Machine Learning with Applications 22 (2025) 100751

Contents lists available at ScienceDirect

Machine Learning with Applications

journal homepage: www.elsevier.com/locate/mlwa

Estimating vessel arrival times in global supply chains

Mathijs Pellemans [a],*, Jesper Slik [b], Sandjai Bhulai [a]

[a] Vrije Universiteit Amsterdam, De Boelelaan 1111, 1081 HV, Amsterdam, Netherlands
[b] Pon Holdings B.V., Stadionplein 28, 1076 CM, Amsterdam, Netherlands

ARTICLE INFO

Keywords:
AIS data
Vessel destination prediction
Trajectory compression
Waypoint extraction
Classification-based ETA prediction
Machine Learning

ABSTRACT

The global maritime industry, which facilitates around 90% of the world's trade, faces operational inefficiencies due to inconsistent Automatic Identification System (AIS) data and scalability challenges, limiting the reliability of current tracking systems. This study aims to improve the reliability of both Estimated Time of Arrival (ETA) and vessel destination predictions. We apply Machine Learning (ML) techniques to predict ETAs through both regression and classification models, while leveraging high-order Markov chains for destination prediction based on sequential maritime route patterns. To ensure the practical applicability of these models, we developed an end-to-end pipeline that incorporates waypoint-based trajectory compression, optimizing the handling of satellite-based AIS data (s-AIS). Using a real-world dataset of global shipping routes, our ML models, particularly Support Vector Regression (SVR), achieved a lower mean absolute error than captain-provided ETAs (16.01 vs. 22.15 h). When the time to arrival was more than 75 h, SVR outperformed captain-provided ETAs, whereas captain-provided ETAs were more accurate in the final three days before arrival, due to frequent manual updates. High-order Markov chains achieved a near-perfect accuracy of 99.00% (std: 1.19%) in destination prediction, confirming the regularity of cargo ship routes. These findings demonstrate the potential of combining ML models with Markov chains to enhance the accuracy and reliability of long-term maritime logistics forecasting, transforming raw s-AIS data into actionable insights for improved operational decision-making.

1. Introduction

1.1. Motivation

The international maritime sector facilitates around 90% of the world's trade, involving many companies participating in or depending on this sector (Pallotta et al., 2013). The container transport process involves multiple entities, including carriers, port authorities, terminal operators, hinterland transportation companies, and commercial importers, all of whom must collaborate and share information. However, the existing operational system faces challenges in transparency due to its competitive nature, with stakeholders often relying on each other for essential information and having limited visibility into the entire supply chain (Meijer, 2017; Parolas, 2016). Such fragmentation frequently results in missed transshipment connections, bottlenecks at ports, and misalignment with inland transport schedules, raising logistics costs and reducing overall reliability of global supply chains (Yazar Okur & Tuna, 2022).

Recent technological advancements indicate promising solutions to enhance transparency in the maritime sector (Viellechner, 2022). Emerging technologies such as the Internet of Things (IoT) transform supply chain management by improving data visibility and real-time stakeholder communication (Aslam et al., 2020; Grover et al., 2024; Valero et al., 2021). These technologies help overcome some limitations of the current operating systems. For example, IoT-enabled intelligent tracking technologies provide real-time container tracking, helping mitigate risks encountered during shipments (Ouedraogo et al., 2022).

Vessel tracking depends on real-time positional data, with most parties relying on the Automatic Identification System (AIS), one of the most widely utilized sources. However, AIS data often suffers from issues such as incorrect entries, missing values, inconsistent data formats, or delayed updates (Aslam et al., 2020; Yang et al., 2024). The reliability of the AIS data, particularly regarding the destination port and Estimated Time of Arrival (ETA), is frequently corrupted by manual input and incomplete information. Moreover, proprietary systems, data ownership issues, and a lack of standardized interfaces further complicate the quality of AIS data, resulting in inconsistencies that necessitate significant pre-processing before becoming useful (Aslam et al., 2020).

Satellite-based AIS data (s-AIS) is available, which extends coverage beyond coastal zones and improves overall data quality (Feng et al.,

* Corresponding author.
E-mail addresses: m.j.pellemans@vu.nl (M. Pellemans), jesper.slik@pon.com (J. Slik), s.bhulai@vu.nl (S. Bhulai).

https://doi.org/10.1016/j.mlwa.2025.100751
Received 15 July 2025; Received in revised form 28 September 2025; Accepted 6 October 2025
Available online 18 October 2025

2020), but nonautomated data entry still introduces errors, occasionally intentional due to commercial considerations (Aslam et al., 2020). This inconsistency in AIS data, combined with the lack of widespread adoption of advanced arrival prediction methods by shipping lines (Veenstra & Harmelink, 2021) continues to pose challenges for accurate and reliable ETA predictions. These technical limitations also have direct economic implications: inaccurate ETA's can worsen berth congestion, force costly rescheduling, and trigger delay penalties, while also undermining coordination with hinterland haulage and rail services (Chu et al., 2024; Qu et al., 2024).

Consequently, improving the automation and accuracy of AIS data remains critical to enhancing the reliability of vessel tracking systems and ensuring efficient port operations management.

1.2. Relevance of ETA prediction models

Vessel ETA prediction, heavily reliant on predicting a vessel's destination, is essential for optimizing port operations and improving supply chain management. Accurate ETA predictions allow ports to streamline berth allocation (Ji et al., 2023; Kolley et al., 2023), minimize delays (Pani et al., 2014; Viellechner, 2022), and increase overall terminal productivity (El Mekkaoui et al., 2022; Yu et al., 2018), facilitating better coordination between stakeholders. Machine Learning (ML) models and other Artificial Intelligence (AI) approaches can significantly improve the accuracy of ETA forecasts by learning from historical data and incorporating real-time inputs (El Mekkaoui et al., 2023).

To ensure these ETA prediction models are fully effective, reliable data inputs and integration into existing operational workflows are essential. Building a robust decision support system that meets industry demands requires improving the quality management of incoming AIS messages, extracting waypoints and routes from large datasets, and increasing accuracy by incorporating behavioral patterns (Dobrkovic et al., 2016).

Waypoint extraction. Predicting complex, long-distance vessel trajectories demands substantial computational resources (Suo et al., 2020), making real-time application difficult. Waypoint extraction addresses this challenge by identifying key locations from vessel movement observations to create a simplified representation of a vessel's path, keeping only the essential structure. This approach not only decreases data storage and computational resources, but also improves the efficiency of predictive models by focusing on the most relevant information.

Several methodologies demonstrate the effectiveness of waypoint extraction for constructing and analyzing maritime paths to improve ETA prediction. For example, the Traffic Route Extraction and Anomaly Detection (TREAD) methodology extracts waypoints to map navigation paths, which supports better anticipation of a vessel's next destination and possible routes (Pallotta et al., 2013). Building on this, both Lee et al. (2022) and Xu et al. (2022) used the Ramer-Douglas-Peucker (RDP) algorithm combined with density-based spatial clustering of applications with noise (DBSCAN) to identify vessel waypoints and connect them to create representative routes. Similarly, Karataş et al. (2021) developed a trajectory extraction algorithm inspired by TREAD, followed by DBSCAN clustering to group waypoints. Altan et al. (2023) employed Dynamic Time Warping (DTW) to measure trajectory similarity, then interpolated and identified waypoints from an average trajectory. In another approach, Dobrkovic et al. (2018) applied a Genetic Algorithm (GA) to detect waypoints. These diverse methods show how identifying waypoints improves constructing paths, leading to more accurate ETA predictions.

Although many ETA prediction techniques are in development, they have yet to see adoption in shipping operations. Practical use cases are essential in encouraging their uptake (Veenstra & Harmelink, 2021). Our study provides such a use case, contributing to this effort.

1.3. Contributions

Prior studies on vessel ETA and destination prediction face several recurring challenges, which we elaborate on further in the literature review (Section 2). AIS data is often affected by manual, error-prone reporting, which undermines the reliability of ETA broadcasts. Existing work has mainly focused on regression-based ETA prediction, while other approaches such as classification with strong operational relevance have been less explored. Much of the literature concentrates on local or regional contexts, limiting the applicability of findings to long-haul and international shipping routes. Finally, the scale and noisiness of s-AIS data create challenges for practical adoption in industry, as processing large volumes of raw AIS records is computationally demanding.

To address these issues, this study develops regression and classification models that refine the ETA inputs provided by the captain and employs Markov models to predict the vessel's next destination. The inclusion of classification for ETA bucketing also contributes to filling a research gap, as this approach has received limited attention in the existing literature. In addition, we introduce a waypoint-based trajectory compression procedure to reduce data volume ensuring that the data remains both manageable and practical for real-world industry applications. This compression preserves essential navigational information, enabling stakeholders to work with streamlined datasets.

The contributions of this study are fourfold. First, it improves the accuracy of ETA and next-destination predictions relative to existing approaches. Second, by addressing these data quality and prediction challenges, our work attempts to improve the technical accuracy of AIS-based ETA forecasts, with the potential to help mitigate costly supply chain inefficiencies. Third, another contribution of this study is the focus on identifying behavioral patterns in cargo vessel movements across international shipping routes. Whereas much of the existing literature concentrates on local patterns within ports or regional clusters, our analysis considers long-haul trajectories that span multiple regions and shipping routes, thereby offering insights into regularities in cargo vessel behavior at an international scale. Fourth, by integrating large datasets from global shipping networks, this study develops methods that improve AIS data reliability, predict vessel destinations, and provide ETA forecasts that could support operational decision-making within a decision support system.

2. Literature review

2.1. Data-driven approaches for vessel ETA prediction

Vessel ETA prediction has been extensively explored through various data-driven approaches applied to historical AIS data, especially through ML and Deep Learning (DL) methods. This review summarizes key research on these methods, case study insights, and challenges, with particular focus on ETA prediction for container vessels on a global scale. Table 1 shows a summary of studies in vessel ETA prediction, highlighting diverse methodologies, vessel types, geographical areas and prediction horizons that illustrate the field's evolution.

ETA prediction is typically framed as either a regression or classification task. Regression models predict ETA as a continuous variable, providing precise arrival estimates useful for real-time monitoring. Classification models, on the other hand, group ETA predictions into categories such as *on-time* or *delayed*, which can facilitate quick, actionable insights for operational planning (Kim et al., 2017; Yu et al., 2018). Despite the relevance of both approaches, most existing studies frame ETA prediction as a regression task, while only a few explicitly consider classification strategies such as ETA bucketing or vessel delay classes. Recent reviews confirm this imbalance, noting that classification remains comparatively underexplored (Jiang et al., 2025; Noman et al., 2025).

Table 1
Literature overview of vessel ETA prediction studies.

Study	Ship type	Geographical area	Prediction horizon	Best performing model
Regression				
Abdi and Amrit (2024)	Various incl. container	North Sea	Entire route	Deep-learning architecture (VATP)
Lloret-Batlle et al. (2024)	Container	Cross-Pacific	Entire route	XGBoost
Zhong Chu and Wang (2024)	Various incl. container	Hong Kong port	Up to 17 h, between 18 and 36 h, and greater than 36 h	Random forest
Zygouras et al. (2024)	Passenger, Container	Global	Underway short-term, Underway long-term, Entire route	Path-based clustering
Bourzak et al. (2023)	Merchant (cargo)	Saint Lawrence River, Canada	Near-future arrival times	BiLSTM
Chondrodima et al. (2023)	Various incl. container	Aegean Sea, U.S. West Coast and Brest	Up to 60 min into the future	VLF-LSTM
El Mekkaoui et al. (2020, 2022, 2023)	Tanker, Container	Jorf Lasfar port, Morocco, Mediterranean Sea	0 h to 7 days before arrival, Entire route	LSTM, NN
Kolley et al. (2023)	Container	Coast of Miami	0 to 24 h before arrival	k-NN, NN
Yoon et al. (2023)	Container	Busan New Port, South Korea	Near-future arrival times	Path-based interpolation
Valero et al. (2022)	Container	*Not reported*	Near-future arrival times	k-NN
Kwun and Bae (2021)	*Not reported*	Eastern Asia	0 to 800 km before arrival	Path-finding algorithm
Noman et al. (2021)	Container	Bremerhaven to Branschweig port	Near future arrival times	GRU
Park et al. (2021)	*Not reported*	Mediterranean Sea, South Korea	0 to 600 km before arrival	Reinforcement learning
Lechtenberg et al. (2019)	*Not reported*	*Not reported*	Up to 30 days into the future	XGBoost
Alessandrini et al. (2018)	Container	Mediterranean Sea	Up to 10 h, between 10 and 38 h, and greater than 40 h	Path-finding algorithm
Bodunov et al. (2018)	*Not reported*	Mediterranean Sea	Near-future arrival times	Ensemble model
Hardij (2018)	Tanker	Bay of Biscay to Rotterdam	Entire route	NN
Jahn and Scheidweiler (2018)	*Not reported*	German North, Baltic Sea	0 to 24 h before arrival	NN
Meijer (2017)	Various incl. container	Global	Entire route	k-NN
Parolas (2016)	Container	Global	0 h to 5 days before arrival	SVR
Classification				
Yu et al. (2018)	Container	Ningbo-Zhoushan port	1, 2, 3, 4, 6, 8 h delay	Random forest
Kim et al. (2017)	Container	Global	Entire route	Case-based reasoning

Key methodologies. Studies on vessel ETA prediction predominantly use ML and DL models, leveraging large amounts of AIS data to improve accuracy over various routes and prediction horizons. Among ML methods, XGBoost and Random Forest (RF) are frequently used. For instance, XGBoost has shown high reliability for full-route predictions on cross-Pacific container ship routes (Lloret-Batlle et al., 2024), while RF has been effective for shorter, segmented prediction intervals near Hong Kong port (Zhong Chu & Wang, 2024). Path-based approaches, such as those by Zygouras et al. (2024) and Yoon et al. (2023), capture regional and port-specific behaviors for global and short-term predictions.

