LEGION

WHITE PAPER

Achieve New Levels of Accuracy with Legion Demand Forecasting

July 2021

CONTENTS

sales@legion.co | +1.650.980.4873 | legion.co | @legiontech

Introduction

The advanced machine learning (ML) based demand forecasting engine is a key pillar and a foundational element of Legion's Workforce Management (WFM) platform. Legion's ML-based demand forecasting has several advantages over other approaches:

- Learns more subtle patterns
- Incorporates different types of data
- Scales and produces consistent and accurate forecasts week over week for hundreds of thousands of data sets.

There are a lot of claims around accuracy, but what exactly is an accurate forecast?
And why does the way you report and measure it matter? It's easy to measure accuracy – you get a forecast, wait for the actuals to come in, then compute the difference, and perhaps you express the difference as a percentage. That's the easy part. But, it leaves many open questions:

- Over what time period should you measure?
- If you have many data sets, how do you aggregate the results?
- And most importantly, if you have a number representing accuracy, how do you know if it's good or bad?

This white paper provides insights about demand forecasting approaches and best practices. It describes how Legion WFM applies answers to the questions above and how we measure and report the accuracy of our demand forecasts.

Measuring Accuracy

Three main principles must be applied when measuring demand forecast accuracy.

1. Consider the predictability level of the forecasted data set

   First, data sets differ widely in how predictable they are. Some data sets have strong and regular patterns; they are more predictable. Others are more random; they are less predictable. We know this fact intuitively. For example, we know the traffic flow in an office park is more predictable than in an open shopping plaza. But we need to find a way to quantify this reality. That means any statement about the forecast accuracy has to consider the predictability level of a forecasted data set. But, there's no universal number that represents an "accurate" forecast. And, there are no industry benchmarks for forecast accuracy, just as there are no industry benchmarks for predictability.

2. Look at results over an extended period of time

   Any measurement of forecast accuracy needs to be done over an extended period. Measuring accuracy day to day or even over a week or two isn't informative. Most forecast models wait before responding to deviations from their forecasts to see whether or not a deviation is, in fact, an established pattern – because they learn that gives the best results over the long run. So, to get a true measure of forecast accuracy, we need to look at results over several weeks.

3. Compare the results of a forecast model to a target based on the data's predictability

   The best way to measure demand forecast accuracy is to compare the results of a forecast model not to an arbitrary number but a target based on the data's predictability. One way to obtain a target is to use a Naive model(s) because both models will be equally affected by variations in the data set predictability. For example, random walk or moving averages are commonly used as baseline models for generating targets.

To measure the accuracy of our forecast models, we compare our forecast accuracy averaged over a trailing 8-week period to the target accuracy. That ensures we measure the overall accuracy of our forecast models instead of short-term trends that may not persist.

sales@legion.co | +1.650.980.4873 | legion.co | @legiontech

Forecast Accuracy Metrics

Now, we've established that forecast accuracy should consider the predictability of a data set, and it should be averaged over several weeks. So let's turn to what specifically should be measured.

In the end, all forecast accuracy metrics measure how forecasts compare to actuals. The only difference is in the details of how the metrics are aggregated, averaged, or the audience they are designed to serve. For example, some metrics are designed to give data scientists deep insights into the performance of a forecasting model. Others are developed to help managers make higher-level operational or business decisions.

Our preferred metric for detailed performance reporting is the weighted mean absolute percentage error (WMAPE). We calculate and report this metric for each data set because each one has its unique forecast model, and data scientists typically want insights at this level. We calculate WMAPE over a trailing 8-week period using this formula. For each day in this period, we compute the absolute percentage error as:

- The absolute value of (forecast minus actuals) divided by actuals
- Then we compute the weighted average of the absolute percentage error using the actuals as weights

WMAPE is a commonly used metric with several desirable mathematical properties and is very familiar to data scientists. For operational and management reporting, though, it has some drawbacks – mainly that although it's a percentage, there's no particular significance to a 0% WMAPE or a 100% WMAPE. Plus, it's not very intuitive. So for management reporting, we use Forecast Accuracy (FA), a metric commonly used in management reports.

FA is very similar to WMAPE, except that the daily absolute percentage error is calculated as:

- The absolute value of (forecast minus actuals) divided by the maximum of the forecast and actuals
- Using the maximum ensures the result will always be between 0% and 100%
- Then we compute the weighted average, as discussed above, to get the forecast error
- The forecast accuracy is 100% minus the forecast error
- As with WMAPE, we calculate FA over a trailing 8-week period

Weighted MAPE

$$\frac{\Sigma\,(forecast - actual)}{\Sigma\,actual} \times 100\%$$

Weighted FA

$$\frac{\Sigma\left(1 - \frac{abs(forecast - actual)}{max(forecast,\ actual)}\right) \times actual}{\Sigma\,actual} \times 100\%$$

For management reporting, the advantage of FA over WMAPE is that it's always between 0% and 100%; we also aggregate the FA across all data sets by computing the weighted average over all data sets with the same demand driver to get a single overall number. Then, managers can use that number to make business decisions. Conversely, WMAPE is reported at each data set level, and data scientists use it to get performance insights for each forecast model.

