TPG PARTNERS
Optimizing Supply Chains Globally

TPG PARTNERS TECHNOLOGY WHITE PAPER

Agile Inventory Management (AIM™)

TPG Partners – Dallas / Fort Worth – www.TPG-partners.com

September 20, 2021

Table of Contents

Executive Summary

Amazon. A name that has struck fear and awe into retailers for two decades - and had so much impact on eCommerce shopping expectations that there's a name for it: The Amazon Effect. At its heart, it means a 'frictionless' shopping experience; low prices, intelligent and automatic displays of competing and parallel products, and immediate gratification in terms of shipping. And for retailers that haven't found a way to find the agility necessary to compete, it's meant record brick-and-mortar closures and outright bankruptcies of some of America's classic retailers.

But there's a catch in those declining brick-front numbers; for some retailers (including Amazon itself) there's a movement back into physical stores. Why? Because that Amazon Experience has continued to evolve, and now fast shipping isn't enough – buyers want to order online and find their product available at a local store for immediate pickup. Today. So, the new mantra is having a seamless synergy between your eCommerce presence and your physical stores.

Doing that isn't simple – and failure is expensive, perhaps terminally so. To succeed in that synergistic model, inventory management needs to be incredibly agile – somehow, we need to know ahead of time what the demand for a particular product will be at a particular store, then use that information on a timely basis to get products to the right place at the right time. To say this sounds like a tall order is an understatement!

TPG Partners' Agile Inventory Management (AIM™) product will provide just such a solution. Using state-of-the-art Big Data and a Machine Learning predictive model that is configured to a particular retailer's products, customers and industry profile, AIM™ can predict future orders with high accuracy – and on a timeframe that provides the means to move product to where it needs to be with minimal overhead. And the beauty of a Machine Learning system is that it continues to evaluate its own model, adapting to changing conditions and trends. In effect, AIM™ continues to learn.

AIM™ is based on a highly successful model developed by experts in the field – a model that our own designers – in partnership with participating retailers – will continue to evolve through a series of trials.

The system will provide for exhaustive display and output capability, allowing easily configurable feeds into inventory and ordering systems, as well as browser-based inquiry and data views for humans. This allows a retailer to generate and view dashboards and reports describing product and store performance and customer shopping behaviour.

Implementing AIM™ is designed to be as non-intrusive as possible, with both data storage and Black Box components staged in the Cloud, providing for scalable resources as demand grows. No need for a bigger IT footprint – focus on the job of selling and let AIM™ provide you with the means to be successful.

Introduction

Over the past two decades, Amazon has transformed itself from a fledging eCommerce player to an overwhelming global juggernaut that has affected, either directly or indirectly, virtually every retailer on the planet. In doing so, it has changed the way people think about shopping, introduced leading edge technologies, impacted pricing, and put intense pressure on every one of its competitors.

Need more? By the end of 2018, Amazon alone will account for *half* of all retail growth in the United States. It will account for one third of all on-line purchases. Let's repeat that – one half and one third of a 2.5 *trillion* dollar industry.

More numbers? Eighty million shoppers in the Unites States have purchased Amazon Prime memberships. That means one in every four Americans has put their hard-earned money not directly into a product but into a de facto pledge of buying loyalty!

But arguably its greatest mark has been made on people. The impact of low prices, extensive (and proactive) comparison shopping and express shipping have changed the expectations of every shopper. There's even a name for it: the Amazon Effect. Thanks to Amazon, shoppers now insist on a 'frictionless' buying experience – from the very beginning of the process (show me what I want before I tell you) to the very end (I want it *now*) – the expectations that a potential buyer of your product has are phenomenally high. And of course, they expect all that for a price that leaves you, as a retailer, with an alarmingly small margin.

Through price and technology, Amazon has – in twenty short years – transformed the industry. And what has this transformation meant to retailers?

The simple answer for many is devastation. Toys-R-Us: bankrupt. Radio Shack: bankrupt. Bon-Ton, Claire's, Nine West: bankrupt. Sears: closing 103 stores. Sam's Club: closing 63 stores. Macy's: closing 11 stores. And most of these closures are at 100,000 square foot-plus footprint locations. In fact, more than seventy million square feet of retail space is expected to close in 2018 alone. Looking at this from the perspective of year-to-year comps, it's estimated that between 30,000 and 80,000 stores will need to be closed by 2025 in order to maintain even low-single digit comps at remaining locations.

What's a retailer to do? Well, obviously, develop an effective, compelling eCommerce presence. Oh, and while you're at it, buy or develop the technological expertise to provide some of that 'frictionless' buying experience that the Amazon Effect has inculcated into your prospective customers. But what about your brick and mortar, all the investment and equity in the infrastructure that has been your traditional sales locus; are they destined to become little more than warehouses? Or should you just find a buyer?

