WHITE PAPER

Using machine learning and demand sensing to enhance short-term forecasting

Contents

Introduction

Effective supply chain management is critical to ensure the timely replenishment of products by properly managing the movement and storage of raw materials and finished goods from the point of origin to the point of consumption.

A key component of supply chain management is forecasting short-term consumer demand accurately to ensure that products are available to consumers when they need them in the near term. The negative impacts of incomplete or inaccurate data on consumer demand was one of the key takeaways from a famous supply chain simulation game developed by MIT Sloan and taught in its MBA and executive education program for nearly 40 years.

As supply chain management evolved in the 1980s and 1990s, the visibility of consumer demand – always of importance to retail and CPG companies – became an important factor for many manufacturing companies, particularly in the high-tech sector. The personal computer industry became a hotbed of supply chain innovation and was a sector where demand sensing evolved into demand shaping.

Implementing a short-term (one to eight week) forecast is critical to understanding and predicting changing consumer demand associated with sales promotions, events, weather conditions, natural disasters and other unexpected shifts (anomalies) in consumer demand patterns. Short-term demand sensing allows manufacturers, retailers and CPG companies to predict and adapt to those changing consumer demand patterns. Traditional time-series forecasting techniques that model patterns associated with trend and seasonality are typically used for demand forecasting at manufacturing companies, retailers and CPG companies. These models can uncover those two historical demand patterns and provide an estimate of demand into the future. In addition to the historical demand data, organizations can use other data feeds – such as point-of-sale (POS) information, future firm open orders and promotional events – to create a collective picture of the demand signal.

The key benefits for using demand sensing include:

- Increasing sales revenue by improving sensing capability to drive an agile supply chain response to meet customer and consumer demand needs.
- Improving transportation planning with preferred carriers, cutting execution costs by reducing redeployment and lowering inventory carrying costs.
- Improving customer service levels and on-shelf availability of products, ensuring consumers find the products that they want and need.
- Improving revenue/profit through enhanced replenishment efficiencies and fewer stock-outs on shelf.

SAS now has a patented machine learning (ML) approach to creating one-to-eight-week (weekly and daily) demand forecasts. This new approach, combining machine learning and traditional time-series forecasting models, allows business analysts to generate improved weekly and daily forecasts by using the historical supply signal (shipments) data in combination with point-of-sale data (demand signal).

A practical application of demand sensing using machine learning

Consumer packaged goods companies (CPG) account for some of the biggest industries in the world and provide items that are used regularly by consumers, including food, beverages and other household products. Since many products provided by CPG companies have short shelf lives and are intended to be used quickly, companies need to routinely replenish products on store shelves to meet consumer needs.

Working with a large global CPG company, we were provided with seven years of order shipment history for triplets of Product (Prod), Shipping Location (ShipLoc) and Customer Location (CustLoc) at a daily resolution as shown in Figure 1.

Shipping locations corresponded to the CPG distribution centers (DCs) and retail customer locations were customer DCs. Thus, the order history corresponded to the number of daily product shipments from ShipLoc to CustLoc. We were also provided future open orders for both data sets at the lowest (Prod – ShipLoc – CustLoc) level. The shipment order history and future open orders were the main inputs in our forecasting advanced analytics models. We were also provided with point-of-sale data and inventory data for one specific CPG retail customer, that is (Prod – ShipLoc – CustLoc = customer), for less than one year of history.

