FIRST ANALYTICS®

Machine Augmented Demand Planning

A Framework for Classifying AI/ML and Analytical Applications

AUGUST 2019

FIRST ANALYTICS®

Contents

FIRST ANALYTICS®

Introduction and a Framework

The skepticism about the hype surrounding machine learning and AI is justified. This is especially true when it comes to forecasting and demand planning, as these domains do not fit the typical use cases proposed for these technologies.

Experienced forecasters have always been incredulous that technologies like neural networks, which traditionally require tens or hundreds of thousands of data points for training, can work on time series where there is not much data. Time series forecasting models go back decades, and studies and competitions have shown that often a simpler model, versus an elaborate one, is the most accurate. It is impressive how many times a simple exponential smoothing model gets the job done.

However, forecasters overlook the potential for advanced analytics in aspects of demand planning that do not relate to the statistical modeling itself. These present areas of opportunities for machine algorithms to improve overall demand planning accuracy and productivity.

What Do We Mean by "Machine Augmented"?

By using the word "augmented" we propose that machine learning and other algorithms can assist demand planners in doing their job. We are not suggesting that artificial intelligence (AI) and machine learning (ML) are going to replace demand planners. Rather, demand planners can rely on these technologies to help them do their jobs better.

Furthermore, we recognize that machine learning, per se, is only one approach to analytics amongst a broader set of methodologies, which could include traditional statistical modeling. There are traditional methods which, though not as seductive as machine learning, have been underutilized in demand planning.

A Framework for Classifying AI and Analytical Applications

To help understand where these technologies can make an impact, we have established a framework to classify them. In doing so, we note that no framework can be comprehensive or complete. Furthermore, in the current days of rapidly developing technologies, we realize that the framework can become outdated very quickly.

At a high level we divide the applications into three broad categories: data engineering, forecasting models, and planner decisions. These categories follow, sequentially, the workflow of the planner. It starts with data, which goes into a forecasting model that is then acted upon by the planner.

Within each of these categories we have placed applications into two sub-categories, for a total of six groupings of analytical applications. In the sections that follow we describe each of these categories and subcategories and provide selected examples for each. An overview of this framework is shown in figure 1.

Figure 1- Application Framework for Machine Augmented Demand Planning

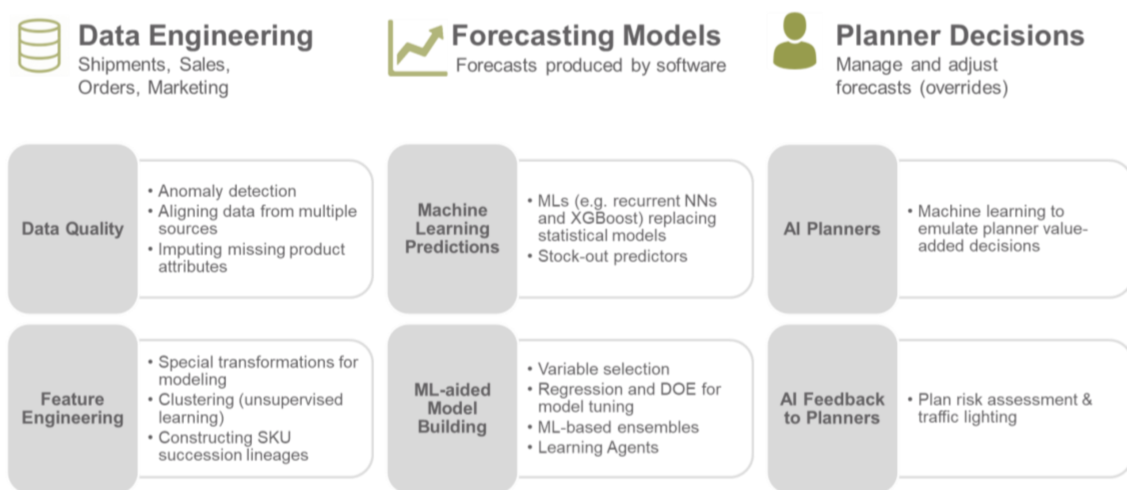

Category 1: Data Engineering

This refers to the data used to support demand planning and is upstream in the process. The data can be shipments, sales, or orders, and may include data on marketing, weather, or other events.

Data Quality

The (lack of) quality of the data used in the demand planning process can significantly impact the results. Models and algorithms can help with:

- **Anomaly detection**

  Finding records that are unusual or are outliers. Statistical analysis has long been used for this, but often without rigor. Now, machine learning classifiers may be used to flag records for inspection.

