A Startup path to ML success

How to leverage Machine Learning to power your Startup growth

Start your engines

When deployed with the right strategies, machine learning (ML) can increase agility, streamline processes, boost revenue by creating new products and improving existing ones, and enable better, faster decision-making, all critical elements as your startup is in its earliest stages.

There's no doubt machine learning and artificial intelligence (AI) can help startups achieve more—in a recent survey by McKinsey, 63 percent of respondents reported revenue increases from AI adoption in the business units where their companies use AI.[1] It's also clear that adopters continue to have confidence in AI technologies' ability to drive value and advantage. According to Gartner, 75 percent of companies will shift from piloting AI to operationalizing AI by the end of 2024,[2] and according to Deloitte, 57 percent say AI will transform their organization over the next three years.[3] While machine learning has been around for decades, its accessibility as a tool to transform startups is relatively new.

And the lack of a singular, proven path to machine learning success is keeping some waiting on the sidelines, unsure of how to take the next (or even the first) step on the journey. This eBook is designed to help startups forge ahead, outlining a proven path—from the first step to measuring results—with insights from AWS's own machine learning heritage and its experience helping thousands of customers realize their own initiatives.

It's time to overcome your machine learning worries, stop playing catch-up, and forge ahead with confidence. No matter their size or where they are on your machine learning paths, startups will find the guidance they need to take the next step to machine learning success.

[1] https://www.mckinsey.com/featured-insights/artificial-intelligence/global-ai-survey-ai-proves-its-worth-but-few-scale-impact
[2] https://www.gartner.com/en/newsroom/press-releases/2020-06-22-gartner-identifies-top-10-data-and-analytics-technolo
[3] https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html

What are artificial intelligence and machine learning?

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

Artificial intelligence (AI) is a way to describe any system that can replicate tasks that previously required human intelligence.

Almost always, this is related to some kind of 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 machine learning. Machine learning 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, machine learning is the underlying technology that powers intelligent systems.

AI can be created without machine learning, but right now, machine learning is the primary method for creating AI systems. Similarly, machine learning can be used for more than AI, but right now, the majority of machine learning is AI-related.

Why machine learning?

Before digging into the steps of the machine learning journey, let's explore why startups 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 unwavering dedication. Founders and technical decision makers will need to regularly remind yourselves what you're fighting for—keeping your eyes on the precise business benefits that can be unlocked by fully leveraging machine learning technology.

Startups already realize the impact of:

1

Optimizing business with new efficiencies

Machine learning can be used to improve the customer experience and create greater efficiency through sophisticated demand planning and forecasting models. While this is true in almost every industry, retail provides some specific evidence. AI-based forecasting is reducing lost sales due to product unavailability by up to 65 percent and resulting in two million fewer product returns per year.[4] And using machine learning tools from Amazon Web Services (AWS), augmented reality startup PulpoAR allows customers to try on makeup looks virtually, streamlining the online shopping experience and overcoming a barrier that keeps many customers from clicking the check-out button.

2

Smarter, faster decision-making

Informed by data and analytics sources that grow smarter through machine learning, startups and their workforces can make more informed, faster decisions that allow them to act on opportunities sooner and get better results. Legal technology startup DISCO uses cutting-edge AI to analyze data quickly, identify key information, and free up resources for tasks that require legal judgment — providing a competitive advantage in an industry where speed and accuracy are critical.

[4] https://www.mckinsey.com/~/media/mckinsey/industries/advanced%20electronics/our%20insights/how%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/mgi-artificial-intelligence-discussion-paper.ashx

3

Adapting to a fast-changing world

When circumstances change, startups need to be able to react quickly and pivot to new strategies. Machine learning can help startups adapt at the speed of innovation. During the COVID-19 pandemic, for example, lengthy lines and unfamiliar screening processes at many airports have made it clear that existing systems simply can't cope with the "new normal" of traveling. Anticipating the need to navigate this complex and high-stakes landscape, border security experts Travizory developed a world-leading secure SaaS border security and management platform using cutting-edge biometrics, AI and machine learning technologies. These tools have allowed the startup to develop features faster and deploy them in weeks or even days — instead of months or years — enabling countries to safely welcome back visitors more quickly.

