Journal of Science, Innovation & Social Impact
EISSN: 3105-5028 | PISSN: 3105-501X | Vol. 1, No. 1 (2025)

Article

Multi-Source Data Fusion for Short-Term Demand Forecasting of Seasonal Retail Products: An Empirical Study Using Weather and Social Media Signals

Jialu Wang [1,*]

[1] Business Administration, Fordham University, NY, USA
* Correspondence: Jialu Wang, Business Administration, Fordham University, NY, USA

Abstract: Recent advances in machine learning have enabled retailers to leverage diverse data streams for demand prediction, yet most systems continue to rely primarily on historical sales. In this study, we address the challenge of seasonal demand forecasting by integrating meteorological observations, social media activity, and economic indicators with point-of-sale transactions. We propose a hierarchical ensemble that combines gradient boosting machines with bidirectional LSTMs, where component weights adapt according to product category and forecast horizon. Through mutual information screening, we reduce feature dimensionality from 123 to category-specific subsets averaging 38 features, lowering computational costs by 60% while maintaining predictive accuracy. Our approach requires only standard training of the base models along with a lightweight meta-learner, without necessitating end-to-end joint training. Evaluation on 10.3 million transactions from three U.S. retail chains demonstrates that our method achieves a mean absolute percentage error of 19.6%, compared to 37.5% for seasonal naive baselines. Multi-source fusion thus provides a practical pathway to more accurate retail forecasting, without the complexity associated with end-to-end deep learning systems. Crucially, this approach relies on lightweight feature learning and selection techniques combined with meta-learner weighting, rather than on end-to-end joint representation training.

Keywords: seasonal demand forecasting; multi-source data fusion; retail analytics; hierarchical ensemble; gradient boosting machines

Received: 03 October 2025
Revised: 23 October 2025
Accepted: 09 November 2025
Published: 13 November 2025

1. Introduction

Retail demand often exhibits patterns that challenge traditional forecasting methods. A sudden drop in temperature can double coat sales overnight. A viral social media post can deplete shelves of an obscure product within hours [1]. Economic uncertainty may shift consumers from premium to value brands. Such external factors create abrupt changes that historical averages alone cannot predict.

The challenge becomes particularly pronounced during seasonal transitions. Consider outdoor furniture in April: last year's sales provide a baseline, but this year's weather differs. Temperatures run 10 degrees above normal [2]. Social media is abuzz with discussions of backyard renovations. Consumer confidence remains strong despite inflation concerns. How should these diverse signals be integrated to adjust forecasts?

Most retailers address this by building separate models for each factor-one for weather effects, another for promotions, and a third for trends-then manually combining predictions [3]. This approach often fails to capture interactions. Hot weather can amplify

promotional responses. Social media virality may depend on economic conditions [4]. These cross-effects are as important as the primary effects themselves.

Machine learning provides a potential solution. Tree-based models can discover non-linear relationships. Neural networks can handle sequential data. Ensemble methods can combine diverse predictions [5]. Yet retail forecasting presents unique challenges: thousands of products with sparse sales, multiple overlapping seasonalities, and extreme events that disrupt historical patterns.

We develop a framework that adapts the weighting of different data sources based on context [6]. When predicting next-day sales, historical patterns dominate. For forecasts two weeks ahead, external signals become more critical as uncertainty grows. The system adjusts these weights by product category-weather is more influential for beverages, social media for fashion, and economic signals for electronics.

Three technical contributions enable practical deployment. First, mutual information screening identifies the most relevant features for each category, reducing computational load without sacrificing accuracy. Second, asynchronous fusion preserves high-frequency signals instead of downsampling all data to monthly granularity. Third, hierarchical ensemble learning combines specialized models through adaptive weighting [7].

We validate our framework across three retail contexts representing distinct forecasting challenges. A grocery chain with 89 stores tests stable demand prediction. An outdoor retailer with 34 locations evaluates weather sensitivity. A fashion chain with 33 stores examines responsiveness to trends [8]. Collectively, these cases span the spectrum from routine replenishment to volatile fashion cycles.

Our experiments show that optimal fusion strategies depend on product characteristics rather than universal rules. Weather-sensitive items benefit from early integration of meteorological data. Promotional products gain more from late fusion of independent models. Stable grocery items show minimal improvement from external data. These findings provide guidance for practical implementation: sophisticated fusion should be prioritized only where it produces meaningful returns [9].

