Detecting Anomalous Vessel Trajectories: A Collaborative Clustering-Based Approach*

Claudio Vasconcelos Ribeiro[1], Marcos Bedo[1], Ronaldo Mello[2]
Aline Paes[1], Daniel de Oliveira[1]

[1]Universidade Federal Fluminense (UFF), Niterói, RJ, Brazil

[2]Universidade Federal de Santa Catarina (UFSC), SC, Brazil

{marcosbedo, alinepaes, danielcmo}@ic.uff.br, r.mello@ufsc.br

Abstract. *Thousands of vessels operating worldwide may present issues related to anomalous trajectories, often characterized by unexpected changes in course, speed, or navigation through restricted areas. Manually detecting these anomalies is impractical, evidencing the need for automated support for maritime surveillance agents. Although Automatic Identification Systems (AIS) enhance situational awareness, they are insufficient for detecting trajectory anomalies. Existing approaches already use AIS data to detect anomalous vessel trajectories, but they do not consider contextual variables (e.g., oil spills) or the domain expertise of maritime surveillance agents. This paper introduces* `INSTRUCTOR`*, a solution that combines AIS data with clustering algorithms to group vessel trajectories based on similar directions and movement patterns. This clustering output assists experts in collaboratively identifying anomalous trajectories, a process validated through the collective judgment of multiple domain specialists. Experiments conducted in partnership with the Brazilian Navy corroborate the effectiveness of* `INSTRUCTOR` *in detecting anomalous maritime trajectories.*

1. Introduction

Maritime transportation plays an important role in the global economy, accounting for approximately 90% of international trade by volume [UNCTAD 2019]. It is the dominant transportation type regarding trade volume and monetary value [Riveiro et al. 2018]. In 2019, the global fleet had nearly 52,000 commercial vessels [UNCTAD 2019]. That year, 3,174 maritime incidents were reported, including 95 severe casualties, 53 fatalities, and over 900 injuries [EMSA 2019]. Although incident rates have decreased in recent years, further reduction presents significant operational and safety challenges [Weng et al. 2018].

Maritime incidents encompass various types, with collisions, sinkings, and groundings being the most common [Ribeiro et al. 2023]. Other events, including illegal activities such as piracy and terrorism, also pose risks to maritime operations [Zor and Kittler 2017]. A common feature among these incidents is anomalous vessel behavior, typically characterized by deviations from standard routes. Such anomalies may involve unexpected changes in course or speed, navigation through restricted areas, or other irregular patterns

*The authors would like to thank the financial support from the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001, the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) – 311898/2021-1, and the Fundação de Amparo à Pesquisa do Estado do Rio de Janeiro (FAPERJ) – Grants E-26/204.238/2024 and E-26/204.544/2024.

[Riveiro et al. 2018]. For example, Figure 1 presents real vessels' positions data obtained from Automatic Identification System (AIS) transmissions. AIS operates through transponders installed on vessels, enabling the automatic exchange of navigational data between ships and stations. Each vessel transmits and receives information such as position, speed, and heading. AIS data constitutes the primary source for maritime traffic monitoring and situational awareness [Radon et al. 2015]. Figure 1(a) shows AIS data along the Brazilian coast, where a distinct marker represents each vessel. Route patterns emerge as dense lines formed by repeated trajectories, and deviations from these patterns can indicate anomalous behavior. Figure 1(b) depicts irregular vessel movements caused by limited visibility due to heavy fog and high speed as a vessel approached the Gulf of Mexico, USA.

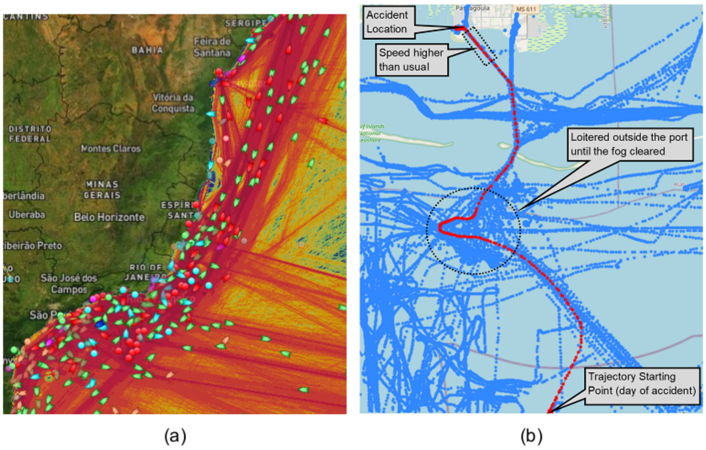

Figure 1. (a) Density map (over one year AIS plots) from part of the Brazilian Coast, (b) Accident tracking was reconstructed from AIS data in the Gulf of Mexico.

AIS data has become valuable due to the expansion of terrestrial and satellite-based AIS infrastructure [Pallotta 2014], enabling access to large volumes of maritime traffic information unavailable a decade ago. However, AIS data alone is insufficient for reliably detecting trajectory anomalies. Given the scale of maritime traffic, manual identification of anomalous behavior is tedious and error-prone. Therefore, automated anomaly detection methods are essential. These systems can support maritime surveillance by reducing operator workload and minimizing the risk of oversight [Kowalska and Peel 2012]. One approach for identifying anomalous vessel trajectories is to train Machine Learning (ML) models using AIS data. Clustering algorithms, *e.g.*, DBSCAN [Dobrkovic et al. 2015, Wang et al. 2014], are commonly applied for this purpose. For instance, [Dobrkovic et al. 2015] apply DBSCAN using only positional data (latitude and longitude) from AIS to detect trajectory anomalies, whereas [Wang et al. 2014] extend this approach by incorporating new features such as Course over Ground (COG) and Speed over Ground (SOG).

Although these methods represent a step forward, their outputs alone may be insufficient for reliable decision-making in maritime surveillance. Additional contextual variables such as wind speed, dynamic navigational restrictions (*e.g.*, oil spills), or temporary maritime regulations are often not included in clustering models but can influence vessel trajectories. Also, each clustering algorithm presents advantages and limitations (*e.g.*, the number of clusters k in K-Means is a sensitive parameter), which may influence the decision. Furthermore, surveillance operators use visual analysis to interpret route clusters and identify anomalies based on domain expertise. Existing solutions do not provide approaches for visual comparison between a target trajectory and clustered routes in a collaborative way, since, due to the complexity of maritime behavior analysis, evaluation based on a single operator may not be recommended, underscoring the need for collaborative, expert-driven decision support.

