A Multi-Layer Machine Learning Approach for Inventory Optimization: Case Study of Product Shipment Trend Forecasting

Junli Shi, Youpeng Fan*, Yue Lv

Dalian Polytechnic University, Dalian, 116034, China
*Corresponding author: qwe18396728207@163.com

Abstract: As market competition intensifies and technological advancements progress, the time and labor costs associated with individual components have risen accordingly. Simultaneously, corporate inventory levels have increased due to various factors, resulting in higher operational expenses. As a result, effective inventory management has become a critical concern for businesses. This paper presents a machine learning-based solution for inventory optimization, addressing the challenges faced by a specific company. By analyzing historical shipment data, evaluating sales representative performance, assessing client purchasing power, and forecasting product shipment trends, the solution effectively reduces inventory costs while optimizing production planning. The proposed model has been shown to deliver high accuracy and feasibility in both shipment forecasting and inventory management.

Keywords: Inventory Optimization, Machine Learning, Shipment Trend Forecasting, Sales Representative Performance, Production Planning

1. Introduction

In many modern enterprises, the wide variety of similar products (e.g., bolts) often leads to production losses, such as defective parts. Additionally, most clients tend to purchase the same products within a two-month period. As a result, sales representatives frequently place multiple production orders out of habit, causing production volumes to exceed actual purchase quantities. This, combined with the inherent variation in order sizes, leads to the accumulation of excess stock for each product, resulting in an overall inventory level that becomes disproportionately high. Maintaining an optimal inventory level is essential to meet market demand while avoiding issues such as increased costs from overstocking, production disruptions from insufficient inventory, and reduced customer satisfaction.

Chen et al. observed that with the rapid advancement of information technology and the application of big data techniques, machine learning-based methods have demonstrated significant potential in fields such as time series forecasting and pattern recognition [1]. Zhang et al. research indicated that big data-driven demand forecasting methodologies could assist enterprises in better navigating market fluctuations and mitigating risks, thereby enhancing the overall competitiveness of supply chains [2]. Additionally, findings from Yue et al. suggested that in inventory management, utilizing historical data to forecast future product shipment trends and optimizing production schedules to reduce inventory levels have become key strategies for modern enterprises pursuing efficient operations [3].

Building on the previous analysis, this paper presents a time series forecasting method that integrates business rule feature extraction with machine learning models. The approach aims to reduce inventory levels and improve operational efficiency and competitiveness by forecasting future shipment trends and optimizing production planning. First, by analyzing the number of clients managed and shipment volumes per sales representative, a comprehensive ranking of the 12 sales representatives' capabilities is developed. Clients are then categorized into three purchasing power tiers—high, medium, and low—before using the Prophet additive model to predict future shipment volumes. This strategy seeks to optimize production planning, thereby reducing inventory levels and enhancing operational efficiency.

2. Model development

After conducting a thorough assessment of sales representatives' capabilities and clients' purchasing power, product shipment trends are predicted. A multiple linear regression analysis is used to evaluate

the capabilities of sales representatives, while a three-tier cluster analysis is performed to assess the purchasing power of clients.

2.1. Sales representative competency analysis

This analysis of salesperson performance seeks to objectively quantify individual contributions, providing a solid foundation for decision-making in order allocation. The study employs a dual-dimensional assessment framework that integrates multivariate linear regression with weighted business metrics.

During the data preparation phase, the study utilized the corrected sales representative identifiers through the function to retrieve the relevant data, grouping by 'sales representative'. This process accurately extracted each representative's total order count (O_j), total shipment volume (S_j), and total customer count (C_j). To address the scale differences between these metrics, Min-Max normalization was applied to the three key indicators using a utility class. The formula used is as follows:

$$x_{\text{norm}} = \frac{x - x_{\min}}{x_{\max} - x_{\min}} \tag{1}$$

This results in the normalized variables $O_j^{\text{norm}}, S_j^{\text{norm}}, C_j^{\text{norm}}$.

Building on this foundation, a multiple linear regression model was developed with 'shipment volume' as the dependent variable. Dummy variables representing sales representatives were generated using a function (while ensuring the avoidance of multicollinearity) as independent variables, and the regression equation was subsequently fitted:

$$y = \beta_0 + \beta_1 x_1 + \cdots + \beta_{11} x_{11} + \epsilon \tag{2}$$

The model quantifies the independent contribution of each sales representative to dispatch efficiency, with its validity confirmed by an $R^2 > 0.7$ and core regression coefficients with $p < 0.05$.