4

Adding new capabilities to existing products

Machine learning can enrich existing products, improving customer engagement and attracting new users through deeper experiences. For instance, ICONY partners with more than 200 media companies to create bespoke online dating platforms. To combat the growing threat from scammers, bots, and other bad actors, the startup uses machine learning to detect and remove fraudulent profiles from the network. This leads to a better user experience, fewer abuse reports, and higher retention. With less harassment from fake accounts, users feel more comfortable on the platform and enjoy using it. It's also helped the startup as well — leveraging machine learning meant the ICONY team could spend less time dealing with fake accounts, and dedicate more resources to improving quality on the platform.

5

Inventing net-new products

Machine learning is not simply for established companies. It can also play a key role in generating exciting new ideas and disrupting slow-to-change industries. In its quest to create a wearable device for visually-impaired people, .lumen relied on AWS's AI tools and services in order to train its models, accelerate its research pace, and begin testing its product with hundreds of blind individuals. "Everyone told us it would take 10 years, and we did it in six months," said .lumen CEO and founder Cornel Amariei. "We can do everything — it's all possible."

6

Accelerating growth

Leveraging the power of machine learning can also help startups scale up quickly. For example, the healthcare industry is built on fragmented, outdated technology. To address this, Olive is using AI to build the Internet of Healthcare, connecting the industry's disparate technology and reducing the costs and burdens of the healthcare system. The results have been game-changing: Olive is connecting over 675 hospitals across the U.S. and has improved the healthcare experience for millions of patients.

Now that we've covered the "why" of machine learning, it's time to explore the "how." The next sections will demonstrate the steps of the machine learning journey, using Amazon's own path and those of AWS customers to exemplify the necessary changes that must take place in order to successfully implement, deploy, and scale machine learning.

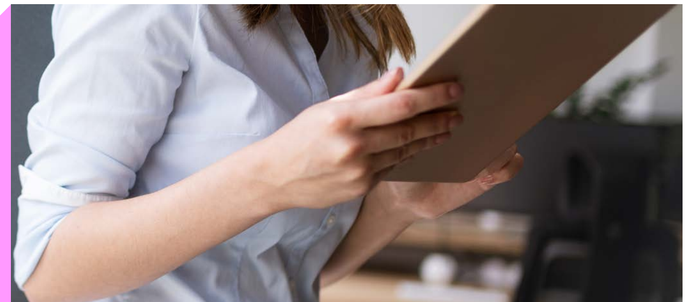

The machine learning journey

The machine learning journey is not always a straight path. Achieving success with machine learning 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 machine learning in efficient, sustainable ways.

