

Article

Predicting and Mitigating Delays in Cross-Dock Operations: A Data-Driven Approach

Amna Altaf ^{1,2}, Adeel Mehmood ^{2,*} , Adnen El Amraoui ¹, François Delmotte ¹ and Christophe Lecoutre ³

¹ UR 3926 Laboratoire de Génie Informatique et d'Automatique de l'Artois (LGI2A), University of Artois, F-62400 Béthune, France; amna.altaf@univ-artois.fr or a.altaf@hull.ac.uk (A.A.); francois.delmotte@univ-artois.fr (F.D.)

² School of Computer Science and Technology, Faculty of Science and Engineering, University of Hull, Hull HU6 7RX, UK

³ CRIL-CNRS, UMR 8188, University of Artois, F-62307 Lens, France; christophe.lecoutre@univ-artois.fr

* Correspondence: a.mehmood@hull.ac.uk

Abstract: Cross-docking operations are highly dependent on precise scheduling and timely truck arrivals to ensure streamlined logistics and minimal storage costs. Predicting potential delays in truck arrivals is essential to avoiding disruptions that can propagate throughout the cross-dock facility. This paper investigates the effectiveness of deep learning models, including Convolutional Neural Networks (CNN), Multilayer Perceptrons (MLPs), and Recurrent Neural Networks (RNNs), in predicting late arrivals of trucks. Through extensive comparative analysis, we evaluate the performance of each model in terms of prediction accuracy and applicability to real-world cross-docking requirements. The results highlight which models can most accurately predict delays, enabling proactive measures for handling deviations and improving operational efficiency. Our findings support the potential for deep learning models to enhance cross-docking reliability, ultimately contributing to optimized logistics and supply chain resilience.

Keywords: cross-dock; truck to door assignment; temporary storage; fuzzy chance constrained optimization; uncertainty; optimization programming language



Academic Editor: Wei Zhu

Received: 22 November 2024

Revised: 20 December 2024

Accepted: 23 December 2024

Published: 20 January 2025

Citation: Altaf, A.; Mehmood, A.; El Amraoui, A.; Delmotte, F.; Lecoutre, C. Predicting and Mitigating Delays in Cross-Dock Operations: A Data-Driven Approach. *Stats* **2025**, *8*, 9. <https://doi.org/10.3390/stats8010009>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Cross-docking is a pivotal logistics strategy designed to minimize storage time and maximize operational efficiency by enabling the direct transfer of goods from inbound to outbound trucks. In this tightly coordinated system, the timely arrival of trucks at cross-docking facilities is crucial for maintaining a seamless flow of goods. Delays in truck arrivals can disrupt carefully planned schedules and resource allocations, leading to increased handling and temporary storage costs. These disruptions can cascade through the supply chain, causing delays in deliveries, increased operational expenses, and diminished customer satisfaction [1].

While substantial advancements have been made in optimizing cross-docking workflows, the inherent uncertainty and variability in truck arrival times remain significant challenges [2,3]. Addressing this issue is critical for ensuring the reliability and efficiency of cross-docking operations. Figure 1 provides an illustrative layout of a cross-docking facility, showcasing the coordination required between inbound and outbound operations and highlighting the importance of synchronized scheduling [4].

Recent research has focused on enhancing cross-docking operations through a variety of optimization methods. Studies in this area have primarily concentrated on

scheduling [5], vehicle routing [6,7], and inventory management [8], with the goal of minimizing operational costs and improving throughput efficiency [2,9]. However, these approaches frequently rely on predictable static schedules, overlooking the variability introduced by factors such as traffic conditions, weather disturbances, and unexpected logistical issues. This limitation constrains the effectiveness of traditional optimization models, which lack mechanisms to accommodate real-time changes in arrival schedules and as such are unable to proactively address delays [10,11].




Figure 1. Layout of a cross-docking facility.

Given the critical importance of timely truck arrivals for cross-docking efficiency, this paper introduces a predictive model designed to anticipate potential delays while allowing for preemptive adjustments to mitigate operational disruptions. Because empirical data specific to arrival delays in cross-docking facilities are often limited, we constructed a representative hypothetical dataset that simulates realistic truck arrival patterns, incorporating factors such as historical delay trends, traffic conditions, and random disturbances. The primary objective of this predictive model is not only to forecast arrival delays but also to serve as a proactive tool for operational continuity, thereby enabling cross-docking facilities to adapt their processes in real-time and minimize the impact of potential delays.

In this study, we assess the performance of three advanced deep learning architectures: Multilayer Perceptron (MLP), Recurrent Neural Network (RNN), and Convolutional Neural Network (CNN). Each model offers distinct advantages in predictive modeling. MLP, as a feed-forward neural network, is effective in learning from historical data patterns. RNNs, which are inherently suited for sequential data, are adept at capturing temporal dependencies within time-series data—a critical feature when modeling truck arrival patterns [12]. In contrast, CNNs are traditionally used for spatial data processing, and have shown promise in feature extraction when applied to time-series data [13,14]. CNNs can detect latent patterns that may influence arrival times, making them a powerful tool for prediction in complex environments [15,16]. By evaluating the strengths and limitations of each model, we aim to identify the most suitable architecture for predicting late arrivals and generating timely alerts for operational adjustments.

To enhance the reliability of our predictions, we employ a rigorous hyperparameter optimization process across all three models. Fine-tuning these parameters is essential for achieving high accuracy and model stability, particularly in the dynamic cross-docking environment, where operational efficiency is sensitive to real-time changes in truck arrival schedules.

The primary contributions of this research are as follows: (1) introducing a novel predictive approach for forecasting late truck arrivals in cross-docking facilities; (2) implementing and comparing the performance of MLP, RNN, and CNN models on a simulated

dataset tailored to reflect cross-docking arrival patterns; and (3) performing systematic hyperparameter optimization to maximize prediction accuracy and model applicability. Our findings provide actionable insights for logistics managers, enabling them to proactively manage delays and improve the overall efficiency of cross-docking operations.

The remainder of this paper outlines our methodological approach, describes the experimental setup, and presents an analysis of the models' performance results. Through this research, we bridge the gap between advanced deep learning techniques and logistical optimization, contributing to the growing body of literature on dynamic data-driven solutions in supply chain management.

