Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 238 (2024) 503–510

Procedia Computer Science

www.elsevier.com/locate/procedia

The 15th International Conference on Ambient Systems, Networks and Technologies (ANT)
April 23-25, 2024, Hasselt, Belgium

Drop Trailer Forecasting in Volatile Networks

Alex Carroll[a], Troy Egar[a], Thomas Koch[a], Elenna R. Dugundji[a,*]

[a]Center for Transportation and Logistics, Massachusetts Institute of Technology, 1 Amherst St., Cambridge, MA, 02142, USA

Abstract

This paper delves into the critical aspect of demand forecasting within the broader context of optimizing drop trailer management in volatile networks, with a specific focus on a large pallet manufacturer's supply chain operation. The study underscores the importance of accurate demand forecasting as a foundational element for informing subsequent optimization models. The main objective is to enhance our sponsor company's supply chain efficiency by accurately predicting future trailer requirements, which is crucial for the subsequent development of an effective inventory control and optimization model. This research utilizes forecasting methods like Gradient Boosted Trees and highlights the challenges of traditional forecasting methods in the context of our sponsor company's complex and dynamic supply chain network. The demand forecast is meant to inform the development of optimization models crucial for ensuring the effective allocation and management of trailer assets, ultimately leading to cost reductions and improved service levels within our sponsor company's network. The study contributes significantly to supply chain management literature by showcasing the application of sophisticated forecasting techniques in a real-world context and setting the stage for the development of robust optimization models in the domain of drop trailer management.

Keywords: Drop Trailer Management, Machine Learning, Demand Forecasting, Gradient Boosted Trees, Inventory Management

1. Introduction

Drop trailer programs are an important transportation modality for retailers, manufacturers, and carrier partners within the supply chain. A typical "live load" shipment involves a carrier arriving at a shipper with its own empty

* Corresponding author. Tel.: +1-617-258-6048
E-mail address: e.r.dugundji@mit.edu

10.1016/j.procs.2024.06.053

trailer, loading that empty trailer upon arrival, and then delivering that cargo to a receiver that unloads the freight upon arrival (Liu and Ker, 2022). In contrast, drop trailer programs have different requirements, benefits, and challenges. With a drop load, the cargo is loaded into a trailer that is already at the shipper's premises prior to any carrier coming to pick it up. The driver picking up the load arrives without any trailer attached (or simply exchanges trailers in a "drop and hook" procedure), hooks up to the preloaded trailer, and then takes the cargo to the delivery location. This modality can speed up the loading process and reduce dwell time for the carrier partner at pickup, helping to reduce detention charges, which are additional payments to the carrier partner for loading or unloading times that exceed an industry standard. This approach promotes "shipper of choice" status, which directly impacts a company's bottom line (Liu and Ker, 2022).

Our sponsor company utilizes such a program and has a unique position in the supply chain space that touches most shippers and retailers in the United States. The sponsor company is a pallet service and rental company that provides rentable pallets for shippers, and then picks up those pallets from an end delivery location, such as a retailer, for repair, reuse, and recycling. The physical flow of goods starts at a service center, where pallets are made and then delivered to manufacturers of goods in the United States. From there, the manufacturer packages and loads its goods onto those pallets, which are then shipped to its own customers.

The sponsor company then picks up the used pallets at the manufacturer's customer's location via its carrier partners and delivers them to a service center to be inspected and recycled into the system for further use. If done correctly, this "pooling" method, where pallets are shared with other companies and reused in a circular model, increases asset utilization and profitability for our sponsor. Concurrently, this saves manufacturers from holding excess pallet inventories, and they additionally benefit from the lower per-unit cost for pallets that our sponsor's scale can provide.

Larger retailers and distributors dealing with a high volume of pallets typically utilize a drop trailer program (Fankhauser and Li, 2019). In this program, our sponsor rents the use of dedicated carrier fleets and places their trailers at a customer's facility. The customer, or a third-party logistics provider (3PL), then loads these trailers. Once loaded, a request for freight pickup is generated, and the cost of each pickup is covered by the sponsor. This cost is influenced by fluctuations and seasonal variations in the US truckload transportation market, at both a nationwide and a regional level. Due to these factors, there exists a risk of an imbalance between supply and demand, leading to more freight shipments than the available carriers can handle. These imbalances can result in varying costs, service levels, delivery times, and pallet availability for customers across different US regions (Pickett, 2018).