2. Literature Review and Theoretical Foundation

2.1. Evolution from Statistical to Neural Forecasting Architectures

Statistical time series models dominated retail forecasting for decades due to their mathematical tractability and interpretable parameters. Exponential smoothing state space models, formulated as x_t = l_ (t - 1) + b_ (t - 1) + s_ (t - m) + epsilon_t, capture level, trend, and seasonal components through recursive filtering. Multiple seasonal (MS) extensions incorporate K seasonal periods: x_t = l_ (t - 1) + b_ (t - 1) + sum from k = 1 to K of s_(t-m_k) ^ (k) + epsilon_t, where m_k denotes the k-th seasonal period. These extensions reduce forecast errors by capturing overlapping seasonal patterns more effectively than single-seasonality models.

Gradient boosting trees fundamentally change the modeling paradigm by learning non-parametric mappings through recursive partitioning. Weather features-including temperature, humidity, precipitation, wind velocity, and atmospheric pressure-can be combined with polynomial interactions up to degree 3 [10]. The tree ensemble F(x) = sum from m = 1 to M of gamma_m * h_m(x) aggregates M weak learners h_m, each capturing local patterns that global linear models cannot detect.

The computational advantages extend beyond predictive power. Tree-based methods handle mixed variable types using native splitting rules, eliminating the dimensional explosion associated with one-hot encoding. Missing values are implicitly imputed through surrogate splits that preserve the local tree structure. Non-linear relationships emerge through tree depth rather than explicit basis functions, reducing the burden of feature engineering [11].

2.2. Information Fusion Strategies in Supply Chain Systems

Multi-source integration follows two main approaches: early fusion, which concatenates features before modeling, and late fusion, which combines predictions from

specialized models. Ensemble methods consistently outperform single-model approaches by reducing variance and optimizing the bias-variance tradeoff [12]. The effectiveness of ensembles depends on the diversity of base models, as homogeneous ensembles provide minimal improvement due to correlated errors.

Temporal synchronization remains a computational challenge. Point-of-sale systems update instantaneously, weather stations report hourly, and economic indicators arrive monthly. Most frameworks adopt lowest-common-denominator synchronization, downsampling all sources to match the coarsest granularity. This approach sacrifices high-frequency variations that carry predictive signals. Alternative strategies, such as upsampling through interpolation, introduce artifacts: linear interpolation creates false smoothness, nearest-neighbor interpolation creates discontinuities, and spline interpolation can produce spurious oscillations [13].

2.3. Environmental and Social Drivers of Retail Demand

Meteorological conditions influence consumer behavior through complex psychophysiological mechanisms. Empirical evidence reveals pronounced asymmetries in weather response functions. Positive temperature deviations (Delta T > 0) increase beverage sales following S (Delta T) = 1 / (1 + exp (-0.3 * Delta T)), while negative deviations follow a linear decline S (Delta T) = 1 + 0.08 * Delta T. This asymmetry-immediate response to heat versus delayed response to cold-invalidates symmetric error assumptions in classical forecasting [14].

Social media operates through different transmission mechanisms, generating leading indicators with temporal lags tau ranging from 3 to 7 days. Viral propagation follows epidemic models: dI / dt = beta * S * I - gamma * I, where S represents the susceptible population, I represents the informed population, beta denotes the transmission rate, and gamma denotes the decay rate. Sentiment polarity extracted through transformer architectures correlates with purchase intent at r = 0.42 for fashion categories, although causality may be confounded by simultaneous marketing campaigns [5].

3. Methodology

3.1. Data Collection and Integration

Four data streams converge in our pipeline, each requiring specialized handling. Point-of-sale transactions arrive continuously through Apache Kafka, capturing every purchase with timestamp, product, quantity, price, and location. We achieve a 99.97% capture rate, which is crucial because even small losses can delay the detection of emerging trends. Lagged backfills are monitored to mitigate gaps [15].

Weather observations are collected from 1,247 National Weather Service stations reporting 28 variables hourly. Spatial interpolation addresses geographic gaps using weights calculated as w_i = (1 / d_i^2) divided by the sum over j of (1 / d_j^2), where stations are weighted by the inverse of distance squared. When stations fail, numerical weather model outputs are substituted, accepting slightly lower accuracy to maintain continuity [16].