2. Problem Formulation and Predictive Modeling Methodology

In this section, we provide a structured approach to tackling the challenge of delayed truck arrivals in cross-docking facilities. We first outline the problem formulation, followed by a detailed description of the predictive modeling methodology encompassing data generation, model selection, and evaluation.

2.1. Problem Description

Timely arrival of trucks at cross-docking facilities is essential to maintaining synchronized operations. Any deviation from scheduled arrival times can cause significant operational disruptions, leading to resource inefficiencies and increased handling costs. Efficient cross-docking operations require precise synchronization between incoming and outgoing trucks. Any delays in truck arrival disrupt this synchronization, leading to potential congestion, underutilized resources, and increased operational costs. This study addresses the problem of predicting late arrivals of trucks at cross-docking facilities in order to preemptively mitigate the operational disruptions caused by such delays. Given the absence of publicly available datasets for cross-docking truck arrival times, we have generated a hypothetical dataset designed to capture the complexities of real-world conditions. This dataset includes key features such as historical arrival times, traffic conditions, seasonal and weather-related delays, and other unforeseen events that commonly impact truck arrival schedules.

The primary goal of our study is to develop a reliable predictive model that can forecast late arrivals, enabling cross-dock operators to take proactive measures. By achieving accurate predictions, this model can facilitate better scheduling adjustments, improve resource allocation, and reduce bottlenecks, ultimately contributing to smoother cross-docking operations.

2.2. Methodology

To address this issue, we employ deep learning models capable of capturing complex patterns within arrival time data. This section describes the data preparation steps, model selection process, hyperparameter tuning, evaluation metrics, and experimental setup used to ensure accurate predictions of late arrivals.

1. **Data Generation and Preprocessing.** Given the absence of empirical data, a synthetic dataset was created to simulate real-world scenarios. The data generation process incorporated multiple variables that are known to influence arrival times, such as:
 - Historical Arrival Patterns: Time series data reflecting previous truck arrivals, delays, and patterns.
 - Traffic Conditions: Indicators of traffic density, road congestion, and expected delays based on common traffic patterns.
 - Unforeseen Delays: Data points representing unpredictable factors such as road closures or breakdowns modeled to simulate real-world disruptions.

These data were then preprocessed to ensure compatibility with deep learning models. Features were normalized to a common scale and missing data points were imputed where necessary. Finally, the dataset was divided into training, validation, and testing sets while ensuring that the time series nature of the data were preserved, particularly for models such as RNNs that rely on sequential patterns.

2. **Predictive Model Selection.** To identify the most effective model for predicting truck arrival delays, we explored three prominent deep learning architectures:

- **Multilayer Perceptron (MLP):** This feed-forward neural network is designed to capture historical patterns and relationships in the data. We used MLP as a baseline, serving as a control model for benchmarking against more advanced architectures.
- **Recurrent Neural Network (RNN):** Given the sequential nature of arrival times and delay patterns, the RNN architecture is well-suited for our task. RNNs excel at learning from time-dependent data, capturing temporal relationships that help forecast future arrival times based on past sequences.
- **Convolutional Neural Network (CNN):** Although CNNs are traditionally used for spatial data, they have shown promise in feature extraction from time series data by identifying underlying patterns. A CNN model was adapted for this application to capture spatial and sequential features.

3. **Hyperparameter Optimization.** To enhance the model's predictive accuracy, we performed systematic hyperparameter optimization across all experimental configurations. Techniques such as grid search and random search were employed to identify the optimal settings that maximized model performance. This rigorous process involved testing over 2000 experimental arrangements to assess the impact of various hyperparameters on the model's predictive performance. Given that this project involves a classification task, our optimization approach targeted maximizing accuracy with a discrete categorical outcome.

Hyperparameter optimization was conducted in two stages:

- (a) **Basic Network Optimization (BNO):** In this stage, fundamental hyperparameters were tuned. Basic parameters included those essential for defining the model structure and training behavior, such as batch size, number of epochs, neuron count, and maximum input length. Optimizing these foundational parameters ensures the network's structural integrity, as these values have a direct impact on the model's convergence and generalization ability.
- (b) **Advanced Network Optimization (ANO):** In this second stage, additional more subtle parameters were analyzed to further refine the model's performance. These advanced features included weight initialization techniques, batch normalization, and varying the model architecture. By examining these parameters, we aimed to improve training stability and potentially speed up convergence. Each advanced feature's impact was tested across diverse datasets to assess its effect on the network's performance under varying data conditions.

4. **Evaluation Metrics.** Model performance was assessed using standard metrics relevant to prediction accuracy, including:

- **Mean Absolute Error (MAE):** Used to capture the average deviation between predicted and actual arrival times.
- **Root Mean Squared Error (RMSE):** Used to penalize larger errors more heavily, offering insight into significant deviations.
- **Accuracy:** Used for overall model efficacy, particularly in binary classification (on-time vs. delayed).

Employing these metrics ensures a comprehensive evaluation of each model's effectiveness in capturing arrival delays.

5. **Experimental Setup.** All models were implemented using Python and TensorFlow, with experiments conducted on a machine equipped with high-performance GPUs to support deep learning training. A consistent experimental environment was maintained across all models to allow for fair and accurate comparisons.

This methodology establishes a framework for evaluating the predictive capabilities of MLP, RNN, and CNN models, and the results demonstrate which architecture is most effective for forecasting truck delays in a cross-docking context. The insights gained from this study aim to offer actionable intelligence for logistics operators, facilitating proactive management of cross-dock operations amid unpredictable arrival schedules.

3. Model Architecture and Experimental Configurations

This section describes the model architectures and configurations designed for predicting late truck arrivals. Our objective is to identify a robust and adaptable configuration of deep learning algorithms that achieves stable and reliable results across varied datasets. We first provide an overview of the three key processing units within the predictive system, followed by details of each experimental configuration involving the MLP, RNN, and CNN architectures.