In this paper, we propose a novel approach named `INSTRUCTOR` (IdeNtification of anomalous veSsel TRajectory Using ClusTers Of tRajectories). `INSTRUCTOR` integrates AIS and contextual data preprocessing with trajectory clustering using various clustering algorithms, including DBSCAN, K-Means, and Spectral Clustering. After clustering, the trajectory under analysis is visually compared against identified route patterns to support expert-driven anomaly detection. The system enables multiple experts to independently evaluate and categorize a trajectory as anomalous. Final decisions are derived by aggregating individual evaluations, weighted by each expert's level of domain expertise. Experiments conducted with the Brazilian Navy showed that `INSTRUCTOR` enhanced experts' ability to detect trajectory deviations more effectively.

The remainder of this paper is organized into four sections. Section 2 reviews related work and provides background information. Section 3 describes the proposed approach. Section 4 presents the evaluation results. Finally, Section 5 concludes the paper and outlines directions for future work.

2. Background and Related Work

2.1. AIS data in a nutshell

AIS data are a key source for analyzing anomalous vessel trajectories. AIS transponders are mandatory for vessels over 300 Gross Tonnes (GT) on international voyages, cargo vessels over 500 GT, and all passenger ships [Zhen et al. 2017]. Messages are encoded using standard protocols (NMEA AIVDM/AIVDO) [Silveira et al. 2013] and transmitted periodically over two VHF channels (161.975 and 162.025 MHz). These transmissions are publicly accessible, enabling wide-scale data collection and analysis [Pallotta 2014].

Such data include both static and dynamic information about vessels [Wang et al. 2014]. There are 23 AIS message types [Riveiro and Falkman 2009]. Static messages are typically broadcast every six minutes, while dynamic messages are transmitted at intervals ranging from two seconds to three minutes, proportionally to vessel speed [Soleimani et al. 2015]. Dynamic position-related data include: timestamp, Maritime Mobile Service Identity (MMSI), navigational status, Rate of Turn (RoT), Speed over Ground (SOG), position accuracy, latitude, longitude, Course over Ground (COG), and heading. The MMSI is the unique vessel identifier commonly used as the primary key in AIS datasets [Zhen et al. 2017]. Static information includes MMSI, IMO number, vessel type, dimensions, and name. These static attributes are essential for vessel identification.

Despite its importance, the usage of AIS data presents limitations. Since the vessel's crew manually informs some attributes, AIS messages are susceptible to human error and even intentional falsification [Osekowska et al. 2013]. For example, destination ports and Estimated Time of Arrival (ETA) are often incorrectly or inconsistently configured, leading to unreliable information [Nishizaki et al. 2018]. Additionally, the positional accuracy of most AIS transponders is limited to approximately 10 meters [Osekowska et al. 2013]. Vessel positions are typically derived from onboard GPS receivers integrated with the AIS system.

A key challenge in developing methods for detecting anomalous vessel trajectories is the efficient processing of large-scale AIS data. AIS systems can produce over 70 GB of data annually from thousands of vessels. This volume is due to the slow movement of ships combined with high-frequency AIS transmissions, producing data redundancy [Fu et al. 2017]. To address this, data reduction and modeling techniques are often employed. One common strategy involves discretization using *bins*. For example, vessel course can be categorized into cardinal and intercardinal directions (N, NE, E, SE, S, SW, W, NW), while time can be segmented into *bins* such as morning, afternoon, evening, and night [Osekowska et al. 2013]. Similar discretization can be applied to Course over Ground (COG), Speed over Ground (SOG), and spatial data using latitude/longitude grids. Additionally, depending on application requirements, precision reduction—such as truncating decimal digits in speed values—can be adopted without significant loss of relevant information [Nguyen et al. 2018].

2.2. Related Work

This subsection presents a simplified systematic mapping on the topic of Maritime Traffic Anomaly Detection Based on AIS Data and ML techniques, aiming to categorize and analyze relevant publications in the field. The snowballing technique was employed as the primary search strategy, following the procedures for backward and forward iterations described by [Wohlin 2014]. As the initial Seed Set, three representative papers were selected based on relevance and coverage: [Riveiro et al. 2018, Sidibé and Shu 2017, Smith et al. 2012].

The following criteria guided the paper selection process: *(i)* use of specific keywords: ML-based maritime traffic anomalous behavior, maritime surveillance systems, vessel movement patterns, and Automatic Identification System (AIS); *(ii)* exclusion of studies with a primary focus on military or defense applications; and *(iii)* exclusion of theses and dissertations, as recommended by [Wohlin 2014]. To minimize publisher bias, the primary search engine used was Google Scholar, with results from 10 major digital libraries: IEEE Xplore, Elsevier, Springer, ACM Digital Library, Cambridge University Press, MDPI, Wiley Online Library, Taylor & Francis, IOS Press, and DiVA.

In total, 18 papers were selected based on the defined inclusion criteria and full-text analysis. A common approach to maritime anomaly detection involves learning a model of normal vessel behavior from historical movement data. New vessel movements are then submitted to the trained model, and any significant deviation is interpreted as anomalous. In [Zhao and Shi 2019], the authors present a trajectory prediction model based on Long Short-Term Memory (LSTM) neural networks, trained with historical AIS data preprocessed by DBSCAN. The DBSCAN parameters are estimated using inverse Gaussian fitting, mitigating the sensitivity to input settings. Vessel trajectories are divided into sliding input-output sequences, and deviation between the predicted and actual positions, considering speed, course, and route, is used to flag anomalies.

