Improving Delivery Performance Through Predicted Transit Times

by

Debra Lee
Bachelor of Science, Industrial Engineering and Operations Research, University of California, Berkeley

SUBMITTED TO THE PROGRAM IN SUPPLY CHAIN MANAGEMENT
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF APPLIED SCIENCE IN SUPPLY CHAIN MANAGEMENT
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

Signature of Author: ____

Department of Supply Chain Management
May 6, 2022

Certified by: ____

Dr. Milena Janjevic
Research Scientist
Capstone Advisor

Accepted by:

____

Prof. Yossi Sheffi
Director, Center for Transportation and Logistics
Elisha Gray II Professor of Engineering Systems
Professor, Civil and Environmental Engineering

Improving Delivery Performance Through Predicted Transit Times

by

Debra Lee

Submitted to the Program in Supply Chain Management
on May 6, 2022 in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management

ABSTRACT

Maintaining high on-time delivery performance is critical for third-party logistics providers. Not only does reliable on-time delivery function as a competitive edge and allow logistics providers to serve their customers more effectively, but it can also improve operational efficiency for the provider itself. Our sponsor company is a third-party logistics provider that must balance a growing portfolio of shipments that leverage external less-than-truckload (LTL) carriers. We therefore proposed and validated a two-pronged approach that utilized machine learning to improve LTL transit time predictions and then used these predictions in an integer programming model to maximize on-time orders. Independently, the prediction model improved the transit time RMSE from 3.07 days to 1.97 days. However, we obtained the most improvement in delivery performance through the optimization problem. By investigating the effect of lengthening the buffer days, or additional days added to the lead time beyond the predicted transit time over a rolling weekly basis, we obtained up to a 41% improvement in on-time deliveries over the status quo. Overall, the research demonstrates the strength of this mixed approach and provides flexibility to expand to other modes of transportation or a variety of objectives that may arise when planning shipments.

Capstone Advisor: Dr. Milena Janjevic
Title: Research Scientist

ACKNOWLEDGEMENTS

I would like to thank Dr. Milena Janjevic for her endless expertise and guidance throughout the year. This capstone would not have been possible without her support. I would also like to thank all of my family and friends who have been so encouraging throughout this entire journey. This experience would not have been the same without you all.

TABLE OF CONTENTS

LIST OF FIGURES

1 INTRODUCTION

1.1 Company Background

ABC Transportation Group's Retail Supplier Solutions division operates as a hybrid between a freight consolidator and a traditional third-party logistics provider (3PL). Their competitive edge comes from the high-performing and low-cost consolidation services offered through partnerships with major retailers such as Walmart, Target, and Kroger. While many of ABC Transportation Group's customers partner with them to participate in these consolidation programs, a growing share also turns to ABC Transportation Group as a full-outsource solution to manage all of their less-than-truckload (LTL) and truckload (TL) freight.

Balancing these discrete yet highly interconnected areas of the business is a challenge. Both consolidation and full-outsource customers share the same warehouse labor pool and have the same high expectation for on-time delivery performance. As more full-outsource opportunities arise, it is essential that this additional freight not hinder the existing consolidation programs and is individually successful in delivery performance and profitability. Therefore, ABC Transportation Group's focus is to optimize its LTL shipment scheduling process to ensure efficient growth.

1.2 Problem Statement

In this capstone project, we will use three years of data from ABC Transportation Group's Illinois warehouse to determine the optimal LTL shipment scheduling process for its operation. Illinois is one of ABC Transportation Group's most complex operations. It provides all of the available consolidation services and manages the highest volume of LTL freight. The largest gain could be obtained by optimizing this facility.

By nature, LTL shipments are less predictable than TL shipments. Instead of delivering directly from point A to point B, an order passes through several terminals before arriving at its final destination. Barcos et al.'s (2010) research explains that at a minimum, an LTL shipment will go to an origin end-of-line terminal (EOL) after being picked up regionally, ship to a destination EOL, and then travel to its final destination. LTL carriers may also utilize break-bulk terminals to further consolidate local freight before making the haul to the destination EOL (Barcos et al., 2010). Upon reaching the delivery terminal, depending on the retailer, the LTL carrier will still have to obtain a delivery appointment or deliver the freight based on some pre-set delivery schedule, often known as a Drop Trailer Program. All of these steps lead to increased opportunities for delays. ABC Transportation Group's data support this observation, as 74% of shipments took longer to deliver than the estimated transit time. Figure 1 contains histograms showing the distribution of actual transit time versus estimated transit time.

Figure 1 Actual Transit Time Versus Estimated Transit Time, 2019 - 2021

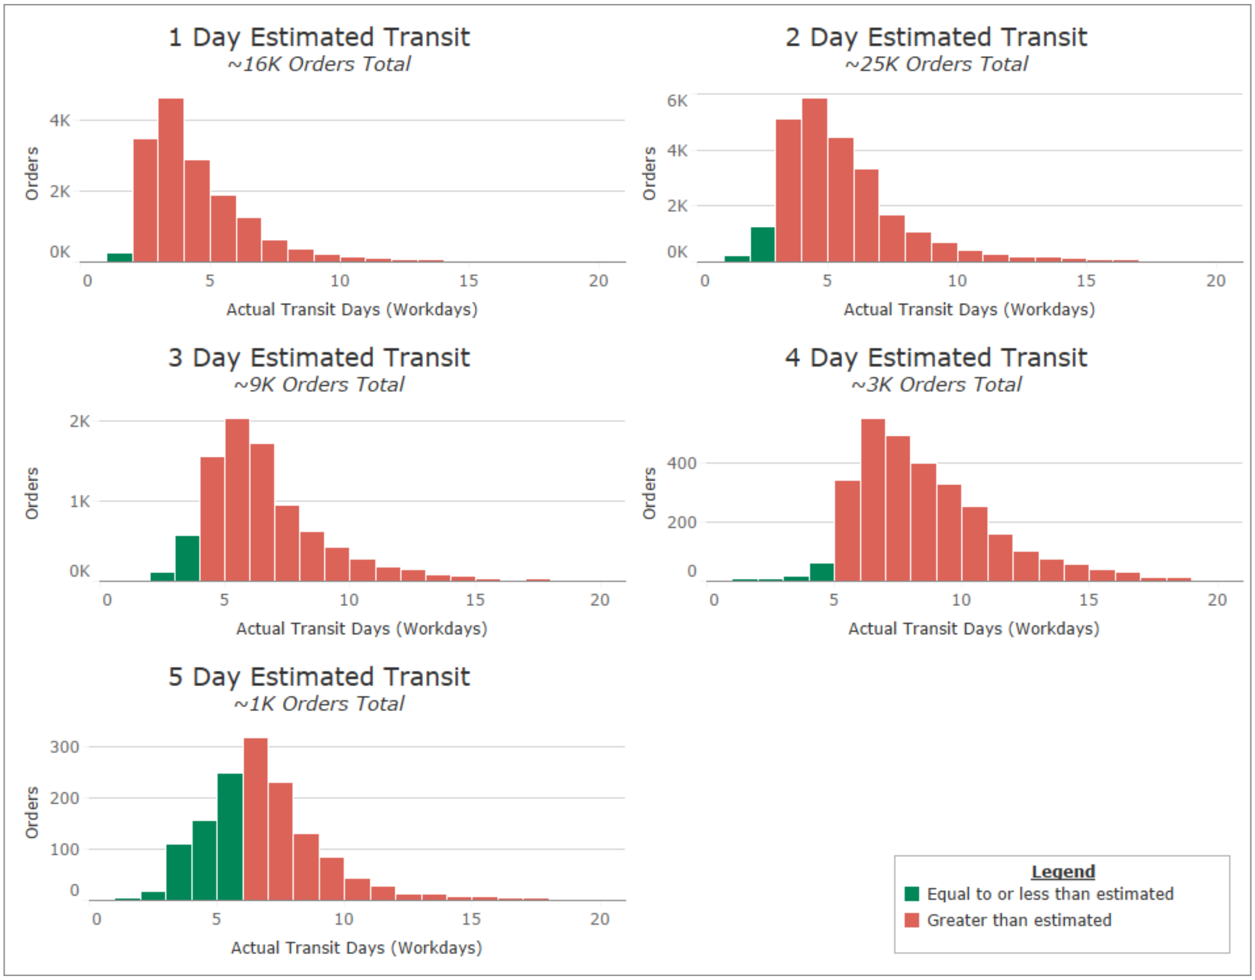

While transit time is undoubtedly a factor in determining on-time delivery, it is not the only one. Figure 2 displays the reason codes associated with all non-compliant orders in 2020. Ignoring external, uncontrollable causes, 6% of non-compliant orders were caused by internal planning issues, and 47% were related to LTL carrier operations. Although we cannot expect to change LTL carrier processes, we can anticipate delays and plan for them. We, therefore, could potentially impact 53% of non-compliant orders by implementing a more sophisticated shipment scheduling method.

Figure 2 ABC Transportation Group's Non-Compliant Delivery Categories and Underlying Reason Code Occurrence Frequency, 2020

Non-Compliant Category	Reason Code	
External Issue	Order Entered with Short Lead Time	23.6%
	Inventory Shortage	16.1%
	Extreme Weather	3.0%
	Customer Issue	2.0%
	Force Majeure	1.6%
	Error with Order	0.8%
Internal Planning Issue	Planning Error	6.1%
	Warehouse Delay	0.2%
LTL Carrier Operations	Appointment Set for After Due Date	27.8%
	Transit Delay	12.1%
	Drop Trailer Program	6.6%

Considering these reasons for delayed orders, ABC Transportation Group's current method for shipment scheduling presents an opportunity for improvement. Supply Chain Analysts (SCA) at ABC Transportation Group schedule shipments using estimated transit time information only, adjusting the lead times based on their subjective knowledge of the lanes. Because each SCA operates independently, overscheduling is common and goes unnoticed until later in the process. Labor availability also fluctuates depending on the volume of consolidated shipments that need to be picked and processed, so some days have lower capacity than others. When there are inevitably more orders than can be processed, SCAs must choose between paying for overtime labor or rescheduling the shipment to a later day, risking late delivery. Figure 3 illustrates ABC Transportation Group’s current LTL scheduling process. Each segment also lists the most common reasons for non-compliant delivery associated with it.

Figure 3 ABC Transportation Group's Order Scheduling Process

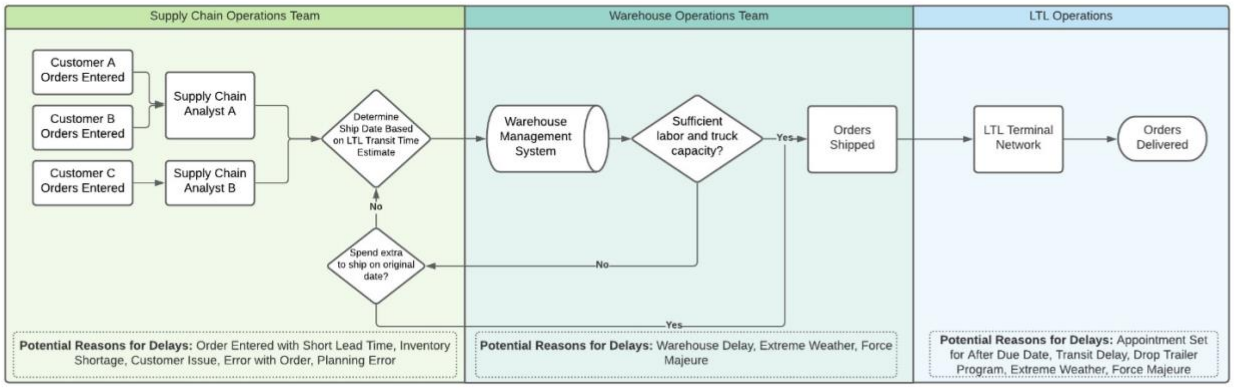

1.3 Methodology

To help solve ABC Transportation Group's LTL scheduling problem, we will take a two-pronged approach. We will first create a machine learning model to predict transit time, replacing the estimated transit times the organization is currently using. This model will utilize predictors such as origin and destination zip code, distance, retailer, time of year, geographical region, and carrier. Through this research, we will first determine which of these predictors influence transit time the most. Then, we will create a scheduling algorithm that maximizes projected delivery performance using predicted transit time as an input to a linear program.