- **Aligning data from multiple sources**

  Sometimes combining data with different hierarchies, granularities, and time intervals can be challenging. Machine learning has been shown to automate this process, whereas a human had to intervene in the past.

- **Imputing missing product attributes**

  Many item master files, which contain the attributes of the product, have missing data. Some new item forecasts depend on reliable attribute data. Statistical imputation routines for categorical data can be used to fill in those holes.

FIRST ANALYTICS®

Feature Engineering

"Feature Engineering" is a term from the machine learning community that refers to transformations of data inputs. Examples include:

- **Special transformations for modeling**

  Sometimes better improvements in forecast accuracy can be realized through transformations of the inputs, rather than by selecting better models. There are now heuristics and algorithms to help with this transformation process.

- **Clustering**

  Referred to in the machine learning community as "unsupervised learning", this is a way of grouping items (e.g. product/location time series) into clusters with similar attributes, patterns, and behaviors. Treating these clusters similarly with respect to the forecasting models employed has been shown to increase accuracy. It is worth noting that some unsupervised learning methods are statistical methods dating back many decades. Such is the case with K-means clustering.

- **Constructing SKU succession lineages**

  SKUs evolve over time. New SKU variants are released, usually not very different than the prior version in its attributes, yet assigned a new SKU number. Planners like to treat these variants as a single item. But often the history of how the SKU has evolved is not well-maintained, or otherwise poses challenges to constructing a long, consistent history. This is especially true with seasonal fashion products, where a new SKU is established each season for essentially the same product. Manual mapping of past SKUs to the current one can be labor intensive, and ML can help by suggesting which past variants might be associated with the current.

Category 2: Forecasting Models

Most demand planning processes rely on (usually statistically based) forecasting models. When we hear the latest claims about machine learning we most often assume they are being positioned to replace statistical models. In some cases, this may be true. But almost always, machine learning models can be used alongside traditional models, or may help in the model building itself. Thus, these two subcategories:

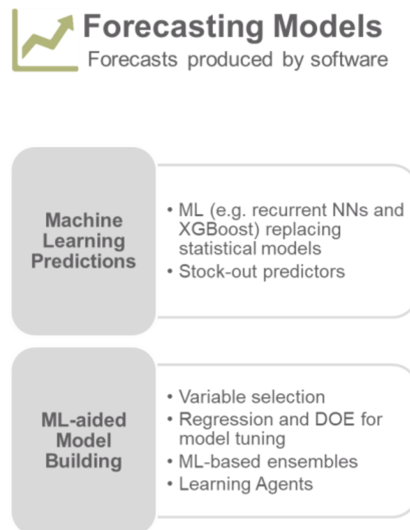

Machine Learning Predictions

This is the case where a machine learning model takes the place of what has typically been a statistical or time series model.

- **Machine Learning Replacing Statistical Models**

  Time series forecasting is amongst the last domains where machine learning models have begun to take hold. Two examples include recurrent neural networks, such as LSTMs, and XGBoost. These in fact are beginning to show promise, yet still complex to implement.

- **Stock-out predictions**

  Statistical models have been used with success in predicting out-of-stocks in advance at the item, selling-location level. ML classifiers can produce probabilities of stockouts, often capturing “non-linearities” that traditional statistical methods with distributional functional forms and parameters cannot.

FIRST ANALYTICS®

Machine-learning Aided Model Building

A challenge for model builders lies in making decisions about what kinds of models to build, and how to set their estimation parameters. Here are four examples of how algorithms (speaking in broader terms) can help.

- **Variable selection**

  Both traditional statistical methods and more recent variable importance routines in modern algorithms are available to help a modeler decide what variables to include in a time series forecasting model.

- **Design of Experiments and regression for model tuning**

  Most forecasting software systems provide many control settings and parameters for the modeler to evaluate and set. But with this flexibility and power comes the challenge of finding the best "set points" for the parameters. The possible combinations can be considered infinite. Design of Experiments, often coupled with regression analysis, is a rigorous, scientific way to understand the sensitivity of the model forecast error to each of the settings in a very efficient way.

- **ML-based ensembles**

  We have known for a long time that combining models, or "model ensembles," can provide slightly better forecasts than a single model alone. Traditional methods for combining models include simple or weighted averages, or even regression-base combinations. Now, machine learning techniques offer, in some cases, a more accurate approach to combining models.

- **Software learning agents for model tuning**

  Reinforcement learning is an example of a system where a software agent examines a state, takes an action, and perceives a reward based on the action. Thus, the agent learns which actions result in the best outcome. In the case of establishing forecasting model parameters, the actions would be the setting of those parameters, and the reward would be some accuracy measure, like MAPE. This is an area that has not been fully proven but conceptually holds promise as an augmentation aid to modelers who would do this manually.

