Student thesis series INES nr 669

Utilizing AI to enhance monitoring of fishing activities via cooperative and non-cooperative methods

Ingthor Ingason

2024
Department of
Physical Geography and Ecosystem Science
Lund University
Sölvegatan 12
S-223 62 Lund
Sweden

Ingthor Ingason (2024).
Utilizing AI to enhance monitoring of fishing activities via cooperative and non-cooperative methods
Master degree thesis, 30 credits in *Subject of degree*
Department of Physical Geography and Ecosystem Science, Lund University

Level: Master of Science (MSc)

Course duration: January 2024 until June 2024

Utilizing AI to enhance monitoring of fishing activities via cooperative and non-cooperative methods

---

Ingthor Ingason

Master thesis, 30 credits, in *GIS and Remote Sensing*

Supervisor:
Pengxiang Zhao
Dep. of Physical Geography and Ecosystem Science, Lund University

Co-supervisor:
Josefine Larsson
Marint centrum, Simrishamn Kommun, Simrishamn, Sweden

Exam committee:

Zheng Duan, Dep. of Physical Geography and Ecosystem Science, Lund University
Marko Scholze, Dep. of Physical Geography and Ecosystem Science, Lund University

Acknowledgements

I would like to express my gratitude to friends and family who have supported me during the development of this thesis. Special thanks to Josefine Larsson at Marine Center for sharing her expert knowledge and providing invaluable insights throughout my research. I am also grateful to my thesis supervisor, Pengxiang Zhao, for helpful suggestions, guidance and constructive feedback.

Abstract

The fishing industry is vital for global food security, providing nutrient-rich food, supporting livelihoods and contributing to economic growth. As the global population continues to grow, it is essential to approach fishing in a sustainable way and preserve marine ecosystems. The United Nations has recognized overfishing as a significant issue and for 2017 it was estimated that about 34% of fish stock was overfished worldwide. However, overfishing isn't the only threat to our oceans. Illegal, unreported, and unregulated (IUU) fishing also poses a major global challenge. While overfishing depletes fish populations through unsustainable practices, IUU fishing further exacerbates the problem by undermining regulations and conservation efforts. One way to help solve these problems is with effective monitoring of fishing efforts. The two main ways to do so are via cooperative methods, where fishing vessels are responsible for broadcasting their own data, and non-cooperative methods, which use remote sensing technology to monitor vessels without any involvement from the vessels themselves. This project investigates whether fishing operations in the southern Baltic Sea can be effectively monitored using publicly available data. More specifically, it utilizes the Automatic Identification System (AIS) which is a cooperative method where ships broadcast information on their location and movements via radio signals and Synthetic Aperture Radar (SAR) images, an active remote sensing technique used in non-cooperative monitoring to detect ships from satellites. This project utilized historical AIS data for the year of 2018 to demonstrate the effectiveness of AIS in monitoring fishing activities on a small scale. A machine learning model was used to predict fishing events from individual ship paths. On a larger scale, spatiotemporal analyses were performed which gave insight into fishing trends and patterns throughout the year. Additionally, a deep learning model was employed to detect ships on SAR images which in turn was used to get a sense of the level of AIS uptake among the fishing fleet in the Southern Baltic Sea. The results demonstrate that both AIS and SAR data can be effectively used in the context of monitoring fishing activities. The combined use of these methods revealed that the vast majority of boats do in fact transmit their AIS data in compliance with European regulations.

Keywords: Fishing, Monitoring, Remote sensing, Machine Learning, Deep Learning, Synthetic Aperture Radar, Automatic Identification System

Contents

1. Introduction

The fishing industry plays a crucial role in ensuring food security worldwide. It provides people with nutrient rich food, supports livelihoods and sustains the economy of many countries. With this in mind and considering that over three billion individuals rely on the ocean for their livelihoods (Hendriks, 2022) as well as the fact that the world population continues to grow, it is of utmost importance that we approach this vital resource in a sustainable way.

In the last half-century, we have seen the global population double while the fish and seafood production has had a fourfold increase (Ritchie & Roser, 2021). This indicates that the average consumption of fish per person has also risen alongside the global population which further underscores the need for sustainable fishing practices to ensure the long-term health of marine ecosystems and secure food sources for future generations. The increasing adoption of aquaculture in recent years has helped fill the global demand for seafood and helped alleviate pressure on wild fish stocks. However, aquaculture is not a perfect alternative and we still rely heavily on wild fish populations and therefore practising sustainable fishing is needed. In this context, sustainable fishing refers to catching the optimal amount of fish to prevent stock depletion while maximizing food resources and income. Maintaining this balance can be challenging. Capturing too many fish from a particular stock leads to overfishing, which is a significant global issue and one that needs to be eradicate. For 2017 it was estimated that about 34% of fish stock was overfished worldwide (Ritchie & Roser, 2021). The United Nations (UN) has recognized overfishing as a worldwide issue that diminishes food production, impairs ecosystem functions and decreases biodiversity (Snapir, Waine, & Biermann, 2019).

Another issue facing the fishing industry is illegal, unreported, and unregulated fishing (IUU) that contributes significantly to overfishing. IUU is a term defined by the Food and Agriculture Organization (FAO) of the United Nations and it encompasses a range of activities violating national and international fishing laws such as fishing without permission, misreporting catches, operating in protected areas and more (Food & of the United Nations, 2024). This in turn undermines efforts to manage and conserve fish stocks making it harder to tackle. The economic motivations behind IUU are easy to see, indicating that such practices will not cease by themselves and some framework must be established to counter illegal fishing. One study that investigated the relationship between local situational factors and the prevalence of illegal fishing across territorial waters of 53 countries concluded that the capacity for monitoring, control, and surveillanc (MCS) was the most effective predictor which underscore the necessity of enhancing MCS frameworks to counter illegal fishing (Petrossian, 2015). Here, monitoring refers to closely observing fishing activities using remote sensing technology such as satellite systems or onboard trackers that transmit the vessels location frequently. Control refers to implementing rules and regulations to ensure lawful fishing. This can include setting fishing boundaries and quotas for instance. And lastly, surveillance ensures compliance with fishing regulations with violations leading to penalties like fines or

license revocation (Cochrane, 2002). All of these components work together to support sustainable fishing practices. However, for the purposes of this work, focus was directed solely towards the monitoring aspect.

The EU is actively combating IUU in several ways. The cornerstone of the EU's efforts against IUU fishing is the IUU Regulation which has been in effect since 2010. This regulation establishes strict controls on the seafood entering the EU market by requiring that all imports are certified as legally caught (European Commission, 2021). Countries failing to cooperate with the EU will be issued a warning which can escalate to a ban where the country in question is not allowed to export it's fisheries products to the EU mark. Regarding fishing in EU waters, vessels longer than 12m are required to be equipped with a Vessel Monitoring System (VMS) (Thoya, Maina, Möllmann, & Schiele, 2021). This system was specifically designed to oversee fishing operations by requiring ships to broadcast their location continuously. This transmission is encrypted and only delivered to the relevant government agencies meaning public access is very limited (Thoya et al., 2021). This ensure privacy and security but hinders open research on the data and minimizes transparency for the public. An alternative to VMS is the Automatic identification system (AIS), this system was originally introduced by the International Maritime Organization (IMO) and was meant to enhance nautical safety (Thoya et al., 2021). Similar to VMS, AIS operates by requiring vessels to continually broadcast their locations and these signals are receivable by neighboring vessels, ground-based receivers and satellites. This increases navigational awareness and helps ships avoid collisions with each other. The EU mandates that all fishing vessels exceeding 15 in length shall be equipped with the AIS system (Bunwaree, 2023). AIS data is openly accessible to anyone and features high temporal resolution where ships frequently send out signals. This has meant that its application has broadened significantly beyond merely enhancing safety and made it increasingly popular for monitoring fishing activity and tracking movements across the oceans in general (Bunwaree, 2023).

Both AIS and VMS operate on a cooperative basis where the ships are in charge of broadcasting their data. This approach presents a vulnerability as ships with malicious intent can take action to prevent their location being recorded. Fortunately, the advent of satellite technology offers an alternative way to monitor non-cooperative ships. Both satellite that are equipped with optical sensors and those that carry active remote sensing instruments are important tools for the comprehensive monitoring of the world's oceans. Synthetic Aperture Radar (SAR) is a type of active remote sensing technique that has been particularly useful in this field because it can penetrate cloud cover and operate independent of the day-night cycle (Yasir et al., 2023). This makes SAR data very constant and reliable. By combining both cooperative (AIS,VMS) and non-cooperative (SAR) methods, it becomes possible to get comprehensive overview of fishing efforts.

1.1 Aim

Although AIS has proven to be a valuable tool when it comes to monitoring fishing activities, it is important to reiterate that this was not its original intended purpose and there are some complications to address. The aim of this project is to answer the following question: Can fishing operations in the southern Baltic Sea be effectively monitored using publicly available data? This project will aim to do so both with cooperative (AIS) and non-cooperative (SAR) methods. Specifically, historical AIS data spanning the year of 2018 will be used to analyze the trajectory of fishing boats. A logistic regression machine learning model will then be implemented to predict periods of fishing along a ship's path. For non-cooperative analysis, a deep learning object detection model will be developed to automatically locate ships in SAR images. This model will be applied to 10 SAR images from January 2018, and the resulting detections will be matched with the aforementioned AIS data. Doing so will give an insight into the level of AIS uptake among the fishing fleet operating in the southern Baltic Sea. Unmatched vessels are of particular interest as they indicate a lack of AIS transmission and possible illegal operation. A machine learning model is then utilized to classify all unmatched boats into fishing and non-fishing since the focus of this project is centered around fishing efforts.

The novelty of this master thesis lies in the integration of both cooperative (AIS) and non-cooperative (SAR) methods to monitor fishing activities. While AIS is commonly used for vessel tracking, applying it specifically to fishing operations in the southern Baltic Sea alongside SAR detection is novel. This dual-method approach offers a comprehensive view of fishing activities and goes beyond studies that typically rely on a single data source.

2. Background

2.1 Cooperative vessel tracking

As mentioned in the introduction, cooperative vessel tracking refers to when boats are responsible for reporting their own data to some overseeing entity or in compliance with a broader legal framework. This data is an important tool for authorities to managing maritime activities like fishing and to ensure safety out at sea. When it comes to fishing there are 2 main cooperative boat tracking systems relevant: VMS and AIS (Thoya et al., 2021). This section will provide an overview of these essential systems.

2.1.1 Vessel Monitoring System

Vessel Monitoring System (VMS) is a satellite-based monitoring system primarily used for tracking the location and movement of fishing vessels. VMS is designed to provide regulatory bodies and fisheries with data such as location, course and velocity for the vessels that adhere to them. This helps oversee fishing activities and ensures compliance with regulations which in turn protects resources and fights against IUU. In 1998, the European Commission enacted a legislation to utilize the VMS system to monitor European fishing and in 2004 it became compulsory for all fishing boats exceeding 15 meters in length to carry a VMS transmitter on board (Mills, Townsend, Jennings, Eastwood, & Houghton, 2006). Vessels equipped with VMS send out their data at equal intervals making the data very continuous and therefore beneficial for a variety of spatial temporal analysis. It is evident that VMS is an invaluable tool for monitoring fishing activity and to this day it continues to be used extensively around the world. However it does have some drawbacks, VMS signals are encrypted to ensure that only the relevant government agencies receive the data which means uninvolved 3rd parties can not follow fishing efforts (Thoya et al., 2021). Furthermore, the frequency of signals is comparatively low meaning that analysis requiring high temporal resolution becomes challenging (Thoya et al., 2021).

2.1.2 Automatic Identification System

An alternative to VMS is the Automatic Identification System (AIS). First introduced by the IMO to improve navigation and to enhance the safety of vessels out at sea by aiding in collision avoidance (Natale, Gibin, Alessandrini, Vespe, & Paulrud, 2015). Similarly to VMS, AIS is a broadcasting device installed onboard ships that transmit information relevant to a ship's voyage. VMS however is built around a point-to-point communication whereas AIS sends out the signal in all directions enabling any nearby ship, ground-based receiver, or satellite to receive the message (Natale et al., 2015). AIS signals operate primarily within a line-of-sight owing to the nature of the signal. This means that depending on the elevation of both the receiver and transmitter, AIS typically has a range of just a few nautical miles up to approximately 100 natutical miles (Natale et al., 2015). This means that AIS

can have significant blind spots in the ocean where there are no AIS receivers to log a ship movement making the data not as reliable and predictable as VMS. Additionally, AIS devices are vulnerable to being altered or even switched off if a vessel decides to partake in illegal activities like unauthorized fishing (Mazzarella, Vespe, & Santamaria, 2015). This weakness is less prevalent for the VMS. One important advantage of AIS over VMS lies in its temporal resolution. While VMS transmits a signal every 2 hours AIS does so every few seconds up to a few minutes. This renders AIS much more suitable for accurate spatiotemporal analysis (Natale et al., 2015). But the most impactful advantages that AIS possesses and the reason why it has the potential to become the key source of data for mapping fishing activity globally is its low operational costs and accessibility (Thoya et al., 2021).

The IMO has been the main driving force behind mandating the usage of AIS data and as mentioned in the introduction, carrying AIS has become compulsory within Europe for fishing vessels exceeding 15 meters in length since May 2014. A study published a year later revealed that 75% of fishing vessels complied with this regulation (Natale et al., 2015). Another more recent study which cross-referenced AIS data with SAR data over a span of 4 years found that 61% of industrial fishing vessels are publicly tracked in Europe (Paolo et al., 2024). This indicates a lack of compliance. Although some of the mismatch can be explained by expected limitations in the study, it is nevertheless apparent that the level of AIS uptake in Europe is not meeting expectations. A big reason for this is minimal enforcement and while some EU member states implement annual inspections to ensure that ships have AIS onboard, there is no comprehensive system in place to verify that these ships are actively using it (Natale et al., 2015).

Given that AIS data is inherently public and open-access in nature, numerous websites have sprung up that provide access to this information to anyone interested and several commercial services offering historical AIS data have gained prominence. This development has been a cause for concern to some as it raises some questions regarding the potential misuse and privacy law violations. In 2012 the European Data Protection Supervisor (EDPS) stated that tracking ships with AIS data becomes problematic if it can be connected to specific individuals (Natale et al., 2015). The IMO has strongly discouraged using AIS data for commercial purposes in light of potential risks to the safety and security of ships (Natale et al., 2015). Furthermore, simply equipping ships with AIS transmitters can compromise their safety. Pirates or other hostile groups can locate and target ships with the help of AIS data. IMO recognizes this threat and allows ships to switch off their AIS broadcasting in case the ship's safety and security is not guaranteed even though the ship is required to carry AIS under the International Convention for the Safety of Life at Sea (SOLAS). This exception has been exercised recently where vessels sailing in the Red Sea opt to go dark to evade potential attacks (Bartlett, 2024).

As mentioned, AIS was initially created for improving navigation safety. However, as its adoption has grown over the years, more and more data is being generated. This has made it increasingly popular for other kinds of use cases and academic researchers continue to find new ways of using it. For example, AIS has proven useful in assessing trade flows and analyzing economic trends by tracking international commercial vessels. It has helped in

environmental studies by monitoring the effects of shipping on nature (D. Yang, Wu, Wang, Jia, & Li, 2019) . By combining AIS data with information from shipping agencies and port authorities it is possible to enhance the management of ships in ports. This helps in scheduling their arrival times and reducing how long they stay docked (D. Yang et al., 2019). While there are numerous other examples to consider, focusing specifically on monitoring fishing efforts, the use of AIS data has made a significant impact. Global Fishing Watch (GFW) is a non-profit and independent organization. Their mission is to promote ocean sustainability and ensure responsible fishing practices worldwide (Taconet, Kroodsma, & Fernandes, 2019). GFW receives more than 50 million AIS messages every day from a constellation of satellites and a network of ground based AIS receivers. In one year, signals from about 300,000 vessels are processed. These boats represent a variety of vessel types, with a specific focus on identifying and monitoring fishing activities. GFW uses 2 core models. One to distinguish between fishing and non-fishing vessels as well as a gear type prediction and another to determine when a fishing boat is likely engaging in fishing (Taconet et al., 2019). For the latter, GFW developed a sophisticated convolutional neural networks (CNN) model that considers multiple types of gears for optimal performance. These models allow GFW to obtain a comprehensive global overview of fishing activities in a transparent way. On their website they give free access to an interactive global map that shows fishing activity derived from different sources.

Some limitations need to be addressed when discussing the usage of AIS data to identify fishing activity. Poor AIS reception is one issue that has already been touched on above. Due to limited cover by land based receivers and satellites, blind spots in the ocean exist where ships go undetected. Noise can cause further difficulties because of poor error checking in AIS transmission protocols and data loss during transmission. This sometimes leads to ships having unrealistic movement but it can be countered by algorithms that filter out nonsensical data (Taconet et al., 2019). Additionally, issues with segmenting and spoofing arise when the same unique vessel identifier number (MMSI) is used by multiple vessels. This as well can often be tackled by specialized algorithms and is most commonly observed in the Chinese fishing fleet (Taconet et al., 2019). Lastly, offsetting refers to ships that broadcast their location far from their actual positions. This error can be identified when ships broadcast their location to a receiver that is beyond its range (Taconet et al., 2019). In some cases this offsetting can be unintentional but in other cases vessels falsify their AIS message on purpose. For instance to avoid sanctions and monitoring by international authorities (Bergman, 2023). Overall these challenges highlight the complications in accurately tracking fishing activity.

2.2 Non-cooperative vessel tracking

Although cooperative vessel tracking is extremely useful for monitoring fishing activity it has the inherent weakness of being managed by the vessels themselves. In many cases this does not pose a problem but occasionally, if a vessel is engaged in illegal activities it could potentially switch off or alter their tracking data (Paolo et al., 2024). To address this issue, non-cooperative vessel tracking methods can be used and they do not depend on vessels to transmit their own data. Instead, they utilize some form of remote sensing technology to

detect vessels from a distance. Satellites in particular with their global coverage are well suited for this task and especially SAR satellites (Paolo et al., 2024).

