



Enterprise AI Guide

The Executive's Playbook
for Business Success with AI

Artificial intelligence is quickly becoming a fundamental building block of business operations, providing real gains such as improved processes, increased efficiency, and accelerated innovation. Advances in machine learning technology have combined with high-performance compute options and an abundance of data to create a perfect storm for AI to transform organizations of all sizes.

For enterprise leaders, the need to embrace AI capabilities is immediate. IDC estimates that 40% of digital transformation initiatives will be supported by AI in 2019¹ and 42% of executives say AI will drive innovation in their organization.² As enterprises accrue data from numerous sources—ranging from customer feedback, sales logs, internal processes, and financial forecasting—there is opportunity to operationalize this often “dark data” into insights that can improve decision-making, create new revenue streams, and optimize the business value chain.

¹ “Worldwide Spending on Cognitive and Artificial Intelligence Systems Will Grow to \$19.1 Billion in 2018,” March 22, 2018, IDC

² “Reworking the Revolution: Future Workforce,” 2018, Accenture





Impacting business goals

Early adopters are weaving AI across the organization to address business priorities. Some machine learning projects tackle incremental gains to automate processes to create efficiencies. Others are transformational initiatives aimed at innovation and competitive differentiation. While there are numerous applications of AI, enterprises already realize great value from use cases that provide new experiences for their customers and drive business growth.

Goals for AI innovation

- 1 Enable product and service innovation
- 2 Drive research and discovery
- 3 Enhance the customer experience
- 4 Improve customer service
- 5 Improve healthcare outcomes
- 6 Increase efficiency and productivity
- 7 Improve security and compliance
- 8 Optimize supply chain operations
- 9 Improve decision making

1 Enable product and service innovation

While finding new ways to innovate is top of mind, more mature AI enterprises already use machine learning pervasively. They inform their product roadmap through intelligence extracted from customer-feedback; drive the product development lifecycle, including DevOps and quality assurance through automation and intelligence; and infuse machine learning capabilities directly into new products that benefit the end user.

2 Drive research and discovery

AI-driven research unlocks new ideas in diverse industries such as automotive, manufacturing, and energy and utilities, where machine learning capabilities help uncover new formulas, new materials, and new ways of thinking about age-old problems. It is especially relevant in industries such as healthcare and life sciences, where machine learning models can make predictions that drive research toward breakthroughs that lead to new and improved patient treatments.

3 Enhance the customer experience

Across industries including financial services, healthcare, retail, and media and entertainment, targeted recommendations have become important for personalization that delights and engages customers. Epsilon research indicates 80 percent of consumers are more likely to make a purchase when brands offer personalized experiences. Using voice-of-customer analytics to understand customer behavior and preferences, as well as recommendations engines to predict best matches, enterprises can push customized offers to location-based mobile services that offer real-time convenience and choice.

4 Improve customer service

Critical to every business is ensuring that various customer touchpoints—both internal and external—provide timely, accurate, and meaningful customer service. AI technology such as conversational agents bridge the gap, modernizing the contact center to provide quick responses to customer calls and chat requests. In addition, text-to-speech and natural language processing capabilities can be applied to the streams of incoming customer data such as texts, voice messages, and customer service logs to better understand customer needs and sentiment in order to help improve the quality of customer service.

5 Improve healthcare outcomes

Today AI is being used in healthcare to help patients quickly locate available in-network doctors. Behind the scenes, machine learning is being used to analyze everything from X-rays to patient-history data. As well, it can make predictions on patients at risk, and help physicians and hospital personnel determine the best course of action.

6 Increase efficiency and productivity

Across business operations, pervasive use of machine learning to automate and streamline processes has resulted in increased efficiency and reduced costs. For example, using automated media tagging and large-scale document recognition and analysis, machine learning can eliminate countless hours of labor. Machine learning can also analyze manufacturing incidents in real time, providing early warning of potential problems. HR departments are even turning to machine learning to assess job applicants more efficiently.

7 Improve security and compliance

Use of machine learning across security, risk, and compliance use cases is a fast-growing trend especially in the financial services sector. Fraud-detection models help to keep consumer data safe and prevent malicious attacks against bank accounts and even mobile endpoints. On the other end of the spectrum, machine learning automates mundane tasks such as financial-document analysis, reducing manual effort and allowing the workforce to focus on higher-value tasks.

8 Optimize supply chain operations

From point-of-sale to freight-delivery management to forecasting demand, AI plays a pivotal role across the supply chain. In demand planning for example, machine learning algorithms can look at historical data as well as demand and other relevant data—such as product schedules, weather patterns, and competitor pricing—to determine when it's time to restock or end of life a product. Sophisticated machine learning models can predict weekly, daily, and even down-to-the-minute forecasts to help companies streamline supply chain management, save cost, and increase efficiency.

9 Improve decision-making

Machine learning-based predictions augment decision making across all functions of the organization and across use cases. These predictions produce continuous actionable insights that help leaders and teams tackle operational and business challenges. This can range from forecasting that impacts pricing optimization, to predictions for autonomous vehicles to make informed choices.