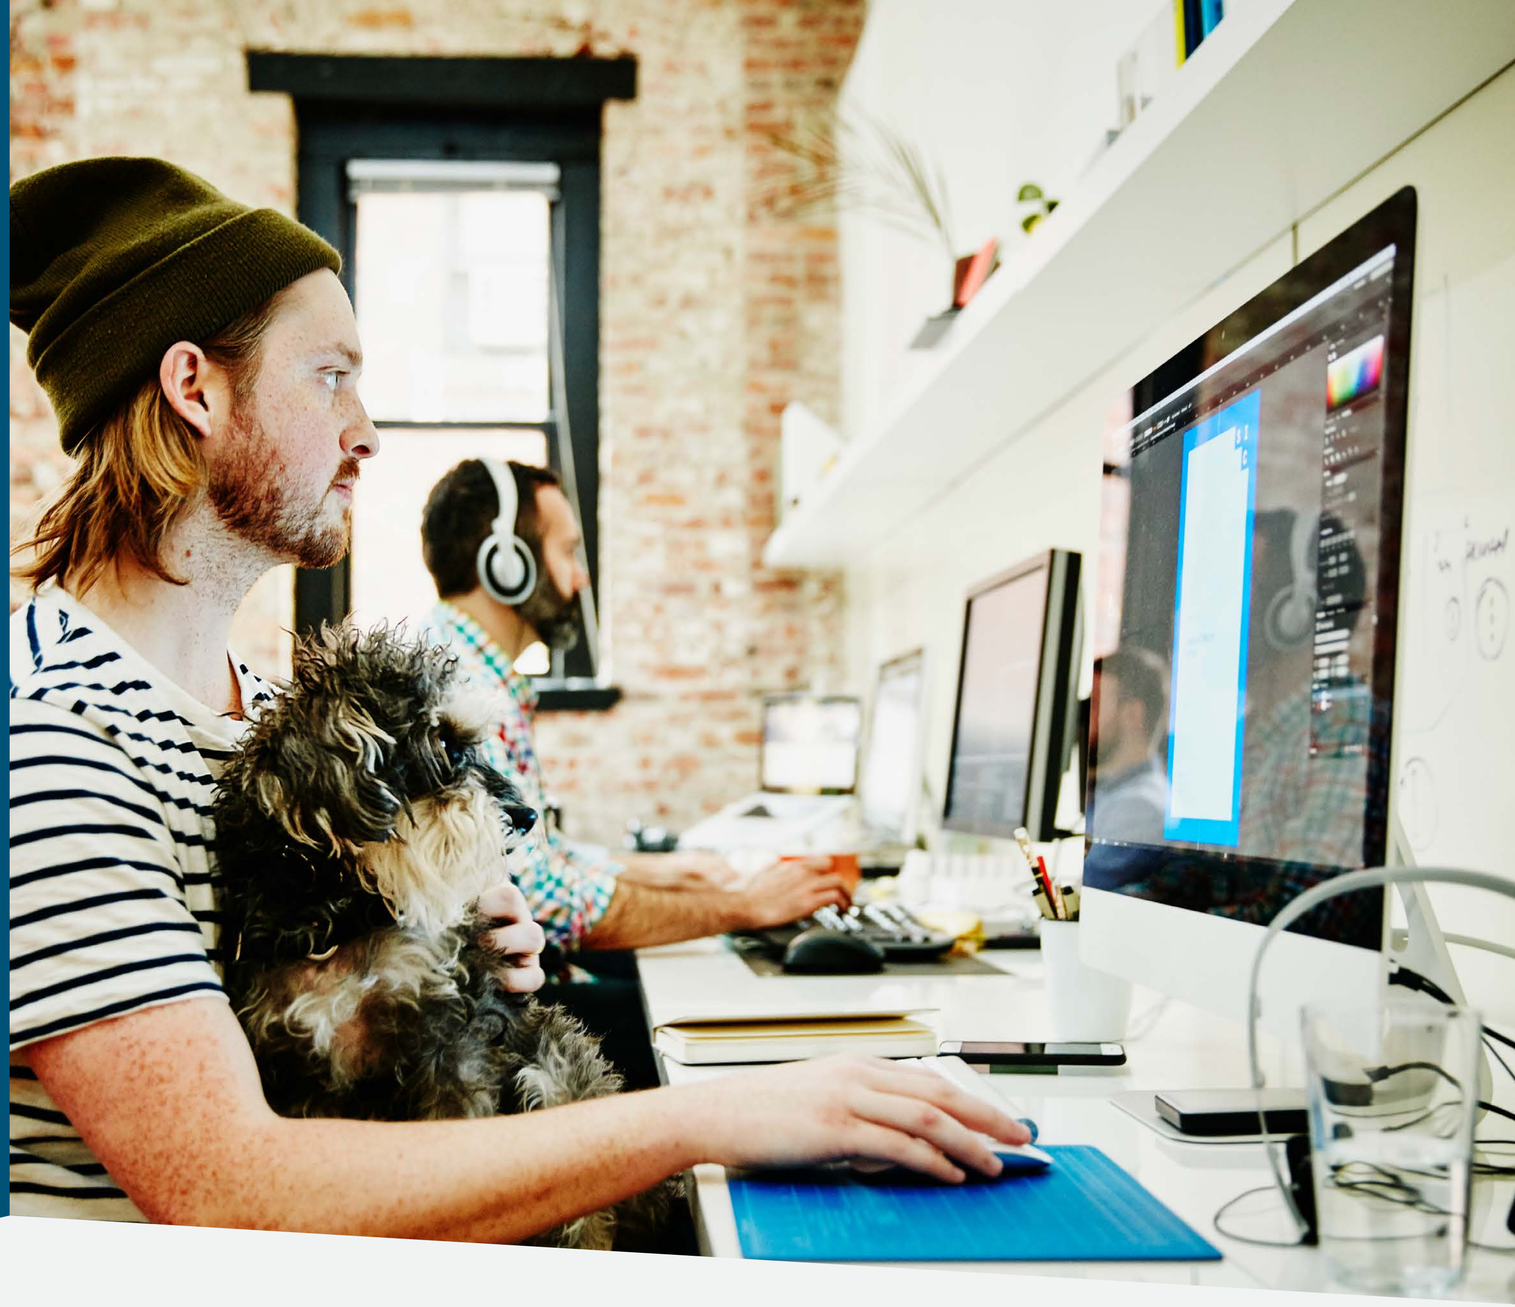

1

Championing a machine learning culture

According to Gartner, the global business value derived from artificial intelligence is projected to reach $3.9 trillion in 2022.[5]

Unlocking the full business potential of machine learning will require cultural changes in team organization, objectives, and outlook.

In order for machine learning to proliferate through the organization, both business and technical teams must work together and share the same priorities. To achieve this at the outset, the machine learning effort must be supported from the highest levels, with goals set by executive champions and an investment in the technology and the processes that enable success.

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

Startups often utilize the opportunity inherent in mistakes, and this is particularly true when it comes to machine learning. Machine learning is an iterative process, one which can only succeed through constant experimentation. More often than not, 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.

[5] https://www.gartner.com/en/newsroom/press-releases/2018-04-25-gartner-says-global-artificial-intelligence-business-value-to-reach-1-point-2-trillion-in-2018

How Amazon did it

Amazon has been using machine learning in the business for over 20 years. But more than 10 years ago, to further adoption of machine learning across the company, Amazon's leadership team asked every business leader in the organization—irrespective of whether they ran a research team, a fulfillment center, or an HR department—to answer the question of how they planned to leverage machine learning in their businesses.

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

In addition to hiring data scientists externally, Amazon also created the Machine Learning University, which trained many of its developers to use machine learning more effectively. The company built tools like Amazon SageMaker, which simplifies the process of creating models and lowers the barrier to entry—so machine learning could scale more effectively.