DL methods, such as Long Short-Term Memory (LSTM) networks and reinforcement learning, are widely used for modeling sequential data and capturing complex temporal patterns in vessel trajectories. LSTM is particularly popular for ETA predictions across regions such as Moroccan ports (El Mekkaoui et al., 2023) and the Mediterranean Sea (El Mekkaoui et al., 2020), as well as for varying prediction horizons. More sophisticated architectures, including BiLSTM, VLF-LSTM, and attention mechanisms, have further improved prediction accuracy (Abdi & Amrit, 2024; Bourzak et al., 2023; Chondrodima et al., 2023).

While the literature highlights a shift toward exploring advanced DL models due to their higher accuracy, practical adoption remains challenging. Stakeholders often favor interpretable models, such as decision trees or simpler ML approaches, as these provide clearer insights into prediction outcomes and enable informed, actionable decision-making. This trade-off between model complexity and interpretability underscores the need to balance performance with transparency in operational settings.

Methodological differences. The prediction horizon for vessel ETA varies significantly across studies, which require distinct methodological approaches for ETA prediction depending on the timeframes involved. For longer horizons, such as multi-day or entire-route predictions on global shipping lanes, models like XGBoost (Lloret-Batlle et al., 2024) and support vector regression (SVR) (Parolas, 2016) are often employed, as they can manage complex interactions over long-range distances and incorporate multiple external factors. These models are well-suited for large-scale, high-traffic routes, where vessels encounter diverse conditions.

In contrast, shorter prediction horizons, such as those under 24 h before arrival, are particularly relevant in coastal or port areas, where trajectory data changes rapidly and predictions must adapt to the dynamic conditions near ports such as port congestion and weather variability. For instance, studies at Hong Kong port (Chen et al., 2022) and the coast of Miami (Kolley et al., 2023) utilize methods such as RF and k-Nearest Neighbors (k-NN) to efficiently capture short-term trajectory shifts. These shorter horizons typically yield more precise predictions as they approach their destination, supporting more immediate operational decisions in port logistics.

2.2. Vessel destination prediction techniques

Accurate vessel destination prediction is essential in maritime logistics as it directly affects ETA forecasts (Alessandrini et al., 2018). Industries relying on the global shipping network use destination predictions to help stakeholders make decisions about cargo and shipping schedules (Zhang et al., 2020). With growing maritime traffic, relying on human-entered AIS input becomes impractical due to variations in how captains enter destination details, such as city names, port names, or terminals, leading to inconsistencies and lack of standardization (Fu et al., 2021). These inconsistencies complicate large-scale data processing across thousands of vessels. Data-driven approaches efficiently standardize and interpret AIS data to address these challenges. The ability to predict destinations more reliably than relying on raw AIS data alone thus improves the overall reliability of AIS data (Yin et al., 2022).

Markov models, particularly n-order Markov chains, have been widely used in vessel destination prediction because of their effectiveness in capturing the sequential dependencies in ship movements.

Table 2
Summary of vessel destination prediction models from the literature, including reported accuracy, study size, and best-performing models. Direct comparison of accuracy across studies should be approached with caution, as differences in geographic coverage, dataset size, and model complexity can influence results. Smaller study areas often yield higher accuracy due to lower problem complexity, while larger coverage areas introduce greater variability and predictive challenges. The scope of our study aligns most closely with Lechtenberg et al. (2019), who also conducted a global investigation of cargo ships.

Study	Study size	Best performing model	Accuracy
Lechtenberg et al. (2019)	44 regions	Markov model	0.974
Bodunov et al. (2018)	40 ports	Voting ensemble learner	0.97
Wang et al. (2022)	128 ports	LSTM	0.87
Karataş et al. (2021)	8 ports	RF	0.86
Barreiro Lindo (2020)	55 ports	RF	0.792
Zhang et al. (2020)	9493 ports	RF	0.666
Gouareb et al. (2022)	391 ports	Graph based model	0.647
Yin et al. (2022)	9 regions	Bayesian neural network	0.562

Considering the last n ports visited, these models can capture the regularity in maritime routes, which is especially relevant for cargo ships that often follow regular, repeating routes (Kaluza et al., 2010; Pani et al., 2015; Singh et al., 2022). In the literature, Markov models have outperformed ML models such as RF in vessel destination prediction (see Table 2).

Beyond traditional models, recent studies have introduced more complex architectures such as LSTM-based hybrids and Graph Neural Networks (GNNs). LSTMs capture long-range temporal dependencies in vessel trajectories (Wang et al., 2022), while GNNs leverage the relational structure of ports to model network-wide patterns (Gouareb et al., 2022). These architectures are well suited for capturing complex dependencies, but they are typically data-hungry and computationally demanding. In contrast, n-order Markov chains are computationally efficient, fast to train, and allow direct control of historical context through the parameter n. Important to note is that a trade-off arises when increasing the order n of a Markov chain. While higher orders capture longer historical context, the number of possible n-port sequences grows exponentially, leading to sparsity in the transition matrix and an increased risk of overfitting. This issue is well studied in sequence modeling, and mitigation strategies include smoothing and back-off (Chen & Goodman, 1999; Katz, 1987), and variable-order Markov models (e.g. Begleiter et al. (2004), Dimitrakakis (2010)) that adapt their memory length to the data.

One of the advantages of Markov models in vessel destination prediction is their interpretability compared to more complex models such as LSTM networks. Markov models provide clear, probabilistic insights into how past port visits influence future destinations, making them easier to understand and implement in operational settings. This transparency is most valuable in maritime logistics, where human decision-makers benefit from interpretable models to inform their decisions (Yin et al., 2022).

3. Stakeholder insights and data description

3.1. Case validation

We interviewed diverse stakeholders in the mobility industry supply chain to assess the potential of AIS data in predicting vessel movements. Each stakeholder provided insights based on their roles and use cases. Many saw positive potential in AIS data. For instance, a maritime technology provider uses AIS transponders to determine ship locations and predict ETAs, balancing fuel efficiency with optimal arrival times.

A cycling accessories company sees potential for AIS data in predicting vessel movements in combination with Purchase Order (PO) level tracking to better handle ship and product delays. Insights into commodity flow could also help them plan their order picking more efficiently. A logistics optimization firm, working with shipping partners to optimize schedules, showed interest in predicting container rolling and blank sailings. A distribution center manager sees promise in integrating AIS data with their Enterprise Resource Planning (ERP) system for better shipment tracking at both vessel and PO levels.

Conversely, some stakeholders face challenges or were skeptical about AIS data. A power solutions provider, though having used AIS data before, has deprioritized its use due to reliability issues. A bicycle manufacturer is enthusiastic but struggles with data availability and reliability. Companies involved in port planning and maritime services acknowledge AIS data's potential but remain cautious due to similar reliability concerns. Other stakeholders were more reserved or indifferent. A cycling product manufacturer did not prioritize AIS data, and an automotive industry stakeholder did not expect significant supply chain changes from AIS data predictions.

Overall, while AIS data's potential and challenges varied among stakeholders, common concerns about communication and data reliability align with existing literature themes.

3.2. Dataset

We obtained data from an interviewed stakeholder who sources AIS information from the S&P Global Maritime Portal.[1] This dataset includes Sea-web™ and AIS Live Ship Tracking Intelligence data. *Sea-web™* provides data fields on 220,000+ ships of 100 gross tonnage and above, including detailed ship information and technical specifications under *Sea-web™ Ships* (hereafter, *ship information dataset*), and information on ship arrivals, departures, and port and berth calls under *Sea-web™ Movements* (hereafter, *port visits dataset*). *AIS Live* offers real-time satellite ship movement data. To ensure alignment with the stakeholder's ERP system data, we selected a specific sample to validate our techniques, despite having access to a larger dataset. Our sample consists of 43 international cargo container vessels, including their engine specifications. The live s-AIS data in our sample is available at 10 min intervals spanning from March 1, 2022, to August 1, 2023. For a summary of the data sources and features see Supplementary Material Appendix A.

4. Methodology

Research on ETA prediction often employs methodological approaches involving data pre-processing, destination prediction, trajectory extraction and waypoint detection, as seen in studies by Kwun and Bae (2021), Lechtenberg et al. (2019), and Yoon et al. (2023). In line with these practices, this section outlines our approach. An overview of the methodology is provided in Fig. 1, with explanations in the following subsections.

4.1. Initial analysis

To better interpret the data from our case study, Fig. 2 visualizes the pre-processed s-AIS data along with the most frequented ports in each geographical sub-region. The figure emphasizes numerous international routes.

Investigating the port visits dataset, we observe substantial missing values in the destination data, with an overall percentage of 53.05%. Missing values upon arrival reach a high percentage of 88.52%, while the percentage upon departure is lower at 17.67%. The accuracy of the destination port information varies: 30.51% of entries in the dataset are

[1] https://www.spglobal.com/marketintelligence/en/mi/products/maritime-ship-tracker-ais/-live-ship-data-seaweb.html.

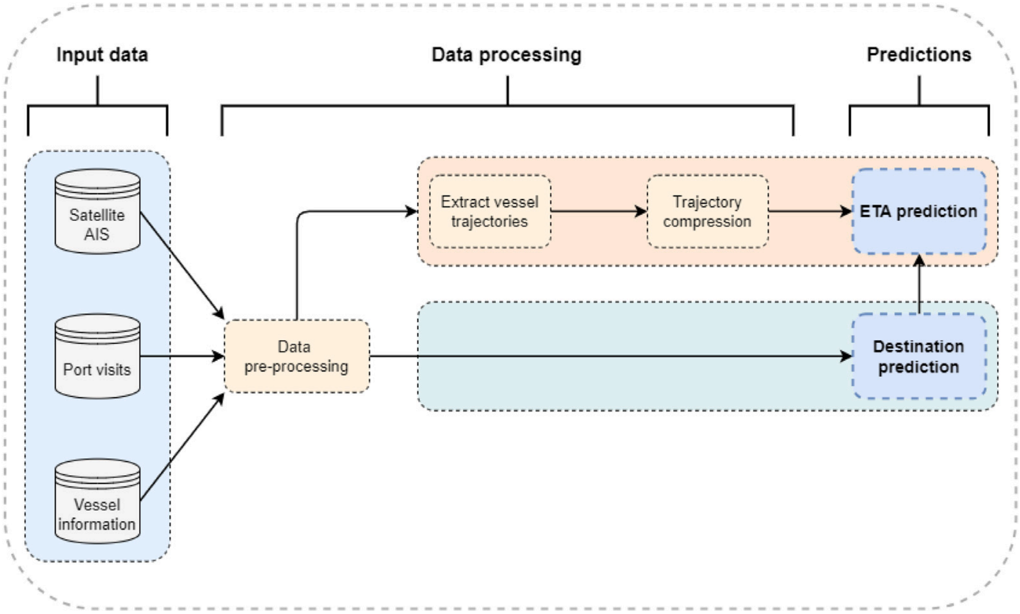

Fig. 1. Flowchart of the methodology for vessel destination and ETA prediction.

