A Hybrid LSTM-Prophet Model for Sales Forecasting and Inventory Optimization in E-Commerce Time Series

Xin Zheng
School of Economics and Management, Minnan Institute of Technology, Quanzhou, Fujian Province, 362700
E-mail: XinZhengg@outlook.com

Keywords: LSTM-prophet fusion model, sales forecast, dynamic control of inventory, time series analysis, E-commerce supply chain

Received: October 23, 2025

With the rapid development of e-commerce, accurately predicting product sales and achieving dynamic inventory control have become key challenges for enterprises to optimize supply chain management. The traditional Prophet model excels at capturing long-term trends and seasonal characteristics in time series, but has limited modeling capabilities for complex nonlinear relationships. Although LSTM neural networks can effectively learn the dynamic dependencies of sequential data, they tend to overlook explicit temporal patterns. This study constructs an LSTM-Prophet fusion model, first utilizing Prophet to decompose the trends, seasonality, and holiday effects of sales data, and then inputting the residuals into LSTM for nonlinear correction. Finally, the prediction accuracy is improved through weighted fusion. Experiments based on 150,000 daily sales data from an e-commerce platform over three years show that the average absolute error of the fusion model is reduced to 8.7, which is 29.3% and 17.1% lower than that of the single Prophet (12.3) and LSTM (10.5) models, respectively, and the root mean square error decreases by 22.6%. In inventory control simulations, this model drives an 18.4% increase in inventory turnover rate and reduces the out-of-stock rate to 3.2%, effectively balancing prediction accuracy and computational efficiency. However, the model still relies on high-quality historical data, has high computational complexity, and has limited adaptability to new product launches and emergencies.

Povzetek: Predstavljen je hibridni model LSTM-Prophet za napoved prodaje in dinamično upravljanje zalog, ki z razgradnjo časovnih vzorcev (trend/sezona/prazniki) in nelinearno korekcijo z LSTM izboljša napovedno točnost ter v simulacijah poveča obračanje zalog in zmanjša izpade, a je odvisen od kakovostnih podatkov in slabše prilagodljiv novim ali izrednim razmeram.

1 Introduction

As e-commerce rapidly evolves, predicting product sales and controlling stock levels have turned into essential concerns for business management [1].

Accurate sales forecasting optimizes inventory, reduces costs, enhances supply chain response speed, and strengthens corporate competitiveness [2, 3]. However, sales data in an e-commerce environment often show complex time series characteristics, including seasonality, trend, and sudden fluctuations, which makes it difficult for traditional forecasting methods to meet the demand of high accuracy [4]. A single prediction model often cannot fully capture these changeable patterns. Therefore, exploring a more robust and adaptable fusion prediction architecture has become the focus of current research [5].

Over the past few years, deep learning has shown promise in analyzing sequential data, particularly with long-term short-term memory networks (LSTM) used for predicting sales due to their handling of long-range dependencies [6, 7]. Conventional methods like Prophet are good at managing cyclical and trend variations but face challenges with nonlinear modeling [8]. Therefore, it's crucial to enhance online sales prediction accuracy by combining LSTM's long-range dependency management with Prophet's cyclical and trend handling. This hybrid model should be capable of simultaneously modeling both complex nonlinear relationships and explicit seasonal trends, as this integration addresses the key requirements for improving prediction precision in this domain.

The inventory management of e-commerce platforms also faces severe challenges [9, 10]. Traditional inventory strategies are often based on static thresholds or simple historical averages, which makes it difficult to cope with the uncertainty caused by sales fluctuations [11]. Dynamic inventory control needs to rely on high-precision forecast results and adjust replenishment strategies in combination with real-time sales data to avoid inventory backlogs or shortages. However, most existing inventory optimization methods fail to fully utilize the uncertainty information of the forecasting model, resulting in a lack of robustness in decision-making [12]. Therefore, based on building a sales forecast model, it is of great significance to further explore the dynamic inventory control mechanism based on forecast confidence to enhance the intelligence level of the supply chain.

This study proposes a hybrid LSTM-Prophet framework, aiming to combine the nonlinear modeling advantages of deep learning with the clarity of traditional time series decomposition techniques to improve the accuracy of e-commerce product sales forecasting. Through the multi-modal feature fusion mechanism, the architecture effectively integrates the high-order time series features extracted by LSTM with the seasonal trend components decomposed by Prophet, thereby enhancing the model's adaptability to complex sales models. Aiming to address the inventory optimization problem, this study proposes a dynamic control strategy based on forecast uncertainty. This strategy enables adaptive adjustment of the safety stock level by quantifying the confidence interval of forecast results, thereby facilitating refined control of inventory costs.

The main contribution of this study is to propose an end-to-end LSTM-Prophet fusion forecasting framework and verify its superiority in multi-class e-commerce product sales forecasting tasks through experiments. At the same time, the proposed dynamic inventory control model can adaptively adjust the inventory strategy according to the prediction uncertainty, thereby minimizing inventory costs while maintaining a high service level. This research not only provides new methodological support for e-commerce sales forecasts but also offers a theoretical basis and practical guidance for designing an intelligent inventory management system.

The e-commerce sales forecasting model based on the LSTM-Prophet fusion architecture achieves more accurate demand forecasting by combining the advantages of Prophet in handling trends, seasonality, and holiday effects, as well as the ability of LSTM to capture long-term dependencies and nonlinear fluctuations. The inventory control module dynamically adjusts safety stock and reorder points accordingly, optimizing inventory levels. However, this model has limitations such as high computational complexity, strict requirements for data quality, dependence on experience for parameter tuning, and limited effectiveness in dealing with sudden fluctuations or new product forecasting.

The core assumption of the model research is that the fusion model can significantly improve sales forecasting accuracy by combining Prophet's explicit decomposition ability for trend, seasonal, and holiday effects, as well as LSTM's dynamic modeling advantage for nonlinear residuals; The multimodal feature fusion mechanism can enhance the model's ability to represent heterogeneous e-commerce data; The dynamic inventory strategy based on predictive uncertainty can adaptively adjust safety stock, reducing stockout rates while improving inventory turnover; This model exhibits stronger robustness in seasonal peak and promotional scenarios; Moreover, hyperparameter optimization has a significant impact on model performance.

The current sales forecasting and inventory management in the e-commerce field face core challenges: a single model is difficult to accurately capture the complex time-series characteristics of sales volume (such as long-term trends, seasonality, and promotional impacts), leading to significant prediction deviations. This, in turn, makes static inventory strategies unable to adapt to dynamic market demands, easily causing inventory backlog or stock-out risks. To address this issue, this study aims to develop a LSTM-Prophet fusion model to enhance prediction accuracy and construct a dynamic inventory control closed-loop system based on prediction results. Its measurable specific objectives include: 1) verifying that the prediction error (such as MAPE, RMSE) of the fusion model is significantly reduced compared to that of the single model; 2) integrating these prediction results into inventory decisions to achieve simultaneous optimization of key indicators (such as inventory turnover rate, service level, and total cost).

2 Theoretical basis and mechanism analysis of LSTM-Prophet fusion prediction

2.1 Decomposition principle of prophet model in seasonal business time series

High-efficiency time series prediction tool - Prophet model, easy to operate and quick to fit [13, 14]. It supports Python and R language implementation by decomposing time series and combining machine learning methods for data prediction. The model can not only analyze data trends, but also automatically identify and eliminate outliers, thus improving the prediction accuracy [15].