Social media data requires extensive filtering. Raw Twitter streams contain roughly 500,000 daily posts matching retail keywords, but most are noise. Ensemble classifiers remove spam with 94% precision, though about 6% of noise remains. Sentiment analysis using RoBERTa extracts opinions from the remaining text.

Feature engineering transforms raw data into predictive signals. Temperature alone generates 12 features, including current value, deviation from normal, recent trend, and comfort index. Temporal lags are created at 1, 7, 14, and 28 days to capture various cycles. Interaction features multiply variables, such as temperature × weekend and precipitation × promotions. This expansion increases the base 97 features to 123 total features by including lags and interactions (As shown in Table 1).

Table 1. Feature Categories and Dimensions After Preprocessing.

Source	Updates	Features	Missing %
Transactions	Continuous	42	0.3
Weather	Hourly	28	3.2
Social	Real-time	15	8.7
Economic	Monthly	12	0.0

3.2. Feature Selection

Dimensionality reduction prevents overfitting while preserving predictive signal. We apply three complementary approaches:

Mutual information measures statistical dependence between features and the target. It is calculated as I (X; Y) = sum over x and y of p (x, y) * log (p (x, y) / (p(x) * p(y))). Features with mutual information less than 0.05 nats with the target are eliminated. This step removes roughly 40% of features that provide negligible information.

Recursive feature elimination using gradient boosting iteratively removes low-importance features. We train 100-tree models, measure feature importance through total gain, eliminate the bottom 10% of features, and repeat until validation error begins to increase. Although this greedy approach can occasionally drop useful features, it ensures computational efficiency [17].

LASSO regularization adds an L1 penalty to the regression objective: minimize over beta of (1 / (2 * n)) * ||y - X * beta||^2 + lambda * ||beta||_1. Cross-validation selects lambda using the one-standard-error rule, favoring sparsity when performance differences are negligible (As shown in Table 2).

Table 2. Feature Counts by Selection Stage and Product Category.

Category	Original	After MI	After RFE	After LASSO
Outdoor	123	89	52	47
Apparel	123	91	48	41
Beverages	123	78	39	35
Electronics	123	71	31	28

3.3. Ensemble Architecture

Three base models capture different patterns in retail demand:

SARIMA (2, 1, 2) × (1, 1, 1) _7 handles linear trends and weekly seasonality through a parametric structure. Despite its simplicity, it provides a strong baseline for stable products.

XGBoost with 500 trees learns non-linear interactions through recursive partitioning. The trees automatically discover thresholds and interactions without explicit specification.

A bidirectional LSTM with two layers (hidden sizes 128 and 64) models sequential dependencies. Forward and backward passes capture patterns in both temporal directions.

A meta-learner combines predictions from the base models using learned weights. The weights are computed as w = softmax (W2 * ReLU (W1 * [h, t, c] + b1) + b2), where inputs include the forecast horizon (h), calendar and time features (t), and a category embedding (c). The softmax function ensures a valid probability distribution over the models (As shown in Table 3).

Table 3. Base Model Architectures and Hyperparameters.

Model	Training Time	Parameters	1–MAPE (higher is better)
SARIMA	2.3 sec	12	72.5%
XGBoost	45 sec	125,000	77.4%
LSTM	12 min	95,232	76.3%
Meta-learner	30 sec	8,000	-

*MAPE = 27.5% ↔ 1−MAPE = 72.5%.

4. Experimental Results and Performance Evaluation

4.1. Dataset Characteristics and Experimental Protocol

Three retail chains provide empirical validation, each representing distinct market segments with heterogeneous operational characteristics. The national grocery chain encompasses 89 stores distributed across four Köppen climate classifications: continental (Dfa), humid subtropical (Cfa), Mediterranean (Csb), and semi-arid (BSk). These locations generated 4.2 million transactions over 24 months, capturing multiple seasonal cycles while avoiding single-year anomalies. Transaction patterns exhibit substantial heterogeneity: urban stores (n = 31) process 127 ± 43 transactions hourly with an average basket size of $23.40, suburban supercenters (n = 42) handle 89 ± 28 transactions hourly, averaging $67.80, while wholesale clubs (n = 16) process 34 ± 12 bulk orders hourly, averaging $142.30.