[Wang et al. 2014] extended DBSCAN by introducing DBSCAN_SD, which incorporates COG and SOG into the clustering process. A labeled dataset, built with expert input, is used to train a Parallel Meta-Learning (PML) algorithm on Apache Hadoop, showing scalability and improved accuracy with complex base models and larger datasets. [Mantecón et al. 2019] propose a supervised deep learning approach using Convolutional Neural Networks (CNNs) to infer navigation status (stationary, cruising, fishing) from AIS data. The method assumes that navigation status and ship type are available in AIS messages, which is not always true. [Nguyen et al. 2018] proposed a VRNN-based architecture that encodes noisy, irregular AIS data into embeddings. These embeddings support trajectory reconstruction, anomaly detection, and vessel type classification without requiring clustering. [Venskus et al. 2017] developed a self-learning classifier combining Self-Organizing Maps (SOMs) with virtual pheromones. The method is extended in [Venskus et al. 2019] to improve retraining efficiency for real-time data via batching strategies. A cloud-based neural network system is proposed for short- and long-term trajectory prediction by [Zissis et al. 2015]. The model supports port scheduling and route planning and emphasizes the difficulty of capturing abrupt speed or course changes, often associated with port maneuvers. [Nishizaki et al. 2018] applied a Support Vector Machine to forecast vessel courses at traffic exits in Tokyo Bay.

[Mazzarella et al. 2014] identified fishing areas using AIS data and DBSCAN clustering of stop and move segments, based on [Spaccapietra et al. 2008]. Similarly, [Yan et al. 2016] detected Areas of Interest (AOIs) and major traffic routes by clustering speed and directional behavior patterns. [Dobrkovic et al. 2015] evaluated clustering methods for waypoint extraction, which is extended in [Dobrkovic et al. 2018] to construct directed graphs for anomaly detection. [Zhang et al. 2018] applied ACO to identify turning points in compressed trajectories and infer optimal routes.

[Liu et al. 2015] proposed a point-based anomaly scoring model, where deviations in speed and direction of vessels result in scores interpretable by humans. [Laxhammar and Falkman 2010] applies conformal prediction theory for online anomaly detection, treating full trajectories as non-i.i.d. instances and evaluating anomalies via simulated paths. [Fu et al. 2017] implemented a modular detection system combining data preprocessing, DTW-based similarity measurement, and outlier scoring, with time-aware trajectory imputation and reference trajectory construction for anomaly evaluation. [Radon et al. 2015] incorporated contextual factors, such as weather conditions, into a clustering-based framework to minimize false alarms. [Lei 2016] ranked vessel trajectories using spatial, sequential, and behavioral outlier metrics. Similarly, the TREAD system [Pallotta et al. 2013a, Pallotta et al. 2013b] extracts traffic patterns via a point-based approach and detects anomalies through sliding window analysis. This is extended in [Pallotta 2014] with Ornstein-Uhlenbeck processes to predict vessel motions. [Zhen et al. 2017] proposed a Naïve Bayes classifier for anomaly detection in coastal traffic, utilizing hierarchical and K-Medoids clustered trajectories. Although these approaches represent a step forward, they do not allow for collaborative trajectory analysis to be solely based on trained data, which may be insufficient, as not all variables can be accounted for during training.

To the best of the authors' knowledge, no existing method enables a collaborative and visual comparison between a trajectory under analysis and clusters produced by multiple clustering algorithms for maritime surveillance operators, who still rely on the visual analysis of trajectory clusters and anomalies to support decision-making. The approach proposed in this paper aims to address this gap by enabling maritime traffic operators to visually evaluate trajectory clusters derived from historical AIS data and compare new trajectories with those generated by multiple ML algorithms. Armed with this visual comparison and contextual information (*e.g.*, sea conditions), operators can assess whether a trajectory is anomalous.

3. Proposed Approach: `INSTRUCTOR`

The architecture of `INSTRUCTOR` (Figure 2) consists of six components: *(i)* AIS Processor, *(ii)* AIS Database, *(iii)* Cluster Generator, *(iv)* External Data Extractor, *(v)* Structured External Database, and *(vi)* Web Portal. `INSTRUCTOR` receives as input (step 1) a set of files containing AIS data, which record ship positions over time for multiple vessels. The *AIS*

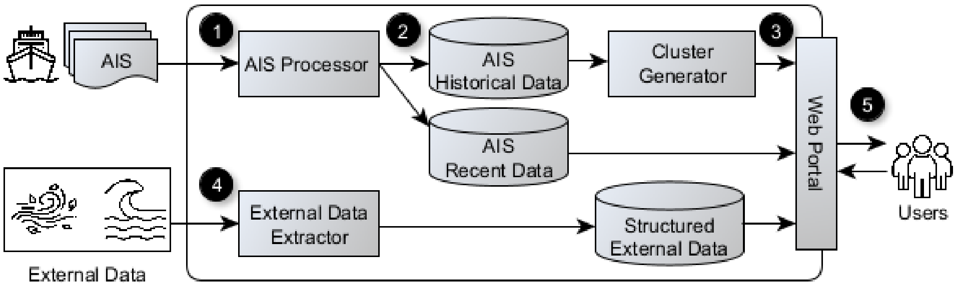

Figure 2. Architecture of INSTRUCTOR.

Processor handles these files and enables users to select historical or recent data to store in the *AIS Database* (step 2). This database includes timestamp, latitude, longitude, speed, and heading attributes.

The *Cluster Generator* queries the AIS Database to perform clustering (step 3), using a catalog of standard algorithms (*e.g.*, DBSCAN, K-Means, Spectral Clustering, *etc.*). Unlike the current mainstream, `INSTRUCTOR` applies a preprocessing step before executing clustering with each algorithm in the catalog. In existing vessel anomalous trajectory detection approaches, clustering algorithms are typically executed solely based on the vessel's location points, which collectively form the trajectories. These algorithms group the points into non-contiguous areas or distribute them spatially based on distances to identify clusters.

However, using solely latitude and longitude, the generated clusters do not always represent navigation flows, as presented in Figure 1. Although high-density regions can be visually identified, they primarily reflect vessel traffic concentration rather than meaningful maritime flows. Therefore, uncovering the semantics behind the data points is important so that the identified flows correspond to an actual trajectory. To address this, the *Cluster Generator* applies the aforementioned preprocessing stage to categorize historical trajectories with similar directions and transit locations. The clustering algorithms then process these categorized trajectories. This approach supports the definition of accurate vessel trajectory patterns, easing collaborative analysis by experts.

