A Mariner White Paper

MARINER

Demand Forecasting using Machine Learning to Reduce Working Capital

By: Colby Ford, Data Scientist

2719 Coltsgate Road • Charlotte, NC 28211
tel. 704.540-9500 • fax. 704.540-9501
www.mariner-usa.com

Contents

Abstract

A daunting problem for any business that carries and sells a physical product is understanding how much or how little of said product needs to be on hand. In a perfect world, companies would love to know that they will sell *x* number of an item, having *x* on hand, and having none leftover.

The popular solution to this tell-tale dilemma is forecasting; either by using sophisticated mathematical techniques or by tribal knowledge to better predict what products will sell and how to best meet consumer demand. Predict too few and customers go away empty handed, making the company miss revenue. Predict too many and see your clearance shelves overflow with unsellable product, missing out on gross margin.

Although forecasting is far from a new problem in the business world, we luckily have a new, automated approach. By employing Microsoft Azure Machine Learning, we can build a solution to process the sales history for all of a company's products, understand the trends, and better forecast how much of each product will be needed in the future.

The Problem

A large, global distributor purchases items from various manufacturers and then stocks the shelves of big-box home improvement retailers as well as smaller, independent shops with their products. Currently, the distributor utilizes a third-party forecasting solution to help decision makers make more accurate purchase orders from manufacturers. Although their current system is performing well, there is always room for improvement in any forecast unless it is 100% accurate all of the time.

This distributor also has another company goal: the reduction of working capital. Part of the agreement between this company and their retailers is to have a certain percentage of guaranteed stock on hand at all times. To be able to back this guarantee, they need to keep some extra items on hand in a warehouse in the event the retailer needs more stock in a pinch. However, with better forecasting, this distributor could reduce the amount of excess stock on hand, thereby reducing the working capital or money tied up unnecessarily in potentially unneeded items.

The remainder of this white paper is a real-world example of how we were able to improve demand forecasting. While the methodology, tools, and approach are accurate, the customer, data, and other information has been changed.

The Process

Our solution is to use the power of Azure to run a forecasting script in the machine learning (ML) to process the sales history of this distributor.

The process:

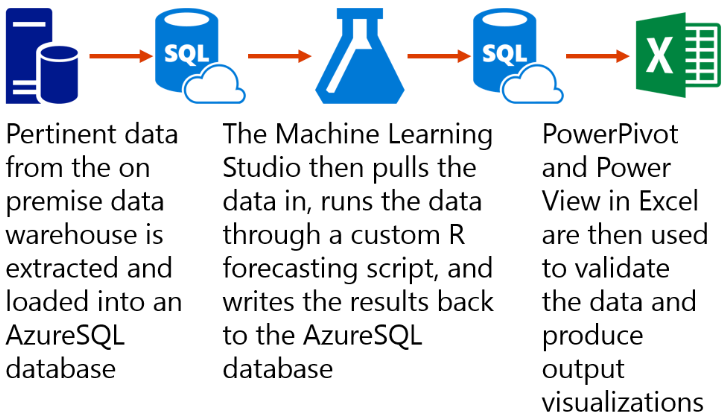

The "pertinent data" used for forecasting included Customer, SKU, Lead Time (the number of months ahead an item needs to be ordered), Date, and Quantity. To measure working capital performance and fill rate consequences, Product Line, Product Category, Unit Cost, and Unit Sales Price were also included. This generated a forecasting dataset of ~20 million rows for all products sold from January 2011 to January 2015, a sizeable dataset for forecasting. (In machine learning, quality, abundant data is always better.)

To choose a forecasting algorithm, a subset of the data was put through each of the algorithms. The results were tested to determine which of the algorithm choices worked best out of several options: Autoregressive Integrated Moving Average (ARIMA), Exponential Smoothing (ETS), and Naïve Bayes. Also there are two other components that have to be specified: test length (how many periods do we need to look back to make the forecast?) and seasonality (how long might a seasonal trend last in that length of time?). To determine which of the choices, either algorithm or test length and seasonality, work best, a measure called Mean Absolute Percentage Error (MAPE) was used (where the smaller the MAPE, the better).

MAPE is calculated by:

$$M = \frac{1}{n}\sum_{t=1}^{n}\left|\frac{A_t - F_t}{F_t}\right|$$

Where A is the actual quantity of a product sold and F is the forecasted quantity, each at time t to time n.

After running the experiment on an ensemble of parameter options from different forecasting algorithms and different test lengths (24 months, 12 months, etc.), along with multiple seasonality lengths (3 months, 6 months, etc.), a combination was chosen which yielded the lowest MAPE. Also, Lead Time was used to determine the number of months to forecast ahead. For example, if it takes 4 months to get a particular product in from the manufacturer, we needed to know how many will be needed 4 months from now. Finally, the experiment generated one large, final output full of forecasts.

The Outcome

The closer a forecasting system can get to the actual number of units sold at any given time, the better the total outcome will be. Having an accurate prediction can provide better insight into future product demand, leading to reduction of working capital by keeping less money tied up into excess products on hand.

