International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 07, July 2026)

AI-Driven Demand Forecasting and Inventory Optimization in Consumer-Packaged Goods Supply Chains

Asha Kiran Ganesna

Abstract— **Consumer packaged goods (CPG) firms face a highly dynamic business landscape where: Stock-keeping unit (SKU) levels are high; Promotions are complicated; Product life cycles are short; consumer behavior is unpredictable. Many of the traditional statistical forecasting methods, which are largely based on sales data extrapolation from a single previous period, are finding it difficult to keep up with such complexity. This article discusses the evolution of classic time-series approaches to machine learning techniques such as random forests, gradient boosting, support vector regression, artificial neural networks, and deep learning systems, e.g., long short-term memory networks, gated recurrent units, and transformers, and their application to demand forecasting and inventory optimization in CPG supply chains. It also explores the role of AI in optimizing inventory with dynamic reorder points, demand sensing, multi-echelon inventory optimization, reinforcement learning policies for inventory replenishment, and digital twin simulation.**

Based on reported industry practice examples from Procter & Gamble, Unilever, and Walmart, the article shows examples of how these approaches have led to important operational results that can be measured, such as better forecast accuracy, higher fill rates, less safety stock, and less waste. It also discusses real-world challenges to adoption, such as data quality and governance problems, model interpretability, scalability, organizational change management, and ethical issues related to bias and privacy. The article wraps up with four trends: generative AI, large language models, autonomous supply chains, and sustainable inventory optimization. The results indicate that AI can significantly enhance the accuracy of predictions and reduce inventory costs, but its success will not just be determined by the complexity of the algorithms but also by the readiness of the data infrastructure, governance, and organizations.

Keywords-- **artificial intelligence, consumer packaged goods, demand forecasting, inventory optimization, supply chain management, machine learning, multi-echelon inventory.**

I. Introduction

A. Background

Consumer packaged goods (CPG), which include food and beverage, household, and personal care, are one of the biggest and most complex areas of the global economy.

CPG manufacturers typically have tens of thousands of stock-keeping units (SKUs) moving across dozens of markets via multi-echelon distribution channels that feature co-manufacturers, regional distribution centers, retailer warehouses, and, more recently, direct-to-consumer e-commerce channels [1]. Forecasting product cycle changes, such as in CPG, is incredibly challenging due to the many moving parts and the fact that it has a lot of similarities across different categories, so even small miscalculations can have a dramatic impact on bottom-line performance.

The last decade has witnessed increased volatility in the market environment of the CPG supply chain [2]. There has been an increasing number of channels and fulfillment models to cater to these consumers, and e-commerce promotions have become more frequent, and social media can create a viral cycle that changes the lifecycle of a niche product in days, rather than years. Not to mention, macroeconomic shifts in consumer spending can have a surprising impact on the flow of demand in a category that can't be predicted by historical trends alone [3]. The uncertainty and risks from the markets' dependence on assumptions that the future will be like the past have been heightened in recent years by a range of supply shocks, including pandemic-induced changes in demand, shipping bottlenecks, and raw-material shortages.

B. Problem Statement

All of these factors put pressure on CPG supply chain management. Overstocking uses up working capital, drives up warehousing expenses, and, in the case of perishables, results in loss of sales due to markdowns and waste. Understocking or stockouts lead to lost sales, lost retail/brand loyalty, and damaged relationships with retailers. This dichotomy is compounded by the bullwhip effect, where small variations in end-consumer demand get magnified as they go further up the chain, resulting in over- or under-production by the manufacturers.

This is a particular problem because of several structural characteristics of the CPG industry. Several historical demand patterns at and under the SKU level are too brief to be relied upon for meaningful classical forecasting.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 07, July 2026)

There is a large degree of promotional uncertainty, which causes sudden but short-lived peaks in demand, which can affect price, seasonality, and competitor behavior in a nonlinear manner. Seasonal variations are those that are related to holidays, weather, and cultural events, and occur differently across geography and type of seasonality than can be captured in a single national-level model. All of this implies that the task of CPG demand forecasting goes beyond simply fitting a trend line to historical sales; it's about weaving through numerous uncertainties to generate a reliable, ever-evolving forecast.

C. Why AI?

Traditional forecasting techniques continue to be popular: moving average and exponentially weighted moving average (EWMA) approaches, autoregressive integrated moving average (ARIMA) models, and classical linear regression models, which are low-cost, interpretable, and work well when the product is stable, has a high volume, and has a recognizable seasonal pattern [4], [5]. But these approaches have one major structural flaw: they do not include many of the factors that interact to shape CPG demand and are predominantly univariate, based almost exclusively on a product's past sales.

Moving averages and simple exponential smoothing are based on the assumption that the most recent history will be most predictive of the near term and can fail for promotions, new product introduction, or demand shock. ARIMA models work well for series that are stable and contain trend and seasonality, but they require extensive historical data and are difficult to apply to irregular series whose demand is influenced by promotions [6]. With an additional variable, classical regression methods are available, but they also require that the relationships between the predictors and the demand be linear, which is often not true in CPG markets, where the price, promotion, weather, and competitor activity interact in a non-linear way.

To overcome these challenges, AI and machine learning are capable of tackling demand forecasting as a multi-dimensional predictive modeling task [7], [8]. Machine learning models can learn from a variety of different features, potentially learn nonlinear relationships between the features, and — in the case of a global or cross-series model — borrow statistical strength across thousands of related SKUs to enhance out-of-sample forecasts for new or less-popular products. But this ability to merge various and high-frequency data sets into a unified forecast is ultimately what's made AI the most compelling advancement in demand planning in the CPG industry.

