Short-term vessel trajectory and manoeuvre prediction

Predicting vessel trajectories including swing manoeuvres in a scoped area of the Port of Antwerp-Bruges

Joaquin van Loon

Delft University of Technology

TU Delft

Short-term vessel trajectory and manoeuvre prediction

Predicting vessel trajectories including swing manoeuvres in a scoped area of the Port of Antwerp-Bruges

by

Joaquin van Loon

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday June 14, 2023 at 10:00 AM.

Student number:	4896793	
Project duration:	November 14, 2022 – June 14, 2023	
Thesis committee:	Dr. N. Yorke-Smith	TU Delft, supervisor
	Dr. R.T. Rajan	TU Delft
	Timo Koppenberg	Macomi

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TU Delft

Abstract

Situational awareness within port areas is crucial to avoid collisions, navigate efficiently and reduce congestion. Maritime-traffic controllers constantly monitor the situation in the port and intervene when needed. This study proposes a deep learning model that predicts future vessel positions to assist in this process. The model employs a target conditioned trajectory prediction component composed of two recurrent neural networks arranged in an encoder-decoder structure that utilizes historical data points to forecast future trajectories. The model considers multiple factors, including vessel speed, location, length, depth, draught, and the tide. Additionally, this study addresses the prediction of swing manoeuvres, which are special U-turn-like manoeuvres executed during terminal arrival or departure. These manoeuvres can block a significant portion of the waterway and, as such, are essential to consider when gaining a complete understanding of future situations within the port. An integration of both models is applied to a use case study in a scoped area of the Port of Antwerp-Bruges. The models were trained using AIS and VTS data collected at 30-second intervals. Swing manoeuvres are predicted with an accuracy of 84%, the locations of these manoeuvres are predicted with an average deviation of 212 meter and the duration error is 1.6 minutes on average. The complete predicted trajectories, including potential swing manoeuvres, have an average displacement error and final displacement error of 147 and 117 meter on average, respectively. Overall, the study demonstrates the potential of deep learning models for improving situational awareness within port areas and assisting traffic controllers in making informed decisions.

Preface

Presented before you is my master's thesis, the result of a six-month research study I conducted to conclude my academic journey at the Technical University of Delft. This thesis marks the final step towards obtaining my master's degree in Computer Science. After five amazing years at the TU Delft, I am excited to now put the knowledge and skills I have acquired into practice in my future endeavors.

My interest in algorithms and artificial intelligence led me to approach Dr. N. Yorke-Smith in early 2022. I was fortunate that Neil agreed to supervise my thesis and arrange a research internship for me at Macomi. Working alongside my colleagues at Macomi has been inspiring and I could not have wished for a better environment to conduct my research. I am grateful for the warm welcome and excellent guidance I received at Macomi throughout my research. I want to thank all my colleagues at Macomi for this opportunity. Special appreciation must be given to Timo Koppenberg. Timo was my daily supervisor and point of contact at Macomi. He had a huge impact at every step of my thesis and always helped me out when needed. I am grateful for his invaluable support and devotion to my thesis.

Additionally, I had periodic meetings with Dr. N. Yorke-Smith, who played a significant role in my thesis by providing insightful comments and advice on my progress and next steps. I would like to express my gratitude to Neil for his guidance. I also extend my appreciation to Tessy Vanhoenacker for providing me with real-world data and the opportunity to conduct my thesis in collaboration with the Port of Antwerp-Bruges.

Lastly, I would like to thank my parents for their unconditional love and support throughout my entire life and academic career at the Technical University of Delft. I am forever grateful to both of you.

Joaquin van Loon
Rotterdam, June 2023

Contents

Abstract i

Preface ii

1 Introduction 1
- 1.1 Port of Antwerp-Bruges . . . 1
  - 1.1.1 Scoped area . . . 1
- 1.2 Research objective . . . 3
- 1.3 Thesis contributions . . . 3
- 1.4 Thesis outline . . . 3

2 Data 5
- 2.1 Features . . . 5
- 2.2 Preprocessing . . . 6
  - 2.2.1 Positional data and meta data . . . 6
  - 2.2.2 Tide data . . . 6
  - 2.2.3 Trajectory extraction and interpolation . . . 6
  - 2.2.4 Trajectory destinations . . . 9

3 Trajectory Prediction 12
- 3.1 Problem formulation . . . 12
- 3.2 Background . . . 12
  - 3.2.1 Modelling and data based approaches . . . 13
  - 3.2.2 Deep learning-based approaches . . . 13
- 3.3 Neural networks for sequential data . . . 14
  - 3.3.1 Recurrent Neural Network . . . 14
  - 3.3.2 Long Short-Term Memory . . . 14
  - 3.3.3 Gated Recurrent Unit . . . 15
- 3.4 Target conditioned trajectory prediction . . . 16
  - 3.4.1 Encoder-decoder structure . . . 16
  - 3.4.2 Complete model . . . 16
  - 3.4.3 Loss function . . . 18

4 Swing manoeuvres 20
- 4.1 Swing definition . . . 20
- 4.2 Background . . . 20
- 4.3 Swing detection . . . 21
- 4.4 Swing extraction . . . 23
- 4.5 Historical swings . . . 26
- 4.6 Swing prediction . . . 29
  - 4.6.1 Swing occurrence prediction . . . 29
  - 4.6.2 Swing location prediction . . . 31

5 Trajectory and manoeuvre prediction model 32
- 5.1 Vessel route options . . . 32
- 5.2 Model overview . . . 33
  - 5.2.1 Docking and undocking trajectories . . . 34
  - 5.2.2 Finishing swing manoeuvres . . . 36

6 Results 38
- 6.1 Model training . . . 38
  - 6.1.1 Trajectory prediction training . . . 38

6.1.2 Swing prediction training . . . 38
6.2 Performance metrics . . . 39
6.3 Target conditioned trajectory prediction results . . . 41
6.4 Swing prediction results . . . 46
6.5 Full model results . . . 46

7 Conclusion **53**
7.1 Summary . . . 53
7.2 Conclusions . . . 54
7.3 Discussion and future work . . . 54
7.3.1 Model generalization . . . 55

References **56**

A Paper **59**

1 Introduction

The Port of Antwerp-Bruges is a busy hub of maritime activity with a complex nautical situation that demands continuous monitoring. Given the large volume of vessels that navigate the port's waterways, ensuring safety is of crucial importance. In addition to preventing potential hazards and reducing congestion, maintaining situational awareness can also help to minimize carbon emissions by optimizing vessel traffic flows and reducing idle time. Monitoring the live positions of vessels contributes to the situational awareness of maritime traffic controllers. Additionally, the implementation of a live prediction model, as detailed in this study, can further improve the understanding of future situations in the port.

1.1. Port of Antwerp-Bruges

This thesis conducts a case study at the Port of Antwerp-Bruges. The port is the second largest seaport in Europe and handled around 250 million tonnes of maritime freight volume in the year 2021 [5]. The port has five deep-sea container terminals, which are used to load and unload large container ships arriving and departing from the port.

Since these container vessels arrive and leave the port through one waterway (see Figure 1.1), they need to make a U-turn somewhere within the port. Such a U-turn manoeuvre is done right before a vessel arrives at the terminal location or just after it leaves the terminal place. The potential locations for a large container vessel to execute a swing manoeuvre are dependent upon various factors, including but not limited to the vessel's draught and the tide conditions. Chapter 4 will explain and investigate these swing manoeuvres in more detail, but the main takeaway is that these swing manoeuvres are required, block a large part of the waterway, and can take some time.

1.1.1. Scoped area

This study narrows its focus to a small, scoped area of the port, see red rectangle in Figure 1.1. This scoped area is of particular interest, because it connects the southern parts of the port to the North Sea and it contains two terminals and two locks in a relatively small area. A bird eye view of this area can be seen in Figure 1.2. The Europa Terminal (EUT) and Noordzee Terminal (NZT) are used to dock container vessels and load and unload them. Furthermore, the Zandvliet Lock and Berendrecht Lock are used to enter or leave the inland waters of the port. As can be seen in Figure 1.1, the waterway which goes through the scoped area connects the port with the sea. This means that all vessels leaving the port to the sea or arriving from the sea will pass this scoped area. Due to this high traffic intensity and the presence of two terminals and locks, this area is considered of critical importance.

Swing manoeuvres especially have impact in the scoped area of the port. In this area there are namely two terminals which are connected to a busy waterway. Swinging vessels blocking the waterway in front of the Noordzee Terminal restrict traffic to more southern parts of the port and to the locks. Moreover, swinging vessels in front of the Europa Terminal block traffic to the southern parts of the port. Since it is crucial to avoid collisions in the port and minimize congestion and delays, these swing manoeuvres play an important role in the management of traffic.

Due to its high traffic volume and bottlenecking characteristics, the scoped area serves as a crucial point in the Port of Antwerp-Bruges. As such, it is critical for the Port of Antwerp-Bruges to maintain an attentive watch over both current and future traffic situations in the area.

Figure 1.1: Location of the Port of Antwerp-Bruges and the Scheldt river. The red rectangle contains the area of the port which this study focuses on. Map data @2023, Google.

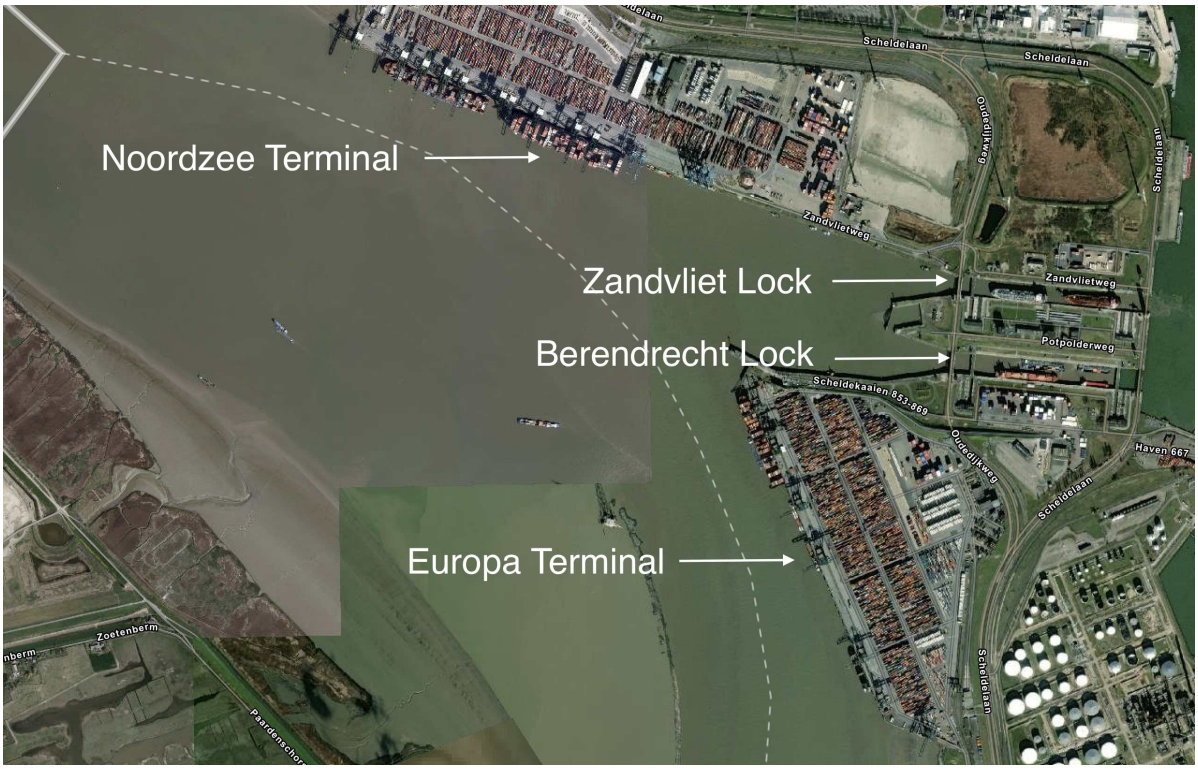

Figure 1.2: A bird eye view of the scoped area of the Port of Antwerp-Bruges

1.2. Research objective

According to a study that analyzed marine collision reports, 71% of human error is caused due to lack of situational awareness. Specifically, of the errors related to situational awareness, around 9% were attributed to a lack of awareness regarding future actions in the marine environment [20]. The objective of this study is to assist maritime traffic controllers in obtaining situational awareness. Traffic controllers must be aware of the current situation and look ahead in time to foresee congestion. To aid them in this process, this study aims to create a model that can predict future vessel trajectories. Since swing manoeuvres are crucial parts of these trajectories, the model must also take them into account.

Given this objective, the main research question is:

> *How can a vessel trajectory prediction model that accounts for swinging manoeuvres be developed and applied to a case study at the Port of Antwerp-Bruges?*

This question is further broken down in the following sub questions.

RQ1. How is the problem of trajectory prediction formalized?

RQ2. What methods can be used to tackle the trajectory prediction problem?

RQ3. How are swing manoeuvres defined and extracted from historical trajectories?

RQ4. How effective are the found trajectory prediction methods in predicting simple trajectories without swinging manoeuvres?

RQ5. Can these same trajectory prediction methods predict trajectories with swinging manoeuvres or is a special model needed for this purpose?

RQ6. How is a single model constructed which can accurately predict complete vessel trajectories, including swinging manoeuvres?

1.3. Thesis contributions

Trajectory prediction is a well-known area of study and quite some research has been done in this domain. These studies do not only concern vessel trajectory prediction, but similar methods are used in other domains, such as car, aircraft or pedestrian trajectory prediction. A more detailed overview of the state-of-the-art methods for trajectory prediction can be found in Chapter 3. Although substantial literature is present on trajectory prediction, limited research is done on trajectory prediction including swing manoeuvres. The present literature only analyses swing manoeuvre with the objective to improve port design. However, to the best of the author's knowledge, predicting swing trajectories is not done before.

As mentioned in section 1.1, swing manoeuvres do form a critical part of the trajectories of vessel in the Port of Antwerp-Bruges. Therefore, in this case study they must be incorporated in the trajectory predictions.

Therefore, the contribution of this study is two-fold:

- This study extends state-of-the-art vessel trajectory prediction methods to incorporate swing manoeuvres.
- This study applies trajectory prediction models on a specific use case scenario at the Port of Antwerp-Bruges.

1.4. Thesis outline

The structure of the remainder of this thesis is guided by the research questions. First of all, Chapter 2 introduces the data used in this thesis and gives some useful insights in this data. Chapter 3 formulates the problem of trajectory prediction and provides an overview of existing methods to tackle this problem. At the end of this chapter the trajectory prediction model used for this study is introduced. The next chapter, Chapter 4, elaborates on the topic of swing manoeuvres and introduces the swing prediction

model. Next, Chapter 5 presents the full model which combines the trajectory prediction model from Chapter 3 and the swing prediction model from 4 to predict complete future trajectories including swing manoeuvres. After all models have been discussed, Chapter 6 discusses the evaluation and performance of these models. This document ends with a summary, conclusion, discussion and pointers for future work in Chapter 7.

2 Data

With around 90 percent of global trade conducted via sea transport and continued growth in marine traffic [24], ensuring the safety and security of ships and their cargo is of great importance. To this end, the International Maritime Organization (IMO) has introduced Automatic Identification Systems (AIS) that periodically transmit data for vessels at sea, with the goal of improving navigation efficiency and safety [23]. AIS data provides a real-time view of vessel positions and can enhance situational awareness, making it a valuable tool for maritime traffic controllers. When vessels enter a port, a vessel traffic service (VTS) communicates with the vessels with the goal of improving the safety and efficiency of the vessels within the port. A VTS typically uses radar to keep track of the locations of vessels within the port. The data collected by the VTS combined with the AIS data can be used for reporting and analysis purposes. In this study this data is used to train neural networks.

This chapter presents the two years of historical vessel data from the Port of Antwerp-Bruges which is used in this research. Section 2.1 provides an overview of the data structure, while section 2.2 describes the necessary preprocessing steps.

2.1. Features

The Port of Antwerp-Bruges continuously records the positional data and meta data of the vessels entering the port. For this study, The Port of Antwerp-Bruges provided two years of vessel data. Most data come from the AIS source, but a part originates from the VTS. The data was recorded in the scoped area of the port as described section 1.1. However, complete data covering two years is only available for container vessels. For other vessel types, only data from September and October 2022 is present. The availability of a significant amount of container vessel data is crucial, as trajectories of container vessels interacting with the Noordzee Terminal or Europa Terminal are less frequent than those bypassing the scoped area of the port. The Port of Antwerp-Bruges provided the data in three parts:

1. Vessel positional data
2. Vessel meta data
3. Vessel trajectory data

The positional data records were captured at approximately 30-second intervals and include the relevant features shown in Table 2.1. Additionally, static data for each vessel that passed through the port was provided in the vessel meta data file. The relevant features from the meta data are listed in Table 2.2. The vessel origin and destination contain the exact name of a specific location, such as ”Loodskotter West” or the terminal location ”S853”. The vessel trajectory data file consists of a list of vessel trajectories, each characterized by a start time (in case of departure) or end time (in case of arrival), along with additional meta data fields. Notably, the draught of a vessel during its trajectory, which is the distance between the waterline and the bottom of the vessel, was included as a useful feature that is not present in the meta data file.

MMSI	The unique identification number of the vessel
Timestamp	An epoch timestamp of the time at which this data point was recorded.
Latitude	The latitude position of the vessel
Longitude	The longitude position of the vessel
Speed over ground	The ground speed of the vessel in knots
Heading	The heading of the vessel in degrees, where zeros degrees is the north.

Table 2.1: Useful sequence features from AIS data

MMSI	The unique identification number of the vessel
Vessel type	The type of vessel
Depth	The vessel's depth
Length	The vessel's length
Width	The vessel's width
Destination	The vessel's destination
Origin	The vessel's origin

Table 2.2: Useful meta data features from AIS data

2.2. Preprocessing

Before the data can be used in the models described in Chapters 3, 4, and 5, several preprocessing steps are necessary. This section provides a detailed description of each of these steps. Firstly, subsection 2.2.1 explains how all the positional data is combined. Subsection 2.2.2 outlines how tide data is integrated, followed by subsection 2.2.3, which details the process of extracting and interpolating vessel trajectories. Lastly, subsection 2.2.4 explains how the destinations of these trajectories were determined.

2.2.1. Positional data and meta data

The first step in the data preprocessing is to merge the vessel meta data with the positional data using the vessel's unique Maritime Mobile Service Identity (MMSI) number. The merged vessel data contains all the features presented in Tables 2.1 and 2.2 for each vessel. In total, the dataset contains $6,613,409$ positional data records within the scoped area. Meta data is available for 5060 different vessels. However, not all vessel types are relevant for the scope of this research. The Port of Antwerp-Bruges is primarily interested in seagoing vessels, as manoeuvres performed by these large vessels can significantly affect the flow of traffic within the port. Therefore, only data from five vessel types is retained: Dry, Liquid, RoRo, Container, and General Cargo. These vessel types are not present directly in the data, so a mapping must be applied to correctly categorize the vessels. Table 2.3 displays how the vessel types are mapped to the correct vessel category as used by the Port of Antwerp-Bruges. After filtering for the vessel categories of interest, meta data for 2203 unique vessels are in the dataset.

2.2.2. Tide data

As tide data could be an important factor in the models according to domain knowledge, it is beneficial to add it to the existing AIS data. The original delivered data does not contain tide data. However, the Flemish government provides tide tables on their website [37], which indicate the water levels at different locations in the Belgium waters. The tide data in the port was estimated by using the water levels at "Prosperpolder", which is the closest location to the port where tide data is available.

2.2.3. Trajectory extraction and interpolation

Once the vessel positional data, vessel meta data, and tide data are combined, the vessel trajectories file is processed to group the data points per trajectory. Since the departure time or arrival time is known and the interval between two data points is assumed to be approximately 30 seconds, the process of

Vessel type	Vessel category
container ship (full)	Container
chemical tanker	Liquid
gastanker	Liquid
liquefied gas tanker	Liquid
tanker	Liquid
LNG gastanker	Liquid
LPG gastanker	Liquid
general cargo	GeneralCargo
general cargo/container ship	GeneralCargo
general cargo/tanker	GeneralCargo
heavy load carrier	GeneralCargo
refrigerated cargo	GeneralCargo
special cargo	GeneralCargo
ore/bulk/oil carrier	Dry
bulk carrier	Dry
bulk carrier/chemical tanker	Dry
ore/oil carrier	Dry
roro cargo	RoRo
roro cargo/container ship	RoRo
roro cargo/general cargo	RoRo
roro cargo/vehicles carrier	RoRo
roro containers	RoRo
vehicles carrier	RoRo

Table 2.3: Mapping of AIS data Vessel type to vessel category used by the Port of Antwerp-Bruges

grouping the points by trajectory is possible.

The output of this extraction step is a list of data points per trajectory. However, since the data points are not uniformly spaced in time, an interpolation step is necessary to ensure that the interval between adjacent data points is always 30 seconds. Specifically, linear interpolation is applied to the longitude, latitude, and speed over ground, while nearest neighbor interpolation is used for the heading. All these attributes are interpolated individually. However, to avoid inaccurate interpolations, trajectories with more than 15 missing points are removed from the dataset. The number of missing points is calculated by subtracting the expected number of points, i.e. a point every 30 seconds, from the actual number of points in the data.

To limit the scope of this research, trajectories which only consist of internal port movements are excluded from the data. These movements are quite rare and are mostly executed by smaller vessels within the port. Additionally, outlier trajectories are removed from the data to smooth the training of the model. Outlier trajectories are identified manually and contain invalid data points (e.g., outside the sailable area or impossible positional or heading changes). These outliers are a result of imperfections in the vessel positional data.

After filtering a total of 16238 are extracted. All these remaining trajectories are shown in Figure 2.1. Figure 2.2 provides an overview of the number of trajectories per vessel category. As can be seen, the majority of the trajectories are from container vessels.

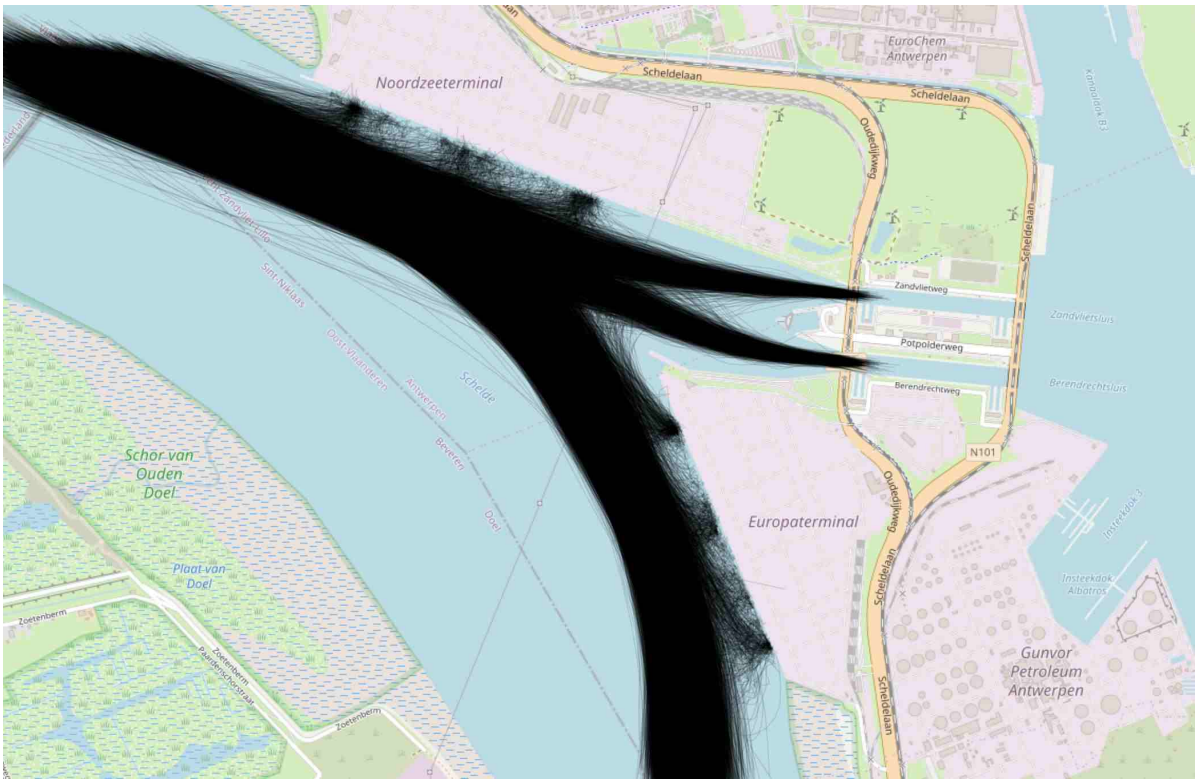

Figure 2.1: All trajectories in the scoped area of the Port of Antwerp-Bruges

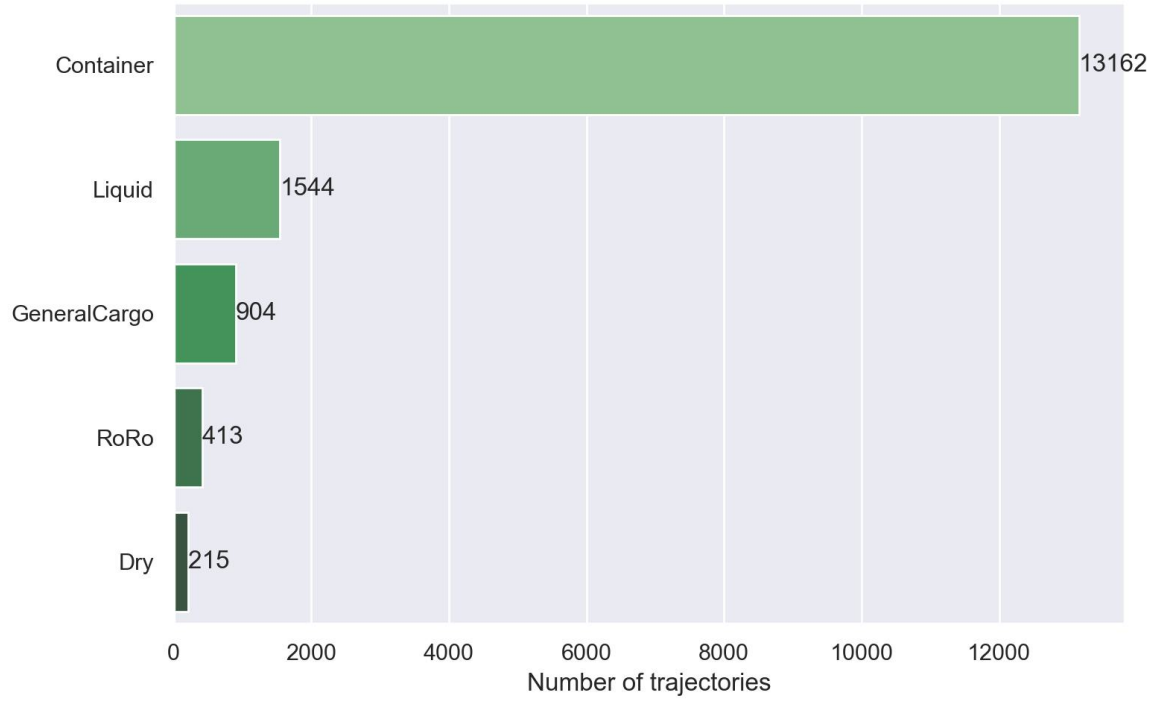

Figure 2.2: Trajectories per vessel category

2.2.4. Trajectory destinations

The vessels passing the area of interest (as defined in section 1.1), have more than 200 different unique destinations. Since most of these final destinations are outside the area of interest, a categorization of the final destinations is conducted. The trajectories are grouped in six distinct categories based on their final destination, see Table 2.4.

New route destination	Group of vessels
EUT	Vessels that have a destination at the Europa Terminal
NZT	Vessels that have a destination at the Noordzee Terminal
NORTH	Vessels that leave the port, i.e. are heading north
SOUTH	Vessels continuing upriver past EUT
ZVS	Vessels with a destination behind the locks which are passing the Zandvliet Lock to reach this destination
BES	Vessels with a destination behind the locks which are passing the Berendrecht Lock to reach this destination

Table 2.4: The new route destination categories

The categories EUT, NZT, NORTH, SOUTH, ZVS and BES are considered to be the new destinations of the vessel. Figure 2.3 illustrates the vessel trajectories split by this new destination and Figure 2.4 shows the number of trajectories per destination type. The start positions of vessels are marked in blue, and the end positions in red. In most cases, mapping the actual final destination to one of these categories is straightforward. For instance, if the final destination is a place at the Europa Terminal such as ”S853”, the category is EUT. However, when the final destination is behind locks, the specific lock through which the vessel will pass cannot be determined in advance. Nevertheless, by utilizing the full historical trajectory it can determined by geographical inspection through which lock the vessel went. In live scenarios, available lock schedules can be used to deduct the lock the vessel will pass through. Thus, it can be assumed that this method of mapping final destinations to the six categories can be applied to each vessel during both the training and inference phases.

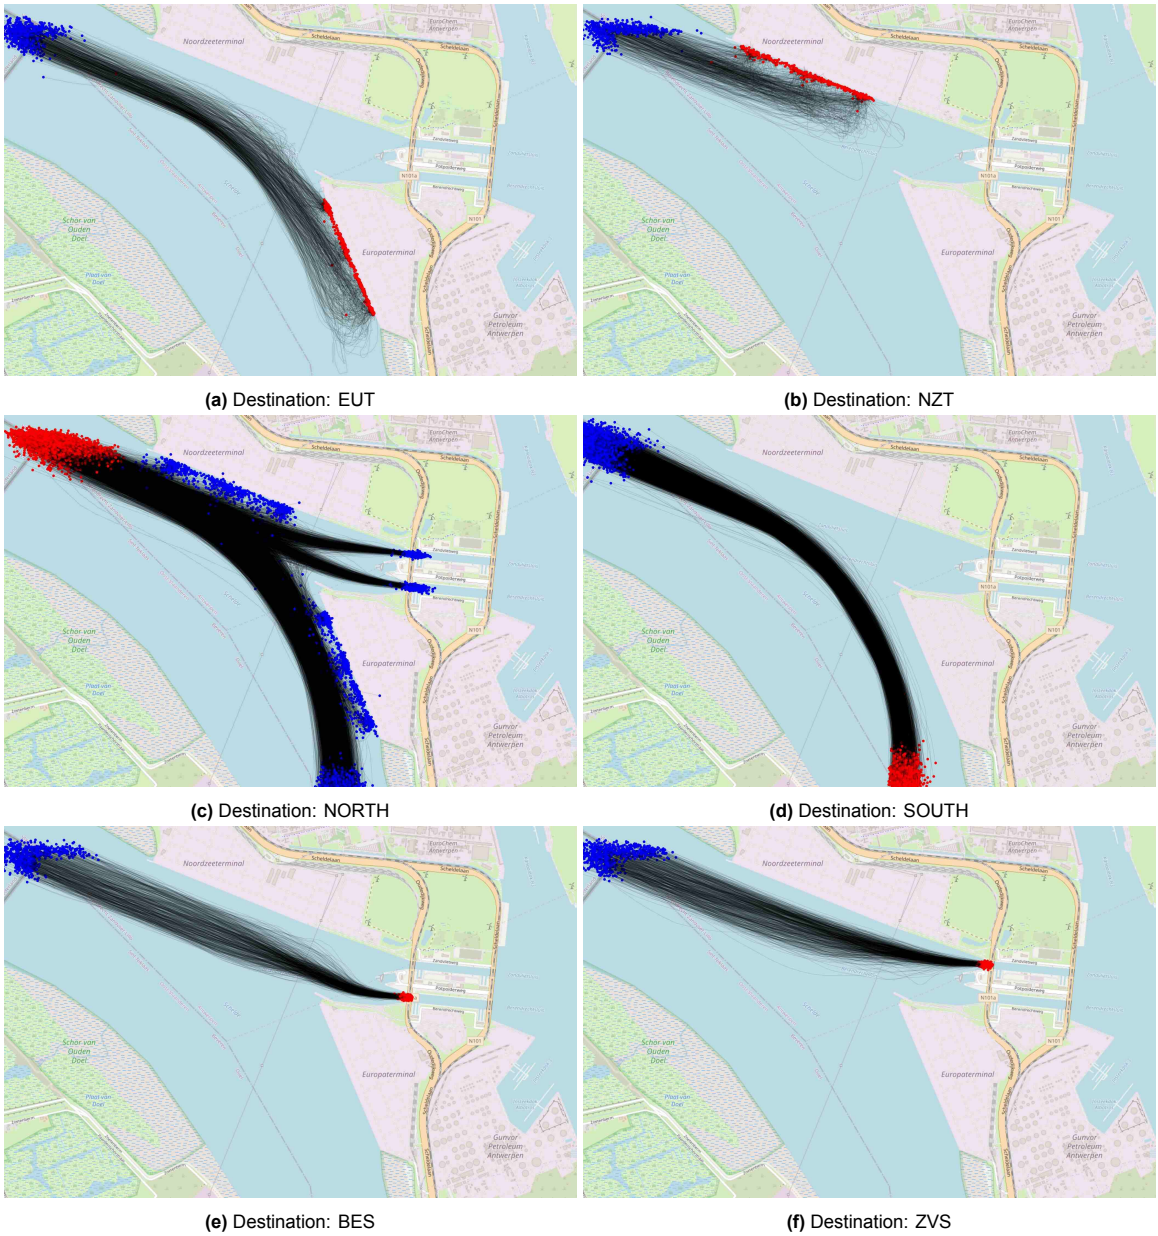

(a) Destination: EUT

(b) Destination: NZT

(c) Destination: NORTH

(d) Destination: SOUTH

(e) Destination: BES

(f) Destination: ZVS

Figure 2.3: All vessel trajectories categorized by destination. Blue points indicate the start of a trajectory, red points indicate the end

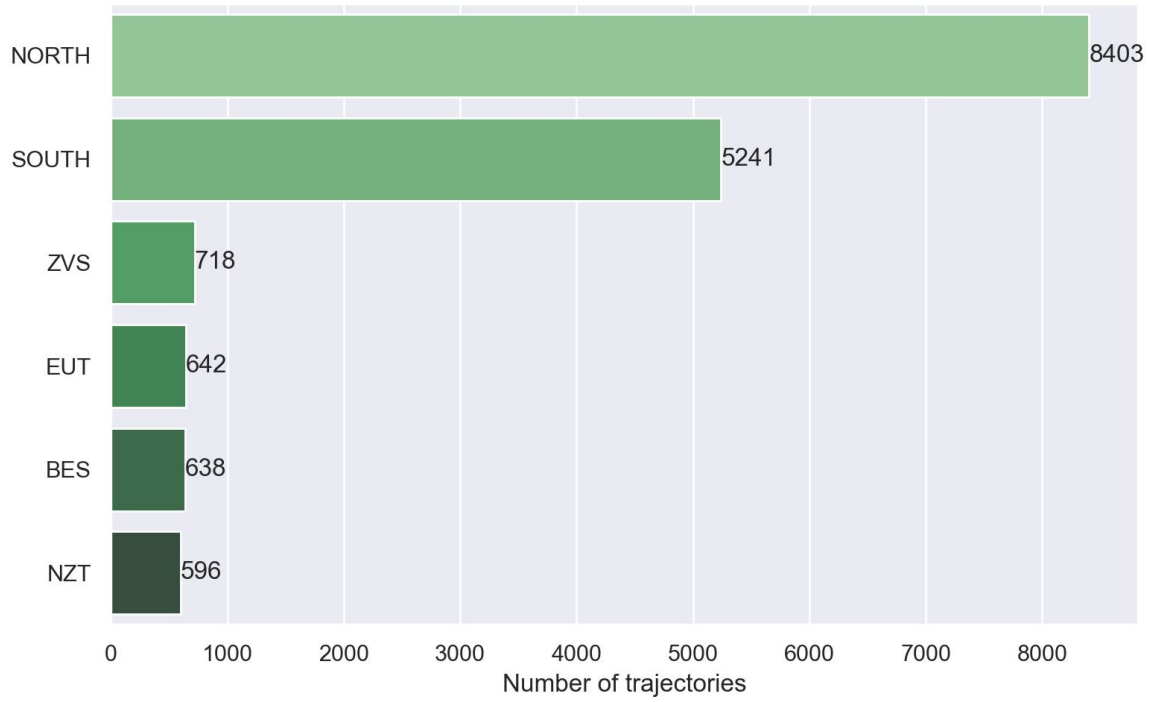

Figure 2.4: Number of trajectories per destination

3

Trajectory Prediction

The goal of this study is to develop a model for predicting the future trajectory of a vessel, including swing manoeuvres, based on its past trajectory. This chapter focuses on the methods for predicting trajectories, before delving into swing manoeuvres in Chapter 4. The problem formulation is presented in section 3.1, while section 3.2 provides an overview of the relevant literature on this topic. In section 3.4, the model utilized in this study to predict simple trajectories (i.e., those without swing manoeuvres) is introduced.