The fact is, it's not that simple.

A new trend is developing, a counter-movement to closing physical stores, to abandoning a brick-and-mortar presence. There are a number of retailers rushing to fill those thousands of square feet of empty mall space left behind.

A few examples:

TJX Companies – aggressively expanding into an additional 1200 stores

Ross Stores – currently plans expansions amounting to 2500 new physical locations

Target – opened 25 'small format' locations in 2017, plans 35 more in 2018

Kohl's – investing in resizing current stores, adding a dozen 'small format' locations

Nordstrom's – Opening dozens of 'Nordstrom Local' locations with no inventory

Five Below – opened 100 new stores in 2017, planning similar openings in 2018

In the current environment of draconian closures and sagging brick-and-mortar sales, how can these companies be expanding *into* physical locations?

Here's the reality – what Amazon began twenty years ago continues to evolve. 'Frictionless' buying now often means no delivery at all. Instant gratification. Buyers want order on line and pick their items up at a store nearby. Today. Or perhaps they go to the same store and handle two different styles of the product they've viewed online.

So, the traditional brick-and-mortar is, in the new world of retail, no dinosaur at all. In fact, it becomes a critical half of an integrated whole; it partners with the retailer's eCommerce presence to provide a seamless buying experience for its customers. A frictionless experience, in which the buyer can visit a local store, perhaps see an array of products tailored to the locality (no swimsuits to look at in the Anchorage store), and have whatever they've ordered there waiting for them.

And Amazon? Moving into brick-and-mortar themselves. Whole Foods. An expanding physical bookstore presence. And - a small-footprint concept called Amazon Locker. Order a product; go pick it up a Locker near you. Sound familiar?

It all makes sense. There's a role – an important one – for brick-and-mortar locations after all.

But to achieve this seamlessly integrated frictionless environment, there are monumental challenges.

The Challenge

Consider the goals – we need knowledge of our products, our sales, our customers, the demographics we operate within, trends in our industry and a whole host of seemingly random variables affecting customer buying patterns. We need to be able to combine all those data into a cohesive – and reliable – calculation that predicts our inventory replenishment in unprecedented detail. And by the way, the prediction needs to be sufficiently timely – and frequent - to allow us to get shipments in specific stores 'just in time'.

What challenges do we face?

Challenge	Explanation
Having enough Information	Knowledge is power. Every data point regarding our potential buyers, our industry trends, demographics of our store locations, weather conditions and forecasts, events in the area, online activity, and dozens of other seemingly trivial data points - all of it combined holds the key.
Understanding what's Important	Even when we gather all these data points, how do we know what relationships between them are important? How could any human evaluate all of these potential conditions and accurately forecast what is likely to be ordered in a specific location?
Adaptability as Conditions Change	Even if we overcome the two challenges above, life is never static – conditions change. We need to be able to instantly adapt to changes in the relative importance of all those data points, and use that to continue to make accurate predictions of inventory at each location.

These challenges may seem insurmountable. While customer, product and sales data may seem obvious, and your eCommerce site is probably tracking and storing online browsing and buying, the rest of it may seem elusive. And even if you knew where to get it, what would you do with it?

Could a human possibly evaluate hundreds of data points, ranking their importance in relationship to each other and using those relationships to make detailed predictions? Sure, we'd probably know not to replenish our Texas stores with wool sweaters in July, but in today's retail world, we need to go far, far beyond that.

The Solution

The challenges are significant, but the solution to meet them lies within two powerful domains of modern information technology: Big Data and Machine Learning.

Big Data	Big Data means much more than simply "lots of it." It signifies the ability to acquire and store different types of data within the same repository – text, numbers, dates, even images! Even more importantly, it involves the ability to "clean" data coming in so that it can be used reliably (and even fill in missing data), and carries with it the ability to relate different types of data to one another. Beyond eCommerce data, and information we can obtain from 3rd-party providers, what are some examples of important information that we, as retailers, can begin acquiring to enhance our replenishment and pricing capabilities? • In-store video traffic analysis • POS logs • Beacon traffic • Product sensors
Machine Learning	Wikipedia tells us that Machine Learning is: "a subset of artificial intelligence in the field of computer science that often uses statistical techniques to give computers the ability to 'learn' with data, without being explicitly programmed." OK, what does that mean? Without Machine Learning, a human needs to know the solution in order to program it. Within Machine Learning, a human is simply providing the software with tools for obtaining facts on its own. The software learns by using these statistical tools to exhaustively examine hundreds or thousands of data points (from our Big Data), looping through *millions* of times, evaluating the impact of tiny changes in relationships on known optimal outcomes (like perfect replenishment orders!) As data points change, the software continues to learn.