Unlike the more widely recognized optical imagery, Synthetic Aperture Radar (SAR) uses active remote sensing which means it does not depend on external light sources such as the sun for illumination. Instead, SAR systems actively emit their own energy towards the Earth's surface and measure the energy that is reflected back (Earthdata, N/A). The wavelength of the signals emitted by satellites plays a huge role in determining their interaction with the Earth's surface. For example, shorter wavelength (around 1cm) experience limited penetration through vegetation and mostly capturing surface-level details. In contrast, signals with longer wavelengths (around 30cm) can penetrate deeper into vegetative layers allowing for the collection of data from beneath the canopy. By picking the right wavelength scientists can better target the specific environmental features related to their investigations making SAR a very powerful tool within remote sensing (Earthdata, N/A). Another aspect of SAR is its capability to operate with signals in different polarizations. SAR satellites are equipped to both emit and capture signals in either horizontal (H) or vertical (V) orientations allowing for 4 possible combination HH, HV, VV and VH where the first letter denotes the orientation of the transmitted signal and the second one denotes the orientation of the received signal. Different polarizations lead to distinct scattering behaviors from objects on the ground which enhance its usability and gives scientists even more tools to tailor their analyses (Earthdata, N/A). In the context of ship detection, the choice of polarization impacts the detection performance. It has been found that the HH channel provides the highest detection probability for single polarization but utilizing dual polarization proved to be the most successful (Mahgoun, Chaffa, Ouarzeddine, & Souissi, 2020).

2.2.1 Ship detection with SAR

SAR has proven to be a highly effective method for monitoring maritime activities owing to the fact that SAR images can be taken regardless of weather conditions and the day-night cycle making them highly consistent (Yasir et al., 2023). Furthermore, SAR images maintain consistent resolution regardless of the distance from the targets being observed (Wang, Wang, Zhang, Dong, & Wei, 2019). This makes SAR images well suited for ship detection and with and with the recent surge in the availability of free high-resolution SAR data, development in this field has grown fast. Ships at sea are easily identifiable as they reflect satellite signals far more strongly than the surrounding ocean causing them to prominently stand out in SAR imagery. This is because of ships being solid metal objects with completely different scatter properties than water. Figure 2.1 illustrates the appearance of ships on a standard SAR image. It can clearly be seen that ships appear bright compared to the surrounding water, making them easy to find.

Ever since the United States deployed the first SAR satellite in 1978, numerous methods for detecting ships have been introduced (Zhang et al., 2021). One might assume that systematically detecting ships out at sea from SAR images would be a straightforward process given their distinct appearance against the oceanic background but in fact it presents quite a challenge. In figure 2.2 an overview of common methods is presented.

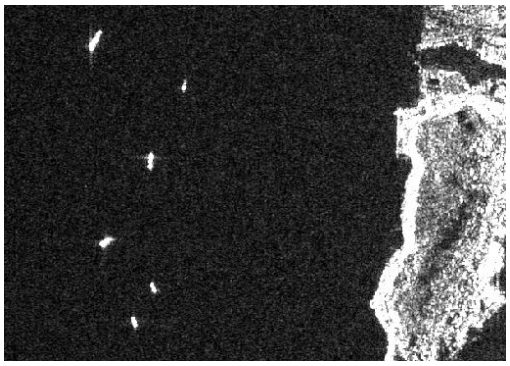

Figure 2.1: An example of how ships appear on SAR images with six vessels visible on the left and landmass situated on the right. The image was taken on April 2024 and it was sourced from a Sentinel-1 (IW swath, Level 1 GRD processed image)

Traditional techniques for separating ship targets from the background relied on artificially creating features (Yasir et al., 2023). These features are crafted based on the characteristics of ships compared to the sea surface. In other words, how ships reflect the signal of the SAR satellite compared to the ocean. Constant false alarm rate (CFAR) is an algorithm that has frequently been used for this task and it simply classifies pixels into two categories: ships or non-ships. It does this by statistically modeling the noise in the image and then it establishes a threshold value, pixels above that threshold will be classified as ships while those below are identified as non-ships. The final step involves amalgamating the pixels that together represent the same ship (Li, Xu, Su, Gao, & Wang, 2022). Variety of other algorithms with the same objective have been developed over the years. These include the Generalized-Likelihood Ratio Test (GLRT) detector, the Spatially Enhanced Pixel Descriptor (SEPD) and more (Yasir et al., 2023). These traditional methods identify ships by employing some statistical model or pre-picked features that are designed to be adaptable across a broad range of images and in scenarios with straightforward scenes they often preform well. However, adaptability over different environments is limited and difficulties associated with accurately modeling noise limit their performance. This often results in missing targets especially in complex scenes (Yasir et al., 2023).

Ever since deep learning (DL) was introduced around 2012 it has increasingly been adopted in the field of ship detection from SAR imagery with significant success (Li et al., 2022). In fact, DL outperforms the traditional methods in both speed and accuracy (Li et al., 2022). In recent years DL in ship detection has been gaining traction and it continues to be a popular topic for researchers. One study that was comparing the performance of the conventional CFAR method with DL models discovered that the traditional approach achieved an average precision of 27.1%, while DL models had around 80% precision (Sun et al., 2019).

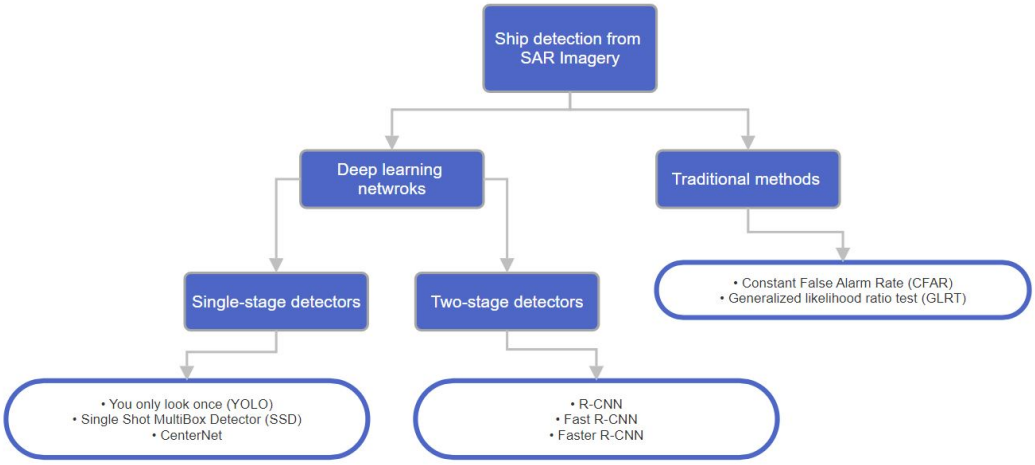

Figure 2.2: Well known methods for detecting ships out at sea from SAR (Li et al., 2022; Yasir et al., 2023; Zhang et al., 2021)

Object detection leveraging deep learning can be broadly divided into two main categories: Two-stage detectors and one-stage detectors (Yasir et al., 2023). The former is designed to prioritize speed and efficiency while sacrificing accuracy to some degree. They rely entirely on convolutional layers to classify and regress anchor boxes in a single step to obtain detection results which are a set of bounding boxes outlining the object within a given image (Yasir et al., 2023). Popular one-stage detectors include: You Only Look Once (YOLO), Single Shot MultiBox Detector (SSD). Two-stage detectors on the other hand begin by using convolutional neural networks (CNN) to identify regions of interest within a given image. Then the same logic from the one-stage detectors (classify and regress anchor boxes to get object detection) is applied to all the previously designated areas of interest. This approach yields higher accuracy but is slower and more computationally heavy (Yasir et al., 2023). Region-based Convolutional Neural Networks (R-CNN) and Fast R-CNN are well known examples of Two-Stage Detectors. It works by generating region proposals and then classifying each region using a CNN. Fast R-CNN is an improved version that speeds up the process by sharing computation for overlapping regions.

In recent years transformers have become very prominent deep learning models and shown incredible performance in large language models (LLM). Transformers have also shown promising results in the world of object detection. In recent years, the research team at Meta published the DEtection TRansformer (DETR) model, a relatively simple architecture compared to many other models that uses the transformer encoder-decoder architecture after an initial CNN backbone (Carion et al., 2020). This model showed promising results when compared to other classical object detection models like Faster R-CNN. Specifically for boat detection on SAR images, one study demonstrated that their transformer-based method outperformed traditional CNN-based methods on 2 SAR ship datasets (Shi, Chai, Wang, & Chen, 2022).

However, there exists another type of single-stage detectors known as anchor-free detectors. Unlike other methods, these detectors do not depend on predefined anchor boxes to determine object locations but instead identify objects as a collection of key points (Li et al., 2022). This method simplifies the detection pipeline and reduces the computational complexity. Anchor-free detectors are well suited to address several crucial challenges that other DL models struggle with such as detecting small size objects and sparse distribution of objects (Li et al., 2022). Boats frequently appear small on SAR images since the spatial resolution typically spans tens of meters meaning that ships are present in only a handful of pixels and smaller vessels may not be detected at all. Additionally, given the vastness of the ocean, boats often appear far apart. This makes anchor-free DL models prove particularly effective. In fact, it was an anchor-free DL model that won the xView3 competition (model is available here). The aim of this competition was to create a tool to efficiently detect IUU on SAR images and it was organized by the Defense Innovation Unit and Global Fishing Watch with support from some agencies including the U.S. Coast Guard and NOAA (*xView3 Competition on Detecting Dark Vessels to Combat Illegal Fishing*, 2023).

Despite the success of DL models in distinguishing ships from SAR images there is still room for improvement and complications to overcome. Distinguishing ships from other sea-based objects like icebergs and windmills has proven difficult because of similar scattering properties these objects share with ships (Li et al., 2022). Another significant challenge to resolve is the lack of large SAR datasets for training DL models and this scarcity makes it hard to build a model from the ground up (H. Yang, Kang, Liu, Liu, & Huang, 2023). In the field of object detection, it is common practice to leverage the power of transfer learning and especially when training data is in short supply. Transfer learning refers to the process of taking a pre-trained model that has been trained on a large, diverse dataset and tailoring it to a more specific task. This pre-trained model has learned basic pattern recognition and feature extraction that can be repurposed in a quick way. The problem is that these pre-trained models are trained on optical images which are inherently different from SAR images and this difference has resulted in limited success when fine-tuning with SAR datasets (H. Yang et al., 2023). Making the SAR images look more like standard optical images would help bridge this gap and potentially achieve a better result when using transfer learning. One way is to generate some kind of a pseudo-color image where the SAR data is assigned to the conventional RGB channels found in standard optical images. For example, the different polarization states of a SAR data can be mapped to RGB channels as follows:

$$Red = VV \qquad Green = VV/VH \qquad Blue = VH$$

With appropriate normalization a colored image can be obtained. Admittedly, this is a simple example and would not significantly alter the overall appearance of the image. Other more sophisticated techniques to make SAR images appear more life like exist but they are often computationally heavy and generally not used in the context of object detection to the to the authors knowledge.

2.2.2 Matching SAR with AIS for monitoring fishing

Monitoring fishing activity with just SAR data is challenging because of the large temporal resolution of satellite systems. Instead of tracking a boat's path, SAR only provides a single snapshot of the sea. This as well as the limited number of features extractable from a ship in a SAR image complicates its standalone use for monitoring fishing activity. However, integrating SAR data with AIS data has proven to be useful and is currently a key focus for many researchers. One study (Kurekin et al., 2019) focusing on combating IUU along with its associated environmental and economic impacts found that this integration was very advantageous. This study used the Search for Unidentified Maritime Objects (SUMO) tool which is an open source project developed by the European Commission's Joint Research Centre (JRC) that uses CFAR to identify ships in SAR images (Kurekin et al., 2019). Then to match the detections with corresponding AIS data, the authors implemented a technique called Munkres assignment algorithm that interpolates AIS data based on historical patterns found in AIS to enhance its alignment with the SAR data (Kurekin et al., 2019). Finally, the authors designed the workflow to be automatic and continuously monitoring Ghana's oceanic regions. Suspicious vessels get flagged and relevant stakeholders informed by the system. Preliminary results indicate that over 75% of detected boats got classified as non-cooperative which highlights the scale of the issue (Kurekin et al., 2019). Another study (Galdelli, Mancini, Ferrà, & Tassetti, 2021) investigating the matching of SAR-AIS data for the purposes of maritime montering specifically within the Adriatic Sea found that AIS-SAR matching could effectively be used to spot suspicious behaviors around protected areas. The researchers also utilized the SUMO software to handle the ship detections and then matched them with AIS data either with a straightforward point-to-point match or a more sophisticated point to line association which was designed to track vessels that might not be broadcasting their AIS data intentionally or unintentionally (Galdelli et al., 2021). This is particularly valuable for investigating suspicious or non-compliant activities close to regulated marine areas. Furthermore, the study utilized a Fast Fourier Transform (FFT) to analyze the position and course data which in turn significantly improved the feature extraction process from a SAR image. These extracted features were then put into a machine-learning model designed to classify various types of vessel trips (Galdelli et al., 2021).

One study (Snapir et al., 2019) employed SAR data from the North Sea alongside AIS data to monitor fishing activity. Firstly, the researchers used the AIS data to train a random forest (RF) model to classify ships identified on SAR images as either fishing or non-fishing vessels. The features used for the RF model were the ship's coordinates, its length, distance to shore and time of day. With an overall classification accuracy of 91% the authors showed that a RF can offer an effective way to distinguish between fishing and non-fishing vessels (Snapir et al., 2019). The study then goes on to use SUMO to detect ships across a vast number of SAR images which in turn allowed for comprehensive spatial and temporal analysis of fishing activities across the North Sea. Even though this study did not directly match SAR and AIS data, it showed that with a pre-trained machine learning model it is quite doable to assess fishing efforts solely with SAR data.

3. Data and methodology

3.1 Study area and data

The study area for this project is depicted in figure 3.1 which covers a sizeable part of the southern Baltic Sea. This location was chosen because of data availability and proximity to Sweden. Additionally, this region is of interest for a plethora of reasons. Firstly, it hosts a diverse range of marine species making it a prime location for fishing activities (Ojaveer et al., 2010). It includes vital reproduction zones for cod, a highly sought-after fish known for its commercial value. Vessels from multiple nations operate in these waters adding a complex layer of regulatory and economic interactions (Ojaveer et al., 2010). Furthermore, the study area receives intensive amounts of traffic from boats that are simply traversing the area, sailing between the inner Baltic Sea and the North Sea. This traffic may cause ecological disturbances and limit fishing opportunities.

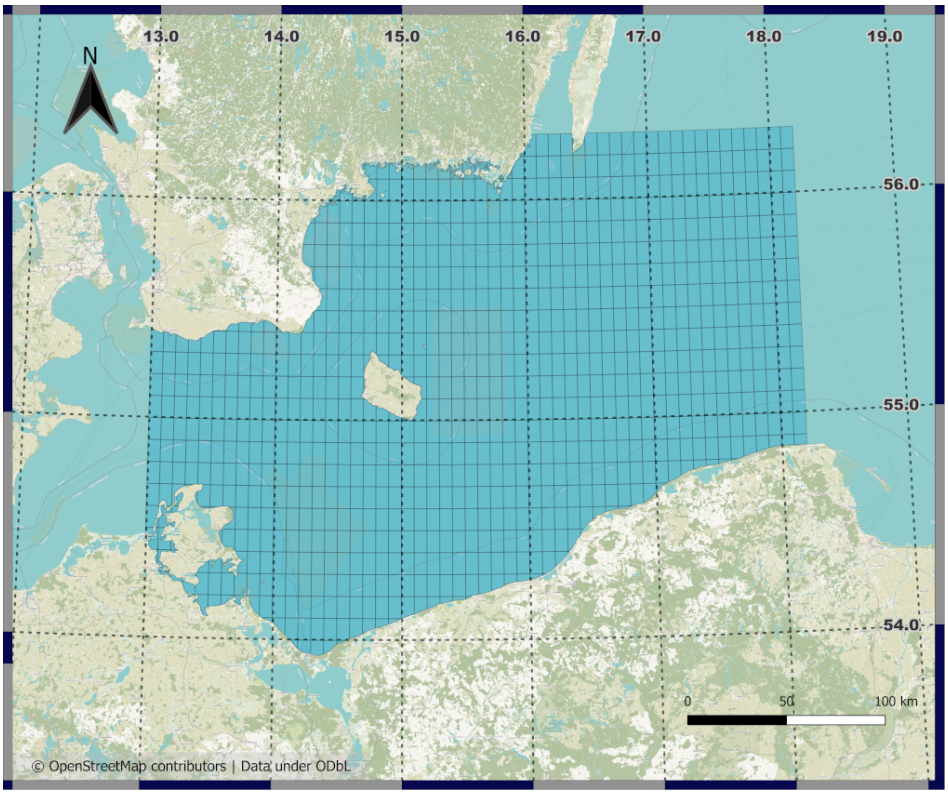

Figure 3.1: A map of the study area

This project leveraged two main types of data: SAR and AIS. More information regarding the data will be provided in the bullet points below. Additionally, two other data sources were utilized in this project. The first is a detailed polygon representation of the ocean, obtained from OpenStreetMap (OSM), which was used to define the study area. The second

is a dataset for training and evaluating the deep learning model which will be discussed in detail in section 3.4.

- **SAR** data was obtained via the Copernicus Browser which is part of the Copernicus program that aims to provide continuous, independent, and reliable access to earth observation data. This initiative is led by the European Commission in partnership with the European Space Agency (ESA) and they provide their data entirely free of charge allowing anyone access for analysis and visualization (*Copernicus Programme*, 2024). There are many satellite constellations under the Copernicus umbrella but the Sentinel-1 mission is the one producing SAR data and therefore the most relevant for this work. Sentinel-1 currently had 2 operational satellites: Sentinel-1A and Sentinel-1B. They work together in tandem to provide a comprehensive coverage of the Earth's surface and continuously produce about 6 TB of data every day (Snapir et al., 2019). Table 3.1 in conjunction with figure 3.2 provides a list of all the SAR images downloaded from the Copernicus Browser. One thing worth mentioning is that all the SAR images downloaded are classified as Single Look Complex (SLC) which indicates they are not raw satellite data but instead have bee processed to account for sensor geometry and earth curvature. Importantly, unlike other common processing levels, SLC retains both amplitude and phase information of the signal which allows for a better pseudo-color image to be generated (Hu, Li, & Pan, 2021).