Now, more than10 years later, there isn't a single department at Amazon that hasn't been touched by machine learning. Amazon's personalization technology that provides recommendations to customers has significantly improved since its first model 20 years earlier and has been applied to other areas of the business.

The company uses machine learning throughout its fulfilment 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."

Amazon has also developed technology to give consumers an entirely new way to interact with technology through Alexa. The company has developed groundbreaking technology with autonomous flight via Amazon Prime Air drones and uses robotics in its fulfillment centers to get packages to consumers faster. Plus, Amazon is using machine learning to minimize the amount of packaging that customers have to dispose of, reducing the weight of their outbound packaging by 33 percent and eliminating 915,000 tons of packaging material worldwide.

Achieving these successes required great investments in technology, research, and talent. But those investments would have gone to waste without the cultural change 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 machine learning journey can truly begin.

2

Reinventing data strategy

Machine learning successes are 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 machine learning should aim to break down silos, enabling 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 to store the data—whether structured or unstructured—in a centralized repository.

Once stored, the data can be analyzed by many types of analytics and machine learning 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, learn how to put data first at your startup.

How claimsforce did it

In 2021, a 100-year flood devastated property and killed more than 180 people across Europe. For InsurTech startup claimsforce, this resulted in a sudden spike in insurance claims — a situation that could have strained the company's resources at a time when customers were in dire need of help. By leveraging machine learning and AWS data analytics tools, however, claimsforce was able to handle the increased demand with ease.