A further complication is our sponsor's contractual obligations to customers. Due to the variety of manufacturers and customers our sponsor serves, there is a large variation in demand and service requirements. In many cases, the sponsor is charged a penalty by the minute if a customer stocks out on pallets during a planned manufacturing run. This can lead to a tendency to overstock customers with pallets, especially at the end of a week, to cover demand over the weekend. Driven in part by a misalignment in service center operating hours, the customer's operating hours, and carrier availability, the sponsor will overstock a customer on Friday to ensure they have enough pallets to make it through a weekend with high certainty. Overpooling, which refers to excessively supplying a customer with pallets, can result in shortages throughout the network. This may leave other customers without the pallets they need.

Given that there are over 500 service centers across North America and thousands of retailers, manufacturers, distributors, and other partners within the sponsor's supply chain network, optimizing their network for efficiency is a difficult, large-scale, and constantly evolving process. Extended delays experienced by retailers in the removal of large trailers and pallets from their facilities have the potential to adversely affect the relational dynamics between the retailer and the manufacturers; a customer's problem is a problem for the sponsor. The trailers our sponsor provides are a key asset necessary to maintain the integrity of their pooling model and become costly liabilities if not efficiently utilized. By effectively managing the allocations of drop trailers at a regional and customer level, ensuring sufficient levels of pallet inventory during customer operating hours, efficiently scheduling carriers, and accurately tracking costs, our sponsor can increase revenue, because it will be able to rent more pallets, avoid contractual penalties, reduce transportation costs, and improve relationships with both its customers and carrier partners.

2. Problem Statement and Research Questions

Our primary goal in this project is to refine our sponsor company's supply chain performance by enhancing the precision of demand forecasting within their drop trailer network. A key objective is to develop a sophisticated demand forecasting model that accurately captures the dynamic movement of goods throughout the network. This model will serve as the cornerstone for understanding and managing fluctuations in trailer requirements, which are critical for optimizing asset efficiency in the supply chain.

Achieving our goal of enhancing supply chain efficiency through innovative forecasting methods entails several crucial aspects:

- Developing an accurate predictive model that can anticipate future demand for trailers at various nodes in the supply chain, considering factors such as seasonal variations, market trends, and customer behavior.
- Integrating this demand forecast into our sponsor company's operational framework to enable more informed decision-making in areas such as trailer allocation, inventory management, and scheduling of pickups and deliveries.
- Utilizing the insights from the demand forecasting model to identify potential efficiencies and cost savings, particularly in reducing dwell times, minimizing idle pallets, and avoiding stock outs.

We hypothesize that a Gradient Boosting demand forecasting model (XGBoost), tailored to the sponsor company's operational context, will be capable of predicting future trailer requirements with precision and provide essential insights for improving key operational metrics within their supply chain. By comparing XGBoost with traditional models, we aim to demonstrate its superiority in predicting trailer requirements, which is crucial for optimizing our sponsor's asset allocation and operational efficiency. Our research aims to answer the following questions: How can XGBoost be effectively applied to forecast drop trailer demand in volatile networks? How does the predictive performance of XGBoost compare to traditional models like ARIMA and Simple Naïve in this specific context?

3. Scope: Project Goals and Expected Outcomes

The Gradient Boosting demand forecasting model (XGBoost) we propose is designed to provide the inputs necessary for a catered optimization model meant to minimize costs while respecting our sponsor's service requirements. The success of this model hinges on the quality of the data at our disposal and its ability to reflect the true dynamics of the sponsor's network operations. The insights generated by this model are expected to inform strategic decision-making, enhancing operational efficiency and effectiveness.

The process flow of our engagement with our sponsor company begins with defining the scope and problem, followed by meticulous data gathering and analysis. These initial stages are critical, as the model's recommendations will be heavily influenced by the quality of the input data. Following data preparation and validation, our focus will shift exclusively to developing the demand forecasting model. The development phase will be iterative, involving continuous interactions with the sponsor company. This collaboration is vital for refining the model, adjusting assumptions, and ensuring the relevance of the data included. This model will enable our sponsor company to respond swiftly and effectively to fluctuations in demand and supply, particularly concerning trailer availability and carrier scheduling.