The solution should ultimately provide several benefits – better warehouse efficiency by reducing overtime labor, decreased time spent by personnel manually scheduling and rescheduling orders, and improved delivery performance.

1.4 Relevance

As logistics have become more complex over the years, there has been a collective shift in interest from managing logistics in-house to investing in a professional logistics partner (Sheffi, 1990). Customers of 3PLs, therefore, expect their logistics providers to produce the best possible results at an affordable price while implementing innovative solutions (Premkumar et al., 2020).

Although predicting transit time is not a novel idea and has been replicated in many supply chain industries, its application to the complex shared warehouse environment of a 3PL provides a significant opportunity. Here, prediction accuracy impacts not only delivery performance but also overall operational efficiency. It embodies the type of advanced methods customers are looking for when selecting a logistics provider.

Ultimately, by improving scheduling methods in this way, logistics providers can continue to grow their operations and be assured to respond to any changes in the landscape. Although this capstone examines only LTL shipments, the general approach can be scaled to include multiple modes of transportation.

2 LITERATURE REVIEW

The lack of reliability in arrival times has consequences for buyers, suppliers, and carriers alike, leading to the burgeoning interest in using data-driven systems in supply chain management to handle uncertainty (Urciuoli, 2018). This capstone will contribute to this field by augmenting a classic scheduling optimization problem with predicted transit times so that ABC Transportation Group can increase its on-time delivery performance and leverage its LTL network to the fullest. This chapter will first explore the field of optimization and its applications to order management to understand the overall approach to the problem. Then, we will examine the research done regarding transit time prediction to find the models and factors that are most relevant and useful. These predicted transit times will ultimately become a component of the optimization problem but separately also offer insight into ABC Transportation Group's LTL network that they can apply elsewhere.

2.1 Scheduling Optimization

Scheduling optimization consists of determining a shipping plan from an origin to multiple destinations for a pre-determined period that minimizes shipping costs while adhering to feasibility constraints such as time, labor, or fleet size. This type of problem is well suited for and commonly solved using linear programming (LP). In this section, we will provide an overview of LP and its application to this capstone.

Linear programming, which has existed since the 1940s, was first used with great success on a large scale to create schedules for trains, buses, and airlines (Salvendy, 2007). It is, therefore, no surprise that linear programming has become a staple in supply chain management, particularly for transportation and logistics design. In an industry where time is one of the scarcest resources, linear programming provides a way to reduce time and cost while also responding to the ever-increasing demand for perfection from customers and competition within the market (Salvendy, 2007).

Simply put, an LP consists of a linear objective function to be minimized or maximized that includes decision variables and a set of constraints expressed as linear equations or inequalities (Salvendy, 2007). Discrete optimization is an extension of LP, which Salvendy (2007) characterizes to deal with decisions that have a logical or countable nature. This class of problems, typically referred to as integer programming or mixed-integer programming, restricts some or all variables in the LP to be integers (Williams, 2009). For simplicity, we will refer to this class of problems as IP going forward. The ABC Transportation Group order scheduling problem involves deciding which orders to allocate to specific shipping days, modeled using decision variables that are either 1 or 0 -- to ship or not to ship. The integer components of this model make it a classic IP problem.

The overall goal of discrete optimization is different from that of other types of optimization problems. IP problems are notoriously difficult to solve. The most elementary method for solving this class of problems is enumeration, where the user identifies all possible solutions and then selects the one that provides the best objective value (Conforti et al., 2014). While feasible in theory, the set of potential solutions can get unwieldy quickly. Conforti et al. (2014) use a straightforward assignment problem to illustrate the point. This example aims to match *n* employees to *n* tasks, which means the number of possible assignments is *n!*. When n = 10, which would be a small IP problem, the value of *n!* and thus the number of potential solutions is already 3.6×10^6.

Fortunately, finding a mathematically optimal solution for an IP is neither practical nor necessary since the model is simply a representation of an existing system (Salvendy, 2007). Instead, Salvendy (2007) states that the emphasis is on identifying feasible solutions that provide a desirable objective function value. This approach suits the ABC Transportation Group problem since the focus is on distinguishing ways to schedule orders to attain high delivery performance. There is no need to maximize the delivery performance as long as it surpasses a certain threshold.

A common way to simplify solving IP problems is to use relaxation, a methodology that substitutes an easier-to-solve approximation for the original IP. Linear programming relaxations are a popular approach because linear programs can be efficiently solved and lend themselves well to creating increasingly tighter approximations (Conforti et al., 2014). Conforti et al. (2014) mention Branch and Bound and Cutting Plane as two such methods that are the basis for modern IP software. Since this capstone will be using such software, describing all the methods for solving IP problems is out of the scope of this literature review. The reader can find more information on them in *Integer Programming* (Conforti et al., 2014) or *Logic and Integer Programming* (Williams, 2009).

2.2 Transit Time Prediction

Estimated transit time is a critical component of the optimization problem addressed in this capstone. The delivery performance projected by the solution can only be achievable in reality if the transit time inputs are accurate. The motivation chapter has already demonstrated that the estimated transit times used by ABC Transportation Group now are insufficient. Therefore, we seek to augment those estimates using predictive analysis.

Poschmann et al. (2019) posit that transit time prediction is best solved using either a model-based approach (simulation models or analytical models like graph theory, queueing theory) or a data-based approach (statistical methods, machine learning). Since the ABC Transportation Group problem aims to understand an LTL network that the company neither has control over nor has detailed information on, a model-based approach is not suitable. Instead, we aim to infer a relationship between actual transit time and various shipment-related data points. This strategy aligns with a data-based approach, specifically machine learning, which excels at finding patterns in a large amount of data. All the literature reviewed in this capstone utilizes machine learning, which we will also adopt for this capstone. In the upcoming sections, we will discuss the components of machine learning we are most concerned with -- model and variable selection.

2.2.1 Model Selection

Machine learning is categorized into three distinct types -- supervised, unsupervised, and reinforcement learning. This capstone will focus on supervised learning, also known as predictive learning, which uses past data to train a machine to learn about a topic and make predictions (Chandramouli et al., 2018). Many supervised machine learning models exist. Chandramouli et al. (2018) mention Naive Bayes, K-Nearest Neighbor, and decision tree as some oft-used models.

While machine learning models used in the studies examined for this literature review varied, the random forest model was the most commonly used and found to be accurate. Regardless of the network, whether it was intermodal freight transportation with multiple nodes (Balster et al., 2020), an open-pit mining operation (Sun et al., 2018), or even an airline (Alla et al., 2021), random forests provided the best accuracy in terms of root mean square error. The success of this technique makes inherent sense as the benefit of random forests is its tendency to avoid overfitting while taking in many variables and a high level of noise (Breiman, 2001). Considering the variability seen in ABC Transportation Group's LTL network and the number of factors in determining actual transit time, we will also adopt the random forest technique in this capstone.

Among the studies that did not use a random forest model, Truong's (2014) paper on estimated outbound linehaul transit times from an Amazon DC closely matches the ABC Transportation Group problem in scope and goal, so it is worth exploring further. Although linehaul inherently has less variability than the LTL network examined in the ABC Transportation

Group problem, the Amazon DC's usage of non-asset-based carriers also resulted in less availability of granular transit data. Because of this lack of data, this Amazon DC routinely determined lead times based on a combination of historical data and subjective inputs from their planners and carrier partners, which is very much like the current ABC Transportation Group process (Truong, 2014). Since the DC wanted solutions at different performance levels, Truong (2014) used a quantile regression forest model that provided estimated transit times for a given probability of on-time arrival, *p*. Although the model was successful, Truong (2014) noted that it produced conservative estimates for high *p*'s and short hauls more affected by variability. Since the bulk of ABC Transportation Group's LTL shipments are considered short hauls, this phenomenon is something that we should also pay attention to, despite using a different model.

2.2.2 Variable Selection

Regardless of a predictive model's inherent effectiveness, its success depends on variable selection. Shmueli and Koppius (2011) assert that "predictors are chosen based on a combination of theory, domain knowledge, and empirical evidence of association with the response" (p. 564). However, simply identifying predictors with a causal relationship with the response studied is not sufficient. Predictors must be available at the time of prediction and of high measurement quality, which often means that a proxy variable can be a better choice for the model than a causal variable that might be more intuitive (Shmueli & Koppius, 2011). In the ABC Transportation Group problem, for instance, late delivery reason codes would fall in this category. Although they have a relationship with the actual transit time and are very descriptive, they are not suitable for the model because they do not exist before an order is shipped.

A range of approaches are used to identify variables. Van der Spoel et al. (2017) utilized surveys to determine predictors empirically. Poschmann et al. (2019) took a similar approach and interviewed the personnel at each node within the intermodal transport chain they modeled. Several studies used Pearson correlation coefficients (PCC) to understand the relationships between different predictors to avoid using too many overlapping variables in their models. Van der Spoel et al. (2017) and Truong (2014) used PCCs on all possible predictors, where Sun et al. (2018) solely employed this method on weather-related variables like humidity and wind speed. Instead of directly finding correlations between variables, Balster et al. (2020) iteratively reduced the list of predictors by removing the least impactful variable one at a time when running the prediction model. Figure 4 summarizes the types of predictors used in different research papers. Based on these studies and conversations with the sponsoring company, it should be helpful to include destination, temporal, and weather data into the ABC Transportation Group model, among other variables to be determined.

Figure 4 Predictors Investigated by Papers in this Literature Review

Predictors	Alla et al., 2021 *(airline study)*	Balster et al., 2020	Poschmann et al., 2019	Simorth & Zahle, 2011	Sun et al., 2018	Truong, 2014	van der Spoel et al., 2017
Origin/Destination	x	x	x	x	x	x	x
Origin/Destination Density	x						
Departure Delay Reasons						x	
Lane Type						x	
Specific Temporal (Arrival/Departure Times)	x				x	x	x
Broad Temporal (Season, Holiday, Weekday)	x	x	x			x	
Order Composition		x	x				
Weather		x	x		x		x
Vehicle Features		x	x				x
Infrastructure			x				
Staff Quality			x				
Load Characteristics			x		x		x
Road Type					x		
Carrier	x				x		x
Traffic							x
Distance	x			x			x
Route				x	x		

2.3 Conclusion

This literature review explored scheduling optimization and the model and variable selection process for predicting transit time using machine learning. While optimization and transit time prediction are both well-studied areas of supply chain management, they both employ certain assumptions when solved separately. In the current body of literature, not many papers thoroughly explore these assumptions. Scheduling optimization models assume a high accuracy of inputs and mainly focus on the linear programming formulation. Transit time prediction problems provide predictions and sometimes present results of pilot programs using the data but typically do not explore its use in the scheduling process. This capstone will bring additional perspective and transparency by merging these problems to assess the relationship between the two distinct but interconnected areas of study.