Challenges of adoption

Machine learning is frequently the catalyst that turns business data into accurate predictions and actionable information, but as with many emerging technologies, there are challenges to adoption including data, complexity, cost, and lack of skills.


Data ambiguity: Enterprises can struggle with various issues related to data. First and foremost, many are unaware of all their possible data sources that may hold hidden insights. Even when they've identified data, there's a lack of labeled data ready for machine learning. Furthermore, even labeled data can prove to be an issue where integrity is in question since data can often have hidden biases from human labelers. Finally, enterprises often struggle with ensuring the right data management and governance policies are in place to allow the right people and processes to securely access, store, and manage the data.

Complexity: The machine learning workflow can be time-consuming and iterative, which leaves many organizations and developers thinking machine learning is complex and difficult to use. There are many steps involved from

prepping data, choosing algorithms, building, training, and deploying models ... and iterating over and over again. There are decisions to be made about infrastructure—selecting the right compute for training and inference, considerations for cloud, on premise, and edge deployments.

Cost inhibitors: Machine learning training and inference can be expensive, especially since models require iterations to improve the accuracy of predictions. Because embarking on AI initiatives is new to many companies, they also don't have the experiences or skills in-house, and often have to rely on external resources to kick-start projects.

Lack of skills: Even when companies embrace new technologies like AI to drive business transformation, not having the right skills is often a road block to getting started. AI projects require machine learning expertise to build and train ML models—this includes the skills of machine learning developers, data scientists, and researchers to build algorithms and train models. These skills are not in great supply and are often unavailable in-house which leaves businesses struggling to train or supplement their skills with consultants and partners.



Getting started with AI

AI presents new opportunities to realize foundational gains such as efficiency and cost savings, as well as higher-value gains such as product innovation and spurring discovery and research. But how do organizations get started? For many, AI adoption begins by identifying workflows and business processes that suffer from low efficiency or where human mistakes abound. They consider all their data sources and existing data strategy. They determine the best cloud-based infrastructure and tools to scale AI. And last, they ensure that the right skills are on board for machine learning projects to be successful.

Key considerations

- Understand business objectives
- Advance your data strategy
- Leverage the cloud
- Enable your organization





- Identify desired outcomes and deliver better experiences for your customers

Understand business objectives

Understanding the business benefits of AI adoption—in particular, the specific benefits relevant to your organization—is critical for enterprise success with AI. Once objectives are identified, it's important for business and technical leaders to understand and champion their role.

- **Select a targeted use case:** When choosing a pilot, consider use cases where AI can have the most impact, and those from which you can learn to scale enterprise deployment. Focus on how you can deliver a better experience for your customers and identify the business and operational outcomes desired. Then establish one or two high-value proofs of concepts (PoC) that can really make a difference to your organization and quickly demonstrate results. For the PoC to succeed, it's critical to have the right resources in place, including infrastructure, data, and capabilities.
- **Understand the impact:** From the outset, consider the operational effect of new AI solutions. AI can have transformative impacts, so it's important to plan in advance for what you want to achieve and measure. This can also help you determine how to measure success. When considering ROI, ensure you have value checkpoints early in the project lifecycle. This lets you adapt the approach before you've scaled investment.
- **Iterate and learn:** Once you've proven the potential of machine learning, the next step is to move from pilot to production, which may include integrating the machine learning capability into a larger IT system. This move typically takes longer than the pilot process and can vary depending on the complexity of the overall system and how large-scale the production deployment will be.

Advance your data strategy


Data is gold for leaders who are looking to disrupt their industries with machine learning, but many organizations don't have machine learning-ready data. Recognizing the importance of data and developing a plan to collect and use that data is critical for successful machine learning adoption, even at the PoC stage.

All sources of data need to be uncovered, from structured data like billing and CRM to unstructured data like social media feeds, images, and forms. Then that data needs to be evaluated for quality and usefulness. Finally, data needs to be cleaned and accurately labeled for machine learning models to transform it into valuable insights.

Leverage the cloud

Successful machine learning initiatives need more than just the right tools. A comprehensive platform brings together data store, security, and analytics services, as well as compute resources for training and deployment. Turning to the cloud for these services brings a wide range of benefits, including speed, scalability, flexibility, resilience, security, and reduced cost.

As well, the cloud offers the widest range of high-performance CPU and GPU processor types, which are essential for large-scale training and for deployment in a production environment. Using cloud-based data lakes and storage also ensures that you can easily access and manage data so that machine learning initiatives are seamless, repeatable, and scalable.



› Data is gold for leaders who are looking to disrupt their industries

Enable your organization

Along with the right use cases, having the right skills to build machine learning applications and systems, as well as the right process and operating model, are essential to getting pilots off the ground and scaling enterprise AI.

