Using Machine Learning to Forecast Demand – *Looking Back to the Future*

BY JACOB JHAVERI & MÅNS SKOGMAN (JUNE 2026)

A challenge that every company faces is the attempt to predict the future. This issue transcends most divisions, including sales and operations planning. Sadly, fortune tellers are of short supply. Thus, an array of methods to achieve accurate predictions have been developed over the years, including statistical models and machine learning algorithms.

The presence of machine learning (ML) tools in supply chain management has increased significantly over the past couple of years, with demand forecasting being the most common application. However, despite the common view on artificial intelligence as Mr Know-it-all, using ML to forecast demand has its limitations.

Asking an employee to predict the sales quantity for the next month will activate a reasoning that takes a large amount of factors into account. How much was sold last month? Do we have any new customers in the pipeline? Are there any issues in our production/delivery process? Is there a new competitor on the market? The questions go on and on. A highly skilled employee will deduce a reasonable prediction, but the risk of missing the mark remains.

Tree-based ML models, on the other hand, base their predictions solely on the historical information used to construct them. Most commonly, the historical sales data represents the main information source on which the models are trained. Thus, looking back to predict the future.

Using the rearview mirror creates an obvious vulnerability. What if our mirror is foggy? What if it is angled incorrectly? The forecasts are only as good as the data the models are trained on. The data available in practice offers an unlimited amount of flaws and inconcistencies. So, how do we handle imperfect data?

Picture a demand forecast through a magnifying glass. Is the demand forecasted on a monthly or weekly basis? Are we forecasting every unique article or larger product groups? When zoomed in, we get more precise information about when and how great the demand is, including the variance, but also the noise. On the other hand, zoomed out will look more stable and easier to predict, but result in less useful information for the demand planning operations.

Photo by Hasnain Sikora on Unsplash

The foundation of the thesis consisted of reviews of literature within demand forecasting and machine learning. Four ML models were developed and tested iteratively through design science research methodology. The findings in the thesis showed that the structure of the data as well as the data granularity had a larger impact on the results than the choice of model. Using methods such as first-order differencing and scaling allowed the ML models to generate forecasts that exceeded the performance of the current forecast at the case company.

However, the findings also showed the limitations of predicting through the rearview mirror. The questions posed through human intelligence offer useful insights. Insights that cannot be replaced by simple ML models looking back at historical data. Despite outperforming the alternatives, ML is not a fortune teller on its own. To create a fortune teller, one must therefore combine human intelligence and artificial intelligence.

This popular scientific article is derived from the master thesis: Machine Learning as a Tool for Demand Forecasting: *A Study at HMS Networks*, written by Jacob Jhaveri and Måns Skogman at the Department of Industrial and Mechanical Sciences, Division of Engineering Logistics (2026).