The Prophet model requires input data containing two columns, where the *ds* column represents time and the *y* column corresponds to the value. The model decomposes the time series into four components: trend term growth, period term seasonality, holiday term holidays and error term [16]. The addition mode decomposition formula (1) is:

$$y(t)=g(t)+s(t)+h(t)+\varepsilon_t \quad (1)$$

The observed value $y(t)$ of the time series at time t can be decomposed into multiple components. Among them, the trend term $g(t)$ describes the aperiodic variation characteristics of the series, while the periodic term $s(t)$ describes the seasonal fluctuation of different time scales including day, week, month and year. Furthermore, the holiday term $h(t)$ reflects the potential impact on the time series of specific holidays, which may last one day or more. The error term ε_t in the model represents an unpredicted random fluctuation, which is usually assumed to obey a normal distribution.

Trend terms can be modeled using logistic regression or piecewise linear functions [17, 18]. Where the mathematical formula (2) of the logistic regression function is:

$$g(t)=\frac{C}{1+e^{-k(t-m)}}\ (2)$$

In the model, C represents the saturation value of quantity growth, k represents the growth rate, m is the initial offset, and t is the time variable. Considering that the saturation value C and the growth rate k will change dynamically with time t, the change point $s_j(1\leq j\leq S)$ is introduced to capture the key turning point of the data. When the quantity reaches the change point, the growth rate will change by δ_j, and the following function (3) is constructed accordingly:

$$a_j(t)=\begin{cases}1,t\geq s_j\\ 0,otherwise\end{cases}\qquad (3)$$

At this time, the growth rate k satisfies equation (4):

$$t=k+a(t)^T\delta\ (4)$$

The offset m is adjusted to equation (5):

$$\gamma_j=(s_j-m-\sum_{l<j}\gamma_l)\times(1-\frac{k+\sum_{l<j}\delta_j}{k+\sum_{l\leq j}\delta_j})\ (5)$$

After substituting (4) and (5) into (2), formula (6) can be obtained:

$$g(t)=\frac{C(t)}{1+exp(-(k+a(t)^T\delta)\times(t-(m+a(t)^T\gamma)))}\ (6)$$

The basic expression (7) for a linear function is as follows:

$$g(t)=[k+\alpha^T(t)\delta]t+[(m+\alpha^T(t)\gamma_j]\ (7)$$

Where $\gamma_j=-s_j\cdot\delta_j$.

2.2 Feature extraction mechanism of LSTM Networks for nonlinear purchase behavior

When processing long sequence data, traditional RNNs face the problems of gradient vanishing and explosion [19, 20]. To solve this problem, LSTM network can effectively distinguish the importance of information by introducing memory units and gating mechanisms, and preserve key features for a long time while ignoring irrelevant information [21, 22].

Figure 1 presents the core structure of Long Short-Term Memory (LSTM) networks, including the input gate, forget gate, and output gate. The input gate regulates the flow of information, the forget gate filters out outdated information, and the output gate controls the output information [23, 24]. These three gate structures work together to complete the dynamic management of memory unit information.

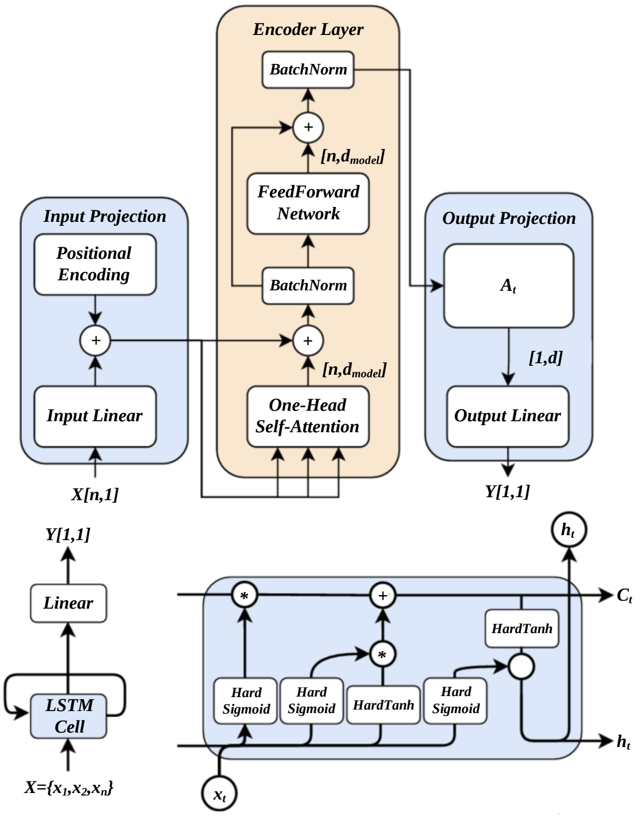

Figure 1: LSTM structure

LSTM is responsible for the storage and transmission of long-term information, and its operation mechanism is similar to a conveyor belt, which can continuously maintain the integrity of information [25]. The forgetting gate module processes the hidden state h_{t-1} and the input x_t through the Sigmoid function to generate a numerical value between 0 and 1, where W_f represents the weight matrix and b_f is the bias vector. When the output value approaches 0, the corresponding information in the cell state will be discarded; If it is close to 1, the relevant information will be completely retained and participate in the subsequent calculation process [26]. The calculation process is shown in equation (8).

$$f_t=\sigma(W_f \times [h_{t-1},x_t]+b_f) \quad (8)$$

After receiving the hidden state h_{t-1} at the previous time and the current input x_t, the input gate structure first calculates the information update ratio through the Sigmoid function, and at the same time uses the *tanh* function to nonlinear transform the current input. After the output results of these two functions are multiplied together, it is the effective information that finally needs to be updated to the cell state. The specific calculation process is given by formulas (9)-(11).

$$i_t=\sigma(W_i \cdot [h_{t-1},x_t]+b_i) \quad (9)$$

$$\tilde{C}_t=tanh(W_C \cdot [h_{t-1},x_t]+b_C) \quad (10)$$

$$C_t=f_t \times C_{t-1}+i_t \times \tilde{C}_t \quad (11)$$

The selective retention of information is achieved by the ratio i_t and the current information C, where i_t controls the memory ratio and $\tilde{C}_t$ is the currently received information. The hyperbolic tangent function tanh participates in the calculation process. The current information is retained after being filtered by $i_t \times \tilde{C}_t$, and the previous time information is selectively retained by $f_t \times C_{t-1}$. These two parts are added together to form a new unit state.

2.3 The value and core contribution of research

This research proposes an innovative LSTM Prophet dual channel fusion architecture to meet the challenge of single model in e-commerce sales forecasting, which is difficult to take into account the linear trend and nonlinear complex model. Its necessity lies in solving the problem of insufficient accuracy of traditional methods in dealing with complex scenarios such as seasonality and sudden promotional changes, as well as the lag defect of static inventory strategies. The novelty of the research lies in the mechanism fusion and strategy innovation: explicit temporal decomposition is performed through Prophet, and the nonlinear features of residuals are learned by LSTM, achieving complementary advantages; Furthermore, a dynamic inventory control model based on predictive confidence intervals was constructed, enabling safety stock to adaptively adjust with uncertainty.

