Getting Started with Machine Learning in Retail

How retailers can improve customer experience, increase revenue, and boost operational efficiency using machine learning on AWS

Contents

“By 2024, 40 percent of the top 500 retailers will use AI-enabled decision-support to drive improvements in the new KPIs of omnichannel retail, including customer lifetime value (CLV), productivity, and profitability.”

Source: IDC

Introduction

Machine learning (ML) continues to draw more attention from organizations that need to make quick, accurate, and automated decisions based on the many data points available. Accuracy requires creating models that must be trained, validated, and tested using vast amounts of data before deploying them to a production environment.

For example, retailers need an automated way of making good product recommendations to help customers find relevant products and content more efficiently. These recommendation solutions leverage historical data, such as items searched for, browsed, selected, or bought to determine the relevance of items to an individual customer, thereby making it more likely the shopper will add the item to their cart.

The retail industry—especially ecommerce retailers—has compelling reasons to investigate and try ML. As in other sectors, however, retailers either aren’t aware of these opportunities or, more commonly, don’t have the expertise or resources to take advantage of them.

At JBS Custom Software Solutions (JBS Solutions), we’re committed to helping retailers overcome their hesitancy to tap into the opportunities that ML offers. Part of the problem is a lack of proper understanding of what ML is. Another is that there’s simply no one-size-fits-all approach to implementing ML in the enterprise, making it difficult to get started.

This whitepaper explores:

- What ML is and common misconceptions about it
- The most significant opportunities for retailers to boost the customer experience, revenue, and operational efficiencies
- Challenges in getting started
- Seven steps for selecting your first ML project

What is machine learning?

When people think of ML—or artificial intelligence (AI) for that matter—they may think of a silver bullet that will suddenly produce magical insights, solve all manner of problems, and make lots of money. They believe ML will tell them what they need to do or instantly give them some critical information they would not ordinarily have.

Those new to the technologies often use the terms ML and AI synonymously, but they are not the same. AI enables a machine to simulate human thinking and behavior. **ML is a type of artificial intelligence that uses data and algorithms to train systems how to learn for themselves and make predictions and decisions—continuously improving through experience automatically, without explicit programming.**

Indeed, once you reach a certain level of maturity with ML, it can provide insights that you can then use to automate processes, improve the customer experience, and make more accurate and timely decisions. However, getting started with ML is not quite as simple as flipping a switch, to say nothing of deploying, refining, expanding, and maturing it. First, you need to decide whether ML is the tool you need.

ML is a type of artificial intelligence that uses data and algorithms to train systems.

Misconceptions about machine learning

With the growing buzz around ML, it's tempting to assume that it should solve all our business inefficiencies faster and better than our existing people, processes, and technology. That isn't necessarily the case. ML isn't the answer to every problem. It may turn out that the answer or insight you need lies in something other than an ML solution.

Before arbitrarily deciding ML is the key to solving an issue, it's critical that you truly understand the specific question or problem you're trying to solve. For example, one JBS Solutions client was willing to make a significant monetary investment in an ML solution for insights into its sales pipeline. Although the client insisted on ML, it was quicker, easier, and far less expensive to provide a relatively simple sales report from an existing database.

If you do determine that the right approach is to incorporate ML into your processes, applications, or websites, the next misconception you may encounter is the answer to the question: **Where will you spend your time? It's not writing the code.** Creating the algorithm that leads to enlightenment, accuracy, and automation will actually be a smaller part of the project. Instead, if you're to be successful, you'll spend the bulk of your time in data acquisition and preparation, and then in the train-and-test cycle.

At its core, ML requires one or more models, but the **datasets to build, train, validate, and test the model are most important.** In many cases, the algorithms may have already been written for you. Take fraud detection, for example. There are well-defined fraud detection models out there and available in SaaS services like Amazon Fraud Detector. You can tweak the behavior of the black boxes with mere configuration changes, so there is little need for you to write your own, especially when you are just starting.

What makes an individual implementation and deployment of a given model successful is the data you make available.

Training data

With this dataset, you can fit the parameters of the model and then train it by examining the output of the model based on the input vector and compared with the expected outcome. You can then tune the parameters and their respective weights until the model produces the desired result.

Validation data

Independent of the training data, this is used to get an unbiased assessment of how well the model parameters were tuned during training. It can often identify holes in the training data, where it might not truly represent the population.