- **Assemble the team:** Consider appointing a Chief Data Officer (CDO) to lead the charge on data strategy and governance, bring together interdisciplinary teams, and streamline data processes. Assemble a team of machine learning developers and data scientists essential for a successful PoC, and train teams for future deployments. It's also important to involve subject matter experts who understand your business vernacular, especially for industry domains, to help you get to ground truth with your data.
- **Create the process:** AI may not bring the expected value if the results are not integrated with other areas of the organization. Operationalizing machine learning models is hard—as many as half of PoCs don't get deployed into production. Therefore executive sponsorship to change business processes and alignment with applications development is key. Successful teams create processes to align AI experts, data scientists and developers with key business stakeholders. A well-defined process also helps ensure the final output is well integrated into business processes.
- **Build the culture:** To help realize its potential, there needs to be cultural acceptance that AI is an important part of business and operations. Some initiatives may require information from across these domains, so it's important to understand all the stakeholders who need to be involved, and bring together stakeholders who can champion adoption.

- › Enlist the right stakeholders and skills to get pilots off the ground and scale




Getting started with AWS AI

AWS has the broadest and deepest set of AI and Machine Learning services for your business. AWS and Intel have been working together for over a decade to deliver the most comprehensive set of resources, tools, training, and services. On behalf of our customers, we are focused on solving some of the toughest challenges that hold back AI adoption.

Choose from pre-trained AI services; Amazon SageMaker to build and scale machine learning; or build custom models with support for all the popular open-source frameworks. AWS AI capabilities are built on a comprehensive cloud platform, optimized for machine learning with high performance compute, security and analytics.

AI Services: AWS pre-trained AI Services provide ready-made intelligence for your applications and workflows. AI Services easily integrate with your applications to address common use cases such as personalized recommendations, modernizing your contact center, improving safety and security, and increasing customer engagement. Because we use the same deep learning technology that powers Amazon.com and our Machine Learning Services, you get quality and accuracy from continuously-learning APIs. And best of all, AI Services on AWS don't require machine learning experience.

ML Services: Amazon SageMaker provides every developer and data scientist with the ability to build, train, and deploy machine learning models quickly without needing data engineers or DevOps. Amazon SageMaker, with optimizations from Intel, is a fully-managed service that covers the entire machine learning workflow to label and prepare your data, choose an algorithm, train tune and optimize models for deployment, make predictions, and take action. Your models get to production faster with much less effort and lower cost.

A decorative background on the right side of the page, consisting of a dark blue area filled with numerous small yellow dots. A faint, glowing yellow line or path winds through the dots, starting from the bottom left and moving towards the top right.

200+
New ML features and
services since Jan 2018

75%
Lower ML inference cost

One-click
Model training
and deployment

10x
Better ML algorithm
performance

Amazon ML Solutions Lab: The Amazon ML Solutions Lab combines hands-on educational workshops with advisory professional services to help you "work backwards" from business challenges, and then go step-by-step through the process of developing machine learning-based solutions. You'll be able to take what you have learned through the process and use it elsewhere in your organization to apply machine learning to business opportunities.

AWS Machine Learning Competency Partners: AWS Machine Learning Competency Partners offer a range of services and solutions to help you create intelligent solutions for your business. From concept development through execution, partner solutions offer tools and expertise to help with the entire machine learning workflow.

Training and Education: Start training on machine learning on AWS with courses based on the same material used to train Amazon's developers through the combination of foundational knowledge and real-world application. Developers, data scientists, data platform engineers, and business decision makers can use this training to learn how to apply machine learning, AI, and deep learning to their businesses.

➤ To learn how to put AI to work for your business, visit [AWS.ai](https://aws.ai)



Build your organization's AI with AWS and Intel


More machine learning is built on AWS than anywhere else, and getting started has never been easier.

AWS provides a smooth and simple on-ramp for companies looking to adopt AI and ML:

- Support for all major ML frameworks, including Intel-optimized versions of Tensor Flow and MXNet, giving engineers and developers built-in performance and freedom of choice
- Pre-built algorithms and model creation tools with Amazon SageMaker, including reinforcement learning tools like Intel® RL Coach and Ray RLib
- Partner-built models in AWS ML Marketplace that can quickly get your organization started with the most common use cases
- Optimized infrastructure—such as the Amazon EC2 C5 and C5n instances—running on custom Intel® Xeon® Scalable Processors for top performance on compute intensive ML training and inference workloads
- Software and hardware tools—including AWS DeepLens cameras and AWS DeepRacer—featuring Intel Atom® processors and Intel OpenVINO™ software that empower your developers to learn about ML hands-on
- AWS and Intel offer a full portfolio of business development and technical training resources, giving your organization the skills it needs to push AI forward with confidence



- > **Amazon SageMaker** enables one-click training, model tuning, and one-click deployment to the cloud or edge aws.amazon.com/sagemaker
- > **Amazon EC2 C5 and C5n** instances offer the lowest price per vCPU in the Amazon EC2 family and are ideal for running advanced compute intensive workloads like ML training and inference aws.amazon.com/ec2/instance-types/c5
- > **AWS DeepLens** is a fully programmable video camera, tutorials, code and pre-trained models designed to expand deep learning skills aws.amazon.com/deeplens
- > **AWS DeepRacer** gives developers a way to get hands-on with machine learning, experiment, and learn through autonomous driving aws.amazon.com/deepracer
- > **AWS ML Marketplace**—over 150 pre-built models and algorithms—including 14 from Intel—to get you started on the most common ML projects aws.amazon.com/marketplace/solutions/machinelearning



aws AI