3.1. Problem formulation

A trajectory is a path that a vessel sails within a certain time period. Since this study only focuses on a scoped area of the Port of Antwerp-Bruges, trajectories only contain vessel positions within this area. Figure 2.3 visualises some example trajectories.

A trajectory can formally be defined as a set V of n data points $V = \{v_1, v_2, \ldots, v_{n-1}, v_n\}$. Each data point v_i is a vector $v_i = (p_i, q_i, t_i, s_i, h_i)$, representing the current state of the vessel. p_i holds the longitude position of the vessel, q_i the latitude position, t_i the epoch timestamp in milliseconds, s_i the speed over ground in knots and h_i the heading of the vessel in degrees. The time between two consecutive data points is assumed to be constant. This constant time window is established through an interpolation preprocessing step, described in section 2.2. Additionally, the vector m holds meta data about the vessel and its journey. Specifically, m holds the following information: destination of the journey, category of the journey (i.e. arrival or departure), the vessel type, the vessel's length, the vessel's width, the vessel's depth and the draught of the vessel during the journey.

Given a meta data vector m and a historical trajectory which can be of varying length, the task is to predict a future trajectory. Since this study only considers a scoped area of the Port of Antwerp-Bruges in this study, the predicted future trajectory should stay inside this area. Therefore, the length of the future trajectory is automatically limited by this requirement. This formalization of the problem provides an answer to research question RQ1: How is the problem of trajectory prediction formalized?

3.2. Background

The large availability of AIS data records has enabled the possibility of predicting trajectories in the domain of vessels. However, the task of trajectory prediction is not unique to the prediction of vessel movements and is relevant in more domains, such as human, aircraft or car trajectory prediction. This section reviews the approaches used in the existing literature on agent trajectory prediction, where an agent can be a vessel, human, car, etc.

3.2.1. Modelling and data based approaches

A basic and commonly used model for vessel trajectory prediction is to assume near constant speed and course [28]. Improvements on such a model are made by modelling the speed of a vessel using an Ornstein-Uhlenbeck process [38]. The Kalman filter is a different algorithm which can also be used to model and predict the movement of objects [42]. However, these physic-based approaches do not utilise the knowledge of historical AIS data. Prediction methods which are based on historic AIS data often cluster historical trajectories, classify incoming AIS data points to one of these clusters and then construct a prediction [21, 3]. These methods can differ in clustering algorithms and trajectory distance measures [29, 15, 22]. Besides such a trajectory-based similarity search prediction model, Alizadeh et al. presents a point-based similarity search prediction approach [3]. Here, each AIS record is treated as a singular point, and spatial, speed, and course variables are used to calculate the distances between the vessel's most recent point and all the historical points measured. By identifying the most similar historical point to the target point, the subsequent locations of the vessel can be predicted based on that point. A similar approach is used by Hexeberg et. al [21].

3.2.2. Deep learning-based approaches

The success of deep neural networks has led to a significant advancement in the field of trajectory prediction in recent years. Deep learning approaches used in literature take a historic agent state representation, pass this to a deep learning model and output a certain representation of the predicted future trajectory. The agent state is often represented as a vector of features, which include the timestamp, speed, heading and position under a certain coordinate system [9, 35, 6, 27, 2]. The input of the learning model then consists of a vector of these historic agent states. Next to this agent state, some researches use additional scene context representations to give the model more information. A popular scene context representation are images [9, 36, 40, 11, 30].

The future trajectory, i.e. the model's output, can simply be given as a point set [45, 35, 27, 2]. This single modality representation aligns with the input representation. Outputs can also be represented as probability heat-maps [31] or probabilities can be attached to each prediction [36, 30]. Instead of outputting exact positional values, a model could also predict a mean position μ and variance σ around it. Controlling uncertainty in this manner is done in [10].

In the literature, several neural network architectures are proposed for the actual deep-learning model. Since the input representation is often a vector of trajectory points, recurrent neural networks can be used to encode the historic trajectory [9, 45, 6, 34, 1]. Given this encoding, a second recurrent neural network can decode the vector to a future trajectory, this encoder-decoder architecture is also used in other studies [16, 45, 8]. Image inputs are often processed using convolutional neural networks (CNN) [36, 31] and graph neural networks (GNN) can be used to model interaction effects [41, 32, 13].

Supervised learning is the most common method used to train these deep learning models and the loss function commonly uses a mean-square-error (MSE) loss to compare the predicted trajectory against the truth [46]. Additional components can be added to the loss function to smooth the prediction, such as an inconsistency penalty, penetration penalty or dispersion penalty as used in [12].

A good trajectory prediction model should be able to capture the multimodality of the task, i.e. the model should be able to handle distinct end locations properly. However, models tend to suffer from mode collapse, which means average and non-realistic trajectories are predicted instead of diverse predictions. Dendorfer et al. give an excellent example of a neural network suffering from mode collapse when predicting pedestrian trajectories near a crossroad [14]. A way to counter this problem is to first use a model to cluster trajectories and then use a trajectory prediction model per cluster [17, 43]. By first clustering the trajectories, the neural network needs to learn less trajectory variations and can potentially generalize better.

In literature, an alternative method that is employed involves making a prediction of the goal location of the trajectory, followed by a prediction of the trajectory itself. Instead of asking a model to extrapolate a past trajectory, this method changes the problem to an interpolation task between the predicted goal position and the historic trajectory. Such approaches are commonly two staged, a target location is predicted based on the historical trajectory and scene context and this target prediction assists later trajectory forecasting [12, 18, 46, 47]. Target conditioned trajectory prediction has shown to perform better on diverse trajectories [14].

3.3. Neural networks for sequential data

A possible way to solve the challenge of predicting a vessel's future trajectory from its historical positions is by formulating it as a sequence-to-sequence task, which can be tackled using recurrent neural networks. Using recurrent neural networks for trajectory prediction is often done in other studies (3.2.2). This section will explain neural network structures which can be used to process sequential data. A good resource for more in-depth explanations of mathematics and ideas behind the different neural networks can be found in the book "Deep Learning" by Ian J. Goodfellow, Yoshua Bengio and Aaron Courville [19].

3.3.1. Recurrent Neural Network

A recurrent neural network (RNN) is a type of neural network specifically designed to handle sequential data and to memorize previously seen inputs. Its architecture differs from that of a normal feedforward neural network, as depicted in Figure 3.1a. In an RNN, the input is processed sequentially, and a hidden state is maintained throughout the process. The next hidden state is calculated as a function of both the current input x_i and the previous hidden state h_i. The current output is computed as a function of the current hidden state. To perform these computations, the model uses weight vectors w_x, w_h, w_y, and bias b_h as its parameters. These parameters are optimized during the training phase. The computation of the next hidden state is represented by the formula $h_{i+1} = f(w_x x_t + w_h h_t + b_h)$, while the current output at y_i is computed as $y_i = f(w_y h_t + b_y)$, where f is an activation function such as the sigmoid function. It's worth noting that each unit in an RNN uses the same set of parameters w_x, w_h, w_y, and b_h. This parameter sharing approach not only reduces the number of model parameters, but also enables RNNs to handle variable-length sequential inputs, which is a key advantage of this type of network.

3.3.2. Long Short-Term Memory

A Long Short Term Memory neural network is a variant of the vanilla recurrent neural network as described in the previous section. It is considered an improvement, since it is capable of modelling long-term dependencies in the input sequence better. Although a vanilla RNN can theoretically keep track of both long-term and short-term dependencies in input sequences, a problem arises during the back-propagation algorithm. During training gradients are backpropagated through the unfolded network. Since input sequences can be long, the network can contain a lot of RNN units. Therefore, gradients are multiplied quite often. If small gradients are repeatedly multiplied, the final gradient used in the parameter update step tends to go to zero and if large gradients are repeatedly multiplied then the gradients used in the update step will be extremely large. Both a vanishing gradient and an exploding gradient do no good to the learning process.

A LSTM network tries to solve this problem of vanishing and exploding gradients and is therefore considered to be better in keeping track of long-term dependencies in input sequences than the vanilla RNN. It does this by not only keeping track of a hidden state h_i, but also a cell state c_i, see Figure 3.1c. This cell state is updated in each unit of the network. A LSTM has three gates to control this update of the cell state: a forget gate, input gate and output gate. The forget gate controls what information in the cell state can be forgotten after processing the current input. The input gate controls what new information will be encoded in the cell state given the current input. Finally, the output gate controls what information in the cell state is used for the hidden state of the next unit. By using these three gates, the LSTM can update its cell state depending on whether the currently processed input is important or not.

3.3.3. Gated Recurrent Unit

A gated recurrent neural network is again an advancement over the vanilla RNN discussed in section 3.3.1. Similar to the LSTM network, GRUs can address the issue of vanishing or exploding gradients. A GRU architecture has a hidden state and internally updates this hidden state using an update gate and reset gate. The update gate determines which past information should be carried forward to the output, i.e. the next hidden state, while the reset gate decides what past information to discard. By utilizing these gates, a GRU can retain essential information from the distant past and discard irrelevant information.

Both a LSTM and GRU network can handle long-term dependencies better than a vanilla RNN. GRUs require fewer computations and memory, making them faster than LSTMs. However, LSTMs may perform better on datasets with long sequences. Ultimately, the choice between using a LSTM or GRU depends on the specific use case [44, 7].

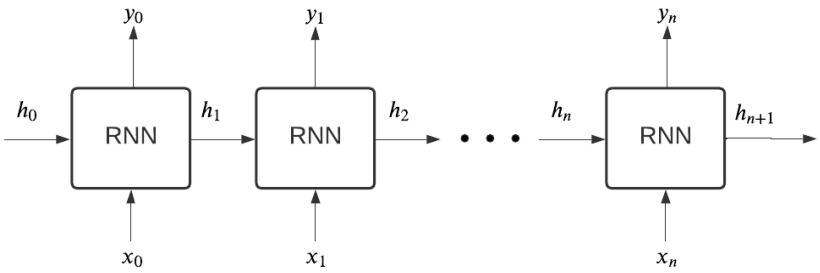

(a) (Vanilla) Recurrent Neural Network

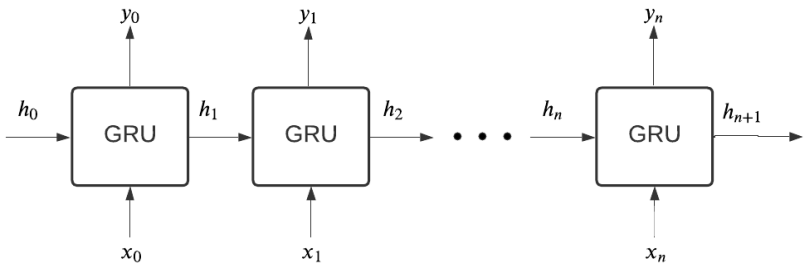

(b) Gated Recurrent Neural Network

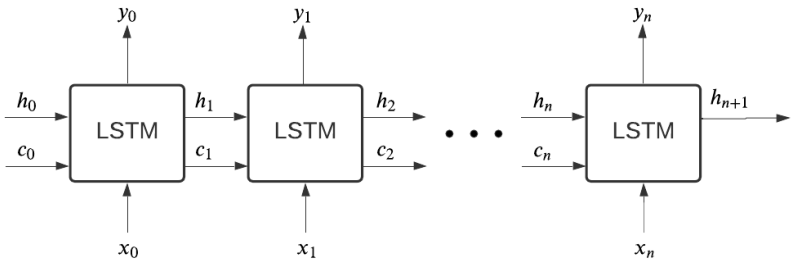

(c) Long Short-Term Memory Network

Figure 3.1: Schematic overview of neural networks for sequential data

3.4. Target conditioned trajectory prediction

Predicting the future trajectory of a vessel based on its historical positions can be treated as a sequence-to-sequence task, which can be addressed using recurrent neural networks as discussed in the previous section. This section delves into the model utilized for predicting future trajectories based on past data points.

3.4.1. Encoder-decoder structure

An encoder-decoder architecture with two recurrent neural networks can be utilized to predict future vessel trajectories based on historical positions. The first recurrent neural network encodes the historic trajectory into a vector embedding, while the second recurrent neural network generates a future trajectory based on this vector embedding. The vessel meta data can also be incorporated into the decoder network. This model can be trained end-to-end on the historical trajectories. A schematic representation of this model utilizing gated recurrent units can be observed in Figure 3.2. In this encoder-decoder setup, the encoding network essentially summarizes the historic trajectory in a vector embedding and gives this summary to the decoder network to generate a future trajectory.

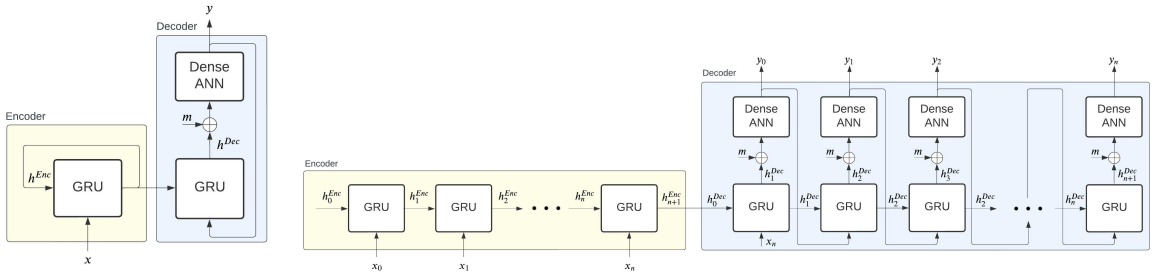

(a) Overview **(b)** Unfolded

Figure 3.2: Schematic overview of a GRU encoder-decoder network

3.4.2. Complete model

During model training, the decoder can be constrained to produce a future trajectory of the same length as the training sample. However, during inference, it is uncertain when the decoder should terminate the generation of trajectory points. As explained in section 3.1, the length of the predicted future trajectory is unknown. The only known requirement is that the predicted trajectories must stop once the vessel leaves the scoped area. Formulating a clear stopping criterion for the decoder would be beneficial, and this can be achieved if the target position, which is the last position of the future trajectory, was known. If the target position is known, a stopping criterion that uses the distance to the target position can determine whether the decoder should generate more future trajectory points or not. However, the nature of the trajectory prediction problem does not provide access to the last future trajectory point since the future trajectory is something that needs to be predicted. Despite the lack of knowledge about the target point, it can still be predicted using the output of the encoder, which is a summary of the historic trajectory, and the meta data. Given a predicted target position and the known historic trajectory, the decoder can construct the future trajectory. The task for the decoder now has now changed from extrapolation to interpolation given this target position.

A complete trajectory prediction model can now be constructed by incorporating the concepts from the encoder-decoder network structure and the target prediction process. Figure 3.3 illustrates the complete model, where the yellow portion represents the encoder. The encoder uses a recurrent neural network to encode the historical data points into a vector embedding. Before the encoder does this, it embeds the input in a higher dimension using a multilayer perceptron. The encoder output vector is combined with the meta data and fed through a feedforward neural network to predict the target position, which consists of a latitude, longitude and heading. Once the target position is predicted, it is concatenated with the encoder output and then passed through a multilayer perceptron to produce the

initial hidden state of the decoder. The decoder sequentially generates future points by taking predicted future points as input for the next block. For the first input, the last point in the historical trajectory is used. Besides these future points the decoder also uses the target position as input. The future point and target position are concatenated and passed through a multilayer perceptron to form the input of the decoder. The decoder is given the target position at each time stamp to help remember this important piece of information.

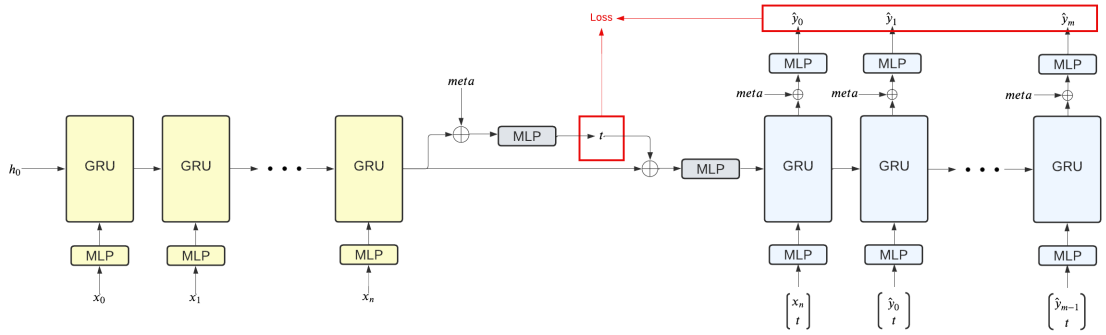

Figure 3.3: Complete target conditioned trajectory prediction model network structure using gated recurrent units

During training a fixed prediction window is used, i.e. the decoder will return the correct number of future trajectory points. However, during inference a stopping criterion is still required. Since a target point is now predicted, a stopping criterion can be formulated based on this point. The decoder should stop generating points if one of these criteria is met:

1. The current point is less than 500 meters from the target point and the decrease in distance to the target position between the current point and the previous point is less than δ
2. The future trajectory contains more than F points

The first criterion that the vessel comes closer to the target position every timestamp, i.e. 30 seconds and the second criterion provides an upper bound on the trajectory length. This upper bound is set to 100, as the majority (97%) of the historic trajectories contain less than 100 data points.

Figure 3.4 shows a distribution of the decrease in distance to the target point. A distinction is made between trajectories with destination NORTH or SOUTH and trajectories going to the locks or terminals (without swing). These latter categories of trajectories will approach their target point with a lower speed namely. It can be concluded that most of the historical data is in line with the assumption that the distance to the target point keeps decreasing over time. A reasonable value to use for δ would be that of the first percentile. Given the historical data, this would mean $\delta = 3$ and $\delta = 65$ are used for trajectories with destination ZVS, BES, NZT or EUT and destinations NORTH or SOUTH, respectively.

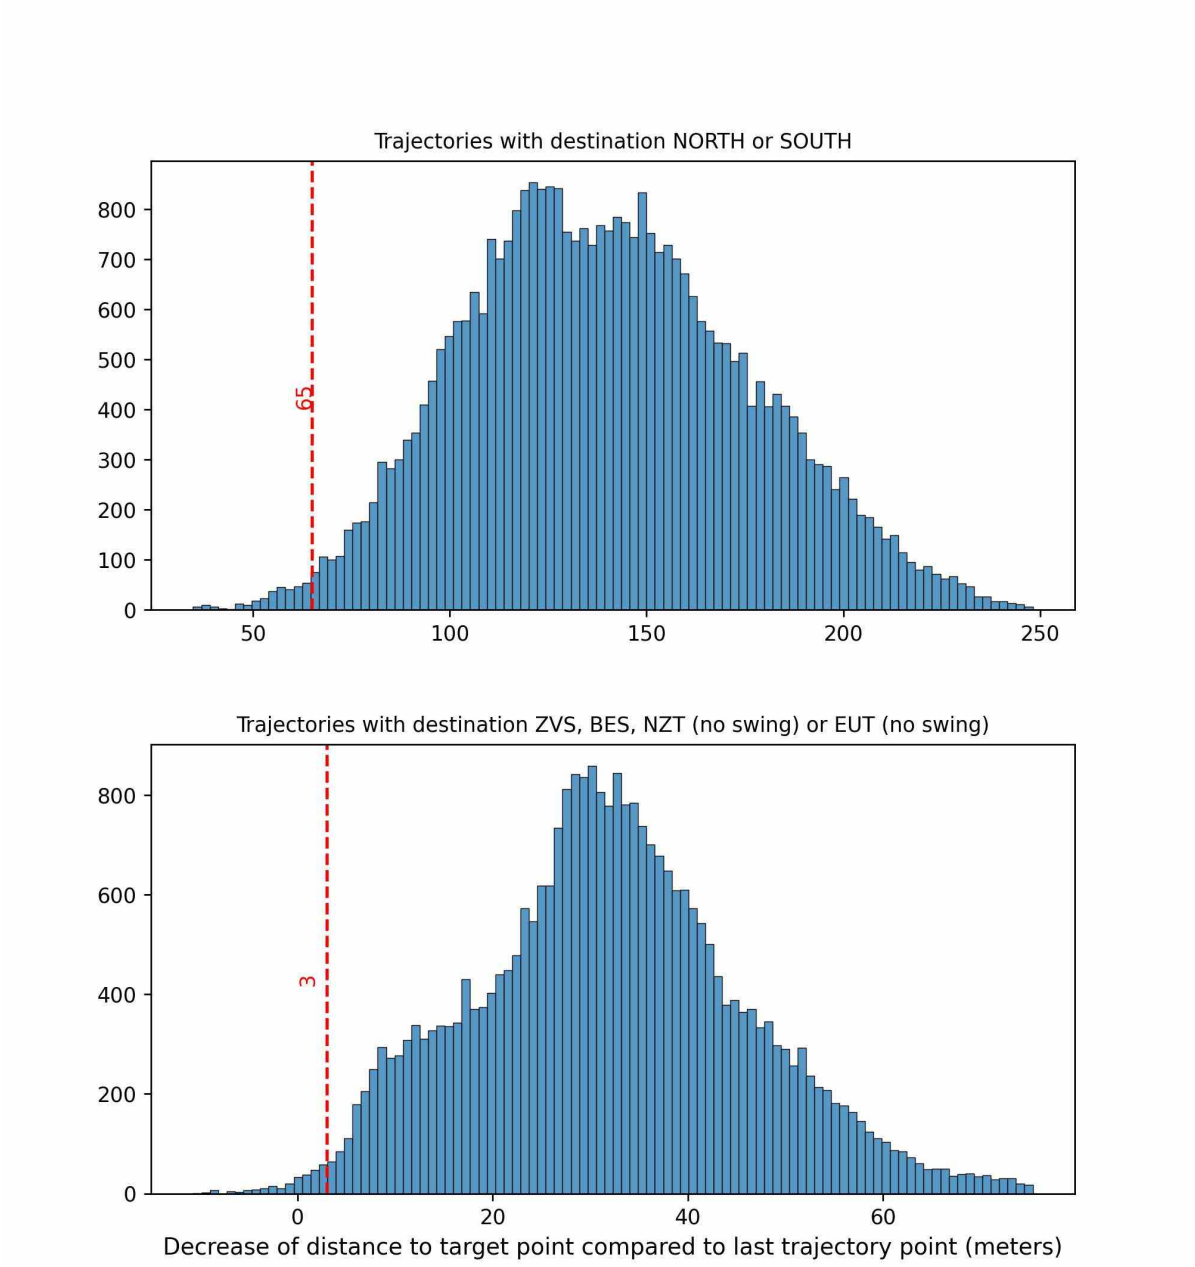

Figure 3.4: Distribution of decrease in distance to target for all points within 500 meters of the final location

3.4.3. Loss function

The final goal of the complete model is to predict the future trajectory as accurately as possible by using a predicted target point that should be as close as possible to the actual target point. Therefore, the loss of the model should align with this goal. Given a historic trajectory $\{x_0, x_1, \ldots, x_m\}$, a future trajectory $\{y_0, y_1, \ldots, y_m\}$, a predicted target point t and a predicted future trajectory $\{\hat{y}_0, \hat{y}_1, \ldots, \hat{y}_m\}$, the loss of the model is calculated with equation 3.1

$$L = c_1 * \frac{1}{m}\sum_{i=0}^{m}(y_i - \hat{y}_i)^2 + c_2 * (t - \hat{y}_m)^2 + c_3 * (y_m - \hat{y}_m)^2 + c_4 * \frac{1}{m}\sum_{i=0}^{m}(y_i - \hat{y}_{i-1})^2. \tag{3.1}$$

In this equation, c_1, c_2, c_3, c_4 are constants. In the last term of the formula x_n is used for $\hat{y}_{-1}$ (i.e.

when $i = 0$ in the sum).

The first term of the loss function is a mean squared error component on the full future trajectory. The second and third terms are both squared errors on the target prediction and the last predicted trajectory point. Finally, the last term is a mean squared error between the predicted trajectory and the predicted trajectory shifted back one time stamp. By adding this term, the model is restricted from taking unrealistic speed and position changes between adjacent frames. Such an inconsistency penalty term to smooth the predicted trajectories by encouraging temporal consistency between trajectory points is used more often in trajectory prediction models [12].

4

Swing manoeuvres

Swings are U-turn-like manoeuvres which occur when container vessels arrive at or depart from terminals. Container vessels interacting with terminals are required to make a swing manoeuvre on arrival or departure, which means about half of the container trajectories will contain swing manoeuvres. Since swing manoeuvres are such a big part of container vessel trajectories, they cannot be ignored in trajectory predictions. This chapter delves deeper into swing manoeuvres.

Section 4.1 defines what a swing manoeuvre exactly is and section 4.2 explores the current literature around swing manoeuvres. The process of detecting swings from trajectories is discussed in section 4.3, followed by section 4.4 explaining how swing manoeuvres can be extracted from the full trajectory. Once the swing manoeuvres are detected and extracted, properties of the historical swing manoeuvres in the Port of Antwerp-Bruges can be investigated. This is done in section 4.5. Lastly, in section 4.6, a model for predicting swings is presented.

4.1. Swing definition

The Port of Antwerp-Bruges is situated far inland, making it Europe's most inland port [5]. The journey to the port area for container vessels arriving from the sea involves navigating 80 kilometers through The Scheldt river, which is the only waterway linking the port to the North Sea. As there is only one entry and exit point for sea vessels, a U-turn manoeuvre is necessary for vessels on a round-trip, commonly known as a swing manoeuvre. Typically, swing manoeuvres are executed near the terminal location of the vessel, either before it docks on arrival or after it is released from the quay during departure. These manoeuvres are time-consuming and require tugboat assistance. Moreover, for safety reasons, a buffer zone must be maintained around the swinging vessel, which can block parts of the waterway, affecting the passage of other vessels. Environmental factors such as tide conditions and wind strength also play a crucial role in deciding to swing or not. If a captain decides not to swing on arrival, it essentially postpones this manoeuvre to departure. Ultimately, the captain makes the decision to swing or not while taking all environmental and traffic related factors into account.

4.2. Background

Swing manoeuvres are not unique for the Port of Antwerp-Bruges. More ports use swing manoeuvres to align container vessels with their desired direction. Especially in ports with only one waterway for entering and leaving the port, vessels need to make such a manoeuvre. Since it often raises difficulties to find a suitable location for container vessel to execute a swing manoeuvre, as is the problem in the Port of Antwerp-Bruges (section 1.1), studies have been conducted on designing port areas in an optimal manner that considers swing manoeuvres [26, 25]. However, changing existing port layouts is difficult and the fact that container vessels have become longer over time adds to the problem. Paulauskas discusses the influence of these big container vessels on ports designed using old safety and navigation standards [39]. With his analysis on turning basins, he claims that due to modern accu-

rate navigational and measurement equipment, the old safety standards implemented in several ports should be reviewed.

The trajectory prediction model presented by Hexeberg et. al obtained some large errors on certain trajectories, which were caused by swing manoeuvres as reported by the authors [21]. Besides this notion of swing manoeuvres in the vessel trajectory prediction literature, no literature has been found on trajectory prediction methods which incorporate swing manoeuvres.

4.3. Swing detection

Swing detection is the process of determining whether a historical trajectory of a vessel contains a swing manoeuvre or not. Such manoeuvres are distinguished by a considerable change in the vessel's heading over a short duration of time. As the vessel executes a U-turn, the heading change is typically around 180 degrees over the entire manoeuvre. Additionally, swing manoeuvres usually take around ± 20 minutes to complete. Based on these observations, swing detection logic can be created to identify trajectories with swing manoeuvres.

Given is a trajectory representation V of n data points $V = \{\vec{v}_1, \vec{v}_2, \ldots, \vec{v}_{n-1}, \vec{v}_n\}$. Each data point $\vec{v}_i$ is a vector $\vec{v}_i = (p_i, q_i, t_i, s_i, h_i)$, where $p_i \in P$ represents the longitude position, $q_i \in Q$ the latitude position, $t_i \in T$ the epoch timestamp in milliseconds, $s_i \in S$ the speed over ground in knots and $h_i \in H$ the heading of the vessel in degrees.

Now define a function $f : H \times H \rightarrow [0, 180]$ which takes two headings and outputs the smallest difference in degrees between these headings. The smallest difference between the headings is considered, because of the modular 360 nature of circles. As an example, take $h_i = 40$ and $h_{i+1} = 345$, visualised in Figure 4.1. By using the formula $|h_i - h_{i+1}|$ to calculate the difference between headings, the difference would be $|h_i - h_{i+1}| = |40 - 345| = 305$ degrees. This angle is visualised in blue in Figure 4.1. Although this is not a wrong answer to the question "what is the difference between angle h_i and h_{i+1}?", it is more likely that the vessel made a anticlockwise turn and the actual change in heading is $\alpha = 65$ degrees.

Therefore, to make sure the detection process does not extract false positives, i.e. trajectories without a swing indicated as having a swing, f returns the smallest difference between two headings. Thus, the following formula is used for f:

$$f(h_i, h_j) = 180 - ||h_i - h_j| - 180| \tag{4.1}$$

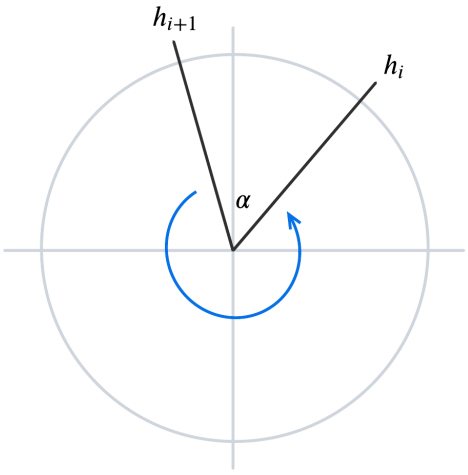

Figure 4.1: Two headings of an anticlockwise swinging vessel

By combining domain knowledge and this function, it is possible to formalize a definition for a trajectory that includes a swing manoeuvre:

V contains a swing manoeuvre $\Leftrightarrow \{(i,j)|1 \leq i,j \leq n \wedge j > i \wedge t_j - t_i \leq 25 * 60000 \wedge f(h_i,h_j) \geq 125\} \neq \emptyset$ (4.2)

This formalisation means that a vessel makes a swing manoeuvre if and only if there exist two data points on the trajectory which are no more than 25 minutes apart and have a difference in heading larger or equal to 125 degrees, where the difference in heading is calculated using equation 4.1. Note that a factor 60000 is used to convert the 25 minutes to milliseconds, which is the unit of t_i and t_j.