Test data

A final dataset provides another independent, unbiased assessment of the model and its parameters.

ML systems personify the computer science concept of “garbage in, garbage out,” coined by early IBM programmer and instructor George Fuechsel. Given nonsense data as input, a computer program will almost certainly produce a nonsensical result. Most ML models may be well-written sets of instructions, but if the data used to train them is inadequate, inaccurate, or simply unrepresentative of the actual population, you’ll find yourself back at the drawing board.

High-quality, representative data is what drives successful ML, so never underestimate the time and resources required for data acquisition and preparation for the training, validation, and final test data phases.

Is machine learning practical for the average enterprise?

There are many examples of ML applications across every industry—retail, financial services, entertainment, medicine and healthcare, manufacturing, and many others. And there have been many benefits. Among them:

- **Retail:** Thirty-five percent of all Amazon sales come from ML-driven product recommendations.
- **Entertainment:** Eighty percent of all Netflix programs are streamed due to the entertainment giant's personalized, ML-driven content recommendations.
- **Healthcare:** ML helps researchers diagnose more rapidly and accurately using medical diagnostic and digital medical imaging data (x-rays , MRIs, and CAT scans).
- **Financial, retail, and others:** ML can fight fraud in many ways, not only by detecting suspicious transactions but also by using sophisticated video, facial, and handwriting recognition algorithms.
- **Education:** ML can detect subtle changes in a student's performance and recommend tutoring, remediation, or counseling to increase their likelihood of staying on track.
- **Consumer-oriented apps:** ML even powers apps that can analyze data from video cameras and doorbells for security or the contents of your refrigerator to find the perfect recipe for you!

But is ML only for Fortune 500–level companies? From some of these examples, you may think that ML is great for big enterprises like Amazon, Netflix, and Goldman Sachs. They have large teams of IT personnel, programmers, database administrators, data engineers, and data scientists. Where does that leave typical enterprises or small and medium-sized businesses (SMBs)?

"Five to ten years ago, I would have said 'good luck' to most companies considering using ML," says Philip Horwitz, chief architect at JBS Solutions. "You're going to need an army of data scientists, PhDs, and all sorts of highly experienced, highly trained, and specialized personnel to help you develop scalable, bullet-proof, and productionized ML solutions."

A lot has changed since then. ML is becoming commoditized, where services and the underlying ML models (algorithms) are available as SaaS on cloud providers, such as Amazon Web Services (AWS). This commoditization has significantly lowered the barrier of entry, not only in terms of infrastructure but also in building the models themselves.

Now, with just one or two specialists, almost any organization can start incorporating ML into the enterprise. And they can get the same kinds of results that years ago would have required a multitude of specialists. Of course, you shouldn't expect to jump in and use all these tools without having any idea what you're doing. But with a couple of savvy programmers who understand your enterprise data, you can start implementing (and benefiting from) ML.

Let’s talk about retail

ML has the potential to transform enterprises in almost all industries. **In retail, there are opportunities to boost customer experience, revenue, and operational efficiencies.** Ordered from the least to the most complex, here are just a few areas where retailers of any size can apply ML to yield measurable benefits.

JBS Solutions values its partnerships with major cloud solution providers and their platforms—including AWS. Ninety percent of our engineers have extensive experience building and architecting AWS cloud-based solutions, many for more than 10 years. So, as we discuss various key opportunities where retailers can leverage ML, we’ll highlight some of the AWS services that make getting started even easier.

Fraud detection

In a connected world, detecting fraudulent transactions is no longer a matter of requesting a customer's ID at the checkout. Nor do the fraud detection systems that card-issuing banks have in place always catch everything. Cleaning and reconciling fraudulent transactions is time-consuming and expensive, which is why retailers need a solution of their own.

Traditional fraud detection mechanisms are usually rules-based. They are effective to a degree, but they present retailers (and their customers) with a few serious issues:

- False positives: Rules will often block genuine transactions by mistake.
- Fixed outcomes: Thresholds for rules can change over time, which can invalidate them.
- Maintenance: Rules (and the rule engine) can be challenging to manage and scale.

For ML, fraud detection is a "classification" problem, similar to spam detectors or loan default prediction. It can classify a transaction as legitimate or fraudulent based on amount, merchant, location, time, and other factors. ML can complement rules-based fraud detection to provide a more holistic fraud detection strategy.

