Proceedings of the 4th International Conference on Machine Learning and Automation
DOI: 10.54254/2755-2721/2026.39040

Multi-Source Data Fusion Based Market Demand Forecasting and Intelligent Business Decision Optimization

Yuwei Li

School of Applied Science and Engineering, University of Toronto, Toronto, Canada
wellington589125@gmail.com

Abstract. New retail demand is determined by unpredictable sales dynamics, inventory level, promotion actions, weather conditions, and time factors, whereas stockouts may lead to demand underestimation due to distorted sales and negative replenishment decision-making. In order to solve the issue, a multi-source forecasting and business decision system has been designed with the help of FreshRetailNet-50K. First, stockout-affected demand was restored, followed by encoding and dynamic fusion of historical demand, inventory, promotion, weather, calendar, and product data via gated attention network to obtain seven days point and quantile forecasting. After that, forecasting was integrated in the rolling replenishment strategy with regard to ordering, holding, spoilage, stockout, capacity, shelf-life, and lead-time costs. According to experimental results, the proposed multi-source model yielded an MAE score of 6.84, an RMSE score of 10.27, and WAPE value of 29.36% along with the cost index reduction of the replenishment strategy from 424 to 326. The stockout ratio was lowered to 5.8%, spoilage ratio – to 3.7%, inventory turnover became equal to 8.6, and demand fulfillment to 94.2%. Consequently, one may state that incorporation of heterogeneous market data into probabilistic forecasting may lead to better demand estimation and operational efficiency. Thus, the system offers practical grounds for intelligent retail decision optimization, maintaining the link between forecast accuracy and inventory management metrics.

Keywords: Multi-source data fusion, market demand forecasting, latent demand recovery, gated attention network, replenishment optimization

1. Introduction

With constant generation of sales, inventory, promotions, weather, and calendar variables by digital retail systems, there is increased reliance on heterogeneous data for effective demand forecasting beyond historical transactions. Machine learning applications in the retail industry have established that model effectiveness varies widely depending on different demand characteristics and types of goods and services, necessitating the use of temporal and feature heterogeneities appropriately [1]. There is another type of distortion caused by stockouts where the true sales get censored by non-availability of inventory, thus leading to underestimation of the latent demand [2]. Deep learning techniques also suggest the joint representation of local exogenous dynamics and temporal dynamics in order to improve retail sales forecasts [3]. All this calls for a framework that can reconstruct

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.39040

demand distorted by stockouts, fuse multiple information sources on a real-time basis, deliver seven day forecasts of demand both in terms of point and quantiles, and translate such forecasts into rolling replenishment orders.

2. Literature review

2.1. Market demand forecasting

Sales demand forecasting relies more and more on the integration of statistical learning and neural structures since sales time-series exhibit nonlinear seasonality, promotional effects, and distinct properties of products. Neural networks augmented with attention mechanisms could focus on relevant data points and ensure robustness against recent noise events that may hinder the traditional sequential learning [4]. Ensemble hybrid frameworks provide additional evidence that products exhibiting distinct time-dependent patterns require distinct feature extraction and model specification, not one unified approach [5]. Although such frameworks offer more flexible data representation, the forecast quality is also conditioned on the extent to which the target variable represents actual demand. In perishing retail, stockouts affect sales numbers; thus, even a precise model could learn a biased demand process.

2.2. Multi-source data fusion

The retail demand signal depends on cues whose significance varies according to product and time, meaning that static weights are not ideal for capturing feature contributions in a market environment where there are many factors at play. Multi-source deep learning has found that historic demand, economic state, searches and contextual cues could all be sources of information that complement each other as long as interaction between sources is considered [6]. Likewise, studies in multi-channel forecasting have shown that preference information could boost forecasting as long as time and channel-dependent structures are simultaneously captured [7]. The two discoveries reveal a key dilemma: more information means broader coverage, but too much could add noise to a volatile set of relationships.