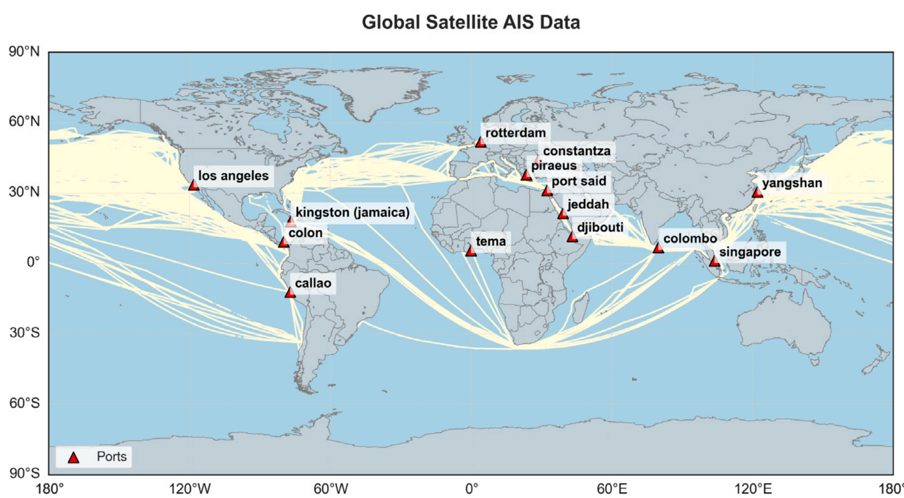

Fig. 2. Vessel trajectories from s-AIS data and most visited ports per sub-region.

correct, which drops to 6.70% upon arrival but rises to 54.25% upon departure.

In the AIS data, we observe that the destination information has a missing value percentage of 19.62% with a relatively low accuracy of 61.38% among the entries that are not missing. The ETA data has a missing value percentage of 26.65%, and questions remain about its reliability. Previous studies have shown that AIS and port-call data are prone to inconsistencies such as manual entry errors, missing transmissions, and log discrepancies (Harati-Mokhtari et al., 2007; Wang et al., 2021). To mitigate these challenges, validation approaches such as cross-checking vessel identifiers against registries, filtering implausible positions or speeds, and reconstructing missing trajectory segments through interpolation or repairing frameworks have been proposed (Liang et al., 2024; Zaman et al., 2023). In our dataset, standardized UN/LOCODE identifiers for ports are not included, and their availability would further improve the reliability of port-call validation. While our dataset has undergone provider-level curation, it shares the common limitations of large-scale AIS data, and awareness of these issues is essential.

4.2. Vessel destination prediction

In our study, we compare the performance of Markov chains in predicting vessel destinations across 120 ports distributed over four continents. To investigate how much we can learn from historical data, we use n-order Markov chains and consider the last n ports visited by a vessel.

We compute the transition probabilities in the Markov chain by analyzing historical vessel movement patterns across ports. We calculate the conditional probability for each state transition based on prior port visits. In an n-order Markov chain, the probability of transitioning from a sequence of states $S_{t-n}, \ldots, S_t$ to the next state S_{t+1} is given by the conditional probability:

$$P\left(S_{t+1} \mid S_t, S_{t-1}, \ldots, S_{t-n}\right) = \frac{C\left(S_{t-n}, \ldots, S_t, S_{t+1}\right)}{C\left(S_{t-n}, \ldots, S_t\right)} \tag{1}$$

where:

$C\left(S_{t-n}, \ldots, S_t, S_{t+1}\right)$ is the count of observed transitions from the sequence of states $S_{t-n}, \ldots, S_t, S_{t+1}$.

$C\left(S_{t-n},\ldots,S_t\right)$ is the count of the sequence $S_{t-n},\ldots,S_t$ occurring in the dataset.

The transition probabilities are thus estimated from empirical data by counting the occurrences of state sequences and transitions. We make predictions by selecting the port with the highest transition probability, ensuring the model identifies the most likely destination based on the vessel's recent history.

Given the general vulnerability of Higher-order Markov to matrix sparsity, this issue is less pronounced in our dataset and domain. Container shipping services typically operate along structured, repeating loops with limited variation in port sequences. Major liner routes follow stable rotations that are replicated across many voyages and vessels, ensuring that even high-order chains encounter sufficient recurrence of multi-port sequences to estimate reliable transition probabilities rather than isolated one-off events. These regularities were also clearly observed in our dataset, consistent with prior findings on the stability of liner shipping networks (Kaluza et al., 2010; Pani et al., 2015; Singh et al., 2022).

4.3. Trajectory extraction

Predicting ETAs using historical AIS data requires identifying port arrival and departures. Typically, trajectory extraction algorithms are needed as these events are not explicitly indicated in s-AIS streams. This study employs the *port visits dataset*, which already includes these events, making additional techniques unnecessary. Nonetheless, understanding how to extract trajectories is vital for handling incomplete, noisy, or missing port visit data and for better analyzing vessel behavior.

To extract vessel trajectories, voyage creation or voyage splitting algorithms are commonly applied (Gouareb et al., 2022; Yoon et al., 2023). These methods rely on geofencing techniques, where port boundaries are represented as port polygons - geospatial areas surrounding ports. The entry and exit of vessels from these polygons are identified by monitoring when a vessel's AIS signal is detected at these boundaries. This approach divides AIS data into distinct voyages, allowing for analysis in predictive tasks such as ETA forecasting.

We define a *route* and a *trajectory* as distinct concepts. A route R_{ij} is defined as a general sequence of coordinates (latitude and longitude) connecting two ports, P_i and P_j. A trajectory $T^v_{i,j}(t)$ is the unique, time-dependent path a vessel v travels along a route R_{ij}, capturing the specific details of its movement. Mathematically, a trajectory is expressed as:

$$T^v_{i,j}(t) = \left[(x_1,y_1,t_1),(x_2,y_2,t_2),\ldots,(x_n,y_n,t_n)\right], \tag{2}$$

where (x_k,y_k) denotes the vessel's position (latitude and longitude) at timestamp t_k. A route R_{ij} may contain multiple trajectories, as vessels traveling between the same ports can take varied paths.

To process our data, we first select only the fully available trajectories in the port visits dataset, ensuring they include departure and arrival instances along with their corresponding locations and timestamps. Next, we merge the s-AIS data with the trajectory information, retaining only those AIS data rows that fall between the port departure and arrival timestamps. We further refine our dataset by keeping only the statuses labeled *Under way using engine* and *Under way sailing*, thus removing entries related to *moored*, *anchored*, or similar statuses. The excluded statuses, though beyond this study's scope, could offer further insights into vessel behavior.

From the cleaned dataset, we build trajectories by linking consecutive departures and arrivals at various ports, assigning a unique identifier to each trajectory. For further analysis, we compute the cumulative distance using the haversine formula (see Supplementary Material Appendix B) and the cumulative travel time for each trajectory ID, along with additional features for ETA prediction. Moreover, we assign the *Not a Timestamp (NaT)* value to ETA instances where the ETA timestamp precedes the current timestamp, maintaining consistency in the temporal data. Finally, we keep only the trajectory ID's that consist of more than 10 s-AIS entries on a trajectory to ensure the robustness of our analysis.

4.4. Trajectory compression

This subsection outlines our methodology for extracting waypoints from s-AIS data streams, a key step in compressing vessel trajectories while keeping their essential navigational characteristics. We follow a multi-step methodology, detailed as follows.

Step 1: Mean Path Computation. We compute a mean path for each route R_{ij} to represent a general navigational pattern between ports P_i and P_j. This is achieved by applying the DTW algorithm (see Supplementary Material Appendix C) to align and average multiple trajectories $T^v_{i,j}(t)$ associated with the same route.

DTW is applied to all trajectory pairs $\left(T^{v1}_{i,j}(t),T^{v2}_{i,j}(t)\right)$ within the route R_{ij}. For each pair, DTW calculates an optimal alignment by finding a warping path π, which minimizes the distance between corresponding points of the two trajectories.

Once the warping paths are aligned across all trajectory pairs, the mean trajectory $\bar{T}_{i,j}$ is constructed by averaging the aligned points along the warping paths. Formally, for each time step t, and alignment π_k over K trajectory pairs, the mean trajectory is given by:

$$\bar{T}_{i,j}(t) = \frac{1}{K}\sum_{k=1}^{K}\pi_k(t), \tag{3}$$

where $\pi_k(t)$ represents the aligned points for time t in the kth pair.

DTW ensures that trajectories of varying lengths can be aligned and that local temporal misalignments are corrected through non-linear warping. In our methodology, this is complemented by removing loops (Step 2), which reduces structural noise from detours or abnormal data patterns. Together, these steps yield mean trajectories that are robust to length differences and large-scale inconsistencies. Residual issues such as positional inaccuracies, however, are not explicitly corrected here and may still influence the computed mean path. Prior studies have proposed methods to address such challenges, including interpolation-based reconstruction (Guo et al., 2021; Zaman et al., 2023), AIS trajectory cleaning (Liang et al., 2024), and trajectory repairing frameworks (Yu et al., 2025). While we do not implement these complementary approaches here, they provide promising avenues for enhancing AIS data quality in conjunction with our compression and prediction framework.

Step 2: Removing Loops. The computed mean path $\bar{T}_{i,j}$ is refined by detecting and eliminating intersecting line segments using geometric analysis. This process eliminates loops or self-intersections, resulting in a simplified trajectory $\bar{T}^{\text{simplified}}_{i,j}$. This ensures a realistic representation of a habitual route without distortions or observed manoeuvres.

Step 3: Waypoint Extraction. After removing loops, the RDP algorithm is applied to $\bar{T}^{\text{simplified}}_{i,j}$. The RDP algorithm (see Supplementary Material Appendix D) selects m waypoints, where $m \ll n$, based on a tolerance parameter ϵ controlling the level of detail. The waypoints $W^v_{i,j}(t)$ are expressed as:

$$W^v_{i,j}(t) = \text{RDP}\left(\bar{T}^{\text{simplified}}_{i,j},\epsilon\right), \tag{4}$$

yielding $W^v_{i,j}(t) = \left\{\left(x_{w_1},y_{w_1},t_{w_1}\right),\ldots,\left(x_{w_m},y_{w_m},t_{w_m}\right)\right\}$ which preserves the trajectory's structure.

Step 4: Enforcing Maximum Distance D Between Waypoints. To meet industry demands for providing consistent updates on a vessel's ETA, we make sure that no two consecutive waypoints are more than D km apart. For any pair of consecutive waypoints $\left(x_{w_k},y_{w_k}\right)$ and $\left(x_{w_{k+1}},y_{w_{k+1}}\right)$ in $W^v_{i,j}$, if their haversine distance d exceeds D, additional waypoints are iteratively inserted as follows:

- The point on the trajectory $\bar{T}_{i,j}^{\text{simplified}}$ closest to the midpoint of the segment between $\left(x_{w_k}, y_{w_k}\right)$ and $\left(x_{w_{k+1}}, y_{w_{k+1}}\right)$ is selected as a new waypoint to add.
- Then, we recalculate distances until all consecutive points in $W_{i,j}^{v}$ satisfy $d\left(\left(x_{w_k}, y_{w_k}\right), \left(x_{w_{k+1}}, y_{w_{k+1}}\right)\right) \leq D$.

This procedure reduces the number of added waypoints to a minimum while ensuring that the distance D between consecutive waypoints is not exceeded.

Step 5: Mapping Waypoints to Original Trajectory. To align the extracted waypoints with the original trajectory $T_{i,j}^{v}(t)$, we map each waypoint $\left(x_{w_k}, y_{w_k}, t_{w_k}\right) \in W_{i,j}^{v}(t)$ to its closest point in $T_{i,j}^{v}(t)$ based on the haversine distance d.

Baseline construction. Step 5 of the trajectory compression procedure maps each waypoint $\left(x_{w_k}, y_{w_k}, t_{w_k}\right)$ to its closest AIS record $\left(x_a, y_a, t_a, ETA_a\right)$ in the original trajectory $T_{i,j}^{v}(t)$. Using this mapping, we compute the baseline and true Remaining Travel Time (RTT):

$$RTT_a^{\text{base}} = ETA_a - t_{\text{port-entry}}, \qquad RTT_a^{\text{true}} = t_{\text{port-entry}} - t_a, \tag{5}$$

where ETA_a is the captain-provided ETA broadcast in the matched AIS record, t_a is the timestamp of that AIS record, and $t_{\text{port-entry}}$ is the provider-supplied port entry (arrival) timestamp. All differences are expressed in hours.

