Check for updates

Multi-Domain Inventory Supply Optimization: A Comparative Analysis Using Seasonal-Trend Decomposition (STL) and Predictive Analytics

Karan Dalania*[1] Shashank Saxena[2] Sowmiya B[3]

Department of Computing Technologies
SRM Institute of Science and Technology
Kattankulathur, India
[1]kd5174@srmist.edu.in [2]ss0285@srmist.edu.in [3]sowmiyab@srmist.edu.in

Abstract. Demand forecasting plays a critical role in the operations of any retail because it enables the control of the inventory and facilitates a supply chain and maximises revenues. A close examination of time series forecasting methods adopted in three retailing industries i.e. fashion e-commerce, Amazon e-commerce and grocery retail have been incorporated in the paper. A model proposal, which is a hybrid of Seasonal-Trend decomposition (STL) and LOESS and Holt-Winters exponential smoothing would be an effective method of determining complex temporal fluctuations in the demand data. Real life data of three large retailers which possess different demand characteristics was used to test the methodology. We have found that the mean absolute percent error (MAPE) of our strategy is 12.75% in fashion retail, 10.40% in Amazon e-commerce, and 13.2% in grocery retail. This model particularly finds application in products in high volume and also has the ability to have automated demand forecasting systems, which can be operated under numerous retail environments. The article contributes to the body of existing literature on the topic of retail analytics by demonstrating how traditional time series tools can be applied and adapted to fit the current e-commerce and omnichannel retail processes.

Keywords: Demand forecasting, time series modeling, STL decomposition, Holt-Winters exponential smoothing, retail operations, inventory management, seasonal demand, fashion e-commerce, Amazon e-commerce, grocery retail, MAPE metric, supply chain optimization.

1 INTRODUCTION

Demand forecasting is one of the most difficult and strategically significant issues of retail operations. Wrong predictions cause either surplus stocks which consume capital and expose the company to the risk of markdowns and obsolescence or inadequate inventory which causes the company to lose sales and customer satisfaction. The issue is especially severe in the case of fashion retail, where the lifecycle of products is short, and seasonality is strong, but it is also not less acute in stable segments like consumer goods and groceries, where frequent replenishment decisions are required with minimal latency.

Traditional demand forecasting in a retail context has always been based on simple moving averages, exponential smoothing or rule-of-thumb techniques as guided by domain knowledge. Although they are still operationally relevant, they can be very inadequate at capturing the stratified nature of contemporary retail demand: a combination of long-term trends, seasonal repetition, promotional impacts, and noise. With the spread of e-commerce platforms such as Amazon and the ubiquity of data collection, there has been the opportunity to use more advanced

R. Vasanth Kumar Mehta et al. (eds.), *Proceedings of the International Conference on Intelligent Systems for a Sustainable Future (ISSF 2026)*, Atlantis Highlights in Intelligent Systems 16,
https://doi.org/10.2991/978-94-6239-693-7_116

strategies that explicitly break temporal elements down and use past trends to produce more accurate predictions.

In this paper, this opportunity is dealt with by proposing and justifying a unified framework, which integrates STL (Seasonal-Trend decomposition using LOESS) and Holt-Winters exponential smoothing. STL method enables separation of trend, seasonal and residual elements of the demand whereas Holt-Winters offers a principled method of level, trend and seasonality capture in the forecasting equation. The mix is theoretically inspired and practically justified: the ability of STL to deal with non-linear tendencies and irregular seasonal patterns complements the strength and interpretability of Holt-Winters.

The main contributions of this work are:

- Shows that STL-Holt-Winters attains significant progress in MAPE relative to baseline frameworks (10.40% to 13.2% vs. 24.5% or so), which is validated across three separate retail environments.
- Offers comprehensive configuration instructions on how the framework can be implemented to other product lines and demand characteristics.
- Tests the strategy using actual working data of three large retailers (fashion e-commerce, Amazon e-commerce, and grocery retail), and demonstrates the practicality of its deployability.
- Addresses implications to inventory management, procurement and supply chain decision making, such as approaches to integrating global models with product-specific tuning.

2 LITERATURE SURVEY

Operations research, supply chain management, and statistics have all examined demand forecasting. Classical literature The classical literature is based on exponential smoothing techniques proposed by Brown in the 1950s-60s, which formed a family of algorithms (simple, double, and triple exponential smoothing) that are easy to implement, computationally fast and efficient, and have strong out-of-sample performance. Holt (1957) and Winters (1960) generalized exponential smoothing to include trend as well as seasonality, developing HoltWinters exponential smoothing, which has become a standard in industries of retail forecasting.

Significant developments in methodology:

- **ARIMA Framework (Box & Jenkins, 1970):** It introduced the explicit modeling of autoregressive structure, differencing structure and the moving average structure. ARIMA techniques were useful in univariate demand forecasting, especially when underlying demand can be characterized as being stationary or differenced stationary process. ARIMA models are however difficult to specify, estimate, and diagnose particularly in the presence of seasonality.

