Forecast-Driven Inventory Management for the Fast-Moving Consumer Goods Industry

by

Abdulelah S. Al Mesfer

B.Sc. in Industrial and Systems Engineering, King Fahd University of Petroleum and Minerals (2015)

Submitted to the System Design and Management Program in partial fulfillment of the requirements for the degree of

Master of Science in Engineering and Management

at the

Massachusetts Institute of Technology

June 2023

Authored by: Abdulelah S. Al Mesfer
System Design and Management Program
May 12, 2023

Certified by: Joan S. Rubin
Senior Lecturer, System Design & Management Program
Thesis Supervisor

Accepted by: Warren Seering
Weber-Shaughness Professor of Mechanical Engineering

Forecast-Driven Inventory Management for the Fast-Moving Consumer Goods Industry

by

Abdulelah S. Al Mesfer

Submitted to the System Design and Management Program

on May 12, 2023, in partial fulfillment of the

requirements for the degree of

Master of Science in Engineering and Management

Abstract

This thesis investigates the development and evaluation of various demand forecasting models for the Fast-Moving Consumer Goods (FMCG) industry on real-world data, to devise an inventory control policy for a third-party logistics provider. Demand forecasting is crucial in the retail industry, influencing supply chain management, inventory control, and pricing strategies. Accurately predicting demand is essential for optimizing resource allocation, reducing stockouts, and minimizing holding costs. In this study, we employ several time series models, including traditional time series models, such as ARIMA and SARIMA, and machine learning techniques, such as Random Forests, XGBoost, and Prophet, to forecast retail demand. The performance of these models is assessed using time series cross-validation techniques and accuracy measures, such as RMSE, MAPE, and MAE. Data preprocessing steps, including resampling, imputation of missing values and outliers, SKU prioritization, and feature engineering, are performed to enhance the reliability of the forecasting models. The results indicate that XGBoost outperforms the other models, showcasing its ability to generate accurate FMCG demand forecasts. Based on the forecasting error, a continuous review (s,Q) policy is formulated to improve inventory management for the third-party logistics provider. The proposed inventory control policy demonstrates the potential to minimize holding costs for the FMCG industry. Future research directions include the investigation of additional forecasting models, the integration of external factors, and the extension of the study to other retail contexts.

Thesis Supervisor: Joan S. Rubin

Title: Senior Lecturer, System Design & Management Program

Acknowledgments

I would like to express my deepest appreciation to my supervisor, Joan Rubin, for her unwavering guidance, support, and encouragement throughout my research. I am grateful for her dedication and mentorship, which played a critical role in my academic and personal growth.

I would also like to express my sincere thanks to Dr. Ilya Jackson for his continuous support and invaluable advice throughout the research process. His willingness to impart his knowledge and expertise significantly contributed to the success of my work. I am thankful for the countless hours he dedicated to assisting me.

Finally, my profound gratitude goes to my father, Saad Al Mesfer, and my mother, Basmah Dhaifallah. Since my earliest childhood memories, their faith in my abilities has been the constant wind beneath my wings. Their enduring support and motivation throughout my life have been the driving forces propelling me toward this accomplishment. My heartfelt thanks to my siblings, Faris, Ghaidaa, Khaled, Lamees, Layan, and Saud, who have stood by my side, offering encouragement and companionship throughout this journey. It is to my family, the pillar of my life, that I sincerely dedicate this thesis.

Contents

List of Abbreviations

Abbreviation	Definition
3PL	Third-Party Logistics
EOQ	Economic Order Quantity
CV	Coefficient of Variation
KPI	Key Performance Indicator
FMCG	Fast-Moving Consumer Goods
LSTM	Long Short Term Memory
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
RMSE	Root Mean Squared Error
SKU	Stock Keeping Units
OTIF	On Time Demand In Full
XGB, XGBoost	eXtreme Gradient Boosting

List of Figures

List of Tables

Chapter 1

1. Introduction

This chapter explains the motivation and purpose of this thesis. It presents an overview of the main and secondary objectives this research aims to explore. In addition, it provides an outline of the thesis chapters, offering a roadmap for the reader to better understand the key components and results of this academic work.

1.1. Motivation

Effective inventory management is critical to organizations operating in today's complex and competitive global supply chains. Striking a balance between maintaining optimal stock levels and avoiding stockouts is a significant challenge, as both have financial and operational implications (Silver et al., 2016).

In this context, machine learning and data-driven approaches present an opportunity to revolutionize inventory management by leveraging historical demand data to create more accurate demand forecasts. These forecasts can inform tailored inventory control policies, optimizing stock levels and improving overall supply chain performance.

This thesis investigates applying actual historical demand data and machine learning techniques to design a forecasting-based inventory control strategy. The research will cover data preprocessing, Pareto analysis for SKU prioritization, customization of forecasting models for each SKU, development of inventory control policies, and evaluation of the proposed strategy. By exploring these areas, the thesis aims to contribute to the existing body of knowledge and provide a practical framework for implementing data-driven inventory management strategies in real-world settings.

1.2. Thesis objectives

The primary objective of this thesis is to investigate the potential of leveraging historical demand data and machine learning to design a forecasting-based inventory control strategy. To accomplish this, the study seeks to achieve several secondary objectives.

First, a robust pipeline will be developed to preprocess the dataset for analysis, ensuring data quality and consistency. This involves addressing missing values, outliers, and mismatching data types, and transforming the data as needed. Second, Pareto analysis will focus on the dataset's highest priority items, allowing for a more targeted approach to inventory control. Third, various demand forecasting models selected from the literature review will be tested to select the most appropriate ones for the problem.

Finally, inventory control policies will be tailored to incorporate forecasted demand, forecasting errors, lead time, safety stock levels, and order quantities, to maximize inventory efficiency and minimize costs associated with stockouts and excess inventory.

By addressing these objectives, this thesis aims to contribute to the existing body of knowledge in inventory management and demand forecasting and provide a practical framework for implementing data-driven inventory management strategies in real-world settings.

1.3. Thesis outline

Chapter 1 begins with an introduction that sets the context and highlights the importance of demand forecasting in retail, emphasizing the need for accurate models to inform inventory management and supply chain operations. The objectives of the thesis are also outlined in this chapter, focusing on developing and evaluating various forecasting models and formulating an inventory control policy for a third-party logistics provider.

Following the Introduction, chapter 2 delves into a comprehensive review of the relevant literature. It starts with an overview of the FMCG industry, discussing its significance, challenges, and trends. The chapter then moves on to explore the role of demand forecasting in retail and presents various methodologies and techniques that have been employed in the field.

Chapter 3 outlines the methodologies employed in the study, from data preprocessing to forecasting model selection, cross-validation, and accuracy measurements. It details the various preprocessing techniques applied to the data, such as resampling, imputation of missing values and outliers, SKU prioritization, and feature engineering. The chapter also presents the forecasting models used in the study, ranging from baseline approaches to statistical and machine learning models. It discusses cross-validation techniques and accuracy measures employed to evaluate the performance of these models.

In chapter 4, the results of the forecasting models are presented along with their implications for inventory management. The development of a continuous review (s,Q) policy based on forecasting error is also discussed in this chapter. The performance of the various forecasting models is compared, and their accuracy in predicting FMCG demand is assessed.

The thesis concludes with chapter 5, which summarizes the key findings of the study and suggests avenues for future work.

Chapter 2

2. Background and literature review

This chapter presents the background necessary to understand the FMCG and 3PL industry alongside the importance of forecasting in inventory management. Moreover, it highlights the related work that inspired the methods used in this thesis.

2.1. Industry overview

The Fast-Moving Consumer Goods (FMCG) industry comprises products sold quickly and at relatively low cost, including items such as packaged foods, beverages, toiletries, and household cleaning products. These goods are characterized by their high turnover rates and short shelf lives, necessitating efficient and effective inventory management strategies. The following background discusses several key aspects of the FMCG industry, providing context for its unique challenges and opportunities.

Margins in the FMCG industry are typically thin due to intense competition and price sensitivity among consumers. Companies operating in this sector must focus on cost reduction and operational efficiency to maintain profitability (Narula, Anupam, 2012). The volume of orders in the FMCG industry is generally high, driven by the short shelf lives of the products and the constant need for replenishment. This presents challenges in managing inventory levels, as companies must balance the need for adequate stock to meet demand while avoiding overstocking, which can lead to waste and increased holding costs.

The FMCG industry is marked by its fast-paced nature, with rapid changes in consumer preferences and market trends. Companies must be agile and adaptable, responding quickly to shifts in demand and adjusting their inventory management strategies accordingly. This may involve adopting advanced forecasting techniques and utilizing real-time data to inform decision-making.

Seasonality has a considerable impact on the FMCG industry. Demand for certain products varies throughout the year due to factors such as holidays, weather patterns, and cultural events. Companies must account for these seasonal fluctuations in their inventory management strategies, ensuring that sufficient stock is available to meet demand during peak periods while avoiding excess inventory during periods of lower demand. Various factors, including promotional activities, competitor actions, and macroeconomic conditions, can drive demand fluctuations in the FMCG industry. Companies must be prepared to manage these fluctuations and adjust their inventory levels accordingly to avoid stockouts and lost sales opportunities. This may involve the use of sophisticated demand forecasting models that can incorporate external factors and provide more accurate predictions of future demand (“Global Executives Survey,” 2017).

3PL providers play an essential role in this industry. These providers offer a range of services, including transportation, warehousing, distribution, and inventory management, allowing their stakeholders to focus on their core business operations while benefiting from the resources of the 3PL providers. In this context, consolidators work with different retailers and suppliers. One of the most significant effects of retail consolidation is the increased demand for integrated and comprehensive logistics services, as retailers seek to optimize their supply chain operations and reduce costs through economies of scale. This has led to a growing need for 3PL providers that can offer a full suite of services, encompassing transportation, warehousing, distribution, and inventory management. Providers that can deliver these end-to-end solutions are well-positioned to capitalize on the opportunities presented above, as they can serve as a one-stop-shop for retailers' logistics needs (Wagner & Sutter, 2012).

