A COMPREHENSIVE ANALYSIS AND NOVEL METHODS FOR ON-PURPOSE AIS SWITCH-OFF DETECTION

by

Matthew R. Brousseau

Submitted in partial fulfillment of the requirements
for the degree of Master of Computer Science

at

Dalhousie University
Halifax, Nova Scotia
September 2021

Table of Contents

List of Tables

List of Figures

Abstract

Maritime security, as well as monitoring of illegal and illicit activities such as smuggling and Illegal, Unreported and Unregulated (or IUU) Fishing are reliant on Automatic Identification System (AIS). As AIS is a self-reporting system, determining where and when suspect activities are taking place is difficult when vessels are simply able to decide to switch off their AIS device on purpose. The AIS on-purpose switch-off problem attempts to differentiate between typical loss of signal between a vessel and base station with an intentional shut down of the on-board AIS transponder.

Previous works, while few, have attempted various strategies for solving this problem. From statistical analysis, to machine learning, to reconstructing message signal strength from base stations. All of these works, however, have some area in which they lack.

In this work we provide a comprehensive analysis and propose improvements to what is considered the state-of-the-art approach to solving this problem. The improvements on the previous works' extrapolation algorithm showed a statistically significant increase to the F1-Score ranging from 0.03 to 0.051 at an α of 0.05 as well as an almost halving of the standard deviation. Digging deeper into the results of the various parameters and the relationship between their values and the overall performance provides even more promising results, showing exactly where it is that the original baseline fails and how the newer approaches proposed here are able to compensate for those original deficiencies.

List of Abbreviations and Symbols Used

AIS Automatic Identification System

ARR Averaged RSSI Raster

AUC Area Under the Curve

BS Base Station

COG Course Over Ground

DBSCAN Density-Based Spatial Clustering of Applications With Noise

GAM Generalized Additive Model

IMO International Maritime Organization

ITU International Telecommunication Union

IUU Illegal, Unreported and Unregulated (Fishing)

MMSI Maritime Mobile Service Identity

ROT Rate of Turn

RSSI Received Signal Strength Information

SOG Speed Over Ground

SVM Support Vector Machine

T-AIS Terrestrial AIS

UTM Universal Transerve Mercator Coordinate System

Acknowledgements

I would like to begin by thanking Dr. Stan Matwin for all of the support that he has shown me over the years. He is a professor for whom I have the deepest respect and who has helped me greatly over the years. His care for the institute, as well as his students and their progress has helped push me to do better.

I would also like to thank Dr. Amilcar Soares. I would not be here right now if it was not for him. Dr. Soares has been there for me academically, professionally, and personally for years and I am eternally grateful for all that he has done.

And finally, I would like to thank my wife Ruusu-Maaria Brousseau for being with me through this entire journey. World events have made it that it was not the easiest time to go through such an ordeal, but through all of that, she has been here to support me and help me complete this work. At the end of it all, this feels as much our work as it does mine.

Chapter 1

Introduction

The maritime domain is of critical importance globally. From international trade to the transportation of people, the ocean connects the world. Over 90% of global trade is still done by ship [25], while over 400 million passengers travel by water every year in Europe alone [11]. In addition to the transportation of goods and persons, the ocean is the primary source of protein for over 4.3 billion people [12]. Therefore, the protection of goods, people and the conservation of fish stocks are of the utmost importance. Great effort has been made to ensure the safety and efficiency of ocean-based travel and the regulation of fishing in national and international waters. One of the significant components of this safety and regulation is the Automatic Identification System (AIS). The AIS is a self-reporting system that broadcasts static and positional data from vessels at fixed intervals [25]. This allows for governing authorities to keep track of global as well as local marine traffic for safety, efficiency, and environmental protection.

Though AIS was originally utilized as a collision avoidance tool, its usage has evolved with the increase of computing power and availability. Today, AIS is used by various companies, government and non-government organizations to perform several tasks. These include detecting piracy, illegal and unregulated fishing, smuggling, illegal immigration, pollution, and terrorism. However, as the developers of the AIS protocol did not originally envision these use cases, it was not built to address these tasks to its best. One major concession that was made in the development of AIS was that it was created to be a self-reporting system. This means that control of the transponder is entirely within the realm of the crew physically aboard the vessel. Because of this, those aboard can turn the transponder on and off whenever they see fit and can change the values that are being broadcast from their transponders, such as their identification numbers, destination, and identification characteristics of the vessel.

There are various reasons why these concessions were made in the development of AIS. Vessels traveling in areas of the world known for piracy wanted the ability to hide their location to keep themselves safe, fishing vessels did not want to broadcast fishing locations that their families may have used for generations, and navies did not want other nations to always know the location of their ships. Though many of these reasons are valid, it has created both issues and opportunities to research potential use cases of AIS.

As vessels are able to turn their transponders on and off whenever they see fit, smuggling and piracy become much more challenging to detect. What if we were able to detect with a high degree of certainty when a vessel turns off its transponder on purpose? At first, this may seem a trivial task. One could simply check to determine when a vessel is no longer sending a signal and then raise an alarm. While this would solve part of the problem, it would not work in all scenarios for numerous reasons. The main problem with this approach is that signal loss is not an uncommon occurrence in AIS. Signals are sent from AIS transponders that are held aboard vessels and received by either satellite or terrestrial-based stations. Whether said stations receive the signals depends on numerous variables. These include the weather, coverage of a particular geographical area, current formation of the satellites, as well as vessel congestion. Each of these can affect the propagation of the signal and prevent it from being received.

Currently, there are a number of previous works that have attempted to solve this problem. While some have generated some positive results, the research is still in its infancy. This problem has not been thoroughly studied and each attempt has come with a number of pitfalls. Some examples of those pitfalls include requiring previous routes of a given vessel to compare to, requiring models to be rebuilt if there is ever a new vessel as well as using only satellite or terrestrial-based AIS messages, and being inflexible with using the other.

One of the most recent attempts to solve this problem [19] has provided what is, in our opinion, the best overall method at solving this problem. The work of [19] uses a general extrapolation algorithm to predict vessel positions once a signal dropout occurs. Their algorithm, however, does not work well in all situations and makes several assumptions that turn out to fail once put to the test. Therefore, this work employs some other extrapolation algorithms that improve both the accuracy and the variability of the models.

1.1 Research Questions

The focus of this work is not only to be able to detect when a signal from a vessel is lost but to be able to differentiate, with some degree of certainty, between signal loss and an intentional switch-off of an AIS transponder. Therefore, the following research questions are addressed in this work:

- Given an AIS dataset containing various gaps throughout the trajectories, are we able to differentiate between an intentional switch off of the AIS transponder and signal loss from natural causes?
- Can the current state-of-the-art implementation of the algorithms used to solve the AIS intentional switch-off problem be improved by using other extrapolation algorithms that are better able to take the variability of a vessel's heading into consideration?

1.2 Contributions

The contributions of this work in regards to the on purpose AIS switch-off problem are the following:

- We provide an in-depth look at the on purpose AIS switch-off problem and literature. We go step-by-step through each previous work that has been done in the field, the strategies implemented, as well as the strengths, weaknesses, and areas of improvement for each.
- We propose a new algorithm to work as an improvement on the most recent and state-of-the-art on purpose AIS switch-off strategy [19]. These improvements focus on the extrapolation algorithm employed. The new extrapolation algorithms are better able to take into account the current behaviour of a vessel and can better predict the locations of vessels that are not travelling in a straight line as previously done.

- Since real data for detecting intentional AIS switch-off is not publicly available, we carefully design experiments to artificially create such situations and test the baseline algorithm and our modifications. The results indicate that our modifications indeed improve the the detection of on purpose AIS switch-off AIS.

1.3 Thesis outline

This thesis is organized as follows. Chapter 2 provides a background to various terminologies and technologies that are vital to this work. Chapter 3 looks at previous works that have attempted to solve the AIS Switch-Off problem. These previous works are broken down into various categories based on the way in which they approached the problem. Chapter 4 exhibits the methods employed to better solve the on-purpose AIS switch-off problem. The models, how we have evaluated them, and the data used are discussed in this Chapter. Chapter 5 presents the experimental setup, the results of both the baseline experiments as well as our own. Finally, Chapter 6 discusses the limitations of our approach, as well as potential future work in the field.

Chapter 2

Background

This chapter details the technology of the Automatic Identification System, how it works, its issues, as well as current and future research in the field. Section 2.1 shows an overall look at AIS and how it works. It also provides a background into the types of messages, the type of information provided, its original usage, as well as current research topics in the field. After, Section 2.2 looks at the problem of the intentional AIS switch-off, why the problem is even possible, as well as some other pitfalls in AIS research.

2.1 Automatic Identification System (AIS)

The Automatic Identification System (AIS) is a standardized and unencrypted self-reporting maritime surveillance system. The protocol functions by sending 1 or more of 27 different types of messages from an onboard AIS transponder at regular intervals.

The intervals at which messages are sent are based on the current state of the vessel. Vessels that are anchored or moored will send their messages every 3 minutes, while vessels travelling at a top speed will send messages every 2 seconds [1].

AIS messages fall into one of two groups; they can either be static or dynamic [24]. Dynamic messages refer to the changing state of the vessel from signal to signal. These messages include those that provide values for the Speed Over Ground (SOG), the Course Over Ground (COG), the Rate of Turn (ROT), as well as the current positional information for the vessel (latitudes and longitudes). The static messages contain values for the vessel that do not change every few minutes. These values include the Maritime Mobile Service Identity (MMSI) number, the International Maritime Organization (IMO) number, the vessel's name, call sign, vessel type, vessel dimensions, and destination.

VHF Radio messages from vessels are sent from transponders located aboard the vessels and then picked up by either satellite or terrestrial networks, allowing for maritime traffic to be understood to a greater extent and analyzed. While capacity has increased globally to collect this data, it is still a common occurrence for signals from vessels to be lost. Terrestrial base stations are limited by their physical range, while satellite AIS receivers are limited based on their position globally.

Originally, AIS was developed to prevent vessel collisions. By tracking a vessel's location, Operators would be able to help guide vessels down difficult waterways safely. Because of this, many of the current uses of AIS and much of the research were not thought of as a possibility when AIS was being developed. The technology has now evolved beyond its original purpose, and various fields of research have grown from it. Commonly uses and fields of research now include vessel tracking [23, 29], vessel behaviour analysis [26, 25], data fusion and enrichment [8, 34, 16], network analysis [37, 6, 36], knowledge extraction [30] and anomaly detection [24, 18, 22] and visual analytics [17, 2, 3].

2.2 AIS on Purpose Switch-Off

One of the major issues surrounding AIS is that it is a self reporting system. This means that while having an AIS transponder aboard a vessel is compulsory, having it turned on is not. While every vessel over 300 gross tonnes must have the hardware [27], they are not required to actually share their location and are able to turn it off and on whenever they desire. Often, there are valid reasons for this. In many parts of the world, piracy is still a significant issue and keeping an AIS transponder turned on in these regions could be considered putting a target on one's back. As there are many websites where all AIS messages and vessel positions can be viewed for free, many fishing vessels are hesitant to turn on their AIS transponders when they are fishing. The reason for this is that there are often areas that they have fished for generations and they do not want to broadcast those places to the world. The self reporting aspect to AIS, however, has a number of issues with it.

Some of the major issues regarding AIS include:

- Data can often be incomplete or misleading. As vessels can turn off their transponder, critical information can be missing. As such, illegal activities can be easily hidden.

- As a lot of the static data about a vessel is manually inputted by the crew, data can be erroneous. This can be accidental, intentional, or default values may simply not have been changed.

- AIS data is easily spoofed. There are examples on the internet that can be found where AIS signals have been spoofed to create drawings in the ocean.

While the second and third problems are serious, and research has gone into both, the focus of this paper is the first problem. As vessels are able to switch off their AIS transponders whenever they desire, being able to detect illicit activities such as illegal smuggling and fishing become much more challenging. While satellite imaging could be used in order to detect vessels that attempt to hide their signals, this strategy is more complicated and expensive. AIS data is abundant, and being able to real time, up to date, and ubiquitous data, rather than imaging data that is dependent on satellite positions at a given time is preferable.

When looking at an AIS dataset in which a vessel has turned off its transponder, the loss of signal will be seen as a gap in the timestamp of the data. Attempting to find these gaps is trivial. The challenge is that not all gaps in AIS transmission data are significant. There are numerous external variables that can cause a signal dropout where the vessel is not engaging in illicit activity. The challenge is not only to determine whether there has been a loss of signal from a vessel but to also be able to differentiate between a loss that is caused by external events such as lack of coverage by satellite or base station, interference by weather, poor quality of equipment aboard the vessel, or interference by congestion of other vessels.

Chapter 3

Related Works

While both the technology of AIS and the problem of AIS switch-off are not new or novel, little research has been done in finding a solution to the problem. There have been, however, in recent years a number of proposed frameworks put into place in order to tackle the problem. While each solution is unique in its own way, there are commonalities between them. These previous works have therefore been divided into various groupings in order to directly compare the strategies employed by the authors of each. The rest of this section will discuss the papers, their strategies, and various comparisons of strategies and the papers that fall into each.

