Beyond Traditional Forecasting: The Hybrid AI Advantage

Manhattan

Contents

The Evolution of Demand Forecasting

Developing accurate demand forecasts for large volumes of diverse products with potentially complex and continuously evolving demand patterns has always been a challenge. For most suppliers, manufacturers and retailers, it's educated guesswork at best. But the job has gotten measurably harder over the past decade, as organizations deal with extraordinary unpredictability.

Consumer behavior has grown more volatile, driven by rapidly changing preferences, shorter product life cycles, the rise of omnichannel commerce and an explosion of pickup and delivery options. Globalization, accelerating market changes and more frequent external influences have added complexity to extended supply chains, making forecasting more susceptible to disruptions and demand shifts. Additionally, the proliferation of data sources from external demand factors, ranging from social media to economic indicators, have overwhelmed traditional forecasting methods.

These challenges often lead to unfavorable outcomes. Lost sales, increased costs and poor customer service due to inaccurate forecasts, out-of-stocks, overstocks, unsuccessful promotions and markdowns are detrimental to any company's bottom line. It's difficult to remain profitable when writing off hundreds of thousands – or even millions – of dollars' worth of merchandise on an annual basis or watching a customer base erode due to inadequate service.

Many demand forecasting solutions and tools exist in the marketplace today, leveraging a variety of mechanisms to help solve these problems. These mechanisms use traditional time series forecasting methods or perhaps some more advanced machine learning (ML) methods. As we'll see, each method has benefits and challenges.

To maximize the benefits of traditional forecasting methods and ML, while minimizing the challenges of each, Manhattan has exclusively developed a new generation of hybrid AI demand forecasting that combines advanced, unified, broad-spectrum, statistical demand modeling with the extensible and adaptive predictive power of ML.

It's called the Unified Forecasting Method with Artificial Intelligence, or UFM.ai for short.

Forecasting Complex Supply Chains Is Hard

"It's tough to make predictions, especially about the future."

– Yogi Berra

Given the negative and costly effects of poorly forecasted demand, why aren't organizations more focused on this aspect of their supply chains? Quite simply, it's hard.

Consumer behavior can be unpredictable. Customer preferences, tastes and purchasing habits change rapidly and impulsively. Trends, promotions, seasonality and external events like weather or economic changes can cause sudden spikes or drops in demand.

Correctly modeling variables like seasonality, price changes, promotions, competitor actions, economic conditions and all their interactions is complex. And any missing or incorrectly weighted factors can skew forecasts.

Demand patterns might be continuous, erratic, intermittent, lumpy, seasonal or any assortment of these. And they often shift continuously or evolve in their rates and patterns due to trends, product life-cycle stages or even disruptions.

Unanticipated factors like governmental regulation, global strife, union strikes and impactful social milestones add volatility and less certainty.

Operational issues also come into play for retailers and wholesale distributors, who often need to manage, analyze and forecast very large SKU volumes.

Facing these challenges, organizations often focus on finding easier efficiency gains across labor, distribution and transportation on the execution side of the supply chain. However, no matter how good the execution systems and teams are, they cannot overcome the lack of inventory in the right place to begin with. The single largest capital investment by most organizations is tied to their inventory, and that is also where to find the largest efficiency gains left in the supply chain.

Traditional Demand Forecasting Methods Are Limited

Statistical time series models have been used for demand forecasting for a long time, with proven and established methods and practices. They work well for short-term forecasting in stable environments where demand patterns are consistent and external influences are minimal. They rely on historical data and model recurring patterns effectively, which is useful for products with stable demand histories, because they assume that past trends and seasonality are good indicators of future demand. Statistical time series models generally require fewer data points than ML models. They often perform quite well, even when some external factors (like promotional or macroeconomic data) are not available, since they rely primarily on internal historical data. They are generally easy to implement and computationally efficient and interpretable, making them effective for products with clear trends or stable seasonal patterns. Many industries have successfully implemented these models, making them a reliable and trusted option. However, there are limitations.