As an example, the methodology is visualized in Fig. 3, demonstrating its application to extract waypoints for vessel trajectories between Antwerp and Rotterdam. By applying the trajectory compression procedure to our dataset ($\epsilon = 0.1, D = 220$ km), we reduced the number of s-AIS entries from 1072,684 to 20,022, achieving a compression ratio of 98.13%. The tolerance parameter was set to $\epsilon = 0.1$ to ensure a high degree of simplification while retaining the core geometric structure of the routes. The maximum distance between waypoints was set to $D = 220$ km to align with industry requirements for regular ETA updates. The compressed dataset includes 1865 trajectories across 279 routes, with a total of 20,022 s-AIS records. This dataset is then used as input for training ML models to predict vessel ETAs.

Sensitivity analysis. To evaluate how the compression parameters (D, ϵ) affect data reduction, we varied $\epsilon = \{0.5, 0.1, 0.05, 0.01, 0.005, 0.001\}$ and $D = \{20, 50, 100, 220\}$. For each parameter pair, we applied the compression pipeline to the AIS trajectories and computed (i) the compression rate (percentage of data reduction compared to the uncompressed dataset) and (ii) the compressed dataset size. This analysis gives insights in the trade-off between stronger compression and the dataset size. The procedure to investigate the implications of this compression for predictive model performance is described in the following section subsection. The full results for all parameter settings are reported in Supplementary Material Appendix E.

4.5. Vessel ETA prediction

Given that the vessel's destination is determined and known, we developed and evaluated ML models for ETA prediction following using two complementary approaches: regression and classification. The regression models estimate the remaining travel time to the next destination port as a continuous variable, measured in hours, using historical s-AIS data and vessel-specific features. This fine-grained prediction is particularly useful for operational scenarios requiring precise arrival estimates to support port scheduling, resource allocation, and supply chain management.

To further align with practical maritime logistics needs, we extended the regression approach into a classification task, where the continuous RTT predictions were converted into discrete time intervals, such as 12, 24, or 72 h, up to two weeks. This transformation provides broader arrival windows, which are often more practical for port operations and cargo handling than exact arrival times. The thresholds were chosen to reflect operational planning horizons and stakeholder needs. Short-term and multi-day intervals (12–96 h) align with horizons used in prior ETA studies (Kim et al., 2017; Meijer, 2017; Parolas, 2016; Yu et al., 2018), while the weekly horizons (168–336 h) extend these to the two-week forecasts requested by stakeholders and are consistent with longer-term planning discussions in port logistics (Ji et al., 2023; Kolley et al., 2023). In this way, the ETA bucket design spans short-, medium-, and longer-term planning horizons and enables systematic evaluation of prediction accuracy across different temporal scales.

To improve predictive performance, we expanded the initial feature set with additional variables. These include cumulative metrics such as distance traveled, travel time, and average speed; the mean latitude and longitude of the arrival port; the mean historical trajectory distance for each route and IMO number; and the latitude and longitude of the previous waypoint. We applied the k-Best feature selection method to identify the 10 most statistically significant features for predicting the target variable. Retaining this subset consistently improved model accuracy while reducing computational complexity compared to using the full feature set.

To ensure a robust evaluation of model performance, the data was divided into training, validation, and test sets using an 80/10/10 ratio, stratified by route to preserve balanced representation. Splits were generated chronologically, keeping waypoints from the same trajectory within a single subset to prevent data leakage. The training set consists of 16,137 rows from 1491 trajectories across 273 routes. The validation set, used for hyperparameter tuning, includes 2058 rows from 186 trajectories spanning 103 routes. Finally, the test set, reserved for an unbiased performance evaluation, consists of 1827 rows from 187 trajectories over 96 routes.

The ML models employed in this study include Quantile Regression (QR), SVR, RF, XGBoost, and LightGBM. Hyperparameters were optimized using grid search, and model validation was performed through five-fold cross-validation. Hyperparameter configurations for all models are detailed in the Supplementary Material Appendix F. The performance of these models is evaluated against the baseline, where the baseline RTT is defined as the difference between the captain-provided ETA broadcast and the actual port-entry time (see Section 4.4). This provides a benchmark for assessing whether the models improve upon the real-time ETA information available to operators in practice.

We assessed the regression model performance using Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Devation (MAD) and the Coefficient of Determination (R^2). We evaluated the classification models using accuracy, precision, and recall to measure how effectively they categorized vessel arrivals within defined time intervals. To quantify uncertainty in the performance metrics, we applied a bootstrapping procedure resampling the test set with replacement ($n = 2000$) to compute mean values and standard deviations for the regression error metrics. To further assess whether observed differences between models were statistically significant, we performed pairwise non-parametric Wilcoxon signed-rank tests on the paired error distributions for each metric ($p < 0.05$). To account for multiple comparisons, Holm-Bonferroni correction was applied.

In addition to performance evaluation, we conducted a permutation importance analysis to assess the relative influence of the selected features. This approach measures the increase in prediction error (MAE) when the values of a single feature are randomly permuted, thereby quantifying its contribution to model performance.

Runtimes and scalability. To analyze computational trade-offs between models, we benchmarked training and inference at increasing training sizes, using the compressed dataset fixed at $\epsilon = 0.1$ and $D = 220$. We computed the training time (wall-clock time for model fitting on each training subset) and batch inference time (time to predict the full test set in a single call) from which we derive the throughput (samples/s).

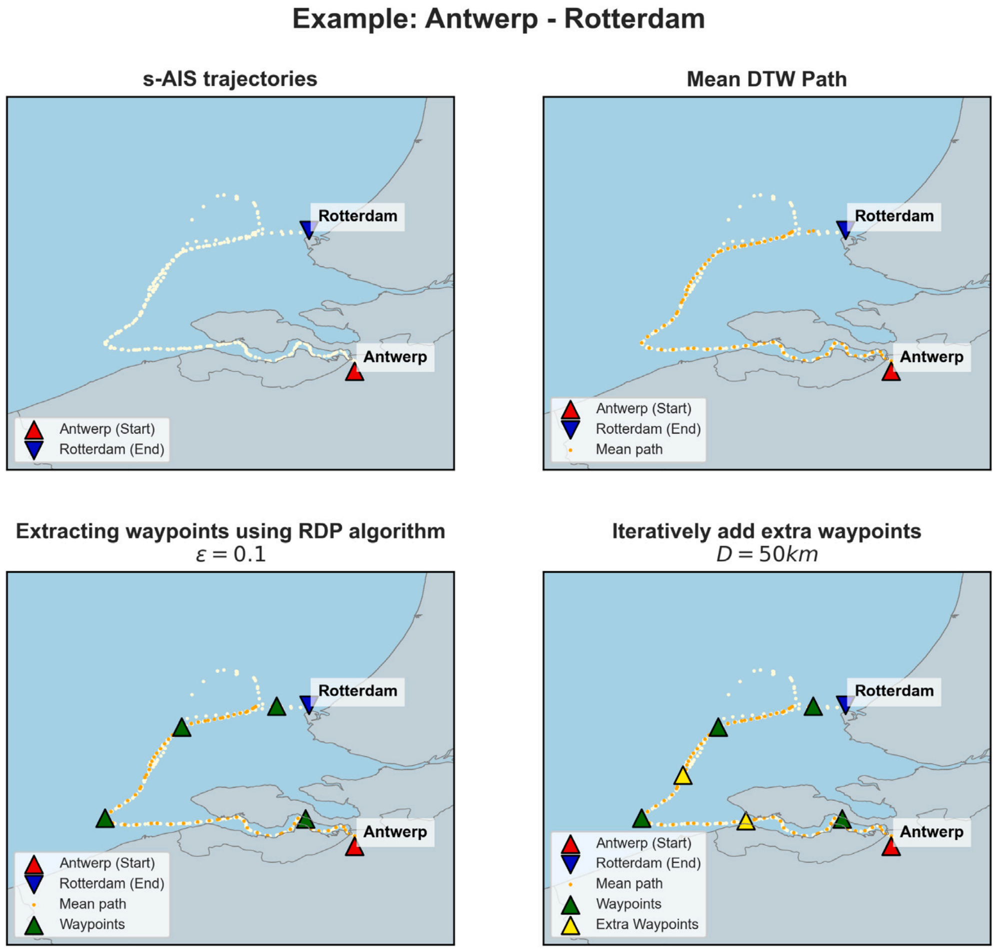

Fig. 3. Visualization of the waypoint extraction methodology applied to the Antwerp - Rotterdam route. The s-AIS trajectories illustrate the variability in paths taken by vessels. The computed mean trajectory represents a generalized path for the route based on the observed trajectories. Extracted waypoints (green triangles) are extracted using the RDP algorithm to compress the mean trajectory. Additional waypoints (yellow triangles) are inserted to make sure that the distance between consecutive waypoints does not exceed a predefined threshold D aligning with operational requirements for regular ETA updates. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

We evaluated scaling at training fractions $(0.05, 0.10, 0.25, 0.50, 1.00)$. Each experiment was repeated three times to mitigate system variability. Predictive performance metrics are reported on the test set with 95% bootstrap confidence intervals. This design allows us to investigate the cost of periodic retraining (training time scaling) from the demands of real-time decision support (inference throughput).

4.6. Computational resources

We conducted all experiments on a Windows laptop with a 12th Generation Intel® Core™ i7-12800H processor, operating at 2400 MHz with 14 cores. We computed the results using Python 3.11 with a x64 GCC C/C++ compiler to improve processing speed.

5. Results

In this section, we present an evaluation and performance analysis of the models used to predict a vessel's next destination and ETA.

5.1. Destination prediction

During the study period, vessels visited an average of 61 ports, with the number of ports per vessel ranging from 38 to 95. As shown in Fig. 4 and Supplementary Material Appendix G, the performance of n-order Markov models ($n = 1, 2, \dots, 10$) improves progressively with larger n, both in terms of mean accuracy and reduced variability. The baseline model, which indicates the AIS-declared destination, achieves a mean accuracy of 53.48% (std: 11.71%). Incorporating historical data from the two most recent ports ($n = 2$) increases the mean accuracy to 89.34% (std: 3.86%), reflecting a significant improvement of 67.05% over the baseline. By leveraging a history of up to 10 ports ($n = 10$), the model achieves a near-perfect mean accuracy of 99.00% (std: 1.19%), marking an 85.12% improvement over the baseline.

Each successive increase in n demonstrates a statistically significant improvement over the preceding model, as shown by paired t-tests with a p-value < 0.05. Despite the high overall accuracy, vessel-specific behavior and route patterns still affect prediction performance.

5.2. ETA prediction

Table 3 summarizes the regression outcomes for ETA prediction, comparing the ML models to the baseline. The baseline demonstrates poor performance, with an RMSE of 159.36 h, an MAE of 22.15 h, and an R^2 value of 0.28. The relatively low MAD of 3.17 h indicates that predictions are typically near the median ETA, though outliers significantly affect the overall performance of the baseline. All ML models outperform the baseline's results. Particularly, QR and SVR yield the most favorable outcomes, with RMSE values around 27 h, an MAE of 16 h, and an R^2 of 0.96, reflecting both high precision and explanatory power. Meanwhile, RF, XGBoost, and LightGBM also demonstrate

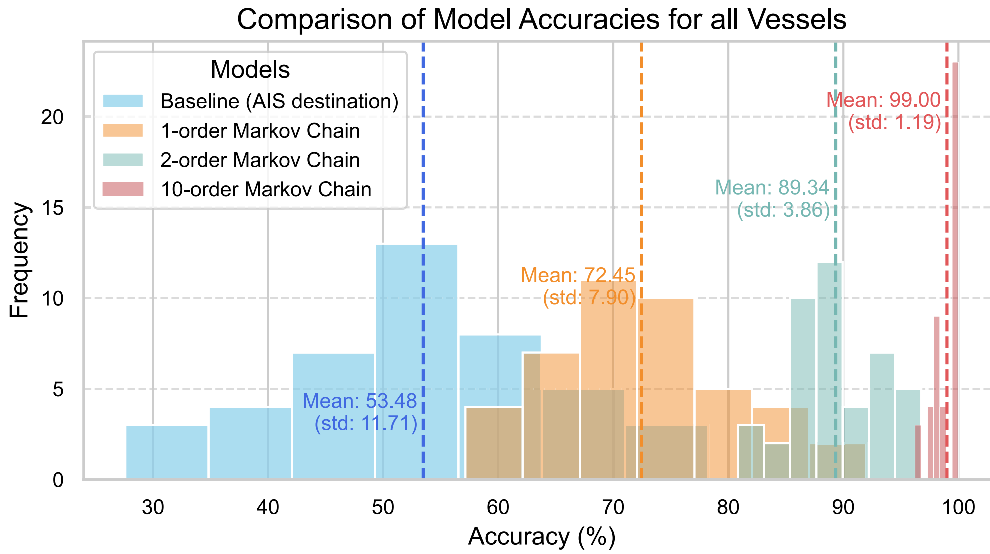

Fig. 4. Performance comparison of n-Order Markov Models and the baseline approach for vessel destination prediction: accuracy and variability.