The system architecture for predicting late truck arrivals at cross-docking facilities is organized into a multistage process, with each component tailored to ensure efficient data handling, model training, and outcome generation. This architecture is divided into three main units:

- Input Preprocessing Unit
- Main Processing Unit
- Output Processing Unit

Each unit plays a specific role within the prediction system, transforming raw data into actionable insights for logistics operators. The overall system is presented in Figure 2, reflecting the use of MLP model; the architectures of the CNN and RNN models are presented in Figure 3.




Figure 2. General representation of the system with MLP.




Figure 3. General representation of the RNN and CNN model architectures.

3.1. Input Preprocessing Unit

This unit begins with synthetic data generation, given the absence of real-world datasets; the generated dataset includes variables such as historical arrival times, traffic conditions, weather data, and time-based patterns. Each data point undergoes preprocessing, which involves several critical steps.

First, to normalize the range of variables and prevent features with larger scales from disproportionately impacting model performance, features are scaled to a common range, such as $[0, 1]$. This can be achieved using Min–Max Scaling, mathematically defined as

$$X_{\text{scaled}} = \frac{X - \min(X)}{\max(X) - \min(X)},$$

where X represents the feature vector, $\min(X)$ is the minimum value, and $\max(X)$ is the maximum value of X .

Where data points are incomplete (e.g., missing traffic or weather data), values are imputed using statistical or synthetic methods. A common method is mean imputation.

$$X_{\text{imputed}} = \begin{cases} X_i & \text{if } X_i \text{ is observed} \\ \frac{\sum_{i=1}^n X_i}{n} & \text{if } X_i \text{ is missing} \end{cases}$$

Given the sequential nature of time series data, particularly for models such as RNNs, the data points are organized chronologically, letting $\{x_t\}_{t=1}^T$ represent a sequence of input features over time t . This enables the models to learn temporal dependencies crucial for accurate predictions.

3.2. Main Processing Unit

This is the core unit where predictive modeling occurs. We tested three primary deep learning architectures.

MLP captures general patterns without accounting for temporal dependencies. The mathematical representation of an MLP with one hidden layer is

$$\mathbf{h} = \sigma(\mathbf{W}_1 \mathbf{x} + \mathbf{b}_1), \quad \hat{y} = \mathbf{W}_2 \mathbf{h} + \mathbf{b}_2,$$

where $\mathbf{x}$ is the input vector, $\mathbf{W}_1, \mathbf{W}_2$ are weight matrices, $\mathbf{b}_1, \mathbf{b}_2$ are biases, and σ is an activation function (e.g., ReLU or sigmoid).

Given the time-sequential nature of truck arrival data, RNN models (especially LSTM models) capture time-dependent trends and patterns. An RNN updates its hidden state $\mathbf{h}_t$ as follows:

$$\mathbf{h}_t = \sigma(\mathbf{W}_h \mathbf{h}_{t-1} + \mathbf{W}_x \mathbf{x}_t + \mathbf{b}).$$

For LSTM models, the cell state $\mathbf{c}_t$ and hidden state $\mathbf{h}_t$ are updated using gating mechanisms:

$$\begin{aligned} \mathbf{f}_t &= \sigma(\mathbf{W}_f[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_f), \quad \mathbf{i}_t = \sigma(\mathbf{W}_i[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_i) \\ \mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}_c[\mathbf{h}_{t-1}, \mathbf{x}_t] + \mathbf{b}_c), \quad \mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \end{aligned}$$

where $\mathbf{f}_t, \mathbf{i}_t, \mathbf{o}_t$ are the forget, input, and output gates, respectively.

Known for their feature extraction capabilities, CNNs are used to capture localized patterns. For time series data, a 1D convolution is applied:

$$\mathbf{h}_t = \sigma \left(\sum_{k=1}^K \mathbf{w}_k \cdot \mathbf{x}_{t+k} + b \right)$$

where K is the kernel size, $\mathbf{w}_k$ are filter weights, and b is the bias term.

Hyperparameters such as the learning rate (α), number of layers (L), and number of units (U) are systematically tuned. For instance, optimizing α might involve minimizing the loss L using gradient descent:

$$\theta \leftarrow \theta - \alpha \nabla_{\theta} L.$$

3.3. Output Processing Unit

This final unit transforms model outputs into meaningful predictions for logistics operators.

The output from the deep learning models is a continuous score p , representing the probability of delay. For binary classification, a threshold τ is applied.

$$y_{\text{pred}} = \begin{cases} \text{Delayed} & \text{if } p \geq \tau \\ \text{On-Time} & \text{if } p < \tau \end{cases}$$

Additional metrics are generated as well, such as confidence intervals and predictions over time; for instance, the confidence score can be expressed as

$$\text{Confidence} = |p - \tau|.$$

These insights allow for real-time adjustments to schedules, resource allocation, and operational strategies.

Experimental Configuration

- **Experimental Configuration I:** Our first configuration utilizes the Multilayer Perceptron (MLP), a feedforward neural network. The MLP is structured with three layers, namely, input, hidden, and output, where each neuron in the hidden layer aids in classification tasks. We leverage MLP's strength in handling random variables and categorizing complex inputs to examine its effectiveness in predicting late arrivals.
- **Experimental Configuration II:** In this setup, we apply Recurrent Neural Networks (RNNs) capable of sequentially processing data and retaining “memory” of previous states. Given the sequential nature of truck arrival times, RNNs, particularly their

Vanilla and LSTM variants, are explored to analyze their ability to capture time-dependent patterns in arrival data.

- **Experimental Configuration III:** The final configuration employs Convolutional Neural Networks (CNNs). Known for their efficiency in minimal preprocessing and spatial pattern recognition, CNNs provide a robust option for handling diverse data inputs. We evaluated CNNs on their ability to recognize relevant patterns that could contribute to predicting truck arrival delays.

3.4. Data Synthesis for Experimentation

The prediction of truck arrival times at cross-docking facilities is critical for optimizing logistics operations and enhancing supply chain efficiency. However, the lack of available real-world data poses significant challenges in developing robust predictive models. To address this gap, we employed a data synthesis approach that enabled the generation of a realistic dataset capturing various factors influencing truck arrival times.

Data synthesis involves creating a dataset that mimics the characteristics and distributions of real-world data to enable the training and validation of predictive models even in the absence of historical data. Our synthesis process comprised the following steps:

- **Defining Key Variables.** To effectively simulate truck arrival times, we identified the key variables that are known to influence delivery performance, as presented in Table 1.