To achieve this, claimsforce built a data Lake House on AWS to handle the increasing volume of data it was collecting as a fast-growing startup. That allowed it to perform efficient queries on large amounts of data and identify regions where more resources were needed. Using Amazon SageMaker, Amazon's fully managed machine learning platform, allowed the company to make better assignment decisions and deploy resources more efficiently — helping to deliver assistance to people when they needed it most.

Read the full story ›

3

Finding the right business problem to address

For startups, embracing machine learning early on can often be challenging. Most early stage startups are focused on developing their minimum viable product, and without operable or sizeable data sets, it can be difficult for new companies to demonstrate the value of machine learning projects to investors. This can stall or even stop progress on machine learning initiatives.

Here are some important questions founders should ask before embarking on their machine learning journeys:

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

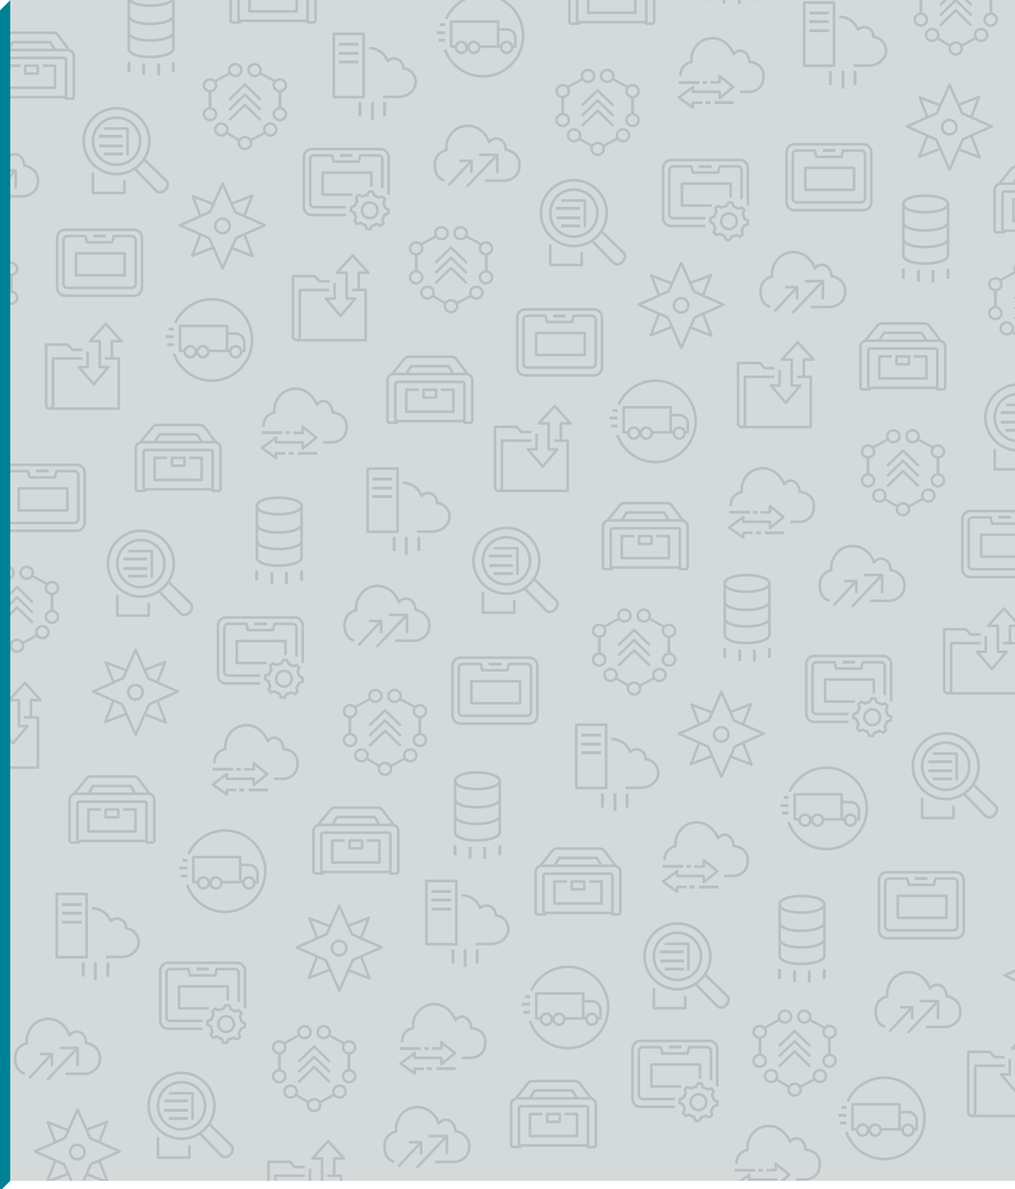

“A first step is to identify a problem that is rich in data, but (one that) you haven’t been able to solve through traditional methods,” Freshwater said.

In a successful machine learning journey, startups create machine learning teams that are built 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 machine learning 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 machine learning 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 machine learning and to implement the appropriate pilot programs. AWS offers the Amazon Machine Learning Solutions Lab to help customers “work backward” from business challenges—and then go step-by-step through the process of creating machine learning projects to solve them.

How Common Room did it

Communities are vital to the health of individuals and companies, but they can be sprawling, disjointed, and difficult to grasp—especially online. Enter Common Room, a community-intelligence platform that uses machine learning to provide a single view into everything that's important in your online community.

Community exchanges—from complaints, to requests for help, to feedback—often occur across many separate platforms, like Twitter, Slack, Intercom, Shopify, and more. Common Room compiles these interactions and mines them for what matters most, providing a deeper and clearer perspective on the health and metrics of a community, and the individuals or organizations that are engaging with it.

To do this, the startup sorts through an immense and diverse array of data—Tweets, GitHub comments, Slack messages—to homogenize and unify everything into one place.

AWS machine learning services allow the company to analyze this data and surface signal and meaning from the community. This allows Common Room to identify not only the negative comments in the community—who's angry? who's upset?—but also the positive feedback, which can be super motivating for the people at the company.

Read the full story ›