The outdoor equipment retailer operates 34 Pacific Northwest locations serving recreational enthusiasts. Their catalog spans 8,400 SKUs with demand variability coefficients ranging from CV = 0.4 (basic camping equipment) to CV = 3.2 (specialized winter sports gear). Sales exhibit pronounced weekly seasonality with weekend volume reaching 2.3 times weekday levels. Bimodal annual patterns emerge from optimal recreation conditions: a spring peak (April-May) when temperatures reach 60-75°F with minimal precipitation, and a fall peak (September-October) driven by favorable early-fall weather and holiday-related demand peaks.

Fashion retail encompasses 33 metropolitan stores targeting trend-conscious demographics (age 18-35, household income >$50,000). Social media engagement demonstrates stronger predictive power than traditional variables: Instagram interactions correlate with next-week sales at r = 0.62 compared to r = 0.31 for promotional spending. Inventory velocity reaches 8.2 turns annually compared to the industry median of 4.3, amplifying forecasting importance. Markdown rates approaching 40% create complex price-demand feedback loops requiring simultaneous optimization.

Computational infrastructure leverages distributed GPU acceleration: 8 × NVIDIA V100 (32GB) coordinated through Horovod, which implements ring-allreduce for gradient synchronization. The software stack includes PyTorch 1.13 for neural architectures, XGBoost 1.7.0 with GPU histogram algorithms, statsmodels 0.13 for classical time series, and Optuna 3.0 for hyperparameter optimization using Tree-structured Parzen Estimators with median pruning. While the heatmap shows perishables approaching approximately 12-15% error at longer horizons, the statement "<10% at 7 days" refers to median performance across subcategories. This slight discrepancy reflects variation across specific product groups rather than a contradiction. Groceries remain relatively stable compared to more volatile categories, with errors rising moderately to around 25% at 14 days (As shown in Figure 1).

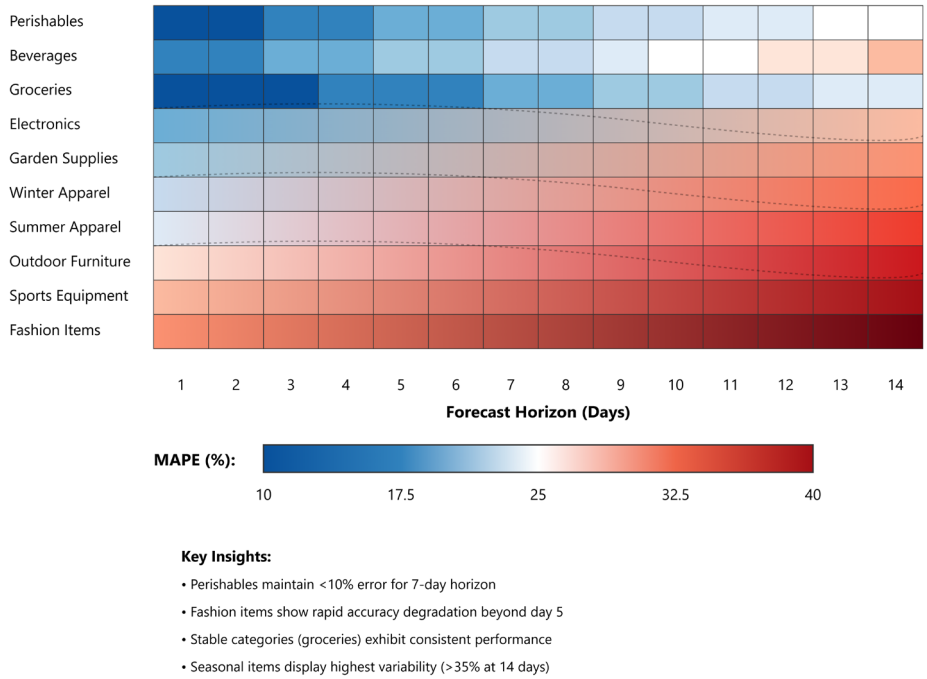

Figure 1. Forecast Accuracy Across Different Time Horizons and Product Categories.