- Time series models struggle with outside-in demand planning objectives that incorporate external data, are forward-looking and are more dynamic.
- Time series models have difficulty with demand volatility, such as sudden spikes or drops from one-time events or disruptions, as they assume past patterns will continue indefinitely.
- These models have difficulty incorporating market dynamics like consumer sentiment, economic indicators, weather conditions and competitor actions, instead of solely relying on past sales trends.
- Time series models are typically individually tailored for specific forecasting problems, which can require multiple, different models to address the variety of demand patterns across a product portfolio. Managing and optimizing these models becomes complex as the portfolio grows.
- When demand dynamics change, new statistical time series models must be created, especially for large portfolios. Automation can help but it often struggles with changing demand patterns, leading to errors when switching models.

  For example, demand for winter parkas fluctuates dramatically between seasons, requiring different models to forecast demand for off-season (heavily sparse, intermittent) and in-season (fast-moving, seasonal).
- Automated model-switching strategies often accumulate forecasting errors before switching methods, creating a cycle of accuracy loss, which negatively impacts planning and inventory optimization.
- Time series models generally assume linearity, making it difficult to capture complex relationships or sudden changes in demand from external factors.

Machine Learning: Powerful, But Not a Panacea

While statistical time series models have been used for a long time due to their intuitiveness and effectiveness in capturing trends and seasonality, their potential limitations quickly become evident in today's more complex and dynamic supply chains. Machine learning, however, thrives in environments where demand is volatile, nonlinear and influenced by numerous external factors. It excels at uncovering intricate patterns and relationships in the data that traditional models often miss, making it a powerful tool for more accurate and adaptive demand forecasting.

One of the key strengths of ML models is their ability to find nonlinear patterns within demand, which are often driven by factors such as promotions, pricing, weather, customer preferences and broader economic conditions. Unlike statistical models that assume demand closely follows historical patterns and trends, ML models can dynamically learn and evolve, continuously improving their performance as new data comes in. This makes them particularly valuable in fast-changing markets where demand conditions can shift rapidly.

Machine learning models can seamlessly integrate a variety of external data sources, from competitor actions and market trends to social media sentiment and global events. For example, a sudden shift in consumer preferences due to a viral trend or a disruption in the supply chain from geopolitical events can be quickly incorporated into the model's learning process, leading to forecasts that are more aligned with real-time conditions. This proactive responsiveness offers demand planners an added agility and even a competitive edge by helping to anticipate demand changes that traditional models might miss.

Another advantage of ML models is the ability to automatically identify and scale the most relevant features influencing demand. Through techniques like feature selection, ML algorithms can continuously and autonomously assess and weight key demand factors that either influence or correlate with demand. And they can do this without requiring specialized or extensive manual intervention. This capability reduces the reliance on domain expertise and allows businesses to uncover insights from vast amounts of data that may not be immediately apparent, even to seasoned analysts.

