INAWISDOM

Machine Learning on AWS with Inawisdom

EBOOK

aws partner network

Table of contents

Introduction

Machine learning is quickly becoming a fundamental building block of business operations, providing real impact 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 to transform organizations of all sizes.

For enterprise leaders, the need to embrace machine learning capabilities is immediate. As enterprises accrue data from numerous sources—ranging from customer feedback and sales logs to internal processes and financial forecasting—there is an opportunity to operationalize this often "dark data" into insights that can create new revenue streams.

Impacting business goals

Companies are using machine learning across the organization to address business priorities. Some machine learning initiatives 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 machine learning, enterprises already realize great value from use cases that provide new experiences for their customers and drive business growth.

Goals for machine learning innovation

1 Enable product and service innovation

2 Drive research and discovery

3 Enhance the customer experience

4 Modernize customer service

5 Increase efficiency and productivity

6 Improve security and compliance

7 Optimize operations

8 Enable agile decision-making

Impacting business goals (cont.)

1
Enable product and service innovation

Enterprises are already using machine learning pervasively to enable product and service innovation. 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

In industries such as automotive, manufacturing, and energy and utilities, machine learning helps 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 towards 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. 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
Modernize customer service

Critical to every business is ensuring that various customer touch points—both internal and external—provide timely, accurate, and meaningful customer service. Machine learning 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 to help improve the quality of customer service.

Impacting business goals (cont.)

5

Increase efficiency and productivity

Across business operations, the 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.

6

Improve security and compliance

Leveraging 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 or 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.

7

Optimize operations

From point of sale to freight delivery management to forecasting demand, machine learning 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 of 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.

8

Enable agile 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. In healthcare, behind the scenes, machine learning is being used to analyze everything from x-rays to patients' history data, helping clinicians make better and faster decisions.

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 ambiguity, 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 might 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 based on 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 and choosing algorithms to 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-premises, 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 machine learning initiatives is new to many companies, they also don't have the experiences or skills in-house and often have to rely on costly external resources to kick-start projects.

Lack of skills

Even when companies embrace new technologies like machine learning to drive business transformation, having the right skills is often a roadblocker for getting started. Machine learning initiatives require machine learning expertise to build and train machine learning 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.

Getting started with machine learning

Machine learning presents new opportunities to realize foundational gains such as efficiency and cost saving, as well as higher-value gains such as product innovation and spurring discovery and research. But how do organizations get started? For many, machine learning adoption begins by considering all data sources and existing data strategies. They identify workflows and business processes that suffer from low efficiency. They consider all their data sources and existing data strategy. They determine the best cloud-based infrastructure and tools to scale machine learning. And last, they ensure that the right skills are onboard for the machine learning initiative to be successful.

Key considerations

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

Key considerations

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 proof of concepts (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.

Understand business objectives

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

Select a targeted use case

When choosing a pilot, consider use cases where machine learning 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 PoCs 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.

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

Key considerations (cont.)

Understand the impact

From the outset, consider the operational effect of new machine learning solutions. Machine learning can have a transformative impact, 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. Rather than aiming for a target financial ROI within a given time frame, you'll find more success by measuring the impact of machine learning initiatives in terms of business agility, competitive advantage, and risk tolerance.

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.

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.

Additionally, 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.

> Identify desired outcomes and deliver better experiences for your customers

Key considerations (cont.)

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, is essential to getting pilots off the ground and scaling enterprise machine learning.

Assemble the team

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 the ground truth with your data. Consider appointing a Chief Data Officer (CDO) to lead the charge on data strategy and governance, bring together interdisciplinary teams, and streamline data processes.

Create the process

Machine learning 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 application development is key. Successful teams create processes to align machine learning 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 machine learning 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

Machine learning on AWS with Inawisdom

Inawisdom is a leading AI/ML data expert, focussed on delivering you valuable insights from the data you already hold. We know that for many businesses, a successful ML deployment may require skills and expertise beyond that of their existing team and we are here to make sure you can still take advantage of this transformational technology.

Depending on the size of your business and the in-house resources, we can supplement the team with bespoke expertise and experience at different stages of the journey – during the model creation phase or to support the MLOps process, for example. You may require support for the whole project, from strategic development through to the productionisation of the models, to determine the vital insights that can impact day-to-day business operations.

Whatever the journey looks like for your business, Inawisdom can offer the expertise you need to make AI/ML a success. As an AWS Premier Partner and the AWS Machine Learning Partner of the Year 2020, we're perfectly placed to help businesses unlock the value in their data with ML, to drive digital transformation and enable data-driven decision-making.

As proven AWS experts, we provide full-stack Amazon Web Services (AWS) Cloud and Data Services, including Data Lake and Data Platform through to Engineering, Predictive Analytics and IoT. Our best-in-class team have deployed these skills for AWS customers across the globe, helping them to achieve their goals.