3 METHODOLOGY

We have determined that the best way to help ABC Transportation Group improve its LTL scheduling process is by using predicted transit times in a scheduling optimization problem. To achieve this goal, we took a three-step approach. First, we cleaned and augmented the data provided by ABC Transportation Group. We then used this data to create and refine an ML model to predict transit times. Finally, we formulated a IP to generate a shipment schedule that maximizes the number of on-time orders using these transit times.

3.1 Data

ABC Transportation Group provided nearly three years' worth of shipment data from its Illinois warehouse to use in our transit time prediction model. This data is output from their proprietary enterprise resource planning system and contains all shipment-related details for LTL shipments delivered between January 2019 and September 2021. In the following sections, we will introduce the data and explain the cleaning steps taken. The data analysis and cleaning took place in Excel and Tableau.

3.1.1 Data Types

The overarching data types provided by ABC Transportation Group are as follows:

- Order ID: A unique identifier per order.
- Dates: Includes ship, delivery, and requested delivery (due) date.
- Origin: Includes the origin name, state, city, and postcode. Since we are only analyzing one specific warehouse, these fields are uniform for the entire dataset.
- Destination: Includes the destination name, state, city, and postcode.
- Distance: The distance in miles between the origin and destination location.
- Order Size: Includes pallet count, weight, cube, and case count.
- Estimated Transit Time: The estimated workdays to deliver from an origin to a destination by postcode. This figure does not include the day of pickup and is what ABC Transportation Group currently uses when determining lead times. They are an average of the published lead times of the carriers in their LTL network.
- Late Reason Code: The reason an order was delivered late. These are tagged manually by an employee after a shipment is delivered.
- Carrier: A unique identifier of an LTL carrier assigned to an order.
- Order Mode: The specific type of LTL service used, such as refrigerated, guaranteed, or expedited.

To aid in the analysis, we created some additional data points derived from the provided data.

- Actual Transit Days: The number of days, excluding weekends and holidays, that a shipment was in transit. This number includes the date the order was delivered, but not the pickup date, to match the logic used for the estimated transit time provided by the company. LTL carriers do not operate over the weekends and select national holidays, so this figure provides the most accurate representation of the days a shipment was in transit.

- Holiday: A binary indicator for whether a shipment fell between one week before and one week after a holiday week. The sponsoring company believed that the reduced shipment days tied to holiday closures would cause more congestion within the LTL network, leading to delays.
- Quarter: The fiscal quarter, based on the Requested Delivery Date, is used to capture seasonality.
- Day of the Week of the Requested Delivery Date: Delivery locations are known to have recurring drop appointments scheduled with LTL carriers. The day of the week of the requested delivery date could indicate which scheduled drop appointment the shipment would deliver on.
- Three-Digit Zip Code: The three-digit zip code is the first three digits of a five-digit zip code and ties to a geographical region, typically a large city (ZIP Code™ - The Basics, 2021). Based on the data, 56% of the five-digit destination zip codes had less than ten shipments which would likely not provide enough data to make a conclusive prediction. Three-digit zip codes are a way to generalize to a larger region that is still more granular than simply using the destination state.
- Retailer: The destination names on orders are tied to ABC Transportation Group's customers' systems, and hence have a lot of variability. We manually tagged the retailer names based on the destination address information and verified them with the sponsor.

3.1.2 Data Cleaning

Data cleaning is necessary to create an accurate predictive model. We implemented the adjustments below following discussions with the sponsor:

- Remove Cross-Docked Orders: Most of ABC Transportation Group's LTL shipments originate from their warehouse, so the order size details are known. However, ABC Transportation Group also operates as a cross-docking point, which means clients may deliver freight picked and palletized elsewhere to be shipped LTL from ABC Transportation Group's warehouse. In these situations, the exact case count, pallets, and weight are manually entered based on information provided by the client, and therefore often has missing values due to human error. Cube is not available because pallet dimensions are not collected. Since these types of orders only made up 2% of the records in the dataset, we elected to exclude them.
- Remove Orders with Impossible Cube or Weight: The maximum capacity for a trailer is approximately 3000 cubic feet or 42,000 pounds. We excluded shipments that had a cube or weight higher than the maximum allowable, as well as null values.
- Remove Orders with Zero or Negative Transit Time: The ship and delivery dates have a high level of accuracy because ABC Transportation Group's LTL carriers communicate shipment statuses, including activity dates, via Electronic Data Interchange. However, we identified orders with delivery dates on or before their ship dates, which is impossible.
- Remove Orders Shipped with Non-Standard Services: Orders can be shipped with guaranteed or expedited services at a premium price to ensure a transit time. The LTL processes surrounding these elevated shipments are different, as they are given priority and

must meet the service level agreed upon during the quoting process. Refrigerated freight uses a separate network from standard LTL. Since the total number of orders for these three service types makes up less than 1% of the dataset, there would not be enough representation to use as a variable. Moreover, in the case of the premium services, one could assume that the agreed-upon transit time would be the best estimation of the transit time already, making any predictions redundant.

- Remove Outliers Above the 99^{th} Percentile: The observed transit days displayed a normal distribution with a long right tail, as can be seen in Figure 5. To balance between capturing the underlying variability but excluding true outliers, we assessed removing orders that were above the 90^{th} (7 Transit Days), 95^{th} (9 Transit Days), and 99^{th} (15 Transit Days) percentiles. To ensure we were not removing valid data for key delivery locations, we calculated the average transit times for all locations tied to the top 10 retailers by volume. Since several commonly shipped lanes had transit times that routinely exceeded 11 days, we only excluded orders above the 99^{th} percentile. The company sponsor reviewed the internal order notes for these excluded orders and confirmed that most had rare issues during transit that caused the delivery time to be extended. These should therefore be considered true outliers and would not add value to a prediction model.

Figure 5 Number of Orders by Total Transit Days, with 90th, 95th, and 99th Percentiles

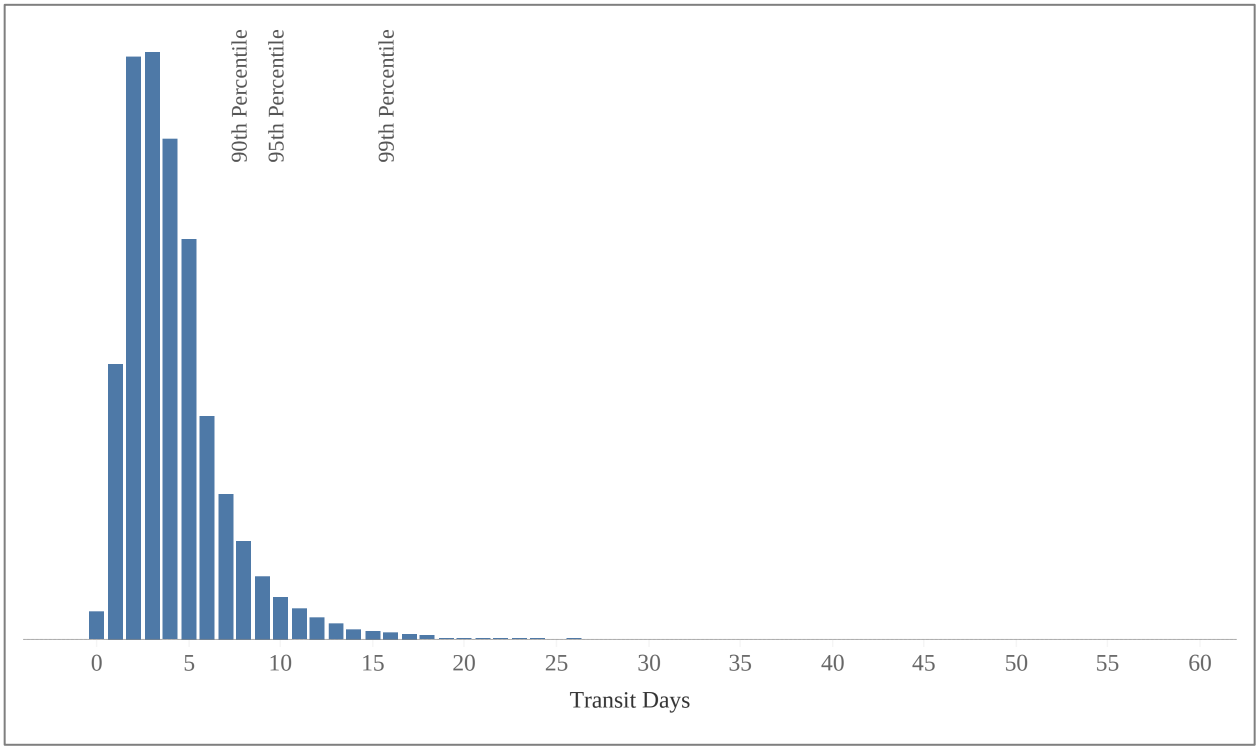

Interestingly, we expected COVID-19 to negatively impact transit times during the initial breakout period in early 2020, but there was no evidence of this in the data. Instead, we see a jump in average transit days in 2021. Both of these observations can be seen in Figure 6. Analysts seem to attribute this change to two factors: struggles with the receiving process at delivery locations due to COVID-19 issues and the increase in LTL usage in the growing e-commerce industry leading to tightened overall capacity (Shulz, 2021; *United States less-than-truckload (LTL) market growth, trends, covid-19 impact, and forecasts 2021-2026 - researchandmarkets.com*., 2021). We will attempt to capture these trends in the prediction model instead of excluding any data points. We will discuss this further in the next section on transit time prediction.