3.1 A description of the analysed aspects

For the comparison of previous works, common aspects have been produced in order to compare and contrast these papers. These aspects are compared to each other and papers are categorized based upon them.

3.1.1 Real-Time vs Historical

Real-Time vs Historical refers to the data that is required in order to generate the model and its ability to handle new data. Some models take large amounts of historical data and generate results that they are able to present. If new data were to arrive, they would need to train the entire model again in order to generate results. While they may be able to find anomalies, they are unable to handle new data in real time. Other models are able to do this. They are able to create fairly complex models which are able to find anomalies but are also able to take in real time data in order to determine whether anomalies are ongoing.

3.1.2 Stream vs Batch

How do models handle new data? The way that this occurs is divided into two different groups. The first is the group of models that are able to take in live streams of data (in the case of the paper that does this it is actually able to handle distributed streams on a cluster of machines) and consistently provide a response as it happens. The second group consists of those that perform batch processing. As mentioned in Section 3.1.1, some models are not able to handle new data without rebuilding their entire model. Those papers have been also classified within the batch group as technically they do provide batch processing, just where the batch is the entire dataset. Others, however, are able to take in new data that comes in real time and provide batch processing on it in order to determine anomalous behaviour.

3.1.3 Spatial Density vs Vessel History

How are we able to differentiate between a vessel turning off their AIS transponder and what could simply be a loss in transmission? This is not an easy question to answer. There are many reasons as to why there could be signal dropout. These reasons include, but are not limited to poor weather, vessel congestion in a given area, as well as possible equipment failure aboard the vessel or at the receiver. As such, finding AIS switch-offs is not as simple as finding the gaps in the data and reporting them, false positives are a real challenge in this problem.

In order to determine whether there is an AIS dropout or whether there is potentially something more nefarious, there are a couple of strategies implemented. For this Thesis, they are broken into two groups based on the type of data that they use.

The first group, which for the purposes of this Thesis are referred to as those that utilize Spatial Density, are those that divide the ocean up in sections (typically squares of some kind) and when there is a potential dropout, they check the general area where either the dropout occurred or where they believe the vessel was going. If there happen to be signals coming from other vessels in those areas, then the likelihood that the missing signal is due to something such as poor weather would be lower. These strategies are based simply on vessel congestion and the signals of other vessels in a given area, they do not rely on any information about the vessel in question.

On the other hand, certain strategies look at the historical behaviour of the vessel. Some may attempt to determine whether gaps are a common occurrence in their history, they may use the length of the gap, or other characteristics of previous voyages. The main characteristics of Vessel History strategies are that they both analyze the trajectories of the vessels themselves and utilize previous trips of a given vessel to determine future behaviour.

3.1.4 Terrestrial vs Satellite

When AIS first became a standard, transponders were put aboard vessels and the signal was captured by land based AIS receiving stations. While these stations could be reliable, they were limited by their need to be placed ashore, making open ocean monitoring more difficult. At the beginning, AIS was better used near ports, rivers, etc. However, satellites have changed the way that AIS operates. More and more AIS satellites are being launched, providing better coverage across the globe. While satellite AIS provides greater coverage, the data is more sparse than that of terrestrial AIS [33]. Because of these characterization differences, some models are better suited for one or the other, while for some there is no real distinction and both are utilized.

3.1.5 Ranking

All of the papers in this review implement some form of anomaly detection. As we do not have an abundance of labelled data, we often treat this problem as either a form of one-class classification or some kind of outlier detection. We have many instances of vessels that continue to run normally, with no interference in their AIS signal, but few exceptions. Because of this, the strategy typically implemented is more about finding new data that does not conform to the norm and presenting that as a potential problem.

Here, however, is where a number of papers differ. Some papers perform a hard classification. They state that there is an outlier or anomaly and tell the user that they have found an AIS switch-off. Other papers go a little further. Rather than simply provide a result, they attempt to rank vessels on either the likelihood or severity of that anomaly. By providing a ranking system, they tell the user not only that there are problems, but which problems are the greatest and which need to be addressed first. This has the potential to give much more information to the user though they are more difficult to validate as finding some form of ground truth for a ranking system in this context is incredibly difficult.

3.1.6 A summary of the aspects evaluated

In order to better compile each related work into a more digestible format, the works have been categorized based on their strategies. Tab. 3.1 contains a summary of the aspects of all works evaluated in this work. The most considerable commonality regarding the aspects is in the Stream/Batch category. Most of the papers have worked with batch processing. They are not able to take in new data or vessels without rebuilding the models. These approaches work for both academic research and for various forms of post-hoc analysis where reacting in real time is not considered essential. However, for practical real-world use, users would want to determine whether some switch-off event had occurred in real-time. Working on live AIS streams, while more difficult, is much more useful in actuality, and only one paper has attempted it. With some thought, some of these papers could be transformed to self-update and adapt at real time, though that is a work outside the scope of this one.

The other categories are fairly evenly split, and research has been done to some degree in all of them. Most of these categories do not have one value that is inherently better than the other (aside from Terrestrial/Satellite where all else is equal, a model that can utilize any form of AIS data is superior to one that is limited by its source).

	Real-Time/ Historical	**Stream/ Batch**	**Spatial Density/ Vessel History**	**Terrestrial/ Satellite**	**Ranking**
Mazzarella et al. (2016) [25]	Real-Time	Batch	Spatial Density	Terrestrial	No
Mazzarella et al. (2017) [24]	Real-Time	Batch	Spatial Density	Terrestrial	No
d'Afflisio et al. (2018) [9]	Historical	Batch	Vessel History	Both	No
Ford et al. (2018) [13]	Historical	Batch	Vessel History	Satellite	Yes
Shahir et al. (2019) [33]	Historical	Batch	Vessel History	Terrestrial	Yes
Kontopoulos et al. (2020) [18]	Real-Time	Stream	Spatial Density	Both	No

Table 3.1. Comparison between all analysed works.

3.2 Literature review

3.2.1 AIS Reception Characterisation For AIS On/Off Anomaly Detection

AIS uses VHF Radio signals which are picked up by either satellite or terrestrial networks, allowing for global marine traffic to be understood to a greater extent and analysed. There are, however, a number of issues with AIS which have been mentioned previously. AIS results can be spoofed. By this, we mean that vessels can broadcast falsified or erroneous data. Vessels also have the power to turn off their transponders so that they cannot be tracked. This is a huge problem in illegal activities such as smuggling or illegal fishing.

The goal of this paper[25] is to find a solution to the AIS switch-off problem. In order to do this, the authors asks a simple question: How can we differentiate between a communication dropout and an anomalous event (for example an AIS switch-off)? As the probability of communication dropout is dependant on other factors (for example the position of the vessel in relation to base stations or the congestion of vessels in a given area), this is a non-trivial task.

The authors propose an architecture that will account for electromagnetic propagation in order to determine whether the event is anomalous or whether it is a case of normal dropout caused by external events. This paper utilizes historic AIS collected from various Base Stations and attempts to use the Received Signal Strength Information or RSSI. The idea behind the RSSI is pretty straightforward. It is the measure of how strong the signal is when it is received by the base station.

In order to test this, the authors use the AIS data that has been collected from a single base station as training data. This training data is used to generate a Knowledge-based model for base station coverage. This is done by using the relationship between signal strength and the distance between a vessel and the base station. Using this training data, two values are generated. The first is the track that the vessel takes which is compared to the RSSI at a given point. The second is a partition that is created where on one side of the track represents normal behaviour and the other represents an anomaly.

Fig. 3.1 represents a visualization of the system. What we have discussed so far

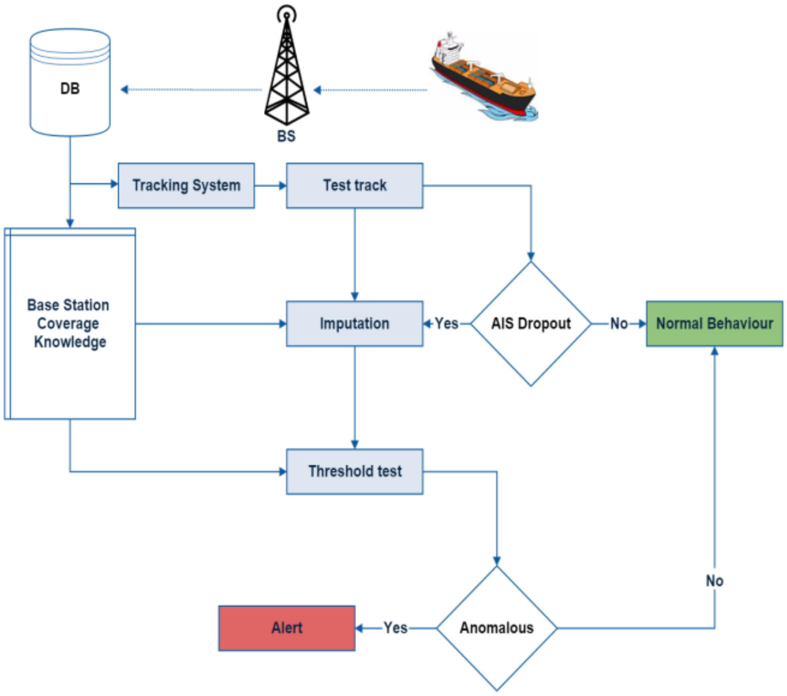

Figure 3.1. Architecture of strategy implemented by [25]

is the Knowledge based model which is represented as the Base Station Coverage Knowledge. This generates the threshold test as can also be seen in the flow chart. There are, however, other parts to this architecture. First we have the tracking system which, as the name suggests, keeps track of the current trajectory of a vessel.

This then goes into the test track. The test track periodically checks the tracking system for dropouts. If there is no dropout, it is marked as Normal behaviour. However, if it is determined that there was a dropout, it moves to the Imputation step. In the imputation step, it is determined what the RSSI value for the vessel should be based on its current position. That reconstructed RSSI is then given as input to the Threshold Test which determines which side of the threshold the vessel falls on.

$$Range \approx 4.131 * (\sqrt{h_T} + \sqrt{h_R}) \tag{3.1}$$

As AIS was originally developed for collision avoidance, the range of the signal was not the greatest consideration. Because of this, the range is fairly limited. The typical range of an AIS signal is around 40 nautical miles, though this is dependant on weather. There is a formula for estimating the range of an AIS signal which is given in Eq. 3.1 where h_T and h_R represent the heights of the antennae used in transmission and reception respectively.

The next goal is to be able to generate an estimate for the RSSI. In order to do that, an estimate of the power of the signal between the vessel and the Base Station needs to be calculated. However, there are numerous variables that can have an effect on these calculations.

Firstly, diffraction of the radio waves around the curvature of the earth can create an amplification effect and extend the range of the signal Because we are dealing with open ocean, the signal can travel in various ways and across various paths which can create variability in the RSSI. This is affected by how the antenna is mounted, the wind speed and the physical characteristics of the terrain of where the Base Station is located. And finally, depending on where in the world it is located and the weather, the refractivity of the air can also extend the transmission range.

$$P_R = \frac{P_T * G_T * G_R}{L_{FSL}} \tag{3.2}$$

$$L_{2-ray} = (\frac{\lambda}{4\pi * d})^2[1 + \rho * exp(j * k\frac{2h_R * h_T}{d})] \tag{3.3}$$

In order to generate an estimate of the power of the signal, the Friis equation can be used which is defined in Eq. 3.2. In this first equation, the Power of the receiver is equal to the Power of the transmission times the antennae gains in both the transmitter and receiver, divided by what is called the Free Space Path Loss which is defined Eq. 3.3.

In Eq. 3.3, we have ρ which represents a reflection coefficient from the ocean and k represents the wave. Unfortunately, the rest of the variables in the equation were not defined in the paper.This equation works assuming that we are working on a plane, which is fine given that the distance between the Base Station and the Vessel are not terribly far.

$$d_c = \frac{4 * h_T * h_R}{\lambda} \tag{3.4}$$

How do we define "not terribly far"? The critical distance is defined as the distance until which, the Friis model can be used and after which the Round Earth Model must be used and is shown in Eq. 3.4. This is calculated by multiplying the heights of the transmitter and receiver by 4 and dividing by lambda which represents the wavelength of the signal.