In brief, the FMCG industry is characterized by its thin margins, high volume of orders, fast-paced nature, seasonality, and demand fluctuations. These factors present unique challenges for inventory management, requiring companies to adopt innovative strategies and leverage advanced technologies to remain competitive and maintain profitability. In this context, using data-driven approaches and machine learning techniques for demand forecasting and inventory control offers significant potential for improving efficiency and performance within the FMCG sector.

2.2. Demand forecasting in retail

Demand forecasting is a necessity in retail, as it informs various aspects of supply chain management, inventory control, and pricing strategies. This literature review presents an overview of the key studies and methodologies in demand forecasting for retail, providing a foundation for selecting the approach to tackle the research objectives of this thesis.

ARIMA (AutoRegressive Integrated Moving Average) models have long been a cornerstone in time series analysis and have been widely applied to various retail forecasting problems (Fildes et al., 2022). The popularity of ARIMA models in retail forecasting can be attributed to their ability to model and predict linear relationships in time series data, effectively capturing seasonality and trends through their autoregressive (AR) and moving average (MA) components (Hyndman & Athanasopoulos, 2018). Several studies have demonstrated the effectiveness of ARIMA models in retail sales forecasting, showcasing their ability to provide accurate and reliable predictions. (Veiga et al., 2014) argued that analytical time series methods, such as ARIMA, deliver precise outcomes more straightforwardly compared to other complex techniques for perishable food demand forecasting. (Falatouri et al., 2022) found that SARIMA outperformed the deep learning method (LSTM) for retail products that exhibit seasonal behavior in their data. Despite the emergence of more advanced machine learning techniques, ARIMA models remain a relevant and valuable approach in retail forecasting due to their simplicity, interpretability, and ease of implementation.

In recent years, machine learning techniques have gained prominence in demand forecasting, offering the potential for improved accuracy and adaptability. Supervised machine learning techniques have been applied to retail demand forecasting with promising results, demonstrating their ability to model complex relationships and capture non-linear patterns.

Prophet, an open-source forecasting tool developed by Facebook, has gained traction in the academic community as a robust option for time series forecasting, including applications in retail forecasting. As a flexible and scalable algorithm, Prophet is adept at addressing the intricacies and peculiarities often present in time series data, such as seasonality, trend shifts, and irregularities that are frequently observed in retail demand patterns. (Kumar Jha & Pande, 2021) examined Prophet's performance against ARIMA for supermarket sales forecasting and concluded that Prophet outperformed ARIMA. (Kolari & Sanz, 2022) employed Prophet to predict bank capital ratios, finding that it outperformed the ARIMAX model. These studies underscore the potential of Prophet as a valuable and effective tool for time series forecasting in various domains, including retail forecasting.

Random Forests have gained considerable attention as a powerful machine learning technique for demand forecasting. The non-parametric nature of Random Forests allows for the modeling of complex relationships between features and the target variable without the need for explicit assumptions about the underlying data distribution (Breiman, 2001). Random Forests have been successfully applied to various retail forecasting problems. (Vairagade et al., 2019) highlight that Random Forests outperformed Neural Networks in forecasting grocery item sales. (Papacharalampous & Tyralis, 2018) compared the performance of Prophet and Random Forests in predicting daily river streamflow and found that Random Forests outperformed Prophet. These studies highlight the ability of Random Forests to capture intricate patterns in the data, including seasonality, trends, and interactions between features, leading to improved forecasting accuracy compared to traditional time series models.

XGBoost, short for eXtreme Gradient Boosting, is a machine learning algorithm that has gained significant popularity in recent years for its ability to generate accurate forecasts across various domains, including retail forecasting. As an ensemble method, XGBoost builds upon the principles of decision trees and gradient boosting, allowing it to effectively model complex relationships and handle various data types and structures. (Wang et al., 2020) suggests using XGBoost for long-term forecasting compared to LSTM for thermal load prediction. (Ibrahem Ahmed Osman et al., 2021) compared XGBoost, SVR, and ANN performance in forecasting groundwater level in Malaysia and found XGBoost to outperform both models. Collectively, these studies highlight the versatility and effectiveness of XGBoost as a promising forecasting tool across diverse applications, warranting its consideration for retail forecasting tasks.

Demand forecasting is of paramount importance in the retail industry, as it impacts various aspects of supply chain management, inventory control, and pricing strategies. This literature review has provided a comprehensive overview of key studies and methodologies in retail demand forecasting, laying the groundwork for selecting the appropriate approach to address the research objectives of this thesis.

Traditional time series models, such as ARIMA, have maintained their relevance in retail forecasting due to their simplicity, interpretability, and ability to capture linear relationships, seasonality, and trends. Meanwhile, machine learning techniques have emerged as promising alternatives, offering enhanced accuracy and adaptability in modeling complex relationships and capturing non-linear patterns. Tools such as Prophet, Random Forests, and XGBoost have demonstrated their effectiveness in various forecasting applications, including retail demand.

Based on the literature examined in this review, it is evident that forecasting is a complex undertaking, with each method possibly emerging as the most effective approach under certain conditions. Consequently, this research employs real-world data to evaluate which method yields the highest accuracy concerning the 3PL provider. The investigation considers both conventional time series models and modern machine learning techniques, recognizing their potential to enhance the accuracy and reliability of retail demand forecasting for the 3PL provider.

3. Methods

This chapter presents the data processing, forecasting, and inventory control methods used for the analysis. After conducting the necessary literature review, a robust pipeline for preparing historical demand data is built, ensuring its quality and consistency by addressing missing values, outliers, and other issues and transforming and aggregating the data as required. Next, Pareto analysis is implemented to focus on the most significant items, enabling a targeted approach to inventory control and identification of high-impact SKUs. Various demand forecasting models are then explored, with the most suitable ones selected for the problem and customized for each SKU. Lastly, inventory control policies are tailored to each SKU to optimize inventory efficiency and minimize costs (see figure 3-1).

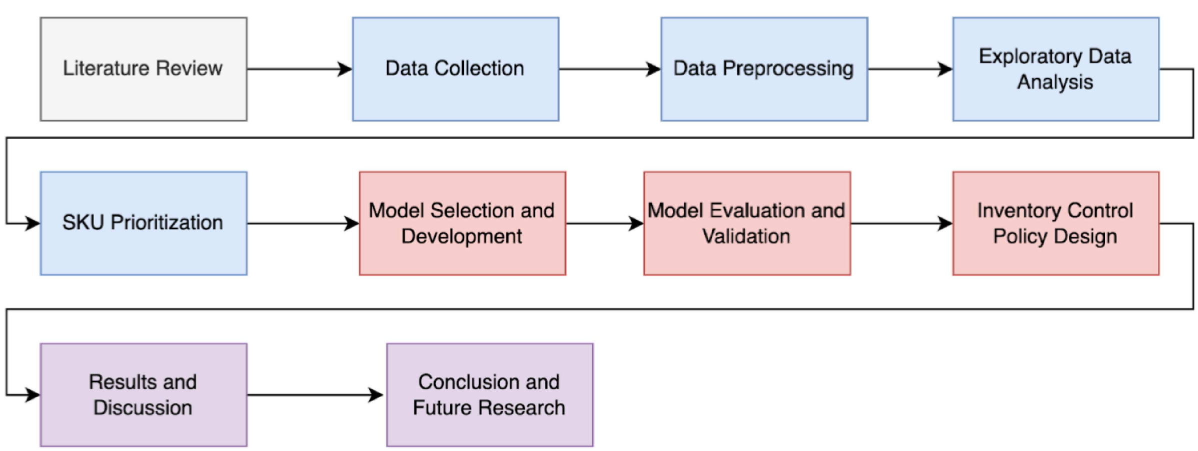

Figure 3-1 Methods Overview

3.1. Data preprocessing

This sub-section presents the various methodologies employed to prepare and process the data for subsequent analysis. This includes examining the techniques used to address data quality concerns, such as missing values, outliers, and inconsistencies. A brief description of the dataset's characteristics and relevant features is provided to offer context for the preprocessing steps and their significance in the overall forecasting process.

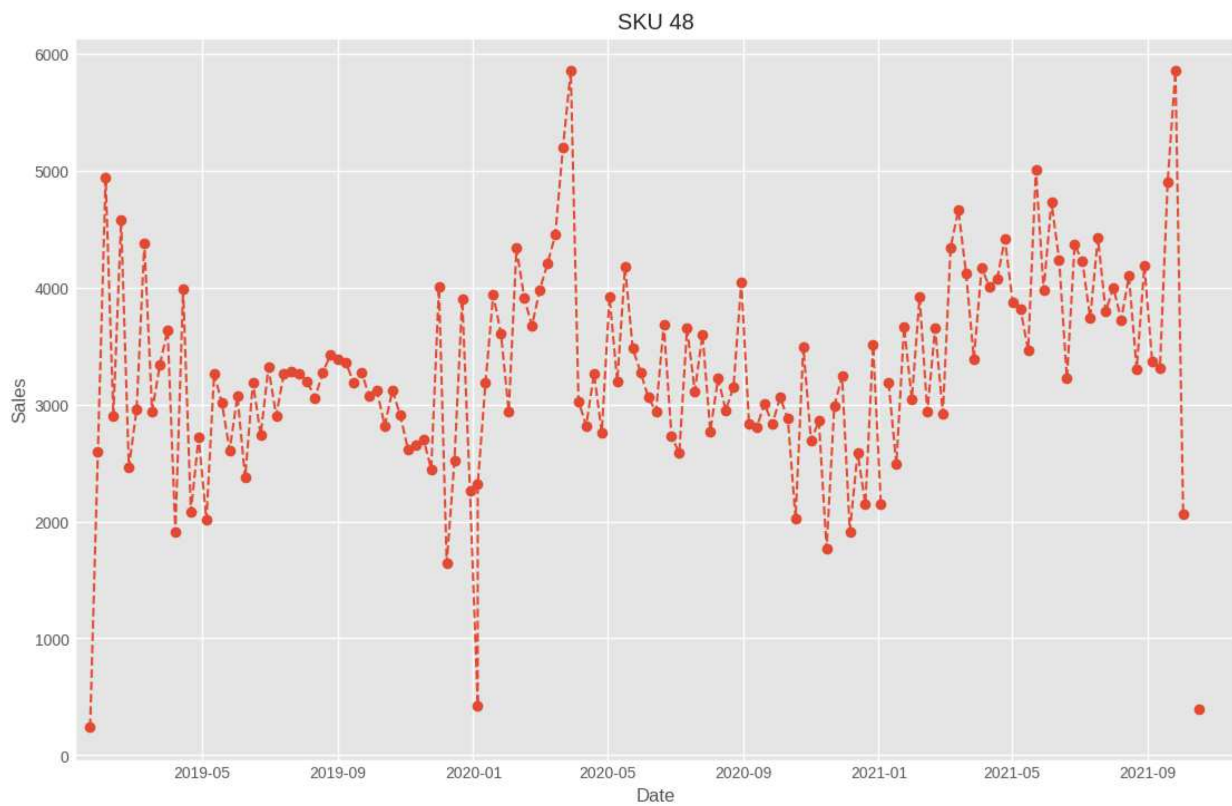

Figure 3-2 Example of historical demand data before preprocessing

The dataset was compiled from the 3PL provider and contained historical sales figures for the SKUs that the brand manages. After removing duplicate records and cleaning the data, the dataset contained 405 SKU sales over the period spanning from 2019 to 2021. It is worth noting that the SKUs were anonymized to maintain the confidentiality of the data. An example of historical data for SKU48 can be seen in Figure 3-2, where outliers can be observed in certain periods alongside missing values in the end of 2021.