D. Objectives

This article aims to compare the key AI-driven and machine learning techniques to the traditional methods with which they complement; to explore the further applications of AI in inventory optimization, such as dynamic reorder points, demand sensing, multi-echelon optimization, reinforcement learning, and digital twins; to discuss the challenges of practical implementation, like data quality issues, interpretability, scalability, and challenges of ethics; and to look at trends that are stirring the next generation of CPG demand forecasting.

II. Consumer Packaged Goods Supply Chains

A. Characteristics

CPG supply chains have some unique structural features that set them apart from other industrial supply chains. Products move quickly, which means that inventory turns are high, and the cost of predicting wrongly becomes high with a large number of replenishments. With sizes, flavors, seasonal packaging, and private label agreements, the count of SKUs (e.g., product sizes) is large, creating what amounts to a new forecasting problem at each SKU location [9], [10]. Distribution networks are far from simple and usually involve multiple manufacturing sites, regional distribution centers, and thousands of retail delivery locations, each of which has its own lead times and replenishment policies. Retail complexity (such as several store models or e-commerce fulfillment processes) entails having to run numerous parallel forecasting and replenishment processes for each retail partner—particularly a single manufacturer.

B. Sources of Demand Uncertainty

Uncertainty about demand in CPG categories has several sources, many of which are interrelated. Baseline demand can be multiplied several times by promotions, whether they come in the form of price discounts, displays, or end cap placement, for the length of time the promotion runs, after which comes a post-promo decline as consumers deplete pantry-covered stocks. Some categories are directly impacted by weather, like beverages, ice cream, and seasonal snacks, and can see a change in regional demand within days of a temperature change. These holidays and cultural events cause geographically varying, recurring pressures in demand, which should be modeled locally, not nationally [11]. Customer buying behaviour trends that impact category-level demand do so slowly but steadily.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 07, July 2026)

Changes in market share that are not reflected in a brand's own sales history can be caused by competitor behaviour, such as the introduction of a new product or the brand's price cutting. Discretionary spending for premium versus value tiers is affected by macroeconomic factors such as inflation and consumer confidence.

C. Inventory Management Challenges

Making demand forecasts work in an effective inventory policy comes with a different set of challenges. In addition, safety stock, the inventory kept to account for stock-out risks caused by variations in demand and lead time, must be carefully tuned: lower safety stock levels put the business at risk of running out of stock, while higher safety stock levels consume working capital and waste, for perishable goods [13]. Replenishment processes need to decide, at each replenishment node, the amount and frequency of replenishment, taking into account the costs of ordering and holding inventories. The service level (the target probability of not experiencing a stockout) depends on the importance of the SKU and any contractual agreements with the retailer, and needs to be translated into quantifiable inventory targets [14]. The holding costs include warehousing, capital, and obsolescence risk, and these have a direct impact on profitability and need to be balanced against the cost of lost sales. Minimizing waste, especially for perishable products, is a direct reduction in product costs and their associated emissions, making it an economical and sustainability-driven concept.

III. AI-DRIVEN DEMAND FORECASTING

A. Traditional Forecasting Approaches

It is helpful to initially look at the classical approaches that AI methods are based on before diving into the specifics of the methods. Naive forecasting means taking the most recent demand observed as a starting point and extending it out into the future. Moving averages are calculated as the average value of the demand over a backwrapping window, thereby eliminating the short-term noise from the demand [15]. The calculation is simple, but it does not reflect real changes or trends in the demand. Exponential smoothing methods help by giving more weight to the more recent data and by explicitly representing the trend and seasonality of the data. The ARIMA and seasonal ARIMA models go further by including an explicit model for the autocorrelation structure of a time series, and quite often they are a good fit to a well-behaved time series with a long enough historical record.

All of these have the same drawback—they require a history of their own product and struggle to integrate external factors like promotions, weather, competitor action, etc., except for a lot of manual feature engineering. They are also facing intermittent demand, new product launches, and structural changes due to external events, which are becoming commonplace in today's CPG markets. Below (Table 1) is a comparison of classical statistical forecasting techniques and AI approaches on the dimensions that matter most to CPG demand planning. It emphasizes why AI replaces univariate, low frequency methods with high dimensional, real-time models that are better suited to volatile, promotion-heavy demand.

TABLE I
TRADITIONAL VS. AI FORECASTING TECHNIQUES

Dimension	Traditional Forecasting	AI-Driven Forecasting
Primary input	Own historical sales (univariate)	Hundreds of internal and external variables
Handles nonlinearity	Limited or none	Learns complex nonlinear interactions
New product / cold-start performance	Poor — no history to extrapolate from	Improved via cross-series learning and analog mapping
Granularity	Typically category or regional level	SKU-store-day level feasible
Update frequency	Periodic (weekly / monthly)	Near real-time / daily demand sensing
Interpretability	High	Variable — depends on model choice and XAI tooling
Computational cost	Low	Moderate to high

B. Expected Trends in the Application of AI and Machine Learning to Forecasting

In the application of machine learning (ML) methods for forecasting demand, there are usually three learning paradigms that are widely referred to. The prevailing method for production forecast systems is known as "supervised learning," which involves training a model using past examples in which the input is a set of features like price, promotion status, seasonality, etc., and the outcome of the production is known, so that the model can be trained to learn a mapping between features and future production [16].