- **AIS** AIS data was sourced through the Danish Maritime Authority. They oversee maritime activities in Denmark and provide a plethora of useful information on navigation, safety regulations, and environmental protection. They also manage and provide access to a comprehensive and continuously updated repository of AIS data that goes as far back as 2006. This data is raw/large and covers all Danish territorial waters and more. For the purposes of this project, AIS data spanning the entire year of 2018 was downloaded.

Table 3.1: SAR images processed and used for this project. All images have VV+VH polarisation, IW acquisition mode, and SLC is the product type

	Area	**Sensing Time**
1	B	2018-01-04 16:43:39
2	A	2018-01-04 16:44:04
3	B	2018-01-10 16:44:21
4	A	2018-01-10 16:44:46
5	B	2018-01-16 16:43:39
6	A	2018-01-16 16:44:04
7	B	2018-01-22 16:44:21
8	A	2018-01-22 16:44:45
9	B	2018-01-28 16:43:38
10	A	2018-01-28 16:44:03

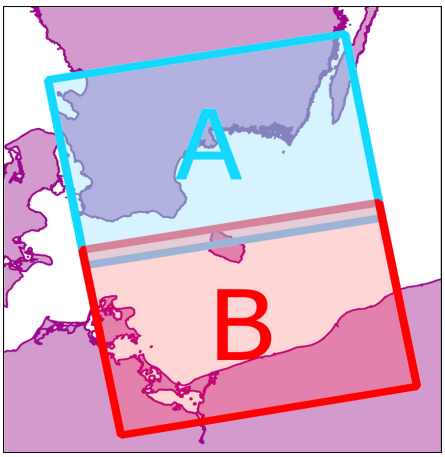

Figure 3.2: Coverage of the SAR images processed (see table 3.1)

3.2 Filter data by study area

Filtering data by the study area was handled with a function called: IsWithinOcean. This function plays a pivotal role in both AIS and SAR analyses. Written in Python, it takes in a point (latitude, longitude) and outputs whether or not it falls within the project's study area. It works by loading in a vector file that delineates the study area and then by using inbuilt functionality from the *Shapely* library it can return whether or not a given point is contained within the ocean polygon. This ocean polygon was obtained from OpenStreetMap here. To prepare the data for this project's analytical needs, *QGIS* was utilized and the processing steps taken there are shown in figure 3.3. The Ocean vector data cover the whole earth so the initial step is to clip the data to the confines of the study area. Secondly, the data is shrunken by 200 meters with the buffer tool since the focus is exclusively on ships out at sea. This allows for distinguishing between docked ships and ships sailing in the study area, a useful feature as will become apparent in later sections. The final step is to subdivide the polygon into multiple discrete cells. This IsWithinOcean function needs to be called extremely often (order of millions) so optimizing it for speed becomes very beneficial. One way to do so is by constructing these cells, the algorithm can than swiftly identify which cell a point is contained in and then run the *Shapely* toolbox on that targeted cell. This is much quicker then passing the entire polygon through *Shapely* every time the function is called since the speed of a point-in-polygon is linearly related to the number of vertices in a polygon. The final polygon representing the ocean study area is shown in figure 3.1.

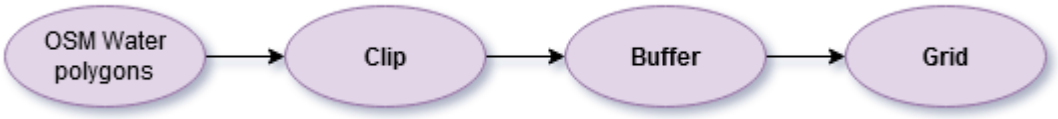

Figure 3.3: The processing steps taken in *QGIS* to filter data for the study area

3.3 Vessel trajectory from AIS

This section explores the capabilities of AIS data for monitoring fishing activities. In particular it discusses a script written in Python that takes in a ship's MMSI number (ID) as input and outputs all the found trajectories transmitted by the vessel within a specified time window. It does this by connecting individual AIS messages together based on a time threshold and then extracting useful information from that path. This section will also explore techniques for determining periods of fishing activity along these paths which in turn allows for the creation of heat maps that provide a clear visualization of the primary fishing grounds.

3.3.1 AIS data pre-processing

The AIS dataset provided by the Danish Maritime Authority consists of multiple large CSV files. Each row in this file stores an unique AIS transmission which has a plethora of data

Table 3.2: An overview of the key Information contained in an AIS Signal

Attribute	**Description**
Timestamp	The exact date and time the AIS message was recorded.
MMSI	A unique identifier assigned to each AIS-transmitting vessel.
Latitude	Coordinate of the ship
Longitude	Coordinate of the ship
Navigational status	The operational condition of the vessel, such as 'Engaged in fishing' or 'Under way using engine'.
COG	The direction in which the vessel is currently moving relative to the ground, measured in degrees.
SOG	The vessel's speed relative to the ground, typically measured in knots.
Ship type	The classification of the vessel according to its AIS signal, such as 'Fishing'
Length	The overall length of the vessel.

associated with it and the key ones (the ones used in this project) are highlighted in Table 3.2. This dataset encompasses all Danish territorial waters which means it needs to be filtered for only the study area. For that the *IsWithinOcean* function (detailed in section 3.2) was utilized which also has the added benefit of removing signals from vessels docked at port and only focusing on those ones that are actively navigating the ocean. B-class AIS signals were also filtered out because they originate from boats not required to carry AIS transmitters and are less reliable. The CSV file is ordered chronologically which is suitable for some applications but less so when extracting specific data like a specific boat by MMSI number. In those cases one would always have to increment through the whole file to extract the needed data. A more efficient solution would be to store the data in a database which enables quicker data retrieval through writing queries. This solution would also significantly decrease the storage space needed by eliminating redundancies. An attempt was made to create a relational database in SQLite which included a table for static boat data (e.g., name and length) and a dynamic table for navigational updates. Ultimately, transferring the millions of lines of data from the CSV file to the SQLite database proved too time-consuming so the idea was abandoned. Although reading the CSV files was somewhat slow, it was sufficiently efficient for the purposes of this project.

3.3.2 AIS data analysis

This section provides a more detailed information of the vessel trajectory analysis script introduced in the beginning of section 3.3. Given some arbitrary segment of AIS data, the

scripts connects individual ship transmissions together and if no new messages are received from a ship over a period of 6 hours the trajectory is considered to be finished. Once a ship's trajectory is fully established the script calculates the distance to land from both the starting and ending points with the help of the *shapely* library in python. This is done because one would expect all vessel routes to begin and conclude at ports but in practice that is not the case. Ship trajectories often start far from any land and computing these metrics helps understanding how frequent it is for ships not to broadcast their whole journey. Occasionally it was noticed that information on ship speed and course contained within a AIS message was missing so the script manually calculates that from the coordinates of the boat. That also enables a comparison between the transmitted and calculated values for speed and course.

Information on when along its path a boat is engaged in fishing, is a key element in monitoring. Knowledge about where and for how long a vessel is fishing can help regulators ensure sustainable fishing practices. AIS systems include navigational status in the signal sent out so vessels can disclose whether they are actively fishing or not. However, some exploring of the data revealed that most ships do not update their navigational status frequently enough or effectively use this feature. So in order to get a better picture of when a ship is fishing another method is needed. One common approach is to simply filter out all ships traveling faster than 4 knots. Including this threshold removes all the boats that are merely transiting from point A to point B and keeping those that are fishing since boats tend to go slower when fishing. Another more sophisticated approach that was be used in this project is to train a machine learning model to predict when along its trajectory a boat is fishing. The Global Fishing Watch (GFW) has published several such models on GitHub (Möller & Hochberg, 2020). The features used by those models to run predictions is average speed and deviation in both speed and course over a time window. The models were trained on a hand-labeled dataset containing 29 unique boats equipped with four distinct types of gear. The particular model used in this project is the generic logistic regression model, which has a 12-hour time window to calculate these features. This model can handle non-linear relationships well (Möller & Hochberg, 2020). The model was evaluated with a few different gear types and it arguably did the best with trawlers where it had a precision of 93% and a recall of 91% (Möller & Hochberg, 2020). The final output of the script once it has fully processed a single ship trajectory are the following 6 files:

- A GEOJSON file that contains the trajectory of the ship, including computed metrics such as journey duration and distance from land.
- A map highlighting the trajectory with marked starting and ending points of the path.
- A figure that shows the speed of the vessel as a function of time. Both the speed contained in the AIS message and the calculated speed. The background in the figure shows the navigational status.
- A figure that shows the course of the vessel as a function of time. Both the course contained in the AIS message and the calculated course.
- A figure that displays the output of the logistic model, depicting the predicted fishing score as a function of time.

➢ A map of the logistic model displaying the fishing probability along the ship's route.

The vessel trajectory script focuses solely on the paths taken by individual fishing boats. While that provides a very detailed view of fishing activities it was decided to also capture a broader perspective by conducting both spatial and temporal analysis for the entire study area. This will involve creating heatmaps and visualizing the temporal distribution of fishing activity throughout the year 2018.

3.4 SAR ship detection tool

This section outlines the inner-functionality of a Python-based script that was created for the purposes of this project. It takes in a processed SAR image and outputs a list of detected ships. This script utilizes a state of the art deep learning object detection model. More specifically, it uses a CenterNet architecture with a ResNet-50 backbone. Details on how this model was trained and tailored to SAR images will be presented as well as general information regarding other steps involved in the detection process.

As was discussed in section 2.2.1 the availability of datasets with a large number of annotated ships in SAR imagery is limited but some do exist. One noteworthy dataset is called DSSDD which was created by Yuxin Hu, Yini Li, and Zongxu Pan (Hu et al., 2021) and it will be used for this project. DSSDD contains in total 50 dual-polarimetric Sentinel-1 SLC images which have undergone particular processing. This processing results in the dataset being deviated into a total of 1236 distinct images each with dimensions of 256x256 pixels and they are provided as 16-bit TIFF files that contain the covariance matrix of the signal (Hu et al., 2021). Utilizing the covariance matrix enables the full exploitation of dual-polarimetric SAR data and offers a convenient way to merge complex data into a more manageable format making visualization of the data more easy. The covariance matrix equation can be seen below (S_{VV} and S_{VH} denote the complex scattering coefficients for VV and VH polarizations, respectively) (Hu et al., 2021).

$$C_2 = \begin{bmatrix} \langle |S_{VV}|^2 \rangle & \langle S_{VV} S_{VH}^* \rangle \\ \langle S_{VH} S_{VV}^* \rangle & \langle |S_{VH}|^2 \rangle \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix}$$

Table 3.3: The equations used to convert the 16-bit covariance matrix values to 8-bit RGB channels

Red:	$C_{11} \cdot \frac{256}{0.4}$
Green:	$\lvert C_{12} \rvert \cdot \frac{256}{0.1}$
Blue:	$C_{22} \cdot \frac{256}{0.04}$

In this equation, C_{11} and C_{22} denote the energy amounts in the 2 polarizations while C_{12} and C_{21} are in fact conjugated complex numbers which means they will always possess equal magnitudes. Consequently, the covariance matrix offers 3 unique values that conveniently can be assigned to the 3 RGB channels for a pseudo-color image (Hu et al., 2021). For that there are numerous approaches possible and the author explored several options but ultimately, a simple linear scaling of C_{11}, C_{22} and C_{12} was employed. The exact equations are visible in table 3.3 and they show how the 16-bit values are scaled to the more standard 8-bit range for each color channel. The somewhat arbitrary constants in the equations are obtained after extensive testing and it was concluded by the author that these constants yield the most visually pleasing images. Although that is somewhat subjective the goal was to have a pseudo-color image that highlighted well the boats in the ocean and made sure the background noise did not overwhelm. After converting all the provided 256x256 TIFF images containing the covariance matrix to the 8-bit RGB format, the dataset was primed for training the deep learning model. It is worth pointing out that this conversion to 8-bit pseudo-color is necessary to utilize the power of transfer learning. The deep learning model deployed in this project was pre-trained on an extensive dataset of optical images and matching the SAR images to that format is crucial. A handful of converted SAR images from the DSSDD dataset are shown in figure 3.5. For comparison, an unprocessed SAR image is presented in figure 3.4. It is important to note, however, that it is impossible to fully display such an image because it contains both phase and amplitude data for each polarization totaling four channels that cannot be visualized all togather. Instead, figure 3.4 shows the amplitude for the VV polarization in grayscale. Possessing a deep learning model trained on the DSSDD dataset is one thing but the objective here lies in creating a versatile tool capable of processing new data. That means creating a streamlined process that can seamlessly intake any given SAR data (conforming to the same format as specified in Table 3.1) and output a processed 8-bit pseudo-color image. It is thus crucial to follow the same processing steps outlined in the DSSDD documentation (Hu et al., 2021) and then further performing the same pseudo-color conversion as talked about above. This ensures consistency in the appearance of new SAR images compared with those the model was trained on and that is a critical factor for achieving optimal performance out of the deep learning network. To clearly state the steps needed to use this ship detection tool the following bullet points are given:

1. Download a SAR image from the Copernicus Browser (ensuring it matches the types specified in Table 3.1)

2. Utilize SNAP software to convert the SAR image into an RGB pseudo-color TIFF file (section 3.4.1).

3. Run the ship detection script using the processed TIFF file (section 3.4.3).

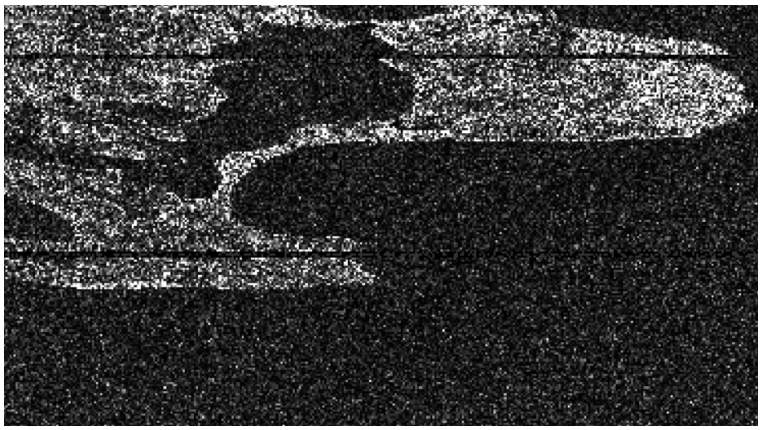

Figure 3.4: A small segment of an unprocessed SAR image

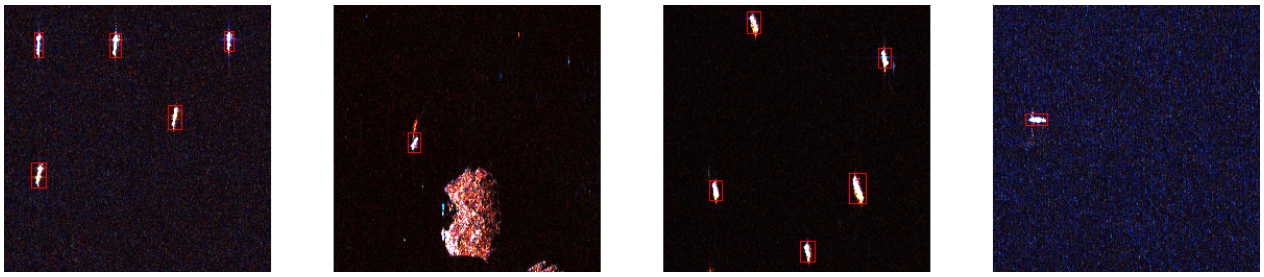

Figure 3.5: An example of how images from DSSDD appear when converted into 8-bit pseudo-color images (ships are labeled in red)

3.4.1 SAR image processing

The Sentinel Application Platform (SNAP) software will be used to preform all the processing steps in a streamlined workflow. SNAP is an open-source platform developed by the European Space Agency (ESA) to analyze, process and visualize remote sensing data. It is especially well equipped to handle data from ESA's Sentinel missions such as the Sentinel-1 mission.

As stated in section 3.1 the SAR images downloaded have undergone initial processing but they need further processing to ensure usability with the ship detection tool. SLC images are generally considered to be in a raw format as they have complex cell values, substantial file sizes and large amount of noise. Therefore it is important to implement extra processing steps. Figure 3.6 shows a flowchart illustrating the processing sequence needed where the input is a SAR image and the output is an RGB pseudo-color image exported as a TIFF file which is ready for subsequent analysis using the object detection tool. The remaining text in the section will proved a short description about each processing step shown in figure 3.6 (all

the information will be taken from the software documentation and it can also be accessed here). Lastly, a fully processed image has the dimensions of approximately 35933 pixels in width by 16283 pixels in height and a spatial resolution of about 10x10 meters. An example is shown in figure 3.7.

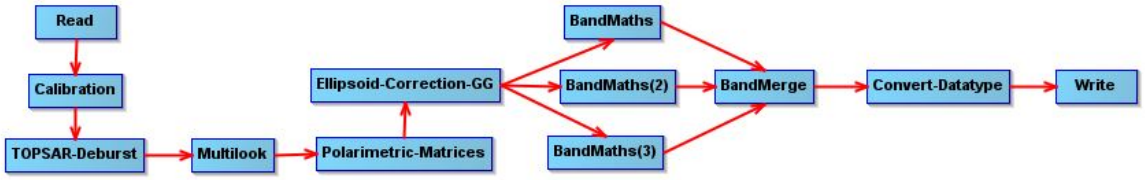

Figure 3.6: Flowchart of the SAR processing steps. (screenshot from SNAP)

- **Calibration:** This step ensures that the pixel values in a SAR image are comparable across different SAR images. It converts the raw data into meaningful quantities by analyzing the radar backscattering and performing corrections.
- **TOPSAR-Deburst:** Here the SAR image is simplified and combined into a unified images. The 3 SAR bursts are added together to ensure seamless coverage of the whole area both in the range and azimuth directions.
- **Multilook:** To minimize this speckled effect found in SAR images, the Multilook process averages multiple looks (independent observations) of the same scene. This process reduces the spatial resolution but simultaneously it enhances the clarity of the image.
- **Polarimetric-Matrices:** This step calculates the covariance matrix of each pixel using the equation shown in section 3.4.
- **Ellipsoid-Correction-GG:** Here the image is georeferenced by using an ellipsoidal model of the Earth.
- **BandMaths:** The aim of this step is to scale the 16-bit pixel values to a more manageable 8-bit range. The equations from table 3.3 are used to normalize the values so they can be better assigned to the RGB color channels.
- **BandMerge:** This step is simply merging the 3 color bands (red, grenn and blue) togather to a single raster.
- **Convert-Datatype:** The last step is to actually transform the raster data type from 16-bit to 8-bit. The 8-bit values have already been calculated in a previous step.