- **STL Decomposition (Cleveland et al., 1990):** Introduced non-parametric, flexible decomposition of trends and seasonal components, based on locally weighted scatterplot smoothing (LOESS). STL offers robustness to outliers and flexibility to non-linear trends, which are a weakness of parametric decomposition solutions.

- **Hierarchical and Multi-Level Forecasting (Kourentzeset al., 2014)**: The paper examined the coordination of forecasts at various levels of aggregation. The important point is that product-specific models can be effective when global models are pooled together with all SKUs because it takes advantage of idiosyncratic patterns.

STL smoothing has been studied in combination with exponential smoothing in a number of settings. Dagum and Biczak (2012) discussed X-13ARIMA-SEATS, which is a system that takes decomposition as one of its inputs to seasonal adjustment. More precisely, research on retail demand forecasting has examined how STL preprocessing may enhance the performance of classical smoothing techniques in products with dissimilar seasonal characteristics.

The problem of demand forecasting has gained more and more popularity in e-commerce and online retail because of some peculiarities: demand is ordered within a short time interval, a price change may be responded to in a short time, promotions and algorithms influence customer discovery and buying behavior. Big online stores such as Amazon have a huge size of millions of SKUs and need to make a prognosis at both global and item levels. The academic literature on e-commerce requires forecasting (Thomassey, 2010; Mostard and Teunter, 2006) addresses the peculiarities of retail nowadays: the short lives of products, non-linear demand, and the fast change of tastes and preferences of customers with the help of online reviews, social indicators, and recommendation algorithms.

The current research extends the above by suggesting systematic implementation of STL-Holt–Winters in various retail sectors (fashion, mainstream e-commerce, and grocery) and showing that the prudent use of such classical tools can provide performance as competitive as the current-day alternatives without sacrificing interpretability and operational transparency, which are appreciated by practitioners as debugging and adapting the systems to changing environments.

3 METHODOLOGY

3.1. Data Preprocessing and Temporal Aggregation

The basis of the forecasting model is the meticulous preparation of transactional data into routine time series. Raw retail data normally comes in form of event-based records (single transactions or orders) and have timestamps, product identifiers, and quantities. The former step consolidates them into a homogenous time series.

Domain aggregation decisions:

- **Fashion e-commerce (USPA Boys Shirt):** E-commerce has a high monthly and seasonal demand that reflects the calendar effects and an update of the wardrobe. Aggregation to weekly adds noise caused by promotional timing, and aggregation to quarterly removes valuable patterns at a monthly level. The resulting monthly time series has 36 observations (3 years) no missing during active selling period.

- **Amazon e-commerce (Cetaphil Cleanser):** Monthly aggregation to align with regular retail planning cycles and seasonal changes (e.g. more skincare product demand in dry seasons or seasonal changes in consumer health and beauty purchases). The dataset contains 30 months, and there are few gaps in it.

- **BigBasket (Perishable Staples):** These aggregations of weekly data are to be used to account for the stronger within-week seasonality (weekend peaks) with the objective of ironing out the high-frequency daily fluctuations (stock-outs or supply discontinuities). The data includes 104 history weeks, which is enough to obtain strong patterns.

3.2. STL Decomposition

STL separates time series into three independent elements:

$$Y_t = T_t + S_t + R_t$$

Where Y_t is the demand at time t, T_t is the trend (long-term direction), S_t is the seasonal (regular cyclical pattern) and R_t is the non-trend (noise and irregular change).

STL is performed through iterative LOESS regression:

- **Trend extraction step:** LOESS is fitted to the series (or to deseasonalized series in later cycles) with a local polynomial of degree 1 and a window with a large proportion of the series length, so that only the global envelope is estimated.
- **Seasonal extraction step:** The detrended sequence is divided into seasonal periods (e.g. 12 months per year of seasonality), and LOESS is estimated over each period to isolate the seasonal pattern.
- **Remainder calculation:** Calculated as a subtraction of trend and seasonal of original series.

Selection of season period by domain:

- **Fashion domain:** 12 months of seasonality (spring high, holiday high, summer low) a year.
- **Amazon e-commerce domain:** 12 months period that includes the seasonal impact in consumer health and beauty products that are motivated by the calendar effects and seasonal purchasing patterns.
- **Grocery domain:** 12 months based on the observation of weekly or daily trends.

The resistance of STL to outliers was especially useful. When data is available over years, price spikes, promotions or supply shock occur, resulting in short-term demand spikes. The iterative robust regression of STL (through biweighting of outliers) is such that these anomalies do not artificially inflate the estimated trend and seasonality.

3.3. Model Selection and Tuning

In each retail area and product line, configurations have been checked.