Our MAPE for this forecast was ~7%, an improvement over the third party's MAPE over the same time period of ~16%.

	Total Units	Excess Purchase	Error
Actual Units Sold	180,080,135		
Third Party Forecast	152,748,391	28,027,796	-27,331,744
Azure Machine Learning Forecast	168,042,508	9,004,997	-12,037,627

**This is a comparison between forecasts on all products in sold in 2014.*

Simply looking at excess purchases (which only counts where a forecast is higher than the actual amount sold) shows that the Azure Machine Learning would greatly reduce working capital by decreasing the excess quantity ordered for each product.

Splitting this output by product line, we can see how each forecast performs. In general, both forecasts behave similarly with a few exceptions (such as the machine learning performance in Product A versus the third party forecast).

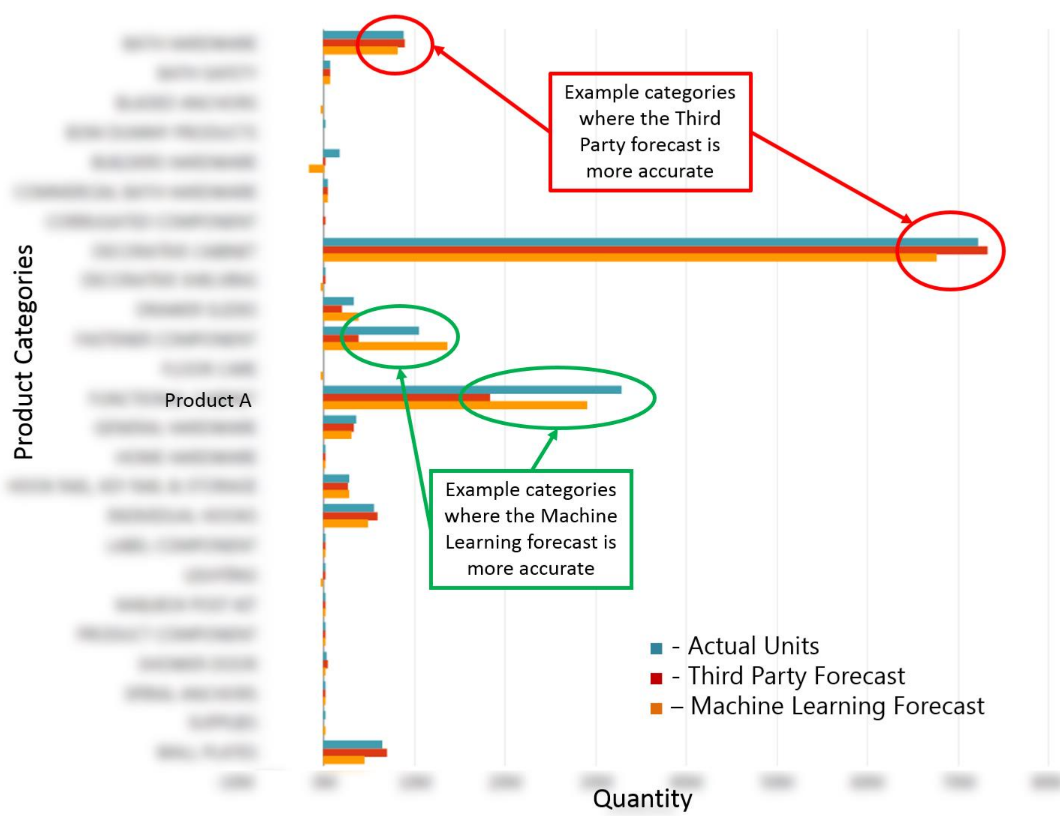

The Impact

The financial impact of a forecast is directly related to the working capital reduction measures taken by using a given forecast. Once the result for the Azure Machine Learning forecast was generated, the data was then joined back into the extra data that was pulled earlier. Unit Cost and Unit Sales Price were multiplied to the actual quantities sold, the third party's forecast, and the ML forecast. This gives a total overview of the financials for all products sold in 2014.

	Total Cost	Error	Total Sales	Error
Actual Units Sold	$184,532,956		$322,288,453	
Third Party Forecast	$156,525,383	-$28,007,573	$273,372,979	-$48,915,474
Azure Machine Learning Forecast	$170,391,396	-$14,141,560	$297,590,093	-$24,698,360

Using the Azure Machine Learning Forecast carries about half the error of their current system.

To further measure financial impact, we looked at inaccuracy in two different lights, Missed Gross Margin and Excess Purchase. Missed Gross Margin is when a forecast tells the distribution company to buy too little of a product, which results in customers not having product to buy (or retail stores not having enough product to sell). Excess Purchase is where a forecast tells the company to buy too much of a particular product. When this occurs, the customers only buy (or retailers only sell) what they need, which leaves us with an excess of products which will have to be sold at a discount to get rid of or be disposed of altogether.