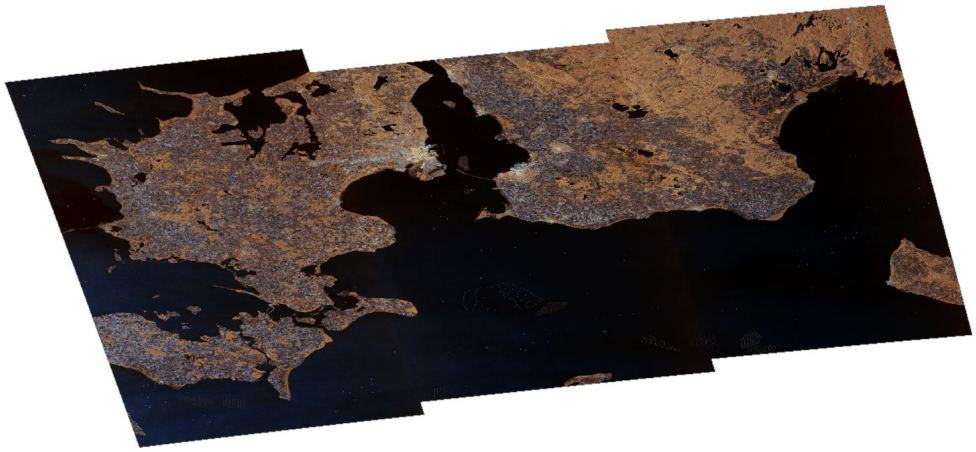

Figure 3.7: Fully processed SAR image of Sjaelland, Denmark, and Skane, Sweden (centered at approximately 55.98°N, 12.37°E)

3.4.2 Ship detection model training

For training and working with deep learning model TensorFlow will be used. TensorFlow is an open-source machine learning (ML) environment developed by Google that allows its users to create deep learning algorithms. It offers extensive libraries and tools for tasks such as classification, regression, prediction, language models, image recognition to name a few. In the context of this project, object detection is the most relevant and TensorFlow is well equipped for such tasks. The TensorFlow Object Detection API is an excellent framework built on top of TensorFlow specifically designed to create, develop, train, and use object detection models. It contains a plethora of pre-trained models that can be used for a wide range of applications. The API facilitates ease of use and simplifies the steps involved in fine tuning a model making it more accessible to researchers. Pre-trained models provided by the API can be found at TensorFlow's official GitHub page under the name TensorFlow 2 Detection Model Zoo. A handful of optimized models can be found here all striking some balance between speed and accuracy. All models have in common being pre-trained on the COCO 2017 dataset which is a database containing about 330 thousand optical images with 80 objects types.

From the previously mentioned TensorFlow 2 Detection Model Zoo, the decision was made to fine-tune the *CenterNet Resnet50 V2 512x512* model. This models is classified as an anchor-free model and as was discussed in section 2.2.1 they are particularly effective when it come to ship detection in SAR images. ResNet50, a variant of the Residual Network architecture has the capability to train an extremely deep neural network. Those networks often suffer from either gradient vanishing or gradient explosion which refers to the weights of the network fluctuating or failing to converge in the training process (Wang, Zhao, He, Zhu, & Wei, 2021). ResNet circumvents this by implementing skip connections which allow gradients to jump between layers of the network. This helps minimize the negative effects seen in deep neural networks and allows ResNet to maintain a high level of accuracy making it an excellent backbone for object detection models that need to handle complex feature extraction in

images (Wang et al., 2021). These feature extractions are then fed into CenterNet which in turn is responsible for the actual predictions. Unlike other object detection, CenterNet does not represent objects as bounding boxes but rather it uses keypoints (Wang et al., 2021). This approach simplifies the detection process and gets rid of post-processing steps making it faster than many other models. Lastly, to predict object dimensions, CenterNet employs a straightforward regression (Wang et al., 2021). Ultimately, *CenterNet Resnet50 V2 512x512* is an object detection model that balances a slight preference for speed over accuracy making it ideal for ship detection. Speed is important due to the vastness of the ocean where large areas are devoid of ships which in turn means more data needs to be processed.

When fine-tuning the *CenterNet Resnet50 V2 512x512* model, configuring the training process is a crucial first step. The batch size was set to 16 and the number of training steps was set to 2000. Since the model is exclusively detecting ships the number of classes was reduced to 1. The Adam optimizer was used and appropriate loss functions where assigned. To compensate for the limited training data, various data augmentation techniques were implemented like random cropping, padding and flipping allowing the model to better generalize with the data it has. Color based augmentation was avoided since objects on SAR images have relatively predictable appearance unlike optical images where things such as lighting can greatly alter the appearance of a given object in an image. As outlined earlier, the DSSDD dataset was utilized to fine-tune the model. About 70% of it was used for training while 30% was for testing. A python script was written to transform the *.png* image files to a tfrecord format. A common practice when training with TensorFlow that coverts the images to a binary form which increases the efficiency of data processing. An important thing to point out here is that the model requires an input size of 512x512 pixels while the images in DSSDD have a resolution of 256x256 pixels. This discrepancy calls for upscaling and since the DSSDD images are exactly half the required dimensions a rapid upscaling algorithm that simply duplicates the pixels in the horizontal and vertical directions can be used. This has the added benefit of enlarging the appearance of ships which is advantageous because it addresses the issue discussed in section 2.2.1. Some ships span only a few pixels in a SAR image making it hard for object detection models to spot them.

The training process itself took about24 hours on an older laptop model with an Intel Core i5-7200U CPU (2.50GHz) and 8GB of RAM. Figure 3.8 shows the total loss as a function of training steps. Total loss is the summation of 3 separate loss indicators which are loss for box offset (how well the model predicts the position of the bounding box), loss for box scale (how well the model predicts the size of bounding boxes) and loss for object center (how well the model predicts the central point of bounding boxes). Combined they give a comprehensive indication to the models performance. By revisiting Figure 3.8, it can be seen that there is a sharp decrease in total loss during the initial few hundred training steps but then it seems to stabilize and fluctuates around a value of 0.9. This suggests that the model is learning fast in the beginning but at a certain point it gets stuck and further training would most centrally be in vain. The training process could have been halted after approximately 400 steps but the model was allowed to complete the full 2000 steps as originally planned. A value of 0.9 is alarmingly large and means that the fine-tuning process did not yield the anticipated success. The reason for this suboptimal performance is not clear to the author

Figure 3.8: Total loss from the fine tuning process of the deep learning model as a function of training steps

and experiments with other models as well as different training configuration did not result in a significantly different outcome. For the purposes of this project it was decided to move forward with this model even though it is imperfect.

3.4.3 Output from the Ship Detection Tool

The deep learning model was designed to only handle images with the size of 256x256 pixels. However, the SAR images are significantly larger (around 36,000x17,000 pixels). To handle this mismatch a Python script was written that fully loads the SAR image with the *cv2* library and methodically loops across it in a grid pattern predicting a window of 256x256 as it goes. To go between local pixel coordinates and the actual geographical coordinates, the *rasterio* library is used to load in the georeferenced SAR file and then extract the inherent affine transformation which in turn can be applied to any point. One challenge that was encountered is that a ship has the chance of being divided between two adjacent windows and that causes the model to miss the ship detection. To mitigate this, the script jumps by 246 pixels instead of the full 256 pixels and thus introduces a 10 pixel overlap between windows to ensure no ship is missed. Looping over every possible 256x256 window is time consuming and unnecessary since most of the SAR image extends beyond the study area. The IsWithinOcean function (described in section 3.2) was used to exclude all windows that do not fall within the study area which significantly enhances the script's efficiency.

The tool's output is a comprehensive CSV file listing all detected ships along with the models prediction confidence, estimated ship size and coordinates. An optional feature of the script is to save all windows containing a ship as PNG files which is helpful for visually examining the detected ships. As an example, figure 3.9 zooms in on the ocean between the Danish island of Bornholm and mainland Sweden and it shows detected ships by the model. Although the model is quite proficient in detecting ships it is not flawless and it may not spot every single one. Such is the case on figure 3.9 where the arrow is pointing to a ship

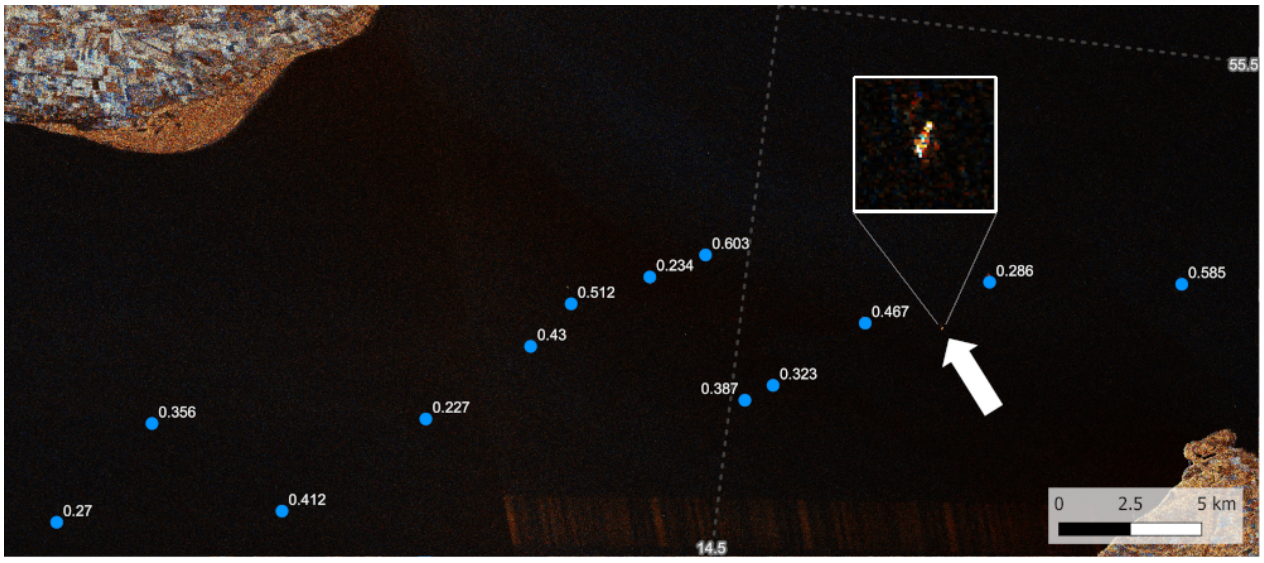

Figure 3.9: Predicted ships on a fully processed SAR image. The labels indicate predicted confidence by the model and the arrow points to a missed detection

that the model overlooked. The ship is difficult to detect due to its small size relative to the vast expanse of the ocean. However, upon zooming in, the details become clear allowing the author to identify this particular ship.

3.5 Random forest classifier for unmatched vessels

Ships detected in SAR images have limited distinguishable features which makes it impossible to identify if there is no corresponding AIS signal. However, If linking the boat with an AIS signal is successful all needed information about the vessel, such as whether it is a fishing boat or not becomes readily available. For the purposes of this study it would be very meaningful to determine if a boat detected on a SAR image is a fishing boat especially if it can not be linked with AIS data as those boats could potentially be engaged in illegal fishing. In an attempt to tackle this issue, a random forest (RF) machine learning model will be created to classify detected boats as either fishing vessels or other types of ships. The approach here will be implemented similarly to the method used by Snapir, Waine, and Biermann (Snapir et al., 2019), who predicted a vessel type from only 3 features; latitude, longitude and length of a detected ship.

Creating the dataset is the initial step for making the RF model. This process involved iterating through all the AIS data within the study area and randomly selecting 1,000 boats. To ensure the model did not become biased, an equal number of fishing and non-fishing boats were included (with 500 of each type). For each selected boat, both its location and length were expected and stored in a dedicated file alongside whether or not it was a fishing boat or not. Although more boats could have been used for the RF model training the author found that 1000 boats had a good balance between workability and model generalization. Once the data was ready it was split into 80% training and 20% testing. The *scikit-learn* library

in Python was utilized to both train and evaluate the Random Forest model. All default hyperparameters were maintained except for the number of trees which was adjusted to 100. This decision was based on experimental results. Despite the potential for a more thorough exploration the model performed well meaning that further optimization was unnecessary.

3.6 SAR and AIS data matching

As discussed in the previous section the output of the SAR ship detection tool consists of a list that holds all identified ships within a given SAR image. To correlate them with the AIS data, it is necessary to develop a script that selects all AIS transmissions within exactly the same timeframe as the SAR image. Since ships transmit their data semi-irregularly it is important to interpolate between the 2 nearest transmissions to accurately align ship locations to the right spot. After executing this script a list is generated that holds all ships that transmitted AIS data close to a particular moment in time. Armed with this list and the one generated by the detection operation it becomes possible to initiate the matching process. Theoretically points in both list should line up flawlessly but in practice that is not the case. Instead the points are close to each other but do not align perfectly. This is caused by inherent inaccuracies in both lists but fortunately this difference is small and can easily be accounted for. The last step involves developing a script that compares all the points in both lists. If any pair of points falls within a threshold of 800 meters they are deemed a match. This threshold was picked after considerable manual examination of pairs of lists. Figure 3.10 shows an overview of the key steps involved in matching AIS and SAR data to determine the number of fishing boats broadcasting their AIS data.

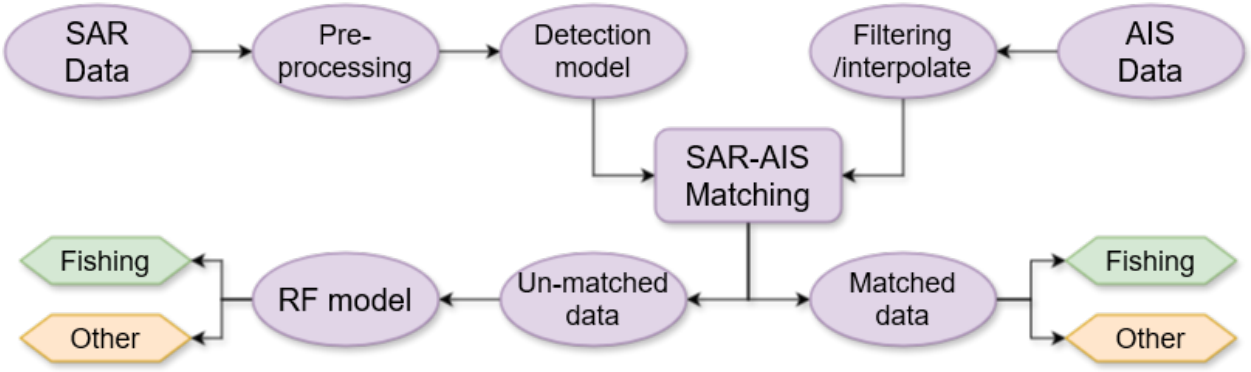

Figure 3.10: Flowchart illustrating the key steps in matching AIS and SAR data

3.7 Evaluating performance of the models

To evaluate the performance of both the random forest model and the deep learning model, it is important to use some metrics that give insights into how well the models are performing. One common practice is to construct a confusion matrix from the test data, which serves as the foundation for calculating several key performance metrics. Precision, recall, F1 score, and accuracy. A confusion matrix is a table used to describe the performance of a

classification model on a set of test data where the true values are known. The matrix shows the counts of true positive (TP), true negative (TN), false positive (FP), and false negative (FN) predictions. While the confusion matrix is valuable on its own, further analysis can yield more concrete performance metrics. These values can be calculated with the following equations:

$$\text{Precision} = \frac{TP}{TP+FP} \qquad \text{Recall} = \frac{TP}{TP+FN}$$

$$\text{F1 Score} = \frac{2 \cdot (\text{Precision} \cdot \text{Recall})}{\text{Precision} + \text{Recall}} \qquad \text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN}$$

4. Results

4.1 Metrics for the ship detection tool

To get a better understanding of how the *CenterNet Resnet50 V2 512x512* model is performing it is necessary to conduct a thorough evaluation. This entails passing the 30% test dataset through the detection model and comparing its predictions against the actual data. This information is then showcased in a confusion matrix which can be further used to extract metrics providing a good indication of the model's effectiveness. The comparison between predicted and actual bounding boxes is not a straightforward process when it comes to object detection. It requires some logic that attempts to align the predicted bounding box with an actual one. There are then 3 distinct outcomes: a match is found (TP), the model predicted a ship that was not present (FP) or the model failed to detect a ship that was present (FN). These 3 values make up the confusion matrix and from it precision, recall and F1 can be derived. Intersection over Union (IoU) is the key component of the matching process. IoU calculates how much a predicted bounding box area overlaps with the actual bounding box area, compared to the total area both cover. Boxes that overlap each other completely have an IoU of 1 while completely separate boxes have an IoU of 0. For this evaluation an IoU threshold of 0.5 and a prediction confidence threshold of 0.275 which resulted in the confusion matrix seen in table 4.1 and table 4.2 shows the metrics derived from it. From the tables it becomes evident that the fine-tuning effort had a massive impact and the model is performing adequately. With a precision, recall and F1 score all about 0.9 the model is able to confidently identify the majority of ships.

Table 4.1: The confusion matrix (Left: before fine tuning) (Right: after fine tuning)

	Predicted Positive	Predicted Negative
Actual Positive	31 (TP)	38 (FP)
Actual Negative	1172 (FN)	

	Predicted Positive	Predicted Negative
Actual Positive	1088 (TP)	112 (FP)
Actual Negative	115 (FN)	

Table 4.2: Metrics derived from the confusion matrix

	Precision	Recall	F1
Before fine tuning	0.449	0.026	0.049
After fine tuning	0.907	0.904	0.906

4.2 Metrics for the vessel classification model

Training the random forest model to classify boats in SAR images as either fishing or non-fishing was a quick process. Using primarily the default hyperparameters the model performed remarkably well. The performance metrics are shown in Table 4.3 and they suggest that the model is quite capable of accurately distinguishing between fishing and non-fishing boats. Recall has the highest score and that indicates that the model correctly identifies 98% of all actual fishing boats. That is particularly advantageous since in the actual data fishing boats represent a small part of the whole data so not missing any is key. Overall, all metrics are large so it is fair to say the model is a success.