Category 3: Planner Decisions

This area deals with the day-to-day job of planners in working with statistical base forecasts. It is the closest to AI that we have in demand planning.

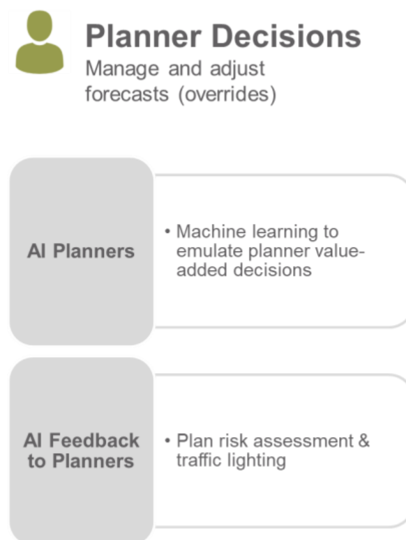

AI Planners

The question here is: can machine learning create a sort of artificial intelligence that emulates the actions of demand planners?

- **Value-added overrides**

  Planners regularly override baseline statistical forecasts. Sometimes they help (and add value), but sometimes they do not. ML algorithms can be shown the good versus the bad historical decisions that planners make, and emulate their behavior, suggesting possible good overrides that the planner may want to make.

AI Feedback to Planners

In this case, AI observes the proposed action a planner intends to take and provides advice as to the outcome.

- **Plan risk assessment**

  As an example, an AI can be trained on the outcome of past promotion plans and assess whether a new proposed plan has a likelihood of executing well, given the inputs. Sometimes, planners are overly optimistic about the level of retailer support they will get. A traffic-lighting system can be established to provide red/yellow/green feedback to the planner, which tells them the reality of their plan coming to fruition, and thus mitigates the risk of overly optimistic plans and their forecasts.

FIRST ANALYTICS®

Conclusion: Automation vs. Augmentation

Our firm, First Analytics, has long used machine learning and other advanced algorithms, in areas where there were better fits for the technologies and were mostly outside of demand planning. We have always been skeptical of their usefulness in this domain. But we have changed our opinion in recent years, as we are starting to see success.

That is not to say that other approaches will be failures, but it is worth examining how machine augmented demand planning can help you. We have many case studies we can show you for the examples illustrated in this paper. Most of the applications are in areas ancillary to statistical forecasting of demand.

This all leads to the question of whether AI will eventually replace demand planners. Should demand planners feel threatened?

Our co-founder, Tom Davenport addresses the issue of augmentation vs. automation in a general sense in his book, *Only Humans Need Apply*[1]. He finds that nearly all machine learning implementations have augmented, not replaced, humans. We are confident that AI and algorithms will make planners more productive and help them achieve better demand planning performance.

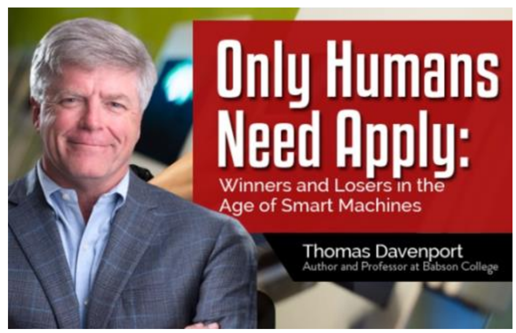

[1] Thomas H. Davenport, Julie Kirby (2016). Only Humans Need Apply: Winners and Losers in the Age of Smart Machines. Harper Business.

FIRST ANALYTICS®

About First Analytics

Analytics will play an ever-increasing role in helping businesses be more accurate and effective, overcome their greatest challenges, and compete more effectively. We have seen the power and potential of advanced analytics firsthand, and know exactly how to help companies turn vast amounts of raw, unsifted data into rich and tangible intelligence they can use. To what end? To anticipate risk and drive safety; to save dollars and lives; to optimize capacity, inventory and performance; and to enable change. One-time analysis or real-time, end-to-end applications. Big data or small. On-premises or in the cloud. We help make analytics a vital element of your operations and a key to sustainable growth and success.

Spanning multiple industries and applications, our data scientists and analytical experts turn information into insight, knowledge into know-how and evidence into action. Simply put, we leverage the most up-to-date analytical tools, strategies and technologies out there in order to help improve your business and optimize operations.

We have a long and successful history of helping our clients with supply chain optimization, always with a data analytics approach. Visit our website to view our case studies and to learn more about how we can help your company leverage the power of advanced analytics with machine learning and artificial intelligence.