The data preprocessing approach can be visualized in a sequence of steps, as observed in Figure 3-3, where each stage contains multiple steps that the following sections will describe in detail.

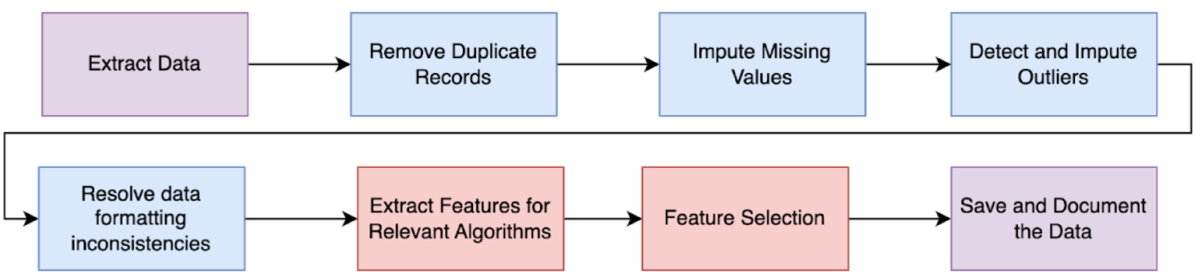

Figure 3-3 Data Preprocessing Approach

3.1.1. Resampling

Resampling is a vital technique to harmonize datasets with disparate time frequencies, rendering them suitable for comprehensive analysis. This process facilitates the completion of missing data points and effectively smooths out noise within the data. By employing resampling, the resulting models' accuracy can improve while circumventing issues arising from imbalanced data (Moniz et al., 2017). In the case of the dataset under consideration, a weekly resampling strategy was executed, accompanied by a corresponding aggregation to further refine the data for subsequent analysis.

3.1.2. Imputing missing values and outliers

The Interquartile Range (IQR) method was employed to identify and manage outliers within the dataset. It is a robust technique for detecting aberrant data points in the presence of a potential skewness (Jones, 2019). The process commenced with sorting the dataset in ascending order, followed by calculating the 1st and 3rd quartiles (Q1 and Q3) to define the central 50% of the data. The IQR was then computed by subtracting Q1 from Q3 ($IQR = Q3 - Q1$). Subsequently, the lower and upper bounds were determined using the formulae: $lowerbound = (Q1 - 1.5 \times IQR)$ and $upperbound = (Q3 + 1.5 \times IQR)$. These bounds served as the threshold values to identify potential outliers within the dataset. In iterating through each data point, values falling below the lower bound or above the upper bound were flagged as outliers, effectively allowing for their management in subsequent analysis (Dash et al., 2023). After detecting those outliers, they are removed from the dataset and replaced with the median of the data. The scale factor (1.5) is selected to approximate a 6σ range and is commonly used in practice (Chaudhary, 2020). Figure 3-4 showcases SKU1 historical data after applying the outliers and missing values imputation approach.

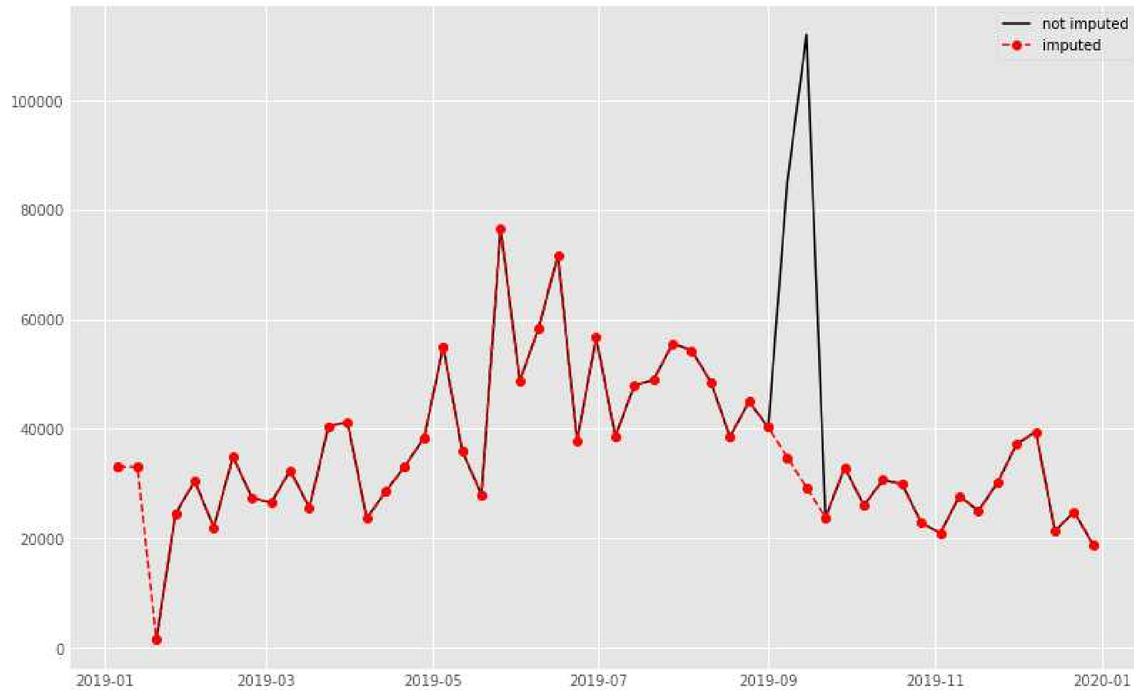

Figure 3-4 Illustration of outliers and missing value imputation on SKU1

3.1.3. SKU prioritization

Owing to constraints in the computational capacity that precluded the forecasting of all SKUs, a Pareto analysis was utilized to establish a hierarchy among the SKUs chosen for this study. This approach, grounded in the Pareto principle, focuses on identifying the top-selling SKUs that account for approximately 80% of the total sales, allowing for a more targeted and efficient forecasting effort (Y. Chen et al., 2008). The implementation of Pareto analysis in this study ensures that the demand forecasting is primarily directed toward the most significant SKUs in terms of sales contribution, thereby maximizing the effectiveness of the inventory control policy devised for the 3PL provider. 96 out of 405 SKUs contributed to 80% of the total sales.

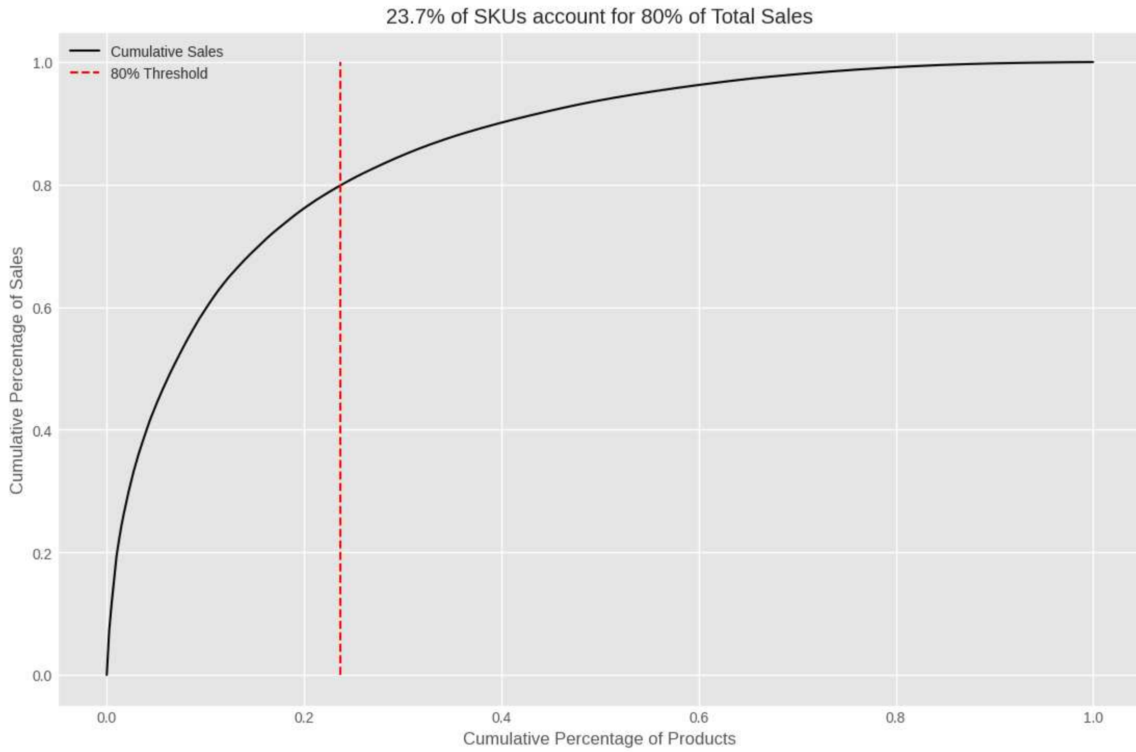

Figure 3-5 SKU Prioritization

3.1.4. Feature engineering

Feature engineering plays a crucial role in enhancing the forecasting capabilities of the selected models. A set of features was generated from the historical data, including order_date, target, quarter, month, dayofyear, dayofmonth, lag1, and lag2. These features aim to capture various aspects of the time series data, thereby providing valuable input for the forecasting models (Joseph, 2022).