To provide a comprehensive evaluation of sales representatives' performance, a composite score was calculated based on business weightings: order volume (0.4), shipment volume (0.2), and customer count (0.4).

$$\text{Overall Score}_j = 0.4 \times O_j^{\text{norm}} + 0.2 \times S_j^{\text{norm}} + 0.4 \times C_j^{\text{norm}} \tag{3}$$

A function was employed to rank sales representatives and stratify their performance into three tiers: high (≥ 0.8), medium (0.3–0.8), and low (≤ 0.3).

This tiered classification directly informs business decisions: orders from high-purchasing-power clients are prioritized for top-performing sales representatives, while targeted training programs may be developed for underperforming staff based on their specific shortcomings across various metrics, thereby enhancing the overall capabilities of the sales team.

2.2. Purchasing Power Analysis for Party A

The purchasing power analysis for Client A aims to categorize all 22 clients into high, medium, and low segments using the RFM model and K-means clustering algorithm, thereby facilitating the precise identification of customer value.

To construct the RFM characteristics, 'Client A' was used as the grouping key to extract three-dimensional features: Recency (number of days since the last purchase), Frequency (purchase frequency), and Monetary (purchase volume). The formula for Recency is as follows [4]:

$$Recency_j = ma - min(day_i) \tag{4}$$

Here, max_d denotes the maximum number of days across the entire observation period, whilst $min(day_i)$ represents the index of the most recent procurement date for customer j. Frequency is defined as the number of shipment records for the customer, whereas Monetary signifies the customer's total shipment volume.

Due to the substantial differences in magnitude among the three RFM characteristics, the features were subsequently standardized using the StandardScaler method. The transformation formula is as follows:

$$x_{\text{std}} = \frac{x - \mu}{\sigma} \tag{5}$$

Here, x denotes the original feature value, μ represents the mean, and σ signifies the standard deviation.

After standardization, cluster analysis was performed [5]. The model minimizes the following objective function to ensure minimal intra-cluster distance:

$$\arg \min_{S} \sum_{i=1}^{3} \sum_{x \in S_i} |x - \mu_i|^2 \tag{6}$$

Here, μ_i denotes the $i-th$ cluster centre. The validity of the clustering results was validated using the contour coefficient (P > 0.5). Ultimately, the labels obtained from clustering (0/1/2) were mapped to 'low/medium/high purchasing power' labels based on the principle of descending *Monetary* order combined with ascending *Recency* order.

2.3. Product Shipment Trend Forecast

The inventory adjustment strategy is based on forecasts from the Prophet model, with the goal of aligning current stock levels with safety stock to ensure dynamic synchronization between demand, production, and inventory. The Prophet formula is as follows:

$$\hat{y}(t) = g(t) + s(t) + h(t) + \epsilon_t \tag{7}$$

Here, $\hat{y}(t)$ denotes the forecasted shipment volume at time t, this study first calculated the product's safety stock using the following formula:

$$I_{\text{safety}} = 1.2 \times \sigma \tag{8}$$

Here, σ denotes the standard deviation of product shipments over the preceding ten weeks, calculated as such, whilst 1.2 represents the predetermined safety factor. Integrating the future shipment volume predicted by the Prophet model($\widehat{y_{\text{future}}}$)alongside the current inventory level(I_{current})this study determines the optimal production quantity according to the following formula:

$$\text{Production volume} = ma\,\mathrm{x}\big(\widehat{y_{\text{future}}} - I_{\text{current}} + I_{\text{safety}}, 0\big) \tag{9}$$

2.4. Inventory Adjustment Strategy

The inventory adjustment strategy is built around the core framework of 'demand forecasting – inventory alignment – production optimization,' establishing a quantitative decision-making model driven by future shipment volume projections from the Prophet additive model. This approach incorporates assessments of sales representative capabilities and client purchasing power segmentation, aiming to create a closed-loop operational cycle of 'forecasting – execution – validation.

