WHITE PAPER

Assisted Demand Planning Using Machine Learning for CPG and Retail

Contents

We are entering the second Machine Age. This first began with the Industrial Revolution, which was driven primarily by technology innovation – the ability to generate mechanical power to make humans more productive. Rather than the steam engine that started the Industrial Revolution, the second Machine Age uses computers and other digital advancements to help our brains understand and shape our environments.

Impressive advances in artificial intelligence (AI) and machine learning (ML) during the past decade are supported by supervised deep learning to train ML algorithms to perform narrow, single-domain tasks. The learning is supervised because you're telling the algorithm the correct answer as it is exposed to many examples using big data sets. But we're now seeing unsupervised learning systems that learn faster, require less data and achieve impressive performance. These supervised and unsupervised intelligent automation techniques can help humans automate tasks. That doesn't eliminate the need for humans, it just allows them to do their work more effectively.

Intelligent automation techniques can be applied to all kinds of activities across your organization to reduce the everyday repetitive work while uncovering key insights to improve the effectiveness of your processes, as well as your workforce.

How intelligent automation enhances existing processes

Intelligent automation driven by AI and ML are disrupting the way companies do business. The rapid deployment of automation is helping us set new standards of efficiency, speed and functionality.

Instead of being replaced, humans will see unprecedented job creation and new opportunities to add more value. Intelligent automation techniques can be applied to all kinds of activities across your organization to reduce the daily repetitive work while uncovering key insights to improve the effectiveness of your processes and your workforce.

Applications can range from routine to groundbreaking, such as collecting, analyzing and making decisions about textual information to guiding demand planners to anticipating consumer purchasing behavior. It is already helping companies surmount conventional performance tradeoffs to achieve unprecedented levels of efficiency that reduce costs while increasing profitability.

The variety of business challenges to which intelligent automation can be applied is expanding as technologies for voice recognition, natural language processing and ML improve. These technologies are becoming increasingly available as IoT devices capture streaming information and open-source cloud-based services become more widespread.

> "The challenge is that people have not developed the level of trust in artificial intelligence and machine learning that they have in other technologies that automate tasks. People sometimes confuse automation with autonomy."
>
> Oliver Schabenberger,
> COO and CTO, SAS

Leading organizations are driving more of their processes into smarter machines. They are rethinking what they do across every area of the enterprise – from their business processes to the customer experience. Some activities where intelligent automation is helping companies are:

- Data collection.
- Security and systems monitoring.
- Transaction management with ERP systems.
- Scheduling and staffing.
- Accounting and finance.
- Business planning.
- Customer experience.
- Marketing and communications.

The levels of automation include:

- Basic automation of frequent repetitive simple tasks.
- Advance automation that orchestrates workflows across departments and applications.
- Intelligent automation that mimics complex research and expert decision making.

The foundational elements for intelligent automation are:

- Institutionalized business process.
- Reducing effort and increasing accuracy.
- Centralized application with auditing and instrumentation.
- Quantifiable metrics to measure and inform model improvements.
- Analytics infrastructure to support machine learning.
- APIs to allow software agents to mimic and drive tasks.
- Continuous monitoring and automation governance.

Intelligent automation is empowering humans with advanced smart technologies and agile processes for faster, more intelligent decisions. The key benefits of intelligent automation in business include:

- Improved productivity.
- Increased process efficiency.
- Improved customer experience.
- Unprecedented value (ROI).

The advantages of using intelligent automation tools can lead to new business strategies that could not have been conceived of previously. Companies who have invested in intelligent automation have been able to automate about half of their tasks, reducing process times by as much as 50 percent. Completing tasks more quickly means companies address more complex tasks without spending additional time (and revenue). Depending on the industry, faster job completion with improved precision can mean increased revenue.

Practical application of intelligent automation: A day in the life of a demand planner

Demand planning at most companies is a weeklong, multistep process that includes an array of manual workflow processes (see Figure 1). The workflow starts with a review of the latest statistical forecast created by manually cleansing the demand history (sales orders or shipments) by removing outliers and removing promotion volumes, then using the most effective statistical models to forecast the baseline volume. The promotion volumes that were removed from the demand history are manually reviewed and revised by the commercial team (sales and marketing), and then blended into the baseline forecast to create a demand plan.