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 07, July 2026)

Besides, unsupervised learning is used in a supportive role, such as grouping SKUs into clusters by similar demand patterns or detecting anomalies to identify unusual spikes in demand [17]. Ultimately, reinforcement learning is used less for forecasting itself and more for downstream inventory decisions, which is where the forecast is used, which is done by simulating trial and error or with a rule [18]. This diagram shows how raw data flows through integration, modeling, and output stages to produce a demand forecast

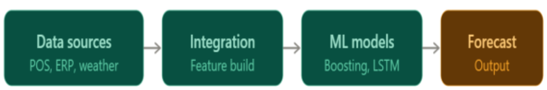

FIGURE I AI-BASED DEMAND FORECASTING ARCHITECTURE

C. Machine Learning Algorithms

Random forest models develop an ensemble of decision trees, where each of the trees is trained on a bootstrapped sample of the data and a random subset of features, and prediction is done by taking the average across the trees. Random forests are used in CPG forecasting for mid-frequency forecasting by SKU and feature-importance analysis, which enables planners to see which drivers have the greatest impact on the demand for a specific category. They offer good resistance to noisy data, support mixed data types, and have a low tuning burden. The biggest drawback is that, while they are a collection of mostly independent trees, they do not capture the sequential dependencies in a time series as well as architectures designed for them.

The gradient boosting algorithms like XGBoost, LightGBM, and CatBoost sequentially fit decision trees, and each tree is trained to minimize the residual errors of the last ensemble [19]. Notably, gradient boosting has emerged as a "workhorse" of production CPG forecasting because of its ability to achieve good accuracy at scale and for relatively small features, and to work well with mixed, tabular, feature-rich forecasting problems like those in demand forecasting [20]. It is commonly applied to large-scale SKU/location forecasting and also applied to modeling the promotional uplift, where discount depth, type of displays, and seasonality all play non-linear roles.

The support vector regression (SVR) function is an extension of support vector machines to continuous prediction, which aims to find a function that is within a certain margin of tolerance of the data. It has the advantage of having a decent performance in smaller datasets with high dimensions, being fairly robust to outliers, and being comparatively well-suited to solving problems with very large datasets, but it requires some care in selecting its kernel. Its popularity is limited compared to gradient boosting and deep learning in production.

Feed-forward artificial neural networks are composed of layers of interconnected nodes, which learn the non-linear transformation of the features of the input in an iterative fashion. They build complex models between demand and the drivers of demand, like the effect of price elasticity, whose sensitivity varies across different depths of promotions, in demand forecasting. One great feature is their flexibility as universal function approximators, although it demands larger training sets and less interpretability compared to tree-based approaches.

Numerous deep learning architectures have been modified to be applied to time-series forecasting. Recurrent architectures like long short-term memory networks and gated recurrent units are suitable for storing information for longer sequences, such as networks and multi-week promotional lead and lag effects. Temporal convolutional networks are an alternative approach that can be trained in parallel and that also captures long-range dependencies by using causal dilated convolutions on time-series data. Transformer models, which were developed in the field of NLP, have been particularly successful, and they can be trained as a global model for thousands of similar SKUs at once, leveraging statistical power from other similar products when forecasting new or less popular products due to their “attention” mechanism. Technically speaking, these architectures work well on large, high-frequency datasets not only because CPG demand series are so numerous and related and exogenous signals are abundant, but also because they actually do. The table below summarizes the main machine learning algorithms used in CPG demand forecasting and their typical applications, strengths and limitations. It is intended to assist planners and analysts in choosing the appropriate model for the particular forecasting requirement – whether that’s new-item forecasting, promotional uplift or long-horizon sequential demand.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 07, July 2026)

TABLE II
COMPARISON OF MACHINE LEARNING ALGORITHMS FOR DEMAND FORECASTING

Algorithm	Typical Applications	Key Strengths	Key Limitations
Random Forest	Mid-frequency SKU forecasting; feature importance analysis	Robust to noise; handles mixed data types; low tuning burden	Can underperform on strong sequential/temporal patterns
Gradient Boosting (XGBoost, LightGBM, CatBoost)	Production-scale SKU-location forecasting; promotional uplift modeling	High accuracy; captures nonlinear interactions; efficient at scale	Requires careful tuning; less naturally sequential than RNNs
Support Vector Regression	Smaller datasets; stable, low-volatility categories	Effective in high-dimensional spaces; robust to outliers	Scales poorly to very large datasets; sensitive to kernel choice
Artificial Neural Networks	Nonlinear demand-price-promotion relationships	Flexible function approximation; captures complex interactions	Requires more data; less interpretable
LSTM / GRU (Deep Learning)	Sequential demand with long promotional lead/lag effects	Captures long-range temporal dependencies	Computationally intensive; needs substantial training data
Temporal Convolutional Networks	High-frequency, granular time series	Parallelizable; captures long receptive fields efficiently	Less established in production tooling than RNNs/GBMs
Transformer Models	Global forecasting across thousands of related SKUs simultaneously	Learns shared patterns across series; strong new-item performance	

D. Data Sources for AI Forecasting

The accuracy of any AI forecasting tool is limited by the scope and accuracy of the information used in forecasting. The most direct demand signal is available based on point-of-sale data, with consumer purchases creating the signal (fewer distortions), whereas retailers order based on shipment data, which is subject to the bullwhip effect [21]. ERP systems offer visibility into internal shipment and inventory data, and customer relationship management data offers details about ordering context at the account level. For temperature- and season-sensitive categories, companies need weather data to support the categories; weather data is critical for temperature- and season-sensitive categories. Social media signals can give early ideas about the emerging trend, especially when it comes to viral items. Real-time, store-level inventory visibility is now gained from Internet of Things sensor data such as smart shelf and freezer telemetry. The promotion calendars and pricing information are crucial to extracting promotional uplift, and macroeconomic indicators assist in forecasting changes in category-level demand based on overall economic factors.