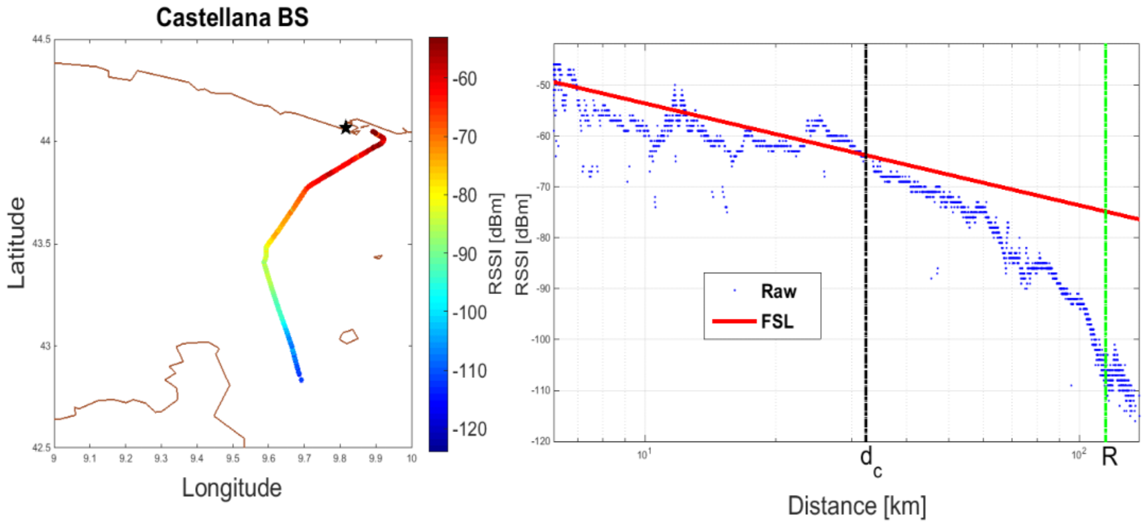

Figure 3.2. Signal Example [25]

A visual example of a single vessel is shown in Fig. 3.2. Shown on the left is a Base Station represented by a star with the coloured line representing the trajectory of a vessel. The colour of the trajectory represents the RSSI of the vessel at that given point. On the right is similar. The x-axis shows the distance between the vessel and the base station while the y-axis shows the RSSI. On the far right, we have a green line which represents the range of the signal between the vessel and the base station, while the d_c value represents the critical distance value, or the value after which a round earth model needs to be implemented. Finally, the red line represents the loss function that was found using the Friis equation. After the critical distance line, we are able to see why a different model is required. This paper does not implement the entire architecture itself. This is rather a more general overview of the approach that the authors would like to take. They do not explain what kind of threshold test they want to use (though they do in their next paper, described in Section 3.2.2) but rather show how they plan to generate the features and model that they are going to use.

3.2.2 A Novel Anomaly Detection Approach To Identify Intentional AIS On-Off Switching

A Novel Anomaly Detection Approach To Identify Intentional AIS On-Off Switching [24] is a continuation of the paper described in Section 3.2.1.

The basic idea of the paper was that the authors wanted to use the Received Signal Strength Information (or RSSI) in order to differentiate between signal dropout and anomalous events. The RSSI is a measure of how strong the signal from the vessel is when it arrives at the base station. The idea was that if the signal "should" have been strong, then the chance of the base station not picking up the signal should be low and therefore it could be classified as anomalous. The previous paper was mostly theoretical, with a number of examples included in order to illustrate what it was that they hoped to accomplish. This paper on the other hand is an implementation of those ideas.

In order to build this architecture (visualized in Fig. 3.3), there are two models created. The first is used to test normality at the base station. While the second is to test for normality at the vessel signal.

If the model detects a potential dropout, represented as a time threshold where no new signal has been received, it then moves to the Feature Extraction phase. At this point, the position that the vessel should be is interpolated. Once the interpolation is calculated, the distance between the vessel position and the Base Station is calculated. With that distance, the potential RSSI can be calculated with the equations mentioned in Section 3.2.1.

Once the Feature Extraction module has reconstructed the message of the potential dropout, the message is then passed to the two normality models, one to compare the reconstructed message to the Base Station's model and one for the Single Vessel model. Likelihood values are generated by each and passed to a threshold test which then determines whether the missing message was due to dropout or due to an anomalous event.

In order to perform Feature Extraction, five features are extracted or generated from historical AIS data. They are: Vessel Position (in Latitude and Longitude), The distance between the vessel and the base station, the Reconstructed RSSI value, the antenna heights for both the base station and the vessel, and finally the length of the

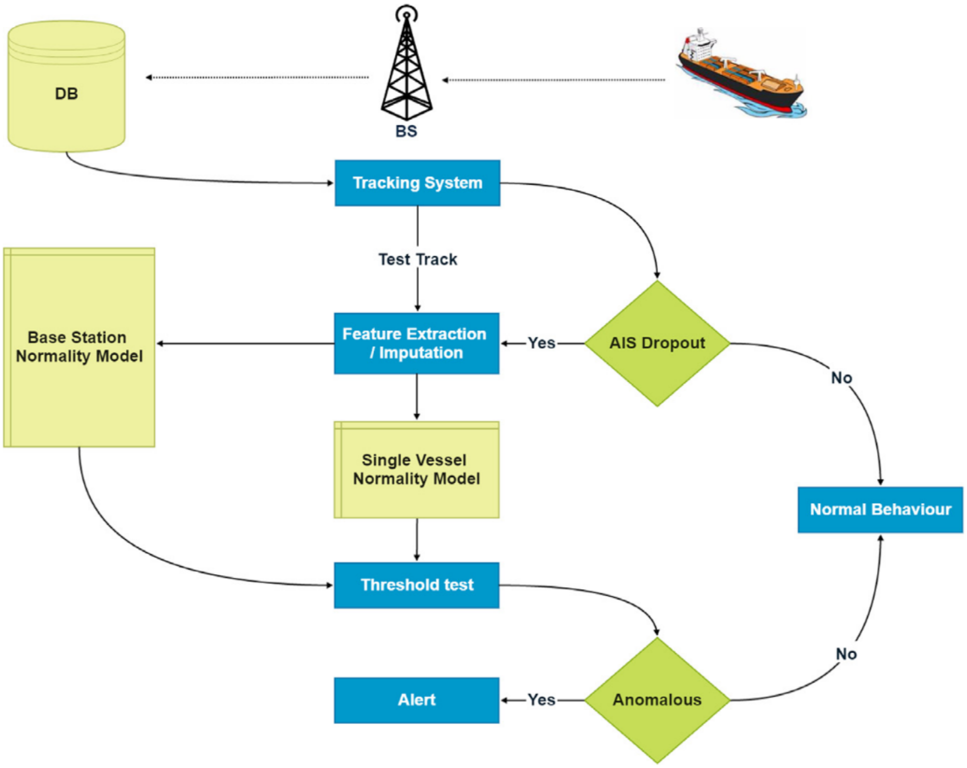

Figure 3.3. Architecture of strategy implemented by [24]

vessel. The vessel position, the Reconstructed RSSI and the vessel length are used for the Base Station's Normality model while the distance and the antenna heights are used for the single vessel normality model.

For the Base Station's normality model which attempts to determine whether the dropout given the Reconstructed RSSI value should be classified as anomolous, two algorithms were tested. The first is a One-class SVM while the second is called the Averaged RSSI Raster (or ARR).

When performing classification, one can often run into the problem where one is only able to obtain data for a single class. In this case, there are algorithms that are able to perform a One-class classification. Essentially, the way that it works is that the training data is passed in and a decision boundary is generated around the training data that determines what is "normal behaviour". If then, some result arrives that is outside of that decision boundary, then it can be flagged as an anomaly. So in this example, the one class SVM is trained with the spatial coordinates, the RSSI and the length of the vessel.

The idea behind the Averaged RSSI Raster, as the name suggests, is to treat the Base Station model as a geo-referenced raster. A raster is a rectangular grid that represents data over a geographical space. The raster is divided into rows and columns with cells throughout. Each cell in the raster has a value that is specific to the geographic area that it covers. In this example, each cell holds the median value of the RSSI received from that area. Median is used as this type of data can be noisy and the median is resistant to outliers. These median values are then normalized so that each cell holds a value between 0 and 1. The generated RSSI value is then compared to the corresponding normalized median value of the given square.

The Base Station Normality model is what determines the likelihood results. With the One-class SVM, the positional data, the RSSI and the length of the vessel are passed in as test data. Once they are passed in, the SVM will output either a -1 or a 1. -1 means that the test feature does not belong to the same set as the training data. A value 1 on the other hand means that the test feature belongs in the same set as the training data. In this case a 1 is what we will define as an anomaly. The reason for this is that since it is a 1, the value should have appeared but didn't. Therefore, it is an anomaly.

This output is then transformed into a probability using a type of soft classification that is referenced in the paper. Secondly, we have the ARR. Unfortunately, the authors did not really go into detail on how they performed these calculations. What they did say though is that the test value is compared to the corresponding values in the raster cell. The likelihood is then calculated based on the median values in that raster cell. The authors then move onto the Single Vessel Normality model. The reason that two models are needed is that the way that the signal propagates from the vessel to the base station is different between vessels. Therefore, it needs to be taken into consideration. A lot of how this was accomplished was discussed Section 3.2.1, though certain aspects were made more explicit in this paper. Previously, we had discussed two functions. The first was called the Free Space Path Loss. This function allows us to estimate the power of the signal given that the signal is travelling across a flat plane. Up until a certain distance we can use this. The point after which this algorithm starts to break down depends on other factors. Namely, the height of the antenna on the vessel and the base station, as well as the length of the wavelength being sent. Using these factors, we can generate a critical distance value that will determine the distance after which the function must change. After this critical distance value is reached, a Round Earth Loss model must be used. For that model, the authors used the ITU-R P.1546 which is the ITU recommended algorithm. The algorithm itself is not defined in this paper, though it is referenced.

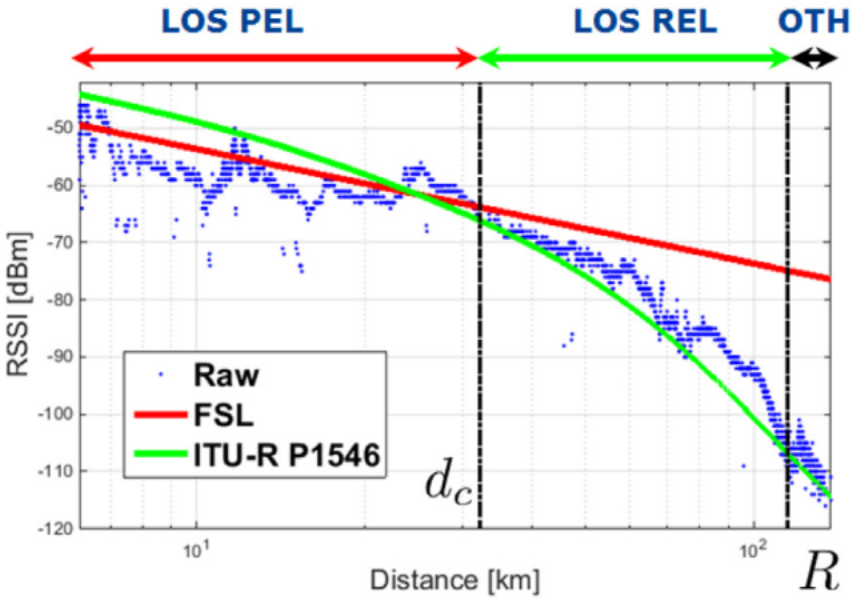

Figure 3.4. Comparison of Loss Functions given distances [24]

Fig. 3.4 shows how the two algorithms perform given the distance between the vessel and the base station. It is similar to Fig. 3.2 though it now includes a comparison between the Free Space Path Loss and the Real Earth Loss models. The red line represents the Free Space Loss while the green represents the Real Earth Loss. The first vertical black line represents the critical distance after which the function used changes and the second vertical line represents the maximum range of the signal. According to this, the strength of the signal does seem to change after that critical distance and the Real Earth Loss does seem adequate in compensating for that change.

$$\begin{gathered} d_d \leq d_c \Rightarrow \alpha \leq r_v < 1 \\ d_c < d_d \leq R \Rightarrow \beta \leq r_v < \alpha \\ d_d > R \Rightarrow 0 \leq r_v < \beta \end{gathered} \tag{3.5}$$

In order to actually calculate the Normality Likelihood from the Single Vessel model, we need to look at three different scenarios which are defined in Eq. 3.5. The first is the scenario where the distance between the vessel and the base station is less than the critical value.

In this case, the risk value will be set to be between α and 1 where alpha equals the RSSI at the critical distance divided by the maximum RSSI. The second scenario is for where the distance between the vessel and base station is greater than the critical distance, but less than the maximum theoretical range of the signal. In this case, the risk value will be between β and α where β equals the RRSI at the maximum distance divided by the max RSSI value. And finally, we have the scenario where the distance between the vessel and base station is greater than the maximum theoretical range. In this case, the risk value is set to be between 0 and β. Unfortunately, the paper does not go into more detail to determine exactly how that value is calculate other than to give the ranges of them dependant on the scenario. It seems as though it is calculated by taking the interpolated RSSI and dividing by the max RSSI though this is not stated explicitly.

So, now that the models for the Base Stations and the single vessels have been defined, we now are able to discuss how the computations are performed. Firstly, a value of σ is set. σ represents the number of minutes between each iteration of the algorithm. Once σ minutes arrives, the algorithm checks to see if all the vessels that were in range of the base station previously have sent a message.

If so, they are ignored. However, if a vessel had not sent a message during the previous interval, it will be checked for dropout. Once this happens, the feature extraction algorithm mentioned previously will run in order to obtain the positional data as well as the estimated RSSI and the other features mentioned. Once these features are generated, they are passed to the normality models for the Base Station (so either an SVM or the ARR) or the Single Vessel model (which uses one of two loss functions depending on where the vessel is in relation to the base station). These models will then generate two normality likelihood values which we have previously discussed that will be passed on to the next part of the algorithm.

