3.2. Maritime risk assessment systems and methods

3.2.1. Formal Safety Assessment

In the maritime context, there is a rational and systematic risk-based approach for safety assessment—Formal Safety Assessment (FSA) (Berle, Asbjørnslett, & Rice, 2011; Trucco, Cagno, Ruggeri, & Grande, 2008). FSA was developed by the International Maritime Organization, which is the basic international institution responsible for developing and maintaining a comprehensive regulatory framework for shipping, and thus for providing maritime security and safety.

FSA can be applied to specific maritime safety issues in order to identify cost-effective risk reduction options. The FSA process consists of five steps (Berle et al., 2011; Ellis et al., 2008):

(1) hazard identification: identification of all hazards related to the activity / ship;
(2) risk assessment: building a risk model and determining probabilities and consequences for all branches of the risk model;
(3) risk control options: identification of measures to control and reduce the identified risks;
(4) cost benefit assessment: determining cost effectiveness of each risk mitigation option and preparing a ranking for them;
(5) recommendations for decision making: deciding and making a plan of future activities, based on the results of previous steps.

FSA is commonly seen as the premier scientific method for maritime risk analysis and for formulating maritime regulatory policy (Goerlandt & Montewka, 2015). Therefore it was selected as a foundation for the risk and reliability assessment method that will be presented in Chapter 7.

3.2.2. Maritime risk assessment approaches

In the literature, there are many different analysis techniques and models that have been developed to aid in conducting risk assessments in the maritime domain and which are dedicated to the different steps of FSA.

With regard to the first step of FSA—identification of threats and risk variables—the commonly used methods are: literature review, brainstorming, methods for analysis of possible threats, and unwanted events (e.g., Hazard Identification Study, Hazard and Operability Study Failure Mode and Effect Analysis) (ABS, 2020; Ellis et al., 2008).

The second step—risk assessment—concerns mainly building a risk model. The methods here can be divided into qualitative and quantitative ones. The quantitative methods include: statistical analysis (based on historical records) (Blaich, Köhler, Reuter, & Hahn, 2015; Gerigk, 2012; Soares & Teixeira, 2001), Bayesian Networks (Berle et al., 2011; Gyftakis et al., 2018; Trucco et al., 2008), correlation analysis, Fuzzy Logic (Balmat, Lafont, Maifret, & Pessel, 2009; Elsayed, 2009; Johansson & Falkman, 2007), simulation-based methods (Blaich et al., 2015), or a combination of several methods (Eleye-Datubo, Wall, & Wang, 2008; Tu, Zhang, Rachmawati, Rajabally, & Huang, 2017).

With regard to qualitative risk assessment, the common methods are: Fault Tree Analysis (Hahn, 2014), Event Tree Analysis (Berle et al., 2011), risk matrixes, risk profiles, F-N curves, and relative ranking/risk indexes (ABS, 2020).

There are also risk assessment methods with a differentiation of critical factors which influence the overall risk level more heavily. They include either weights (Balmat et al., 2009; J. Liu, Yang, Wang, & Sii, 2005) or assume that only these risk variables are taken into account for which the probability of their occurrence is above a defined threshold (Trucco et al., 2008).

From the point of view of information systems, risk models are developed based on various artificial intelligence and machine learning methods. They focus mainly on modeling a "normal behavior of a ship by application of supervised and unsupervised techniques, such as classification, SVM, clustering, neural networks, or rule-based systems" (Chandola, Banerjee, & Kumar, 2009; Laxhammar & Falkman, 2010; Lee & Lee, 2006). Besides, test beds for assessment of new safety and risk applications are used (Hahn, 2014).

Table 3.1 presents a summary of popular methods for risk assessment, which are applied in the maritime domain.

The presented summary shows that there is a number of methods that can be applied to conduct maritime risk assessment. Therefore, the key issue is to choose the right method (or a combination of methods) which best matches the analyzed situation. The selected approach must also take into consideration that estimation of the probability of an adverse event and its effects. In relation to maritime transport this estimation may depend on various factors such as: itinerary, cargo size and volume, type of cargo and its properties (see Section 3.3 for a detailed overview of risk factors). One of the methods presented in this research (Chapter 7) assumes utilization of Bayesian Network (BN). The method for punctuality prediction, in turn (Chapter 8), uses concepts of a route prediction, ETA estimation, ship's density and various hazard in the maritime operational environment, including geopolitical risk. Therefore, these methods are presented in more detail in the next section.

Table 3.1. The selected risk analysis methods used in the maritime domain

Category	Method	Description and application	References
Qualitative	Expert-based analysis	Method in which risk factors, risk scenarios, their probability and impact are determined by subject matter experts; examples are brainstorming, interviews and Delphi method; used as a descriptive risk assessment; applicable to any type of risk	(Başhan, Demirel, & Gul, 2020; Ellis et al., 2008; Riveiro, 2011; van Laere & Nilsson, 2009; Wan, Yan, Zhang, & Yang, 2019b)
	Risk catalog	Collection of risks, defined using the common language; generic in nature; all items are potential risks that have been identified; each item is defined by risk type, scope and risk factor; used to generate possible risk scenarios	(Choi, Pelinovsky, Lee, & Woo, 2005; Miler, 2015).
	Fault Tree Analysis	Graphical, deductive technique for identification and analysis of risk; it starts from an undesired event and shows logical relationships between equipment failures, human errors and external events which might cause a specific event; applicable for almost every type of analysis, mostly used to address the fundamental causes of specific system failures; often used for complex electronic, control and communication systems	(ABS, 2020; Arici, Akyuz, & Arslan, 2020; Berle, Asbjørnslett, & Rice, 2011; Cem Kuzu, Akyuz, & Arslan, 2019)
	Event Tree Analysis	Graphical, deductive technique for identification and analysis of risk; it analyzes possible outcomes of an initiating event, capable of producing a mishap; applicable for almost every type of analysis, mostly used to address possible outcomes of an initiating event; often used for analysis of vessel movement mishaps and propagation of fire/explosion	(ABS, 2020; Arici, Akyuz, & Arslan, 2020; Berle, Asbjørnslett, & Rice, 2011; Ellis et al., 2008; Endrina, Rasero, & Konovessis, 2018; Hahn, 2014)
	Root cause analysis	Set of analysis tools, such as Event charting, 5 Whys technique, Root Cause Map, used to systematically discover how a mishap has occurred and the underlying root causes of the key contributors; applicable to any type of risk	(ABS, 2020)
	Risk matrix	Method used to rank the risk criticality of the failure modes; each risk scenario is evaluated taking into account its likelihood and consequences using a matrix; for each combination of likelihood-consequence, a level of risk is defined; used for evaluation of the impact of different scenarios with respect to different consequences	(ABS, 2020; Elsayed, 2009; Endrina, Rasero, & Konovessis, 2018; Wan, Yan, Zhang, & Yang, 2019b)