E. Forecast Accuracy Metrics

Because CPG products have intermittent demand and high SKU counts, metrics for evaluating the performance of a CPG forecast need to be different. The mean absolute error is the average of the absolute difference between forecast and actual demand, and it is easy to interpret and has a low weight given to large errors. Root mean square error (RMSE) will square the errors before averaging, giving greater weight to larger errors. The mean absolute percentage error represents error as a percent of the actual demand, which makes it easy to compare errors across SKUs, and is unstable for the intermittent items that have near-zero demand. To overcome this, weighted absolute percentage error (WAPE) is embraced, being the recommended measurement for portfolio reporting.

WAPE can be calculated by;

$$WAPE = \frac{\sum_{i=1}^{n} |A_i - F_i}{\sum_{i=1}^{n} |A_i} X100\% \qquad \text{……………….. (1)}$$

Where;

A_i is actual demand;
F_i is forecasted demand, and;
n is the number of periods or SKUs being evaluated.

WAPE weights forecast errors by demand volume, making it the preferred accuracy metric for portfolio-level reporting in CPG settings where MAPE becomes unstable for low-volume or intermittent SKUs.

In addition to these magnitude-related metrics, the forecast bias is a measure of whether forecasts are biased to overestimate or underestimate, which is crucial because a biased good model can still lead to chronic overstock or understock.

IV. AI-Based Inventory Optimization

A. *Inventory Optimization Concepts*

In this context, AI optimization of inventory is an extension of, not a replacement for, a series of basic premises. If the demand and delivery lead times remain fairly constant, the economic order quantity model helps to decide the quantity to order to minimize the ordering and inventory-holding costs. The target inventory level for a new order, based on anticipated demand over a reorder period plus a safety stock, is the reorder point. Safety stock is the buffer stock that is maintained above the expected demand because of variability and lies at the heart of the service level/holding cost trade-off. Safety stock can be calculated through;

$$SS = z \; X \; \sigma LT \; X \; \sqrt{L} \quad \ldots\ldots\ldots\ldots\ldots\ (2)$$

Where;

z is the service-level factor (from the standard normal distribution, corresponding to the target probability of avoiding a stockout);

σ_LT is the standard deviation of demand during lead time, and;

L is the average lead time in periods.

This is the classical safety stock formula that AI-driven systems extend by replacing the normal-distribution assumption with the actual shape of the probabilistic demand distribution generated by the forecasting model.

ABC analysis divides the inventory into segments according to its value or volume, then applies differing policies to each of the segments, tightening control of low-velocity items versus high-velocity items [22]. Service level optimization establishes explicit target levels for the probability of not having a stockout, which is a risk tolerance made concrete in inventory parameters.

B. *AI for Inventory Decisions*

Each of these rudimentary concepts is improved by the use of AI by substituting simple, periodically replaced parameters with dynamic ones that are constantly updated.

Dynamic reorder points are new calculations in near real time based on new demand and lead-time information, rather than being established at a fixed level and reviewed quarterly. The term 'demand sensing' is popular in the CPG industry and is essentially forecasting for a short time horizon, but combining data from the near real-time point-of-sale system and the channel inventory with the underlying statistical/machine learning forecast [23]. By doing this, you can capture demand shifts weeks before this information would be apparent in a traditional, monthly forecast. Multi-echelon inventory optimizes the distribution network as a whole, rather than optimizing each node, and can usually result in a lower overall inventory for a business with the same service levels. Predictive replenishment takes it one step further by leveraging AI-powered forecasts to generate replenishment orders and order sizes automatically, thus freeing up planner time for exception management of items with higher uncertainty. Diagram II traces how a demand forecast is translated step by step into an executed replenishment order

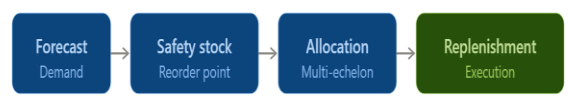

Figure II Inventory Optimization Workflow

C. *Reinforcement Learning*

The approach of reinforcement learning is very different from the analytical formulas cited above to make inventory decisions. An agent interacts with a simulated representation of the supply chain by making actions, e.g., submitting a new order for a certain item and duration, and receives a reward signal, which is based on performance relating to both service level and inventory cost. The agent learns a policy that maps the current state of the system to an action that should maximize the long-term cumulative reward in the system over many simulated episodes. These agents learn from experience, not by a preprogrammed rule, so they can, in principle, capture more complex, sequential compromises, like the interplay between order timing, the size of the orders, and the uncertainty of the demand over multiple periods, that closed-form safety stock formulas only approximate [24]. The approach is still not as fully developed as supervised forecasting when applied to production deployments, but it is an active field of application, especially for complex multi-echelon networks where the state space is too large to be handled with full analytical optimization.

D. Digital Twins

In the context of the supply chain, a digital twin is a simulation model of the supply chain that closely reflects its structure and dynamics, such that it can be used to test scenarios prior to trying a policy change in the real supply chain. Digital twins can also simulate what would happen if a disruption, such as a supplier outage or unexpected high demand, took place, and test different policies without real-world risk to the service level or costs [25]. This capability is especially meaningful in the face of recent years of supply chain perturbations, where companies are looking for supply chain resilience to tail risk events, and not only efficiency in the middle of the bell curve, but also utilization in sales and operations planning to consider tradeoffs before committing to them.

E. Benefits

AI-driven stock management, coupled with the right strategy, can result in cost-effective and service-oriented scenarios. The cost-side benefits usually include lower holding costs because the performance of accurately calibrated safety stock means that the total amount of inventory for a given service level decreases: fewer stockouts. After all, more frequent replenishment can be based on demand sensing (early indicators of demand shifts), and there is less waste (a smaller number of products marked down or off-sales as less product is discarded or discounted due to the closer match between short-horizon forecasting and proactive markdown timing). In addition, the organizations report higher customer satisfaction on the service side due to on-shelf availability, better fill rate, and better inventory turnover due to better movement of products and reduced capital in inventory. Table III links AI applications to the particular supply chain functions they support, from demand forecasting through replenishment and planner support. It shows how the same underlying AI toolkit—machine learning, simulation and generative models—gets applied differently depending on where in the planning cycle a decision is being made.