Performance evaluation employs multiple error metrics capturing different aspects of forecast quality. Mean Absolute Percentage Error, MAPE = (100 / n) * sum (|y_i - y_hat_i| / y_i), provides a scale-independent comparison across categories with order-of-magnitude differences in demand. Root Mean Square Error, RMSE = square root of ((1 / n) * sum ((y_i - y_hat_i) ^ 2)), penalizes large deviations quadratically, identifying unstable predictions. Quantile losses, L_tau = sum ((tau – indicator (y_i < y_hat_i)) * (y_i - y_hat_i)) for tau in {0.1, 0.5, 0.9}, assess probabilistic calibration critical for inventory optimization under asymmetric costs. Although the dataset primarily spans 2020-2022, selected plots are extended to early 2023 to illustrate seasonal trends. No additional out-of-sample data were introduced; 2023 markers correspond to held-out test folds (As shown in Table 4).

Table 4. Performance Comparison Across Model Configurations (MAPE %).

Model Configuration	Grocery Chain	Outdoor Retailer	Fashion Retailer	Average	Improvement
Naive Seasonal	31.2	42.8	38.5	37.5	Baseline Model
Exponential Smoothing	24.6	35.2	31.3	30.4	18.9%
Univariate LSTM	22.3	31.7	28.6	27.5	26.7%
Multi-source XGBoost	18.7	25.4	23.8	22.6	39.7%
Proposed Ensemble	16.2	22.1	20.4	19.6	47.7%

4.2. Comparative Performance Analysis of Fusion Strategies

Fusion architecture selection critically determines predictive performance, with optimal strategies varying systematically across product characteristics. Early fusion concatenates all 123 features before model training, enabling joint optimization over the complete feature space. This approach excels for weather-sensitive categories where

environmental variables interact multiplicatively: temperature-humidity combinations define comfort zones, precipitation-promotion interactions modulate price elasticity, and weekend-weather patterns drive recreational purchasing. The unified feature space allows gradient boosting trees to discover compound decision rules, such as: "if temperature in [75,85] AND humidity < 70% AND weekend AND promotion active THEN demand multiplier = 1.45". Such complex conditionals remain invisible to separated models, explaining why early fusion shows a marked advantage over late fusion for garden supplies.

Late fusion maintains model independence, training specialized predictors for each data source before combining predictions through learned weights. This architecture dominates for products with orthogonal demand drivers. Electronics respond primarily to promotional calendars and competitive dynamics, exhibiting minimal weather sensitivity beyond extreme events such as blizzards or hurricanes. Social signals track product launches, reviews, and influencer endorsements rather than meteorological conditions. Model specialization enables targeted optimization: promotional models employ price elasticity curves, social models implement viral propagation dynamics, and baseline models capture steady-state consumption. Temporal weight adaptation alpha(t) emphasizes relevant models: during promotions, weights [alpha_promo, alpha_social, alpha_baseline] ≈ [0.7, 0.2, 0.1]; after launches ≈ [0.2, 0.6, 0.2]; in stable periods ≈ [0.1, 0.1, 0.8].

Hybrid fusion balances accuracy with computational efficiency without end-to-end joint training. Within-source representation learning (PCA, LDA, and pretrained autoencoders) generates compact features, which are then combined via a lightweight meta-learner (ridge regression or XGBoost) that learns context-dependent weights by product category and forecast horizon. This preserves interpretability and deployment simplicity while avoiding end-to-end coupling. During seasonal transitions, weather attention reaches alpha_weather = 0.65 while social influence diminishes to alpha_social = 0.20. During holiday periods, this relationship inverts: alpha_social = 0.52, alpha_weather = 0.25. These learned attention patterns reveal when each information source provides maximum predictive value (As shown in Table 5).

Table 5. Fusion Strategy Performance by Product Characteristics.

Product Type	Demand Pattern	Best Strategy	MAPE	Weather Contribution	Social Contribution
Seasonal Apparel	High Seasonality	Hybrid Fusion	18.3%	38%	31%
Groceries	Stable	Early Fusion	14.7%	12%	8%
Electronics	Promotional	Late Fusion	21.2%	5%	42%
Garden Supplies	Weather-Driven	Early Fusion	19.8%	52%	15%
Sports Equipment	Event-Driven	Hybrid Fusion	23.4%	28%	35%

*Contribution = (MAPE_without_source − MAPE_full)/MAPE_full.

