TURUN YLIOPISTO
UNIVERSITY OF TURKU

DIGITAL MARITIME MONITORING: ENHANCING SITUATIONAL AWARENESS IN SHIPPING TRAFFIC USING AI-BASED MODELS

Farshad Farahnakian

TURUN YLIOPISTON JULKAISUJA – ANNALES UNIVERSITATIS TURKUENSIS
PUBLICATION SERIES INFORMATION SHOULD GO HERE

University of Turku

Faculty of Technology
Department of Computing
Doctor of Science Information and Communication Technology
Doctoral Programme in Technology (DPT)

Supervised by

Professor, Jukka Heikkonen
Department of Computing
Faculty of Technology
University of Turku (UTU)

University Teacher, Dr C. Raduly-Baka
Department of Computing
Faculty of Technology
University of Turku (UTU)

Reviewed by

Professor, Edmund Førland Brekke
Department of Engineering Cybernetics
Norwegian University of Science and Technology (NTNU)

Professor, Roberto Setola
Department of Engineering
Università Campus Bio-Medico di Roma (UCBM)

Opponent

Professor, Ari Visa
Department of Computing Sciences
Faculty of Information Technology and Communication Sciences
Tampere University

The originality of this publication has been checked in accordance with the University of Turku quality assurance system using the Turnitin OriginalityCheck service.

ISBN 978-952-02-0526-3 (PRINT)
ISBN 978-952-02-0527-0 (PDF)
ISSN 2736-9390 (PRINT)
ISSN 2736-9684 (ONLINE)
Grano Oy, Lemminkäisenkatu 50, 20520, Turku, Finland

I dedicate this doctoral thesis to my parents, Akbar Farahnakian and Fakhri Beigi, for their unwavering love and support throughout my life. This achievement wouldn't be possible without your sacrifice.

UNIVERSITY OF TURKU
Faculty of Technology
Department of Computing
Doctor of Science Information and Communication Technology
FARAHNAKIAN, FARSHAD: Digital Maritime Monitoring: Enhancing Situational Awareness in Shipping Traffic using AI-based Models
Doctoral dissertation, 180 pp.
Doctoral Programme in Technology (DPT)
February 2026

ABSTRACT

Maritime Situational Awareness (MSA) is crucial for the safety and security of maritime operations. It depends on advanced sensing technologies, such as the Automatic Identification System (AIS), which enables continuous tracking data. However, the large volume of AIS messages transmitted by numerous vessels poses challenges for traditional analysis methods, which are often inefficient and unable to provide real-time monitoring and accurate insights. To address these challenges, it is necessary to integrate digital services equipped with AI-based predictive models and data-driven solutions to automate maritime monitoring tasks.

This thesis presents the development and evaluation of two key AI-driven services: (1) Ship Abnormal Behavior Detection (ShABD) and (2) Ship Movement Prediction (ShMP). The ShABD service targets a range of crucial maritime scenarios, including detecting dark ships, identifying unexpected changes in movement, recognizing spiral maneuvers, and uncovering potential smuggling activities. Since these behaviors threaten maritime security, effective, data-driven approaches are critical for timely detection and response. To address these threats, this work uses clustering-based algorithms, thereby eliminating the need to label millions of AIS observations manually. For the ShMP service, a variety of advanced Machine Learning (ML) approaches are investigated, including sequence-to-sequence models for short-term forecasting and similarity-based measures for predicting long-term trajectories. The design of the ShMP frameworks enables two types of predictions, each catering to different stakeholder requirements. For example, port authorities benefit from comprehensive forecasts of ship movements to streamline loading and unloading operations, while coast guards require accurate short-term predictions to support timely and effective interventions.

The research presented in this thesis utilizes historical AIS data from ships operating in the Baltic Sea between 2022 and 2024. For data, two reliable sources were used: (1) HELCOM and (2) Digitraffic APIs. The main findings can be summarized as follows: K-Means outperformed DBSCAN, Gaussian Mixtures, and Affinity Propagation in detecting dark ships and spiral maneuvers, with an average silhouette coefficient of 0.755. The use of 3D input features (latitude, longitude, speed, or course) enhanced anomaly detection by more effectively separating normal and anomalous vessel behaviors. Furthermore, the proposed ship-trajectory forecasting models demonstrated strong performance. The short-term ShMP model, which uses a feed-forward neural network, achieved mean absolute errors (MAEs) of 0.05-0.13

for three-hour predictions and accuracies of 82–99% across different ship types. For long-term forecasts, the ShMP model utilized the Symmetrized Segment-Path Distance (SSPD) to find historically similar routes, achieving similarity scores of 0.1 and 0.11 for 10-hour predictions.

Moreover, integrating the long-term ShMP model with the AI-ARC maritime surveillance system improved smuggling detection rates, particularly in complex scenarios such as ship-to-ship transfers and unexpected route deviations. Although reliable smuggling detection cannot rely solely on AIS positional data, distinct features suggesting smuggling activity have been observed. In addition, a new version of the ShMP service was successfully evaluated with optimized look-back window sizes for 1-hour and 5-hour forecasting intervals in the Baltic Sea, demonstrating an improved ability to select key AIS messages for accurate ship movement prediction. The service leverages a look-back window size determination algorithm to train a Deep Learning (DL) model—specifically, the Temporal Convolutional Network (TCN)—thus optimizing the computational resources required for sustained operational efficiency.

The AI-based services developed in this thesis were further validated during the AI-ARC project demonstration event, utilizing real-time AIS data from ships operating in the Baltic Sea, particularly near Malmö, Karlskrona, and Copenhagen. Both the ShMP and ShABD services are poised to deliver significant benefits to maritime operations by providing actionable insights and advanced predictive analytics for maritime traffic management.

KEYWORDS: maritime situational awareness; maritime traffic monitoring, AIS data; data-driven methods; machine learning

TURUN YLIOPISTO
Tiedekunta
Laitos
Oppiaine
FARAHNAKIAN, FARSHAD: Digital Maritime Monitoring: Enhancing Situational Awareness in Shipping Traffic using AI-based Models
Väitöskirja, 180 s.
Tohtoriohjelma
Kuukausi 2026

TIIVISTELMÄ

Merenkulun tilannetietoisuus (Maritime Situational Awareness, MSA) on keskeinen tekijä sekä siviili- että sotilasmerenkulun turvallisuudessa ja tehokkuudessa. Tilannetietoisuus perustuu vahvasti kehittyneisiin tunnistusteknologioihin, kuten automaattiseen tunnistusjärjestelmään (AIS), joka mahdollistaa alusten jatkuvan seurannan ja navigointitiedon välittämisen. AIS-viestien valtava määrä kuitenkin kuormittaa perinteisiä analyysimenetelmiä, jotka eivät kykene tehokkaasti käsittelemään dataa reaaliajassa eivätkä tuottamaan kokonaisvaltaista tilannekuvaa. Näiden haasteiden ratkaiseminen edellyttää digitaalisten palveluiden ja tekoälypohjaisten ennakoivien mallien integrointia merenkulun valvonnan automatisoimiseksi ja tilannetietoisuuden parantamiseksi.

Tämä opinnäytetyö keskittyy kahden tekoälypohjaisen palvelun kehittämiseen ja arviointiin: (1) alusten epänormaalin käyttäytymisen havaitsemiseen (ShABD) ja (2) alusten liikkeen ja reitin ennustamiseen (ShMP).

ShABD-palvelu tutkii tiettyjä skenaarioita, mukaan lukien tunnuksettomien laivojen havaitseminen, odottamattomat muutokset liikekuvioissa, spiraaliliikkeet ja salakuljetustoiminta. Kyseisten poikkeamien havaitseminen vaatii datapohjaisia menetelmiä oikea-aikaiseen reagointiin. Tässä työssä niiden tunnistamiseen käytetään ryhmittelyyn perustuvia algoritmeja, mikä poistaa tarpeen luokitella manuaalisesti miljoonia AIS-havaintoja.

ShMP-palvelun osalta työssä tarkastellaan edistyneitä koneoppimismenetelmiä-alusten sekä lyhyen (sequence-to-sequence-mallit) sekä pitkän aikavälin reittien arviointiin samankaltaisuusmittareiden avulla. Näiden osalta mm. satamaviranomaiset hyöty-vät pitkän aikavälin ennusteista, joiden avulla voidaan optimoida satamatoimintoja, kun taas rannikkovartiosto tarvitsee tarkkoja lyhyen aikavälin ennusteita operatiivisen päätöksenteon tueksi. Lisäksi palvelu sisältää menetelmän optimaalisen tarkas-teluikkunan koon määrittämiseksi, mikä mahdollistaa tehokkaan mallin koulutuksen ja laskennallisten kustannusten minimoinnin.

Opinnäytetyön keskeiset tulokset ovat seuraavat:

- Klusterointipohjaiset poikkeamien havaitsemismenetelmät osoittivat, että K-Means algoritmi suoriutui parhaiten verrattuna DBSCAN, Gaussian Mixture Models ja Affinity Propagation menetelmiin erityisesti tunnuksettomien alusten ja spiraaliliikkeiden havaitsemisessa. K-Means-malli saavutti keskimääräi-

sen siluettikertoimen 0,755, mikä osoittaa selkeitä klusterirakenteita. Lisäksi havaittiin, että lisäsyötteiden (leveysaste, pituusaste sekä nopeus- tai kurssitiedot) hyödyntäminen parantaa merkittävästi poikkeamien havaitsemisen tarkkuutta.

- ShMP-palvelun kaksitasoiset ennustemallit osoittivat vahvaa suorituskykyä. Lyhyen aikavälin ennustemalli, joka perustuu myötäkytkentäiseen neuroverkkoon, saavutti kolmen tunnin ennusteissa keskimääräisen absoluuttisen virheen (MAE)- välillä 0,05–0,13, ja kokonaisennustetarkkuus eri alustyypeillä vaihteli 82–99%. Pitk-än aikavälin ennustemalli hyödyntää Symmetrized Segment-Path Distance (SSPD) -samankaltaisuusmittaria tunnistaakseen historialliset reitit, jotka muistuttavat- kohdealuksen liikkeitä. Menetelmä saavutti samankaltaisuusarvot 0,1–0,11 kymmenen tunnin ennusteissa.

- Pitkän aikavälin ShMP-mallin integrointi merenkulun valvontajärjestelmään (AI-ARC) paransi salakuljetukseen viittaavien tapahtumien havaitsemista, erityisesti tilanteissa, joissa alusten välillä siirrettiin tavaraa tai henkilöitä tai joissa havaittiin merkittäviä poikkeamia odotetuilta reiteiltä. Vaikka AIS-tieto ei yksin riitä salakuljetuksen varmaan tunnistamiseen, malli pystyy tunnistamaan siihen liittyviä indikaattoreita. Lisäksi ShMP-palvelun optimaalinen aikaikkunakoko todettiin toimivaksi kahdessa Itämeren alueella suoritetussa testiskenaariossa (1 h ja 5 h ennusteet).

Kehitetyt ShABD- ja ShMP-palvelut tarjoavat merkittäviä käytännön hyötyjä suomalaisessa merivalvonnassa. Ne tuottavat arvokasta tilannekuvaa ja ennakoivaa analytiikkaa Itämeren meriliikenteestä, tukien MSA:n päätöksentekoa. ShABD-palvelu tehostaa poikkeamien tunnistamista ja vapauttaa Merivartioston resursseja etsintä- ja pelastustehtäviin sekä kriittisen offshore-infrastruktuurin suojaamiseen. Tämä on erityisen tärkeää Suomen strategisen sijainnin ja Itämeren geopoliittisten jännitteiden vuoksi.

ASIASANAT: merenkulun tilannetietoisuus; AIS-data; tekoälypohjaiset palvelut; datapohjaiset menetelmät; koneoppiminen; syväoppiminen

Acknowledgements

I am incredibly fortunate to have had the opportunity to conduct my doctoral research at the University of Turku in Finland. I would like to extend my sincere thanks to all those who have contributed to my academic and personal growth.

I am profoundly grateful to my principal supervisor, Prof. Jukka Heikkonen. I will always cherish the memory of our first meeting when I was a recent master's graduate, filled with hope and uncertainty. Your unwavering belief in me ignited a passion that motivated me to embark on this doctoral journey. I feel fortunate to have been guided by such an exceptional mentor. Your boundless enthusiasm for research, a keen eye for detail, and commitment to excellence have made a lasting impression on me. I am also indebted to my second supervisor, Dr. Csaba Raduly-Baka, for his guidance and support throughout my doctoral studies. His expertise in machine learning and pattern recognition has been invaluable in helping me navigate the challenges of this research. His constructive feedback and encouragement have significantly enabled me to produce a high-quality dissertation.

I would also like to express my sincere gratitude to the reviewers of my thesis, Prof. Edmund Førland Brekke and Prof. Roberto Setola, for their insightful comments and constructive criticism. Their valuable feedback has significantly improved the quality of this work. I am also indebted to the opponent, Prof. Ari Visa, for his thoughtful questions and for the challenging discussion during the defense. His examination has helped me to strengthen my understanding of the research topic.

I am deeply grateful to Dr. Paavo Nevalainen for his inspiring passion for data analysis and mathematics. Our insightful discussions on maritime safety and security challenges have driven me to strive for excellence. I am proud to call you a mentor and friend. I also want to thank my incredible sister, Dr. Fahimeh Farahnakian. Your unwavering support has been invaluable, from taking care of me in countless ways to sharing your expertise. Your selfless sacrifices and steadfast belief in me have helped me overcome numerous challenges. I am eternally grateful for your constant motivation and insightful advice throughout my doctoral journey.

During my doctoral studies, I had the privilege of collaborating with esteemed experts on the AI-ARC project. I would like to extend my sincere gratitude to Prof. Isto Mattila, Dr. Pontus Sevenson, Ms. Johanna Karvonen, and the Fraunhofer Research Organization for their invaluable contributions to my research. Their guidance and expertise greatly enriched my work.

Analyzing and processing a comprehensive AIS dataset presented a significant challenge in my research. However, with the support of the Baltic Marine Environment Protection Commission – also known as the Helsinki Commission (HELCOM), especially from Nicolas Florent, I was able to overcome these obstacles. His expertise in sharing valuable insights was essential for successfully integrating the AIS data into ship movement prediction and detecting abnormal ship behavior.

I would like to express my gratitude to the Nokia Foundation for its financial support and for awarding me this prestigious prize for my dissertation. This award was motivating and provided me with the energy to complete my studies.

My doctoral journey was greatly influenced by the support and friendship of my colleagues in the Department of Computing, particularly the Algorithmic Computational Intelligence (ACI) research group. I am very thankful to Pouya Jafarzadeh, Dr. Luca Zelioli, Javad Sheikh, Adrian Borzyszkowski, Tanja Vähämäki, Nina Reini, Rinna Aho, Mohammad Feli, and Kianoosh Kazemi for making this experience memorable. We spent countless hours discussing research ideas, debating different arguments, and celebrating our successes. These moments helped us build lasting friendships. Their helpful feedback, strong support, and enthusiasm have played a key role in my academic and personal growth.

Last but not least, I am eternally grateful to my brother, Dr. Farshid Farahnakian, my sister, Farnoosh Farahnakian, and my brother-in-law, Nima Ahmadzadeh. Your encouragement has propelled me forward, even during the most challenging times. I will never forget the memories of our shared laughter, heartfelt conversations, and unconditional love.

January 25, 2026
Farshad Farahnakian

FARSHAD FARAHNAKIAN

Improving shipping safety is crucial because it directly impacts the global economy. By utilizing AI-based services, such as ship movement prediction and ship abnormal behavior detection, users can gain accurate insights into future situations at sea. These systems employ Machine Learning algorithms on historical AIS data to interpret features and extract patterns in maritime traffic. The primary objective of this thesis is to enhance situational awareness for stakeholders, minimize operational disruptions, and foster sustainable maritime practices.

Table of Contents

Abbreviations

MSA	Maritime Situational Awareness
AI	Artificial Intelligence
ML	Machine Learning
ANN	Artificial Neural Network
DL	Deep Learning
IoT	Internet of Things
IMO	International Maritime Organization
AIS	Automatic Identification System
VTS	Vessel Traffic Service
VHF	Very High Frequency
RQ	Research Question
GMM	Gaussian Mixture Model
ReLU	Rectified Linear Unit
CNNs	Convolutional Neural Networks
RNNs	Recurrent Neural Networks
LSTM	Long Short-Term Memory
TCN	Temporal Convolutional Network
MLP	Multi-Layer Perceptron
MMSI	Maritime Mobile Service Identity
SOG	Speed Over Ground
COG	Course Over Ground
EU	European Union
NATO	North Atlantic Treaty Organization
EMSA	European Maritime Safety Agency
SOLAS	Safety of Life at Sea
OoW	Officer of the Watch
UNODC	United Nations Office on Drugs and Crime
HELCOM	Baltic Marine Environment Protection Commission
API	Application Programming Interface
DBSCAN	Density-Based Clustering of Applications with Noise
AF	Affinity Propagation
SSPD	Symmetrized Segment-Path Distance
LLM	Large Language Model

List of Original Publications

This dissertation is based on the following original publications, which are referred to in the text by their Roman numerals:

I F. Farahnakian, J. Heikkonen, and P. Nevalainen. ”Abnormal Behaviour Detection by Using Machine Learning-Based Approaches in the Marine Environment: A Literature Survey.” In 2022 International Conference on Electrical, Computer and Energy Technologies (ICECET), pp. 1-11. IEEE.

II F. Farahnakian, F. Nicolas, F. Farahnakian, P. Nevalainen, J. Sheikh, J. Heikkonen, and C. Raduly-Baka. ”A comprehensive study of clustering-based techniques for detecting abnormal vessel behavior.” Remote Sensing 15, no. 6 (2023): 1477.

III F. Farahnakian, F. Farahnakian, J. Sheikh, P. Nevalainen, and J. Heikkonen. ”Short and Long Term Vessel Movement Prediction for Maritime Traffic.” In International Conference on Critical Information Infrastructures Security, pp. 62-80. Cham: Springer Nature Switzerland, 2024.

IV P. Svenson, A. Holst, A. Wallberg, P. Nevalainen, F. Farahnakian, Á. Alfonso, G. Vincenzo et al. ”AI-ARC Baltic Demo: Detecting Illegal Activities at Sea.” In 2024 27th International Conference on Information Fusion (FUSION), pp. 1-8. IEEE, 2024.

V F. Farahnakian, P. Nevalainen, F. Farahnakian, T. Vähämäki, and J. Heikkonen. ”Maritime Vessel Movement Prediction: A Temporal Convolutional Network Model with Optimal Look-back Window Size Determination.” Multimodal Transportation. 2025 Mar 1;4(1):100191.

List of Co-authored Publications

The list of co-authored publications that are not included in this dissertation is as follows:

- Javad Sheikh, Fahimeh Farahnakian, Farshad Farahnakian, and Jukka Heikkonen. ”Ice-Water Segmentation Using Deep Convolutional Neural Network-Based Fusion Approach.” In 2023 28th International Conference on Automation and Computing (ICAC), pp. 1-6. IEEE, 2023.

- Javad Sheikh, Fahimeh Farahnakian, Farshad Farahnakian, and Jukka Heikkonen. ”Sea Ice Concentration Estimation Via Fusion of Sentinel-1 and AMSR2 Based on Encoder-Decoder Architecture.” In 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC), pp. 5989-5994. IEEE, 2023.

- Vähämäki, Tanja, Farshad Farahnakian, Paavo Nevalainen, and Jukka Heikkonen. ”Wide-Area Ship Movement Prediction Using Random Forests.” In International Symposium on Intelligent Technology for Future Transportation (ITFT), pp. 220-245. Cham: Springer Nature Switzerland, 2024.

1 Introduction

This chapter offers an overview of the research conducted in this thesis. It begins by presenting the background and motivation for the study. Following that, the research objectives of the thesis are introduced. Next, the main research contributions of each paper are briefly explained. Finally, an outline of the remaining chapters is provided.

1.1 Complex and dynamic maritime environment

Five hundred years ago, shipping was a far more risky venture than it is today. Shipping company owners often walk a fine line between fortune and disaster [1]. When luck was on their side, they could earn a substantial amount of money, but they could just as easily lose everything. The risks were ever-present; ships were frequently wrecked or met with accidents that spelled significant financial losses for the companies. Moreover, weather conditions were unpredictable, technology was limited, and the oceans were filled with hidden dangers, making each journey an unsafe adventure. Shipping was not just a business; it was a gamble on the seas, and only the bravest seafarers dared to navigate these treacherous waters.

The shipping industry experienced a remarkable transformation at the beginning of the 19th century. The introduction of steamships changed the game for maritime trade [2]. Unlike traditional sailing ships, which were completely dependent on the wind and often faced delays, steamships provided reliability and smoother navigation. Shipping companies quickly embraced this innovation, as it allowed them to have better control over their routes and schedules. This shift allowed businesses to expand with newfound confidence, knowing they could rely on consistent transportation for goods across the world. The steamship era marked the beginning of a more structured, trustworthy industry, leading to significant economic growth and heightened international connectivity.

The shipping industry has continued to grow rapidly due to advancements in shipbuilding and power systems. Modern motor ships, particularly those with hybrid engines, are now capable of sailing more smoothly and efficiently than ever before [3]. Today's ocean-going vessels are designed to withstand high waves, strong winds, and other extreme maritime conditions that would have overwhelmed older ships. Container-carrying capacity has increased by approximately 1,500% since 1968 and has nearly doubled in the last decade [4][5]. With innovations in materials, navigation, and engine technology, contemporary ships not only offer increased reliability and sustainability but also enhanced safety for both cargo and crew, enabling them to travel the globe with fewer interruptions. New vessels known as "green ships" even produce lower carbon dioxide (CO2) emissions, contributing to the fight against global warming [6].

The advancements in maritime technology extend beyond just shipbuilding; they also encompass sensor technologies that significantly enhance navigation and safety at sea. Key systems such as the Global Navigation Satellite System (GNSS), Radio Detection and Ranging (RADAR), and Automatic Identification System (AIS) have been developed to ensure smooth and secure shipping operations [7]. GNSS provides accurate positioning data that is crucial for navigation, while Radar allows ships to detect obstacles and other vessels in their vicinity, even in challenging weather conditions. Meanwhile, AIS, the latest maritime sensor, serves as an automatic tracking system that transmits and receives information about a ship's identity, position, speed, and course. These technologies have become essential tools for modern maritime operations, greatly contributing to the overall safety and efficiency of the shipping industry.

Machine learning (ML) is a subfield of Artificial Intelligence (AIS) and is a computational process that focuses on algorithms and techniques for extracting features and building predictive models from datasets [8]. ML models have the potential to learn and improve over time by processing large volumes of information, making ML particularly suited for handling the complexity of maritime data. In recent years, advancements in ML have been further accelerated by developments in Deep Learning (DL), a specialized subfield that utilizes Artificial Neural Networks (ANN) inspired by the structure and functioning of the human brain [9]. These networks are designed to adapt and optimize their parameters through data-driven learning, making them powerful tools for uncovering hidden patterns in maritime traffic and turning manual shipping monitoring into automatic. The availability of AIS data has encouraged researchers and data analysts to apply ML and DL techniques across various applications in the maritime domain. These applications include shipping route optimization [10], estimation of ship arrival times [11], collision risk analysis [12], ship movements prediction [13], and ship abnormal behavior detection [14].

However, despite these progresses in shipbuilding, sensor technologies, and ML methods, it would be a mistake to think that today's shipping industry has left all its struggles behind. While technology has improved dramatically, the maritime world still swings between huge profits and losses. Even in the last decade, with all the sophisticated tools and large vessels, incidents continue to highlight the industry's vulnerability. For instance, between 2014 and 2022, the European Maritime Safety Agency (EMSA) reported 23,814 marine incidents involving 26,108 ships and 604 fatalities, primarily due to collisions, flooding, and groundings [15]. Another striking example was when the massive Ever Given container ship got stuck across the Suez

Canal on March 23, 2021, blocking a critical artery of trade [16]. The incident held up an estimated $9.6 billion goods each day or $400 million an hour in trade. It took a week of complex, coordinated efforts involving tugs, dredging, and a specialized team of experts to finally free the 220,000-tonne vessel [17]. Such incidents remind us that while modern technology has smoothed out many of shipping's rough edges, the industry remains dynamic, challenging, and is still far from predictable.

