

Technical Case Study: Advanced Retail Forecasting Using AI

Introduction:

Our customer, a leading retail company, had a critical need to improve their day-to-day forecasting system. With a vast array of 500 item ranges available in more than 1000 stores, accurate demand forecasting was crucial for efficient inventory management, optimizing cost, and enhancing customer satisfaction.

Challenge:

The existing forecasting system lacked efficiency, robustness, and flexibility. It was unable to accurately factor in all the critical variables that significantly influence retail demand, such as payday weekends, bank holidays, promotions, and weather conditions. All of which vary on an individual Item and Store level.

Solution:

Two proof-of-concept projects were undertaken to demonstrate the effectiveness of a new forecasting. The initial proof-of-concept was to take a sample of 4 months of historical Sales data across 20 stores and see if we could build a forecasting algorithm which improves on their existing one.

Over a few weeks of trial and error, testing different ML software and a myriad of variables we managed to come up with a new algorithm incorporated a multitude of parameters, including day, week of the month, month, payday weekend, bank holiday, promotion, and weather. Furthermore, it employed time-series forecasting, allowing it to consider the previous month's trend when forecasting for the current period. This enabled us to use all the historical data to train the model and not be hindered by Sales patterns for Items and Stores changing year upon year.

We eventually settled on predictive agents, which were built specifically for the task, and perfect for the challenge facing our customer. This software is unlike most other ML software's as it is suited and optimised for small data. This allowed us to start predicting with as little as 14 data points, so 2 weeks after a new Item hit the shelves, we could start producing forecasts for it. It was also extremely quick and lightweight, which would enable us to be able to produce the number of forecasts within the timeframe available each day.

Following the initial proof-of-concept, they were happy with the results, but wanted to know how it would perform during, historically, the most difficult periods to forecast (i.e., Christmas), so we performed a similar process and built a model that produced accurate results for seasonal periods (models that were unique to Christmas, that included different parameters such as 'Days to Christmas' etc).

Happy with these results, the Customer then gave the green light to go with a full build and rollout which would lead up to 500 Item ranges and 1000 Stores nationwide.

A distinctive feature of the solution was its full configurability with a user-friendly front-end application that enabled real-time analytics, error logging, model health/diagnostics, and full configurability. The process, while being completely automated, allowed granular adjustments such as the length of the forecast period, the length of the Trend, the days that go into training the model (extreme periods can be removed that may prove to be unreliable predictions for the future such as certain times during COVID). Providing the Retail company with maximum control and flexibility.

Implementation:

The algorithm was applied across the range of 500 items in 1000 stores. Each Item in each Store being individually modelled as every single combination has a unique Sales pattern, which, if not modelled individually would sacrifice accuracy. Despite the vast volume of data, the automated system could produce 7 million forecasts in a mere 2 hours, a remarkable feat that underscored the power and efficiency of the new solution.

At every turn, the solution was built with best practice, and data optimisation in mind. The sheer number of forecasts and data produced meant that it was critical that there was no redundant logic or data.

The whole process is fully automated and requires no manual intervention (though is configurable if necessary).

Results:

The implementation of the AI forecasting agent led to significant improvements in the customer's forecasting system. The enhanced accuracy in demand prediction led to:

- Optimized inventory management, leading to a decrease in costs associated with overstocking and understocking.
- Increased customer satisfaction due to better availability of products.
- Greater operational efficiency with the automation and configurability of the forecasting process.
- Improved decision-making ability due to the incorporation of real-time analytics in the front-end application.

The improvements noted during the two proof-of-concept projects translated into real world gains when the solution is being implemented across all stores, confirming the effectiveness of the new algorithm and its adaptability to complex, real-world retail operations.

Conclusion:

Prophesy transformed the Retail company's approach to demand forecasting. With the new solution, the company is now able to react quickly and intelligently to shifts in demand patterns, effectively optimizing inventory, enhancing customer satisfaction, and driving overall business growth. This case study underscores the power of advanced algorithms and predictive analytics in revolutionizing retail operations.