4. State of the Practice

The forecasting and optimization of trailer assets play a critical role for shippers engaged in drop trailer circular economies. In the flow from a recycling center, to manufacturer, to retailer, and back to the recycling center, proper trailer allocation determines asset efficiency. Effective forecasting allows shippers to anticipate demand and allocate trailers at each node, ensuring that their goods are moved in the most efficient manner possible. Optimizing trailer assets involves strategically deploying these resources where they are most needed within a network. This process is largely dependent on accurate demand forecasting, which is essential for predicting future freight flows. Working in tandem, improvements in these techniques can lead to tangible improvements such as reduced transportation costs, improved delivery times, and enhanced service quality.

This fact is recognized by our sponsor company, which utilizes a drop trailer program within a circular economy but does not own trucking assets for moving its products outright. Instead, it outsources the movement of its freight to dedicated trucking companies, which provide x number of trucks and trailers on a contractual basis solely servicing the sponsor company's freight. The outsourcing of these assets is costly, and there can be pressure to have an oversupply of loaded trailers at manufacturers; if a manufacturer does not have adequate pallets to ship its goods, its entire operations will shut down. To prevent capital being tied up in an oversupply of outsourced assets, having an accurate, dynamic demand forecast can give confidence and ensure that assets are being utilized efficiently. An optimal scenario where demand and supply are perfectly matched is shown in Figure 1.

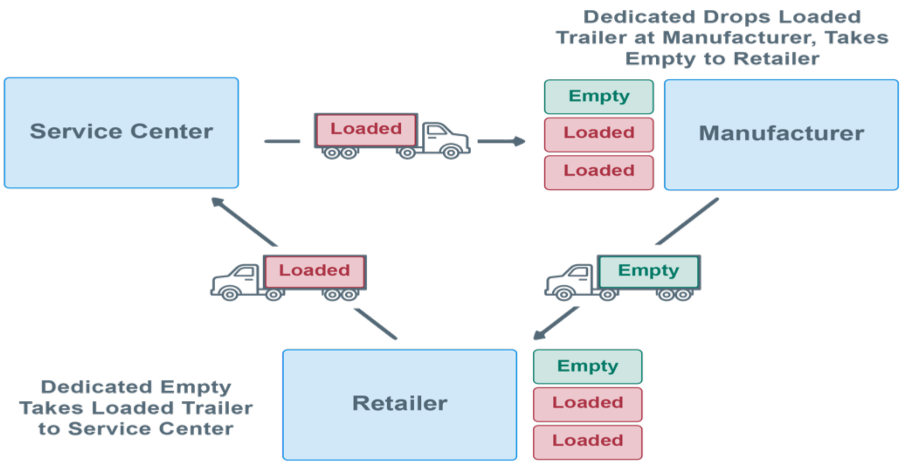

Fig. 1. Dedicated Drop Trailer Flow from our Sponsor Company

Within supply chain management, shippers commonly employ a blend of quantitative and qualitative forecasting methods to predict demand for their products. The quantitative approach often involves time series analysis, where historical sales data are scrutinized to identify patterns and trends that might predict future demand. This analysis may include techniques like moving averages, exponential smoothing, or more complex statistical models, such as Auto Regressive Integrated Moving Average (ARIMA), for more stable and predictable markets (Caplice, 2023a). On the qualitative side, shippers often rely on market research, expert opinions, and customer feedback, especially when launching new products or entering untested markets where historical data are scarce or non-existent. They may also use Delphi techniques, gathering and synthesizing insights from panels of experts, to forecast demand under conditions of high uncertainty (Caplice, 2023b).

Understanding customer demand amid this uncertainty is essential for businesses to maintain resilience and adaptability, make informed decisions, and strategically align resources. In the last decade, highly sophisticated and computationally intensive methods have emerged in the field of machine learning (ML) to hedge against this risk. Various subcategories of ML, such as neural networks, deep learning, reinforcement learning, and deep reinforcement learning, aim to go beyond conventional methods by extracting unseen relationships among variables that impact demand (Seyedan & Mafakheri, 2020). This approach has led to a reduction in uncertainty and an increase in prediction accuracy within supply chain management. Regardless of type, these models are often judged according to their predictive accuracy and bias, utilizing metrics such as Root Mean Squared Error (RMSE) and Mean Squared Error (MSE), respectively (Caplice & Ponce, 2023)

A variety of limiting factors in adopting some of these more complex forecasting techniques exist. First, these models require vast amounts of data that often are spread across disparate sources; therefore, care must be taken to integrate as well as clean the data for usability and accuracy. This data must also be collected, transferred, and analyzed at high speed, often continuously, requiring large amounts of computational power and labor to maintain. Additionally, while older, more traditional models are more transparent, in advanced ML models it can be difficult to trace how input data is transformed into predictions. Low levels of model interpretability, when there are high financial and operational implications, can be a major roadblock to adoption. Lastly, it goes without saying that the data flowing into the model must also be accurate and conform across time and formats (Seyedan & Mafakheri, 2020).