Table 1. Key Variables and Description for the Dataset.

Variable	Description
Departure Time	The timestamp when the truck leaves its origin.
Distance	The total distance (in miles) from the origin to the destination, impacting travel time.
Average Speed	The speed of the truck during transit, influenced by traffic conditions and road types.
Traffic Conditions	Variations in traffic density (e.g., Light, Moderate, Heavy) affecting travel speed.
Weather Conditions	External environmental factors (e.g., Clear, Rain, Snow) influencing driving conditions and speed.
Time of Day	Different traffic patterns associated with morning, afternoon, and evening hours.
Route Characteristics	Attributes of the route, such as urban versus rural areas, and the number of stops along the way.

- **Generating Synthetic Data.** The synthetic dataset was generated using a Python-based simulation approach. We designed a program to create data points based on random selections and distributions for each variable. The following considerations were taken into account during the simulation:
 - Distances were randomly generated within a specified range (e.g., 10 to 500 miles) to reflect typical logistics operations.
 - A baseline average speed of 60 mph was established, with adjustments made based on traffic and weather conditions; for instance, under heavy traffic, the speed was reduced by 50%, while adverse weather conditions further decreased speed by 15–40%.
 - The travel time was computed as the ratio of distance to adjusted speed, resulting in realistic transit times.

- Departure times were randomized within a 30-day window, and arrival times were calculated by adding the travel time to the departure time.

Throughout optimization, we tracked the highest scores of validation loss and accuracy for evaluation purposes. Cross-entropy was used as the primary loss function, and accuracy served as the principal performance metric. For model implementation, Keras version 2.15.0 was employed, with layer structures configured to ensure modular, adaptable design. Sigmoid layers were used for binary classification, while softmax layers were used to address multiclass problems. Additionally, we applied an early stopping mechanism to terminate training if the validation accuracy and loss showed no improvement over five consecutive epochs, thereby preventing overfitting and unnecessary iterations.

3.5. Performance Measure Analysis

Upon data preparation and model configuration, we analyzed algorithmic performance through network optimization. This analysis was split into two phases:

- **Basic Parameter Variation:** We first tested the effect of basic hyperparameters, including the number of epochs, dropout rate, and batch size. Adjustments were made systematically to identify the configuration yielding the highest prediction accuracy and stability. Fine-tuning these parameters allowed for baseline performance insights before adding complexity with advanced features.
- **Advanced Feature Variation:** Following basic optimization, we introduced advanced network features such as weight initialization methods, batch normalization, and model variants in order to evaluate their impact on model accuracy and robustness. These refinements allowed us to assess the model's responsiveness to deeper architectural changes.

For more robust accuracy, instead of a fixed 80–20 training–testing split, we implemented k -fold cross-validation using the scikit-learn library. This method shuffles the dataset and splits it into k sections, allowing each subset to serve as both training and testing data across multiple iterations. This process reduces variability, offers insights into model consistency, and helps to normalize outliers by calculating an averaged final accuracy score. Tolerance levels for accuracy deviation across iterations were also evaluated, allowing us to gauge the model's performance range against specific hyperparameter settings.

4. Model Evaluation and Results

We conducted simulation studies using Python 3.13.1 in the Google Colab environment. Several key libraries were utilized to perform data analysis, modeling, and computational tasks, including NumPy 2.2.0, SciPy 1.15.1, Pandas 2.2.3, Matplotlib 3.9.0, and TensorFlow 2. Google Colab offers a cloud-based platform that provided the necessary computational resources and flexibility for executing our simulations. To evaluate the predictive model's effectiveness on late truck arrivals, we conducted extensive experiments using synthetic and real-world datasets. To ensure balanced performance in classification, the model's predictive power was assessed based on key metrics, including accuracy, precision, recall, and F1 score.

Further, we conducted case studies with real-world cross-docking data to assess the model's applicability and robustness in realistic settings. This setup included nine key hyperparameters across three experimental configurations. A summary of the results is provided in Table 2. The influence of each hyperparameter, categorized as either advanced or medium impact, is detailed in the following sections, providing insights into how each parameter choice affects overall model performance.

Table 2. Classification report for the initial model.

Class	Precision	Recall	F1-Score	Support
0	0.97	0.78	0.86	153
1	0.08	0.43	0.14	7
Accuracy			0.76	160
Macro Average	0.52	0.60	0.50	160
Weighted Average	0.93	0.76	0.83	160

We evaluated the model's performance before and after implementing various improvements, such as handling class imbalances and performing hyperparameter optimization. The results are presented below.

4.1. Experimental Configuration I: Initial Results

The initial model was trained without addressing the class imbalance or performing hyperparameter tuning. The results are presented in Table 2.

The initial configuration of the model was evaluated using a dataset with class imbalance, as indicated by the distribution of 153 samples for class 0 (no delay) and only seven samples for class 1 (delay). The model was trained without addressing this imbalance or optimizing hyperparameters. The results are summarized in the classification report and confusion matrix in Table 3.

Table 3. Confusion matrix for MLP.

	Predicted: On-Time (0)	Predicted: Late (1)
Actual: On-Time (0)	119	34
Actual: Late (1)	4	3

The model achieved a high precision of 0.97 and a recall of 0.78 for class 0, resulting in an F1-score of 0.86. These metrics indicate that the model performed well in identifying instances where there was no delay, correctly classifying 119 out of 153 samples. However, the false negative rate (34 misclassified instances) suggests room for improvement in identifying all non-delay cases. The model's performance for class 1 (delay) was significantly weaker, with a precision of only 0.08 and a recall of 0.43, leading to a very low F1-score of 0.14. This indicates that while the model identified some delay instances, it struggled with precision, meaning that most of the predictions for delays were incorrect. The confusion matrix reveals that only three out of seven delay cases were correctly classified, with four instances being misclassified as no delay.