Domain knowledge indicates that swing manoeuvres can take up to 20 minutes, therefore 25 minutes is a reasonably safe upper bound. The choice for 125 degrees is based on visual inspection. In the histogram in Figure 4.2 the red dotted line at $x = 125$ clearly splits two clusters of trajectories. To the left of this line are trajectories who do not contain a swing manoeuvre and too the right of this line are those who do. An expected peak around 180 degrees is visible, which represents full U-turn manoeuvres. There are also some swings outside this peak, i.e. the smaller cluster of swings around 150 degrees. These originate from vessels arriving to and departing from the Europa Terminal (EUT). Figure 4.3 shows an example of a container vessel departing from the EUT using a anticlockwise swing. Since the EUT is not parallel to the waterway which container vessels use to arrive from and depart to the sea, a swing of around 150 degrees is sufficient.

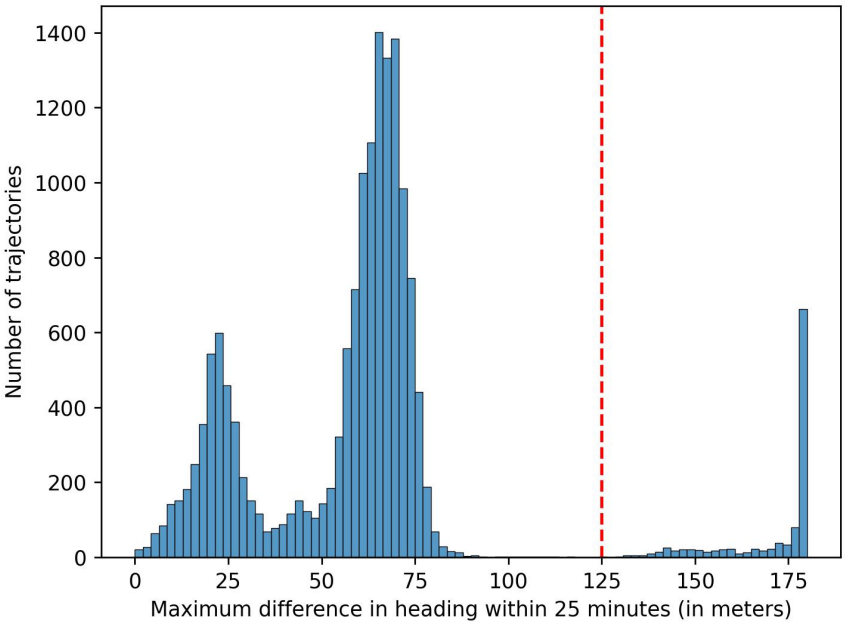

Figure 4.2: Distribution of maximum directional change in container vessel trajectories

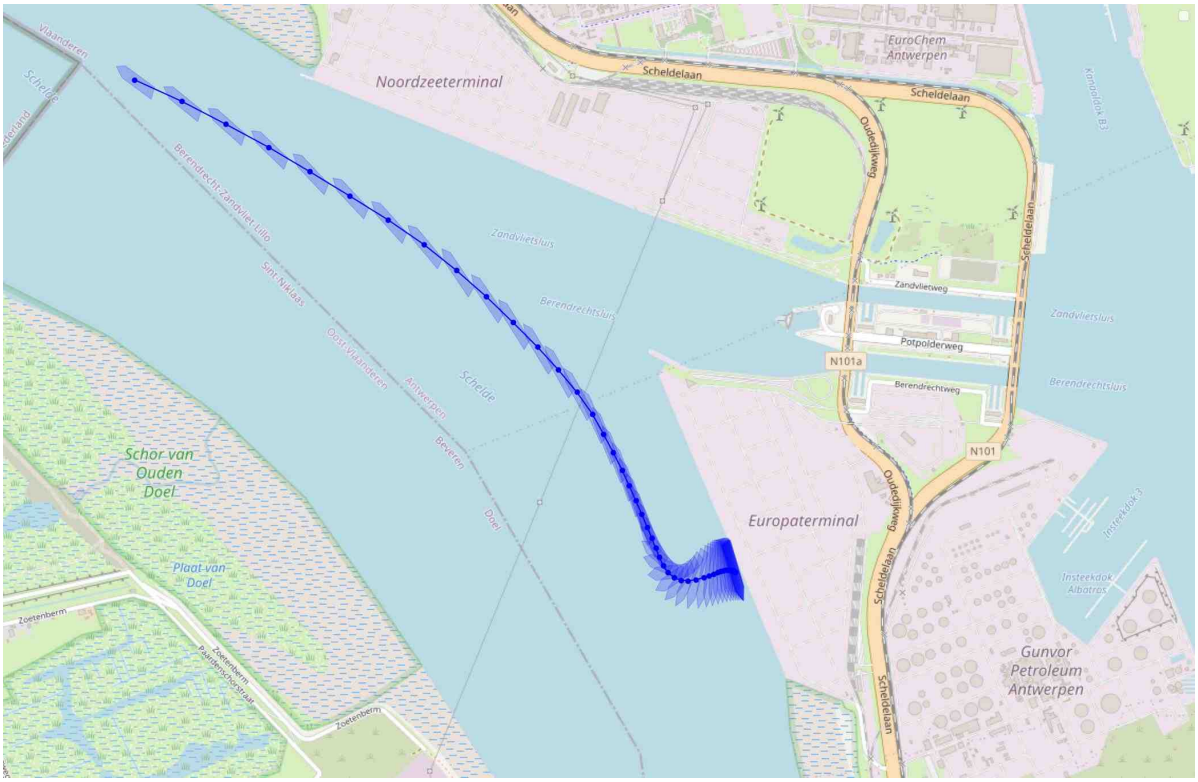

Figure 4.3: Example of vessel leaving the EUT with an anticlockwise swing

4.4. Swing extraction

Once swing manoeuvres are detected from historical trajectories, it is important to extract them from the entire trajectory. This means that only the segment of the trajectory in which the vessel performs a swing manoeuvre is kept. For instance, if a trajectory lasts for 40 minutes, the relevant portion of the trajectory is only the 10 minutes in which the vessel is performing the swing manoeuvre.

Swing manoeuvres are often done approximately in place, which means the speed of a vessel during its trajectory is an important factor to look at. Furthermore, the change of the vessel's heading over time is something to look at. This is expressed by the yaw rate, which is the directional change in heading in degrees per minute. For the directional change between two headings equation 4.1 can be used again.

Using the same trajectory definition as in section 4.3, a swing window $w = [b_{swing}, e_{swing}]$ can be defined where $b_{swing}, e_{swing} \in \{1, \ldots, n\}$. This window indicates the part of the trajectory in which the vessel is swinging.

Further let $\gamma : \{1, \ldots, n-1\} \to \mathbb{R}$ be a function which, given $i \in \{1, \ldots, n-1\}$, computes the yaw rate in the interval $[i, i+1]$.

$$\gamma(i) = \frac{f(h_i, h_{i+1})}{(t_{i+1} - t_i) * \frac{1}{30000}} \tag{4.3}$$

Next, let the indicator function $\psi : \{1, \ldots, n-1\} \to \{0, 1\}$ indicates if a data point $\vec{v}_i$ is during a swing manoeuvre or not. If $\psi(i) = 1$, then the vessel is performing a swing manoeuvre. However, the inverse is not necessarily true.

$$\psi(i) = \begin{cases} 1, & \text{if } 1 \leq i \leq n \wedge \gamma(i) \geq 10 \wedge s_i \leq 2 \\ 0, & \text{otherwise} \end{cases} \tag{4.4}$$

Hence, a data point is classified to be during a swing manoeuvre if the change in heading is currently larger or equal to 10 degrees per minute and the speed is lower or equal to 2 knots, which is approxi-

mately 3.7 kilometers per hour. The motivation behind these numbers come from domain knowledge and visual inspection.

Using the indicator function ψ a window $[b'_{swing}, e'_{swing}]$ can be defined in which the vessel is definitely swinging. An assumption is made that each trajectory comprises a single swing manoeuvre, which is reasonable.

$$b'_{swing} = \min_i[\psi(i) = 1] \text{ where } i \in \{1, \ldots, n-1\} \tag{4.5}$$

$$e'_{swing} = \max_i[\psi(i) = 1] \text{ where } i \in \{1, \ldots, n-1\} \tag{4.6}$$

Vessels can have a lower yaw rate than 10 degrees per minute in the beginning or ending of a swing. Moreover, small vessels could swing at larger speeds than 2 knots, especially in front of the EUT. This means that the use of the indicator function in equation 4.5 and 4.6 might mean that the window $[b'_{swing}, e'_{swing}]$ is smaller than the actual swinging window $[b_{swing}, e_{swing}]$, i.e. $b'_{swing} > b_{swing} \wedge e'_{swing} < e_{swing}$. Using a lower yaw rate bound and higher speed bound in the definition of ψ to solve the problem is not possible. ψ will violates its meaning then, as it will be too sensitive and falsely classify small course changes during the trajectory of a vessel as being part of a swing manoeuvre.

To still be able to extract the full swing window, the window is extended on both sides by including data points where the yaw rate is still above 6 degrees per minute. This procedure is written down in equation 4.7 and 4.8.

$$b_{swing} = \min\{i \in \{1, \ldots, b'_{swing}\} | (\forall k \in \{i-1, \ldots, b'_{swing}\})[\gamma(k) \geq 6]\} \tag{4.7}$$

$$e_{swing} = min(n, 1 + \max\{i \in \{e'_{swing}, \ldots, n\} | (\forall k \in \{e'_{swing}, \ldots, n\})[\gamma(k) \geq 6]\}) \tag{4.8}$$

An example of swing extraction for a specific trajectory using the procedure described above is shown in Figure 4.4. The points in this graph are plotted in one-minute intervals for clarity. In this example $[b'_{swing}, e'_{swing}] = [6, 10]$ and $[b_{swing}, e_{swing}] = [4, 13]$. The horizontal lines are plotted at a yaw rate of 6 and 10.

Figure 4.5 shows the vessel heading, yaw rate and speed of four more example trajectories. The extracted swing window using equation 4.7 and 4.8 is indicated using two dotted black lines. In these plots the time between two adjacent points is 30 seconds.

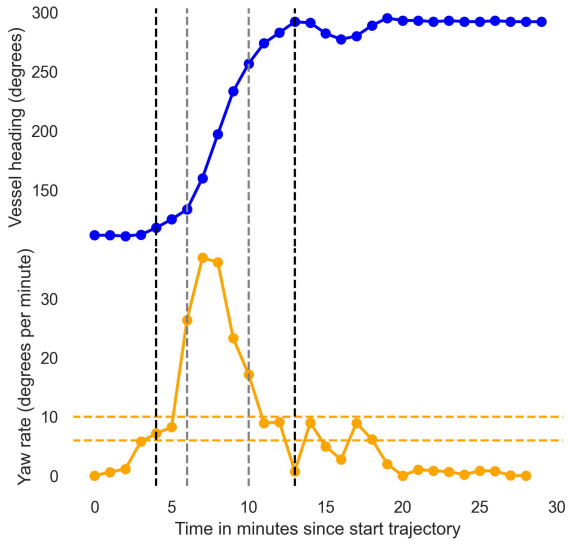

Figure 4.4: A tight swing window (grey dotted line) and extended swing window (black dotted line). Interval between points is one minute.

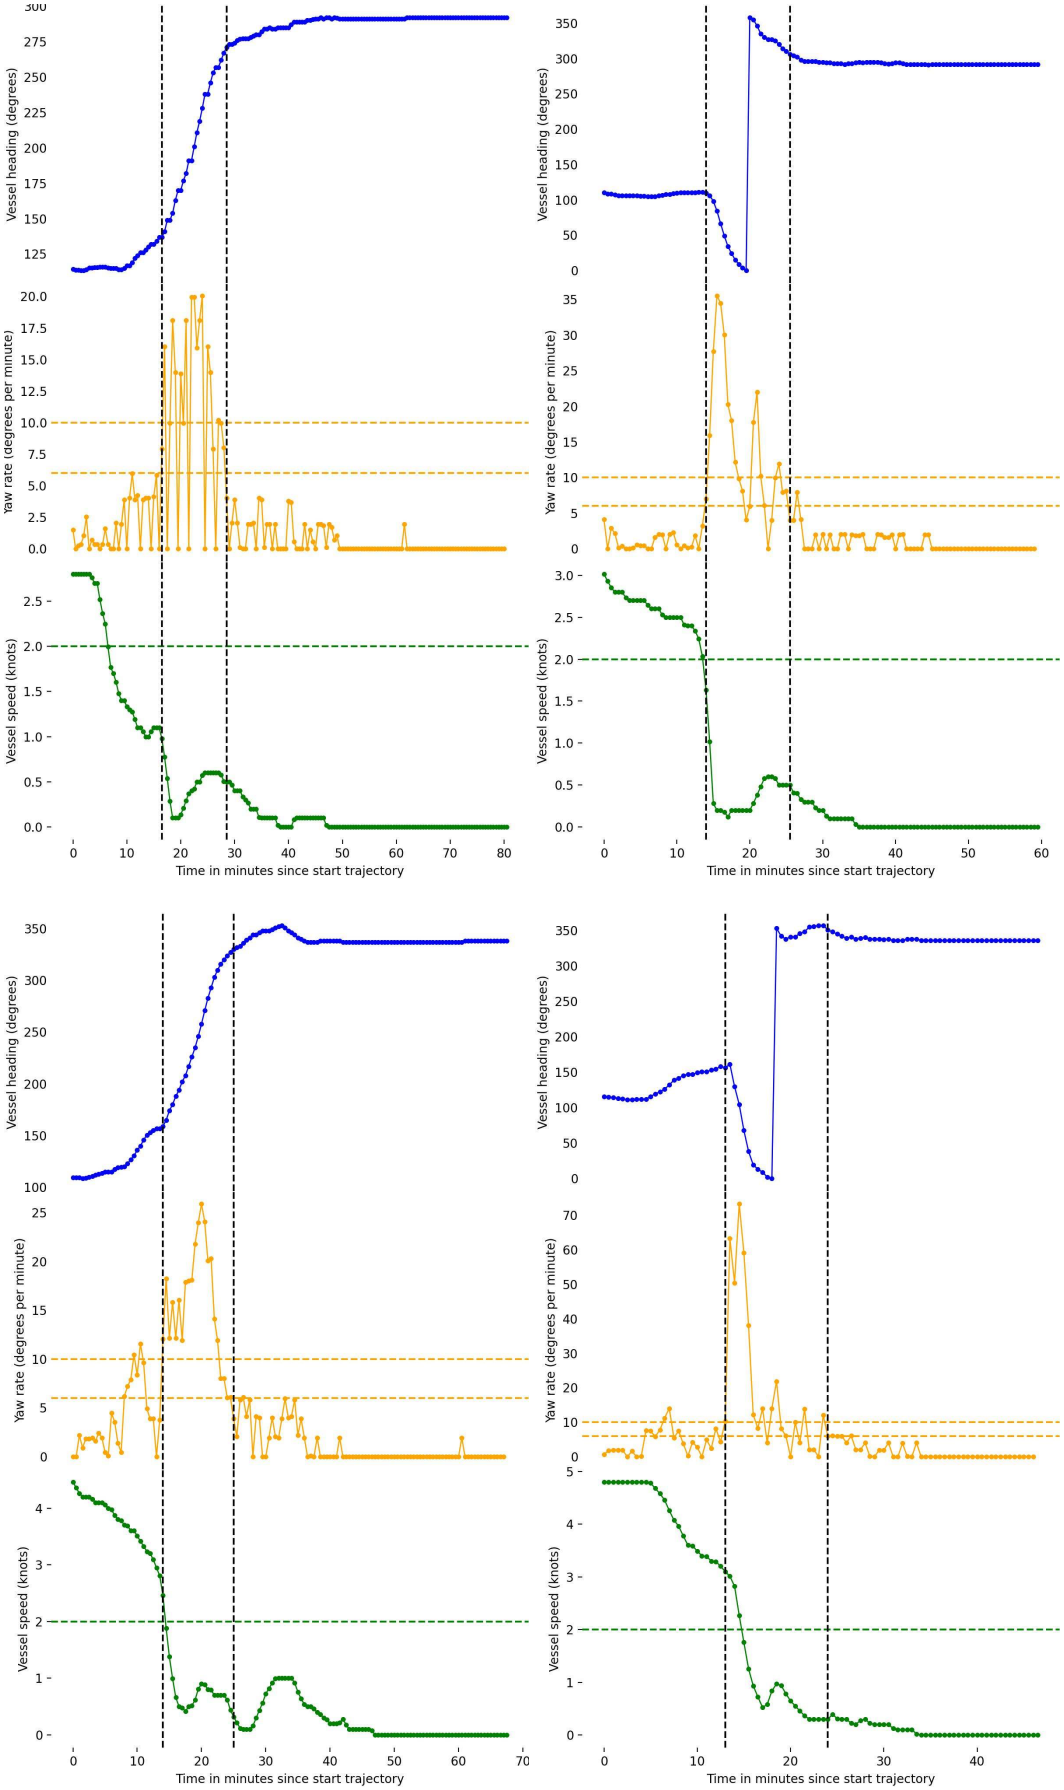

Figure 4.5: The vessel heading, yaw rate, speed and swing window during four different trajectories (interval is 30 seconds)

4.5. Historical swings

Given the historical AIS data (chapter 2) and the swing detection (section 4.3) and swing extraction (section 4.4) procedures, an analysis on historical swing manoeuvres is possible.

The Port of Antwerp-Bruges provided all container vessel trajectory AIS data during the years 2021 and 2022. Figure 4.6 shows an overview of the number of swing manoeuvres per route type and terminal within these years and Figure 4.7 visualises the swing duration per terminal. As can be seen from the latter figure, swing manoeuvres took between 6 to 15 minutes. The swing manoeuvres at the Noordzee Terminal seem to take a bit longer on average. Given that it takes longer to swing large vessels, one possible explanation for the latter observation is that swinging vessels at the Noordzee Terminal have a mean length of 258 meters, while those at the Europa Terminal have a mean length of 232 meters.

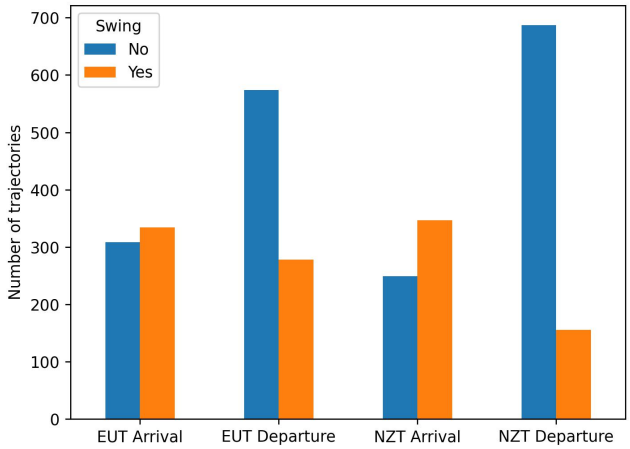

Figure 4.6: The number of swing manoeuvres per route type and terminal

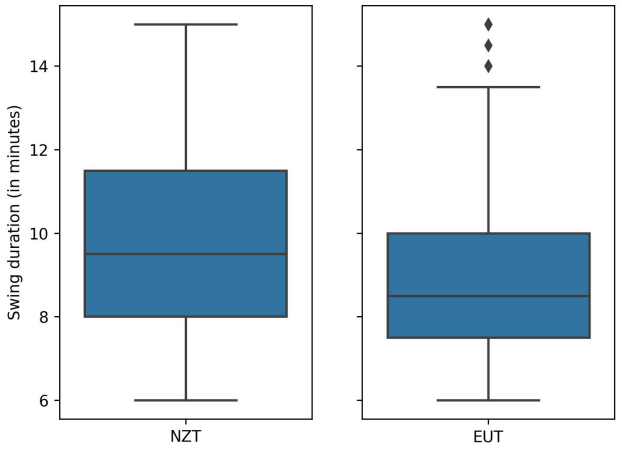

Figure 4.7: The duration of swing manoeuvres per terminal. Outliers outside 3 standard deviations are removed.

The length of a vessel and the draught of a vessel can influence the duration of a swing manoeuvre. Small vessels can make a swing manoeuvre relatively quickly on their own, but larger vessel need tugboats. Figure 4.8 shows the impact of vessel length and draught on the swing manoeuvre duration.

As expected, a positive relationship between vessel length and swing duration and vessel draught and swing duration is observed.

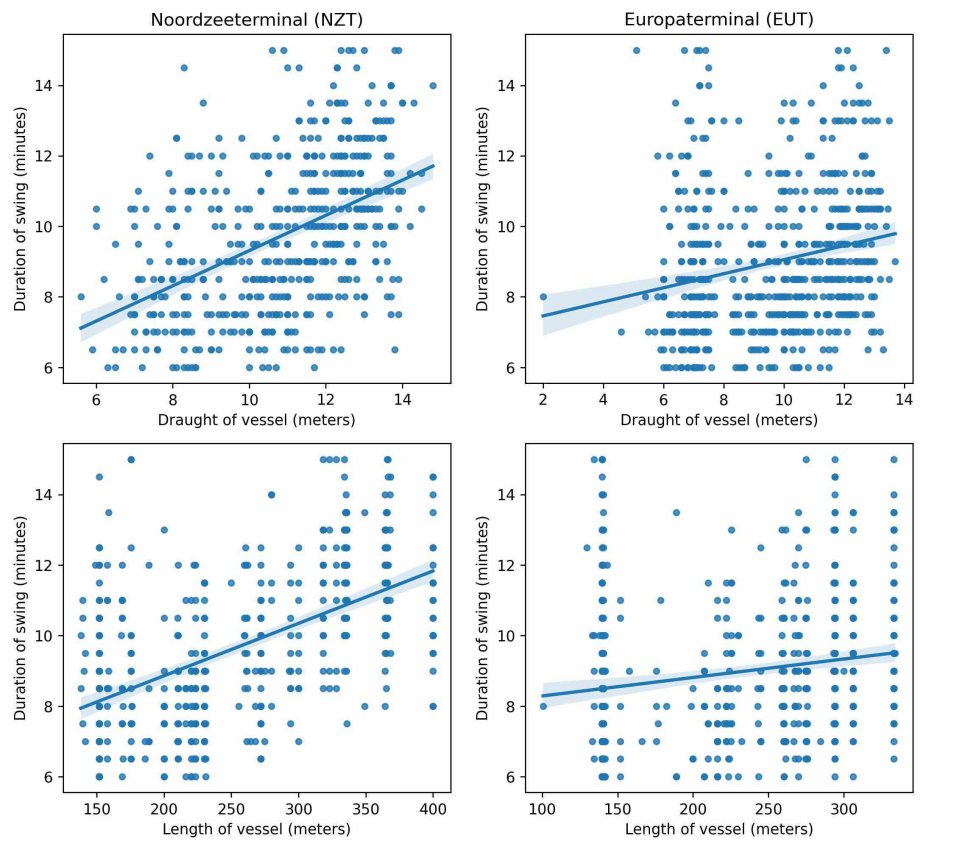

Figure 4.8: The influence of the vessel length and draught on the swing manoeuvre duration

As discussed in section 1.1, the location of the swing manoeuvre is hard to pick, as all locations block different parts of the waterway. Therefore, it is useful to gain some insights in the location of historical swing manoeuvres. Figure 4.9 visualises the trajectory during all swing manoeuvres and the location for all historic swing manoeuvres can be seen in Figure 4.10. The location of a swing is defined as the location halfway through the swing manoeuvre. It seems that smaller vessels swinging at the Europa Terminal often do this near their terminal location, whereas larger vessel swing more in front of the locks. At the Noordzee Terminal both small and large vessel swing approximately in front of their terminal location. However, smaller vessels tend to swing a bit closer to the terminal than larger vessels. At the Noordzee Terminal, small vessels made a swing manoeuvre 294 meters from the terminal location on average, while large vessels did this 389 meters from the quay. For the Europa Terminal, smaller vessels also made the swing manoeuvre closer to the terminal, namely 308 meters compared to 650 meters for large vessels.

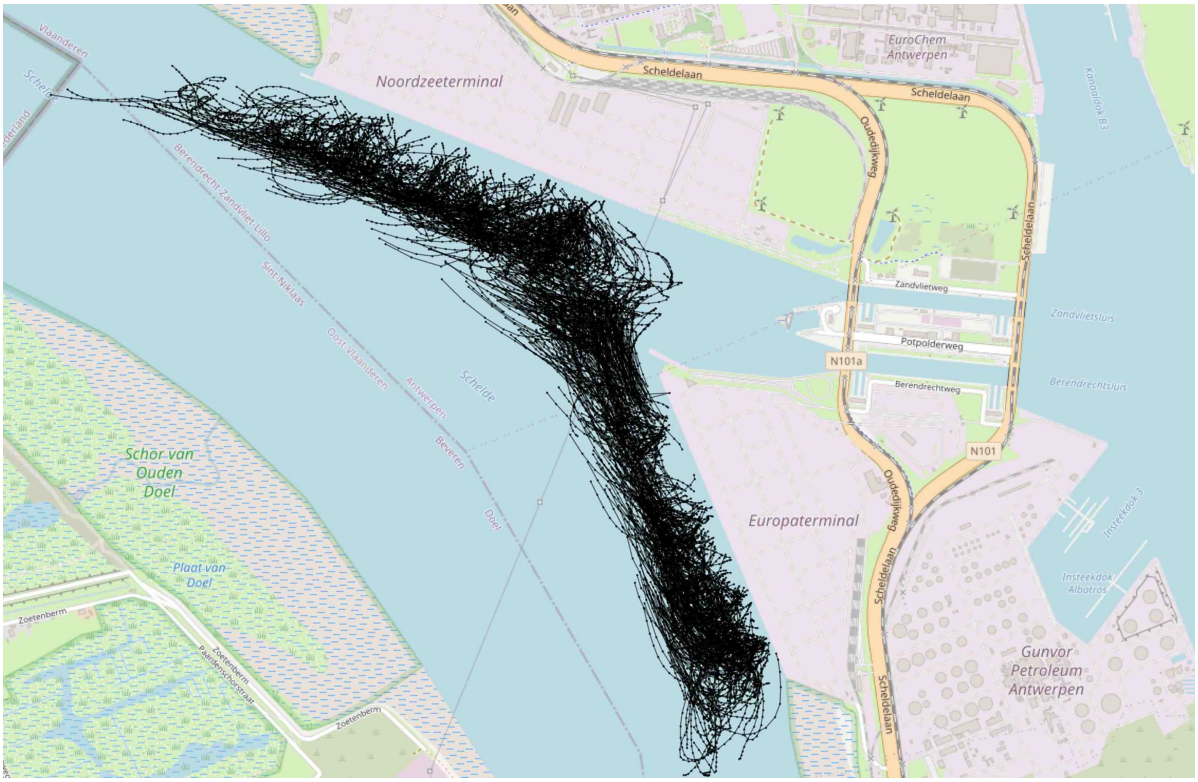

Figure 4.9: Historic swings of all container vessels in the year 2021 and 2022

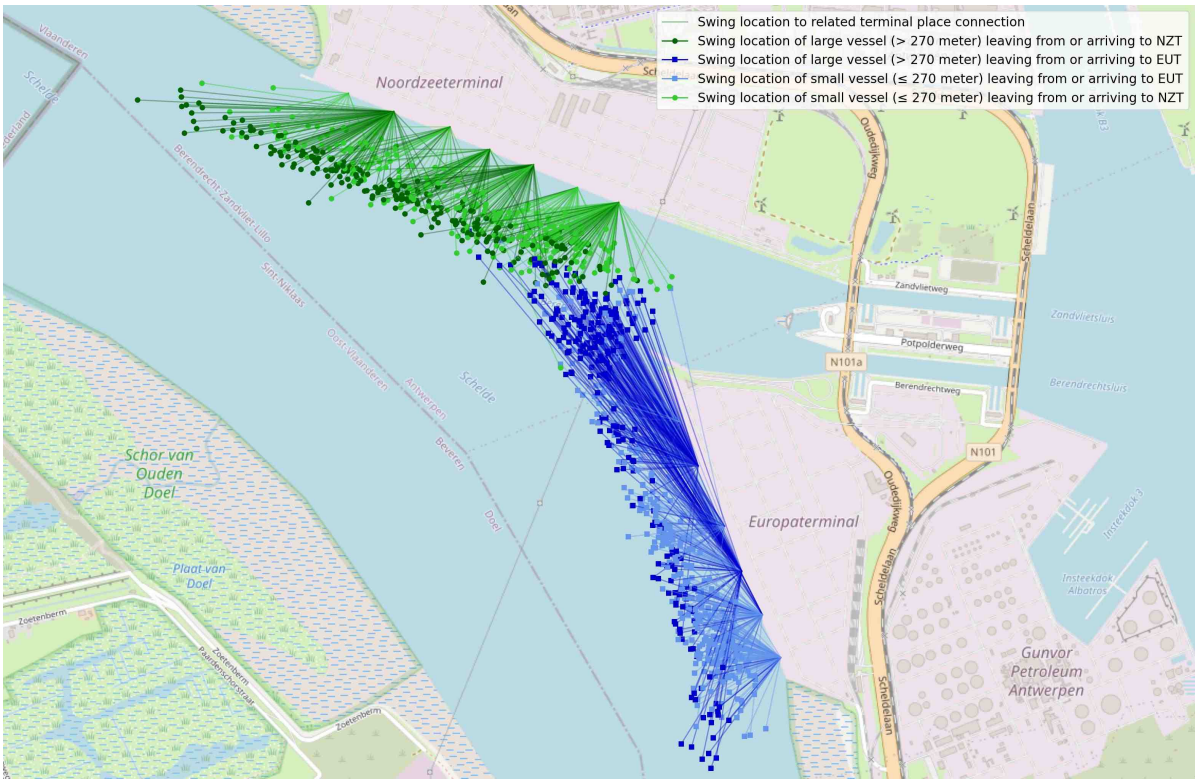

Figure 4.10: Historic swing locations of all container vessels in the year 2021 and 2022

4.6. Swing prediction

Swing manoeuvres are crucial for determining the paths of container vessels. In order to have a precise understanding of upcoming situations in the port, it is essential to consider these manoeuvres. As a result, the need to anticipate the occurrence and movement of swing manoeuvres becomes apparent. Predicting swing manoeuvres can be done in different levels of details. An option is to predict the whole swing manoeuvre, i.e. the location and headings during the whole manoeuvre. However, prediction of the full swing manoeuvre is a difficult problem to solve since these manoeuvres are relatively rare and can vary a lot in their exact details. The goal of the study is to make a capacity prediction of the waterways. Exact heading and position of the vessel during the swing manoeuvre is less important than the occupation of the waterway resources and the duration of this occupation. Therefore, an alternative could be to only model the location and duration of the swing and assume the vessel will swing in place. However, the downside of this approach is that it ignores the movements which the vessel sometimes makes during the swing manoeuvre. Therefore, this study uses an in-between approach by predicting the start location, end location and duration of the swing. This abstraction keeps the relevant details of the swing while not modelling unnecessary details. Figure 4.11a, 4.11b and 4.11c visualise these three different ways of modelling a swing manoeuvre.

(a) Option 1: full swing manoeuvre (reality)

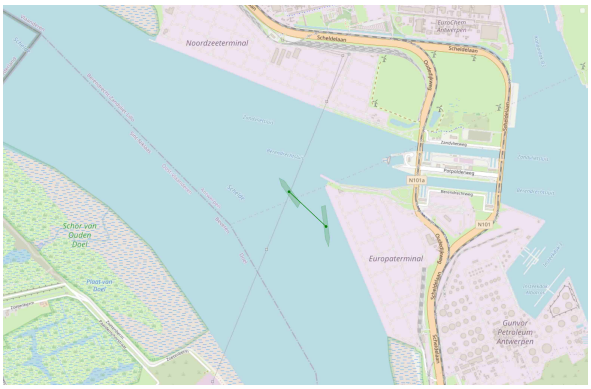

(b) Option 2: start and end of swing manoeuvre

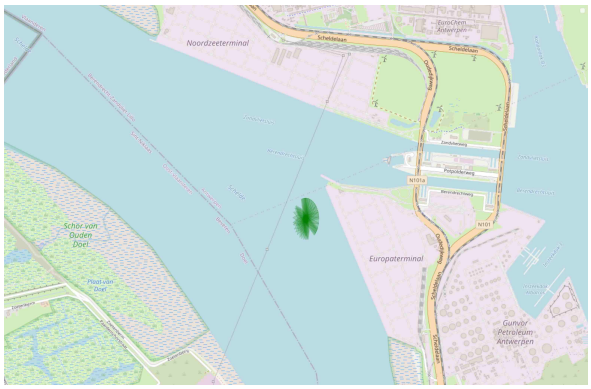

(c) Option 3: in place swing manoeuvre

Figure 4.11: Different options for modelling a swing manoeuvre

4.6.1. Swing occurrence prediction

Container vessels interacting with terminals need to swing either on arrival or departure. The captain's preference plays a crucial role in the decision whether to swing or not when arriving in the port. The captain bases its choice mainly on environmental factors, of which the tide condition is really important, and traffic conditions.

Determining whether a departing vessel will swing is easy. The way these vessels are oriented against the terminal reveal whether they require to swing or not. At the Europa Terminal, vessels of

which the heading is pointing to the south need to swing (see Figure 4.12a). Vessels with a heading pointing to the north can immediately leave without a swing manoeuvre (see Figure 4.12b). At the Noordzee Terminal, vessels with a heading pointing to the east need to swing (see Figure 4.12c) and vessels with a heading pointing to the west do not require a manoeuvre (see Figure 4.12d).