ML benefits of fraud detection

- Good first ML solution for retailers
- Speed and scale
- Less expensive than more traditional rules-based systems
- More accurate and effective

Fraud detection

With tools like Amazon Fraud Detector, retailers can automate the build, train, and deploy lifecycle.
Additional benefits:

- Allows model customization based on the retailer's own datasets
- Leverages Amazon's 20+ years of ecommerce data for model development
- Avoids significant upfront costs
- Enhances model with custom rules that trigger actions
- Easily integrates into retail checkout flows

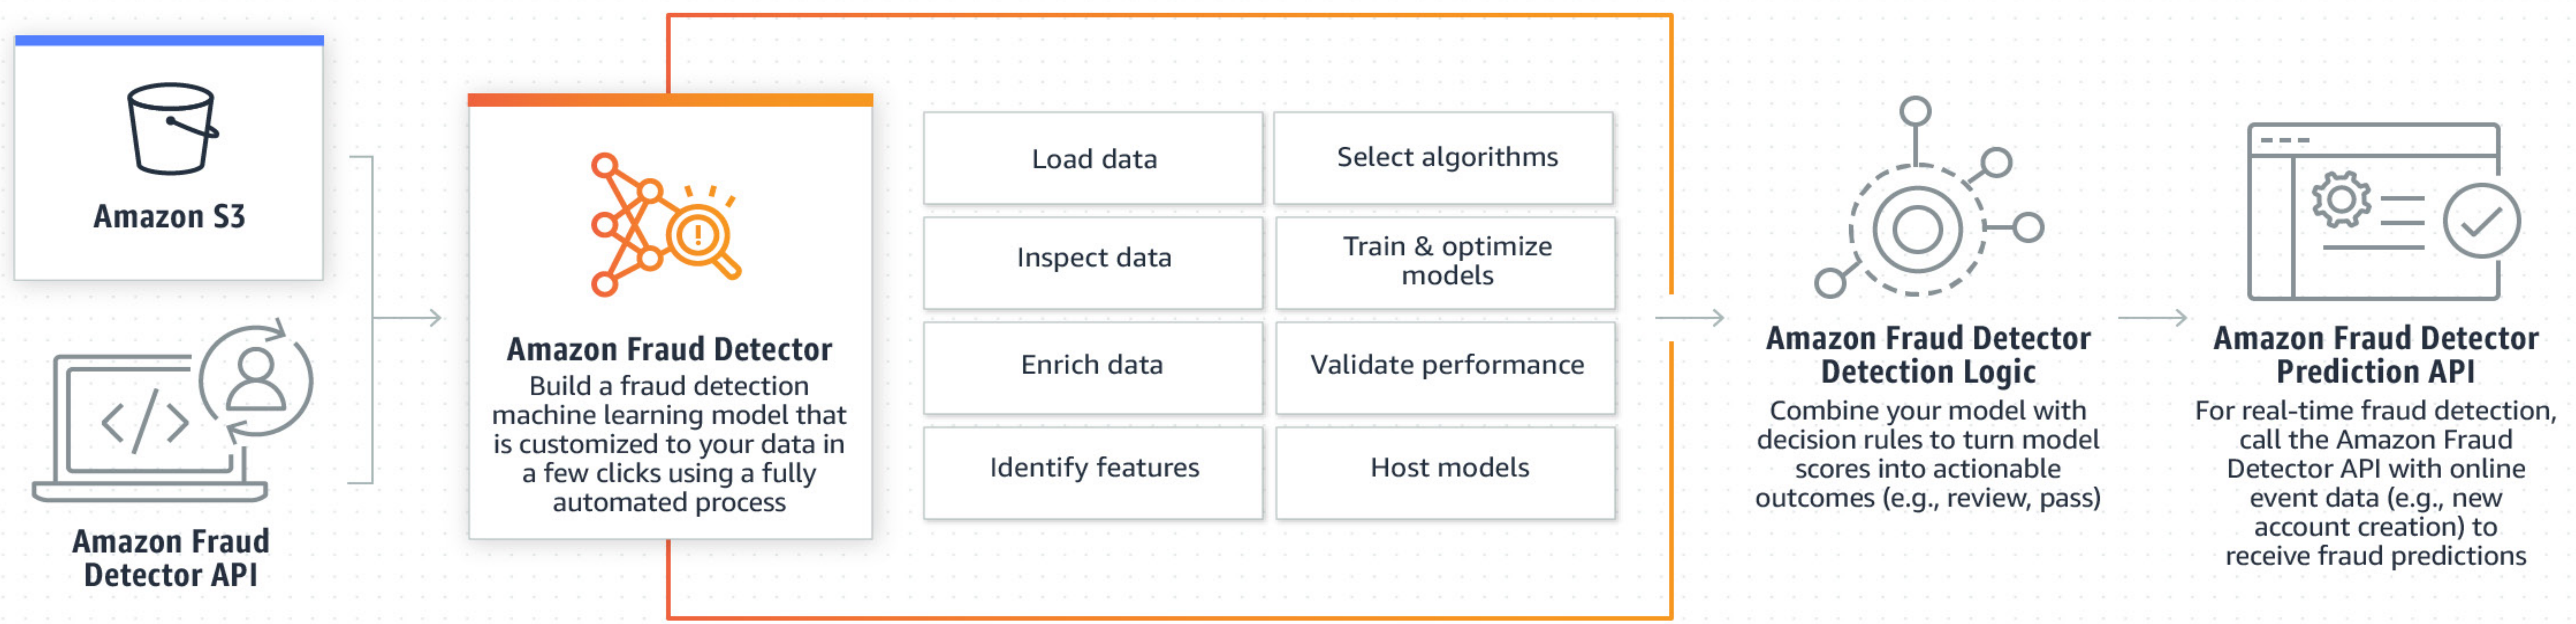

Buyer segmentation

By dividing customers into smaller groups based on shared characteristics, retailers can better tailor and target their marketing efforts to the appropriate segments.

Traditional segmentation uses demographics and other characteristics like age, gender, location, prior product purchases, and marital status. The results are useful, but this approach presents retailers with several problems:

- Categories are too broad/fixed: Not every 40-year-old likes the same thing.
- Easily outdated: People move, they marry or divorce, and their tastes change.
- Not dynamic: They can't take in-the-moment data (e.g., clickstream data) into account to maximize current relevance.

The key to optimal buyer segmentation is to leverage **customer behavioral data** in addition to demographics and other characteristics. This data might include recency, frequency, and monetary (RMF) data about a customer's purchases, plus real-time data both online (clickstream) and in the store (video and location). ML algorithms can divide buyers into groups that otherwise would be difficult or impossible to find via traditional means.

ML benefits of buyer segmentation

- Good first ML solution for retailers
- Less complex than typical supervised learning models (no training data required)
- Automatic and dynamic behavioral connections between customers (clusters)
- Connections/groupings always up to date

With tools like [Amazon SageMaker](), retailers get a wide variety of ML and deep learning algorithms that they can use for buyer segmentation, including:

- Fully managed service to build, train, and deploy ML models
- Out-of-the-box tools for creating and managing training datasets
- Automatic hyperparameter tuning
- Automatic model monitoring to determine effectiveness

Product recommendations

A good recommendation solution can make finding relevant products and content more efficient, save customers time, and result in more sales. According to a McKinsey study, 35 percent of Amazon sales come from ML-driven product recommendations. And ML-driven content recommendations account for 80 percent of what Netflix customers watch. The better the recommendations, the more (and faster) customers will buy.

Product recommendation solutions use a class of algorithms and techniques that suggest the relevance of items to an individual customer. These algorithms determine relevance based primarily on historical data, such as items searched for, browsed, selected, or bought. The presentation layer can then use relevance ranking to determine what product suggestions to display. The goal is to make it more likely for the shopper to take the next step—to watch another show or add the item to their cart.

There are two broad categories of algorithms for ML product recommendation solutions:

- **Collaborative filtering models** are based on past interactions between customers and target items. These algorithms work on the idea that historical data alone should be enough to make recommendations.
- **Content-based systems** use historical data and may include other detailed information about the customer (location, demographics, and behavior) and your product catalog (geographic availability, product metadata, rankings, and return rates). The idea is that the better you know your customers and your products, the better recommendation you can make.

Each approach has its benefits and drawbacks. The former requires fewer data points about a customer, but its rankings lack the quality and personalization of content-based solutions. The latter requires a much more sophisticated data pipeline.