This is a manually intensive workflow during the first week of each month. Demand planners typically develop the baseline forecasts or analysts, who may be regional or centrally located. They work under the guidance of the demand planning manager. Baseline forecasts are communicated to members of the demand management team. This usually includes regional sales leaders, brand managers and product managers. Additional manual adjustments are made throughout the process because of changes in marketing programming, corporate strategies and market dynamics along with changing consumer preferences and behavior. This adds additional complexities as demand planners try to manually incorporate these changes into the demand plan.

More than 40 percent of a demand planner's time is spent managing information and data. Another 30-40 percent is spent managing and fine-tuning the demand forecast based on new market and customer information, changes in marketing programming (tactics) and coordinating the consensus forecast (plan). Finally, creating and updating KPI reports represents about 10 percent of a demand planner's time.

With the introduction of intelligent automation using ML, a large portion of the manual, repetitive activities can be automated (see the highlighted activities in Figure 1) allowing demand planners to be more productive and adding real value to the overall process.

Monday	Tuesday	Wednesday	Thursday	Friday
Generate and review forecast exceptions from previous week.	Import and apply promotion uplifts created by sales/marketing.	Review forecasting parameters for A class and strategic items.	Lead consensus planning meeting with sales and marketing.	Gather inputs from S&OP meeting.
Manually run data cleansing routines (separate historical demand into baseline and promoted volumes).	Apply new product volumes where needed.	Review forecast at higher aggregations.	Apply any changes arising from the consensus planning meeting to the forecast.	Apply any changes arising from the pre-S&OP and/or S&OP meeting to the forecast.
Resolve any issues related to last forecast submitted to supply.	Review statistical baseline/promoted forecasts generated on exception basis.	Coordinate with sales and marketing to capture any changes to targets.	Prepare and submit reports for pre-S&OP (weekly) and S&OP (monthly) meetings.	Last-minute checks at the aggregate level.
Generate/review performance metrics (MAPE, WMAPE, FVA) report for previous week.	Manually, blend promotion uplifts into the baseline statistical forecast.	Reconcile top-down forecast changes (apply mass aggregations) to lower levels.		Approve and submit forecast to supply plan.
Run and review statistical baseline forecast generated based on new history on exception basis.		Manually, tune lower levels based on sales/marketing input.		

Demand Planner's Workload

- 40% Data Issues/Setup/Cleansing
- 30% Forecast Management/Tuning
- 20% Coordinate/Achieve Consensus
- 10% Reports/KPIs and Learning

Figure 1: Demand planner weekly activities.

What is forecast value added?

Companies have been searching for a performance measurement that can effectively measure and improve the demand forecasting process, reduce cycle time and minimize the number of touch points. The best approach is to implement a methodology called forecast value added (FVA), or lean forecasting.

FVA is a metric for evaluating the performance of each step and participant in the forecasting process. Simply put, FVA is the change in forecast accuracy before and after each touch point in the process. The process can be based on any specific forecast performance measurement, such as percentage error (PE), absolute percentage error (APE), mean absolute percentage error (MAPE) or weighted mean absolute percentage error (WMAPE).

FVA determines whether any manual adjustments to the forecast have added any value to the accuracy of the demand forecast. If those manual adjustments increase the accuracy of the statistical baseline forecast, then those changes should remain in the process. However, if the manual adjustments do not improve the accuracy, they should be eliminated or minimized (simplified) to reduce cycle time and resources, thereby improving forecast process efficiency (see Figure 2).

Process Step (1)	MAPE (2)	Naïve (3)	Statistical (4)	Override (5)	Consensus (6)
Naïve	50%	-	-	-	-
Statistical Forecast	45%	5%	-	-	-
Analyst Override	40%	10%	5%	-	-
Consensus Forecast	35%	15%	10%	5%	-
Approved Forecast	40%	5%	5%	0%	-5%

Notes: 1. Column 2 gives MAPE of each set of forecast. For example, 50% MAPE is of naïve forecasts, 45% of statistical forecasts, and so on.
2. Other columns give percentage point improvement made by one set of forecasts over the other. For example, statistical forecasts improved over the analyst override by five percentage points, and 10% over the consensus forecasts.

Figure 2: FVA Performance Metrics Comparison

FVA is a common-sense approach that is easy to understand. The idea is simple – it's basic statistics. What are the results of doing something versus doing nothing? FVA can be either positive or negative, telling you whether your efforts are adding value by making the forecast more accurate, or whether you are making things worse. FVA analysis consists of a variety of methods that have been evolving through industry practitioners' applications around this innovative performance metric.