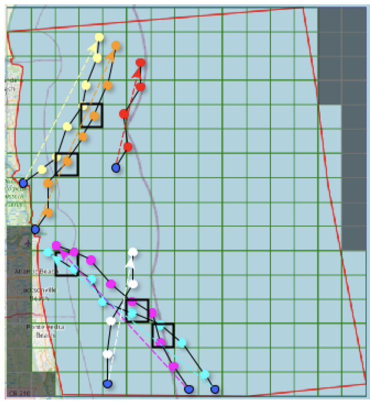

Figure 3. Data preprocessing in *Cluster Generator* component.

Initially, `INSTRUCTOR` divides the Area of Interest (AOI) into cells. This spatial discretization provides adequate granularity for preprocessing and clustering historical AIS navigational data and computing trajectory similarity. `INSTRUCTOR` defines s as the grid cell size, and Δx and Δy as the distances along the x and y axes. The total number of grid cells, denoted by X_t and Y_t, is given by: $X_t = \left\lceil \frac{\Delta x}{s} \right\rceil + 1$ and $Y_t = \left\lceil \frac{\Delta y}{s} \right\rceil + 1$. The grid's origin, Grid$_{RP}$, corresponds to the coordinate with the lowest latitude and longitude. All distance calculations between two points in `INSTRUCTOR` use the Haversine distance. The Mercator projection is used for coordinate representation and distance calculations, and each pair (lat/long) in a vessel trajectory is mapped to a grid cell.

Once the grid is defined, the AIS data can be effectively preprocessed. Figure 3 shows the preprocessing stage, where the results serve as input for the clustering algorithms. This procedure consists of the following steps: *(i)* the start and end points are identified for each vessel trajectory in the historical AIS database and mapped to a grid cell (the blue dots in Figure 3 represent the starting points of the trajectories), *(ii)* each trajectory creates a vector based on the trajectory points, indicating its direction toward the destination (dashed arrows in Figure 3), *(iii)* trajectories are then grouped based on the angle of their vectors, *i.e.*, if a trajectory's direction falls within the same angular range (sector), it is assigned to that sector (in Figure 3, trajectories with yellow, orange, red, and white dots fall within the 0–30 degree sector, while those with pink and cyan dots fall within the 300–330 degree sector).

The preprocessed data is used to identify clusters, and the resulting clusters, representing trajectory patterns, are saved for future inference. In addition to clusters generated from AIS data, `INSTRUCTOR` also supports the integration of external sources, such as wind and sea condition data. The *External Data Extractor* collects data from a predefined catalog of external sources and periodically stores it in the *Structured External Database* (step 4). Users interact with `INSTRUCTOR` through the *Web Portal* (step 5), which displays clustered trajectories and associated external data to support the analysis. Considering the clustered trajectories (*i.e.*, patterns) displayed over a map (Figure 4), and additional data such as Rosewind information, the user can classify the selected trajectory as either normal or anomalous. Based on all classifications submitted by users for a given trajectory, `INSTRUCTOR` provides a classification decision using majority voting, weighted by the level of confidence of each user.

The following workflow defines the usage of `INSTRUCTOR` web portal, presented in Figure 4. After selecting the AOI and setting the grid size, the user loads historical data from AIS files. The selected file is loaded (area ❶), and the user chooses a clustering algorithm to execute (area ❷). Available algorithms include: Agglomerative Clustering, DBSCAN, HDBSCAN, K-Means, Gaussian Mixture, Spectral Clustering, and Ensembles. After performing the clustering, the clustered trajectories are displayed using distinct colors (area ❸), where black is used for points not assigned to any cluster. These clusters can be saved for later use. To classify a new trajectory, the user selects a file containing AIS data for vessels in transit (area ❹) and compares the new trajectory against the identified clusters. The similarity between the new trajectory and each cluster is computed and displayed in a panel (area ❸). If the file includes data from multiple vessels, the user can choose the vessel of interest to compare. The user can inspect the vessel's heading and speed at each location point (area ❻) and compare this with wind data, shown in a wind rose (area ❼). A play button (area ❹) enables playback of the entire trajectory.

After the trajectory analysis, the user can classify the trajectory and export the results to a database or a CSV file (area ❺). In the collaborative analysis setting, once users submit their classifications, a final decision can be "*normal*" or "*anomalous*" and is made through majority voting, weighted by user expertise. For trajectories classified as "*anomalous*", `INSTRUCTOR` supports labeling based on the following classes: *(i)* unusual speed, *(ii)* U-turn, *(iii)* prohibited area, *(iv)* uncertain destination, *(v)* loitering, *(vi)* drifting, *(vii)* outside shipping lane, *(viii)* heading into danger, *(ix)* regulatory infraction, and *(x)* threat to infrastructure. `INSTRUCTOR` was developed using JavaScript and Flask, using scikit-learn as an ML library. The source code of `INSTRUCTOR` is available in GitHub: `https://github.com/UFFeScience/instructor`.

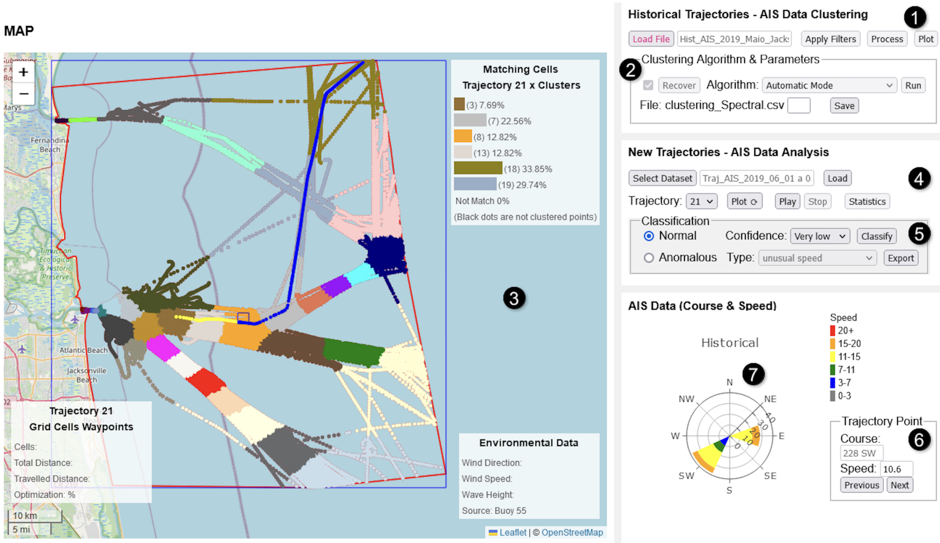

Figure 4. INSTRUCTOR Web portal.

