



Democratized, operationalized, responsible: the 3 keys to successful AI and ML outcomes

Leverage the full power of artificial intelligence
and machine learning



Table of contents

PART I - THE MARKET

Introduction	4
Overview	5
Go time	6

PART II - THE KEYS TO SUCCESS

Key #1: Democratization	8
Key #2: Operationalization	11
Key #3: Responsible AI	13
Successful outcomes start here	15

PART I - THE MARKET

The rapid development of AI and ML technologies is fueling ever-growing investments and innovations

INTRODUCTION

Why artificial intelligence and machine learning matter more than ever

Artificial intelligence (AI) and machine learning (ML) are among the most transformative technologies of all generations. That helps explain why AI and ML adoption is on the agenda of so many business and technical leaders. Organizations everywhere are investing in AI and ML for a variety of reasons. Most notably, they fuel rapid innovation. They also drive competitive advantage by adding intelligence to existing processes, automating time-intensive manual tasks, and accelerating the creation of new products and services.

Although specific estimates on the growth of AI and ML vary, there is no debate about the impact of these technologies on the global economy. **Gartner predicts** that the size of the AI and ML market will reach \$134.8 billion by 2025, considerably outpacing the growth of the worldwide software market. Likewise, **IDC forecasts** that AI and ML investments will jump from \$85.3 billion in 2021 to \$204 billion in 2025, resulting in a compound annual growth rate (CAGR) of 24.5 percent for the 2021–2025 period. The effects of these numbers are (potentially) staggering—both for the organizations that make the investments and for those that do not.

Whether you are expanding your existing AI and ML initiatives or just getting started, a clear reference guide can help you develop your strategy and help ensure successful outcomes. This eBook outlines the three strategic pillars to success and provides practical recommendations that you can apply to your organization.



OVERVIEW

What AI and ML mean for business

The language of AI and ML reflects distinctions between the technologies themselves. Understanding these distinctions will help you know how to talk about the relevant technologies with those both inside and outside of your project teams.

At the most basic level, AI is a way to describe any system that can replicate tasks previously requiring human intelligence. Most AI systems look at probabilistic outcomes based on incoming real-world data to predict the most likely scenarios.

AI systems may use any combination of techniques to arrive at predictive decisions, including “if/then” statement logic. However, most AI systems are created using ML, which ingests large amounts of data to create and validate decision logic (also known as a model).

Deep learning is the part of ML that uses a technique known as deep neural networks to mimic the way biological neurons signal to one another. Deep learning is typically used in complex AI systems, such as autonomous vehicles or natural language processing (NLP).

A subset of deep learning called reinforcement learning is a technique that enables the system to learn in an interactive environment by simulated trial and error using feedback from its own actions and experiences.

Considering the recent rapid growth of artificial intelligence and machine learning

AI and ML have experienced highs and lows since emerging in the latter half of the twentieth century, but the technologies have become firmly established and market-facing over the past decade. A primary factor contributing to the recent upsurge in ML is the unprecedented utility of the cloud. Cloud computing brings together robust AI and ML services along with high performance computing (HPC), such as CPUs and GPUs, networking, and storage for model training and deployment. Additionally, the modern cloud continually protects the data and applications to support rigid security, compliance, and privacy requirements. These and other cloud-enabled building blocks make ML available to any organization.

The continued proliferation of unstructured data (which makes up over 80 percent of the world's data) has also helped ML gain converts and momentum. ML is unique in that it can use unstructured data to make predictions, unlike traditional analytics, which usually rely on structured data in a tabular format. To support the explosive growth in data, cloud services provide efficient ways to access, process, analyze, and store any amount of data. That makes it possible to realize value from the massive volumes of data created daily.

AI and ML models can leverage much of the world's available data. These models are equipped to inform predictions and enable the inclusion of video-, audio-, image-, and text-based data in a range of applications—from contact center intelligence and document processing to personalization and content moderation. For example, think of the cancer specialist seeking patient insights from X-rays, MRIs, or even handwritten prescriptions. Neither apps nor analytics can reliably extract predictive information from this kind of data; the physician (and her patient) need advanced intelligence for that.

With the help of AI and ML technologies, decision makers can now uncover patterns in unstructured data by utilizing AI and ML to decipher video and image content. The tools are ready, and the time is right to open the door to meaningful AI and ML innovation. Now, all you need are the keys to delivering successful outcomes.