If there's power here, there's also complexity. Developing an integrated solution to achieve the kind of agility that retailers need today to acquire that 'integrated whole' that combines brick-and-mortar with eCommerce is not within the scope of a typical retail IT department.

That's where TPG Partners' Agile Inventory Management (AIM™) software enters the picture. Leveraging recent research in Machine Learning studies at several American universities, it will offer an easily implemented, yet highly sophisticated solution – providing everything a retailer needs to achieve the goal of just-in-time replenishment that anticipates ordering behaviour on a product-by-product, store-by-store basis.

Here's a quick introduction to the capabilities we're designing:

Feature	Detail
A Complete package	Big Data acquisition. Data cleansing. Data storage. Machine Learning. Testing. Delivery of outcomes. There are a lot of moving parts in a Machine Learning system. AIM™ provides a comprehensive solution that includes every aspect, leaving nothing to chance and nothing to add.
Easy Integration and Use	Even the best of solutions, if they present challenges to implement and/or use, can have their usefulness compromised. AIM™ is designed to utilize contemporary web services architecture and cloud-based storage and functionality, putting the smallest possible demand on your IT staff and systems, and a minimal footprint on your resources.
Each Implementation Unique	Most software packages offer a single solution that all clients must be shoe-horned to fit within. AIM™ is designed from the ground up to be implemented as a unique solution for each customer. This means your data, your customers, your products and your demographics will drive the outcomes – not some generic standard.
Designed to Keep Learning	Once AIM™ is implemented and delivering, it will automatically continue to learn. As buying patterns change, as industry trends develop and fade, and demographics of the population around your stores change, AIM™ will continue to re-evaluate the data, refining and altering its learned behavior to provide constantly refreshed intelligence.

None of these technologies, taken as an area of focus within Computer Science, is brand new. From Big Data to using Machine Learning in retail, the capabilities provided by AIM™ are all individually well established and completely proven.

What AIM™ does uniquely is focus these individual technologies on the specific problem of providing for an agile inventory – utilizing them to drive intelligent replenishment at a level of insight, speed and detail never seen before. And yes, that can include pricing recommendations as well, if desired!

Benefits

If the capabilities of TPG Partners' AIM™ design are comprehensive, the benefits they deliver are equally so. Remember the goal: provide for a cohesive integration between the retailer's online presence and his brick-and-motor. Accomplish that by knowing our customers, products and buying patterns so well that we can predict in advance where and what we need stocked in our stores.

AIM™'s benefits provide the means for that to be accomplished.

Know your Customers	The insights provided by AIM™ analytics as your repository of Big Data is churned to reveal demand patterns will be of use to more than the Machine Learning statistical routines – they'll also be of immense use to your marketing team, as AIM™ provides web-based views into the factors that drive an agile inventory.
Optimize Inventory / Out of Stock Management	Once you begin viewing inventory as a dynamic, short-timeframe mechanism used to support specific products, customers and store locations, you will be able to optimize inventory management. Pre-positioning warehouse contents, just-in-time delivery, and minimizing shipping and handling overhead are just a few of the benefits that come from an optimized inventory. One of the most immediate benefits from AIM™ is reducing the occurrence of out of stock inventory.
Intelligent Replenishment	How quickly would you like to be able to respond to anticipated orders at the specific product and store level? Weekly? Monthly? No problem with AIM™! Make sure that when your best customers log into your web site hoping to find what they want at their local store, that product is already there waiting. And by the way, it wouldn't hurt to have them guided to similar products that they may be interested in looking at when they go pick their order up.
Respond in Time	Things change. People change. Areas change. Trends start and stop. If Machine Learning were static, it would soon be out of touch. But AIM™ is designed to keep on learning, continually monitoring, testing and updating itself so that your inventory not sonly stays agile, it stays accurate and timely.
Optimize Pricing	As an optional aspect of the kinds of agile inventory management capabilities that AIM™ brings to the table, it's also able to recommend optimal pricing for every product and store in your environment, helping you not only with availability, but with competitiveness.

Specifications – Core Architecture

AIM™ is designed as a system of layers, with each layer carrying out a role in the overall processing. The layers themselves are built as a collection of modules that communicate with each other via service calls. The goal is to provide for a lightweight system that is highly scalable, easily updated and supportive of deep customization.

The software comprising each layer is designed to be developed via a combination of languages and utilities, with the goal of utilizing the best tool for the respective function being performed.