2.3. Forecast-driven inventory decisions

Forecasting accuracy is not operationally meaningful unless the forecasted demand affects inventory and replenishment decisions. Research on perishable supply chains indicates that heterogeneity in customers and products may significantly influence the forecasting accuracy and inventory-related decisions, implying a need for context-aware demand forecasts [8]. The modeling of inventory policies for finite lifetime products also indicates that age-based demand, clearance, and expiration should be considered in replenishment policies since new arrivals can help in reducing shortages at the cost of increased spoilage risk [9]. This presents a trade-off for decision-makers between safeguarding against service and inventory decay risks. Point forecasts only provide expected demand values but offer little insight regarding tail risk, but quantile demand can enable risk-based ordering decisions.

DOI: 10.54254/2755-2721/2026.39040

3. Experimental methods

3.1. Data acquisition and preprocessing

FreshRetailNet-50K is chosen as the experimental data source because it offers stockout-aware high-frequency data for the censored-demand problem [10]. The dataset includes 50,000 store-product time series over 90 days from 898 stores in 18 cities and over 860 perishable products. The available features are hourly sales, hourly stock status, discount rate, promotions, precipitation, temperature, humidity, wind level, holidays, store IDs, and hierarchical product categories. Stockout hours are defined via the inventory status, and censored sales are imputed by taking the median demand of similar non-stockout hours that match on the store-product pairs, day-of-week, and week-to-week cycle. Aggregated demand values are recovered to daily frequency. Continuous features are scaled using training-period statistics, while categorical store, city, category, and calendar features are one-hot encoded. Each supervised learning example is based on the previous 28 days of data and aims to predict demands of the subsequent seven days. The dataset is split chronologically into 70%, 15%, and 15% train, validation, and test subsets, respectively, and the parameters of all preprocessing steps are trained only on the training data.

3.2. Multi-source forecasting network

Models such as seasonal naïve, XGBoost, and LSTM are used to set up baselines, and the fusion network is applied to distinguish between past demand, stock, promotion-weather, and categorical data. One-D CNN is employed for capturing the short-term dynamics of the demand, while GRU is used for encoding the sequence of stock, discounts, rainfalls, and temperatures. Dynamic source importance is calculated through gated attention, as shown in Equation (1), allowing the contribution of each information stream to vary by forecast origin and product.

$$e_{s,t}=v^{T}\tanh\left(W_h h_{s,t}+W_c c_t+b\right)$$

$$\alpha_{s,t}=\frac{\exp(e_{s,t})}{\sum_{j=1}^{S}\exp(e_{j,t})} \tag{1}$$

Here, h is the hidden representation of each source, c is the contextual vector, and α is the normalized attention weight. The fused vector generates the seven-day point forecast together with 0.1 and 0.9 demand quantiles. Training jointly controls magnitude error and interval calibration through Equation (2).

$$\mathcal{L}=\lambda\frac{\sum_{i=1}^{N}w_i\left|y_i-\hat{y}_i\right|}{\sum_{i=1}^{N}w_i}+(1-\lambda)\sum_{\tau\in\{0.1,0.9\}}\frac{1}{N}\sum_{i=1}^{N}\rho_\tau\left(y_i-\hat{y}_i^{(\tau)}\right) \tag{2}$$

Here, y is observed demand, ŷ is predicted demand, w is the sample weight, λ balances the objectives, and ρ denotes quantile loss. Ablation tests separately remove inventory, promotion-weather, and demand-recovery inputs.

DOI: 10.54254/2755-2721/2026.39040

3.3. Rolling replenishment optimization

Seven-day forecasts from Section 3.2 enter a store-product rolling replenishment procedure updated once per day. At each decision point, current inventory is observed, expired units are removed, demand forecasts are refreshed, and the order quantity is recalculated for the remaining horizon. The cost objective minimizes ordering, holding, spoilage, and unmet-demand penalties, as shown in Equation (3), following the rolling-horizon logic used in recent perishable-inventory optimization [11].

$$\min_{\{q_t\}} C = \sum_{t=1}^{H} \left(c_o q_t + c_h I_t^+ + c_w W_t + c_s U_t \right) \tag{3}$$

Here, q is the order quantity, I is positive ending inventory, W is expired inventory, U is unmet demand, and the four c terms represent unit costs. Inventory transition and capacity feasibility are imposed through Equation (4).

$$I_{t+1} = \max\left\{0, I_t + q_{t-L} - \hat{D}_t - W_t\right\}, \quad 0 \le q_t \le K - I_t \tag{4}$$

Here, L denotes replenishment lead time, K is storage capacity, and $\hat{D}$ is predicted demand. Shelf-life expiration is tracked by inventory age before each transition. Historical replay compares fixed-cycle ordering, observed-sales forecasting, recovered-demand forecasting, and multi-source probabilistic forecasting under identical cost parameters.