Category	Method	Description and application	References
Qualitative	Risk matrix	Method used to rank the risk criticality of the failure modes; each risk scenario is evaluated taking into account its likelihood and consequences using a matrix; for each combination of likelihood-consequence, a level of risk is defined; used for evaluation of the impact of different scenarios with respect to different consequences	(ABS, 2020; Elsayed, 2009; Endrina, Rasero, & Konovessis, 2018; Wan, Yan, Zhang, & Yang, 2019b)
	Preliminary risk assessment / risk profiles	Mashup-based approach used to characterize the risk associated with significant loss scenarios; relies on subject matter experts; used for generating risk profiles across broad range of activities (e.g., port-wide risk assessment)	(ABS, 2020)
	Influence diagrams	Graphical technique used to show interrelations between regulatory, operational and organizational factors; it models series of possible events and allows to define the critical events; applicable for various hazard categories like grounding, collision, fire, loss of propulsion and steering	(Berle, Asbjørnslett, & Rice, 2011; Cross & Ballesio, 2003)
Quantitative	Risk index	Method that uses various attributes (e.g., features of vessel/port) to calculate indexes that are further useful for making relative comparisons of various alternatives; generally applicable to any type of analysis situation (especially when only relative priorities are needed); extensively used to establish priorities for boarding and inspecting vessels	(ABS, 2020; Balmat, Lafont, Maifret, & Pessel, 2009; Ellis et al., 2008)
	Statistical analysis	Statistical methods like correlation, analysis of histograms, sensitivity, standard deviation; based mainly on historical data (e.g., accident statistics); used to estimate the probability of an event or risk scenario	(Blaich, Köhler, Reuter, & Hahn, 2015; Eiden & Martinsen, 2010; Endrina, Rasero, & Konovessis, 2018; Gerigk, 2012; Li, Meng, & Qu, 2012; Soares & Teixeira, 2001)

Category	Method	Description and application	References
Quantitative	Bayesian Networks	Graph-based technique, where nodes are risk variables with defined probabilities; these probabilities can depend on other node(s), through connections made by arcs; used for probability estimation of various risk scenarios, often in the causation analysis	(Berle, Asbjørnslett, & Rice, 2011; Gyftakis et al., 2018; Trucco, Cagno, Ruggeri, & Grande, 2008; Tu, Zhang, Rachmawati, Rajabally, & Huang, 2017; Wan, Yan, Zhang, Qu, & Yang, 2019a)
	Fuzzy Logic, IF-THEN rules	Fuzziness is a type of deterministic uncertainty that describes the event class ambiguity (i.e., outcomes that belong to several event classes at the same time but to different degrees); it measures the degree to which an event occurs, not whether it occurs; used when there exists an uncertainty for a risk factor (factor is vague, ambiguous, or fuzzy) and thus cannot be represented precisely by a probability distribution; allows for incorporation of human factors in risk analysis; often used with if-then rules which map probability, consequences and risk value	(Arici, Akyuz, & Arslan, 2020; Balmat, Lafont, Maifret, & Pessel, 2009; Eleye-Datubo, Wall, & Wang, 2008; Elsayed, 2009; Gaonkar et al., 2011; J. Liu, Yang, Wang, & Sii, 2005; Markowski, Mannan, & Bigoszewska, 2009; Tu, Zhang, Rachmawati, Rajabally, & Huang, 2017; Wan, Yan, Zhang, Qu, & Yang, 2019a)
	Geometrical estimation	Mathematical methods used to estimate a probability of an event; applicable for modeling risk of grounding or collision	(Li, Meng, & Qu, 2012)

Source: Own work.

interpreted. Therefore, information systems with advanced processing, analysis, and reasoning capabilities are required, which would provide a fast assessment of the situation and support users in decision-making (Pallotta, Vespe, & Bryan, 2013). It concerns real-time identification of potential maritime threats in particular.

Other reasons why the AIS suffers from some data quality problems and needs further improvements are as follows:

- Along coastal regions, ships are tracked using the terrestrial station network, offering update frequency of ship positions within 15 minutes.
- Satellite AIS reception in coastal regions, especially in areas of high vessel density such as the North Sea or the Baltic Sea, is relatively poor due to the limited storage capacity of satellites. Still, there are maritime regions where AIS coverage is limited (see an example of the Baltic Sea with poor AIS coverage, i.e., in Bothnian Bay, the East Gotland Basin, and the Bornholm Basin) (Figure 4.3).
- Despite the growing satellite constellation (totaling to 60 at the end of 2020, provided by different companies like ORBCOMM, exactEarth, or Spire Global), AIS reception on the open seas (outside of terrestrial coverage) may still be

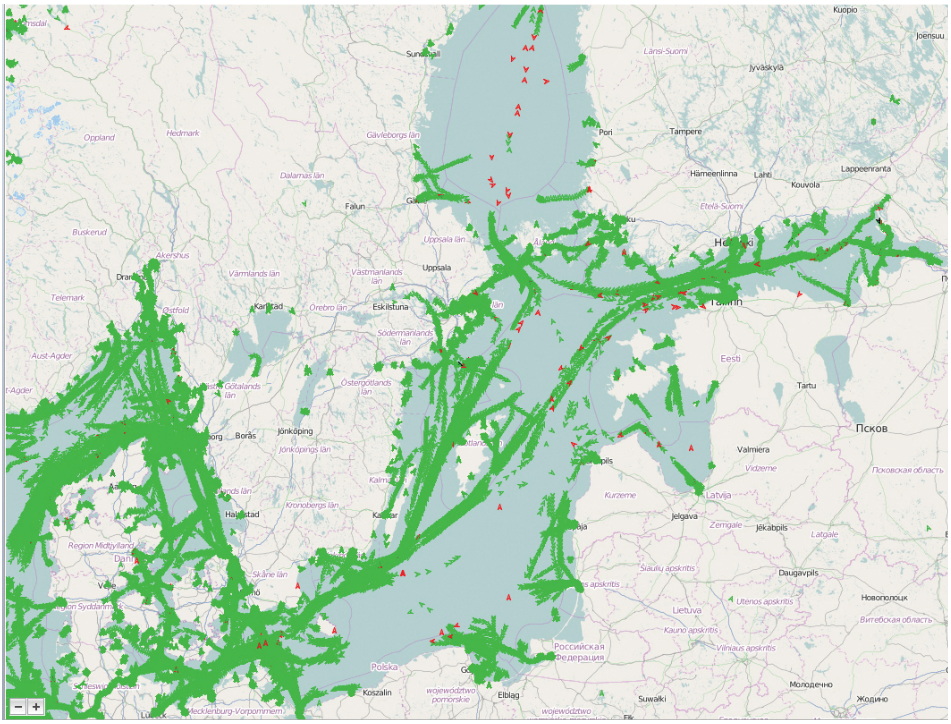

Figure 4.3. An example of AIS coverage on the Baltic Sea

Source: The SimmoViewer application developed within the SIMMO project.

limited. As a result, access gaps, i.e., time periods when a ship is not in view of an AIS satellite and no vessel position can be acquired, still happen.

To sum up, the current capabilities in the area of AIS data provision and utilization are still under development. This especially concerns the integration of data about ships from various sources and the use of intelligent data analysis tools. Even when it comes to AIS data, the usage of terrestrial and satellite-based AIS has not yet been fully exploited. As a result, there are some challenges with regard to the capabilities of maritime surveillance systems.

Long Range Identification System (LRIT)

The Long Range Identification System (LRIT) is another international tracking and identification system incorporated by the IMO under its SOLAS convention to ensure a monitoring system for ships across the world. The LRIT is required of all passenger ships, cargo ships of 300 gross tonnage and above engaged in international voyages, and mobile offshore drilling units. These ships must send reports to their flag administration at least four times a day (i.e., every 6 hours). A vessel transmits its identity, position (latitude and longitude), and the date and time of the position. The system consists of shipborne satellite communications equipment, like INMARSAT or IRIDIUM, and is a point-to-point communication system. The data transmitted within the LRIT is stored in national or regional LRIT Data Centers, which are managed by contracting governments. In the case of the European Union (EU), LRIT data are stored in the EU LRIT Data Centre and are managed by the European Maritime Safety Agency. The data are available only to authorized entities of the Member States.

Vessel Monitoring System (VMS)

Vessel Monitoring Systems (VMS) are tracking systems that are used to track and monitor the activities of commercial fishing vessels. They are mainly used for fisheries management by ensuring proper fishing practices to prevent illegal fishing. A VMS usually covers the territorial waters of a country or Exclusive Economic Zone. Unlike the AIS, it is not standardized globally. Therefore, the functionality of a VMS varies according to the requirements of the nation to which vessel is registered and the regional or national waters on which the vessel is operating. However, the ships under EU flags must send reports based on the EU standard, EU-VMS. The main disadvantages of VMS is that a VMS data are private and not publicly available.