Furthermore, the maritime industry has increasingly faced threats from abnormal behaviors and maritime crimes, including the operation of ”dark ships” and smuggling activities. Dark ships, which turn off their AIS transponders to avoid detection, pose significant risks to maritime security and safety. These vessels often engage in illegal activities, such as unauthorized fishing, human trafficking, and the transport of drugs, weapons, and sanctioned goods. According to the United Nations Office on Drugs and Crime (UNODC), a significant portion of cocaine destined for European markets is smuggled via ocean shipments, particularly through dark ships that evade conventional tracking systems [18]. Additionally, the shipbroker BRS reported that 787 oil tankers are operating as shadow or dark vessels, which accounts for 8.5 percent of the world's total fleet [19]. A notable incident involved the crude oil carrier Andromeda Star, which collided with another vessel in the Øresund Strait in 2024 while trying to evade detection [19]. The increase in smuggling activities and dark ships has not only impacted shipping safety but has also led to significant economic losses and environmental harm. In 2023 alone, the global shadow fleet was estimated to have caused approximately $10 billion in lost revenue due to smuggling [20].

1.2 Motivation of the research

Human errors and manual analysis of complex maritime networks remain significant causes of these incidents, underscoring the critical need for smarter, more automated, and digital systems within the shipping industry [21]. The availability of AIS data collected from both advanced offshore antennas and satellites presents new opportunities to implement ML methods that enhance safety, efficiency, and reliability across various maritime operations. This thesis aims to develop and evaluate ML models trained on the AIS dataset collected in the Baltic Sea for two applications: (1) ship movement prediction and (2) ship abnormal behavior detection.

ML models offer significant advantages in addressing the limitations of traditional and manual analysis of AIS data for both ship movement prediction and abnormal behavior detection. Traditional methods often rely on fixed rules or statistical models that may not capture the complexity of maritime traffic dynamics, leading to reduced accuracy in predicting movements and identifying abnormal patterns [22][23]. ML models, by contrast, can learn from large and diverse historical AIS records to uncover main maritime routes and movement patterns, enabling more accurate predictions of ship trajectories and detection of anomalies such as illegal activities, erratic maneuvers, or AIS spoofing.

Despite extensive research on the use of ML models for ship movement prediction and abnormal behavior detection, several critical challenges persist. A significant limitation is the reliance on small and geographically constrained datasets for evaluating abnormal behavior detection systems. For example, a system tested exclusively on AIS data from a small coastal region may perform well in that specific area but fail to generalize to diverse maritime environments, such as open oceans or international shipping lanes [24][25]. This raises concerns about the scalability and reliability of these systems when deployed in real-world settings. Furthermore, supervised ML models for abnormal behavior detection often require labeled AIS datasets, which are both resource-intensive and time-consuming. The process of categorizing data into ”normal” and ”abnormal” behavior involves substantial manual effort, making large-scale implementation costly [26]. For this reason, exploring and evaluating unsupervised ML methods on a wide research area is necessary.

In the context of ship movement prediction, existing models frequently rely on fixed look-back window sizes without accounting for the data requirements of specific scenarios. For instance, in [27], the first 20-min sequence of historical AIS data was used for training models to predict the trajectory future 1-min, 5-min, 10-min, and 15-min, respectively. In another [28], the proposed prediction model gets a fixed 1 hr historical AIS dataset as the train sets. This can result in suboptimal forecasts, either overlooking critical short-term patterns or introducing unnecessary complexity through redundant long-term data. A model employing a fixed look-back window might fail to capture dynamic short-term changes or incorporate irrelevant long-term trends, leading to reduced accuracy. Additionally, training DL models on large AIS datasets demands significant computational resources, which may pose practical limitations for smaller organizations or research institutions.

Another notable gap lies in the inability of current systems to provide both short-term and long-term prediction views simultaneously. This dual capability is essential for meeting diverse operational requirements. For instance, port authorities rely on long-term predictions for strategic decisions, optimizing docking schedules and managing port operations efficiently, while coast guards need accurate short-term predictions to make tactical decisions and respond promptly to potential threats and emergencies.

1.3 Research objectives

The research presented in this doctoral thesis is part of the AI-ARC Horizon 2020 project (`ai-arc.eu`), a collaborative initiative aimed at advancing maritime surveillance applications through innovative AI-driven solutions. The primary goal of AI-ARC is to develop a flexible, micro-service-based system that enhances the capabilities of maritime surveillance operators. By integrating data from multi-modal sources (sensor fusion) and visualizing the obtained results from analyzing data in 2D and 3D formats, the project enables operators to make informed, actionable decisions. This system, called the AI-ARC Virtual Control Room, consolidates insights across different platforms, facilitating real-time surveillance and response. The AI-ARC system framework, shown in Figure 1, allows users to configure service chains tailored to specific operational needs. It also includes services that detect intentions and send alerts to prevent potential illegal activities. Finally, the AI-ARC project was successfully showcased as an advanced maritime surveillance system during two demonstration events held in Malmo, Sweden, and Reykjavik, Iceland. Officers from the Swedish and Icelandic Coast Guards attended these events, as both organizations are being considered as potential users of the system.

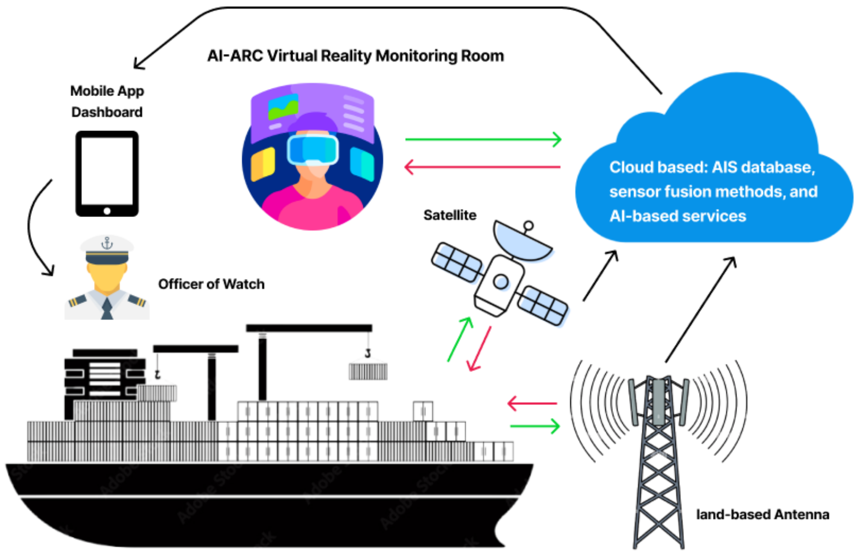

Figure 1. An overview of AI-ARC system.

This thesis leverages the significant advancements achieved in the AI-ARC Horizon 2020 project, focusing on enhancing situational awareness in the maritime environment through powerful AI-driven predictive models. Specifically, it presents the research conducted by the University of Turku for the AI-ARC project, focusing on two critical tasks: ship movement prediction and ship abnormal behavior detection systems. Consequently, the general objective of this thesis aligns with that of the AI-ARC project: to enhance situational awareness using AI-based predictive models for various stakeholders, from port authorities to Coast Guards. More precisely, this thesis has four main objectives: (1) To develop dynamic and scalable ML models for ship movement prediction tasks with low computational demands, (2) To explore unsupervised learning approaches to reduce dependency on manually labeled datasets for abnormal behavior detection systems, focusing on smuggling, dark ship activities, and unexpected ship movement scenarios, (3) To create predictive models capable of providing both short-term and long-term predictions for maritime traffic, and (4) To evaluate the ML-based models for both applications across diverse maritime environments to ensure generalization and robustness.

1.4 Research questions

This thesis addresses the following key research question for ship movement prediction and abnormal behavior detection tasks. These questions aim to address the identified research gaps and contribute to more robust and scalable solutions.

- **RQ1:** How are the current abnormal behavior detection methods for ships developed, utilized, and validated to improve situation awareness? What enhancements should be made to advance these techniques?

- **RQ2:** How can AIS data be integrated with clustering-based and unsupervised anomaly detection methods to reliably identify patterns and predict illegal activities, particularly in high-density maritime regions like the Baltic Sea, while minimizing the need for manual labeling?

- **RQ3:** What framework can be developed to provide both long-term and short-term predictions within a single system, addressing the operational requirements of diverse stakeholders such as port authorities and Coast Guards?

- **RQ4:** Can a ship movement prediction system integrate with other AI-based services to identify more complex abnormal behaviors, particularly smuggling in maritime traffic?

- **RQ5:** How does dynamically optimized look-back window sizing improve the accuracy of ship movement predictions across various time intervals and user-specific needs while reducing computational demands?

The first research question (RQ1) is addressed in Publication I, which reviews various studies that have employed ML models to detect abnormal behavior at sea. This publication identifies the gaps and limitations of existing methods and suggests future directions for improvement. The second research question (RQ2) is explored in Publication II, where different clustering-based methods are evaluated to detect two critical anomalies: dark ships and unexpected ship movements in the Baltic Sea. To enhance model accuracy, this publication utilizes statistical methods and experimental analyses to optimize the hyperparameters of the clustering techniques. In Publication III, RQ3 is tackled, focusing on the need for both short-term and long-term predictions by proposing two distinct prediction frameworks. RQ4 is addressed in Publication IV, which integrates long-term prediction methods with other AI-based services in the AI-ARC architecture to detect smuggling activities. Finally, Publication V explores the last research question (RQ5), introducing a ship movement prediction framework with dynamically optimized look-back window sizes. This framework aims to reduce computational costs and tailor prediction services to seasonal requirements.

1.5 Structure of the thesis

The structure of this dissertation is organized as follows:

- **Chapter 2** provides an overview of situational awareness in the maritime environment, with a detailed explanation of the AIS sensor and its role in maritime traffic monitoring.
- **Chapter 3** describes the methodology employed for both short-term and long-term ship movement prediction frameworks. This chapter introduces similarity measurement techniques and sequence-to-sequence models while also detailing the preprocessing stages involved.
- **Chapter 4** offers a brief review of current approaches for detecting ship anomalies at sea. It then delves into data-driven methods, focusing on supervised and unsupervised ML techniques, with a particular emphasis on clustering-based methods.
- **Chapter 5** summarizes the research contributions and results from the included publications, highlighting their relevance to the thesis objectives.
- **Chapter 6** concludes the thesis, presenting the conclusion, a discussion of the findings, and suggestions for future research directions.

2 Foundation

2.1 Maritime situational awareness

Enhancing maritime situational awareness (MSA) is essential for ensuring safe and efficient shipping operations while minimizing human errors. MSA is a multifaceted concept that has been emphasized by key organizations such as the European Union (EU), the International Maritime Organization (IMO), and the North Atlantic Treaty Organization (NATO) for its critical role in maritime safety and security. It encompasses real-time monitoring, analysis, and proactive management of maritime risks, enabling the prevention of accidents and the ability to respond effectively to unusual or threatening situations [29][30][31][32]. Simply put, MSA involves understanding a ship's surroundings and interpreting current information to anticipate future implications [33]. Based on this definition, MSA is commonly divided into three levels, as illustrated in Figure 2.

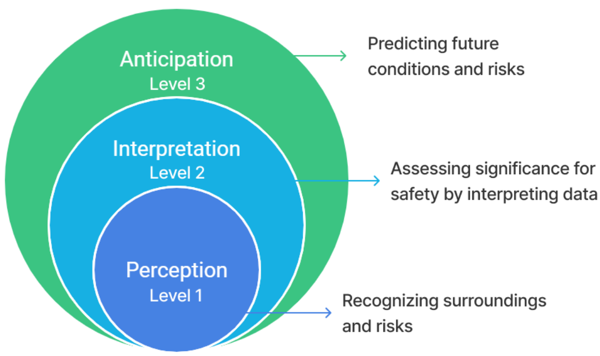

Figure 2. Levels of situational awareness in maritime navigation.

2.1.1 Level-1 maritime situational awarness

The first level of MSA involves perceiving the status, attributes, and dynamics of relevant elements in the maritime environment. All information for Level-1 MSA is provided by sensors and radars installed on board, and this key information, such as the ship's speed, position, and the location of nearby vessels, as well as environmental factors like weather and sea conditions, must be monitored constantly. These data points are grouped into categories such as ship status, equipment status, route information, traffic, and weather conditions. This level focuses on gathering and understanding basic but critical information for safe navigation.

2.1.2 Level-2 maritime situational awarness

The second level of MSA builds upon Level 1 by synthesizing and interpreting the gathered information to understand its impact on navigational goals and objectives. This involves integrating diverse data points through pattern recognition, evaluation, and contextual understanding. For example, navigators assess deviations between planned and current routes, the ship's speed and position, and the impact of traffic, weather, and maneuvers on the ship. To achieve Level-2 MSA, navigators evaluate the differences between the ideal system state and the current state while considering how external factors may affect navigation.

2.1.3 Level-3 maritime situational awarness

The third and highest level of MSA focuses on anticipating future conditions based on current data. This includes projecting the ship's future position, predicting the movements of nearby vessels, and assessing future weather conditions. Level-3 MSA requires comprehending the current situation (Levels 1 and 2) and using this understanding to predict and prepare for future operational states. For instance, it involves simulating future ship traffic patterns and analyzing behaviors to detect abnormal activities, such as illegal operations or navigational hazards.

2.1.4 Challenges in achieving MSA

Maintaining MSA throughout a voyage is primarily the responsibility of the Officer of the Watch (OoW) [34]. This involves constant monitoring of the maritime environment, including the presence and movements of nearby vessels, sea depth, and weather conditions, to identify and respond to potential risks. The OoW must also have an understanding of the ship's characteristics, such as its construction, propulsion system, and onboard equipment, to support effective decision-making. Additionally, the OoW is responsible for tracking the ship's position and preparing for unforeseen events or emergencies. However, maintaining the required focus and concentration under such demanding conditions for extended periods can be extremely tedious and challenging, especially in high traffic density, such as congested ports. This exhaustion can lead to human errors while trying to keep MSA throughout the journey. Therefore, to support safe and efficient operations, it is crucial to develop AI-based systems that can automate some monitoring tasks, such as simulating the future movement of ships and detecting unexpected and illegal activities of ships.

Recent advancements in automation have highlighted the development of AI-based systems that provide Level 1 situational awareness. These systems assist the OoW in better understanding the objects surrounding a ship. Technologies such as computer vision and image analysis are used effectively to detect these objects [35][36]. These tools deliver real-time information about nearby objects and conditions, making it easier for operators to remain aware of their environment. However, experts suggest that merely recognizing the current situation is insufficient. The most effective way to support operators is by helping them predict what might happen next.

2.2 Sensors in the maritime domain

Modern maritime navigation relies on a combination of sensors to ensure situational awareness and reduce response times [37]. These sensors include Global Navigation Satellite System (GNSS), RADAR, AIS, Light Detection And Ranging (LiDAR), visual cameras, and Sonar, each providing unique data to enhance navigational decisions. GNSS systems, such as GPS, GLONASS, Beidou, and Galileo satellites at around 20,000 km altitude, deliver precise positioning and timing information, enabling accurate ship localization even under challenging conditions [38]. Radar, a long-standing staple of maritime navigation, excels in detecting objects and obstacles over long distances, regardless of weather conditions. However, its resolution is limited in distinguishing fine details [7]. In contrast, LiDAR provides high-resolution mapping of nearby objects; however, it has some limitations. Its range is restricted due to eye safety constraints, and it is less effective on darker targets or in adverse weather conditions due to its reliance on light waves. Additionally, commercially available models often have a lower angular resolution, which makes them unsuitable for effectively detecting large or distant vessels [7]. Visual cameras, including RGB and infrared, complement these systems by offering object classification and environmental monitoring capabilities [39]. The Sonar sensor contributes to situational awareness by utilizing sound waves to detect objects underwater [40].

Table 1 provides an overview of the applications of maritime sensors and their contributions to various levels of MSA. Most sensors, including GNSS, Radar, LiDAR, visual cameras, and Sonar systems, mainly enhance Level 1 (Perception) by providing real-time data, as well as Level 2 (Interpretation) through contextual understanding. However, AIS is unique in its capability to contribute to Level 3 (Anticipation). AIS enables the projection of future vessel behavior by broadcasting real-time information about ship movements. Since this thesis focuses on Level 3 MSA, specifically on predicting ship movements and detecting abnormal ship behaviors, AIS data has been selected as the primary sensor of the research due to its effectiveness in supporting these advanced predictive tasks.

Table 1. Applications of maritime sensors and their corresponding contributions to MSA levels.

Sensor	Applications and MSA Contributions
GNSS	GNSS systems provide precise positioning and timing information, enabling accurate ship localization and route tracking. These sensors enhance Level 1 (Perception) by delivering real-time data on ship position, speed, and route status.
RADAR	RADAR excels in detecting objects and obstacles over long distances, even in poor visibility conditions, making it invaluable for collision avoidance and weather monitoring. It contributes to Level 1 (Perception) by detecting objects and Level 2 (Interpretation) by assessing obstacle impacts on navigation.
LiDAR	LiDAR offers high-resolution mapping of nearby objects and supports tasks like obstacle detection in ports and docking assistance. It enhances Level 1 (Perception) through detailed object detection and Level 2 (Interpretation) by providing context for obstacle proximity and alignment.
RGB Cameras	RGB cameras provide detailed imagery for object classification (e.g., vessel types) and environmental monitoring. They contribute to Level 1 (Perception) by capturing visual cues and Level 2 (Interpretation) by enabling object identification and classification.
Infrared Cameras	Infrared cameras are critical for nighttime navigation and low-visibility object detection, enhancing Level 1 (Perception) by monitoring the environment in challenging conditions.
Sonar	Sonar detect underwater objects by capturing acoustic signals, contributing to Level 1 (Perception) through data acquisition and to Level 2 (Interpretation) by analyzing patterns to identify anomalies.
AIS	AIS provides real-time information on vessel identification, movement prediction, and anomaly detection (e.g., dark ships, smuggling). It supports Level 1 (Perception) by broadcasting vessel data, Level 2 (Interpretation) by identifying anomalies, and Level 3 (Anticipation) by predicting vessel behavior and future traffic.

2.2.1 Automatic Identification System (AIS)

AIS has received considerable interest among maritime researchers due to its potential to significantly enhance Level-3 of MSA. AIS provides real-time information about vessels, which assists the OoW in making informed decisions and maintaining situational awareness. This system employs a Very High Frequency (VHF) transmitter, two VHF receivers capable of operating on any AIS channel, and a GPS receiver for synchronizing time, tracking and monitoring ship movements [41]. Essentially, AIS operates on two VHF channels: 161.975 MHz (Channel 87B) and 162.025 MHz (Channel 88B) [41]. According to IMO, AIS was designed to identify ships, track targets, and facilitate more efficient vessel traffic flow [42]. Since 2004, it has been mandatory for all ships weighing over 300 tons to be equipped with AIS transponders, as stipulated by the Safety of Life at Sea (SOLAS) convention [43].

AIS messages are structured data packets, each designed to serve specific purposes [44]. Messages 1, 2, and 3 are real-time position reports sent by vessels at regular intervals—typically every 2 to 10 seconds, especially when the ship's speed is up to 14 knots [45]. These messages provide critical information, including the vessel's position (latitude and longitude), speed over ground (SOG), course over ground (COG), heading, and timestamp. These messages are essential for collision avoidance, vessel tracking, and traffic monitoring in real time.

AIS Message 4 is transmitted by shore-based AIS base stations to synchronize time and provide reference positions for vessels in the area. This message ensures the proper functioning of the Self-Organizing Time Division Multiple Access (SOTDMA) protocol, which manages communication between vessel-to-vessel and vessel-to-shore [46]. AIS Message 5 is sent by vessels every six minutes or upon request. This message provides static and voyage-related information about the vessel, including its IMO number, Maritime Mobile Service Identity (MMSI), ship name, call sign, type, dimensions (length and width), draft, destination, and estimated time of arrival (ETA). This message is crucial for identification and operational decision-making. Together, these messages enable seamless communication between ships.

However, AIS messages are not always fully reliable due to various sources of error. Input errors can occur when manually entered fields, such as navigation status, destination, or ETA, contain mistakes [47]. For example, a study found that nearly 40% of voyage destination reports are incorrectly entered, either intentionally or unintentionally [48]. Additionally, static data—such as ship type, length, and beam dimensions—can often be inaccurate. Furthermore, the accuracy of dynamic data, including position and speed, relies on the quality and calibration of onboard equipment, which can introduce variability and potential inaccuracies. These issues underscore the need for caution and the importance of conducting a preprocessing stage before using AIS data. The preprocessing stage can help identify and correct errors, thereby improving the reliability of predictive models.

3 Data collection and data preprocessing

3.1 Digitraffic API

Digitraffic (`digitraffic.fi`) as the main AIS data source for this research, is a service operated by Fintraffic, providing real-time traffic information for road, marine, and rail transportation. For maritime traffic, the data is sourced from professional maritime systems managed by Vessel Traffic Services (VTS) Finland and the Finnish Transport Infrastructure Agency. AIS data is accessible through two distinct Application Programming Interfaces (APIs): one focused on dynamic information, such as real-time vessel positions and movements, and the other dedicated to static metadata, including ship-specific information like dimensions and identification details. Figure 3 illustrates the workflow of AIS data collection via these APIs. The data obtained from the APIs is systematically stored in our research group's database, where it undergoes further analysis and preprocessing to support research objectives.

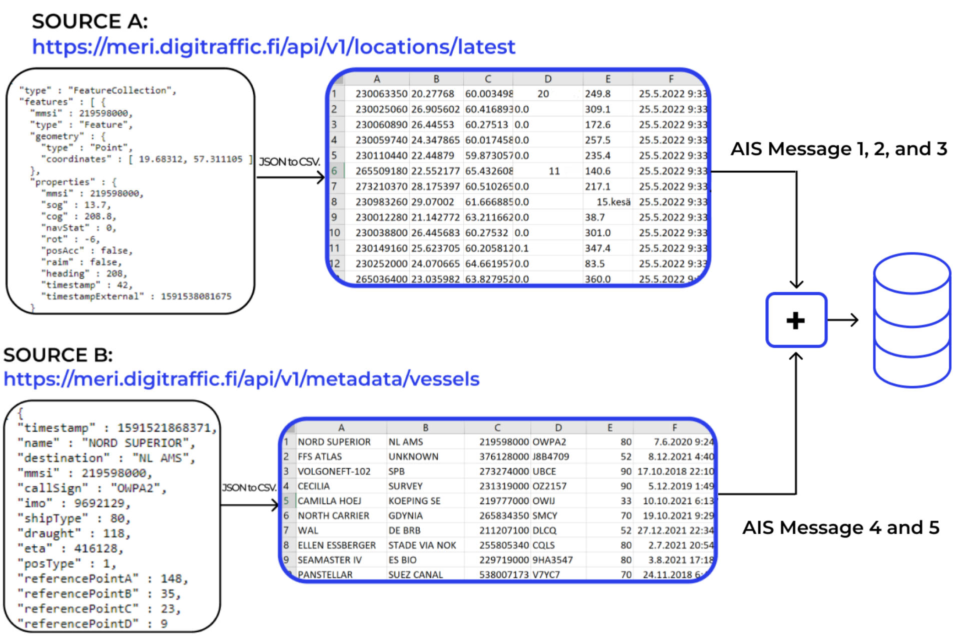

Figure 3. AIS data collection process from Digitraffic APIs.

3.2 Outlier analysis and preprocessing

The AIS dataset often contains outliers or data points that deviate significantly from the typical patterns or behaviors observed in the rest of the data. Outliers can result from environmental factors, human errors, and sensor faults [49]. The impact of outliers should not be underestimated, and identifying them is essential; they can significantly skew dramatically the analysis, model, or outcome, even with a small sample size [50]. Many data analysis techniques treat these outliers as noise or exceptions and discard them. However, outlier analysis and the treatment in this thesis depend on the application. It means that outlier analysis for AIS data is tailored to the specific needs of the two key applications: ship movement prediction and abnormal behavior detection. Some preprocessing stages, such as handling missing data and removing invalid, erroneous, and duplicated records, are common to both applications. However, for abnormal behavior detection, the focus is on retaining and analyzing outliers, as they often represent unusual or unexpected movements critical to identifying anomalies. Consequently, different preprocessing methods were applied and selected based on each application's requirements to ensure the prepared data's reliability and relevance. In the following, all techniques used for outlier analysis and data preprocessing used for both applications are explained in more detail.