4. Results

4.1. Forecasting performance

Table 1 indicates that the multi-source fusion network attained the best forecasting accuracy in terms of MAE of 6.84 units, RMSE of 10.27 units, and WAPE of 29.36%, compared with those of the LSTM algorithm. With regard to MAE, there was a reduction of 17.2%, while for RMSE, there was also a reduction of 17.2%, implying that the heterogeneous context improved not only the control of average errors but also large errors. Moreover, there was WPE of 0.74% in the fusion algorithm, while in the non-fusion model, which lacked latent demand recovery, it was −5.83%. This implies that the reconstruction of the stockout periods mainly fixed systematic underestimation rather than decreasing random error. In addition, when promotion-weather inputs were stripped off, there was an increase in WAPE from 29.36% to 31.12%, meaning that the exogenous factors provided additional short-term information. The prediction interval of 0.1−0.9 gave an empirical coverage of 78.6% instead of the nominal 80%.

Table 1. Forecasting performance on the chronological test set

Model	MAE	RMSE	WAPE	WPE	0.1–0.9 Coverage
Seasonal Naïve	9.71	14.63	35.92%	−7.44%	—
XGBoost	8.54	12.88	31.64%	−4.12%	—

DOI: 10.54254/2755-2721/2026.39040

Table 1. (continued)

LSTM	8.26	12.41	30.87%	−3.21%	—
Fusion without Demand Recovery	7.53	11.20	31.05%	−5.83%	76.9%
Fusion without Promotion-Weather	7.31	10.98	31.12%	−0.91%	77.4%
Multi-source Fusion	6.84	10.27	29.36%	0.74%	78.6%

4.2. Replenishment performance

As shown in Figure 1, the cumulative cost of operations under the multi-source forecasting-based replenishment policy grew at the slowest pace, resulting in an index of 326 after four weeks of replays, as opposed to 355 with the recovered-demand forecasting, 382 with the observed sales forecasting, and 424 with the fixed cycle ordering. This means that the cost saving is equal to 23.1%, demonstrating that the benefits of forecasting have been transferred into the operation domain. The multi-source policy led to a reduction of the stockout rate from 10.9% to 5.8%, a reduction of the spoilage rate from 5.4% to 3.7%, an increase of the inventory turnover rate from 7.1 to 8.6 cycles, and an increase of demand satisfaction from 89.1% to 94.2%. Both the stockout and the spoilage rates were lowered at the same time; thus, the improvement was made due to something else except for a large amount of safety stock.

