at scale can be labor-intensive and complex, which can slow innovation.

Many organizations are solving scalability with **Amazon SageMaker**, an end-to-end solution that covers the entire ML workflow to prepare data and build, train, and deploy ML models. By using SageMaker, an organization can get its models into production faster and at a lower cost, enabling sustainable expansion of ML initiatives beyond pilot projects.

Other organizations are scaling through **AWS AI services**, a set of pretrained and managed services that can be used as building blocks to address common use cases, including personalizing recommendations, modernizing their contact centers, improving safety and security, and increasing customer engagement.

There are several ways companies approach the cultural shift necessary to scale ML. Some might find success by creating a center of excellence that rallies the community and continues to push for new initiatives. Or, like Amazon, organizations can make ML an integral part of yearly planning processes, continuously bringing domain and technical experts together to brainstorm and determine their next steps.

How Intuit did it

Using Amazon SageMaker, **Intuit** reduced ML deployment time by 90 percent, from six months to one week. By centralizing its ML initiatives, Intuit fosters innovation and deploys AI and ML techniques at speed and scale, achieving business value that goes beyond its products and services.

"AWS gives people within Intuit a common platform to share and collaborate with data in a secure environment," said Ashok Srivastava, senior vice president and chief data officer at Intuit. "For example, Amazon SageMaker gives us the platform and infrastructure we need to apply our sophisticated AI and machine learning technologies."

"In the future, we're really looking forward to leverage the full power of (Amazon) SageMaker…to allow us to fundamentally change the product experience for our users."

Nhung Ho, VP of AI, Intuit

Watch the video ›

STEP 6

Measuring the results

When measuring the results of ML efforts, the traditional "project ROI" viewpoint—where a project has a defined start and end point, a budget, and a return—is reductive and can be detrimental to the initiative's success. If the project doesn't generate a positive return within the given time frame, the business may lose interest and miss out on critical opportunities down the line.

Instead, executives and IT alike must measure ML efforts based on what success means for their businesses with regard to the processes being optimized. In addition, they must view ML efforts as long-term investments, acknowledging that a true "return" may not be realized for several years and after countless iterations.

When planning ML initiatives, it's better to view the process through the lenses of agility, competitive advantage, or risk tolerance rather than expected return. An organization will have greater success if it disregards the question of "What will be my return on investment in X months?" in favor of something more like "If we don't invest in this now, will we fall behind our competitors in X years when the technology matures?"

While traditional ROI metrics may not be the best approach, the business impact of ML initiatives can still be measured—it just requires a different outlook.

ML results can be measured through something resembling a "value tree," where the main trunk of the tree represents the traditional "revenue return" and branches extending from the trunk recognize the value of other business outcomes.

The specific branches of the value tree will depend on the organization, the industry, and the initiative. But they might be things like time saved through automated processes, new leads, markets, and opportunities identified, customer service improvements, or increases in upsells.

Measuring the success of ML through a more holistic and long-term model will keep organizations focused on the best outcomes for their business future.

How Lotte Mart did it

Compared to the prior recommendation approach Lotte Mart used, **Amazon Personalize** eliminated the need for tedious and complex manual data analysis and reduced development time by 50 percent. This saved time because Lotte Mart only had to provide pre-defined interactions, users, and item datasets. The engineering team could generate test results in half the time compared to the prior approach. With Amazon Personalize, Lotte Mart could cost-effectively recommend new products that were difficult to promote and drive demand through traditional methods.

"To enable us to be more customer centric, scale our reach, and increase uptake by users, we turned to Amazon Personalize to enable over 600,000 users of our M Coupon mobile app to save on their in store shopping experience. By using Amazon Personalize, we have seen a 5x increase in response to recommended products compared to our prior big data analytics solution resulting in increased revenue per month. In particular, Amazon Personalize has increased the number of products that the customer has never purchased before up to 40%. The new recommendation service powered by AWS is the first of a much broader roll-out of AI technologies across our organization."

Jaehyun Shin, Big Data Team Leader, Lotte Mart

Read the blog ›

Taking the next step with AWS

No matter where your organization is in its ML journey, you can take the next step with AWS solutions built on the most comprehensive cloud platform and optimized for ML with high-performance compute, security, and analytics. Featuring the world's broadest and deepest set of AI and ML services, over 100,000 customers are running their AI and ML workloads on AWS.

To help you identify the most relevant AI and ML use cases and solutions, we have built the **AI Use Case Explorer**, an easy-to-use tool that will suggest the best use cases based on your industry, function, and business priority, including:

Intelligent contact center: Enhance the contact center experience to personalize and streamline customer interactions, improve agent productivity, and increase business process efficiency by integrating ML into your contact center.