The layers of the AIM™ Architecture:

Data Layer	It is data that drives the system – without abundant, accurate and clean data there can be no analytics and no prediction. The role of the data layer is to provide the means to enter various types of data (in both real-time and batch modes), evaluate and cleanse it, and store it in a repository in a fashion that ensures its usability by the Model. Data are provided both by the customer (for things like order history, customer information, products, etc.) and by third-party providers for things like regional demographics, industry trends, weather, etc.
Model Layer	The Model is the heart of the AIM™ system – where the data analytics are carried out and the Machine Learning takes place. The term 'model' refers to the highly individualized analytical approach that is used to establish the relationship values among the many data points being considered. The outcomes from the model are the detailed predictions of orders for stores, by product, for specific time periods. These are what drive replenishment. Learning itself is a separate process that the model utilizes to 'train' itself, establishing the foundation that ultimately processes data and makes the desired predictions.
Reporting Layer	Reporting provides access for delivering the outcomes (the predictions for replenishment), including sophisticated set of data visualization and the ability to carry out 'what if' scenarios that provide the means for the customer to evaluate things like pricing and stock changes. These deliverables take a variety of forms, from batch-based file output for integration with ordering systems to human interactive web-based applications for visualization and model comparisons.

The Data Layer

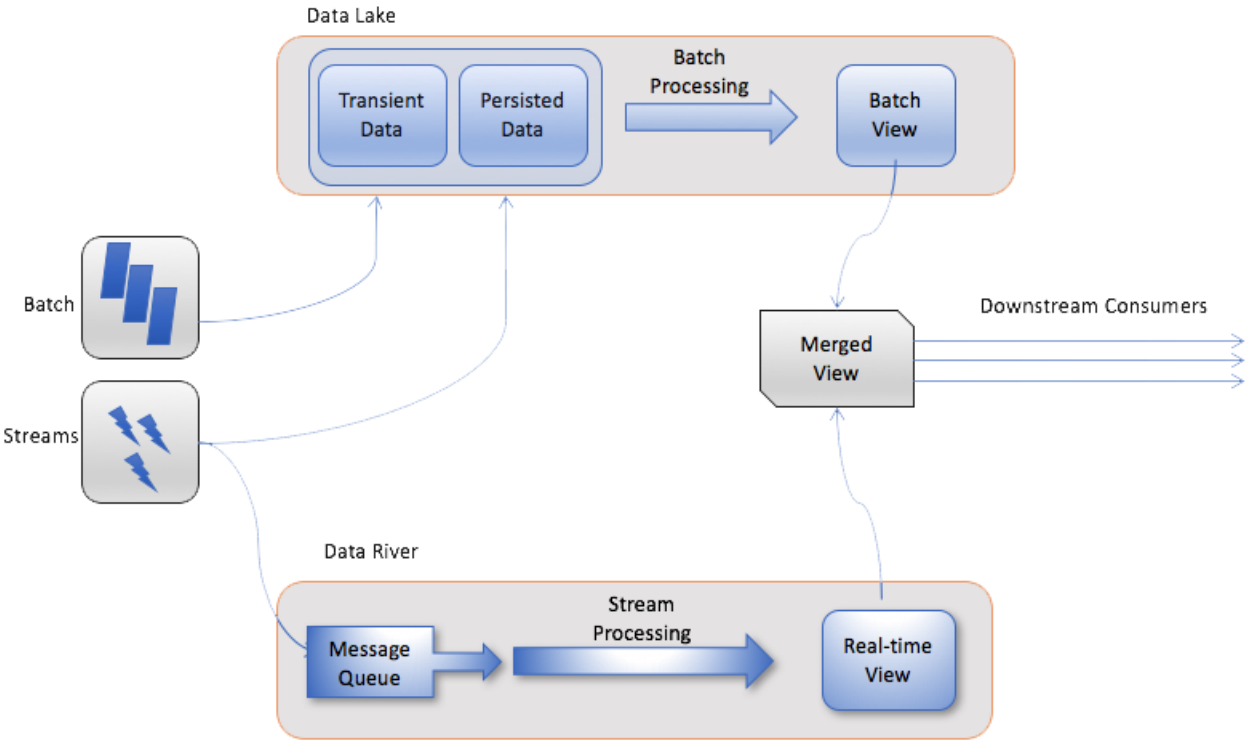

Key elements:

- Data input for both stream and batch
- Big Data capability
- Data processing operations include evaluation and cleansing
- A 'merged view' of data for the Model and Reporting

The goal of the data layer is to provide a comprehensive ability to handle any and all manner of data, provide the sophistication necessary to evaluate it for reasonableness and clean it where necessary, then persist those various data in a form that directly provides the data points to the Model.