In order to make the decision on whether the dropout event is due to actual dropout or whether it is an anomalous event, the user first decides on a threshold after which the event will be determined to be anomalous. Once this value is decided, the risk scores for both the single vessel and base stations are multiplied together and if the that final risk score is greater than the threshold value, it is classified as an anomaly. Else, it is not.

In order to perform their experiments, a single day's worth of AIS data was used for both training and testing. In order to generate both true and false anomalies, points were taken out of the test data at various points of the track. Unfortunately, I do not know how much data was used in training and testing. I do know now how many anomalies were present in the data and I don't actually know the results of the experimentation. For our results, we were given a ROC curve for the results using the one class SVM and another for the ARR. The first had an AUC score of 0.99 while the ARR algorithm had an AUC score of 0.97 which implies that the SVM worked better. False positives were said to be below 10% and detection rate was said to be above 90%. However, there were no tables showing any sort of results.

3.2.3 Detecting Suspicious Activities At Sea Based on Anomalies in Automatic Identification Systems Transmissions

The goal of [13] is to expand upon previous attempts to solve the AIS off problem by allowing for the usage of Satellite AIS data. In order to accomplish this, the authors used a Generalized Additive Model.

These Generalized Additive Models are used because they work when there is no known underlying function in the data. These GAMs are used to detect spatial and temporal patterns of the frequency of vessel transmissions as well as receiver coverage. By doing this, they are able to calculate expected transmission rates in a given region. The idea behind the Generalized Additive Models is to find a way to model the relationship between a predictor and response. In this example, we are looking at a binomial GAM where we are testing the probability of the detection of a successful AIS transmission. When a binomial distribution is used, we are talking about a series of Bernoulli events.

In this example, the event is whether a transmission will be sent during a one hour interval and the result of the event is either a successful transmission or a failure of receipt. There are two intervals being used. The first is the summary of the data as a whole which is divided into 24 hour trajectories. The second is the gaps in the trajectories themselves, which are calculated using 1 hour intervals. A gap is then determined by a series of 1 or more missed hourly intervals in the 24 hour trajectory interval.

$$\begin{aligned} PoG &= p(1-p)^L \\ StdPoG &= \frac{PoG_i}{\sum_i PoG_i/N} \end{aligned} \tag{3.6}$$

The assumption that is being made is that AIS switch-off will result in long runs of non-transmission gaps. In order to test this, the equations defined in Eq. 3.6 are used. The first equation is called the Probability of Gap function. This equation will take the length of the gap, calculated as the number of consecutive one hour gaps which resulted in no transmission, and calculate the probability of that gap occurring. The probability of transmission value mentioned previously is used here and is multiplied by $(1-p)^L$ where L equals the length of the gap.

In order to test their statistical models, various gap and trajectory lengths were artificially added to the data. Gap lengths of 1, 2, 4, 6, 8, 10, and 12 hours were used. For trajectory lengths, 1, 1.5, 2, 3 and 5 days were used. For every combination of both trajectory length and gap length, 100 simulations were run for a total of 3500 simulations. Unfortunately, in regards to their Experimental Setup, little was said.

What is known is that the data came from the Arafura Sea region. This region is the sea that is bordered by Indonesia, Papua New Guinea and Australia. It is a major shipping lane for that part of the Pacific Ocean. Unfortunately, we do not know how much data was used or the number of messages. It is mentioned that around 400 vessels and 2000 segmented trajectories were tested. Finally, it was mentioned that gaps of 5 days or more were removed. No reason is given for this and it seems fairly suspect. Perhaps their statistical models did not handle outliers particularly well so they were removed.

The authors state that the vessels were ranked based on their risk, which was calculated using the standardized probability of gap measure that was shown in Eq. 3.6. However, they do not mention what was used as the cutoff of what they considered at risk. It is also unclear whether if they actually detect AIS off or if they simply determine which vessels are most at risk of turning off their transponders which is much less informative as there is always the possibility that no vessels have turned off their transponders, which would make any ranking pointless.

The authors state that they were able to determine that high risk vessels had a probability of gap value of $5e-2$ and a median gap of 36 hours. They also state that their model is less sensitive to smaller gaps and is better at finding the larger ones (gaps of 10 to 12 hours). This seems like a major drawback as the larger gaps are going to be easier to detect than the smaller ones. Not being as successful at finding the smaller gaps would be a large downside to this implementation.

The authors also stated that they picked up vessels with known associations with Illegal Unreported and Unregulated fishing activity though they say nothing more about it. How many of these vessels are there? How many did they pick up? How many did they miss? How did they have a list of these vessels? Are there other vessels that they believe they found that could have been fishing illegally that were not previously on the lists? None of these questions are answered.

Finally, one interesting finding was that they claimed that they were able to find an AIS switch-off hotspot on the border of Australia and Papua New Guinea. Vessels, the authors state, were likely to turn off their AIS as they were leaving Australian waters.

3.2.4 Mining Vessel Trajectories For Illegal Fishing Detection

Unlike previous papers which dealt with the broader issue of the AIS switch-off problem, this paper [33] attempts to deal specifically with Dark Fishing. Dark Fishing refers to the AIS switch-off problem where fishing vessels turn off their transponders in order to hide illegal fishing activities. This paper attempts to deal with the problem by generating profiles and ranks for each vessel based on their trajectory movement patterns.

For the experiments run, 5 years of data was used. This data ranged from the years 2000 to 2014. It included 330 million data points that were segmented into 250 000 trajectories which came from 799 unique vessels.

The output of their algorithm is a score that the authors call the Vessel Suspiciousness Score. Out of the 799 vessels in their dataset, 753 had a Vessel Suspiciousness Score above 0. The other 46 vessels, which had a score of 0, either did not move for the entire duration of the trajectory or were in full regulatory compliance during their trip (though the authors do not state the ratios for each).

In order to perform their experiments, the data was split into 5 groups where each group represented a single year. Of the 753 vessels that had suspiciousness scores of above 0, only 218 had scores above 0 for each year. Of the 218 vessels that had a suspiciousness score of above 0 for all 5 years, 55% of those vessels always ranked in the top 20% of suspicious vessels.

The solution proposed by the authors has for its input, historical AIS data of vessels that the authors state routinely operate in coastal waters, though they do not define what they mean by that. The output of the solution then is a ranking which is based on the suspiciousness of the vessel. Worth noting is that a weak order is used in the ranking as vessels can be tied in their risk level.

The algorithm itself is broken into three steps. The first is endpoint identification, the second is end-to-end trip segmentation, and finally we have the ranking of suspicious activity.

The first of these steps, the endpoint identification, is essentially the beginning of the segmentation phase of the algorithm. In order to accomplish this task, the authors wanted to determine where a voyage begins and ends. To do this, a two step process was implemented. Firstly, DBSCAN was used in order to cluster the Latitudes and Longitudes of the vessels. The assumption being made is that if an area that has a high concentration of points from a vessel, it will most likely be a place where it is at anchorage.

By doing this, the authors are able to extract a series of anchorage sites. Once anchorage points are extracted for all the vessels, they are then clustered in order to determine the end points which are the points where trips by vessels either begin or end. These end points include ports as well as other unknown harbours.

Once the places where we the trips begin and end are determined, the data itself can be segmented. For the segmentation phase, two steps are performed. Firstly, interpolation is done to fill in some gaps in the trajectory. The authors state that they do this because things like data corruption, noise and inconsistencies in AIS data exist. They state that they do this when there are reasonably short gaps, though they do not define what they mean by this. Once interpolation has been completed, the newly interpolated trajectory is segmented into end to end trips by port location.

Then, for each trajectory that has been created, four scores are generated. One is a Global Anomaly Score. This score measures the deviation of the current trip relative to similar trips performed by other vessels in the past. Secondly, is the Local Anomaly Score. This score measures the behaviour of the particular trip in question. Finally, by combining the Global and Local Anomaly Scores, a Trip Suspiciousness Score is calculated. A value between 0 and 1 that tells us how suspicious the vessel was on that particular trip. Finally, we receive a Vessel Suspiciousness Score which is an aggregation of all Trip Suspiciousness Scores.

How is the Global Score calculated? Firstly, all end to end trips are grouped together based on their endpoints. The idea is that the endpoints will tell us something about the trip itself. For example, if the first end point is one port and the second is another port and there are known fishing grounds between them, that would say something about the probable behaviour of the vessel. However, if one of the endpoints is a harbour that is known for unloading, then the vessel may go in a straight line to get there.

According to the authors, there are three kinds of trips, transition trips, fishing trips and searching trips. Trajectories can be categorized into one of these groups based on their spatio-temporal characteristics. Though how they determine which trajectories go in which category, why they do this, or even what transition and searching trips are is not defined in the paper.

In order to calculate the Global Score itself, we start with a set of all trajectories that are clustered together based on their origin and destination. The only feature that is taken from this is the length of the trajectory as the authors claim that it is the most important indicator of trip type though they do not say why this is the case. Then, using the length of the trip, they cluster each trip type using a Kernel Density Estimation clustering method. The Global Anomaly Score is calculated as the distance from the trajectory to the centre of the cluster.

The Local Anomaly Score is calculated based on the behaviour of a particular vessel for a particular trip. It is then calculated by then summing the total length of dark periods and dividing by the total length of the trip. This term is then multiplied by a factor called λ. λ is simply a multiplicative term that increases the risk if it is considered a fishing trip or does not change the anomaly score if it is not.

The Trip Suspiciousness score is simply the Local Anomaly Score and Global Score each multiplied by user determined variables and then added together. Vessels are then ranked by their total score.

3.2.5 Real-time Maritime Anomaly Detection: Detecting Intentional AIS Switch-Off

Firstly, this paper attempts to formally define the AIS switch-off problem. It begins with an AIS dataset that we will call D. D is the set of spatio-temporal points. p are the temporally sorted points where $p \in D$.

Each p consists of:

- A timestamp which is represented as the number of seconds that have passed since January 1st 1970
- A Longitude
- A Latitude
- A speed that is measured in knots
- And finally the Course Over Ground or COG which is the direction of the vessel relative to True North. It is represented as a value between 0 and 359

Next, comes the definition of an AIS gap. If we want to be able to determine whether a vessel has turned off their transponder, we need to be able to recognize the gap in data that will occur. This paper defines a gap as an absence of points p from D for an unjustified period of time t. In order to determine what unjustified means, we need to be able to know things such as the last known location, speed and COG. The reason for this is that the amount of time between transmission is dependent on these factors.

When a gap in AIS transmission data is found, there are two possibilities. The first is that a vessel has turned off their AIS transponder. The second is that the signal was not picked up by an AIS receiver. With that in mind, the paper then next defines network coverage. That is, a measure of whether an AIS receiver should be able to pick up the signal of the vessel given the geographic area. In order to accomplish this, the geographical area of interest is divided into a grid of equal sized squares. We then define C as the set of all squares, or cells. c is each individual cell where $c \in C$.

Each c consists of:

- A minimum Longitude
- A minimum Latitude
- Size of cell measured in degrees

Each cell, as well as the physical characteristics such as the Latitudes and Longitudes, also contain a series of points A such that $A \subseteq D$ which is defined as the entire dataset of points defined previously. For a given time period, defined as $T_{start} - T_{end}$, we refer to a square as a Dark Cell if and only if the subset of points for that cell at a given time period T is null.

The dark cells are represented as the variable B such that $B \subseteq C$ and that for each cell c that is an element of B, the set of points in that cell, represented as A_c is null.

The final definition needed is the Projection. The extrapolation as shown in Eq. 3.7 and 3.8, is defined as the variable P is the projected destination of a trajectory given the most recent positional point, represented as p, as well as the the distance travelled and the bearing of the vessel represented as δ and Θ.

$$P.latitude = asin(sin(p.latitude) * cos(\delta) + cos(p.latitude) * sin(\delta) * cos(\Theta)) \tag{3.7}$$

$$\begin{gathered} P.longitude = p.longitude + atan2(sin(\Theta) * sin(\delta) * cos(p.latitude), \\ cos(\delta) - sin(p.latitude) * sin(P.latitude)) \end{gathered} \tag{3.8}$$

As to the approach taken by the authors, the area that is being covered is firstly divided into equally sized cells. Each cell is going to receive varying numbers of AIS messages. Lack of coverage, or dark cells, can be caused by either weather conditions, lack of base station in the area, or the receivers could be jammed. Unfortunately, only the locations of the base stations can be known a priori. However, we are able to infer problems caused by weather or jamming by the number of AIS messages per cell which are visualized in Fig. 3.6. Another challenge to add to this is that while there will be cells that are dark, they will not always be so. The status of a cell is not static as things such as the weather change over time, making the challenge greater.

The idea is that network coverage is used to infer the type of message loss. In order to do this, Akka is used. Akka is a toolkit that is used for concurrent and distributed computing and its use will be described in greater detail later. But essentially, an actor, or a worker node is assigned to each vessel. If a timeout occurs where an actor node does not receive a new message, the last known location of the vessel is used in order to calculate the extrapolation. Once the extrapolation is calculated, it is checked to see if the cell that the extrapolation falls in is a dark cell or not. If it is a dark cell, it is not considered a switch-off. Else, it is. First the extrapolation is calculated, then it is checked to determine which cell the extrapolation falls into. Then it is seen whether the cell is dark or not.