Table 3
Overall performance metrics for ETA prediction models with their standard deviations between brackets.

Model	RMSE (h)	MAE (h)	MAD (h)	R^2
Baseline	159.36	22.15	3.17	0.28
Quantile Regression[a]	27.02 (0.89)	16.56 (0.50)	8.14 (0.30)	0.96 (2e−3)
SVR[a]	27.15 (1.04)	16.01 (0.51)	7.61 (0.30)	0.96 (2e−3)
Random Forest[b]	40.30 (1.75)	23.36 (0.77)	10.44 (0.38)	0.92 (8e−3)
XGBoost[b]	39.96 (1.64)	23.90 (0.74)	11.90 (0.42)	0.92 (1e−2)
LightGBM[b]	37.51 (2.13)	21.60 (0.74)	10.80 (0.37)	0.93 (9e−3)

[a,b] Models with the same superscript letter are not significantly different. Group (a) performed significantly better than group (b) across all error metrics (Wilcoxon signed-rank test with Holm-Bonferroni correction, $p < 0.05$).

reliable performance, showing strong predictive capabilities, although slightly less accurate than QR and SVR. The pairwise Wilcoxon signed-rank tests confirm that QR and SVR do not differ significantly from one another, but both perform significantly better than RF, XGBoost, and LightGBM across all metrics ($p < 0.05$). Furthermore, variability is the lowest for QR and SVR, indicating stable performance across resamples, while the ensemble models show wider error distributions.

An examination of errors shows notable variability in model performance across different trajectories, with the size of errors heavily affected by both the length of the journey and the path complexity (see Supplementary Material Appendix H). Longer trips, such as those from Oakland to Busan and Savannah to Qingdao, display increased error variability, with deviations exceeding 300 h early in the journey. This greater dispersion might sometimes be attributed to vessel captains delaying updates to the initial ETA. Nevertheless, this pattern does not hold consistently for all trajectories, as the baseline ETA often yields accurate forecasts, especially on shorter, less complex paths. For shorter routes, such as Savannah to Charleston, there is a marked reduction in error variability with more consistent predictions across different models. In particular, the baseline model shows superior performance, frequently outperforming the SVR model, especially during the later stages of the trajectory where error variance tends to stabilize.

When the outcomes of the regression analysis are mapped to a classification problem, a distinct shift in performance is observed between the best-performing ML model and the baseline ETA. The SVR model demonstrates superior predictive accuracy compared to the baseline up to roughly 74.82 h prior to arrival, indicating enhanced performance in the earlier phases of a vessel's journey (see Fig. 5). However, in the final three days leading up to arrival, the baseline ETA outperforms all prediction models, showing an improvement in accuracy as the vessels near their destination. For a summary of all classification performance results, see Supplementary Material Appendix I.

5.3. Sensitivity analysis

The sensitivity analysis demonstrates a systematic relationship between the compression parameters and the resulting dataset size. Both the RDP tolerance ϵ and the maximum waypoint distance D influence compression, but their effects differ in magnitude. The tolerance parameter is the dominant factor; increasing ϵ from 0.001 to 0.5 raises the compression rate from roughly 65%–75% to more than 99%, effectively reducing the dataset from nearly 400k points to fewer than 10k. By contrast, changes in D primarily fine-tune the compression achieved at a fixed ϵ with larger distances consistently leading to smaller compressed datasets.

Fig. 6 illustrates that compression increases almost monotonically with ϵ accross all D values. Furthermore, extreme parameter combinations ($\epsilon = 0.5$, $D = 220$ eliminates most of the data, whereas strict tolerances retain a much larger fraction of points. Since the choice of compression parameters directly determines the size of the dataset, it also has important consequences for model scalability, as training and inference runtimes increase with the number of data points.

5.4. Runtimes & scalability

All models improve as the training fraction increases, but the shape of the gains differs (see Fig. 7). SVR shows the most pronounced and consistent reductions in error with more data and achieves the best performance across all metrics at the largest fractions. QR is already competitive at small fractions and changes comparatively little as data increases, indicating early stability. XGBoost, LightGBM and RF improve steadily from small to moderate fractions. For R^2, SVR and QR reach higher levels at larger fractions, while the tree-based models continue to climb with additional data but remain at comparatively lower levels; variability across splits is modest at larger fractions for all models.

Training time increased with dataset size for all models, though the rate of increase varied significantly (see Fig. 8). SVR and QR scaled the steepest, requiring substantially longer training times as the fraction grew. XGBoost, LightGBM and RF trained fastest, with training times remaining comparatively low even at the largest fraction. Inference throughput remained stable across training fractions for all models

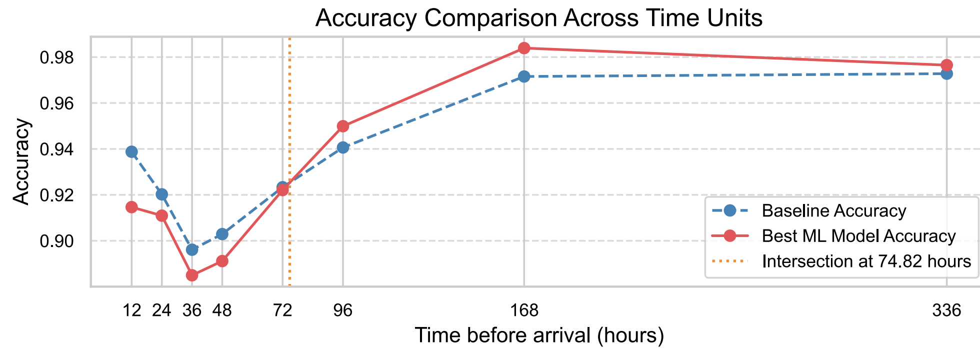

Fig. 5. Accuracy comparison between the baseline ETA and the best ML model across different time-segments before arrival.

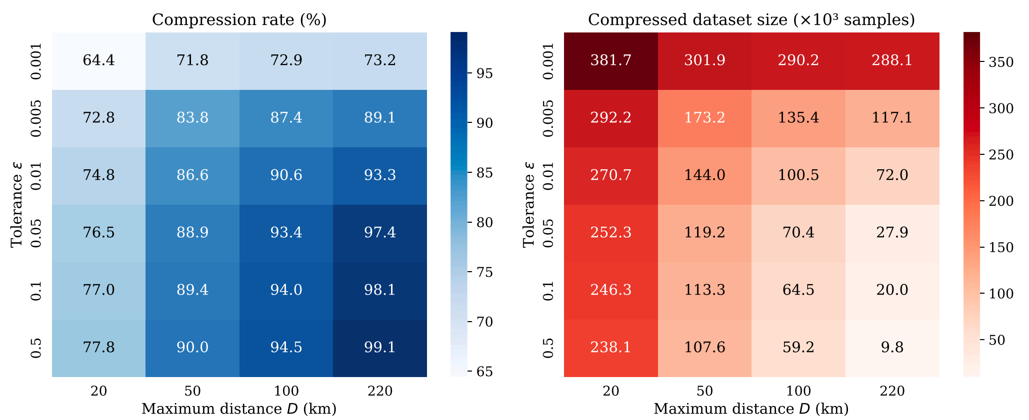

Fig. 6. Compression rate (%) and resulting dataset size ($\times 10^3$ samples) for different settings of the RDP tolerance (ϵ) and maximum waypoint distance (D).

except SVR, indicating that inference cost is not strongly dependent on training set size for these models. For SVR, the inference throughput also increased as the training size increased.

5.5. Model interpretability

Fig. 9 presents the permutation importance plots for the prediction models. Across all models, voyage progress variables are the most influential. In particular, *mean_distance_route_imo* and *cumulative_distance* consistently show the largest increase in MAE when permuted, confirming that spatial progress along the route is the strongest determinant of ETA accuracy. A secondary cluster of predictors, including *mean_lon_arrival* and *longitude (lon)*, also accounted for a noticeable share of predictive power, while *draught* and *speed over ground (SOG)* contributed at a smaller scale.

Model families differ in how broadly they distributed importance. Ensemble methods (RF, XGBoost, LightGBM) attributed weight to a wider set of contextual features, including *cumulative_travel_time* as well as vessel characteristics such as *draught* and *SOG*. QR and SVR relied almost exclusively on distance-based variables, with QR showing the sharpest concentration on route progress. This contrast highlights complementary strategies: ensembles emphasize broader operational context, while regression-based models emphasize stable predictive signals from spatial progress.

6. Conclusion and future work

6.1. Interpretation of the results

This study successfully validates the effectiveness of ML techniques, particularly SVR and QR, in improving the reliability and accuracy of vessel ETA predictions using historical s-AIS data. The results demonstrate that ML models excel in predicting ETAs in early-voyage predictions, while the baseline ETA, benefiting from frequent manual updates, becomes increasingly reliable as vessels near their destinations. This pattern is further supported by the classification results, where ML models outperform the baseline up to 75 h before arrival, while the baseline ETA dominates in the final three days of a voyage. This *three-day threshold* is particularly insightful as it aligns with the common maritime practice where vessel captains provide their final, high-accuracy ETA to port authorities approximately 72 h before arrival. In this sense, this crossover point should be interpreted as an overall pattern rather than a universal threshold. For instance, in the Savannah-Charleston case presented in the supplementary material, no crossover occurs: baseline captain-provided ETAs outperform SVR predictions throughout the full voyage, illustrating that threshold dynamics may depend on specific operational settings. Regional or seasonal operating conditions could likewise affect threshold behavior.

This implies that our models' primary operational advantage is for strategic, long-range planning. For supply chain managers, cargo owners, and logistics planners, receiving a more accurate ETA forecast weeks in advance is critical for optimizing downstream operations, managing inventory, and arranging hinterland transportation. The findings highlight the complementary strengths of both approaches and suggest opportunities for hybrid modeling frameworks to optimize predictions across different voyage phases. A practical implementation of such a hybrid framework could follow a simple rule-based strategy driven by the predicted RTT. For horizons beyond ~75 h, the system would present the ML forecast. For horizons below ~75 h, it would defer to the captain's ETA; and within an overlap band (e.g., 72–84 h), a smooth blending of the two could prevent abrupt switching. To illustrate, the following pseudo-logic (Algorithm 1) sketches how this could be operationalized in a Decision Support System (DSS). In practice, such a system should also apply temporal smoothing to reduce fluctuations