The internal architecture is designed to be table-driven, meaning that configuration changes for each customer require no software customizations at all, and are easily altered as customer environments and needs change.

The Model Layer

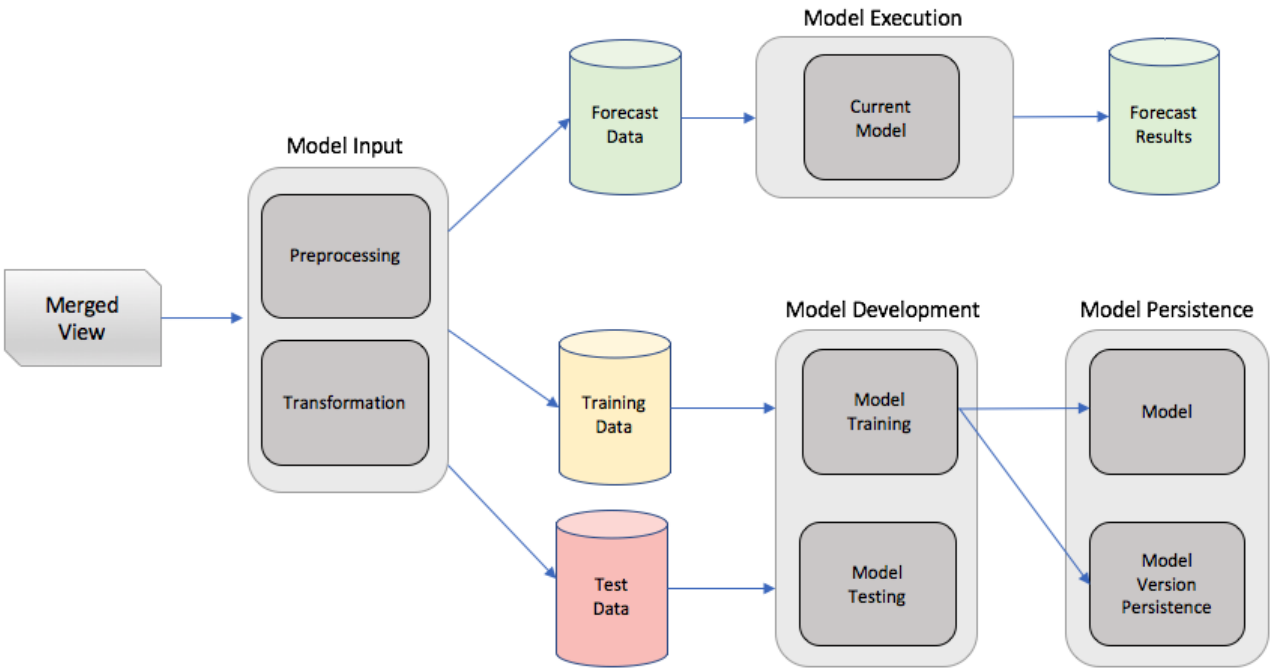

Key elements:

- Configurable pre-processing and Transformation of incoming data
- Training and re-Training of Model
- Creation of multiple models for comparative results

Note that the diagram above is a high-level, logical representation of the general functionality being provided by the Model Layer.

The Model itself comprises a very complex set of software routines that need to carry out a great variety of functions, consisting primarily of highly sophisticated statistical algorithms. The Model has to be able to train itself, store the iterations of itself for comparison, continually test its own performance, and provide the means for human beings to tune it.

AIM™ uses a complex approach that is characterized by its ability to evaluate a wide variety of seemingly unrelated data and discover relationships between them.