The lag features, lag1, and lag2, represent the sales from the previous week and two weeks prior. Incorporating these lagged variables allows the models to account for potential short-term dependencies or autocorrelations in the data, which is particularly important for time series forecasting. By considering recent sales history, the models can better understand the data dynamics and make more accurate predictions for future demand.

Other features, such as quarter, month, dayofyear, and dayofmonth, capture the cyclical patterns and seasonality in the data, which may be driven by factors such as holidays, promotions, or consumer behavior. Including these features enables the models to learn and exploit these patterns to generate more accurate forecasts. Table 3-1 provides a snapshot of one of the SKUs feature sets alongside the corresponding observations.

Table 3-1 Snapshot of features of SKU1 post engineering

order_date	target	quarter	month	dayofyear	dayofmonth	lag1	lag2
2019-02-17	34936	1	2	48	17	27372	26496
2019-02-24	27372	1	2	55	24	26496	32400
2019-03-03	26496	1	3	62	3	32400	25584
2019-03-10	32400	1	3	69	10	25584	40488
2019-03-17	25584	1	3	76	17	40488	41244
2019-03-24	40488	1	3	83	24	41244	23640

For tree-based models, such as Random Forests and XGBoost, importance scores are derived from the number of times the feature is used to split the data across all trees. In Figure 3-6, a snapshot of the feature importance for SKU53 was selected to showcase its impact on feature selection. Exploring feature importance is beneficial for several reasons. It allows researchers to identify and focus on the most relevant features, potentially leading to simpler, more interpretable models with reduced computational complexity. Further, it enables the discovery of previously unknown relationships between features and the target variable, which can inform further research or domain-specific insights. In this dataset, the feature importance changes from one SKU to another, which indicates the inherent differences in the nature of the SKU demand patterns.

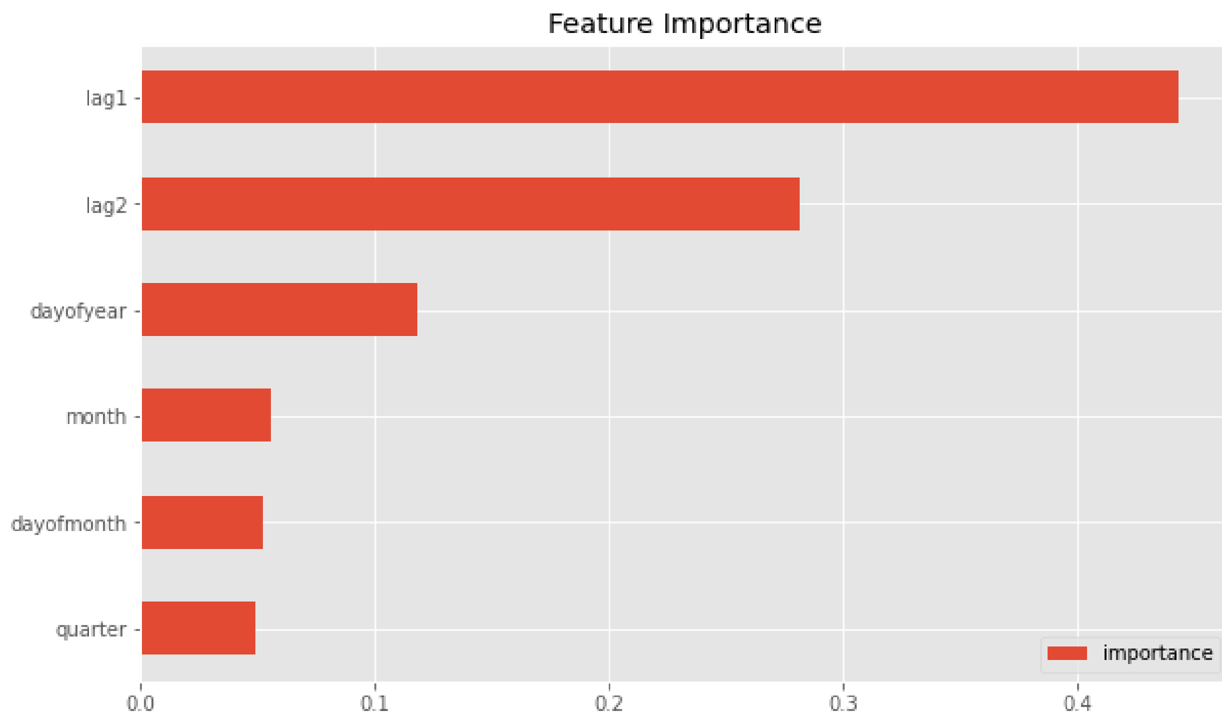

Figure 3-6 Feature Importance for SKU53

3.2. Forecasting

This section describes the methods employed to create the forecasting models. A discussion of the fundamental concepts behind these models is included, along with an analysis of the properties they offer in the context of forecasting.

3.2.1. Baseline benchmark

In demand forecasting, the Naïve Mean method represents a straightforward and expedient approach that can be employed for generating predictions (Hyndman & Athanasopoulos, 2018). As one of the most elementary techniques in the field, this method entails calculating the mean of all preceding observations in the dataset to forecast the subsequent data point. The simplicity and speed of the Naïve Mean method rendered it an excellent baseline to use in the context of this research.

3.2.2. Statistical models

AutoRegressive Integrated Moving Average (ARIMA) and Seasonal AutoRegressive Integrated Moving Average (SARIMA) models were employed to generate accurate and reliable demand forecasts. ARIMA models, a class of linear models, are widely recognized for their ability to capture the underlying structure of time series data by incorporating autoregressive (AR), moving average (MA), and differencing components (I). These components collectively allow ARIMA models to account for trends, seasonality, and noise within the data, resulting in improved forecasting performance (Peixeiro, 2022). The mathematical formulation of the ARIMA(p,d,q) model is given by the following equations (Adhikari & Agrawal, 2013):

$$\left(1-\sum_{i=1}^{p}\varphi_i L^i\right)(1-L)^d\, y_t = \left(1+\sum_{j=1}^{q}\vartheta_j L^j\right)\varepsilon_t$$

Where:

- y_t is the actual value, and ε_t is the random error at time period t.
- $p, d, and\ q$ refer to the order of the model's AR, I, and MA parts, respectively.
- L refers to lag operator and is defined as $Ly_t = y_{t-1}$.
- φ_i and $\vartheta_j$are model parameters $(i = 1,2,\dots,p)$, $(j = 1,2,\dots,p)$.

SARIMA models extend the capabilities of ARIMA models by explicitly incorporating a seasonal component. In time series data exhibiting strong seasonality, SARIMA models can more effectively capture the recurring patterns, enhancing forecast accuracy. The mathematical formulation for $SARIMA(p,d,q)\times(P,D,Q)^s$ is given as (Adhikari & Agrawal, 2013):

$$\phi_P(L^s)\varphi_p(L)z_t\, y_t = \theta_Q(L^s)\theta_q\varepsilon_t$$

Where $z_t = y_t - y_{t-s}$ which refers to the seasonality differenced series, and s refers to the seasonality order of the series. For example, for a monthly time series, $s = 12$.

By integrating both ARIMA and SARIMA models into this research, a comprehensive analysis of the time series data can be conducted, accounting for various components that might influence the demand patterns. AutoARIMA was utilized to streamline the model selection process and ensure that the best-fitting ARIMA or SARIMA model was chosen for this study, AutoARIMA is an automated approach that identifies the optimal ARIMA or SARIMA model for a given time series by iteratively testing various combinations of AR, I, MA, and seasonal components. It selects the model with the Akaike Information Criterion (AIC) as the best fit for the data. This automated procedure simplifies the model selection process, minimizes the risk of human error in choosing appropriate model parameters, and may enhance the overall forecasting accuracy (Hyndman & Khandakar, 2008).

3.2.3. Machine learning models

Prophet, Random Forests, and XGBoost are machine learning techniques in time series forecasting that are widely adopted in the industry and don't require complex approaches in modeling the data.

Prophet is a flexible and robust time series forecasting tool developed by Facebook, designed to effectively handle data exhibiting seasonality, trends, and holidays or special events. As an additive model, Prophet combines multiple components, including trend, seasonality, and holiday effects, to generate accurate forecasts. It is designed to generate accurate forecasts with minimal user input, as it automatically identifies and adjusts for seasonality, trends, and holidays or special events. This is achieved through its additive model, which combines these components while requiring only a few essential parameters. Moreover, Prophet employs a Bayesian approach, combining default priors and data-driven parameter estimation, enabling it to produce reliable forecasts without the need for extensive parameter tuning. As a result, Prophet offers an efficient forecasting tool suitable for a wide range of applications in this study. It can deliver precise demand forecasts without requiring exhaustive parameter specifications. Prophet's mathematical formulation is given below (Taylor & Letham, 2018):

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

Assuming a constant growth rate (i.e. linear trend with changepoints), $g(t)$ is given as:

$$g(t) = (k + a(t)^{\intercal}\delta)t + (m + a(t)^{\intercal}\gamma)$$

Where:

- k is the growth rate
- $\delta \in \mathbb{R}^S,\ assuming\ S\ changepoints\ at\ times\ s_j, j = 1, \dots, S$
- $a(t) \in \{0,1\}^S, such\ that\ a_j(t) = \begin{cases} 1, & if\ t > s_j, \\ 0, & otherwise \end{cases}$
- m is the offset parameter
- $\gamma_j = -s_j\delta_j$

The seasonality component $s(t)$ is given as:

$$s(t) = \sum_{n=1}^{N} \left(a_n \cos\left(\frac{2\pi\, nt}{P}\right) + b_n \sin\left(\frac{2\pi\, nt}{P}\right)\right)$$

Where:

- The Smoothing prior $\beta \sim Normal(0, \sigma^2), where\ \beta = [a_1, b_1, \dots, a_N, b_N]^{\intercal}$

For the holidays and events component, $h(t)$, such as Thanksgiving or the Superbowl, Prophet approaches modeling these events as the following:

For each holiday i, let D_i represent past and future occurrences. Then:

$$Z(t) = [\mathbf{1}(t \in D_1), \dots, \mathbf{1}(t \in D_L)]$$

And taking $$h(t) = Z(t)\kappa, where \kappa \sim Normal(0, v^2)$$

Random Forests, a powerful ensemble learning technique, was also employed as a forecasting model in this study. The method constructs multiple decision trees, training each on a random subset of the available data and utilizing a random subset of features. The final prediction is derived by aggregating the individual trees' predictions, typically through averaging for regression problems (see Figure 3-7). This approach effectively reduces overfitting and enhances the model's generalization capabilities, improving overall forecasting accuracy.