PART II - THE KEYS TO SUCCESS

3 important keys can help customers bring artificial intelligence and machine learning to more parts of their business

KEY #1

Democratization

An article titled “Data Scientist: The Sexiest Job of the 21st Century” ran in the October 2012 issue of [Harvard Business Review](#). Prescient for sure, but the article may have underestimated the intensity of the demand a decade later—a global need for experienced data scientists that continues to spike repeatedly amid the growing interest in ML applications. Case in point: a 295 percent annual growth rate makes data science the fastest-growing skill set on [ZDNET](#), as measured by technology recruiters.

In response, many organizations are embracing a different approach. Instead of relying strictly on the recruitment of data scientists, they are using tools to empower the people they’ve already hired. By working with existing employees, they are expanding the capabilities and experiences of workers while reducing the need to compete with other companies for the relatively few data scientists who might be available at any one time.

This democratization of AI and ML tools and responsibilities is being widely embraced. It’s one of the three keys to making ML accessible to the many in a scalable way. Organizations that create a formal program for non-coders—such as line-of-business analysts supporting sales, finance, operations, marketing, and other key lines of business—will be far more agile than those that do not. In other words, democratization is good for business.

One way to democratize ML is through no-code tools for business analysts. No-code tools that allow business analysts to make ML predictions even without ML experience are used for common use cases, including churn prediction, lead conversion, price and revenue optimization, delivery time predictions, inventory planning, and more. No-code tools are a way to automate and eliminate dependencies on data science teams. They also give business analysts and data scientists the option of working together and creating opportunities to collaborate and create feedback loops to refine models further over time.

“Recruiters are struggling to find data scientists that can help them move their digital ambitions forward.”¹

[ZDNET](#), 2022

Then again, what's good for business analysts is great for data scientists. Using automated ML, or AutoML, data scientists can automate each step of the ML workflow. It democratizes the process by removing the tedious, iterative, and time-consuming work across the ML workflows—from data acquisition to model operationalization—so data scientists can spend less time on low-level details and more time on improving model accuracy. AutoML tools handle sourcing and preparing data, engineering features, training and tuning models, deploying models, and ongoing model monitoring and updating. AutoML techniques ultimately reduce the burden on data science teams, liberating experts from sacrificing their valuable time on repetitive, low-value tasks.

Democratization also includes support for the developers creating and maintaining critical business applications. While these developers will not typically build, train, and deploy their own models, they will be able to use pre-built models that can easily be integrated into their applications with a single API to add AI capabilities like computer vision, speech recognition, or personalization across a wide range of business applications.

Finally, AI and ML are becoming more readily available thanks to smart processes and tools supporting all lines of business and different job roles, such as AI-powered suggestions for next best action for customer support reps or optimized code suggestions for developers.

Democratized AI automates more of the applications we use to live, work, and play, allowing more time to focus on high-value activities.



A clear path to democratized artificial intelligence

Organizations can democratize AI and ML across a broad range of functions and roles with Amazon Web Services (AWS) solutions. For business analysts, **Amazon SageMaker Canvas** generates accurate ML predictions and boosts collaboration by sharing models and datasets with data scientists.

“With SageMaker Canvas, our business users can easily explore and build ML models to make accurate predictions without writing any code,” says Marc Neumann, product owner, AI Platform, BMW Group. The solution enables business analysts to connect, access, and combine data, create unified datasets, quickly and simply build ML models, analyze these models, and generate accurate predictions with just a few clicks.

According to Davood Naderi, data science team lead at Industrial Applications, Siemens Energy, “The collaboration is important because it helps us productionalize more ML models and ensure all models adhere to our quality standards and policies.”

Ready to remove heavy lifting from the ML process? For ML teams (which can include data scientists, MLOps professionals, and data engineers), there are many ways to leverage AutoML. **Amazon SageMaker JumpStart** offers hundreds of built-in algorithms and pre-built ML solutions that you can deploy with just a few clicks.

Amazon SageMaker Autopilot eliminates the work of building ML models. You simply provide a tabular dataset and select the target column to predict, and SageMaker Autopilot will automatically explore different solutions to find the best model. You then can directly deploy the model to production with just one click or iterate on the recommended solutions to further improve model quality.

“The primary goal in demographic mapping is optimizing across accuracy and scale,” reports Anindya Datta, CEO of Mobilewalla. “While this is generally difficult, we used Amazon SageMaker Autopilot to produce better models that improved our prediction accuracy by 137 percent.”

Developers can leverage a core set of AI services for text and documents, chatbots, speech, and vision—built on 20+ years of Amazon experience in ML—that can be used by any developer to easily add intelligence to any existing business application; no ML skills required. In addition to these core services, you have access to specialized and industry-specific AI services supporting AI and ML use cases, such as forecasting, personalization, anomaly detection, fraud detection, search, DevOps, industrial, and healthcare.