3.2.1 Data cleaning

The AIS dataset utilized in this research was collected over nine months, from June 2022 to May 2023. It includes a total of 33,023,246 samples, each containing 13 features, and covers the whole of the Baltic Sea. To optimize computational efficiency, the dataset was first divided into weekly segments. Following this, rows with invalid MMSI numbers, which must contain exactly nine digits, were removed to ensure data integrity. Any entries with more than five missing values (about 50% of the data in an individual row) were also filtered out. The dataset was then filtered to focus on specific ship types relevant to our research objectives, including cargo, tanker, passenger, fishing, tug, and dredging vessels. To maintain geographical consistency, rows with invalid Longitude values (greater than 180 or less than -180) and Latitude values (greater than 90 or less than -90) were removed.

Furthermore, the COG values, representing the ship's direction of movement, were validated, and rows with values outside the valid range (0–360) were discarded. For the SOG feature, rows with SOG values less than 1 knot (to exclude stationary or moored vessels) or greater than 40 knots (as no super-fast ships are included in the dataset) were also removed. After completing these preprocessing steps, the dataset was reduced to 32,261,712 samples, retaining the original 13 features. These steps are the fundamental and initial data cleaning procedures for AIS data, regardless of their specific application.

3.2.2 Trajectory validation algorithm

The results of visualization of the AIS data on the map revealed a few trajectories with strange patterns. Examples of these irregularities are illustrated in Figure 4 in the form of large gaps and swaps in latitude and longitude values that repeated frequently. A ship may seem to follow a consistent trajectory, but in a particular area, it may deviate erratically, with latitude or longitude values varying significantly. This problem is usually due to onboard malfunctioning equipment or errors in the GPS signals from the ship's AIS system [51][52].

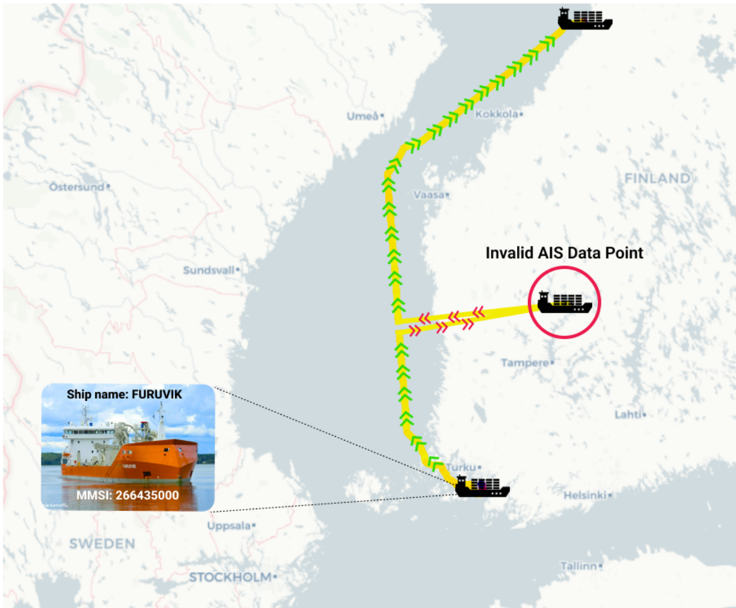

Figure 4. An invalid ship trajectory in the Baltic Sea.

Although, some studies may consider this discrepancy as a result of AIS spoofing or sending false AIS messages [53][54], simply labeling these discrepancies as spoofing with a few conflicting data points does not stand to any reasoning, as spoofing generally involves deliberate manipulation of AIS messages over time. Moreover, AIS spoofing is often in maritime areas rather than offshore places. As a result, these mismatches are primarily due to unintentional errors, which result in unrealistic or noisy detection trajectory data.

To address the issue of noisy or inconsistent AIS trajectories, a trajectory validation algorithm was developed. Algorithm 1 evaluates each point in a vessel's trajectory by calculating the Haversine distance (Formula 1) and time between consecutive points to determine the speed required for the ship to travel between them. The segment is deemed invalid if the calculated speed exceeds a threshold that is the average of ships' speeds in that area (about 15 knots). In such cases, the algorithm removes the intermediate point and rechecks the trajectory between the first and subsequent points. By iterating through the dataset, this approach systematically removes erroneous data points caused by equipment malfunctions or GPS signal issues, ensuring a clean and realistic trajectory dataset for further analysis.

Algorithm 1 Trajectory Validation Algorithm

Input: AIS trajectory data $T = \{p_1, p_2, \ldots, p_n\}$ where $p_i = (lat_i, lon_i, t_i)$
Output: Cleaned trajectory data T'
$T' \leftarrow \{p_1\}$ ▷ Initialize cleaned trajectory with the first point
for $i \leftarrow 2$ to n **do**
    $p_{\text{last}} \leftarrow$ last point in T'
    $d \leftarrow$ Haversine Distance(p_{last}, p_i)
    $\Delta t \leftarrow t_i - t_{\text{last}}$
    if $\Delta t > 0$ **then**
        $v \leftarrow \frac{d}{\Delta t}$ ▷ Calculate speed between points
        if $v \leq$ Average SOG **then**
            $T' \leftarrow T' \cup \{p_i\}$ ▷ Add valid point to cleaned trajectory
        end if
    end if
end for
Return: T'

The Haversine formula calculates the distance between two points as follows:

$$d = 2r \arcsin\left(\sqrt{\sin^2\left(\frac{\phi_2 - \phi_1}{2}\right) + \cos(\phi_1)\cos(\phi_2)\sin^2\left(\frac{\lambda_2 - \lambda_1}{2}\right)}\right) \tag{1}$$

where, d represents the great-circle distance between the two points, and r is the radius of the sphere, which for Earth is approximately 6371 km. The variables ϕ_1 and ϕ_2 denote the latitudes of the two points in radians, while λ_1 and λ_2 represent the longitudes of the two points in radians. It is important to note that all latitude and longitude values are converted to radians before applying the formula.

3.2.3 Simplification ship trajectory algorithm

Simplification of a vessel's trajectory is a vital procedure to decrease the computational complexity of massive AIS datasets while accurately reflecting the overall characteristics of the vessel's journey. To simplify ship trajectories, the Douglas-Peucker (DP) algorithm is employed [55]. This method retains the trajectory shape while effectively reducing the number of data points.

The algorithm begins by drawing a straight line between the first point (P_1) and the last point (P_n) of the trajectory. For each intermediate point (P_i), it computes the perpendicular distance (d_i) to this line using formula 2:

$$d_i = \frac{\left|(y_n - y_1)x_i - (x_n - x_1)y_i + x_n y_1 - y_n x_1\right|}{\sqrt{(x_n - x_1)^2 + (y_n - y_1)^2}} \tag{2}$$

where (x_1, y_1) and (x_n, y_n) are the coordinates of (P_1) and (P_n), and (x_i, y_i) are the coordinates of (P_i). The algorithm identifies (P_m), the point of maximum distance d_{max}. If d_{max} exceeds a predefined tolerance threshold ε, then (P_m) is considered significant and retained. The trajectory is recursively split (P_m), processing the sub-trajectory ((P_1),...,(P_m)) and ((P_m),...,(P_n)). Otherwise, if $d_{max} \leq \varepsilon$, all intermediate points in the segment are removed, as the straight line between (P_1) and (P_n) sufficiently represents the segment.

This process repeats until there are no points whose distance to the corresponding point to the right is greater than ε. It produces a reduced trajectory that retains the main points of the trajectory, ensuring that the overall shape and features of the path remain intact while removing non-essential details. In Figure 5 is an example of a ship trajectory that is smoothed after using the DP algorithm.

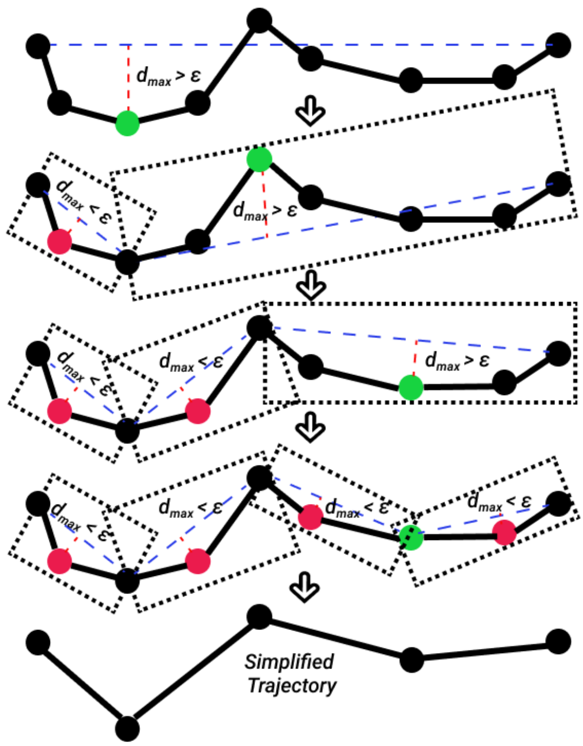

Figure 5. Simplifying a ship's trajectory utilizing the DP algorithm.

To select ε, an algorithm called **Adaptive Core-Threshold Difference** has been used [56][57]. The algorithm begins by testing different thresholds, running the standard DP compression across a range of ε values (e.g., 0.01 × to 10 × ship length), and recording the corresponding number of retained points [57]. These data are then curve-fitted to derive a smooth function $N(\varepsilon)$ that maps threshold values to resulting point counts. The algorithm identifies the "core threshold" at which successive increases in ε yield the largest reduction in points. This corresponds to the knee point of the $N(\varepsilon)$ curve. The difference between consecutive core thresholds, denoted $\Delta\varepsilon_{\text{core}}$, captures the most significant change in compression behavior. Finally, a user-defined compression factor $\rho \in (0, 1)$ scales this difference to determine the optimal threshold:

$$\varepsilon_{\text{opt}} = \Delta\varepsilon_{\text{core}} \times \frac{\rho}{1-\rho}.$$

In fact, the algorithm avoids the arbitrary selection of thresholds by choosing ρ (e.g., $\rho = 0.5$ for moderate compression). Instead, this strategy allows the algorithm to automatically adjust the threshold to suit the complexity of each trajectory, employing a data-driven approach that balances shape fidelity and compression.

Eventually, it should be mentioned that DP algorithm is used for trajectory simplification, but it is not suitable for the abnormal behavior detection system since it relies on subtle deviations in a vessel's movement, and using the DP algorithm would eliminate these valuable anomalies, hindering the detection of unusual actions.

3.2.4 Loitering movement detection

Ship loitering refers to the behavior of vessels, like cargo and tanker ships, that remain in the same position or move only slightly within a restricted area for extended periods. This is often seen around ports or busy maritime zones, characterized by the same AIS messages being sent repeatedly. Reasons for ship loitering include waiting for entry permissions, completing border control procedures, and managing peak traffic times. It can also occur when vessels are loading or unloading cargo at busy ports, or when they're waiting for favorable navigation conditions.

AIS data points related to ship loitering are not useful for ship movement prediction or detecting abnormal behavior. Loitering, characterized by minimal or no movement, adds noise to predictive models, reducing their accuracy and reliability. This type of data lacks variability, making it challenging to identify anomalies such as unexpected movements, smuggling, or unregistered vessels (also known as dark ships). Including loitering data can obscure critical behaviors, complicating detection. Therefore, excluding it helps models focus on relevant patterns.

Algorithm 2 is designed to mitigate the impact of ship loitering on the AIS dataset by reducing spatial redundancy. The DBSCAN clustering algorithm is employed as an effective method for identifying and eliminating dense clusters of AIS points [58]. This algorithm clusters points based on spatial proximity, using the median distance between consecutive AIS points as the ϵ parameter, and setting $min\_samples$ to 1 to allow even isolated points to form individual clusters if no nearby neighbors are found. Clusters exhibiting high spatial density are interpreted as loitering behavior. Once such clusters are identified, each is reduced to a single representative point by selecting its medoid—the point within the cluster that minimizes the average distance to all other points. By choosing the medoid instead of the centroid, the original attributes of the selected AIS point, such as timestamp, SOG, and COG, are preserved. Figure 6 illustrates the trajectory of a ship before and after applying the DBSCAN algorithm with the medoid-based reduction strategy to eliminate the effects of loitering.

Algorithm 2 Eliminating Ship Loitering Using DBSCAN (with Medoid)

Input: AIS data points $\mathcal{D} = \{P_1, P_2, \ldots, P_n\}$ where $P_i = (\text{lat}_i, \text{lon}_i, \text{Timestamp}_i, \text{SOG}_i, \text{COG}_i, \ldots)$
Output: Reduced AIS dataset after removing loitering data points $\mathcal{D}'$
Parameters: ϵ: median distance between consecutive AIS points, $min\_samples = 1$
// Apply DBSCAN clustering
$\mathcal{C} \leftarrow \text{DBSCAN}(\mathcal{D}, \epsilon, min\_samples)$ ▷ Clusters of AIS points
// Initialize reduced dataset
$\mathcal{D}' \leftarrow \emptyset$
for all $C \in \mathcal{C}$ such that $C \neq$ Noise **do** ▷ Ignore noise points
    medoid $\leftarrow \arg\min_{P \in C} \left(\sum_{Q \in C} \text{Distance}(P, Q) \right)$ ▷ Haversine distance
    $\mathcal{D}' \leftarrow \mathcal{D}' \cup \{\text{medoid}\}$ ▷ Preserve full AIS attributes
end for
Return: $\mathcal{D}'$

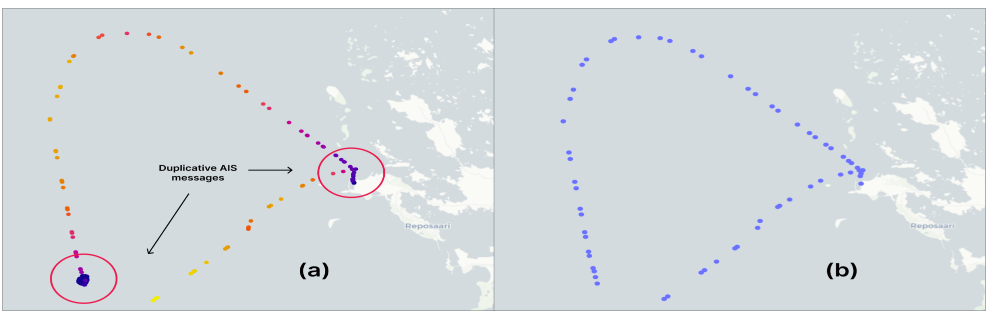

Figure 6. (a) Original AIS data of a ship (b) AIS data of the ship after removing the loitering effect.

3.3 Exploratory data analysis

Exploratory Data Analysis (EDA) is a vital component of any data-driven project. Through systematic exploration, cleaning, and visualization of a dataset, we could gain a comprehensive understanding of its structure, quality, and potential relationships. EDA helps uncover initial insights and identify possible issues, such as missing values or outliers. Ultimately, this foundational understanding allowed us to select the most appropriate analytical or modeling methods, ensuring that our final conclusions or predictive models are built on reliable, well-informed grounds.

In this thesis, we do not treat EDA as a one-time process. Instead, we revisit and update our exploration regularly (weekly) whenever new AIS data becomes available or when additional questions arise during modeling. Figure 7 presents various plots and data visualizations as part of the EDA process.

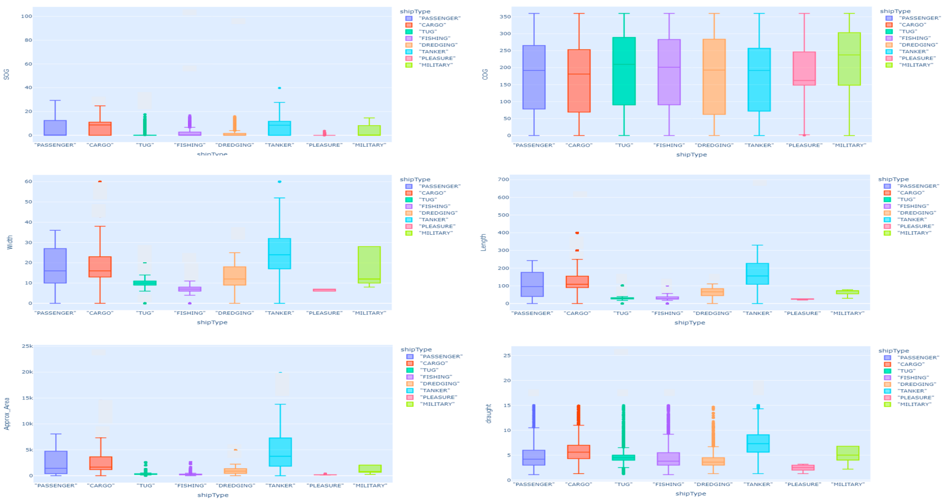

(a)

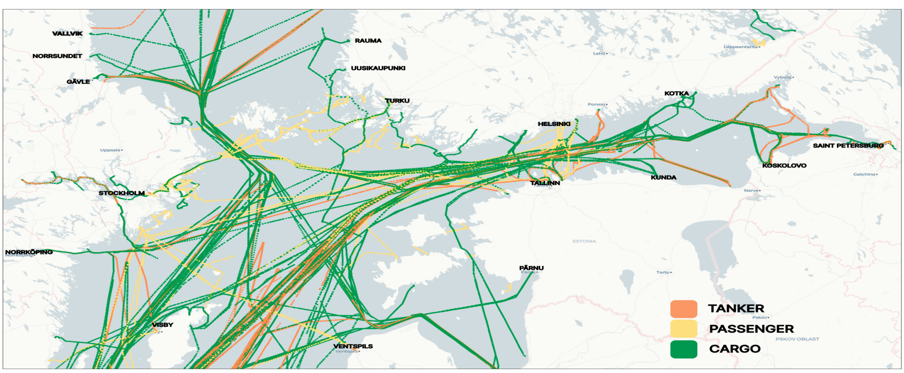

(b)

Figure 7. (a) Box plots for the numerical features of AIS data categorized by ship type. (b) Mapping of AIS data in the Baltic Sea, colored according to different ship types.

4 Ship movement prediction

Statistics and ML offer techniques for spatio-temporal data analysis, providing valuable insights for forecasting future ship movements. This chapter explores two methods: the similarity measurement approach for trajectories and sequential DL algorithms used for both long-term and short-term ship movement prediction.

4.1 Similarity measurement Technique

With the advancement of GPS-enabled devices and mobile computing services, vast amounts of spatio-temporal trajectory data are being collected from moving objects, such as people, ships, and vehicles. A trajectory generally refers to the sequence of positions observed at discrete time intervals [59]. The trajectory analysis enhances our understanding of the movement patterns of objects and individuals. Furthermore, it aids in decision-making, data visualization, and emergency response, enabling more efficient and informed actions across numerous sectors.

Table 2. Comparison of Similarity Measurement Methods.

Method	**Type**	**Noise Sensitivity**	**Unequal Lengths**	**Parameters**
DTW [60]	Spatio-temporal	Sensitive	Yes	No
LCSS [61]	Spatio-temporal	Robust	Yes	Yes
EDR [62]	Spatial	Sensitive	Yes	Yes
Frechet [63]	Spatial	Sensitive	No	No
Hausdorff [64]	Spatial	Sensitive	Yes	No
SSPD [65]	Spatial	Robust	Yes	No

A key feature of trajectory analyses is assessing the relationship or distance between two trajectories, which involves measuring their similarity. Techniques for measuring trajectory similarity can be broadly categorized into three classes: (1) spatial-based measures, (2) temporal-based measures, and (3) spatiotemporal-based measures, each designed for specific use cases [66]. In transportation and urban planning, it helps optimize routes, reduce traffic congestion, and inform infrastructure development [67]. Ecologists utilize trajectory similarity to monitor animal migrations and habitat choices [68]. Law enforcement agencies benefit by tracking the movements of suspects [69]. Additionally, measuring trajectory similarity supports anomaly detection systems to identify suspicious vehicles or aircraft behavior [70][71]. Table 2 summarizes and compares the most common methods for measuring similarity in trajectories.

4.1.1 Symmetrized segment-path distance

The SSPD similarity measurement method is chosen for long-term ship movement prediction for three primary reasons: (1) it is robust against noise, (2) no parameters need adjustment, and (3) it is beneficial in this application where ship trajectories are not temporally aligned but require geometric similarity assessment.

The dissimilarity between two ship trajectories is evaluated by examining spatial alignments (Latitude and Longitude) as follows. Given two trajectories, $A = \{A_1, A_2, \ldots, A_m\}$ and $B = \{B_1, B_2, \ldots, B_n\}$, where each point represents a ship's position in space, SSPD operates in three key steps. First, for each point A_i in trajectory A, the shortest Haversine distance $d(A_i, B)$ is calculated from A_i to all line segments formed between consecutive points in trajectory B. This distance indicates how far each point in trajectory A deviates from the path defined by trajectory B. For any given point on the first trajectory, the method seeks to determine: "What is the shortest distance from a point in trajectory A to any part of trajectory B?" Figure 8 shows this process for a point on a trajectory. Mathematically, it is expressed as:

$$D(A \to B) = \frac{1}{m} \sum_{i=1}^{m} \min_j d(A_i, [B_j, B_{j+1}]) \tag{3}$$

where $d(A_i, [B_j, B_{j+1}])$ represents the shortest distance from point A_i to the segment connecting B_j and B_{j+1}.

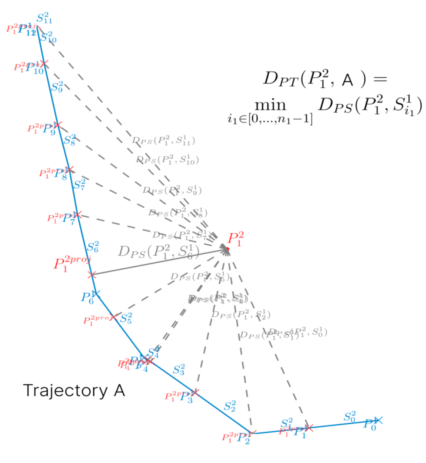

Figure 8. Distance from point p1 to various segments in the trajectory A [65].

In the second step of the SSPD calculation, the process is reversed: for each point B_j in trajectory B, the shortest distance to all segments of trajectory A is calculated. This is denoted as $D(B \to A)$ and can be mathematically expressed as:

$$D(B \to A) = \frac{1}{n} \sum_{j=1}^{n} \min_i d(B_j, [A_i, A_{i+1}]) \tag{4}$$

where $d(B_j, [A_i, A_{i+1}])$ is the shortest Euclidean distance from point B_j to the segment formed between A_i and A_{i+1}. $\min_i$ selects the shortest distance from B_j to any segment of A.

Finally, to ensure symmetry, the SSPD is computed as the average of the distances in both directions:

$$SSPD(A, B) = \frac{D(A \to B) + D(B \to A)}{2} \tag{5}$$

A key property of SSPD is that lower SSPD scores indicate a high level of similarity between the trajectories, as the average distances from points to segments are minimal, reflecting closely aligned paths. Conversely, **higher SSPD scores signify more significant dissimilarity**, corresponding to trajectories that differ significantly in spatial position.

4.2 Deep learning models

DL is a subset of ML with broad applications in computer vision, natural language processing, machine translation, and social network analysis [72]. It uses ANNs with multiple layers that facilitate the learning of data representations at different levels of abstraction. This section highlights DL algorithms designed explicitly for extracting patterns from sequential data.

4.2.1 Recurrent neural network

RNNs are specialized neural network architectures designed for processing sequential data, such as text, audio, and time series [73]. By maintaining a hidden state that carries forward contextual information, RNNs excel in tasks like machine translation, sentiment analysis, and speech recognition [74; 75; 76]. Beyond language and audio tasks, RNNs are particularly effective for time series forecasting and anomaly detection, providing a sequentially aware alternative to Convolutional Neural Networks (CNNs). When combined with CNNs, RNNs enable advanced applications like visual question answering [77], which requires analyzing an image to generate a natural language response. Consequently, RNNs play a vital role in DL, propelling advancements in various fields, including computational biology, recommendation systems, and surgical robotics [78].