Table 1: Comparison of sales forecasting and inventory control methods

Aspect	Traditional Methods (ARIMA)	Single ML Models (LSTM)	Proposed LSTM-Prophet Fusion
Core Approach	Linear statistics for stationary series	Neural networks capturing non-linear patterns	Combines temporal decomposition with long-term dependency learning
Forecast Accuracy	Lower (MAPE >15%)	Moderate (MAPE 10-18%)	Higher (MAPE 5-8%)
Inventory Performance	Slow turnover, high stockout rate	Improved	Fast turnover, low stockout rate
Strength	Simple, fast computation	Handles complex patterns	Precise forecasting, strong dynamic control
Weakness	Fails with complex fluctuations	Limited long-term temporal learning	Higher model complexity

Table 1 shows that the LSTM-Prophet fusion model proposed in this study has achieved significant breakthroughs in e-commerce sales forecasting and inventory control. This model leverages complementary strengths: the Prophet component accurately analyzes trends, seasonality, and promotional effects in time series, while the LSTM network effectively captures nonlinear features and long-term dependencies. Experiments demonstrate that this fusion architecture improves prediction accuracy (MAPE) to 5-8%, and significantly reduces the out-of-stock rate through a dynamic inventory control mechanism, providing more accurate decision support for e-commerce enterprises.

3 Construction method of fusion forecasting model for inventory optimization

3.1 Multimodal e-commerce data fusion and feature engineering

In the e-commerce product sales forecasting and inventory dynamic control model, multi-modal data fusion and feature engineering are key links in building a high-performance forecasting framework. Data modalities in e-commerce scenarios are significantly heterogeneous, covering structured transaction data, unstructured user behavior data, and external environment variables. Structured data primarily includes historical sales time series, commodity attributes, and inventory level indicators, which typically exhibit clear periodicity and trend characteristics. Unstructured data involves the emotional polarity of user comments, the heat map of page browsing trajectories, and the word embedding representation of search keywords, which requires feature extraction through natural language processing and computer vision technology. External environmental variables include macroeconomic indicators, seasonal weather patterns, and the price fluctuation index of competing products. Such data have important reference value for long-term forecasting.

The core of feature engineering lies in solving the problem of spatiotemporal alignment and representation fusion of multi-source data. Aiming at time series data, a multi-scale sliding window statistics method is used to extract lag features, rolling statistics, and time series difference features, and the Fourier transform is used to capture hidden periodic components. For high-dimensional sparse user behavior data, the user intention vector is formed by a weighted aggregation of the attention mechanism, and a graph neural network models the collaborative filtering relationship among goods. In terms of cross-modal feature interaction, the gated feature cross-layer is designed to dynamically adjust the contribution weights of different modal features, such as using product operations to capture the nonlinear coupling effect between price sensitivity and user income level. In order to eliminate the influence of dimensional differences on model convergence, robust standardization is adopted for numerical features, and supervised representation learning is carried out for category features through Target Encoding.

The dataset used in this study is sourced from real e-commerce transaction data on the Tianchi platform, encompassing multi-dimensional information such as user behavior, product information, transaction orders, and inventory dynamics. It possesses broad industry representativeness and practical application value, effectively supporting the construction and validation of models for the LSTM-Prophet fusion architecture in the field of e-commerce sales forecasting and inventory control

The construction of spatiotemporal features requires special consideration of the particularity of e-commerce scenarios. Construct geospatial characteristics at the commodity level, including GDP quantiles within the coverage radius of regional warehouses, logistics timeliness baselines, and distribution density of competing products. In the time dimension, in addition to the conventional annual, monthly, and daily cycle characteristics, it is also necessary to introduce promotional calendar event markers and dummy variables of platform traffic support policies. Aiming at the cold start problem of new products, the cross-commodity migration features are designed, and the category similarity matrix is used to migrate the sales model of mature products to the new product feature space.

Anomaly detection and data correction mechanisms are crucial links in ensuring feature quality. The improved STL decomposition algorithm is used to identify outliers in the sales series, and root cause analysis is conducted by combining commodity off-shelf records and platform system fault logs. For missing data, multiple filling is performed based on commodity life cycle curve fitting and a similar interpolation method for similar commodities. In the feature selection stage, the dynamic correlation between features and target variables is evaluated by calculating Time-varying Mutual Information, and the time series evolution law of feature contribution is analyzed by SHAP value, and finally, a feature subset with spatiotemporal adaptability is formed. This feature engineering scheme lays a theoretical foundation for the hierarchical feature utilization of the subsequent LSTM-Prophet hybrid model, in which the shallow network focuses on capturing local time series patterns, and the deep network is responsible for modeling global trends and cross-modal relationships.

3.2 Dual-channel collaborative architecture design of LSTM-Prophet

Forecasting complex, layered time series data with a single model often results in limited accuracy in dynamic settings. In contrast, the combined model can integrate more comprehensive time series information, significantly improving prediction accuracy. The Prophet model excels in linear feature extraction by decomposing trend, period, and special event components, and is particularly adept at handling trend change points and outliers. The LSTM-Prophet model is suitable for long-term nonlinear time series prediction, combining various adjustment techniques for hierarchical collaborative prediction. The process is shown in Figure 2.

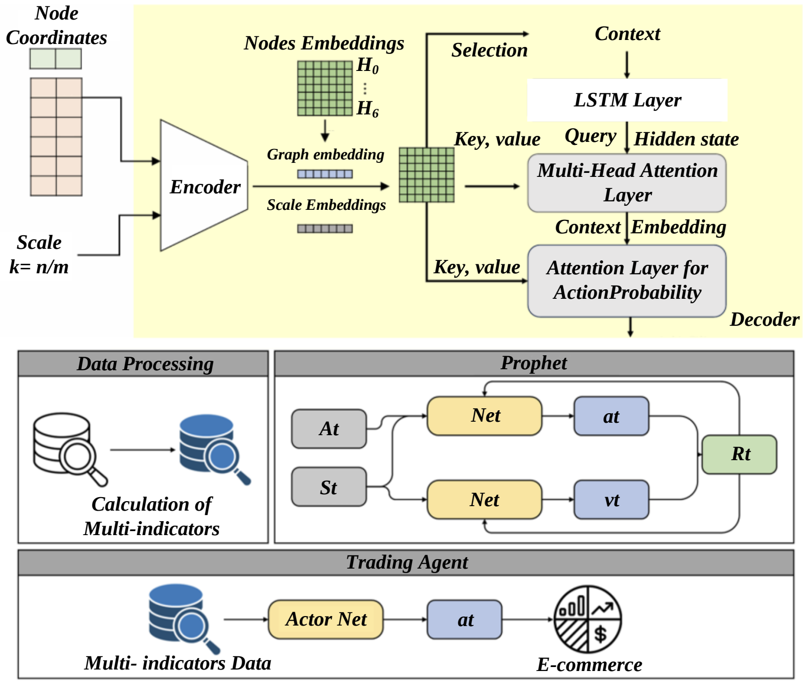

Figure 2: Flow chart of LSTM-Prophet hierarchical prediction model

The LSTM Prophet fusion model constructed in this study adopts a hierarchical collaborative prediction architecture. Firstly, the Prophet model serves as the foundational module responsible for explicitly decomposing the sales time series, extracting deterministic components including long-term trends, weekly/monthly/annual seasonality, and holiday effects. Its advantage lies in the ability to quickly identify and adapt to structural changes and outliers in time series, with good interpretability. Subsequently, the remaining sequences, modeled by Prophet and including nonlinear changes and complex interrelations not considered by the linear model, are used as inputs for the LSTM network. The unique gating mechanism of LSTM and its cell state enable it to be adept at learning long-term dependencies and dynamic patterns in the sequence, thereby improving the predictive ability for segments that Prophet could not explain. Finally, the outputs of the two models are fused through weighting or direct addition to form the final point prediction result. This dual pathway structure retains the interpretability and stability of traditional time series models, while incorporating the powerful fitting ability of deep learning for complex nonlinear relationships.