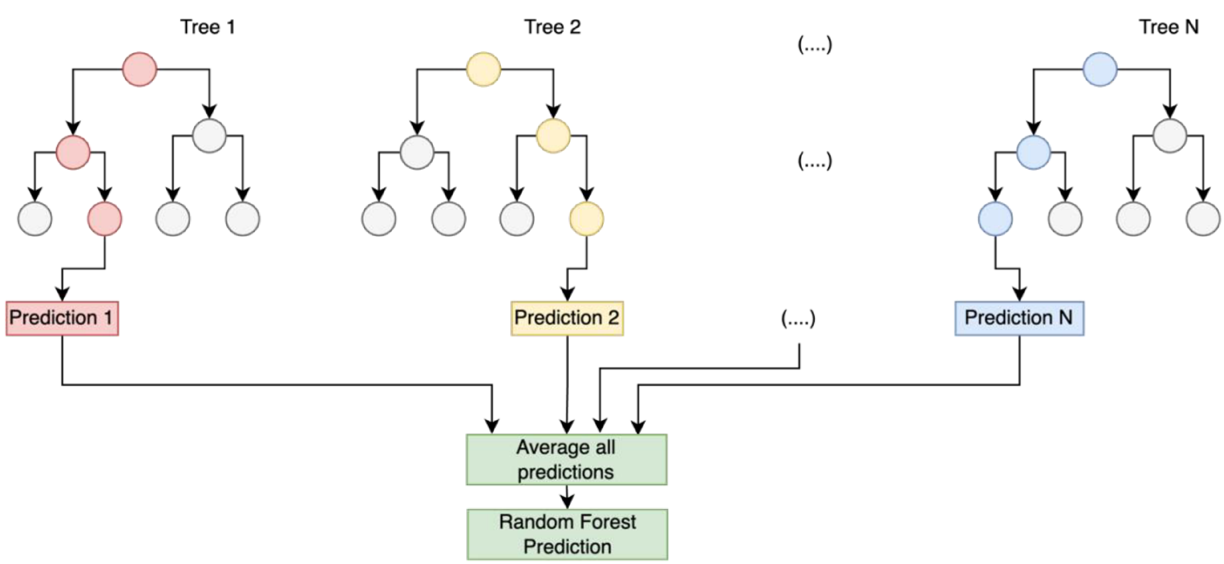

Figure 3-7 Illustration of Random Forests

The Binary Recursive Partitioning algorithm is given by (Zhang & Ma, 2012):

Algorithm 1: Binary Recursive Partitioning

$let\ D = \{(x_1, y_1), \ldots, (x_N, y_N)\}\ denote\ the\ training\ data\ \ with\ x_i = \left(x_{i,1}, \ldots, x_{i,p}\right)^T.$
1. *Start by placing all observations in a single node.*
2. *Recursively repeat the following steps for each unsplit node on all p predictors*
    a. *Find the best binary split among all binary splits on all p predictors.*
    b. *Split the node into two descendant nodes using the best split.*
3. *For predictions at x, pass x down the tree until it lands in a terminal node. Let k denote the training data in node k. Predicted values of the response variable are given by* $h'(x) = \frac{1}{n}\sum_{i=1}^{n} y_{k_i}$

The typical splitting criterion for regression is mean squared error denoted by

$$\mathrm{MSE}(y, \hat{y}) = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y_i})^2$$

The Random Forests algorithm is given by (Zhang & Ma, 2012):

Algorithm 2: Random Forests

$let\ D = \{(x_1, y_1), \ldots, (x_N, y_N)\}\ denote\ the\ training\ data,\ \ with\ x_i = \left(x_{i,1}, \ldots, x_{i,p}\right)^T. For\ j = 1\ to\ J:$
1. *Take a bootstrap sample* D_i *of size* N *from* D
2. *Using the bootstrap sample* D_i *as the training data, fit a tree using binary recursive partitioning (algorithm 1):*
    a. *Start by placing all observations in a single node.*
    b. *Until the stopping criterion is met, repeat the following steps recursively for each unsplit node:*
        i. *Choose m predictors randomly from the total of p accessible predictors.*
        ii. *Identify the most optimal binary split from all the potential binary splits on the m predictors, as determined in step (i).*
        iii. *Divide the initial node into two offspring nodes utilizing the split derived from step (ii) until the maximum number of nodes is reached.*
3. *To predict a new point x,* $f'(x) = \frac{1}{j}\sum_{j=1}^{J} h_j'(x)$

As an alternative to traditional time series models, Random Forests provide a non-linear and non-parametric approach to forecasting, making it a valuable addition to the array of models used in this research for demand prediction.

In addition to Prophet, this study also employs the XGBoost (eXtreme Gradient Boosting) algorithm for demand forecasting. XGBoost is a sophisticated and efficient implementation of gradient boosting machines, which build an ensemble of weak learners, typically decision trees, in a sequential manner. By iteratively fitting new trees to the residual errors of the previous trees, the model progressively improves its prediction performance. XGBoost incorporates regularization terms to control model complexity, preventing overfitting and ensuring robust generalization to new data (T. Chen & Guestrin, 2016). The XGBoost algorithm for regression is given by (Ibrahem Ahmed Osman et al., 2021):

Given a dataset $\{(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)$, the objective of gradient boosting is learn a function $F(x)$ that minimizes a loss function given by $L(y, F(x))$. The objective function is given as:

$$\mathrm{Obj(f)} = \sum_{\mathrm{i=1}}^{\mathrm{n}} L(y_i, F(x_i)) + \Omega(\mathrm{f})$$

Where $\mathrm{L}(\mathrm{y_i}, \mathrm{F}(\mathrm{x_i}))$is the loss function that measures the difference between the true labels $\mathrm{y_i}$ and the predicted labels ,$\mathrm{F}(\mathrm{x_i})$, and $\Omega(f)$is the regularization term defined as:

$$\Omega(f) = \gamma T + \frac{1}{2}\lambda\sum_{j=1}^{T} w_j^2$$

In this equation, T is the number of leaves in the tree, w_j is the prediction score assigned to the $j-th$ leaf, γ is a complexity control parameter, and λ is the L2 regularization parameter.

Gradient approximation in XGBoost is obtained as:

$$g_i = \frac{\partial L\big(y_i, F(x_i)\big)}{\partial F(x_i)}$$
$$h_i = \frac{\partial^2 L\big(y_i, F(x_i)\big)}{\partial F^2(x_i)}$$

Here, g_i represents the first-order gradient of the loss function $L\big(y_i, F(x_i)\big)$ with respect to the model output $F(x_i)$ *and* h_i represents the second-order gradient of the loss function with respect to the model output.

XGBoost trains weak learners iteratively. At each iteration, a new weak learner is added to the model to minimize the objective function.

$$f_m = \arg\min_f \sum_{i=1}^{n} L\big(y_i, F_{m-1}(x_i) + f(x_i)\big)$$

This equation represents the process of adding a new weak learner $f_m$to the model during the $m-th$ iteration to minimize the objective function. $F_{m-1}(x_i)$ is the current model, and $F_m(x_i) = F_{m-1}(x_i) + f_m(x_i)$ is the updated model after adding the new weak learner until a convergence criterion is met. After the training is complete, the final ensemble model is used for making predictions on new data by aggregating the predictions of all the trees in the ensemble:

$$\hat{y}(x) = \sum_{m=1}^{M} f_m(x)$$

Here, M is the total number of trees is in the ensemble, $f_m(x)$ is the prediction of the $m-th$ tree, and $\hat{y}(x)$ is the final prediction of the ensemble model.

XGBoost's ability to capture complex, non-linear relationships in the data and handle diverse feature sets makes it a valuable addition to the range of forecasting models employed in this research for generating accurate and reliable demand predictions (Niu, 2020).

3.2.4. Cross-validation

In this research, rolling K-fold cross-validation for time series is employed as a model validation technique to assess the forecasting performance of the chosen models. This method is specifically adapted for time series data to account for the inherent temporal dependencies and preserve the chronological order of observations.

The rolling K-fold cross-validation procedure involves partitioning the time series dataset into K contiguous, non-overlapping validation folds of approximately equal size. For each fold, the model is trained on a training set that consists of all the observations prior to the validation fold, maintaining the chronological order of the data. The model's forecasting performance is then evaluated on the validation fold. This process is iterated K times, with each fold serving as the validation set once and the model being retrained with an expanding training window that includes all previous observations.

Upon completion of the K iterations, the validation performance metrics, such as Mean Absolute Error (MAE) or Root Mean Squared Error (RMSE), are averaged across the K folds to obtain a single, overall measure of the model's forecasting accuracy. As seen in Figure 3-7, five-fold cross-validation is used on all skus for all of the forecasting models to provide a more robust estimate of the model's generalization performance on unseen data, ensuring that the selected models are well-suited for reliable and accurate demand forecasting.