FVA is used to improve not only forecast accuracy but reduce non-value-added touch points in the demand planning process, thus improving process efficiency. According to a recent analyst report, FVA is the second-most widely used performance metric to measure the effectiveness of a company's demand forecasting and planning process. Weighted MAPE (WMAPE) is the No. 1 performance metric, while the former standard, MAPE, is now the third-most widely used.

Now imagine trying to manage all those touch points using spreadsheets. What if, instead, you could improve your FVA tracking using ML with the goal of reducing the complexity of managing all the FVA information through automation, while providing demand planners with targeted intelligence to pinpoint where, when and by how much to make manual overrides to the statistical forecast?

The importance of intelligent automated

Not all manual overrides, whether fact-based or experience-based, add value. In fact, many add little value. The challenge is to distinguish the non-value-added overrides from those that are value added. This becomes more difficult when you have thousands of SKUs across brick-and-mortar stores, mobile, online, Amazon, and other related e-commerce channels in multiple countries. This could represent millions of forecasts. It would be impossible for demand planners to review and manually adjust all those forecasts using spreadsheets. So they rely on aggregate-level adjustments that are then disaggregated down their business hierarchies. In many cases, mass aggregation of overrides is not an accurate way of manually adjusting statistical forecasts even when adding other sales and marketing activities not considered (or available) for the original forecast.

Using intelligent automation to improve a demand planner's FVA

Assisted demand planning

Recently, SAS tested a new patent-pending intelligent automation technique that uses ML to boost FVA with a large global consumer packaged goods company using ML to learn from past demand planners' manual overrides. The test focused on two main objectives:

- Identify forecast entities that need overrides.
- Provide demand planners with the direction and range of overrides (as to the need to raise or lower statistical forecasts) at various levels of the business hierarchy.

ML analyzes past statistical forecast and consensus forecast overrides to learn from successful and unsuccessful forecast adjustments to identify the best periods to review candidates for overrides. You can then provide guidance to demand planners about where, and by how much, to adjust the business hierarchy forecasts by either raising or lowering the statistical forecast. See Figure 3:

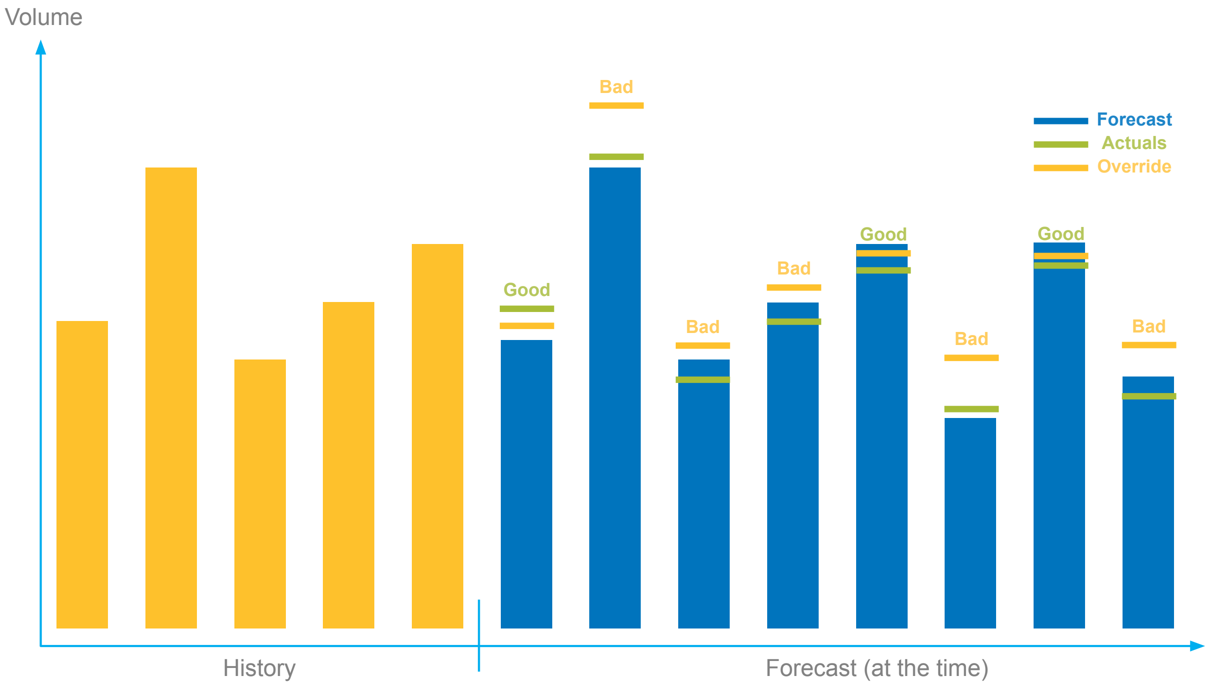

Figure 3: Override only if the forecast can be improved

A process approach

In the example below, a minimum of two-and-a-half years of historical overrides based on an 18-month rolling forecast were collected for five product categories in two geographic areas for more than 700 items. A 60-day future forecast was used for FVA purposes. In-sample and out-of-sample training and validation periods were used in comparison to the FVA analysis to choose the appropriate ML model. A three-step approach was implemented: 1) Enrich 2) model and 3) assess (see Figure 4).

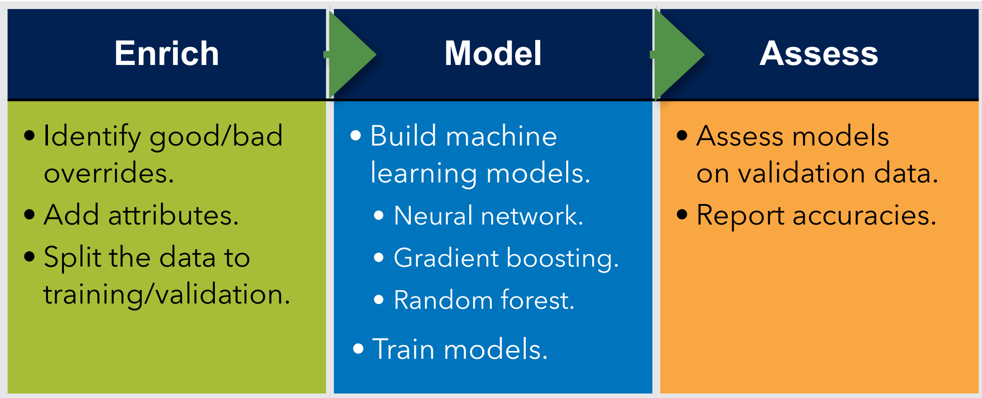

Figure 4: Intelligent automation process approach

Step 1: Enrich the process by identifying value-added and non-value-added overrides made by several demand planners, and add any other attributes that are available.

Step 2: Build ML models using neural networks, gradient boosting and ensemble random forest training models in a competition to determine the champion model.

Step 3: Assess models using the out-of-sample validation data and report the levels of accuracy.

These steps can be enhanced by adding other causal factors like sales promotions, pricing strategies and others. In this example there were no causal factors to enhance the ML models to show proof of value. Smart rules can also be added such as not makeing overrides if the MAPE is less than 10 percent, or that you only consider FVA analysis results if the three-month historical FVA average is greater than 30 percent.

Arriving at a champion model

The assisted demand planning using the intelligent automation process reduced the number of manual overrides by 47 percent, allowing the demand planners to focus only on those products and periods that would benefit the most from overrides. As a result, it improved the value added to the forecast by 6.3 percent.

The user interface in the SAS® Collaborative Planning Workbench guides demand planners in making manual overrides, in what direction (up or down), and within a volume range (see Figure 5). As you can see, green arrows for March, May and June indicate the direction of the override. The yellow highlighted cells indicate the lowest volume range and the white cells above the arrow indicate the maximum overall override volume.

For April there is no arrow up or down (down arrows are highlighted in red), indicating no manual overrides are suggested for that month because the statistically derived forecast is more accurate. Each demand planner can scroll up or down the product hierarchy for planner, customer and product. Note that there is no limitation to the number of levels in the hierarchy. It is based on available data.

If a demand planner chooses to make a manual override for a month not recommended by intelligent automation, a default warning appears to indicate that overriding the statistical derived forecast is not recommended. Demand planners can be blocked from making overrides in those cells, or you can flag those cells with warning messages but still allow the demand planner to make an override.