TABLE III
AI APPLICATIONS ACROSS SUPPLY CHAIN FUNCTIONS

Supply Chain Function	AI Application	Representative Techniques
Demand Forecasting	SKU-location-day level demand prediction	Gradient boosting, LSTM/GRU, transformer models
New Product Forecasting	Cold-start forecasting using analog products	Cross-series / global models, Bayesian analog mapping
Inventory Optimization	Dynamic safety stock and reorder point calculation	Probabilistic forecasting, multi-echelon inventory optimization
Replenishment	Automated, near-real-time order generation	Demand sensing, reinforcement learning-based policies
Markdown / Waste Management	Proactive markdown timing for perishable and slow-moving stock	Predictive analytics, expiration-risk scoring
Scenario Planning	Stress-testing inventory policy against disruptions	Digital twins, Monte Carlo simulation
Planner Support	Natural-language exception summaries and root-cause analysis	Generative AI / large language model copilots

V. INDUSTRY APPLICATIONS AND CASE STUDIES

The following case studies are based on publicly released companies' statements, trade press, and industry case studies. These figures are based on information provided by companies and published in the industry, and should not be viewed as exactly verified or universally applicable measures. Both, however, help provide a realistic sense of scale and impact of the forecasting and inventory optimization that AI brings about in practice.

A. Procter & Gamble

Procter & Gamble (P&G), one of the world's biggest CPG manufacturers, has been heavily investing in AI-powered supply chain planning as part of a larger project dubbed Supply Chain 3.0, which involves the convergence of cutting-edge analytics, digitalization, and organizational restructure to increase planning accuracy across its manufacturing and distribution footprint around the world [26], [27]. P&G's demand forecasting systems rely on point-of-sale, retailer, social, and macroeconomic information to create forecasts down to the product and store level, and on proxy or analog information to forecast demand for new product launches, which do not have historical sales data. The firm has also developed a real-time scenario modeling capability, also known as a control tower, that marries global logistics, enterprise resource planning (ERP), and external data (including weather and information from suppliers) into one view, thereby significantly lowering ineffective empty-truck logistics trips. The company has been using predictive analytics to make incremental year-over-year forecast accuracy gains despite macroeconomic uncertainty, and country-level deployments are correlated with double-digit point improvements in out-of-stock rates in certain geographies, according to its earnings commentary.

B. Unilever

With many brand names across food, home care, and personal care categories, Unilever has applied AI to forecast customer demand and anticipation in many aspects of its business. For ice cream, a product with a highly volatile pattern that is extremely sensitive to weather in a category that serves dozens of countries around the world, Unilever's AI models reportedly raise volume forecast accuracy by about 10 percent, allowing them to achieve more resilience against volume fluctuations caused by small changes in temperature [28].

Another AI-based collaborative planning, forecasting, and replenishment pilot involved Unilever with a major retail partner in Mexico, which incorporates point-of-sale, inventory, and promotional data onto a shared system for a significant portion of Unilever's Mexico operations [29], [30]. Customer on-shelf availability was achieved to around 98 percent, sales improved by double digits, and stock level reduced by 98 percent within the first year, which was published and extended to other markets. Unilever's smart-freezer program also has an important percentage of its display freezers fitted with image-capture artificial intelligence, giving immediate inventory transparency and the ability to automatically reorder; sales results are showing an increase for those markets that have implemented it. Unilever had a similar success story with demand sensing in another North American turnaround when they saw significant improvement in the accuracy of their short-horizon forecasts, resulting in the removal of several days' worth of safety stock as well as reduced freight expense [31].

C. Walmart

Walmart has been a retailer more than a manufacturer, but its size and the degree to which it has mastered the art of AI make it a good reference point from the downstream, retail end of the CPG value chain, where the manufacturer's forecasts are only as good as the downstream information they can get about sell-through and shelf availability [32]. Walmart utilizes machine learning models capable of forecasting demand on the store-SKU-day level, with the models continuously being retrained as they receive new data and use historical sales, the weather, local events, and online search trends. As the company disclosed in the technical texts, a distinguishing aspect is the capability to selectively "discount" anomalies in demand in one-off periods, for example, if the weather in a particular region is unusual, which does not affect future planning.

Walmart's predictions are directly plugged into automated replenishment systems that are used to inform and, in the case of high-confidence items, automatically place reorder orders without much manual oversight, which the company calls hands-free forecasting/replenishment and reorders by exception [33]. Walmart has introduced AI-powered systems to its inventory management that, according to published reports, have helped the company reduce stockouts and increase turnaround of stock, as well as other uses like optimizing delivery routes that have cut tens of millions of unnecessary driving miles a year. Exactly like the other cases mentioned here, the exact numbers are dependent on the source and should not necessarily be considered a true audit; they should only be taken as an indication of direction and approximate size.

D. Cross-Case Observations

Some patterns are duplicated in these cases. Gains are usually greatest when more accurate forecasts are coupled with automated, tightly coupled replenishment execution and a faster response to the operation, and not just when they are more accurate by themselves. The most obvious and measurable AI-driven accuracy increases are seen in the weather categories and the event categories, possibly because these categories have key influencers in their demand drivers that can be easily quantified as model features. The gains in on-shelf availability seem to go hand-in-hand with deeper data integration from retail partners, with direct access to point-of-sale data driving the biggest gains over shipment-based data, confirming that forecasting performance is limited not only by modeling sophistication but also by data access. This diagram shows the continuous cycle of sensing, predicting, deciding, acting, and learning that underlies an AI-driven supply chain in the various organizations.