Figure 6 Average Transit Time by Week Delivered

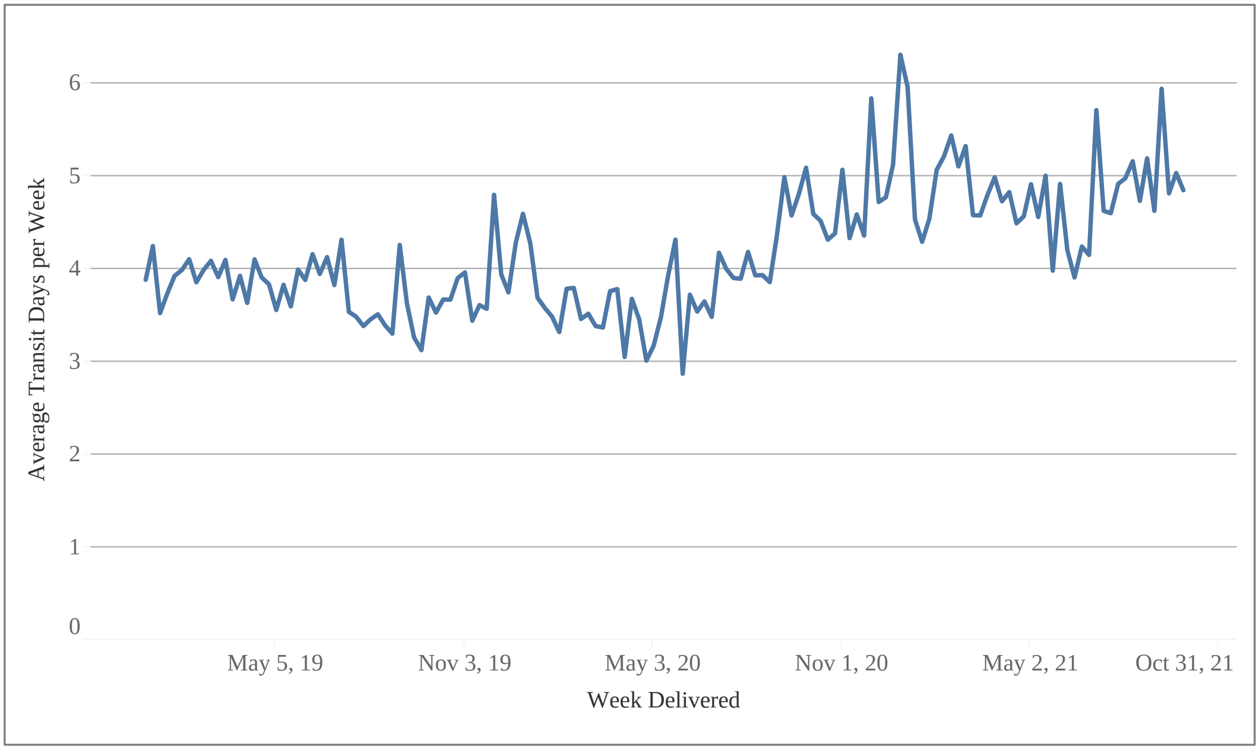

3.2 Transit Time Prediction

We determined that machine learning, and specifically the random forest model, was the best way to predict transit time through the literature review. Random forests are an ensemble machine learning method that utilizes many decision trees based on random independent vectors that then vote for the most popular classification (Breiman, 2001). For this project, all analysis has been done using the sci-kit learn library in Python since it is open-source and could be used by the company after the conclusion of this project. In this section, we will discuss the model formulation and variable selection using this library.

3.2.1 Model Formulation

Random forests are a supervised learning method, so we first split our data into training and testing sets. We utilized 75% of our total data points as the training set and reserved the remaining 25% for testing. To create reproducible results during the formulation process, we set an integer for the random state, which guarantees that the algorithm will select the same values from the dataset each time.

The sci-kit learn classifier has default values for all parameters that provide a good starting point for models but may not be optimal for all scenarios. There are several methods to test and determine which parameters work best. In this study, we elected to use GridSearchCV, which takes the parameters the user wishes to evaluate as inputs and returns the best choice. In our model, the two parameters we wanted to fine-tune were the number of decision trees and the maximum depth of the tree.

In random forests, the more trees there are, the better the prediction strength, but eventually, the reduction in error becomes negligible (Breiman, 2001). Since more computational power is also needed when there are more trees, we aim to find the ideal number that reduces error but runs in a reasonable amount of time. The default setting is 100 trees, and we have chosen to also test the model with 500, 1000, 1500, and 2000 trees.

The maximum depth of a tree indicates the number of splits made within a decision tree. The more splits there are, the higher the computational needs, but a high depth may result in overfitting while a lower depth can cause underfitting (Liu et al., 2017). For this parameter, we have elected to test max depths of 8, 16, and 32.

For this initial test, we utilized all available variables as a baseline and root mean square error (RMSE) to determine the accuracy of fit. We consider parameters x_i, the actual transit time for shipment i, $\hat{x}_i$, the predicted transit time for shipment I, and N as the total number of shipments. RMSE is thus defined as follows:

$$\mathrm{RMSE} = \sqrt{\frac{\sum_{i=1}^{N}(x_i - \hat{x}_i)^2}{N}} \tag{1}$$

The RMSE for the actual transit times compared to the estimated transit times currently used in ABC Transportation Group's scheduling process is 3.07 days. This number is used as the control to evaluate our prediction accuracy. GridSearchCV determined that the best parameters for our prediction model were 1500 trees and a max depth of 16. The RMSE using the default settings was 2.02, and the RMSE using the best parameters produced a slightly better RMSE of 2.01. Both show a considerable improvement over the original estimations. Now that we have developed a working model and identified the optimal parameters, we will next dive into the variable selection process to improve our model further.

3.2.2 Variable Selection

The variables for our model consist of data known about orders before shipment since the sponsor company plans to use this model during the scheduling process. Due to this constraint, we excluded late reason codes since they are only tagged after delivery and cannot be forecasted.

While random forests can handle many variables without overfitting, using strongly correlated factors does not improve prediction accuracy and increases the error rate (Breiman, 2001). Due to the many potential variables, we first created a Pearson Correlation Chart (PCC) to understand the relationships. The PCC is in Appendix A and shows the correlation between all available data points.

Then, we determined the importance of each variable using the feature importance measures within the random forest classifier. For each run, the features are listed from most to least important. Using the PCC, we eliminated less important and highly correlated variables. Each run, the relevant importance factor, and the resulting RMSE are in Appendix B.

The set of variables that led to the lowest RMSE were:

Dependent Variables

1. Actual Transit Time: The number of workdays to deliver from an origin to a destination.

Independent Variables (Listed from most to least, in terms of model importance)

2. Weight: The total weight of a shipment.
3. Distance: The miles between an origin and destination.
4. Month of Requested Delivery Date: Numerical variable denoting the month of the due date.
5. Weekday of Requested Delivery Date: Numerical variable denoting the day of the week of the due date.
6. Destination Postcode: The five-digit postcode of a destination.
7. Retailer ID: Numerical variable unique for each retailer.
8. 3 Digit Zip: The three-digit postcode of a destination.
9. Carrier ID: Numerical variable unique for each carrier.
10. Year of Requested Delivery Date: Numerical variable denoting the year of the due date

11. Quarter of Requested Delivery Date: Numerical variable denoting the quarter of the due date
12. Holiday: Binary variable indicating if a due date falls between one week before and one week after a national holiday

Although we evaluated the RF models based on RMSE, the selected model outperformed the status quo based on other measures beyond just this metric. The model predicted 36% of shipment transit times accurately, compared to just 21% in the company's estimate. Moreover, in the new predictive model, 60% of shipments were accurate or overestimated, compared to 26% for the sponsor company's estimate. This metric is particularly relevant for shipment scheduling, since early and late deliveries do not have the same consequences. Most recipients do not penalize early delivery and thus delivering before the due date is acceptable and even considered on time. Regardless of this loophole, the overestimation only exceeds two days fewer than 10% of the time. A histogram depicting these results is in Figure 7.

Figure 7 Variance in Days Between the Original Estimate and the Actual Transit Time Compared to The Variance in Days Between the Predicted Transit Times and the Actual Transit Time

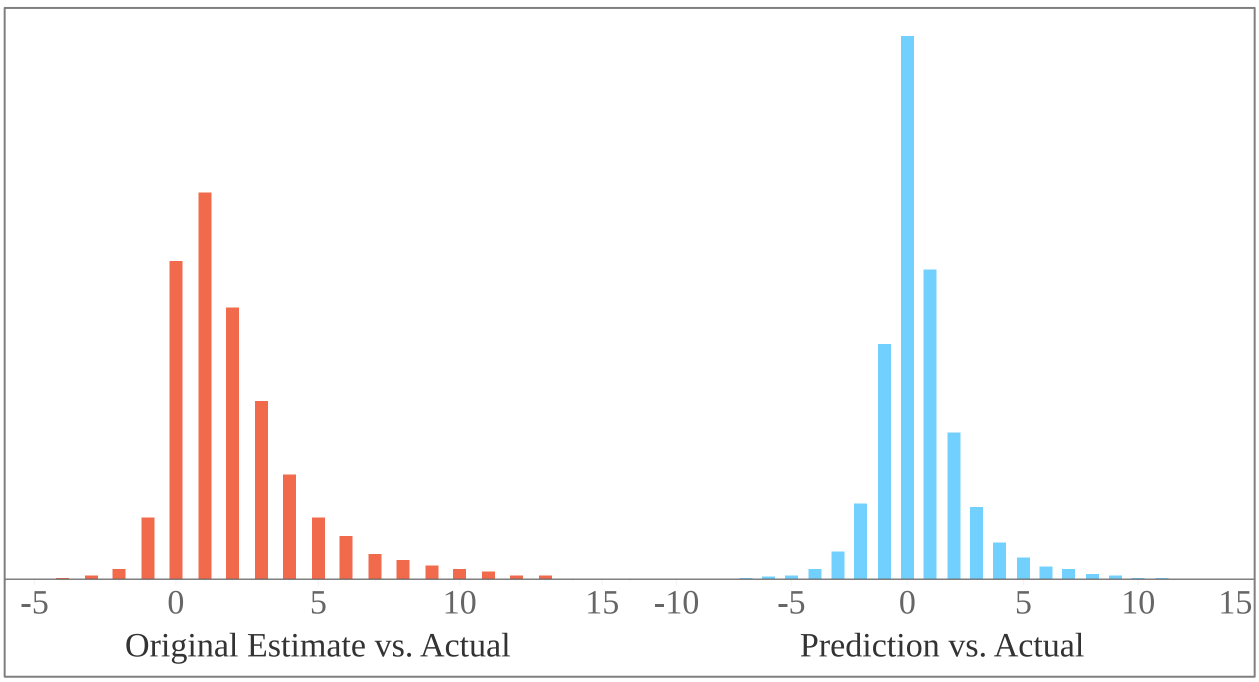

3.3 Shipment Scheduling Optimization

Ultimately, our goal is to help ABC Transportation Group optimize its LTL scheduling process by solving an IP problem. The objective is to maximize the number of orders delivered on time, considering capacity constraints on ship dates.

Now that we have a prediction model that provides improved estimated transit times, we used this information to create inputs to our IP. By knowing the due date of an order and the predicted transit time to deliver to its destination, we can derive which ship dates would allow for on-time delivery. For most consignees, delivering on or before the due date is considered on time, but our model also provides enough flexibility to handle retailers with specific delivery windows.

Based on discussions with ABC Transportation Group's Director of Warehousing, the warehouse labor constraints are straightforward and consist of a contracted number of cases per week for all total TL and LTL activity. Nearly all TL shipments are planned the week before shipment, which allows us to derive the remaining capacity per day to allocate to LTL.

To formulate our optimization model, we consider I, a set of orders that need to be shipped, and J, a set of days on which these orders can be shipped. For every order i in I, we define c_i, which represents the number of cases per order. The case capacity per day is denoted by C_j. The total cases shipped on day j can never surpass C_j.

In addition, we consider o_{ij}, which is computed for every order i in I and every day j in J using the results of the transit time prediction model. o_{ij} denotes whether an order is predicted to arrive on time given the transit time prediction. We can also introduce $\hat{o}_{ij}$, which will take on the value of 1 when the requirements of o_{ij} and any additional rules are met. o_{ij} and $\hat{o}_{ij}$ are meant to be used interchangeably depending on the business need.