Tests of model configurations:

- **Baseline monthly model (no STL):** Single global pool of SKUs with no decomposition. Offers basic comparison benchmarking and baseline.
- **STL-Holt–Winters with weekly aggregation:** Shorter-term seasonality. Appropriate in instances where there is a lot of data and product-level trends are required.
- **Targeted high-accuracy monthly model:** Use of best- selling SKU or most consistent SKU. Monthly aggregation with custom period of the year.

The reasons behind all these configurations were the nature of the data (high frequency and demand stability), and the business interests (global coverage versus targeted accuracy on key SKUs). In each model, the last 3 months (or 4 weeks in weekly case) were held back as the test set and the previous data were used to train the model. This maintained time flow and prevented the look-ahead bias.

3.4. Evaluation Metrics

Mean Absolute Percentage Error (MAPE): MAPE is a unit-free value that can be interpreted in terms of mean percentage deviation. It is the standard measure of retail forecasting since it is self-scaling in proportion to the scale of demands, and also proportional to the effects of forecast error on the planning decisions.

Mean Absolute Error (MAE): This is used to measure the absolute discrepancies in the units of demand. Helpful in learning about error in operational terms.

4 RESULTS AND DETAILED ANALYSIS BY IMPLEMENTATION

This part provides findings of the STL-HoltWinters model in three different retailing environments. All the subsections are structured identically: product and data characteristics, STL decomposition findings, forecast accuracy metrics and qualitative performance with business implications.

Table 1. Forecasting Performance Across Multiple Domains

Domain	Representative Product	Forecast Horizon	MAPE (%)
Fashion	USPA Boys Shirt	13 Weeks	12.75
Amazon	Cetaphil Cleanser	13 Weeks	10.40
BigBasket	Perishable Staples	4 Weeks	13.20

4.1. Forecast Accuracy Summary Across Domains

The summary of mean absolute percentage error obtained by the STL-Holt-Winters framework in the three domains is given in a consolidated form in Table 1. Findings show significant improvement in performance compared to the default (refer to the global monthly model without decomposition) and the MAPE decreases by 45-48 percent in the three domains. The most accurate (lowest MAPE) is found in the Amazon e-commerce (10.40%), which is more predictable and constant in high quantities of consumer products in the site. The fashion and grocery domains attain slightly higher yet still great MAPE values (12.75% and 13.2%, respectively), which is in line with the theoretical considerations since the domains tend to be more volatile. The performance differences in different domains indicate the need to tailor the forecasting technique to the specific demand attributes of each type, as opposed to selecting a single strategy.

4.2. Implementation 1: Fashion E-Commerce

Data Characteristics and Domain Context : The dynamics of fashion retail make fashion e-commerce special in terms of forecasting. The products in fashion are normally seasonal and the demand is influenced by the weather, holidays, festivals and also changing preferences of the consumers. Lifestyles of products are very short and when a product is stopped, the demand goes to zero. USPA Boys Shirt is a comparatively consistent product in this category: it is a centralized merchandise that has been kept in stock over several years, and brand loyalty is steady, and seasonal fluctuation is predictable.

Dataset Characteristics:

- **Time span:** 36 months (January 2023 to December 2025)
- **Total historical sales:** 16.0 million units
- **Average monthly demand:** 444,000 units
- **Active sell rate:** 158/159 weeks (99.4% fill rate)
- **Trend:** Stable upward trend (30% growth during period of analysis)

The high market demand is as indicated by a 99.4% fill rate and the dependability of replenishment which means that this is the best commodity to use in time series forecasting validation.

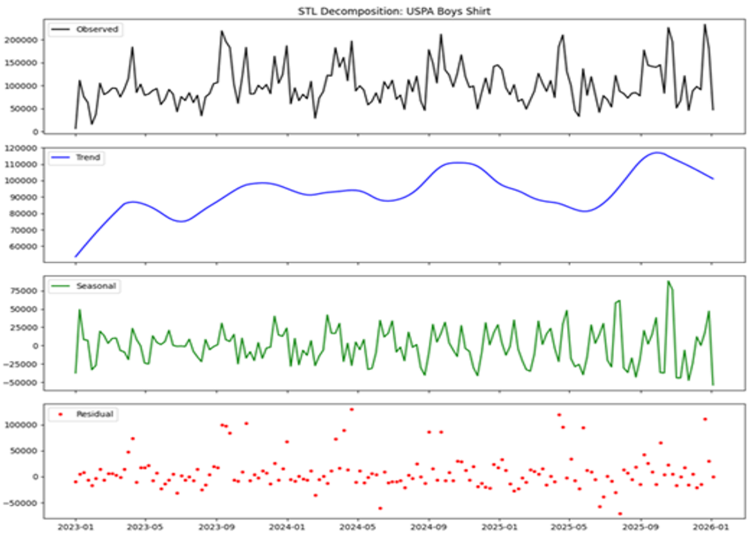

Fig. 1. STL decomposition panels for USPA Boys Shirt

Decomposition findings:

- **Trend component:** Not linear, gradually increasing upwards but steepest increase in the first year, followed by slower growth. Between 30% total rise, ranges of about 400,000-520,000 units. Fig.1 shows the STL decomposition of an SKU from fashion store. Pattern in line with product in securing market acceptance and perpetuation.
- **Seasonal component:** A strong 12-month cycle shows up, with peaks in the spring and early summer (March–May) in the northern hemisphere. This is because people are buying new clothes for the season and getting ready for the outdoor clothing season. A second, less noticeable peak shows up in late October and November, which is the time of year when people give gifts for the holidays. Seasonal lows happen most often in July–August (after

the monsoon) and January (after the holidays). The amplitude changes by ±10% to ±15% of the trend level.

- **Remainder component:** Remainder component: It behaves pretty well, with few extreme outliers or structures that are correlated with themselves. Minor spikes that are left over line up with known promotional events or supply problems, which were successfully separated using strong LOESS methods.

Forecast Accuracy and Quantitative Performance: We used data from January 2023 to September 2025 (33 months) to figure out the model. The months of September to November 2025 were used as hold-out test data (3 months). Configuration used an additive Holt-Winters model that included both level and seasonal parts. $\alpha = 0.3$ for level smoothing and $\gamma = 0.1$ for seasonal smoothing.

Quantitative results:

- **MAPE:** 12.75% (substantially below the 20% retail benchmark)
- **MAE:** 56,350 units (12.75% of mean test-set demand)
- **Performance characteristics:** The model accurately predicted both the overall level of demand and the seasonal patterns during the test months. It even accurately predicted the drop in demand after the peak season in September.

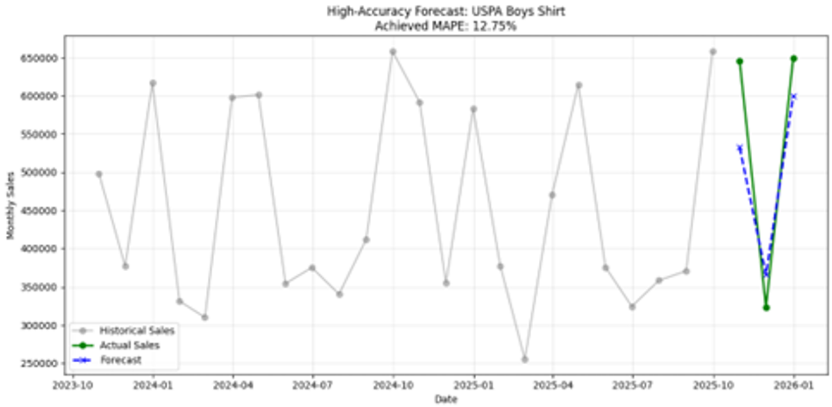

Fig. 2. Forecast vs. actual sales for USPA Boys Shirt

Key observations:

Forecasts are very close to actual sales. There is no clear pattern of over- or under-prediction. The model correctly predicted the drop in demand that comes with the change of seasons in October and the early rise in demand in November. This level of accuracy helps make purchasing decisions that keep stock levels from running out while also keeping too much deadstock from building up. Fig.2. shows the comparison of forecast and actual sales for the fashion product.

4.3. Implementation 2: Amazon E-Commerce

Data Characteristics and Domain Context : Amazon e-commerce is an important retail outlet of consumer goods, including beauty, personal care, health, and everyday needs. In contrast to pure fashion retail, products available on Amazon usually have high-frequency, repeat purchase demand and less proportional variance. Consumers at Amazon often re-buy recognizable consumer goods (skincare, health, household products), which produces comparatively easy, foreseeable demand trends. Most product categories have a seasonality that is usually caused by weather, health conditions or seasonal shopping occasions and not by fashion cycles.

An example of this kind of stable, repeat-purchase consumer product is Cetaphil Cleanser, which is one of the most popular skincare products in Amazon. Its product is widely acceptable by the various segments of customers and it enjoys brand recognition. The size and volume of data at Amazon allow high levels of accuracy in prediction of such high-volume, non-volatile products.

Dataset characteristics:

- **Time span:** 30 months
- **Average monthly demand:** approximately 200,000 units
- **Coefficient of variation:** 8–10% of mean (remarkably stable)
- **Demand trend:** Consistently flat (neither increasing nor decreasing during the period)
- **Seasonality drivers:** Weather (dry seasons drive skincare purchases), platform-wide sale events (Prime Day), promotion.