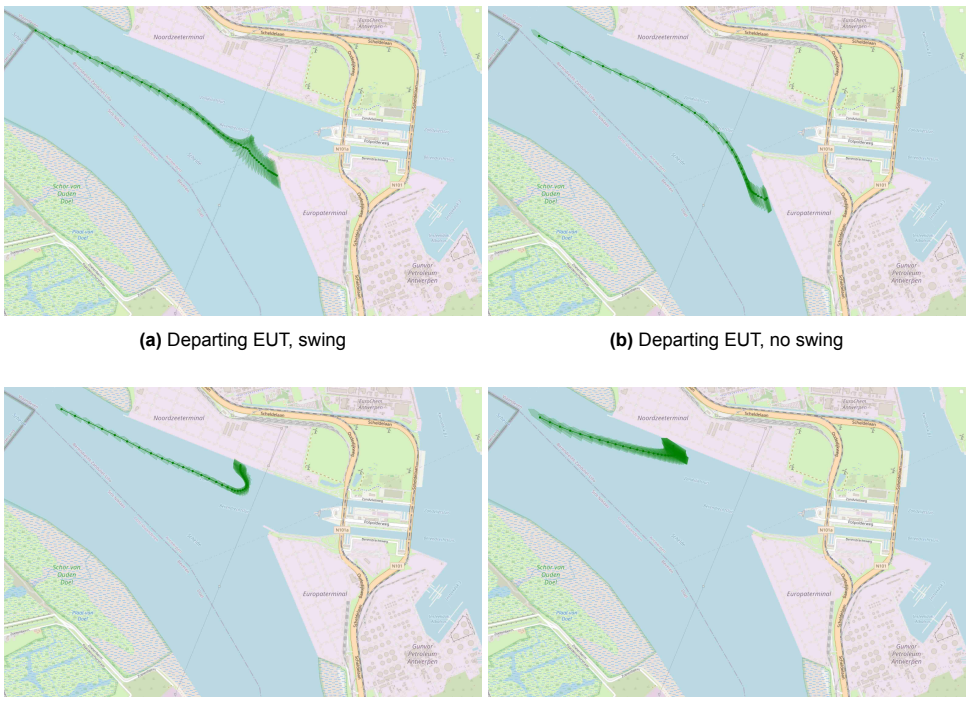

(a) Departing EUT, swing

(b) Departing EUT, no swing

(c) Departing NZT, swing

(d) Departing NZT, no swing

Figure 4.12: Vessels departing from the terminal

For arriving vessels, it is not straightforward whether they will swing or not. If an arriving vessel does not swing, it essentially postpones this manoeuvre to its departure and when the vessel does swing, it can later immediately depart. Given data about the vessel's journey and the environmental conditions, a prediction can be made on whether the vessel will swing on arrival or not. Locations of containers on the vessel are not taken into account and it is assumed that all captains have the same preference profile.

Based on vessel's meta data and on a fraction of the trajectory of varying length before the swing, the goal is to predict whether the vessel will swing or not. For this task the model visualised in Figure 4.13 is used. This model takes the historical trajectory (i.e. before the swing) and encodes it in a vector. This encoding is done in the same manner as historical trajectories are encoded in the target conditioned trajectory prediction model (section 3.4). In Figure 4.13 the historical trajectory is encoded using a GRU, but a LSTM or vanilla RNN could also be used. To this encoded vector, the meta data vector is concatenated and then the result is passed through a final dense neural network. The output layer of this neural network has one neuron with a sigmoid activation function. Therefore, this neuron outputs a value $\hat{y} \in [0, 1]$, $\hat{y} \geq 0.5$ is considered a swing prediction and $\hat{y} < 0.5$ as a prediction without a swing.

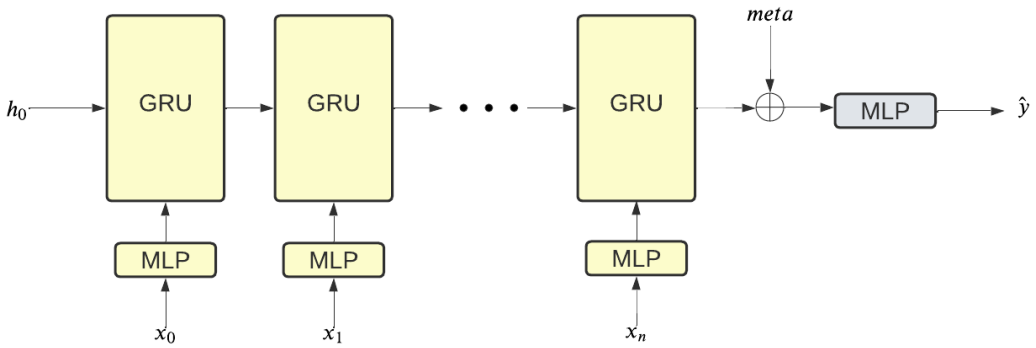

Figure 4.13: Swing occurrence prediction deep learning model architecture using a gated recurrent unit

4.6.2. Swing location prediction

After establishing through the methodology detailed in 4.6.1 that the vessel will undergo a swing manoeuvre, a follow up question is how this swing manoeuvre should look like. For the maritime traffic controllers in the port of Antwerp-Bruges, the most important knowledge to obtain is the location and duration of the swing. Knowing the exact movement during the swing is of less importance. With this requirement in mind, the problem can be simplified from predicting the whole swing manoeuvre to predicting only the location of the manoeuvre. From Figure 4.9, it becomes clear that vessels do not always swing in place: their longitude and latitude position often changes during the manoeuvre. Therefore, predicting only one swing location (as in Figure 4.10) and assuming in place swinging would abstract reality too much and result in inaccurate predictions. Instead, predicting the start and end position of the swing seems a usable approach. The start and end position are characterised by their longitude position, latitude position and heading. Therefore, a swing location prediction model should be able to predict these three characteristics for both the start and end of the swing based on a historic trajectory before the swing manoeuvre.

For this purpose, the same neural network architectural setup as in section 4.6.1 is used again. An encoder network will encode the historic trajectory in a vector and this encoded vector will be concatenated to the meta data vector and passed through a final feedforward neural network. The output layer of this feedforward neural network will contain 9 neurons:

- Neuron 1: the latitude position at the start of the swing
- Neuron 2: the longitude position at the start of the swing
- Neuron 3: the sine of the heading at the start of the swing
- Neuron 4: the cosine of the heading at the start of the swing
- Neuron 5: the latitude position at the end of the swing
- Neuron 6: the longitude position at the end of the swing
- Neuron 7: the sine of the heading at the end of the swing
- Neuron 8: the cosine of the heading at the end of the swing
- Neuron 9: the duration of the swing manoeuvre in minutes.

The first four neurons predict the start of the swing, the next four neurons the end of the swing and the last neuron predicts the duration of the swing. Using these neuron is the most straightforward way to model the start, end and duration of a swing manoeuvre. The sine and cosine of the headings are used instead of the heading directly because of the circular nature of the heading values. A heading change from 2 degrees to 359 degrees seems like a huge difference numerically, but visually these headings are quite close. The cosine and sine values of a heading of 2 degrees and a heading of 359 degrees have the property of being quite close to each other numerically.

5

Trajectory and manoeuvre prediction model

In order to make precise predictions of the future trajectory of a vessel in the Port of Antwerp-Bruges, a model must be capable of handling two crucial aspects of vessel movement: sailing a trajectory and swing manoeuvres. Chapter 3 delved extensively into the topic of trajectory prediction, while Chapter 4 focused on the analysis of swing manoeuvres and the models used to predict their occurrence and location.

This chapter builds upon these concepts by discussing a model that integrates both trajectory prediction and swing manoeuvre handling. By incorporating these two factors into a single model, it becomes possible to predict complete vessel trajectories in the Port of Antwerp-Bruges. Section 5.1 will discuss all possible route options which the model should consider and section 5.2 will give an overview of the model.

5.1. Vessel route options

To be able to predict complete trajectories, the model must be able to handle all scenarios, i.e. all vessel types and route origin and destination combinations in the scoped area of the Port of Antwerp-Bruges must be dealt with. Given the data, as described in section 2, an overview of all vessel route options is visible in Table 5.1.

Identifier	Route type	Route start	Route end	Involved terminal	Swing
R1	Arrival	NORTH	EUT	EUT	No
R2	Arrival	NORTH	EUT	EUT	Yes
R3	Departure	EUT	NORTH	EUT	No
R4	Departure	EUT	NORTH	EUT	Yes
R5	Arrival	NORTH	NZT	NZT	No
R6	Arrival	NORTH	NZT	NZT	Yes
R7	Departure	NZT	NORTH	NZT	No
R8	Departure	NZT	NORTH	NZT	Yes
R9	Arrival	NORTH	ZVS	-	-
R10	Arrival	NORTH	BES	-	-
R11	Arrival	NORTH	SOUTH	-	-
R12	Departure	ZVS	NORTH	-	-
R13	Departure	BES	NORTH	-	-
R14	Departure	SOUTH	NORTH	-	-

Table 5.1: All possible route options in the scoped area of the Port of Antwerp-Bruges

The last six route options, i.e. R9-14 in Table 5.1, do not involve any terminal and will not contain

a swing manoeuvre. Non-container vessels can only do route options R9-R14. The target conditioned trajectory prediction model as presented in Chapter 3 can therefore be used to predict the future trajectory of vessels within these categories. Route options R1 to R8 do involve interaction with one of the terminals. Four of these route options, R1, R3, R5 and R7, do not contain a swing manoeuvre. Therefore, the same trajectory prediction model can be used again to predict the future trajectories for these vessels. Thus, the only route options which need special attention are those with a swing manoeuvre, i.e. R2, R4, R6 and R8.

5.2. Model overview

As mentioned in section 5.1, all route options which do not contain a swing manoeuvre can be dealt with by using the target conditioned trajectory prediction model as presented in Chapter 3. For the route options which do include a swing manoeuvre, a combination of this trajectory prediction model and the swing location prediction model (Chapter 4) can be used. The swing location prediction model will predict the start and end location of the swing manoeuvre and the trajectory prediction model will take care of predicting the trajectory to the start position of the swing or from the end position of the swing onwards. Using this procedure, a model can be constructed which follows the steps as visualized in the diagram in Figure 5.1.

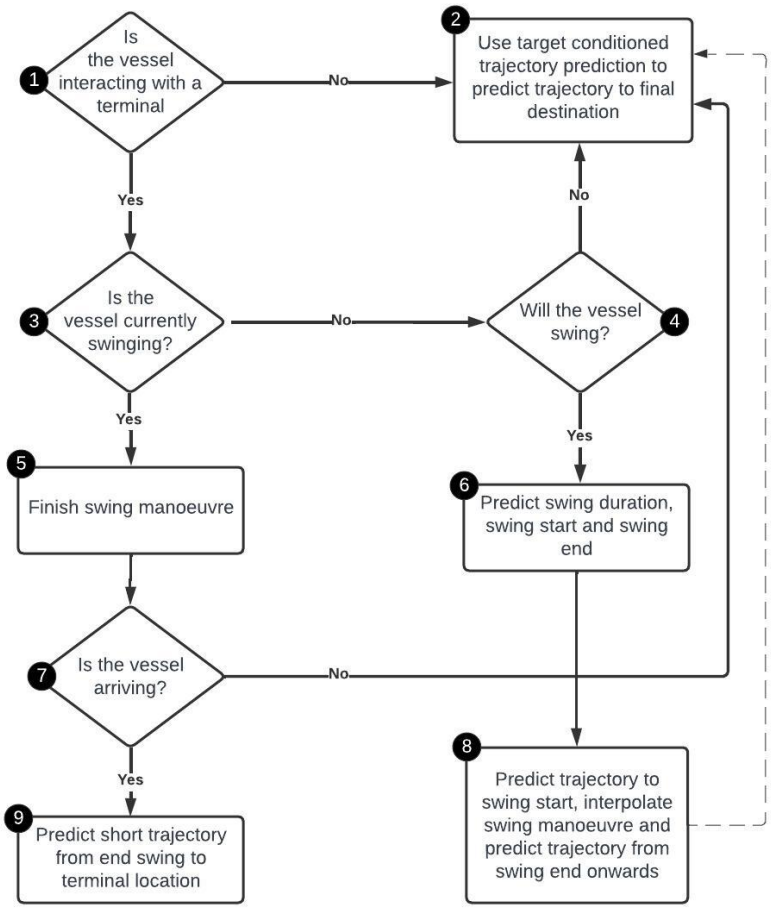

Figure 5.1: Overview of steps within the complete trajectory prediction model

The first step in the model is to determine whether the vessel will interact with a terminal, i.e. depart from or arrive to a terminal location. Determining this is rather straightforward, as the destination and origin of the vessel is known. If the vessel is not interacting with the terminal, the target conditioned trajectory prediction model can be used to predict the future trajectory. If the vessel is interacting with a terminal, the model checks whether the vessel is currently swinging. If so, the model finishes the swing manoeuvre and then predicts the trajectory from the end of the swing onwards. If the vessel is arriving in the port, this latter step boils down to predicting the short trajectory from the end of the swing to the terminal location, i.e. the docking trajectory (see section 5.2.1), and if the vessel is departing, the target conditioned trajectory prediction model is used again to predict the trajectory from the swing end onwards.

For vessels which are currently not swinging, the model checks if the vessel will swing in the future. For this the prediction model presented in section 4.6.1 is used. If the vessel won't swing in the future, the trajectory prediction model is used again. On the other hand, if the vessel does swing in the future, the swing start location, end location and duration is predicted using the model detailed in section 4.6.2. In this latter case, once the swing is predicted the model has to predict the trajectory to the swing start, interpolate the swing and then predict the trajectory from the swing end onwards. Predicting the trajectory to the swing start for departing vessels means predicting the undocking trajectory (see section 5.2.1) and for arriving vessels this means predicting a longer trajectory, for which the trajectory prediction model can be used again. The latitude, longitude and heading during the swing manoeuvre itself are linear interpolated. This is a simplification of the reality, but is sufficient since the duration of the swing manoeuvre and the approximate location are most important to maritime traffic controllers. Knowing the exact position and heading during the swing manoeuvre is of less importance.

Once the swing manoeuvre is interpolated, only the trajectory from the swing end onwards is left to complete the full trajectory. For arriving vessels this means predicting the docking trajectory (see section 5.2.1) and for departing vessels this means predicting a longer trajectory. In the latter case the trajectory prediction model is used again.

The steps in this model can handle all route options as presented in section 5.1. Table 5.2 gives an overview of the steps that the model takes to predict a complete trajectory for each of the route options.

Route option identifier	Steps
R1	$1 \rightarrow 3 \rightarrow 4 \rightarrow 2$
R2	$1 \rightarrow 3 \rightarrow ((4 \rightarrow 6 \rightarrow 8) \vee (5 \rightarrow 7 \rightarrow 9))$
R3	$1 \rightarrow 3 \rightarrow 4 \rightarrow 2$
R4	$1 \rightarrow 3 \rightarrow ((4 \rightarrow 6 \rightarrow 8) \vee (5 \rightarrow 7 \rightarrow 2))$
R5	$1 \rightarrow 3 \rightarrow 4 \rightarrow 2$
R6	$1 \rightarrow 3 \rightarrow ((4 \rightarrow 6 \rightarrow 8) \vee (5 \rightarrow 7 \rightarrow 9))$
R7	$1 \rightarrow 3 \rightarrow 4 \rightarrow 2$
R8	$1 \rightarrow 3 \rightarrow ((4 \rightarrow 6 \rightarrow 8) \vee (5 \rightarrow 7 \rightarrow 2))$
R9	$1 \rightarrow 2$
R10	$1 \rightarrow 2$
R11	$1 \rightarrow 2$
R12	$1 \rightarrow 2$
R13	$1 \rightarrow 2$
R14	$1 \rightarrow 2$

Table 5.2: Overview of the model steps for all route options

5.2.1. Docking and undocking trajectories

The trajectory from the terminal location to the start of a swing manoeuvre is called the undocking trajectory for departing vessels. Similarly, the trajectory from the end of the swing manoeuvre to the terminal location is called the docking trajectory for arriving vessels. The docking trajectory on arrival is not a crucial part of the trajectory. At this point the vessel already sailed through the port area and made its swing manoeuvre namely, which are more important parts of the trajectory. Therefore linear interpolating the vessel position and heading from the end of the swing to the known terminal location is a sufficiently good way to model this docking trajectory. For departing vessel, this undocking

trajectory to the swing start can also be estimated using linear interpolation. In both cases the model does need to know how long the undocking or docking trajectory will take. Especially the duration for undocking trajectories is of importance, as a duration error in this beginning part of the trajectory will accumulate. This could possibly lead to a larger average displacement error and final displacement error between the predicted and true future trajectory. Figure 5.2 and Figure 5.3 show the docking speeds and undocking speeds, respectively. Given the average docking and (un)docking speed, the model can determine how long the (un)docking trajectory will take. Therefore the values in Table 5.3 are used as average speeds in the linear interpolation of docking and undocking trajectories.

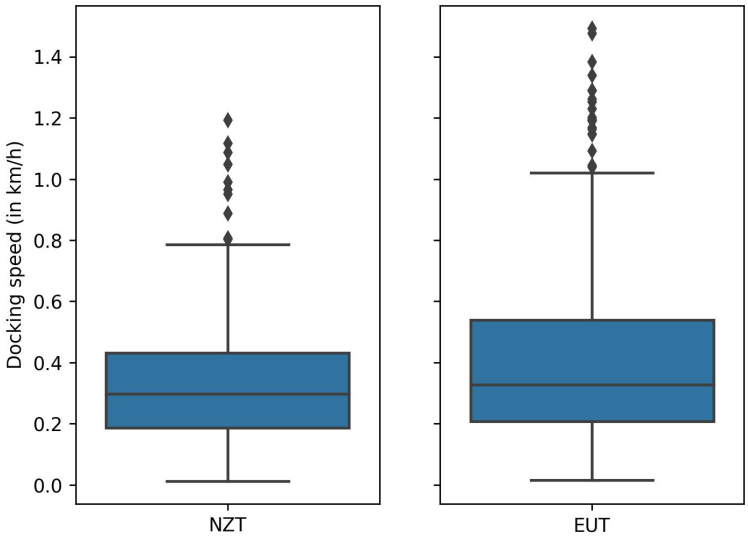

Figure 5.2: Docking speeds at both terminals

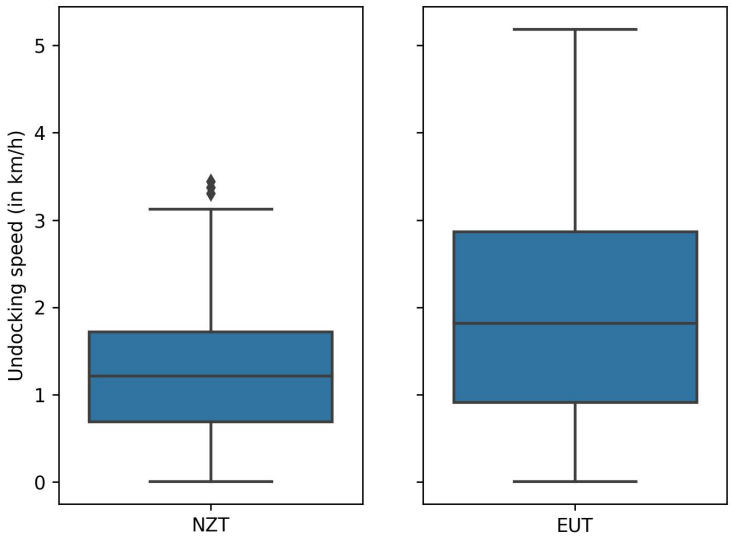

Figure 5.3: Undocking speeds at both terminals

Terminal	Trajectory type	Average speed (km/h)
EUT	Docking trajectory	0.31
NZT	Docking trajectory	0.20
EUT	Undocking trajectory	2.01
NZT	Undocking trajectory	1.30

Table 5.3: Average docking and undocking speeds at both terminals

5.2.2. Finishing swing manoeuvres

In the case a vessel is in the middle of a swing manoeuvre (block 3 in Figure 5.1), this swing manoeuvre should be completed before further parts of the trajectory are predicted. The model can complete the swing manoeuvre by extrapolating the vessel position and heading, until the heading approximately aligns with an expected value. The expected heading value depends on whether the vessel arrives at or departs from the terminal. Figure 5.4 shows all the end heading values of arriving and departing vessel at both terminals. The circular mean of all these end headings, as displayed in Table 5.4, can be used as a target heading when extrapolating swing manoeuvres. The model can stop extrapolating the swing manoeuvre once the difference between the current heading and the target heading is smaller than a constant δ. This difference can be calculated using equation 4.1. A reasonable value of δ is 20. The exact value doesn't matter too much, as trajectory prediction component (i.e. the target condition trajectory prediction model or the docking trajectory interpolation process) after the swing can handle differences in the heading.

Terminal	Route category	Circular mean of end swing heading (degrees)
NZT	Departure	293
NZT	Arrival	299
EUT	Departure	315
EUT	Arrival	343

Table 5.4: Circular mean of end swing headings for all terminal and route category combinations

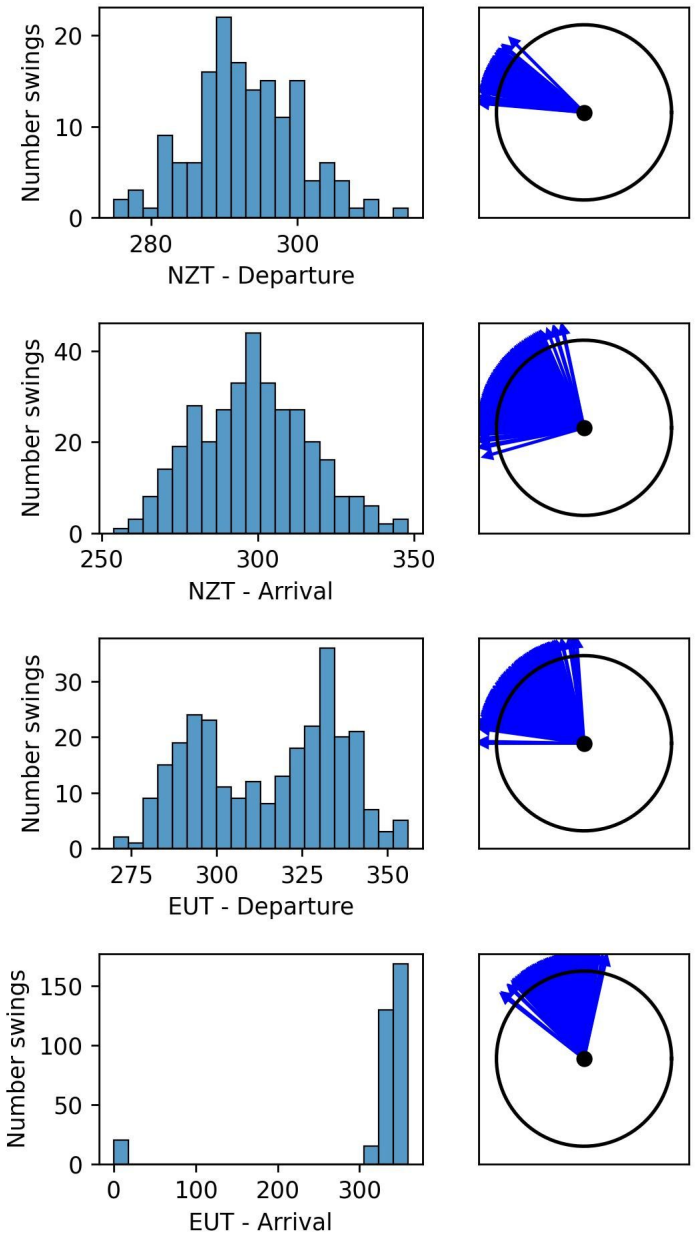

Figure 5.4: Heading at the end of a swing manoeuvre

6 Results

Chapters 3, 4, and 5 presented models for predicting short-term trajectories, including swing manoeuvres, in the Port of Antwerp-Bruges. This chapter discusses how the models were trained, using two years of vessel data, in section 6.1. The performance metrics used to evaluate the models' outputs are discussed in section 6.2, followed by the presentation of the results for the target-conditioned trajectory prediction model, swing prediction model, and integrated model in sections 6.3, 6.4, and 6.5, respectively.

6.1. Model training

The trajectories extracted from the raw vessel data as presented in Chapter 2 are split in an 80% train set and 20% test set. All model components are trained on the train set and evaluated using the test set. This section provides a detailed description of which part of the data the different models are trained on and which preprocessing steps were required.

6.1.1. Trajectory prediction training

To ensure that all models can handle variable-length inputs, they are trained on inputs of different lengths. For a trajectory consisting of n data points, $n-1$ sub-trajectories of lengths $1, 2, 3, \ldots, n-1$ are extracted. Thus, each sub-trajectory always contains the first point and differs in the number of additional historical trajectory points it contains. Sub-trajectories are constructed in this way because in a live scenario you always have all historical points available until the current time. To speed up training and limit the number of gradient steps in each iteration, the sub-trajectories are processed in batches, where each batch contains sub-trajectories of the same length.

Although one trajectory prediction model could be trained given the sub-trajectories and train-test split, it has been suggested in the literature that clustering trajectories and training different models per cluster could improve overall performance [17] [43]. Since the start and end locations of each trajectory are known, a dedicated clustering algorithm is not needed. Instead, trajectories can be clustered based on their origin and destination. A trajectory prediction model is trained for each category of trajectories (i.e., cluster) presented in Table 6.1. It is important to note that M11 and M12 have the ability to use the output of the swing location prediction model as the target location. Additionally, one big model is trained which handles all clusters of trajectories at once.

6.1.2. Swing prediction training

Similar to the process used in section 6.1.1, a sub-trajectory extraction process is utilized for the swing occurrence and swing prediction model. However, only trajectories that interact with terminals and sub-trajectories before the swing manoeuvre are used. For the swing occurrence prediction model, the data is filtered even further, as only arriving trajectories are considered. This filtering is done because swing occurrence prediction for departing vessels is straightforward, as explained in section 4.6.1. Two

Model	Origin	Destination	Trajectories
M1	ZVS	NORTH	Trajectories from ZVS to NORTH
M2	BES	NORTH	Trajectories from BES to NORTH
M3	SOUTH	NORTH	Trajectories from SOUTH to NORTH
M4	NORTH	SOUTH	Trajectories from NORTH to SOUTH
M5	NORTH	ZVS	Trajectories from NORTH to ZVS
M6	NORTH	BES	Trajectories from NORTH to BES
M7	NORTH	EUT	Arriving at EUT, no swing manoeuvre
M8	NORTH	NZT	Arriving at NZT, no swing manoeuvre
M9	EUT	NORTH	Departing from EUT, no swing manoeuvre
M10	NZT	NORTH	Departing from NZT, no swing manoeuvre
M11	NORTH	EUT	Arriving at EUT with swing manoeuvre
M12	NORTH	NZT	Arriving at NZT with swing manoeuvre
M13	EUT	NORTH	Departing from EUT with swing manoeuvre
M14	NZT	NORTH	Departing from NZT with swing manoeuvre

Table 6.1: Trajectory categories used for the target conditioned trajectory prediction models

separate models are employed for the swing occurrence prediction model, one for EUT and one for NZT. The swing location prediction model uses four different models, one for each combination of terminal and route type (i.e., arrival or departure).

6.2. Performance metrics

The most widely used performance metrics in the field of trajectory prediction are the final displacement error and average displacement error [33]. The final displacement error (FDE) is calculated as the Euclidean distance between the last point of the predicted trajectory and the last point of the true trajectory (as shown in equation 6.1). The average displacement error (ADE) is computed as the average pointwise Euclidean distance between the predicted trajectory and ground truth (as shown in equation 6.2). Since the models in this research output latitude and longitude positions, the Haversine distance between locations is used instead of the Euclidean distance to correct for the spherical surface of the earth [4]. In the remainder of this chapter, all ADE and FDE metrics are reported in meters.

$$FDE = \sqrt{(y_n - \hat{y}_n)^2} \tag{6.1}$$

$$ADE = \frac{1}{n}\sum_{i=1}^{n}\sqrt{(y_i - \hat{y}_i)^2} \tag{6.2}$$

In addition to the distance metrics, other metrics are utilized to assess the quality of trajectory predictions. The positional miss rate (PMR) is a performance metric that measures the ratio of unacceptable predictions. Here, an unacceptable prediction is defined as one that involves a positional change greater than 250 meters. Based on two years of data, no such positional changes occurred within a 30-second time frame, as illustrated in Figure 6.1. Hence, it is assumed that such changes are physically impossible. The positional miss rate performance metric has two types, namely the trajectory positional miss rate (TPMR) and point positional miss rate (PPMR). The trajectory positional miss rate measures the ratio of unacceptable trajectory predictions, where an unacceptable trajectory prediction is defined as one that contains at least one impossible positional change. The point positional miss rate, on the other hand, evaluates the proportion of unacceptable point predictions.

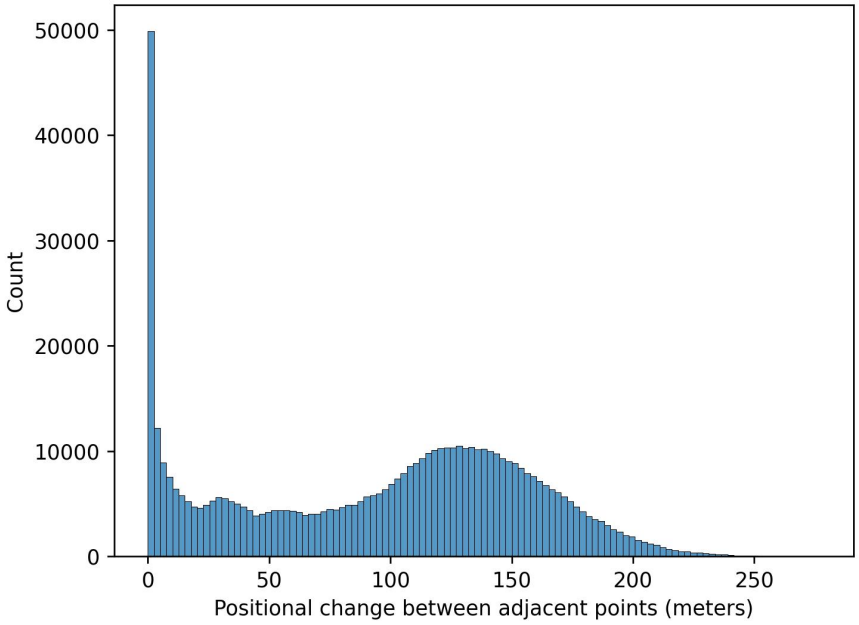

Figure 6.1: Distribution of positional change between two adjacent trajectory points

The evaluation of trajectory predictions involves another important metric known as the sailable area violation rate (SAVR), which measures the proportion of predictions that lie outside the sailable area. The sailable area is a region within the buoy lines and is visualised as the green region in Figure 6.2. Again, two variations of this metric are used for evaluation: the trajectory sailable area violation rate (TSAVR) and point sailable area violation rate (PSAVR).

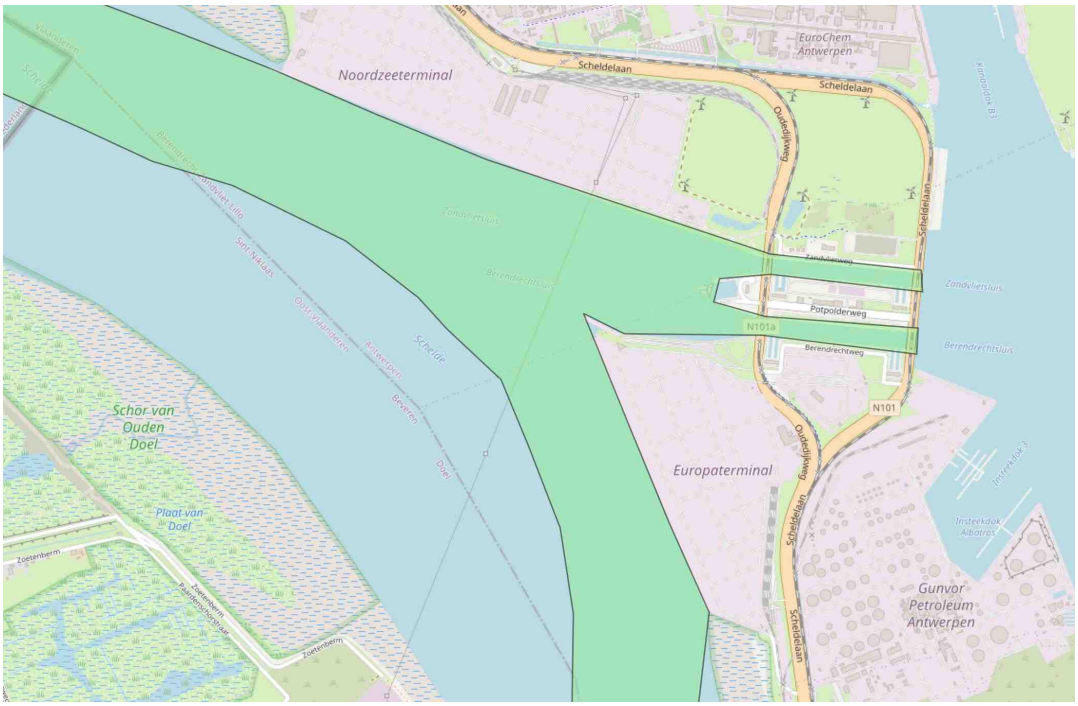

Figure 6.2: Sailable area in the scoped area of the Port of Antwerp-Bruges

6.3. Target conditioned trajectory prediction results

The performance of the target-conditioned trajectory prediction models trained on 14 different trajectory clusters is shown in Table 6.2. These results were obtained using the model outlined in section 3.4, with the configuration specified in Table 6.4. The results of training one model on all the different route options can be found in Table 6.3.