The Reporting Layer

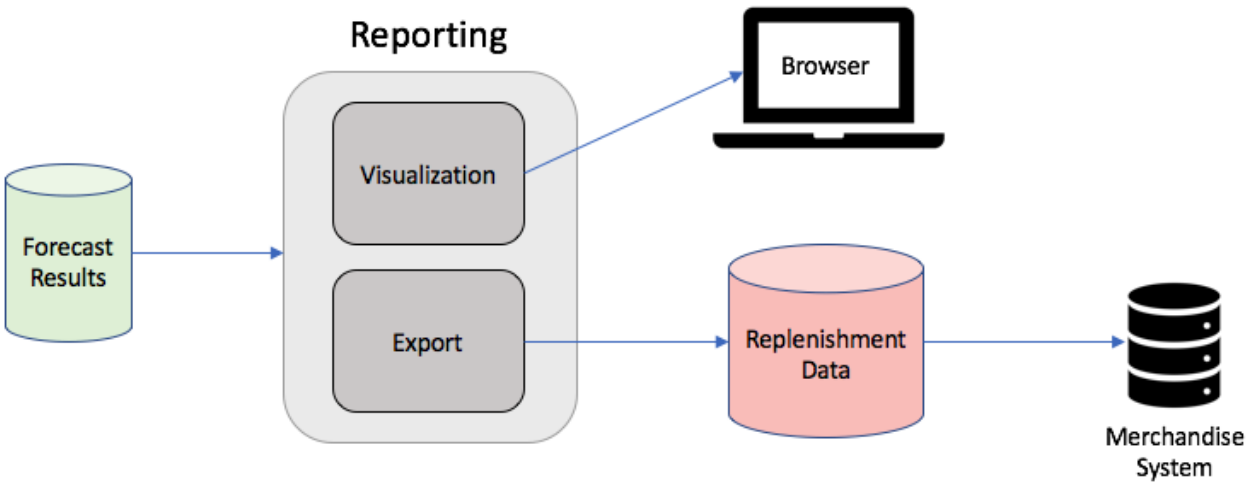

Key elements:

- Visualization via a web browser for data examination
- On-demand comparisons of different models
- Export data for Merchandise Systems

The Reporting Layer provides for two fundamental deliverables; first, a configurable output channel for conveying the outcomes from the model to the downstream systems that will utilize them. Typically, this would be replenishment data going to a merchandise management system.

Additionally, the system provides a sophisticated set of queries, conveyed through a web browser, that allows a human user to view and evaluate data and models, as well as run comparisons between different versions of the model.

Specifications – Machine Learning

The specifics of the Machine Learning algorithms that will yield highly accurate predictive results can only be determined through research and experimentation. There are a number of possible algorithms and variations of algorithms that appear promising based on current research. This list includes: Deep Neural Networks, Support Vector Machines, Genetic Algorithms, Convolutional Neural Networks, and EM Optimization. In addition, the predictive power of multiple machine learning algorithms can be combined using a technique known as ensemble learning.

Specifically, the design of a Deep Neural Network requires knowledge of the input data, desired output variables, and the characteristics of the network, which includes the number of hidden layers and number of nodes in those layers. The nodes operate in a similar manner to human neurons, accepting inputs at the "synapses" and "firing" when the correct threshold of inputs is met.

A Deep Neural Network utilizes training data to create the parameters used by the nodes to predict results. Training data consists not only of all prescribed input data, but also the correct final output. The model is trained to achieve correct results with this data by using what is called back-propagation.

Once successfully trained, testing data is then used to determine the success rate of the predictive capabilities of the model. Testing data contain all the inputs and the correct output that are compared to the outputs generated by the trained model.

This approach allows for multiple versions of the model, which could have different inputs, and different internal structures. Differing models can then be compared to determine the one with optimal results, based on available data.

Data that can prove useful to a model may not always be intuitive. For example, the possible inputs for a model that would predict the sales of a specific item at a specific store for a given week might include:

- Date of the week
- Item
- Price for that week
- Holiday for that week
- Weather for that week
- Community Events for that week
- Store Zip Code
- Store area demographics
- Holiday Days in Week last 52 weeks
- Community Event in week last 52 weeks
- Weather for that Zip Code last 52 weeks

- Number of items sold last 52 weeks
- Number of lost sales last 52 weeks
- Price of item sold last 52 weeks
- Item Placement in Store last 52 weeks

The desired output would be a prediction of the volume of orders for that item for the specified week. The result of the deep learning process results in a predictive procedure which has a high degree of accuracy, but – especially in a Deep Neural Network - the final procedure itself may be difficult to articulate in English. While this can seem opaque (the reason such systems are called 'black boxes'), it means that the network itself captures subtleties and nuances in the data that general pricing rules created by human experts couldn't begin to articulate. The diagram below illustrates a Deep Neural Network with a large number of input nodes, two hidden layers and a large number of output nodes. The output of each node is passed as input to each node in the next layer.