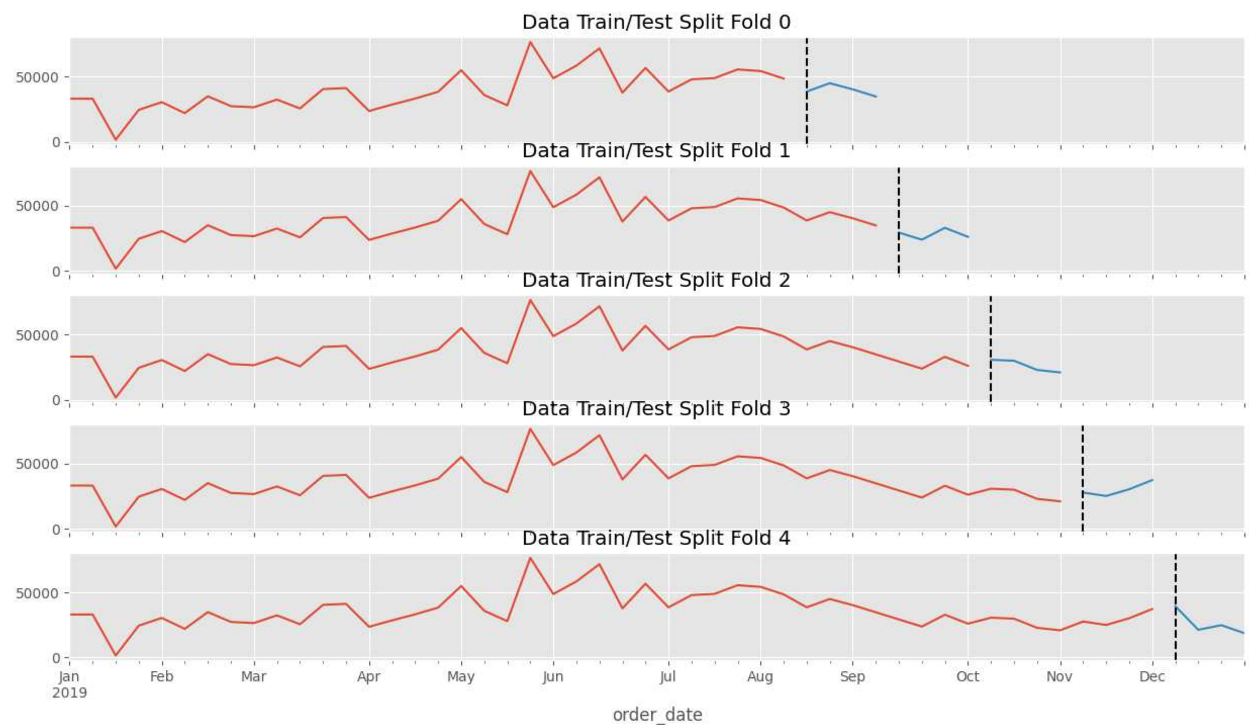

Figure 3-8 Five-Fold Rolling Cross Validation

3.2.5. Accuracy measures

In this study, three common performance metrics are employed to evaluate the forecasting accuracy of the chosen models on time series data: Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and Mean Absolute Error

(MAE). These metrics quantify the discrepancy between the actual and predicted values, providing a basis for model comparison and selection.

RMSE is a widely-used metric that measures the average magnitude of the prediction errors. This is computed by finding the square root of the average of the squared deviations between the real and estimated values. The RMSE penalizes larger errors more severely due to the squaring operation, making it sensitive to outliers and emphasizing the model's overall fit (Adhikari & Agrawal, 2013).

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(Actual_i - Predicted_i)^2}$$

MAPE is a scale-free metric that expresses the average absolute prediction error as a percentage of the actual values. It is calculated by dividing the absolute differences between the real and estimated values by the real values, summing the resulting percentages, and then dividing by the number of observations. MAPE is particularly useful for comparing the performance of models across datasets with different scales or units (Adhikari & Agrawal, 2013).

$$MAPE = \frac{1}{n}\sum_{i=1}^{n}\frac{|Actual_i - Predicted_i|}{|Actual_i|}$$

MAE measures the average magnitude of the prediction errors without considering their direction. It is calculated by taking the mean of the absolute differences between the real and estimated values. Unlike RMSE, MAE does not penalize larger errors as heavily, making it less sensitive to outliers and more focused on the central tendency of the errors.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|Actual_i - Predicted_i|$$

Each of these metrics offers unique insights into the forecasting performance of the models on the time series data, allowing for a comprehensive assessment of their accuracy and suitability for demand forecasting (Adhikari & Agrawal, 2013).

4. Results and discussion

This chapter presents the outcomes derived from the application of forecasting models to the historical data of 96 SKUs. An in-depth discourse on the inventory policy informed by these outcomes is also provided.

4.1. Battle of the models

The analysis of the forecasting models' performance revealed that the XGBoost model consistently outperformed all other models in 2019, 2020, and 2021 (see table 4-1). This finding suggests that XGBoost may be the most suitable choice for demand forecasting for the 3PL company going forward. The superiority of the XGBoost model can be attributed to its ability to capture complex patterns and non-linear relationships in the data and its robustness against overfitting due to the implementation of gradient boosting and regularization techniques. This is evident in the 2020 and 2021 results, where COVID-19 may have presented a significant challenge for other models in detecting the underlying pattern in the data, XGBoost successfully outperformed the rest of the models.

The unprecedented nature of the pandemic increased demand volatility across several categories of products (Moosavi et al., 2022), which makes it difficult for forecasting models to capture sudden shifts in demand, leading to potential inaccuracies in predictions. Additionally, historical data patterns may not provide an accurate representation of the current or future state of affairs, causing forecasting models to struggle in making reliable predictions.

The consistent performance of the XGBoost model across the three years underlines its adaptability and reliability in capturing the underlying trends and fluctuations in the retail demand data. This is especially noteworthy given the potential variations in market conditions, pandemics, consumer preferences, and other factors that could impact demand patterns over time. The strong performance of XGBoost in the face of these changes demonstrates its value as a dependable forecasting tool for FMCG inventory control and management.

Table 4-1 Median MAPE Results

Period	Baseline	Auto ARIMA	Prophet	Random Forest	XGB 0 Lags	XGB 1 Lag	XGB 2 Lags
2019	24.40%	22.81%	26.10%	22.54%	**18.50%**	18.80%	20.00%
2020	30.70%	34.00%	28.50%	32.00%	26.60%	**25.70%**	26.50%
2021	48.20%	38.50%	32.70%	37.54%	**16.50%**	17.80%	17.90%

The analysis of the relationship between the Coefficient of Variation (CV) and the Mean Absolute Percentage Error (MAPE) for all SKUs revealed a consistent linear relationship between these two metrics. This finding provides insights into the underlying dynamics between demand variability and forecast accuracy, and has potential implications for inventory management and control in the retail sector.

Figure 4-1 showcases the relationship between the coefficient of variation ($\frac{\mu}{\sigma}$) The linear relationship between CV and MAPE suggests that as demand variability increases (higher CV), the forecasting error (higher MAPE) also increases proportionally. To test this hypothesis, we perform a regression analysis for all the datasets with the goal of fitting a line that describes the relationship between them.

Using ordinary least-squares, where we estimate a linear model that minimizes the following objective function:

$$\mathrm{MSE}(y, \hat{y}) = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y_i})^2$$

Such that $MAPE = \beta_1 CV$, we get the following results:

Table 4-2 Regression Analysis for CV vs MAPE

Data	R^2	β_1	$p - value$	Correlation Coefficient for CV/MAPE
2019	0.93	0.74	<0.001	0.85
2020	0.97	0.68	<0.001	0.88
2021	0.82	0.57	<0.001	0.71

It is worth noting that the intercept β_0 was statistically insignificant when fitted in the model. Therefore, it was removed from the model. The R^2 value is a measure of how well the linear regression model fits the observed data. It signifies the fraction of variability in the outcome variable that can be accounted for by the predictor variable (Hayter, 2012). For example, in 2019 data, an R^2 value of 0.93 means that 93% of the variability in the dependent variable (MAPE) can be explained by the independent variable (CV). This suggests a strong fit of the model to the data. The β_1 value is the coefficient for the independent variable in the linear regression model. In 2019 data, it means that for every unit increase in the CV, MAPE is expected to increase by 0.74 units, on average, holding all other factors constant. This value gives an understanding of both the intensity and the direction of the connection between the two variables. Moreover, the p-value is used to test the null hypothesis that there is no significant relationship between the independent and dependent variables. A p-value of <0.001 for all years of data indicates that the relationship between CV and MAPE is highly significant, and the null hypothesis can be rejected. In other words, there is strong evidence that CV has a statistically significant effect on MAPE.

Furthermore, the correlation coefficient, also known as Pearson's correlation coefficient, measures the linear relationship's strength and direction between two variables. A correlation coefficient of 0.85 in 2019 data indicates a strong, positive relationship between the independent and dependent variables. This suggests that as CV increases, MAPE also tends to increase, and the relationship is fairly linear. The correlation in 2021 decreases compared to previous years, which indicates a weaker linear relationship between CV and MAPE compared to previous years. Nonetheless, the presence of a consistent linear relationship between CV and MAPE across all SKUs reinforces the notion that demand variability is a key factor influencing forecast accuracy and should be considered when evaluating the performance of forecasting models.