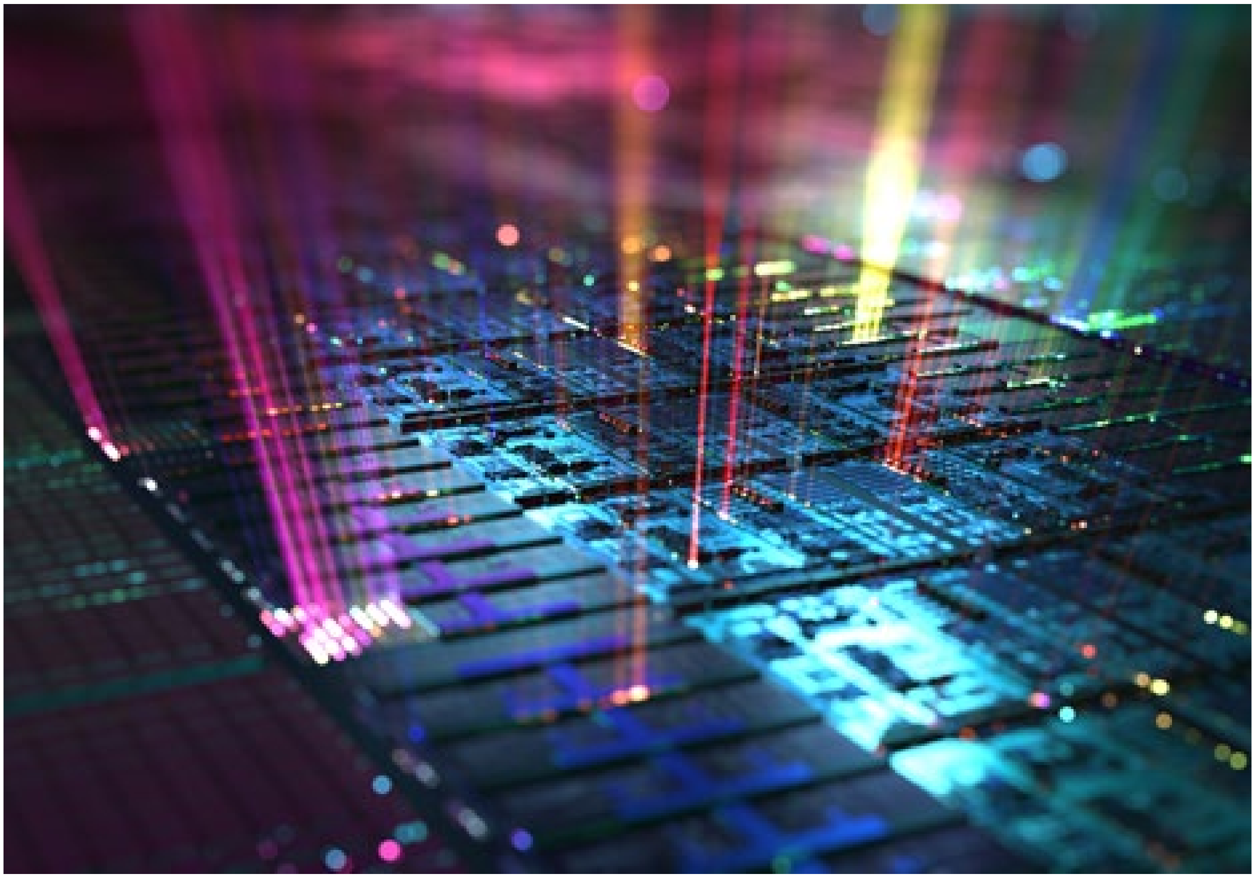

Moreover, ML models can often handle high variability and volatility in demand with greater precision. They excel in scenarios where demand patterns may be erratic or previously unobserved, such as product launches, seasonal spikes during in-season promotions or supply chain disruptions. Their ability to leverage external data sources, along with continuously retraining on new data, ensures that forecasts adapt to quickly evolving market conditions and stay accurate over time. The models do this without the need for frequent manual adjustments or recalibration, which is often necessary with statistical models.

However, ML models can be more complex to implement and maintain than traditional models. They require significant amounts of data (good quality, meaningful data, of course), expensive computational resources, and specialized expertise to set up and monitor properly. One of the key challenges faced by ML models is when input or external data is sparse, limited, incomplete or of poor quality (garbage in, garbage out). In those cases, they struggle to provide accurate or even appropriate predictions. They might overfit the model to irrelevant or noisy data or be misguided due to an incomplete view of data, all of which results in degraded performance. Additionally, most predictive ML models will provide a prediction, no matter what, regardless of the quality within some of the training data or input data passed to the model.

Hybrid AI and the Future of Forecasting

Hybrid AI models combine statistical time series methods with ML algorithms, offering a robust and balanced approach to demand forecasting. These models leverage the strengths of both traditional and modern methods, providing enhanced accuracy, reliability and adaptability in a wide range of forecasting scenarios. By integrating the interpretability, statistical modeling and stability of time series methods with the extensible and adaptive pattern-recognition capabilities of ML, hybrid models can handle a very wide variety of demand data environments more effectively.