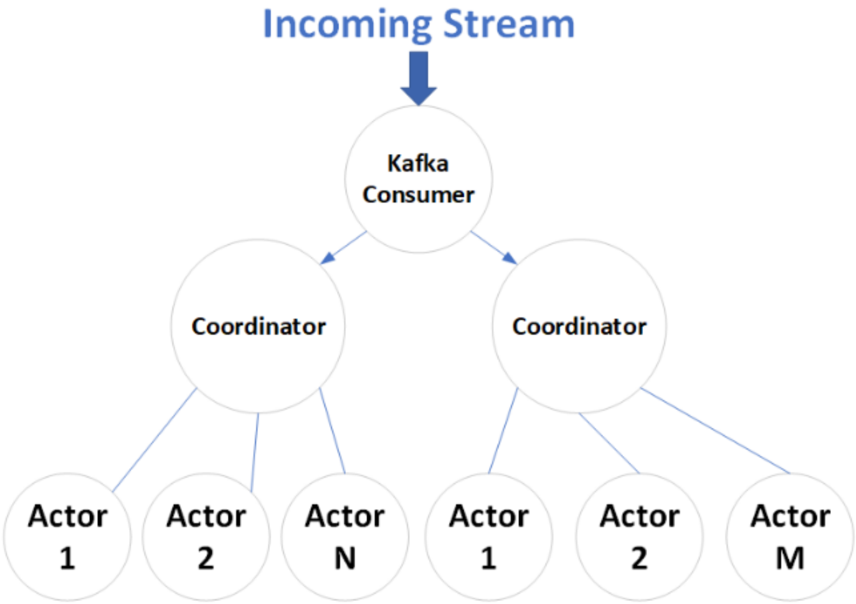

Figure 3.5. System Architecture [18]

Now we move onto the architecture itself. There are three layers. All of which are built upon Akka. Akka was used because of its scalability and configuarability. The base unit of Akka is an Actor. Actors are lightweight classes which have a minuscule footprint. 2.5 million actors can be initialized per GB of heap memory. There are three types of Akka actors used in this work. At the top is the consumer actor. Then below that are the coordinator actors and at the bottom are the worker actors. The general layout can be seen in Fig. 3.5.

The Consumer Node is at the top. It ingests the decoded AIS messages which are stored in Kafka in order to distribute them to the proper Coordinator. It makes sure that each Coordinator receives the same number of ship IDs, and thus performs a form of load balancing. If a Vessel ID arrives that it has seen before, it sends to the proper coordinator and if it receives a new Vessel ID, it sends to the coordinator with the fewest ship IDs.

Secondly, is the Coordinator Node. The Coordinator Node is essentially the router. There is one Coordinator per machine cluster. Once it receives a message, it checks whether a worker has already been created for that vessel ID. If it has, it passes the message on to that worker. If there has not been a worker already created for that vessel ID, a new one is created and the message is passed to the new worker.

Then finally, we have the Worker nodes. Each worker node is responsible for a single vessel. If a new message is not sent to the working for a predefined period of time, a timeout occurs. If that timeout occurs, the algorithm mentioned previously is used to calculate the extrapolation and to determine whether that extrapolation falls in a dark cell or not. If it finds that the extrapolation did not appear in a dark cell, it creates an anomaly event and records that in a Kafka topic that records all anomaly events for all vessels. Fig. 3.5 shows a graph showing the architecture and how everything connects together.

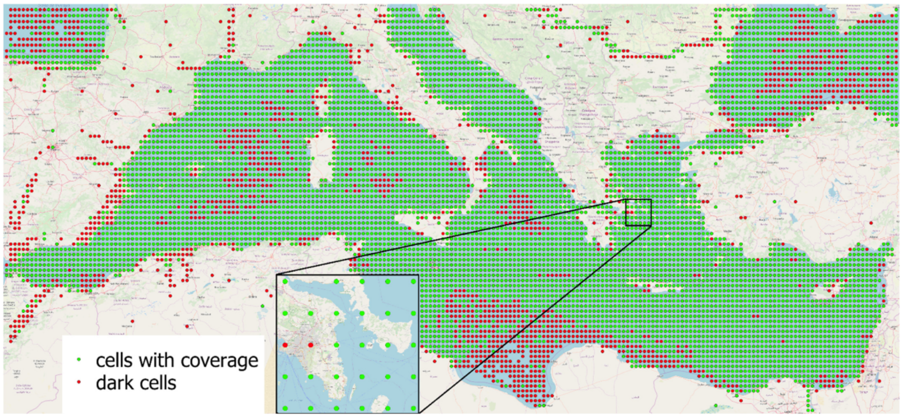

Figure 3.6. Example of System Coverage [18]

3.2.6 Discussion

There are several issues that plague most of the papers discussed in this chapter. Some architectures are only really useful as an academic tool or for some form of post-hoc analysis; some have issues with scaling, while others cannot utilize all kinds of AIS data. The most significant issue, however, for every paper is the lack of labeled data. AIS data can be acquired in abundance. Gigabytes are generated every day. Simply acquiring trajectories is easy, and building models based on those trajectories is not a difficult task. The real challenge comes in the validation of those models. Because the targets cannot be verified by the data itself, it requires either a domain expert who is familiar with smuggling and dark fishing, or a dataset that has already collected said data. Unfortunately, as far as I know, the second has not been accomplished as of yet.

All of the papers presented here have tried to get around this problem by creating artificial gaps in their datasets and then had their models attempt to find those gaps. While this may be sufficient for a proof of concept, there are limitations and potential dangers in this type of strategy. Real world data is messy, and it often does not behave in the way that we assume that it will. By training models on artificial data that was created based on what is assumed would happen in the real world, the models become susceptible to problems once real-world data does not behave in that assumed fashion. Early in their paper [33] proposed a potential solution to this problem where they stated that they could have contacted fishing monitoring organizations in order to obtain a list of vessels that were known to participate in dark fishing. By even having a list of those vessels, a dataset could be generated to train the models. Potentially something similar could be done for vessels that were known to smuggle in the past. It's possible that this data is sensitive and organizations that have this information may be reluctant to hand it over, even for academic purposes, though it would be worth a try and could be a part of some future work.

Another challenge that many papers fell into was the reliance on vessel history. Some papers assumed that previous vessel behaviour could be used in order to predict future behaviour. While this thesis may be sound, it comes with potential pitfalls. Firstly, not all vessels have an extensive history, and these algorithms may not be appropriate for the situation. It is possible that these vessels with limited history could be newly constructed vessels or have changed their MMSI value. This leads us to the second problem. Since AIS is a self-reporting system where the users on the vessel are able to change their identity whenever they want, it becomes difficult to justify the use of historical vessel movements, which will be based on a value such as MMSI or IMO, when a vessel can simply change its identity. This problem could be particularly pronounced for vessels engaged in smuggling as it may be advantageous for them to hide who they really are.

Chapter 4

Problem and Methods

This chapter will delve into the requirements and solutions proposed to advance the on-purpose AIS switch-off problem. Section 4.1 explains the aspects of each work shown in Chapter 3 and determines criteria for choosing a previous work to improve upon. Section 4.2 discusses the issues with the chosen baseline [18], as well as the proposed solutions to these problems. This will be followed by Sections 4.3 which will define the Evaluation Metrics and 4.4 which will discuss the Statistical Analysis used. Finally, Section 4.5 ends the chapter by looking at the data set used in this work to test both the baseline as well as the proposed solutions, where the data came from, as well as the characteristics of the data set.

4.1 Requirements

The efficacy of each solution discussed in Chapter 3 is partly based on the needs of the user and the categories in which each solution falls. If the desire is to create a platform that is able to detect IUU fishing or smuggling in real time in order to prevent it from happening, there are a number of characteristics that would be wanted in a solution.

Firstly, one would want the model to be able to process data in real time. A model that processes only historical data would be beneficial for post-hoc analysis, but it would not help as much with the enforcement of maritime regulations. Secondly, a Spatial Density model would be much more beneficial than a model that is based on Vessel History. The reason for this is that when we are attempting to locate vessels that are taking part in illegal activities, we cannot take for granted that we are going to have the history of the vessel. If a vessel is willing to turn their transponders off, they may be just as willing to change their identification number. We would then run into the issue of needing to find the vessel again under its new MMSI in order to utilize its history. We may also not have the history of the vessel. While AIS data is abundant, we cannot take for granted that all end users are going to have access to all historical AIS data. By using a Spatial Density model, we remove that requirement and make the model more accessible. Finally, a model that is able to use both Terrestrial as well as Satellite AIS is going to be preferable to one that is limited to one or the other. It once again provides flexibility to the user and allows for more data to be used.

Basing a decision on this criteria, there is only one proposed solution that meets these requirements. Real-time maritime anomaly detection: Detecting intentional AIS switch-off [18] meets each of these needs as well as promising computational efficiency as well as dynamic scaling. For these reasons, this approach is used as a baseline.

4.2 Proposed solution overview

As discussed in Section 3.2.6, there are numerous potential issues that arise from the implementation as proposed by the authors of [18]. The proposal of this Thesis is one that will address one of those most fundamental problems.

4.2.1 Motivation

According to the authors' architecture, once a vessel's signal has stopped being received for some time, a test is performed to determine whether a vessel is projected to be in a cell marked as dark or not. In order to do this, an extrapolation is generated based on the final point of the vessel. The last heading and speed are used to determine the best guess of where the vessel should be at the given moment. As has also been mentioned, this approach is lacking in some key ways. The main issue is that this extrapolation model does not consider the vessel's behaviour just before its signal was lost. Vessels behave in all kinds of ways depending on what they are doing and often do not travel in a straight line. As such, we can run into a scenario where a vessel is in the process of turning when it goes dark. This model, however, only looks at the most recent point and draws a straight line based on the direction that it is facing. Though a smaller issue, this approach requires that the heading and speed values within the AIS message be present and valid. However, this is not always a given. In order to address these issues, two solutions are proposed as an alternative to the default extrapolation algorithm.

4.2.2 Heading improvement

In this subsection we present the first proposal of improvement on the original algorithm. Here, we attempt to keep the structure of the formula, but with a change that will take into consideration the previous behaviour of the vessel being analysed.

Algorithm 1 Update to Default Extrapolation Algorithm

```
Input: Previous N points of vessel: $\phi$; Minimum number of points required: n
Output: Latitude and Longitude extrapolation {I}f the size of the dataset is less than n, use the original algorithm
if $N < n$ then
    return DefaultExtrapolation($\phi$, n)
end if
{Calculate the derivative of the heading with respect to time}
$X \leftarrow \frac{d\phi_{heading}}{d\phi_{time}}$
$Y \leftarrow mean(X)$
{Calculate how long since last signal}
$Z \leftarrow currentTime - timeOfLastSignal$
{Multiply the mean difference in heading by the time difference}
$R \leftarrow Y * Z$
{Add R to the heading value of the last point}
$S \leftarrow \phi_{N,Heading} + R$
{With the newly updated heading value, calculate the extrapolation}
$T \leftarrow distanceTravelled(\phi)$
$Extrapolation_{Latitude} \leftarrow asin(sin(\phi_{N,Latitude}) * cos(T) + cos(\phi_{N,Latitude}) * sin(T) * cos(S))$
$Extrapolation_{Longitude} \leftarrow \phi_{N,Longitude} + atan2(sin(S) * sin(T) * cos(\phi_{N,Latitude}), cos(T) - sin(\phi_{N,Latitude}) * sin(Extrapolation_{Latitude}))$
return $Point(Extrapolation_{Latitude}, Extrapolation_{Longitude})$
```

Algorithm 1 contains the pseudo-code for the updated algorithm. For the most part, the algorithm is the same as the one shown in Equations 3.7 and 3.8 with one major difference. Rather than use the heading value present in the final AIS message before the vessel goes dark, multiple heading values are used. The previous N points (where N is a user defined parameter) sent from the vessel are taken and the derivatives of the heading with respect to time are taken. This derivative gives a measure of degrees of change in heading per second. The mean of all of these values is taken, which gives us a mean change in degrees per second for the vessel in question. The number of seconds between the current time and when the last message was sent is calculated and multiplied by our newly acquired mean value. This new value is then passed into the extrapolation algorithm instead of the old heading value.

The idea behind this change is that if we are able to determine an average rate of change in the heading over the vessel's recent history, we can then take a guess as to how that heading should change after the vessel has gone dark. If the vessel is travelling in a straight line, then the derivative of the heading will be 0. If the derivative of the heading is 0, then nothing is added to the default extrapolation algorithm and it is the exact same as before. However, if the vessel had been in the process of turning, we will be able to take that into account in the extrapolation.

4.2.3 Linear Interpolation/Extrapolation

The second approach attempted goes in a completely different direction than the original extrapolation algorithm. Rather than directly generating an extrapolation based on the location, heading and speed of the vessel, a Linear Interpolation algorithm is used. The assumption being made in using this algorithm is that while a vessel is en route, there is going to be some relationship between its Longitude and Latitude.