4

Championing a machine learning culture

In parallel with creating a data strategy, startups must focus on arming their team with the right skills.

Organizations are growing increasingly aware of the machine learning skills gap—the expanding separation between technologies and the ability of internal IT specialists to take full advantage of them. The O'Reilly 2021 AI Adoption in the Enterprise 2021 report, which surveyed more than 3,500 business leaders, found that a lack of skilled people and difficulty hiring topped the list of challenges in AI, with 19 percent of respondents citing it as a "significant" barrier. Closing this gap for machine learning will require a combination of training and recruiting. The reality is, there aren't enough data scientists today to lead the machine learning transformation that is coming. This requires startups that want to leverage machine learning to first invest in developing their talent.

While there is no one-size-fits-all solution to the machine learning 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 pre-trained expert talent. Organizations consistently reach the breakthrough successes waiting on the other side.

These methods include:

Defining the skills gap: Before closing the skills gap, technical decision makers must identify the precise differences between what they need or want employees to do and what its employees currently have the ability to do.

Understanding how skills are mapped: Since machine learning initiatives are interdisciplinary efforts, an organization should map the skills needed across data scientists, machine learning specialists, 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 machine learning needs. Leaders should also investigate pre-trained AI services that provide ready-made intelligence for business applications and workflows.

In addition to training, you'll need to align teams to successfully tackle machine learning problems. This includes:

- Promoting a culture of empowered teams: Machine learning project teams must be cross-functional, possessing the authority to execute individual objectives and the freedom to organically cross-pollinate with other teams as demands dictate and opportunities arise.
- Starting with a pilot team: Establish a pilot team of engineers and task it with a machine learning 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, consider splitting up the team, adding new engineers to create new teams, and tasking 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 startups are finding that the people they currently have actually are the people they need to close their machine learning skills gaps. While some recruiting may still be required, organizational, process, and management changes can do much of the work to upskill talent for machine learning success.

It's also important to ensure that startup leaders are trained to understand machine learning, including what makes a good use case and how to speak the language of machine learning. Many courses are available for business leaders, including Machine Learning Essentials from AWS.