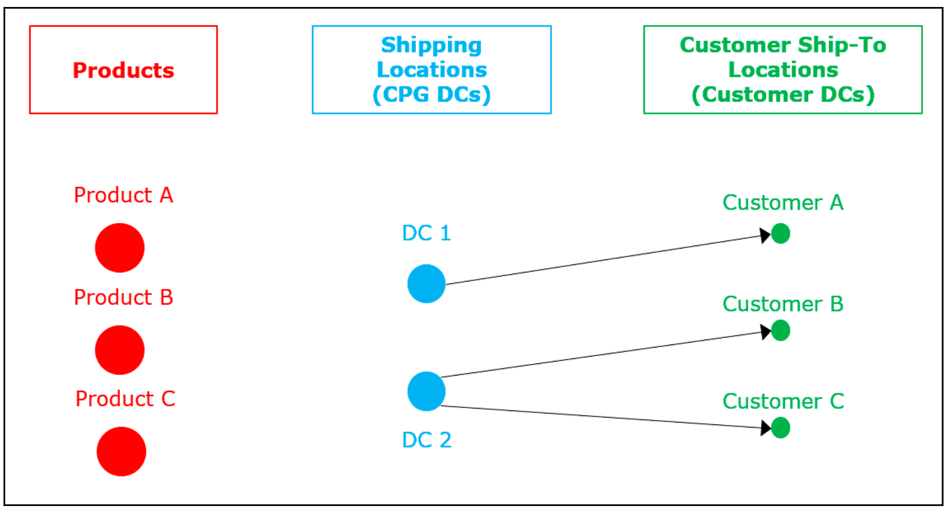

Figure 1. Graphical representation of the provided data. There are three main levels to the data – Product (Prod), Shipping locations (ShipLoc) that are CPG distribution centers (DCs), and Customer locations (CustLoc) that are customer ship-to locations.

In addition to the order history, future open orders and POS data, the customer provided weekly estimated forecasts using its current forecasting procedures and technology. Two such forecast estimates were provided. One was generated using standard procedures based on SAS® (FC-Base), and another was generated by experts who had adjusted FC-Base to further refine the existing forecasts (FC-Base+Expert).

Each of these forecast estimates were provided at the (Prod) and (Prod – ShipLoc) levels. The main goal of this project was to generate better forecast estimates compared to FC-Base and FC-Base+Expert at both (Prod) and (Prod – ShipLoc) levels. Note that the forecasts for comparison were only provided at the weekly level. The CPG company did not provide daily forecasts for comparison.

Rolling weekly forecasts

Several rolling weekly forecasts for different forecast starting dates were generated. Order history data prior to a given forecast start date was used for training, and 12 weeks of data after the start date was used as the holdout data for validation, as demonstrated in Figure 2. For each rolling forecast, we compared performance for the current week (Lag 0) and up to 12 weeks into the future (Lag 11).

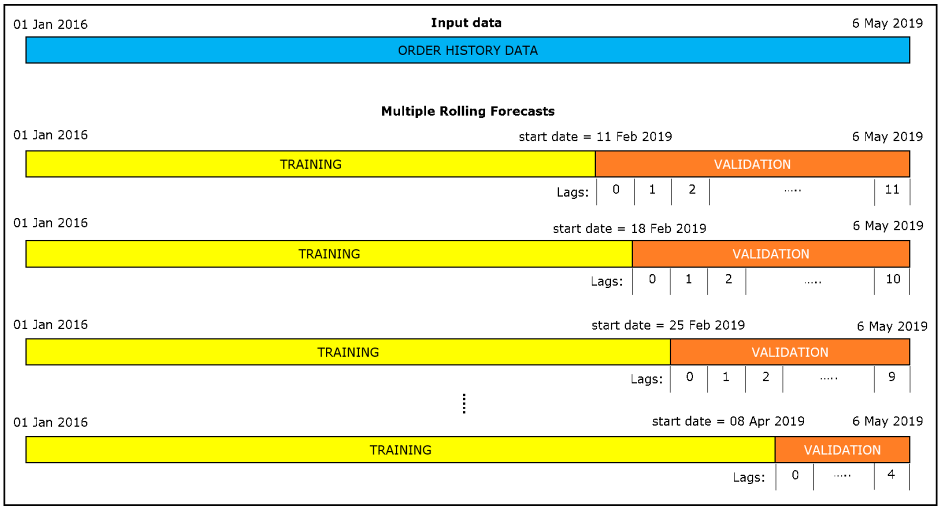

Figure 2. Demonstration of how order history data was split for multiple rolling weekly forecasts. A start date was chosen, and all data prior to that date was used as training data to train the models and the data after the start date was used as the holdout validation data.