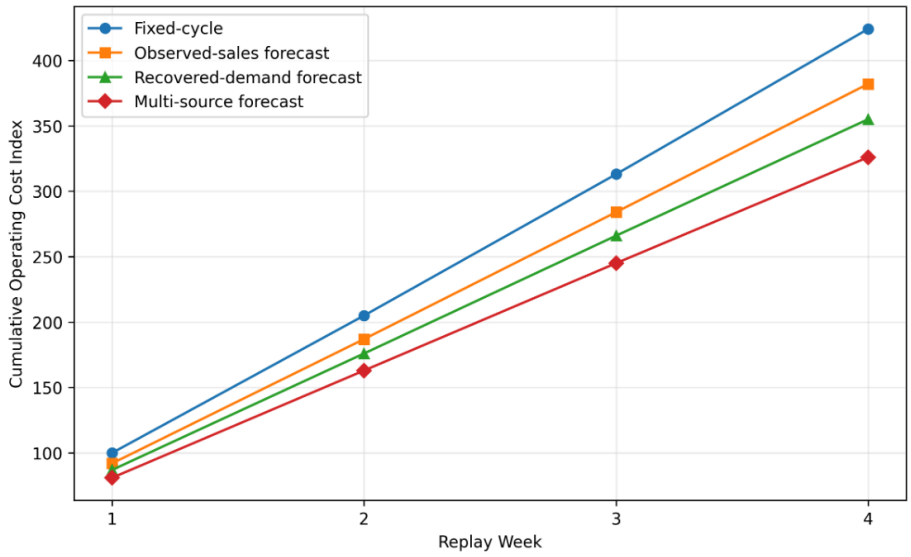


Figure 1. Cumulative cost under rolling replenishment policies

5. Discussion

There are two separate explanations of how these two mechanisms affect performance. Recovery from latent demand is focused on recovery of information losses incurred during stockouts as evidenced by a decrease in the level of negative WPE. Multi-source fusion, on the other hand, is effective in responding to changes in temporary demand caused by promotions and weather. The distinction between the two mechanisms is crucial since a decreased overall error rate is caused by different factors which have their implications for operational management. The replenishment findings indicate that improved forecasts do not necessarily entail increased inventories, as both reductions in stockout and spoilage rates occurred at the same time, thus implying that it was the

DOI: 10.54254/2755-2721/2026.39040

timing and allocation rather than additional safety stock which contributed to better forecasts. The fact of close to nominal quantile coverage is an explanation for that as upper-demand risk was accounted for without considering all periods as those of high demand. However, the residual gap between nominal and empirical coverage still remains because of the uncertainty involved in extreme cases. Cost savings also depend on the values of lead times, shelf lives and penalty coefficients which vary across products.

6. Conclusion

The unified approach to forecasting and replenishment has been developed for perishable goods in retail with consideration of censored sales due to stockouts using the multi-source data approach. The latent-demand recovery mechanism recovered the true value of demand from censored sales, gated attention combined demand, inventory, promotional events, weather, and calendar information, and probabilistic forecasts have been turned into decision-making about orders considering business constraints. The structure of experiments makes it possible to establish the direct link between the quality of forecasting and cost, stockout, spoilage, turnover, and fulfillment results. However, the proposed approach is still limited by only one retailer dataset and artificial cost structure which can reduce its applicability to other businesses and supply chains.

References

[1] Nasseri, M., et al. (2023). Applying machine learning in retail demand prediction—A comparison of tree-based ensembles and long short-term memory-based deep learning. *Applied Sciences, 13*(19), 11112.
[2] Yin, C., Zheng, Y., & Kong, Z. (2026). Demand hidden by stockouts: Censored demand recovery and green waste-reduction forecasting for sustainable fresh-food consumption under emerging-market urbanization. *Sustainability, 18*(15), 7642.
[3] de Castro Moraes, T., Yuan, X.-M., & Chew, E. P. (2024). Hybrid convolutional long short-term memory models for sales forecasting in retail. *Journal of Forecasting, 43*(5), 1278–1293.
[4] Cui, L., et al. (2024). A novel attLSTM framework combining the attention mechanism and bidirectional LSTM for demand forecasting. *Expert Systems with Applications, 254*, 124409.
[5] Mejía, S., & Aguilar, J. (2024). A demand forecasting system of product categories defined by their time series using a hybrid approach of ensemble learning with feature engineering. *Computing, 106*(12), 3945–3965.
[6] Zhang, D., et al. (2024). Forecasting duty-free shopping demand with multisource data: A deep learning approach. *Annals of Operations Research, 339*(1), 861–887.
[7] Wu, J., et al. (2024). Unveiling consumer preferences: A two-stage deep learning approach to enhance accuracy in multi-channel retail sales forecasting. *Expert Systems with Applications, 257*, 125066.
[8] Jauhar, S. K., et al. (2025). Explainable artificial intelligence to improve the resilience of perishable product supply chains by leveraging customer characteristics. *Annals of Operations Research, 354*(1), 103–142.
[9] Goyal, A., & Sapra, A. (2024). Managing replenishment and clearance of perishables: Last-in, first-out (LIFO) issuing policy and age-sensitive demand. *Production and Operations Management, 33*(4), 1031–1052.
[10] Azar, J. (2026). A reproducible, leakage-free pipeline for censored demand forecasting and inventory optimization on FreshRetailNet-50K. *Inventions, 11*(5), 89.
[11] Lejarza, F., Venkatesan, S., & Baldea, M. (2025). Rolling horizon product quality estimation and online optimisation for supply chain management of perishable inventory. *International Journal of Production Research, 63*(10), 3709–3732.