How WOMBO did it

When WOMBO, a Toronto-based synthetic media company, launched its AI-powered lip sync app in 2021, it took off like a rocket ship, becoming one of the fastest-growing consumer apps in history. Its popularity highlights the fact that machine learning for efficiency and productivity won't be the only drivers of its broad-scale adoption. It also helps if it's fun.

For most startups, scaling quickly and becoming extremely popular is a desirable goal. But the rate of WOMBO's growth exceeded even the founders' wildest expectations. Partnering with AWS allowed the company to continue to improve the quality of its algorithms and operate in the cloud, a central part of WOMBO's vision.

AWS also worked closely with WOMBO's engineers to continually optimize the app's data processing to make sure clusters didn't get bottle-necked.

Read the full story ›

> "Our bet on the cloud is a critical part of what synthetic media is. If you compare how our app works to the vast majority of competitor apps in this space, we're the only ones that will bear the much higher costs of processing in the cloud, because we're convinced that the real AI experience people want—these really unique and unexpected visuals—cannot be processed on someone's phone in an appealing manner. For millions of people, we've been able to give them this exposure to AI as something fun."
>
> **Ray Khurana**
> CHIEF OF STAFF, WOMBO

5

Scaling beyond pilot projects

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

Achieving scalability requires organizations to make it easier for their developers to use machine learning. Building machine learning 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 machine learning workflow to build, train, and deploy machine learning models. By using Amazon SageMaker, organizations can get their models into production faster and at a lower cost, enabling sustainable expansion of machine learning initiatives beyond pilot projects.

How Koo did it

Using Amazon SageMaker, together with their own machine learning services, microblogging platform Koo is able to translate content into less widely used languages. To do this, the startup's software must have a rich, complex understanding of each language in order to translate not just literal meaning, but also the sentiment beneath words and phrases. While there's no shortage of data when it comes to Natural Language Processing for English, Koo must pioneer their own models for many of the more obscure, or less widely utilized, languages. SageMaker is a key tool in that endeavour.

"How do we give them a platform so that everyone is able to share their thoughts, their opinions, without any language barrier?" said Phaneesh Gururaj, President of Technology at Koo. "We want to be a part of the journey where we enrich our own language models with the content that we generate and go into that virtuous cycle of ensuring the model accuracy also increases."

Read the full story ›

“How do we give them a platform so that everyone is able to share their thoughts, their opinions, without any language barrier?”

Phaneesh Gururaj

PRESIDENT OF TECHNOLOGY, KOO

6

Measuring the results

When measuring the results of machine learning efforts, the traditional “project ROI” viewpoint—where a project has a defined start and end point, a budget, and a return—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 pivot to other solutions and miss out on critical opportunities down the line.

Instead, startup leaders and technical decision makers alike must measure machine learning efforts based on what success means for their business with regard to the processes being optimized. In addition, they must view machine learning efforts as long-term investments, acknowledging that a true “return” may not be realized for several years and after countless iterations.

When planning machine learning initiatives, it’s better to view the process through the lenses of agility, competitive advantage, and/or risk tolerance rather than expected return. Organizations will have greater success if they disregard 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 machine learning initiatives can still be measured—it just requires a different outlook.

Machine learning 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,” and/or “increases in upsells.”

Measuring the success of machine learning through a more holistic and long-term model will keep your teams focused on the best outcomes for the future of the company.

Taking the next step with AWS

No matter where you are in your machine learning journey, AWS provides products, solutions, and services that can help them take the next step. Featuring the world's broadest and deepest set of machine learning and AI services, AWS has worked with over ten thousand customers to help them successfully implement machine learning.

AWS is dedicated to putting machine learning in the hands of every developer and is working tirelessly to solve the toughest challenges that stand in the way of that goal. AWS capabilities are built on the most comprehensive cloud platform, are optimized for machine learning with high-performance compute, and compromise nothing in security and analytics.

Let’s explore current machine learning offerings from AWS—and see how they can help organizations progress in their journeys.

Intelligent contact center: Enhance your customer service experience and reduce costs by integrating machine learning into your contact center.

Personalization: Improve customer engagement and conversion by creating personalized web 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 siloed and unstructured information sources across the organization.