For each rolling forecast, we used two methods to generate forecasts:

1. Traditional time-series based forecast.
2. Neural Network + Time-Series (NNTS) forecast.

For the time-series based forecast, the SAS traditional time series forecasts are used to diagnose the statistical characteristics of each {Prod – ShipLoc – CustLoc} triplet time series and identify appropriate forecasting models for each data set. The diagnosis of the results was used to generate a time-series forecast estimate and a trend estimate for each {Prod – ShipLoc – CustLoc} triplet.

For the NNTS forecast, we trained a neural network for each Product separately by using the following inputs:

- ShipLoc location.
- CustLoc location.
- Previous four years of order history.
- Forecast estimates.
- Trend estimates.

Next, we incorporated the future open orders data to further refine our enhanced forecast. For each rolling forecast over the validation period, we compared our enhanced forecast estimate with the future open orders and replaced the forecast estimate with the open order quantity, if the open order quantity exceeded the forecast estimate.

To analyze the impact of using point-of-sale data and customer inventory data, we created two versions of the NNTS forecast. One forecast using inputs to the neural network listed above, and another one with two additional inputs corresponding to POS and customer inventory data.

Weekly forecasts to daily forecasts

We generated daily forecasts by disaggregating or breaking down the enhanced weekly forecasts for each {Prod – ShipLoc – CustLoc} triplet into enhanced daily forecasts. We achieved this by estimating the relative daily order proportions to disaggregate the weekly forecast into daily forecast. We estimated these proportions for each {Week – Prod – ShipLoc – CustLoc} combination and then used the proportions to multiply the weekly forecast estimate and obtain a daily forecast estimate. The weekly disaggregation proportions were estimated using three separate models.

1. **Seasonal model**
   This model was used to capture the slow-moving seasonal nature of daily order patterns. For each week, we analyzed the daily order patterns over the last three years and averaged the proportions for each day to estimate the weekly disaggregation proportions for each {Week – Prod – ShipLoc – CustLoc} combination.

2. **Trend model**
   This model was used to capture the more recent daily order trends. For each week, we analyzed the daily order patterns over the previous 13 weeks and averaged the proportions for each day to estimate the weekly disaggregation proportions for each {Week – Prod – ShipLoc – CustLoc} combination.

3. **Neural network model**
   This model was used to estimate daily order quantities based on the previous two years of daily order history using a neural network. For training, we used the daily order history from two years to estimate the daily order quantities of the subsequent year as a holdout using a neural network model. After training the neural network for 100 iterations, we used the daily order quantities from two years prior to predict daily order quantities for the current year. The estimated daily order quantities for the current year were normalized to get weekly disaggregation proportions for each {Week – Prod – ShipLoc – CustLoc} combination.

The weekly forecasts were disaggregated into daily forecasts using the three methods described above. For each forecast, we calculated the mean absolute percentage error (MAPE) for each {Prod – ShipLoc – CustLoc} triplet over the training period and used the forecast with the lowest MAPE as the enhanced daily forecast for a given triplet.

Results

We generated enhanced weekly and daily forecasts for the products that passed our data validation criteria, that is, order shipments in the current year and at least one year of order history. The performance of weekly and daily forecasts was evaluated using two metrics:

1. The absolute error between the forecasted order quantity and the actual order quantity.
2. Bias, which considers the ratio between the actual order quantity and the forecasted order quantity. A negative bias is interpreted as an under-forecast of the actual quantity whereas a positive bias is interpreted as an over-forecast of the actual order quantity.

For weekly forecasts, we evaluated the above measures for our enhanced forecast (FC-Enhanced) and the two forecasts provided by the CPG company – FC-Base and FC-Base+Expert at both {Prod} and {Prod – ShipLoc} levels.