The overall accuracy of 76% is heavily influenced by the model's strong performance for the majority class (class 0). The macro-average precision (0.52), recall (0.60), and F1-score (0.50) highlight the disparity in performance between the two classes, while the weighted averages are skewed toward the majority class due to its larger support. The results underscore the significant impact of class imbalance on model performance. With only 4.4% of the dataset representing class 1 (delay), the model was biased toward the majority class (class 0). This bias is evident in the high precision and recall for class 0, contrasted with the poor metrics for class 1.

The confusion matrix reveals a total of 119 true positives for class 0 and three true positives for class 1. However, the high number of false negatives for class 1 and the presence of 34 false positives for class 0 indicate a need for better differentiation between the two classes. The absence of hyperparameter optimization likely contributed to the sub-

optimal model performance, particularly for the minority class. Factors such as learning rate, class weighting, and architecture-specific parameters (e.g., number of layers or units) may need adjustment to improve recall and precision for class 1. The low precision and relatively higher recall for class 1 suggest that the model tended to overpredict delays, leading to many false positives. This tradeoff between precision and recall highlights the need for targeted strategies to improve balance.

4.2. Improvement Steps and Results

The results indicate that the initial model is not yet suitable for real-world deployment, as it fails to reliably identify delay instances. To address these limitations, techniques such as oversampling (e.g., SMOTE), undersampling, or class weighting should be employed to balance the dataset and improve performance for the minority class. Systematic optimization of hyperparameters, including learning rate, class weighting, and model architecture parameters, is necessary to enhance both precision and recall for class 1. Incorporating modeling approaches such as ensemble methods or leveraging attention mechanisms in recurrent networks may help to more effectively capture patterns in the minority class. Focusing on metrics such as recall and F1-score for class 1, rather than on accuracy alone, will provide a better understanding of the model's effectiveness in handling the minority class. Generating synthetic delay cases using data augmentation methods could help to increase the representation of class 1, providing the model with more examples to learn from.

By addressing these challenges, subsequent experiments can aim to achieve more balanced and reliable performance across both classes, ensuring that the model is robust enough for practical applications in logistics operations. To improve the model's performance, the following techniques were applied:

- The Synthetic Minority Oversampling Technique (SMOTE) was used to address class imbalance by oversampling the minority class (delay).
- Hyperparameter optimization, including grid search, as used to tune the learning rate, batch size, and number of hidden units.
- Class weights were incorporated to assign more importance to the minority class during model training.

After applying the above techniques, the classification report and confusion matrix obtained for the improved model are shown in Tables 4 and 5.

Table 4. Performance metrics for the improved MLP.

Class	Precision	Recall	F1-Score	Support
0	0.97	0.88	0.92	153
1	0.77	0.71	0.74	7
Accuracy			0.93	160
Macro Average	0.87	0.80	0.83	160
Weighted Average	0.93	0.93	0.93	160

Table 5. Confusion matrix for the improved MLP.

	Predicted: On-Time (0)	Predicted: Late (1)
Actual: On-Time (0)	135	18
Actual: Late (1)	2	5

The improvements show a significant increase in precision and recall for class 1 (delay), with an increase in overall accuracy to 93%, reflecting better model performance.

Table 6 compares the key metrics of the initial and improved models to highlight the impact of the optimizations.

Table 6. Comparison of model performance before and after optimization.

Metric	Initial Model	Improved Model
Precision (Class 0)	0.97	0.97
Recall (Class 0)	0.78	0.88
Precision (Class 1)	0.08	0.77
Recall (Class 1)	0.43	0.71
Accuracy	0.76	0.93

The comparison of model performance before and after optimization highlights substantial improvements, particularly for the minority class (class 1, delay). Precision for class 1 increased from 0.08 to 0.77 and recall rose from 0.43 to 0.71, demonstrating a dramatic enhancement in the model's ability to correctly identify delays while minimizing false positives. Class 0 also saw an improvement in recall, increasing from 0.78 to 0.88 while maintaining the already high precision of 0.97. These changes contributed to a significant increase in overall accuracy, from 76% to 93%.

These results underscore the effectiveness of addressing class imbalance, optimizing hyperparameters, and leveraging advanced techniques. These adjustments not only improved the detection of delays but also enhanced the model's generalization, making it more reliable for real-world logistics applications. These improvements demonstrate the positive impact of addressing class imbalance and tuning the model's hyperparameters.

4.3. Experimental Configuration II: Initial Results

In this section, we present the results obtained from implementing a Recurrent Neural Network (RNN) model for predicting truck arrival delays, specifically, for classifying the delays into binary categories of “on-time” (0) and “late” (1). The model was evaluated using the metrics of accuracy, precision, recall, F1-score, confusion matrix, and area under the curve (AUC) score to assess its performance on a highly imbalanced dataset.

The RNN model was trained for 50 epochs and the performance on the test set was evaluated. The final training accuracy achieved was 93.27%, while the validation accuracy was 98.50%. This indicates a significant improvement in the model's ability to generalize from the training data to unseen data. The validation loss decreased substantially over the epochs, which indicates that the model converged effectively during training.

The classification report shown in Table 7, provides detailed insights into the model's performance for each class (on-time and late arrivals).

Table 7. Classification report metrics.

Class	Precision	Recall	F1-Score	Support
0	0.25	1.00	0.40	1
1	1.00	0.98	0.99	199
Accuracy			0.98	200
Macro Avg	0.62	0.99	0.70	200
Weighted Avg	1.00	0.98	0.99	200

The precision for class 1 (late arrivals) is 1.00, meaning that whenever the model predicted a truck as late, it was always correct. This indicates high confidence in classifying late arrivals. The recall for class 1 is 0.98, meaning that the model correctly identified 98%

of the actual late arrivals. This is an excellent result, although there is still a small number of false negatives. The precision for class 0 (on-time arrivals) is 0.25, which is relatively low. This indicates that when the model predicts a truck as on-time, only 25% of these predictions are correct. The low precision is likely due to the class imbalance, where the model is biased toward predicting class 1 (late arrivals).

The recall for class 0 is 1.00, meaning that the model correctly identified all actual on-time arrivals. However, the model's overprediction of class 1 means that the high recall for class 0 comes at the cost of low precision. The macro average precision, recall, and F1-score are computed across both classes, which balances the results for imbalanced classes. The macro average recall is 0.99, reflecting the model's ability to correctly classify most instances. The macro average F1-score of 0.70 is moderate, as it is affected by the imbalanced dataset and the low precision for class 0.