Table 4.3: Performance of the random forest model

Precision	95.8%
Recall	98.0%
F1-score	96.9%
Accuracy	96.9%

4.3 AIS-SAR matching

After matching the 10 SAR images with the corresponding AIS data a better picture of the level of AIS uptake among boats can be gotten. Table 4.4 presents the results and in total 643 boats were detected but only 395 successfully matched with AIS data. Initially, this mismatch might appear alarming since it suggests that a great deal of ships in the South Baltic Sea do not broadcast their AIS data. However, by visually inspecting the spatial distribution of the matches it becomes somewhat clear what happened. A great deal of unmatched boats tend to cluster around particular areas. Specifically, the coastline near Germany and Poland as well as the area east of the 16 degrees longitude line had significantly fewer matches compared to areas near Bornholm. This is clearly visible on figure 4.2 which shows both the matched and unmatched boats across the 10 SAR images. With this in mind, a more plausible explanation for the high number of unmatched boats is that while the majority of boats in fact are broadcasting AIS data they are not being logged since the network of AIS receivers simply has a poor reception in those areas. In regions where AIS reception appears to be robust, nearly all boats matched successfully, with only a handful of exceptions. So instead of looking at the numbers seen in table 4.4 it would be more reasonable to exclude the areas with poor reception and recount the matches but that is not so straightforward since there are no clear boundaries to divide the study area up.

Two columns in table 4.4 shows how many of the matched or unmatched boats are fishing vessels. For the matched boats counting how many of them are fishing boats is easy since the information is readily available in the data. For unmatched boats it is difficult however but that is where the random forest model described in section 4.2 comes in. Among other information the detection tool outputs an estimated height and width of each detected boat in pixels. A simple Pythagoras operation can convert those numbers into length and lastly that number is multiplied with the spatial resolution of the raster which is just about 10x10

meters. This will not result in a precise length of the boat but it gives an estimation for the random forest model. All unmatched detections are run through the model and 12 out of the 248 boats were predicted as being fishing. That is a fairly similar ratio to the matched boats or about 4% which is reasonable. This ratio of about 4% is consistent with the one observed for the matched boats which suggests that the model's estimates are reasonable. Verifying these values is not viable but keeping in mind the model's evaluation metrics gives some reassurance of accuracy.

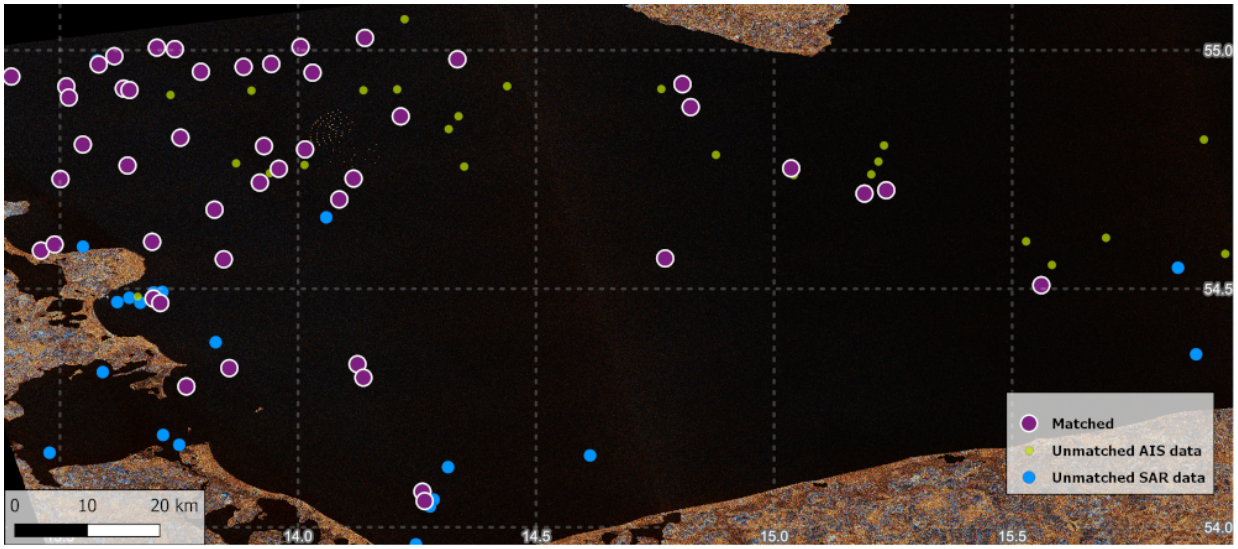

Figure 4.1: Findings from a single SAR Image in Table 3.1, showing the vessels that could be successfully linked between SAR and AIS datasets, as well as those that could not be matched

Table 4.4: The outcomes of the AIS-SAR matching process for each SAR image

SAR image	**Area**	**Matched**		**Unmatched**	
		Other	Fishing	Other	Fishing
1	B	46	5	21	2
2	A	41	0	23	2
3	B	28	0	26	0
4	A	45	3	19	1
5	B	35	2	32	1
6	A	49	1	17	0
7	B	31	0	29	1
8	A	34	1	11	3
9	B	26	1	31	0
10	A	44	3	27	2
Total		379	16	236	12

Figure 4.1 shows the matching results for a segment of a single SAR image. Ideally, one would like to see all boats get matched between the datasets but that is not always the case. The yellow points on figure 4.1 represent boats that did not get detected on the SAR image of which there are quite a few. At first this may seem like the detection tool performed very poorly at identifying boats but a closer inspection reveals that smaller boats are almost indistinguishable from background noise in the SAR image. The spatial resolution of the SAR image is approximately 10x10 meters so some boats may only occupy a pixel or two making them very hard to spot. It should be noted however that this is not always the case and occasionally the tool simply fails to detect boats due to its limitations. Figure 4.2 shows the same thing as 4.1 but instead of just showing the results from a single SAR image is shows results from all SAR images on a single map. Unmatched boats from AIS data are not shown on figure 4.2 due to constraints on the map. As discussed before in this section unmatched ships are predominantly located near the shorelines of Germany and Poland and that is visible on figure 4.2.

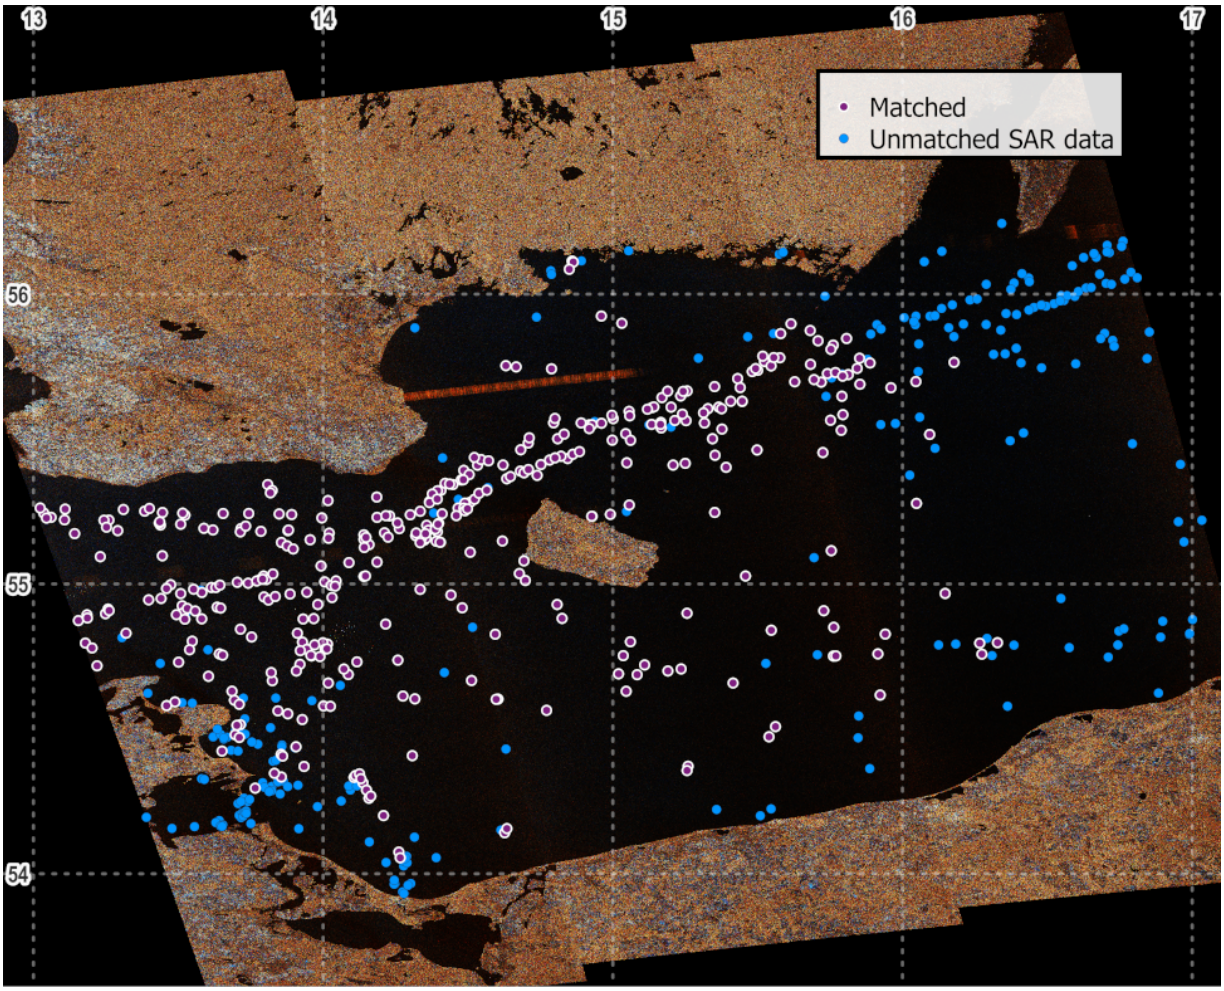

Figure 4.2: Matching results from all SAR images in Table 3.1, showing the vessels that could be successfully linked between SAR and AIS datasets, as well as those that could not be matched

4.4 Vessel trajectory analyses

The script described in section 3.3.2 was used with data encompassing the entire year of 2018. It generated detailed outputs for each route sailed by a fishing vessel during this period. Figure 4.3 shows a specific output from this script, representing a single journey by a single boat that lasted approximately 47 hours. This ship started its journey near Karlskrona, Sweden but the recorded path is incomplete since it does not show the vessel's return to port. Although not ideal for monitoring, the majority of the path is believed to be there. The lower three plots in figure 4.3 detail the vessel's speed, course, and the results from the GFW logistic regression model over time. The background color in these plots indicates the navigational status of the vessel where light purple indicates active fishing, light red indicates transit without fishing activity and white represents periods where no navigational status is recorded. By examining the plots, it can be seen that when the ship moves slowly, it typically indicates that it is fishing and that is also in line with what the GFW model predicts. As the vessel decelerates, it switches its navigational status to active fishing and when it accelerates again it switches back again. There is an instance along the ship's path where no navigational status is transmitted which could simply be due to human error or some technical issue. But by looking at the speed patterns of the boat one could assume ongoing fishing activity during this period. A noticeable uptick in the GFW model prediction can be seen in that time period indicating that there is indeed a high probability of fishing at that moment as well. The connection between the ship's course and fishing activity is less clear. Extensive review of other outputs from the script suggests that speed is a more reliable indicator of fishing efforts. It was also observed that many ships do not update their navigational status consistently making those ships that do more valuable to compare with the GFW model predictions. On figure 4.3 it can also be seen that there is a strong correlation between both the calculated and broadcasted speeds and courses, and that holds true even when exploring many other examples from the script despite occasional anomalies. The uppermost plots on figure 4.3 show the actual path taken by the vessel. The left plot shows its location within the study area while the right plot focuses on the GFW model predictions along the path instead of as a function of time.

After analyzing a full year's worth of AIS data there were 5304 individual fishing boat journeys recorded taken by 284 fishing boats. Given the extensive volume of the data it is impractical to show it all in this document but figure 4.3 displays a single example. However, by examining all 5304 trajectories together additional insights can be had. To get a sense of how many of the fishing boat paths are incomplete, the proximity of each journey's start and end points to the nearest land was looked at. A fully connected path both begins and ends next to land one can assume. For the purposes of this project, the boundary of the study area was considered as land since paths passing through the study area boundaries can not be classified as incomplete. The results can be seen on the three bar graphs of figure 4.4. A 500-meter threshold to classify the proximity of journey endpoints to land was used. If a journey's end point is within this distance it is marked as "yes", otherwise as "no". It can be seen that for both the starting and final points, less than half are within this established threshold, with only about 21% of journeys having both endpoints within this 500-meter

range. The choice of the 500-meter threshold is somewhat arbitrary but it was picked after a thorough review of the data. Adjusting the threshold would have slightly changed the number of ships classified as having an incomplete path, but 500 meters was an optimal choice. Another insightful metric to analyze is the duration of each journey undertaken by fishing boats. Figure 4.5 presents a histogram showing these durations and revealing that the majority of fishing boats are typically tracked for around 19 hours while the longest continuously recorded journey of a fishing vessel spanned a total of 188 hours.

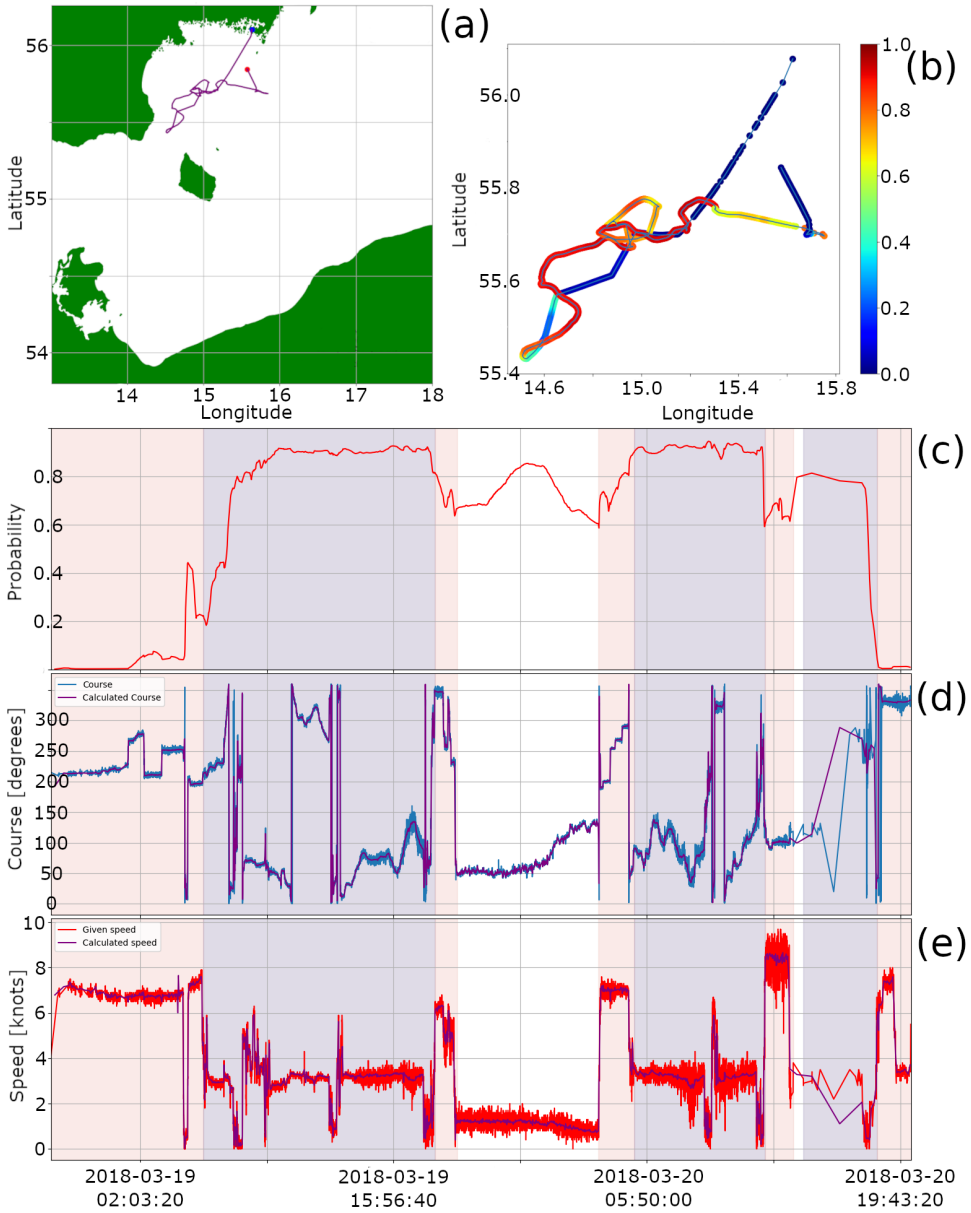

Figure 4.3: Analysis of a Single Fishing Boat's Trajectory (from AIS data): (a) Trajectory mapped, (b) Fishing activity predictions along the path, (c) Model predictions over time, (d) Course over time, (e) Speed over time

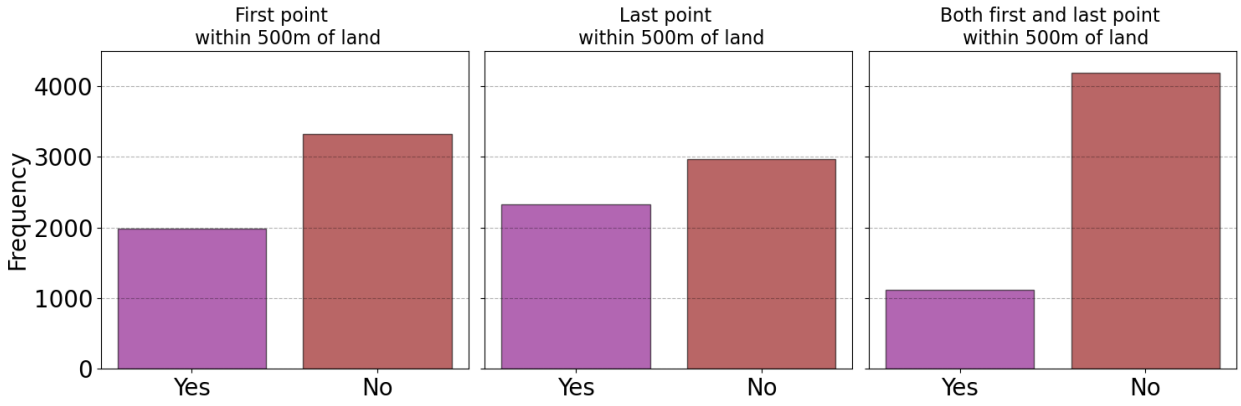

Figure 4.4: Distances from both the start and end points of all recorded fishing routes in 2018 to the nearest land (from AIS data)