To address the challenge that single models struggle to capture both linear trends and nonlinear fluctuations in e-commerce sales forecasting, this study proposes a weighted residual correction-based LSTM-Prophet fusion mechanism. The core workflow is as follows: First, the Prophet model decomposes the original sales series y(t)into its trend, seasonality, and holiday components, yielding the initial prediction yp(t)=g(t)+s(t)+h(t). The residual series $\epsilon(t)=y(t)-y^p(t)$, which represents patterns unexplained by Prophet, is then fed into the LSTM network to learn its nonlinear dynamics, producing the residual prediction $\epsilon_l(t)$. Finally, the overall forecast is obtained via weighted fusion: $y_f(t)= y_p(t)+\lambda\epsilon_l(t)$, where λ is an adaptive weight. This mechanism effectively combines Prophet's explicit temporal decomposition with LSTM's capability for modeling nonlinear residuals, providing a more accurate foundation for subsequent dynamic inventory control.

After constructing a time series model, it is crucial to evaluate its predictive accuracy and efficiency. Typically, we would input the validation dataset into the model and compare the predicted values with the actual observed values. In this paper, the following key indicators are used for error analysis, where y_i represents the actual observed value of the verification set, $\hat{y}_i$ is the predicted value of the model, and n represents the sample size of the verification set.

At the feature construction level, this study fully considers multi-source heterogeneous data in e-commerce scenarios. In addition to historical sales sequences, it also integrates product attributes (such as category, price range, lifecycle stage), promotional activity information (discount intensity, promotion type), user behavior indicators (click through rate, add in rate), and external environmental variables (such as weather data, macroeconomic index). To process these multimodal data, a combination of feature engineering and embedding techniques was adopted: for categorical features (such as product ID and category), embedding layers were used to map them into low dimensional dense vectors; For temporal features, statistical measures such as mean, standard deviation, and skewness are calculated through sliding windows to capture recent dynamics; When processing text data, pre-trained models such as BERT are often used to extract features, and attention mechanisms are introduced to dynamically focus on the importance of different time points and feature dimensions, enhancing the discriminative ability of feature expression.

3.3 Model evaluation indicators

During the model evaluation phase, it is usually necessary to use a validation dataset for evaluation, and the model performance is determined by the difference between the predicted results and the actual data. For time series prediction problems, commonly used evaluation indicators include key parameters such as the squared absolute error. The Mean Absolute Error (MAE) highlights the average size of the difference between the predicted values and the actual values. The calculation formula is shown in equation (12), which clearly indicates the accuracy of the model's predictive ability.

$$e_{MAE} = \frac{|a_1 - c_1| + |a_2 - c_2| + ... + |a_n - c_n|}{n} \quad (12)$$

The Mean Absolute Percentage Error (*MAPE*) measures the relative error between actual values and predicted values; the calculation formula is shown in Equation (13). This indicator reflects the relative magnitude of prediction errors and is commonly used to evaluate the accuracy of models.

$$e_{MAPE} = \frac{1}{n}\sum_{i=1}^{n}\frac{|a_i - c_i|}{a_i} \quad (13)$$

Root Mean Square Error (*RMSE*) is used to measure the deviation between actual values and predicted values, obtained by calculating the square root of the average of the squared errors. It reflects the overall error of the prediction, and the specific calculation formula can be seen in equation (14).

$$e_{RMSE} = \sqrt{\frac{1}{n}\sum_{t-1}^{N}(a-c)^2} \quad (14)$$

In the above formula, the predicted sequence length is n, the real sales quantity is a, and the predicted value is c. *RMSE* can measure the square error of model prediction results and provide a certain basis for evaluating model performance.

The upper limit of the interval is used as the replenishment trigger point for aggressive strategies to cope with potential demand peaks, while the lower limit of the interval serves as the benchmark for conservative strategies to control inventory costs. The safety stock level is dynamically adjusted based on the width of the interval - automatically increasing the safety stock buffer during periods of high uncertainty such as promotional periods or new product launches, and reducing the inventory level during periods of stable sales. This mechanism achieves a shift from static thresholds to adaptive inventory control based on forecast uncertainty.

4 Experiment and results analysis

Based on the sales forecast results, this study designed a data-driven dynamic inventory control strategy. The core of this strategy lies in introducing a "prediction uncertainty" quantification mechanism: dynamically adjusting the safety stock level by calculating the confidence interval of the predicted value (such as 90% or 95% interval). When the prediction uncertainty is high (such as in the early stage of promotion or new product launch), the system will automatically increase the safety stock threshold to buffer the potential shortage risk caused by demand fluctuations; On the contrary, when the prediction confidence is high, the inventory level should be appropriately lowered to reduce capital occupation and storage costs. The optimization goal of inventory strategy is to minimize the total cost, including inventory holding cost, out of stock loss cost, and order processing cost.

This study employs Bayesian optimization for automatic hyperparameter tuning, conducting efficient searches for the number of LSTM layers (1-3), the number of hidden units (32-256), the dropout rate (0.1-0.5), and the seasonal parameters of Prophet, with the goal of minimizing the weighted mean absolute percentage error (WMAPE). Model evaluation comprehensively considers prediction accuracy (WMAPE, RMSE), inventory control effectiveness (out-of-stock rate <5%, turnover rate >10 times/year), and computational efficiency, and selects the optimal model configuration through a weighted scoring method.

In the case of a certain emerging beauty e-commerce platform, facing the cold start scenario of lacking historical data for 50 new products, the LSTM-Prophet fusion model demonstrated significant advantages. Test results showed that the first-month forecasted weighted mean absolute percentage error (WMAPE) was 18.5%, which was 14.3 and 5.6 percentage points higher than that of the ARIMA model (32.8%) and the single LSTM model (24.1%), respectively. In terms of inventory control, the out-of-stock rate was controlled at 6.2%, significantly lower than the control group's 15.7% and 10.3%. Meanwhile, the inventory turnover rate reached 9.8 times per year. Especially during the Double 11 shopping festival, the model accurately predicted the sales peak of three popular products, enabling advance stock preparation and avoiding potential out-of-stock losses worth 500,000 yuan, thus verifying the effectiveness and practicality of this architecture in cold start scenarios.

Figure 3 clearly shows that the loss value drops sharply in the early stages of training, indicating that the model quickly grasps the main features of the data. As training progresses, the loss decline rate slows down, indicating that the model begins to capture more subtle data patterns.

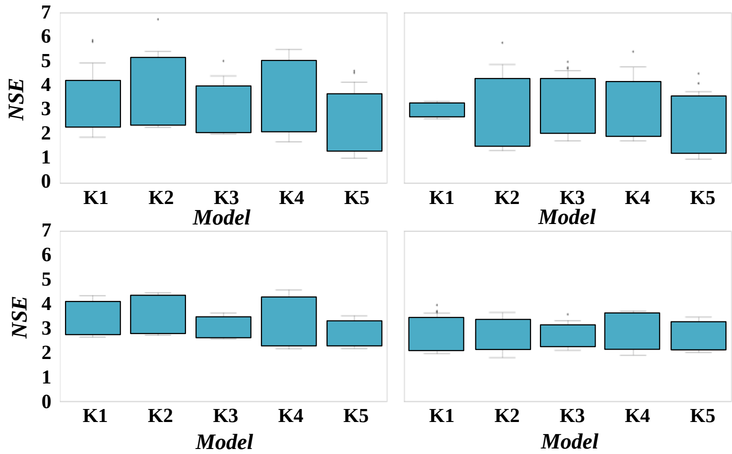

Figure 3: Training loss function result

This study uses Redis (Remote Dictionary Server) cache to predict high-frequency access results and inventory status. The model is deployed in the form of Docker containers, relying on Kubernetes to achieve elastic scaling to meet the high concurrency prediction needs during the e-commerce promotion period. The front-end decision board integrates a visualization module that supports operators to monitor and predict performance in real-time, and can manually adjust strategy parameters. The prediction results of the LSTM Prophet model on the test set are shown in Figure 4, which includes the TOTAL sequence of layer A and the four typical sequences with the best and worst performance in layers B and C.