Model	Origin	Destination	Swing	ADE (m)	FDE (m)	TPMR	PPMR	TSAVR	PSAVR
M1	ZVS	NORTH	No	125	183	0.14	0.01	0.00	0.00
M2	BES	NORTH	No	130	206	0.12	0.01	0.01	0.01
M3	SOUTH	NORTH	No	99	140	0.13	0.01	0.01	0.01
M4	NORTH	SOUTH	No	78	85	0.03	0.00	0.03	0.01
M5	NORTH	ZVS	No	117	64	0.01	0.00	0.01	0.01
M6	NORTH	BES	No	155	180	0.04	0.00	0.01	0.01
M7	NORTH	EUT	No	87	74	0.01	0.00	0.01	0.01
M8	NORTH	NZT	No	76	77	0.00	0.00	0.01	0.01
M9	EUT	NORTH	No	155	114	0.22	0.02	0.02	0.01
M10	NZT	NORTH	No	128	117	0.01	0.00	0.01	0.01
M11	NORTH	EUT	Yes	149	274	0.01	0.00	0.01	0.01
M12	NORTH	NZT	Yes	99	212	0.01	0.00	0.00	0.00
M13	EUT	NORTH	Yes	128	150	0.05	0.01	0.00	0.00
M14	NZT	NORTH	Yes	65	115	0.03	0.00	0.00	0.00

Table 6.2: Results of the 14 target conditioned trajectory prediction models

Origin	Destination	Swing	ADE (m)	FDE (m)	TPMR	PPMR	TSAVR	PSAVR
ZVS	NORTH	No	115	127	0.18	0.02	0.02	0.01
BES	NORTH	No	111	131	0.20	0.02	0.02	0.01
SOUTH	NORTH	No	109	134	0.22	0.03	0.01	0.01
NORTH	SOUTH	No	100	117	0.23	0.02	0.10	0.03
NORTH	ZVS	No	182	423	0.01	0.00	0.01	0.01
NORTH	BES	No	185	422	0.00	0.00	0.01	0.01
NORTH	EUT	No	200	398	0.01	0.00	0.01	0.01
NORTH	NZT	No	187	336	0.00	0.00	0.00	0.00
EUT	NORTH	No	149	146	0.16	0.02	0.08	0.01
NZT	NORTH	No	154	216	0.01	0.00	0.03	0.02
NORTH	EUT	Yes	123	261	0.01	0.00	0.01	0.01
NORTH	NZT	Yes	150	301	0.01	0.00	0.00	0.00
EUT	NORTH	Yes	106	143	0.02	0.00	0.00	0.00
NZT	NORTH	Yes	73	108	0.00	0.00	0.00	0.00

Table 6.3: Results of using one target conditioned trajectory prediction model

Configuration	Value
Encoder input features	Latitude, longitude, heading, speed over ground
Meta input features	Tide, vessel category, length, depth, width and draught
Encoder input dimension	64
Encoder hidden dimension	64
Encoder input MLP layers	8, 16, 32
Target prediction MLP layers	64, 32, 16, 8
Decoder hidden input MLP layers	64, 64
Decoder input features	Latitude, longitude, heading, target
Decoder input dimension	64
Decoder hidden dimension	64
Decoder input MLP layers	8, 16, 32
Decoder output MLP layers	32, 16, 8
Recurrent unit type	Gated Recurrent Unit (GRU)
Optimizer	Adam
Learning rate	0.001 (Adam default)
Batch size	10000
Loss function	$c_1 = 1$, $c_2 = 1$, $c_3 = 1$, $c_4 = 0.4$

Table 6.4: Target conditioned trajectory prediction model configuration

The average displacement error and final displacement error are acceptable for all models. Moreover, the sailable area violation rates and point miss rates are near perfect for all models. However, the trajectory positional miss rate is high for some models, especially M1, M2, M3 and M9. This low point positional miss rate and high trajectory positional miss rate indicate that the point misses are not clustered in certain trajectories but instead a lot of trajectories contain a small number of these misses. M1, M2 and M9 contain trajectories in which a vessel needs to accelerate in order to leave the port. It seems that the model has some difficulty in avoiding too large accelerations.

No extremely large differences are found when comparing the results of training a separate model per cluster (Table 6.2) and the results of training one model to handle all route options (Table 6.3). The ADE, FDE, positional miss rates and sailable area violation rates are just slightly higher for most route options when using one model. Therefore, it seems that the benefits of using a different neural network per trajectory cluster is minimal.

Visualisations of a few random example trajectory predictions are presented in Figure 6.3, 6.4, and 6.5. The predicted points are shown in blue, the ground truth is shown in green and the historical trajectory is shown in black.

The results of the target conditioned trajectory prediction model, given the configuration from Table 6.4, are satisfactory. The current configuration utilises GRUs, since the training time is lowest when using these kind of units compared to LSTM units. Zhao et al. employed a comparable setup where they set the hidden state of the encoder and decoder to 128, and allowed the MLPs surrounding it to either double or halve the number of neurons in each subsequent hidden layer [47]. A similar approach is used in this study, but a hidden state of size 64 is used to reduce training time. Reasonable values are picked for the constants in the loss function. All terms have equal weight except the somewhat less important inconsistency component. It is possible that additional performance improvements could be achieved by adjusting the configuration values in Table 6.4. However, given that the primary objective of this research is to determine the feasibility of creating a trajectory prediction model, further optimization of these parameters is not pursued. Devoting a lot of resources and training time to finding better parameter configurations is a significant investment with limited return, as it would not alter the answer to the research question.

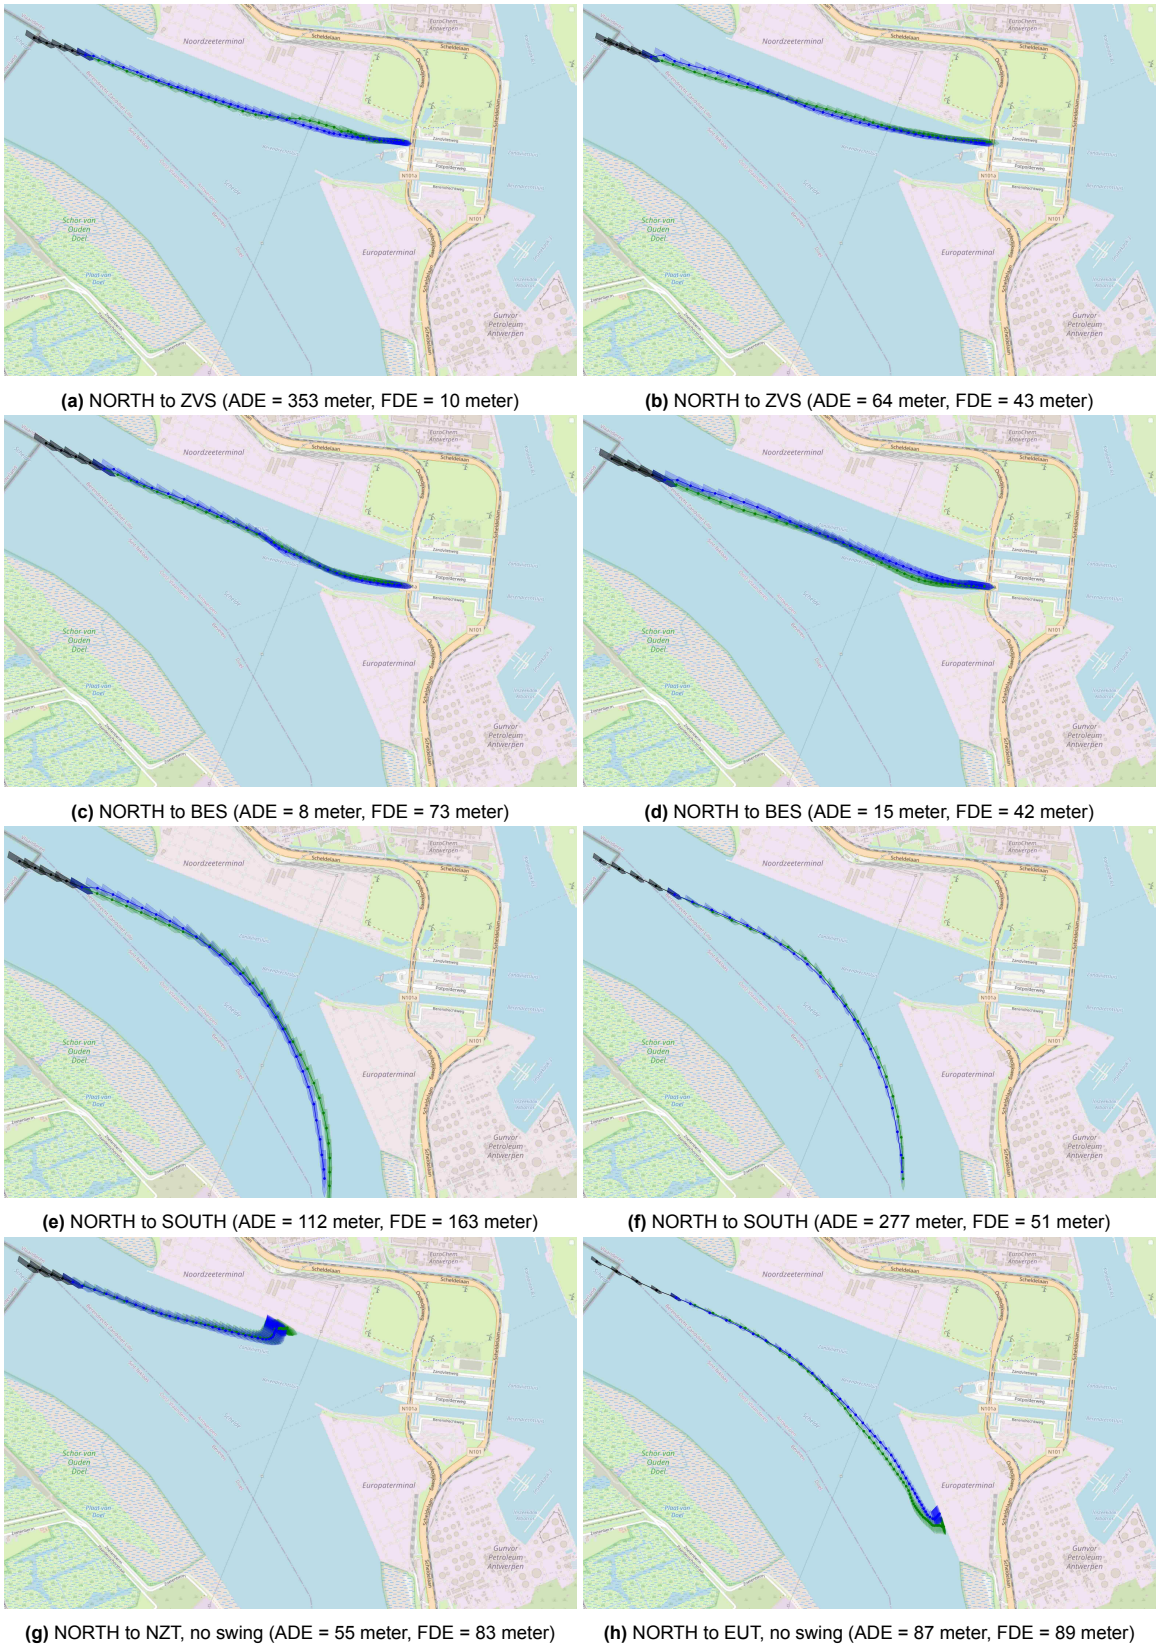

(a) NORTH to ZVS (ADE = 353 meter, FDE = 10 meter)

(b) NORTH to ZVS (ADE = 64 meter, FDE = 43 meter)

(c) NORTH to BES (ADE = 8 meter, FDE = 73 meter)

(d) NORTH to BES (ADE = 15 meter, FDE = 42 meter)

(e) NORTH to SOUTH (ADE = 112 meter, FDE = 163 meter)

(f) NORTH to SOUTH (ADE = 277 meter, FDE = 51 meter)

(g) NORTH to NZT, no swing (ADE = 55 meter, FDE = 83 meter)

(h) NORTH to EUT, no swing (ADE = 87 meter, FDE = 89 meter)

Figure 6.3: Several arrival example trajectory predictions using the target conditioned trajectory prediction model. Blue points indicate the predicted trajectory, green points the ground truth and black points the historic trajectory.

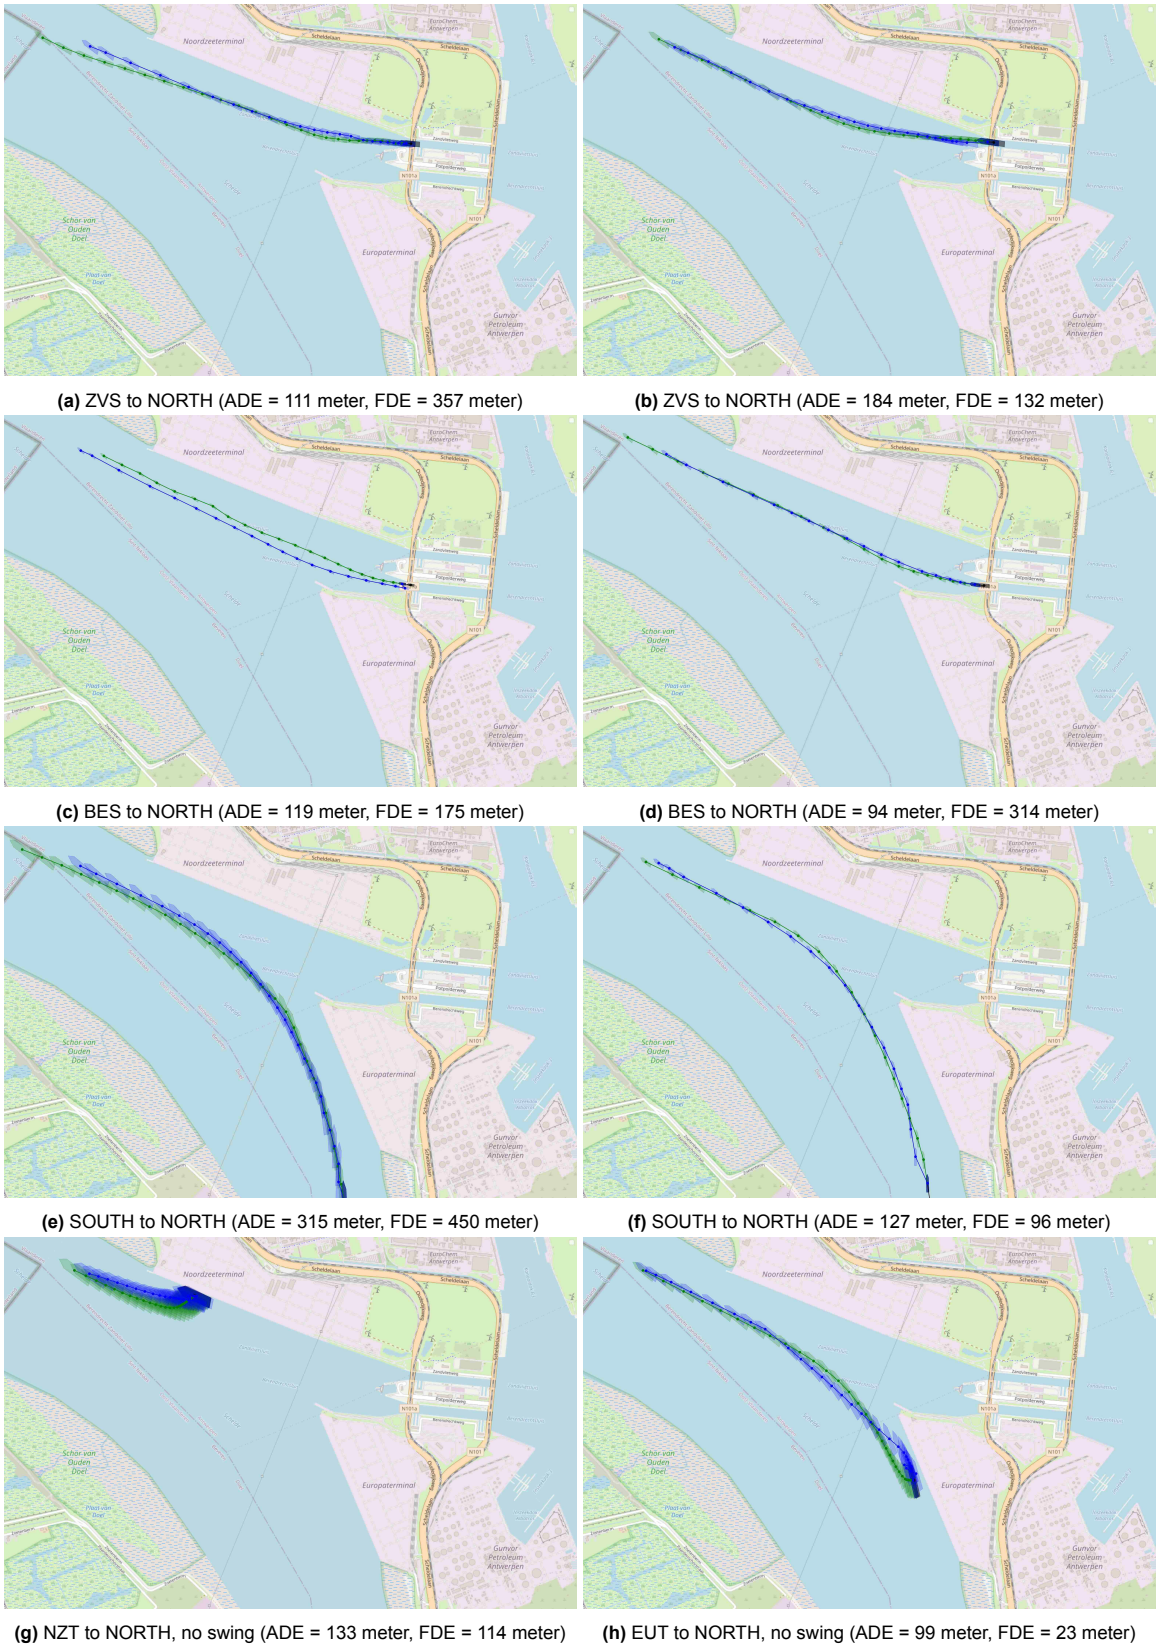

(a) ZVS to NORTH (ADE = 111 meter, FDE = 357 meter)

(b) ZVS to NORTH (ADE = 184 meter, FDE = 132 meter)

(c) BES to NORTH (ADE = 119 meter, FDE = 175 meter)

(d) BES to NORTH (ADE = 94 meter, FDE = 314 meter)

(e) SOUTH to NORTH (ADE = 315 meter, FDE = 450 meter)

(f) SOUTH to NORTH (ADE = 127 meter, FDE = 96 meter)

(g) NZT to NORTH, no swing (ADE = 133 meter, FDE = 114 meter)

(h) EUT to NORTH, no swing (ADE = 99 meter, FDE = 23 meter)

Figure 6.4: Several departing example trajectory predictions using the target conditioned trajectory prediction model. Blue points indicate the predicted trajectory, green points the ground truth and black points the historic trajectory.

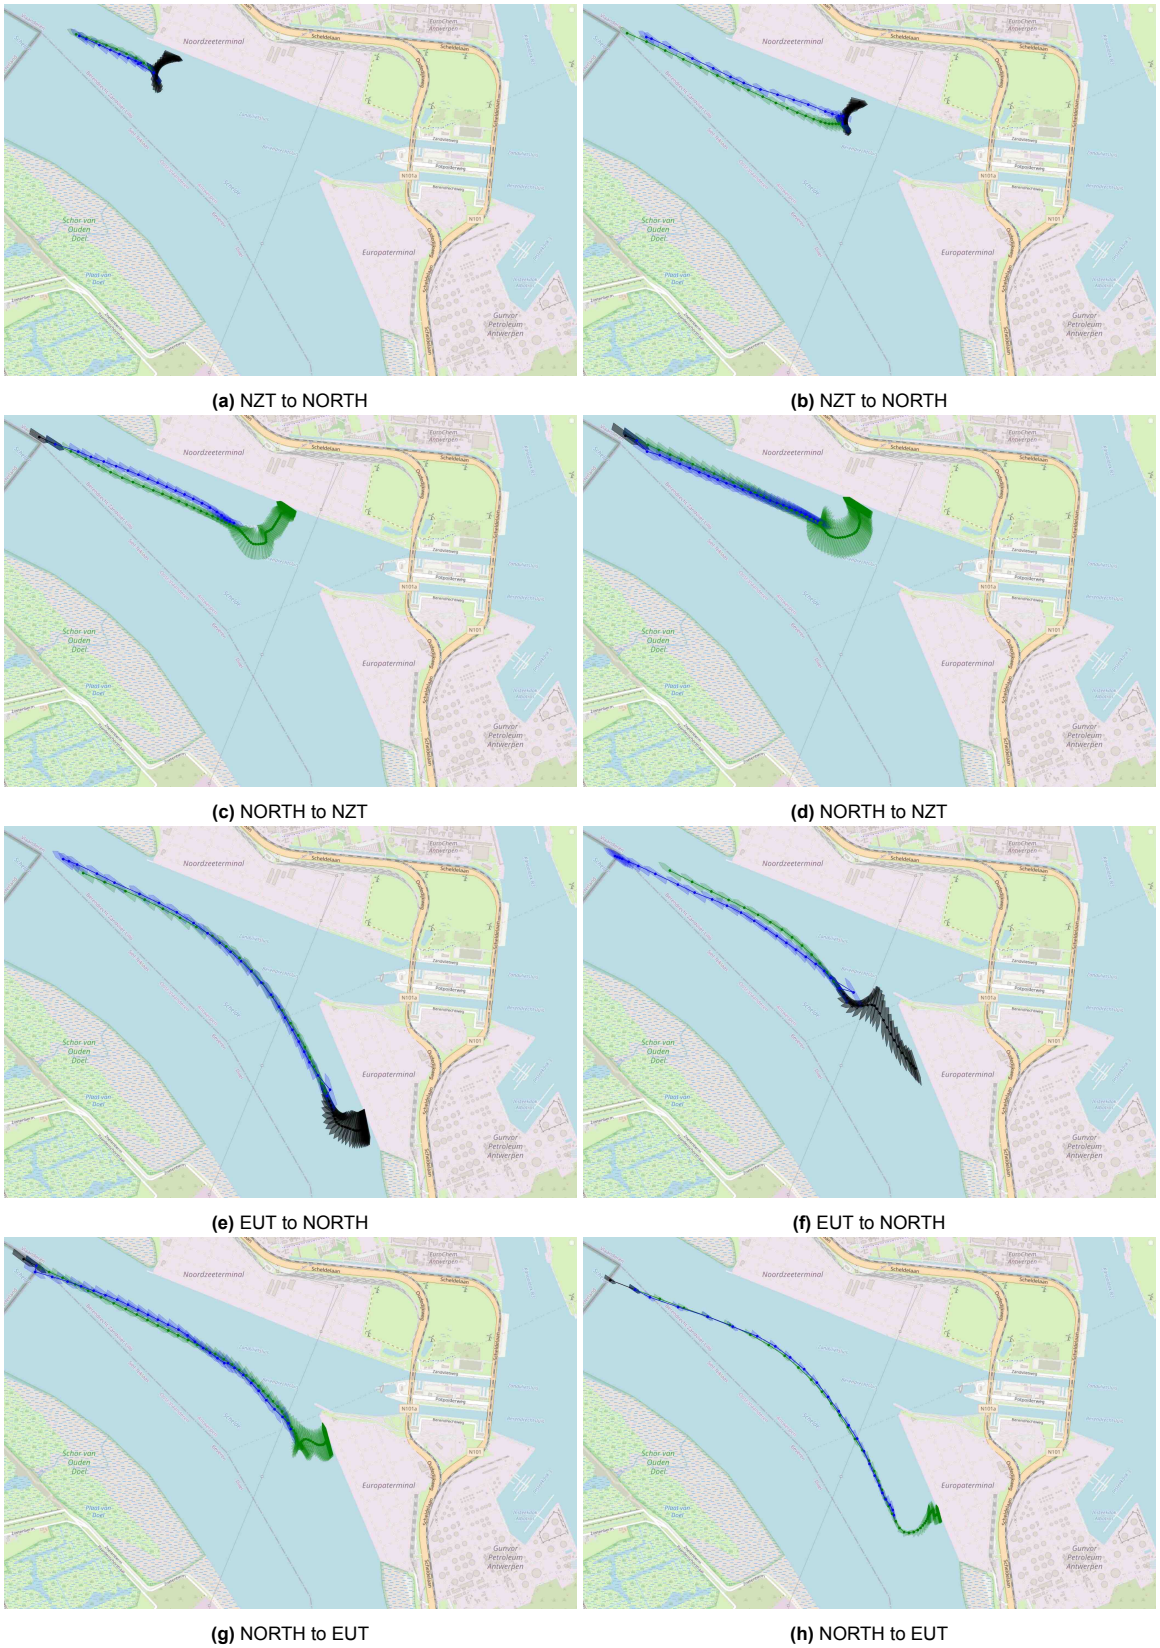

(a) NZT to NORTH **(b)** NZT to NORTH

(c) NORTH to NZT **(d)** NORTH to NZT

(e) EUT to NORTH **(f)** EUT to NORTH

(g) NORTH to EUT **(h)** NORTH to EUT

Figure 6.5: Several example trajectory predictions using the target conditioned trajectory prediction model before or after a swing manoeuvre. (a), (b), (e) and (f) show a trajectory prediction after a swing manoeuvre. (c), (d), (g) and (h) visualise the trajectory prediction to the swing start location. Blue points indicate the predicted trajectory, green points the ground truth and black points the historic trajectory.

6.4. Swing prediction results

The encoder structure and configuration of the swing occurrence prediction model are the same as used in the target conditioned trajectory prediction model, which can be found in Table 6.4. The final hidden state of the encoder has a length of 64, and the multilayer perceptron that produces the swing probability has hidden layers of sizes 32, 16, and 8, respectively. Table 6.5 shows the accuracy results of the swing occurrence prediction model for arriving vessels. No model is necessary for departing vessels, as their heading alone indicates whether the vessel will swing or not.

Terminal (arrival)	Accuracy
NZT	0.84
EUT	0.83

Table 6.5: Accuracy results of the swing occurrence prediction model

The swing location prediction model again uses the same encoder structure and configuration as the target conditioned trajectory prediction model. However, it uses a multilayer perceptron with hidden layers of sizes 64, 32, and 16 to embed the encoder hidden state of dimension 64 to the output layer consisting of nine neurons (as explained in section 4.6.2). The outcomes of this model can be found in Table 6.6.

Terminal	Route type	Positional error (m)		Heading error (degrees)		Duration error (min.)
		Start	End	Start	End	
NZT	Arrival	237	216	7.4	12.8	1.4
NZT	Departure	132	140	15.3	6.7	1.9
EUT	Arrival	254	217	9.9	15.7	1.7
EUT	Departure	199	301	9.6	14.2	1.5

Table 6.6: Results of the swing location prediction model. Positional errors are in meters, heading errors in degrees and duration errors in minutes.

6.5. Full model results

The previous two sections provided separate evaluations of the individual components of the complete model. This section presents the evaluation of the full model, as described in 3.4.2. Four metrics are used to compare a complete trajectory prediction with the ground truth. The first metric is the duration error, which is simply the time difference between the predicted and actual trajectory in minutes. Additionally, the final displacement error (FDE) and sailable area violation rate (SAVR) can be used as metrics. The average displacement error (ADE) can also be used but requires some adjustment since the length of the predicted and actual trajectory can differ. To calculate the ADE, the length of the ground truth trajectory is used to compare against the predicted trajectory. If the predicted trajectory is longer, some points are not considered in the metric. On the other hand, if the predicted trajectory is shorter, the last predicted point is used multiple times.

Using the given performance metrics, Table 6.7 presents the results of the complete model. The results are categorized into trajectories interacting with terminals and trajectories just passing through the scoped area.

Terminal interaction	Duration error (min.)	FDE (m)	ADE (m)	TPMR	PPMR	TSAVR	PSAVR
No	1.3	102	149	0.01	0.01	0.01	0.01
Yes	11.0	194	146	0.06	0.01	0.25	0.01
Combined	3.2	117	147	0.05	0.01	0.06	0.01

Table 6.7: Results of complete model

A large difference in duration error is visible between the two categories of trajectories. The duration errors for the terminal trajectories are mainly made in the undocking and docking portion of the trajectory.

These parts of the trajectory take relatively long, meaning there is room for large duration errors. The model currently uses an average for the (un)docking speed. However, since these speeds seem to vary for each trajectory it leads to the large average duration error. For docking trajectories the impact of these duration errors is minimal, since the crucial part of the trajectory is already done. However, for the trajectories leaving the terminal, this duration error has accumulating impact resulting in an inaccurate view of the situation in the future. The combined duration error is quite low, because there are more non terminal trajectories in the test set. For both categories the point positional miss rate and point sailable area violation rate are acceptable. Interesting is to see that the small portion of sailable area violation errors for terminal trajectories are spread over a lot of trajectories, the TSAVR is namely 0.25.

While an average FDE of 117 and average ADE of 147 may appear high, these predictions seem acceptable on a complete map, particularly given the average vessel length of 227 meters. To put these numbers in a bit of context, Figure 6.3 and Figure 6.4 report the ADE and FDE of some plotted trajectories. The terminal trajectories show a slightly higher final displacement error since the final location at the terminal is harder to predict than the final location of a simple trajectory going to, for example, the Zandvliet Lock. The port is satisfied with the the presented results of the model.

Figure 6.7 presents a visualisation of several random examples of complete trajectory predictions. Additionally, to provide more context on the quality of the predictions, Figure 6.8 and Figure 6.8 show a distribution of the average displacement error of the trajectories per destination. For this calculation all trajectories in the test set are evaluated using one historic point, i.e. the most difficult case. Moreover, Figure 6.10 and 6.11 show the predictions with the highest average displacement error for each route option in this same scenario.

Interesting to analyse is how the performance of the complete trajectory prediction models evolves when more historic data points are used. Figure 6.6 visualises the ADE and FDE when using a different amount of historical data points. For all trajectories, the average displacement error seems to decrease when more historical data points are used. The final displacement error doesn't seem to relate too much to the number of historical data points for the trajectories not interacting with the terminals. However, for trajectories which do interact with the terminals, the number of historic data points does seem to make an impact on the final displacement error. Intuitively this makes sense, since for arriving vessels more historic data points give a better idea on where the final terminal location will be.