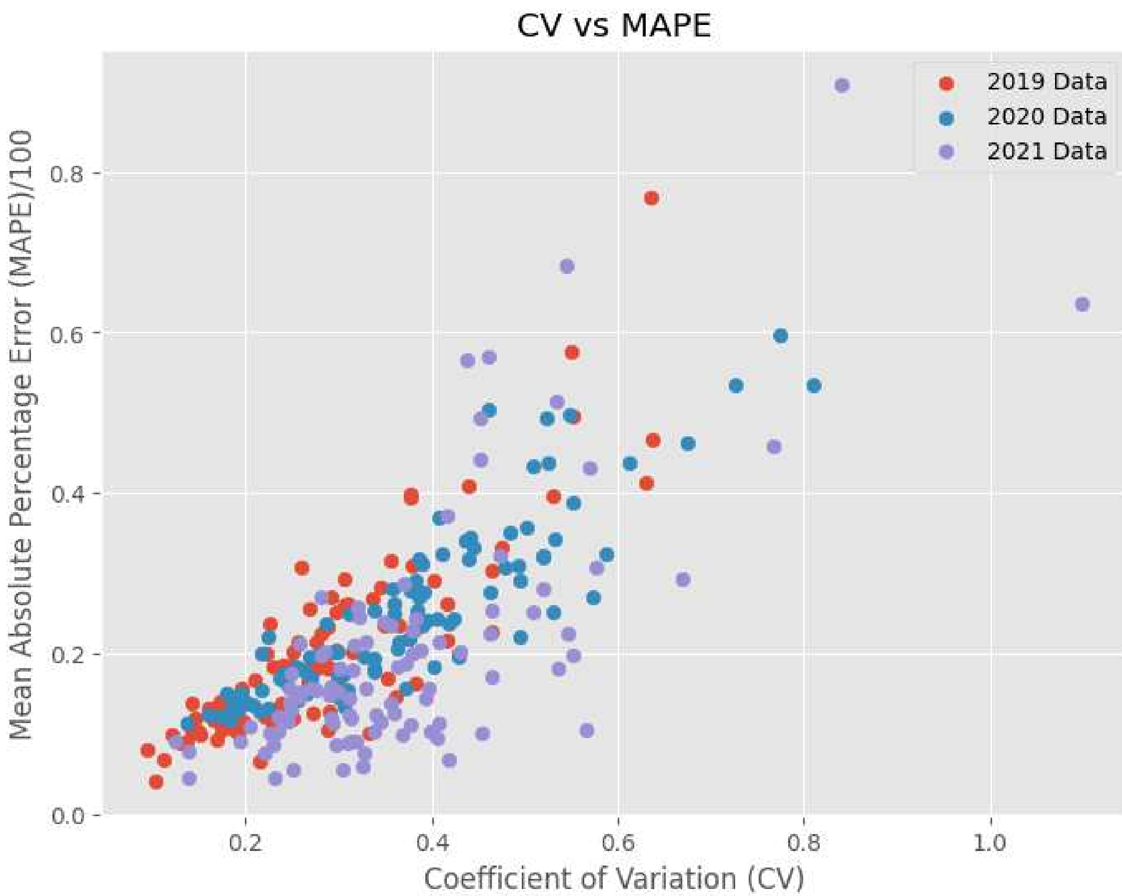

Figure 4-1 Coefficient of Variation vs. MAPE

This contributes to the understanding of the interplay between demand variability and forecast accuracy in the context of this 3PL supply chain demand. The findings highlight the importance of accounting for demand variability when designing inventory control strategies, as higher demand variability can lead to increased forecasting errors and consequently, suboptimal inventory decisions. By identifying the relationship between CV and MAPE, practitioners can better assess the performance of their forecasting models in the face of varying demand patterns and make informed adjustments to their inventory control policies. Future research could explore methods to improve forecast accuracy for SKUs with high demand variability, as well as the potential benefits of incorporating demand variability measures into inventory control policies to enhance their effectiveness.

4.2. Steering inventory policy with a forecast-driven approach

Given the results of the demand forecasts, tailoring an effective inventory policy for the FMCG sector entails taking into consideration the unique characteristics and requirements of the industry. Several key factors should be considered in the development of an inventory policy, such as forecast accuracy, safety stock levels, and order quantities. These factors play a critical role in determining the overall performance and efficiency of the inventory management system.

4.2.1. The interplay of demand forecasting and total cost

The total cost equation is denoted as the following:

$$TC(Q) = cD + c_t\left(\frac{D}{Q}\right) + c_e\left(\frac{Q}{2} + k\sigma_L + DL\right) + B_{SO}\left(\frac{D}{Q}\right)\Pr[SO]$$

Where:

- c is the unit cost
- D is the annual demand
- Q is the order quantity
- c_t is the setup cost
- $c_e$is the inventory holding cost
- k is the safety factor
- $\sigma_L$is demand variability
- L is lead time
- B_{SO} is stockout cost
- $\Pr[SO]$ is the probability of stockout

The terms in the total cost equation can be summarized as Purchase cost (cD), the setup cost ($c_t\left(\frac{D}{Q}\right)$), inventory holding cost ($c_e\left(\frac{Q}{2}+k\sigma_L+DL\right)$), and is the stockout cost ($B_{SO}\left(\frac{D}{Q}\right)\mathrm{P\,r}[SO]$)

The interdependence of the total cost equation and demand forecasts can be observed in each term of the equation. The setup cost experiences a substantial increase as a consequence of deviations in demand forecasts. With regard to inventory holding cost, the accuracy of demand forecasts plays a pivotal role in determining the safety stock levels, which subsequently leads to a proportional increase in holding costs. Concerning stockout cost, the quantities ordered are contingent upon the demand estimations generated by the forecasting models, which may subsequently result in larger and more frequent stockouts.

The sensitivity of the total relevant costs is estimated with the formula seen below:

$$\frac{TRC(Q_F^*)}{TRC(Q_A^*)}=\left(\frac{1}{2}\right)(\sqrt{\frac{D_A}{D_F}}+\sqrt{\frac{D_F}{D_A}}$$

Where D_A refers to the actual annual demand for the product. While D_F denotes to the incorrect forecast which is estimated by the forecasting models. In the context of the EOQ model, which operates at an annual level, it is important to address the discrepancy between the model's time frame and the weekly forecasting accuracy used in this study. To bridge this gap, an assumption is made that the forecasting error remains consistent across the year. This assumption is both realistic and reasonable since forecasting errors can be both positive and negative, which tend to offset each other over time. Assuming a consistent forecasting error across the year implies that these offsets are taken into account, resulting in a more accurate representation of the overall demand uncertainty.

Moreover, making this assumption simplifies the calculation of the EOQ and provides a more straightforward way of incorporating the weekly forecasting accuracy into the model. While this may introduce some level of approximation, it is a reasonable trade-off to make in order to facilitate the practical application of the model in the context of inventory management.

An illustration of sensitivity of the total cost for SKU96 is presented in Table 4-2, where a comparison between the best performing model and the baseline reveals a 10.4% discrepancy in the total relevant cost. This finding underscores the influence of forecasting forecasting accuracy on the total relevant cost.

Table 4-3 TRC for SKU96

SKU96	RMSE	MAPE	TRC
Baseline	492	61.80%	1.118
XGB	129	9%	1.001

As consolidators do not exert direct control over the order quantity from suppliers, the level of safety stock significantly affects their fulfillment of the On-Time In-Full (OTIF) Key Performance Indicator (KPI) with retailers. Consequently, this study prioritizes the Continous Review Policy for investigation. This policy facilitates the establishment of safety stock levels and determines the threshold to trigger the replenishment order. Hence, it provides a benchmark for the 3PL research partner to ascertain the optimal safety stock and order quantity necessary to achieve a targeted OTIF of 98%.

4.2.2. Tailoring the continuous review (s,Q) policy

The Continuous Review Policy, also referred to as the (s,Q) policy, is a widely recognized inventory control approach employed in various industries, including retail. Under this policy, inventory levels are continually monitored, and replenishment orders are initiated when the inventory position reaches or falls below a predefined reorder point s. Upon reaching the reorder point, an order of a fixed quantity 'Q' is placed to replenish the inventory, where the formula for the reorder point is given as the following:

$$s = \mu_{DL} + k\sigma_{DL}$$

Where μ_{DL} is the average demand over lead time, and $k\sigma_{DL}$represents the safety stock level.

The order quantity 'Q' in the (s,Q) policy is typically determined based on the Economic Order Quantity (EOQ) model, which considers factors such as ordering cost, holding cost, and demand rate. The order quantity is estimated in practice with the economic order quantity (EOQ), which is given by the formula below:

$$Q^* = \sqrt{\frac{2c_t D}{c_e}}$$

As evident from the formula, the order quantity is directly dependent on demand estimation, which consequently impacts each component of the total cost equation.

The reorder point s is determined by considering the lead time demand, which accounts for both the average demand during the lead time and the safety stock required to mitigate the risk of stockouts. The safety stock level depends on the desired service level and the variability of both demand and lead time. Assuming a lead time of one week, and a service level target of 98% based in the OTIF KPI, the safety stock term of that equation is estimated with the forecast error (RMSE) (Silver et al., 2016) of the SKU forecast with the following equation:

$$Safety\ Stock = k \times RMSE$$

In order to demonstrate the influence of forecasting accuracy on the continuous review policy, the application of the XGBoost model forecast to SKU96 results in a substantial reduction of holding cost by 73%, as depicted in Table 4-3. This exemplifies the potential benefits of employing accurate forecasting models in the context of inventory management under the (s,Q) policy.

Table 4-4 Forecasting accuracy's impact on Safety Stock holding cost

SKU96	RMSE	MAPE	Safety Factor	Safety Stock	SS Holding Cost
Baseline	492	61.8%	2.326	1144	$34,332
XGBoost	129	9%	2.326	300	$9,002

The Continuous Review Policy offers several advantages, including real-time inventory monitoring, responsiveness to demand fluctuations, and the ability to adapt to changes in lead times. However, the policy also requires more frequent inventory tracking and administrative efforts compared to periodic review policies. Despite these trade-offs, the (s,Q) policy remains an important inventory control approach for businesses seeking to optimize their inventory management and maintain high service level

5. Conclusion and future work

5.1. Conclusion

This thesis has developed an effective inventory control policy for a third-party logistics (3PL) provider by forecasting demand using historical data and incorporating the forecasting error into a continuous review inventory policy. A range of forecasting models were employed, including traditional time series models such as ARIMA, SARIMA as well as advanced machine learning algorithms like Prophet, Random Forests and XGBoost. These models were chosen for their ability to capture complex patterns and trends in the data, ensuring accurate and reliable demand forecasts.

Through rigorous model validation using rolling K-fold cross-validation for time series, the study identified the most suitable models for this task, based on their generalization performance on unseen data. The forecasting accuracy of the models was further evaluated and compared using performance metrics such as RMSE, MAPE, and MAE, providing a comprehensive assessment of their suitability for demand forecasting in the context of 3PL providers.