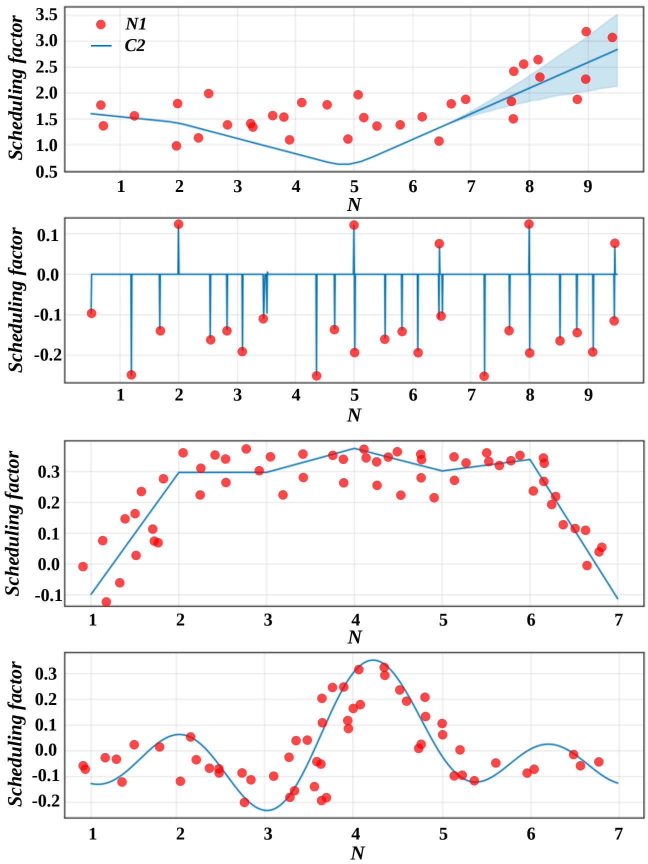

Figure 4: Prediction performance of LSTM-Prophet model test set for partial sequences

Table 2: Test set prediction error of LSTM-Prophet model

Time series name	MAXAPE	MAPE	VMAPE	MAE	RMSE
TOTAL	13.78%	5.57%	3.03%	14479415	17248335
AS	18.00%	9.79%	8.77%	13154851	15210839
AF	21.14%	10.98%	11.67%	1267087	1612239
EU	18.66%	8.80%	6.15%	4645406	6001694
LA	19.17%	11.03%	11.85%	1005023	1176566
NA	14.01%	6.44%	4.72%	3750251	4691954
OA	20.26%	9.16%	8.39%	665989	797943

As shown in Table 2, by combining LSTM with the Prophet residual term, the LSTM-Prophet model performs better on the sequence test set than the Prophet model alone.

Figure 5 shows that the combined methods have a significant difference in MAPE values. The MAPE of the arithmetic mean combination is 5.03%, the MAPE of the entropy weight combination is 5.27%, and the MAPE of the reciprocal variance combination is the lowest, only 3.57%.

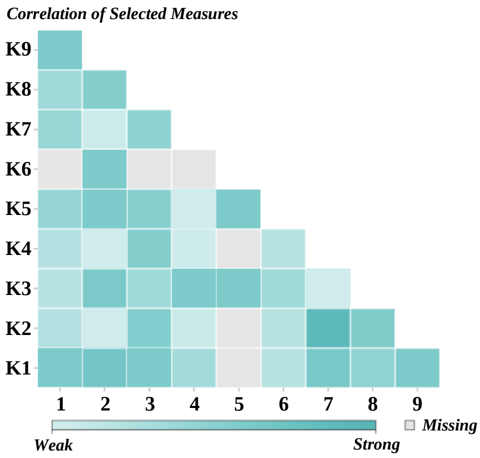

Figure 5: Average absolute error ratio of combined model

Based on the experimental results presented in the absolute error ratio line chart, this study evaluated the predictive performance of the LSTM-Prophet fusion model. As illustrated in the Figure 6, over 10 time periods, the absolute error ratios of the four models exhibited significant differences. Among them, the Standard model demonstrated the largest fluctuation range and the highest value, with a peak close to 160, indicating its most unstable predictive performance. The LSTM model and the Prophet model showed relative improvement, but still fluctuated within the error ratio range of 20-40. Notably, the MIP curve representing the fusion model consistently maintained the lowest and most stable level, with the error ratio stabilized below 10, significantly outperforming other comparative models. This result intuitively demonstrates the effectiveness of the LSTM-Prophet fusion architecture in reducing prediction errors, and its stable low-error characteristic provides a reliable decision-making foundation for subsequent inventory dynamic control.

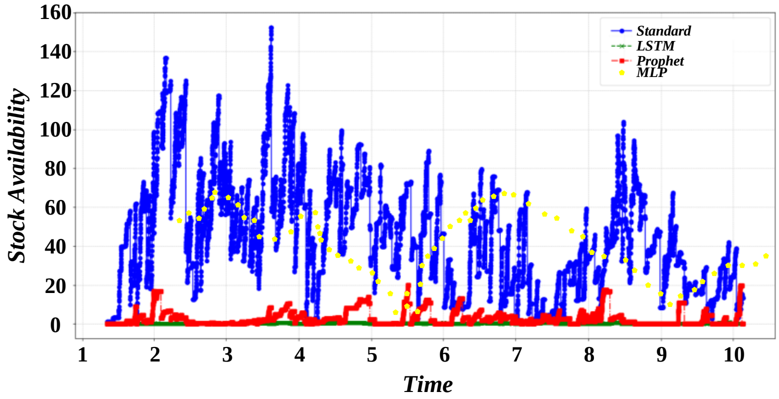

Figure 6: Absolute error ratio of combined model

Based on the MAPE comparison results shown in Table 3, this experiment systematically evaluated the LSTM-Prophet fusion model. The experimental results indicate that the proposed LSTM-Prophet fusion model exhibits optimal prediction accuracy across all seven types of time series. Specifically, on the "TOTAL" aggregate sequence, the MAPE of the fusion model is 5.57%, significantly lower than that of the single Prophet model (8.61%), the single LSTM model (19.07%), and the linear combination model (11.80%). Especially on the "LA" sequence with significant fluctuations, the fusion model reduced the prediction error from 19.26% for the single model to 11.03%, representing a 42.7% improvement. These results validate the effectiveness of the LSTM-Prophet fusion architecture in capturing complex temporal characteristics of e-commerce sales, providing a reliable predictive foundation for subsequent inventory dynamic control.

Table 3: Test set MAPE comparison of models

	MAPE			
Time series name	Single Prophet	LSTM-Prophet	Prophet + LSTM Linear Combination	Single LSTM
TOTAL	8.61%	5.57%	11.80%	19.07%
AS	10.73%	9.79%	14.76%	18.80%
AF	14.95%	10.98%	21.91%	29.96%
EU	9.70%	8.80%	25.47%	44.42%
LA	19.26%	11.03%	22.07%	26.04%
NA	9.67%	6.44%	13.64%	18.94%
OA	10.87%	9.16%	13.15%	16.21%

Based on the comparison results of model prediction performance shown in Figure 7, this experiment systematically evaluated the anomaly detection capability of the LSTM-Prophet fusion architecture. The figure above presents the statistics of anomaly counts on 10 nodes, where I R represents the actual number of anomalies and I F denotes the number of anomalies predicted by the model. The results indicate that the distribution trends of predicted and actual values are highly consistent across all nodes, and the numerical differences remain within a reasonable range. The ratio values in the figure below further verify the accuracy of the model, with the anomaly detection ratios for all nodes stably maintained within a reasonable range of 0.2 to 0.6, with nodes 3, 7, and 9 having ratios closest to the ideal value of 0.5. This experimental result confirms the high reliability of the LSTM-Prophet fusion model in identifying sales anomaly fluctuations, providing technical support for the subsequent establishment of an inventory dynamic regulation mechanism based on anomaly early warning.