Multi-sensor contact data

A multi-sensor signal generates contact-level data for all available sensors, such as coastal radar, SAR, video, IR, etc.

Coastal High Frequency Radar (HFR) provides regular, high-quality information on ocean surface currents. The HF-Radar provides real-time observational data of the surface currents via coastal stations. Understanding marine currents is of great importance for the development of activities related to maritime transport, since it provides information about the trajectory of a vessel or drifting object. Thus, it allows vessels in the radar range to be tracked.

Synthetic-aperture radar (SAR) is a form of radar that is used to create two- or three-dimensional images of objects, such as maritime areas. SAR uses the motion of the radar antenna over a target region to provide better spatial resolution than conventional radars. An SAR is typically mounted on a moving platform, such as an aircraft or spacecraft.

SAR images have a wide scope of applications in remote sensing and mapping the surface of the Earth. It is also a useful technology in environmental monitoring, foer example oil spills, flooding, urban growth, and global change. Measurements that cover an ocean area can be used to deduce surface waves or to detect and analyze surface features such as fronts, eddies, and oil slicks. SAR can also be implemented as inverse SAR in order to observe moving targets over time (e.g., ships). In the maritime domain, apart from mapping the surface of the sea and oceanography, it is used to detect objects in open seas. Some SAR images are published by the European Space Agency, but access to the data requires prior registration and the submission and approval of a proposal.[7]

Signal Intelligence refers to the capability to detect, characterize, and geolocate various types of radio frequency emitters. Specifically, in the context of maritime surveillance and the detection of non-cooperative ships, signal intelligence data are key. Signal intelligence data are commonly collected by various military stakeholders, but recently private entities are also offering such capabilities, for example HawkEye 360 (US) or Kleos (UK).

Other sensor data include cameras, closed circuit television (CCTV), infra red imaging, and underwater sensors.

Geographic Information System data

A Geographic Information System (GIS) is a system designed to capture, store, manage and analyze spatial and geographic data. GIS datasets can be used in various applications, especially for locating all kinds of phenomena, especially those which vary over time, and for further visualizing them on maps. In the

7. https://earth.esa.int/web/guest/data-access/products-typology/radar-imagery/

maritime domain examples of GIS data are port locations, maritime protected areas, ocean fishing regions, fish species habitat distribution, political national borders and Exclusive Economic Zones, bathymetry, etc. Much of the data is freely available for potential users.

4.1.2. Weather data

There are several sources that provides weather data for maritime areas on a regular basis. They can be grouped into two categories:

- sources providing only forecast data, for example, windy.com, predictwind.com, NOAA;
- sources providing forecast data and historical weather data, for example, yr.no, Copernicus, or the European Centre for Medium-Range Weather Forecasts (ECMWF).

The first group provides only forecast data for a defined number of days in advance (e.g., 5- or 10-day), while the second one additionally offers information about actual weather in the past in the form of daily, monthly, or yearly means. The available weather data sources also differ with respect to the area covered (global or selected local areas), data resolution (from 30 km up to 7 km), update frequency (once or several times a day), the forecast model used, and the scope of the data (the set of weather parameters that can be observed). Moreover, the technical parameters of the available data may vary with regard to the data format (the most popular are grib or NetCDF files, though JSON/XLM formats are also supported), how the data are shared (via API, a webservice, or ftp), and data accessibility (there are fully open and free data sources, such as Copernicus, yr.no, or NOAA National Weather Service), commercial sources with free and paid options available (e.g., windy.com, or predictwind.com), as well as sources available only to authorized users (e.g., ECMWF).

In the study presented in this book historical weather data from Copernicus were used. Therefore, this data source is described in more detail.

Copernicus[8] is the European Union program aimed at developing European information services based on satellite Earth observation. It is managed by the EU and the European Space Agency (ESA). Within this program vast amounts of global data from satellites and seaborne measurement systems are provided. The content is freely and openly accessible to users.

The information services offered by Copernicus can be grouped into six main themes: land, ocean, emergency response, atmosphere, security, and climate

8. http://www.copernicus.eu

change. For the scope of this research, the ocean topic is highly relevant. Copernicus offers sea status observation and forecast information for various parameters like wind, temperature, ice cover, salinity, or chlorophyll. These datasets can be downloaded in an automatically from the data hub.[9]

The main source of maritime weather data is the Copernicus Marine Environment Monitoring Service (CMEMS).[10] The service provides information from both satellite and in situ observations, daily state-of-the-art analyses and forecasts daily, and historical weather data for different maritime areas. The data are available through the CMEMS services that are open, free, reliable, and sustainable.[11]

The Copernicus weather data are stored in NetCDF files—the Network Common Data Form. This is a file format dedicated to sharing array-oriented scientific data. It is also the standard of the Open Geospatial Consortium (Opengeospatial.org, 2018). Version 4.0 (released in 2008) allows for the HDF5 data file format. Hierarchical Data Format (HDF) is a file format designed to store and organize large amounts of data.

The characteristic thing about NetCDF is its capability of self-description. The header of the file describes the layout of the rest of the file, in particular the data arrays. It can also provide arbitrary file metadata in the form of name-value attributes. The NetCDF format is platform independent and there libraries available for all major programming languages.

For the research presented further in Chapter 9, from all available Copernicus services we used only those that provide parameters of interest to our analysis, that is, data about wind (speed and direction), wave height, sea currents and tides, ice coverage, and covering selected maritime areas (i.e., the Baltic Sea, the North Sea, and the Norwegian Sea in the Arctic Ocean). The process of acquiring and extracting weather data from Copernicus is elaborated in Section 4.6.3.

4.1.3. Internet sources

Sensor data, like the AIS, provide only basic information about a given ship. In order to complement the sensor data with relevant information about ships, external sources and databases can be used. A great example might be various Internet sources that publish maritime-related data.

Open Internet data sources can provide general ship data (flag, detailed type, length, gross tonnage, capacity, technical specifications, and construction details),

9. http://marine.copernicus.eu/services-portfolio/access-to-products/
10. http://marine.copernicus.eu
11. The detailed catalogue of services is available at http://marine.copernicus.eu/wp-content/uploads/2016/06/r2421_9_catalogue_services.pdf

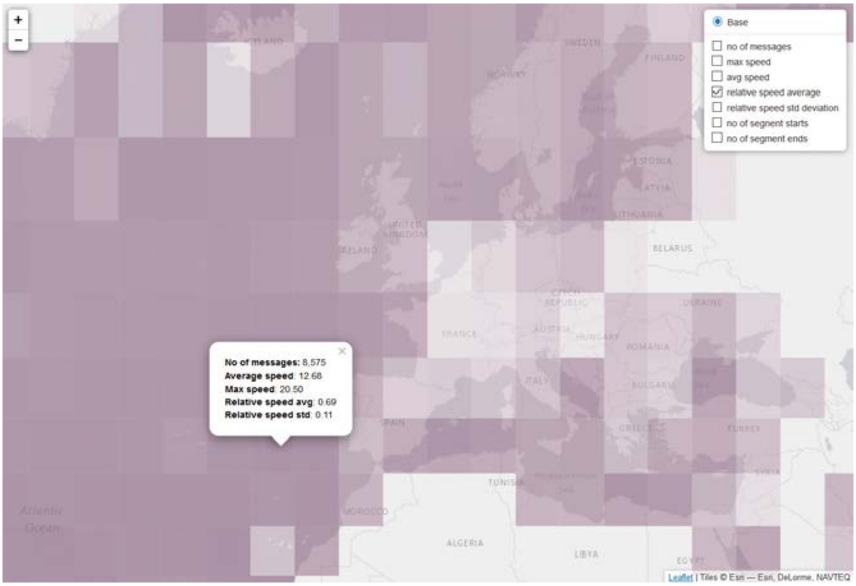

Figure 6.5. Average relative speed of vessels in a given segment, Europe

Source: Own work.