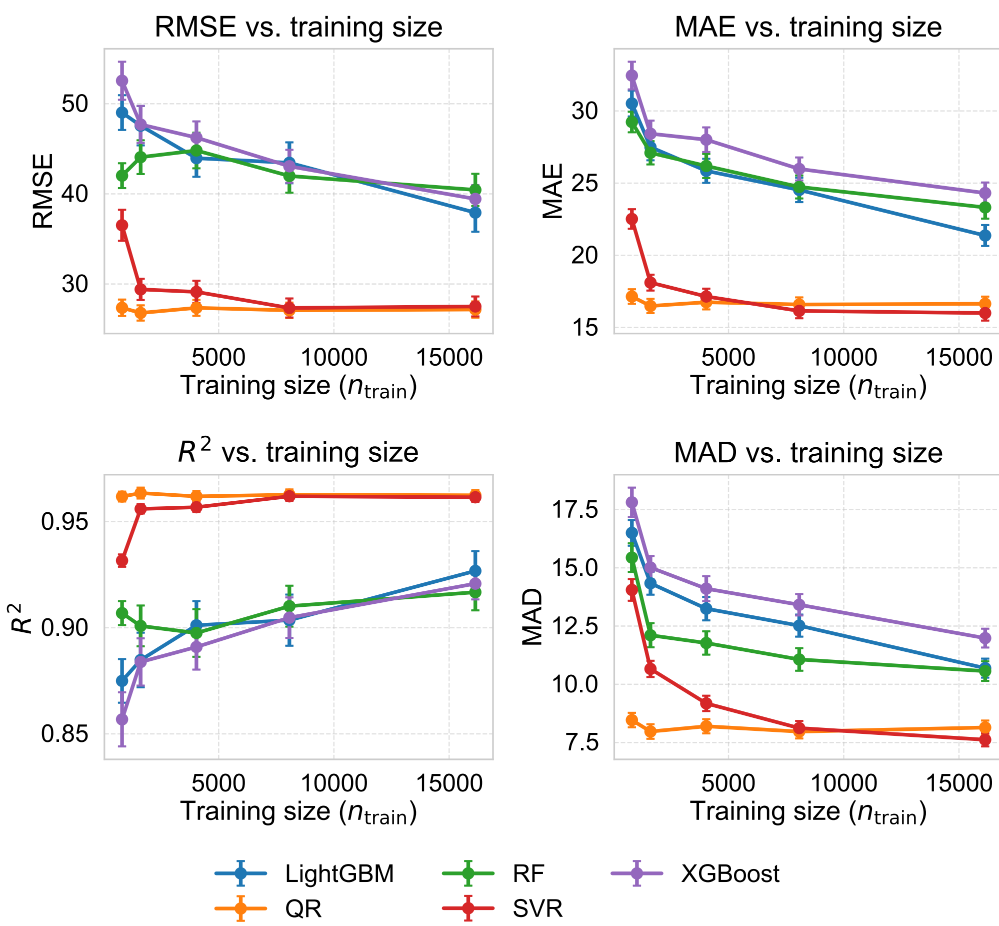

Fig. 7. Predictive performance metrics (RMSE, MAE, R^2, MAD) with bootstrapped 95% confidence intervals as a function of training set size.

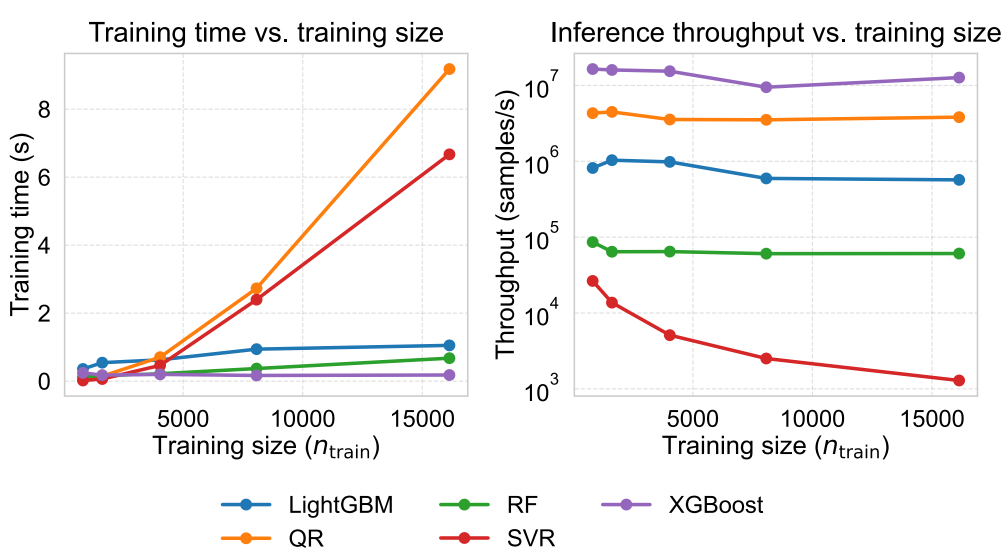

Fig. 8. Training time (left) and inference throughput (right, log scale) of all models as a function of training set size.

from small forecast updates, attach uncertainty bands so users can judge the reliability of each estimate, and trigger alerts when large deviations occur. These measures ensure that hybrid forecasts are more stable, interpretable, and actionable for planners.

The sensitivity analysis shows how the choice of compression parameters directly shapes dataset size and, consequently, computational demands. Higher tolerance values yield dramatic reductions in trajectory points, easing storage and processing requirements but at the risk of discarding relevant detail. In operational terms, this means that compression settings must be chosen to balance efficiency with predictive fidelity: overly strict compression inflates runtimes unnecessarily, while overly aggressive compression may undermine model accuracy. These

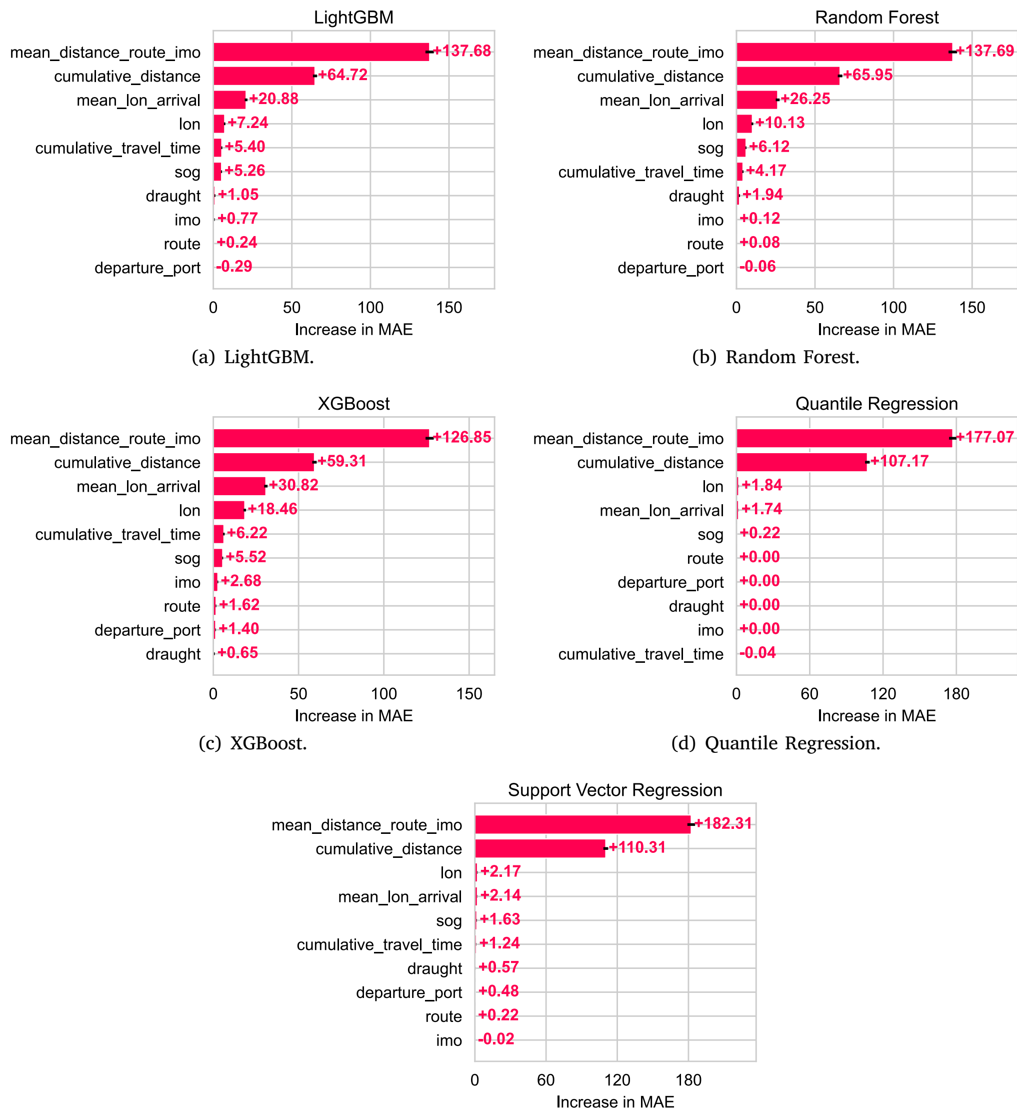

(a) LightGBM.

(b) Random Forest.

(c) XGBoost.

(d) Quantile Regression.

(e) Support Vector Regression.

Fig. 9. Permutation importance plots for the prediction models.

trade-offs connect directly to the scalability results, where dataset size proved a key driver of training and inference costs.

The runtime and scalability experiments highlight a trade-off between predictive accuracy and computational efficiency with direct implications for industry practice. SVR achieved the strongest predictive performance but at the cost of substantially longer training and slower inference. This makes high-accuracy models such as SVR and QR more suitable for periodic offline retraining or scenarios where precision outweighs latency. In contrast, boosted-tree models offer fast predictions at scale, making them well suited for real-time decision-support systems where sacrifices in accuracy are acceptable.

When contextualizing our findings, a direct numerical comparison of performance metrics with prior studies is challenging due to variations in datasets, geographical scope, and vessel types. However, the performance of our best models is competitive to for example Parolas (2016) whose SVR was the top-performing model for predicting the ETA for ships, including container vessels, on global routes. Additionally, the use of high-order Markov chains for destination prediction achieves near-perfect accuracy, aligning with prior research (Lechtenberg et al., 2019) and reinforcing the reliability of sequential models in identifying destination ports. The analysis of the dataset reveals consistent patterns across the international routes sailed by the studied cargo vessels, providing insights into recurring shipping behaviors and contributing to the improvement of s-AIS data quality.

A likely explanation for the consistently stronger performance of SVR and QR relates to the specific nature of the RTT target variable. Reported ETAs and s-AIS records contain noise from manual reporting, and further variability stems from unobserved influences such as weather, port congestion, and routing practices. Because a portion of the RTT variance cannot be explained by the available features, ensemble models are prone to fitting residual noise. In contrast, SVR benefits from the robustness provided by kernel regularization (Xu et al., 2009), while QR is resilient to heteroskedasticity and outliers (Zhao et al., 2015). These robustness properties might explain why SVR and QR delivered superior and more stable performance compared to the ensemble models. The interpretability analysis reinforces this view: models that concentrate on voyage progress appear less exposed to noisy or weakly informative signals, while ensembles distribute importance across a broader feature set, some of which may capture residual noise. This distinction has practical relevance: regression-based

Algorithm 1: Hybrid ETA forecast selection combining ML predictions and broadcasted captain ETAs. Here, ETA_{ML} is derived from RTT_{ML} as t_{now} + RTT_{ML}. The linear weight w ensures a smooth transition in the overlap window (e.g., 72–84 h). While a simple linear rule is shown here, more adaptive strategies could assign weights dynamically based on factors such as voyage duration, vessel type, or route complexity.

Input: ML-predicted ETA (ETA_{ML}), Broadcasted captain ETA (ETA_{CAP})
Output: Published ETA forecast
if *captain ETA is invalid* **then**
  publish ← ETA_{ML}
end
else if $ETA_{CAP} - t_{now} \leq 72\ h$ **then**
  publish ← ETA_{CAP}
end
else if $RTT_{ML} \geq 84\ h$ **then**
  publish ← ETA_{ML}
end
else
  Compute weight $w = \dfrac{RTT_{ML} - 72}{12}$ `// Linear transition between 72 h-84 h`
  publish ← $w \cdot ETA_{ML} + (1 - w) \cdot ETA_{CAP}$
end

models may be more reliable in settings where data quality is uneven, whereas ensembles could offer advantages when richer and more reliable operational features are available.

This research extends the current body of literature by addressing the full prediction pipeline, from raw s-AIS data processing to actionable insights, in contrast to prior studies that primarily focused on isolated components such as waypoint detection (Altan et al., 2023; Lee et al., 2022) or model optimization (Yoon et al., 2023). While our evaluation metrics focus on accuracy, overall s-AIS data reliability is improved by developing a pipeline that mitigates issues from inconsistent s-AIS data streams, such as high volumes of missing destination entries. The waypoint detection method developed in this study enables the creation of efficient, streamlined datasets that retain practical applicability for industry implementation. Furthermore, this study's global scope - spanning both short-haul regional and long-haul international routes - distinguishes it from regionally focused works (Li et al., 2023; Zhang et al., 2020), offering broader insights into shipping behavior and operational complexities. Finally, the integration of classification predictions alongside the regression approach adds a layer of operational utility, facilitating the development of practical tools for port resource allocation and voyage planning, an area less frequently explored in maritime literature.