Weekly forecasts

This approach demonstrated a significant improvement in forecasting accuracy using FC-Enhanced compared to both FC-Base and FC-Base+Expert, as shown in Figure 3, with results summarized in Table 1. At the {Prod} level we observed that FC-Base+Expert is better than FC-Base for lags 0-5; however FC-Enhanced is better than both the forecasts across all eight lags. At the {Prod – ShipLoc} level, we observed a dip in accuracy for both the FC-Base and FC-Base+Expert forecast. However, we did not observe such a dip in accuracy using FC-Enhanced forecast. Note lags 0 and 7 are missing in {Prod – ShipLoc} level because forecasts for these lags were not available in the data provided by the CPG company.

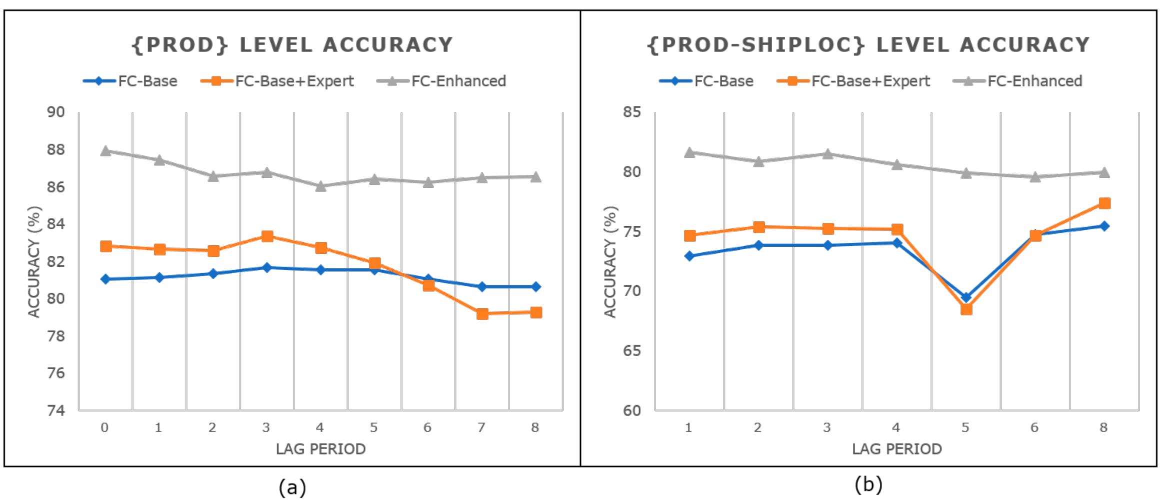

Figure 3. Comparison of forecast accuracy using three forecasts at (a) {Prod} level and (b) {Prod-ShipLoc} level. Significant accuracy improvement across all lags at both levels of the forecast are seen using FC-Enhanced.

Forecast level	FC-Base	FC-Base+Expert	FC-Enhanced
{Prod}	81.18 ± 0.36 %	81.70 ± 1.48 %	86.71 ± 0.57 %
{Prod – ShipLoc}	73.49 ± 1.80 %	74.46 ± 2.56 %	80.58 ± 0.74 %

Table 1. Summary of weekly forecasting accuracy at {Prod} and {Prod – ShipLoc} levels from Figure 3 summarized using mean ± standard deviation.

In terms of forecast bias, you can see that both FC-Base and FC-Enhanced have a consistent positive bias across all lags at the {Prod} level. FC-Base+Expert has a consistent negative bias across all lags at both forecasting levels as shown in Figure 4 and summarized in Table 2. A positive bias indicates that there is a systematic over-forecast of order quantities compared to actual order quantities, whereas negative bias indicates an under-forecast. A non-zero bias can lead to an over- or under-supply of shipments from ShipLoc DCs to CustTo DCs. The consistent positive bias using FC-Enhanced can be attributed to the use of a neural network procedure. We found that the neural network tended to produce positive non-zero outputs for zero valued inputs.