Figure 6.6. Standard deviation of the relative speed of vessels in a given segment, Europe

Source: Own work.

(5) Compare the relative speed with the average relative speed characteristic for a given region. The allowed deviation is determined by standard deviation.

The algorithm required some tuning, i.e., it had to be decided what the reasonable deviation was. Figure 6.7 compares two variants: with 1-σ and 2-σ, where σ is the standard deviation. The latter seems to return fewer false positives, therefore in further experiments this value was kept. However, this value can be further parameterized, if needed. The meaning of the colours is as follows:

- red: the ship is traveling at the relative speed lower than the average relative speed *minus* 2 times standard deviation (2σ);
- green: the ship is traveling at the relative speed higher than the average relative speed *plus* 2 times standard deviation.

The red-marked messages are considered as loitering.

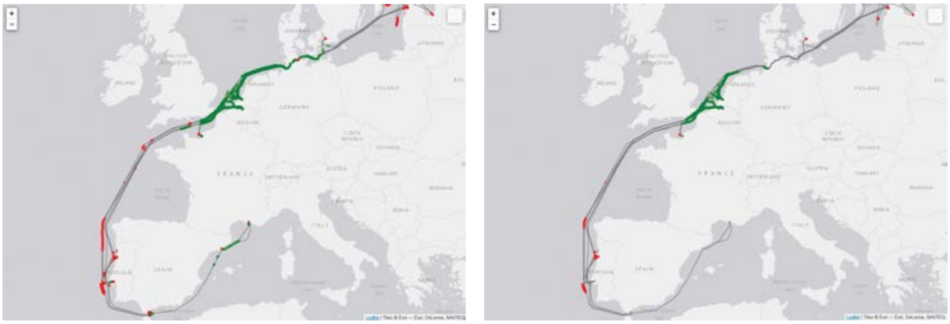

Figure 6.7. Relative speed anomalies with two deviation variants, MMSI 210688000

Source: Own work.

Route anomaly. The route anomaly is defined as an unpredictable movement, i.e., not following a trend or a pattern. The most typical examples are a sudden change of speed or course over ground. In this approach a prediction is made based on a current trajectory of an analysed vessel. We analyse the trajectory and based on the last three locations (from AIS) we extrapolate the next location.

During this analysis we came across several sub-types of how anomaly can be discovered:

- average speed anomaly: speed higher than possible for a ship; this way we also clean incorrect AIS data readings,
- location anomaly: a ship is found in another location than inferred from the previous course,
- triangle anomaly: a ship is traveling along the longer edges of a triangle instead of the shorter, e.g., making a zig-zag or traveling back and forth,

- angle anomaly: change of course over 90 degrees; we assume that a ship should not change course rapidly; if this is the case, it should be interpreted as loitering.

Unpredictable location anomaly. In this method the following algorithm is used:

(1) Take two preceding locations along with timestamps.
(2) Based on speed and timing predict the next location.
(3) If the real position is different from the one predicted, raise an issue.

In order not to raise too many warnings, we allow the deviation from the predicted position of 3 miles (the tolerance). We also do not try to predict if the time intervals between positions are longer than a specific amount of time (here 1 hour). Prediction is also not conducted at the beginning of the travel segment, when the necessary number of measurements is not yet available. Sample anomalies using the method are presented in Figure 6.8.

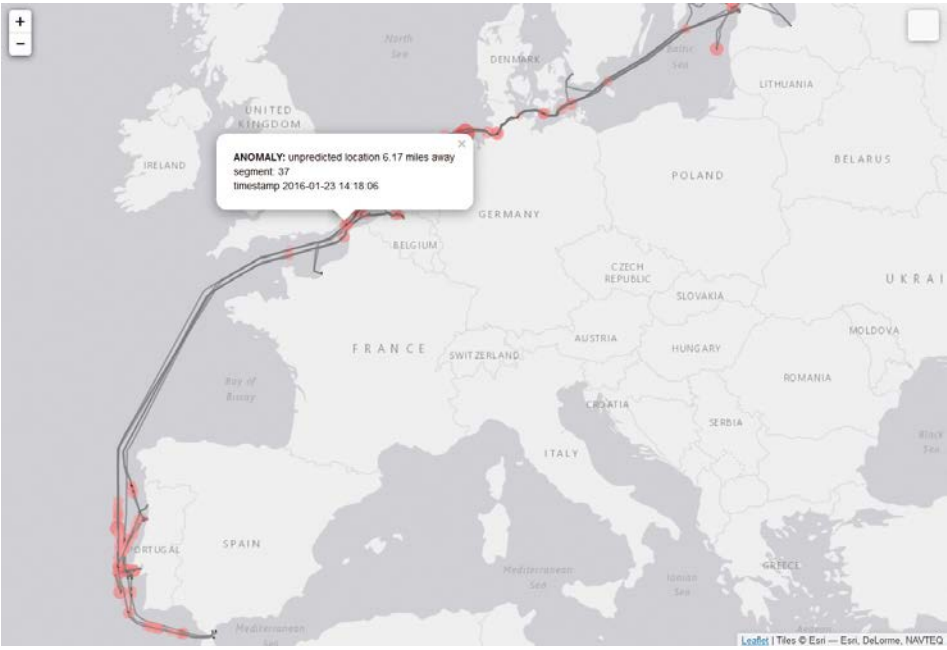

Figure 6.8. Trajectory of ship Amazonith (MMSI: 210688000) with unpredictable location anomalies

Source: Own work.

What can be concluded from the mentioned figure is that anomalies have the tendency to focus around certain areas. These are the regions where heavier marine traffic can be expected, for example in the English Channel and around Portuguese ports. Another explanation is that more anomalies occur around destination ports. This can be connected with waiting for the permission to enter the port.

Sharp change of course. In this case the following heuristics is used: if the ship changes the course more than 90 degrees, as measured between three consecutive messages, then the issue is raised.

Thanks to the proposed method it was possible to discover quite interesting angle anomalies. For example, one vessel, while waiting for the entry to the port, was traveling in circles (see Figure 6.9). More rational behaviour would be rather to stop on the high sea, so it was another reason why such an example should be treated as a loitering anomaly.

Some of the discovered anomalies seemed to be false positives and required a more detailed analysis. For example, in some cases anomalies were discovered on seemingly straight course trajectories (see Figure 6.10 left). In these particular cases the turn was almost 180 degrees. However, what was peculiar, it always occurred two times in a row. Later on, we identified the source of the problem: messages