In the forecasting phase, weekly shipment volumes for the next 5–7 weeks, generated by the Prophet model in Section 3.3 (which includes overall trends and forecasts for the top five core products), are used as the predicted shipment volume $\widehat{y_{\text{future}}}$. Current inventory ($I_{\text{current}}$) is retrieved in real-time from the enterprise WMS system, excluding non-usable stock such as damaged goods or items pending dispatch. Safety stock (I_{safety}) is calculated based on the historical variability of product shipments. First, the standard deviation σ of actual shipment quantities y_t over the past 10 weeks is computed using the formula:

$$\sigma = \sqrt{\frac{1}{n-1} \sum_{t=1}^{n} (y_t - \bar{y})^2}$$

where $\bar{y}$ is the mean shipment quantity over the past 10 weeks [6], and $n = 10$. This standard deviation is then combined with a safety factor of 1.2 (designed to balance service levels and inventory costs) to determine the safety stock:

$$I_{\text{safety}} = 1.2 \times \sigma \tag{10}$$

Based on the aforementioned parameters, the optimal production volume is determined by the following formula:

$$\text{Production volume} = ma\,\mathrm{x}\big(\widehat{y_{\text{future}}} - I_{\text{current}} + I_{\text{safety}}, 0\big) \tag{11}$$

If the calculated result is greater than zero, production should be scheduled 1-2 weeks in advance to

meet future demand and maintain an appropriate safety reserve. Conversely, if the result is less than or equal to zero, production should be suspended, with priority given to utilizing existing stock to avoid overstocking.

During the strategy implementation phase, the sales representative competency scores from Section A were correlated with client purchasing power classifications from Section B. In practice, large-value orders from high-purchasing-power clients (defined as Recency < 5 days and Monetary ≥ 1.5 times the mean) were prioritized for allocation to high-performing sales representatives, whose composite score was no less than 0.8. This approach aimed to reduce the order deviation rate from 15% to below 8%. For fragmented demand from low-purchasing-power clients (Recency > 10 days and Frequency ≤ 2 times), a 'small-batch, high-frequency' production model [7] was implemented. Additionally, inventory verification was recommended to begin two weeks in advance for demand peaks predicted by the Prophet model.

3. Actual production validation

3.1. Data Preprocessing and Feature Engineering

The data for this study was sourced from a Chinese enterprise, which provided shipment records covering a period of nearly 20 weeks (approximately five months) during a specific quarter.

The raw data, with the week as the smallest unit of observation, underwent comprehensive cleansing and consolidation to create a standardized analytical dataset. This dataset includes over twenty client organizations, each uniquely identified by a single letter. More than 65% of these clients are long-term partners with whom we have collaborated for over three years, while the remaining clients are temporary, having engaged with us for less than one year. The product dimension consists of more than fifty standardized components, each with clearly defined specifications and assigned a unique integer code, facilitating detailed analysis of shipment trends for individual items. As a key metric, the shipment volume field utilizes the interquartile range (IQR) method to identify and exclude extreme outliers. This method calculates the 25th percentile (Q_1) and 75th percentile (Q_3) of shipment volumes, defining the valid data range using the following formula.

This dataset is structured around three core analytical dimensions: time, salesperson-customer relationships, and product-customer interactions, allowing for multi-dimensional insights. All sensitive information (e.g., actual customer names) has been anonymized (salespersons are represented by integers 1–12). Additionally, records with missing key fields were excluded, and invalid data was filtered based on business logic, resulting in a dataset validity rate exceeding 90%. For numerical features such as 'shipment quantity' and 'salesperson order count,' Min-Max normalization was applied. This method linearly transforms values to the [0, 1] range using the conversion formula. This thorough preprocessing process provides a robust foundation for subsequent feature engineering and modeling analysis.

At the feature engineering level, this study further extracted key business insights. Based on the RFM model (Recency: number of days since last purchase, Frequency: purchase frequency, and Monetary: purchase volume), core features reflecting the client's purchasing power were constructed.

3.2. Model Selection and Training

This study employed distinct models to address three core tasks: quantifying sales representative performance, segmenting clients, and forecasting shipment trends. To assess the contribution of individual sales representatives to shipment efficiency, a multiple linear regression model was used, with 'sales representative identifier' as the independent variable and 'shipment volume' as the dependent variable. The model was constructed by incorporating a constant term. For client segmentation, the K-means clustering algorithm was applied, using standardized Recency, Frequency, and Monetary (RFM) features as inputs. The number of clusters was determined, and the objective function was optimized accordingly. To forecast shipment trends, the Prophet additive model was employed, with weekly total shipment volume as the target variable to predict aggregate volumes for the upcoming two weeks. The training set included 70% of the entire observation period, ensuring the model captured long-term trends and cyclical variations. Additionally, to improve prediction accuracy for core products (the top five by shipment volume), a simple exponential smoothing model was incorporated. By minimizing the sum of squared residuals to optimize smoothing coefficients, more accurate forecasts for key products were achieved.