While implementing a product recommendation solution is probably not the first ML project you should tackle, the benefits of a high-quality solution are apparent. (Amazon and Netflix—and their customers—would agree.)

Case study summary

Wine Access Modernizes and Sees 400% Site Traffic Growth Using AWS Partner JBS Solutions

Wine Access and JBS Solutions are using Amazon Personalize to create dynamic website content and deliver highly customized wine recommendations based on users' browsing activity. Wine retailer Wine Access modernized while saving costs and increasing performance by using a serverless architecture on AWS. The company migrated from on-premises servers to AWS, where AWS Retail Competency Partner JBS Solutions helped it transform its architecture on AWS Lambda. The new solution scaled automatically to meet a 400 percent increase in website traffic. The company experienced no outages, cut compute costs significantly, and decreased database costs by about 50 percent. Now Wine Access has an agile ecommerce solution that supports its pursuit of new business ideas and innovations, improves customer experience, saves hundreds of in-house work hours, and launches new revenue streams.

Read the full case study here »

Although a retail product recommendation is a complex ML solution, tools like Amazon Personalize can make the journey more manageable. Benefits include:

- Delivers out-of-the-box, real-time or batch recommendations on products or content
- Powered by the same algorithms as Amazon retail business
- Uses API-based interface to simplify integration into existing systems
- Automates the building, tuning, training, and deployment of recommendation models
- Handles the "cold start" problem of collaborative filtering (new users or no historical data)
- Incorporates user activity with user and product information in real time
- Integrates with websites, apps, SMS, and email marketing systems
- Includes user personalization/segmentation, similar item recommendations, and customized rankings

Inventory and supply chain optimization

Inventory and supply chain management has never been a simple problem because it involves so many moving parts in a retail organization. Retailers have become pretty good at forecasting, but without a real supply chain optimization solution in place, it's an educated guessing game. An ML solution can take away a lot of the guesswork.

A comprehensive inventory and supply chain optimization solution is one of the more complex ML initiatives a retailer can begin. However, it can also be one of the most rewarding. **A McKinsey study claims ML-enabled supply chain management can reduce forecasting errors and help scale back inventory by up to 50 percent and reduce lost sales by up to 65 percent.**

ML solutions for retail inventory and supply chain can analyze and leverage customer behavioral data to:

- More accurately predict near-future demand for optimal inventory levels
- Better understand how price changes impact sales
- Minimize ideal stock and optimize inventory storage
- Predict how cannibalization, the halo effect, and deferred demand are affected by the promotion of a particular item
- Incorporate more factors than just sales data to deliver much more accurate results
- Use third-party data in their ML solution (e.g., using weather data to forecast when you will need more products for family outings and cookouts)
- Leverage the power of complex algorithms, including random forests and deep learning (neural nets)

Despite the tremendous benefits, however, inventory and supply chain optimization is not the best candidate for a retailer just starting out with ML, unless it addresses only a small component of the overall system.

Amazon SageMaker is a great tool for those who want to incorporate ML solutions into their business—including integrating complex models into their supply chain optimization strategy.

The use cases above are only four of many possible applications of ML to optimize retail sales, marketing, and operations. Other retail applications for ML include:

- Pricing optimizations
- Optimized search results for improved product discovery
- Image recognition for automatic number plate recognition for curbside pickup
- Computer vision for frictionless checkout solutions

Challenges retailers face starting with machine learning

For organizations new to ML, getting a solution into production usually takes longer and requires far more resources than expected—if they make it to production at all. According to Forbes, a **shocking 87 percent of ML models never make it out of the experimental stage and into production**.

Why? Organizations tend to underestimate the time and effort required for data acquisition and preparation, not to mention the train-and-test cycles. Understanding the problem and knowing the questions you want to answer are a close second. And there are other common pitfalls for retailers to avoid.

Data bias

Besides having a large enough dataset to train (and improve) your models, your data must accurately represent the audience you're trying to serve. **Bias can be introduced into the dataset in many ways**, or it may be there from the beginning. If this is the case, your model may not give the answer you want or need.

Retailers should recognize that a certain amount of bias is inevitable and tune the model using hyperparameters to overcome that bias. Also, depending on the quantitative maturity of your organization, there are other more sophisticated statistical methods of removing bias you can use.