Despite these challenges, the potential of machine learning in demand forecasting is substantial, particularly in areas where traditional methods fall short. Machine learning algorithms excel in identifying complex, non-linear relationships within large datasets that conventional models often miss. This capability becomes particularly advantageous in short- to medium-term forecasting, where the dynamic nature of market conditions, consumer behavior, and external factors like economic shifts or seasonal trends rapidly change (Ferreira et al., 2016). ML models, with their capacity to process and learn from vast volumes of data from varied sources, can adapt more quickly to these changes, providing more accurate and timely predictions. This foresight is invaluable for strategic planning and decision-making, allowing the identification and use of demand drivers that may have previously been underutilized. As advancements in technology continue to address issues of interpretability and data management, the adoption of ML in demand forecasting is poised to offer transformative benefits, enhancing accuracy, efficiency, and strategic agility in supply chain management (Seyedan & Mafakheri, 2020).

The XGBoost (eXtreme Gradient Boosting) model is a comprehensive machine learning algorithm built upon decision trees, which are central to the gradient boosting framework. Developed by Tianqi Chen, this algorithm utilizes second-order information to create an efficient and adaptable system. It combines multiple regression trees to form a robust classifier and effectively counters the overfitting issue that is common in tree models. With parallel processing, XGBoost operates significantly faster than comparable algorithms under similar conditions (Chen & Guestrin, 2016). Its exceptional performance in analyzing high-dimensional data demonstrates its strong capability in handling complex processes. XGBoost's high efficiency and minimal requirements have led to its widespread application in various fields, including disease prediction, credit risk assessment, driving evaluation, and route planning (Zhang, et al., 2021).

At its core, XGBoost iteratively updates the parameters of its preceding classifier to minimize the loss function's gradient, thereby producing a new classifier. This process enhances prediction accuracy by utilizing several regression trees, ensuring maximum generalization capability for the tree group. The model's loss function includes a regularization term, and the solution is derived using the second-order Taylor expansion, focusing on minimizing the loss function for optimal node splitting. The incorporation of second-order derivative information and regularization techniques significantly boosts the model's generalization and computational performance (Zhang, et al., 2021). To our knowledge, this is the first attempt to apply XGBoost in drop trailer forecasting specifically, representing a novel approach in addressing one of the commonly encountered complexities of supply chain management.

5. Methodology

Our primary goal is to refine our sponsor's existing demand forecast and produce an accurate prediction of future trailer requirements at each manufacturing site served by the sponsor's service centers, which will inform strategic trailer allocation to maintain operational efficiency while avoiding resource surplus or deficits. We selected the XGBoost algorithm for its effectiveness in handling sparse data and its capacity for parallel processing, which is ideal for forecasting trailer demand. Unlike traditional models like ARIMA, XGBoost excels in capturing non-linear patterns and interactions within large datasets (Sagi & Rokash, 2021), making it particularly suitable for the volatile nature of drop trailer forecasting. To illustrate the efficacy of the XGBoost model, we conducted a comparative analysis with ARIMA and Simple Naïve models, focusing on key error measurements.

We sourced our data from lane-level transactions recorded between the sponsor's service centers and shippers, spanning from January 2018 to December 2023. This dataset includes historical time series data on empty and loaded trailer volumes, fleet utilization rates, and operational hours at various supply chain nodes, such as service centers and manufacturing sites. We initiated our analysis with an intensive data preparation phase. Our team actively cleaned and processed the data, eliminating duplicate records, rectifying inconsistencies in shipment logs, and discarding incomplete or irrelevant entries. Furthermore, we normalized variables like shipment times and distances for consistency. After refining the data, we amassed a comprehensive dataset of approximately 10,000 transactions, providing a solid foundation for subsequent analysis.

To develop our forecasting model, we allocated 70% of the weekly aggregated data for training and reserved 30% for testing. This split was carefully chosen to ensure enough data points for the model to learn underlying patterns while also allowing for a robust validation of its predictive capabilities on unseen data. We configured the model with a squared-error loss function, reflecting our focus on minimizing squared errors, which is a common practice for regression problems. A total of 1,000 trees were included in the model to balance the trade-off between learning capability and computational efficiency. To mitigate overfitting, we employed early stopping, ceasing training if the model's performance on the test set did not improve for 50 consecutive rounds. This strategy ensures that our model generalizes well to new data and does not learn noise from the training dataset.