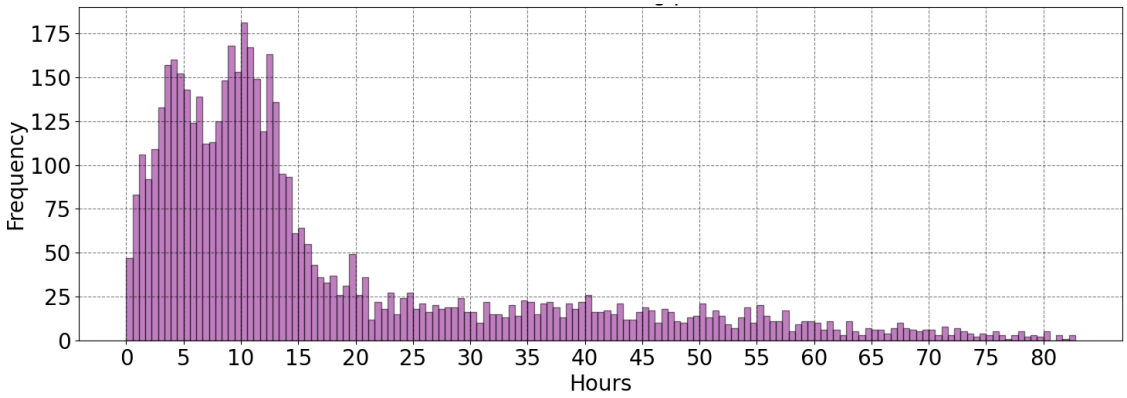

Figure 4.5: Histogram of fishing boat path durations (from AIS data)

4.5 Spatiotemporal analysis of fishing activity

After a comprehensive analysis of AIS data from 2018 several noteworthy findings surfaced. Firstly, figure 4.6 shows the spatial distribution of fishing activity by month. It's important to clarify that in this context fishing is defined as a fishing boat going slower than 4 knots. A more sophisticated approach would be to use the prediction model from GFW but the analyses in section 4.4 revealed that just using speed to predict fishing effort yields sufficiently accurate results. Although the GFW model might give better results, using speed to filter data is considerably more efficient for processing large amounts of data. Figure 4.6 shows that the area southeast of Bornholm is among the most significant fishing grounds in the study area. Although not a precise point, this general region consistently exhibits high activity throughout most of the year. In general however, the clustering of fishing activities shifts significantly over the course of the year. For example, the fishing hot-spots in November differ quite a bit from those in March. This seasonal variation in fishing hot-spots is primarily caused by targeted fishing practices where fishermen pursue a certain type of

fish species based on their availability at a given time and that leads to dynamic changes in fishing hot-spots.

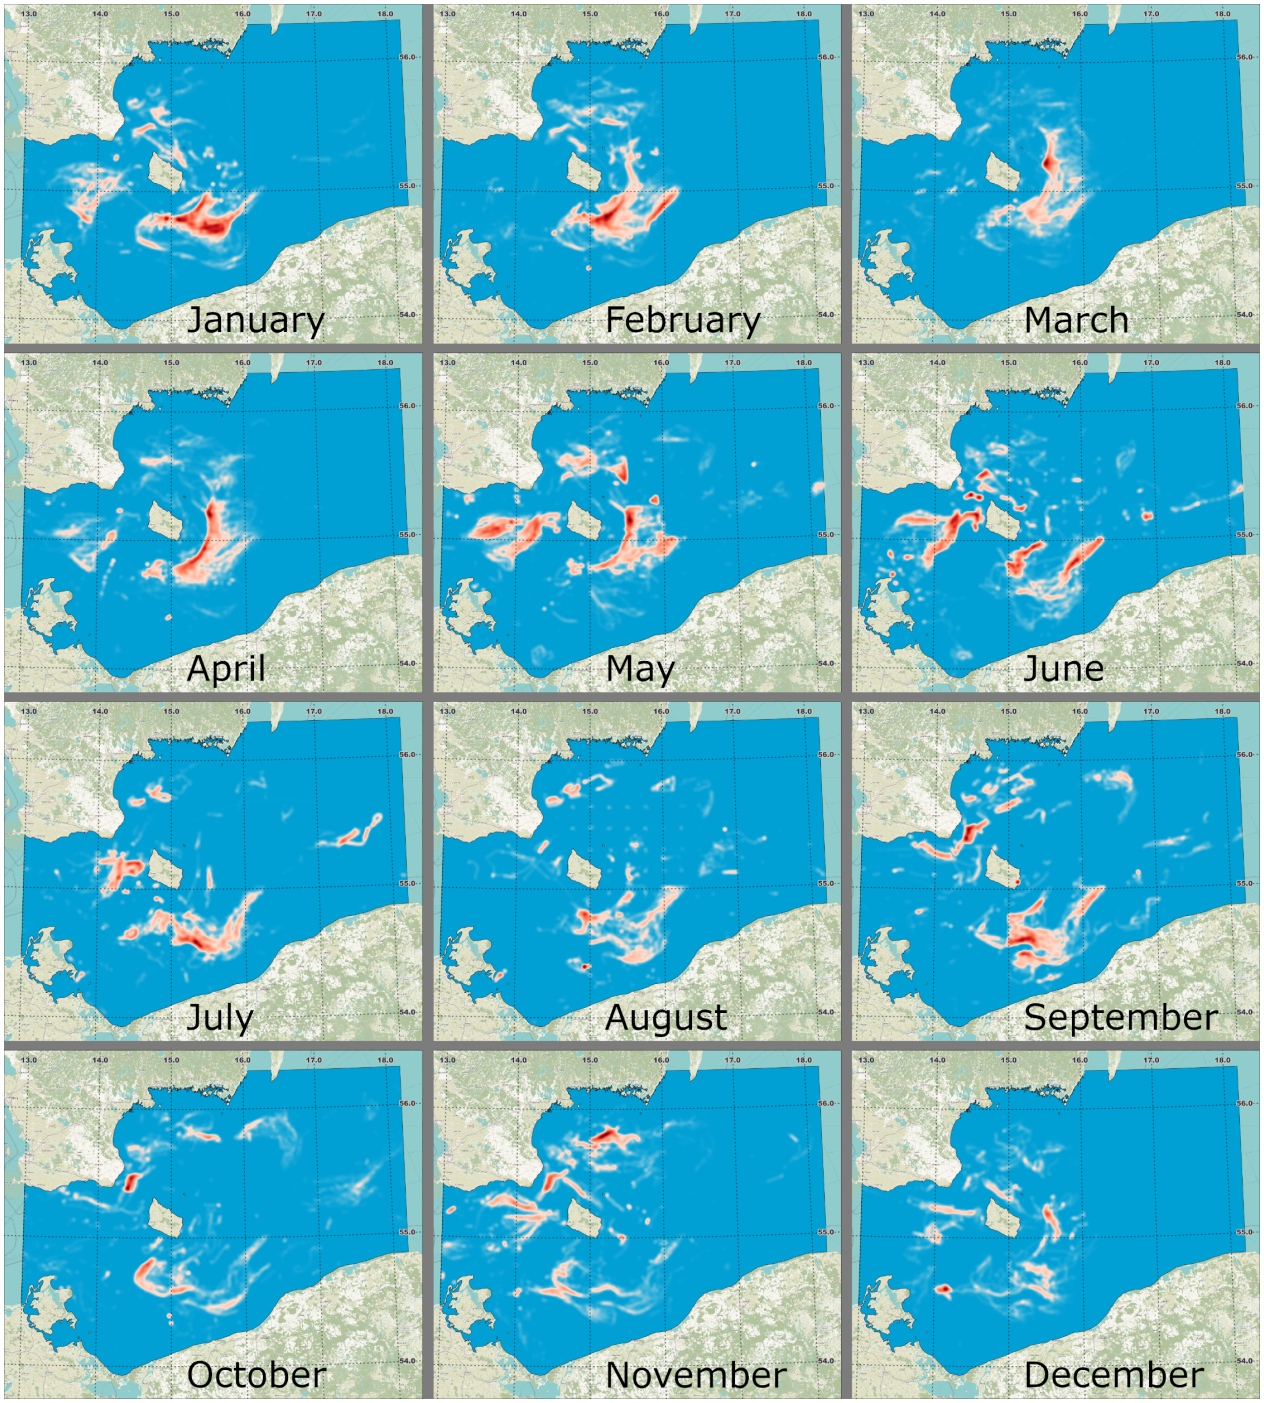

Figure 4.6: A heatmap of fishing activity by month in 2018 (from AIS data)

The temporal analysis from the same data is shown in figure 4.7 and there the y-axis represents the frequency of AIS signals reporting speeds below 4 knots within the study

area. It is clear from the figure that fishing efforts reached their peak in March, followed by a gradual decline to the annual minimum in August. The following months see a slight increase but it still remains significantly lower than the March levels. A more detailed figure showing temporal distribution within each month can be viewed in appendix A. The analysis indicates that fishing activities increase during the colder months due to favorable conditions for certain species. For instance cod fishing peaks during these months. Moreover, quota limitations significantly impact fishing activities. Quotas can cause fluctuations in fishing intensity as fishermen strive to maximize their catches within the permitted limits. This often leads to a surge in activity at the beginning of the year when new quotas are released. In appendix B fishing effort by country can be seen.

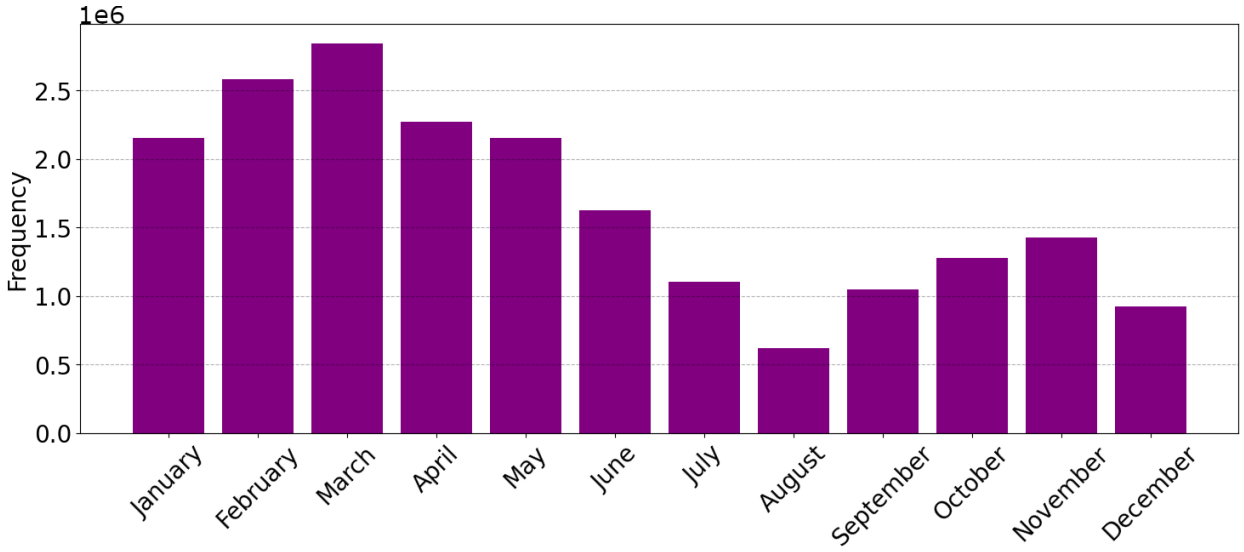

Figure 4.7: The number of AIS transitions during fishing by month in 2018 (from AIS data)

5. Discussion

5.1 Fishing boat analysis from AIS

This project involved several analyses that exclusively used AIS data. Specifically, the vessel trajectory script examined individual voyages, while both the spatial and temporal analyses focused on broader aspects of fishing activity within the study area. This section will provide discussion on these topics.

5.1.1 Possible improvements to the vessel trajectory analyses

Numerous improvements could be made to the AIS trajectory analysis script. Given that the data was stored in a CSV file and not a database, directly querying for a specific boat was not feasible and instead the script had to iterate over each line in all CSV files to locate all instances of the desired ship. A better approach would have been to conduct trajectory analyses for all boats simultaneously, thus requiring only a single pass through all the data. While this modification would have significantly sped up the process, its implementation is not straightforward since loading all the data into memory would have too much. Nevertheless, alternative approaches that bypass these limitations do exist. There are also undoubtedly other ways to optimize the script but for the purposes of this project it was sufficient. In the script a time threshold of 6 hours was set to determine when a ship's path had concluded. If no new AIS data was received within a window of 6 hours the path was deemed finished but otherwise the AIS message was considered to be a part of the ongoing path (even though 5 hours had passed for example). This threshold of 6 hours is somewhat arbitrary but it was selected after careful inspection of the data which showed that fishing boats sometimes disappear from AIS tracking for a few hours before reappearing. It was also noted however that occasionally boats are docked for a handful of hours and then start another voyage but the script treated these instances as one continuous journey even though these are two distinct paths. A more sophisticated approach to segmenting fishing boat journeys could be for example to define an "in port" status for ships and utilizing this status to segment paths. This could help differentiate between separate trips and enhance the analysis of fishing activities.

5.1.2 Limitations of the AIS data

A crucial question when it comes to the effectiveness of using AIS for monitoring fishing activities concerns the system's reliability. As discussed in section 2.1.2, AIS was originally designed for increasing maritime safety although in practice it has proven to be very useful for monitoring purposes. However, its effectiveness for monitoring is not flawless since there are instances where ships lawfully broadcasting their AIS data are not consistently tracked. The primary reason for this is that the ship's AIS signal may not be captured by any receivers. This is particularly relevant in the context of this study which uses AIS data provided by the

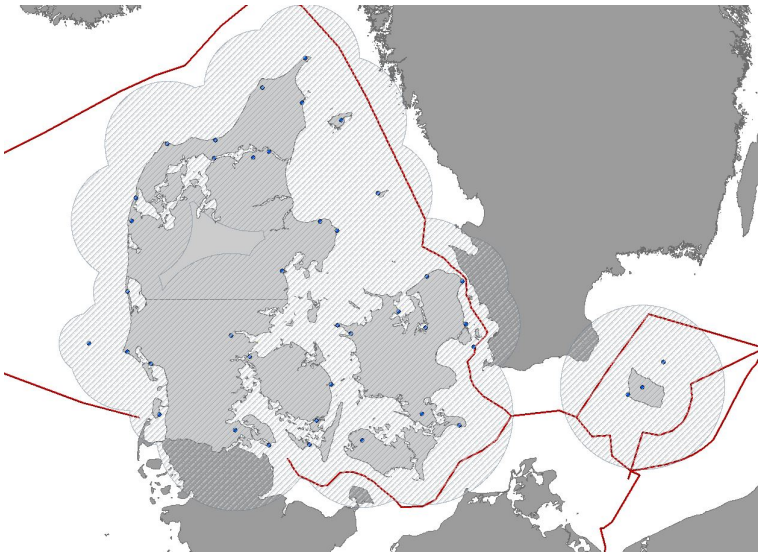

Figure 5.1: Map showing the network of land-based AIS receivers in use by the Danish Maritime Authority. Permission to use this map was granted by the authority. Note: The map is not updated and should be used solely for illustrative purposes

Danish Maritime Authority, whose network of AIS receivers is intended solely for covering Danish waters. Although it often captures signals from vessels that are far outside this area that is much less reliable data. Figure 5.1 shows the network of AIS receivers and the coverage they provide. This goes far in explaining the prevalence of incomplete paths as shown in figure 4.4. This alone however does not fully account for the incomplete paths observed because some paths were seen starting or ending within well covered areas. An obvious conclusion in those cases would be that those particular boats are actively disabling their AIS data for whatever reason but that is not a guarantee. Given that AIS signals are inherently radio-based (Taconet et al., 2019), many other factors play a role in determining whether a not a signal gets picked up by a receiver such as the height and power of the transmitter onboard the ship, atmospheric conditions, terrain, radio interference and more. These elements contribute to the observed dynamic range of the AIS receivers and occasionally allow for the detection of distant boats. This also has an impact on the findings in section 4.5 since areas close to the island of Bornholm get a better representation. Given these considerations it would have been more reasonable to limit the study area to regions with more reliable AIS data coverage from the beginning.

5.1.3 Further research

The research done in this project could be taken further in several ways. Expanding a longer time frame could reveal how fishing trends evolve over the years and provide valuable insights into the fishing industry. Additionally, increasing the size of the study area would also be an option although this would require an alternative source of data. In this project, a single logistic regression model from Global Fishing Watch (GFW) was used to monitor fishing activity on a small-scale and detect when fishing occurs along a vessel's path. It would have been beneficial to test and compare this model with other models. One way to approach this

could involve creating a custom machine learning model tailored to the project's study area which would potentially provide a more accurate prediction. Obtaining the data needed for such an improvement however could prove to be difficult but one way around this could be to use the navigational status from AIS data. Although one of the findings in this project was indeed that many fishing boats do not consistently update their navigational status message while operating, but some do, and a few handpicked vessels with reliable navigational status messages could possibly provide quality training data for a machine learning model specifically tailored to the study area.