	Missed Gross Margin	+	Excess Purchase $	=	Total Inaccuracy $
Third Party Forecast	$42,002,918.57	+	$28,027,796.66	=	$70,030,715.23
Azure Machine Learning Forecast	$11,502,109.20	+	$9,004,997.56	=	$20,507,106.76

Note the large decrease in the total inaccuracy between the current forecast system used by the company and the Azure Machine Learning result. Using Azure Machine Learning, the distributor could reduce working capital by millions of dollars by having a better forecaster on their side, in turn increasing the accuracy of what products they order and sell to their clients and also reducing the excess product they keep on hand.

Benefits and Future Work

Full Script Control

There are a multitude of benefits to using the Azure Machine Learning solution for forecasting, starting with having full control of the script that performs the forecast. You can directly make changes to forecasting length, seasonality, etc. along with changing the algorithm and even the format of the output. A proprietary ERP system may or may not allow for such flexibility.

Ensemble Modeling

It was mentioned earlier that many combinations of algorithm-test length-seasonality parameters were tested. The ML system does all of this for you. Plus, it chooses a parameter combination for each product line based on which yields the most accurate forecast. Ensemble modeling increases the overall accuracy by avoiding the "one size fits all" mentality for every product, a downfall of many ERP systems. Why should we assume that one product would have a similar sales trend as any other product? Thus, their models should be different and this is a task that Azure Machine Learning handles very well (whereas many ERP systems today cannot).

Simulated Scenarios

The Machine Learning experiment used a sliding window technique to simulate having been running the forecaster for all of 2014. A set of data is pushed through the forecaster for each time frame in 2014 to gain an individual forecast for each month. This provided the expected performance over an entire year rather than just running the forecast in a one-time snapshot. Using this method, comparisons can be made between the entire year's worth of forecasts from the third party tool against the Azure Machine Learning tool. The next step would be to set the machine learning forecaster live alongside the existing third party solution and allow them run for a while together to truly see how they compare with new incoming data. This way, there is no danger in setting the ML experiment up to run to see how it behaves over time before making a drastic implementation decision.

Future Work

Future work will include full implementation into the distribution company's operational supply chain system and demand requirements planning process. For the places where their current forecasting system performed better than Azure Machine Learning, much of it can be explained by "insider knowledge" that their system had that ML did not.

For example, take this client's purchases in 2014 for a particular item:

Month	Jan	…	Jul	Aug	Sep	Oct	Nov	Dec
Actual Units	38,894	…	51,077	31,003	0	0	0	0
3rd Party Forecast Units	58,340	…	42,990	28,112	0	0	0	0
ML Forecast Units	33,553	…	54,274	51,668	60,658	67,987	40,156	45,629

The sum of actual units sold is 355,536 whereas the sum of their third party forecaster was 377,662 (a difference of 22,126). The ML total, 559,414 (an over forecast of 203,878 units) is much farther off than the third party's, but their system was connected to the distributor's system, which noted that this particular item was discontinued after August, 2014. The ML system was ignorant of this information and therefore attempted to forecast past the discontinued date. If you take out the discontinued months, the machine learning forecast was 344,984, an under forecast of only 10,552 units. This is just one example of how the machine learning system will be even smarter once fully integrated.

A big goal for this global distributor is to reduce working capital. As stated before, the easiest way to do this is to avoid excess product purchases. With Microsoft Azure Machine Learning on board, demand forecasting will be more accurate and help them reach their working capital reduction goal with ease.

About the Author

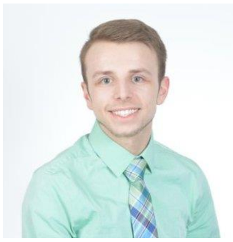

Colby Ford, M.Sc., Data Scientist

Colby is a Data Scientist at Mariner. Using a background in mathematics and statistics, he puts the Azure Machine Learning system to work along with R to gain insight from data. Outside of work, Colby is a pianist, Netflix junkie, bioinformatics researcher and Ph.D. student at the University of North Carolina at Charlotte.

Connect with Colby Ford

LinkedIn: www.linkedin.com/in/colbytylerford/
Twitter: @colbytylerford
Website: www.colbyford.com

About Mariner

Microsoft Partner
Gold Data Analytics
Silver Data Platform
Silver Cloud Platform
Silver Intelligent Systems

If you're looking for a business partner that can harness all types of information to enhance customer experiences, optimize internal processes or create new business models, then look no further. Leveraging big data, cloud, social, predictive analytics, Internet of Things, decision management and mobile technology, Mariner quickly draws a target on the results you want and then delivers your digital business solution.
Geared toward helping the leaders of mid-market companies, we help you drive automated decision making and empower decision makers with actionable, trusted information for your company, your department and you. Mariner is a Microsoft Certified Gold Partner with competencies in Data Analytics, Data Platform and Intelligent Systems Service. Mariner-Insight to Achieve. www.mariner-usa.com. For more information about our Digital Business solutions, visit http://www.mariner-usa.com/digital_business

Copyright Information