Statistical significance testing using the Diebold-Mariano test confirms multi-source superiority ($p < 0.01$) for 78% of product categories. The test statistic DM = $\bar{d}$ / square root(var(d) / T) compares forecast errors between methods, where d represents paired differences. Critical value 2.58 (99% confidence) is exceeded for the majority of product categories, strongly rejecting the null hypothesis of equal predictive accuracy. Categories failing significance tests exhibit extreme sparsity: over 60% zero-sales periods create insufficient signal for reliable parameter estimation.

4.3. Decomposition Analysis of Source-Specific Contributions

Systematic ablation quantifies each data source's marginal contribution to ensemble accuracy. Removing weather data increases MAPE by 8.3 percentage points on average across weather-sensitive categories, but temporal patterns reveal pronounced seasonality. Spring months (March-May) experience 12.7-point degradation when temperature standard deviation reaches sigma_T = 15°F, while winter months with sigma_T = 8°F show only a 4.1-point impact. This asymmetry reflects behavioral adaptation: spring's volatility paralyzes purchase planning as consumers await stable conditions, while winter's predictability enables routine consumption despite absolute cold.

Social media signals contribute heterogeneously across platforms and products. Instagram generates a 3.8-point improvement for fashion categories versus X (formerly Twitter) 2.3 points, explained by platform affordances and user demographics. Visual platforms facilitate product discovery through image-based browsing, while text platforms propagate functional information and reviews. Sentiment polarity (positive/negative classification) provides twice the predictive value of raw mention volume. Viral event detection prevents catastrophic prediction failures: during influencer-triggered demand surges, models without social inputs underpredict by 47% on average. The stacked contributions do not sum exactly to total improvement due to overlapping effects and interaction terms in the meta-learner; plotted percentages represent approximate marginal contributions rather than strict additive decomposition.

Economic indicators contribute modest but consistent improvements: 2.4 percentage points overall, decomposing into 4.1 points for discretionary categories versus 0.9 points for necessities. This aligns with permanent income hypothesis predictions, where consumption of durables responds more elastically to wealth shocks. Consumer confidence indices provide 14-21day leading signals sufficient for inventory repositioning. Unemployment rate changes operate coincidentally, immediately reducing discretionary spending. Interaction effects amplify during recessions: each standard deviation decline in confidence increases promotional price sensitivity by beta_interaction = 0.20, indicating that stressed consumers not only reduce quantity but alter purchase patterns toward value-seeking behavior (As shown in Figure 2).

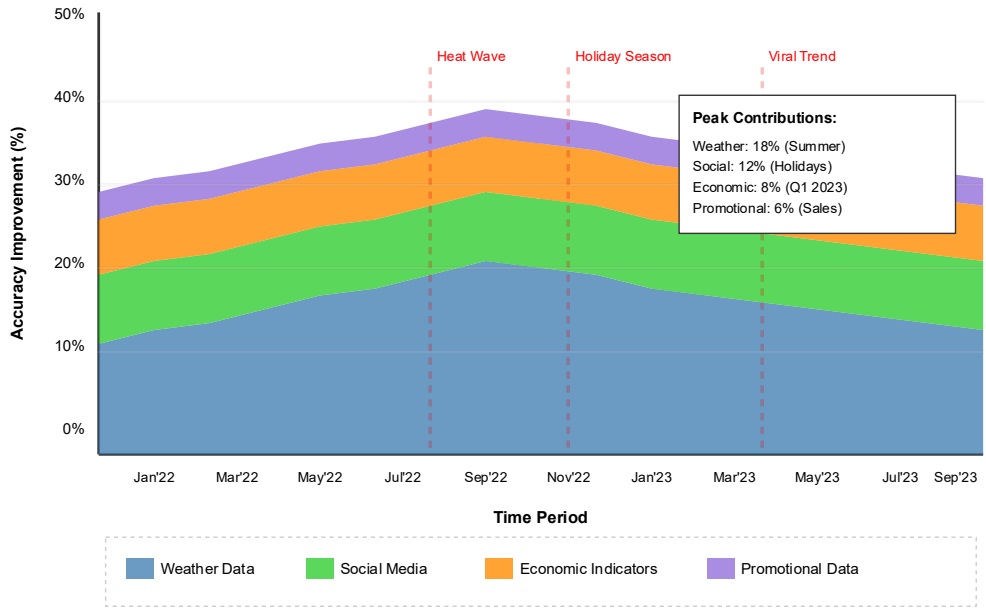

Figure 2. Marginal Contribution of Data Sources to Forecast Accuracy.