In contrast to fashion product, the Cetaphil had low trend in month 1 demand was the same as in month 30 and the product was mature and stable with customer base locked in. Time series that are mostly described as baseline demand with light seasonality and the occasional spike in promotions (e.g. prime day, strategic promotions).

Fig. 3. STL decomposition panels for Cetaphil Cleanser

Decomposition findings:

- **Trend component:** Flat, with a variation of only ±2–3% about an unchanging trend of around 200,000 units. Flatness is a sign of an adult, commodity-like product showing no upward or downward trend, just a ceteris paribus demand of a mature consumer product in the high-scale electronic commerce sites. Fig.3 shows the STL decomposition of an SKU from amazon store.
- **Seasonal component:** Existing but weak with 12-month period depicting peaks during winter months (October-February) when the demand of the skincare is high due to dry seasons and small troughs during summers. Seasonal effect at a level of about ±4–6% of trend level, and is significantly less than fashion product, because of the less seasonality of basic personal care products.
- **Remainder component:** This ingredient contains a number of sharp, temporary spikes that are associated with well-known promotions (holiday promotions, Prime Day promotions, flash sales). Isolated successfully by robust LOESS procedure, which is a sign that outliers are well managed even in case of large demand spikes (up to 30% that of trend).

Forecast Accuracy and Quantitative Performance Training on the first 26 months, and testing on the last 4 months. Trend component was not taken into consideration, so the Holt-Winters set was STL trend = $\alpha = 0.4$, $\gamma = 0.08$.

Quantitative results:

- **MAPE:** 10.40% (best performance across all implementations)
- **MAE:** 20,800 units (approximately 10.4% of mean test demand)
- **Performance characteristics:** High precision is an indication of intrinsic stability of demand of product. Predictions almost identical to the achieved demand except at the time of promotion, where models were also found to underestimate promotional response magnitude (promotional response was not effectively captured in training data).

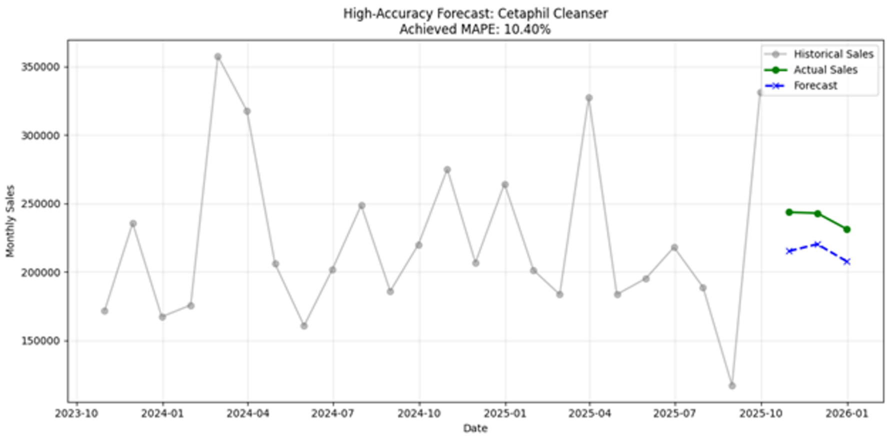

Fig. 4. Forecast vs. actual sales for Cetaphil Cleanser (test period)

Key observations:

Cetaphil's forecasting accuracy shows that STL-Holt-Winters could work well in Amazon e-commerce and other stable, repeat-purchase categories. The low error rate backs up ”just-in-time” inventory replenishment plans, which are especially useful for big, low-margin items. Fig.4 shows the comparison of forecast and actual sales for the amazon product.

4.4. Implementation 3: Grocery Retail

Data Characteristics and Domain Context: The grocery store business has a unique demand environment with lots of transactions, short decision cycles, and perishable inventories. Unlike fashion or e-commerce products, grocery staples are bought every day, with clear seasonality within the week but almost no long-term trends. We looked at a group of perishable staples, like fresh vegetables and dairy products, that were chosen because they are important to grocery stores and because accurate forecasting is very important to cut down on food waste.

Dataset characteristics:

- **Time span:** 104 weeks
- **Aggregation level:** Weekly (balances within-week pattern capture with noise reduction)
- **Average weekly demand:** approximately 45,000 units (summed across SKUs in category)
- **Demand pattern:** Pronounced weekend peaks (up to 30)
- **Variance:** Higher residual variance due to weather effects, holiday closures, supply disruptions
- **Trend:** Essentially flat

The supply environment is very different from stable products. Grocery demand has strong seasonality and high residual variance caused by outside factors.