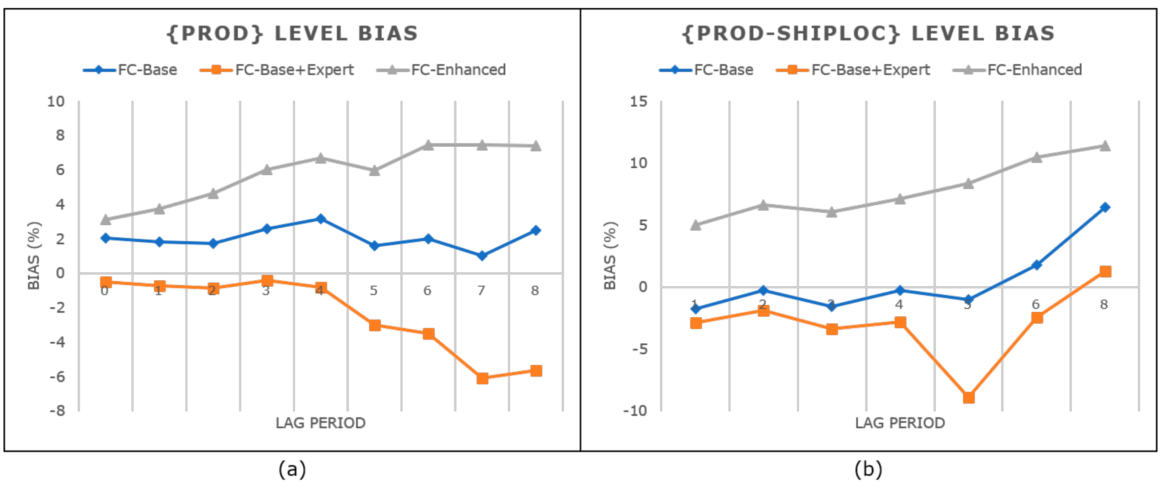

Figure 4. Comparison of forecast bias using three forecasts at (a) {Prod} level and (b) {Prod-ShipLoc} level. FC-Enhanced is more positively biased compared to FC-Base and FC-Base+Expert.

Forecast level	FC-Base	FC-Base+Expert	FC-Enhanced
{Prod}	2.06 ± 0.59 %	-2.39 ± 2.14 %	5.83 ± 1.54 %
{Prod – ShipLoc}	0.47 ± 2.67 %	-3.01 ± 2.79 %	7.87 ± 2.17 %

Table 2. Summary of weekly forecasting bias at {Prod} and {Prod – ShipLoc} levels from Figure 4 summarized using mean ± standard deviation.

In Figure 5 the results demonstrate the comparison of forecast accuracy, including point-of-sale and customer inventory data along with order history data. Both forecasts were generated using the NNTS method. The CPG company only had limited point-of-sale and inventory data for the current year. Recall that the order history data provided was available for seven years. Thus, in order to generate comparable forecasts, we used limited order history data over the same time period based on the availability of POS and inventory data.

The results clearly demonstrate an incremental improvement in weekly forecast accuracy using the additional POS and inventory data. For lags 0-5, we see around 1%-2% accuracy improvement by including the POS and inventory data. However, the improvement is less than 1% for lags 6-8. Furthermore, there is a sharp dip in accuracy for lags 7-8 for both the forecasts. We believe that this incremental improvement in forecast accuracy using POS and inventory data is due to the limited data availability.