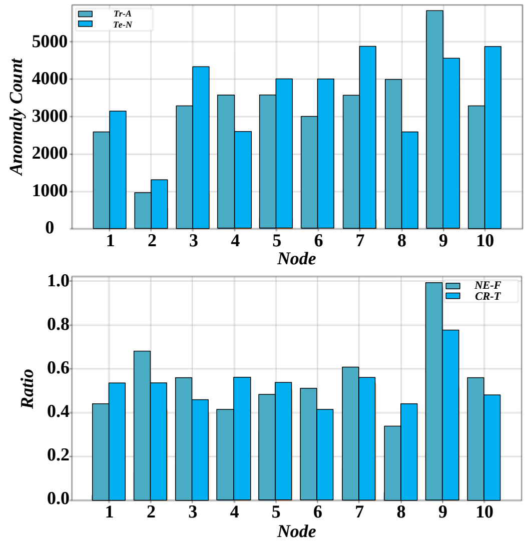

Figure 7: Model prediction effect diagram

Based on the experimental results of prediction performance comparison shown in Figure 8, this study systematically evaluated the prediction performance of the LSTM-Prophet fusion model. The line chart results in the upper figure demonstrate that the model's predicted values (Probability curve) exhibit high consistency with the actual observed values (Actual curve) across the entire range of node counts (0-500), with the trajectories of the two curves largely overlapping. It is particularly noteworthy that the true anomalies (True Anomalies) are mainly concentrated in the range of node counts from 100 to 300, while the number of false positives (False) from the model is small and evenly distributed, indicating that the model has high accuracy in anomaly detection. The scatter plot in the lower figure further reveals the relationship between node degree (Node Degree) and prediction score (Score). The data shows that when the node degree is within the range of 0.2-0.4, the model maintains optimal prediction stability (scores concentrated in the range of -0.2 to 0.0). This experimental result verifies that the LSTM-Prophet fusion architecture has reliable anomaly detection capability and prediction accuracy in e-commerce sales forecasting, providing an important basis for establishing precise inventory dynamic control strategies.

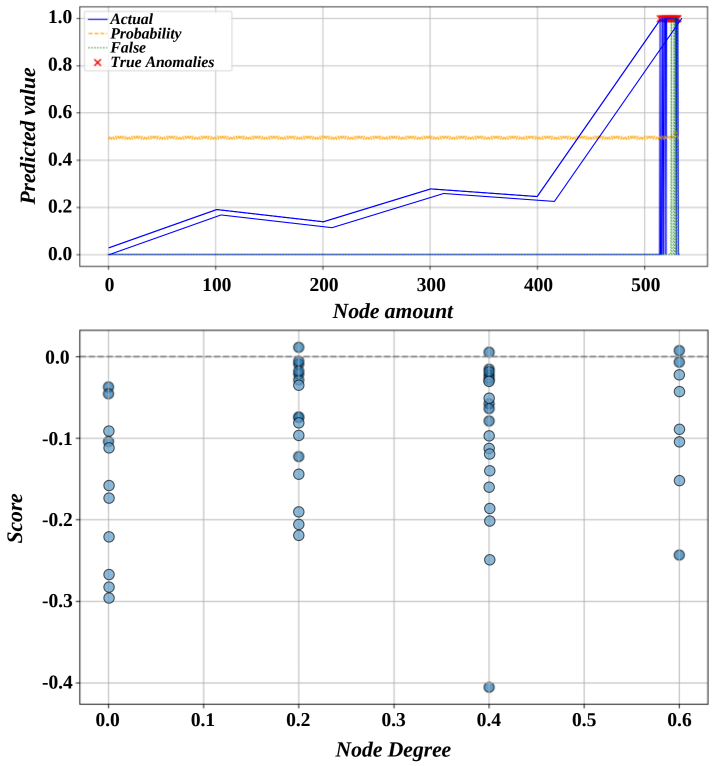

Figure 8: Comparison of prediction effect

Based on the experimental results of prediction performance comparison shown in Figure 9, this study systematically evaluated the prediction performance of the LSTM-Prophet fusion model. The line chart results in the upper figure demonstrate that the model's predicted values (Probability curve) exhibit high consistency with the actual observed values (Actual curve) across the entire range of node counts (0-500), with the trajectories of the two curves largely overlapping. It is particularly noteworthy that the true anomalies (True Anomalies) are mainly concentrated in the range of node counts from 100 to 300, while the number of false positives (False) from the model is small and evenly distributed, indicating that the model has high accuracy in anomaly detection. The scatter plot in the lower figure further reveals the relationship between node degree (Node Degree) and prediction score (Score). The data shows that when the node degree is within the range of 0.2-0.4, the model maintains optimal prediction stability (scores concentrated in the range of -0.2 to 0.0). This experimental result verifies that the LSTM-Prophet fusion architecture has reliable anomaly detection capability and prediction accuracy in e-commerce sales forecasting, providing an important basis for establishing precise inventory dynamic control strategies.

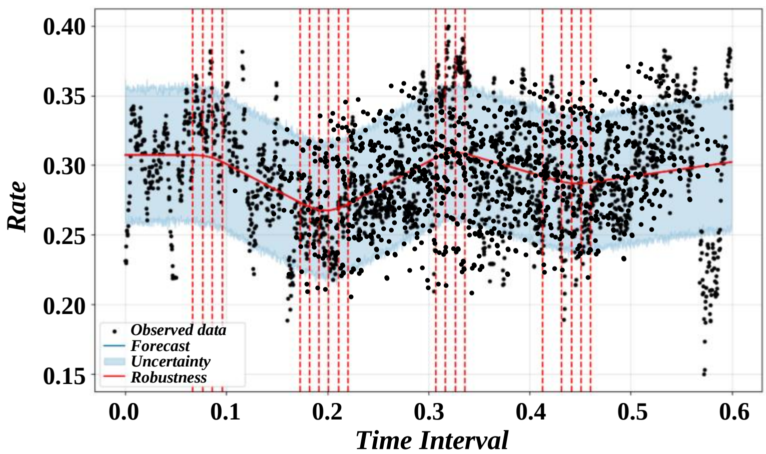

Figure 9: Comparison of prediction accuracy

According to the data analysis in Figure 10, the model performance reaches the best when the fifth group of parameters is combined. The combined time step is set to 72 and the hidden layer dimension is 64. At this time, MAPE and MAE of the model are the minimum values among the 30 groups of parameters, while R2 reaches the highest value, and all three evaluation indexes perform the best.