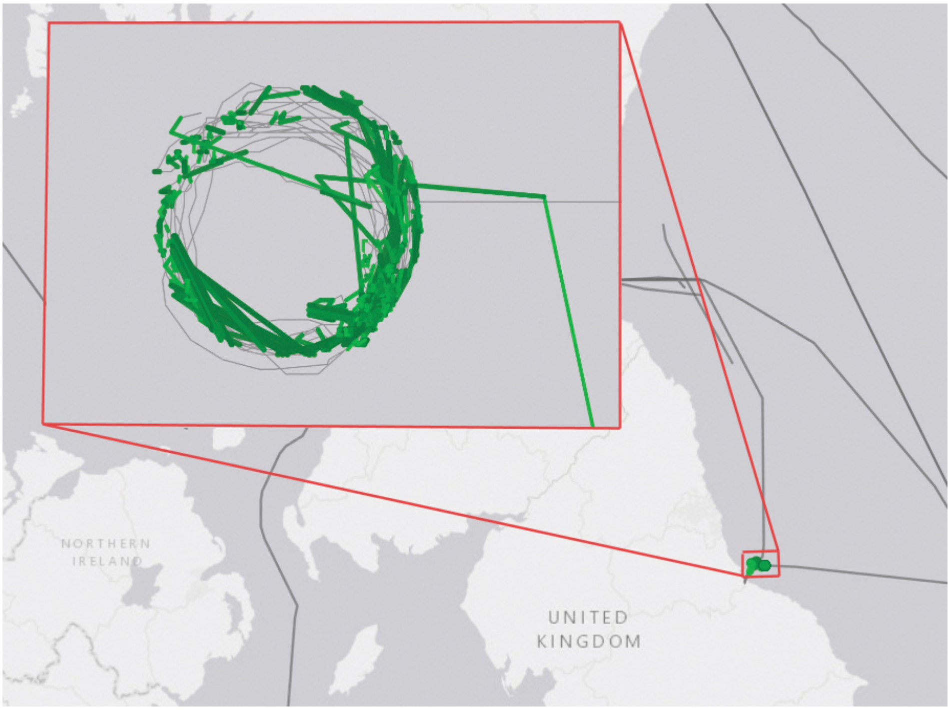

Figure 6.9. Angle anomaly—a vessel traveling in small circles

Source: Own work.

were received by various AIS devices, which had unsynchronized clocks. Thus, the problem resulted from the incorrect ordering of the points forming the trajectory. Fortunately, the device responsible for incorrect timestamps was responsible for less than 0.5% of messages. Another source of false positives can be manoeuvres close to ports (see Figure 6.10 right).

Figure 6.10. Trajectories with marked angle anomalies. Left: anomalies on straight trajectories. Right: false positives around ports

Source: Own work.

Travel-time anomaly. We also proposed a method to discover loitering by looking at the longer segments of ships trajectories, not only at single messages. We tried to estimate the typical travel time between certain areas. Loitering would be discovered when a non-typical travel time was detected. More specifically, it would happen when a vessel was not following a normal or historical route: different times of travel when compared to its own historical routes or routes of a similar ships (type, size, cargo).

For the detection of this type of loitering we needed typical travel times between trajectory segments. In our database, after execution of previous algorithms, we had already had trajectory segments, i.e., parts that have the same navigational status and contain consecutive locations of a ship. It was then possible to measure the travel time and distance between the starts and ends of many segments. If normally the travel takes 12 days and we observe 17 days, then the whole track can

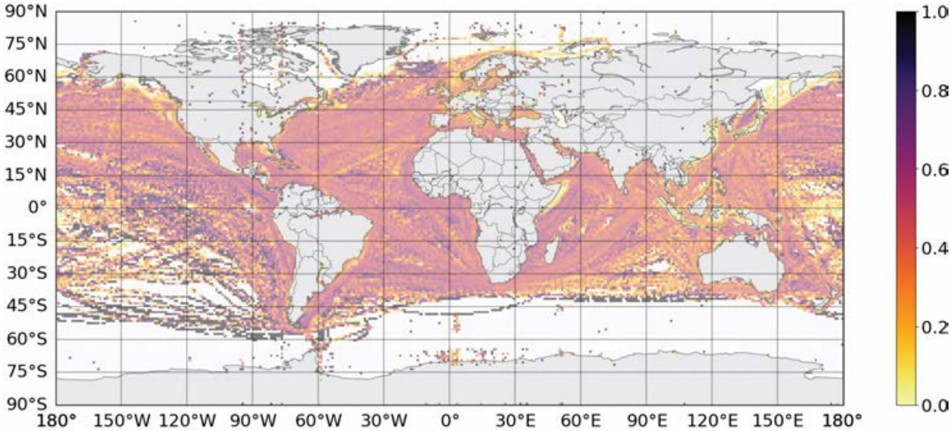

Figure 9.7. Anomaly S3—AIS position reports sent by FoC tankers (relative)

Source: (Filipiak et al., 2018).

Further on, we analysed the spatial distribution of tankers being on a detention or banned ships list. A ship can be subject to Port State Control (PSC), after which, in the case of an occurrence of any deficiencies that are clearly hazardous to the safety of a state or to the environment, a ship can be detained (Anomaly S5). If a ship was detained three or more times by a maritime authority during the last 12 or 24 months, it is classified as banned or added to the list of under-performing ships by a given MoU (Anomaly S4). In the course of the analysis just a single banned tanker was found in the area of the Gulf of Oman (Figure 9.8). On the other hand, detained tankers were found across the whole globe. However, it seems that they were active mostly near Micronesia and the Marshall Islands (Figure 9.10).

Then, we analysed classification certificates issued by the so-called low-performing Recognized Organizations / classification societies (RO). Classification societies are non-governmental organizations that establish and maintain technical standards for construction and operation of marine vessels. The primary role of a classification society is to validate if a design and technical equipment of a ship are in accordance with the published standards. If a ship meets all the requirements, a classification society issues a classification certificate. However, among the classification societies, there are ones that do not perform a minimum number of inspections in a 3-year period and are called Recognized Organizations (RO). If ROs do not meet the criteria for their ships to qualify as Low Risk Ships, they are listed as low-performing ROs (Anomaly S6). Thus, ships having a classification certificate issued by a low-performing RO are potentially dangerous. Our analysis showed that in 2015 such tankers concentrated mostly at the Chinese coast and particularly near Taiwan (Figure 9.9).

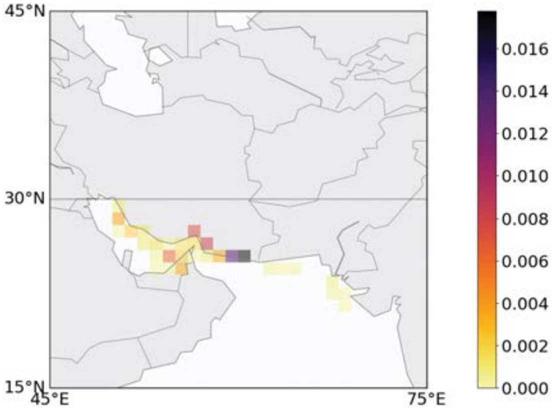

Figure 9.8. Anomalies S4 / S8—AIS position reports sent by a banned and withdrawn or suspended tanker (relative)

Source: (Filipiak et al., 2018).

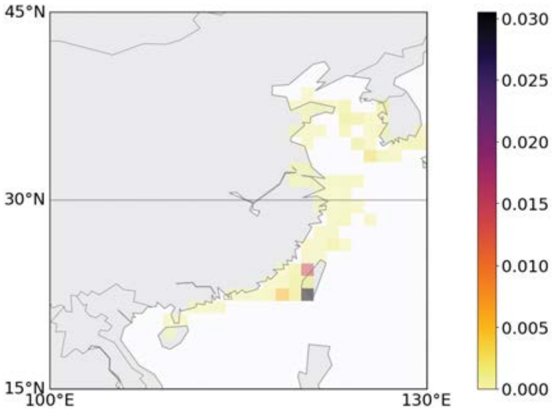

Figure 9.9. Anomaly S6—AIS position reports sent by tankers belonging to low performing ROs (relative)

Source: (Filipiak et al., 2018).

Classification societies are also responsible for granting a classification status for ships. This status is designated based on a periodical survey of a ship and it ensures that a ship meets the classification standards. There are five classification statuses that may be granted: delivered, suspended, reinstated, withdrawn, or reassigned. The ships with the withdrawn and suspended status may be regarded as an anomaly (Anomaly S8). We detected only one tanker that matched this criterion—it was the same vessel as presented in Figure 9.8.