For time series models, the mean absolute error and root mean square error were calculated using the test set to forecast the overall trend of core products. For multiple linear regression models, the coefficient of determination assessed model explanatory power, while the significance of regression coefficients ($p<0.05$) validated the reliability of quantifying salesperson capabilities [8]. For K-means clustering models, the rationality of Party A's stratification is validated through the contour ratio and the intra-cluster/inter-cluster distance ratio. Concurrently, inventory levels before and after production plan adjustments are compared.

Calculate the safety stock $I_{safety} = 1.2 \times \sigma$, where σ denotes the standard deviation of historical weekly shipment volumes for the product, and 1.2 represents the reference safety factor. Then, based on the forecast shipment volume $\widehat{y_{future}}$ for the core product and the current inventory level $I_{current}$.

The optimal production volume is determined by the formula: Production Quantity $= \max(\widehat{y_{future}} - I_{current} + I_{safety}, 0)$ If the result is negative, production is suspended to prevent inventory accumulation. If the result is positive, production is scheduled 1-2 weeks ahead of the forecast cycle. Concurrently, order allocation is optimized based on sales representative capability scores and client purchasing power classifications. High-value orders from clients with strong purchasing power are prioritized for high-performing sales representatives (e.g., those with a composite score > 0.8) in order to reduce order deviation rates. For fragmented demand from low-purchasing-power clients, a 'small-batch, high-frequency' production model is implemented, aligning with their non-cyclical procurement patterns. This approach ensures dynamic alignment between production planning and forecasted demand.
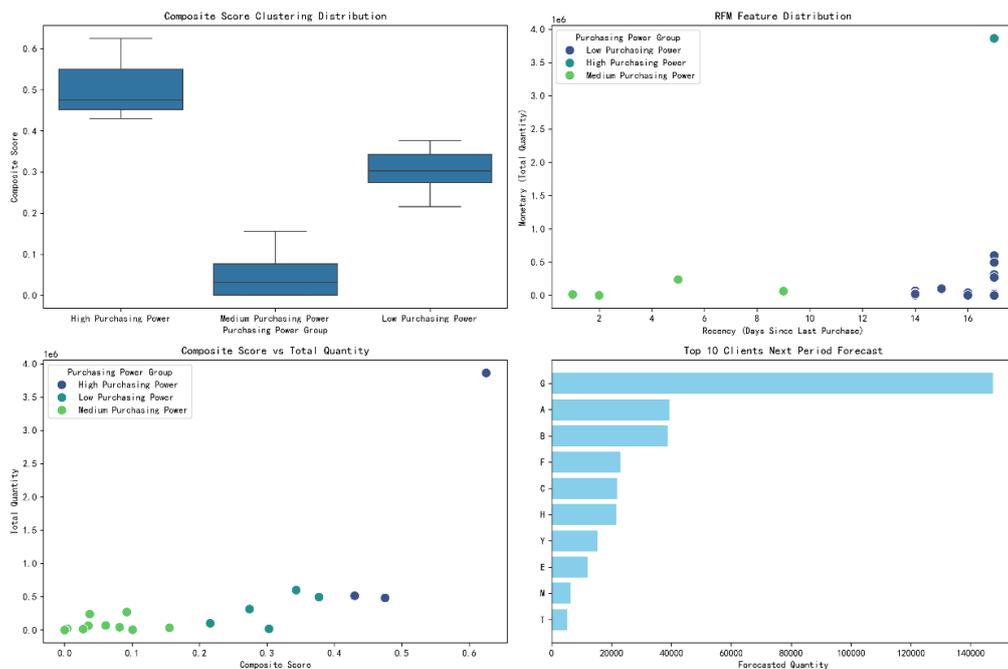
3.3. Analysis of Prediction Results

Fig.1. The relationship between customer purchasing power and behavioural characteristics