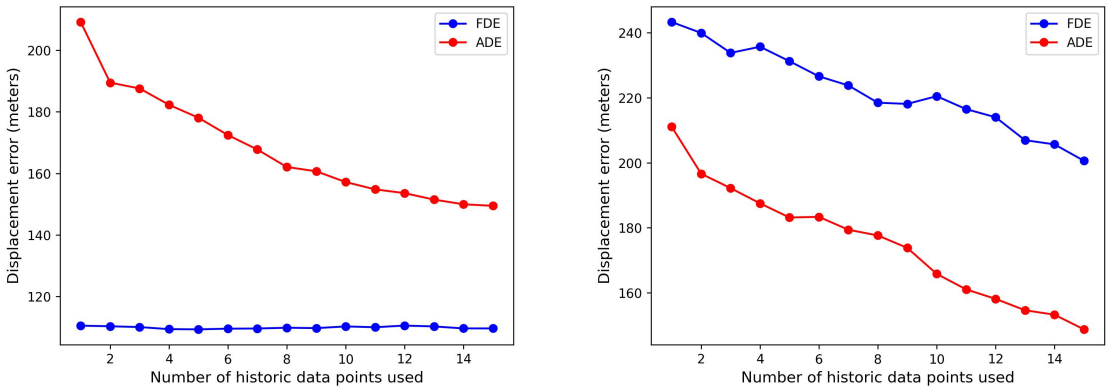

(a) Trajectories not interacting with a terminal

(b) Trajectories interacting with NZT or EUT

Figure 6.6: Performance of complete model for different number of historical data points used

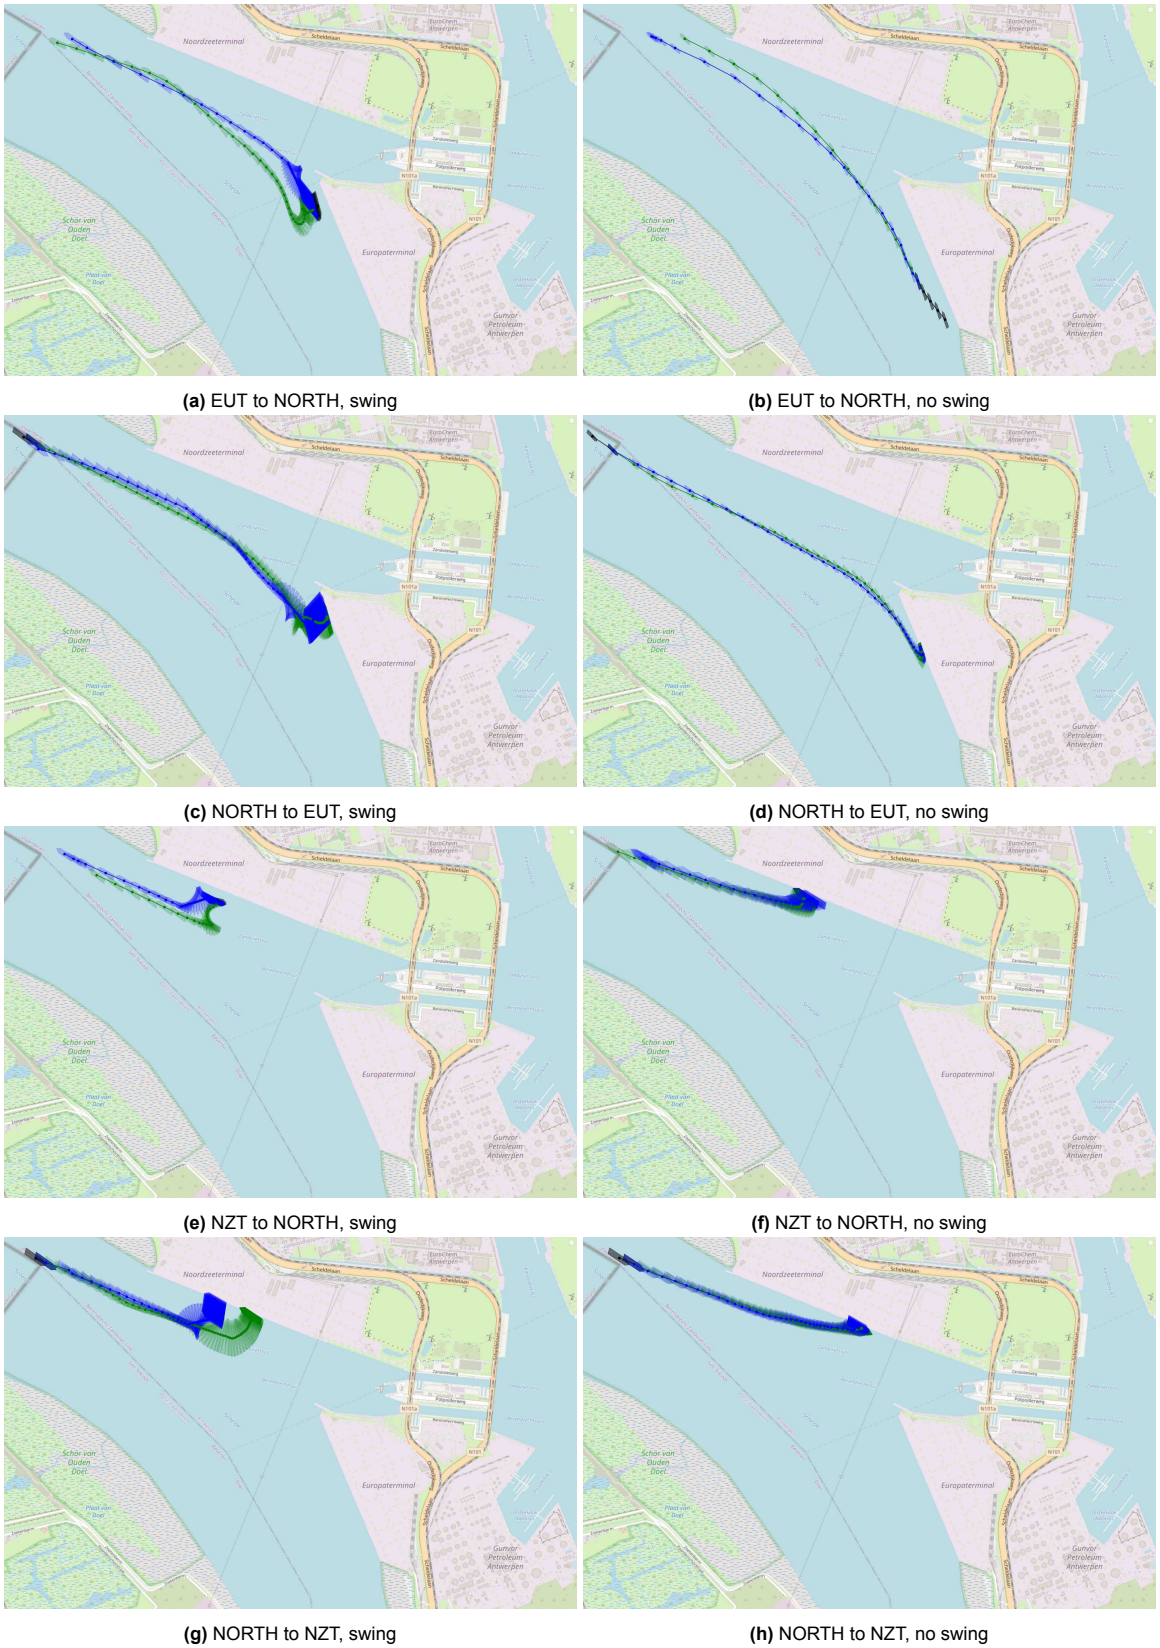

(a) EUT to NORTH, swing

(b) EUT to NORTH, no swing

(c) NORTH to EUT, swing

(d) NORTH to EUT, no swing

(e) NZT to NORTH, swing

(f) NZT to NORTH, no swing

(g) NORTH to NZT, swing

(h) NORTH to NZT, no swing

Figure 6.7: Several complete trajectory predictions example. Blue points indicate the predicted trajectory, green points the ground truth and black points the historic trajectory.

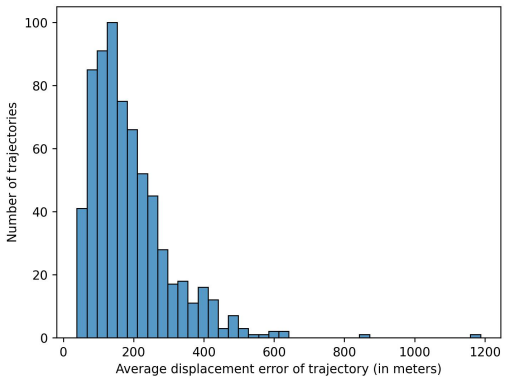

(a) ZVS to NORTH

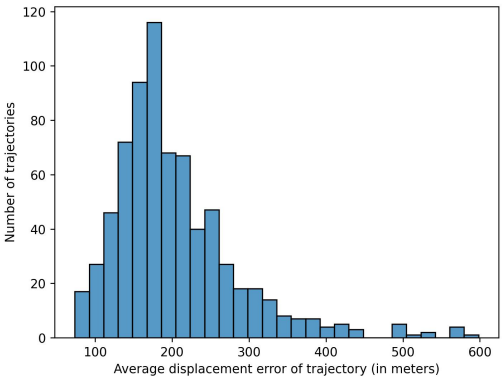

(b) NORTH to ZVS

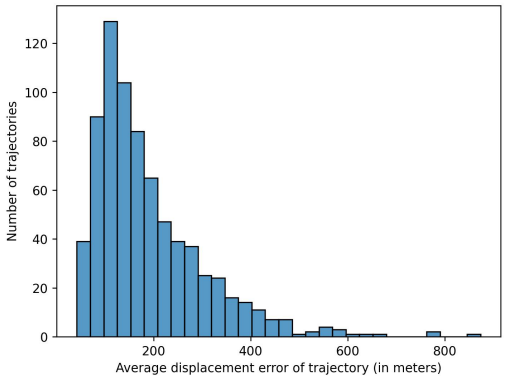

(c) BES to NORTH

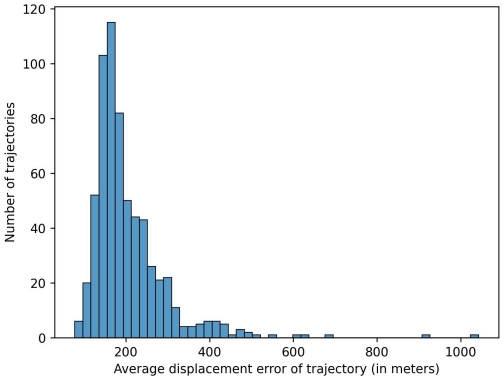

(d) NORTH to BES

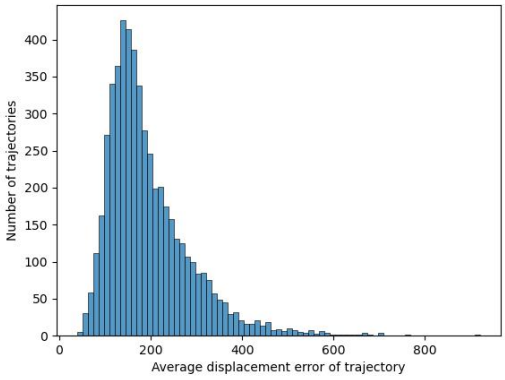

(e) SOUTH to NORTH

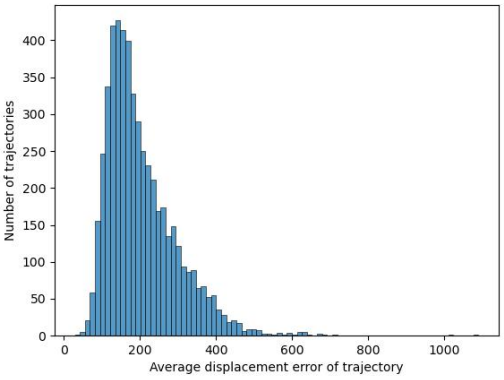

(f) NORTH to SOUTH

Figure 6.8: The distribution of average displacement errors per route option which does not involve any terminal interaction. All trajectories in the test set are evaluated using one historic data point.

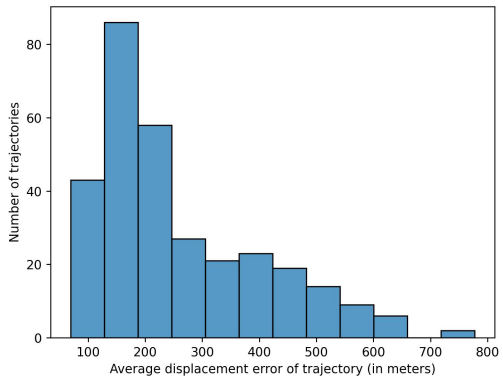

(a) NORTH to EUT, no swing

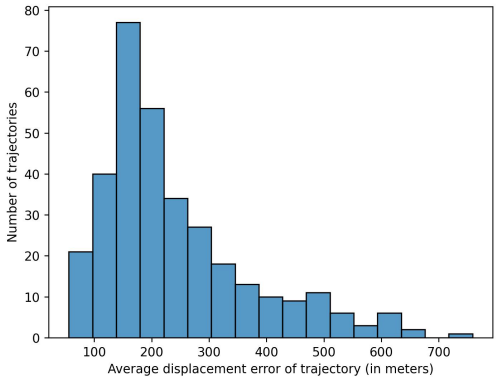

(b) NORTH to EUT, swing

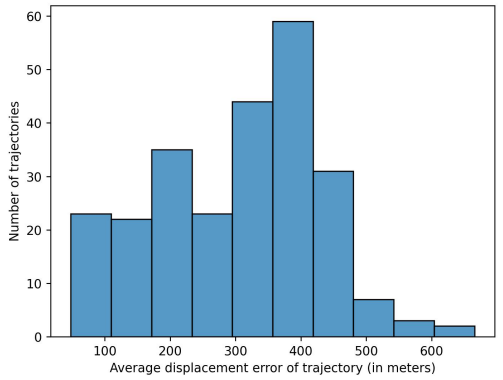

(c) NORTH to NZT, no swing

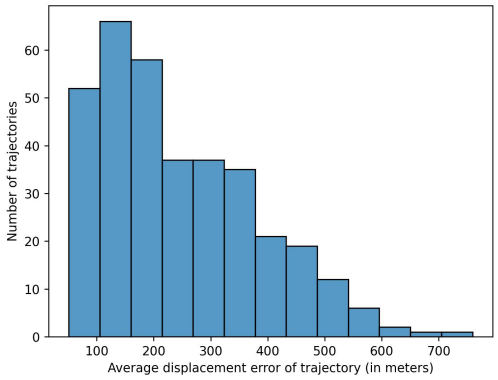

(d) NORTH to NZT, swing

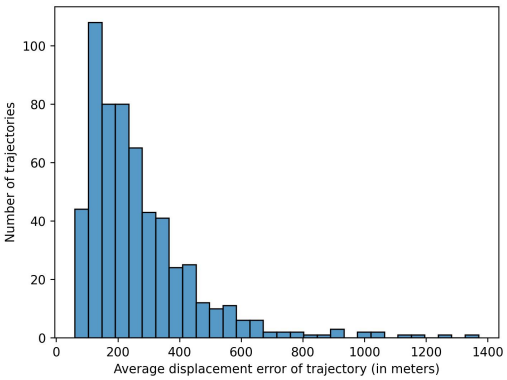

(e) EUT to NORTH, no swing

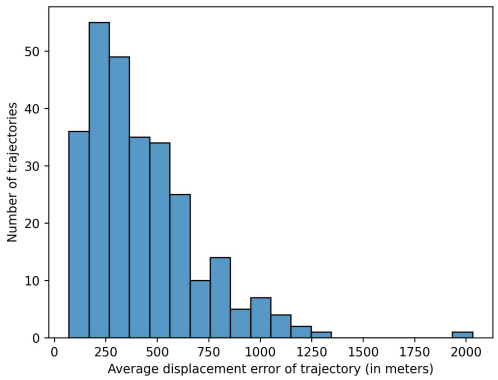

(f) EUT to NORTH, swing

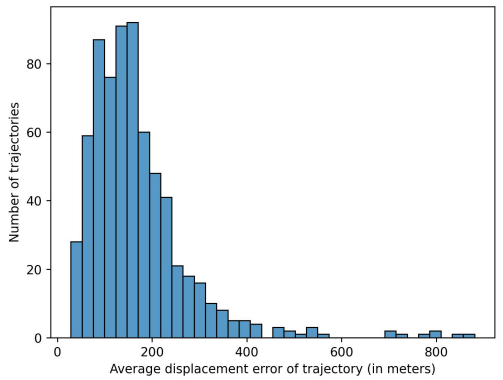

(g) NZT to NORTH, no swing

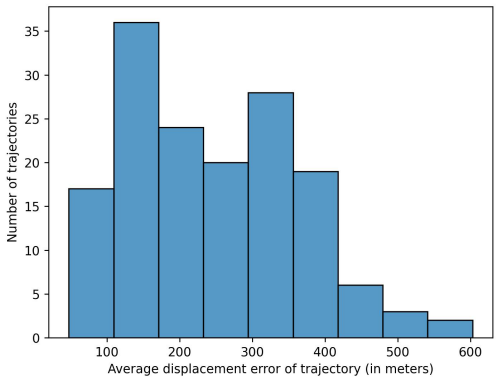

(h) NZT to NORTH, swing

Figure 6.9: The distribution of average displacement errors per route option which involve a terminal interaction. All trajectories in the test set are evaluated using one historic data point.

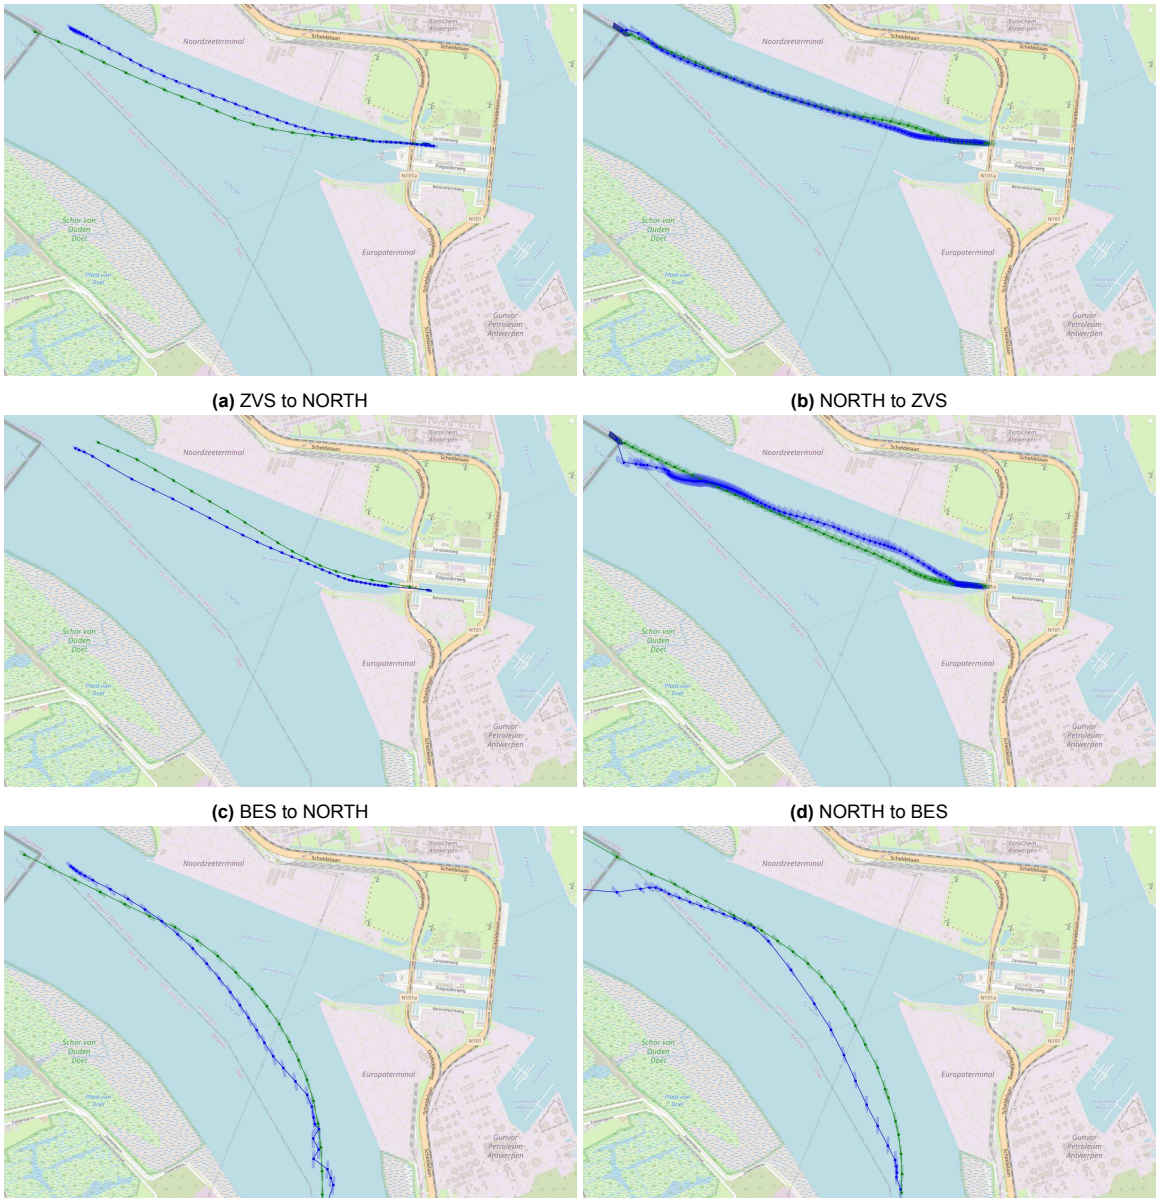

(a) ZVS to NORTH **(b)** NORTH to ZVS

(c) BES to NORTH **(d)** NORTH to BES

(e) SOUTH to NORTH **(f)** NORTH to SOUTH

Figure 6.10: The worst trajectory predictions, in terms of average displacement error, for each route option not interacting with a terminal. Blue points indicate the predicted trajectory, green points the ground truth and black points the historic trajectory.

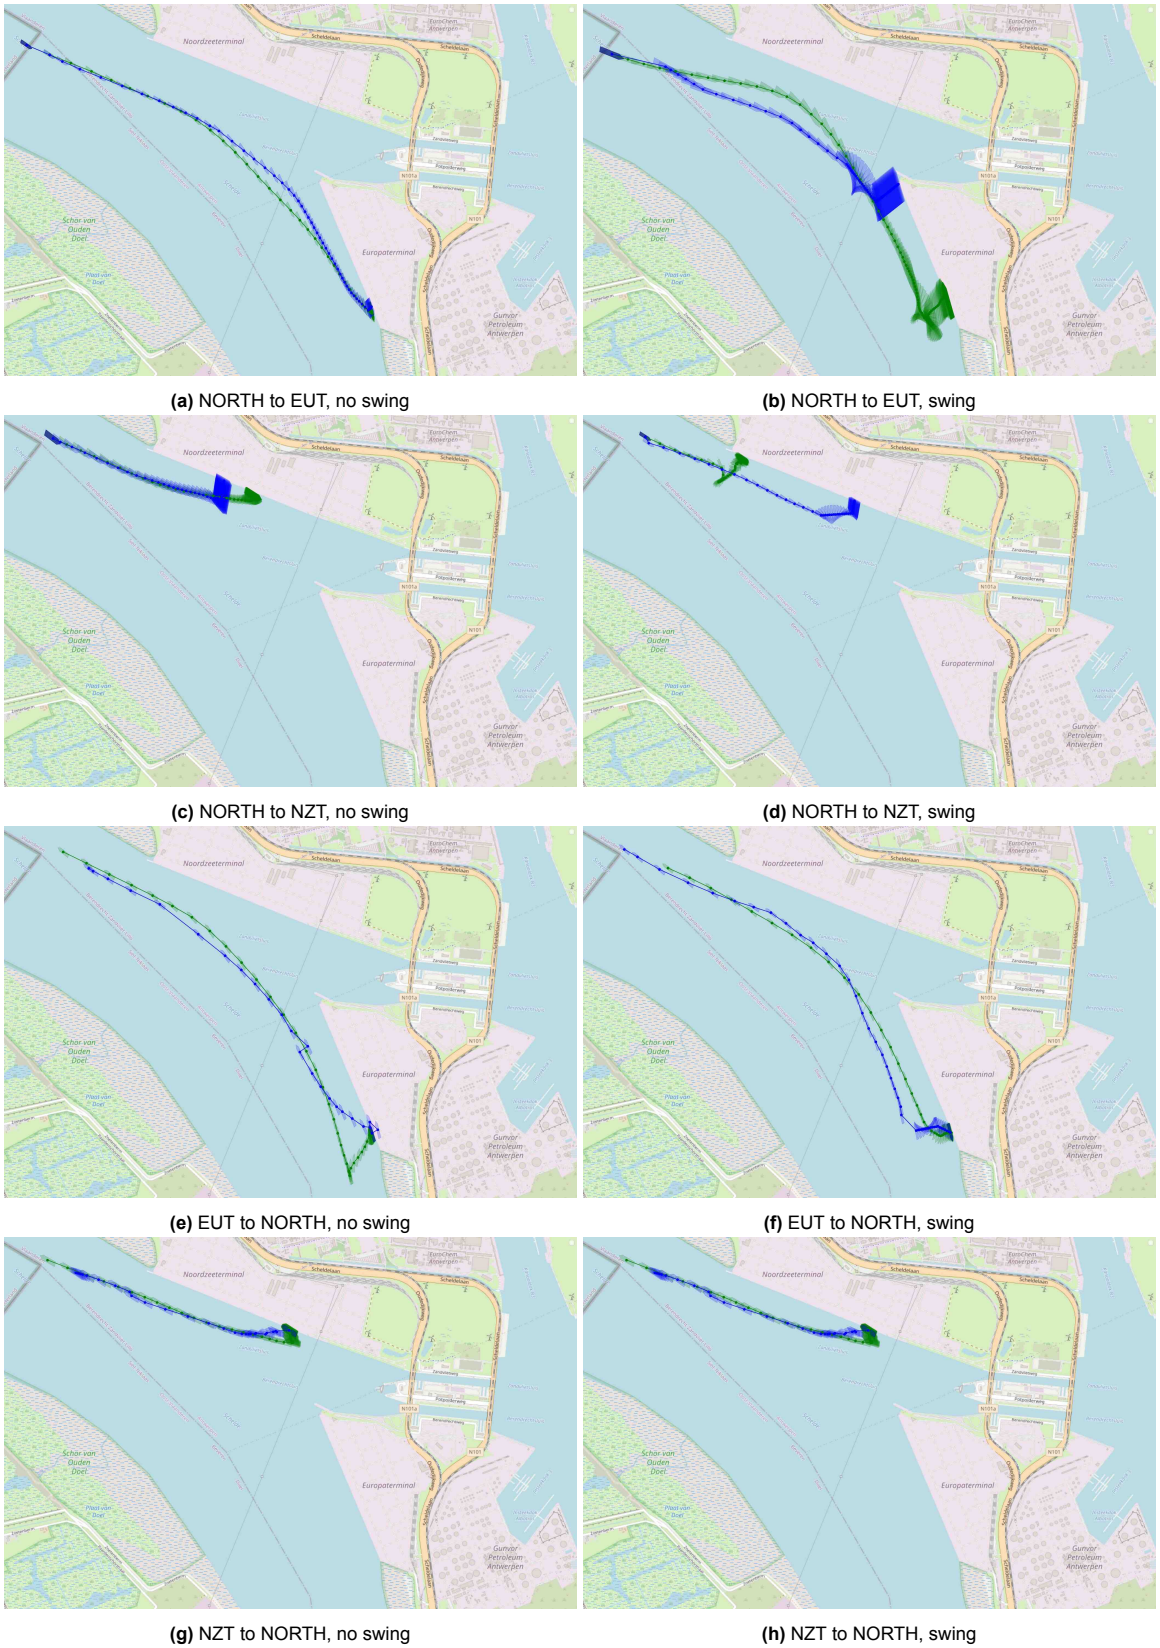

(a) NORTH to EUT, no swing

(b) NORTH to EUT, swing

(c) NORTH to NZT, no swing

(d) NORTH to NZT, swing

(e) EUT to NORTH, no swing

(f) EUT to NORTH, swing

(g) NZT to NORTH, no swing

(h) NZT to NORTH, swing

Figure 6.11: The worst trajectory predictions, in terms of average displacement error, for each route option interacting with a terminal. Blue points indicate the predicted trajectory, green points the ground truth and black points the historic trajectory.

7

Conclusion

This final chapter presents a comprehensive overview of the work done throughout the project, highlighting the main findings and conclusions. In section 7.1, a summary of the key aspects of the research is provided, including the research questions, methodology, and results.

Following this, in section 7.2, the main conclusions and answer to the research question are presented. This section provides an overview of the research findings and identifies the most significant contributions made by the study.

Finally, in section 7.3, the research is reflected upon and suggestions for future directions are given. This section discusses the limitations and strengths of the research, along with potential areas for future investigation. It also highlights the impact of implementation of the developed models in real-world scenarios.

7.1. Summary

Vessel safety and navigation are essential components of maritime operations. Proper navigation practices and safety procedures are critical to ensure that vessels and their crews reach their destinations safely without causing significant congestion. The maritime industry continues to invest in technology to improve vessel safety and navigation practices. This study takes the initial step towards the integration of an artificial intelligence tool that can aid maritime traffic controllers at the Port of Antwerp-Bruges in their work. The study is conducted as a case study at the Port of Antwerp-Bruges and focuses only on a scoped area of the port. The area under consideration in this study has the characteristic of a bottleneck, as it serves as a connection point between the entire port and the North Sea. Moreover, it attracts a lot of traffic due to the presence of two terminals and two locks. The goal of this study is to examine the feasibility of developing a model that can predict vessel position in a short time horizon in this scoped area of the port. Such a complete model for vessel trajectory prediction should have the ability to predict trajectories while taking swing manoeuvres into account. Swing manoeuvres are U-turn manoeuvres executed by container vessels upon arrival or departure at a terminal. Due to the presence of the Europa Terminal (EUT) and Noordzee Terminal (NZT) in the scoped area, swing manoeuvres play a significant role. The Port of Antwerp-Bruges provided two years of historical vessel trajectory data. The structure of this data along with relevant preprocessing steps are examined in Chapter 2. The problem of trajectory prediction is formulated in Chapter 3, which also delves in the historical method to utilize the historical vessel data to predict trajectories. From the literature study it becomes apparent that deep learning-based approach are popular methods to use on trajectory prediction related problems. Therefore, a target conditioned trajectory prediction model is presented, which takes a variable number of past positions of a vessel and predicts a future trajectory. This model uses an encoder-decoder structure, where the encoder is a recurrent neural network which encodes the historic trajectory and the decoder is again a recurrent neural network which generates a future trajectory based on this encoding. As an intermediate step, the model predicts a target location for the future trajectory. This target location is passed to the decoder along with the encoding of the historic trajectory. The model uses the distance between the predicted trajectory and actual future trajectory as loss

function. Additionally, penalty terms are added to the loss function to smooth the prediction. Multiple trajectory prediction models are trained on different clusters of trajectories. These models obtain a final displacement error of 142 meters and average displacement error of 91 meters on average.

However, this target conditioned trajectory prediction model is not sufficient, as it doesn't consider these crucial swing manoeuvres. Chapter 4 delves into the definition, detection and extraction of swing manoeuvres. Next, this chapter presents an approach to predict swing manoeuvres. The prediction is split in two steps. First of all, the occurrence of a swing manoeuvre is predicted and secondly the start and end location of the manoeuvre are predicted along with the duration. The heading and position of the vessel are interpolated in between the predicted start and end position. Both the swing occurrence and swing location prediction model use an encoder recurrent neural network to encode the historical trajectory in a vector embedding. A multilayer perceptron is used to convert this vector embedding to a swing occurrence or swing location prediction. Predicting swing manoeuvres for departing vessels is trivial, so the model only considers arriving vessels. On these arriving vessels the model obtains an accuracy of 84% at the Noordzee Terminal and 83% at the Europa Terminal. The swing location prediction model predicts the start of a swing with an error of 206 meters and the end of a swing with an error of 219 meters on average.

Given the target conditioned trajectory prediction model and the swing manoeuvre prediction model, Chapter 5 presents an integration of the two models. The resulting complete model can handle the prediction of complete trajectories, including swing manoeuvres.

7.2. Conclusions

The main goal of this research was to explore the feasibility of constructing a model that could accurately predict short-term vessel trajectories in a specific area of the Port of Antwerp-Bruges. The study identified and investigated the components of such a model and developed a complete model that achieved an average displacement error of 149 meters on trajectories interacting with terminals and 146 meters on trajectories that did not interact with terminals. The findings indicate that it is possible to construct a model that can accurately predict short-term vessel trajectories using swing manoeuvre prediction and target conditioned trajectory prediction as subroutines.

Notably, the study highlights that swing manoeuvres are a crucial component of vessel trajectory prediction. Swing manoeuvres are modelled by predicting their start position, end position, and duration. The research shows that while the model can predict these factors accurately, there are relatively large duration errors during the undocking and docking phases of terminal trajectories.

Additionally, the study found that training several models per trajectory cluster did not significantly impact overall model performance compared to using one large model for all trajectories. This finding could have practical implications for the future scalability of the model to cover larger areas of the port, where the number of available route options could increase.

This research is considered to be an initial pilot for the Port of Antwerp-Bruges. The port has been satisfied with the obtained results so far and now the next step in the pilot is to integrate the model with their live data sources.

The primary contribution of this research is the addition of the swing manoeuvre extension to the trajectory prediction model and the application of this model in a real-world setting. The model shows that incorporating swing manoeuvres can lead to accurate short-term trajectory predictions.

7.3. Discussion and future work

Even though the predicted trajectories are of sufficient quality, additional performance could be obtained by incorporating certain changes.

First of all, wind direction and strength are not incorporated into the analysis; including these factors may be beneficial. Based on domain knowledge, these factors are namely believed to be linked with the trajectory path and the location of the swing manoeuvre. Additional accuracy could probably also be obtained by using the exact planned final location of the vessel. Currently the terminal number is used as the destination for vessels arriving at terminals. However, the mooring bollards that the vessel will use are also known in advance. These bollards give a more accurate indication of the destination

of the vessel along the terminal.

Another key aspect which is omitted in this research are interaction effects. Trajectory, swing occurrence and swing location prediction are currently all based on solely the historical trajectory and meta data factors. However, in reality a swing decision or the sailing course could be affected by the current or even future traffic condition. If the traffic keeps increasing in the Port of Antwerp-Bruges, taking interactions effects into account becomes even more important. The decision to not take interaction effects into account in this study was based on the limited presence of vessel interaction in the given data. A suggestion for future research would be to use graph neural networks (GNN) to include interaction effects once more of these effects are visible in the data.

The miss rate and sailable area compliance rate are two metrics used to evaluate the performance of predicted trajectories. Given the fact that a good output has a low number of misses and number of points outside the sailable area, various correction layers could be built to address these points. An example correction layer could simply map a point outside the sailable area to the closest point within the sailable area polygon.

Furthermore, several data preprocessing steps could potentially reduce the complexity of the prediction and therefore improve performance. A suggestion could be to use the position of the vessel at the last historical data point as the origin of the coordinate system for representing trajectories. Similarly, the heading of the vessel could be represented such that the heading at the last data point aligns with the x-axis.

Lastly, important to mention is that extensive parameter and architectural configuration optimisation is left out of this study. Additional performance can possibly be gained by tweaking the neural network architectures or other parameters listed in Table 6.4. The search for better parameters is omitted, as the goal of this thesis was to show the feasibility and method for creating a model which could predict short term trajectories including swing manoeuvres rather than creating the most optimal model. Further research could improve on the models presented in this thesis. An interesting suggestion would be to retrain the models using LSTM units instead of GRU. LSTMs could perform better on datasets with long sequences.

7.3.1. Model generalization

The model as presented in this thesis is created with the scoped area of the Port of Antwerp-Bruges in mind. Therefore, it contains some assumptions which are specific to the use case in this area. The exact model cannot not be transferred directly to other uses cases. However, the main ideas behind the models can be used on other port locations. To facilitate the adaptation to other use cases, one has to identify the possible route options first. If the possible route options are limited, as is the case in the scoped area under study, a similar approach of training a model per route option can be used (see section 5.1). However, if the number of possible route origin and route destination combinations is too large, a single model can be used. Probably a bit of performance is then traded off against the ability to handle all route options with one model. Moreover, the swing detection and extraction process contains some constants which are partly use case specific. Equation 4.2 uses a heading difference of 125 degrees to distinguish between swing manoeuvres and normal trajectories. This value is based on historical data of the scoped area and could be different for other ports. Moreover, it could be that the swing detection and extraction process should be restricted even more in other use cases. Sudden large turns could namely be identified as swing manoeuvres by the current detection and extraction logic. A straightforward way to restrict swing detection is to define only certain areas in the port where swing manoeuvres are expected.

References

[1] Alexandre Alahi, Vignesh Ramanathan, and Li Fei-Fei. “Socially-aware large-scale crowd forecasting”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2014, pp. 2203–2210.

[2] Alexandre Alahi et al. “Social lstm: Human trajectory prediction in crowded spaces”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 961–971.

[3] Danial Alizadeh, Ali Asghar Alesheikh, and Mohammad Sharif. “Vessel trajectory prediction using historical automatic identification system data”. In: *The Journal of Navigation* 74.1 (2021), pp. 156–174.