AIS messages serve as a valuable source of sequential data, providing an ideal application domain for RNNs. By learning from past ship trajectories, RNN-based models can forecast the future positions and headings of vessels, even in complex maritime environments affected by factors such as changing weather conditions and traffic density [79]. In RNNs, the input is represented as a sequence of values $(x^1, x^2, \ldots, x^t)$, where x^t corresponds to a vessel's navigational state (e.g., position, speed, course) at time t based on AIS data, and T denotes the total number of time steps analyzed. This sequential representation allows the network to model how the vessel's behavior evolves over time, enabling the prediction of future positions [72; 80]. In an RNNs architecture, as shown in Figure 9, each input x^t at time t is projected into the hidden state h^t through the input-to-hidden weight matrix U, while the hidden state from the previous time step h^{t-1} is fed back into the network via the hidden-to-hidden weight matrix W. Only the last hidden state h^T is directly connected to the output through the hidden-to-output weight matrix V, resulting in an output $\hat{y}^T$. Formally, the hidden state update is given by

$$h^t = f\big(Ux^t + Wh^{t-1} + a\big) \tag{6}$$

and the final output is obtained as

$$\hat{y}^T = g\big(Vh^T + b\big) \tag{7}$$

where $f()$ and $g()$ are non-linear activation functions (e.g., sigmoid, ReLU, or softmax), and a and b are bias terms. This design captures temporal dependencies across consecutive time steps and highlights how network outputs can optionally be generated at any desired time step or only at the final time step.

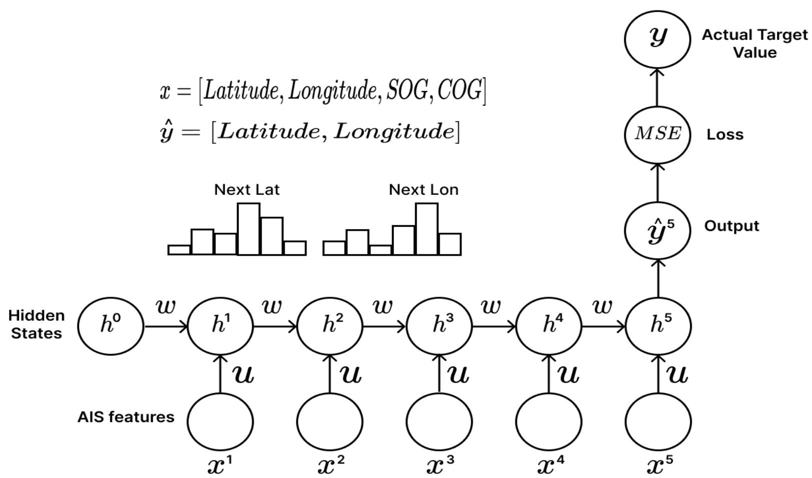

Figure 9. An example of RNNs, consisting of five input units, five hidden units, and one output unit.

In the basic RNNs, information is passed from one step to the next through a special memory component. This design works well for predicting over short spans of data (e.g., a few time steps). However, in practical situations, the length of the sequence is often unknown, potentially leading to issues like gradient vanishing or explosion during learning [79; 81]. Specifically, when you train RNNs on a long sequence, the gradients (the gradient feedback the network uses to adjust its parameters) can either become extremely small (the vanishing gradient problem) or extremely large (the exploding gradient problem). Both situations make it difficult for an RNN to learn how to handle very long sequences effectively.

4.2.2 Long short-term memory

The advanced architecture, LSTM (Long Short-Term Memory), is designed to handle longer look-back windows more effectively [73]. An LSTM introduces an internal ”cell state” that can carry information across time steps with minimal changes. This is often referred to as a “constant error carousel,” indicating that it allows the gradient to flow back through time without diminishing too quickly. This design enables the error signal (used for learning) to move more freely through numerous time steps, reducing the risk of gradients becoming excessively small. Consequently, LSTMs can retain relevant information across longer sequences and make more accurate predictions over extended periods.

The gating structure, shown in Figure 10, enables LSTM networks to dynamically determine when to update the cell state and when to ignore or forget specific information. Consequently, they can maintain vital context over extended periods, enhancing their effectiveness for tasks where insights from earlier in the sequence are essential (e.g., long text passages and extensive time-series data). While this does not entirely resolve all challenges associated with very long sequences, it is significantly more robust than a traditional RNN or a simple variant of RNN.

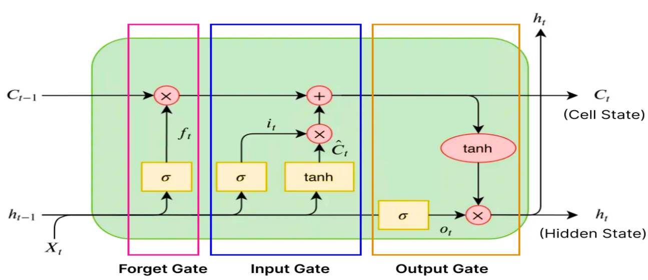

Figure 10. LSTM cells consist of fundamental components, including forget gate, input gate, output gate, and a cell state.

4.2.3 Temporal convolutional network

A Temporal Convolutional Network (TCN), one of the members of CNN, is a fully convolutional architecture designed for processing sequential data [82]. Its defining features are: (1) causal convolutions, which prevent "leakage" from future to past by ensuring each output at time t depends only on current and previous inputs; and (2) the ability to handle input sequences of any length and produce outputs of the same length, just like RNNs. TCNs often serve as an effective alternative to recurrent models, providing robust performance and efficient parallelization. In contrast to RNNs, where later timesteps are processed only after earlier ones, TCNs apply convolutional filters across the entire input sequence simultaneously, accelerating both training and inference [82].

A key strength of TCNs, as illustrated in Figure 11a, lies in the combination of causal and dilated convolutions with residual connections. Causal convolutions ensure that the model considers only current and past inputs, preventing any unintentional use of future information. However, using solely causal convolutions for very long sequences poses challenges, as each additional layer only increases the model's accessible history, leading to a large network for extensive look-back windows. To address this issue, dilated convolutions provide an exponentially larger receptive field by skipping certain inputs [83]. For instance, a dilation factor of 2 processes every other data point, and stacking multiple layers with progressively increasing dilation factors (e.g., 1, 2, 4, 8) enables the network to capture long-range dependencies without requiring oversized filters. Finally, residual connections—similar to those in ResNet [84]—improve training stability by allowing information flow across deeper layers and lengthy sequences.

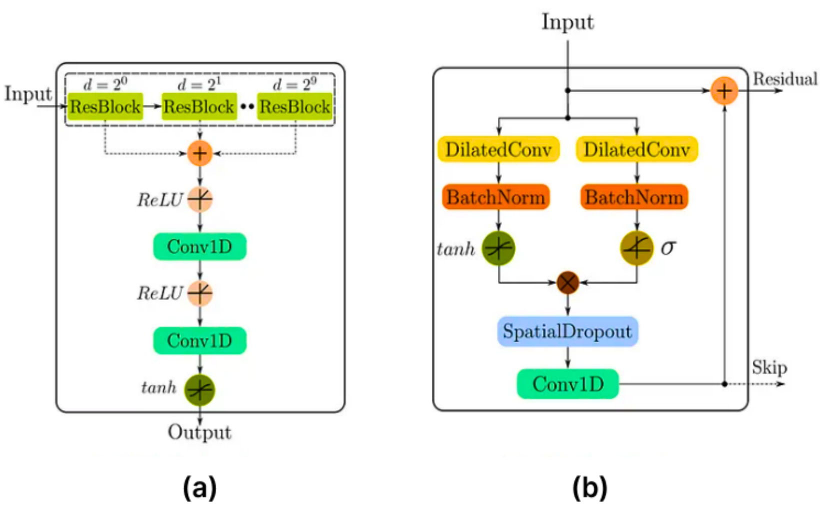

Figure 11. (a) TCN block architecture and (b) ResBlock architecture [85].

Table 3. Comparison of RNN, LSTM, and TCN.

Feature	RNN	LSTM	TCN
Architecture	Maintains a hidden state updated at each time step.	Employs a memory cell with gating mechanisms (input, forget, output).	Uses 1D dilated causal convolutions with skip connections.
Memory	It does not have a specialized memory cell; it relies on the hidden state.	Features a dedicated memory cell for storing long-term information.	No explicit cell; uses convolutions and receptive fields to capture context.
Long-Term Sequences	Prone to vanishing or exploding gradients, limiting long-term learning.	Gating mitigates vanishing gradients, thereby improving the capture of long-term dependencies.	Dilated convolutions offer large receptive fields, effectively handling long-range dependencies.
Training	Conceptually simpler but less effective on extended sequences.	More complex structure (gates) but generally better at learning over long sequences.	Often stable to train due to residual connections and fully convolutional design.

4.2.4 Hyperparameter selection

Selecting appropriate hyperparameters is essential for maximizing the performance of deep learning (DL) models [86]. Among these, dropout rate and learning rate play particularly important roles in determining how well a model generalizes and converges. The dropout rate helps mitigate overfitting by randomly zeroing out a portion of inputs during training, thereby improving the model's ability to generalize. The learning rate controls how much the network's weights are updated in response to the loss gradient, balancing faster convergence with training stability.

In this thesis, a grid search strategy is employed to find the optimal values for these hyperparameters [87]. Specifically, dropout rates [0.1, 0.2, 0.3, 0.4] and learning rates [0.001,0.01,0.1] are systematically tested in different combinations. This comprehensive exploration of the hyperparameter space ensures that the final model configuration achieves the best possible performance on the task at hand.

4.2.5 Evaluation metrics

In ship trajectory prediction tasks, it is common to assess model performance using standard error-based metrics. In this thesis, we focus on four widely used measures: *Mean Squared Error (MSE)*, *Root Mean Squared Error (RMSE)*, *Mean Absolute Error (MAE)*, and the *coefficient of determination* (R^2). Let N denote the total number of observations, y_t be the true value at time step t, $\hat{y}_t$ the predicted value and $\bar{y}$ the mean of all y_t values. The metrics are defined as follows:

Mean Squared Error (MSE)

$$\text{MSE} = \frac{1}{N}\sum_{t=1}^{N}\left(y_t - \hat{y}_t\right)^2. \tag{8}$$

This metric penalizes large errors more heavily by squaring the difference between the actual and predicted values.

Root Mean Squared Error (RMSE)

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{t=1}^{N}\left(y_t - \hat{y}_t\right)^2}. \tag{9}$$

RMSE is simply the square root of MSE and is expressed in the same units as the original data, which can make interpretation more intuitive.

Mean Absolute Error (MAE)

$$\text{MAE} = \frac{1}{N}\sum_{t=1}^{N}\left|y_t - \hat{y}_t\right|. \tag{10}$$

Unlike MSE, MAE does not square the errors, making it less sensitive to outliers and easier to interpret in certain contexts.

Coefficient of Determination (R^2)

$$R^2 = 1 - \frac{\sum_{t=1}^{N}\left(y_t - \hat{y}_t\right)^2}{\sum_{t=1}^{N}\left(y_t - \bar{y}\right)^2}. \tag{11}$$

The R^2 metric measures how well the predictions explain the variance in the actual data, with values closer to 1 indicating better predictive performance.

5 Ship abnormal behavior detection

5.1 Ship Behavior Analysis

The integration of AIS data and data-driven methods can perform many tasks in the maritime industry, such as optimizing shipping routes [88], port management [89], environmental monitoring [90], fleet management [91], collision avoidance [92], and cargo tracking [93], however, one stands out as the most important, abnormal ship behavior detection. This task is crucial for enhancing MSA because it directly impacts the safety and security of maritime operations, allowing the OoW and local authorities to have a clear understanding of forthcoming situations.

According to [94], these irregular movements can be classified into five categories: (1) positional anomalies, which refer to deviations from expected positions; (2) contextual anomalies, which involve atypical navigation times or vessel types in specific locations; (3) kinematic anomalies, characterized by unusual speeds, courses, turns, and stops; (4) complex anomalies, which require multiple anomaly detectors to identify specific behaviors; and (5) data-related anomalies, such as incomplete trajectory information.

Due to the substantial number of vessels globally, the manual detection of anomalous behaviors may prove to be a laborious and error-prone endeavor. Hence, the implementation of automated detection methods is advisable in this context. The automatic identification of anomalous vessel behavior could significantly assist operators of maritime surveillance systems, particularly in light of time constraints, cognitive overload, and the potential for human error associated with the inspection of a vast number of vessels [95].

Generally, methods for detecting abnormal ship behavior can be categorized into three main groups: (1) rule-based methods, (2) data-driven methods, and (3) hybrid methods. Rule-based methods, also referred to as knowledge-based methods, depend largely on expert insights and clear-cut rules [96]. For example, anomalies are detected when a ship's speed or course strays beyond a specific threshold compared to average or historically typical values. While these techniques are low in computational cost and straightforward to interpret, they struggle to adapt to complex behaviors or shifting maritime conditions [96]. Consequently, they frequently generate false positives. In hybrid methods, AIS integrates with other sensor modalities, making it possible to detect suspicious activities such as ship-to-ship transfers or loitering inconsistent with the declared destination [97]. Hybrid methods are significantly more accurate at detecting anomalies, particularly complex ones. However, this method requires careful optimization, as it combines various sensors and demands powerful computational hardware [97]. This chapter aims to focus on data-driven techniques for positional, kinematic, and complex anomalies, as they have been used to develop abnormal ship behavior detection systems in this thesis. The summary of all sections in this chapter is presented in Figure 12.

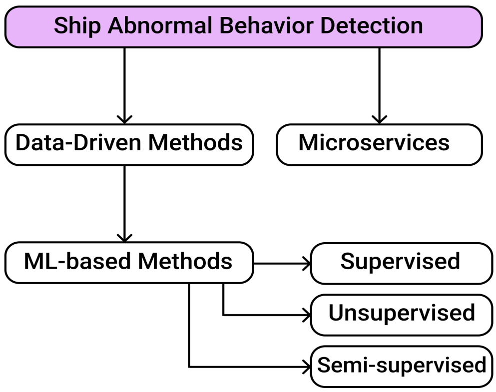

Figure 12. An overview of sections included in chapter 5.

5.2 Data-driven approach

Over the past decade, there has been a significant increase in studies focusing on data-driven methods for monitoring maritime traffic and managing AIS as big data. These methods aim to replace manual analyses with automated processes and efficiently process and analyze large-scale, complex data generated in the maritime domain. Among the most promising methods is ML, which has transformed how data is analyzed and insights are extracted. ML, a branch of AI, allows computers to identify patterns and make decisions based on data rather than relying on explicit programming. This capacity to learn and improve over time by processing vast amounts of information makes ML particularly well-suited for managing the complexity of maritime data.

AIS data provides both dynamic and static information about vessels. By training machine learning algorithms on historical AIS records of ships in normal status, they can reveal primary maritime routes and typical movement patterns, such as speed and course. The knowledge extracted can help identify unusual changes in movement patterns, like high-speed maneuvers, unexpected turns, or spiraling movements [98]. ML algorithms utilized for detecting anomalies in AIS data can be classified into two primary categories: (1) supervised, (2) unsupervised, and (3) Semi-supervised.

5.2.1 Supervised machine learning method

Supervised learning employs labeled datasets to learn input-output relationships [99]. Data scientists create these datasets with input data and corresponding labels. It trains models to apply correct outputs to new input data in real-world scenarios. During training, the algorithm processes large datasets to identify correlations. Model performance is then evaluated using test data to determine training success.

Supervised learning primarily classified ship movements into two categories: (1) irregular and (2) regular. These classification models were applied to analyze simpler abnormal behaviors and intents, such as irregular movements [26; 100]. These learning methods demonstrated acceptable accuracy. However, this approach often demands labeled AIS data, which can be difficult and costly to acquire [99]. Furthermore, maritime networks are dynamic and continuously evolving, making it challenging to maintain updated labeled AIS data and requiring frequent model retraining to capture newly emerging traffic patterns and vessel behaviors.

5.2.2 Unsupervised machine learning method

Unsupervised learning techniques detect abnormal ship behaviors without relying on labeled examples [101]. Instead, they look for natural groupings or patterns within the data to define what is "normal." Anything that does not fit these normal patterns is treated as an anomaly. This makes unsupervised methods valuable in maritime environments, where large quantities of unlabeled AIS data are available, but obtaining labeled examples of illicit or unexpected behavior is difficult [26].

One common approach is clustering, which groups data objects (e.g., vessel routes, speeds, and headings) into clusters so that those within a cluster are highly similar to each other yet clearly distinct from those in other clusters [102]. These similarities or dissimilarities typically rely on distance measures computed from the attributes describing each vessel. As an example, if most ships traveling between two ports follow a typical route, a vessel deviating substantially from that route would be classified as anomalous. Clustering algorithms generally can be categorized into four groups: (1) Partitioning methods, (2) Hierarchical methods, (3) Density-based methods, and (4) Graph-based methods.

Partitioning methods group data into a specified number of clusters, typically using iterative refinement to minimize an objective function [103]. Among these, k-Means is common, assigning each data point to the cluster with the nearest centroid and updating centroids iteratively [104]. Figure 13 shows how each data point is assigned to the nearest cluster during different iterations. K-Medoids (or Partitioning Around Medoids) is a robust alternative that uses actual data points (medoids) as cluster centers, which can be more resilient to outliers [105]. In maritime analysis, partitioning methods can segment vessel trajectories into major shipping lanes and reveal route deviations that merit closer attention [106].

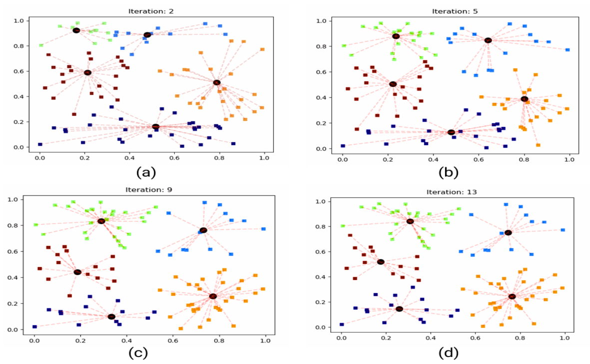

Figure 13. Visualization that shows K-means algorithm's iterative process in action. Iterations: (a) 2, (b) 5, (c) 9, and (d) 13.

Hierarchical clustering groups data in stages, either starting from individual elements and merging them step-by-step (agglomerative) or beginning with a single large group and progressively splitting it (divisive) [107]. The outcome is often presented as a dendrogram, which visually displays how data points combine or separate at different levels [108]. Researchers can then "cut" the dendrogram at various heights to form a chosen number of clusters [109]. This flexibility is particularly valuable when one is uncertain of the optimal number of clusters or seeks to explore how clusters grow or divide over time. In the maritime setting, hierarchical clustering can reveal nested vessel groupings—such as routes within larger operational regions—thus enabling more detailed insights into unusual journeys or suspicious route patterns [110].

Density-based approaches, such as DBSCAN [58] and HDBSCAN [111], identify clusters by finding regions of high data density that are separated by areas of lower density. Outliers in sparse regions are naturally flagged as anomalies, which is particularly beneficial when analyzing irregular vessel trajectories or shipping patterns [111]. On the other hand, **grid-based methods**, such as STING [112] and CLIQUE [112], divide the data space into cells and then aggregate these cells into clusters based on predefined density thresholds. These methods can efficiently handle large maritime datasets, enabling analysts to uncover complex behaviors, such as route convergence in port approaches or unusual vessel groupings in near real-time [113].

Graph-Based Clustering Algorithms treat each data point as a node, with edges representing the similarity or distance between these nodes [114]. Unlike traditional methods that require a predetermined number of clusters or partition the data into fixed cells, graph-based approaches focus on local neighborhood relationships and connectivity patterns to identify natural groupings [115]. A notable example is Affinity Propagation (AF), which considers each data point as a potential exemplar or cluster center [116]. It iteratively ”passes messages” along the edges to identify the most representative nodes [117]. This message-passing framework allows AF to adapt to various traffic patterns and effectively reveal unusual vessel behaviors. By minimizing reliance on strict global parameters, AF would adjust well to heterogeneous maritime datasets, uncovering hidden outliers such as ’dark ships’ or vessels displaying suspicious movement patterns.

Although unsupervised approaches offer significant advantages, changing patterns, such as the emergence of new shipping routes or seasonal variations, can lead to false alarms [26]. To improve the accuracy of the detection process and minimize errors, it is often helpful to incorporate domain-specific knowledge, such as maritime regulations, navigational constraints, and established seasonal behaviors.

5.2.3 Semi-supervised machine learning method

Semi-supervised learning is a branch of ML that integrates both supervised and unsupervised learning by utilizing labeled and unlabeled data [118]. In the context of detecting ship anomalies, a clustering-based algorithm is typically employed to group similar abnormal behavior patterns based on their feature similarities. Meanwhile, an ML-based method trains a classification model using the features extracted from these clusters [14; 119].

In addition to providing a balance between labeled and unlabeled data, semi-supervised approaches can iteratively refine classification boundaries as more examples become available [118]. For instance, a small number of known anomalous trajectories can guide the clustering process by “label spreading,” ensuring that similar unlabeled instances receive provisional labels [120]. These types of ML methods, like others, require encountering new or evolving traffic patterns; these provisional labels can be updated or validated by domain experts, progressively strengthening the classifier’s accuracy [118]. In fact, this feedback loop is particularly valuable in maritime scenarios, where obtaining new labels is challenging and costly, but large amounts of AIS data are continuously generated.

5.3 Micro-services

In modern maritime anomaly detection, a micro-services architecture effectively addresses complex behaviors like smuggling or AIS spoofing by breaking the application into smaller, independent units that collaborate to form a cohesive system [121]. Unlike traditional monolithic designs, where all functionalities reside in a single large application, micro-services divide the detection process into focused services—each running its own process and communicating with others over a network [122]. For example, one micro-service might use predictive models on historical

AIS data to forecast ship trajectories; significant deviations could indicate abnormal behavior. This approach allows for collaboration between services monitoring sharp turns, unusual loitering, or suspicious rendezvous points, providing a robust and scalable framework for detecting various abnormal behaviors in real-world scenarios.

Implementing a weighted decision-making mechanism across micro-services is beneficial [123]. Each service provides a confidence score, adjusted by a weight that reflects its accuracy for specific scenarios [123]. Critical services, such as those analyzing high-impact anomalies like AIS spoofing, should have higher weights, while less critical services receive lower weights. The final anomaly score is calculated by using Equation 12. This approach allows maritime operators to generate a unified decision for detecting and responding to abnormal ship behaviors, ensuring that the most relevant analyses inform the final assessment.

$$\text{Final Anomaly Score} = \sum_{i=1}^{n} w_i \times S_i \tag{12}$$

where w_i is the weight assigned to each micro-service and S_i is the output score produced by that micro-service.

Figure 14 shows an example of how three micro-services: (1) Trajectory prediction, (2) Proximity detection, and (3) Sharp turn detector- contribute to a weighted decision-making mechanism for ship anomaly detection. Each micro-service outputs a score multiplied by a pre-assigned weight. These weighted scores combine to create a final anomaly score, and if it exceeds a set threshold, the ship's movement is flagged as anomalous.

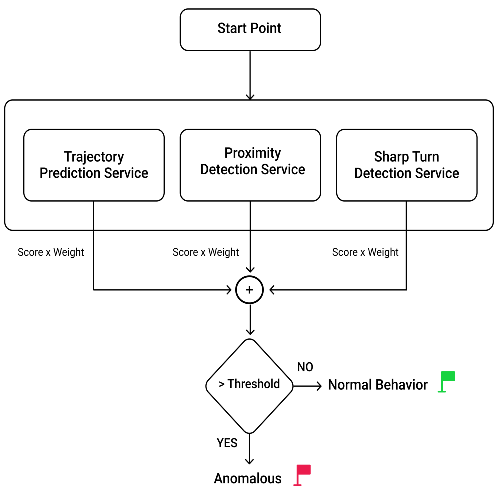

Figure 14. A flowchart shows the collaboration of three micro-services, working in harmony to unveil intricate instances of abnormal behavior.

6 Research studies and results

6.1 Summary of publications

This section reviews the five original publications in Part II of this thesis. For each publication, we summarize the study objectives, motivations, methods, results, conclusions, and contributions to the research questions in Section 1.4. Additionally, we will highlight the author's contributions to each publication.