The study identified distinct characteristics across purchasing power segments in Fig.1: High-purchasing-power Category A clients typically exhibit Recency < 5 days, Frequency ≥ 8 times, and Monetary value ≥ 1.5 times the average. Such clients predominantly comprise long-term automotive manufacturing partners. Mid-purchasing-power clients exhibit Recency between 5-10 days, Frequency between 3-7 times, and Monetary value approximately 0.8-1.5 times the mean. These primarily comprise stable clients within the mechanical processing sector. Low-purchasing-power clients, conversely, demonstrate Recency > 10 days, Frequency ≤ 2 times, and Monetary value < 0.8 times the mean. These typically represent temporary collaborative clients. The relationship between sales representatives and clients is illustrated in Fig.2.
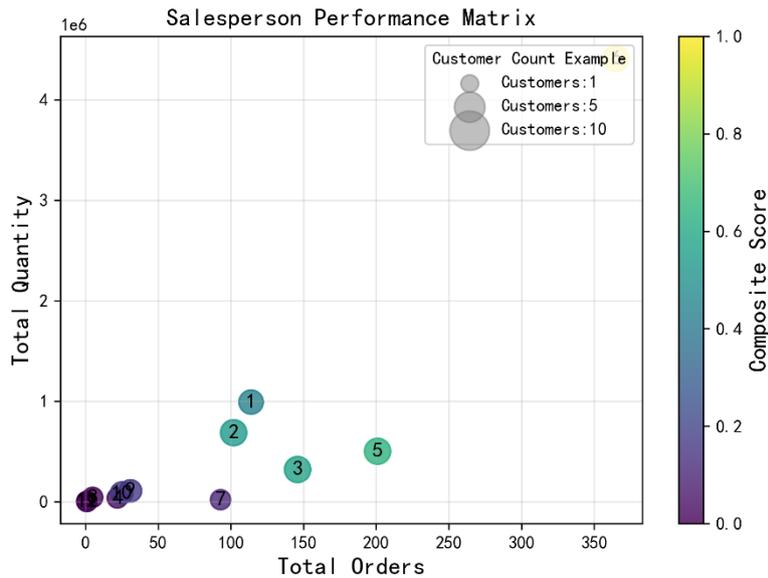
a. Salesperson Performance Matrix
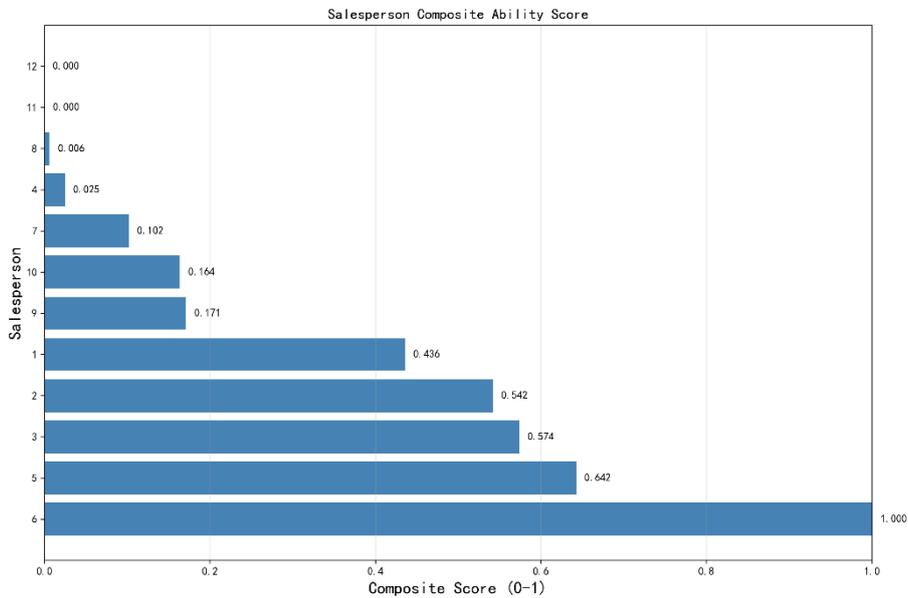
b. Salesperson Composite Ability Score

Fig.2. Salesperson-customer relationship

The scatter plot is used to assess the performance of sales representatives. High-performing representatives (e.g., Representative 6) manage a larger volume of orders and clients, while those with lower performance show weaker results. This visualization helps identify top performers as well as representatives who may require additional training, thereby facilitating more effective resource allocation and refinement of sales strategies.

Another chart, 'Sales Representative Comprehensive Capability Scores', illustrates each representative's capability strengths. Salesperson 6 demonstrates notably strong overall capabilities. This chart facilitates analysis of interaction intensity between sales personnel and clients, alongside market opportunities.