A more sophisticated approach to predict when along its path a boat is engaged in fishing would be to do so by gear type. Different fishing activities employ various types of gear and creating a machine learning model for each gear type could improve accuracy. As it turns out, GFW offers a dataset specifically for this purpose which can be accessed here. This data provides extensive training data for gear types such as drifting longlines, fixed gear and trawlers to name a few. Developing a model that predicts fishing activity based on gear type could yield a detailed view of fishing practices and provide an effective tool to monitor fishing activity. An important consideration here is that this approach assumes prior knowledge of a vessel's gear type, which is not always available. When a gear type of a tracked ship is unknown one would need some way of classifying it based on its movement in the ocean and make sure it is even a fishing vessel to begin with. Fortunately, machine learning solutions are available to address this challenge.

The study area experienced a significant decline in fishing activities following the 2019 ban on cod fishing and the reduction of quotas for herring and other species (The Fisheries Secretariat, 2019). However, the data used in this project is all from 2018, when there was a significant amount of fishing activity compared to the following years. Exploring this slowdown by analyzing more years would have been interesting but that falls out of the scope of this project.

5.2 Fishing boat analysis from SAR

The heart of all the SAR analyses done in this project is undoubtedly the deep learning model itself. The *CenterNet Resnet50 V2 512x512* to be specific and there is certainly room for improvement. As became apparent in section 4.1 the model's performance fell short of expectations and how it can be improved is a bit of a complex issue. There are a lot of nuances and many factors that play into the performance of this deep learning model. This section will highlight some key talking points and mention possible improvements as well as provide insight into the decision making process.

5.2.1 Pre-processing and training

One major topic of discussion is the choice of training data. The dataset used in this project (DSSDD) was limited and required extensive pre-processing to be used. The training images from DSSDD were derived from the SLC level of processing which is quite a raw format and

requires significant processing to work with. The fact that the covariance matrix is computed for each pixel further adds to the processing requirements. The most time-consuming part of the SAR image pre-processing however was the georeferencing step. This step is essential to get coordinates from detected ships but proved to be a major bottleneck. A lot of effort was spent trying to circumvent this step by looking into alternative tools and trying to simply use the rudimentary georeferencing that was stored within the SCL images. Ultimately this effort was in vain and the georeferencing step remained in the workflow. As a result, processing a single SAR image took approximately 90 minutes on a reasonably fast computer which severely limited the number of images that could be processed for this project. Using an already georeferenced image would have been a good option but it required designing a completely new workflow. A new dataset and retraining of the deep learning model would have been needed which would be too time-consuming.

The DSSDD training data was specifically chosen to leverage the power of transfer learning by fine-tuning a pre-existing model and those models are typically trained on a large set of common RGB images. The DSSDD provides two alternative formats for the images, 16-bit files containing 4 bands for each cell in the covariance matrix and a more conventional 8-bit RGB file. It would have been ideal to use the latter but the exact approach used to obtain the RGB files was not clearly documented. This knowledge is crucial since the final deep learning model must receive data consistent with its training input. Because the author was unable to recreate the RGB files provided it was decided to use the 16-bit files and perform a custom scaling instead to have reproducible 8-bit RGB files. The author experimented with several methods to achieve this and ultimately used the equations shown in Table 3.3. The chosen approach for this scaling significantly impacts the model's performance and the objective should be to make the RGB images resemble real-life scenes as closely as possible to effectively utilize transfer learning. This is because transfer learning relies on pre-trained models that have learned features from real-world images, so having similar data ensures these features remain useful for the new SAR images. A simple linear scaling method was used in this project, although there are many alternatives, such as logarithmic scaling. However, increasing the complexity of the conversion process also increases the processing time. Additionally, experimenting with different scaling methods can be very time-consuming since training a deep learning model is a very lengthy process. The equations in table 3.3 were found to assign colors well since the resulting images were not completely dominated by one color and boats at sea were easily distinguishable from the ocean background to the naked eye.

The training process of the deep learning model could benefit from some hyperparameter tuning. Trying different loss functions, optimizers and data augmentation would have been interesting to explore but since it takes such a long time to train a single model a thorough hyperparameter tuning was not feasible. In total the author ran the training process 4 times with. Twice with the *CenterNet Resnet50 V2 512x512*, one time with another similar CenterNet based model and one time with an SSD. The limited testing conducted did not yield significant improvements and more extensive testing is necessary to develop a model that substantially outperforms the one used in this project.

5.2.2 Challenges in utilizing the deep learning model

Once the deep learning model was fully trained, it underwent testing with processed SAR images specifically acquired for this project. During this phase a few unexpected obstacles emerged that needed attention. One such issue was the model's tendency to misclassify windmills as boats which is perhaps not all too surprising given their similar appearances in SAR imagery. To address this the author drew polygons around the 2 windmill farms that were in the study area and updated the IsWithinOcean function accordingly. That ment that all predictions made by the model within those areas get filtered out. Another problem that became apparent when testing was a peculiar anomaly seen only on some SAR images. A thick horizontal band made up of numerous vertical lines. This artifact varied in transparency and sometimes it completely obstructed anything under it but other times it was semi-transparent. An example of this can be seen just under the 56th latitude line on figure 4.2. This issue is likely related to the satellite itself and does not significantly impact the overall image. Fortunately the deep learning model did not get misled these anomalies so it was not of high importance to fix.

The performance of the deep learning model observed in practice aligns closely with the metrics presented in table 4.2. Although the model successfully detects the majority of boats it sometimes overlooks very obvious instances of boats without any apparent reason. It was also noted that the model occasionally might incorrectly detect a single boat as multiple ones. Instead of outlining a single boat with one bounding box the model divides it into multiple sections. To address this a minimum threshold was established where detections within 100 meters of each other get reduced to a single detection. This is not a perfect solution however since the bounding box that is left does not cover the boat fully which results in an incorrect estimation of a boat's length. This has no effect on the AIS-SAR matching process but it does impact the random forest classifier. Improving this aspect would require creating a new deep learning model as was just discussed, that is neither quick nor straightforward.

5.2.3 The random forest classifier

The random forest classifier was trained to categorize unmatched boats from SAR as either fishing vessels or other types. The model was trained using only three features: latitude, longitude, and estimated length but despite its simplicity, the model performed well as can be seen in table 4.3. Adding more features, such as the distance to land and time of day could have made the model even better. Hyperparameter tuning was also a potential area for improvement but it was decided to prioritize other tasks. The model's strong performance can be understood to some degree by looking at the characteristics of the study area. Most boats in the region are simply navigating between the inner Baltic Sea and the North Sea, passing through Denmark. These transit boats follow a clear path through the study area. Additionally, fishing vessels are generally smaller than other types and these factors make it easier for the random forest model to classify the boats accurately.

5.2.4 The AIS-SAR matching

Matching boats between data sources was achieved using a proximity threshold which worked well. However, the threshold had to be set to 800 meters which is a relatively large distance. This was needed because the boats often did not perfectly align between the 2 data sources. It was initially thought that a small time shift in the AIS data could improve alignment but some testing showed this was not helpful. Poor georeferencing or slight errors in the AIS messages could contribute to this mismatch but further investigation would be needed. Fortunately, because of the vastness of the ocean, the risk of misidentification due to the large threshold was low. Boats rarely come so close to one another out at sea. Extensive analysis of the matching results confirmed that this slight offset had no significant impact on the outcomes.

By examining the spatial distribution of the unmatched boats in figure 4.2, it is clear that the majority of these boats cluster in specific areas. This behavior is primarily due to poor AIS reception in certain regions which figure 5.1 highlights. To address this issue, one could either utilize an alternative AIS data source or limit the study area to regions with consistently reliable AIS reception. Disregarding AIS coverage compromises the reliability of the matching process results. The purpose of matching AIS data to SAR data is to estimate the number of ships actively using the AIS system which means adequate AIS coverage is needed for meaningful analysis. With that in mind, on figure 4.2, focusing solely on areas with good AIS coverage results in nearly all boats being matched.

5.2.5 Further research

There are several ways to build upon the work done in this project. As was the case in section 5.1.3, a logical next step would be to process more SAR images over an extended time period to gain a deeper understanding of fishing activity and trends. Analyzing only 10 images from a single month provided limited data making it challenging to draw reliable conclusions. Furthermore, it is essential for future research to take into account AIS coverage to ensure reliable findings. Although the deep learning model used in this project performed reasonably well, further research might explore alternative approaches for detecting ships at sea. The SUMO software, a simpler yet well-established tool often used by researchers could be considered. However, using deep learning models has been the trend in recent years and if that approach is preferred an alternative training data source could be explored. Given the niche nature of ship detection from SAR images there are few alternatives available but some do exist and examples can be found here. Most training sets use Ground Range Detected (GRD) images instead of SLC since GRD is a more processed and standardized product with smaller file sizes making it easier to work with. In particular, the xView3-SAR dataset, which contains a larger volume of GRD data could be a suitable alternative to the DSSDD dataset. This extensive data could be used to train a deep learning model from scratch and avoid the need to fine-tune pre-trained models that have been trained on different image types (common RGB images instead of SAR). When it comes to the deep learning model itself, one could explore further the usage of different models like You Only Look Once (YOLO) or transformers which have proven useful.

6. Conclusion

This project set out to determine whether fishing operations in the southern Baltic Sea could be effectively monitored using publicly available data, specifically through cooperative (AIS) and non-cooperative (SAR) methods. The results conclusively show that both AIS and SAR data can be effectively utilized to monitor fishing activity in this region.

The analysis of historical AIS data for the year of 2018 provided a comprehensive overview of fishing patterns and trends across the study area. Both temporal and spatial distributions were explored in great detail, offering valuable insights into the behaviors of fishing vessels over the whole year. On a smaller scale, the trajectory of individual fishing boats was analyzed, demonstrating that AIS data can effectively monitor the activity of a single vessel and potentially identify suspicious activity. This project shows that even though many ships do not disclose periods of active fishing in their AIS signals, a machine learning model can predict when a boat is fishing based on its trajectory. However, this analysis is limited to vessels that actively broadcast their AIS data and in fact not all boats do so. In some cases it is unintentional but other times boats intentionally turn off their AIS data indicating they might be engaged in illegal fishing. These non-broadcasting vessels can still be detected using satellite imagery. In particular this project used SAR images which proved effective in identifying boats that were not transmitting AIS data. The deep learning model fine-tuned for this project was a useful tool in this regard. By matching SAR data to AIS data, the level of AIS uptake among fishing vessels was assessed which revealed that most ships in the study area do indeed transmit their AIS signals with only a handful of exceptions. However, the few boats detected in areas with good AIS reception but without corresponding AIS signals are of particular interest because they might be engaged in illegal activities. Distinguishing between fishing boats and other types of ships is necessary to determine the amount of fishing boats not transmitting their AIS signal. For that, a random forest model was trained. This model showed promising results but due to the limited dataset (only 10 SAR images from one month) it generated limited results. Processing more SAR images over a larger area and a longer timeframe would have been ideal.

In conclusion, the integration of AIS and SAR methods offers a viable and effective solution for monitoring fishing activities in the southern Baltic Sea. As remote sensing technology continues to advance, the combined use of cooperative and non-cooperative monitoring tools will play an increasingly important role in sustainable fishing practices and the enforcement of maritime regulations.

References

Bartlett, C. (2024). *Vessels go dark to avoid houthi attacks, but may still be vulnerable.* Retrieved 2024-04-22, from https://theloadstar.com/vessels-go-dark-to-avoid-houthi-attacks-but-may-still-be-vulnerable/

Bergman, B. (2023, May 31). *Tankers falsify ais tracking positions to hide entry into russian black sea ports.* Global Fishing Watch. Retrieved from https://globalfishingwatch.org/data/tankers-falsify-ais-tracking-positions-to-hide-entry-into-russian-black-sea-ports/

Bunwaree, P. (2023). The illegality of fishing vessels 'going dark' and methods of deterrence. *International and Comparative Law Quarterly*, *72*(1), 179–211. doi: 10.1017/S0020589322000525

Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., & Zagoruyko, S. (2020). End-to-end object detection with transformers. In *European conference on computer vision* (pp. 213–229).

Cochrane, K. L. (2002). *A fishery manager's guidebook: Management measures and their application* (No. 424). Rome, Italy: Food and Agriculture Organization of the United Nations. (Series number: 0429-9345 - T424. Job number: Y3427E)

Copernicus programme. (2024). https://sentiwiki.copernicus.eu/web/copernicus-programme.

Earthdata. (N/A). *What is synthetic aperture radar?* Retrieved 2024-03-02, from https://www.earthdata.nasa.gov/learn/backgrounders/what-is-sar

European Commission. (2021). *Highlights, press releases and speeches.* https://ec.europa.eu/commission/presscorner/detail/en/qanda_21_2747. (Accessed: March 6, 2024)

Food, & of the United Nations, A. O. (2024). *What is iuu fishing?* Retrieved 2024-02-29, from https://www.fao.org/iuu-fishing/background/what-is-iuu-fishing/en/

Galdelli, A., Mancini, A., Ferrà, C., & Tassetti, A. N. (2021). A synergic integration of ais data and sar imagery to monitor fisheries and detect suspicious activities. *Sensors*, *21*(8). Retrieved from https://www.mdpi.com/1424-8220/21/8/2756 doi: 10.3390/s21082756

Hendriks, S. L. (2022, June 22). Sustainable small-scale fisheries can help people and the planet. *Nature*. Retrieved from https://www.nature.com/articles/d41586-022-01683-2 (Accessed: 2024-01-25) doi: 10.1038/d41586-022-01683-2

Hu, Y., Li, Y., & Pan, Z. (2021). A dual-polarimetric sar ship detection dataset and a memory-augmented autoencoder-based detection method. *Sensors*, *21*(24). Retrieved from https://www.mdpi.com/1424-8220/21/24/8478 doi: 10.3390/s21248478

Kurekin, A. A., Loveday, B. R., Clements, O., Quartly, G. D., Miller, P. I., Wiafe, G., & Adu Agyekum, K. (2019). Operational monitoring of illegal fishing in ghana through exploitation of satellite earth observation and ais data. *Remote Sensing*, *11*(3). Retrieved from https://www.mdpi.com/2072-4292/11/3/293 doi: 10.3390/rs11030293

Li, J., Xu, C., Su, H., Gao, L., & Wang, T. (2022). Deep learning for sar ship detection: Past,

present and future. *Remote Sensing*, *14*(11). Retrieved from https://www.mdpi.com/2072-4292/14/11/2712 doi: 10.3390/rs14112712

Mahgoun, H., Chaffa, N. E., Ouarzeddine, M., & Souissi, B. (2020, November 6). Application of polarimetric-sar decompositions on radarsat-2 fine quad-pol images to enhance the performances of ships detection algorithms. *Sensing and Imaging*, *21*(1), 56. Retrieved from https://doi.org/10.1007/s11220-020-00321-3 doi: 10.1007/s11220-020-00321-3

Mazzarella, F., Vespe, M., & Santamaria, C. (2015). Sar ship detection and self-reporting data fusion based on traffic knowledge. *IEEE Geoscience and Remote Sensing Letters*, *12*(8), 1685-1689. doi: 10.1109/LGRS.2015.2419371

Mills, C. M., Townsend, S. E., Jennings, S., Eastwood, P. D., & Houghton, C. A. (2006, 12). Estimating high resolution trawl fishing effort from satellite-based vessel monitoring system data. *ICES Journal of Marine Science*, *64*(2), 248-255. Retrieved from https://doi.org/10.1093/icesjms/fsl026 doi: 10.1093/icesjms/fsl026

Möller, E., & Hochberg, T. (2020). *Vessel-scoring.* https://github.com/GlobalFishingWatch/vessel-scoring/tree/master. (Accessed: 27 April 2024)

Natale, F., Gibin, M., Alessandrini, A., Vespe, M., & Paulrud, A. (2015). Mapping fishing effort through ais data. *PloS one*, *10*(6), e0130746.

Ojaveer, H., Jaanus, A., MacKenzie, B. R., Martin, G., Olenin, S., Radziejewska, T., Telesh, I., Zettler, M. L., & Zaiko, A. (2010, 09). Status of biodiversity in the baltic sea. *PLOS ONE*, *5*(9), 1-19. Retrieved from https://doi.org/10.1371/journal.pone.0012467 doi: 10.1371/journal.pone.0012467

Paolo, F., Kroodsma, D., Raynor, J., Hochberg, T., Davis, P., Cleary, J., Marsaglia, L., Orofino, S., Thomas, C., & Halpin, P. (2024). Satellite mapping reveals extensive industrial activity at sea. *Nature*, *625*(7993), 85–91. Retrieved from https://doi.org/10.1038/s41586-023-06825-8 doi: 10.1038/s41586-023-06825-8

Petrossian, G. A. (2015). Preventing illegal, unreported and unregulated (iuu) fishing: A situational approach. *Biological Conservation*, *189*, 39-48. Retrieved from https://www.sciencedirect.com/science/article/pii/S0006320714003140 (Detecting and Understanding Non-compliance with Conservation Rules) doi: https://doi.org/10.1016/j.biocon.2014.09.005

Ritchie, H., & Roser, M. (2021). Fish and overfishing. *Our World in Data*. (https://ourworldindata.org/fish-and-overfishing)

Shi, H., Chai, B., Wang, Y., & Chen, L. (2022). A local-sparse-information-aggregation transformer with explicit contour guidance for sar ship detection. *Remote Sensing*, *14*(20). Retrieved from https://www.mdpi.com/2072-4292/14/20/5247 doi: 10.3390/rs14205247

Snapir, B., Waine, T. W., & Biermann, L. (2019). Maritime vessel classification to monitor fisheries with sar: Demonstration in the north sea. *Remote Sensing*, *11*(3). Retrieved from https://www.mdpi.com/2072-4292/11/3/353 doi: 10.3390/rs11030353

Sun, X., Wang, Z., Sun, Y., Diao, W., Zhang, Y., & Fu, K. (2019). Air-sarship-1.0: High-resolution sar ship detection dataset (in english). *Journal of Radars*, *8*(R19097). Retrieved from https://radars.ac.cn/en/article/doi/10.12000/JR19097 doi: 10.12000/JR19097

Taconet, M., Kroodsma, D., & Fernandes, J. (2019). *Global atlas of ais-based fishing activity*

(1st ed.). Rome, Italy: FAO.