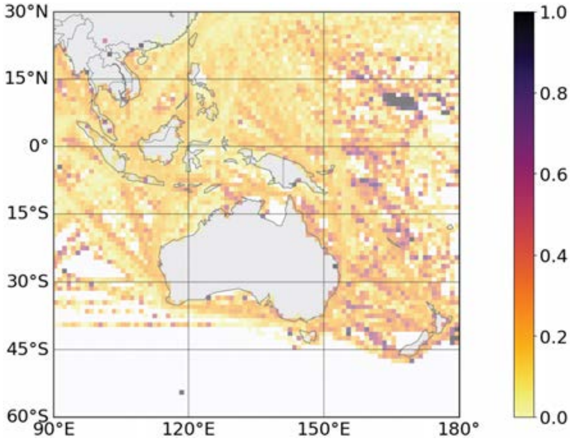

Figure 9.10. Anomaly S5—AIS position reports sent by tankers marked as detained (relative)

Source: (Filipiak et al., 2018).

The final static anomaly concerned tankers being owned / managed by a poor--performing company. The European Maritime Safety Agency (EMSA) publishes a list of such poor-performing companies (Anomaly S7). In the course of the analysis, 24 tankers matching that criterion were identified. They were particularly active in some parts of the Pacific Ocean, south of Hawaii (Figure 9.11).

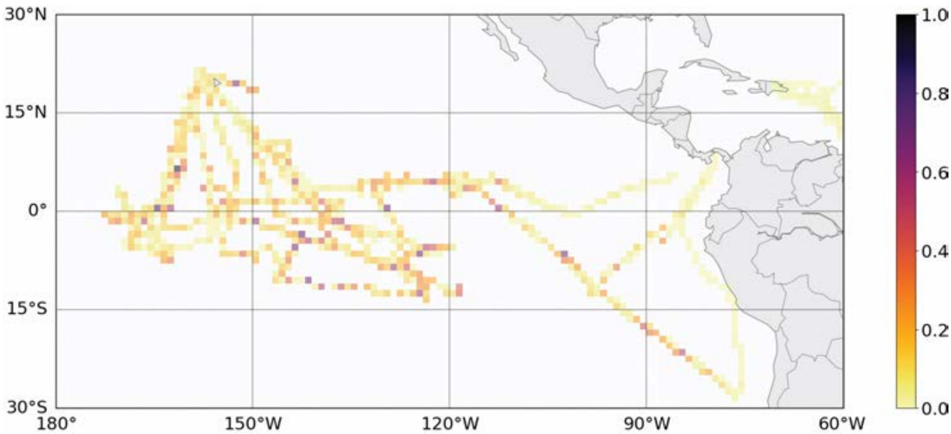

Figure 9.11. Anomaly S7—AIS position reports sent by tankers belonging to low performing companies (relative to the number of all considered position reports in a segment)

Source: (Filipiak et al., 2018).

Table 9.2 summarizes the detected static anomalies for tankers.

Table 9.2. Static anomalies related to tankers in 2015

ID	Anomaly	No. of tankers	%
S4	IMO in banned list	1	0.003
S8	Withdrawn or suspended	1	0.003
S6	Low performing RO	5	0.014
S7	Low performing company	24	0.069
S1	Black-listed flag	1512	4.362
S2	Gray-listed flag	1521	4.388
S5	IMO in detention list	1983	5.721
S3	Flag of Convenience	7097	20.475
	Tankers without static data anomalies	24345	70.235
	Tankers total	34662	100.000

Source: (Filipiak et al., 2018).

9.1.5. Loitering detection

As already indicated in Section 6.4, loitering is mainly related to an anomalous speed of a vessel. In our research, loitering-related anomalies were divided into seven categories: invalid coordinates (L1), location (L2) or speed (L3), sharp change of course (L4), unpredicted location (L5), and unusually low (L6) or high speed (L7).

The first three types of anomalies (L1–L3) result from the verification of the correctness of AIS data values. First, we checked if correct coordinates are provided in an AIS message. If not, Anomaly L1 is reported. Then, whether the reported speed over ground is within expected limits (Anomaly L2). We set a threshold at 25 knots, meaning that a speed above this value will be perceived as an anomaly. Thanks to this, segments with the highest relative number of reports of invalid speeds were identified (Figure 9.12). In the next step, we checked whether an actual position of a ship is reliable considering its potential speed over ground (Anomaly L3). This method makes it possible to eliminate problems with incorrect AIS reading since it filters out cases of sudden *teleportation* of a ship (Figure 9.13).

The next method concerns an angle anomaly (Anomaly L4), which detects a sharp change of course (over 90 degrees). If a ship changes its course so rapidly, it might be interpreted as loitering (Figure 9.14).

Anomaly L5—unpredicted location—concerns a situation when a ship is found in another location than inferred from its previous course. The expected location is predicted based on two previous locations of a ship (points and times), assuming that a vessel should continue its trajectory. A location other than the predicted

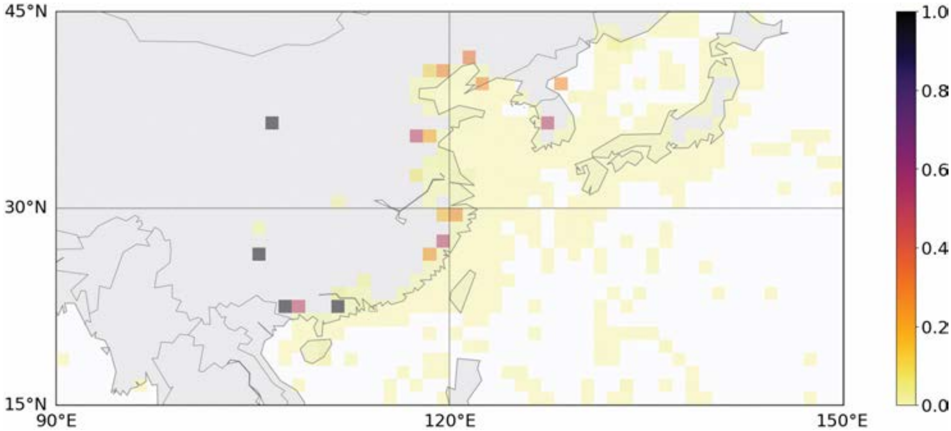

Figure 9.12. Anomaly L2—AIS position reports with an invalid speed (relative)

Source: (Filipiak et al., 2018).

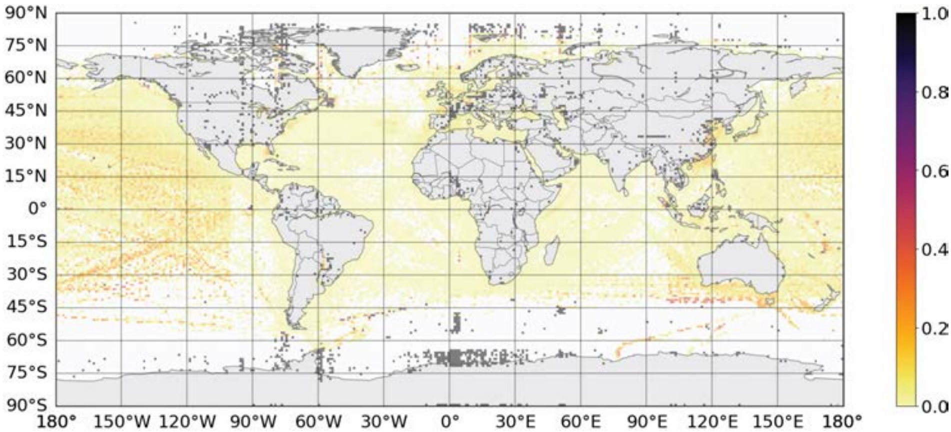

Figure 9.13. Anomaly L3—AIS position reports with an invalid location (relative)

Source: (Filipiak et al., 2018).

one with a margin of 3 miles is considered as anomalous. However, ships that do not move for over 1 hour are excluded. The detected anomalies with regard to unpredicted location are presented in Figure 9.15.