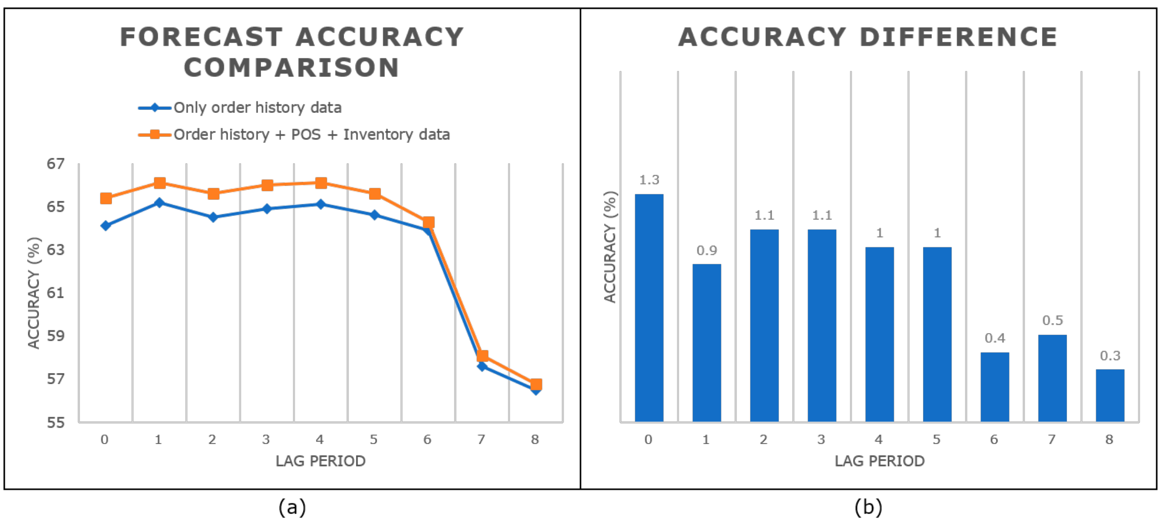

Figure 5. Comparison of weekly forecast accuracy using only order history data and using order history, point-of-sale and customer inventory data. (a) is a plot of accuracy of individual forecasts, and (b) highlights the lift in accuracy by including the POS and customer inventory data. Note the forecasts are at the (Prod – ShipLoc – CustLoc = customer) level.

Daily forecasts

Next, we enhanced daily forecast results, which were generated by disaggregating the improved weekly forecasts into daily forecasts using three separate models. Daily forecast accuracy and bias over 30 lags, that is, 30 days into the future, are shown in Figure 6 and summarized in Table 3. The CPG company did not provide any daily forecasts for comparison; hence we have presented only FC-Enhanced results.

Comparing Tables 1 and 3, we observe that compared to weekly forecasts, there is a dip in daily forecast accuracy at both {Prod} and {Prod – ShipLoc} levels. This is to be expected since we are breaking down weekly forecasts into daily forecasts. At both {Prod} and {Prod – ShipLoc} levels, there is a noticeable cyclical pattern in the daily forecast accuracy that repeats every seven days. For the majority of the lags, the {Prod} level forecast is more accurate than the {Prod – ShipLoc} level forecast, but we observe some exceptions, for example, at lags 5, 12, 19, 26. In terms of forecast bias, on average, the {Prod} level forecast has a positive bias, whereas the {Prod – ShipLoc} level forecast has a negative bias.