6. Initial Findings

The sponsoring organization has provided a dataset that includes detailed information on customer accounts and trailer movements. Our initial analysis focuses on the largest client of the sponsor company, with a specific emphasis on the trailer movements from the primary fulfillment center to this client. This dataset contains records of both outbound movements, which involve trailers leaving the fulfillment center, and inbound movements, where trailers are received from the client's site. To reduce the impact of time delays and potential order cancellations, our analysis has been specifically limited to examining inbound movements.

To ensure relevance and accuracy, the data has been filtered to exclude movements prior to January 1, 2018. We then aggregated this data on a weekly basis, categorizing it according to the origin, destination, and the number of trailers received (expressed as loads). Although our goal is to develop a forecast for daily customer requirements, it is important to note that the sponsor company's service center does not operate on a 24/7 basis. Consequently, the company maintains a buffer stock of trailers at their clients' sites to enable operations during weekends without the need for continuous replenishment. This strategic approach justifies our decision to aggregate the data on a weekly basis, as it offers a more accurate representation of the actual demand.

Our analysis revealed that XGBoost significantly outperformed the ARIMA and Simple Naïve models. The Mean Squared Error (MSE) for XGBoost was 360.16, which is considerably lower than 556.96 for Simple Naïve. In terms of the Root Mean Squared Error (RMSE), XGBoost achieved 18.98, showcasing superior accuracy compared to ARIMA's 1335.27 and Simple Naïve's 23.60. Furthermore, the Mean Absolute Percentage Error (MAPE) for XGBoost stood at 10.53%, which was better than ARIMA's 18.55% and Simple Naïve's 12.67%, a performance that indicates higher precision in forecasting. However, while XGBoost showed promising results, it is important to acknowledge its limitations. The model's performance is heavily dependent on the quality and granularity of the input data, and it may not generalize well to scenarios that are vastly different from the training dataset. Additionally, XGBoost, like other machine learning models, can struggle with extremely volatile or unpredictable market conditions, which may impact the accuracy of its predictions.

Figures 2 and 3 show visualizations of the weekly load demand for Service Center A and actual, predicted, and forecasted load demand using XGBoost, respectively. Table 1 shows the Mean Squared Error, Root Mean Squared Error, and Mean Absolute Percentage Error for the XGBoost training and testing data from Figure 3.

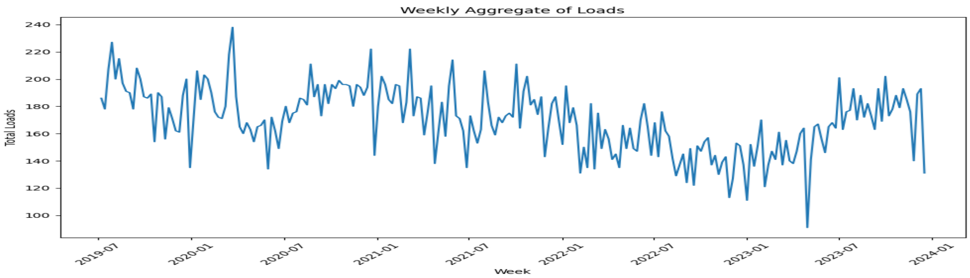

Fig. 2. Weekly Load Count, Service Center A

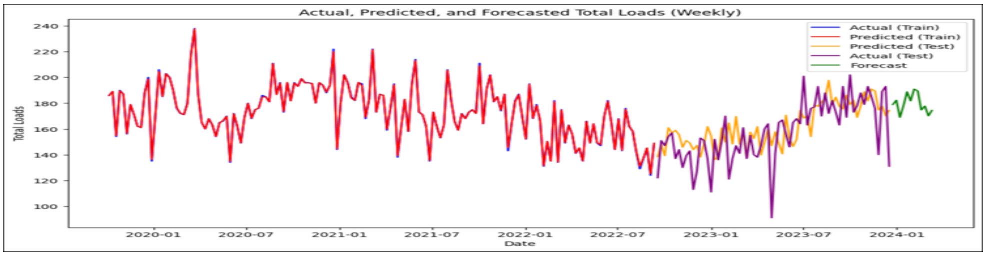

Fig. 3. Actual, Predicted, and Forecasted Total Loads, Service Center A

Table 1. XGBoost Forecasting Metrics

Error Measurement	XGBoost	ARIMA	Simple Naïve
MSE Test	360.16	36.54	556.96
RMSE Test	18.98	1335.27	23.60
MAPE Test	10.53%	18.55%	12.67%