Insufficient data collection architecture/pipeline

Sometimes the problem a retailer is trying to solve simply requires more data than they are collecting. For example, if you are trying to provide in-the-moment product recommendations to an online customer, you need a data pipeline to collect clickstream data. This data includes the product they are looking at, how long they've spent on the page, how they are scrolling, and much more.

You can provision servers with all the ML services you like, but until you provide your model with the data it needs to solve the problem, it's not going to give you the answers you need. To provide the necessary data requires building a real-time system to grab, ingest, process, and transform it, and then inject it into the ML system to produce the recommendations.

Limiting the data to transactional sales data

Sometimes sales data—which many retail ML models rely on—isn't enough to produce the accuracy of results you need. **If you can augment static sales data with customer behavioral data, you can make better recommendations.** For example, if you detect that an online customer went to the same product page three times without clicking "Add to Cart," perhaps your model can recommend something to entice the buyer to take the final step.

Speaking of product recommendations, you can also leverage third-party or local data to enhance the quality and effectiveness of the results. For example, local weather data isn't stored in most order entry systems, but using it could boost sales for a clothing store or sporting goods ecommerce site. The lack of this type of "extra" data might not cause a model to fail, but it can limit its effectiveness.

Lack of the right personnel and expertise

Cloud providers like AWS offer prebuilt, ready-to-use ML services, which can relieve the need for software developers to build algorithms from scratch. However, **any company that wants to leverage ML has to deal with data engineering and data architecture.** On top of that, they may simply lack familiarity with planning, executing, and deploying an ML solution. Hiring a technology partner to train up your staff to own and maintain the solution can fill in these gaps during the project.

Tackling the most complex problem first

It's tempting for retail IT shops to tackle the most significant problem first to demonstrate this new technology's value. But remember that "new" is the operative word. **Start small and simple** until you've developed some experience with ML. For example, with its single yes-no answer, fraud detection is far simpler than creating, testing, and deploying a product recommendations solution.

Building on a shaky application foundation

If your application architecture isn't well thought out, you're going to run into issues. Integrating ML into a poorly designed piece of software is like building a house on a shaky foundation. Figuring out the source of problems—is it the old code or the new?—is difficult in the best case and impossible or cost-prohibitive in the worst. If things do work, you'll run into performance issues, poor reliability, and lack of scalability. **Building new features—including adding ML—is simpler and more organic with proper application architecture.**

The once-and-done assumption

It's not recommended to assume that you are "done" once an ML solution goes into production. **You must continuously monitor the model's output**, retune it as needed, and change it altogether if you cannot correct drift using hyperparameters. Be aware that a significant shift in customer behavior—such as that caused by a global pandemic—can render an otherwise accurate model inaccurate without adjustment or outright change.

Seven steps for selecting your first machine learning project

Given those challenges, exactly how should a retailer start to take advantage of ML?

First, understand that **there are no one-size-fits-all solutions** a retailer can use. All retailers have some similar challenges, yet each one has different requirements for how to solve them. Every retailer does checkout, but each may require different kinds of data. And despite a certain commonality among datasets—products, orders, and customers—the data resides in a variety of systems from different vendors, some open source and some proprietary, some standardized and others entirely customized.

While there is a baseline of functionality and data that every retailer shares, there are enough differentiating factors to not assume a model that works for one will work for another. Nor will the same data points make sense for both a large retailer and a small one.

Many retail ML projects fail, even for large retailers, because they try to do too much too fast. The best approach is to scope a minor problem to solve or a specific question you want ML to answer. Rather than dive in headfirst, here are seven steps you should take:

- **1. Select a trial use case** to lower the barrier to entry. For example, of the four use cases we described earlier, fraud detection and buyer segmentation are the simplest. Fraud detection addresses a simple question with a simple answer. Buyer segmentation involves unsupervised ML with little to no training required—you simply throw the data at it and analyze the patterns that come out.

- **2. Wherever possible, leverage SaaS services** such as those from AWS that provide numerous retail ML algorithms right out of the box. This lets you avoid coding and lets you concentrate your resources and effort on the data prep, training, and testing stages.

- **3. Start with a small dataset**, and then prep it, transform it, and run it through the model. Go through the entire lifecycle to get a feel for how this process works.

- **4. Examine the results.** Did you get the answer you expected? If not, adjust the model's hyperparameters to determine if you can increase the accuracy of the output.