To illustrate the idea behind o_{ij} and $\hat{o}_{ij}$, let us consider Order 1 that may ship any time between Day 1 and Day 5. If the predictive model indicates that Order 1 must ship on Day 3 to allow enough lead time to deliver, o_{11}, o_{12}, and o_{13} would all take on values of 1, and o_{14} and o_{15} would be 0. If we enforce an additional requirement that at least 1 extra day needs to be provided to the lead time beyond what is predicted, now $\hat{o}_{11}$and $\hat{o}_{12}$ would be 1, and $\hat{o}_{13}$, $\hat{o}_{14}$, and $\hat{o}_{15}$ would be 0.

The formulation is therefore as follows:

Notation:

$x{:}_{ij}$: Binary variable, indicates whether an order *i* is shipped on day *j*

o_{ij}: Parameter indicating whether an order *i* would be on time if shipped on day *j*

$\hat{o}_{ij}$: Parameter indicating whether an order *i* would be on time if shipped on day *j* and if it satisfies additional user-defined rules

c_i: Cases per order *i*

C_j: Case capacity per day *j*

Formulation:

$$Maximize: \sum_{i \in I} \sum_{j \in J} x_{ij} o_{ij} \tag{2}$$

Subject to:

$$\sum_{i \in I} x_{ij} c_{ij} \leq C_j, \qquad \forall\, j = 1, \dots, m, \tag{3}$$

$$\sum_{j \in J} x_{ij} = 1, \qquad \forall\, i = 1, \dots, n, \tag{4}$$

$$x_{ij} = \{0,1\}. \tag{5}$$

The model maximizes (2), which is the total number of on-time shipments. This objective function is a simple approach but adequately captures the main goal to deliver the most shipments on time from the facility. Future models being discussed add factors to this equation, such as penalties for orders over a certain number of days late or customers whose on-time performance is below a threshold.

Equations (3) – (5) consist of the constraints of the model. The daily case capacity is captured in (3), where all cases shipped in a day must be less than the amount of available capacity that day. Equation (4) ensures all orders are shipped once, and (5) enforces that the variable is binary.

A notable exclusion from this model is carrier selection. Carriers are pre-assigned systemically based on cost and compliance considerations before scheduling, which is already optimal. Also, since LTL trailers are pre-loaded and pickup windows are scheduled for the same time daily, all scheduled shipments are typically picked up without issue, regardless of volume. However, considering the changing LTL landscape as mentioned in prior sections, this may be a valuable addition to a future iteration of this model. In the next chapter, we will review the results from applying our optimization model to a test scenario.

4 RESULTS

The proposed ML and optimization model must produce superior on-time delivery results for ABC Transportation Group to replace its current method of shipment scheduling. We, therefore, simulated a real use-case scenario to determine the magnitude of the improvement.

ABC Transportation Group's supply chain analyst team currently schedules shipments manually, using the estimated transit time, typically taking thirty seconds to process each order. On this front, the optimization model already shows an advantage over the original process since the model can schedule a week's worth of orders in, on average, .38 seconds.

Depending on their personal preference, analysts may also add days to the lead time if they feel the original estimate is too short. This manual addition of buffer time makes it challenging to compare the performance of the proposed model with the status quo on an equal playing field since the buffer added can vary greatly. Therefore, we assess our model based on two factors: the model's on-time performance compared to the actual on-time performance (including any buffer) and the model's on-time performance compared to the projected on-time performance using estimated transit times only (no buffer). This approach strikes a balance between assessing the model's strength compared to the base case, which is the delivery performance based on the current estimated transit times, and what happens in practice. A successful model should perform better in both situations.

While building the original model, the program randomly selected data to use as the training and testing sets. This approach provides the most robust model for ABC Transportation Group's use in the future because it utilizes the entirety of the data we acquired. However, since we could not access data for new, unshipped orders, we simulated the experience by changing the makeup of the training and testing sets. Instead of a random selection, we now split the data chronologically. The oldest 90% of data, consisting of 46,730 data points, served as the training set, while the newest 10% of data, consisting of 5,193 data points, served as the testing set. This test run took 40.57 seconds and produced an RMSE of 2.45 days, slightly higher than the original RMSE of 1.97. We anticipated the RMSE for this model to be higher than the original since the training set did not span the entirety of the time covered by the provided data.

Since ABC Transportation Group will ultimately apply this model on a rolling weekly basis, our simulation also considered one week as the planning period. In this experiment, we focused on weeks of the test set with order volumes deemed to be representative by the sponsor, which narrowed the number of weeks down to 13 from the 19 available. The test weeks also encompass the Independence Day holiday and will be illustrative of the model’s resilience when transit times are expected to be longer due to nationwide closures.

In the next sections, we introduce variations of the original optimization model and discuss the performance and merits of each.

4.1 Baseline Optimization Model - No Buffer

The baseline optimization model is described in the Methodology chapter and seeks to maximize the number of on-time orders, given warehouse labor constraints. Since the warehouse labor availability for LTL shipments varies depending on the business need of other parts of the business, we took a straightforward approach that allocated labor evenly for the test run. We gave each day a capacity of 1/5 of the total number of cases per work week.

The baseline model indicated, on average, a 31% on-time delivery performance increase over the on-time delivery performance based on the estimated transit time but showed no gains versus the actual on-time score. The distribution of the results is displayed in Figure 8.

Figure 8 Projected Performance of the Baseline Optimization Model Compared to ABC Transportation Group's Actual Past Performance and Projected Performance if the Original LTL Estimate Was Used

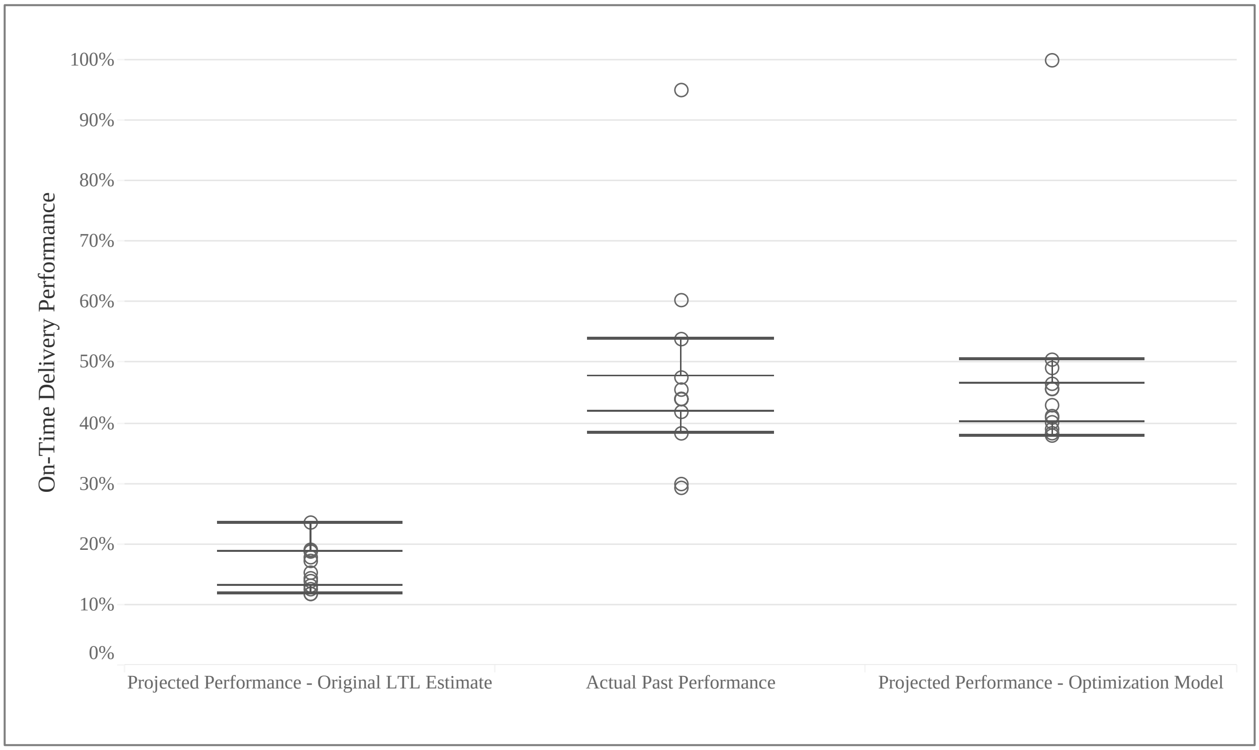

4.2 Optimization Model with One-Day Buffer

This optimization model builds on the baseline model by enforcing a minimum one-day buffer, if possible, which forces the ship date to be at least one day before the ship date derived from the predictive model. When adding a buffer day is not possible, the earliest ship date in the shipping period is considered the only choice. Within the model, this is done by adjusting o_{ij} to take into account the additional day of transit required.

Here, the on-time delivery performance increased by, on average, 50% versus the on-time delivery performance based on the estimated transit time and by 18% versus the actual on-time performance achieved by the company. The distribution of the results is displayed in Figure 9.

Figure 9 Projected Performance of the One-Day Buffer Optimization Model Compared to ABC Transportation Group's Actual Past Performance and Projected Performance if the Original LTL Estimate Was Used

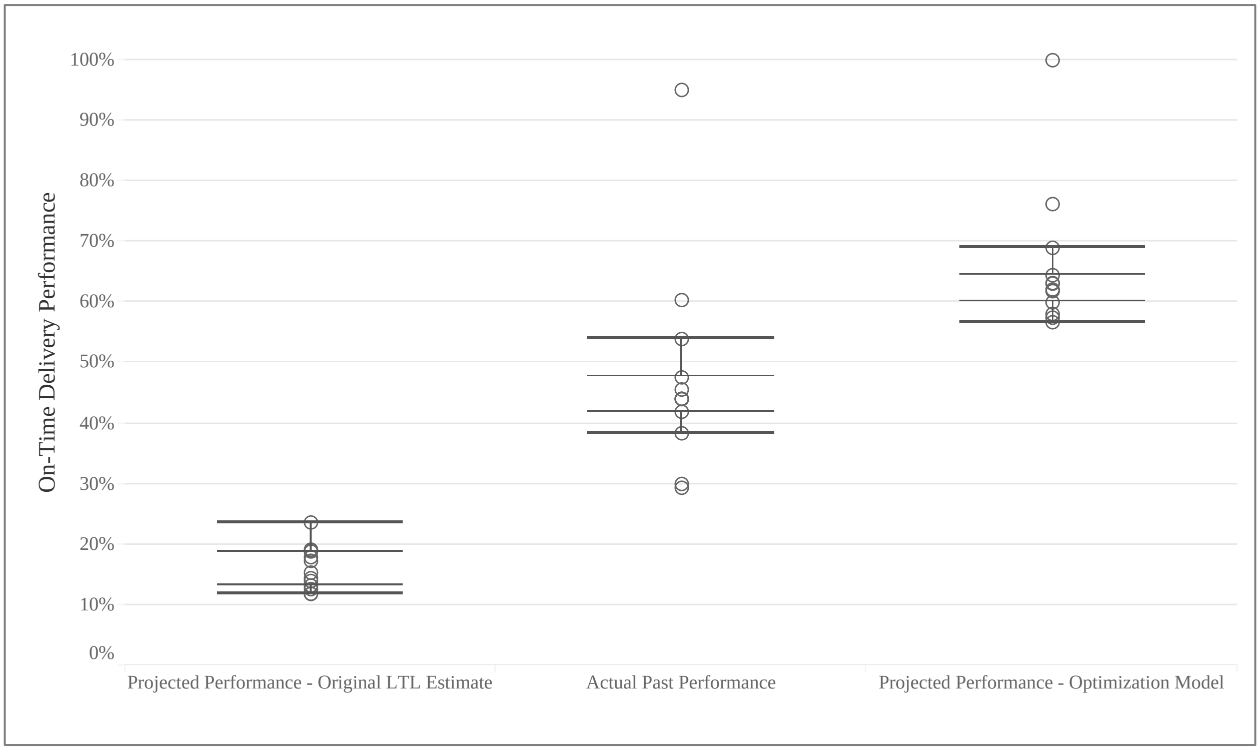

4.3 Optimization Model with Two-Day Buffer

The two-day buffer model builds on the previous model by increasing the buffer by one day. Again, when it is not possible to add buffer days, the earliest ship date in the shipping period is indicated as the only option, and o_{ij} is adjusted accordingly to include the extra lead time.