Using various forms of interpolation in the context of AIS has been done numerous times in the past [15, 19, 40]. In these previous works, authors have attempted to reconstruct AIS trajectories using various forms of interpolation. This approach, while similar, differs in one key aspect. When reconstructing routes, the gaps that are being filled tend to be in the middle of two separate trajectories. This problem lends well to the idea of interpolation where we can think of a trajectory as almost a graph with the Longitude on the x axis and the Latitude on the y. A relationship between the two can often be ascertained and given a Longitude in a gap, the corresponding

Latitude can be guessed.

With this problem, there is no gap that is in the middle of a larger trajectory. We are attempting to fill in a trajectory with an extrapolation, but that trajectory is at the edge. This means that if the vessel is travelling in a straight line that the values may fall outside of the set that the algorithm has seen. This poses a problem to many interpolation algorithms that extrapolate incredibly poorly. Because of this, a general extrapolation algorithm is used for when data falls outside of the scope that the algorithm has seen.

As can be seen in the strategy shown in Alg. 2, the original Extrapolation algorithm is used to define the Latitude extrapolation. That extrapolation is then passed into a general linear extrapolation algorithm which will then return the Longitude extrapolation.

Algorithm 2 Linear Interpolation Extrapolation Algorithm

```
Input: Previous N points of vessel: $\phi$; Minimum number of points required: n
Output: Latitude and Longitude extrapolation
if $Length(N) < n$ then
    return DefaultExtrapolation($\phi$, n)
end if
$X \leftarrow LinearInterpolation(\phi_{Latitude}, \phi_{Longitude})$
$Extrapolation_{Latitude} \leftarrow asin(sin(\phi_{N,Latitude}) * cos(T) + cos(\phi_{N,Latitude}) * sin(T) * cos(S))$
$Extrapolation_{Longitude} \leftarrow X(Extrapolation_{Latitude})$
return $Point(Extrapolation_{Latitude}, Extrapolation_{Longitude})$
```

4.3 Evaluation Metrics

In order to evaluate the algorithms for detecting on-purpose AIS switch-off, various measures are used. As we are able to generate both positive and negative classes, more robust measures can be taken. Various measures were captured in order to properly compare and contrast models. The major ones, the data from which will be shared later in Section 5, include:

- Mean Accuracy

- Standard Deviation of Accuracy
- Mean F1 Score
- Standard Deviation of F1 Score

As will be seen later on, the results of these experiments are typically skewed in one direction or another. As such, the mean accuracy itself is an inappropriate measure by itself as these values can be greatly affected by outlier values and an incomplete picture can be painted with those values in isolation. Because of this problem, the mean values will be shared in conjunction with various other measures that will paint a much more complete picture of the true distribution of the results.

4.4 Statistical Analysis

In order to validate the results of these models, best practice is that some form of statistical analysis is performed to determine whether the results are statistically significant. Because there is always variance within the results, simply taking the top-line numbers may not tell us the entire picture. Therefore, some form of statistical test will be run to better understand the results.

Often, t-tests are performed to determine whether there is a statistical difference between the means of two groups. However, this statistic is often misused as there are many assumptions made by using the test that is not accounted for. These assumptions include normality of the data and similarity in variance between groups if a basic t-test is being used. Because these assumptions cannot be met with any degree of certainty, a nonparametric test is used.

For these tests, the Mann-Whitney U Test [20] is implemented to test whether the distribution of F1-Scores between each interpolation method is the same or not. Mann-Whitney is a nonparametric statistical test that is used to measure whether the distribution between two groups is statistically different or not. The main benefit of using the Man-Whitney U Test is that the assumptions being made are few. Firstly, there is an assumption that the two groups are independent. This assumption can be met. While the same data is being used, the models that we are testing are different, and a change to one model will not affect the results of another. And secondly, the values must be ordinal. As we measure F1-Scores, this assumption is also met. Finally, a choice must be made as to an α value. As 0.05 is a standard value used in most works found in literature, this is the one that was chosen in our experiments.

4.5 Data source

The data utilized for these experimentations is an open-source AIS dataset that has been collected by the United States Coast Guard[21]. This data, from January 2017, contains all AIS messages from what is referred to as Zone 3 in the Universal Transverse Mercator Coordinate System (UTM) coordinate system[38]. UTM is a coordinate system for dividing the globe into a series of squares. Zones are divided vertically every 6 degrees of longitude, creating 60 zones, each beginning at 84°N and ending at 80°S. Global shipping data is often broken into UTM Zones to better break massive datasets into smaller sets that are divided by geographic area. Zone 3, the geographical region that this represents, is bounded by longitude values -162 to -168 and contains the western coast of Alaska. The area used in these experiments can be seen in Fig. 4.1.

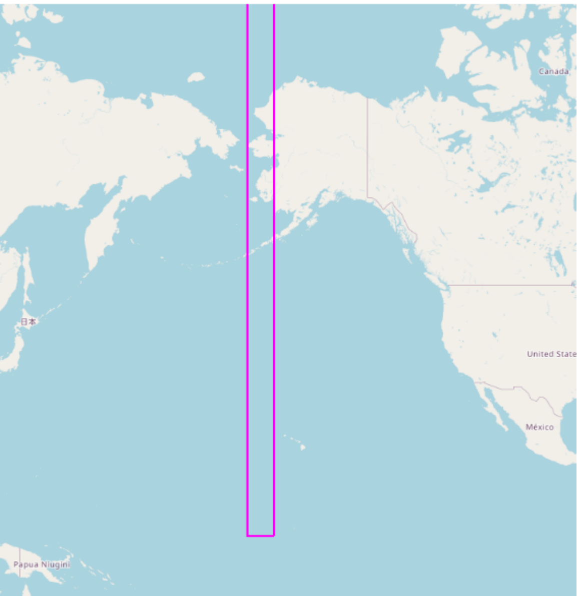

Figure 4.1. Area of interest

There are a couple of reasons why this particular dataset and this region were chosen. The main reason is that fishing in Alaska accounts for over half of all fishing in the US [31]. This fact led us to want to test these methods out in this region. Dutch Harbor, a harbor that is situated in Unalaska, is the largest fishing port in the United States and has been for almost 25 years [31]. Dutch Harbor, along with the AIS data for the given time period, can be seen in Figure 4.3. The dataset itself contains over 2.9 million AIS messages and 823 different trajectories of ships which can be seen in Fig. 4.2.

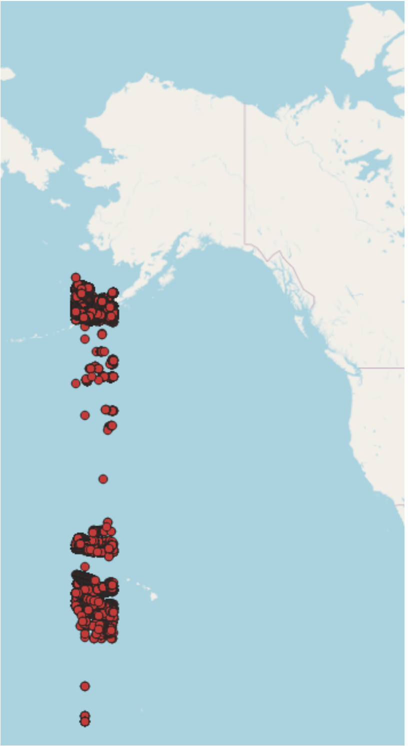

Figure 4.2. Area of interest with data visualized

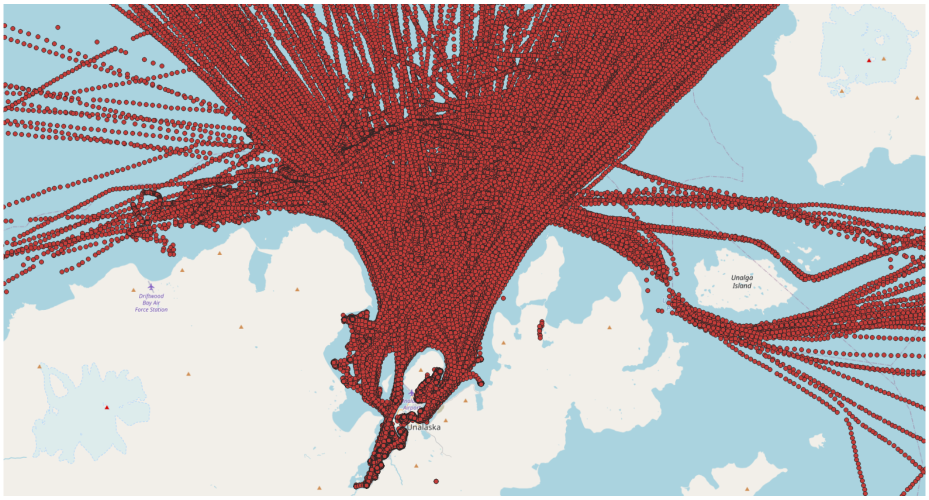

Figure 4.3. Dutch Harbor and surrounding areas

Chapter 5

Experimentation and Evaluation

This chapter looks at the results of the experiments conducted in this work. The chapter begins in Section 5.1 by looking at the setup of the experiments that were done. This includes data preparation, gap generation, as well hyperparameter selection. After, Section 5.2 will look at the results of the experiments run. It will show why the approaches taken in this paper were justified and dig deeper into the various results grouped by parameter. Finally, in this same section, each of the results will be discussed and analyzed in detail.

5.1 Experimental Setup

5.1.1 Data preparation and gap generation

Unfortunately, like all of the other papers that have attempted to solve this same problem before, we could not obtain a real-world dataset containing known AIS switch-off events. We believe that we came close and that doing so would provide an interesting future work.

As such, the gaps that are used in this work are artificially generated. The way that this was done is similar to the method as described in the baseline paper [18]. The area of interest (the area in which the vessels are found) is divided up into a grid of equally sized squares. The size of those squares is a user-defined parameter. Once the grid is in place, a proportion of those cells are marked as dark. These cells represent those in which there is some form of interference. All points in those squares are removed and the gaps created by the Dark Cells are marked as the negative class in our experiments. As the Dark Cells are chosen at random, and the proportion of cells marked as dark is one of the parameters being tested, the number of gaps in the negative class is not consistent though typically not much different from the size of the positive class.

After this, 240 vessels (30% of the total number of vessels) that do not pass through the squares that are marked as dark are sampled with replacement, and gaps of 2 hours are created in each. The start points of each of these gaps are recorded as the positive class in the experiments. Having only vessels that did not pass through Dark Cells was a choice that the original authors made. This was probably to ensure that a scenario did not occur where gaps were on top of each other that belonged to both the positive and negative classes.

5.1.2 Experiments

The experimentation was done in the form of a grid search. Various hyperparameters were chosen for each user definable parameter and the results of each extrapolation model are compared.

The user defined parameters include:

- Dark Cell Proportion - The proportion of cells to mark as ”Dark”
- Cell Length - The length of squares in the grid
- Time Delta - The length of time between signals after which a timeout occurs

For each run of the grid search, both new Dark Cells and gaps were generated. The idea behind this was that a certain combination of Dark Cell and gap results might create a model that is artificially too difficult or too easy. By generating everything again at random at every run of the model, we can hopefully avoid artificial results and obtain more well-rounded results that better reflect how well the extrapolation algorithms actually work.

A minimum of 36 runs was completed for each extrapolation algorithm with 4 Dark Cell Proportion values tested and 3 of each of the other two parameters tested as well. The choice of hyperparameters was mostly based on the hyperparameters chosen by the original authors for proper reproduction of their results. The full breakdown of the parameters, as well as the values chosen for each, can be found in Tab. 5.1.

Parameters	**Values**
Projection type	[Baseline, Heading, Interpolation]
Dark Cell Proportion	[0.01, 0.05, 0.1, 0.2]
Time Delta (seconds)	[3600, 5400, 7200]
Cell Length (degrees)	[3, 5, 7]

Table 5.1. Parameters and chosen values

5.2 Results

This section is divided into three parts. Section 5.2.1 looks at some of the high level results of the improved heading update algorithm. It shows a number of figures showing the reasoning behind its conception and why it works. Secondly, Section 5.2.2 presents an overall comparison of each approach, looking at the evaluation metrics of the runs as a whole. Finally, Section 5.2.3 shows the results of each parameter grouped by hyperparameter values and a comparison will be made between the baseline and the two other extrapolation algorithms. The idea behind digging into each parameter and how the models behave is two-fold. Firstly, we are able to determine which parameters have the largest impact on the models themselves. This can show us where the architecture is both weakest and strongest. Secondly, it allows us to better determine how well the models are working. It is possible to have one of the extrapolation algorithms work very well with some hyperparameters and very poorly with others. While at the same time, another extrapolation algorithm could be middling for all hyperparameter values. Without looking in depth at the parameter breakdown, these results may look quite similar, while in reality there are start differences which could greatly affect future experimentation.