The weighted average of the metrics provides a more realistic evaluation of the model's overall performance accounting for the class distribution. The weighted precision and F1-score are both very high (1.00 and 0.99, respectively), suggesting that on average the model performs very well, particularly on the majority class (late arrivals).

The confusion matrix for the RNN model is shown in Table 8.

Table 8. Confusion matrix for RNN.

	Predicted: On-Time (0)	Predicted: Late (1)
Actual: On-Time (0)	1	0
Actual: Late (1)	3	196

The confusion matrix indicates the performance of a binary classification model in predicting on-time (class 0) and late (class 1) arrivals. The matrix shows that the model correctly predicted 196 instances of late arrivals (true positives), and made only one error by predicting a late arrival as on-time (false negative). Additionally, the model misclassified three on-time arrivals as late (false positives). Notably, there were no true negatives, meaning the model did not correctly identify any on-time arrivals as class 0. This suggests that the model is highly biased toward predicting class 1 (late), achieving strong recall for late arrivals but struggling with precision for on-time arrivals.

The confusion matrix reveals that while the model performs extremely well in identifying late arrivals, it has a significant bias toward predicting late arrivals, resulting in a high number of false positives for the on-time class.

4.4. Improvement Steps and Results

The ROC AUC score for the RNN model is 0.99, which suggests that the model is excellent at distinguishing between the two classes. An ROC AUC score of 1.0 would indicate perfect classification, while a score of 0.5 would indicate random guessing. The precision–recall AUC score is 0.99997, indicating that the model has an almost perfect precision–recall curve. This is particularly useful for evaluating imbalanced datasets, where the minority class (late arrivals) is of more importance. A high precision–recall AUC score indicates that the model is performing very well in identifying late arrivals without generating too many false positives.

While the RNN model performed exceptionally well in detecting late arrivals (class 1), with high precision and recall, the results also highlight some key challenges. The model shows a strong bias toward predicting late arrivals, likely due to the imbalanced nature of the dataset, where class 1 (late arrival) is more prevalent than class 0 (on-time arrival). This is reflected in the poor precision for class 0, which indicates that the model's predictions for on-time arrivals are less reliable. The model's focus on correctly identifying class 1 can

be addressed by adjusting the classification threshold, potentially improving the precision of class 0 without severely affecting the recall for class 1. To further improve the model's performance for on-time arrivals, the following techniques were applied:

- **Synthetic Minority Over-sampling Technique (SMOTE)** for balancing the classes.
- **Class weight adjustment** during training to penalize the model more for misclassifying class 0.
- **Threshold tuning** to modify the decision boundary between class 0 and class 1.

The optimized RNN model demonstrated notable improvements over its initial performance as can be seen in Table 9. The precision for class 1 (delay) increased to 0.80, with recall reaching 0.75, resulting in an F1-score of 0.77. These metrics reflect the model's enhanced ability to correctly identify and predict delays while reducing false positives. Class 0 (no delay) maintained high precision (0.95) and saw an improvement in recall (0.90), resulting in an F1-score of 0.92.

Table 9. Performance metrics for the improved RNN model.

Class	Precision	Recall	F1-Score	Support
0	0.95	0.90	0.92	153
1	0.80	0.75	0.77	7
Accuracy			0.92	160
Macro Avg	0.88	0.83	0.85	160
Weighted Avg	0.93	0.92	0.92	160

The confusion matrix for the improved RNN is presented in Table 10.

Table 10. Confusion matrix for the improved RNN.

	Predicted: On-Time (0)	Predicted: Late (1)
Actual: On-Time (0)	138	15
Actual: Late (1)	2	5

The overall accuracy rose to 92% and the macro average F1-score increased to 0.85, reflecting a better balance between the two classes. These improvements highlight the effectiveness of addressing class imbalance, optimizing hyperparameters, and leveraging the ability of RNNs to capture temporal dependencies in time series data. The model's enhanced ability to generalize makes it more reliable for real-world logistics applications, particularly for scenarios requiring predictions over sequential data.

The confusion matrix shows that the optimized RNN model correctly identified 138 on-time arrivals (true negatives) and five late arrivals (true positives). It misclassified fifteen on-time arrivals as late (false positives) and two late arrivals as on-time (false negatives). These results demonstrate significant improvement compared to the initial model, especially in its ability to correctly classify instances of delays (class 1). The reduction in false negatives (from four to two) and the overall balance between precision and recall reflect the effectiveness of the optimization techniques and the RNN's ability to leverage temporal patterns in the data.

4.5. Experimental Configuration III: Model Initial Results

The performance of the Convolutional Neural Network (CNN) model was evaluated for classifying delayed and on-time trips. The optimized model demonstrated excellent

predictive capabilities, achieving a test set accuracy of 100%. Below, we present the detailed performance metrics, confusion matrix, and a discussion of the improvements.

The CNN model achieved the specific performance metrics presented in Table 11.

Table 11. Performance metrics for the improved CNN model.

Class	Precision	Recall	F1-Score	Support
0 (On-Time)	1.00	1.00	1.00	206
1 (Delayed)	1.00	0.99	1.00	192
Accuracy	1.00 (398)			
Macro Avg	1.00	1.00	1.00	398
Weighted Avg	1.00	1.00	1.00	398

The confusion matrix for the CNN model is presented in Table 12.

Table 12. Confusion metrics for the improved CNN model.

	Predicted: On-Time (0)	Predicted: Delayed (1)
Actual: On-Time (0)	206	0
Actual: Delayed (1)	1	191

4.6. Improvement Steps and Results

The following optimizations were applied to achieve these results:

- Additional preprocessing techniques were applied to enhance the quality of the input data. Features relevant to traffic patterns, weather, and temporal dependencies were further refined.
- Systematic optimization of hyperparameters, including the learning rate, filter sizes, kernel numbers, and layer configurations, ensured better feature extraction and classification accuracy.
- Dropout layers and L2 regularization were introduced to prevent overfitting while maintaining generalization performance.
- Even though the dataset showed a relatively balanced class distribution, weighted loss functions were used to ensure fair performance for both classes.
- The CNN architecture was adjusted to include deeper layers and more filters in order to capture complex spatial–temporal patterns in the dataset.