[4] Huseyin Alkan and Hasari Celebi. “The implementation of positioning system with trilateration of haversine distance”. In: *2019 IEEE 30th annual international symposium on personal, indoor and mobile radio communications (PIMRC)*. IEEE. 2019, pp. 1–6.

[5] Port Of Antwerp-Bruges. *Booklet with facts and figures of 2021 (Antwerp branch)*. 2022. URL: `https://www.portofantwerpbruges.com` (visited on 04/04/2022).

[6] Holger Caesar et al. “nuscenes: A multimodal dataset for autonomous driving”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2020, pp. 11621–11631.

[7] Roberto Cahuantzi, Xinye Chen, and Stefan Güttel. “A comparison of LSTM and GRU networks for learning symbolic sequences”. In: *arXiv preprint arXiv:2107.02248* (2021).

[8] Samuele Capobianco et al. “Deep learning methods for vessel trajectory prediction based on recurrent neural networks”. In: *IEEE Transactions on Aerospace and Electronic Systems* 57.6 (2021), pp. 4329–4346.

[9] Sergio Casas, Wenjie Luo, and Raquel Urtasun. “Intentnet: Learning to predict intention from raw sensor data”. In: *Conference on Robot Learning*. PMLR. 2018, pp. 947–956.

[10] Yuning Chai et al. “Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction”. In: *arXiv preprint arXiv:1910.05449* (2019).

[11] Ming-Fang Chang et al. “Argoverse: 3d tracking and forecasting with rich maps”. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2019, pp. 8748–8757.

[12] Chiho Choi et al. “Drogon: A trajectory prediction model based on intention-conditioned behavior reasoning”. In: *Conference on Robot Learning*. PMLR. 2021, pp. 49–63.

[13] Biϕrnar R Dalsnes et al. “The neighbor course distribution method with Gaussian mixture models for AIS-based vessel trajectory prediction”. In: *2018 21st International Conference on Information Fusion (FUSION)*. IEEE. 2018, pp. 580–587.

[14] Patrick Dendorfer, Aljosa Osep, and Laura Leal-Taixé. “Goal-gan: Multimodal trajectory prediction based on goal position estimation”. In: *Proceedings of the Asian Conference on Computer Vision*. 2020.

[15] Martin Ester et al. “A density-based algorithm for discovering clusters in large spatial databases with noise.” In: *kdd*. Vol. 96. 34. 1996, pp. 226–231.

[16] Nicola Forti et al. “Prediction oof vessel trajectories from AIS data via sequence-to-sequence recurrent neural networks”. In: *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2020, pp. 8936–8940.

[17] Shaojun Gan et al. “Ship trajectory prediction for intelligent traffic management using clustering and ANN”. In: *2016 UKACC 11th International Conference on Control (CONTROL)*. IEEE. 2016, pp. 1–6.

[18] Harshayu Girase et al. “Loki: Long term and key intentions for trajectory prediction”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 9803–9812.

[19] Ian J. Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. `http://www.deeplearningbook.org`. Cambridge, MA, USA: MIT Press, 2016.

[20] Michelle R Grech, Tim Horberry, and Andrew Smith. “Human error in maritime operations: Analyses of accident reports using the Leximancer tool”. In: *Proceedings of the human factors and ergonomics society annual meeting*. Vol. 46. 19. Sage Publications Sage CA: Los Angeles, CA. 2002, pp. 1718–1721.

[21] Simen Hexeberg, Andreas L Flåten, Edmund F Brekke, et al. “AIS-based vessel trajectory prediction”. In: *2017 20th International Conference on Information Fusion (Fusion)*. IEEE. 2017, pp. 1–8.

[22] Chen Jiashun. “A new trajectory clustering algorithm based on TRACLUS”. In: *Proceedings of 2012 2nd International Conference on Computer Science and Network Technology*. IEEE. 2012, pp. 783–787.

[23] Anish Joseph and Dimitrios Dalaklis. “The international convention for the safety of life at sea: highlighting interrelations of measures towards effective risk mitigation”. In: *Journal of International Maritime Safety, Environmental Affairs, and Shipping* 5.1 (2021), pp. 1–11.

[24] Pablo Kaluza et al. “The complex network of global cargo ship movements”. In: *Journal of the Royal Society Interface* 7.48 (2010), pp. 1093–1103.

[25] Se-Won Kim, Yoon-Suk Lee, and Young-Soo Park. “A Study on the Size of Turning Basin for Vessels of Arrival & Departure in the Berths”. In: *Journal of Fisheries and Marine Sciences Education* 24.6 (2012), pp. 872–883.

[26] J Kornacki and W Galor. “Analysis of ships turn manoeuvres in port water area”. In: *TransNav, International Journal on Marine Navigation and Safety od Sea Transportation* 1.1 (2007).

[27] Parth Kothari, Sven Kreiss, and Alexandre Alahi. “Human trajectory forecasting in crowds: A deep learning perspective”. In: *IEEE Transactions on Intelligent Transportation Systems* 23.7 (2021), pp. 7386–7400.

[28] Yoshiaki Kuwata et al. “Safe maritime autonomous navigation with COLREGS, using velocity obstacles”. In: *IEEE Journal of Oceanic Engineering* 39.1 (2013), pp. 110–119.

[29] Jae-Gil Lee, Jiawei Han, and Kyu-Young Whang. “Trajectory clustering: a partition-and-group framework”. In: *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*. 2007, pp. 593–604.

[30] Jiachen Li et al. “Rain: Reinforced hybrid attention inference network for motion forecasting”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 16096–16106.

[31] Matteo Lisotto, Pasquale Coscia, and Lamberto Ballan. “Social and scene-aware trajectory prediction in crowded spaces”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*. 2019, pp. 0–0.

[32] Bingbin Liu et al. “Spatiotemporal relationship reasoning for pedestrian intent prediction”. In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 3485–3492.

[33] Jianbang Liu et al. “A survey on deep-learning approaches for vehicle trajectory prediction in autonomous driving”. In: *2021 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE. 2021, pp. 978–985.

[34] Hengbo Ma et al. “Continual multi-agent interaction behavior prediction with conditional generative memory”. In: *IEEE Robotics and Automation Letters* 6.4 (2021), pp. 8410–8417.

[35] Srikanth Malla, Behzad Dariush, and Chiho Choi. “Titan: Future forecast using action priors”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, pp. 11186–11196.

[36] Karttikeya Mangalam et al. “It is not the journey but the destination: Endpoint conditioned trajectory prediction”. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*. Springer. 2020, pp. 759–776.

[37] Agentschap voor Maritieme Dienstverlening en Kust. Afdeling Kust. Vlaamse Hydrografie. *Getijtafels 2022 voor Nieuwpoort, Oostende, Zeebrugge, Vlissingen, Prosperpolder, Antwerpen en Wintam*. 2021. URL: https://www.vlaanderen.be/publicaties/getijtafels-2022-voor-nieuwpoort-oostende-zeebrugge-vlissingen-prosperpolder-antwerpen-en-wintam (visited on 04/12/2022).

[38] Leonardo M Millefiori et al. "Modeling vessel kinematics using a stochastic mean-reverting process for long-term prediction". In: *IEEE Transactions on Aerospace and Electronic Systems* 52.5 (2016), pp. 2313–2330.

[39] Vytautas Paulauskas. "Ships turning basins in ports for big container vessels". In: *Zeszyty Naukowe Akademii Morskiej w Szczecinie* 20 (92 (2010), pp. 102–106.

[40] Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. "Are they going to cross? a benchmark dataset and baseline for pedestrian crosswalk behavior". In: *Proceedings of the IEEE International Conference on Computer Vision Workshops*. 2017, pp. 206–213.

[41] Andrey Rudenko et al. "Human motion trajectory prediction: A survey". In: *The International Journal of Robotics Research* 39.8 (2020), pp. 895–935.

[42] Greg Welch, Gary Bishop, et al. "An introduction to the Kalman filter". In: (1995).

[43] Cheng-Hong Yang et al. "Deep Learning for Vessel Trajectory Prediction Using Clustered AIS Data". In: *Mathematics* 10.16 (2022), p. 2936.

[44] Shudong Yang, Xueying Yu, and Ying Zhou. "Lstm and gru neural network performance comparison study: Taking yelp review dataset as an example". In: *2020 International workshop on electronic communication and artificial intelligence (IWECAI)*. IEEE. 2020, pp. 98–101.

[45] Lan You et al. "St-seq2seq: A spatio-temporal feature-optimized seq2seq model for short-term vessel trajectory prediction". In: *IEEE Access* 8 (2020), pp. 218565–218574.

[46] Hang Zhao et al. "Tnt: Target-driven trajectory prediction". In: *Conference on Robot Learning*. PMLR. 2021, pp. 895–904.

[47] He Zhao and Richard P Wildes. "Where are you heading? dynamic trajectory prediction with expert goal examples". In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, pp. 7629–7638.

A

Paper

Short-term vessel trajectory and swing manoeuvre prediction

Joaquin van Loon[2] , **Timo Koppenberg**[1] , **Neil Yorke-Smith**[2]
[1]Macomi B.V., [2]Delft University of Technology
jlkvanloon@student.tudelft.nl, t.koppenberg@macomi.nl, n.yorke-smith@tudelft.nl

Abstract

Situational awareness within port areas is crucial to avoid collisions, navigate efficiently and reduce congestion. Maritime traffic controllers constantly monitor the situation in the port and intervene when needed. This study proposes a deep learning model that predicts future vessel positions to assist in this process. The model employs a target conditioned trajectory prediction component composed of two recurrent neural networks arranged in an encoder-decoder structure that utilizes historical data points to forecast future trajectories. The model considers multiple factors, including vessel speed, location, length, depth, draught, and the tide. Additionally, this study addresses the prediction of swing manoeuvres, which are special U-turn-like manoeuvres executed during terminal arrival or departure. These manoeuvres can block a significant portion of the waterway and, as such, are essential to consider when gaining a complete understanding of future situations within the port. An integration of both model is applied to a use case study in a scoped area of the Port of Antwerp-Bruges. The models were trained using AIS and VTS data collected at 30-second intervals. Swing manoeuvres are predicted with an accuracy of 84%, the locations of these manoeuvres are predicted with an average deviation of 212 meter and the duration error is 1.6 minutes on average. The complete predicted trajectories, including potential swing manoeuvres, have an average displacement error and final displacement error of 147 and 117 meter on average, respectively. Overall, the study demonstrates the potential of deep learning models for improving situational awareness within port areas and assisting traffic controllers in making informed decisions.

1 Introduction

The Port of Antwerp-Bruges is the second largest seaport in Europe and handled around 250 million tonnes of maritime freight volume in the year 2021 [4]. It is a busy hub of maritime activity with a complex nautical situation that demands continuous monitoring. Given the large volume of vessels that navigate the port's waterways, ensuring safety is of crucial importance. In addition to preventing potential hazards and reducing congestion, maintaining situational awareness can also help to minimize carbon emissions by optimizing vessel traffic flows and reducing idle time.

The Antwerp port area has five deep-sea container terminals, which are used to load and unload large container ships arriving and departing from the port. Since these container vessels arrive and leave the port through one waterway, they need to make a U-turn somewhere within the port. Such a U-turn manoeuvre is tugboat assisted and is performed either on arrival or departure. Swinging is performed at the terminal or in a dedicated turning basin. The exact location of a swing is dependent upon various factors, including but not limited to the vessel's draught and the tide conditions. These swing manoeuvres are required, block a large part of the waterway, and take substantial time.

This study narrows its focus to a small, scoped area of the Antwerp port area. This scoped area is of particular interest, because it connects the southern parts of the port to the North Sea and it contains two terminals and two locks in a relatively small area. A bird eye view of this area can be seen in Figure 1. The Europa Terminal (EUT) and Noordzee Terminal (NZT) are used to dock container vessels and load and unload them. Furthermore, the Zandvliet Lock (ZVS) and Berendrecht Lock (BES) are used to enter or leave the inland waters of the port. All vessels leaving the port to the sea or arriving from the sea will pass this scoped area. Due to this high traffic intensity and the presence of two terminals and locks, this area is considered of critical importance.

The presence of the Europa Terminal and Noordzee Terminal significantly affects the area of the port, particularly when it comes to swing manoeuvres. Swinging vessels blocking the waterway in front of the Noordzee Terminal restrict traffic to more southern parts of the port and to the locks. Moreover, swinging vessels in front of the Europa Terminal restrict traffic to the southern parts of the port. Since it is crucial to avoid collisions in the port and minimize congestion and delays, these swing manoeuvres play an important role in the management of traffic.

According to a study that analyzed marine collision reports, 71% of human error is caused due to lack of situational awareness. Specifically, of the errors related to situational awareness, around 9% were attributed to a lack of awareness

regarding future actions in the marine environment [17]. The objective of this study is to assist maritime traffic controllers in obtaining situational awareness. To aid them in this process, this study aims to create a model that can predict future vessel trajectories. Since swing manoeuvres are crucial parts of these trajectories, the model must also take them into account. Given this objective, the main research question is:

> *How can a vessel trajectory prediction model that accounts for swinging manoeuvres be developed and applied to a case study at the Port of Antwerp-Bruges?*

By creating such a model, the contribution of this case study is two-fold:

- This study extends state-of-the-art vessel trajectory prediction methods to incorporate swing manoeuvres.
- This study applies trajectory prediction models on a specific use case scenario at the Port of Antwerp-Bruges.

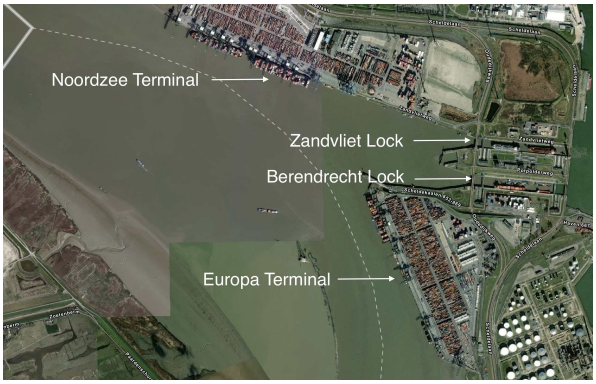

Figure 1: A bird eye view of the scoped area of the Port of Antwerp-Bruges

2 Background

The large availability of AIS data records has enabled the possibility of predicting trajectories in the domain of vessels. However, the task of trajectory prediction is not unique to the prediction of vessel movements and is relevant in more domains, such as human, aircraft or car trajectory prediction.

A basic and commonly used model for vessel trajectory prediction is to assume near constant speed and course [23]. Improvements on such a model are made by modelling the speed of a vessel using an Ornstein-Uhlenbeck process [31]. The Kalman filter is a different algorithm which can also be used to model and predict the movement of objects [36]. However, these physic-based approaches do not utilise the knowledge of historical AIS data. Prediction methods which are based on historic AIS data often cluster historical trajectories, classify incoming AIS data points to one of these clusters and then construct a prediction [3, 18]. These methods can differ in clustering algorithms and trajectory distance measures [13, 19, 24]. Besides such a trajectory-based similarity search prediction model, Alizadeh et al. presents a point-based similarity search prediction approach [3]. Here, each AIS record is treated as a singular point, and spatial, speed, and course variables are used to calculate the distances between the vessel's most recent point and all the historical points measured. By identifying the most similar historical point to the target point, the subsequent locations of the vessel can be predicted based on that point. A similar approach is used by Hexeberg et. al [18].

The success of deep neural networks has led to a significant advancement in the field of trajectory prediction in recent years. Deep learning approaches used in literature take a historic agent state representation, pass this to a deep learning model and output a certain representation of the predicted future trajectory. The agent state is often represented as a vector of features, which include the timestamp, speed, heading and position under a certain coordinate system [1,5,7,22,29]. The input of the learning model then consists of a vector of these historic agent states. Next to this agent state, some researches use additional scene context representations to give the model more information. A popular scene context representation are images, used in for example [7, 9, 30, 33] and [25].

The future trajectory, i.e. the model's output, can simply be given as a point set [1, 22, 29, 38]. This single modality representation aligns with the input representation. Outputs can also be represented as probability heat-maps [26] or probabilities can be attached to each prediction [25, 30]. Instead of outputting exact positional values, a model could also predict a mean position μ and variance σ around it. Controlling uncertainty in this manner is done in [8].

In the literature, several neural network architectures are proposed for the actual deep-learning model. Since the input representation is often a vector of trajectory points, recurrent neural networks can be used to encode the historic trajectory [2, 5, 7, 28, 38]. Given this encoding, a second recurrent neural network can decode the vector to a future trajectory, this encoder-decoder architecture is used in [14, 38] and [6] for example. Image inputs are often processed using convolutional neural networks (CNN) [26, 30] and graph neural networks (GNN) can be used to model interaction effects [11, 27, 34].

Supervised learning is the most common method used to train these deep learning models and the loss function commonly uses a mean-square-error (MSE) loss to compare the predicted trajectory against the truth [39]. Additional components can be added to the loss function to smooth the prediction, such as an inconsistency penalty, penetration penalty or dispersion penalty as used in [10].

A good trajectory prediction model should be able to capture the multimodality of the task, i.e. the model should be able to handle distinct end locations properly. However, models tend to suffer from mode collapse, which means average and non-realistic trajectories are predicted instead of diverse predictions. Dendorfer et al. [12] give an excellent example of a neural network suffering from mode collapse when predicting pedestrian trajectories near a crossroad. A way to counter this problem is to first use a model to cluster trajectories and then use a trajectory prediction model per cluster [15, 37]. By first clustering the trajectories, the neural network needs to learn less trajectory variations and can potentially generalize better.

In literature, an alternative method that is employed involves making a prediction of the goal location of the trajectory, followed by a prediction of the trajectory itself. Instead of asking a model to extrapolate a past trajectory, this method changes the problem to an interpolation task between the predicted goal position and the historic trajectory. Such approaches are commonly two staged, a target location is predicted based on the historical trajectory and scene context and this target prediction assists later trajectory forecasting [10, 16, 39, 40]. Target conditioned trajectory prediction has shown to perform better on diverse trajectories [12].

As shown, there is a lot of literature present on plain trajectory prediction. However, no relevant literature has been found which involves swing manoeuvres in these predicted trajectories. Swing manoeuvres are not unique for the Port of Antwerp-Bruges. More ports use swing manoeuvres to align container vessels with their desired direction. Especially in ports with only one waterway for entering and leaving the port, vessels need to make such a manoeuvre. Since it often raises difficulties to find a suitable location for container vessel to execute a swing manoeuvre, as is the problem in the Port of Antwerp-Bruges, studies have been conducted on designing port areas in an optimal manner that considers swing manoeuvres [20, 21]. However, changing existing port layouts is difficult and the fact that container vessels have become longer over time adds to the problem. Paulauskas discusses the influence of these big container vessels on ports designed using old safety and navigation standards [32]. With his analysis on turning basins, he claims that due to modern accurate navigational and measurement equipment, the old safety standards implemented in several ports should be reviewed.

The trajectory prediction model presented by Hexeberg et. al obtained some large errors on certain trajectories, which were caused by swing manoeuvres as reported by the authors [18]. Besides this notion of swing manoeuvres in the vessel trajectory prediction literature, no literature is available on trajectory prediction methods which incorporate swing manoeuvres to the best of the authors' knowledge.

3 Data

The Port of Antwerp-Bruges continuously records the positional data and meta data of the vessels entering the port. For this study, The Port of Antwerp-Bruges provided two years of vessel data. This study only considers data of the seagoing vessels, internal port movements are ignored. Most data come from the AIS source, but a part originates from the VTS. Complete data covering two years is only available for container vessels. For other vessel types, only data from September and October 2022 is present. The availability of a significant amount of container vessel data is crucial, as trajectories of container vessels interacting with the Noordzee Terminal or Europa Terminal are less frequent than those bypassing the scoped area of the port. The Port of Antwerp-Bruges provided the data in three parts: vessel positional data, vessel meta data and vessel trajectory data

The positional data records were captured at approximately 30-second intervals and include a timestamp, latitude position, longitude position, speed over ground (in knots) and heading of the vessel. Additionally, static data for each vessel that passed through the port was provided in the vessel meta data file. The relevant features from the meta data are the type, depth, draugt, length, width, origin and destination of the vessel. The vessel origin and destination contain the exact name of a specific location, such as "Loodskotter West" or the terminal location "S853". A pre-processing step is applied to classify these origins and destinations into these categories: ZVS, BES, EUT, NZT, NORTH and SOUTH. The vessel trajectory data file consists of a list of vessel trajectories, each characterized by a start time (in case of departure) or end time (in case of arrival), along with additional meta data fields. Notably, the draught of a vessel during its trajectory, which is the distance between the waterline and the bottom of the vessel, was included as a useful feature that is not present in the meta data file. The trajectory data is used to extract trajectories from the positional data. An interpolating pre-processing step is applied to these trajectories such that the time difference between successive points is always exactly 30 seconds.

As tide data is be an important factor in vessel traffic planning according to domain knowledge, it is beneficial to add it to the existing AIS data. The original delivered data does not contain tide data. However, the Flemish government provides tide tables on their website [35], which indicate the water levels at different locations in the Belgium waters. The tide data in the port was estimated by using the water levels at "Prosperpolder", which is the closest location to the port where tide data is available.

4 Swing manoeuvres

Since there is only one entry and exit point for sea vessels, a U-turn manoeuvre is necessary for vessels on a round-trip, commonly known as a swing manoeuvre. Typically, swing manoeuvres are executed near the terminal location of the vessel, either before it docks on arrival or after it is released from the quay during departure. These manoeuvres are time-consuming and require tugboat assistance. Moreover, for safety reasons, a buffer zone must be maintained around the swinging vessel, which can block parts of the waterway, affecting the passage of other vessels. Environmental factors such as tide conditions and wind strength also play a crucial role in deciding to swing or not.

4.1 Swing detection

Swing detection is the process of determining whether a historical trajectory of a vessel contains a swing manoeuvre or not. Such manoeuvres are distinguished by a considerable change in the vessel's heading over a short duration of time. As the vessel executes a U-turn, the heading change is typically around 180 degrees over the entire manoeuvre. Additionally, swing manoeuvres usually take around ± 20 minutes to complete. Based on these observations, swing detection logic can be created to identify trajectories with swing manoeuvres.

Given is a trajectory representation V of n data points $V = \{\vec{v}_1, \vec{v}_2, \ldots, \vec{v}_{n-1}, \vec{v}_n\}$. Each data point $\vec{v}_i$ is a vector

$\vec{v}_i = (p_i, q_i, t_i, s_i, h_i)$, where $p_i \in P$ represents the longitude position, $q_i \in Q$ the latitude position, $t_i \in T$ the epoch timestamp in milliseconds, $s_i \in S$ the speed over ground in knots and $h_i \in H$ the heading of the vessel in degrees.

Now define a function $f : H \times H \to [0, 180]$ which takes two headings and outputs the smallest difference in degrees between these headings. The smallest difference between the headings is considered, because of the modular 360 nature of circles.

$$f(h_i, h_j) = 180 - ||h_i - h_j| - 180| \tag{1}$$

By combining domain knowledge and this function, it is possible to formalize a definition for a trajectory that includes a swing manoeuvre:

$$\begin{aligned} V \text{ contains a swing manoeuvre} \Leftrightarrow \\ \{(i,j) | 1 \leq i, j \leq n \wedge j > i \wedge t_j - t_i \leq 25 * 60000 \\ \wedge f(h_i, h_j) \geq 125\} \neq \emptyset \end{aligned} \tag{2}$$

This formalisation means that a vessel makes a swing manoeuvre if and only if there exist two data points on the trajectory which are no more than 25 minutes apart and have a difference in heading larger or equal to 125 degrees, where the difference in heading is calculated using equation 1. Note that a factor 60000 is used to convert the 25 minutes to milliseconds, which is the unit of t_i and t_j.

Domain knowledge indicates that swing manoeuvres can take up to 20 minutes, therefore 25 minutes is a reasonably safe upper bound. The choice for 125 degrees is based on visual inspection. In the histogram in Figure 2 the red dotted line at $x = 125$ clearly splits two clusters of trajectories. To the left of this line are trajectories who do not contain a swing manoeuvre and too the right of this line are those who do. An expected peak around 180 degrees is visible, which represents full U-turn manoeuvres.

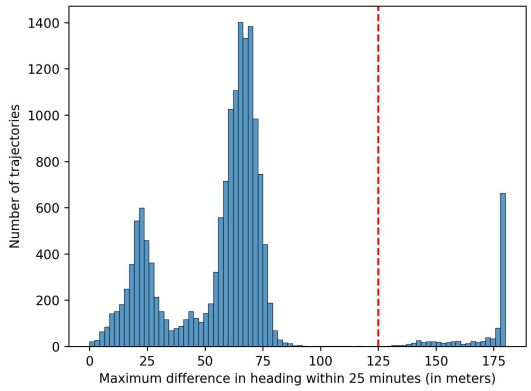

Figure 2: Distribution of maximum directional change in container vessel trajectories

4.2 Swing extraction

Once swing manoeuvres are detected from historical trajectories, it is important to extract them from the entire trajectory. This means that only the segment of the trajectory in which the vessel performs a swing manoeuvre is kept. Swing manoeuvres are often done approximately in place, which means the speed of a vessel during its trajectory is an important factor to look at. Furthermore, the change of the vessel's heading over time is something to look at. This is expressed by the yaw rate, which is the directional change in heading in degrees per minute. For the directional change between two headings equation 1 can be used again.

A swing window $w = [b_{swing}, e_{swing}]$ can be defined where $b_{swing}, e_{swing} \in \{1, \ldots, n\}$. This window indicates the part of the trajectory in which the vessel is swinging. Further let $\gamma : \{1, \ldots, n-1\} \to \mathbb{R}$ be a function which, given $i \in \{1, \ldots, n-1\}$, computes the yaw rate in the interval $[i, i+1]$.

$$\gamma(i) = \frac{f(h_i, h_{i+1})}{(t_{i+1} - t_i) * \frac{1}{30000}} \tag{3}$$

Next, let the indicator function $\psi : \{1, \ldots, n-1\} \to \{0, 1\}$ indicates if a data point $\vec{v}_i$ is during a swing manoeuvre or not. If $\psi(i) = 1$, then the vessel is performing a swing manoeuvre. However, the inverse is not necessarily true.

$$\psi(i) = \begin{cases} 1, & \text{if } 1 \leq i \leq n \wedge \gamma(i) \geq 10 \wedge s_i \leq 2 \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

Hence, a data point is classified to be during a swing manoeuvre if the change in heading is currently larger or equal to 10 degrees per minute and the speed is lower or equal to 2 knots, which is approximately 3.7 kilometers per hour. The motivation behind these numbers come from domain knowledge and visual inspection.

Using the indicator function ψ a window $[b'_{swing}, e'_{swing}]$ can be defined in which the vessel is definitely swinging. An assumption is made that each trajectory comprises a single swing manoeuvre.

$$b'_{swing} = \min_i [\psi(i) = 1] \text{ where } i \in \{1, \ldots, n-1\} \tag{5}$$

$$e'_{swing} = \max_i [\psi(i) = 1] \text{ where } i \in \{1, \ldots, n-1\} \tag{6}$$

Vessels can have a lower yaw rate than 10 degrees per minute in the beginning or ending of a swing. Moreover, small vessels could swing at larger speeds than 2 knots, especially in front of the EUT. This means that the use of the indicator function in equation 5 and 6 might mean that the window $[b'_{swing}, e'_{swing}]$ is smaller than the actual swinging window $[b_{swing}, e_{swing}]$, i.e. $b'_{swing} > b_{swing} \wedge e'_{swing} < e_{swing}$. Using a lower yaw rate bound and higher speed bound in the definition of ψ to solve the problem is not possible. ψ will violates its meaning then, as it will be too sensitive and falsely classify small course changes during the trajectory of a vessel as being part of a swing manoeuvre.

To still be able to extract the full swing window, the window is extended on both sides by including data points where the yaw rate is still above 6 degrees per minute. This procedure is written down in equation 7 and 8.

$$b_{swing} = \min\{i \in \{1, \ldots, b'_{swing}\}| \\ (\forall k \in \{i-1, \ldots, b'_{swing}\})[\gamma(k) \geq 6]\} \quad (7)$$

$$e_{swing} = min(n, 1 + \max\{i \in \{e'_{swing}, \ldots, n\}| \\ (\forall k \in \{e'_{swing}, \ldots, n\})[\gamma(k) \geq 6]\}) \quad (8)$$

An example of swing extraction for a specific trajectory using the procedure described above is shown in Figure 3. The points in this graph are plotted in one-minute intervals for clarity. In this example $[b'_{swing}, e'_{swing}] = [6, 10]$ and $[b_{swing}, e_{swing}] = [4, 13]$. The horizontal lines are plotted at a yaw rate of 6 and 10.

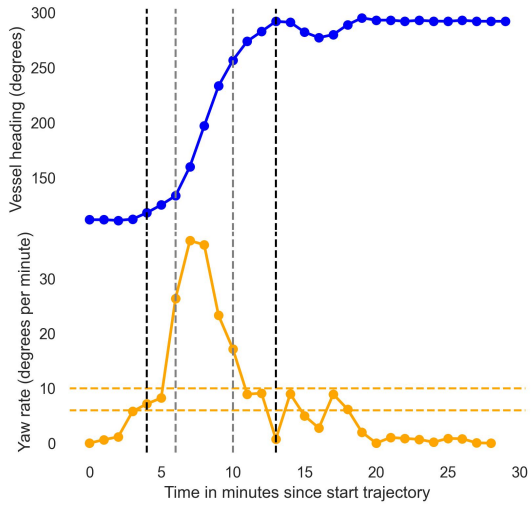

Figure 3: A tight swing window (grey dotted line) and extended swing window (black dotted line). Interval between points is one minute.

4.3 Swing modelling

Predicting swing manoeuvres can be done in different levels of details. An option is to predict the whole swing manoeuvre, i.e. the location and headings during the whole manoeuvre. However, prediction of the full swing manoeuvre is a difficult problem to solve since these manoeuvres are relatively rare and can vary a lot in their exact details. The goal of the study is to make a capacity prediction of the waterways. Exact heading and position of the vessel during the swing manoeuvre is less important than the occupation of the waterway resources and the duration of this occupation. Therefore, an alternative could be to only model the location and duration of the swing and assume the vessel will swing in place. However, the downside of this approach is that it ignores the movements which the vessel sometimes makes during the swing manoeuvre. Therefore, this study uses an in-between approach by predicting the start location, end location and duration of the swing. This abstraction keeps the relevant details of the swing while not modelling unnecessary details. Figure 4, 5 and 6 visualise these three different ways of modelling a swing manoeuvre.

Figure 4: Option 1: modelling a swing manoeuvre fully (reality)

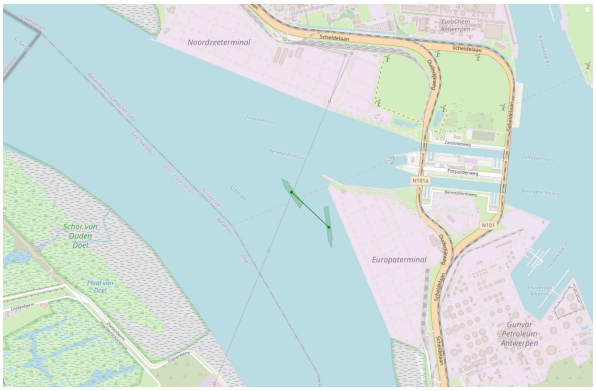

Figure 5: Option 2: modelling only the start and end of a swing manoeuvre)

5 Model

The model which makes short term vessel trajectory predictions consists of two parts: a trajectory prediction component and a swing manoeuvre prediction component.

5.1 Trajectory prediction

Predicting the future trajectory of a vessel based on its historical positions can be treated as a sequence-to-sequence task, which can be addressed using recurrent neural networks. An encoder-decoder architecture with two recurrent neural networks can be utilized to predict future vessel trajectories based on historical positions. The encoder uses a recurrent neural network to encode the historical data points into a vector embedding. Before the encoder does this, it embeds the input in a higher dimension using a multilayer perceptron. The encoder output vector is combined with the meta data and fed through a feedforward neural network to predict the target position, which consists of a latitude, longitude and heading. Once the target position is predicted, it is concatenated with the encoder output and then passed through a multilayer perceptron to produce the initial hidden state of the decoder. The

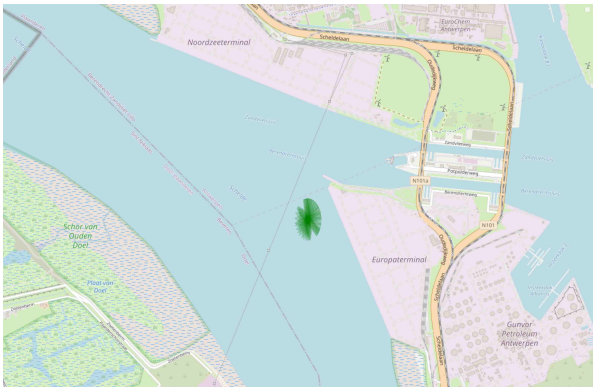

Figure 6: Option 3: modelling the swing manoeuvre in place

decoder sequentially generates future points by taking predicted future points as input for the next block. For the first input, the last point in the historical trajectory is used. Besides these future points the decoder also uses the target position as input. The future point and target position are concatenated and passed through a multilayer perceptron to form the input of the decoder. The decoder is given the target position at each time stamp to help remember this important piece of information. A schematic overview of this target conditioned trajectory prediction model can be seen in Figure 7.

The final goal of the complete model is to predict the future trajectory as accurately as possible by using a predicted target point that should be as close as possible to the actual target point. Therefore, the loss of the model should align with this goal. Given a historic trajectory $\{x_0, x_1, \ldots, x_m\}$, a future trajectory $\{y_0, y_1, \ldots, y_m\}$, a predicted target point t and a predicted future trajectory $\{\hat{y}_0, \hat{y}_1, \ldots, \hat{y}_m\}$, the loss of the model is calculated with equation 9

$$L = c_1 * \frac{1}{m}\sum_{i=0}^{m}(y_i - \hat{y}_i)^2 + c_2 * (t - \hat{y}_m)^2 + c_3 * (y_m - \hat{y}_m)^2 + c_4 * \frac{1}{m}\sum_{i=0}^{m}(y_i - \hat{y}_{i-1})^2. \quad (9)$$

In this equation, $c_1 = 1, c_2 = 1, c_3 = 1, c_4 = 0.75$ are reasonable values for these constants. In the last term of the formula x_n is used for $\hat{y}_{-1}$ (i.e. when $i = 0$ in the sum).