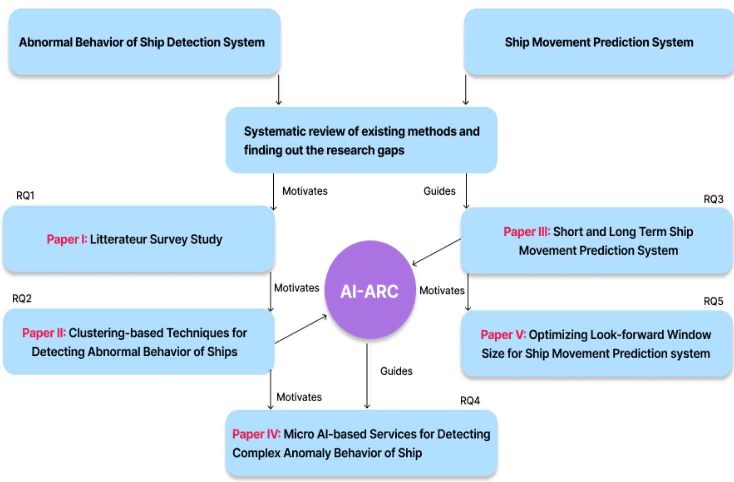

Figure 15. Logical relationships between the RQs and research papers.

Figure 15 also highlights the iterative nature of this doctoral research, showing how each paper not only addresses a specific research question (RQ) but also guides or motivates the next stage of the study. After the systematic review surfaces knowledge gaps, Paper I (RQ1) provides a broad literature survey, establishing both baseline findings and open challenges. These insights then motivate the more targeted Paper II (RQ2), which focuses on clustering-based techniques for detecting abnormal vessel activities. In turn, outcomes from these initial studies feed into the AI-ARC framework at the center, shaping the methodological foundations for subsequent system-oriented investigations. Paper III (RQ3) explores short- and long-term ship movement prediction, while Paper V (RQ5) refines these predictive models by optimizing the look-forward window size. Meanwhile, Paper IV (RQ4) leverages the AI-ARC blueprint to propose micro-AI-based services for complex anomaly detection. Through this interconnected progression, each paper informs and refines the path forward, ensuring that RQs are addressed logically and comprehensively. In this thesis, the preprocessing stages outlined in Section 3.2 have been applied according to the requirements of each research study.

6.1.1 Publication I

Abnormal Behavior Detection by Using Machine Learning-Based Approaches in the Marine Environment: A Literature Survey

Summary: In this paper, we present a detailed literature survey on ML techniques for abnormal vessel behavior detection in maritime environments. We emphasize the significant role of maritime monitoring for trade, safety, and security, noting the challenges posed by vast AIS data, radar signals, and satellite imagery. By reviewing 45 peer-reviewed articles, we categorized existing detection methods based on their use of supervised, unsupervised, or hybrid ML approaches. We also distinguish between trajectory-based (analyzing vessel paths) and point-based (focusing on single data points) methods. Our goal is to highlight the need for robust, scalable, and multi-sensor solutions that enhance MSA in real time.

Objectives: In Publication I, we established two main objectives. First, we aimed to identify and summarize the various methods for detecting abnormal behavior using machine learning algorithms. Second, we intended to provide a comprehensive overview of these existing methods, highlighting their limitations. Our ultimate goal is to empower researchers in the field and illuminate potential pathways for future exploration and innovation.

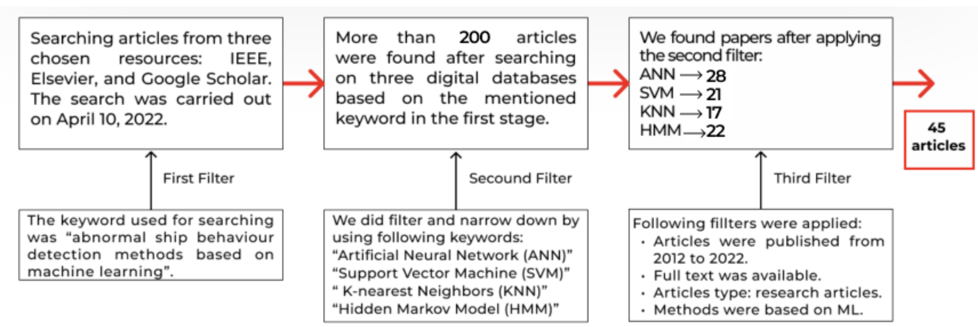

Figure 16. Methodology used for selecting articles reviewed in Publication I.

Methods: We conducted a systematic search using IEEE, ScienceDirect (Elsevier), and Google Scholar from 2012 to 2022. Figure 16 provides an overview of the steps we took to find relevant articles for our review. We started with general queries, yielding over 200 matches, and then applied relevant filters (e.g., research focus, publication date range) to narrow our selection down to 45 articles. Our review centers on two major data types for maritime anomaly detection: (1) AIS (providing vessel positions and identifying attributes) and (2) Satellite Imagery (particularly SAR), enabling vessel tracking even when AIS transponders are switched off. Finally, we surveyed key ML algorithms: Support Vector Machines [124], K-Nearest Neighbors [125], Hidden Markov Models [126], and Neural Networks (including CNNs and RNNs) to see how they were applied to maritime data. We evaluated each approach's advantages, limitations, and real-world deployment considerations.

Results and contribution: This literature survey identified several limitations in current methods for detecting abnormal behavior in maritime environments. One of the primary issues is the absence of a unified framework, as evidenced by the varying definitions and inconsistent methodologies found across the articles we reviewed. Additionally, we noted that most studies utilized relatively small and constrained datasets, which complicates the validation of system performance in real-world scenarios. Therefore, we suggest that future research should focus on larger and more diverse data sources to enhance both robustness and scalability.

Another significant finding is that existing systems seldom incorporate predictive insights. To address this gap, we recommend developing methods that can forecast vessel behaviors over user-defined time intervals, which would facilitate proactive measures and improve maritime situational awareness. The main Contributions of Publication I are as follows:

- Comprehensive Literature Review: We provide a thorough analysis of existing machine learning-based methods for detecting abnormal vessel behaviors in maritime environments. It categorizes the reviewed articles into two main groups: (1) methods and (2) data sources, which helps in understanding the landscape of current research.

- Highlighting Research Challenges: We discussed various challenges and limitations associated with existing methods, such as the focus on offline anomaly detection and the need for improved data labeling processes. By identifying these gaps, we encouraged future research to address these issues for enhancing maritime safety.

- Analysis of ML Techniques: This paper not only reviews existing ML techniques but also analyzes their advantages and disadvantages. This analysis serves as a guide for researchers looking to develop more advanced frameworks and tools in the field of maritime anomaly detection.

- Focus on Recent Developments: This literature survey emphasizes recent advancements in ML methods and their applications in maritime environments, providing a current perspective on the state of research in this area.

Author's Contribution: The author of this thesis conducted a systematic literature review on ML-based methods for detecting abnormal behavior using maritime data, specifically AIS and SAR imagery. Additionally, he contributed to the writing process of the article and proposed a possible multi-modal ML-based framework aimed at addressing gaps and improving the accuracy of anomaly detection in the maritime environment.

6.1.2 Publication II

A comprehensive study of clustering-based techniques for detecting abnormal vessel behavior

Summary: We present a comprehensive comparison of well-known clustering algorithms: (1) K-means, (2) Density-Based Spatial Clustering of Applications with Noise (DBSCAN), (3) AF, and (4) GMM for detecting abnormal vessel behaviors in this study. To conduct the analyses, we first collected a three-month AIS dataset covering the Baltic Sea and subjected it to extensive preprocessing to ensure data reliability. By systematically testing multiple input combinations (latitude, longitude, speed over ground, and course over ground), we highlight the unique strengths and challenges of each clustering technique. Their findings show that K-means achieves the highest average silhouette coefficient (approximately 0.755), although the results obtained from the experimental part show that four clustering algorithms demonstrate strong potential for uncovering anomalous maritime behaviors. Furthermore, the study emphasized the importance of parameter selection (e.g., number of clusters, radius for DBSCAN) and proposed an algorithmic strategy to automate this step for DBSCAN and AF. Overall, the authors conclude that clustering-based approaches, when configured properly, represent a robust, cost-effective (using unlabeled data) solution for real-time maritime anomaly detection. Figure 17 provides an overview and a summary of Publication II.

Objectives: The main objective of Publication II is to delve into the exciting world of clustering methods and their effectiveness in detecting two types of abnormal ship behaviors: (1) dark ships and (2) spiral movements. By exploring these techniques, we aimed to deepen our understanding of their applications in maritime settings. Additionally, this study focused on analyzing three months of AIS data from the Baltic Sea, highlighting AIS as a promising sensor for developing accurate automated anomaly detection systems.

Methods and data: We initiated our study by collecting three months of AIS messages using publicly available APIs. This process ensured that we retrieved both static data (such as ship identifiers and vessel types) and dynamic data (including latitude, longitude, SOG, and COG). The initial dataset contained 11,023,246 AIS messages. To optimize processing time, we first divided the data into weekly segments. Next, we implemented several cleaning steps (Sections 3.2.1 and 3.2.2). These in-

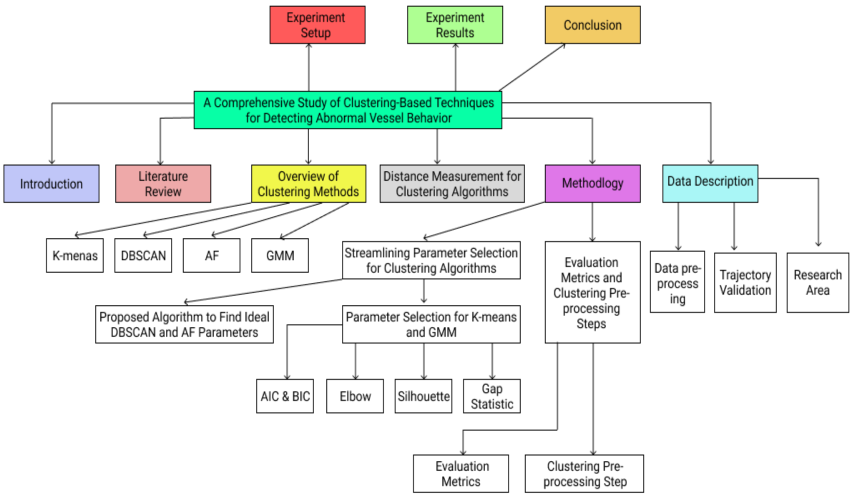

Figure 17. The structure of Publication II.

cluded removing incomplete entries and discarding invalid ship identifiers to create an accurate and consistent dataset. We also filtered out irrelevant vessel types, focusing exclusively on the categories most relevant to our abnormal behavior detection task, such as cargo ships, tankers, passenger vessels, fishing boats, tugs, and dredgers. After completing the preprocessing steps, the final dataset, which covered the entire Baltic Sea, includes 10,261,712 samples.

Next, we identified the key features necessary for detecting abnormal vessel behavior. These included both two-dimensional (e.g., latitude and SOG) and three-dimensional inputs (e.g., latitude, longitude, and COG), reflecting a variety of possible configurations for the clustering analysis. Depending on each clustering algorithm's sensitivity to scale, we performed feature normalization to align numerical ranges. This ensured that certain attributes (e.g., SOG) did not inadvertently dominate the clustering outcomes due to differences in magnitude. The following combinations of 2D and 3D inputs have been investigated: (1) Latitude and SOG (2D); (2) Longitude and SOG (2D); (3) Longitude and COG (2D); (4) Latitude and COG (2D); (5) Latitude, Longitude, and SOG (3D); (6) Longitude, SOG, and COG (3D); (7) Latitude, Longitude, and COG (3D); (8) Latitude, SOG, and COG (3D).

We then selected and configured four clustering algorithms: K-means, DBSCAN, Affinity Propagation, and Gaussian Mixture Models, each of which required distinct parameter-tuning strategies. For K-means, we employed the elbow, silhouette, and gap statistic methods to determine the optimal number of clusters. For DBSCAN, we proposed a novel parameter-selection technique that integrates k-distance plots and grid searching to identify suitable values for Eps and MinPts. In the case of

GMM, the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) were used to ascertain the best number of Gaussian components. Silhouette scores guided our final parameter choices, helping to ensure that data points formed cohesive and well-separated clusters.

Finally, we assessed each clustering model's capacity to discern anomalies by calculating silhouette coefficients, which quantify how distinctly different one cluster is from another. Additionally, we visualized the clusters to confirm whether "dark ship" intervals (where AIS data were disabled or suspiciously missing) and "spiral" vessel movements (indicating possible illicit behavior) were correctly identified. This combination of quantitative metrics and qualitative inspection helped us validate the overall effectiveness of the framework for real-world maritime anomaly detection.

Results and contribution: The obtained results demonstrated that clustering can effectively uncover hidden patterns in AIS data, especially if appropriate parameter tuning and dimensional inputs (2D vs. 3D) are carefully selected. Among the four approaches, K-means achieved the highest overall silhouette coefficient, hovering around 0.755 for certain three-dimensional inputs. The results of the K-means algorithm from the experimental phase are presented in Figure 18. Additionally, the identified anomalies are mapped in Figure 19, which indicates where the detected ship turned off its AIS and where the other ship began to exhibit spiral movements. Nevertheless, the other methods also performed well for specific scenarios, underscoring the importance of a systematic parameter-selection strategy.

The proposed algorithm incorporating statistical methods like the elbow [127], gap statistic [128], AIC [129], and BIC [129], proved instrumental in identifying optimal hyper-parameters, notably Eps and $MinPts$ for DBSCAN and damping values for Affinity Propagation. This parameter tuning was crucial not only for isolating suspicious clusters but also for reducing false positives. The results regarding the optimal number of clusters and other necessary parameters for clustering algorithms are presented in Figure 20.

Significantly, the experiments revealed that both dark ships and spiral behaviors are detectable within even a relatively small subset of Baltic Sea data. By matching clusters with time stamps, we pinpointed prolonged intervals lacking AIS transmissions or featuring unusual course changes, clearly indicative of potentially illegal activities. These anomalies were further validated by visualization, offering insights that could benefit maritime authorities and port operators.

Key Contributions of Publication II are as follows:

- Comprehensive Evaluation of Clustering Algorithms and Parameter Selection: Publication II begins by conducting an extensive evaluation of four widely used clustering algorithms, K-means, DBSCAN, AP, and GMM, tailored to the tasks of detecting dark ships and spiral vessel movements. Rather than re-

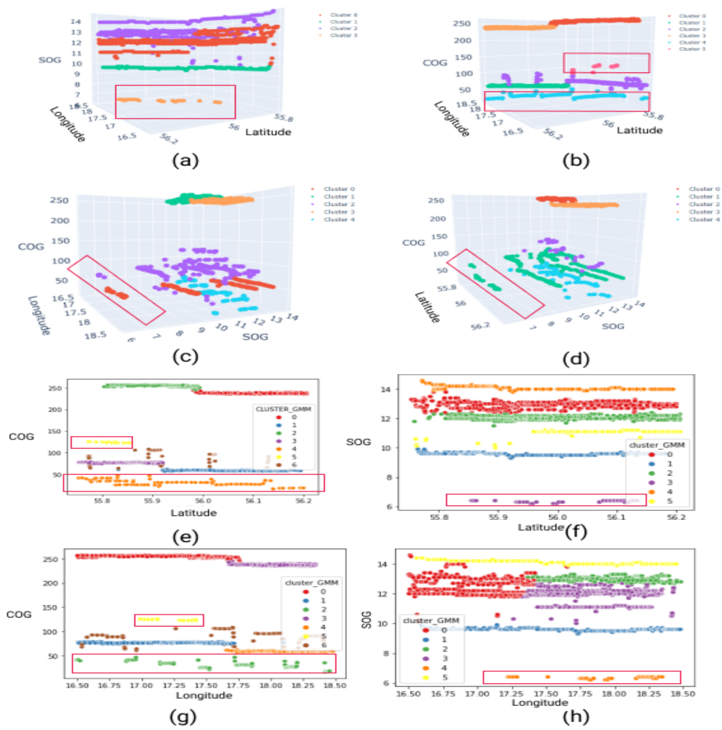

Figure 18. The results obtained from the K-means algorithm applied to AIS messages of ships that sailed in the targeted research area.

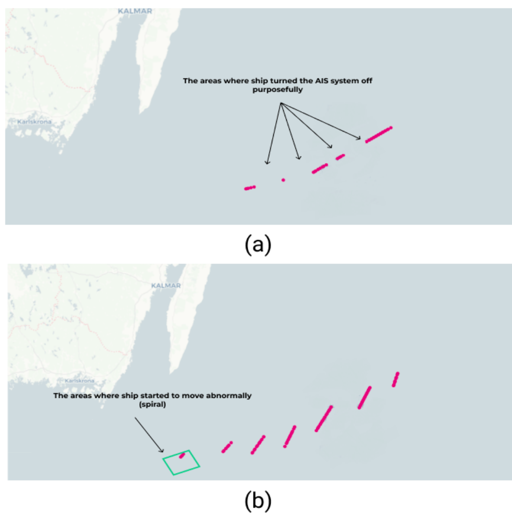

Figure 19. (a) The ship disables the AIS system, and crews stop transmitting any AIS messages. (b) The cargo ship's spiral movements have been detected.

Input	K-Means (Number of Clusters)	DBSCAN (*Epsilon*, *MinPts*)	AF (Dumping)	GMM (Number of Components)
Latitude, *Longitude*, *COG*	Elbow method: 6, Silhouette method: 4, Gap stastistic: 4	*Epsilon*: 0.91897, *MinPts*: 6	Dumping: 0.98367, Silhouette: 0.87062	BIC method: 7, AIC method: 7
Longitude, *COG*	Elbow method: 6, Silhouette method: 4, Gap stastistic: 6	*Epsilon*: 0.90223, *MinPts*: 6	Dumping: 0.72061, Silhouette: 0.91174	BIC method: 6, AIC method: 6
Latitude, *COG*	Elbow method: 4, Silhouette method: 4, Gap stastistic: 6	*Epsilon*: 0.98348, *MinPts*: 7	Dumping: 0.73448, Silhouette: 0.91201	BIC method: 6, AIC method: 6
Longitude, *SOG*	Elbow method: 8, Silhouette method: 6, Gap stastistic: 8	*Epsilon*: 0.59135, *MinPts*: 4	Dumping: 0.72069, Silhouette: 0.45281	BIC method: 8, AIC method: 8
Latitude, *SOG*	Elbow method: 6, Silhouette method: 6, Gap stastistic: 6	*Epsilon*: 0.31475, *MinPts*: 4	Dumping: 0.77586, Silhouette: 0.52233	BIC method: 6, AIC method: 6
Longitude, *SOG*, *COG*	Elbow method: 6, Silhouette method: 4, Gap stastistic: 6	*Epsilon*: 0.99887, *MinPts*: 5	Dumping: 0.85862, Silhouette: 0.82800	BIC method: 5, AIC method: 5
Latitude, *SOG*, *COG*	Elbow method: 6, Silhouette method: 4, Gap stastistic: 6	*Epsilon*: 0.93469, *MinPts*: 4	Dumping: 0.88622, Silhouette: 0.83040	BIC method: 5, AIC method: 5

Figure 20. Results of the parameter selection process for clustering algorithms based on different inputs: two-dimensional and three-dimensional data.

lying on conventional benchmarks alone, the study incorporates an innovative algorithm and robust statistical methods to determine optimal input parameters and the ideal number of clusters. Together, these contributions help data scientists, ensuring accurate anomaly detection with minimal trial and error.

- Focus on Specific Maritime Anomalies and Data Utilization: A key strength of this research is its direct attention to two particularly problematic anomalies: dark ships, which deactivate their AIS to hide illicit activities, and spiral movements, indicative of suspicious maneuvers. By honing in on these targeted behaviors, the paper underscores the urgent need to address distinct illegal practices in maritime contexts. Moreover, the study leverages a substantial AIS dataset, comprising over ten million samples from Baltic Sea traffic. This rich and expansive dataset furnishes a reliable groundwork for analyzing vessel behavior patterns and ensures that the findings remain robust across varying maritime conditions.

- Spatiotemporal Features and Practical Implications: The method outlined in Publication II further integrates spatial and temporal features—an approach that captures vessel behaviors over different time intervals and locations. As a result, maritime agents can better detect and track suspicious vessels, even when they alter their routes or speeds. This heightened detection capability translates directly into improved maritime safety and security, with authorities better equipped to identify and intervene in illegal or potentially hazardous situations. Notably, the high citation count over the past two years signifies the study's visibility and influence among peers, underscoring its importance within the domain of maritime anomaly detection.

Author's Contribution: The author of this thesis was responsible for curating the Baltic Sea AIS dataset, which included carrying out spatiotemporal filtering to prepare the data for clustering-based anomaly detection. Additionally, the author developed Python code to implement both the parameter selection algorithm and the multi-dimensional clustering pipeline, evaluating its effectiveness in detecting dark ships and spiral movements. All co-authors contributed to interpreting the experimental results. Furthermore, the author played a key role in writing the paper, creating detailed tables that compared the effectiveness of different clustering algorithms, thereby ensuring the clarity and accuracy of the final presentation.

6.1.3 Publication III

Short and Long Term Vessel Movement Prediction for Maritime Traffic

Summary: This article explores a dual-approach framework for predicting vessel movements in maritime traffic. The authors introduce two complementary methods: (1) short-term and (2) long-term—to accommodate the contrasting requirements of stakeholders such as port authorities (who often need swift, detailed predictions) and Coast Guards (who may need broader, longer-range forecasts). The short-term method applies a feed-forward neural network to user-defined intervals and excels at precise trajectory estimates in real time, while the long-term method leverages a distance-based similarity measure to predict vessel paths over extended durations. These approaches were tested on real-world AIS data from the Baltic Sea. Experimental findings suggest that the short-term model produces highly accurate forecasts within a short window size, whereas the long-term model is both computationally efficient and well-suited for strategic planning.

Objectives: Our main goal is to seamlessly integrate short-term and long-term prediction horizons into one exciting solution that meets the diverse needs of maritime stakeholders, such as port authorities, Coast Guards, and shipping companies. We're also eager to introduce a fresh evaluation metric that complements established metrics for short-term predictions, effectively showcasing how accurately we can forecast a vessel's future movements while offering invaluable insights.

Methods and data: The short-term vessel movement system leverages a feed-forward neural network (NN) architecture specifically developed for predicting latitude and longitude coordinates. The model comprises three hidden layers, each containing 32 neurons, and uses ReLU activation, with dropout set to 0.2 after the first two layers to curb overfitting. This configuration proved effective in handling the five selected features (timestamp, latitude, longitude, SOG, and COG), ultimately producing a two-neuron output layer for the vessel's future coordinates. The Adam optimizer (learning rate 0.001) was chosen following comparative experiments with different optimizers, while MSE served as the primary loss function in tandem with MAE and accuracy for thorough evaluation. The model underwent training for 100 epochs, setting aside 20% of the overall data for validation to monitor performance. Continuous integration of newly arriving AIS messages ensures up-to-date predictions, thus making the system resilient to evolving environmental factors or emergent traffic patterns. This short-term prediction approach can reliably forecast upcoming vessel trajectories in a fast-changing maritime environment by seamlessly blending robust hyperparameter choices, careful data preprocessing, and supervised learning.

The long-term vessel prediction approach relies on identifying similarities between a target ship's partial route and an extensive database of historical trajectories. Specifically, it employs the SSPD similarity measurement method, a purely spatial metric that gauges how closely two vessel paths match based on latitude and longitude. In practice, SSPD divides each route into smaller segments and then compares these segments pairwise in both forward and backward directions. The final similarity score is the average of the two, ensuring a robust measure of geometric alignment regardless of travel direction. By ranking all historical trajectories in terms of SSPD, the prediction system isolates the route or routes that best resemble the target ship's current position data. It subsequently infers that the target vessel will follow a path similar to the remainder of the highest-ranked historical trajectories. This method focuses entirely on shape similarity, neglecting time and speed—and is therefore well suited for making longer-term forecasts under the assumption that geometric patterns can predict a vessel's eventual path.

AIS data (both static and dynamic) was collected from two open APIs provided by the Finnish transport infrastructure agency (digitraffic.fi). These APIs supply marine warnings, harbor schedules, and vessel location AIS messages for Baltic Sea traffic. Between June 1 and July 31, 2022, about 11 GB of data were gathered, merging vessel features such as latitude, longitude, SOG, COG, and ship identifiers like MMSI and IMO numbers. The collected AIS data has been processed using the methods mentioned in Sections 3.2.1 and 3.2.3. This study mainly targeted cargo, tanker, and passenger ships due to their high representation in the dataset and their regular movements.