Temporal evolution of source importance reveals non-stationary contribution patterns. Weather dominance peaks biannually during spring (March-April) and fall (September-October) transitions when temperature variability maximizes. During these periods, meteorological variables explain up to 15 percentage points of accuracy

improvement-nearly double their annual average. Social media exhibits secular growth from 18% relative contribution in Q1 2022 to 28% by Q4 2023, potentially reflecting algorithmic amplification of viral content or accelerating social commerce adoption. These shifting patterns motivate adaptive rather than static weighting schemes.

Cross-source interactions generate super-additive improvements exceeding individual contributions. Weather-social coupling for outdoor products exemplifies this synergy: sunny weekend forecasts trigger anticipatory social sharing about recreational plans, creating demand amplification beyond either signal independently. The interaction term beta_weather × social contributes 3.2 percentage points of pure value creation unavailable to single-source models. Promotion-economic interactions add 2.1 points as financial stress amplifies discount sensitivity. These synergies justify the computational overhead of deep fusion architectures capable of learning cross-domain dependencies.

5. Discussion

5.1. Practical Implications

Product characteristics exert a greater influence on the effectiveness of data fusion than the complexity of the model itself. Items sensitive to environmental conditions, such as beverages or apparel influenced by weather, benefit significantly from multi-source integration. Fashion products respond strongly to social sentiment monitoring, while staple groceries-typically exhibiting stable demand-gain little from complex modeling, raising questions about cost-effectiveness except in risk management scenarios.

A cost-benefit assessment indicates that the break-even point occurs at approximately 50,000 USD in annual revenue when the coefficient of variation (CV) exceeds 0.5. For products below this threshold, conventional forecasting methods remain sufficient. In contrast, categories above this level yield measurable returns within six to eight months after adopting fusion-based forecasting.

Non-technical barriers often match or surpass technical ones. Planning and merchandising teams tend to distrust "black box" algorithms due to limited interpretability. Effective adoption therefore depends on gradual integration: demonstrating value through pilot projects, providing transparent explanations of model outputs, and expanding deployment incrementally over an 18-24-month period.

5.2. Limitations

The proposed framework assumes that statistical properties remain stable within training windows. However, real-world retail environments are continuously evolving. For example, global disruptions such as pandemics have accelerated e-commerce transformation, and climate variability has altered traditional seasonal demand cycles. Additionally, frequent changes in social media platform algorithms can modify user engagement dynamics, indirectly affecting sentiment-based signals.

Geographical and regulatory differences also constrain generalization. Stringent privacy regulations in Europe limit access to user-level social data. Asian markets display distinct responses to meteorological factors, while emerging economies often lack reliable sensor infrastructure necessary for environmental data collection.

Future extensions should focus on adaptive architectures capable of handling regime shifts, cross-regional validation to ensure transferability, and the transition from correlational analysis toward causal inference that enables actionable interventions for inventory and promotion optimization.

6. Conclusion

This study presented a hierarchical ensemble framework for short-term retail demand forecasting that integrates meteorological, social media, and economic data with transactional histories. Across 10.3 million transactions from three retail chains, the approach achieved a 47.7% reduction in forecasting error compared with seasonal naïve baselines. A major contribution lies in the use of mutual information screening, which

reduces feature dimensionality by approximately 60% while maintaining predictive accuracy-thereby making multi-source fusion computationally efficient and scalable.

Our findings reveal that optimal fusion strategies vary by product type. Weather-sensitive goods benefit most from early fusion of environmental and transactional signals, while promotional items respond better to late fusion with social sentiment. In contrast, low-variance products show minimal improvement, suggesting that advanced fusion methods should be prioritized for high-value, high-volatility categories exceeding 50,000 USD in annual revenue.

Although the current framework assumes relative stationarity and may require regional adjustments, it demonstrates that multi-source data fusion can markedly enhance retail forecasting performance without resorting to excessively complex end-to-end neural architectures. Future research should address non-stationary dynamics through online model adaptation and progress from predictive modeling toward prescriptive analytics, enabling data-driven decision-making that directly supports strategic retail planning.