The first term of the loss function is a mean squared error component on the full future trajectory. The second and third terms are both squared errors on the target prediction and the last predicted trajectory point. Finally, the last term is a mean squared error between the predicted trajectory and the predicted trajectory shifted back one time stamp. By adding this term, the model is restricted from taking unrealistic speed and position changes between adjacent frames. Such an inconsistency penalty term to smooth the predicted trajectories by encouraging temporal consistency between trajectory points is used more often in trajectory prediction models [10].

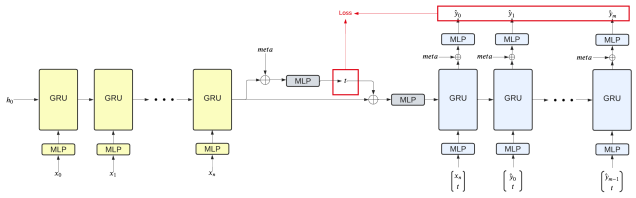

Figure 7: Target conditioned trajectory prediction model network structure using gated recurrent units

5.2 Swing prediction

Container vessels interacting with terminals need to swing either on arrival or departure. The captain's preference plays a crucial role in the decision whether to swing or not when arriving in the port. The captain bases its choice mainly on environmental factors such as tide conditions and on traffic conditions.

Determining whether a departing vessel will swing is easy. The way these vessels are oriented against the terminal reveal whether they require to swing or not. For arriving vessels, it is not straightforward whether they will swing or not. If an arriving vessel does not swing, it essentially postpones this manoeuvre to its departure and when the vessel does swing, it can later immediately depart. Given data about the vessel's journey and the environmental conditions, a prediction can be made on whether the vessel will swing on arrival or not. Locations of containers on the vessel are not taken into account and it is assumed that all captains have the same preference profile.

Based on vessel's meta data and on a fraction of the trajectory of varying length before the swing, the goal is to predict whether the vessel will swing or not. For this task the model visualised in Figure 8 is used. This model takes the historical trajectory (i.e. before the swing) and encodes it in a vector. This encoding is done in the same manner as historical trajectories are encoded in the target conditioned trajectory prediction model. To this encoded vector, the meta data vector is concatenated and then the result is passed through a final dense neural network. The output layer of this neural network has one neuron with a sigmoid activation function. Therefore, this neuron outputs a value $\hat{y} \in [0, 1]$, $\hat{y} \geq 0.5$ is considered a swing prediction and $\hat{y} < 0.5$ as a prediction without a swing.

After establishing that the vessel will undergo a swing manoeuvre, a follow up question is how this swing manoeuvre should look like. For the maritime traffic controllers in the port of Antwerp-Bruges, the most important knowledge to obtain is the location and duration of the swing. Therefore, predicting the start and end position of the swing is a usable approach. The start and end position are characterised by their longitude position, latitude position and heading. A swing location prediction model should be able to predict these three characteristics for both the start and end of the swing based on a historic trajectory before the swing manoeuvre. For this purpose, the same neural network architectural setup as before is used again. An encoder network will encode the historic trajectory in a vector and this encoded vector will be concatenated to the meta data vector and passed through a final feedforward neural network, see Figure 8. The

output layer of this feedforward neural network will contain 9 neurons:

- Neuron 1: the latitude position at the start of the swing
- Neuron 2: the longitude position at the start of the swing
- Neuron 3: the sine of the heading at the start of the swing
- Neuron 4: the cosine of the heading at the start of the swing
- Neuron 5: the latitude position at the end of the swing
- Neuron 6: the longitude position at the end of the swing
- Neuron 7: the sine of the heading at the end of the swing
- Neuron 8: the cosine of the heading at the end of the swing
- Neuron 9: the duration of the swing manoeuvre in minutes.

The sine and cosine of the headings are used instead of the heading directly because of the circular nature of the heading values. A heading change from 2 degrees to 359 degrees seems like a huge difference numerically, but visually these headings are quite close. The cosine and sine values of a heading of 2 degrees and a heading of 359 degrees have the property of being quite close to each other numerically.

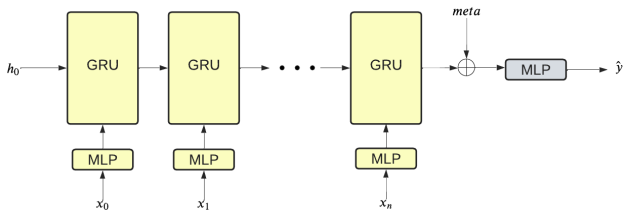

Figure 8: Swing prediction deep learning model architecture using a gated recurrent unit

5.3 Complete model

Given the target conditioned trajectory prediction model and the swing prediction model, a complete model can be constructed which integrates both components and can make full short term trajectory predictions.

All route options which do not contain a swing manoeuvre can be dealt with by just using the target conditioned trajectory prediction model. For the route options which do include a swing manoeuvre, a combination of this trajectory prediction model and the swing location prediction model can be used. The swing location prediction model will predict the start and end location of the swing manoeuvre and the trajectory prediction model will take care of predicting the trajectory to the start position of the swing or from the end position of the swing onwards.

6 Results

The trajectories extracted from the raw vessel data are split in an 80% train set and 20% test set. To ensure that all models can handle variable-length inputs, they are trained on inputs of different lengths. For a trajectory consisting of n data points, $n-1$ sub-trajectories of lengths $1, 2, 3, \ldots, n-1$ are extracted. Thus, each sub-trajectory always contains the first point and differs in the number of additional historical trajectory points it contains.

The performance of the target-conditioned trajectory prediction model can be found in Table 1. The average displacement error (ADE) and final displacement error (FDE) are reported for different trajectory options, identified by their origin and destination.

Origin	Destination	ADE	FDE
Zandvliet Lock	North	115	127
Berendrecht Lock	North	111	131
South	North	109	134
North	South	100	117
North	Zandvliet Lock	182	423
North	Berendrecht Lock	185	422
North	Europa Terminal	200	398
North	Noordzee Terminal	187	336
Europa Terminal	North	149	146
Noordzee Terminal	North	154	216

Table 1: Results of the target conditioned trajectory prediction model. ADE and FDE are both in meters.

Table 2 shows the accuracy results of the swing occurrence prediction model for arriving vessels. No model is necessary for departing vessels, as their heading alone indicates whether the vessel will swing or not. Table 3 presents the results of the swing location prediction model.

Terminal (arrival)	Accuracy
Noordzee Terminal	0.84
Europa Terminal	0.83

Table 2: Accuracy results of the swing occurrence prediction model

Terminal	Route type	P_e		H_e		D_e
		Start	End	Start	End	
Noordzee	Arrival	237	216	7.4	12.8	1.4
Noordzee	Departure	132	140	15.3	6.7	1.9
Europa	Arrival	254	217	9.9	15.7	1.7
Europa	Departure	199	301	9.6	14.2	1.5

Table 3: Results of the swing location prediction model. P_e indicate the positional errors n meters, H_e the heading errors in degrees and D_e the duration errors in minutes.

The results of the complete model are presented in Table 4. These results are categorized into trajectories interacting with terminals and trajectories just passing through the scoped area.

A large difference in duration error is visible between the two categories of trajectories. The duration errors for the terminal trajectories are mainly made in the undocking and docking portion of the trajectory. These parts of the trajectory take relatively long, meaning there is room for large duration errors. The complete model uses an average for the

Terminal interaction	D_e	FDE	ADE
No	1.3	102	149
Yes	11.0	194	146
Combined	3.2	117	147

Table 4: Results of complete model. ADE and FDE are both in meters. D_e is the duruation error in minutes

(un)docking speed as approximation. However, since these speeds seem to vary for each trajectory it leads to the large average duration error. For docking trajectories the impact of these duration errors is minimal, since the crucial part of the trajectory is already done. However, for the trajectories leaving the terminal, this duration error has accumulating impact resulting in an inaccurate view of the situation in the future. The combined duration error is quite low, because there are more non terminal trajectories in the test set. The terminal trajectories show a slightly higher final displacement error since the final location at the terminal is harder to predict than the final location of a simple trajectory going to, for example, the Zandvliet Lock.

While an average FDE of 117 and average ADE of 147 may appear high, these predictions seem acceptable on a complete map, particularly given the average vessel length of 227 meters. To put these numbers in a bit of context, Figure 9 and Figure 10 report the ADE and FDE of some plotted trajectories. Additionally, to give some insights in how the trajectories with swing predictions look like, Figure 11 illustrates several complete trajectory predictions including swing manoeuvres.

7 Discussion

Even though the predicted trajectories are of sufficient quality, additional performance could be obtained by incorporating certain changes.

First of all, wind direction and strength are not incorporated into the analysis; including these factors may be beneficial. Based on domain knowledge, these factors are namely believed to be linked with the trajectory path and the location of the swing manoeuvre. Additional accuracy could probably also be obtained by using the exact planned final location of the vessel. Currently the terminal number is used as the destination for vessels arriving at terminals. However, the mooring bollards that the vessel will use are also known in advance. These bollards give a more accurate indication of the destination of the vessel along the terminal.

Another key aspect which is omitted in this research are interaction effects. Trajectory, swing occurrence and swing location prediction are currently all based on solely the historical trajectory and meta data factors. However, in reality a swing decision or the sailing course could be affected by the current or even future traffic condition. If the traffic keeps increasing in the Port of Antwerp-Bruges, taking interactions effects into account becomes even more important. The decision to not take interaction effects into account in this study was based on the limited presence of vessel interaction in the given data. A suggestion for future research would be to use graph neural networks (GNN) to include interaction effects once more of these effects are visible in the data.

Furthermore, several data pre-processing steps could potentially reduce the complexity of the prediction and therefore improve performance. A suggestion could be to use the position of the vessel at the last historical data point as the origin of the coordinate system for representing trajectories. Similarly, the heading of the vessel could be represented such that the heading at the last data point aligns with the x-axis.

Lastly, important to mention is that extensive parameter and architectural configuration optimisation is left out of this study. Additional performance can possibly be gained by tweaking the neural network architectures or hyperparameters parameters, such as the loss function constants. The search for better parameters is omitted, as the goal of this study was to show the feasibility and method for creating a model which could predict short term trajectories including swing manoeuvres rather than creating the most optimal model. Further research could improve on the models presented in this thesis. An interesting suggestion would be to retrain the models using LSTM units instead of GRU. LSTMs could perform better on datasets with long sequences.

8 Conclusions

The main goal of this research was to explore the feasibility of constructing a model that could accurately predict short-term vessel trajectories in a specific area of the Port of Antwerp-Bruges. The study identified and investigated the components of such a model and developed a complete model that achieved an average displacement error of 149 meters on trajectories interacting with terminals and 146 meters on trajectories that did not interact with terminals. The findings indicate that it is possible to construct a model that can accurately predict short-term vessel trajectories using swing manoeuvre prediction and target conditioned trajectory prediction as subroutines.

Notably, the study highlights that swing manoeuvres are a crucial component of vessel trajectory prediction. Swing manoeuvres are modelled by predicting their start position, end position, and duration. The research shows that while the model can predict these factors accurately, there are relatively large duration errors during the undocking and docking phases of terminal trajectories.

The primary contribution of this research is the addition of the swing manoeuvre extension to the trajectory prediction model and the application of this model in a real-world setting. The model shows that incorporating swing manoeuvres can lead to accurate short-term trajectory predictions.

References

[1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.

[2] Alexandre Alahi, Vignesh Ramanathan, and Li Fei-Fei. Socially-aware large-scale crowd forecasting. In *Pro-*

ceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2203–2210, 2014.

[3] Danial Alizadeh, Ali Asghar Alesheikh, and Mohammad Sharif. Vessel trajectory prediction using historical automatic identification system data. *The Journal of Navigation*, 74(1):156–174, 2021.

[4] Port Of Antwerp-Bruges. Booklet with facts and figures of 2021 (antwerp branch), 2022.

[5] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.

[6] Samuele Capobianco, Leonardo M Millefiori, Nicola Forti, Paolo Braca, and Peter Willett. Deep learning methods for vessel trajectory prediction based on recurrent neural networks. *IEEE Transactions on Aerospace and Electronic Systems*, 57(6):4329–4346, 2021.

[7] Sergio Casas, Wenjie Luo, and Raquel Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *Conference on Robot Learning*, pages 947–956. PMLR, 2018.

[8] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. *arXiv preprint arXiv:1910.05449*, 2019.

[9] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3d tracking and forecasting with rich maps. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8748–8757, 2019.

[10] Chiho Choi, Srikanth Malla, Abhishek Patil, and Joon Hee Choi. Drogon: A trajectory prediction model based on intention-conditioned behavior reasoning. In *Conference on Robot Learning*, pages 49–63. PMLR, 2021.

[11] Biϕrnar R Dalsnes, Simen Hexeberg, Andreas L Flåten, Bjϕrn-Olav H Eriksen, and Edmund F Brekke. The neighbor course distribution method with gaussian mixture models for ais-based vessel trajectory prediction. In *2018 21st International Conference on Information Fusion (FUSION)*, pages 580–587. IEEE, 2018.

[12] Patrick Dendorfer, Aljosa Osep, and Laura Leal-Taixé. Goal-gan: Multimodal trajectory prediction based on goal position estimation. In *Proceedings of the Asian Conference on Computer Vision*, 2020.

[13] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996.

[14] Nicola Forti, Leonardo M Millefiori, Paolo Braca, and Peter Willett. Prediction oof vessel trajectories from ais data via sequence-to-sequence recurrent neural networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8936–8940. IEEE, 2020.

[15] Shaojun Gan, Shan Liang, Kang Li, Jing Deng, and Tingli Cheng. Ship trajectory prediction for intelligent traffic management using clustering and ann. In *2016 UKACC 11th International Conference on Control (CONTROL)*, pages 1–6. IEEE, 2016.

[16] Harshayu Girase, Haiming Gang, Srikanth Malla, Jiachen Li, Akira Kanehara, Karttikeya Mangalam, and Chiho Choi. Loki: Long term and key intentions for trajectory prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9803–9812, 2021.

[17] Michelle R Grech, Tim Horberry, and Andrew Smith. Human error in maritime operations: Analyses of accident reports using the leximancer tool. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 46, pages 1718–1721. Sage Publications Sage CA: Los Angeles, CA, 2002.

[18] Simen Hexeberg, Andreas L Flåten, Edmund F Brekke, et al. Ais-based vessel trajectory prediction. In *2017 20th International Conference on Information Fusion (Fusion)*, pages 1–8. IEEE, 2017.

[19] Chen Jiashun. A new trajectory clustering algorithm based on traclus. In *Proceedings of 2012 2nd International Conference on Computer Science and Network Technology*, pages 783–787. IEEE, 2012.

[20] Se-Won Kim, Yoon-Suk Lee, and Young-Soo Park. A study on the size of turning basin for vessels of arrival & departure in the berths. *Journal of Fisheries and Marine Sciences Education*, 24(6):872–883, 2012.

[21] J Kornacki and W Galor. Analysis of ships turn manoeuvres in port water area. *TransNav, International Journal on Marine Navigation and Safety od Sea Transportation*, 1(1), 2007.

[22] Parth Kothari, Sven Kreiss, and Alexandre Alahi. Human trajectory forecasting in crowds: A deep learning perspective. *IEEE Transactions on Intelligent Transportation Systems*, 23(7):7386–7400, 2021.

[23] Yoshiaki Kuwata, Michael T Wolf, Dimitri Zarzhitsky, and Terrance L Huntsberger. Safe maritime autonomous navigation with colregs, using velocity obstacles. *IEEE Journal of Oceanic Engineering*, 39(1):110–119, 2013.

[24] Jae-Gil Lee, Jiawei Han, and Kyu-Young Whang. Trajectory clustering: a partition-and-group framework. In *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*, pages 593–604, 2007.

[25] Jiachen Li, Fan Yang, Hengbo Ma, Srikanth Malla, Masayoshi Tomizuka, and Chiho Choi. Rain: Reinforced hybrid attention inference network for motion

forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16096–16106, 2021.

[26] Matteo Lisotto, Pasquale Coscia, and Lamberto Ballan. Social and scene-aware trajectory prediction in crowded spaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019.

[27] Bingbin Liu, Ehsan Adeli, Zhangjie Cao, Kuan-Hui Lee, Abhijeet Shenoi, Adrien Gaidon, and Juan Carlos Niebles. Spatiotemporal relationship reasoning for pedestrian intent prediction. *IEEE Robotics and Automation Letters*, 5(2):3485–3492, 2020.

[28] Hengbo Ma, Yaofeng Sun, Jiachen Li, Masayoshi Tomizuka, and Chiho Choi. Continual multi-agent interaction behavior prediction with conditional generative memory. *IEEE Robotics and Automation Letters*, 6(4):8410–8417, 2021.

[29] Srikanth Malla, Behzad Dariush, and Chiho Choi. Titan: Future forecast using action priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11186–11196, 2020.

[30] Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, Kuan-Hui Lee, Ehsan Adeli, Jitendra Malik, and Adrien Gaidon. It is not the journey but the destination: Endpoint conditioned trajectory prediction. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 759–776. Springer, 2020.

[31] Leonardo M Millefiori, Paolo Braca, Karna Bryan, and Peter Willett. Modeling vessel kinematics using a stochastic mean-reverting process for long-term prediction. *IEEE Transactions on Aerospace and Electronic Systems*, 52(5):2313–2330, 2016.

[32] Vytautas Paulauskas. Ships turning basins in ports for big container vessels. *Zeszyty Naukowe Akademii Morskiej w Szczecinie*, (20 (92):102–106, 2010.

[33] Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. Are they going to cross? a benchmark dataset and baseline for pedestrian crosswalk behavior. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 206–213, 2017.

[34] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrila, and Kai O Arras. Human motion trajectory prediction: A survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020.

[35] Agentschap voor Maritieme Dienstverlening en Kust. Afdeling Kust. Vlaamse Hydrografie. Getijtafels 2022 voor nieuwpoort, oostende, zeebrugge, vlissingen, prosperpolder, antwerpen en wintam, 2021.

[36] Greg Welch, Gary Bishop, et al. An introduction to the kalman filter. 1995.

[37] Cheng-Hong Yang, Guan-Cheng Lin, Chih-Hsien Wu, Yen-Hsien Liu, Yi-Chuan Wang, and Kuo-Chang Chen. Deep learning for vessel trajectory prediction using clustered ais data. *Mathematics*, 10(16):2936, 2022.

[38] Lan You, Siyu Xiao, Qingxi Peng, Christophe Claramunt, Xuewei Han, Zhengyi Guan, and Jiahe Zhang. St-seq2seq: A spatio-temporal feature-optimized seq2seq model for short-term vessel trajectory prediction. *IEEE Access*, 8:218565–218574, 2020.

[39] Hang Zhao, Jiyang Gao, Tian Lan, Chen Sun, Ben Sapp, Balakrishnan Varadarajan, Yue Shen, Yi Shen, Yuning Chai, Cordelia Schmid, et al. Tnt: Target-driven trajectory prediction. In *Conference on Robot Learning*, pages 895–904. PMLR, 2021.

[40] He Zhao and Richard P Wildes. Where are you heading? dynamic trajectory prediction with expert goal examples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7629–7638, 2021.

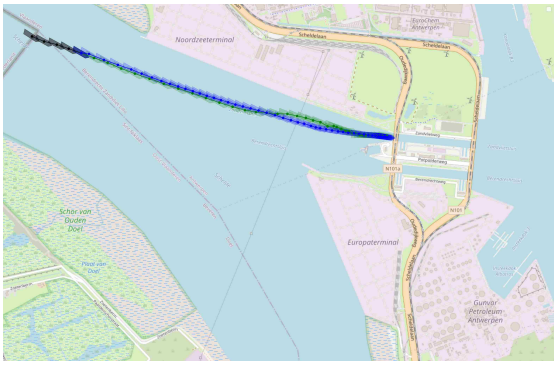

(a) NORTH to ZVS (ADE = 353 meter, FDE = 10 meter)

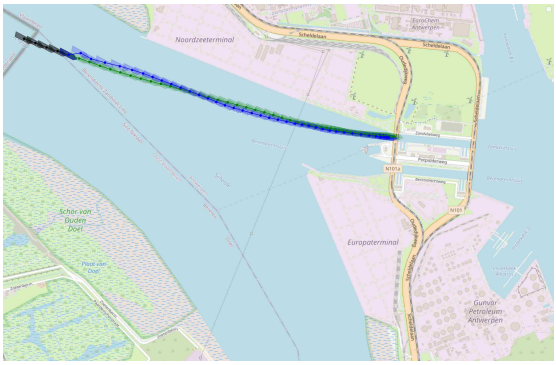

(b) NORTH to ZVS (ADE = 64 meter, FDE = 43 meter)

(c) NORTH to BES (ADE = 8 meter, FDE = 73 meter)

(d) NORTH to BES (ADE = 15 meter, FDE = 42 meter)

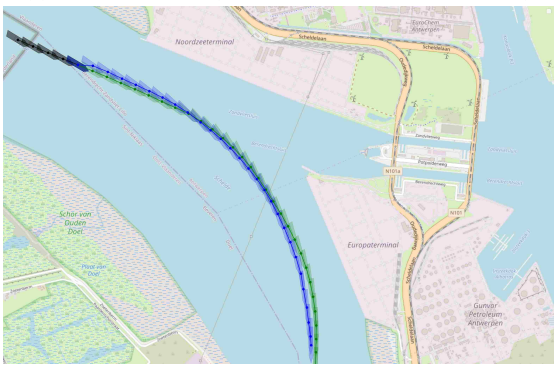

(e) NORTH to SOUTH (ADE = 112 meter, FDE = 163 meter)

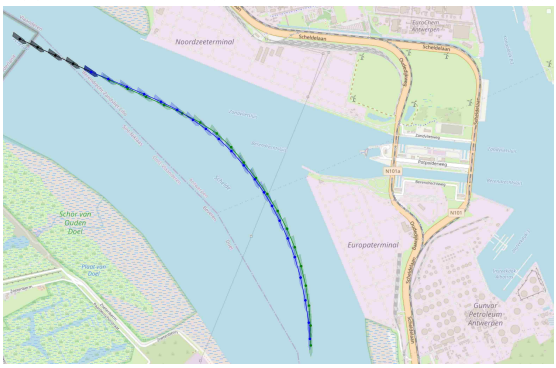

(f) NORTH to SOUTH (ADE = 277 meter, FDE = 51 meter)

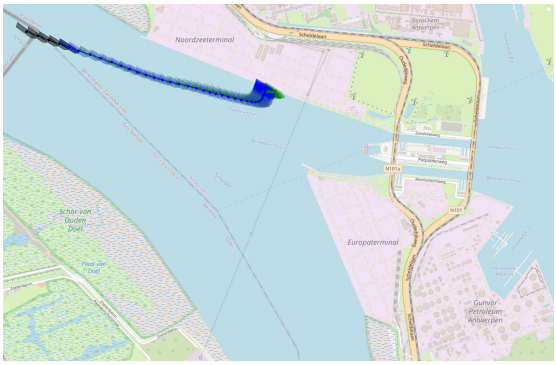

(g) NORTH to NZT, no swing (ADE = 55 meter, FDE = 83 meter)

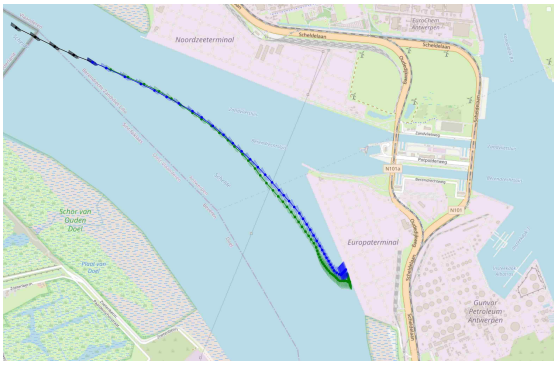

(h) NORTH to EUT, no swing (ADE = 87 meter, FDE = 89 meter)

Figure 9: Several arrival example trajectory predictions using the target conditioned trajectory prediction model. Blue points indicate the predicted trajectory, green points the ground truth and black points the historic trajectory.

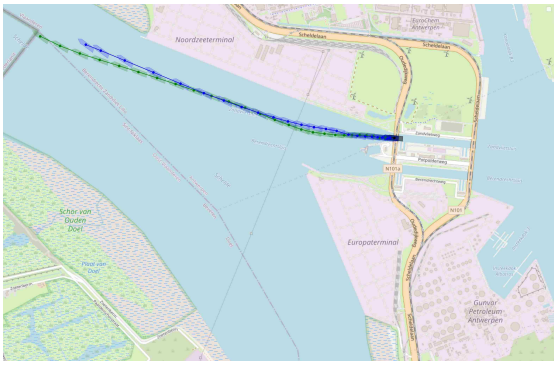

(a) ZVS to NORTH (ADE = 111 meter, FDE = 357 meter)

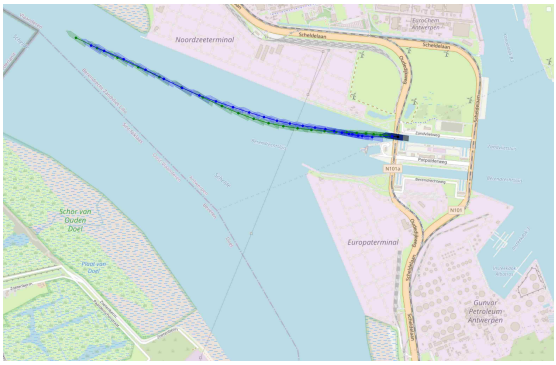

(b) ZVS to NORTH (ADE = 184 meter, FDE = 132 meter)

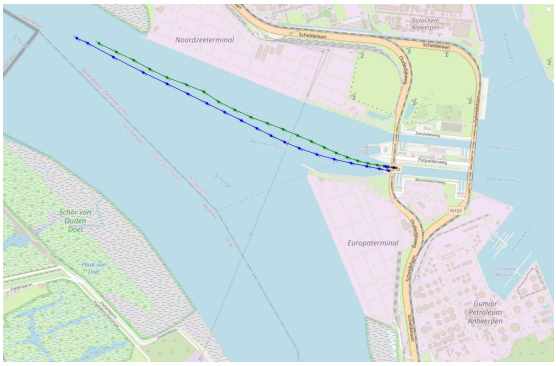

(c) BES to NORTH (ADE = 119 meter, FDE = 175 meter)

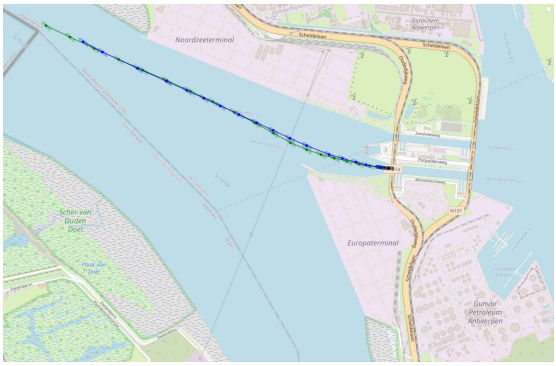

(d) BES to NORTH (ADE = 94 meter, FDE = 314 meter)

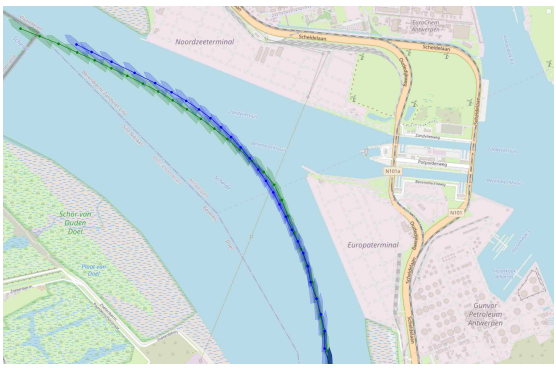

(e) SOUTH to NORTH (ADE = 315 meter, FDE = 450 meter)

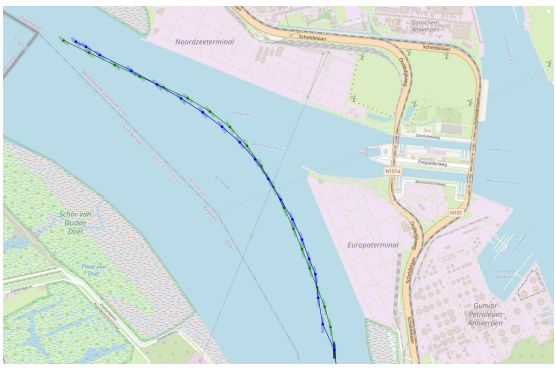

(f) SOUTH to NORTH (ADE = 127 meter, FDE = 96 meter)

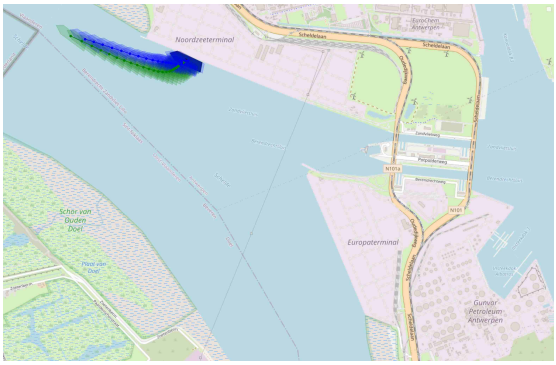

(g) NZT to NORTH, no swing (ADE = 133 meter, FDE = 114 meter)

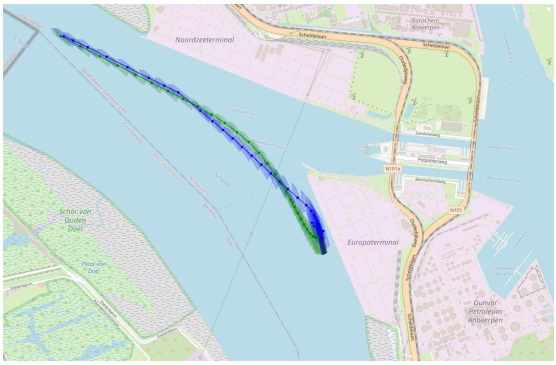

(h) EUT to NORTH, no swing (ADE = 99 meter, FDE = 23 meter)

Figure 10: Several departing example trajectory predictions using the target conditioned trajectory prediction model. Blue points indicate the predicted trajectory, green points the ground truth and black points the historic trajectory.

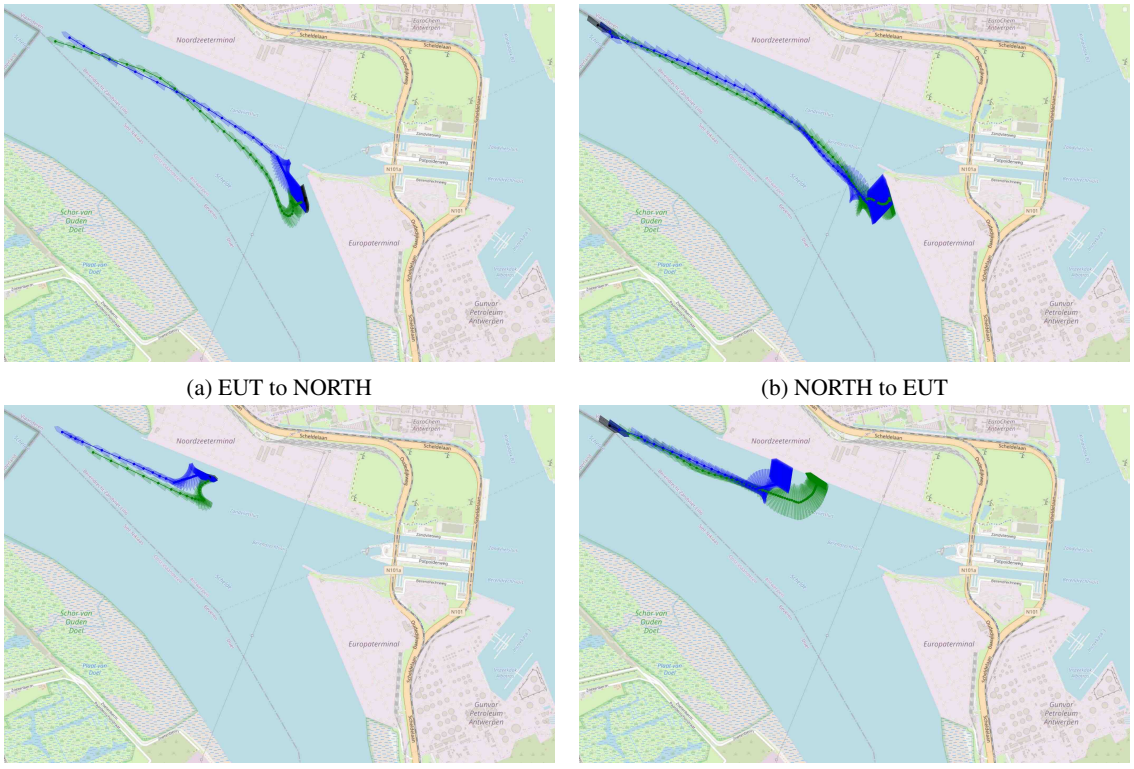

(a) EUT to NORTH

(b) NORTH to EUT

(c) NZT to NORTH

(d) NORTH to NZT

Figure 11: Several complete trajectory predictions examples which include swing manoeuvres. Blue points indicate the predicted trajectory, green points the ground truth and black points the historic trajectory.