4. Experimental Evaluation

This section presents the evaluation of INSTRUCTOR. The experiments were conducted with three maritime surveillance experts from the Brazilian Navy and two novice users. First, we discuss the impact of cell size on identifying anomalous trajectories. Next, we analyze the effect of data pre-processing. Finally, we present an evaluation using a labeled dataset, comparing the labeled trajectory classification with the evaluations made by experts. For the experiments, an AOI has been delimited in by an area of 70km × 80km of ocean near Jacksonville on the east coast of the USA (AOI presented in area 3 in Figure 4), which was chosen based on suggestions from Brazilian Navy experts. In the experiments, only cargo ships and tanks with a length of more than 99 meters were selected. Other ship classes were not considered, avoiding short and random trajectories of tugboats, sailboats, and speedboats.

The AIS data was the primary source of positional and dynamic information for vessel trajectories in the experiments. The AIS dataset is derived from the Nationwide Automatic Identification System (NAIS) maintained by the U.S. Coast Guard. The data files, organized by day, contain vessel position updates recorded at one to two-minute intervals and are provided in CSV format. Based on this update frequency, any interval greater than five minutes without a position update was considered the start of a new trajectory. Additionally, if a vessel's speed dropped below 1.5 knots, it was considered a stop condition, and a new trajectory was considered. Each data file is approximately 700 MB and includes 17 attributes related to vessel static and dynamic data, timestamps, and transceiver class type.

For the experiments presented in this section, we used 14 days of historical AIS data from May 2019 to identify clusters and 12 trajectories (Figure 5) from June 2019 to be classified. Another type of data not applied in the clustering process during the experiments is related to weather and oceanic conditions. This data supports maritime experts in analyzing vessel trajectories, as certain deviations may be explained by adverse environmental factors. In the experiments, we used data from oceanic buoys[1] provided by the National Data Buoy Center. The buoy data is available in daily files, updated hourly, and provided in plain text.

[1]`https://www.ndbc.noaa.gov/`

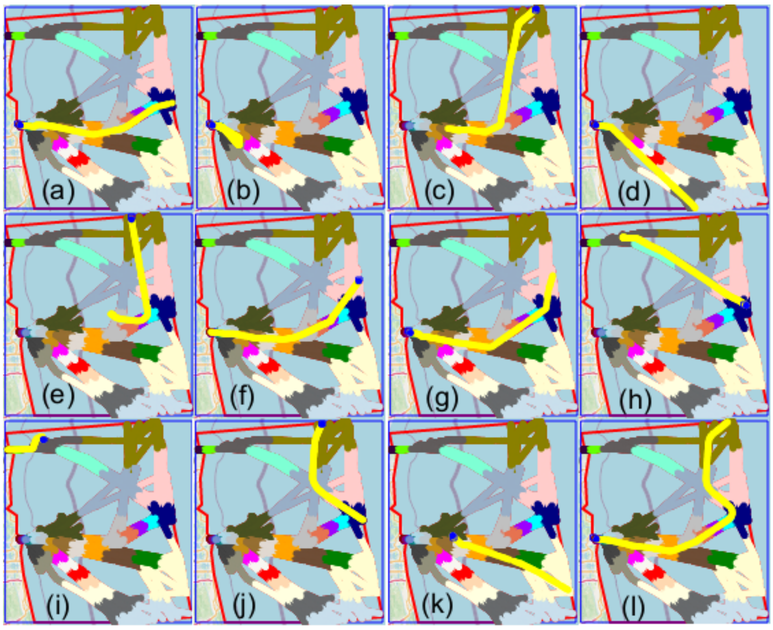

Figure 5. Visualization of the selected trajectories for analyses. The letters identify the IDs of the 12 trajectories highlighted with a yellow line. *(a)* 13, *(b)* 20, *(c)* 21, *(d)* 31, *(e)* 37, *(f)* 43, *(g)* 58, *(h)* 69, *(i)* 70, *(j)* 74, *(k)* 82, and *(l)* 88.

The first experiment evaluated the impact of cell size on the clustering results. Tables 1 and 2 show the results of the Spectral Clustering algorithm (due to space restrictions, the results of other clustering algorithms are not presented) applied with different grid cell sizes. The evaluation metrics include: Max Cluster Similarity Value (MS), which indicates the percentage of similarity between the new trajectory and the clustered trajectories; Clusters Not Matching (NM), which represents the percentage of the trajectory that does not match any cluster; and the number of clusters that match at least a portion of the trajectory.

Based on these results, it can be observed that large cell sizes (*e.g.*, 5000m) result in 0% Not Matching (NM%), meaning every point of the new trajectory aligns with points in the clustered trajectories. However, this often leads to overly generalized patterns, as many clusters match and high similarity scores are observed. On the other hand, a small cell size (*e.g.*, 100m) yields fewer matching clusters and higher NM%, which may hinder effective pattern detection. Although NM% appears useful, its consistently low values across most trajectories make it less reliable for identifying patterns, particularly anomalies. Therefore, a 2km cell size was selected for all experiments, balancing visual clarity in grid construction and clustering performance within the AOI.

The second experiment evaluated the impact of data preprocessing in `INSTRUCTOR`. For clustering without preprocessing, a conventional method was applied using only vessel location points (latitude and longitude). Due to space limitations, Spectral Clustering was selected as a representative algorithm. When running `INSTRUCTOR` with preprocessing, the area was divided into 12 sectors, each covering an angle of 30 degrees. To ensure a fair comparison, the Spectral Clustering configuration with preprocessing was set to produce 3 clusters per sector, resulting in 36 clusters from 12 independent executions. Figure 6 shows the clustering results from both approaches.