The statistical aspects of hybrid models are based on well-established time series methods for capturing trends, seasonality and linear relationships within demand data. That makes them highly effective when forecasting in relatively stable demand data environments where historical patterns are expected to extrapolate well to future demand patterns. Also, given that these methods are based on decades of proven demand forecasting science, they often hold up well when data conditions become temporarily limited or incomplete. Machine learning might be challenged with such limited information conditions, because it must build its understanding and representation purely from observed data.

On the other hand, when data conditions are favorable, the hybrid ML component enhances the overall approach by handling the more complex, nonlinear relationships within the data. This allows the model to incorporate external variables such as weather, competitor actions and economic indicators – factors that traditional time series models would usually fail to consider. Machine learning algorithms within a hybrid model can learn from real-time data, tuning the forecasts dynamically to account for sudden demand shifts, such as those caused by market disruptions or external shocks. By combining these approaches, hybrid models are the best of all worlds, effectively overcoming the limitations inherent in using either method alone.

Hybrid models are understandably the most complex to develop, requiring expertise in both the statistical and ML methods and the ability to cleverly combine them into a unified model. However, the payoff is often significant in terms of improved forecast accuracy and adaptability, particularly in volatile environments where demand is influenced by a mix of historical patterns and external variables. As a result, hybrid AI demand forecasting enables planners to predict demand more accurately, especially in complex, dynamic environments.

Unified Forecasting Method with Artificial Intelligence

Manhattan Active Supply Chain Planning exclusively includes UFM.ai, the first-ever, commercial, hybrid AI demand forecasting technology. It uses a unified statistical modeling system built from well-proven statistical time series methods, infused with the powerful, data-driven capabilities of ML.

By unifying key aspects of well-established statistical models such as Croston's, Variable Response Exponential Smoothing, Holt-Winters and ARIMA/SARIMA with the adaptive, scalable and predictive strengths of ML, UFM.ai delivers a comprehensive solution to modern demand forecasting challenges. This unique, groundbreaking, hybrid approach is more accurate, robust and versatile than any individual forecasting method, making it particularly well-suited for today's complex and dynamic markets.

UFM.ai uses a unique comprehensive statistical foundation that is a refactored combination of principal components from leading statistical forecasting methods into a single formulation. The result is a unified system of equations, allowing it to holistically capture trends, seasonality and intermittent demand patterns – or any combination thereof – within one, versatile statistical model.

This statistical modeling aspect of UFM.ai allows it to model (analyze) local demand dynamics exceptionally well across a diverse range of demand patterns, where historical data provides a reliable basis for predicting future demand. The ML aspect of UFM.ai significantly enhances the model further by incorporating extended internal factors (expanded item and location characteristics and relationships) and external factors (economic indicators, weather, social influence). The identification of nonlinear patterns in data that traditional statistical methods often miss makes UFM.ai more adaptable to real-world market dynamics.

Since ML models have the potential for poor performance when faced with sparse, limited, incomplete, inaccurate or noisy demand factor data, UFM.ai leverages its unified statistical foundation to bolster the predictive model to provide reliable forecasts when external data is unreliable or unavailable. In situations where standalone ML models might struggle with poor-quality data, UFM.ai's statistical capabilities continue to offer robust and consistent predictions derived from well-established statistical methods and internal historical data.

Combining a statistical time series model with machine learning into a single integrated model was proven in the Makridakis M4 Competition to be the most accurate and robust forecasting method, while also providing a fail-safe mechanism that ensures reliable performance even in the face of imperfect training or input data. This makes UFM.ai far more versatile and dependable than pure machine learning approaches.