Results and contribution: The short-term neural model exhibits precise, fine-grained forecasts, often deviating by less than 0.3 km from vessels' actual positions in tests. Table 4 shows the results obtained from training the model for short-term prediction, reflecting high accuracy for the selected ships. Meanwhile, the long-term similarity approach offers stable and efficient extended predictions, trading minimal detail in exchange for gains in computational speed and memory usage. The paper underscores how real-time and long-term forecasts can coexist to meet varied operational demands by housing these complementary algorithms within a single framework. A novel metric called the Error Rate of Prediction (ERP) has been introduced to evaluate short-term model accuracy. The ERP metric is applied after each prediction to calculate the Haversine distance (Equation 1) between actual and predicted points, demonstrating the model's accuracy in forecasting ship locations. This ongoing validation process allows us to gauge the model's effectiveness and refine the prediction strategy based on empirical results. Overall, it provides a quantifiable measure of prediction accuracy. This metric can enhance decision-making processes and improve safety standards in maritime operations by ensuring that ship movements are predicted with a high degree of reliability. Figure 21 illustrates the prediction results based on the ERP metrics for the selected ships.

Table 4. Performance metrics for short-term vessel movement prediction.

Ship	Ship Type	MSE	MAE	Accuracy
A	Cargo	0.03	0.13	0.86
B	Passenger	0.01	0.08	0.98
C	Tanker	0.006	0.60	0.89
D	Tanker	0.002	0.04	0.99
E	Passenger	0.005	0.06	0.82
F	Cargo	0.04	0.05	0.98

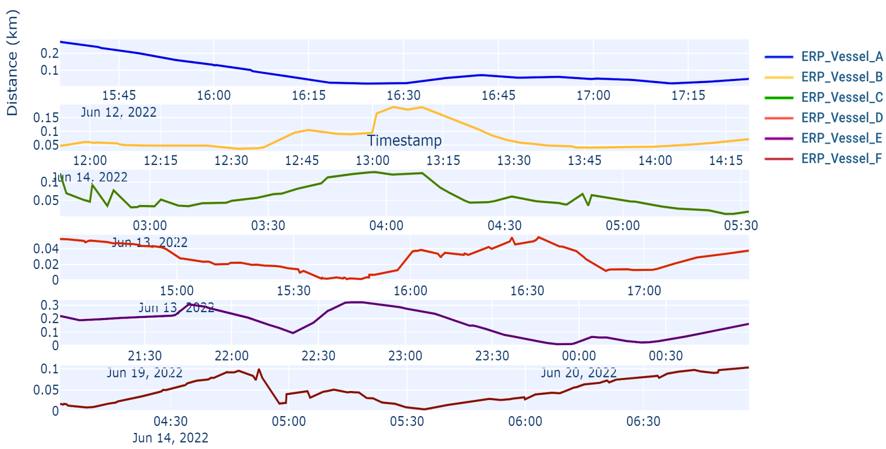

Figure 21. ERP values collected during the experimental stage.

Author's Contribution: The author of this thesis contributed extensively to the core methodology of the dual ship movement prediction system—both the short-term and the long-term approaches. He implemented the Python code underlying data preprocessing, feature extraction, and the dual prediction pipeline. This coding effort included establishing a reliable process for cleaning AIS messages, selecting the relevant variables, and integrating the final predictions of each method. Furthermore, he participated in writing the methodology section, clarifying how the short-term model was trained and updated, as well as how the long-term similarity measure was used for longer-term prediction.

6.1.4 Publication IV

AI-ARC Baltic Demo: Detecting Illegal Activities at Sea

Summary: This paper presents the AI-ARC system, which demonstrates how an advanced microservices architecture can detect and analyze complex illegal activities at sea. The system combines multiple AI-driven services with our route deviation detection (RDD). It utilizes AIS data, satellite imagery, environmental information, and anomaly detection pipelines for different use cases, such as grounding, oil spills, critical offshore infrastructures, illegal fishing, and smuggling. In a demonstration conducted in the Baltic Sea, realistic criminal incidents were tested using historical data to showcase how AI-driven microservices, such as ship movement prediction, anomaly recognition, and satellite-based services, can work together to enhance MSA. Operators interact through a visualization platform (map-based) called DigLT, which integrates and displays relevant alerts, enabling law enforcement, border security, and other maritime stakeholders to make decisions.

Objectives: A primary goal is to determine whether our ship movement prediction system can seamlessly integrate with other AI-based services for more complex anomaly detection, with an emphasis on smuggling in maritime contexts. The project specifically aims to evaluate the RDD system while fusing with other anomaly detection services (e.g., route deviations, suspicious rendezvous) to create a more robust detection process. Furthermore, we intend to demonstrate that multiple microservices, ranging from sensor-level data processing to high-level behavioral inference, can be coordinated effectively. Eventually, provide maritime security operators with an AI-fused platform that can handle events from environmental violations to covert smuggling activities, all displayed in a clear, centralized, and user-friendly interface.

Methods and data: Among these use cases, the smuggling service acts as a key aggregator, synthesizing suspicious indicators, referred to as ”weak signals”, from multiple specialized modules (mirco-services), as illustrated in Figure 22. For example, RISE's Meeting Detector identifies close-proximity encounters between vessels, particularly when one is not transmitting AIS data or during unusually prolonged rendezvous events. The Destination Detector, also developed by RISE, monitors deviations between a ship's actual path and its declared destination. Meanwhile, TREE's Vessel Behavior Anomaly Detector uses machine learning techniques, such as isolation forests and local outlier factors, to detect unusual speed or course patterns in AIS data. The results from this detector are further evaluated by FhG-EMI's reliability layer, which estimates the confidence level of the anomaly classification.

Our UTU RDD system, designed based on the concept of long-term ship movement prediction (Publication III), offers crucial insights through a sequential process. It begins by loading a three-day AIS dataset while setting aside the target vessel's AIS records for specialized analysis. The pipeline creates a dictionary of trajectories for the remaining vessels and builds a similar dictionary for the target vessel. It then divides the target vessel's coordinates, collected every three minutes, into 30-minute segments, effectively breaking down 12 hours of data into six segments. At each step, the LCSS similarity measurement method is applied to identify the closest historical trajectory for the relevant segment of the target vessel's route. Once all segments have been compared, the system assigns the most similar historical trajectory as the predicted future path and expected future behavior for the target vessel. Simultaneously, real-time data from the target vessel is cleansed, resampled, and compared using the Haversine distance to the predicted endpoint. If the distances consistently increase as new real-time data points are received, it indicates that the vessel is likely deviating from its forecasted path, suggesting a possible irregularity. On the other hand, if the distances steadily decrease, it implies that the vessel is following its expected route.

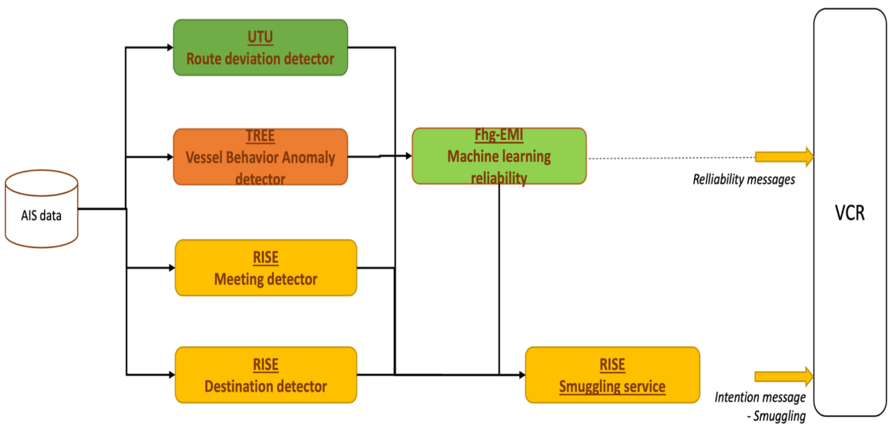

Figure 22. Micro-services used in smuggling detection system.

Once UTU RDD service alerts are generated, they feed into the smuggling aggregator alongside other microservices' signals, such as close rendezvous events or abrupt course changes—each accompanied by a confidence score. In accordance with the weighted decision-making mechanism, these confidence scores are multiplied by weights reflecting the relative importance of each microservice. A service identifying high-impact anomalies (e.g., RDD) may thus receive a higher weight, whereas signals deemed lower risk receive a smaller weight. The aggregator then computes a final anomaly score (Equation 12), and if that total surpasses a threshold, the system issues a high-priority "intention" alert. For example, a commercial vessel unexpectedly deviating from its declared route and subsequently meeting a smaller, non-AIS boat can accumulate sufficient weighted evidence to be flagged as potential smuggling activity. This weighted approach ensures that the most critical indicators, particularly those linked to illicit maneuvers, drive the ultimate detection outcome, enabling operators to respond more effectively.

Results and contribution:

The Baltic Sea demonstration in Karlskrona, Sweden, in September 2023, demonstrated that UTU RDD service, combined with a range of detection and interpretation services, significantly improved MSA. User evaluations (Swedish Coast Guards), gathered through questionnaires and operational feedback, indicated that automated alerts for smuggling, illegal fishing, and environmental crimes shorten response times and reduce the manual workload for operators. Additionally, the flexible microservices design proved adaptable to various use cases, from detecting oil spills to identifying suspicious rendezvous at sea. This integrated system fosters collaboration among various AI modules, demonstrating how advanced AI-driven micro-services can enhance anomaly detection systems and underscore its potential for real-world application in maritime security operations.

Table 5. LCSS Scores for the Identified Vessels. The last three digits of ships' MMSI numbers are not displayed for security and user privacy.

Ship with MMSI	LCSS Score
273342XYZ	0.8521
219136XYZ	0.7746
219170XYZ	0.7746
220010XYZ	0.7676
241093XYZ	0.7605
244265XYZ	0.7447
351607XYZ	0.7394
371467XYZ	0.6760

To evaluate the smuggling service, we ran a trial inspired by a high-profile drug-trafficking incident in Denmark. According to Danish media reports [130], authorities intercepted a Bahamas-flagged cargo vessel transferring cocaine to a small boat off the coast of Langeland on February 16, 2020. In our test scenario, we replicated key aspects of this event using AIS data and fed them into the smuggling service. As shown in Figure 23, the red line represents the vessel's real-time trajectory, while the blue line depicts the last known historical route. The system then applied the LCSS similarity measurement algorithm to match the incoming trajectory segments with a pool of historical AIS tracks. Table 5 illustrates how the highest LCSS score belonged to Ship MMSI 273342810 with 0.8521, meaning that its known path closely resembled the vessel under observation. Once flagged, our smuggling aggregator identified a pattern of route deviations and abnormal behavior highlighted by the black circle in Figure 23, thus raising an “intention” alert that could signal potential contraband transfer. This workflow demonstrates how the smuggling service, aided by LCSS-based matching and route deviation detection, effectively pinpoints anomalous maritime operations reminiscent of real-world smuggling cases.

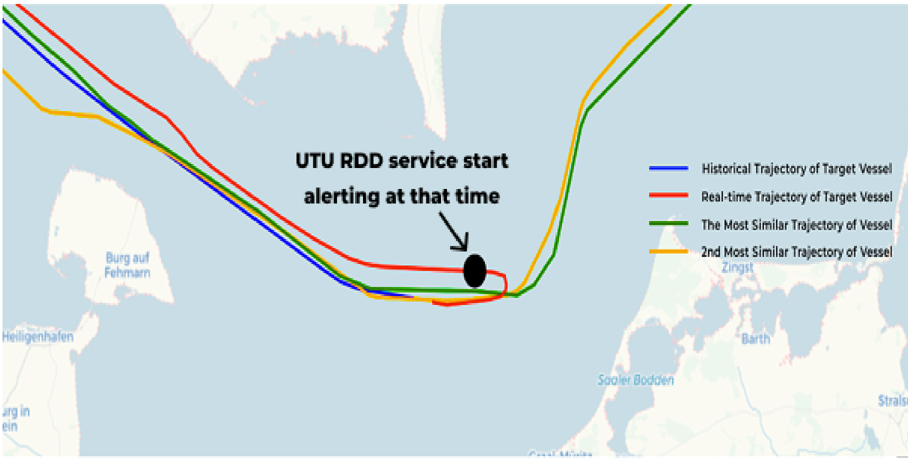

Figure 23. The performance of UTU RDD service on the Danish use case.

Author’s Contribution: The author of this thesis played a role in designing the smuggling detection approach, with emphasis on integrating UTU RDD service into the microservices pipeline. He implemented the Python code for the UTU RDD methodology, overseeing the pipeline’s sequential steps—data loading, dictionary creation, LCSS-based trajectory matching, and real-time Haversine distance evaluation. The author also participated in the Baltic Demo demonstration to ensure that route-deviation outputs were effectively combined with other detection modules to detect the Danish smuggling use case. Furthermore, he contributed to the design, user research study, and testing of the DigLT user interface designed to elegantly show the results of AI-based micro services for end users. Additionally, he was involved in writing the parts connected to the smuggling application and the UTU RDD service for Publication IV.

6.1.5 Publication V

Maritime vessel movement prediction: A temporal convolutional network model with optimal look-back window size determination

Summary: Accurate ship movement predictions are essential for ensuring safety and enhancing situational awareness in complex maritime transportation networks. Many existing AIS-based forecasting methods employ a fixed look-back window, overlooking how much training data is truly required. This paper introduces a novel framework that dynamically chooses an ideal look-back window based on user-specified prediction intervals. The approach begins with applying DBSCAN clustering and various data-cleaning strategies to eliminate irrelevant records and reduce noise in the raw AIS input. Next, we train the TCN algorithm using a one-month AIS dataset from the Baltic Sea (April 2023). We test different look-back window lengths to determine which one best captures the dynamic behaviors of the vessels for two prediction scenarios: 1-hour and 5-hour intervals. The findings indicate that this method effectively determines the minimum AIS sample volume required for accurate forecasting. This leads to more reliable predictions of both short-term and extended ship movements. Additionally, by identifying the optimal window size, the framework significantly reduces computational costs while operating continuously.

Objectives: This paper's objective is to explore how dynamically optimizing the look-back window in AIS-based ship movement prediction can enhance accuracy over varying time horizons, while also managing computational overhead. By systematically adjusting how much historical data is provided to the predictive model, the framework ensures that both short- and long-term forecasts capture enough contextual information to remain precise without overfitting.

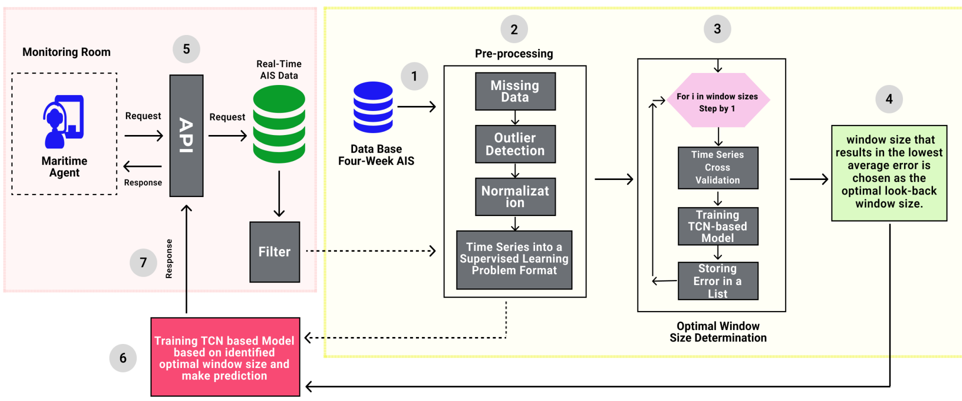

Figure 24. The architecture of the proposed prediction framework outlined in Publication IV.

Methods and data: This study focuses on predicting the target variable y_{t+k} accurately. To do this, we develop a time-series prediction pipeline that is specifically designed for the unique challenges of AIS data. First, the model must handle the non-linear and temporally dependent nature of ship trajectories, where vessel positions and kinematic features (speed, course) can fluctuate in often unpredictable ways [131]. Second, it necessitates selecting an appropriate look-back window (w) that balances the need for sufficient historical context against the risk of overfitting [132]. Finally, the pipeline is designed to adapt to diverse user-specified forecasting intervals, i.e., the different values of $(t+k)$, so that maritime agents with various operational objectives receive predictions aligned to their specific requirements [131].

To address the identified challenges, Publication IV proposed an ML-based prediction framework, illustrated in Figure 24. This framework integrates AIS data collection, preprocessing, model optimization, and real-time prediction into a cohesive pipeline, tackling the critical issues of dynamic window sizing and computational efficiency effectively. It begins with a four-week AIS database (Step 1), from which data undergo pre-processing (Step 2) involving missing data handling, outlier detection, and normalization; the resulting time series are then reformatted into a supervised learning problem. In the optimal window size determination phase (Steps 3 and 4), a TSCV procedure is used to train a TCN-based model with multiple look-back window candidates. Errors from each training run are stored in a list, and the window size yielding the lowest average error is deemed optimal. Meanwhile, real-time AIS data are acquired through an API (Step 5), filtered for the target vessel, and fed into the trained TCN model (Step 6). Once the maritime agent requests short-range or larger-range forecasts, the framework uses the identified smaller and efficient look-back window size to capture sufficient patterns. Ultimately, the predicted results are returned to the maritime agent (Step 7).

We also applied a unique preprocessing method to remove AIS data points generated due to ships' loitering. In particular, cargo and tanker vessels often remained near ports for extended periods, producing frequent, redundant location updates, commonly referred to as loitering behavior . To mitigate these non-essential observations, we employed the DBSCAN clustering algorithm, calibrating ϵ to the median distance between consecutive AIS points and setting $min\_samples = 1$. As a result, each point was either merged into a suitable cluster or treated as its own cluster, thereby minimizing noise while preserving relevant ship movements. Afterward, spline interpolation ensured a consistent 3-minute reporting interval, and the Min-Max scaler normalized all features to a 0- 1 range. This refined dataset, comprising over 3 million entries, proved to be more coherent and balanced for the TCN-based prediction model.

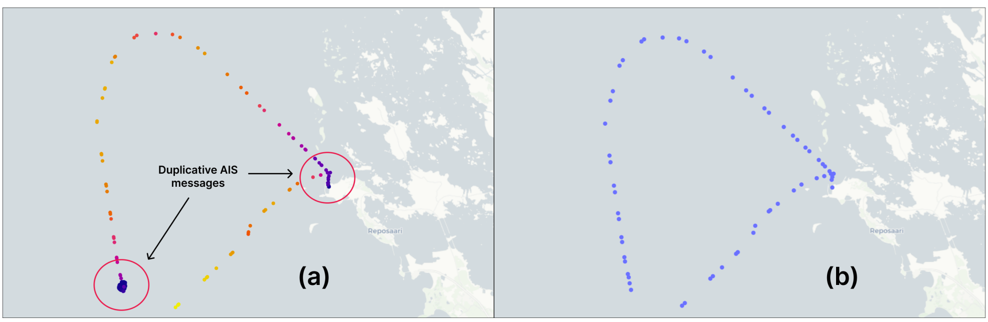

Figure 25. Example of removing loitering AIS data points using the DBSCAN-based approach. (a) Original dataset with large clusters of repetitive AIS messages during port stay, and (b) pruned dataset retaining only the essential route positions.

Results and contribution:

Tables 6 and 7 present a comparative analysis of three prediction models—Linear Regression (LR), LSTM, and TCN, for both 1-hour and 5-hour ship movement predictions under varying look-back window sizes. Table 3 shows that for short-term (1-hour) forecasting, the TCN model attains the lowest RMSE (0.035), MSE (0.0012), and highest R2 (0.98) at a 2-day look-back window, indicating that TCN captures short-range dynamics more effectively than LR or LSTM. In Table 7, aimed at 5-hour predictions, TCN similarly outperforms at a 4-day window (RMSE = 0.118, MSE = 0.0140, R2 = 0.95), underscoring its robustness when handling and learning from long-term time series data dependency.

Table 6. Comparison of LR, LSTM, and TCN model performances for 1-hour ship movement prediction across different look-back windows.

Look-back Window	Metrics	LR	LSTM	TCN
1d	RMSE	0.075	0.052	0.045
	MSE	0.0056	0.0027	0.0020
	R^2	0.89	0.94	0.96
2d	RMSE	0.115	0.041	0.035
	MSE	0.0132	0.0017	0.0012
	R^2	0.85	0.96	0.98
3d	RMSE	0.675	0.058	0.052
	MSE	0.455	0.0034	0.0027
	R^2	0.32	0.95	0.97
4d	RMSE	0.405	0.049	0.046
	MSE	0.164	0.0024	0.0021
	R^2	0.60	0.96	0.97

Figures 26 and 27 show the results for two different look-back window sizes: a two-day window for predicting 1 hour ahead and a four-day window for predicting 5 hours ahead. In Figure 26, the predicted latitude and longitude from the TCN-based model (red lines) closely match the vessel's actual path (blue lines). You can see this alignment in the time-series plots (panel a) and in the tight range of errors shown in panel (b). Even when we extend the prediction to 50 hours (panel c), the model still shows a good overlap between its predictions and the actual positions. This demonstrates that the model can be reliable over longer periods.

Figure 27 focuses on the 5-hour prediction with a four-day window. It also shows strong short-term tracking and very little difference in the error histograms. Even with larger time gaps (panel c), the predicted coordinates still match the actual route closely. In both figures (panel d), the final map displays nearly the same paths for the actual (blue/green) and predicted (red) positions.

This paper offers four main contributions to maritime ship movement prediction.

Table 7. Comparison of LR, LSTM, and TCN model performances for 5-hour ship movement prediction across different look-back windows.

Look-back Window	Metrics	LR	LSTM	TCN
1d	RMSE	0.275	0.210	0.215
	MSE	0.0756	0.0441	0.0462
	R^2	0.70	0.91	0.92
2d	RMSE	0.510	0.130	0.125
	MSE	0.2601	0.0169	0.0156
	R^2	0.64	0.92	0.94
3d	RMSE	0.880	0.140	0.120
	MSE	0.7744	0.0196	0.0144
	R^2	0.34	0.93	0.94
4d	RMSE	0.460	0.125	0.118
	MSE	0.2116	0.0156	0.0140
	R^2	0.22	0.93	0.95

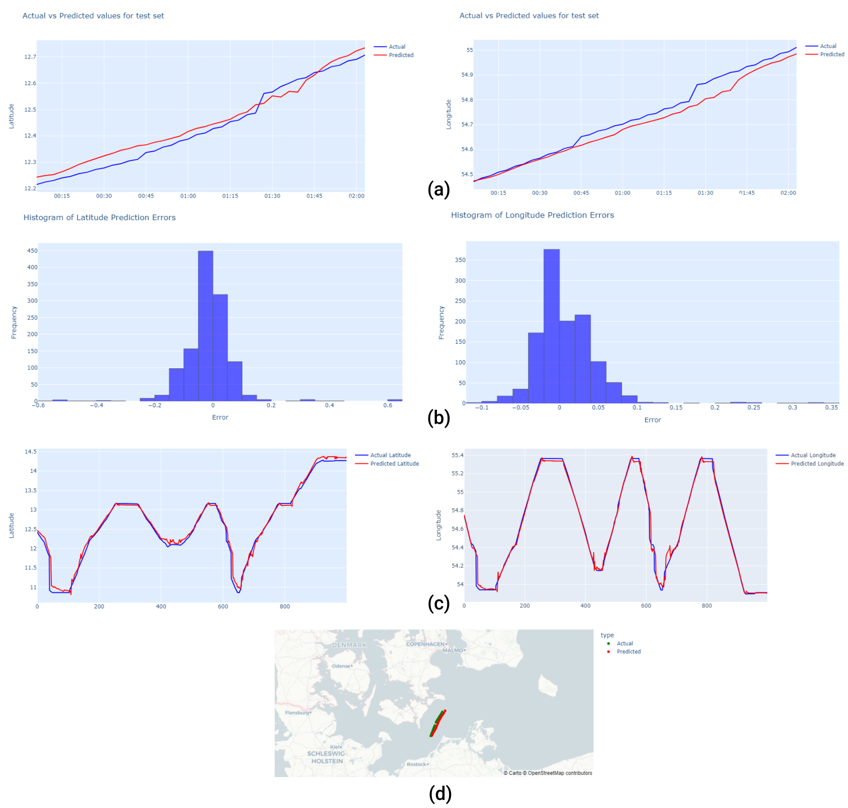

Figure 26. Prediction results for 1-hour interval.