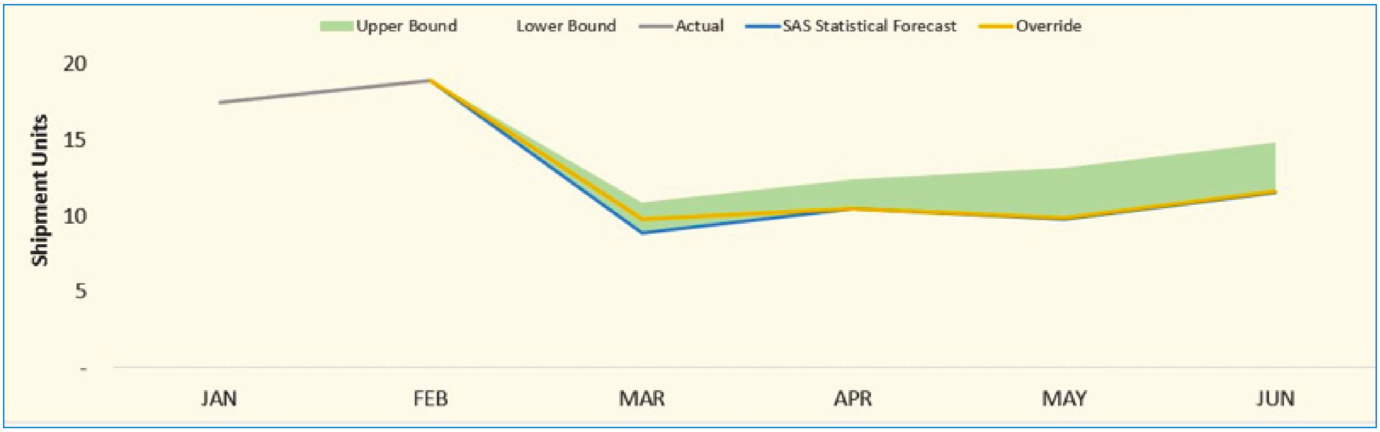

Select the Customer/Product to View or Override:

Planner ... Demand Planner
Customer ... Product Group
Product ... Product A

Override Legend:

- Value entered is out of recommended bounds
- Recommend an override
- No override recommended

	JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
	History	**History**	**Current**	**Future**	**Future**	**Future**	**Future**	**Future**	**Future**	**Future**	**Future**	**Future**
Actual	17.48	18.90	-	-	-	-	-	-	-	-	-	-
SAS Statistical Forecast			8.83	10.48	9.73	11.50	-	-	-	-	-	-
Recommended Adj Direction	-	-	↑		↑	↑						
Override			9.78	10.48	9.87	11.60						
Lower Bound	17.48	18.90	8.83	10.48	9.73	11.60	-	-	-	-	-	-
Upper Bound	17.48	18.90	10.85	12.42	13.17	14.83	-	-	-	-	-	-

✓ Override direction
✓ Range of the overrides

Figure 5: Overriding the user interface

Overall, ML helped to improve the statistical model accuracy (on average) across all the product groups, items and geographies with the champion ML model – in this case, neural networks – as follows:

- Best ML model average accuracy across all product groups, items and geographies was 75 percent.
- Accuracy if ML recommends no overrides: 86 percent.
- Accuracy if ML recommends overrides: 65 percent.

Given these results, using ML to boost the FVA process has demonstrated that it can automate demand planners' repetitive work of managing the enormous amounts of data required to support the FVA process while providing targeted intelligence to pinpoint where, when and by how much to manually adjust the statistical forecast. The real benefits are the ML recommendations to not make manual adjustments to the statistical forecast because it is much more accurate than prior overrides. Intelligent automation reduces complexity and improves accuracy.

In the same way that an automobile's onboard computers can be connected to diagnostic machines allowing a technician to assess the problem within minutes (versus hours manually) to pinpoint the exact failed part, intelligent automation helps the demand planner to work smarter, but not be replaced by the machine. Intelligent automation will help demand planners analyze vast amounts of information to boost the FVA process, guiding them with surgical precision to work smarter. Demand planners will be able to ingest and analyze massive amounts of forecast information, respond quickly to complex inquiries, and make overrides with precision across the entire business hierarchy.

"The conversation we should have is how machines and algorithms can make us smarter, not how smart we can make the machines…"

Tom Gruber, Co-creator of Siri (TED Talk, 2017)

To contact your local SAS office, please visit: sas.com/offices

THE POWER TO KNOW.