The on-time delivery performance increased by 60% versus the on-time delivery performance based on the estimated transit time and by 29% versus the actual on-time performance achieved by the company. The distribution of the results is displayed in Figure 10.

Figure 10 Projected Performance of the Two-Day Buffer Optimization Model Compared to ABC Transportation Group's Actual Past Performance and Projected Performance if the Original LTL Estimate Was Used

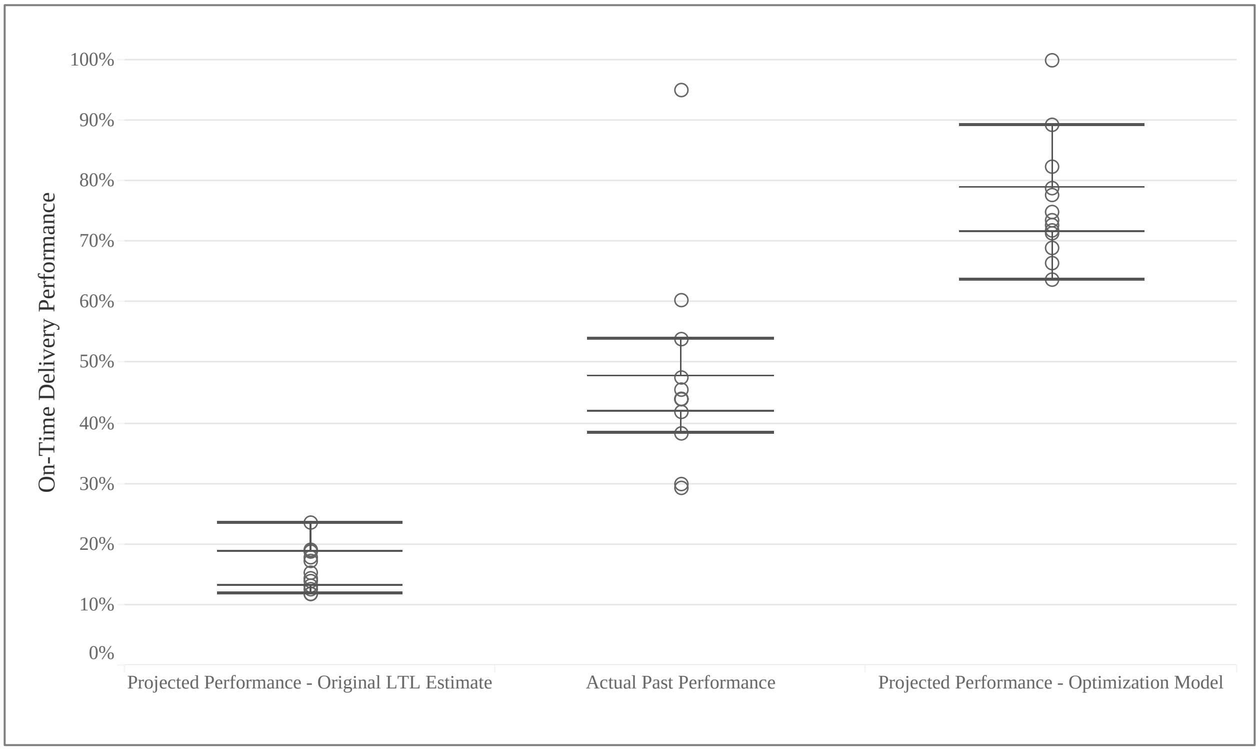

4.4 Optimization Model to Maximize Buffer Days

From the previous two models, it was apparent that an additional buffer provided significant gains in delivery performance. Therefore, the final model uses this concept to adjust the objective function to maximize the buffer time instead of using a fixed number of days as the buffer.

The formulation changes slightly, with the removal of o_{ij} and the addition of b_{ij}. The b_{ij} value can be less than 0 if the selected ship date is after the latest day an order can ship to be on time, per the predictive model.

Notation:

$x\colon_{ij}$: Binary variable, indicates whether an order *i* is shipped on day *j*

b_{ij}: Number of buffer days added to the lead time if shipped on day *j*

c_i: Cases per order *i*

C_j: Case capacity per day *j*

Formulation:

$$Maximize\colon \sum_{i \in I} \sum_{j \in J} x_{ij} b_{ij} \tag{6}$$

Subject to:

$$\sum_{i \in I} x_{ij} c_{ij} \leq C_j, \qquad \forall\, j = 1, \dots, m, \tag{7}$$

$$\sum_{j \in J} x_{ij} = 1, \qquad \forall\, i = 1, \dots, n, \tag{8}$$

$$x_{ij} = \{0,1\}. \tag{9}$$

This model was the most successful and showed a 41% improvement over the actual delivery performance, and a 72% improvement over the original estimate. The distribution of the results is displayed in Figure 11. A comparison of the results of all four of the optimization models introduced in this chapter is shown in Figure 12.

Figure 11 Projected Performance of the Two-Day Buffer Optimization Model Compared to ABC Transportation Group's Actual Past Performance and Projected Performance if the Original LTL Estimate Was Used

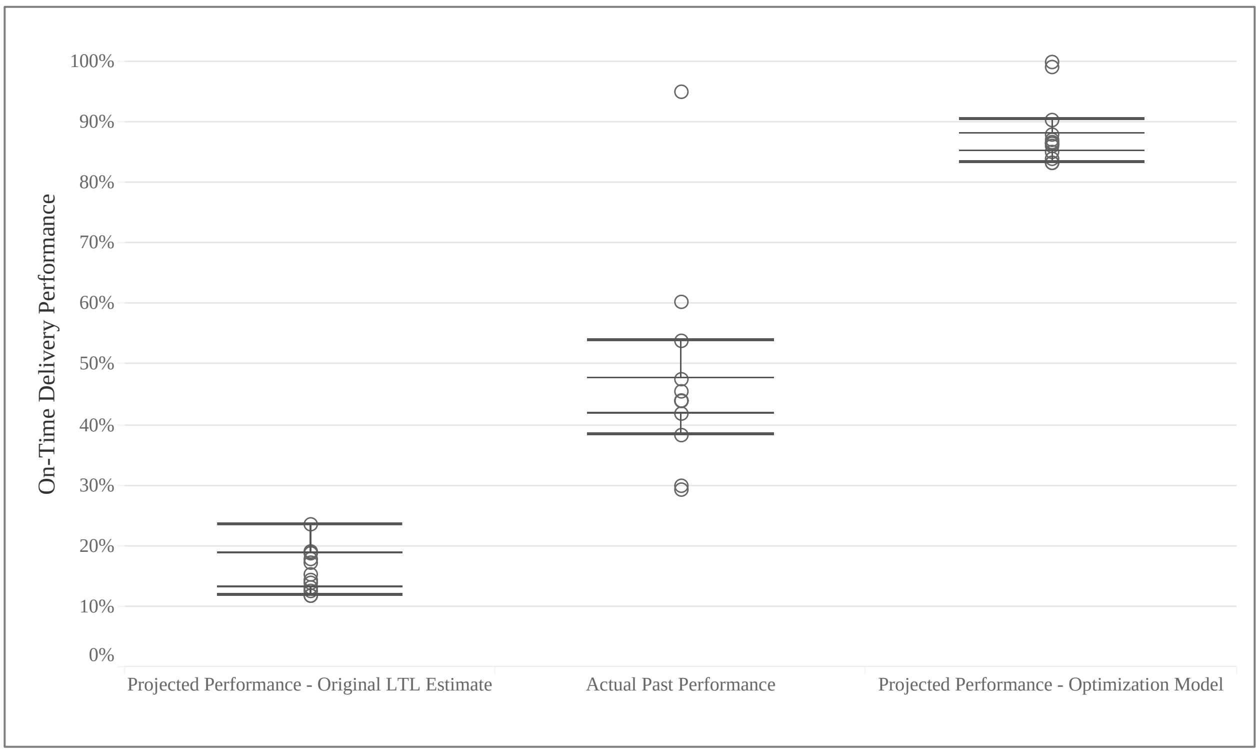

Figure 12 Relative Performance of the Four Optimization Models Introduced

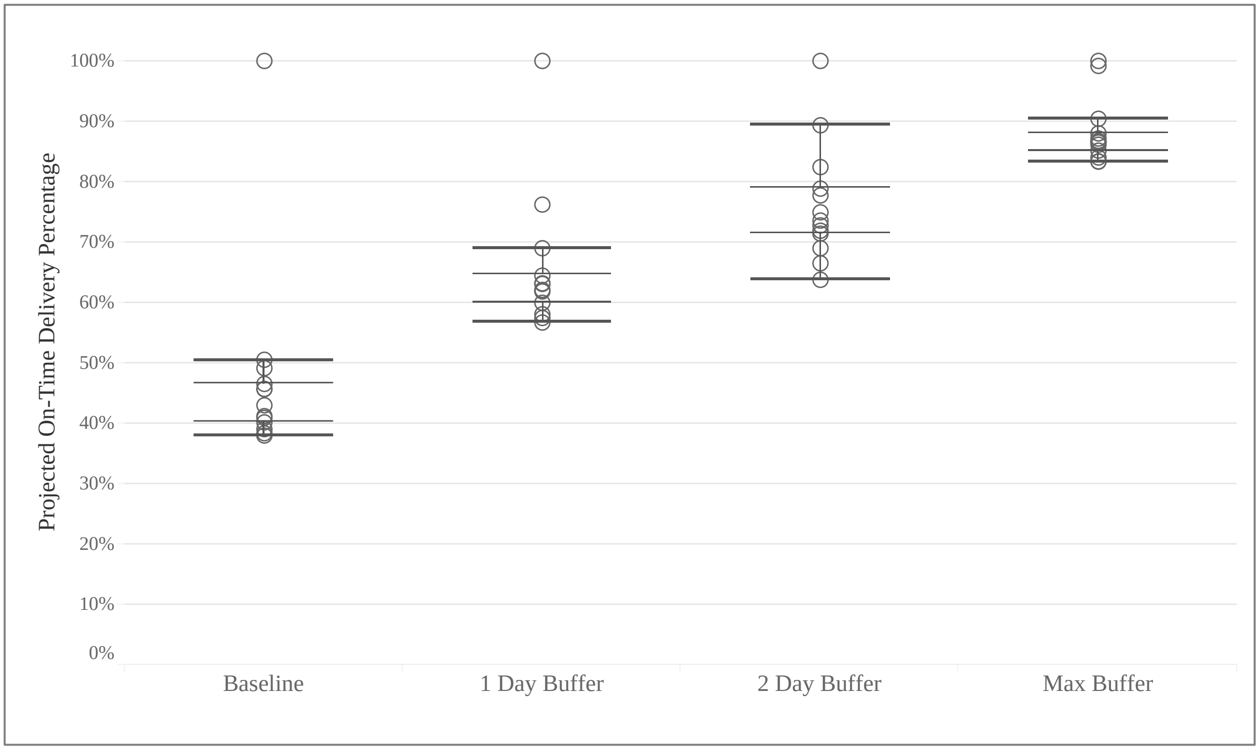

5 DISCUSSION

The results from the four models successfully illustrate the viability of the proposed method of combining a predictive model with an optimization model to improve delivery performance and reduce manual labor. Independently, the prediction model improved the transit time RMSE from 3.07 days to 1.97 days. Moreover, all four scenarios indicated at least a 25% gain in performance versus the delivery performance based on simple estimated transit times, which validates that the transit time predictions are superior to the original transit time estimates. However, by maximizing buffer time in the optimization model, the on-time delivery performance reached an average of 87% -- double the actual delivery performance achieved by ABC Transportation Group. This outcome suggests that by maximizing the buffer time added to the lead time, even on a short planning horizon like one week, there is a significant impact on performance. Furthermore, by automating the scheduling process, we can avoid manual planning errors, which had made up 6% of all late orders in 2020.