This study seeks to evaluate the model's performance in shipment forecasting using real dispatch data. Based on the forecasted results, production planning is optimized, allowing for the observation of subsequent inventory level changes. This analysis aims to determine whether integrating multiple machine learning models significantly improves the accuracy of shipment predictions, thereby contributing to inventory reduction and cost minimization. Additionally, it evaluates the feasibility and effectiveness of implementing this approach within corporate inventory management and sales

forecasting systems. It is important to note that subsequent visual analyses were limited to a subset of product categories within the historical dataset. To provide a focused demonstration of the alignment between forecasts and actual shipments, core products 1002, 1007, and 1014 were selected as key prediction priorities in Fig.3.
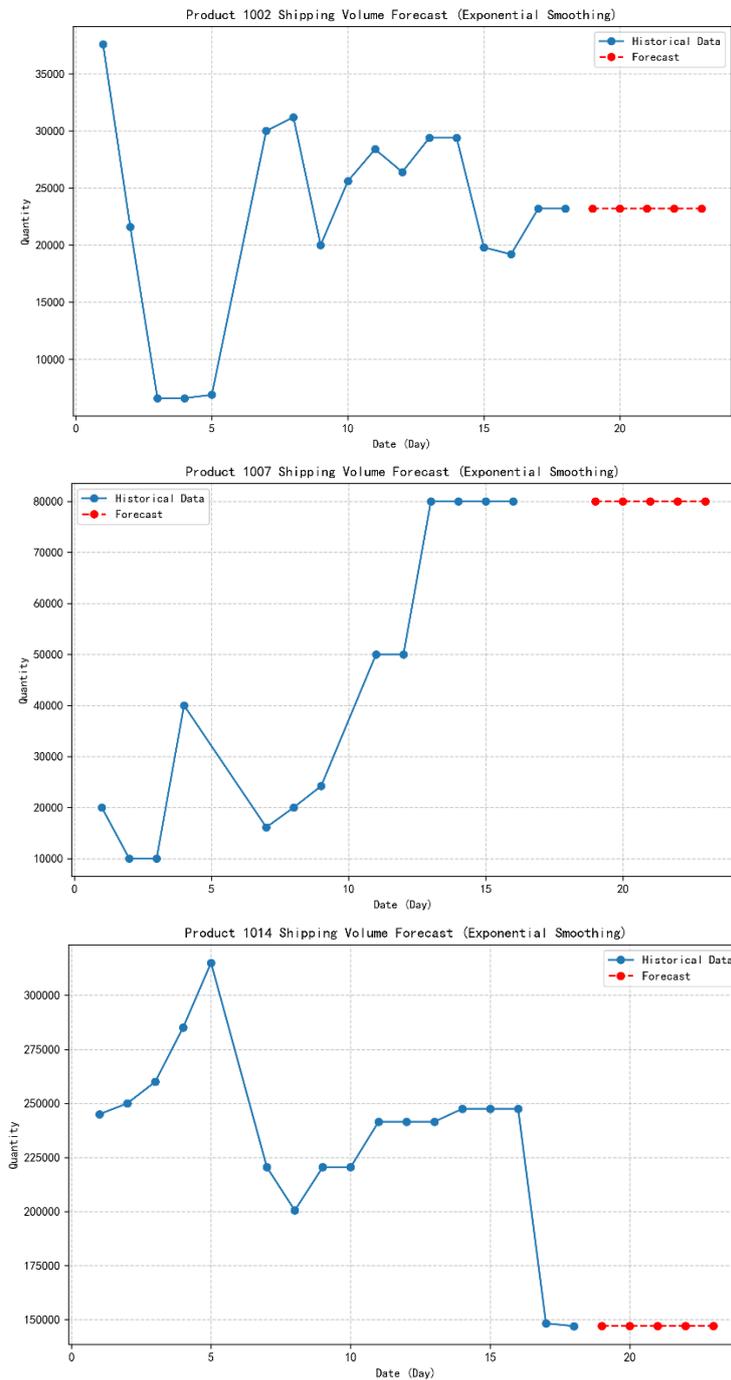

Fig.3. Three-week shipment volume forecast

Utilizing this model for shipment trend forecasting improves the accuracy of shipment predictions, effectively leveraging historical shipment data to forecast future trends. This outcome underscores the significant practical value of the machine learning model in corporate inventory management [9]. It not only offers reliable data support for sales forecasting but also aids companies in making more informed and precise decisions during production planning optimization, addressing the challenges of significant forecasting errors and delayed adjustments commonly associated with traditional inventory management practices [10]. The model smooths historical data fluctuations through exponential smoothing, providing stable projections for future shipment volumes. This, in turn, offers strong support for production

planning, inventory management, and demand scheduling.

