

Artificial Intelligence/Machine Learning + Supply Chain Planning

Summary Report

EXCERPT

Cambridge, Mass.

November 27–28, 2018

Moderated by:

Sergio Caballero PhD
Research Scientist
MIT Center for Transportation & Logistics

James B. Rice, Jr.
Deputy Director
MIT Center for Transportation & Logistics



Table of Contents

Executive Summary	3
The Fundamentals of Machine Learning	4
Machine Learning Methods.....	4
Supervised Learning for Prediction and Classification	4
Unsupervised Learning for Pattern Discovery	5
Reinforcement Learning for Self-Learning Machines	5
Machine Learning with Neural Networks	6
Advanced Visualization	6
Machine Learning in the Supply Chain.....	6
Demand Forecasting.....	7
Context and Goals.....	7
Building a Model.....	7
Getting Results	7
Revenue Management	9
Revenue Forecasting.....	9
Revenue Management: Price Markdown Optimization	9
Revenue Management: Price Markdown Behavior	9
Transportation Planning	11
Predicting Ships' Destinations.....	11
Predicting Asset Availability	11
Predicting Spoilage and Problems in Transit.....	12
Predicting Arrival Times	12
Transportation Document Recognition	13
Predicting Spot Quote Prices	13
Looking Forward: Autonomous Vehicles	14
What Seemed Solved Wasn't.....	14
Who's Minding the Road?	14
Getting Started: People, Methods, Tools, and Data	15
Data Challenges	15
Data Analysts vs. Data Scientists	15
The Quest for Talent.....	16
Tools: Open Source vs. Commercial	16
Putting ML into the Organizational Structure	17
Change Management: Putting ML into Practice	17
Creating Future Supply Chains with Machine Learning.....	18

Executive Summary

MIT's Center for Transportation and Logistics (CTL) held a highly interactive one-and-a-half-day roundtable on the use of machine learning (ML) in supply chains. Representatives from 16 companies in a diversity of industries discussed their organizations' uses of ML in a variety of forecasting, optimization, and management applications in their supply chains. To ensure candor at the event, this report was prepared under the Chatham House Rule of not identifying the specific speakers or affiliations associated with their anecdotes, insights, or recommendations.

During the first half day of the roundtable, presenters from CTL introduced some of the fundamentals of ML. Short tutorials covered supervised learning, unsupervised learning, reinforcement learning, and neural networks. The presenters described several common data-driven algorithms for prediction, classification, and clustering. A later tour of CTL's Computational and Visual Education (CAVE) Lab showed how visualization can enable informed, data-driven decision making.

The second full day of the roundtable focused on specific supply chain applications of machine learning related to demand forecasting, revenue management, and transportation. Discussion of each application began with a kick-off case study presented by one of the industry participants. This was followed by in-depth discussions of the participants' experiences and issues with machine learning for that application. A beverage distributor described how it uses ML for demand forecasting to more accurately plan from the interplay of sales trends, holidays, weather, and promotions on sales volumes. An omnichannel apparel retailer presented its use of ML to optimize price markdowns on fashion items. An ocean freight data company and a 3PL presented half a dozen uses of ML in transportation to predict transportation asset activities, in-transit risks, spot market prices, and other applications. A presentation on autonomous vehicles illustrated the power of ML as well as the weaknesses of ML. This led to a recommendation of creating man-machine collaboration for vehicle operations.

In the final session, the participants discussed many cross-cutting issues relating to how organizations design, develop, and deploy machine learning systems. Key organizational issues included: how to find or create ML talent with the required knowledge of business, math, statistics, and computer science; where to place ML teams in the organization's structure; and how to solve change management issues in deploying data-driven automation. Key takeaways included:

- ML can improve forecasting of supply, demand, pricing, timing, etc. to proactively manage the future.
- ML can cluster and classify supply chain conditions, events, product, and customers, which can help manage complexity through differentiated responses and tailored best practices.
- ML requires data that needs to be gathered, aggregated, cleaned, and manipulated.
- ML requires more math, statistics, and computer science knowledge (and tools) than what most business data analysts and IT professionals have.
- Future supply chain leaders will need to understand enough about what is possible using ML both technologically and organizationally in order to improve business performance.

The Fundamentals of Machine Learning

After a welcome and introductions, Dr. Daniel Merchán began the roundtable by presenting the fundamentals of machine learning (ML). He first asked participants if they were using ML in their workplace, and about half of the roundtable participants raised their hands. Then he asked who was using ML at home, but only four participants raised their hands. Dr. Merchán, however, pointed out that everyone should have raised their hands to both questions. If they used Netflix, Spotify, Uber or Alexa—all of those have ML-enabled applications. Even if a participant only used email, email programs use ML algorithms to detect spam. Everyone in the room had been using ML, whether they knew it or not.