5.1 Managerial Insights

The two-pronged approach to scheduling shipments illustrated in this research paper shows immense promise in improving the current process. Perhaps the biggest takeaway from the results is that maximizing buffer time should be central to the scheduling strategy. Since this idea builds off of existing scheduling processes that sometimes include adding buffer days, this approach is also appealing from a change management standpoint.

Currently, the best model leverages the maximization of buffer time to slot orders subject to warehouse labor constraints. In the short term, future iterations may consider additional customization, such as constraints for minimum performance levels for specific customers or retailers, on top of the buffer maximization model. Further constraints related to warehouse labor capacity may also be included.

In the longer term, model expansions could include additional modes of transportation. Parcel shipments would be simple to include in the optimization model and may become more and more relevant as ABC Transportation Group's B2C or e-commerce divisions grow. Parcel transit times, unlike LTL transit times, have a high level of reliability and can be included in the optimization model without the predictive step. Truckload shipments, however, may not be a match for this model. Due to the additional steps needed to obtain specific pickup and delivery appointments and the need to be cautious of driver hours, scheduling to maximize the lead time does not make sense in this context.

5.2 Limitations

As mentioned throughout this paper, LTL shipments are very unpredictable due to the nature of the network and the delivery process to the consignees. Although time in transit is a significant component of the total transit time, the delivery appointment setting process and delivery schedule of a carrier to the recipient are equally critical to determining the timeliness of delivery. The external factors that lead to variance in the scheduling process are ambiguous and varied. Some features used in the random forest model, such as retailer and holiday, aim to capture this aspect of the process by proxy but can only approximate the phenomena. Moreover, some papers in the literature review included macro factors such as weather, traffic, and route in their predictive models. We could not incorporate these features into our model due to a lack of data, but these may be useful in future iterations.

The optimization model was quite successful and illustrated a significant improvement over the status quo. However, limitations exist in the approach. First, due to a lack of information on the actual case capacity, the constraints are simplistic and solely based on total demand in a week evenly allocated per day. There can feasibly be situations where the case capacity per day is uneven across the week, making it more challenging to schedule shipments optimally. Moreover, we ran the test scenarios on past data and used the original transit time to gauge delivery performance using the proposed schedule. Although this should be a reasonable approximation, the actual transit time may have differed from what we used in the test scenarios due to being shipped on a different day. This difference would result in some variation in the ultimate performance but should directionally be the same.

Finally, and perhaps most importantly, this model does not consider inventory availability. Inventory shortages consisted of 16% of the non-compliant deliveries in 2020. The practice of preventing orders from shipping to wait for inventory is unfortunately unavoidable. ABC Transportation Group should expect a slightly weaker delivery performance when this situation is necessary, and shipments may have to be manually rescheduled by analysts to account for the updated ship date.

6 CONCLUSION

Managing uncertainty well in LTL transportation significantly impacts both on-time delivery performance and operational efficiency for ABC Transportation Group and other companies who utilize this shipping mode. By linking transit time prediction via machine learning with integer programming to create an optimal shipment scheduling solution, this capstone demonstrated a marked improvement in overall delivery performance while allowing the sponsor company more control over resource management. Better usage of labor, whether at the warehouse for picking and shipping purposes or earlier in the process in the planning stage, is possible with this solution.

Ultimately, these two steps generated better results than either approach could accomplish alone. Delivery performance suffered when the estimated transit time was incorrect or if the lead time provided was too short. The solution presented in this capstone necessarily addresses both issues using different methods.

Although transit time prediction is a well-studied topic, the application of the final predictions in business processes is typically not discussed. This study, however, presents a situation where the results from the predictive model are utilized as inputs in a separate optimization model. The success of this two-pronged solution suggests that some overarching goals, such as improving delivery performance, may be best suited to be addressed using multiple methodologies that build off one another.

References

Alla, H., Moumoun, L., & Balouki, Y. (2021). Flight Arrival Delay Prediction Using Supervised Machine Learning Algorithms. In N. Gherabi & J. Kacprzyk (Eds.), *Intelligent Systems in Big Data, Semantic Web and Machine Learning* (pp. 231–246). Springer International Publishing. https://doi.org/10.1007/978-3-030-72588-4_16

Balster, A., Hansen, O., Friedrich, H., & Ludwig, A. (2020). An ETA Prediction Model for Intermodal Transport Networks Based on Machine Learning. *Business & Information Systems Engineering*, *62*(5), 403–416. https://doi.org/10.1007/s12599-020-00653-0

Barcos, L., Rodríguez, V., Álvarez, M. J., & Robusté, F. (2010). Routing design for less-than-truckload motor carriers using Ant Colony Optimization. *Transportation Research Part E: Logistics and Transportation Review*, *46*(3), 367–383. https://doi.org/10.1016/j.tre.2009.11.006

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

Chandramouli, S., Dutt, S., & Das, A. K. (2018). *Machine Learning*. Pearson Education India.

Conforti, M., Cornuéjols, G., & Zambelli, G. (2014). *Integer Programming* (Vol. 271). Springer International Publishing. https://doi.org/10.1007/978-3-319-11008-0

Liu, C. H. B., Chamberlain, B. P., Little, D. A., & Cardoso, Â. (2017). Generalising Random Forest Parameter Optimisation to Include Stability and Cost. In Y. Altun, K. Das, T. Mielikäinen, D. Malerba, J. Stefanowski, J. Read, M. Žitnik, M. Ceci, & S. Džeroski (Eds.), Machine Learning and Knowledge Discovery in Databases (Vol. 10536, pp. 102–113). Springer International Publishing. https://doi.org/10.1007/978-3-319-71273-4_9

Poschmann, P., Weinke, M., Balster, A., Straube, F., Friedrich, H., & Ludwig, A. (2019). Realization of ETA Predictions for Intermodal Logistics Networks Using Artificial Intelligence. In U. Clausen, S. Langkau, & F. Kreuz (Eds.), *Advances in Production, Logistics and Traffic* (pp. 155–176). Springer International Publishing.

Premkumar, P., Gopinath, S., & Mateen, A. (2020). Trends in third party logistics – the past, the present & the future. *International Journal of Logistics Research and Applications*, 1–30. https://doi.org/10.1080/13675567.2020.1782863

Salvendy, G. (Ed.). (2001). *Handbook of industrial engineering: Technology and operations management* (3rd ed). Wiley.

Schulz, J. D. (2021, July 12). *State of logistics 2021: Less than truckload (LTL)*. Recently Filed RSS. Retrieved November 27, 2021, from https://www.logisticsmgmt.com/article/state_of_logistics_2021_less_than_truckload_ltl.

Sheffi, Y. (1990). Third Party Logistics: Present and Future Prospects. *Journal of Business Logistics; Hoboken*, *11*(2), 27.

Shmueli & Koppius. (2011). Predictive Analytics in Information Systems Research. *MIS Quarterly*, *35*(3), 553. https://doi.org/10.2307/23042796

Simroth, A., & Zahle, H. (2011). Travel Time Prediction Using Floating Car Data Applied to Logistics Planning. *IEEE Transactions on Intelligent Transportation Systems*, *12*(1), 243–253. https://doi.org/10.1109/TITS.2010.2090521

Sun, X., Zhang, H., Tian, F., & Yang, L. (2018). The Use of a Machine Learning Method to Predict the Real-Time Link Travel Time of Open-Pit Trucks. *Mathematical Problems in Engineering*, *2018*, 4368045. https://doi.org/10.1155/2018/4368045

Truong, G. (2014). *Forecasting linehaul transit times & on time delivery probability using quantile regression forests* [Massachusetts Institute of Technology]. http://dspace.mit.edu/handle/1721.1/7582

United States less-than-truckload (LTL) market growth, trends, covid-19 impact, and forecasts 2021-2026 - researchandmarkets.com. Business Wire. (2021, June 16). Retrieved November 27, 2021, from https://www.businesswire.com/news/home/20210616005457/en/United-States-Less-than-truckload-LTL-Market-Growth-Trends-COVID-19-Impact-and-Forecasts-2021-2026---ResearchAndMarkets.com.

Urciuoli, L. (2018). An algorithm for improved ETAs estimations and potential impacts on supply chain decision making. *Procedia Manufacturing*, *25*, 185–193. https://doi.org/10.1016/j.promfg.2018.06.073

van der Spoel, S., Amrit, C., & van Hillegersberg, J. (2017). Predictive analytics for truck arrival time estimation: A field study at a European distribution centre. *International Journal of Production Research*, *55*(17), 5062–5078. https://doi.org/10.1080/00207543.2015.1064183

Williams, H. P. (2009). *Logic and Integer Programming* (Vol. 130). Springer US. https://doi.org/10.1007/978-0-387-92280-5

ZIP Code™ - The Basics. USPS.com faqs. (2021, March 20). Retrieved November 27, 2021, from https://faq.usps.com/s/article/ZIP-Code-The-Basics.