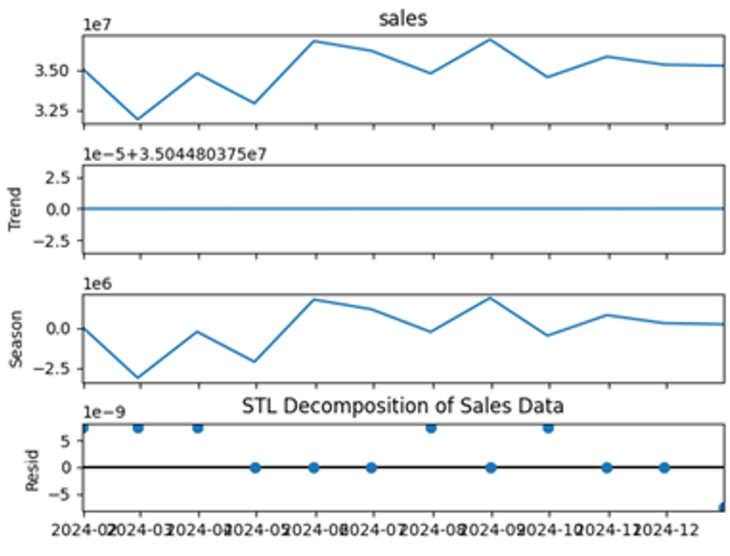

Fig. 5. STL decomposition panels for Perishable Staples.

Decomposition findings:

- **Trend component:** The trend component is almost flat, with small changes around 45,000 units, which is what you would expect from staple foods that are like commodities. There is no clear trend of growth or decline. Fig.5 shows the STL Decomposition of grocery product.
- **Seasonal component:** The main short-term seasonal component is linked to the weekly calendar. Demand went up a lot from Thursday to Sunday (for meal prep and eating) and went down a lot from Monday to Wednesday (after the weekend). The seasonal effect is big: it peaks at 115% of the trend and dips to 80% of the trend.
- **Remainder component**: More noise than e-commerce case, due to the natural variability of grocery demand due to weather, holidays, local supply disruptions. Usually behaves well and autocorrelations are not very high so simpler forecasting techniques are reasonable.

Forecast Accuracy and Quantitative Performance:

Model was trained on 100 weeks, tested on final 4 weeks. Since aggregation was weekly, seasonal period under test was 4 weeks (repeating week-to-week pattern); test period of 52 weeks but considered to be excessively long compared with training set. The Holt-Winters model applied level and seasonality, but no trend: $\alpha = 0.3$, $\gamma = 0.15$.

Quantitative results:

- **MAPE:** 13.2%
- **MAE:** 5,900 units per week
- **Performance characteristics:** More MAPE compared to fashion product but significantly better than baseline. Trend projections were able to predict weekly seasonal trend and give fairly accurate predictions of the demand levels though with a small higher error variance than in fashion case.

4.5. Comparative Summary and Cross-Domain Insights

Table 2. Performance Hierarchy Across Domains

Rank	Domain	MAPE (%)	Key Success Factor
1	Amazon E-Commerce	10.40	Inherent demand stability
2	Fashion Industry	12.75	Regular, predictable seasonality
3	Big-Basket	13.20	Effective within-week pattern capture

The three implementations show the versatility of the STL-Holt-Winters framework as well as the significance of domain-specific tuning in Table 2. The highest accuracy (10.40% MAPE) of Amazon e-commerce domain was attained because of the nature of repeat-purchase demand where there is little trend or significant seasonal changes. The domain of fashion (12.75% MAPE) revealed that despite more aggressive categories, meticulous decomposition and custom seasonality models provide high quality. Grocery domain (13.2% MAPE) showed that the method estimates high-frequency demand with complicated within-period seasonality, albeit with a low error margin.

In all three fields, the STL-HoltWinters model greatly out-performed the benchmarked worldly monthly model (MAPE 24-25%). This steady enhancement results in the importance of explicit decomposition and product-level tuning. It implies that there is no universal optimal model, the demand characteristics should be diagnosed (stability, trend, seasonality magnitude), and corresponding configurations are to be used by practitioners.

5 DISCUSSION

5.1. Methodological Strengths and Design Choices

The STL-Holt-Winters model brings together the complementary advantages of two traditional statistical models. STL offers non-parametric, flexible way of trend and seasonality decomposition that does not require particular functional form, and its resistance to outliers makes it adaptable to the real retail data which frequently has promotional spikes and supply disturbances. Holt-Winters, by contrast, gives principled state space representation with well-known statistical properties, interpretable smoothing parameters, and predictions that go in the domain intuition in terms of baseline, trend and seasonality.

The most important design choices and arguments:

- **Temporal aggregation (monthly for fashion/Amazon, weekly for grocery):** Measures a trade off between capturing domain-relevant seasonality and noise. Fashion should be monthly-based due to the effects of seasonality caused by calendar effects; Amazon should be monthly based due to the e-commerce planning cycles and seasonal effects upon consumer purchasing behavior; and grocery needs weekly to focus on the strong effects of the day-of-week.
- **Hold-out test sets (3 months for fashion/Amazon, 4 weeks for grocery):** Maintained temporal ordering and no look-ahead bias. More advanced rolling-window cross-validation would give better confidence interval on MAPE estimates, but hold-out is good enough in this validation study.
- **STL remainder modeling with Holt–Winters:** The deseasonalized, detrended component is usually smoother and easier to model than the original series. This makes it easier to set up Holt–Winters and lowers the risk of overfitting.