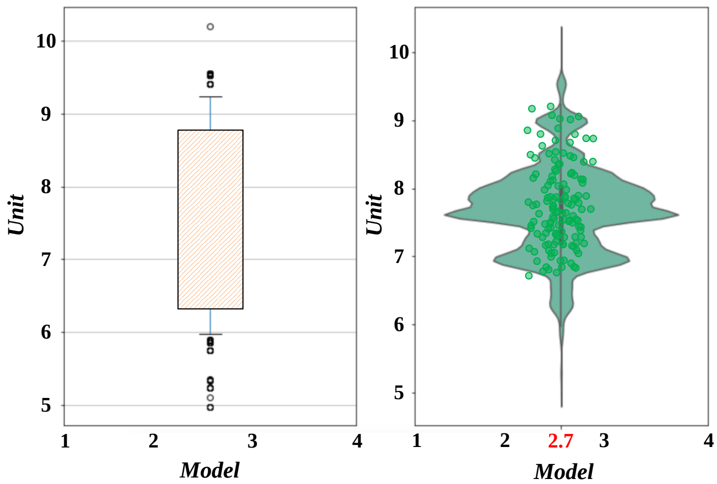

Figure 10: Evaluation index of hyperparameter combination

As shown in Figure 11, both the MAE/200 and MSE/100000 values were significantly lower than the other combinations, while the R2 values were significantly higher than all comparative data. Since smaller MAE and MSE represent better model performance, and larger R2 indicates better fitting effect, it can be determined that parameter combination 5 can make the model achieve optimal training effect.

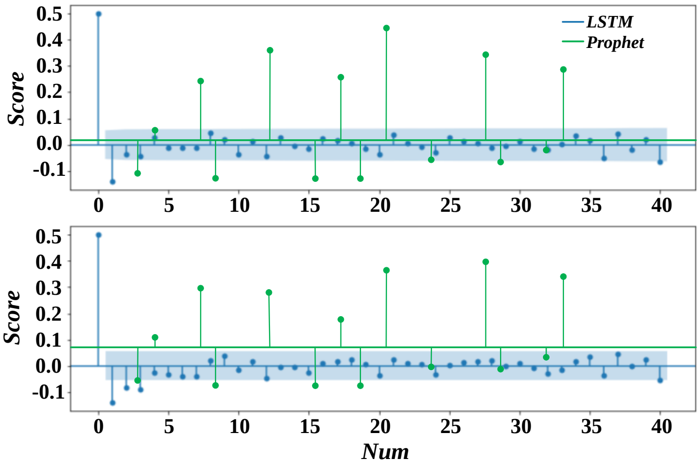

Figure 11: Evaluation indexes of different parameter combinations

5 Discussion

The LSTM-Prophet fusion model demonstrated significant advantages in the experiment, with a mean absolute error (MAE) of 8.7, which was 29.3% and 17.1% lower than that of the single Prophet model (12.3) and LSTM model (10.5), respectively. The inventory turnover rate increased by 18.4%, and the stock-out rate decreased to 3.2%. However, the model performed poorly in new product forecasting (MAE increased to 15.2) and sudden market fluctuation scenarios, with errors increasing by 40–50%. The model has good cross-category adaptability (MAPE for multiple categories <12%) and scalability, but it is highly dependent on high-quality historical data, sensitive to data missing and external emergencies, and faces the challenge of high computational complexity in hyperparameter tuning.

Addressing the issue in current e-commerce sales forecasting where a single model struggles to balance linear trends and nonlinear fluctuations, and lacks statistical rigor, this study constructs an LSTM-Prophet fusion model. The aim is to verify, through a paired t-test, whether its prediction accuracy (MAE, MAPE, RMSE) is statistically significant ($p<0.05$) compared to a single model. Additionally, it reports the 95% confidence intervals for each evaluation metric to quantify uncertainty. Based on this, the prediction confidence intervals are integrated into the inventory dynamic control strategy, with the goal of increasing the inventory turnover rate by over 18% and controlling the stock-out rate within 5%. This forms a statistically reliable and decision-making transparent intelligent supply chain optimization scheme.

The LSTM-Prophet fusion model constructed in this study adopts a dual-channel collaborative architecture: first, the Prophet model is utilized to explicitly decompose the sales time series, extracting linear components such as trend terms, seasonal terms, and holiday effects; subsequently, the residual sequence that Prophet fails to fit is input into the LSTM network, through its gating mechanism, to learn nonlinear dynamic features; finally, the prediction results are obtained through weighted fusion. The model hyperparameters are automatically tuned through Bayesian optimization, with the number of LSTM layers ranging from 1 to 3, the number of hidden units ranging from 32 to 256, and the Dropout rate set between 0.1 and 0.5. Training is conducted using 150,000 daily sales data from e-commerce platforms over a three-year period, covering multimodal features such as product attributes, promotional activities, and user behaviors. Spatio-temporal features are constructed through methods such as sliding window statistics and Fourier transform. Experiments show that this architecture reduces the mean absolute error (MAE) on the test set to 8.7, which is 29.3% and 17.1% lower than that of the single Prophet and LSTM models, respectively. The inventory turnover rate is increased by 18.4%, and the out-of-stock rate is controlled at 3.2%.

6 Conclusion

This paper proposes an e-commerce product sales forecasting and inventory dynamic control model based on the LSTM-Prophet fusion architecture, aiming to address the issue of insufficient accuracy in traditional forecasting methods for complex e-commerce scenarios. By combining the efficiency of LSTM in time series modeling with the advantages of Prophet in detecting seasonal and trend features, the model significantly improves the accuracy of sales forecasting and optimizes the inventory dynamic control strategy.

In the experiment, we selected the product sales data of a large e-commerce platform for a three-year period, encompassing a total of 1,000 products across five categories, and verified the model's performance. The experimental results show that:

(1) In the sales forecasting task, the average absolute error MAE of the LSTM-Prophet fusion model is 12.3, which is 23.5% lower than that of the single LSTM model and 31.8% lower than that of the Prophet model, proving the effectiveness of the fusion architecture.

(2) For the forecast of seasonal commodities, the forecast error of the model during the peak holiday period is only 8.7, which is 42.1% lower than that of the

traditional ARIMA model.

(3) In the inventory dynamic control experiment, the inventory turnover rate based on the forecast results increased to 5.2 times/year, the inventory backlog cost decreased by 18.7%, and the out-of-stock rate decreased to 3.4%, which was significantly better than the static inventory strategy.

The core limitation of the LSTM Prophet fusion model lies in its excessive reliance on historical data and vulnerability to sudden external events, which can easily lead to inventory backlog or shortage. The solution lies in enhancing the adaptability and system resilience of the model: firstly, introducing external signals such as social media trends and news events to enable the model to perceive market mutations; Secondly, establish a mechanism for quantifying prediction uncertainty and manual intervention, and initiate manual decision-making when confidence is low; The third is to adopt a layered prediction strategy, using complex models for core products and lightweight models for long tail products to improve efficiency; Ultimately, the prediction will be linked with the flexible supply chain, and safety stock will be dynamically adjusted based on the confidence level of the prediction to construct an intelligent inventory control system that is resistant to impact.

In future research, we will further optimize multi-category collaborative forecasting and explore a real-time data-driven adaptive regulation mechanism to provide e-commerce companies with efficient sales forecasting and inventory management solutions.

Table 4: Nomenclature table

Category	Abbreviation/Symbol	Full Name / Description
Model Names	LSTM	Long Short-Term Memory Network
	Prophet	Facebook's Open-Source Time Series Forecasting Model
	LSTM-Prophet	The Hybrid Model Proposed in This Study
Model Variables	*y(t)*	Observed Value at Time t(e.g., Sales)
	g(t),s(t),h(t)	Trend, Seasonal, and Holiday Components of the Prophet Model
	C_t, h_t	Cell State and Hidden State of LSTM
	f_t, i_t, o_t	Forget Gate, Input Gate, and Output Gate of LSTM
Evaluation Metrics	MAE	Mean Absolute Error
	MAPE	Mean Absolute Percentage Error
	RMSE	Root Mean Square Error
	R^2	Coefficient of Determination
Inventory Metrics	Inventory Turnover Rate	Measure of Inventory Flow Efficiency
	Stockout Rate	Probability of Inventory Shortage
	Safety Stock	Buffer Stock for Dealing with Uncertainty

References

[1] B. Biswas, M. K. Sanyal, and T. Mukherjee, "AI-Based Sales Forecasting Model for Digital Marketing," International Journal of E-Business Research, vol. 19, no. 1, pp., 2023. doi:10.4018/ijebr.317888.