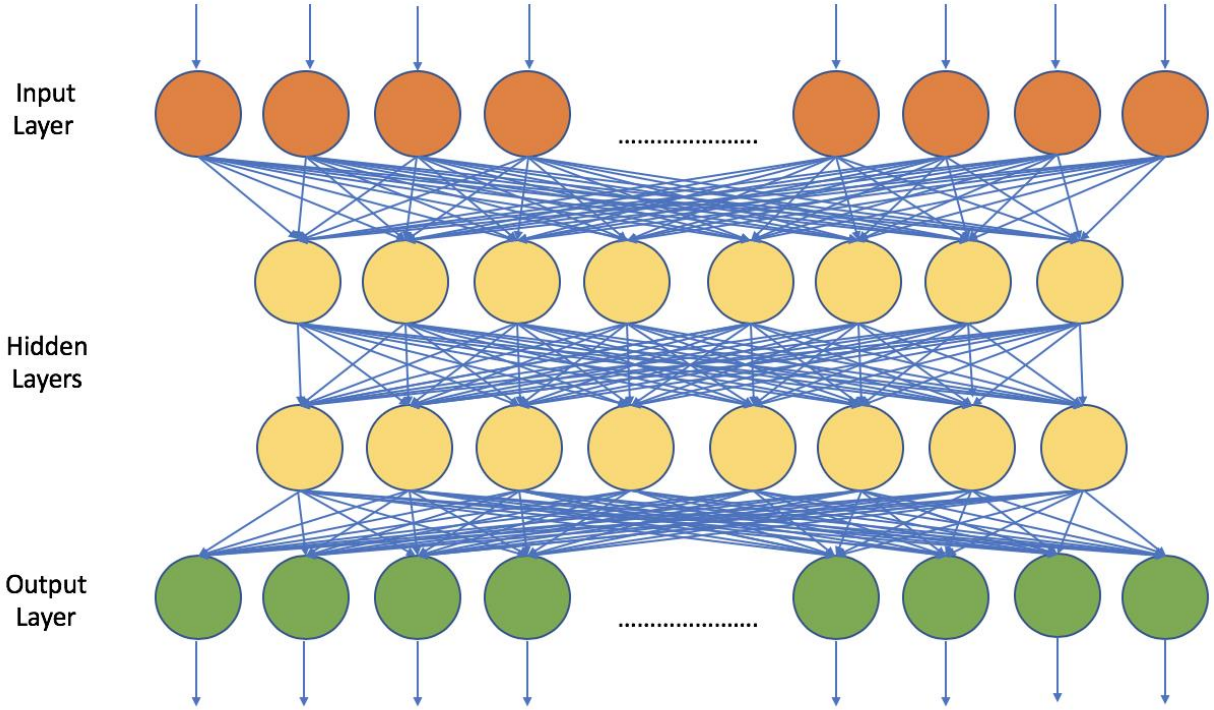

Specifications – Deployment

The AIM™ system is designed for deployment as a cloud-based package for two primary reasons:

- To take advantage of the inherent scalability of cloud-based resources, both in terms of the accumulation of large amounts of data, as well as scaling of processing power. Deep neural networks are resource intensive applications and the ability of today's cloud providers to scale seamlessly is an important consideration in deployment.
- AIM™ is intended to represent the lowest possible intrusion into a customer's IT environment. Through the use of web services and batch interfaces, input and output channels can be provided that present little impact on existing infrastructure and software.