Fraud detection: Improve profitability by automating the detection of potentially fraudulent online activity, such as payment fraud and fake accounts, using machine learning and your own unique data.

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

Business metrics analysis: 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.

Amazon SageMaker: Amazon SageMaker enables developers and data scientists to quickly and easily build, train, and deploy machine learning models—thus simplifying scalability across the entire business. Amazon SageMaker removes the complexity that gets in the way of successfully implementing machine learning across use cases and industries—from running models for real-time fraud detection to virtually analyzing biological impacts of potential drugs to identifying the best driver in F1.

Machine learning with AWS, by the numbers

AWS machine learning solutions:

Reduce training time by **50%**[6]

Provide **90%** scaling efficiency[7]

Deliver **3x faster** network throughput[8]

Improve price and performance by **25%**[9]

91% of cloud-based PyTorch runs on AWS

92% of cloud-based TensorFlow runs on AWS

AI Key Use Cases: Explore the key use cases of machine learning to improve customer experience, optimize business operations, and accelerate innovation. No machine learning experience is required.

- Add intelligence to the contact center
- Personalize customer recommendations
- Automate data extraction and analysis
- Discover accurate information faster with intelligent search
- Identify fraudulent online activities
- Analyze media content and discover new insights
- Improve business operations and forecasting

Machine Learning Frameworks: AWS customers can choose from TensorFlow, PyTorch, Apache MXNet, and other popular frameworks to experiment with and customize machine learning algorithms. They can use the framework of their choice as a managed experience in Amazon SageMaker or use AWS Deep Learning AMIs (Amazon Machine Images), which are fully configured with the latest versions of the most popular deep learning frameworks and tools.

AWS customers also benefit from a broad set of powerful compute options, ranging from GPUs for compute-intensive deep learning to FPGAs for specialized hardware acceleration to high-memory instances for running inference. Amazon EC2 provides a wide selection of instance types optimized to fit machine learning use cases—regardless of whether customers are training models or running inference on trained models.

Implementation Support: The Amazon Machine Learning Solutions Lab pairs your team with machine learning experts to help you identify and build machine learning solutions to address your startup's highest return-on-investment machine learning opportunities. We also offer training to augment the level of machine learning expertise on your team, including developer training, business leader training, and a hands-on event through the Machine Learning Embark Program.

Learning Tools: AWS also offers a number of learning tools and services to help organizations improve their machine learning capabilities, including:

- AWS DeepRacer
- AWS DeepLens
- Machine Learning Training and Certification
- Amazon Machine Learning Solutions Lab

[6] As measured in the ResNet-50 benchmarking test, AWS-optimized TensorFlow recorded the fastest training time, by over 50 percent
[7] Using AWS-optimized TensorFlow allows for near-linear scaling efficiency, up to 90 percent compared to 65 percent using stock TensorFlow
[8] than other providers using P3dn instances
[9] using C5 instances powered by 3.0GHz Intel Xeon compared to previous generation instances

Solving the biggest machine learning challenges

Many startups have made investments in machine learning and are at some stage of the journey. But others find themselves hitting stopgaps along the way, worried that costs and complexities will grow too high as they progress. With help from AWS's diverse and tailored set of machine learning solutions, however, more startups than ever are experimenting and succeeding with machine learning.

In this eBook, we explored the steps toward forging ahead and realizing the full power of machine learning. To recap, let's look at the biggest challenges we identified along the way—with a brief descriptor of how startups can solve them.

To learn more about how startups can overcome obstacles and accelerate their machine learning journeys, visit the AWS machine learning resource hub.

Read the full story ›

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 machine learning skills gap	Adopting new organizational models, processes, and team management philosophies
Sustainably scaling beyond pilot projects	Leveraging end-to-end tools like Amazon SageMaker to simplify machine learning development
Measuring the results	Forgo traditional ROI metrics in favor of agility, competitive advantage, and risk tolerance; use the value tree model

Learn more about Machine Learning for Startups

If you have any questions or want further details, speak to a member of our team

Contact us

Join the AWS Startup Community

Visit the Startup Loft to get started and grow your startup with AWS

Sign up