Next, Dr. Merchán defined artificial intelligence (AI) as “machines capable of performing cognitive functions associated with the human mind.” He positioned machine learning as a subfield of AI, along with robotics, natural language processing, computer vision, and speech recognition. ML is the most dominant subfield of AI, using past data to build models capable of making predictions on future data.

Although AI dates back to the 1950s, ML’s tremendous advances have been achieved only in the past few years due to the increased amounts of computing power and data that were not available before. Indeed, 90% of the world’s data has been produced in the last two years. For example, four million videos are uploaded to YouTube every minute.

Machine Learning Methods

Machine learning uses data, probabilistic models, and algorithms. Because ML uses probabilistic models, the output should be assessed using statistical confidence levels. The machine learning process requires:

- problem identification
- cleaning the data
- implementing the model
- training and testing
- evaluation
- deployment
- updating

Machine learning methods can be classified into three major families. First, **supervised learning methods use labeled training data** to make predictions on future data such as predicting demand, classifying images, detecting fraud, or making medical diagnoses. Second, **unsupervised learning methods find previously unknown patterns in data** and can be used for customer segmentation and product recommendations. Third, **reinforcement learning methods use some notion of reward to guide training** and can be used for skill acquisition. Dr. Daniel Merchán, Dr. Sergio Caballero, and Connor Makowski gave the group tutorials on these different machine learning methods as well as some more advanced neural network methods.

Supervised Learning for Prediction and Classification

Supervised learning algorithms are used for classification and prediction in which the value of the outcome of interest is known in historical data or training data. As Dr. Caballero said, “You know how to label the existing input data and the type of behavior you want to predict, but you want the algorithm to calculate it for you on new data.”

Dr. Caballero briefly described many types supervised learning techniques including linear regression, logistic regression, classification and regression trees (CART), and random forests, explaining the advantages and weaknesses of each. Regression techniques, for example, find the best-fit formula that explains or predicts an outcome, such as level of demand or productivity.

A classification tree, on the other hand, lets a company classify something, such as classifying bank customers as being acceptors or non-acceptors based on various variables such as income, education and credit card expenditure. Trees are good off-the-shelf classifiers and predictors, and they are useful for variable selection but they are sensitive to changes in the data. Slightly different sets of training data could affect the outcome a great deal. A big advantage of decision trees is that they make the logic of the decision easy to see and explain.

Company Example of ML Usage in Last-Mile Productivity

CTL researchers helped a beverage company understand and improve its last-mile productivity. To do this, the researchers developed a regression model that predicts productivity as a function of 18 route, service, and time factors such as: 1) the distance of the route, including the actual distance, the planned distance and deviations (absolute and relative) 2) duration of the route (actual, planned, and deviation) 3) stop sequence (actual, planned, deviation), and 4) vehicle capacity occupation. The analysis determined which variables were good predictors and enabled the researchers to predict whether a planned route would be of low, medium, high, or very high productivity.

Next, the researchers used a classification tree based on the variables to predict the productivity classification of other planned routes. However, they found that the classification tree was quite sensitive to the data, so they switched to a random forest model and also assessed the explanatory and predictive power of the variables. The most relevant factors identified were: route volume, vehicle capacity occupation by percentage, planned service time (hours), average drop size, number of customers, and planned route duration. The team used 110 trees to a depth of nine nodes, yielding a test accuracy of 63 percent.

Classifiers can also be trained to recognize physical objects. For example, the goal might be to classify whether an image shows a picture of a pallet or not. Supervised learning would be used on a database of images that had been labeled as showing or not showing a pallet. That data would then be used to train a model to have the highest possible pallet recognition accuracy. In the end, the trained machine can be presented with a new image, perhaps from a camera on an automated forklift, and the machine will predict whether the image shows a pallet or not.

Unsupervised Learning for Pattern Discovery

Unsupervised learning methods can identify new patterns and categories from data that were not known beforehand. Clustering methods such as k-means clustering, hierarchical clustering, k-medoids clustering, and Gaussian mixture models are common unsupervised learning methods. Each method uses different mathematical functions to aggregate similar data points together and split dissimilar ones into separate groups.

For example, additional work by CTL researchers on behalf of the previously mentioned beverage company looked at ways to tailor the company's last-mile strategies to different urban conditions. But rather than attempt to predefine these conditions for a foreign megacity, the researchers used unsupervised machine learning to discover them in the data. They used principle components analysis and k-means clustering to identify regions of a megacity with similar logistics profiles based on various dimensions of delivery patterns, population density, and road infrastructure properties. From this analysis, the researchers could further analyze critical areas of the city and propose various multi-tier distribution pilots.

Affinity analysis is another kind of unsupervised learning method. For example, Amazon and Netflix use this technique to derive product recommendation "rules" based on co-occurrences of events, such as people who buy one book also bought a certain other book. Amazon looks at conditional probabilities: what is the probability that if a customer buys A s/he will also buy B? If the probability is very high, Amazon can then recommend B to other buyers of A.