General Deployment Model

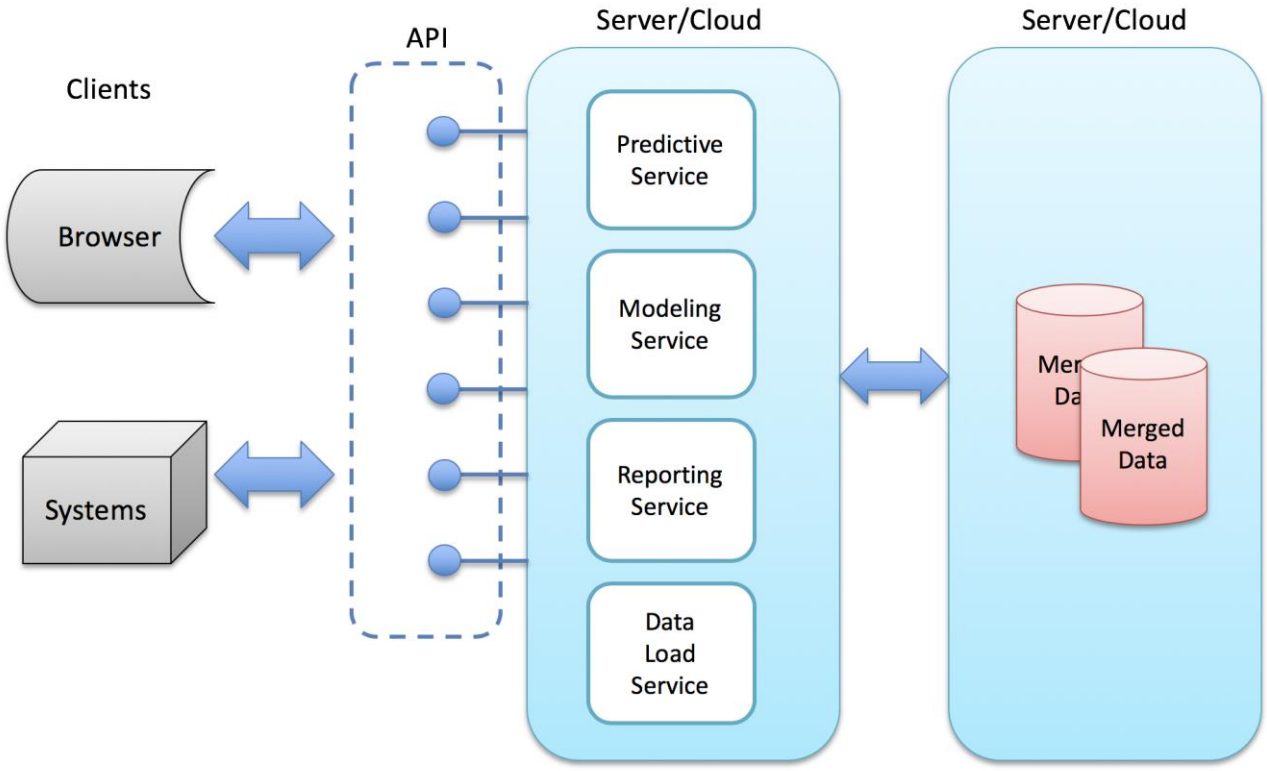

Note that the system provides a logical API layer, through which a variety of system integration and human interface functionality is managed.

Conclusions

We've described the impact that Amazon has had on retailers – and on people in general and their new expectations for 'frictionless' shopping. The Amazon Effect. The combination of price and this frictionless shopping has meant the end for some of America's great retail shops.

We've also described the new movement back toward the brick-and-mortar store, an alteration of the brick-and-mortar role that now needs to work in tight synergy with a retailer's eCommerce presence. The goal is to deliver a shopping experience that has evolved beyond fast delivery; now encompassing a 'zero delivery' aspect – with orders occurring online and buyers expecting to be able to pick up the product at the local store. Today.

That this synergy comes with extreme challenges is undeniable. The requirement to anticipate orders before they occur – and with sufficient time to get the products to where they need to be – is profound. Information about the customer, about industry trends, demographics about the region the customer and store is in, even weather information – all are critical to being able to anticipate ordering patterns. And just having the data isn't enough – we need a facility able to find subtle relationships amongst all these data, relationships that – when taken as a whole – tell us what we need to know.

We've explored how – through the use of Big Data and Machine Learning – TPG Partners has initiated design efforts based on existing research. The TPG Partners Agile Inventory Management (AIM™) system is a cutting edge approach intended to accomplish the kind of predictive analytics necessary in today's retail environment. AIM™ is based on an existing Deep Neural Network model, which along with a Big Data repository to achieve remarkable predictive power. Its outcomes can drive replenishment orders for a virtually unlimited number of products needing routing to a virtually unlimited number of stores – and do so on a timeframe that allows retailers to maximize hits, minimize misses, and keep inventory handling costs to an absolute minimum.

And, as noted, the impact of having this sophisticated capability at your fingertips is kept to a minimum as well, through the deployment of the system on cloud-based resources. Not only does this provide for easily scaled computing and storage power, but it avoids impact on the retailer's IT resources.

Find out more today, and join the movement to an Agile Inventory!

Learn More

Connect with us

https://www.linkedin.com/company/tpg-partners/

TPG Partners, LLC

Dallas / Fort Worth, Texas

marketing@TPG-partners.com

www.tpg-partners.com

Appendix

Definitions, Acronyms and Abbreviations

Term	Meaning
Amazon Effect	An expectation of 'frictionless shopping' including anticipation of interests, parallel products, and immediate availability.
Machine Learning	A sub-field of artificial intelligence in which a computer trains itself using a series of statistical methods to arrive at associations between data points
Deep Neural Network	A sophisticated type of Machine Learning that can evaluate thousands of data points and 'train' itself over time to predict outcomes based on data. Often referred to as a "Black Box" because the associations between data points are formed via deep internal layers of the algorithm.

References

Document Title	Date	Author
Five Big Data Use Cases for Retail	Feb 2018	Erin Hitchcock
How Six Brands are Using Machine Learning to Grow Their Retail	Dec 2017	Bridget Randolph
How Box Retail is Adapting to the Amazon Effect	May 2018	Geoffrey Metz
How Retailers Can Neutralize the amazon Effect	Jan 2018	Faraz Mohammed
Retail Apocalypse - Amazon Accounts for Half of All Retail Growth	Nov 2017	Kate Taylor
What the Amazon Effect Means for Retailers	Feb 2018	Lin Grosman