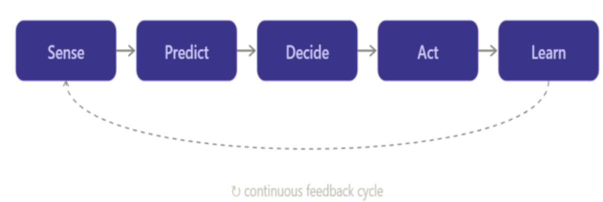

FIGURE III AI-DRIVEN CLOSED-LOOP SUPPLY CHAIN DECISION PROCESS

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 07, July 2026)

VI. Challenges And Limitations

Data Quality. Machine learning models are very sensitive to the data upon which they are trained, and many CPGs struggle to deal with inconsistent product master data, missing product data, siloed systems, and inconsistent data feeds from retailers, many of which vary in format and quality [34]. Often, data governance is not a topic that gets the same level of attention as creating the model itself, meaning that the concepts of ownership, quality standards, and the refresh cadence are not clearly defined for each source.

Model Interpretability. Including complex machine learning models—and deep learning architectures especially—may operate as a black box whose logic is hard for planners to comprehend. If a forecast is significantly out of line with a planner's own intuition, there may be a lack of explainability that would undermine trust in the forecast and result in a forecast being ignored, whether or not it is statistically accurate, undermining the value of the AI investment in the process. This has led to increased interest in EAI methods, which reveal the most important factors involved in a specific prediction that planners can assess.

Scalability. Training and continuously retraining forecasting models on tens of thousands of SKUs can be quite compute-intensive, especially those based on deep learning, which would generally need cloud-based infrastructure with elastic scaling and robust data pipeline engineering. Organizations also need to consider the complexity of the models versus the need for re-training regularly to be useful.

Organizational Issues. Organizational readiness is a crucial factor in successful AI adoption. Here lies the disconnect and the possibility of employee resistance when planners feel AI is replacing their jobs, especially if they aren't informed about how AI will actually replace the exception handling aspect of what they do, instead of eliminating job tasks. One of the most common skills gaps is that data science skill sets differ from traditional demand planning skill sets, which can be resolved through skills training and/or the integration of data science into demand planning skill sets. All costs are a reality, particularly with mid-sized companies, but cloud platforms are beginning to reduce these costs. Seamless integration with a legacy planning system, typically not built for the application of machine learning, takes significant effort, with cybersecurity risk increasing as partner data-sharing becomes more extensive.

Ethical Issues. Ethical concerns that need to be proactively addressed include the use of AI in forecasting and stocking. Proactive measures involve using AI in forecasting and stocking, which poses ethical concerns. Bias can creep into the models when the historical information in the model reflects past inequities in access to or distribution of markets, such as an under-forecast of demand in the period when it was largely unserved, thereby continuing that underservice. Forecasting systems rely on detailed data from the individual consumer level, leading to privacy concerns that require careful attention to principles of data protection and minimization. In a broader sense, being a responsible AI practitioner, such as regularly auditing models and ensuring there is transparency around allocation decisions, is now being looked at as a governance requirement instead of a best practice.

VII. Future Directions

Generative AI. In addition to predictive machine learning tools already integrated within the demand planning systems, generative AI is now starting to change the demand planner-facing component of these systems. Structured forecast exceptions are translated into natural language text to expedite the root cause analysis; demand data and unstructured information like news or conversations with retailers can be brought together in a structured way by using a generative model to create a narrative, avoiding the need to create a formal simulation from scratch to support scenario planning. Some research also focuses on generative models to be used to create plausible scenarios for new products, which have no sales history, i.e., synthetic training data, based on the known history of similar products, to address the aforementioned Cold Start problem.

Large Language Models. Large language models are closely related to generative AI and are being integrated as conversational co-pilots into the planner workflow: A planner can simply type a question into the planner's AI and get a straightforward response from the planner's AI about how much impact a hypothetical promotion or disruption will have on inventory, even when the planner doesn't have any particular expertise in the question. They're also included to provide a summary of forecast review meetings, to capture action items, to correlate qualitative market information (e.g., field notes for a retailer relationship) with quantifiable forecast results, and to help align structured model output with contextual input knowledge from field planners.

Autonomous Supply Chains. One longer-term development is the growing trend towards increasingly autonomous, closed-loop supply chain decision-making, where AI systems forecast and suggest inventory policies and make frequent decisions regarding routine inventory requirements and logistics decisions with minimal human intervention, saving planners for true exceptions.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 07, July 2026)

The hands-free forecasting and replenishment models that have been introduced by some CPG manufacturers and retailers are a sign of this trend and will likely be followed by automated supplier negotiation and self-correcting inventory rebalancing.

Sustainable Inventory Optimization. The effects on carbon footprint and waste are being taken into account alongside the costs and service levels, and the regulations and customer expectations on sustainability are growing stronger. This will likely be reflected in inventory optimization objective functions that include emissions in addition to the cost terms, and in the forecasting system that gives greater weight to minimizing waste, e.g., waste of perishables has a direct environmental as well as economic cost.