Reinforcement Learning for Self-Learning Machines

Reinforcement learning trains a machine through many iterations of decision making and provides reinforcement signals when the machine achieves a good outcome. Reinforcement learning can train a machine to successfully play a game or optimize a task without explicitly encoding the rules of play or strategies for winning. As with unsupervised learning, reinforcement learning can be used when humans don't even know the correct answer. In the case of reinforcement learning, the trainer only needs to be able to recognize a better answer from a worse one.

An example supply chain application of reinforcement learning could be the task of organizing inventory in a warehouse. The machine would try to minimize pick-and-pack labor for complex orders while avoiding congestion in any part of the warehouse. The machine might try various inventory placements and permutations of placements, which are then rewarded or penalized based on the amount of labor hours spent on fulfillment. Reinforcement learning often relies on computer simulations. Simulations are a very inexpensive and fast way to give the machine a lot of experience and a lot of time to try different strategies and tactics.

Machine Learning with Neural Networks

Neural networks are an important class of current-day machine learning algorithms that can be adapted to solve supervised, unsupervised, and reinforcement learning problems. Modeled very loosely on biological nerve tissue, a neural network for machine learning consists of one or more layers of nodes (neurons) connected to each other and to a set of inputs and outputs. The training process for a neural network adjusts the weights on the connections and other parameters to optimize the outputs that the network produces when given a set of inputs. As with other machine learning methods, neural networks come in many types, each suited for different applications such as speech recognition, natural language processing, and image recognition.

Deep learning algorithms use more complex multi-level architectures to model complex relationships. Deep learning requires less feature engineering to pre-structure the inputs and outputs for the learning model. However, deep learning requires both much more data and more computer resources to successfully train the model. Examples of deep learning include voice assistants such as Alexa and Siri, advanced game-playing AI such as AlphaGo, and autonomous vehicle control systems.

Supply chains could also deploy deep learning technologies such as voice recognition that could be used in pick-n-pack, with a voice telling the person what to pick rather than them having to read it. Another application could be voice interfaces for truck drivers, who could communicate with their trucks such as saying, “tell me where the terminal is.”

Similarly, image recognition could be used for product identification (for picking or inventory management) or for detecting cargo damage, such as if a box has been partially crushed. Inside a truck, image recognition could tell whether the truck was empty or full or what it was carrying. Inventory cams could take pictures of the warehouse and estimate how many cases of product there were.

Advanced Visualization

Participants toured the MIT CTL Computational and Visual Education (CAVE) Lab to learn how advanced visualization can accelerate decision making. The lab has a large table-top and floor-to-ceiling touch screen monitors driven by a powerful array of computers. These facilities enable researchers and executives to model and visualize complex supply chain problems through interlinked and color-coded 3-D maps, graphs, diagrams, and animations.

For example, a large retailer and a large chemical company have each sponsored research using the CAVE to help understand and optimize their distribution networks. The visualization takes into account factors such as the numbers and locations of warehouses, proximity to customer populations, inbound and outbound transport costs, market share, sales volume, and profit. The visualization lets an executive, for example, open or close warehouses or change which products are distributed through which warehouses to see how it changes service metrics, market share, and total profit. Interactive systems such as CAVE enable informed, data-driven decision making and bridge the gap between inscrutable optimization algorithms and intuitive images of system performance.

Augmented reality (AR) and virtual reality (VR) are two additional advanced visualization technologies that can enhance supply chain design, operations, and consumer interfaces. AR-goggles could provide a heads-up display for the contents inside vehicles, pallets, and so forth. In picking in a warehouse, an item could be highlighted in green and the picker could say, “I picked the green one.” Or, AR could display schematics for maintenance as needed when looking at a broken conveyor belt, for example, or it could highlight the bolt that needs to be replaced and which way to rotate it, with the end goal of speeding up repair. VR could be used to design new warehouses or train workers to drive forklifts. AR or VR could also help a customer view a product, such as what a couch would look like in their living room.

Machine Learning in the Supply Chain

Machine learning can be used for many categories of supply chain applications. ML can be used for prediction or forecasting of demand, supply, on-time deliveries, and risks. ML can help automate many routine elements of supply chain operations and help detect or predict exceptions to routine operations. ML can be used for planning and design such as of networks, inventories, schedules, and routes. Finally, ML is a key component of autonomous supply chain vehicles such as trucks, ocean freighters, delivery drones, and forklifts.

The second day of the roundtable focused on specific supply chain applications of machine learning. Discussions of each application began with a kick-off case study presented by one of the industry participants. Then participants shared their own experiences with that application of ML, and asked and answered each others' questions about the application.



Report recording prepared by:

Andrea Meyer and Dana Meyer
WorkingKnowledge®
3058 3rd St.
Boulder, CO 80304
workingknowledge.com

Edited by the moderators.