Using our Rapid Analytics and Machine Learning Platform (RAMP) and our Discovery-as-a-Service consulting methodology, we help business quickly prove the value of machine learning and take models from proof-of-concept to production. Based on AWS services and frameworks, our proven approach helps businesses extract valuable insights from their data in as little as 6 weeks.

> **Inawisdom was founded with a simple goal: give our customers the ability to exploit every aspect of their data using AI.**
>
> Neil Miles, CEO, Inawisdom

Benefits of machine learning on AWS with Inawisdom

Choosing the right partner for your AI project is essential. At Inawisdom, we have the deep technical expertise to be able to slot in alongside your team, upskilling and supporting as needed. We have a clear idea of the business value ML will bring your business specifically, so that the project stays on track to deliver tangible value, rather than technology for technology's sake.

At Inawisdom, we've carefully developed our approach to ensure speed to value, actionable insight and a solid foundation for future innovation:

In-depth knowledge of AWS Services: As a Premier Partner, we're closely aligned with AWS frameworks and processes, and hold Competencies in Machine Learning, Data & Analytics, Financial Services, and DevOps. Our experts have a huge range of experience applying and integrating AWS Services, and are continually achieving AWS accreditations and training up on the latest developments. Whether you want support with migration, data lake creation, or running analytics and data exploitation with AI/ML, Inawisdom have the in-depth AWS knowledge to help.

Rapid time to results: With our RAMP platform (built on Amazon Redshift and SageMaker) and our Discovery-as-a-Service, we can help you quickly prove the value of Machine Learning for your business. Because we focus on the business value of ML right from the start, we can prioritise the best use cases and quickly identify the opportunities for transformation. We can take data from multiple sources and in varying formats - from text to speech to images - and turn it in to the visualisations, trends, predictions and opportunities to truly impact your organisation.

Wide-reaching expertise across industries and business challenges: Our expert consultants work closely with business leaders in multiple industries, including Utilities, Manufacturing, Automotive, Transport and Logistics, Finance and Insurance and Retail. We understand the dynamics of different sectors and have deployed tailored AI/ML to uncover insights, across all areas of a business. From customer service and loyalty, to fraud detection, preventative maintenance, ecommerce and supply chain efficiency, AI/ML can dramatically improve the effectiveness of your operation.

Case study: Aramex

Aramex is a leader in the global logistics and transportation industry, and employs more than 18,000 people in 604 locations across 71 countries.

Challenge

With a focus on an ecommerce strategy underpinned by digital transformation, Aramex wanted a way to provide customers with more accurate delivery time estimations, to improve the customer experience and increase efficiency in the delivery process. The solution needed to be secure and scalable, to cope with Aramex's growing customer base and data.

Solution

Using the RAMP platform, Inawisdom implemented a solution on AWS that uses predictive analysis and AI to calculate shipment transit times. The solution included a data lake, web API, DevOps environment and an AI DataOps environment. This involved the integration of several AWS services including Amazon Redshift, Amazon SageMaker, AWS CloudFormation, and Amazon CloudWatch.

Outcomes

In just 8 weeks, Inawisdom delivered an enterprise-grade solution for Aramex. This enabled a complete transformation of the customer experience, with delivery time predictions that were 74% more accurate than previous estimates. These predictions were embedded into Aramex's web and mobile applications, for easy self-service for customers. With access to real-time information and more accurate delivery schedules, customers made 40% fewer shipment-related calls to the contact centre, significantly improving efficiency.

> We have seen a 74 percent increase in the accuracy of our transit-time predictions because of the machine-learning models we developed on AWS with Inawisdom.

Mohammed Sleeq
Chief Digital Officer
Aramex

AWS ML and AI services

AWS has the broadest and deepest set of machine learning (ML) and artificial intelligence (AI) services for your business. On behalf of our tens of thousands of customers, we are focused on solving some of the toughest challenges that hold back machine learning adoption.

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

Machine learning 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 is a fully managed service that covers the entire machine learning workflow to label and prepare your data, choose an algorithm, train the model, tune and optimize it for deployment, make predictions, and take action. Your models get to production faster with much less effort and lower cost.

Artificial intelligence 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.

200+
ML features and services in the last 12 months

75%
Lower ML inference cost

10x
Better ML algorithm performance

1-click
Model training and deployment

About AWS

Learn how to put machine learning to work for your business, visit https://aws.amazon.com/ai/.

To learn more about Inawisdom, visit https://www.inawisdom.com and learn how you can unlock the value in your data..

partner
network

This message is produced and distributed by Inawisdom
Connexions Building, 159 Princes Street
Ipswich, Suffolk, IP1 1QJ
www.inawisdom.com | info@inawisdom.com | 020 3575 1337