5.2.1 Heading Change Results

Figs. 5.1 and 5.2 show a couple of examples comparing the two algorithms while illustrating the reason for the change. In both, we are able to see the original point from which the extrapolation is made. This point is represented by the square. From that point, the original extrapolation algorithm draws a straight line from the origin and ends up predicting that the vessel is in a dark cell (represented by the darker shaded colour). However, the improved heading algorithm is able to look back and determine that the vessel was currently in the process of turning and was able to compensate for that. While the extrapolations were not perfect (falling outside the true trajectory), it was close enough to be able to properly determine which area the vessel was in.

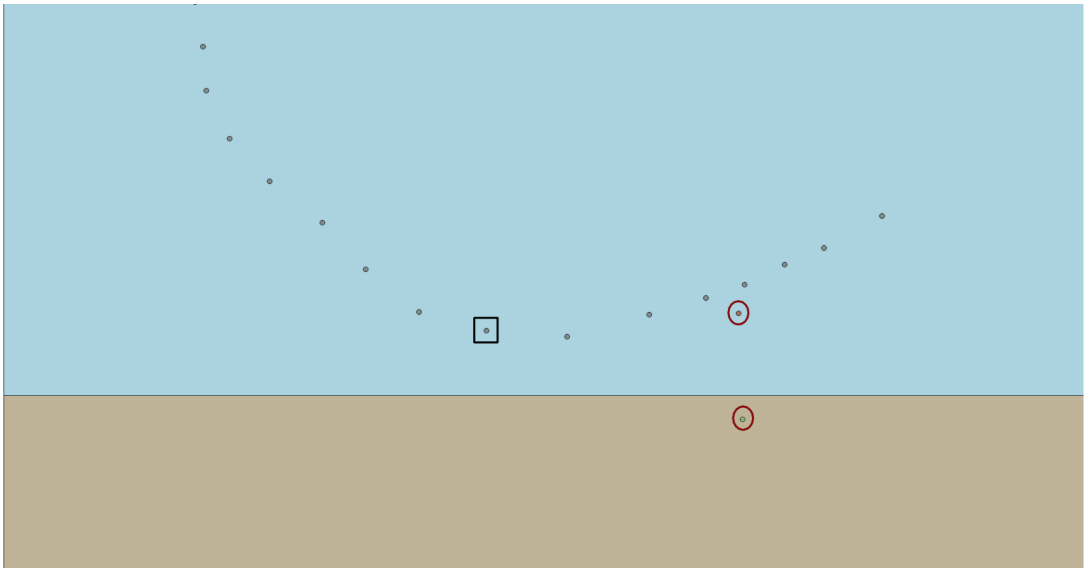

Figure 5.1. Example of issue with original algorithm with square around point extrapolation is made from and the circles around the original extrapolation (bottom) and improved (top)

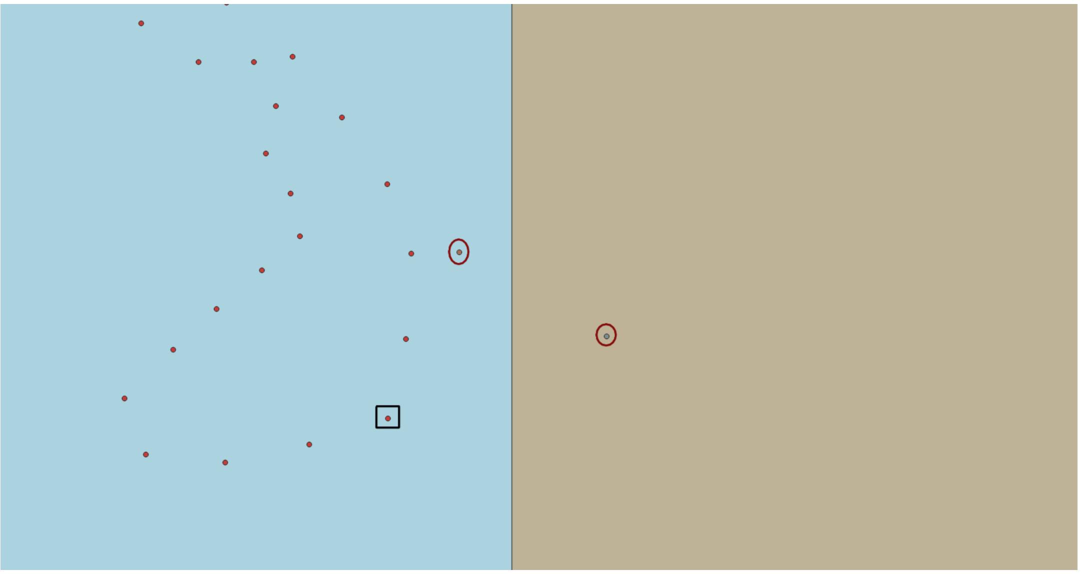

Figure 5.2. Second example of issue with original algorithm with square around point extrapolation is made from and the circles around the original extrapolation (right) and improved (left)

5.2.2 Overall Results

In the subsequent sections, we drill deeply into the various parameters and results of the various extrapolation algorithms in order to determine where the specific models were the strongest and where they were the weakest. In this section however, we take a broader look at the overall results.

Tab. 5.2 provides the top line numbers of the results. From this, we can see a significant improvement on the mean F1-Scores and Accuracies. The Interpolation algorithm provided a mean F1-Score that was, on average, 0.065 better than that of the baseline. The Heading algorithm provided an even stronger result, beating the baseline by 0.079 on average. Potentially even more significant though are the results of the standard deviations. Not only did the new results provide a significant improvement on the baseline, they were able to do so with a standard deviation for the F1-Score of less than half of that of the baseline. While the standard deviations of the means were a little higher, they were still just over half of that of the baseline.

From this, we are able to see that not only did the improvements implemented here provide more accurate results, they did so more consistently and with less variability.

Extrapolation Method	**Mean F1-Score**	**Std F1-Score**	**Mean Accuracy**	**Std Accuracy**
Baseline	0.864	0.153	0.838	0.168
Heading	0.943	0.068	0.902	0.089
Interpolation	0.929	0.076	0.889	0.09

Table 5.2. Overall Results

Firstly, Figs. 5.3 and 5.4 show us the box-plots of the F1-Scores and Accuracies of the various extrapolation algorithms. Our main takeaway from these figures is that we can see a drastic drop in variability in the proposed algorithms. Overall, the proposed algorithms provided better results in every metric tested. Tab. 5.3 shows the differences between the baselines and proposed values. From these, we can see an improvement on both the F1-Score and Accuracy mean values with an improvement ranging between 0.051 and 0.079.

The greatest improvement however, comes in the measure of variability in the models. The heading extrapolation improvement and the interpolation each have

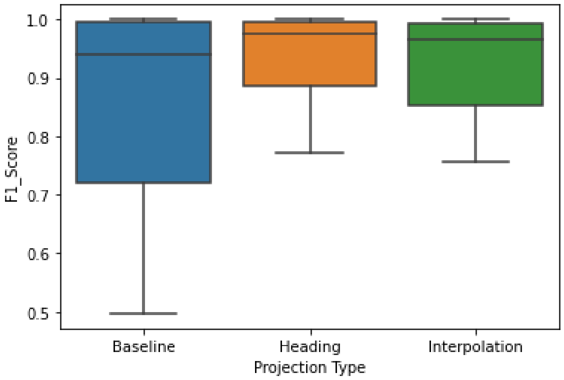

Figure 5.3. F1-Scores organized by extrapolation approach

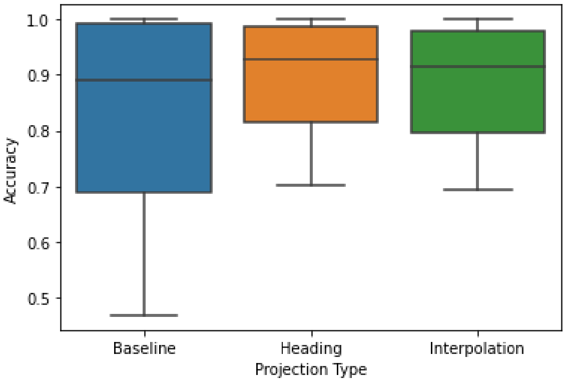

Figure 5.4. Accuracies organized by extrapolation approach

an F1-Score Standard Deviation that is less than half of the corresponding baseline value while the Standard Deviation of the Accuracy values was just shy of half of the corresponding baseline value.

In order to underscore the differences in variability of the models, Fig. 5.5 provides a violin plot so that we can view the distribution of the F1-Scores for each extrapolation model. While there is a smaller cluster of values that are a bit lower than 0.9 in both the heading improvement as well as the interpolation models, the largest cluster, and the vast majority of values fall right around the median values close to 0.97. While the median of the baseline is still high, the distribution of the F1-Scores is quite vast. A significant number of F1-Scores in the baseline fell below 0.7, which never happened with the improved algorithms.

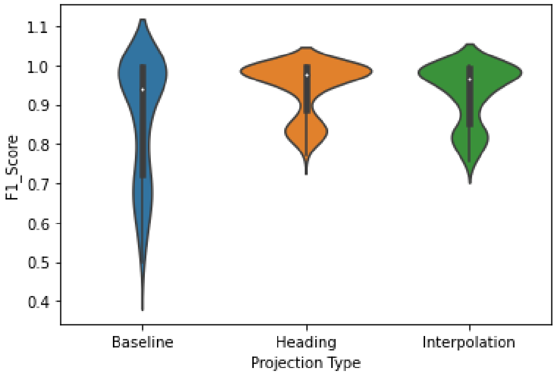

Figure 5.5. Violin Plot of overall F1-Scores compared to extrapolation Type

In the proceeding sections, we will look deeper into the various parameters that were tuned in order to run these experiments. These sections will better demonstrate that the improvements proposed here are vastly superior (more than it may seem while looking at the topline numbers) and will explain why that is the case.

Statistical Analysis

After reporting the overall results, the Mann-Whitney Test can be shown in order to determine whether there is a statistical difference between both groups. Firstly, we can look at Tab. 5.3 in order to see the top-line differences between the Interpolation and Heading algorithms compared to the baseline. Then, the groups themselves can be compared in a one-tailed Mann-Whitney U Test to determine whether the improvements reported are indeed statistically significantly better.

Tab. 5.4 shows both the test statistics as well as the p-value for each algorithm. From this, we can determine that the improvements in the results are indeed significant and the new algorithms perform better than the baseline.

Extrapolation Method	**Mean F1-Score**	**Std F1-Score**	**Mean Accuracy**	**Std Accuracy**
Heading	0.079	-0.085	0.064	-0.079
Interpolation	0.065	-0.077	0.051	-0.078

Table 5.3. Difference between baseline and Proposed Methods

	Test Statistic	**P-Value**
Heading	10661.0	0.00346
Interpolation	7998.5	0.040786

Table 5.4. One-Tailed Mann-Whitney U Test Statistic Results

5.2.3 Input parameter analysis and discussion

Cell Length Results

Figs. 5.6, 5.7, and 5.8 show boxplots for each hyperparameter tested for the Cell Length parameter. The values tested (in degrees) were 3, 5, and 7. Taking a first look at the baseline results, in Fig. 5.6 and Tab. 5.5, we can see a clear relationship between our evaluation metrics and the hyperparameters. The model performed very poorly when the value was small while the performance increased as the cell length value did. The exact reason for this is unclear, however a best guess would be that with a larger grid, vessels that were travelling closer to a boundary between two squares when the cell length was smaller would be less likely to have its extrapolation algorithm fall into the wrong square when the values were larger.

The baseline actually performed very poorly overall when the cell length was 3 degrees with a mean F1-Score of only 0.725, a mean Accuracy of 0.684, and an F1 Standard Deviation of 0.147. From this, we can see that the results were both very poor and highly variable.

While the median F1-Score values in the other two hyperparameters of the baseline results were much higher (0.883 and 0.92), they were still highly variable with Standard Deviation values of 0.163 and 0.112. And while the median values were much better for the other two hyperparameters, we can see that the improvement on the first quartile values is less impressive. The first quartile value when the Cell Length is 5 is still below 0.6 while the first quartile value for a Cell Length value of 7 is barely higher than 0.6.

When looking at the results of the improved heading algorithm and the interpolation algorithm, we can see a large improvement both in accuracy and variability. Tab. 5.5 shows the results for all three methods, while Figs. 5.7 and 5.8 provide visualizations for the heading and interpolation improvement results.

While looking at the results for both, we can see that the same relationship between hyperparameter value and F1-Score/Accuracy exists. The evaluation metrics are lower when the hyperparameter is lower and increases when the cell length increases. However, the starting point of both is quite a bit higher than the baseline. For the heading improvement algorithm, we can see a mean F1-Score starting at 0.912 for a cell length of 3 and growing to a mean F1-Score of 0.982 for a cell length of 7. For the interpolation algorithm we see similar, though slightly lower, results beginning with a mean F1-Score of 0.897 for a cell length value of 3 and growing to 0.977 for a cell length of 7.

While the F1-Score and Accuracy values are better for the heading and interpolation algorithms across the board, another big improvement came in the variability of the models. The Standard Deviation values for the F1-Scores of the heading and interpolation algorithms was on average half of that of the baseline algorithm with the Standard Deviation of a cell length value of 7 going all the way down to 0.016 and 0.022 for heading and interpolation.