Table 1. Spectral Clustering results with Grid Cell values of 100m/200m/300m/500m.

Traj. ID	100 / 554840 MS(%)	100 / 554840 NM(%)	100 / 554840 NrC	200 / 138904 MS(%)	200 / 138904 NM(%)	200 / 138904 NrC	300 / 61901 MS(%)	300 / 61901 NM(%)	300 / 61901 NrC	500 / 22308 MS(%)	500 / 22308 NM(%)	500 / 22308 NrC
13	31.0	41.0	05.0	75.0	11.0	16.0	69.0	07.0	18.0	100.0	00.0	18.0
20	56.8	15.9	12.0	87.9	01.5	13.0	96.2	00.0	13.0	99.2	00.0	13.0
21	08.2	84.6	10.0	18.5	66.7	14.0	21.5	55.9	14.0	34.9	34.9	18.0
31	17.4	77.2	11.0	26.1	67.4	13.0	30.4	59.8	14.0	33.7	50.0	14.0
37	06.0	92.0	03.0	10.0	87.0	05.0	10.0	83.0	08.0	13.0	80.0	08.0
43	26.0	52.1	16.0	45.6	28.4	19.0	52.7	18.3	21.0	74.0	13.0	21.0
58	18.8	50.4	17.0	53.0	18.0	16.0	45.3	07.7	17.0	86.3	01.7	19.0
69	03.5	92.2	06.0	10.6	80.8	09.0	23.4	56.7	11.0	37.6	35.5	11.0
70	17.0	73.6	06.0	30.2	62.3	08.0	34.0	60.4	08.0	39.6	58.5	08.0
74	01.8	95.4	04.0	04.6	86.2	07.0	12.8	61.5	10.0	27.5	42.2	10.0
82	11.1	76.2	07.0	27.0	52.4	08.0	27.0	50.8	09.0	42.9	20.6	10.0
88	16.8	58.9	16.0	38.4	37.3	22.0	38.9	27.0	27.0	58.9	15.7	27.0

Cell size / number of cells

MS ->Max cluster similarity value(%) NM ->Clusters Not Matching value(%) NrC ->Number of clusters with matching trajectory

Table 2. Spectral Clustering results with Grid Cell values of 1Km/2Km/5Km.

Traj. ID	1000 / 5616 MS(%)	1000 / 5616 NM(%)	1000 / 5616 NrC	2000 / 1404 MS(%)	2000 / 1404 NM(%)	2000 / 1404 NrC	5000 / 240 MS(%)	5000 / 240 NM(%)	5000 / 240 NrC
13	99.0	01.0	20.0	79.0	00.0	24.0	82.0	00.0	25.0
20	100.0	00.0	13.0	100.0	00.0	13.0	100.0	00.0	13.0
21	66.7	12.8	19.0	89.7	00.0	20.0	91.3	00.0	22.0
31	43.5	31.5	15.0	54.3	01.1	15.0	67.4	00.0	16.0
37	25.0	59.0	10.0	44.0	38.0	10.0	76.0	09.0	14.0
43	81.1	10.1	23.0	69.2	05.3	25.0	91.1	00.0	25.0
58	94.9	00.0	24.0	68.4	00.0	25.0	73.5	00.0	27.0
69	53.9	12.8	12.0	71.6	04.3	12.0	89.4	00.0	16.0
70	47.2	52.8	08.0	47.2	52.8	08.0	100.0	00.0	08.0
74	48.6	33.9	11.0	58.7	29.4	13.0	94.5	00.0	15.0
82	63.5	06.3	11.0	69.8	01.6	12.0	90.5	00.0	13.0
88	61.6	07.0	28.0	67.3	00.0	30.0	77.3	00.0	31.0

Cell size / number of cells

MS ->Max cluster similarity value(%) NM ->Clusters Not Matching value(%)
NrC ->Number of clusters with matching trajectory

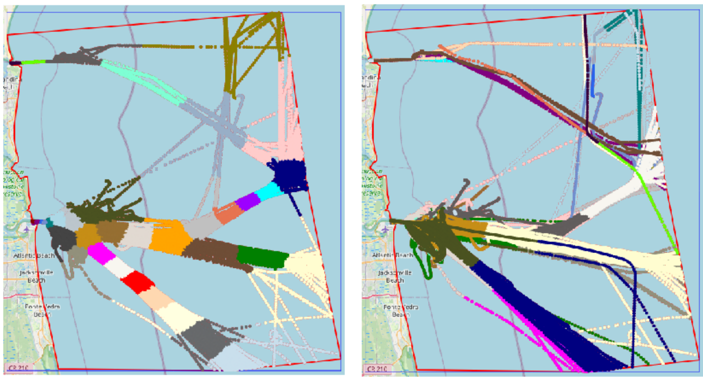

(a) (b)

Figure 6. Trajectory clusterization without and with preprocessing.

The 12 trajectories illustrated in Figure 5 were selected for evaluation of the impact of preprocessing. The selection criteria included longer trajectory lengths and diversity in shape and structure. These experiments were conducted by Brazilian Navy officers, and the results from the procedures without and with preprocessing are summarized in Table 3. The results indicate that the proposed preprocessing approach improves the clustering process, as shown by the higher similarity percentages between the clusters and the analyzed trajectories. For instance, in the case of Trajectory 82, the preprocessing approach yields higher similarity values than the method without preprocessing. The experts could make informed decisions regarding trajectory classification by comparing the similarity metrics with the complementary information from the Rosewind chart and the interactive map. Table 3 compiled and summarized their individual classifications. Based on the expert evaluations presented in Table 3, the trajectories with the highest number of votes for anomalous classification were 37, 70, and 74. These trajectories also presented higher percentages of Not Matching.

Table 3. Clustering results with and without Pre-processing in the dataset.