Chatbots and virtual assistants: Enable omnichannel, 24/7/365 communication with your customers in engaging ways through chatbots, voice assistants, and interactive information kiosks.

Personalization: Improve customer engagement and conversion by creating personalized experiences tailored to individual customer preferences and behaviors across channels.

Intelligent document processing: Instantly extract text and data from virtually any document, such as loan applications and medical forms, without manual effort.

Intelligent search: Boost business productivity and customer satisfaction by delivering accurate and useful information faster from different information sources across the organization using natural language questions.

Identity Verification: Decrease onboarding time, increase user convenience, and reduce fraud while lowering costs by using ML-powered user identity verification.

Content moderation: Create safe online environments, protect your brand, and minimize moderation costs by using ML to moderate large and complex volumes of user-generated content (UGC).

Media intelligence: Maximize the value of media content by adding ML to media workflows, such as search and discovery, content localization, compliance, monetization, and more.

Forecasting: Accurately predict demand forecasting and streamline supply-demand decisions to combine historical time-series data with additional variables such as product features, pricing, and holidays.

ML development modernization: Accelerate innovation while reducing costs by modernizing the ML development lifecycle through scalable infrastructure, integrated tooling, healthy practices for responsible ML use, tool choices accessible to developers of all skill levels, and efficient resource management.

Amazon SageMaker: Create new products and services with Amazon SageMaker. It enables your business analysts, data scientists, and ML engineers to prepare data and easily build, train, and deploy ML models—simplifying scalability across the entire business. SageMaker also removes the complexity that gets in the way of successfully implementing ML across use cases and industries—from running models for real-time fraud detection to virtually analyzing the biological impacts of potential drugs to identifying the best driver in Formula 1 racing.

Machine learning with AWS, by the numbers

100,000+ customers are using AWS for their AI and ML workloads

20+ years of building experience at Amazon

Up to 10x improvement in data scientists' productivity

Hundreds of algorithms and models in Amazon SageMaker JumpStart

ML frameworks: AWS customers can choose from TensorFlow, PyTorch, Apache MXNet, and other popular frameworks to experiment with and customize ML algorithms. You can use the framework of your choice as a managed experience in Amazon SageMaker, or use the **AWS Deep Learning AMIs** (Amazon Machine Images) and **AWS Deep Learning Containers**, which are fully configured with the latest versions of the most popular deep learning frameworks and tools. Amazon Elastic Compute Cloud (Amazon EC2) provides a wide selection of instance types optimized to fit ML use cases—regardless of whether customers are training models or running inference on trained models. These instances range from **GPUs** for compute-intensive deep learning training to **AWS Inferentia** for low-cost inference.

Implementation support: The **Amazon Machine Learning Solutions Lab** pairs your team with ML experts to help you identify and build ML solutions that address your organization's highest ROI ML opportunities.

We also offer training to augment the level of ML expertise on your team, including developer training, business leader training, and a hands-on event through the **AWS Machine Learning Embark Program**.

Learn more about how you can transform the **responsible use of AI and ML** from theory into practice with purpose-built services, resources, and training.

Learning tools: You can improve your ML capabilities with in-depth learning tools, including:

- **AWS DeepRacer**
- **Machine Learning Training and Certification**
- **Amazon Machine Learning Solutions Lab**
- **Amazon SageMaker Studio Lab**

Solving the biggest machine learning challenges

Most organizations have made some investments in ML and have made progress on their ML journeys. But many find themselves hitting stopgaps along the way, worried that costs and complexities will grow too high as they progress.

Throughout this eBook, we explored the steps to forge ahead and realize the full power of ML. To recap, let's look at the biggest challenges we identified along the way, along with a brief recommendation of how your organization can solve them.

To learn more about how you can overcome obstacles and accelerate your ML journey, visit the AWS ML resource hub.

Get started ›

Challenge	Solution
Discouragement from failures	Developing a fault-tolerant culture
Siloed, unprocessed data	Creating a modern data strategy that includes data lakes
Finding the right business problems	Building blended teams that include both technical and domain experts
The ML skills gap	Adopting new organizational models, processes, no-code tools, and team management philosophies
Sustainably scaling beyond pilot projects	Leveraging end-to-end tools like Amazon SageMaker to simplify ML development
Measuring the results	Forgoing traditional ROI metrics in favor of agility, competitive advantage, and risk tolerance; use the value tree model