7. Conclusion

Given our results, this approach shows promising results in optimizing asset utilization within the supply chain of our sponsor company. The comparative analysis underscores XGBoost's advanced predictive capabilities in the context of drop trailer forecasting, demonstrating potential superiority over traditional models like ARIMA and Simple Naïve in handling complex supply chain data. By effectively forecasting trailer demand, we have laid the groundwork for significant improvements in operational efficiency, cost reductions, and service-level enhancement.

For applications within other supply chain contexts in future research, it is important to consider the real-world implications and challenges that may arise when implementing the XGBoost model. First, in a retail supply chain, our model could significantly improve the accuracy of predicting demand spikes during holiday seasons. However, challenges might include integrating the model with real-time inventory systems and adapting it to a diverse product range and fluctuating consumer trends. In a manufacturing context, while our model can optimize raw material procurement based on production schedules, it might face difficulties in scenarios with sudden supply disruptions due to geopolitical or other factors.

In the context of global logistics, implementing our model could enhance route optimization and trailer utilization but could require overcoming practical challenges like varying data standards across countries and the need to handle complex customs regulations. Similarly, in the pharmaceutical supply chain, while our model could streamline drug distribution, it would have to be be tailored to account for stringent regulatory compliance and sensitive product handling requirements. Exploring the integration of real-time data analytics from geospatial trailer data could enhance the model's responsiveness to sudden market changes, further reducing forecasting errors. Lastly, investigating the application of other advanced machine learning techniques, such as neural networks or deep learning, might uncover even more nuanced insights within the data, potentially offering greater predictive accuracy.

Another promising area is the development of comprehensive optimization models that can operate in tandem with our forecasting tool. These models could dynamically adjust trailer allocations and carrier scheduling in real time in response to the predictive insights generated by the XGBoost model. And finally, expanding the scope of our study to include a wider range of supply chain partners, including data from manufacturers and their customers, could provide a more holistic view of the network's dynamics, thus leading to more robust and scalable solutions. Such research endeavors will not only contribute to the academic field of supply chain management, but will also offer practical, impactful tools for businesses navigating complex and ever-changing logistics landscapes.

References

[1] Fankhauser, E., & Li, G. (2019). Decoupled capacity with Powerloop (pp. 8-10) [Capstone Project]. Massachusetts Institute of Technology.

[2] Liu, S., & Ker, S. K. (2022). Uber Freight: Assessment and determination of optimal design features for a drop trailer service offering and network (pp. 2-10) [Capstone Project]. Massachusetts Institute of Technology.

[3] Pickett, C. (2018). Navigating the US truckload capacity cycle: Where are freight rates headed and why? Journal of Supply Chain Management, Logistics and Procurement, 1(1), 57-74. [1] Caplice, C. (2023). Product forecasting [0-23]. Retrieved from [https://canvas.mit.edu/courses/22994/pages/sept-20-session-new-product-forecasting?module_item_id=914396]

[4] Caplice, C. (2023). Product Forecasting [0-23]. Retrieved from [https://canvas.mit.edu/courses/22994/pages/sept-27-sessionadvanced-techniques-arima?module_item_id=914395]

[5] Caplice, C. (2023). ARIMA and Demand Forecasting [5-33]. Retrieved from [https://canvas.mit.edu/courses/22994/pages/sept-27-session-advanced-techniques-arima?module_item_id=914395]

[6] Caplice, C., & Ponce, E. (2023). MITx MicroMasters Program in SCM Key Concepts

[7] Seyedan, M., & Mafakheri, F. (2020). Predictive big data analytics for supply chain demand forecasting: Methods, applications, and research opportunities. Journal of Big Data, 7(1), 53. https://doi.org/10.1186/s40537-020-00329-2

[8] 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

[9] Zhang, M., Hao, S., & Hou, A. (2021). Study on the Intelligent Modeling of the Blade Aerodynamic Force in Compressors Based on Machine Learning. Mathematics, 9(476). https://doi.org/10.3390/math9050476

[10] Sagi, O., & Rokach, L. (2021). Approximating XGBoost with an interpretable decision tree. Information Sciences, 572, 522-542