“Our developers now focus more of their energy on delivering differentiated capabilities and less time investigating problems in our production environment,” says Zak Islam, head of engineering, Tech Teams, Atlassian.

AWS is also adding ML capabilities to other services and products. For example, **Amazon CodeWhisperer** is an ML-powered service that helps improve developer productivity by generating code recommendations based on developer comments made (in natural language and in code) within the integrated development environment (IDE).

“Our developers now focus more of their energy on delivering differentiated capabilities and less time investigating problems in our production environment.”

Zak Islam, Head of Engineering, Tech Teams, Atlassian

KEY #2

Operationalization

With AI and ML adoption growing rapidly, even urgently, business and technical teams are challenged to build more with ML and to build it faster. In this rush to leverage the technologies, organizations rarely stop to set standard tools and processes for ML development. As a result, different teams with different skills and requirements often use completely different and disconnected tools, making collaboration impractical, if not impossible.

For example, an R&D team might be working on a computer vision application with state-of-the-art algorithms and frameworks, while sales and marketing teams are building a linear regression model to forecast customer demand on a locally stored spreadsheet. Or developers could be coding a mobile purchasing app and want to add a recommendation engine to make the customer experience more personal.

Relatively few organizations utilize operational ML tools and practices—such as infrastructure, IDEs, debuggers, profilers, collaboration tools, workflows, and project management tools—that can be connected securely. This reality

complicates management across teams of business analysts, developers, and data scientists and in coordination with existing software tools and processes. In these common scenarios, scaling up or down becomes exceedingly difficult.

The good news is that there's a proven way to minimize the risks and complications of ML while providing straightforward, repeatable practices for teams—by operationalizing ML. Operationalization of ML provides the tools, infrastructure, and operations support to scale.

Operationalizing ML starts with the data acquisition and modeling activities of the data science team being informed by a clear understanding of the business objectives for the ML application and of all governance and compliance issues. MLOps ensures that the data science, production, and operations teams work seamlessly together across a series of ML workflows that are as automated as possible. Human intervention is incorporated as needed, ensuring smooth deployments, data monitoring, and model performance tracking.

How AWS helps customers operationalize machine learning

Amazon SageMaker, which we already know is a powerful service to help democratize ML, is equally suited for operationalization. It automates and standardizes every step of the MLOps workflow to help projects scale without limits. Thanks to SageMaker, AWS customers are running millions of models with billions of parameters and generating hundreds of billions of predictions.

SageMaker also offers an end-to-end ML service for data labeling, data preparation, feature engineering, training, hosting, monitoring, and workflows. Your team can use a single visual inference in Amazon SageMaker Studio. In comparison to self-managed ML environments, the productivity of your data science team can improve by up to 10 times and model development time is reduced from months to weeks. And all SageMaker capabilities are offered on fully managed, low-cost, high-performance infrastructure in the cloud.

AWS customers are realizing massive scale (and savings) with SageMaker tools:

- **Vanguard** has fully automated the setup of its ML environments and now deploys ML models 20 times faster
- **AstraZeneca** can deploy new ML environments in five minutes versus one month to generate insights that improve R&D and accelerate the commercialization of new therapeutics
- **NerdWallet** reduced training costs by close to 75 percent, even while increasing the number of models trained
- **Zendesk** reduced ML inference costs by 90 percent by deploying thousands of models per endpoint using SageMaker multi-model endpoints
- **Mueller Water Products** cut the number of false alerts in half and maximized the potential to identify true leak events

Operationalization also means that we need to deliver the breadth and depth of **AI use cases**, including intelligent contact centers, intelligent document processing, content moderation, personalization, intelligent search, fraud prevention, identity verification, predictive maintenance, AI for DevOps, health AI, and ML-powered BI. AWS offers services for all these use cases and more.



“Being on AWS enables a faster development process for us.”

Kenji Takeuchi, SVP of Technology Solutions,
[Mueller Water Products](#), 2021

KEY #3

Responsible AI

Simply stated, responsible AI is a critical element of any AI and ML project. Encompassing a core set of concepts—fairness, explainability, security, privacy, and governance—responsible AI mitigates risks through the transparent use of data and models. It can be used to enhance model performance, enhance data protection, and enable bias detection and mitigation mechanisms in ML systems to improve fairness.

Responsible AI is an integral part of the complete AI and ML lifecycle, extending from initial design and development to deployment and, ultimately, ongoing use. It is an iterative process that requires ongoing testing and auditing for potential bias and accuracy. While most companies have begun their journey to responsible AI, the majority (94 percent) are struggling to operationalize across all key elements of responsible AI.²

“94% (of companies) are struggling to operationalize across all key elements of responsible AI.”