6.2. Strengths and limitations

There are several strengths and limitations to mention. A key strength of this study lies in its end-to-end approach, which generates predictive outputs that can support operational decisions in maritime logistics and serve as components of a broader decision support system. The waypoint detection method represents an advancement, facilitating efficient trajectory simplification while preserving essential navigational information. This method ensures that the dataset remains manageable for computational modeling, making the approach suitable for real-world implementation where data efficiency is critical. Another notable strength is the ability of the proposed models to capture patterns across diverse international routes represented in the dataset, supporting their applicability beyond purely local contexts and indicating potential for broader generalization.

Several limitations must be acknowledged. First, the generalizability of the ETA prediction models could be further improved through validation on larger and more heterogeneous datasets. The current study is based on a sample of 43 container vessels, which, while demonstrating consistent route patterns, may not capture the full spectrum of operational variability. Applying the models to a more diverse fleet, including different vessel types (e.g., tankers, bulk carriers) or carriers with less regular schedules, would be essential for confirming their broader applicability. A further limitation concerns the Markov chain approach and its reliance on sufficient historical recurrence of port sequences. While the regularity of liner shipping networks in our dataset mitigated sparsity and allowed high-order chains to perform robustly, this strength may not generalize to smaller fleets or more irregular services. In such cases, higher-order models risk overfitting by effectively memorizing rare sequences. Second, local factors such as port-specific congestion or weather conditions were not explicitly modeled, potentially limiting the precision of the predictions in highly localized settings. Addressing these limitations will help refine the framework and improve its practical utility in real-world maritime operations. Despite these limitations, this study lays a strong foundation for future research.

6.3. Future directions

Future research could focus on several key areas. One promising direction is exploring the graph structure of global maritime networks. This approach can uncover both local and global shipping patterns, providing insights into maritime route connectivity, bottlenecks, and vulnerabilities in the global supply chain infrastructure. Another priority is validating models on larger, more diverse datasets. Expanding the dataset scope to include under-documented routes, newly constructed ports, and vessels with limited historical trajectories would improve model robustness and adaptability. Future work could therefore also explore smoothing and variable-order strategies to ensure broader applicability of the Markov model beyond highly structured liner networks.

Future work could investigate methods to strengthen trajectory representation, including interpolation-based reconstruction, trajectory repairing, and denoising techniques that enhance AIS data quality prior to prediction. Generative AI approaches, such as autoencoders and diffusion models, offer additional potential to advance trajectory modeling, supporting more robust forecasting and anomaly detection under noisy or incomplete s-AIS data.

Hybrid models that integrate ML predictions with real-time captain updates could optimize ETA accuracy across all voyage stages. While Algorithm 1 provides a basic sketch of how a hybrid ETA framework could be operationalized, its implementation and validation in real-world DSS settings remain to be developed. Future work could focus on extending this rule-based prototype into a fully operational system, exploring adaptive weighting strategies and integration with industry workflows.

Refining waypoint detection by optimizing the epsilon parameter could balance trajectory simplification and data retention. Destination prediction could be refined by incorporating real-time features like vessel heading and speed at departure, especially for multi-destination voyages.

Another promising avenue is integrating ETA predictions within multi-modal transport networks. Accurate and timely ETA estimates can facilitate coordinated scheduling of cargo transfers across maritime, rail, and road transport. Finally, future work may prioritize improving supply chain visibility by focusing on tracking containerized cargo independently of the vessel, as improved container tracking can streamline transshipment operations.

Abbreviations

AI	Artificial Intelligence
AIS	Automatic Identification System
DBSCAN	Density-Based Spatial Clustering of Applications with Noise
DSS	Decision Support System
DL	Deep Learning
DTW	Dynamic Time Warping
ERP	Enterprise Resource Planning
ETA	Estimated Time of Arrival
GA	Genetic Algorithm
GNN	Graph Neural Network
IMO	International Maritime Organization
IoT	Internet of Things
k-NN	k-Nearest Neighbors
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
ML	Machine Learning
PO	Purchase Order
RDP	Ramer-Douglas-Peucker
QR	Quantile Regression
RF	Random Forest
RMSE	Root Mean Squared Error
RTT	Remaining Travel Time
s-AIS	Satellite-based Automatic Identification System
SOG	Speed Over Ground
SVR	Support Vector Regression
TREAD	Traffic Route Extraction and Anomaly Detection

CRediT authorship contribution statement

Mathijs Pellemans: conceptualization, data curation, Formal analysis, investigation, methodology, project administration, Resources, Software, validation, visualization, Writing – original draft. **Jesper Slik:** conceptualization, data curation, Methodology, Supervision, Writing – review & editing. **Sandjai Bhulai:** conceptualization, Supervision, Writing – review & editing.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used GPT-4o in order to improve the readability and language of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We would like to thank Laith Wehbi and Yasmin Roshandel for their valuable and insightful discussions, which greatly contributed to the development of this work.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.mlwa.2025.100751.

Data availability

The authors do not have permission to share data.

References

Abdi, A., & Amrit, C. (2024). Enhancing vessel arrival time prediction: A fusion-based deep learning approach. *Expert Systems with Applications*, *252*, Article 123988. http://dx.doi.org/10.1016/j.eswa.2024.123988.

Alessandrini, A., Mazzarella, F., & Vespe, M. (2018). Estimated time of arrival using historical vessel tracking data. *IEEE Transactions on Intelligent Transportation Systems*, *PP*, 1–9. http://dx.doi.org/10.1109/TITS.2017.2789279.

Altan, D., Marijan, D., & Kholodna, T. (2023). SafeWay: Improving the safety of autonomous waypoint detection in maritime using transformer and interpolation. *Maritime Transport Research*, *4*, Article 100086. http://dx.doi.org/10.1016/j.martra.2023.100086.

Aslam, S., Michaelides, M. P., & Herodotou, H. (2020). Internet of ships: A survey on architectures, emerging applications, and challenges. *IEEE Internet of Things Journal*, *7*(10), 9714–9727. http://dx.doi.org/10.1109/JIOT.2020.2993411.

Barreiro Lindo, F. (2020). *Predicting ships' path and destination port based on AIS information* [Ph.D. thesis], Haute école de gestion de Genève, https://sonar.ch/global/documents/315176.

Begleiter, R., El-Yaniv, R., & Yona, G. (2004). On prediction using variable order Markov models. *Journal of Artificial Intelligence Research*, *22*(1), 385–421. http://dx.doi.org/10.5555/1622487.1622499.

Bodunov, O., Schmidt, F., Martin, A., Brito, A., & Fetzer, C. (2018). Real-time destination and ETA prediction for maritime traffic. In *Proceedings of the 12th ACM international conference on distributed and event-based systems* (pp. 198–201). New York, NY, USA: Association for Computing Machinery, http://dx.doi.org/10.1145/3210284.3220502.

Bourzak, I., El Mekkaoui, S., Berrado, A., Caron, S., & Benabbou, L. (2023). Deep learning approaches for vessel estimated time of arrival prediction: A case study on the saint lawrence river. In *2023 14th international conference on intelligent systems: theories and applications* (pp. 1–7). http://dx.doi.org/10.1109/SITA60746.2023.10373707.

Chen, S. F., & Goodman, J. (1999). An empirical study of smoothing techniques for language modeling. *Computer Speech & Language*, *13*(4), 359–394. http://dx.doi.org/10.1006/csla.1999.0128.

Chen, X., Wei, C., Zhou, G., Wu, H., Wang, Z., & Biancardo, S. A. (2022). Automatic identification system (AIS) data supported ship trajectory prediction and analysis via a deep learning model. *Journal of Marine Science and Engineering*, *10*(9), 1314. http://dx.doi.org/10.3390/jmse10091314.

Chondrodima, E., Pelekis, N., Pikrakis, A., & Theodoridis, Y. (2023). An efficient LSTM neural network-based framework for vessel location forecasting. *IEEE Transactions on Intelligent Transportation Systems*, *24*(5), 4872–4888. http://dx.doi.org/10.1109/TITS.2023.3247993.

Chu, Z., Yan, R., & Wang, S. (2024). Vessel turnaround time prediction: A machine learning approach. *Ocean & Coastal Management*, *249*, Article 107021. http://dx.doi.org/10.1016/j.ocecoaman.2024.107021.

Dimitrakakis, C. (2010). Bayesian variable order Markov models. In Y. W. Teh, & M. Titterington (Eds.), *Proceedings of machine learning research*: *vol. 9*, *Proceedings of the thirteenth international conference on artificial intelligence and statistics* (pp. 161–168). Chia Laguna Resort, Sardinia, Italy: PMLR, https://proceedings.mlr.press/v9/dimitrakakis10a.html.

Dobrkovic, A., Iacob, M.-E., & van Hillegersberg, J. (2018). Maritime pattern extraction and route reconstruction from incomplete AIS data. *International Journal of Data Science and Analytics*, *5*(2), 111–136. http://dx.doi.org/10.1007/s41060-017-0092-8.

Dobrkovic, A., Iacob, M.-E., van Hillegersberg, J., Mes, M. R. K., & Glandrup, M. (2016). Towards an approach for long term AIS-based prediction of vessel arrival times. In H. Zijm, M. Klumpp, U. Clausen, & M. t. Hompel (Eds.), *Logistics and supply chain innovation: bridging the gap between theory and practice* (pp. 281–294). Cham: Springer International Publishing, ISBN: 978-3-319-22288-2, http://dx.doi.org/10.1007/978-3-319-22288-2_16.

El Mekkaoui, S., Benabbou, L., & Berrado, A. (2020). Predicting ships estimated time of arrival based on AIS data. In *Proceedings of the 13th international conference on intelligent systems: theories and applications* (pp. 1–6). New York, NY, USA: Association for Computing Machinery, ISBN: 9781450377331, http://dx.doi.org/10.1145/3419604.3419768.

El Mekkaoui, S., Benabbou, L., & Berrado, A. (2022). Machine learning models for efficient port terminal operations: Case of vessels' arrival times prediction. *IFAC-PapersOnLine*, *55*(10), 3172–3177. http://dx.doi.org/10.1016/j.ifacol.2022.10.217.

El Mekkaoui, S., Benabbou, L., & Berrado, A. (2023). Deep learning models for vessel's ETA prediction: Bulk ports perspective. *Flexible Services and Manufacturing Journal*, *35*(1), 5–28. http://dx.doi.org/10.1007/s10696-022-09471-w.

Feng, M., Shaw, S.-L., Peng, G., & Fang, Z. (2020). Time efficiency assessment of ship movements in maritime ports: A case study of two ports based on AIS data. *Journal of Transport Geography*, *86*, Article 102741. http://dx.doi.org/10.1016/j.jtrangeo.2020.102741.

Fu, X., Xiao, Z., Xu, H., Jayaraman, V., Othman, N. B., Chua, C. P., & Lind, M. (2021). Ais data analytics for intelligent maritime surveillance systems. In *Progress in IS*, *Maritime informatics* (pp. 393–411). Springer, http://dx.doi.org/10.1007/978-3-030-50892-0.

Gouareb, R., Can, F., Ferdowsi, S., & Teodoro, D. (2022). Vessel destination prediction using a graph-based machine learning model. In P. Ribeiro, F. Silva, J. F. Mendes, & R. Laureano (Eds.), *Network science* (pp. 80–93). Cham: Springer International Publishing, http://dx.doi.org/10.1007/978-3-030-97240-0_7.

Grover, D. V., Balusamy, D. B. B., Milanova, D. M., & Felix, D. A. Y. (Eds.), (2024). *Blockchain, IoT, and AI technologies for supply chain management: apply emerging technologies to address and improve supply chain management*. Berkeley, CA: A Press, http://dx.doi.org/10.1007/979-8-8688-0315-4.

Guo, S., Mou, J., Chen, L., & Chen, P. (2021). Improved kinematic interpolation for AIS trajectory reconstruction. *Ocean Engineering*, *234*, Article 109256. http://dx.doi.org/10.1016/j.oceaneng.2021.109256.

Harati-Mokhtari, A., Wall, A., Brooks, P., & Wang, J. (2007). Automatic identification system (AIS): Data reliability and human error implications. *Journal of Navigation*, *60*(3), 373–389. http://dx.doi.org/10.1017/S0373463307004298.

Hardij, R. (2018). *Predicting arrival times for tankers ships using recurrent neural networks* [Master's thesis], Tilburg School of Humanities and Digital Sciences, https://arno.uvt.nl/show.cgi?fid=147103.