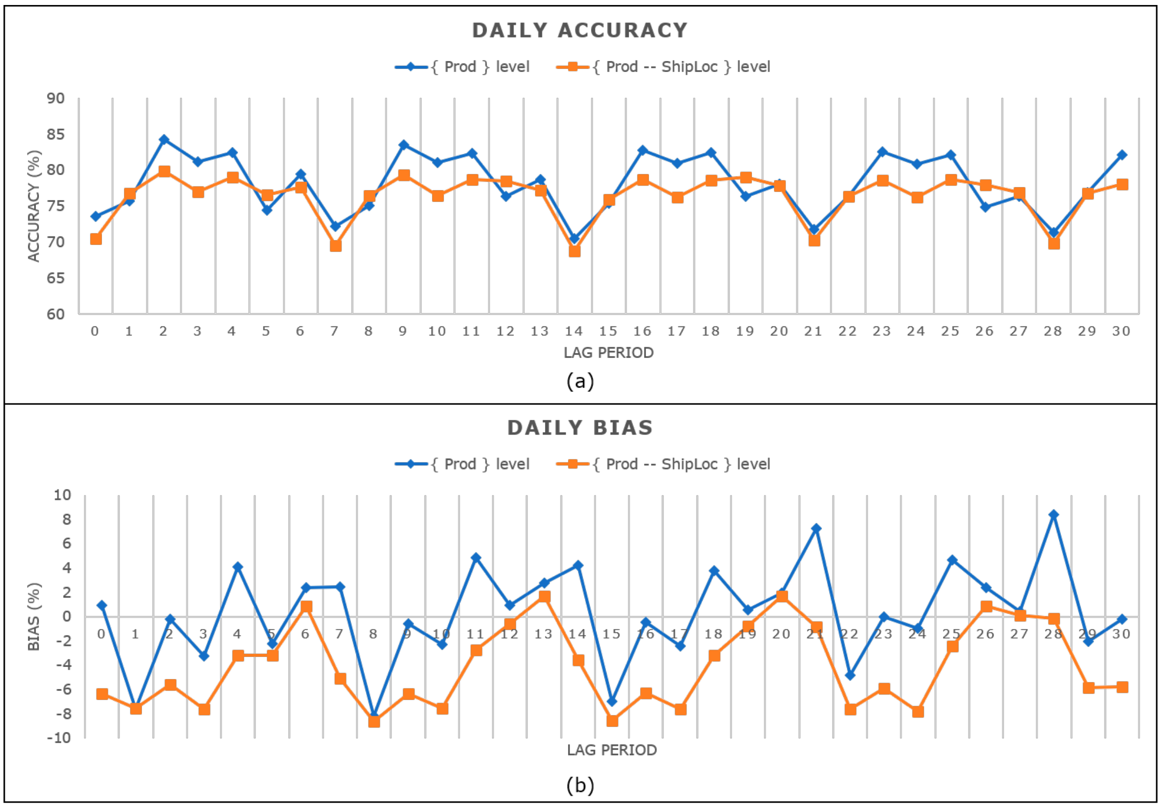

Figure 6. Daily forecasting accuracy at (a) {Prod} level and (b) {Prod-ShipLoc} level.

Forecast level	FC-Enhanced Accuracy	FC-Enhanced Bias
{Prod}	78.13 ± 3.97 %	0.30 ± 3.91 %
{Prod – ShipLoc}	76.39 ± 3.09 %	-4.06 ± 3.28 %

Table 3. Summary of enhanced daily forecast accuracy and bias at {Prod} and {Prod – ShipLoc} levels from Figure 6 summarized using mean ± standard deviation.

Conclusion

The use of machine learning clearly demonstrates how this technique can be used effectively to enhance weekly and daily product demand forecasts. This new weekly forecasting methodology uses a combination of traditional time-series models and machine learning methods to automatically choose the best model for each {Prod – ShipLoc – CustLoc} combination. This new approach using machine learning establishes the efficacy of these methods by improving short-term forecasts for a large CPG company.

At the weekly level, there was a significant improvement in forecast accuracy over existing forecasting procedures across multiple lag periods. The results also demonstrate that point-of-sale and customer inventory data clearly further improve the daily and weekly forecast accuracy. We believe that our methods provide a flexible, transparent and scalable solution for effective supply chain management for manufacturers, CPG companies and retailers.

Contact information

Your comments and questions are valued and encouraged.
Contact the authors at:

Kedar Prabhudesai
Kedar.Prabhudesai@sas.com

Varunraj Valsaraj
Varunraj.Valsaraj@sas.com

Dan Woo
Dan.Woo@sas.com

Jinxin Yi
Jinxin.Yi@sas.com

Roger Baldridge
Roger.Baldridge@sas.com

Tim Fairchild
Tim.Fairchild@sas.com

Learn more about SAS solutions for retail and consumer goods analytics.