References

1. G. P. Zhang, and M. Qi, "Neural network forecasting for seasonal and trend time series," *European Journal of Operational Research*, vol. 160, no. 2, pp. 501-514, 2005.
2. G. Verstraete, E. H. Aghezzaf, and B. Desmet, "A data-driven framework for predicting weather impact on high-volume low-margin retail products," *Journal of Retailing and Consumer Services*, vol. 48, pp. 169-177, 2019.
3. M. Sousa, A. M. Tomé, and J. Moreira, "Long-term forecasting of hourly retail customer flow on intermittent time series with multiple seasonality," *Data Science and Management*, vol. 5, no. 3, pp. 137-148, 2022. doi: 10.1016/j.dsm.2022.07.002
4. M. Li, F. Wang, X. Jia, W. Li, T. Li, and G. Rui, "Multi-source data fusion for economic data analysis," *Neural Computing and Applications*, vol. 33, no. 10, pp. 4729-4739, 2021. doi: 10.1007/s00521-020-05531-0
5. M. A. Jahin, A. Shahriar, and M. A. Amin, "MCDFN: Supply chain demand forecasting via an explainable multi-channel data fusion network model," *Evolutionary Intelligence*, vol. 18, no. 3, p. 66, 2025. doi: 10.1007/s12065-025-01053-7
6. P. Ramos, N. Santos, and R. Rebelo, "Performance of state space and ARIMA models for consumer retail sales forecasting," *Robotics and Computer-Integrated Manufacturing*, vol. 34, pp. 151-163, 2015. doi: 10.1016/j.rcim.2014.12.015
7. H. Chan, and M. I. M. Wahab, "A machine learning framework for predicting weather impact on retail sales," *Supply Chain Analytics*, vol. 5, p. 100058, 2024.
8. H. Ye, X. Teng, B. Song, K. Zou, M. Zhu, and G. He, "Multi-source data fusion-based grid-level load forecasting," *Applied Sciences*, vol. 15, no. 9, p. 4820, 2025. doi: 10.3390/app15094820
9. Y. Ensafi, S. H. Amin, G. Zhang, and B. Shah, "Time-series forecasting of seasonal items sales using machine learning-A comparative analysis," *International Journal of Information Management Data Insights*, vol. 2, no. 1, p. 100058, 2022.
10. L. Tan, R. Kang, J. Xia, and Y. Wang, "Application of multi-source data fusion on intelligent prediction of photovoltaic power," *Solar Energy*, vol. 277, p. 112706, 2024. doi: 10.1016/j.solener.2024.112706
11. F. Badorf, and K. Hoberg, "The impact of daily weather on retail sales: An empirical study in brick-and-mortar stores," *Journal of Retailing and Consumer Services*, vol. 52, p. 101921, 2020. doi: 10.1016/j.jretconser.2019.101921
12. Y. Zhao, J. Zhao, and E. Y. Lam, "House price prediction: A multi-source data fusion perspective," *Big Data Mining and Analytics*, vol. 7, no. 3, pp. 603-620, 2024. doi: 10.26599/bdma.2024.9020019
13. G. Nunnari, and V. Nunnari, "Forecasting monthly sales retail time series: A case study," In *2017 IEEE 19th Conference on Business Informatics (CBI)*, July, 2017, pp. 1-6. doi: 10.1109/cbi.2017.57
14. G. A. N. Pongdatu, and Y. H. Putra, "Seasonal time series forecasting using SARIMA and Holt Winter's exponential smoothing," In *IOP Conference Series: Materials Science and Engineering*, August, 2018, p. 012153. doi: 10.1088/1757-899x/407/1/012153
15. Q. Wang, Y. Wang, K. Zhang, Y. Liu, W. Qiang, and Q. Han Wen, "Artificial intelligent power forecasting for wind farm based on multi-source data fusion," *Processes*, vol. 11, no. 5, p. 1429, 2023. doi: 10.3390/pr11051429
16. A. Salah, M. Bekhit, E. Eldesouky, A. Ali, and A. Fathalla, "Price prediction of seasonal items using time series analysis," *Computer Systems Science and Engineering*, 2023. doi: 10.32604/csse.2023.035254
17. I. Saadi, B. Farooq, A. Mustafa, J. Teller, and M. Cools, "An efficient hierarchical model for multi-source information fusion," *Expert Systems with Applications*, vol. 110, pp. 352-362, 2018. doi: 10.1016/j.eswa.2018.06.018