Accenture, 2022

So, how do organizations transform responsible AI from theory into practice? They begin by educating the next generation of ML leaders to elevate fairness and mitigate bias by bringing more diverse perspectives to the table and providing resources to promote education and training. Responsible AI also requires a multidisciplinary effort by technology companies, policymakers, community groups, scientists, and more to tackle new challenges as they arise and work to share best practices and accelerate research.



Build more responsible, inclusive AI with AWS

AWS is committed to developing fair and accurate AI and ML services and providing customers with the tools and guidance needed to build AI and ML applications responsibly.

As you scale your use of AI and ML technologies, you can leverage AWS resources to help implement responsible AI across the entire AI and ML lifecycle. AWS services help you better detect bias in datasets and models, provide insights into model predictions, and better monitor and review model predictions through automation and human oversight.

You can mitigate bias and improve explainability with AWS purpose-built services. **Amazon SageMaker Clarify** helps you mitigate bias across the ML lifecycle by detecting potential bias during data preparation, after model training, and in your deployed model by examining specific attributes. Similarly, SageMaker Clarify provides greater visibility into model behavior, both overall and for individual predictions, so you can provide transparency to stakeholders, more deeply inform humans making decisions, and track whether a model is performing as intended. Monitoring is also important to maintaining high-quality ML models and ensuring accurate predictions.

Amazon SageMaker Model Monitor automatically detects and alerts you to inaccurate predictions from models deployed in production.

Check out three essential resources to enable more responsible AI:

- **The Responsible Use of Machine Learning guide** provides considerations and recommendations for responsibly developing and using ML systems across three major phases of their lifecycles: 1) design and development, 2) deployment, and 3) ongoing use. [Read the guide >](#)
- Work with experts in responsible AI within our **AWS Professional Services** organization to create an operational approach encompassing people, processes, and technology that maximizes benefit and minimizes risk. The engagement includes the development, deployment, and operationalization of responsible AI principles. [Learn more >](#)
- Continuous education on the latest developments in ML is an important part of responsible use. AWS offers the latest in ML education across your learning journey through programs like the [Machine Learning University \(MLU\)](#), [Training and Certification program](#), [AI & ML Scholarship program](#), and [AWS Machine Learning Embark program](#).

AWS is committed to developing artificial intelligence and machine learning in a responsible way, helping our customers put responsible AI into practice and spurring research and continued development in this area. Our work to build more responsible, inclusive AI is just beginning. [Learn more >](#)

Successful outcomes start here

More than 100,000 customers have chosen AWS for AI and ML to create new customer experiences, optimize their businesses, augment their employees' ingenuity, help improve the quality of their products, and so much more. That's because AWS supports you no matter where you are on your AI and ML journey—with the solutions you need to scale without limits.

Connect with the experts from AWS

- **AWS Professional Services** is a global team of experts that can help you realize your desired business outcomes for the AWS Cloud
- **Amazon ML Solutions Lab** provides a risk-free environment in which to identify and test ML solutions and consider your organization's best ROI-producing opportunities

Collaborate with an official AWS Partner

- **AWS Partners** are uniquely positioned to help AWS customers accelerate the journey to the AWS Cloud

Construct it yourself with proven solutions

- **AWS AI services** let you easily add intelligence to applications—no ML skills required
- **Amazon SageMaker** empowers users to build, train, and deploy ML models for any use case with fully managed infrastructure, tools, and workflows
- **AI Use Case Explorer** helps you discover the top AI use cases, customer stories, and implementation paths based on your business objectives
- **AWS Solutions Library** offers solutions built by AWS and AWS Partners for a broad range of use cases
- **Amazon SageMaker JumpStart** provides SageMaker solutions for the most common use cases



The visionary *Harvard Business Review* article that portrayed the role of the data scientist as “the sexiest job of the 21st century” closed with a universal message aimed at cautious circa-2012 business leaders waiting on the sidelines for the technology to mature. “If companies sit out this trend’s early days for lack of talent, they risk falling behind as competitors and channel partners gain nearly unassailable advantages. Think of big data as an epic wave gathering now, starting to crest. If you want to catch it, you need people who can surf.”³

Fast-forward to today. The data is unquestionably bigger, and the wave more epic—but now it’s a democratized, operationalized wave that all can surf responsibly.

From the world’s largest enterprises to emerging startups, more ML is built on AWS than anywhere else.

**Learn more about how your business can deliver successful
AI and ML outcomes ›**



³ Davenport, H., and Patil, DJ, “Data Scientist: The Sexiest Job of the 21st Century,” *Harvard Business Review*, 2012
© 2023, Amazon Web Services, Inc. or its affiliates. All rights reserved.