The last method tests whether a ship is sailing with an unusually low or high speed (Anomalies L6 and L7). Loitering occurs when a ship being on the high sea starts sailing with a low speed. This method compares the ship's relative speed in a given segment with the average relative speed and its standard deviation calculated for this segment. If the difference exceeds a defined threshold (value of 2 standard deviations), this position report is considered as anomalous. The results for these two types of anomalies are presented on Figures 9.16 and 9.17.

Pandas UDF. Nevertheless, for the sake of precision in the further experiments kNN with the haversine metric and BallTree partitioning were used.

9.2.6. Reconstruction of edges

Having all AIS points annotated with the nearest waypoints, the next step is reconstruction of the edges between these waypoints. The algorithm for this approach is presented in Section 5.3.3. In this case, due to incomplete distribution and often low quality of AIS data, several measures need to be undertaken to achieve good results. Fortunately, generation of edges proved to be a less challenging task from the performance point of view. The only optimization step that had to be applied was materialization of the enriched AIS dataset. For some reason even caching was not helpful—grouping the edges spawned a re-calculation of the closest waypoints. Therefore, our process is divided into two steps:

(1) from raw AIS data to enriched AIS—results are stored in CSV files;
(2) from enriched AIS data (read from the CSV file) to the edges—results are stored in two files: nodes.csv and edges.csv, representing the mesh.

Nevertheless, there were other challenges concerning the output mesh. A visual introspection of maps, which show the generated mesh, proved that the method generated 'impossible' or 'inappropriate' connections between some waypoints which further on had to be eliminated. It was caused partly by the low AIS data quality. However, other means were undertaken to improve the final mesh. Some of the applied techniques are presented below.

For all the tasks presented below we used AIS data from 8 consecutive weeks (2019 w36–w43). AIS data was filtered, so that only AIS from the German Bight for tankers, cargo and passenger ships were included. For this input data, 8,809 waypoints were identified. Input AIS data contained 3,639,631 rows, in which 4,857 distinct MMSIs were found. The key manoeuvre points identified with the CUSUM method contained 414,824 rows, in which 4,609 distinct MMSIs were found.

Edges calculated based on the full AIS data (border points). When a vessel is moving along its trajectory, it passes many waypoints. We know which points are passed by, as AIS data is already annotated with the closest waypoint (see Section 9.2.5). Sometimes there are several consecutive AIS messages with the same waypoint, especially if the distances between the waypoints are long. We need to identify only the places where the 'borders' between affiliation of AIS to different waypoints are crossed, i.e., a given message has a different waypoint from the previous message.

In the implementation of the algorithm, the effect described above is achieved by using the so-called window functions. In these functions it is possible to refer to the previous value with function `lag`. We are then able to identify the 'changed' rows as described in the listing below:

```
from pyspark.sql.window import Window

w = Window.partitionBy("mmsi").orderBy("timestamp_ais")
sdf_ais_with_waypoint_changed = sdf_ais_with_waypoint_idx \
.withColumn("from_waypoint", F.lag('to_waypoint', 1, 0).over(w))
sdf_ais_with_waypoint_changed = sdf_ais_with_waypoint_changed \
.withColumn("changed",
 ↪  (sdf_ais_with_waypoint_changed['from_waypoint'] !=
sdf_ais_with_waypoint_changed['to_waypoint']).cast('int'))
sdf_ais_with_waypoint_filtered = sdf_ais_with_waypoint_changed \
.withColumn("timestamp_delta",
 ↪  sdf_ais_with_waypoint_idx.timestamp_ais-F.lag(
'timestamp_ais', 1, 0).over(w)) \
.where('changed=1') \
.where('from_waypoint<>0')
```

By applying the above procedure, we reduced the initial 3,639,631 messages to the filtered 1,494,227 messages. They contain only the points where a current waypoint (to_waypoint) is different from the previous waypoint (from_waypoint). We can construct a dataset with edges using grouping by from_waypoint and to_waypoint, as illustrated in the code below. We also calculate group statistics like the number of vessels traversing specific edges or time-related stats.

```
sdf_edges = sdf_ais_with_waypoint_filtered \
.groupBy("from_waypoint", "to_waypoint") \
.agg(F.count("*").alias("cnt"),
F.avg("lon").alias("lon"),
F.avg("lat").alias("lat"),
F.avg("timestamp_delta").alias("avg_time"),
F.min("timestamp_delta").alias("min_time"),
F.max("timestamp_delta").alias("max_time"),
F.stddev("timestamp_delta").alias("stddev_time"))
```

In this specific example we generated 170,644 edges between 8,809 waypoints. The visualization of this mesh on the map is presented in Figure 9.38 (p. 294).

Analysis of distance on edges. By looking at the Figure 9.38 (p. 294), we observe a big number of edges that span long distances. Having been visualized on the map, they very often cross the land. Therefore, we decided to study in detail the lengths of the edges to identify and possibly eliminate the problem.

In Figure 9.27 we demonstrate the histogram of edges lengths. Please note that they y-axis is logarithmic. There are almost 50 edges that span two waypoints that are at least 500 km apart. It reveals the weakness of the approach.

Therefore, we had to adjust the approach to eliminate the longest edges. It was done by adding a function FilterEdges (see Algorithm 5.1, as applied for the visualization of the meshes presented in Figures 9.39 (p. 295) and 9.40 (p. 296).

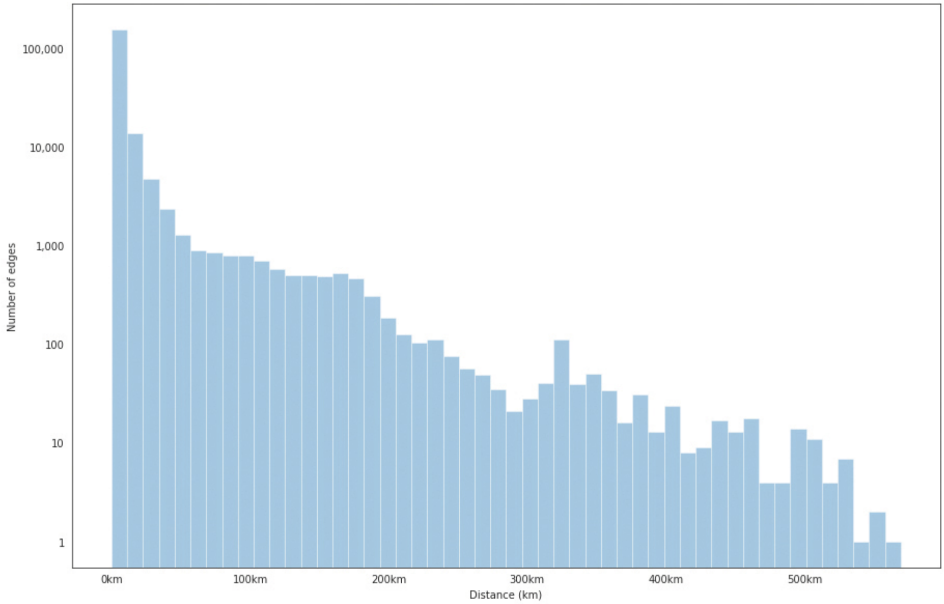

Figure 9.27. Distance between edges in a mesh for all ships in the German Bight

Source: Own work.

Analysis of timestamp delta. AIS data is timestamped. When we analyse a specific trajectory of a given vessel, we can measure the time that passed between two consecutive messages. It is called a timestamp delta, in code referred to as timestamp_delta.