- **5. Experiment with different models** to get an idea of which ones work best for your situation and what kind of data you need.

- **6. Consider what other data you might bring in** from your data pipeline to make the results more accurate and robust.

- **7. Develop a level of familiarity and confidence with the technology. Then you can think bigger.** You can bring in larger datasets, work with new parameters, and add controls for scaling and fault tolerance—but only after getting used to what you're doing.

Ultimately, you can have something in place in which you feel confident when it goes into production. The solution may not be 100 percent successful your first time around, but it should be more robust than if you tried to answer the question without ML.

And remember, ML is not a once-and-done type of project. You should continually monitor and measure the output of the model to ensure it continues to function as expected. You can use sophisticated statistical methods to analyze model performance, but it's often simpler to just look at the data. For example, with a market segmentation ML solution, campaign metrics from your marketing automation system can prove whether your campaigns are generating more sales leads than before.

Let’s get started

When it comes to retail, ML is no longer just a curious technology experiment. There’s no question of the benefits that come with effectively applying ML to a retailer’s ecommerce site and its business and operational systems—the key word being “effectively.” Ecommerce giants have already heavily tapped into this resource, as have many retail chains that operate both online and brick-and-mortar operations. However, unlike large enterprise brands, most retailers don’t have vast programmer resources with numerous data engineers and data scientists to lean on. Even retailers with sizable IT departments often don’t have a sufficient team of ML specialists or the ability to hire them.

Yet, retailers that don’t wade into the ML pond may soon find themselves falling behind the competition. Customers want richer, more innovative, personalized online shopping experiences that make finding what they want faster and more efficient. Retailers need to deliver those experiences, boost revenue, and increase operating efficiencies—all at the same time.

ML can help by:

- Enabling retailers to present product recommendations that are more relevant to each shopper, boosting both customer satisfaction and sales revenue
- Providing more accurate fraud detection, saving retailers money while reducing false positives and the accompanying customer frustration
- Segmenting retail customers using far more than demographics so the retailer can optimize its marketing campaigns (and budget) while showing customers more of what they want
- Optimizing retail inventory and supply chain management, reducing waste while ensuring that stores have the merchandise customers want, when they want it

The fact is that retailers no longer have to hire a staff of ML experts and data engineers. AWS provides ML services that retailers can use right out of the box.

By leveraging a trusted technology partner, you can also accelerate your ML initiatives—and the benefits you and your customers derive from them. A technology partner can help plan, execute, and deploy the solution, not to mention training your existing staff to use and maintain it. You'll still need a couple of savvy programmers that understand your company's retail systems and data so they can create a data pipeline that feeds the ML algorithms. (The right partner technology partner can get them up to speed with how to do that, too.)

By following the steps we've outlined in this paper, avoiding the pitfalls we described, and selecting a savvy partner, any retailer can start enjoying the benefits of ML.

JBS Solutions has years of experience delivering retail ML solutions in the cloud leveraging AWS and using AWS services such as Amazon Fraud Detector, Amazon Personalize, Amazon Forecast , and Amazon SageMaker. For more information on how we can apply this experience to optimize the customer experience, revenue, and operational efficiencies for retailers of any size, contact us today.

About JBS Custom Software Solutions

JBS Custom Software Solutions has been delivering leading software and application-based solutions for our clients for over 20 years. It is important to us, as it is to you, to deliver on time and on budget. Most of the custom development that we've delivered over the years is to support proprietary business models and innovative technological solutions. This is the core of our business, and our experience ranges from conception (design, R&D) to delivery (code complete, deployment). Our rich history is filled with successful proprietary applications and systems developed specifically for our partners, ranging in size from small back-office utilities to large, big data systems with more than 1,000 users.

Because no two projects are alike, we invest considerable time upfront on projects to deeply understand your business challenges and goals. In developing business software solutions for our partners, we focus on your culture, in-house team capabilities, and strategic business goals, always dedicated to delivering appropriate, profit-focused solutions.

JBS Solutions is an **AWS Retail Competency Partner** with validated technical expertise and proven customer experience in delivering retail specific solutions on AWS for retail customers.

Contact us today to learn about the JBS Approach. Our teams, comprised entirely of senior-level experts deployed alongside your existing resources, create a more efficient process to deliver a solution that meets your unique business objectives with a higher success rate than the industry average—and produces a better, faster return on your investment.

powered by aws