Jahn, C., & Scheidweiler, T. (2018). Port call optimization by estimating ships' time of arrival. In *Dynamics in logistics* (pp. 172–177). Cham: Springer International Publishing.

Ji, B., Huang, H., & Yu, S. S. (2023). An enhanced NSGA-II for solving berth allocation and quay crane assignment problem with stochastic arrival times. *IEEE Transactions on Intelligent Transportation Systems*, *24*(1), 459–473. http://dx.doi.org/10.1109/TITS.2022.3213834.

Jiang, S., Liu, L., Peng, P., Xu, M., & Yan, R. (2025). Prediction of vessel arrival time to port: a review of current studies. *Maritime Policy & Management*, 1–26. http://dx.doi.org/10.1080/03088839.2025.2488376.

Kaluza, P., Kölzsch, A., Gastner, M. T., & Blasius, B. (2010). The complex network of global cargo ship movements. *Journal of the Royal Society Interface*, *7*(48), 1093–1103. http://dx.doi.org/10.1098/rsif.2009.0495.

Karataş, G. B., Karagoz, P., & Ayran, O. (2021). Trajectory pattern extraction and anomaly detection for maritime vessels. *Internet of Things*, *16*, Article 100436. http://dx.doi.org/10.1016/j.iot.2021.100436.

Katz, S. (1987). Estimation of probabilities from sparse data for the language model component of a speech recognizer. *IEEE Transactions on Acoustics, Speech and Signal Processing*, *35*(3), 400–401. http://dx.doi.org/10.1109/TASSP.1987.1165125.

Kim, S., Kim, H., & Park, Y. (2017). Early detection of vessel delays using combined historical and real-time information. *Journal of the Operational Research Society*, *68*(2), 182–191. http://dx.doi.org/10.1057/s41274-016-0104-4.

Kolley, L., Rückert, N., Kastner, M., Jahn, C., & Fischer, K. (2023). Robust berth scheduling using machine learning for vessel arrival time prediction. *Flexible Services and Manufacturing Journal*, *35*(1), 29–69. http://dx.doi.org/10.1007/s10696-022-09462-x.

Kwun, H., & Bae, H. (2021). Prediction of vessel arrival time using auto identification system data. *International Journal of Innovative Computing, Information and Control*, *17*(2), 725–734. http://dx.doi.org/10.24507/ijicic.17.02.725.

Lechtenberg, S., de Siqueira Braga, D., & Hellingrath, B. (2019). Automatic identification system (AIS) data based ship-supply forecasting. In *Digital transformation in maritime and city logistics: smart solutions for logistics. Proceedings of the hamburg international conference of logistics (HICL), vol. 28* (pp. 3–24). Berlin: epubli GmbH, http://dx.doi.org/10.15480/882.2487.

Lee, J.-S., Lee, H.-T., & Cho, I.-S. (2022). Maritime traffic route detection framework based on statistical density analysis from ais data using a clustering algorithm. *IEEE Access*, *10*, 23355–23366. http://dx.doi.org/10.1109/ACCESS.2022.3154363.

Li, H., Jiao, H., & Yang, Z. (2023). AIS data-driven ship trajectory prediction modelling and analysis based on machine learning and deep learning methods. *Transportation Research Part E: Logistics and Transportation Review*, *175*, Article 103152. http://dx.doi.org/10.1016/j.tre.2023.103152.

Liang, M., Su, J., Liu, R. W., & Lam, J. S. L. (2024). AISClean: AIS data-driven vessel trajectory reconstruction under uncertain conditions. *Ocean Engineering*, *306*, Article 117987. http://dx.doi.org/10.1016/j.oceaneng.2024.117987.

Lloret-Batlle, R., Lin, S., & Guo, J. (2024). Cross-Pacific vessel estimated time of arrival and next destination prediction with automatic identification system data. *Transportation Research Record*, Article 03611981241275551. http://dx.doi.org/10.1177/03611981241275551.

Meijer, R. (2017). *ETA prediction: Predicting the ETA of a container vessel based on route identification using AIS data* [Master's thesis], Delft University of Technology, https://resolver.tudelft.nl/uuid:cba0ef59-dd23-49aa-91d5-bed239e27395.

Noman, A. A., Heuermann, A., Wiesner, S. A., & Thoben, K.-D. (2021). Towards data-driven GRU based ETA prediction approach for vessels on both inland natural and artificial waterways. In *2021 IEEE international intelligent transportation systems conference* (pp. 2286–2291). http://dx.doi.org/10.1109/ITSC48978.2021.9564883.

Noman, A. A., Heuermann, A., Wiesner, S., & Thoben, K.-D. (2025). A review of vessel time of arrival prediction on waterway networks: Current trends, open issues, and future directions. *Computers*, *14*(2), http://dx.doi.org/10.3390/computers14020041.

Ouedraogo, C., Montarnal, A., & Gourc, D. (2022). Multimodal container ttransportation ttraceability and supply chain risk management: A review of methods and solutions. *International Journal of Supply and Operations Management*, *9*(2), 212–234. http://dx.doi.org/10.22034/ijsom.2022.109139.2201.

Pallotta, G., Vespe, M., & Bryan, K. (2013). Vessel pattern knowledge discovery from AIS data: A framework for anomaly detection and route prediction. *Entropy*, *15*(6), 2218–2245. http://dx.doi.org/10.3390/e15062218.

Pani, C., Fadda, P., Fancello, G., Frigau, L., & Mola, F. (2014). A data mining approach to forecast late arrivals in a transhipment container terminal. *Transport*, *29*(2), 175–184. http://dx.doi.org/10.3846/16484142.2014.930714.

Pani, C., Vanelslander, T., Fancello, G., & Cannas, M. (2015). Prediction of late/early arrivals in container terminals – a qualitative approach. *European Journal of Transport and Infrastructure Research*, *15*(4), http://dx.doi.org/10.18757/ejtir.2015.15.4.3096.

Park, K., Sim, S., & Bae, H. (2021). Vessel estimated time of arrival prediction system based on a path-finding algorithm. *Maritime Transport Research*, *2*, Article 100012. http://dx.doi.org/10.1016/j.martra.2021.100012.

Parolas, I. (2016). *ETA prediction for containerships at the Port of Rotterdam using Machine Learning Techniques* [Master's thesis], TU Delft, https://repository.tudelft.nl/islandora/object/uuid%3A9e95d11f-35ba-4a12-8b34-d137c0a4261d.

Qu, H., Wang, X., Meng, L., & Han, C. (2024). Liner schedule design under port congestion: A container handling efficiency selection mechanism. *Journal of Marine Science and Engineering*, *12*(6), http://dx.doi.org/10.3390/jmse12060951.

Singh, S. K., Fowdur, J. S., Gawlikowski, J., & Medina, D. (2022). Leveraging graph and deep learning uncertainties to detect anomalous maritime trajectories. *IEEE Transactions on Intelligent Transportation Systems*, *23*(12), 23488–23502. http://dx.doi.org/10.1109/TITS.2022.3190834.

Suo, Y., Chen, W., Claramunt, C., & Yang, S. (2020). A ship trajectory prediction framework based on a recurrent neural network. *Sensors (Basel, Switzerland)*, *20*(18), 5133. http://dx.doi.org/10.3390/s20185133.

Valero, C. I., Ivancos Pla, E., Vano, R., Garro, E., Boronat, F., & Palau, C. E. (2021). Design and Development of an AIoT architecture for introducing a vessel ETA cognitive service in a legacy port management solution. *Sensors*, *21*(23), 8133. http://dx.doi.org/10.3390/s21238133.

Valero, C. I., Martinez, Á., Oltra-Badenes, R., Gil, H., Boronat, F., & Palau, C. E. (2022). Prediction of the estimated time of arrival of container ships on short-sea shipping: A pragmatical analysis. *IEEE Latin America Transactions*, *20*(11), 2354–2362. http://dx.doi.org/10.1109/TLA.2022.9904760.

Veenstra, A., & Harmelink, R. (2021). On the quality of ship arrival predictions. *Maritime Economics & Logistics*, *23*(4), 655–673. http://dx.doi.org/10.1057/s41278-021-00187-6.

Viellechner, A. M. (2022). *The new era of predictive analytics in container shipping and air cargo* [Ph.D. thesis], WHU - Otto Beisheim School of Management, https://opus4.kobv.de/opus4-whu/frontdoor/index/index/docId/899.

Wang, S., Bao, Z., Culpepper, J. S., & Cong, G. (2021). A survey on trajectory data management, analytics, and learning. *ACM Computing Surveys*, *54*(2), http://dx.doi.org/10.1145/3440207.

Wang, W., Bin, J., Zaji, A., Halldearn, R., Guillaume, F., Li, E., & Liu, Z. (2022). A multi-task learning-based framework for global maritime trajectory and destination prediction with AIS data. *Maritime Transport Research*, *3*, Article 100072. http://dx.doi.org/10.1016/j.martra.2022.100072.

Xu, H., Caramanis, C., & Mannor, S. (2009). Robustness and regularization of support vector machines. *Journal of Machine Learning Research*, *10*, 1485–1510.

Xu, X., Liu, C., Li, J., & Miao, Y. (2022). Trajectory clustering for SVR-based Time of Arrival estimation. *Ocean Engineering*, *259*, Article 111930. http://dx.doi.org/10.1016/j.oceaneng.2022.111930.

Yang, Y., Liu, Y., Li, G., Zhang, Z., & Liu, Y. (2024). Harnessing the power of machine learning for AIS data-driven maritime research: A comprehensive review. *Transportation Research Part E: Logistics and Transportation Review*, *183*, Article 103426. http://dx.doi.org/10.1016/j.tre.2024.103426.

Yazar Okur, İ. G., & Tuna, O. (2022). Schedule reliability in liner shipping: A study on global shipping lines. *Scientific Journal of Maritime Research (Pomorstvo)*, *36*(2), 389–400. http://dx.doi.org/10.31217/p.36.2.22.

Yin, Z., Yang, D., & Bai, X. (2022). Vessel destination prediction: A stacking approach. *Transportation Research Part C: Emerging Technologies*, *145*, Article 103951. http://dx.doi.org/10.1016/j.trc.2022.103951.

Yoon, J.-H., Kim, D.-H., Yun, S.-W., Kim, H.-J., & Kim, S. (2023). Enhancing container vessel arrival time prediction through past voyage route modeling: A case study of busan new port. *Journal of Marine Science and Engineering*, *11*(6), 1234. http://dx.doi.org/10.3390/jmse11061234.

Yu, C., Jiang, Z., Zhang, X., He, W., & Zhong, C. (2025). A novel trajectory repairing model based on the artificial potential field-enhanced A* algorithm for Small Coastal vessels. *Journal of Marine Science and Engineering*, *13*(7), http://dx.doi.org/10.3390/jmse13071200.

Yu, J., Tang, G., Song, X., Yu, X., Qi, Y., Li, D., & Zhang, Y. (2018). Ship arrival prediction and its value on daily container terminal operation. *Ocean Engineering*, *157*, 73–86. http://dx.doi.org/10.1016/j.oceaneng.2018.03.038.

Zaman, B., Marijan, D., & Kholodna, T. (2023). Interpolation-based inference of vessel trajectory waypoints from sparse AIS data in maritime. *Journal of Marine Science and Engineering*, *11*(3), http://dx.doi.org/10.3390/jmse11030615.

Zhang, C., Bin, J., Wang, W., Peng, X., Wang, R., Halldearn, R., & Liu, Z. (2020). AIS data driven general vessel destination prediction: A random forest based approach. *Transportation Research Part C: Emerging Technologies*, *118*, Article 102729. http://dx.doi.org/10.1016/j.trc.2020.102729.

Zhao, T., Kolar, M., & Liu, H. (2015). A general framework for robust testing and confidence regions in high-dimensional quantile regression. https://arxiv.org/abs/1412.8724.

Zhong Chu, R. Y., & Wang, S. (2024). Evaluation and prediction of punctuality of vessel arrival at port: A case study of Hong Kong. *Maritime Policy & Management*, *51*(6), 1096–1124. http://dx.doi.org/10.1080/03088839.2023.2217168.

Zygouras, N., Troupiotis-Kapeliaris, A., & Zissis, D. (2024). EnvClus*: Extracting common pathways for effective vessel trajectory forecasting. *IEEE Access*, *12*, 3860–3873. http://dx.doi.org/10.1109/ACCESS.2023.3349081.