First, it systematically determines how much AIS data is required for different forecasting intervals, a question often overlooked in prior studies. Second, it integrates a clustering-based approach to identify and remove loitering behaviors, along with a rigorous pre-processing pipeline to manage missing data, outliers, and data normalization, thus improving dataset quality. Third, the framework adopts a user-centric design, allowing maritime agents to specify desired forecasting horizons and then generating custom predictions that leverage an optimally selected look-back window. Finally, its efficacy is illustrated via two distinct forecasting intervals (1-hour and 5-hour), demonstrating practical viability for a variety of maritime operations.

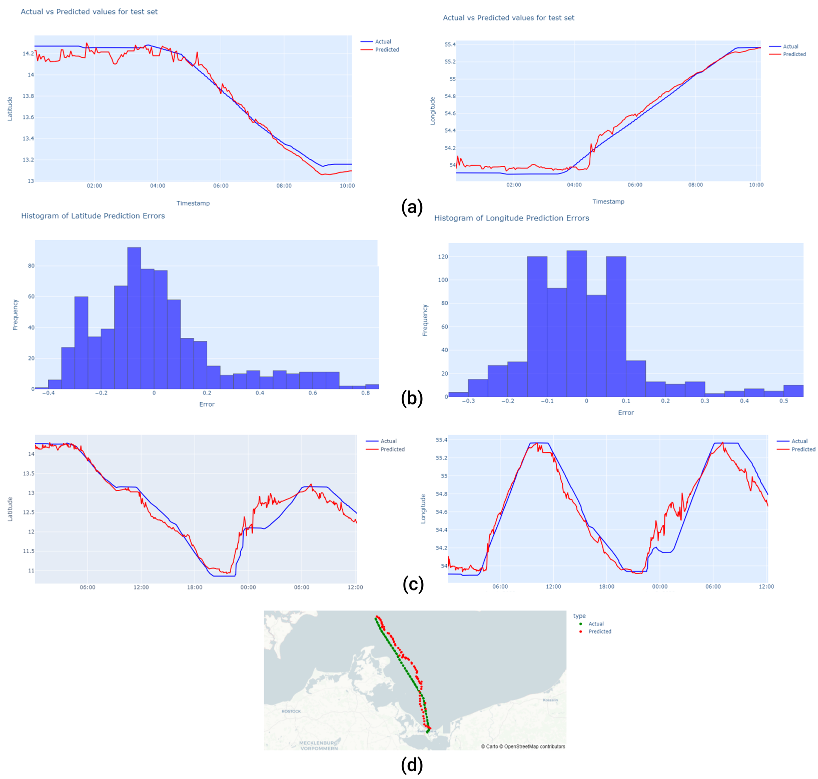

Figure 27. Prediction results for 5-hour interval.

Author's Contribution: The author of this thesis assumed a primary role in conceiving and implementing the TCN-based framework for determining optimal look-back window sizes. Specifically, he designed and coded the Python routines for data filtering, pre-processing, and model evaluation using the TSCV method. Beyond method development, he contributed to shaping the experimental setup and analyzing the results, ensuring that look-back window sizes were rigorously benchmarked for both short- and long-term prediction intervals. Additionally, he contributed to drafting key sections of the paper, including the methodology, results, and discussion.

7 Conclusions

7.1 Summary of the thesis

The timeline of this thesis project, as shown in Figure 28, began in August 2022 with a literature review that established the foundation for investigating maritime prediction methods and anomaly detection. From August 2022 to February 2023, the focus was on detecting abnormal ship behaviors by utilizing unsupervised ML techniques to identify "dark ships" and unexpected movements. The focus then shifted to long- and short-term movement prediction, culminating in a dynamic look-back window framework by January 2024, which adjusts prediction intervals based on user and seasonal needs. Around December 2024, the thesis integrated these findings into a smuggling detection service as a part of the AI-ARC project, to spot suspicious operations in real-time. Finally, as of April 2025, all these components converged into the doctoral writing phase, bringing together the various threads of research into a cohesive dissertation.

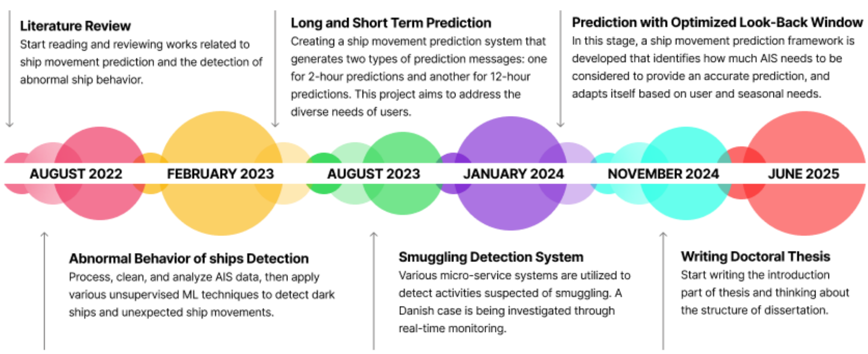

Figure 28. Timeline of the doctoral research outlining key milestones.

Chapter 1 discusses the motivations for improving situational awareness in maritime applications and presents the research questions that guide this study. It provides a broader context for why enhancing predictions of ship movements and detecting anomalies is crucial in today's digital landscape. Additionally, this chapter addresses the challenges posed by large AIS datasets, changing traffic patterns, and complex vessel behaviors.

Chapter 2 explains situational awareness in maritime settings, focusing on the AIS sensor and its essential role in tracking vessel movements. The chapter describes how AIS supports maritime traffic control and safety, highlighting its capabilities and limitations. It also notes the opportunity to gain insights from the data AIS provides. The chapter concludes by demonstrating the importance of AIS data for modern maritime surveillance systems, which vary by region and season.

Chapter 3 presents the process by which AIS data is gathered, refined, and prepared for subsequent modeling tasks. The chapter first introduces the Digitraffic API (Section 3.1), describing how AIS records are requested, stored, and updated in near real-time. Next, it details the outlier analysis and preprocessing routines (Section 3.2), covering four core steps: data cleaning, trajectory validation algorithm, simplification of ship trajectories, and detection of loitering movements. Furthermore, the chapter concludes with an exploratory data analysis (Section 3.3), providing initial insights into core AIS attributes, speed distributions, course patterns, and event frequencies across different vessel types.

Chapter 4 introduces the core methodologies used to model ship movement prediction. It begins by describing a similarity measurement technique (Section 4.1), highlighting the Symmetrized Segment-Path Distance for comparing different vessel trajectories. The chapter then delves into various DL approaches (Section 4.2), starting with overviews of RNNs and LSTM architectures, followed by a focus on TCNs. Practical considerations such as hyperparameter selection (Section 4.2.4) and evaluation metrics (Section 4.2.5) are also discussed, ensuring that the chosen model configurations are systematically tested and validated. Together, these techniques form a robust framework for both short-term and long-term vessel trajectory forecasting, leveraging AIS data in increasingly flexible and precise ways.

Chapter 5 provides a concise review of contemporary approaches to detecting ship anomalies at sea, covering data-driven methods (Section 5.2) in supervised and unsupervised ML. It underscores the importance of clustering-based approaches for identifying suspicious vessel behaviors, such as smuggling or unauthorized route deviations. By comparing different techniques, the chapter clarifies how clustering can serve as either a standalone detector or a complementary step in anomaly detection pipelines. Furthermore, the method of integrating different AI-based micro-services is also presented in Section 5.3 to detect complex abnormal behavior at sea, such as smuggling, dark ships, and oil spills.

Chapter 6 outlines the articles that comprise this thesis. Each article summary encompassed the methods and data used, the results obtained, the article's contributions, and the author's role in the work.

7.2 Discussions and outcomes

7.2.1 Thesis Findings

This section provides a clear summary of the thesis findings and how the research questions defined in Section 1.4 were addressed. The thesis developed data-driven frameworks that (1) flag abnormal behavior in shipping networks and (2) predict vessel movement across various time frames. By addressing the five research questions, it achieved its primary goal of improving MSA in the Baltic Sea.

To address RQ1, a structured literature review analyzed 45 peer-reviewed studies from 2012 to 2022 focused on ML methods for detecting abnormal vessel behaviors at sea. This study is presented in article I, and its findings are as follows. The methods fall into two main categories: trajectory-based approaches, which analyze movement patterns over time, and point-based approaches, which identify outliers, such as sudden speed changes or course deviations. The primary datasets used consist of AIS messages and SAR imagery, known for providing clear images day and night, in all weather, at high resolution. Various supervised and unsupervised ML techniques were applied, including KNN, SVM, and clustering algorithms. Findings reveal that recent studies have integrated DL-based multimodal architectures and clustering algorithms to manage large-scale AIS data and with minimal labeling. The research highlights challenges in the real-time operational use of data-driven models, including a lack of labeled data, limited explainability of predictions, difficulty in detecting complex anomalies, and the need to fuse data from different maritime sensors.

In relation to RQ2, the study (article II) examined the most common clustering-based algorithms to identify abnormal vessel behaviors in congested maritime areas of the Baltic Sea. The main objective was to detect dark ships, a technique frequently used by smugglers to remain undetected while engaging in illegal activities. Due to the complexity of this type of anomaly, key AIS features such as latitude, longitude, SOG, and COG have been integrated in various combinations of 2D and 3D inputs for clustering-based methods. According to the results, the study found that different combinations of these features yield varying results, which can lead to more effective dark ships. The proposed framework effectively identified a suspected dark ship and discovered another vessel exhibiting unusual spiral movements, a behavior that is rare for the targeted area. Additionally, it can provide users with the exact time of the anomalies upon request. All the clustering algorithms tested demonstrated potential in identifying these behaviors, with the K-means algorithm achieving the highest performance. Thus, it is now clear that clustering-based models can be effectively utilized for detecting abnormal behavior in shipping traffic.

Regarding RQ3, article III developed a unified framework capable of producing both short-term and long-term predictions of vessel movement using AIS data from the Baltic Sea. For short-term predictions, a feed-forward neural network was trained using five key AIS features: timestamp, latitude, longitude, SOG, and COG.

Evaluation was based on a custom ERP that calculates the Haversine distance between predicted and actual vessel positions. Results showed that the model achieved an ERP of less than 0.3 km for most vessel types, demonstrating its effectiveness in forecasting near-future positions with high spatial accuracy. For long-term predictions, the framework used a method called SSPD to compare current vessel paths with historical ones from the same port. This technique allowed us to make predictions extending beyond 10 hours and to identify the most likely future routes based on similar past trajectories. We found that tankers, cargo ships, and passenger ships had more predictable routes than fishing vessels, resulting in lower ERP values for long-term predictions. Additionally, we integrated the Apache Kafka API to manage real-time data updates, ensuring continuous scalability during testing in the AI-ARC system. These findings show that combining neural and similarity-based models can effectively meet the short- and long-term needs of maritime stakeholders.

The research presented in article IV addresses RQ4 and shows that a ship movement prediction system can be effectively integrated with other AI-driven services to detect complex maritime crimes such as smuggling. The AI-ARC system employs a modular microservice architecture, integrating data streams such as AIS data, satellite imagery, and anomaly detection services to enhance situational awareness. The UTU RDD module was one of the microservices, utilizing long-term trajectory prediction used the LCSS similarity algorithm. It compares real-time AIS updates of a target vessel to its predicted route, flagging deviations based on spatial distance calculations. A simulated smuggling scenario, based on a 2020 drug-trafficking case off the Danish coast, validated the system's effectiveness, highlighting significant deviations and generating alerts. In addition, a questionnaire completed by law enforcement stakeholders underscored the potential of microservices for addressing complex anomalies at sea, particularly in the Baltic Sea. Ultimately, these findings indicate the potential to improve MSA by integrating a ship movement prediction service within a broader AI framework.

In relation to RQ5, this study (article V) introduced a dynamic look-back window optimization strategy to enhance the performance of ship movement prediction models by tailoring the historical AIS input length according to prediction interval and trajectory complexity. Unlike conventional approaches that use a fixed-size input window (e.g., 30 or 60 minutes), this method applies a data-driven optimization process to identify the optimal historical window length for each prediction case. The optimization strategy was evaluated on a dataset covering one month of Baltic Sea AIS trajectories, segmented by vessel type and prediction interval (ranging from 15 to 120 minutes). Experimental results showed that prediction accuracy significantly improved when using dynamically sized look-back windows. For instance, at a 60-minute prediction horizon, the dynamically optimized model achieved an average error reduction of 16.4% compared to the fixed-window baseline. The greatest gains were observed for fishing and cargo vessels, which exhibited high variability in movement patterns and thus benefited from more adaptive window configurations. In addition to accuracy improvements, the dynamic approach reduced the computational burden by avoiding unnecessarily long input sequences, cutting model training time by up to 23% in some configurations. The study also introduced a user-centered optimization layer, allowing end-users (e.g., port operators, surveillance agents) to define their desired trade-off between accuracy and computational cost. This enabled task-specific model behavior, such as favoring faster processing for near-shore traffic monitoring or higher accuracy for strategic route forecasting. Therefore, dynamic window optimization significantly enhances both prediction quality and system scalability by adapting to vessel behavior, temporal range, and operational goals.

7.2.2 Implications

The findings presented in this thesis have promising implications for both academic research and real-world maritime operations. Academically, the introduced frameworks, particularly adaptive look-back window sizing and novel anomaly detection techniques, extend current literature on time-series modeling in dynamic, high-dimensional environments. These contributions offer a methodological advancement over traditional approaches by enabling more flexible temporal representations of vessel behavior, thereby strengthening the role of ML in maritime analytics and other domains characterized by irregular observational intervals.

From an operational perspective, this research has the potential to improve MSA systems utilized by authorities such as the Coast Guard, port administrators, and marine insurers. Enhanced predictions for both short- and long-term movements can aid in dynamic route planning, congestion management, and the early detection of potentially suspicious activities, such as AIS spoofing and ”dark” ship operations. This can be achieved through the use of similarity measurements and cluster-based techniques that do not rely on labeled data. These capabilities are critical for safeguarding maritime borders, enforcing regulatory compliance, and minimizing economic and security risks. Furthermore, the reduced computational complexity of the proposed models, achieved through selective input optimization and efficient architecture design, makes them more accessible for practical implementation.

Extensively, the ability to detect covert maritime behaviors contributes to global security goals, including the prevention of smuggling, trafficking, and unauthorized environmental violations. The interdisciplinary nature of this work also sets the stage for future collaboration between maritime authorities, ML researchers, and policymakers in building safer, more innovative, and more transparent shipping networks.

7.2.3 Limitations

The research described in this thesis has a few limitations. A fundamental limitation of this research arises from the restricted AIS data coverage, which is confined to the Baltic Sea region. This geographic constraint limits the generalizability of the proposed models to other maritime environments with different traffic patterns, vessel types, and regulatory conditions. Additionally, while smuggling and dark ship scenarios have been prioritized as critical use cases, officially verified examples of such events are rare, which limits the model's exposure to real-world illicit activities and reduces the diversity of training samples for anomaly detection. The absence of RADAR data further constrained the system's ability to perform multimodal sensor fusion, which is essential for detecting non-cooperative or AIS-silent vessels that evade surveillance.

Beyond data limitations, several challenges are related to ML models themselves. AIS data is often noisy and incomplete due to signal loss, timestamp inconsistencies, or hardware malfunction, requiring extensive pre-processing and imputation strategies. Even after cleaning, many ground-truth labels, especially those associated with suspicious or covert behavior, are either unavailable or uncertain, which complicates supervised training and limits the precision of evaluation metrics. From a model-centric perspective, DL architectures used in this research, while powerful, lack inherent interpretability. This black-box nature poses challenges for operational decision-making, particularly in high-stakes maritime security contexts where explainability is critical. Furthermore, these models may be sensitive to unforeseen patterns or behaviors not represented in the training data, leading to poor generalization when confronted with novel tactics employed by adversaries. Ecological variability, such as seasonal changes, variations in sea ice, or shifts in shipping routes, can further degrade model performance unless retraining or domain adaptation techniques are regularly applied.

Finally, practical deployment introduces another layer of complexity. Even if the models perform well under experimental conditions, their integration into maritime surveillance workflows depends on real-time processing capabilities, robust shipboard hardware, and seamless communication channels with command centers. Achieving such integration also requires ongoing stakeholder engagement, standardized protocols, and clear accountability frameworks, all of which are still works in progress by different high-tech companies and port authorities.

7.3 Future work

Three key directions, as shown in Figure 29, are emerging for advancing maritime analytics based on current frameworks. First, multimodal ML frameworks informed by recent breakthroughs in cross-domain data fusion offer a robust paradigm to unify heterogeneous data streams (AIS, RADAR/LiDAR, high-resolution satellite imagery, and probabilistic weather models) into an end-to-end predictive architecture. Various types of neural networks are utilized, such as transformer-based models for AIS data, convolutional networks for images, and graph models for vessel interactions. These systems effectively manage complex relationships among data while reducing noise from sensors. This approach supports efforts to create AI-driven digital twins in maritime environments, enabling the real-time detection of anomalies, such as illegal fishing and near-miss collisions, even in areas with incomplete data.

Second, integrating federated learning and edge computing could address the latency and privacy constraints inherent in centralized maritime surveillance systems. By using simple models on vessels trained with data from AIS and RADAR, predictive models can make predictions without relying on a central system. This is crucial for complying with new regulations like the EU's Data Governance Act. Additionally, explainable AI techniques, such as attention mapping or counterfactual analysis, can enhance trust in the results.

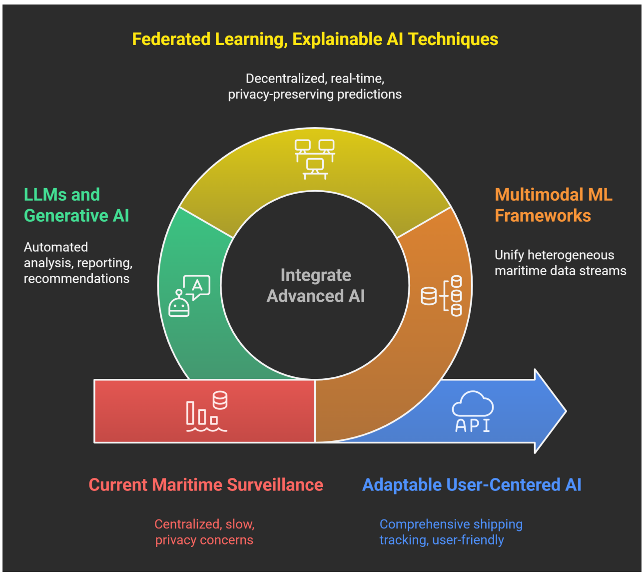

Figure 29. Future research directions for an advanced maritime analysis platform.

Third, the emergence of Large Language Models (LLMs) and generative AI holds great potential for maritime agents. These AI systems can analyze textual inputs and interpret sensor data, generate analyses, and draft comprehensive reports automatically. For instance, if a suspicious route is detected, an LLM-based assistant could summarize relevant regulations, compare the activity to historical cases, and outline potential enforcement actions. By combining multimodal ML with LLM-driven reasoning, future maritime systems could offer promising, adaptable, in-depth, and user-centered AI solutions.

List of References

[1] Ralph Davis. *The rise of the English shipping industry in the seventeenth and eighteenth centuries*. Number 48. Oxford University Press, 2012.

[2] Charles K Harley. The shift from sailing ships to steamships, 1850-1890: a study in technological change and its diffusion. *Essays on a mature economy: Britain after*, 15, 1840.

[3] M Tadros, M Ventura, and C Guedes Soares. Review of current regulations, available technologies, and future trends in the green shipping industry. *Ocean Engineering*, 280:114670, 2023.

[4] Adolf KY Ng. Container liner shipping, port development and competition. *Maritime logistics: Contemporary issues*, pages 5–28, 2012.

[5] Brian J Cudahy. *Box boats: How container ships changed the world*. Fordham Univ Press, 2007.

[6] Jan Otto de Kat and Jad Mouawad. Green ship technologies. *Sustainable Shipping: A Cross-Disciplinary View*, pages 33–92, 2019.

[7] Sarang Thombre, Zheng Zhao, Henrik Ramm-Schmidt, José M Vallet García, Tuomo Malkamäki, Sergey Nikolskiy, Toni Hammarberg, Hiski Nuortie, M Zahidul H Bhuiyan, Simo Särkkä, et al. Sensors and ai techniques for situational awareness in autonomous ships: A review. *IEEE transactions on intelligent transportation systems*, 23(1):64–83, 2020.

[8] Issam El Naqa and Martin J Murphy. *What is machine learning?* Springer, 2015.

[9] Yoshua Bengio, Ian Goodfellow, and Aaron Courville. *Deep learning*, volume 1. MIT press Cambridge, MA, USA, 2017.

[10] Yan Kang He, Di Zhang, Jinfen Zhang, MY Zhang, and TW Li. Ship route planning using historical trajectories derived from ais data. *TransNav, International Journal on Marine Navigation and Safety od Sea Transportation*, 13(1):69–76, 2019.

[11] Sara El Mekkaoui, Loubna Benabbou, and Abdelaziz Berrado. Predicting ships estimated time of arrival based on ais data. In *Proceedings of the 13th International Conference on Intelligent Systems: Theories and Applications*, pages 1–6, 2020.

[12] Andreas Tritsarolis, Eva Chondrodima, Nikos Pelekis, and Yannis Theodoridis. Vessel collision risk assessment using ais data: A machine learning approach. In *2022 23rd IEEE International Conference on Mobile Data Management (MDM)*, pages 425–430. IEEE, 2022.

[13] Da-wei Gao, Yong-sheng Zhu, Jin-fen Zhang, Yan-kang He, Ke Yan, and Bo-ran Yan. A novel mp-lstm method for ship trajectory prediction based on ais data. *Ocean Engineering*, 228: 108956, 2021.

[14] Hao Rong, AP Teixeira, and C Guedes Soares. A framework for ship abnormal behaviour detection and classification using ais data. *Reliability Engineering & System Safety*, 247:110105, 2024.

[15] European Maritime Safety Agency. Annual overview of marine casualties and incidents, 2022.

[16] Aleksander Gerson. Stranding of the mega-ship ever given in the suez canal: Causes, consequences, and lessons to be learned. In *The Suez Canal: Past Lessons and Future Challenges*, pages 231–252. Springer, 2023.

[17] British Broadcasting Corporation (BBC). Suez Canal: Ever Given container ship shifted from shoreline. `https://www.bbc.com/news/world-middle-east-56559904`, 29 March 2021. [Accessed 11-11-2024].

[18] Ileana Citaristi. United nations office on drugs and crime—unodc. In *The Europa Directory of International Organizations 2022*, pages 248–252. Routledge, 2022.

[19] Elisabeth Braw. The threats posed by the global shadow fleet and how to stop it, 2024. URL `https://www.atlanticcouncil.org/in-depth-research-reports/report/the-threats-posed-by-the-global-shadow-fleet-and-how-to-stop-it/`.

[20] Trung Nguyen. The challenges of dark ships to the safety and security of commercial shipping and the way forward. *Asia-Pacific Journal of Ocean Law and Policy*, 8(2):310–328, 2023.

[21] Carine Dominguez-Péry, Lakshmi Narasimha Raju Vuddaraju, Isabelle Corbett-Etchevers, and Rana Tassabehji. Reducing maritime accidents in ships by tackling human error: a bibliometric review and research agenda. *Journal of Shipping and Trade*, 6:1–32, 2021.

[22] Lucas May Petry, Amilcar Soares, Vania Bogorny, Bruno Brandoli, and Stan Matwin. Challenges in vessel behavior and anomaly detection: From classical machine learning to deep learning. In *Advances in Artificial Intelligence: 33rd Canadian Conference on Artificial Intelligence, Canadian AI 2020, Ottawa, ON, Canada, May 13–15, 2020, Proceedings 33*, pages 401–407. Springer, 2020.