Having calculated the time necessary to pass from one waypoint to another (column timestamp_delta), it should be possible to propose the fastest route. Unfortunately, vessels do not go from waypoint to waypoint. Instead, they go between some locations that are nearby the waypoints. Moreover, when aggregated, there is no guarantee that time will be measured between the same points.

We conducted an analysis of timestamp deltas. To present the results of this analysis, below we show a series of histograms, as a single chart is not able to provide enough details. Figure 9.28 presents the overall histogram for all the data from the 8-week period. We see that there are several trajectories that contain gaps of more than 1,000 hours between the messages. The number is not significant but it still can be filtered. The majority of deltas, i.e., more than 1,000,000, still concentrate around zero.

In order to see the details, we need to zoom in the x-axis and show data only for 24 hours. Figure 9.29 presents the results with an increased resolution. We can observe that after filtering longer deltas the remaining messages are not very separate. We need to increase the resolution once more.

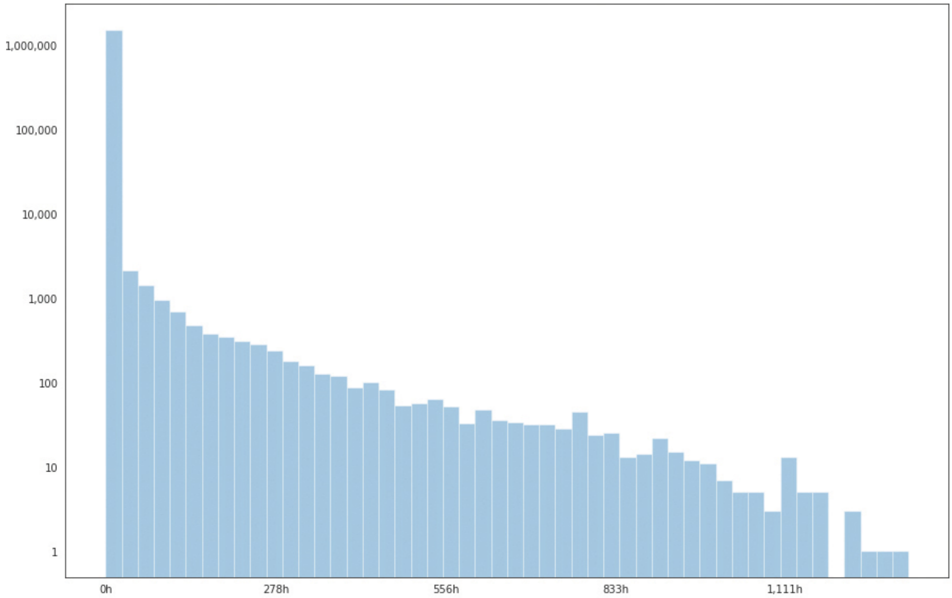

Figure 9.28. Timestamp delta calculated for the whole 8-week period

Source: Own work.

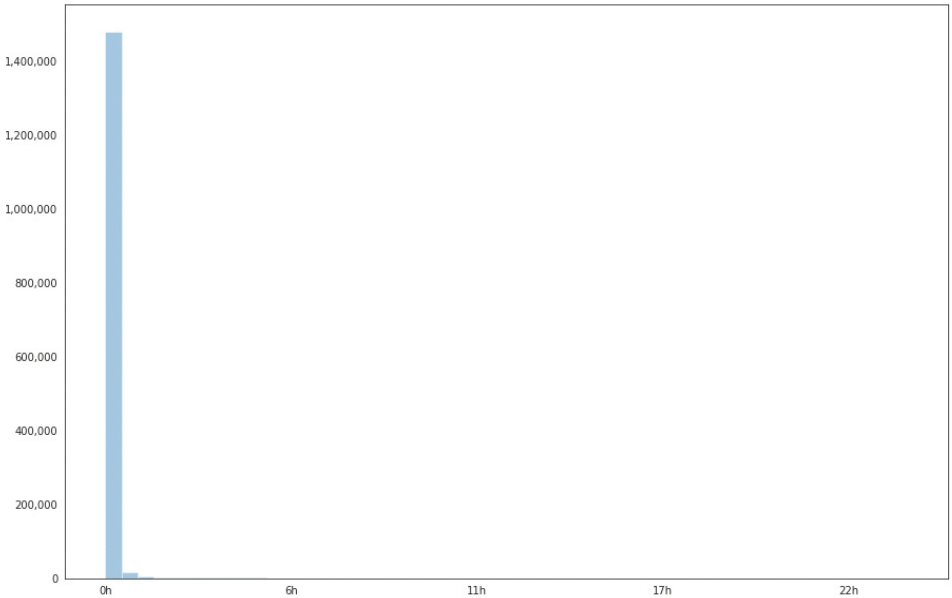

Figure 9.29. Timestamp delta restricted to 24 hours

Source: Own work.

Figure 9.30 presents timestamps deltas for messages that appeared within a period of one hour. The chart reflects the expected distribution of time differences between messages.

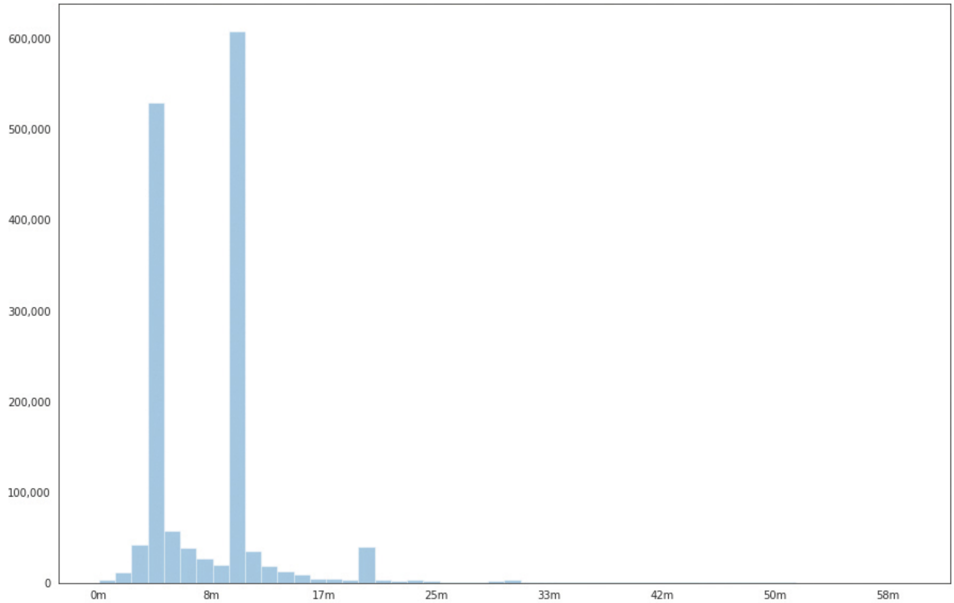

Figure 9.30. Timestamp delta restricted to 1 hour

Source: Own work.

The final chart shows which edges span long distances—see the dark lines in Figure 9.31. These are mostly edges close to the boundaries of the considered area, so it may mean that a vessel left the area and then came back. Thus, a more careful filtering or segmentation is necessary.

To conclude, such a distribution of timestamps suggests that we can safely filter out outliers, i.e., AIS messages that are too far away from each other to form a trajectory. Thus, we can also avoid joining the waypoints that are too far away (or at least are not neighbours).

If we combine two phenomena—imprecise calculation of time deltas and long-distance edges—we also observe anomalies in the average speed as it is calculated as distance divided by time. Figure 9.32 presents the histogram of the average speed. The calculation was conducted for all waypoints.

Edges between minimum distance points (mindist). The analysis conducted in the previous paragraph revealed that more realistic time deltas between waypoints are needed. Our previous approach correctly identified the transition from one waypoint (a waypoint segment to be more precise) to another. We referred to them