THE LEARNING SUPPLY CHAIN

1 1 0 1 0 1 0 0 1

+2a)-(32+.3+2a...)+a

χη**5+x+**

PUSHING OUR LIMITS

As a species, we humans have a lot of limitations. With little hair, poor speed, limited agility and relatively feeble strength, we are fairly defenseless against the elements of nature. But we more than make up for all that with an incredible knack for creativity and ingenuity – thinking outside the box and inventing tools and technology to minimize our limitations. And the human brain is especially good at massively parallel processing (doing a LOT of things at once), and because its heuristics (processing shortcuts that ease the cognitive load of making a decision) are so good, it does not even require much processing power to do it.¹

Despite our highly expansive brains and impressive memory capacity (grossly estimated to be about 10 terabytes), we show poor efficiency when processing a lot of information around a particular series of datapoints.² Our capacity as humans for sheer processing power is impressive, but unfocused. Overall, human brain processing power is estimated to be around 2.2 billion megaflops (millions of operations per second), but as we noted earlier, much of that is spent doing other things like keeping us alive and stressing over who will be kicked off The Bachelor. Computers have a significant advantage in sheer processing speed and focus, putting up to 30 billion megaflops of power to work for a singular purpose.³ As the world shrinks and supply chains accelerate, we must be able to manage, plan and deliver more — more volume and more exceptional experiences than ever before. Responding to these increased demands requires more than humans are capable of, even with traditional computing solutions.

The applied use of artificial intelligence, however, unlocks a new level of potential for us. Artificial-intelligence methods like machine learning offer the opportunity to recast the traditional limitations and resulting boundaries within supply chain execution.

In fact, it has already begun.

Recently, Mike Sparks, the director of supply chain systems at Urban Outfitters, a global apparel retailer, looked for a way to get more out of his distribution centers (DCs). The growing demands of digital commerce and direct-to-consumer shipments have created a significant amount of complexity and volume for distribution and shipping operations, that just a few years ago, were predominantly focused on wholesale and store replenishment. As a result, one of Urban Outfitters' flagship DCs was projected to reach capacity in the near future, and they expected to need another facility soon.

"Even with a million processors, we can only approach 1 percent of the scale of the human brain, and that's with a lot of simplifying assumptions."

- Steve Furber, Professor of Computer Engineering, University of Manchester

In an attempt to change that trajectory, Sparks and his team reached out to Manhattan Associates® to see if leveraging machine learning could help them to better match their demand to their inventory, resources and automation, in order to manage more tasks and extend the use of their DC. Within just a few months, it was clear that the potential of an intelligent optimization technology called "Order Streaming" was significant – it was already generating large reductions in click-to-ship times and remarkable increases in picking throughput. By early 2019, Urban Outfitters was projecting the ability to extend the use of its DC by years. But how was Order Streaming getting these results, and why was it so much more effective than traditional approaches taken by most warehouse management systems?



RISE OF THE MACHINES

"The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."

- John McCarthy, in the proposal for the "Dartmouth Summer Research Project on Artificial Intelligence"

In order to understand why technology like machine learning offers so much potential in the prediction, planning and optimization of supply and demand networks, we need to spend a little time understanding what artificial intelligence is – and is not.

Although the idea of creating machines that could "think" had already been posited by pioneers like Alan Turing (who defeated a Nazi encryption machine called "Enigma" during World War II), the term "artificial intelligence" was not used until 1956.⁴ That summer, Professor John McCarthy, a computer scientist, brought together researchers on language simulation, neuron nets, complexity theory and more to a Dartmouth summer workshop to develop concepts around "thinking machines."⁵ Today, the field of artificial intelligence (AI) is comprised of a number of disciplines, including natural-language processing, speech, vision, expert systems, robotics and machine learning. Although robotics is becoming increasingly useful in supply chain execution, machine learning holds the most potential for supply chain transformation.

Machine learning (ML) is defined as "an application of artificial intelligence, that provides systems the ability to automatically learn and improve from experience, without being explicitly programmed." And like AI, ML has multiple disciplines and methods as well, including neural networks and deep learning.

Technology has created a machine that can learn and adapt on its own, but within the limits it has been given. Machine learning is essentially a combination of three scientific disciplines: data science, computer science and mathematics. Together, they unlock our ability to create machines that can learn, adapt and grow.

HOW MACHINES LEARN

"Predicting the future isn't magic, it's artificial intelligence." - Dave Waters

When ML technology is good at a particular task or set of tasks, we call it "Narrow AI."⁶ This kind of "applied" Al is very good at forecasting, memorizing, reproducing, extrapolating and choosing the best option available given a large number of constraints. What machines cannot do is to become self-aware, get super-smart real fast, do more than they are told to do, or determine humans are useless and destroy us all. The machines we are building today are not sentient, but they are already far better than we are at considering vast amounts of data and instantaneously making an optimal decision.

So, how do the machines actually learn? There are three primary ML strategies or methods: supervised learning, unsupervised learning and reinforcement learning.

We use supervised learning to predict things when data and features are clear. The "supervision" comes from knowing the data we will use and the outcome we desire. For instance, we might use this method to predict the price of a house, since we know things like square footage, number of bedrooms, bathrooms, etc. And because we know the details of other houses nearby and their sale prices, we can use that data as our "training" set for supervised learning to make new predictions about housing prices and market changes. Likewise, machines can be programmed to achieve tasks through supervised learning strategies.