[23] Dominik Filipiak, Milena Stróżyna, Krzysztof Wecel, and Witold Abramowicz. Anomaly detection in the maritime domain: comparison of traditional and big data approach. In *Proceedings of the NATO IST-160-RSM Specialists' Meeting on Big Data and Artificial Intelligence for Military Decision Making*, 2018.

[24] Bohan Zhang, Katsutoshi Hirayama, Hongxiang Ren, Delong Wang, and Haijiang Li. Ship anomalous behavior detection using clustering and deep recurrent neural network. *Journal of Marine Science and Engineering*, 11(4):763, 2023.

[25] Weili Jiang, Ying Zhou, Zikai Song, and Jun Ye. Comprehensive analysis scheme for abnormal behavior of maritime targets. In *International Conference on Cognitive based Information Processing and Applications*, pages 587–601. Springer, 2023.

[26] Claudio V Ribeiro, Aline Paes, and Daniel de Oliveira. Ais-based maritime anomaly traffic detection: A review. *Expert Systems with Applications*, 231:120561, 2023.

[27] Zu Lin, Weiqi Yue, Jie Huang, and Jian Wan. Ship trajectory prediction based on the ttcn-attention-gru model. *Electronics*, 12(12):2556, 2023.

[28] Wentao Wang, Wei Xiong, Xue Ouyang, and Luo Chen. Tptrans: Vessel trajectory prediction model based on transformer using ais data. *ISPRS International Journal of Geo-Information*, 13 (11):400, 2024.

[29] Markus Hansen and Salman Nazir. Situational awareness and decision-making in maritime operations: A cognitive perspective. In *International Conference on Human-Computer Interaction*, pages 165–189. Springer, 2024.

[30] FG Dias, JF Neves, VP da Conceição, and VJAS Lobo. Maritime situational awareness, the singular approach of a dual-use navy. *Scientific Bulletin" Mircea cel Batran" Naval Academy*, 21(1):1–14, 2018.

[31] Andrea Frontini. The european union maritime security strategy: sailing uncharted waters? *European Policy Center (EPC)*, 26, 2014.

[32] Tolga Ahmet Gülcan and Kadir Emrah Erginer. National and international maritime situational awareness model examples and the effects of north stream pipelines sabotage. *International Journal of Critical Infrastructure Protection*, 42:100624, 2023.

[33] Mica R. Endsley and Debra G. Jones. Situation awareness oriented design: Review and future directions. *International Journal of Human–Computer Interaction*, 40(7):1487–1504, 2024. doi: 10.1080/10447318.2024.2318884.

[34] Sergiu Lupu, Andrei Pocora, Florin Nicolae, Cosmin Katona, and Elena Carmen Lupu. The role of the maritime simulators in evaluating the practical competencies of graduated students in order to accomplish officer of the watch duties onboard vessels. In *The International Scientific Conference eLearning and Software for Education*, volume 1, page 216. " Carol I" National Defence University, 2016.

[35] Mohammad-Hashem Haghbayan, Fahimeh Farahnakian, Jonne Poikonen, Markus Laurinen, Paavo Nevalainen, Juha Plosila, and Jukka Heikkonen. An efficient multi-sensor fusion ap-

proach for object detection in maritime environments. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 2163–2170. IEEE, 2018.

[36] Fahimeh Farahnakian and Jukka Heikkonen. Deep learning based multi-modal fusion architectures for maritime vessel detection. *Remote Sensing*, 12(16):2509, 2020.

[37] Adyasha Mohanty and Grace Gao. A survey of machine learning techniques for improving global navigation satellite systems. *EURASIP Journal on Advances in Signal Processing*, 2024 (1):73, 2024.

[38] JJ Gehrt, S Liu, M Nitsch, W Bruhn, S Rohde, R Zweigel, and D Abel. Robust and reliable multi-sensor navigation filter for maritime application. *IFAC-PapersOnLine*, 53(2):14482–14487, 2020.

[39] Defu Yang, Mahmud Iwan Solihin, Yawen Zhao, Benchun Yao, Chaoran Chen, Bingyu Cai, and Affiani Machmudah. A review of intelligent ship marine object detection based on rgb camera. *IET Image Processing*, 18(2):281–297, 2024.

[40] Sarang Thombre, Zheng Zhao, Henrik Ramm-Schmidt, José M. Vallet García, Tuomo Malkamäki, Sergey Nikolskiy, Toni Hammarberg, Hiski Nuortie, M. Zahidul H. Bhuiyan, Simo Särkkä, and Ville V. Lehtola. Sensors and ai techniques for situational awareness in autonomous ships: A review. *IEEE Transactions on Intelligent Transportation Systems*, 23(1):64–83, 2022. doi: 10.1109/TITS.2020.3023957.

[41] Brian J Tetreault. Use of the automatic identification system (ais) for maritime domain awareness (mda). In *Proceedings of Oceans 2005 Mts/Ieee*, pages 1590–1594. Ieee, 2005.

[42] Rosalie Balkin. The international maritime organization and maritime security. *Tul. Mar. LJ*, 30: 1, 2006.

[43] IMO Solas. International convention for the safety of life at sea (solas). *International Maritime Organization, London*, 2002.

[44] U.S. Coast Guard Navigation Center. Class a ais position report (messages 1, 2, and 3), 2024. URL `https://www.navcen.uscg.gov/ais-class-a-reports`. Accessed: 2024-12-24.

[45] Konstantina Bereta, Konstantinos Chatzikokolakis, and Dimitris Zissis. Maritime reporting systems. *Guide to Maritime Informatics*, pages 3–30, 2021.

[46] Shijie Rong and Mao Wang. Multiple access protocol for ad-hoc networks in vhf data exchange system. In *2018 IEEE 4th International Conference on Computer and Communications (ICCC)*, pages 629–633. IEEE, 2018.

[47] Hamed Yaghoubi Shahir, Uwe Glässer, Narek Nalbandyan, and Hans Wehn. Maritime situation analysis: A multi-vessel interaction and anomaly detection framework. In *2014 IEEE Joint Intelligence and Security Informatics Conference*, pages 192–199. IEEE, 2014.

[48] Dong Yang, Lingxiao Wu, and Shuaian Wang. Can we trust the ais destination port information for bulk ships?–implications for shipping policy and practice. *Transportation Research Part E: Logistics and Transportation Review*, 149:102308, 2021.

[49] Mahmood Safaei, Shahla Asadi, Maha Driss, Wadii Boulila, Abdullah Alsaeedi, Hassan Chizari, Rusli Abdullah, and Mitra Safaei. A systematic literature review on outlier detection in wireless sensor networks. *Symmetry*, 12(3):328, 2020.

[50] Jason W Osborne and Amy Overbay. The power of outliers (and why researchers should always check for them). *Practical Assessment, Research, and Evaluation*, 9(1):6, 2019.

[51] Abbas Harati-Mokhtari, Alan Wall, Philip Brooks, and Jin Wang. Automatic identification system (ais): data reliability and human error implications. *the Journal of Navigation*, 60(3):373–389, 2007.

[52] Gabor Visky, Alexander Rohl, Sokratis Katsikas, and Olaf Maennel. Ais data analysis: Reality in the sea of echos. In *2024 IEEE 49th Conference on Local Computer Networks (LCN)*, pages 1–7. IEEE, 2024.

[53] Priank Prasad, Vishesh Vatsal, and Rajarshi Roy Chowdhury. Maritime vessel route extraction and automatic information system (ais) spoofing detection. In *2021 International Conference on*

Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT), pages 1–11. IEEE, 2021.

[54] A Karim, S Utama, and MA Arifin. Algorithms for detecting vessel spoofing of space-based ais data. In *AIP Conference Proceedings*, volume 2366. AIP Publishing, 2021.

[55] Alan Saalfeld. Topologically consistent line simplification with the douglas-peucker algorithm. *Cartography and Geographic Information Science*, 26(1):7–18, 1999.

[56] Jingxian Liu, Huanhuan Li, Zaili Yang, Kefeng Wu, Yi Liu, and Ryan Wen Liu. Adaptive douglas-peucker algorithm with automatic thresholding for ais-based vessel trajectory compression. *IEEE Access*, 7:150677–150692, 2019.

[57] Ting Zhang, Zhiming Wang, and Peiliang Wang. A method for compressing ais trajectory based on the adaptive core threshold difference douglas–peucker algorithm. *Scientific Reports*, 14(1): 21408, 2024.

[58] Dingsheng Deng. Dbscan clustering algorithm based on density. In *2020 7th international forum on electrical engineering and automation (IFEEA)*, pages 949–953. IEEE, 2020.

[59] Haozhou Wang, Han Su, Kai Zheng, Shazia Sadiq, and Xiaofang Zhou. An effectiveness study on trajectory similarity measures. In *Proceedings of the Twenty-Fourth Australasian Database Conference-Volume 137*, pages 13–22, 2013.

[60] Donald J Berndt and James Clifford. Using dynamic time warping to find patterns in time series. In *Proceedings of the 3rd international conference on knowledge discovery and data mining*, pages 359–370, 1994.

[61] Michail Vlachos, George Kollios, and Dimitrios Gunopulos. Discovering similar multidimensional trajectories. In *Proceedings 18th international conference on data engineering*, pages 673–684. IEEE, 2002.

[62] Lei Chen, M Tamer Özsu, and Vincent Oria. Robust and fast similarity search for moving object trajectories. In *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, pages 491–502, 2005.

[63] Thomas Eiter and Heikki Mannila. Computing discrete fréchet distance. 1994.

[64] Daniel P Huttenlocher, Gregory A. Klanderman, and William J Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on pattern analysis and machine intelligence*, 15(9):850–863, 1993.

[65] Philippe Besse, Brendan Guillouet, Jean-Michel Loubes, and Royer François. Review and perspective for distance based trajectory clustering. *arXiv preprint arXiv:1508.04904*, 2015.

[66] Nehal Magdy, Mahmoud A Sakr, Tamer Mostafa, and Khaled El-Bahnasy. Review on trajectory similarity measures. In *2015 IEEE seventh international conference on Intelligent Computing and Information Systems (ICICIS)*, pages 613–619. IEEE, 2015.

[67] Ali Dorosti, Ali Asghar Alesheikh, and Mohammad Sharif. Measuring trajectory similarity based on the spatio-temporal properties of moving objects in road networks. *Information*, 15(1): 51, 2024.

[68] Ian R Cleasby, Ewan D Wakefield, Barbara J Morrissey, Thomas W Bodey, Steven C Votier, Stuart Bearhop, and Keith C Hamer. Using time-series similarity measures to compare animal movement trajectories in ecology. *Behavioral Ecology and Sociobiology*, 73:1–19, 2019.

[69] David A Keatley and David D Clarke. Crime linkage: Finding a behavioral fingerprint using the “path similarity metric”. *Journal of Police and Criminal Psychology*, 35(2):240–246, 2020.

[70] Roniel S De Sousa, Azzedine Boukerche, and Antonio AF Loureiro. Vehicle trajectory similarity: Models, methods, and applications. *ACM Computing Surveys (CSUR)*, 53(5):1–32, 2020.

[71] Marija Todoric and Toni Mastelic. Comparison of similarity measures for trajectory clustering-aviation use case. *Journal of Communications Software and Systems*, 19(2):178–187, 2023.

[72] Jiawei Han, Jian Pei, and Hanghang Tong. *Data mining: concepts and techniques*. Morgan kaufmann, 2022.

[73] S Hochreiter. Long short-term memory. *Neural Computation MIT-Press*, 1997.

[74] Kyunghyun Cho. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.

[75] I Sutskever. Sequence to sequence learning with neural networks. *arXiv preprint arXiv:1409.3215*, 2014.

[76] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6645–6649. Ieee, 2013.

[77] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.

[78] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

[79] 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.

[80] Yongfeng Suo, Wenke Chen, Christophe Claramunt, and Shenhua Yang. A ship trajectory prediction framework based on a recurrent neural network. *Sensors*, 20(18):5133, 2020.

[81] R Pascanu. On the difficulty of training recurrent neural networks. *arXiv preprint arXiv:1211.5063*, 2013.

[82] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv:1803.01271*, 2018.

[83] F Yu. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.

[84] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.

[85] Karim Guirguis, Christoph Schorn, Andre Guntoro, Sherif Abdulatif, and Bin Yang. Seld-tcn: Sound event localization & detection via temporal convolutional networks. In *2020 28th European Signal Processing Conference (EUSIPCO)*, pages 16–20. IEEE, 2021.

[86] Songhua Hu and Chenfeng Xiong. High-dimensional population inflow time series forecasting via an interpretable hierarchical transformer. *Transportation research part C: emerging technologies*, 146:103962, 2023.

[87] Md Riyad Hossain and Douglas Timmer. Machine learning model optimization with hyper parameter tuning approach. *Glob. J. Comput. Sci. Technol. D Neural Artif. Intell*, 21(2):31, 2021.

[88] Xiaoyu Yuan, Jiawei Wang, Guang Zhao, and Hongbo Wang. Comprehensive study on optimizing inland waterway vessel routes using ais data. *Journal of Marine Science and Engineering*, 12(10):1775, 2024.

[89] Wayan Mahardhika Wijaya and Yasuhiro Nakamura. Port performance indicators construction based on the ais-generated trajectory segmentation and classification. *International Journal of Data Science and Analytics*, pages 1–20, 2024.

[90] Tanmay Bisen, Aastha Shayla, and Susham Biswas. Predicting oil spills in real-time: A machine learning and ais data-driven approach. In *Networking and Parallel/Distributed Computing Systems: Volume 18*, pages 99–115. Springer, 2024.

[91] EunSu Lee, Amit J Mokashi, Sang Young Moon, and GeunSub Kim. The maturity of automatic identification systems (ais) and its implications for innovation. *Journal of Marine Science and Engineering*, 7(9):287, 2019.

[92] Arnstein Vestre, Azzeddine Bakdi, Erik Vanem, and Øystein Engelhardtsen. Ais-based near-collision database generation and analysis of real collision avoidance manoeuvres. *The Journal of Navigation*, 74(5):985–1008, 2021.

[93] Kei Kanamoto, Liwen Murong, Minato Nakashima, and Ryuichi Shibasaki. Can maritime big data be applied to shipping industry analysis? focussing on commodities and vessel sizes of dry bulk carriers. *Maritime Economics & Logistics*, 23:211–236, 2021.

[94] Maria Riveiro, Giuliana Pallotta, and Michele Vespe. Maritime anomaly detection: A review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(5):e1266, 2018.

[95] Kira Kowalska and Leto Peel. Maritime anomaly detection using gaussian process active learning. In *2012 15th International Conference on Information Fusion*, pages 1164–1171. IEEE, 2012.

[96] Thomas Stach, Yann Kinkel, Manfred Constapel, and Hans-Christoph Burmeister. Maritime anomaly detection for vessel traffic services: a survey. *Journal of Marine Science and Engineering*, 11(6):1174, 2023.

[97] Youcef Djenouri, Asma Belhadi, Djamel Djenouri, Gautam Srivastava, and Jerry Chun-Wei Lin. Intelligent deep fusion network for anomaly identification in maritime transportation systems. *IEEE Transactions on Intelligent Transportation Systems*, 24(2):2392–2400, 2022.

[98] Abdoulaye Sidibé and Gao Shu. Study of automatic anomalous behaviour detection techniques for maritime vessels. *The journal of Navigation*, 70(4):847–858, 2017.

[99] Amanpreet Singh, Narina Thakur, and Aakanksha Sharma. A review of supervised machine learning algorithms. In *2016 3rd International Conference on Computing for Sustainable Global Development (INDIACom)*, pages 1310–1315. Ieee, 2016.

[100] Chamali Gamage, Randima Dinalankara, Jagath Samarabandu, and Akila Subasinghe. A comprehensive survey on the applications of machine learning techniques on maritime surveillance to detect abnormal maritime vessel behaviors. *WMU Journal of Maritime Affairs*, 22(4):447–477, 2023.

[101] Trevor Hastie, Robert Tibshirani, Jerome Friedman, Trevor Hastie, Robert Tibshirani, and Jerome Friedman. Unsupervised learning. *The elements of statistical learning: Data mining, inference, and prediction*, pages 485–585, 2009.

[102] Daheng Zhang, Yingjun Zhang, and Chuang Zhang. Data mining approach for automatic shiproute design for coastal seas using ais trajectory clustering analysis. *Ocean Engineering*, 236: 109535, 2021.

[103] J Swarndeep Saket and Sharnil Pandya. An overview of partitioning algorithms in clustering techniques. *International Journal of Advanced Research in Computer Engineering & Technology (IJARCET)*, 5(6):1943–1946, 2016.

[104] Kristina P Sinaga and Miin-Shen Yang. Unsupervised k-means clustering algorithm. *IEEE access*, 8:80716–80727, 2020.

[105] Noor Kamal Kaur, Usvir Kaur, and Dheerendra Singh. K-medoid clustering algorithm-a review. *Int. J. Comput. Appl. Technol*, 1(1):42–45, 2014.

[106] Pan Sheng and Jingbo Yin. Extracting shipping route patterns by trajectory clustering model based on automatic identification system data. *Sustainability*, 10(7):2327, 2018.

[107] Fionn Murtagh and Pedro Contreras. Algorithms for hierarchical clustering: an overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 2(1):86–97, 2012.

[108] Ashutosh Karna and Karina Gibert. Automatic identification of the number of clusters in hierarchical clustering. *Neural Computing and Applications*, 34(1):119–134, 2022.

[109] Xingcheng Ran, Yue Xi, Yonggang Lu, Xiangwen Wang, and Zhenyu Lu. Comprehensive survey on hierarchical clustering algorithms and the recent developments. *Artificial Intelligence Review*, 56(8):8219–8264, 2023.

[110] Huanhuan Li, Jasmine Siu Lee Lam, Zaili Yang, Jingxian Liu, Ryan Wen Liu, Maohan Liang, and Yan Li. Unsupervised hierarchical methodology of maritime traffic pattern extraction for knowledge discovery. *Transportation Research Part C: Emerging Technologies*, 143:103856, 2022.

[111] Lianhui Wang, Pengfei Chen, Linying Chen, and Junmin Mou. Ship ais trajectory clustering: An hdbscan-based approach. *Journal of Marine Science and Engineering*, 9(6), 2021. ISSN 2077-1312. doi: 10.3390/jmse9060566. URL `https://www.mdpi.com/2077-1312/9/6/566`.

[112] Pinki Rani et al. A survey on sting and clique grid based clustering methods. *International Journal of Advanced Research in Computer Science*, 8(5), 2017.

[113] Huanhuan Li, Jingxian Liu, Ryan Wen Liu, Naixue Xiong, Kefeng Wu, and Tai-hoon Kim. A dimensionality reduction-based multi-step clustering method for robust vessel trajectory analysis. *Sensors*, 17(8):1792, 2017.

[114] K Kameshwaran and K Malarvizhi. Survey on clustering techniques in data mining. *International Journal of Computer Science and Information Technologies*, 5(2):2272–2276, 2014.

[115] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17:395–416, 2007.

[116] Kaijun Wang, Junying Zhang, Dan Li, Xinna Zhang, and Tao Guo. Adaptive affinity propagation clustering. *arXiv preprint arXiv:0805.1096*, 2008.

[117] Brendan J Frey and Delbert Dueck. Clustering by passing messages between data points. *science*, 315(5814):972–976, 2007.

[118] Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine learning*, 109(2):373–440, 2020.

[119] Ioannis Kontopoulos, Konstantinos Chatzikokolakis, Konstantinos Tserpes, and Dimitris Zissis. Classification of vessel activity in streaming data. In *Proceedings of the 14th ACM International Conference on Distributed and Event-based Systems*, pages 153–164, 2020.

[120] Arnaud Vandecasteele, Rodolphe Devillers, and Aldo Napoli. A semi-supervised learning framework based on spatio-temporal semantic events for maritime anomaly detection and behavior analysis. In *CoastGIS 2013-The 11th International Symposium for GIS and Computer Cartography for Coastal Zone Management*, pages 4–pages, 2013.

[121] João Nobre, EJ Solteiro Pires, and Arsénio Reis. Anomaly detection in microservice-based systems. *Applied Sciences*, 13(13):7891, 2023.

[122] Mark Richards. *Microservices vs. service-oriented architecture*. O'Reilly Media Sebastopol, 2015.

[123] Nicola Dragoni, Saverio Giallorenzo, Alberto Lluch Lafuente, Manuel Mazzara, Fabrizio Montesi, Ruslan Mustafin, and Larisa Safina. Microservices: yesterday, today, and tomorrow. *Present and ulterior software engineering*, pages 195–216, 2017.

[124] Shan Suthaharan and Shan Suthaharan. Support vector machine. *Machine learning models and algorithms for big data classification: thinking with examples for effective learning*, pages 207–235, 2016.

[125] Leif E Peterson. K-nearest neighbor. *Scholarpedia*, 4(2):1883, 2009.

[126] Sean R Eddy. Hidden markov models. *Current opinion in structural biology*, 6(3):361–365, 1996.

[127] Muhammad Ali Syakur, B Khusnul Khotimah, EMS Rochman, and Budi Dwi Satoto. Integration k-means clustering method and elbow method for identification of the best customer profile cluster. In *IOP conference series: materials science and engineering*, volume 336, page 012017. IOP Publishing, 2018.

[128] Amira M El-Mandouh, Hamdi A Mahmoud, Laila A Abd-Elmegid, and Mohamed H Haggag. Optimized k-means clustering model based on gap statistic. *International Journal of Advanced Computer Science and Applications*, 10(1):183–188, 2019.

[129] Eva Patel and Dharmender Singh Kushwaha. Clustering cloud workloads: K-means vs gaussian mixture model. *Procedia computer science*, 171:158–167, 2020.

[130] Fox News. Denmark finds huge cocaine stash on cargo ship, 27 arrested. `https://www.foxnews.com/world/denmark-finds-huge-cocaine-stash-on-cargo-ship-27-arrested`, 2020. [Online; accessed 25 March 2025].

[131] Wenbo Zhao, Dezhi Wang, Kai Gao, Jiani Wu, and Xinghua Cheng. Large-scale long-term prediction of ship ais tracks via linear networks with a look-back window decomposition scheme of time features. *Journal of Marine Science and Engineering*, 11(11):2132, 2023.

[132] Panayu Keelawat, Nattapong Thammasan, Masayuki Numao, and Boonserm Kijsirikul. A comparative study of window size and channel arrangement on eeg-emotion recognition using deep cnn. *Sensors*, 21(5):1678, 2021.

Original Publications

Farshad Farahnakian & Jukka Heikkonen
& Paavo Nevalainen
Title: Abnormal Behaviour Detection by Using Machine Learning-Based Approaches in the Marine Environment: A Literature Survey

2022 International Conference on Electrical, Computer and Energy Technologies (ICECET), Prague, Czech Republic, 2022, pp. 1-11, doi: 10.1109/ICECET55527.2022.9872905.

I

Farshad Farahnakian & Florent Nicolas & Fahimeh Farahnakian & Paavo Nevalainen & Javad Sheikh & Jukka Heikkonen & Csaba Raduly-Baka

Title: A Comprehensive Study of Clustering-Based Techniques for Detecting Abnormal Vessel Behavior

Remote Sensing. 2023, 15, 1477.
https://doi.org/10.3390/ rs15061477

II

Farshad Farahnakian & Fahimeh Farahnakian & Javad Sheikh & Paavo Nevalainen & Jukka Heikkonen
Title: Short and Long Term Vessel Movement Prediction for Maritime Traffic

Springer Nature Switzerland, CRITIS 2023, LNCS 14599, pp. 62–80, 2024.

III

Pontus Svenson & Anders Holst & Anders Wallberg & Paavo Nevalainen & Farshad Farahnakian & Alfonso Alamo & et al.
Title: AI-ARC Baltic Demo: Detecting Illegal Activities at Sea

2024 27th International Conference on Information Fusion (FUSION), Venice, Italy, 2024, pp. 1-8, doi: 10.23919/FUSION59988.2024.10706447.

IV

Farshad Farahnakian & Paavo Nevalainen & Fahimeh Farahnakian & Tanja Vähämäki & Jukka Heikkonen
Title: Maritime Vessel Movement Prediction: A Temporal Convolutional Network Model with Optimal Look-back Window Size Determination

Multimodal Transportation, Volume 4, Issue 1, 2025, ISSN 2772-5863,
https://doi.org/10.1016/j.multra.2025.100191

V