Appendix A – Pearson Correlation Chart

	OID	Y	Q	M	W	RID	DPC	CID	3DZ	H	P	W	CU	CS	AMZ	CHEW	CS	CVS	KEHE	MEI	UNFI	WALG	WALM	WMC	AT	ET	DIST	SID
OID	1.00	0.94	0.12	0.14	-0.03	0.14	0.01	-0.01	0.01	0.03	-0.04	0.02	-0.04	0.01	-0.01	-0.01	0.01	-0.01	0.01	0.01	-0.01	0.04	-0.02	0.01	0.12	0.02	0.02	-0.01
Y	0.94	1.00	-0.22	-0.21	-0.02	0.12	0.01	-0.03	0.01	-0.03	-0.04	0.02	-0.04	0.00	0.00	-0.01	0.01	0.00	0.01	0.01	-0.02	0.03	-0.02	0.02	0.11	0.02	0.02	-0.01
Q	0.12	-0.22	1.00	0.97	-0.02	0.04	0.00	0.06	0.00	0.14	0.02	0.00	0.01	0.01	-0.02	0.01	0.01	-0.03	0.01	0.01	0.02	0.00	-0.02	-0.01	0.03	-0.01	-0.01	0.00
M	0.14	-0.21	0.97	1.00	-0.02	0.04	0.00	0.05	0.00	0.18	0.01	0.00	0.01	0.01	-0.02	0.00	0.01	-0.03	0.01	0.00	0.02	0.00	-0.01	-0.01	0.03	0.00	0.00	0.00
W	-0.03	-0.02	-0.02	-0.02	1.00	0.00	-0.04	0.00	-0.04	-0.03	0.00	-0.01	0.00	-0.01	-0.06	0.06	0.00	-0.02	-0.01	0.00	0.02	0.02	0.02	0.08	-0.02	-0.02	-0.02	-0.02
RID	0.14	0.12	0.04	0.04	0.00	1.00	0.02	0.08	0.02	0.01	0.02	0.04	0.02	0.05	-0.11	-0.06	-0.05	-0.17	-0.13	0.00	-0.11	-0.12	-0.18	-0.05	-0.05	-0.03	-0.04	-0.01
DPC	0.01	0.01	0.00	0.00	-0.04	0.02	1.00	-0.03	1.00	0.00	-0.01	0.03	0.00	0.00	0.05	0.00	-0.17	-0.07	0.07	-0.02	-0.02	-0.01	0.01	-0.01	0.15	0.40	0.49	-0.18
CID	-0.01	-0.03	0.06	0.05	0.00	0.08	-0.03	1.00	-0.03	0.01	0.03	0.00	0.05	0.02	0.06	-0.07	0.07	-0.22	-0.23	-0.24	0.08	0.10	0.03	-0.01	-0.16	0.01	0.00	0.01
3DZ	0.01	0.01	0.00	0.00	-0.04	0.02	1.00	-0.03	1.00	0.00	-0.01	0.03	0.00	0.00	0.05	0.00	-0.17	-0.07	0.07	-0.02	-0.02	-0.01	0.01	-0.01	0.15	0.40	0.49	-0.18
H	0.03	-0.03	0.14	0.18	-0.03	0.01	0.00	0.01	0.00	1.00	0.01	0.01	0.01	0.00	-0.01	0.01	0.01	-0.02	0.01	0.00	0.00	-0.01	-0.01	0.03	0.02	0.00	0.01	0.00
P	-0.04	-0.04	0.02	0.01	0.00	0.02	-0.01	0.03	-0.01	0.01	1.00	0.70	0.83	0.73	-0.04	0.19	-0.02	-0.13	0.04	0.01	0.07	-0.12	-0.07	-0.06	0.01	-0.08	-0.09	0.01
W	0.02	0.02	0.00	0.00	-0.01	0.04	0.03	0.00	0.03	0.01	0.70	1.00	0.67	0.62	-0.03	0.04	-0.03	-0.11	0.02	0.01	0.03	-0.07	-0.03	-0.07	-0.04	-0.11	-0.11	0.00
CU	-0.04	-0.04	0.01	0.01	0.00	0.02	0.00	0.05	0.00	0.01	0.83	0.67	1.00	0.66	-0.06	0.14	-0.02	-0.15	0.04	0.01	0.09	-0.12	-0.05	-0.06	-0.01	-0.08	-0.09	0.01
CS	0.01	0.00	0.01	0.01	-0.01	0.05	0.00	0.02	0.00	0.00	0.73	0.62	0.66	1.00	-0.06	0.09	-0.01	-0.14	0.09	0.01	0.06	-0.12	-0.05	-0.07	0.00	-0.05	-0.06	-0.01
AMZ	-0.01	0.00	-0.02	-0.02	-0.06	-0.11	0.05	0.06	0.05	-0.01	-0.04	-0.03	-0.06	-0.06	1.00	-0.05	-0.05	-0.10	-0.07	-0.06	-0.06	-0.10	-0.12	-0.05	0.12	-0.02	0.00	-0.03
CHEW	-0.01	-0.01	0.01	0.00	0.06	-0.06	0.00	-0.07	0.00	0.01	0.19	0.04	0.14	0.09	-0.05	1.00	-0.02	-0.05	-0.03	-0.03	-0.03	-0.05	-0.06	-0.02	0.14	0.01	0.02	0.02
CS	0.01	0.01	0.01	0.01	0.00	-0.05	-0.17	0.07	-0.17	0.01	-0.02	-0.03	-0.02	-0.01	-0.05	-0.02	1.00	-0.05	-0.03	-0.03	-0.03	-0.05	-0.06	-0.02	-0.02	-0.02	-0.04	0.07
CVS	-0.01	0.00	-0.03	-0.03	-0.02	-0.17	-0.07	-0.22	-0.07	-0.02	-0.13	-0.11	-0.15	-0.14	-0.10	-0.05	-0.05	1.00	-0.07	-0.06	-0.06	-0.10	-0.12	-0.04	0.26	0.03	0.06	0.04
KEHE	0.01	0.01	0.01	0.01	-0.01	-0.13	0.07	-0.23	0.07	0.01	0.04	0.02	0.04	0.09	-0.07	-0.03	-0.03	-0.07	1.00	-0.04	-0.04	-0.07	-0.09	-0.03	0.13	0.04	0.06	-0.07
MEI	0.01	0.01	0.01	0.00	0.00	0.00	-0.02	-0.24	-0.02	0.00	0.01	0.01	0.01	0.01	-0.06	-0.03	-0.03	-0.06	-0.04	1.00	-0.03	-0.06	-0.07	-0.03	-0.15	-0.20	-0.18	0.04
UNFI	-0.01	-0.02	0.02	0.02	0.02	-0.11	-0.02	0.08	-0.02	0.00	0.07	0.03	0.09	0.06	-0.06	-0.03	-0.03	-0.06	-0.04	-0.03	1.00	-0.06	-0.08	-0.03	0.07	0.03	0.02	0.02
WALG	0.04	0.03	0.00	0.00	0.02	-0.12	-0.01	0.10	-0.01	-0.01	-0.12	-0.07	-0.12	-0.12	-0.10	-0.05	-0.05	-0.10	-0.07	-0.06	-0.06	1.00	-0.13	-0.05	-0.22	0.06	0.03	-0.07
WALM	-0.02	-0.02	-0.02	-0.01	0.02	-0.18	0.01	0.03	0.01	-0.01	-0.07	-0.03	-0.05	-0.05	-0.12	-0.06	-0.06	-0.12	-0.09	-0.07	-0.08	-0.13	1.00	-0.06	-0.12	0.02	-0.01	-0.01
WMC	0.01	0.02	-0.01	-0.01	0.08	-0.05	-0.01	-0.01	-0.01	0.03	-0.06	-0.07	-0.06	-0.07	-0.05	-0.02	-0.02	-0.04	-0.03	-0.03	-0.03	-0.05	-0.06	1.00	0.01	0.00	0.01	-0.05
AT	0.12	0.11	0.03	0.03	-0.02	-0.05	0.15	-0.16	0.15	0.02	0.01	-0.04	-0.01	0.00	0.12	0.14	-0.02	0.26	0.13	-0.15	0.07	-0.22	-0.12	0.01	1.00	0.37	0.40	-0.10
ET	0.02	0.02	-0.01	0.00	-0.02	-0.03	0.40	0.01	0.40	0.00	-0.08	-0.11	-0.08	-0.05	-0.02	0.01	-0.02	0.03	0.04	-0.20	0.03	0.06	0.02	0.00	0.37	1.00	0.91	-0.25
DIST	0.02	0.02	-0.01	0.00	-0.02	-0.04	0.49	0.00	0.49	0.01	-0.09	-0.11	-0.09	-0.06	0.00	0.02	-0.04	0.06	0.06	-0.18	0.02	0.03	-0.01	0.01	0.40	0.91	1.00	-0.27
SID	-0.01	-0.01	0.00	0.00	-0.02	-0.01	-0.18	0.01	-0.18	0.00	0.01	0.00	0.01	-0.01	-0.03	0.02	0.07	0.04	-0.07	0.04	0.02	-0.07	-0.01	-0.05	-0.10	-0.25	-0.27	1.00

KEY	
OID	Order ID
Y	Year of Date Req Delivery
Q	Quarter of Date Req Delivery
M	Month of Date Req Delivery
W	Weekday of Date Req Delivery
RID	Retailer ID
DPC	Destination PostCode
CID	Carrier ID
3DZ	3 Digit Zip
H	HolidayNum
P	Pallets
W	Weight
CU	Cube
CS	Cases
AMZ	Amazon?
CHEW	Chewy?
CS	CS?
CVS	CVS?
KEHE	Kehe?
MEI	Meijer?
UNFI	UNFI?
WALG	Walgreens?
WALM	Walmart?
WMC	Walmartcom?
AT	Actual Transit
ET	LTL Estimated Transit
DIST	Distance
SID	State ID

Appendix B – Feature Importance Runs

Features	Run 1	Run 2	Run 3	Run 4	Run 5	Run 6	Run 7
Weight	0.113049	0.195093	0.199563	0.199904	0.202334	0.207079	0.243487
Distance	0.09035	0.104442	0.110628	0.114571	0.121343	0.131136	0.132772
Month of Date Req Delivery	0.0794	0.114657	0.115041	0.11657	0.11614	0.114414	0.159546
Weekday of Date Req Delivery	0.057079	0.087781	0.089039	0.090707	0.090047	0.089825	0.099344
Destination PostCode	0.061967	0.070634	0.078632	0.079245	0.085096	0.094618	0.089262
Retailer ID	0.043221	0.047997	0	0.07437	0.076716	0.081047	0.079475
3 Digit Zip	0.054507	0.061195	0.06654	0.065058	0.07055	0.079989	0.073032
Carrier ID	0.047001	0.062451	0.065426	0.065239	0.064001	0.062769	0.06583
Year of Date Req Delivery	0.041523	0.058642	0.059079	0.056547	0.054586	0.052071	0.057251
Quarter of Date Req Delivery	0.038933	0.052755	0.052132	0.050522	0.048924	0.046416	0
HolidayNum	0.026565	0.039031	0.039111	0.041102	0.041167	0.040635	0
State ID	0.025887	0.030516	0.030536	0.02715	0.029094	0	0
LTL Estimated Transit	0.02004	0.022551	0.023234	0.019014	0	0	0
Cases	0.108086	0	0	0	0	0	0
Cube	0.103292	0	0	0	0	0	0
Pallets	0.041076	0	0	0	0	0	0
Walgreens?	0.014003	0.014569	0.017076	0	0	0	0
Amazon?	0.008453	0.009187	0.011428	0	0	0	0
CVS?	0.005349	0.005614	0.010001	0	0	0	0
Walmart?	0.005252	0.005919	0.00866	0	0	0	0
Meijer?	0.004251	0.004714	0.00583	0	0	0	0
Chewy?	0.003309	0.003671	0.003918	0	0	0	0
UNFI?	0.002558	0.002863	0.005502	0	0	0	0
Kehe?	0.002385	0.002765	0.005003	0	0	0	0
Walmartcom?	0.001256	0.001518	0.001862	0	0	0	0
CS?	0.00121	0.001436	0.001759	0	0	0	0
RMSE	**2.01**	**1.99**	**2**	**1.98**	**1.97**	**1.97**	**1.98**