The Fisheries Secretariat. (2019). *Update on eastern baltic cod.* Retrieved from https://www.fishsec.org/2019/09/28/update-on-eastern-baltic-cod/ (Accessed: 2024-05-20)

Thoya, P., Maina, J., Möllmann, C., & Schiele, K. S. (2021). Ais and vms ensemble can address data gaps on fisheries for marine spatial planning. *Sustainability*, *13*(7). Retrieved from https://www.mdpi.com/2071-1050/13/7/3769 doi: 10.3390/su13073769

Wang, Y., Wang, C., Zhang, H., Dong, Y., & Wei, S. (2019). A sar dataset of ship detection for deep learning under complex backgrounds. *Remote Sensing*, *11*(7). Retrieved from https://www.mdpi.com/2072-4292/11/7/765 doi: 10.3390/rs11070765

Wang, Y., Zhao, Z., He, J., Zhu, Y., & Wei, X. (2021). A method of vehicle flow training and detection based on resnet50 with centernet method. In *2021 international conference on communications, information system and computer engineering (cisce)* (p. 335-339). doi: 10.1109/CISCE52179.2021.9446012

xview3 competition on detecting dark vessels to combat illegal fishing. (2023). https://iuu.xview.us/. (Accessed: 2024-03-12)

Yang, D., Wu, L., Wang, S., Jia, H., & Li, K. X. (2019). How big data enriches maritime research – a critical review of automatic identification system (ais) data applications. *Transport Reviews*, *39*(6), 755–773. Retrieved from https://doi.org/10.1080/01441647.2019.1649315 doi: 10.1080/01441647.2019.1649315

Yang, H., Kang, X., Liu, L., Liu, Y., & Huang, Z. (2023). Sar-hub: Pre-training, fine-tuning, and explaining. *Remote Sensing*, *15*(23). Retrieved from https://www.mdpi.com/2072-4292/15/23/5534 doi: 10.3390/rs15235534

Yasir, M., Jianhua, W., Mingming, X., Hui, S., Zhe, Z., Shanwei, L., Colak, A. T. I., & Hossain, M. S. (2023). Ship detection based on deep learning using sar imagery: a systematic literature review. *Soft Computing*, *27*(1), 63–84. Retrieved from https://doi.org/10.1007/s00500-022-07522-w doi: 10.1007/s00500-022-07522-w

Zhang, T., Zhang, X., Li, J., Xu, X., Wang, B., Zhan, X., Xu, Y., Ke, X., Zeng, T., Su, H., Ahmad, I., Pan, D., Liu, C., Zhou, Y., Shi, J., & Wei, S. (2021). Sar ship detection dataset (ssdd): Official release and comprehensive data analysis. *Remote Sensing*, *13*(18). Retrieved from https://www.mdpi.com/2072-4292/13/18/3690 doi: 10.3390/rs13183690

Appendix A: Temporal fishing patterns

This appendix is an extension to figure 4.7. Instead of showing a single number for each month, this appendix will show a detailed monthly temporal distribution for the year 2018, derived from AIS data.

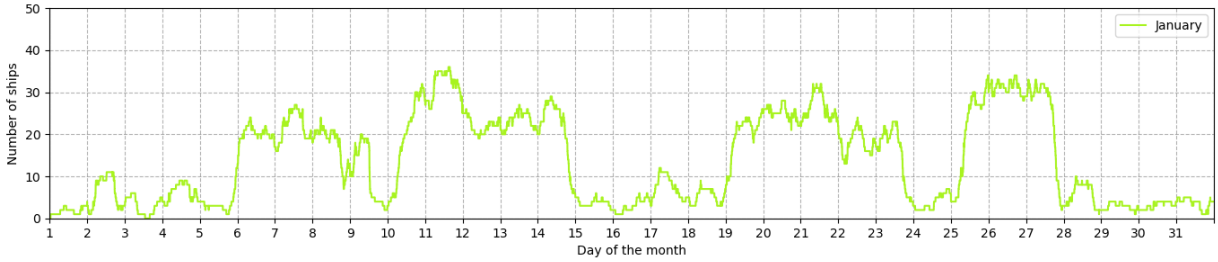

Figure A.1: Amount of fishing effort in January from the AIS data

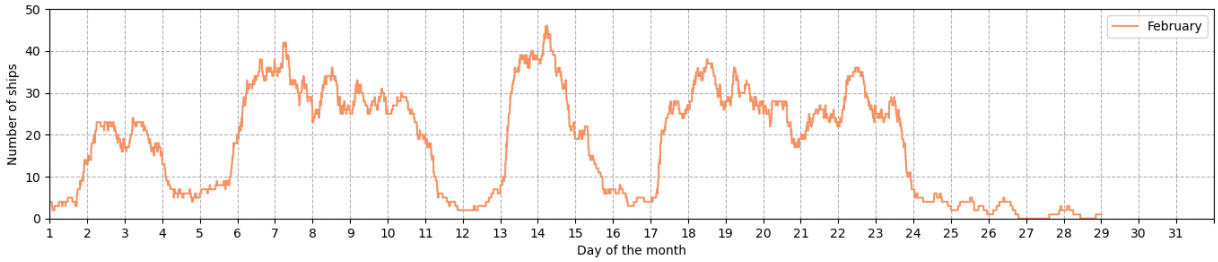

Figure A.2: Amount of fishing effort in February from the AIS data

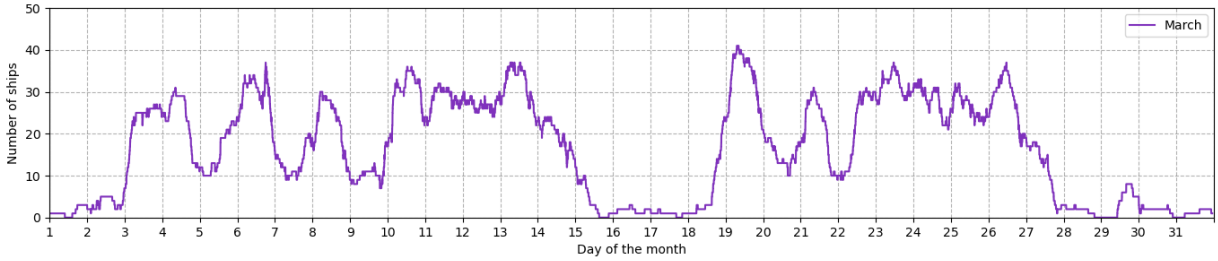

Figure A.3: Amount of fishing effort in March from the AIS data

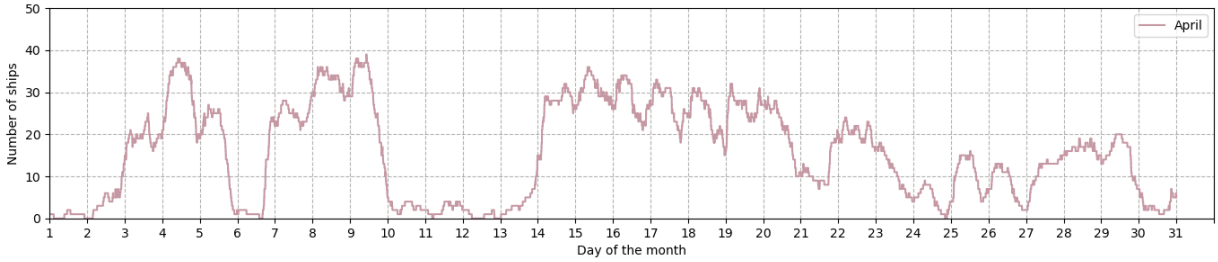

Figure A.4: Amount of fishing effort in April from the AIS data

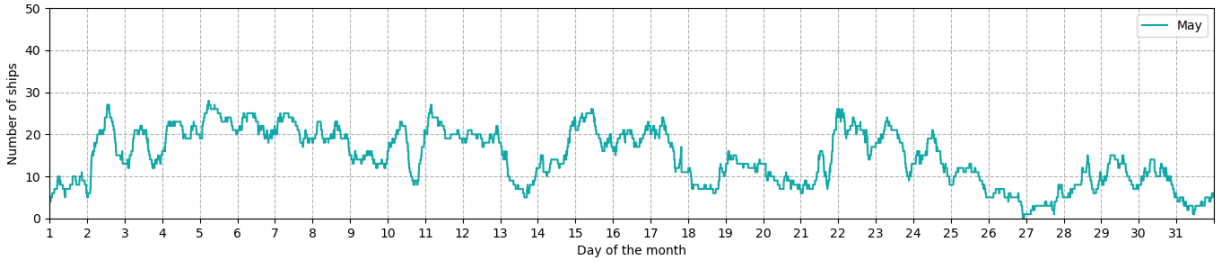

Figure A.5: Amount of fishing effort in May from the AIS data

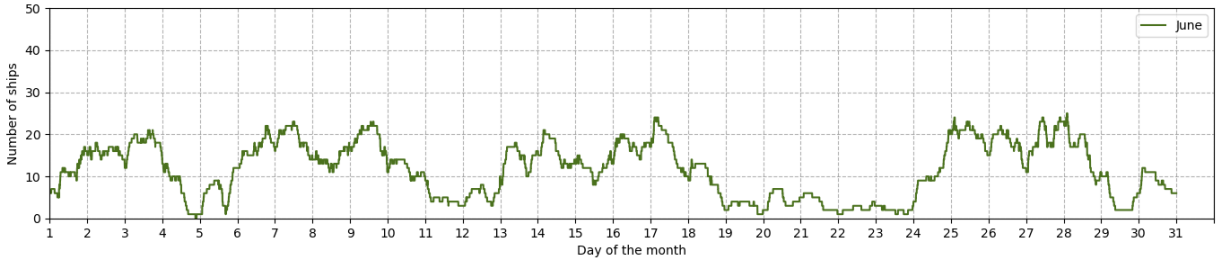

Figure A.6: Amount of fishing effort in June from the AIS data

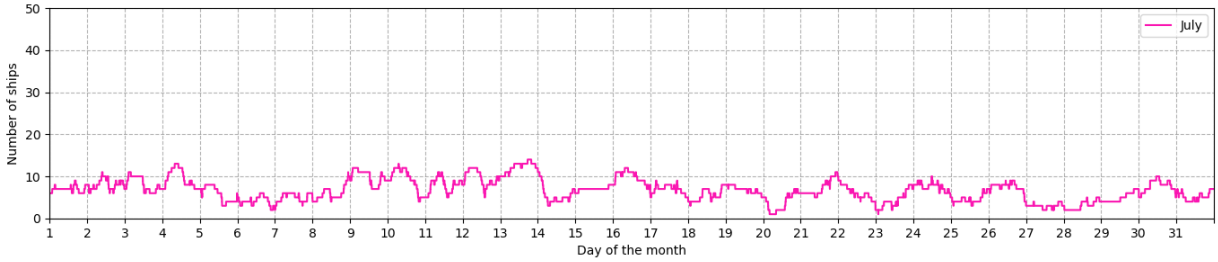

Figure A.7: Amount of fishing effort in July from the AIS data

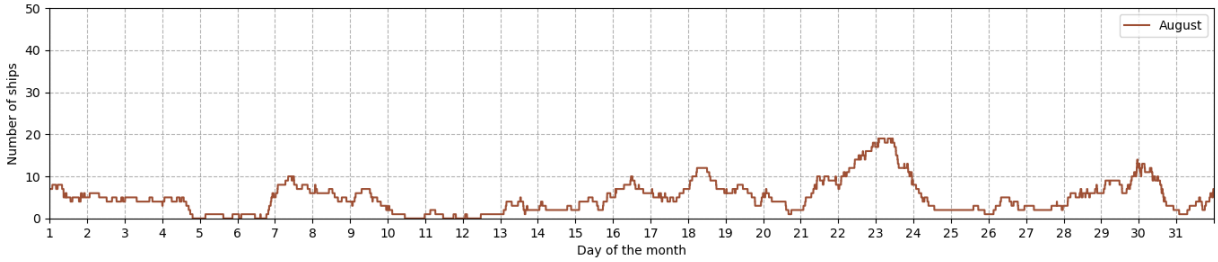

Figure A.8: Amount of fishing effort in August from the AIS data

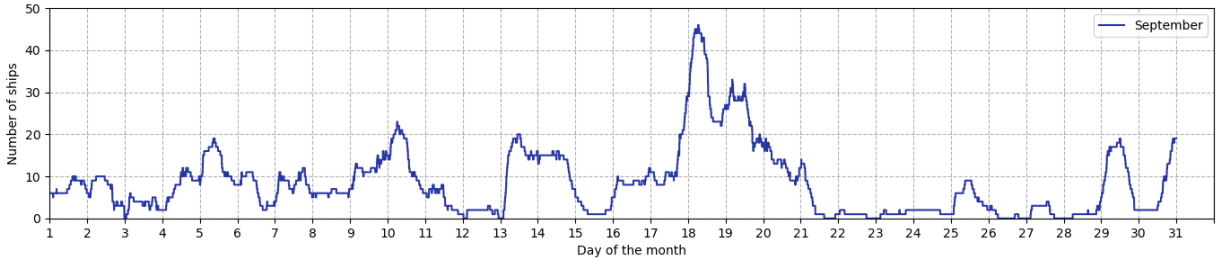

Figure A.9: Amount of fishing effort in September from the AIS data

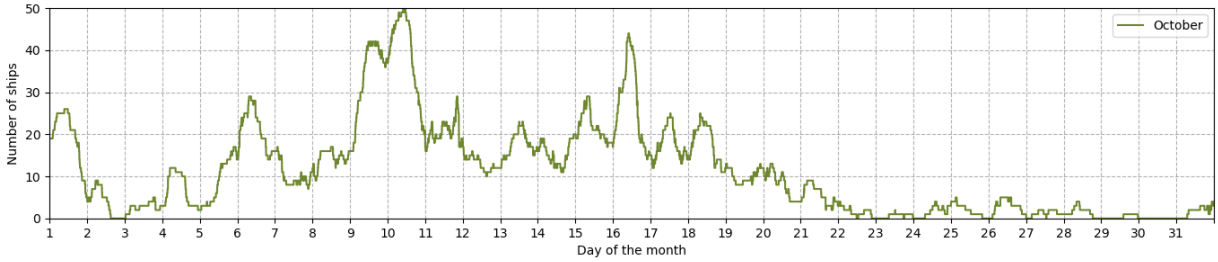

Figure A.10: Amount of fishing effort in October from the AIS data

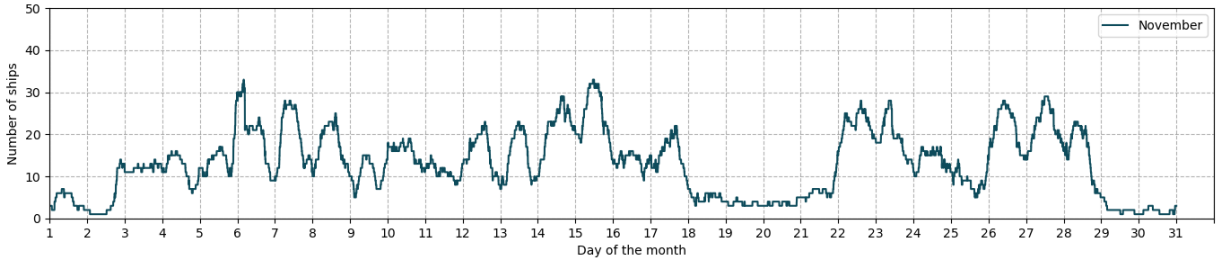

Figure A.11: Amount of fishing effort in November from the AIS data

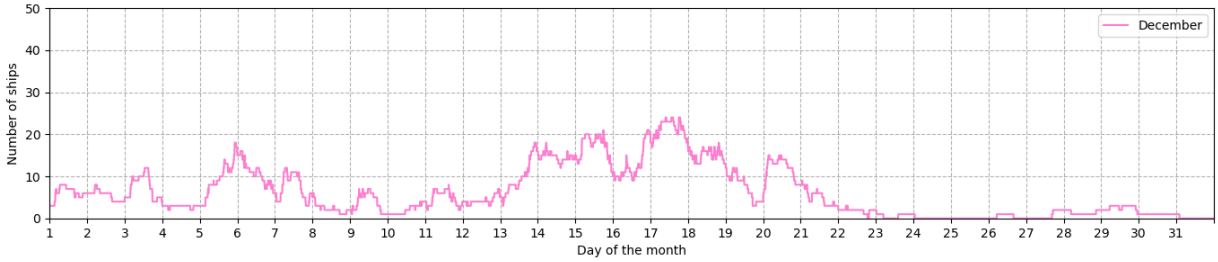

Figure A.12: Amount of fishing effort in December from the AIS data

Appendix B: Fishing effort by country

Among the 5,304 individual fishing boat trajectories identified from AIS data in 2018, the majority were taken by vessels sailing under a Polish flag as is shown in the figure below. To obtain this information, MMSI numbers from the AIS messages were matched with MID table from International Telecommunication Union (ITU).

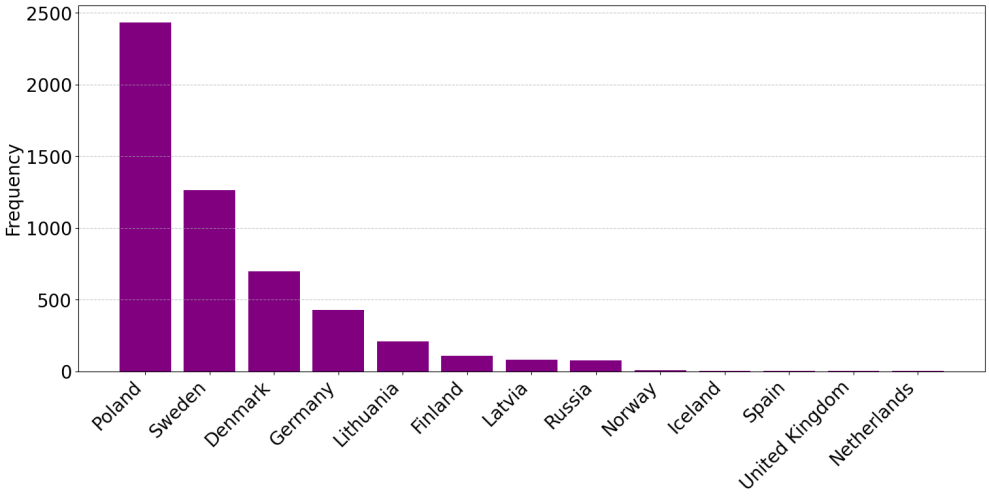

Figure B.1: Fishing effort by country from the AIS data