The COVID-19 pandemic have presented a significant challenge for models in detecting the underlying pattern in the data in 2020. The pandemic has led to abrupt changes in consumer behavior, supply chain dynamics, and market conditions (Moosavi et al., 2022). As a result, historical data patterns may not provide an accurate representation of the current or future state of affairs, causing forecasting models to struggle in making reliable predictions.

Nonetheless, the evaluation of various forecasting models demonstrated that XGBoost consistently outperformed all other models for the years 2019, 2020, and 2021. This suggests that XGBoost may be the optimal choice for demand forecasting in the 3PL company's context. The consistency of XGBoost's performance highlights its adaptability and reliability in detecting underlying trends and variations in

FMCG demand data, even in the face of unprecedented challenges such as the COVID-19 pandemic.

The development of a continuous review policy based on the forecasting error allowed for a more dynamic and responsive inventory management approach. By incorporating the variability in demand predictions, the proposed policy was better equipped to handle uncertainties and fluctuations in the actual demand, ensuring that stock levels were optimized and minimizing the risk of stockouts or excess inventory.

5.2. Future work

Considering the findings from this study, several opportunities for future work were identified to further advance the research in inventory management and demand forecasting for third-party logistics (3PL) providers.

- **Exploration of additional forecasting models:** While this study employed a diverse range of models, including ARIMA, SARIMA, Prophet, Random Forests, and XGBoost, future research could investigate the applicability and performance of other advanced forecasting techniques, such as deep learning models that include TCN (Bai et al., 2018), and N-BEATS (Oreshkin et al., 2020) or hybrid approaches that combine multiple models to enhance prediction accuracy such as LSTM+Prophet (Arslan, 2022).
- **Integration of external factors:** This study primarily focused on historical data to forecast demand. Future work could incorporate external factors, such as

macroeconomic indicators, industry trends, or weather data, to better understand their impact on demand patterns and improve the accuracy of the forecasting models.

- **Investigation of alternative inventory control policies:** While this study developed a continuous review policy based on the forecasting error, future research could explore other inventory control policies, such as periodic review or order-up-to-level policies, and evaluate their effectiveness in different contexts or industries.
- **Longitudinal analysis:** This study relied on a snapshot of historical data to develop the inventory control policy. Future research could benefit from a longitudinal analysis, tracking the performance of the proposed policy over an extended period and identifying potential changes in demand patterns or model performance over time.
- **Evaluation of implementation challenges:** Future work could investigate the practical implications of implementing the proposed inventory control policy, including the required data infrastructure, computational resources, and organizational changes, as well as the potential barriers to adoption and strategies to overcome them.

By addressing these areas of future work, researchers can continue to advance the field of inventory management and demand forecasting for 3PL providers, ultimately contributing to more efficient and effective supply chain operations in both academia and industry.

Bibliography

Adhikari, R., & Agrawal, R. K. (2013). *An Introductory Study on Time Series Modeling and Forecasting* (arXiv:1302.6613). arXiv. https://doi.org/10.48550/arXiv.1302.6613

Arslan, S. (2022). A hybrid forecasting model using LSTM and Prophet for energy consumption with decomposition of time series data. *PeerJ Computer Science*, *8*, e1001. https://doi.org/10.7717/peerj-cs.1001

Bai, S., Kolter, J. Z., & Koltun, V. (2018). *An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling* (arXiv:1803.01271). arXiv. https://doi.org/10.48550/arXiv.1803.01271

Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Chaudhary, S. (2020, August 26). *Why "1.5" in IQR Method of Outlier Detection?* Medium. https://towardsdatascience.com/why-1-5-in-iqr-method-of-outlier-detection-5d07fdc82097

Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. https://doi.org/10.1145/2939672.2939785

Chen, Y., Li, K. W., & Liu, S. (2008). A comparative study on multicriteria ABC analysis in inventory management. *2008 IEEE International Conference on Systems, Man and Cybernetics*, 3280–3285. https://doi.org/10.1109/ICSMC.2008.4811802

Dash, Ch. S. K., Behera, A. K., Dehuri, S., & Ghosh, A. (2023). An outliers detection and elimination framework in classification task of data mining. *Decision Analytics Journal*, *6*, 100164. https://doi.org/10.1016/j.dajour.2023.100164

Falatouri, T., Darbanian, F., Brandtner, P., & Udokwu, C. (2022). Predictive Analytics for Demand Forecasting – A Comparison of SARIMA and LSTM in Retail SCM. *Procedia Computer Science*, *200*, 993–1003. https://doi.org/10.1016/j.procs.2022.01.298

Fildes, R., Ma, S., & Kolassa, S. (2022). Retail forecasting: Research and practice. *International Journal of Forecasting*, *38*(4), 1283–1318. https://doi.org/10.1016/j.ijforecast.2019.06.004

Global Executives Survey: Impact of Seasonality in FMCG industry. (2017). *Market Research Reports & Consulting | GlobalData UK Ltd.* https://www.globaldata.com/store/report/global-executives-survey-impact-of-seasonality-in-fmcg-industry/

Hayter, A. J. (2012). *Probability and statistics for engineers and scientists* (4th ed.). Brooks/Cole, Cengage Learning.

Hyndman & Athanasopoulos. (2018). *Forecasting: Principles and Practice (2nd ed)* (2nd Edition). https://otexts.com/fpp2/

Hyndman, R. J., & Khandakar, Y. (2008). Automatic Time Series Forecasting: The forecast Package for R. *Journal of Statistical Software*, *27*, 1–22. https://doi.org/10.18637/jss.v027.i03

Ibrahem Ahmed Osman, A., Najah Ahmed, A., Chow, M. F., Feng Huang, Y., & El-Shafie, A. (2021). Extreme gradient boosting (Xgboost) model to predict the groundwater levels in

Selangor Malaysia. *Ain Shams Engineering Journal*, *12*(2), 1545–1556. https://doi.org/10.1016/j.asej.2020.11.011

Jones, P. R. (2019). A note on detecting statistical outliers in psychophysical data. *Attention, Perception & Psychophysics*, *81*(5), 1189–1196. https://doi.org/10.3758/s13414-019-01726-3

Joseph, M. (2022). *Modern time series forecasting with Python: Explore industry-ready time series forecasting using modern machine learning and deep learning* (1st ed.). Packt Publishing, Ltd.

Kolari, J. W., & Sanz, I. (2022). *Forecasting Bank Capital Ratios Using the Prophet Model by Facebook* (SSRN Scholarly Paper No. 4141575). https://doi.org/10.2139/ssrn.4141575

Kumar Jha, B., & Pande, S. (2021). Time Series Forecasting Model for Supermarket Sales using FB-Prophet. *2021 5th International Conference on Computing Methodologies and Communication (ICCMC)*, 547–554. https://doi.org/10.1109/ICCMC51019.2021.9418033

Moniz, N., Branco, P., & Torgo, L. (2017). Resampling strategies for imbalanced time series forecasting. *International Journal of Data Science and Analytics*, *3*(3), 161–181. https://doi.org/10.1007/s41060-017-0044-3

Moosavi, J., Fathollahi-Fard, A. M., & Dulebenets, M. A. (2022). Supply chain disruption during the COVID-19 pandemic: Recognizing potential disruption management strategies. *International Journal of Disaster Risk Reduction*, *75*, 102983. https://doi.org/10.1016/j.ijdrr.2022.102983

Narula, Anupam. (2012). *The profit margin squeeze: Structural strategies for consumer product companies*. Deloitte Insights.

https://www2.deloitte.com/content/www/xe/en/insights/industry/consumer-products/the-profit-margin-squeeze-structural-strategies-for-consumer-product-companies.html

Niu, Y. (2020). Walmart Sales Forecasting using XGBoost algorithm and Feature engineering. *2020 International Conference on Big Data & Artificial Intelligence & Software Engineering (ICBASE)*, 458–461. https://doi.org/10.1109/ICBASE51474.2020.00103

Oreshkin, B. N., Carpov, D., Chapados, N., & Bengio, Y. (2020). *N-BEATS: Neural basis expansion analysis for interpretable time series forecasting* (arXiv:1905.10437). arXiv. https://doi.org/10.48550/arXiv.1905.10437

Papacharalampous, G. A., & Tyralis, H. (2018). Evaluation of random forests and Prophet for daily streamflow forecasting. *Advances in Geosciences*, *45*, 201–208. https://doi.org/10.5194/adgeo-45-201-2018

Peixeiro, M. (2022). *Time series forecasting in Python*. Manning Publications Co.

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). *Inventory and Production Management in Supply Chains*. Taylor & Francis Group. http://ebookcentral.proquest.com/lib/mit/detail.action?docID=4771754

Taylor, S. J., & Letham, B. (2018). Forecasting at Scale. *The American Statistician*, *72*(1), 37–45. https://doi.org/10.1080/00031305.2017.1380080

Vairagade, N., Logofatu, D., Leon, F., & Muharemi, F. (2019). Demand Forecasting Using Random Forest and Artificial Neural Network for Supply Chain Management. In N. T. Nguyen, R. Chbeir, E. Exposito, P. Aniorté, & B. Trawiński (Eds.), *Computational Collective Intelligence* (pp. 328–339). Springer International Publishing. https://doi.org/10.1007/978-3-030-28377-3_27

Veiga, C., Veiga, C. R., Catapan, A., Tortato, U., & Silva, W. (2014). Demand forecasting in food retail: A comparison between the Holt-Winters and ARIMA models. *WSEAS Transactions on Business and Economics*, *11*, 608–614.

Wagner, S. M., & Sutter, R. (2012). A qualitative investigation of innovation between third--party logistics providers and customers. *International Journal of Production Economics*, *140*(2), 944–958. https://doi.org/10.1016/j.ijpe.2012.07.018

Wang, Z., Hong, T., & Piette, M. A. (2020). Building thermal load prediction through shallow machine learning and deep learning. *Applied Energy*, *263*, 114683. https://doi.org/10.1016/j.apenergy.2020.114683

Zhang, C., & Ma, Y. (Eds.). (2012). *Ensemble Machine Learning: Methods and Applications*. Springer. https://doi.org/10.1007/978-1-4419-9326-7