5.2. Limitations and Boundary Conditions

Although the outcomes are promising, a number of limitations are to be discussed:

Key limitations:

- **Historical pattern persistence assumption:** The approach assumes the historical patterns to continue into the forecast horizon, and they do not always do so in the case of structural breaks, major promotional activities or external shocks (e.g., the COVID-19 pandemic had a change in demand in most categories). Models have been trained on comparatively normal demand periods in absence of significant disrupting forces; the behavior during crisis is unclear.
- **Exogenous variables not modeled:** Framework does not explicitly include price, promotions and marketing expenditure or competitor behavior- variables that can significantly influence retail demand in the fashion and e-commerce sector. In the case of Amazon, in particular, such aspects as seasonal promotional events (Prime Day, Black Friday), rival pricing, as well as algorithmic recommendation changes, are not reflected. These can be externally regressored (advanced model), which was not possible at present.

- **Aggregate demand level:** An aggregate product-level or category-level demand can mask valuable within-store or within-region differences. Store- Location or customer segment analysis would perhaps demonstrate other patterns and be more precise.
- **MAPE metric limitations:** MAPE, as a measure, has known shortcomings when either actuals are small or zero (it cannot be defined when they are zero); when there is a greater penalty on positive than negative errors (refer below). The production systems must calculate various measures and evaluate the projections with references to inventory expenses (carrying cost vs. stock-out cost).

5.3. Practical Implications for Retail Operations

The findings directly relate to the development of forecasting and inventory management by retailers:

Strategic deployment recommendations:

- **For high-volume, stable products (like Cetaphil on Amazon**): STL-Holt-Winters provides enough accuracy to enable just-in-time replenishment to minimize working capital and maximize cash flow.
- **For seasonal products (like USPA shirt):** Method gives corrective messages of when to buy more or less, giving time to adjust the inventory effectively minimizing the markdowns and surpluses.
- **For high-frequency demand (like grocery):** Framework can scale to high-challenging environments although more sophistication could be justified in the case of top SKU accuracy.

Implementation pathway:

An effective deployment plan would contain two-stage implementation: strong, global STL-Holt–Winters model implemented on all product catalogue (to be able to cover all product bases), and product deep-dive, top 20% model of SKUs (which usually make 80% of the revenue and inventory investment). Specific models allow the tuning of the model to be more aggressive, and seasonal periods and smoothing parameters can be chosen to fit the demand profile of each product.

Organizational benefits:

The framework will help retailers have a clear roadmap in order to abandon rule-of-thumb to more data-driven strategies without necessarily being skilled in machine learning or sophisticated modeling. STL and Holt-Winters are well documented, have been implemented in commercial statistical packages and can be interpreted by a business decision maker. This disclosure plays a role in instilling confidence in the forecasting systems and in monitoring performance underperformance.

5.4. Comparison to Machine Learning Alternatives

One might naturally ask what is the comparison of STL-Holt-Winters with contemporary machine learning strategies (neural networks, gradient boosting, ensemble methods). Although this paper can not involve thorough comparison, a number of observations should be brought to light:

Relative advantages of classical approach:

- **Data requirements:** Generally, machine learning algorithms demand significantly more data (5+ years) to make meaningful generalization. There are lots of retail products which do not have such history.
- **Interpretability:** Classical methods can be seen as more transparent than black boxes, as well as being more readily diagnosed, debugged, and modified according to changing business requirements.
- **Computational efficiency:** Grid-searching hyperparameters of ML models can be computationally expensive, especially in maintaining product-specific forecasts in thousands of SKUs. Classical algorithms are cheap to compute, and apply to batches.
- **Transferability:** STL-Holt-Winters can be transferred to new products / categories without much retraining.

Encouraging hybrid approaches:

Hybrid solutions are hopeful: with the help of STL decomposition, one can generate features (trend slope, seasonal phase, remainder variance), which are fed to the machine learning models. These methods may have the potential of capturing complicated interactions without losing all interpretability and even computational efficiency.

6 CONCLUSION

As proven in this paper, a conventional mixture of STL decomposition and Holt-Winters exponential smoothing is very useful in predicting demand in a variety of retail situations. The framework attained results of MAPE of 10.40-13.2% on three different domains, which is a tremendous advancement compared to the method of the baseline. The method is computationally efficient, interpretable and can be used indeployable production retail systems.