For Product 1002, Salesperson F from Company C recorded actual figures of 3,200 units in Week 6, followed by 3,700 units over the subsequent two weeks. Salesperson I from Company I recorded actual figures of 400 units in Week 6, followed by 400 units over the subsequent two weeks. The total volume for Week 1 was 3,600 units. The second week's total was 4,100 units. Salesperson 6 demonstrated higher overall capability, while Salesperson 7 showed lower overall capability. Both companies were classified as medium-tier clients, with actual forecast data for the subsequent two weeks both at 22,500 units.

For Product 1007, Salesperson 1's actual shipment data for the following week was 0, with the subsequent week showing the same figure. However, Salesperson 1 ranks at an intermediate level, and most clients fall within the medium purchasing power category. This may be attributed to the product's low usage rate, resulting in fragmented actual shipments. Such products require separate consideration;

For Product 1014, actual data for the following week reveals the following sales figures: Salesperson 5 at Company H recorded 5,000 units, Salesperson 6 at Company F recorded 1,000 units, and Salesperson 6 at Company G recorded 56,000 units. Salesperson 2 at Company P recorded 20,000 units, bringing the total to 87,000 units. In the subsequent two weeks, the recorded sales were: Salesperson 5 at Company H, 5,000 units; Salesperson 5 at Company F, 6,000 units; Salesperson 6 at Company G, 61,000 units; and Salesperson 2 at Company P, 20,000 units. The first week's total volume was 87,000 units, and the second week's total was 92,000 units. Sales Representatives 2, 5, and 6 showed consistently high performance. The forecast for Product 1014 in the subsequent week and the following two weeks was 15,000 units.

Data visualizations reveal that previous sales figures exhibited considerable fluctuations, suggesting that straightforward forecasting methods alone cannot resolve practical challenges. These fluctuations highlight the need to account for major societal events, such as the COVID-19 pandemic, and competitive influences. For companies, the demand for streamlined, efficient forecasting methods means that overly complex approaches are not feasible. Simpler methods are insufficient when addressing the issue of overfitting optimization, particularly when data cycles are long, necessitating the readjustment of training datasets only after significant events have occurred.

Integrating salesperson and customer classification with forecasting reveals that when a single salesperson manages multiple companies, extended delivery cycles arise. This constitutes a product issue, amenable to a 'small-batch, high-frequency' production model. Adopt a conservative safety factor of 1.0 as baseline. Adjust this based on customer purchasing power: reduce the safety factor for low-purchasing-power clients, and moderately increase it for high-purchasing-power clients. Where multiple sales representatives handle several companies, delivery cycles tend to be shorter. Changes at individual companies do not typically cause product stagnation. If both sales representative capability and customer classification are at mid-to-high levels or above, production can be significantly scaled up. A higher safety factor, such as 1.3, may be referenced. Where one or two sales representatives manage a single company, delivery cycles are also relatively short but exhibit greater volatility. Sustained high-volume production may lead to significant inventory backlogs and increased production costs. Production should be coordinated using an average value, with a medium safety factor applied (e.g., 1.1 or 1.2).

The effectiveness of the strategy will be assessed by comparing pre- and post-optimization metrics, specifically the inventory turnover rate (with a target improvement of 15% or more) and inventory backlog (with a target reduction of 20% or more). The goal is to achieve precise alignment between demand, inventory, and production, effectively addressing the core operational challenges of overproduction, which leads to inventory backlogs, and underproduction, which results in order delays.

4. Model Discussion and Conclusions

4.1. Discussion and Extension

(1) A comparative analysis of the actual results reveals that relying solely on shipment trend forecasting is insufficient for effectively addressing dispatch challenges. It becomes evident that all three factors must be considered simultaneously. This suggests that, in certain market conditions, a single machine learning approach alone does not efficiently resolve the underlying issues. However, this methodology demonstrates significant practical value in inventory management, enabling organizations to achieve greater scientific rigor and precision in both sales forecasting and production planning optimization.