VIII. Conclusion

AI and machine learning are transforming how consumer packaged goods supply chains make predictions and optimize inventory beyond mere incremental improvements over traditional techniques based on classical statistics. Machine learning models—including gradient boosting ensembles and transformer-based global forecasters—can leverage far more dimensionality of internal and external data signals, including point-of-sale transactions, promotions, and weather and macroeconomic indicators, that are nonlinearly coupled and interacting with each other—signals that are structurally unable to be captured by univariate time-series methods. Complementing these forecasts are AI inventions like dynamic safety stock, multi-echelon network planning, reinforcement learning algorithms for replenishment policies, and digital twin scenario testing, which are proven to deliver tangible benefits in terms of fill rate, on-shelf availability, inventory turnover, and waste reduction—as revealed by the published experiences of Procter & Gamble, Unilever, and Walmart.

For managers, the central implication is that the technology alone does not determine the outcome. The organizations that are getting the most benefits from AI-enabled forecasting have also experienced the most success in implementing deliberate change management, as well as leveraging AI to complement, rather than replace, human planning expertise; and embedding the forecasting process firmly into the larger sales and operations planning program, and not as a standalone technical effort. For the accuracy and efficiency that the algorithms are able to deliver, the technology must be accompanied by preconditions: interpretability, data governance, and trust building within the organization.

There are a number of open research gaps. Beyond vendor- and company-reported case studies, more rigorous, independently verified empirical evidence is needed to determine the actual degree to which AI is impacting the accuracy of forecasts and inventory efficiency within each category and across company sizes. Cross-series and analog approaches to new product forecasting have made a significant improvement in the cold start problem, but the issue is still far from solved. There is also a need to continue research on the ethical aspects of algorithmic allocation, such as the shift of algorithmic bias from historical data to automated decisions impacting market access at different levels in various communities and regions.

Overall, AI-powered predictive demand planning and inventory management are among the more impactful operational bets that can be made today by CPG businesses. As generative AI, autonomous decision-making, and sustainability-linked optimization grow, the difference between companies that have developed a genuine ability to plan with AI and companies still working with outdated statistical models will probably continue to get bigger, and this is something that both the technology and ability to use it well must be maintained if a company is to have true AI power for planning.

REFERENCES

[1] e2open, "Reducing Supply Chain Costs in CPG | Start with Inventory Optimization," E2open.com, 2025. https://www.e2open.com/blog/reducing-supply-chain-costs-cpg-inventory (accessed July 16, 2026).

[2] K. Williamson, "Navigating volatility: How CPG brands can shore up their supply chain in uncertain times," Gen5 SCMR, Aug. 29, 2025. https://www.scmr.com/article/navigating-volatility-how-cpg-brands-can-shore-up-their-supply-chain-in-uncertain-times (accessed July 16, 2026).

[3] T. P. Scholdra, J. R. K. Wichmann, M. Eisenbeiss, and W. J. Reinartz, "Households Under Economic Change: How Micro- and Macroeconomic Conditions Shape Grocery Shopping Behavior," Journal of Marketing, vol. 86, no. 4, p. 002224292110368, Jan. 2022, doi: 10.1177/00222429211036882

[4] T. Corvin, "EWMA Simplified (2025): Definition, Dynamics, Insights," The Trading Analyst, May 31, 2024. https://thetradinganalyst.com/exponentially-weighted-moving-average/ (accessed July 16, 2026).

[5] A. Abbas and Seyed Goosheh, "3.7 Time Series Models," Fundamentals of Operations Management, Aug. 13, 2024. https://ecampusontario.pressbooks.pub/fundamentalsopsmgmt/chapter/3-7-time-series-models/ (accessed July 16, 2026).

[6] B. A. Jagero, M. Thomas, and M. Samuel, "Modelling and Forecasting Inflation Rates in Kenya Using ARIMA Model," European Journal of Statistics and Probability, vol. 11, no. 1, pp. 54–68, Jan. 2023, doi: 10.37745/ejsp.2013/vol11n15468.

[7] S. A. Syed, "Artificial Intelligence–Driven Demand Forecasting Models for Enhancing Supply Chain Planning Accuracy in Saudi Industrial Sectors," International Journal of Artificial Intelligence Engineering and Transformation, vol. 7, no. 1, pp. 10–12, 2026, doi: 10.54660/ijaiet.2026.7.1.10-12

[8] A. Seyam, S. S. Mathew, M. El Barachi, C. Zhang, and J. Shen, "The application of machine learning and deep learning on demand forecasting across time-critical industries: A systematic review," Advanced Engineering Informatics, vol. 74, p. 104625, Sept. 2026, doi: 10.1016/j.aei.2026.104625

[9] WNS, "Demand Forecasting 2.0: How Predictive Models Are Optimizing CPG Supply Chains," Wns.com, June 16, 2025. https://www.wns.com/perspectives/blogs/demand-forecasting-2-0-how-predictive-models-are-optimizing-cpg-supply-chains (accessed July 16, 2026).

[10] I. Khymych, "Inventory Management in CPG: Evolving Best Practices Distributors Should Watch in 2025," SimplyDepo, July 08, 2025. https://simplydepo.com/industry/inventory-management-in-cpg/ (accessed July 16, 2026).

[11] A. Razma and R. Stašys, "THE IMPACT OF CULTURAL EVENTS ON REDUCING THE SEASONALITY OF TOURISM," Tiltai, pp. 227–241, June 2025, doi: 10.15181/tbb.v94i1.2723.

[12] A. Jonker and A. Gomstyn, "Safety stock," Ibm.com, Aug. 21, 2024. https://www.ibm.com/think/topics/safety-stock (accessed July 16, 2026).

[13] I. Al-Khatib, S. Alasheh, and A. Shamayleh, "The Drivers of Complexity in Inventory Management Within the Healthcare Industry," International journal of service science, management, engineering and technology, vol. 15, no. 1, pp. 1–26, May 2024, doi: 10.4018/ijssmet.347332.

[14] A. Jenkins, "How to Calculate Safety Stock," Oracle NetSuite, July 21, 2022. https://www.netsuite.com/portal/resource/articles/inventory-management/safety-stock.shtml (accessed July 16, 2026).