Achieve New Levels of Accuracy

With Legion WFM, you get much higher accuracy levels than other methods can achieve – and it's all automated. There are several reasons why Legion WFM forecasts are highly accurate.

First, Legion's scalable infrastructure can create, train, and manage forecast models at the individual data set level. In other words, each demand category at each enterprise location gets an independent forecast model that is specifically trained to learn the unique patterns of that data set and project these patterns.

At any given time, Legion's infrastructure manages hundreds of thousands of these models. Each of them is continuously evaluated against its target accuracy. Suppose a model strays from its target (as can happen when the external environment changes). In that case, a process is triggered automatically to select a new model.

This level of model specificity is not possible without the scalable, cloud-based machine learning infrastructure that Legion developed. It results in higher accuracy than would be possible with more general-purpose models.

Second, Legion incorporates a variety of external data into its machine learning models. The models learn the impact of these data on demand, incorporating this learning into their forecasts. For example, Legion's infrastructure has data, including holiday calendars, weather, and local events. In addition, customers can provide data, such as sales promotions or other events that could affect their demand. All this improves the accuracy of Legion's forecasts.

Why Legion WFM

Legion WFM is built on a foundation of data science. Legion's AI-powered intelligent demand forecasting engine uses machine learning to compute highly accurate, scalable, and consistent demand forecasts automatically. Legion's approach of using highly specific forecast models tailored to each data set, coupled with external data, leads to higher accuracy than other methods can achieve. And all of this is done within our automated and scalable infrastructure. As a result, you can achieve much higher accuracy levels than other methods can achieve – and it's all automated.

Request a demo to see Legion demand forecasting in action and learn more about how we can help you save time, increase accuracy, avoid errors, and deliver consistent results across your organization.

About Legion

Legion Technologies is the global leader in AI-powered workforce management (WFM). The Legion platform optimizes labor efficiency and employee engagement simultaneously – at scale – enabling companies to significantly reduce labor costs, cut scheduling time in half, and achieve 95% employee engagement. Legion has been recognized as a Gartner Cool Vendor and included on JMP Securities' Hot 100 List of the Best Privately Held Software Companies. The company is backed by Norwest Venture Partners, Stripes, First Round Capital, XYZ Ventures, Webb Investment Network, Workday Ventures, and NTT DOCOMO Ventures. For more information, visit https://legion.co.

Appendix: Forecast Accuracy Reports

The following forecast accuracy reports are added to our reporting infrastructure, and managers can download them weekly from the dashboard.

Data Science Report

This report contains detailed metrics, and it's intended for data scientists or analysts. It has the following columns:

COLUMN NAME	DESCRIPTION
TargetWeek	The week that the report was requested (e.g., Week of 2021/01/07-2021/01/13)
Location	Display name of location
Channel	Display name of channel
Category	Display name of category
DemandDriver	Demand driver type (e.g., Items, Amount, Traffic, Transactions)
ModelType	Type of selected model (e.g., Random Forest, Neural Network, etc.)
ModelSelectionTime	Timestamp of the last forecast generation time
ModelTrainingTime	Timestamp of last model training
ForecastGenerationTime	Timestamp of the last forecast generation time
ForecastAppRefreshTime	Timestamp of the last time the application's forecast was refreshed
ValidationRMSE	Validation set RMSE during training
ValidationWMAPE	Validation set WMAPE during training
ValidationBias	Validation set bias during training
TargetTrailingDailyWMAPE	Target for trailing 8-week daily WMAPE derived from naive model(s)
TargetTrailingDailyFA	Target for trailing 8-week daily FA derived from naive model(s)
TargetTrailingDailyBias	Bias when computing trailing 8-week daily targets from naive model(s)
LegionTrailingDailyWMAPE	Trailing 8-week daily WMAPE of Legion forecast
LegionTrailingDailyFA	Trailing 8-week daily FA of Legion forecast
LegionTrailingDailyBias	Trailing 8-week daily bias of Legion forecast

Management Report

This report contains metrics for management reporting and decision-making. It's derived from the report above and has the following columns:

COLUMN NAME	DESCRIPTION
Date	Date of report
Location	Display name of location
Channel	Display name of channel
Category	Display name of category
DemandDriver	Demand driver type (e.g., Items, Amount, Traffic, Transactions)
Forecast	The forecast demand for the date
Actual	The actual demand for the date
LegionFA	The daily forecast accuracy (FA) of Legion forecast for the date
TargetTrailingDailyFA	Target for trailing 8-week daily FA derived from naive model(s)
LegionTrailingDailyFA	Trailing 8-week daily FA of Legion forecast

Management Dashboard Metric

An overall enterprise-level forecast accuracy metric can be derived from the management report above as follows:

1. The target daily FA and the 8-week Legion FA are averaged across all locations, weighted by actuals for each demand driver type.
2. For enterprises with more than one demand driver type, the target daily FA and the 8-week Legion FA are averaged (without weights) across all demand drivers.

Formulas

Weighted MAPE

$$\frac{\Sigma\,(forecast - actual)}{\Sigma\,actual} \times 100\%$$

Weighted FA

$$\frac{\Sigma\left(1 - \frac{abs(forecast - actual)}{max(forecast,\ actual)}\right) \times actual}{\Sigma\,actual} \times 100\%$$

sales@legion.co | +1.650.980.4873 | legion.co | @legiontech