	Without Pre-processing				With Pre-processing			
Traj.	Max cluster	Clusters	Classification		Max cluster	Clusters	Classification	
ID	similarity	Not Matching	voting	Type*	similarity	Not Matching	voting	Type*
13	21.0	00.0	50.0	01.0	**79.0**	00.0	00.0	-
20	56.1	00.0	50.0	02.0	**100.0**	00.0	50.0	2; 3
21	33.9	00.0	25.0	03.0	**89.8**	00.0	00.0	-
31	23.9	01.1	50.0	01; 04	**54.4**	01.1	25.0	01
37	20.0	38.0	75.0	02; 05	**44.0**	38.0	75.0	01; 02; 05
43	16.6	05.3	50.0	01; 06	**69.2**	05.3	50.0	01; 08
58	14.5	00.0	25.0	07.0	**68.4**	00.0	00.0	-
69	29.1	04.3	25.0	08.0	**71.6**	04.3	25.0	01
70	24.5	52.8	75.0	01; 05; 09	**47.2**	52.8	75.0	01; 02; 05
74	26.6	29.4	75.0	01; 05; 10	**58.7**	29.4	75.0	01; 02; 03
82	38.1	01.6	50.0	01; 02	**69.9**	01.6	50.0	01
88	22.2	00.0	25.0	01.0	**67.0**	00.0	25.0	01

* Type: 01- unusual speed; 02 - unusual course; 03- U-turn; 04 - Prohibit Area; 05 - Uncertain Destiny; 06 - Loitering; 07 - Drifting; 08 - Outside shipping lane; 09 - Heading into danger; 10 - Regulatory infraction; and 11 - Threat to infrastructure.

We also generated simulated data for expert evaluation, aiming to conduct experiments with labeled trajectories. A total of 10 trajectories were created within the same AOI, forming a balanced dataset for classification. The dataset was organized as follows: 5 normal trajectories; 1 anomalous trajectory generated by altering the speed values of a normal trajectory; and 4 trajectories constructed from real location points of sailing and fishing vessels. These last trajectories were considered anomalous due to their distinct navigation patterns and operational purposes compared to cargo and tanker vessels. We used the F-score as a general metric for evaluating the binary classification problem—classifying trajectories as either normal or anomalous. Table 4 shows the classification results of five Brazilian Navy officers on the prepared set of 10 labeled trajectories. Among them, three officers (Expert 1, Expert 2, and Expert 3) have expertise in the maritime domain.

On analyzing the results from Table 4, it can be concluded that `INSTRUCTOR` was effective in supporting the classification of anomalous vessel trajectories in 7 out of 10 cases

Table 4. Results for simulated data.

Traj. ID	Real Label	Expert 1	Expert 2	Expert 3	Novice 1	Novice 2
1	normal	normal	normal	normal	normal	normal
2	anomalous	normal	anomalous	anomalous	normal	normal
3	normal	normal	normal	normal	normal	normal
4	normal	normal	normal	normal	normal	normal
5	normal	normal	normal	normal	normal	normal
6	anomalous	anomalous	anomalous	anomalous	anomalous	anomalous
7	anomalous	anomalous	anomalous	anomalous	anomalous	anomalous
8	normal	normal	normal	normal	normal	normal
9	anomalous	anomalous	anomalous	anomalous	anomalous	normal
10	anomalous	anomalous	anomalous	anomalous	anomalous	normal

(the three cases with discrepancies are highlighted in red). The system's performance is further demonstrated by the average F1-score of 0.89, shown in Table 5, indicating its effectiveness in identifying anomalous trajectories across all users, whether expert or not. Table 5 includes the values for True Positive (TP), True Negative (TN), False Positive (FP), False Negative (FN), Accuracy, Precision, Recall, and F1-score. It is important to note that when $TP = 0$ and $FP = 0$, Precision is undefined. Similarly, when $TP = 0$ and $FN = 0$, Recall is undefined. These metrics are assumed to be 1.00 in these cases, indicating perfect classification of the negative class.

Table 5. Evaluation metrics for each labeled trajectory.

Traj. ID	True Positive	True Negative	False Positive	False Negative	Accuracy	Precision	Recall	F1-score
1	0	5	0	0	1.00	—	—	—
2	2	0	0	3	0.40	1.00	0.40	0.57
3	0	5	0	0	1.00	—	—	—
4	0	5	0	0	1.00	—	—	—
5	0	5	0	0	1.00	—	—	—
6	5	0	0	0	1.00	1.00	1.00	1.00
7	5	0	0	0	1.00	1.00	1.00	1.00
8	0	5	0	0	1.00	—	—	—
9	4	0	0	1	0.80	1.00	0.80	0.89
10	4	0	0	1	0.80	1.00	0.80	0.89
Total	**20**	**25**	**0**	**5**	**0.90**	**1.00**	**0.80**	**0.89**

Although most trajectories were correctly classified using `INSTRUCTOR`, the evaluation with the labeled dataset highlighted that sudden speed changes may require a more detailed speed analysis throughout the entire trajectory by the expert. For instance, in the case of Trajectory 2 from Table 4, this anomalous trajectory was difficult to detect, and only two out of three experts could identify it. The trajectory appears normal for most of its points. While far from the coast, the speed values remain stable within a certain range. As the vessel approaches the coast, especially near channels or docking areas, a sudden increase in speed, as observed in Trajectory 2 in Table 5, is unusual and may indicate an anomalous trajectory.

5. Conclusion

Identifying anomalous vessel trajectories allows authorities to take timely actions to avoid damage or loss of life. This motivates using automatic detection methods, particularly clustering algorithms, to detect anomalous maritime trajectories. The approach proposed in this paper clusters trajectories based on common direction and location, performing clustering on each group and then combining the results, which has been observed to improve the definition of vessel trajectory patterns. The approach `INSTRUCTOR` was developed to aid in clustering and anomaly detection, providing a platform for exploring clusters generated by multiple algorithms without requiring users to manage infrastructure.

Experiments with `INSTRUCTOR` showed improved anomaly detection by refining the identification of maritime traffic patterns, which helped experts detect anomalous trajectories. Since expert analysis can vary depending on experience, a collaborative process was implemented, where individual classifications are aggregated through majority voting to determine whether a trajectory is normal or anomalous. Due to the lack of a benchmark for maritime anomaly detection, simulated anomalies were created with expert support. These simulations involved adding noise, such as variations in speed and course, to real AIS data. Incorporating weather, oceanic, and real-time streaming data could further enhance clustering tasks. Future work will focus on refining the approach with these developments.