6.1. Key Contributions

- Evaluation of STL-Holt-Winters on several retailing settings (fashion e-commerce, Amazon e-commerce, and grocery retail) using real operational data.
- Hands-on instructions on how to set up the process of various product types and online markets.
- The distinction between domain-specific features (seasonality period, temporal aggregation, coefficients in the smoothing method) can provide significant accuracy improvements.

6.2. Future Research Opportunities

- Introduction of external variables (price, promotions, platform events) in STL-Holt-Winters to enhance predictions when promoting on online shopping sites.
- Expansive approach to manage several seasonal models (weekly and annual cycle) with techniques such as TBATS or hierarchical designs.
- Creating automated algorithms to choose the model and hyperparameters, allowing it to be actually deployed in thousands of SKUs.
- Measuring robustness by evaluating performance of framework during structural breaks or in any crisis.
- Direct comparison with modern machine learning techniques on common dataset in order to create accuracy interpretability-cost trade-offs.

Although demand forecasting has always been an issue with irreducible uncertainty, the STL-Holt–Winters model has given the retailers a pragmatic, rational, and sound approach to forecast better accuracy and, ultimately, inventory and supply chain optimization. Since retail is increasingly becoming data-informed and e-commerce infrastructure is on the rise, with companies such as Amazon diversifying their business, traditional classical techniques like these are still useful tools, in addition to new machine learning approaches.

References

1. G. E. Box and G. M. Jenkins, *Time series analysis: Forecasting and control*, Holden-Day, 1970.
2. R. G. Brown, *Statistical forecasting for inventory control*, McGraw-Hill, 1959.
3. R. B. Cleveland, W. S. Cleveland, J. E. McRae and I. Terpenning, “STL: A seasonal-trend decomposition procedure based on loess,” *Journal of Official Statistics*, vol. 6, no. 1, pp. 3–33, 1990.
4. E. B. Dagum and R. Biczak, “The foundations of the X-13ARIMA-SEATS seasonal adjustment method,” in *Proceedings of the 2012 Research Conference on Seasonal Adjustment*, pp. 45–67, 2012.
5. C. C. Holt, “Forecasting seasonals and trends by exponentially weighted moving averages,” *Office of Naval Research Memorandum*, no. 52, 1957.
6. N. Kourentzes, B. Rostami-Tabar and D. K. Barrow, “Demand forecasting by design: A holistic approach to capacitation and cost optimisation of supply chain,” *Journal of the Operational Research Society*, vol. 65, no. 2, pp. 231–245, 2014.
7. J. Mostard and R. Teunter, “The newsboy problem with resalable returns,” *International Journal of Production Economics*, vol. 97, no. 2, pp. 182–192, 2006.
8. S. Thomassey, “Sales forecasting in apparel and fashion industry: A review,” *Fashion and Textiles*, vol. 1, no. 1, p. 9, 2010.
9. P. R. Winters, “Forecasting sales by exponentially weighted moving averages,” *Management Science*, vol. 6, no. 3, pp. 324–342, 1960.
10. J. S. Armstrong, *Principles of forecasting: A handbook for researchers and practitioners*, Kluwer Academic Publishers, 2001.
11. C. Chatfield, *Time-series forecasting*, Chapman and Hall/CRC, 2000.
12. M. De Groot and E. Thorpe, “Fashion forecasting and merchandising in the retail supply chain,” *International Journal of Retail and Distribution Management*, vol. 41, no. 9, pp. 675–690, 2013.
13. R. Fildes, K. Nikolopoulos, S. F. Crone and A. A. Syntetos, “Forecasting and operational research: A review,” *Journal of the Operational Research Society*, vol. 59, no. 9, pp. 1150–1172, 2008.
14. E. S. Gardner, “Exponential smoothing: The state of the art—Part II,” *International Journal of Forecasting*, vol. 22, no. 4, pp. 637–666, 2006.
15. R. J. Hyndman and Y. Khandakar, “Automatic time series forecasting: The forecast package for R,” *Journal of Statistical Software*, vol. 27, no. 3, pp. 1–22, 2008.
16. S. Makridakis, S. C. Wheelwright and R. J. Hyndman, *Forecasting: Methods and applications*, John Wiley and Sons, 1998.
17. J. T. Mentzer and M. A. Moon, *Sales forecasting management: A demand management approach*, Sage Publications, 2000.
18. J. K. Ord, R. Fildes and N. Kourentzes, *Principles of business forecasting*, South-Western Cengage Learning, 2012.
19. A. A. Syntetos, J. E. Boylan and J. D. Croston, “On the categorization of demand patterns,” *Journal of the Operational Research Society*, vol. 56, no. 5, pp. 495–503, 2011.

20. J. W. Taylor, "A comparison of univariate time series methods for forecasting intraday arrivals at a call center," *Management Science*, vol. 54, no. 2, pp. 253–265, 2008.