UFM.ai also offers dynamic adaptability. While traditional time series models require manual recalibration or model-switching when demand patterns shift, UFM.ai's unified statistical model automatically senses and adjusts model parameters. At the same time, the ML component continuously learns and adapts to new data, adjusting predictions in real time as market conditions evolve. This is particularly useful in industries with frequent disruptions, such as retail or manufacturing, where demand can change suddenly due to price adjustments, promotions, economic shifts or external events. UFM.ai was engineered to thrive in these circumstances, with more responsive and accurate forecasts that evolve as new information becomes available. UFM.ai eliminates the need for constant model-switching, simplifying the forecasting process and reducing the likelihood of errors. This makes UFM.ai particularly valuable for businesses managing large product portfolios or operating across multiple regions, where demand patterns can vary widely and change unpredictably.

UFM.ai also delivers advanced automation and scalability, making it suitable for large-scale operations. It automatically incorporates and processes vast amounts of data, from both internal and external sources, scaling seamlessly across complex datasets. Which means UFM.ai is an incredibly powerful tool for organizations dealing with large numbers of products, regions or markets, where forecasting accuracy is critical to maintaining optimal inventory levels and meeting customer demand.

By unifying the best aspects of both an advanced unified statistical model along with sophisticated ML approaches, UFM.ai delivers the most comprehensive, flexible and reliable forecasting solution available today, helping businesses stay ahead in today's fast-moving and unpredictable markets. Examples:

Long Life-Cycle Retail

For auto parts dealers, staying in business means having those hard-to-find items in stock for customers when they need them. However, the business problem in the automotive parts industry, as well as retail in general, is a large volume of the stock-keeping unit (SKU) population has extreme intermittent demand signals. The traditional approach is to throw inventory at the problem; however, more inventory is the wrong answer – revenue will suffer every time. Manhattan's UFM.ai enabled one particular retailer to avoid out-of-stocks while mitigating the risk of having more inventory than required. The resulting success in this specific instance included:

- **9% decrease** in inventory for nonintermittent SKUs
- **1% increase** in service attained for nonintermittent SKUs
- **21% increase** in revenue for intermittent demand SKUs
- **17% increase** in service for intermittent demand SKUs
- **14% increase** in inventory for intermittent demand SKUs

Wholesale Distribution

A very large distribution wholesale enterprise needed to further reduce its inventory investment while maintaining current levels of service. In this case, the wholesaler was already running an extremely fine-tuned version of inventory optimization and had left no stone unturned from a traditional forecasting point of view.

The addition of UFM.ai enabled the wholesaler to lower its overall inventory investment while maintaining service. The resulting success in this instance included:

- **4% overall average decrease** in inventory
- **7% safety stock reduction** across its network
- Service for intermittent SKUs increased slightly, while inventory actually decreased slightly

The Next Generation of Demand Forecasting Is Hybrid AI

An organization's ability to intelligently, holistically and optimally forecast a growing complexity and range of demand signals across the globe has become even more challenging. Hybrid AI demand forecasting, like UFM.ai, provides demand planners, inventory analysts and buyers with a significant blended advantage over traditional statistical or even ML forecasting tools.

Because UFM.ai can adapt to both inside-out and outside-in approaches to modeling and forecasting demand, retailers and wholesalers will gain more responsive and accurate demand forecasts, improved inventory accuracy and enhanced inventory optimization. The results are greater revenue, fewer lost sales and exceptional customer service.

Jeff Beadle is the head of Manhattan's Science Group, where he leads the development of cutting-edge mathematical modeling, optimization, algorithms, heuristics, data science, AI, and machine learning, which form the backbone of Manhattan's supply chain solutions. With deep expertise in data science, AI, machine learning, operations research, advanced analytics, and predictive modeling, Jeff plays a critical role in advancing the sophisticated technologies driving Manhattan's Supply Chain Planning platform.

A seasoned computational and algorithmic data scientist, Jeff brings more than 25 years of experience pioneering innovative solutions to complex optimization and analytical challenges in the supply chain industry. His leadership in R&D ensures that Manhattan's solutions remain at the forefront of supply chain science, continuously pushing the boundaries of what's possible in inventory optimization and forecasting.

Email **information@manh.com**
or call +1 (877) 596-9208

Or visit **manh.com/scp**

Manhattan®