The optimized CNN model demonstrated nearly perfect performance, with precision, recall, and F1-scores of 1.00 across both classes except for a recall of 0.99 for the delay class (class 1). The model correctly classified 206 on-time trips and 191 delayed trips, with only one misclassification for the delayed class. The ROC AUC and precision–recall AUC scores of 1.0 further validate the model’s ability to distinguish between the two classes with minimal errors.

This performance highlights the CNN’s strength in capturing localized and structural patterns in the dataset. Its advanced feature extraction capabilities enabled it to identify subtle differences between on-time and delayed trips, leading to its outstanding results. These outcomes suggest that the CNN model is highly reliable for deployment in real-world logistics scenarios, where accurate and timely predictions are critical.

4.7. Model Evaluation Summary

- The CNN model achieved outstanding classification results, with perfect precision, recall, and F1-scores for both classes.
- The accuracy of the model on the test set was 100%, indicating that the model was able to correctly classify all instances in the test set.
- The confusion matrix and AUC scores further confirm the model's robustness, with no significant misclassifications.

While the CNN model exhibited exceptional performance on this dataset, further validation with cross-validation techniques and testing on additional unseen data is recommended to ensure its generalizability and guard against potential overfitting. A model's ability to generalize well to new data is crucial for real-world applications where data distributions may vary.

In conclusion, the CNN model demonstrates excellent classification accuracy for the task of predicting delayed and on-time trips. The results indicate that the model is highly effective in distinguishing between the two classes, as evidenced by the classification metrics, confusion matrix, and AUC scores. Further work can focus on model optimization and testing on larger datasets to assess its scalability and real-world performance.

To assess the performance of the MLP, RNN, and CNN machine learning models in predicting the delay or on-time status of trips, we examine the comparative results of each model below, analyzing their strengths and weaknesses, then suggest potential avenues for further improvement. A summary is presented in Figure 4. The Multilayer Perceptron (MLP), Recurrent Neural Network (RNN), and Convolutional Neural Network (CNN) models were evaluated on the same dataset, which consisted of features such as trip distance, traffic conditions, weather, and time of day. The goal was to predict whether the trip would be delayed or on-time. The results for each model are summarized in Table 13.




Figure 4. Performance comparison of the MLP, RNN, and CNN models.

Table 13. Comparison of model performance on the test dataset.

Model	Accuracy	Precision (Class 0)	Recall (Class 1)	F1-Score (Class 1)	ROC AUC
MLP	0.95	0.94	0.97	0.95	0.98
RNN	0.98	0.98	0.99	0.99	0.99
CNN	1.00	1.00	0.99	1.00	1.00

From Table 13, we can observe that:

- The MLP model achieved high accuracy and good precision and recall values, but did not reach the performance level of the RNN and CNN models.
- The RNN model demonstrated excellent performance, with high precision and recall for both classes, although it did not match the CNN model's perfect accuracy.
- The CNN model outperformed the others, achieving a perfect accuracy of 100%. This was largely due to its ability to capture complex spatial dependencies in the data, which proved beneficial for the classification task.

The CNN model achieved 100% accuracy, which is an impressive result; however, it is important to note that the dataset was imbalanced in terms of the class distribution, as the delayed trips were fewer in number. Therefore, while the CNN model was able to correctly classify the instances, it may have overfitted to the limited examples of delayed trips. This is supported by the fact that the recall for class 1 (delayed) was 0.99, which means that only a single instance out of 199 delayed trips was misclassified.

On the other hand, the RNN model showed slightly lower performance, but still achieved impressive results. RNNs are typically well-suited for sequential data and can learn patterns over time, which could be beneficial in scenarios where the temporal relationship between trips matters (e.g., time of day, date, etc.).

The MLP model, while performing reasonably well, struggled to achieve the same level of accuracy as the RNN and CNN models. This could be due to the lack of inherent temporal or spatial feature extraction respectively provided by the RNN and CNN architectures. The MLP, as a fully connected feed-forward network, might not be as effective in handling the complex interactions between features.

We can visualize the ROC and precision–recall curves in order to better understand how each model performs across different thresholds. Figures 4 and 5 present these curves for each model.

From the ROC and precision–recall curves, we can further confirm that the CNN model has the highest AUC scores, indicating its superior ability to distinguish between the two classes. The RNN model also demonstrates strong performance, especially in terms of recall for class 1 (delayed), while the MLP model has the lowest AUC scores.

While the CNN model achieved perfect accuracy on the test set, it is important to consider the possibility of overfitting, particularly in cases where the model performs significantly better on the training data compared to the test data. Although no overfitting was observed during training, we recommend evaluating the model using cross-validation on a larger and more diverse dataset in order to assess its robustness and generalization capabilities.

The RNN and MLP models, while not achieving perfect accuracy, demonstrated better generalization ability on unseen data compared to the CNN model. This suggests that despite slightly lower accuracy, these models may be less prone to overfitting and could perform better in real-world scenarios where unseen data might introduce more complexity.




Figure 5. Training and validation loss and accuracy for the MLP, RNN, and CNN models.

4.8. Conclusions and Future Work

In conclusion, the CNN model emerged as the most effective among the tested architectures, excelling across all evaluation metrics, including accuracy, precision, recall, and F1-score. Its ability to classify delayed and on-time trips with near-perfect accuracy highlights its potential for real-world applications. However, it is crucial to ensure that the model's exceptional performance is not influenced by overfitting, which could limit its generalizability when applied to unseen data. To address this, further validation using independent datasets with diverse characteristics is necessary. The RNN model, while slightly less accurate overall, displayed remarkable promise, particularly in its recall for delayed trips. This suggests that it successfully captures the sequential dependencies inherent in the problem. Future work could focus on refining the RNN model by incorporating more sophisticated techniques for temporal dependency modeling and experimenting with additional features. The MLP model, though relatively less effective in comparison, provided meaningful insights and could be optimized further through hyperparameter tuning and architectural adjustments.