References

Dobrkovic, A. et al. (2015). Using machine learning for unsupervised maritime waypoint discovery from streaming ais data. In *Proc. of i-KNOW '15*, pages 1–8.

Dobrkovic, A. et al. (2018). Maritime pattern extraction and route reconstruction from incomplete ais data. *IJDSA*, 5(2-3):111–136.

EMSA (2019). Annual overview of marine casualties and incidents. Technical report.

Fu, P., Wang, H., Liu, K., Hu, X., and Zhang, H. (2017). Finding abnormal vessel trajectories using feature learning. *IEEE Access*, 5:7898–7909.

Kowalska, K. and Peel, L. (2012). Maritime anomaly detection using gaussian process active learning. In *2012 FUSION*, pages 1164–1171. IEEE.

Laxhammar, R. and Falkman, G. (2010). Conformal prediction for distribution-independent anomaly detection in streaming vessel data. In *Proc. of StreamKDD '10*, pages 47–55.

Lei, P.-R. (2016). A framework for anomaly detection in maritime trajectory behavior. *Knowledge and Information Systems*, 47(1):189–214.

Liu, B. et al. (2015). Ship movement anomaly detection using specialized distance measures. In *2015 FUSION*, pages 1113–1120. IEEE.

Mantecón, T. et al. (2019). Deep learning to enhance maritime situation awareness. In *2019 IRS*, pages 1–8. IEEE.

Mazzarella, F. et al. (2014). Discovering vessel activities at sea using ais data: Mapping of fishing footprints. In *17th FUSION*, pages 1–7. IEEE.

Nguyen, D. et al. (2018). A multi-task deep learning architecture for maritime surveillance using ais data streams. In *2018 IEEE DSAA*, pages 331–340. IEEE.

Nishizaki, C. et al. (2018). Development of navigation support system to predict new course of ship. In *2018 WAC*, pages 1–5. IEEE.

Osekowska, E., Axelsson, S., and Carlsson, B. (2013). Potential fields in maritime anomaly detection. In *Proceedings of the 3rd international conference on models and technologies for intelligent transport systems*. TUD Press.

Pallotta, G., Vespe, M., and Bryan, K. (2013a). Traffic knowledge discovery from ais data. In *Proceedings of the 16th International Conference on Information Fusion*, pages 1996–2003. IEEE.

Pallotta, G., Vespe, M., and Bryan, K. (2013b). Vessel pattern knowledge discovery from ais data: A framework for anomaly detection and route prediction. *Entropy*, 15(6):2218–2245.

Pallotta, G. a. (2014). Context-enhanced vessel prediction based on ornstein-uhlenbeck processes using historical ais traffic patterns: Real-world experimental results. In *17th FUSION*, pages 1–7. IEEE.

Radon, A. N., Wang, K., Glässer, U., Wehn, H., and Westwell-Roper, A. (2015). Contextual verification for false alarm reduction in maritime anomaly detection. In *2015 IEEE International Conference on Big Data (Big Data)*, pages 1123–1133. IEEE.

Ribeiro, C. V., Paes, A., and de Oliveira, D. (2023). Ais-based maritime anomaly traffic detection: A review. *Expert Systems with Applications*, page 120561.

Riveiro, M. et al. (2018). Maritime anomaly detection: A review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(5):e1266.

Riveiro, M. and Falkman, G. (2009). Interactive visualization of normal behavioral models and expert rules for maritime anomaly detection. In *CGIV '09*, pages 459–466. IEEE.

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

Silveira, P. et al. (2013). Use of ais data to characterise marine traffic patterns and ship collision risk off the coast of portugal. *The Journal of Navigation*, 66(6):879–898.

Smith, M. et al. (2012). Online maritime abnormality detection using gaussian processes and extreme value theory. In *2012 ICDM*, pages 645–654. IEEE.

Soleimani, B. H. et al. (2015). Anomaly detection in maritime data based on geometrical analysis of trajectories. In *2015 FUSION*, pages 1100–1105. IEEE.

Spaccapietra, S. et al. (2008). A conceptual view on trajectories. *Data & knowledge engineering*, 65(1):126–146.

UNCTAD (2019). *Handbook of Statistics*. UN Conference on Trade and Develop. Available at https://unctad.org/en/PublicationsLibrary/tdstat44_en.pdf.

Venskus, J. et al. (2017). Integration of a self-organizing map and a virtual pheromone for real-time abnormal movement detection in marine traffic. *Informatica*, 28(2):359–374.

Venskus, J. et al. (2019). Real-time maritime traffic anomaly detection based on sensors and history data embedding. *Sensors*, 19(17):3782.

Wang, X. et al. (2014). Vessel route anomaly detection with hadoop mapreduce. In *2014 IEEE Big Data*, pages 25–30. IEEE.

Weng, J. et al. (2018). Combining zero-inflated negative binomial regression with mlrt techniques: an approach to evaluating shipping accident casualties. *Ocean Engineering*, 166:135–144.

Wohlin, C. (2014). Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proc. of EASE'14*, pages 1–10.

Yan, W. et al. (2016). Vessel movement analysis and pattern discovery using density-based clustering approach. In *2016 IEEE Big Data*, pages 3798–3806. IEEE.

Zhang, S.-k. et al. (2018). Data-driven based automatic maritime routing from massive ais trajectories in the face of disparity. *Ocean Engineering*, 155:240–250.

Zhao, L. and Shi, G. (2019). Maritime anomaly detection using density-based clustering and recurrent neural network. *The Journal of Navigation*, 72(4):894–916.

Zhen, R., Jin, Y., Hu, Q., Shao, Z., and Nikitakos, N. (2017). Maritime anomaly detection within coastal waters based on vessel trajectory clustering and naïve bayes classifier. *The Journal of Navigation*, 70(3):648–670.

Zissis, D., Xidias, E. K., and Lekkas, D. (2015). A cloud based architecture capable of perceiving and predicting multiple vessel behaviour. *App. Soft Computing*, 35:652–661.

Zor, C. and Kittler, J. (2017). Maritime anomaly detection in ferry tracks. In *Int. Conf. on Acoustics, Speech and Signal Processing*, pages 2647–2651. IEEE.