(2) The limited data volume may have restricted the model's ability to generalize effectively. The

relatively small scale of the dataset used in this experiment could have constrained the model's capacity to adapt to broader contexts. Future work could expand the dataset by incorporating data from additional enterprises, thereby enhancing the model's applicability across a wider range of scenarios.

(3) After training the relevant models, including the RFM-K-means model for analyzing client purchasing power and the Prophet model for inventory adjustment strategies, these models were deployed for practical inventory management tasks. This implementation allowed for the application of the models in real-world scenarios, supporting more effective decision-making in inventory control. Taking the Prophet model as an example, it combines newly collected data and extracted features to compute future shipment volume predictions (e.g., for the next 5–7 weeks) using its additive formula:$y(t) = g(t) + s(t) + h(t) + \varepsilon$. Finally, based on these forecasts, production order priorities and quantities are adjusted to reduce inventory levels and enhance operational efficiency.

4.2. Conclusion

This work aims to optimise inventory management and production planning by integrating multiple machine learning models within a unified mathematical framework. By analysing complex relationships within business data, the models harness predictive capabilities to help businesses mitigate inventory risks and optimise production schedules.

In the analysis of purchasing power for Party A, the RFM model and K-means clustering algorithm were employed to categorise customers into high, medium, and low purchasing power groups. This segmentation enabled accurate identification of customer value.

For inventory adjustments, a dynamic 'demand-production-inventory' matching mechanism was developed, centred on the Prophet model. This model links current stock levels with safety stock, thereby enhancing the robustness of inventory management.

When implemented for real-world inventory tasks, the model maintained stock levels within optimal bounds by collecting data, extracting relevant features, generating forecasts, and optimising production planning. This approach successfully reduced long-term stockpiling, minimised inventory holding costs, and increased production efficiency.

The model demonstrated strong generalisation capabilities and low computational demands, making it both user-friendly and easy to implement.

Looking ahead, future improvements could focus on several areas. First, augmenting the model with additional enterprise data and longer time-series data would strengthen its generalisation abilities, allowing it to better address diverse business needs. Second, incorporating industry-specific characteristics or external macroeconomic indicators—such as seasonal fluctuations or overall economic growth—could further enhance the model's predictive power, allowing it to more accurately account for market fluctuations and their effects on inventory and production.

In conclusion, the model has proven effective in optimising inventory management and production planning. With further refinements, it has the potential to provide even greater value in a variety of business contexts.

References

[1] Chen, X. G., Sheng, J., & Kong, L. R. (2005). Application of operations research in the inventory management of pharmacy department. Strait Pharmaceutical Journal, (03), 13-15.
[2] Zhang T ,Zhu X ,Gou Q. (2017). Demand Forecasting and Pricing Decision with the Entry of Store Brand under Various Information Sharing Scenarios. Asia-Pacific Journal of Operational Research ,34(2).
[3] Yue, D. Q., Chi, Y. W., & Wei, Y. N. (2025). Performance analysis and inventory strategy optimization control of two-stage production-service inventory system. Chinese Journal of Engineering Mathematics, 42(05), 806-822.
[4] Zhong, F. Y. (2024). Research on material inventory management of traditional foreign-related manufacturing enterprises based on EOQ model. Logistics Sci-Tech, 47(11), 39-42.
[5] Zhang, X. (2024). Research on optimization of enterprise logistics inventory management under the background of supply chain. China Logistics & Purchasing, (10), 95-96.
[6] Zhao, F. Q. (2024). Research on optimization strategy of equipment inventory management based on supply chain. Enterprise Reform and Management, (04), 14-16.

[7] Zhou, J. G. (2023). Research on inventory management problems and countermeasures of manufacturing enterprises. Time-honored Brand Marketing, (11), 169-171.

[8] Zhang, Q., Lu, D., Xiang, Q., Lo, W., & Lin, Y. (2024). Design and optimization of dynamic reliability-driven order allocation and inventory management decision model. PeerJ Computer Science, 10, e2294

[9] Zhang, B.; Tan, W.J.; Cai, W.; Zhang, A.N. Leveraging Multi-Agent Reinforcement Learning for Digital Transformation in Supply Chain Inventory Optimization. Sustainability 2024, 16, 9996.

[10] Yang, H. L., Shi, X. J., & Xin, Y. C. (2024). Optimization research on inventory routing and pricing in omni-channel e-commerce. Journal of Transportation Systems Engineering and Information Technology, 24(04), 223-230.