[15] J. Fernando, "Moving Average (MA): Purpose, Uses, Formula, and Examples," Investopedia, Aug. 06, 2024. https://www.investopedia.com/terms/m/movingaverage.asp (accessed July 16, 2026).

[16] G. Lu et al., "A Survey of Deep Learning for Time Series Forecasting: Theories, Datasets, and State-of-the-Art Techniques," Computers, Materials & Continua, vol. 85, no. 2, pp. 2403–2441, 2025, doi: 10.32604/cmc.2025.068024.

[17] Kumar P, Vijayaraj A, Senthil Pandi. S, and V. P. Balakrishnan, "Sales and Inventory Demand Forecasting Using Machine Learning," pp. 173–178, Nov. 2024, doi: 10.1109/icetems64039.2024.10964960.

[18] J. Mayerhoff and M. Schmidt, "Reinforcement learning for autonomous production planning and control: A systematic literature review," Journal of Manufacturing Systems, vol. 86, pp. 546–568, Apr. 2026, doi: 10.1016/j.jmsy.2026.03.023.

[19] M. Age Coders, "Ensemble Methods - Boosting (XGBoost, LightGBM, CatBoost)," Modern Age Coders, Apr. 10, 2026. https://learn.modernagecoders.com/resources/ai-and-machine-learning/ensemble-methods-and-boosting (accessed July 16, 2026).

[20] P. Supakar, M. Sarkar, and B. Sarkar, "Forecasting grocery item sales using gradient boosting models: A study of GridSearchCV, RandomizedSearchCV, and optuna optimization approaches," Ain Shams Engineering Journal, vol. 17, no. 3, p. 104012, Mar. 2026, doi: 10.1016/j.asej.2026.104012.

[21] H. Song and D. Zhang, "Bullwhip Effect in Supply Chains and Cost Rigidity," Journal of risk and financial management, vol. 18, no. 5, pp. 284–284, May 2025, doi: 10.3390/jrfm18050284

[22] A. Jenkins , "ABC Analysis in Inventory Management: Benefits & Best Practices," Oracle NetSuite, Sept. 06, 2023. https://www.netsuite.com/portal/resource/articles/inventory-management/abc-inventory-analysis.shtml (accessed July 16, 2026).

[23] D. Folinas and S. Rabi, "Estimating benefits of Demand Sensing for consumer goods organisations," Journal of Database Marketing & Customer Strategy Management, vol. 19, no. 4, pp. 245–261, Oct. 2012, doi: 10.1057/dbm.2012.22.

[24] N. M.-J. Park, N. O. Turner, and N. T. Becker, "Inventory Optimization in Retail Supply Chains Using Deep Reinforcement Learning," vol. 1, no. 3, June 2025, doi: 10.62177/amit.v1i3.470.

[25] C. Liu, Y. Wang, L. Purvis, and A. Potter, "Can Digital Twin Technology Enhance Supply-Chain Resilience? A Systematic Literature Review," Sustainability, vol. 18, no. 5, p. 2361, Feb. 2026, doi: 10.3390/su18052361.

[26] R. Jaiswal, "Is PG's Supply Chain 3.0 Redefining Efficiency Through Automation?," Yahoo Finance, Jan. 27, 2026. https://finance.yahoo.com/news/pgs-supply-chain-3-0-142900715.html (accessed July 16, 2026).

[27] A. Gonsalves, "P&G shifts Supply Chain 3.0, other platforms into large-scale rollout," Supply Chain Dive, May 08, 2026. https://www.supplychaindive.com/news/pg-shifts-supply-chain-30-other-platforms-into-large-scale-rollout/819256/ (accessed July 16, 2026).

[28] Unilever, "How AI is transforming Unilever Ice Cream's end-to-end supply chain," Unilever, Jan. 22, 2025. https://www.unilever.com/news/news-search/2025/how-ai-is-transforming-unilever-ice-creams-end-to-end-supply-chain/ (accessed July 16, 2026).

[29] D. Kumar, "How AI Became Unilever's Catalyst for Food Supply Innovation? - Food Innovation Consulting Services," Food Innovation Consulting Services, Apr. 08, 2026. https://food.greyb.com/blog/unilever-ai-food-supply-chain/ (accessed July 16, 2026)

[30] Unilever, "How AI and e-commerce tools are transforming emerging market retail," Unilever, Apr. 16, 2025. https://www.unilever.com/news/news-search/2025/how-ai-and-ecommerce-tools-are-transforming-emerging-market-retail/ (accessed July 16, 2026).

[31] "AI-Driven Demand Sensing: Lessons from Unilever and Amazon for the Supply Chain," AI in the Chain, Oct. 13, 2025. https://aiinthechain.com/2025/10/13/ai-driven-demand-sensing-lessons-from-unilever-and-amazon-for-the-supply-chain/ (accessed July 16, 2026).

[32] R. Molthe, "Walmart: High-tech retailer," NHH, 2025. https://www.nhh.no/en/research-centres/food/food-news/2025/walmart-high-tech-retailer/ (accessed July 16, 2026).

[33] H. Das, "Walmart's AI Inventory Optimization: How Data Keeps Shelves Full - Xaltius Academy," Xaltius Academy, Feb. 24, 2026. https://xaltiusacademy.com/walmarts-ai-inventory-optimization-how-data-keeps-shelves-full/ (accessed July 16, 2026).

[34] Granica, "Navigating Challenges of Data Quality in Machine Learning | Granica Blog," Granica, Jan. 23, 2025. https://www.granica.ai/blog/data-quality-in-machine-learning-grc (accessed July 16, 2026).