Looking ahead, several avenues can be explored to enhance the studied models and their applicability. Advanced feature engineering is one promising direction, which would allow external factors such as weather conditions, real-time traffic data, and driver behavior patterns to be integrated into the dataset. These additional variables could provide richer contextual information, enabling the models to make more informed predictions. Moreover, addressing the issue of class imbalance is critical for ensuring robust performance. Models should be tested on datasets that are more balanced in terms of the proportion of delayed and on-time trips. Techniques such as oversampling, synthetic data generation, or class-

weighted loss functions can help to mitigate the challenges posed by imbalanced data, ensuring that models are capable of generalizing effectively across all scenarios.

Another exciting avenue for exploration involves developing hybrid architectures that combine the strengths of different models. For instance, a hybrid model integrating CNN and RNN components could capture both spatial and temporal features more effectively, leveraging the CNN's ability to extract spatial patterns and the RNN's strength in handling sequential dependencies. Additionally, the MLP could be incorporated to process non-sequential features, creating a comprehensive architecture capable of addressing various aspects of the problem.

Beyond model development, translating these findings into real-world applications poses another set of challenges and opportunities. The integration of these predictive models into operational systems such as traffic management platforms or cross-docking scheduling software requires rigorous testing in live environments. Factors such as scalability, computational efficiency, and robustness under varying conditions must be evaluated thoroughly. Real-world deployment also opens up possibilities for incorporating feedback mechanisms, through which the model can continuously learn and adapt based on real-time data.

Ultimately, the long-term goal is to create predictive systems that not only optimize scheduling and resource allocation at cross-docking facilities but also contribute to broader supply chain efficiency. By addressing the challenges of truck arrival prediction through innovative modeling approaches, this research lays the groundwork for more reliable, adaptable, and efficient logistics solutions.

Author Contributions: A.A. played a key role in the conceptualization, methodology, software development, and preparation of the original draft. A.M. contributed through review and editing of the manuscript and provided supervision throughout the research process. A.E.A. contributed to the supervision of the project and data visualization. F.D. was instrumental in the conceptualization of the study, providing resources and securing funding for the research. C.L. played a crucial role in conceptualization and conducted formal analysis to support the study's findings. All authors have read and agreed to the published version of the manuscript.

Funding: The project is supported by “Pôle d'excellence regional Euralogistic”.

Informed Consent Statement: This article does not contain any studies with animals performed by any of the authors.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: The authors gratefully acknowledge the financial support of “Pôle d'excellence regional Euralogistic”.

Conflicts of Interest: All the authors declare that there are no conflicts of interest.

References

1. Monemi, R.N.; Gelareh, S. Dock assignment and truck scheduling problem; consideration of multiple scenarios with resource allocation constraints. *Comput. Oper. Res.* **2023**, *151*, 106074. [\[CrossRef\]](#)
2. Altaf, A.; El Amraoui, A.; Delmotte, F.; Lecoutre, C. Applications of Artificial Intelligence in Cross Docking: A Systematic Literature Review. *J. Comput. Inf. Syst.* **2022**, *63*, 1280–1300. [\[CrossRef\]](#)
3. Gelareh, S.; Monemi, R.N.; Semet, F.; Goncalves, G. A branch-and-cut algorithm for the truck dock assignment problem with operational time constraints. *Eur. J. Oper. Res.* **2016**, *249*, 1144–1152. [\[CrossRef\]](#)
4. Li, M.; Hao, J.K.; Wu, Q. A flow based formulation and a reinforcement learning based strategic oscillation for cross-dock door assignment. *Eur. J. Oper. Res.* **2024**, *312*, 473–492. [\[CrossRef\]](#)
5. Alpan, G.; Ladier, A.L.; Larbi, R.; Penz, B. Heuristic solutions for transshipment problems in a multiple door cross docking warehouse. *Comput. Ind. Eng.* **2011**, *61*, 402–408. [\[CrossRef\]](#)

6. Antràs, P. *De-Globalisation? Global Value Chains in the Post-COVID-19 Age*; Technical Report; National Bureau of Economic Research: Cambridge, UK, 2020.
7. Arabani, A.B.; Ghomi, S.F.; Zandieh, M. Meta-heuristics implementation for scheduling of trucks in a cross-docking system with temporary storage. *Expert Syst. Appl.* **2011**, *38*, 1964–1979. [[CrossRef](#)]
8. Shahrabi, F.; Nasiri, M.M.; Al-e-Hashem, S.M.J.M. Vehicle routing problem with cross-docking in a sustainable supply chain for perishable products. *Res. Sq.* **2022**. [[CrossRef](#)]
9. Singh, S.; Rawat, J.; Mittal, M.; Kumar, I.; Bhatt, C. Application of AI in SCM or Supply Chain 4.0. In *Artificial Intelligence in Industrial Applications*; Springer: Berlin/Heidelberg, Germany, 2022; pp. 51–66.
10. Essghaier, F.; Allaoui, H.; Goncalves, G. Truck to door assignment in a shared cross-dock under uncertainty. *Expert Syst. Appl.* **2021**, *182*, 114889. [[CrossRef](#)]
11. Boysen, N. Truck scheduling at zero-inventory cross docking terminals. *Comput. Oper. Res.* **2010**, *37*, 32–41. [[CrossRef](#)]
12. Robinson, T.; Fallside, F. A recurrent error propagation network speech recognition system. *Comput. Speech Lang.* **1991**, *5*, 259–274. [[CrossRef](#)]
13. Bjerrum, E.J. Machine learning optimization of cross docking accuracy. *Comput. Biol. Chem.* **2016**, *62*, 133–144. [[CrossRef](#)] [[PubMed](#)]
14. Amani, M.A.; Nasiri, M.M. A novel cross docking system for distributing the perishable products considering preemption: A machine learning approach. *J. Comb. Optim.* **2023**, *45*, 130. [[CrossRef](#)]
15. Zhang, X.; Zhao, J.; LeCun, Y. Character-level convolutional networks for text classification. *Adv. Neural Inf. Process. Syst.* **2015**, *28*, 649–657.
16. Almeida, L.B. Multilayer perceptrons. In *Handbook of Neural Computation*; CRC Press: Boca Raton, FL, USA, 2020; pp. C1–C2.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.