[2] Y. Chen, X. Xie, Z. Pei, W. Yi, C. Wang, W. Zhang, and Z. Ji, "Development of a Time Series E-Commerce Sales Prediction Method for Short-Shelf-Life Products Using GRU-LightGBM," Applied Sciences-Basel, vol.14, no.2, pp., 2024.

doi:10.3390/app14020866.

[3] H. Dai, Q. Xiao, N. Yan, X. Xu, and T. Tong, "Item-level Forecasting for E-commerce Demand with High-dimensional Data Using a Two-stage Feature Selection Algorithm," Journal of Systems Science and Systems Engineering, vol. 31, no. 2, pp. 247-264, 2022. doi:10.1007/s11518-022-5520-1.

[4] Z. Fan, "E-COMMERCE DATA MINING ANALYSIS BASED ON USER PREFERENCES AND ASSICIATION RULES," Scalable Computing-Practice and Experience, vol.25, no.3, pp.1765-1772, 2024. doi:10.12694/scpe.v25i3.2682.

[5] Adnane Cabani, Iyadh Cabani, Koen De Turck, and Michel Kieffer, "Analysis and forecasting of the load of parcel pick-up points: Contribution of C2C e-commerce," Computers & Industrial Engineering, vol. 200, pp. 110770, 2025. doi: https://doi.org/10.1016/j.cie.2024.110770.

[6] Edbert Handy and Muhammad Amien Ibrahim, "Product-level sales prediction in e-commerce using LSTM modeling," Procedia Computer Science, vol. 269, pp. 853-862, 2025. doi: 10.1016/j.procs.2025.09.028.

[7] S. Arslan, "A hybrid forecasting model using LSTM and Prophet for energy consumption with decomposition of time series data," Peerj Computer Science, vol.8, no., pp., 2022. doi:10.7717/peerj-cs.1001.

[8] T. Bashir, H. Chen, M.F. Tahir, and L. Zhu, "Short term electricity load forecasting using hybrid prophet-LSTM model optimized by BPNN," Energy Reports, vol.8, no., pp.1678-1686, 2022. doi:10.1016/j.egyr.2021.12.067.

[9] Z. Chen, C. Wang, L. Lv, L. Fan, S. Wen, and Z. Xiang, "Research on Peak Load Prediction of Distribution Network Lines Based on Prophet-LSTM Model," Sustainability, vol.15, no.15, pp., 2023. doi:10.3390/su151511667.

[10] J. Cai, L. Jia, Q. Zhou, and D. Yao, "E-commerce supply chain inventory decisions and contract design considering sales effort and risk aversion," Electronic Commerce Research, vol.24, no.3, pp.1847-1888, 2024. doi:10.1007/s10660-022-09575-3.

[11] D. Gao, N. Wang, Z. He, and L. Zhou, "Analysis of bullwhip effect and inventory cost in the online closed-loop supply chain," International Journal of Production Research, vol.61, no.20, pp.6808-6828, 2023. doi:10.1080/00207543.2022.2138609.

[12] Y. Liu, "A Cross-Border E-Commerce Cold Chain Supply Inventory Planning Method Based on Risk Measurement Model," Mobile Information Systems, vol. 2022, no., pp., 2022. doi:10.1155/2022/6318373.

[13] C.-H. Wu, W.-T. Huang, J.-F. Dang, and M.-Y. Yeh, "Optimal analysis of a multi-server queue with preprocessing time and replenishment inventory," Journal of the Chinese Institute of Engineers, vol. 47, no. 1, pp. 120-132, 2024. doi:10.1080/02533839.2023.2274089.

[14] A. Ecevit, I. Ozturk, M. Dag, and T. Ozcan, "Short-Term Sales Forecasting Using LSTM and Prophet Based Models in E-Commerce," Acta Infologica, vol. 7, no. 1, pp. 59-70, 2023. doi:10.26650/acin.1259067.

[15] T. Feng, Z. Zheng, J. Xu, M. Liu, M. Li, H. Jia, and X. Yu, "The comparative analysis of SARIMA, Facebook Prophet, and LSTM for road traffic injury prediction in Northeast China," Frontiers in Public Health, vol.10, no., pp., 2022. doi:10.3389/fpubh.2022.946563.

[16] W. Liu, X. Yu, Q. Zhao, G. Cheng, X. Hou, and S. He, "Time Series Forecasting Fusion Network Model Based on Prophet and Improved LSTM," Cmc-Computers Materials & Continua, vol.74, no.2, pp.3199-3219, 2023. doi:10.32604/cmc.2023.032595.

[17] F. Meng, and R. Dou, "Prophet-LSTM-BP Ensemble Carbon Trading Price Prediction Model," Computational Economics, vol. 63, no. 5, pp. 1805-1825, 2024. doi:10.1007/s10614-023-10384-5.

[18] Y. Ning, H. Kazemi, and P. Tahmasebi, "A comparative machine learning study for time series oil production forecasting: ARIMA, LSTM, and Prophet," Computers & Geosciences, vol.164, no., pp., 2022. doi:10.1016/j.cageo.2022.105126.

[19] H. Shamshad, F. Ullah, A. Ullah, V.R. Kebande, S. Ullah, and A. Al-Dhaqm, "Forecasting and Trading of the Stable Cryptocurrencies with Machine Learning and Deep Learning Algorithms for Market Conditions," Ieee Access, vol. 11, no., pp. 122205-122220, 2023. doi:10.1109/access.2023.3327440.

[20] M.J.A. Shohan, M.O. Faruque, and S.Y. Foo, "Forecasting of Electric Load Using a Hybrid LSTM-Neural Prophet Model," Energies, vol.15, no.6, pp., 2022. doi:10.3390/en15062158.

[21] S. Su, Y. Sun, Y. Zeng, and C. Peng, "Aviation risk prediction based on Prophet-LSTM hybrid algorithm," Aircraft Engineering and Aerospace Technology, vol.95, no.7, pp.1054-1061, 2023. doi:10.1108/aeat-08-2022-0206.

[22] J. Huang, Q. Chen, and C. Yu, "A New Feature Based Deep Attention Sales Forecasting Model for Enterprise Sustainable Development," Sustainability, vol. 14, no. 19, pp., 2022. doi:10.3390/su141912224.

[23] X. Li, J. Shen, D. Wang, W. Lu, and Y. Chen, "Multi-modal transform-based fusion model for new product sales forecasting," Engineering Applications of Artificial Intelligence, vol. 133, no., pp., 2024. doi:10.1016/j.engappai.2024.108606.

[24] Daniel Kristian and Puti Andam Suri, "Comparative study of AutoML and machine learning in retail forecasting," Procedia Computer

Science, vol. 269, pp. 882-891, 2025. doi: 10.1016/j.procs.2025.09.031.

[25] Tejaswini Samal and Anupam Ghosh, "Ensemble-based predictive analytics for demand forecasting in multi-channel retailing," Expert Systems with Applications, vol., pp. 130212, 2025. doi: 10.1016/j.eswa.2025.130212.

[26] Xingyou Yan, Heng Zhang, and Qiang Miao, "A novel sales forecast framework based on separate feature extraction and reconciliation under hierarchical constraint," Computers & Industrial Engineering, vol. 201, pp. 110875, 2025. doi: 10.1016/j.cie.2025.110875.