Unsupervised learning is used to look for groupings, patterns or relationships within data, especially when we have little to no real idea of what we are actually looking for. Suppose we need to understand how to group similar stores or items in order to help develop a marketing program using item, store, historical demand behavior and demographical data. Since we don't know which grouping is ideal, we call it "unsupervised." But at least we can measure it to begin learning what grouping is ideal. This strategy can also be applied to ML capabilities. Reinforcement learning is used when we have little or no data but have an environment to interact with. Think of some type of agent navigating and interacting with an environment to try to achieve a particular objective. The environment provides either a reward (e.g., getting closer to the objective) or a penalty (e.g., getting further from the objective) for each agent decision. Using this approach, the machine records and "learns" from its successes and failures. Have you ever watched a baby try to get to a toy he wants on a couch across the room? He will try various techniques to achieve the goal until he is successful. If at first you don't succeed... try, try again. Of course, a baby is a pretty sophisticated "machine," but a robot arm picking and putting away items in a warehouse is doing the same thing: learning about the size and moving speed, and weight of different items to become faster and more efficient.

BUILDING A LEARNING MACHINE

"We know already that although machine learning has huge potential, data sets with ingrained biases will produce biased results – garbage in, garbage out." - Sarah Jeong, journalist and author of "The Internet of Garbage"

So, how do we go about building one of these machines to solve a supply chain problem? There are essentially five steps to creating a machine that can learn.

First, we must understand the problem as completely as possible. Are we trying to predict or estimate something? Or are we perhaps trying to glean insights and make sense of a large amount of data?

Next, we must begin to collect data. We need to collect lots (and lots) of observations about the issue, as much as we can, and as high of quality as possible. It is also important to ensure that the data is valid.

Once we have the data, we must use it to understand the past, through advanced math and algorithms, in order to learn from the historical observations of the problem. After we understand what happened, the systems can begin to simulate millions of "trial-and-error" scenarios in an effort to identify an ideal model to overcome the problem in the future. Then, using the lessons learned from trying to solve the problem, the ML model is groomed and trained. You can think of an ML model as a software program, but instead of humans programming it, the machine can now program itself by using algorithms to learn on its own and determine how to best solve the problem.

Finally, new data about a similar problem is introduced, and the learning machine gets to work, either by providing predictions about something we ask of it, or making discoveries and telling us about something we are interested in knowing about.

A PERFECT MATCH

"All of this is about meeting customer demands, prolonging the life of a major capital investment and positioning us to smoothly transition to the next phase of digital fulfillment."

- Mike Sparks, Director of Supply Chain Systems, Urban Outfitters

Why are these learning machines such a big opportunity for the supply chain? Well, let's take a look at what is causing complexity for shippers today. First, our expectations for exceptional experiences and instant gratification just continue to rise, don't they? In a decade, we have gone from three-to-five-day paid shipping to free two-day shipping, then free next-day delivery and, now, demand for same-day delivery and click-and-collect. In fact, 26 percent of consumers now say they will abandon their cart if they cannot get their item fast enough. Those requirements put tremendous pressure on legacy warehouse operations and systems.

The shift to faster, cheaper delivery means delivering more shipments than ever before. Did you know that over 2,000 parcels are now shipped every second? In fact, global parcel shipments are expected to exceed 100 billion annually by 2020. Because of all that growth, the capacity constraints for traditional carriers is leading to an explosion of alternate shipping partners, including gig-economy freight services, which adds even more complexity to the supply chain, at a time when meeting the customer promise is more important than ever.

With these rapid changes in speed and complexity for distribution and transportation, along with other market forces like exploding ecommerce volumes, rising customer expectations and fracturing competition paradigms in omnichannel market, supply chains are feeling more pressure than ever before.

It is just too much for us to manage anymore. We are no longer fast enough or smart enough to keep up, even with traditional logistics planning and execution software. But as we saw in the earlier example, those types of problems are exactly what autonomous learning systems thrive on: optimizing task-oriented problems, making predictions and uncovering unseen insights and opportunities from vast amounts of data, where traditional analysis is simply impossible to do fast enough.

BETTER TOGETHER

As extraordinary as these machines are, it is when they are paired together with human creativity and ingenuity that we see the real long-term opportunity. It is the augmented intelligence that humans share with the machines that shows the most potential. In fact, the Harvard Business Review found that firms achieve the most significant performance improvements when humans and machines work together.⁷

Combining human capabilities, such as nuances of engagement and thinking "outside the box," with the raw power, the machine's natural capabilities of compute and error-free repetition is better than either working on its own. This pairing creates a scenario where Al technology drives the lower-level, repetitive processes associated with completing a task, while human oversight ensures the timely and accurate completion of that task.

PUSH POSSIBLE

INNOVATING FOR A NEW ERA

At Manhattan Associates, we have been committed to data science since our beginning and innovating for years to create intelligent solutions that increase distribution and transportation speed and efficiency, while reducing cost. We utilize an "applied" approach to AI and ML in order to solve specific challenges, rather than develop a general intelligence function to resolve every situation. We believe there has never been a better time for all companies, big and small, to leverage ML to modernize, strengthen and optimize their supply chain operations and processes to better compete, transform and differentiate themselves.

This isn't science fiction; it is happening right now. And Manhattan Associates is at the forefront of delivering solutions with applied intelligence that are already reshaping the way the supply chain learns, adapts and grows. Nuclei of innovation like Manhattan Active® Warehouse Management use advanced ML to orchestrate unified DC automation, employee engagement and labor, and inventory and orders. In a single solution, all warehouse management aspects are overseen through seamless, enterprise-wide visibility with actionable insights.

Applied intelligence is making our systems, our people and our companies better — faster, smarter and stronger.

Learn more at manh.com/products/warehouse-management

Contact Manhattan Associates today at information@manh.com or +1 (877) 596-9208 and together, we can Push Possible[®].

© 2020 Manhattan Associates, Inc.