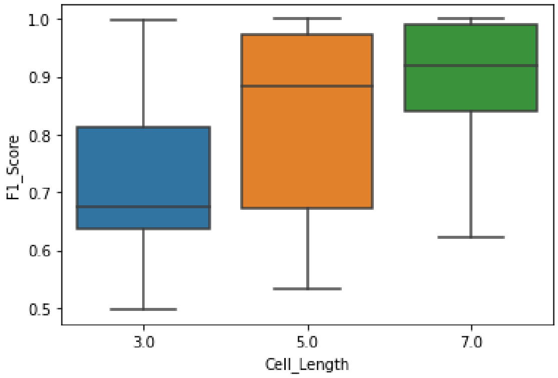

Figure 5.6. Boxplot of Baseline Results using various Cell Length values

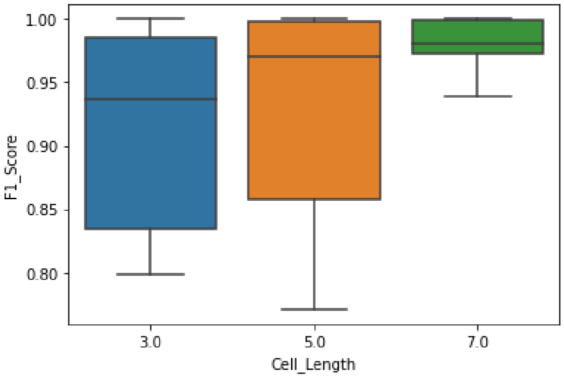

Figure 5.7. Boxplot of Heading Results using various Cell Length values

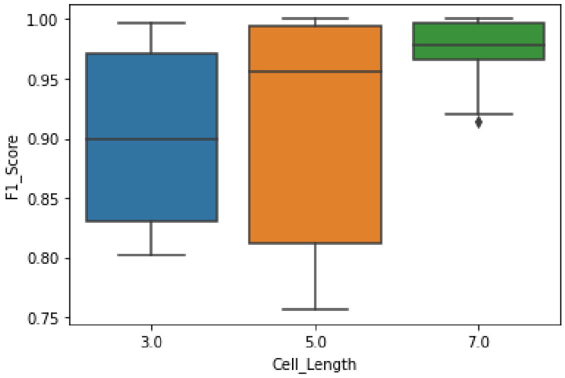

Figure 5.8. Boxplot of Interpolation Results using various Cell Length values

Extrapolation Method	Mean F1-Score	Std F1-Score	Mean Accuracy	Std Accuracy
Baseline (3)	0.725	0.147	0.684	0.153
Heading (3)	0.912	0.076	0.859	0.097
Linear Extrapolation (3)	0.897	0.07	0.845	0.082
Baseline (5)	0.82	0.163	0.795	0.163
Heading (5)	0.935	0.074	0.9	0.093
Linear Extrapolation (5)	0.913	0.092	0.883	0.1
Baseline (7)	0.89	0.112	0.855	0.14
Heading (7)	0.982	0.016	0.947	0.046
Linear Extrapolation (7)	0.977	0.022	0.94	0.057

Table 5.5. Cell Length Results

Dark Cell Proportion Results

In this section, we compare the values of the Dark Cell Proportion parameter. This is the parameter that determines which proportion of cells to mark as dark. In this way, we are able to test various scenarios to see how well the models will behave when large areas suffer from blackout compared to smaller outages.

At first glance, we can see that like the cell length parameter, there is a clear relationship between hyperparameters. Though in this case, the relationship is inversed. This makes intuitive sense. The model has a more difficult time and is more prone to error as outages become greater. This parameter also shows why using a measure such as the F1-Score is important. Because the size of our negative class is determined by which vessels are in cells that have been marked as dark, the inbalance between classes is going to be greater as the proportion of dark cells decreases.

Tab. 5.6 shows the total results for the Dark Cell Proportion parameter with each interpolation algorithm and hyperparameter compared. Fig. 5.9 shows the baseline results for the various Dark Cell Proportion values. From this, we can see a very clear relationship. The models perform very well when the proportion of dark cell values is low. However, the performance drops off a cliff once the number of dark cells increase. When the proportion of dark cells is 0.01, the baseline has a mean F1-Score of 0.965 and very low variability with a Standard Deviation F1-Score of 0.076. However, once we enter the situation where we have an outage that covers 20% of the area of interest, that mean F1-Score drops to 0.733 and the variability doubles to 0.151.

Fig. 5.10 shows the results of the heading improvement for the Dark Cell Proportion parameter while Fig. 5.11 provides the results of the interpolation algorithm for the Dark Cell Proportion parameter. For these extrapolation algorithms, we can see a similar relationship as we did for the baseline extrapolation algorithm. However, while the relationship exists, it is not as pronounced and the higher Dark Cell Proportion values do not have the same drastic drop in performance. While the baseline had a drop in mean F1-Score of 0.262 from a Dark Cell Proportion value of 0.01 to 0.2, the heading improvement and interpolation algorithms had a mean F1-Score value drop of only 0.043 and 0.039 respectively. Between the two algorithms, the mean F1-Score does not drop below 0.935 and the highest variability between the two models (0.084) is close in value to the lowest variability in the baseline (0.076).

From these results, we can see a number of things. Firstly, that having a low Dark Cell Proportion value makes the problem essentially trivial. This makes intuitive sense. However, we can also determine that the new extrapolation algorithms work significantly better than the original baseline when the proportion of dark cells increase. For the times of greatest network blackout, the improvement algorithms have a mean F1-Score more than 0.2 higher than the baseline.

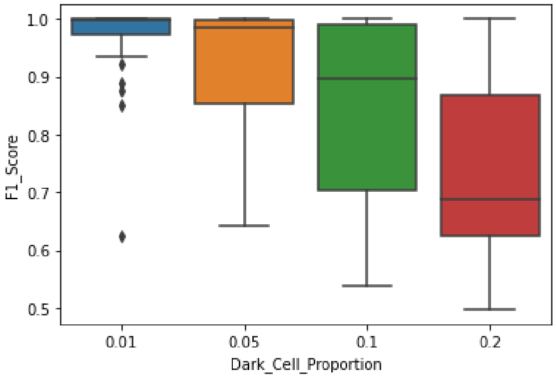

Figure 5.9. Boxplot of Baseline Results using various Dark Cell Proportion Values

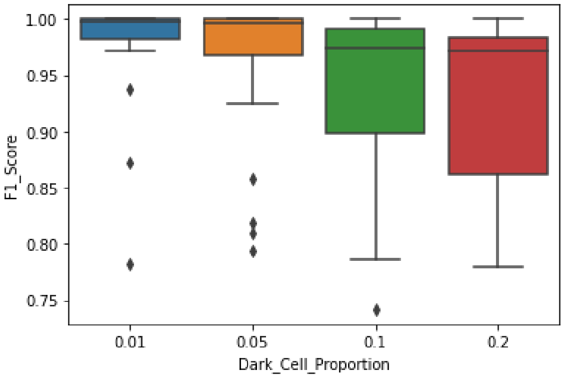

Figure 5.10. Boxplot of Heading Results using various Dark Cell Proportion Values

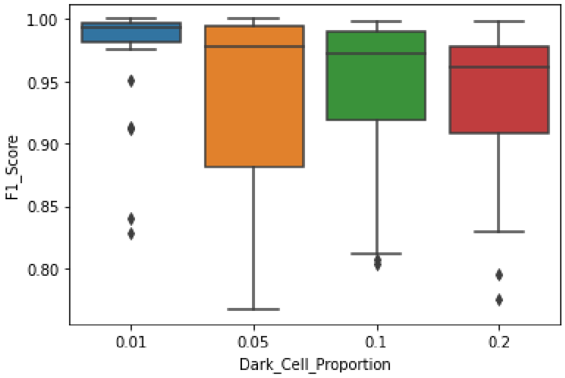

Figure 5.11. Boxplot of Interpolation Results using various Dark Cell Proportion Values

Extrapolation Method	Mean F1-Score	Std F1-Score	Mean Accuracy	Std Accuracy
Baseline (0.01)	0.965	0.076	0.946	0.103
Heading (0.01)	0.985	0.037	0.966	0.062
Linear Extrapolation (0.01)	0.974	0.047	0.95	0.071
Baseline (0.05)	0.914	0.116	0.884	0.141
Heading (0.05)	0.97	0.051	0.939	0.077
Linear Extrapolation (0.05)	0.937	0.084	0.904	0.097
Baseline (0.1)	0.844	0.152	0.807	0.175
Heading (0.1)	0.943	0.072	0.898	0.096
Linear Extrapolation (0.1)	0.937	0.069	0.895	0.083
Baseline (0.2)	0.733	0.151	0.716	0.155
Heading (0.2)	0.942	0.066	0.896	0.084
Linear Extrapolation (0.2)	0.935	0.063	0.883	0.075

Table 5.6. Dark Cell Proportion Results

Time Delta Results

This section looks at the Time Delta parameter. This parameter is the one that determines the timeout after which the model will create an extrapolation algorithm and determine whether an AIS switch-off has occurred or not.

Tab. 5.7 shows the total results for the Time Delta parameter with each interpolation algorithm and hyperparameter compared. The results for this parameter are less interesting than for the other two parameters and we find little new information. Fig. 5.12 provides the baseline results, while Figs. 5.13 and 5.14 provide the improved extrapolation results. From these tables and figures, we can see a very minor negative correlation between median F1-Score and Time Delta value, though it is negligible. What we do see, however, are the same things that we were able to see from the previous parameter values. The improvements proposed provide better mean F1-Scores across the board with half of the variability.

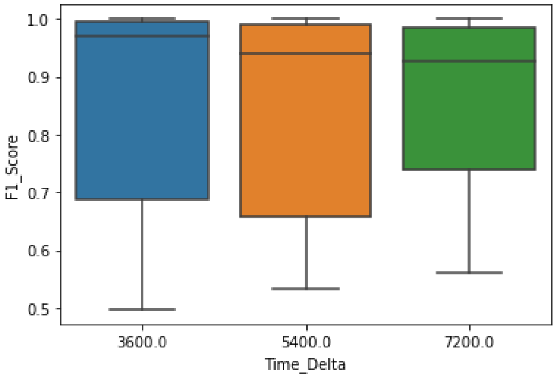

Figure 5.12. Boxplot of Baseline Results using various Time Delta values

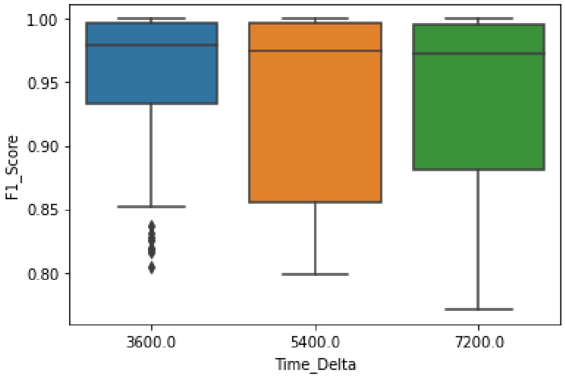

Figure 5.13. Boxplot of Heading Results using various Time Delta values

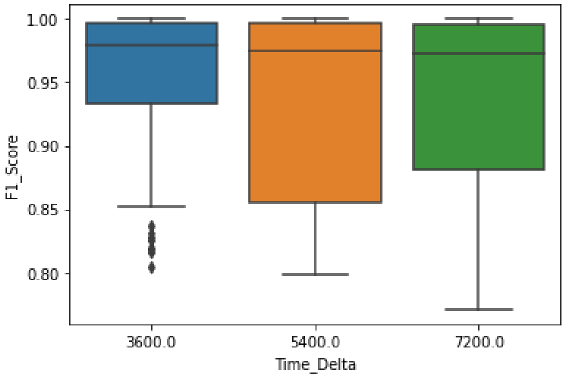

Figure 5.14. Boxplot of Interpolation Results using various Time Delta values

Extrapolation Method	**Mean F1-Score**	**Std F1-Score**	**Mean Accuracy**	**Std Accuracy**
Baseline (3600)	0.849	0.18	0.805	0.215
Heading (3600)	0.949	0.063	0.915	0.081
Linear Extrapolation (3600)	0.943	0.07	0.909	0.088
Baseline (5400)	0.838	0.177	0.796	0.204
Heading (5400)	0.94	0.072	0.899	0.095
Linear Extrapolation (5400)	0.926	0.08	0.892	0.092
Baseline (7200)	0.85	0.154	0.831	0.154
Heading (7200)	0.939	0.069	0.893	0.091
Linear Extrapolation (7200)	0.918	0.077	0.867	0.078

Table 5.7. Time Delta Results

Chapter 6

Conclusions and Future Work

In this work, we provided an in-depth look at the on purpose AIS switch-off problem and literature. We went step-by-step through each previous work done in the field, the strategies implemented, and the strengths, weaknesses, and areas of improvement for each. Then, we proposed a new algorithm to work as an improvement on the most recent and state-of-the-art on purpose AIS switch-off strategy [19]. These improvements focus on the extrapolation algorithm employed. The new extrapolation algorithms can better consider the current behavior of a vessel and can better predict the locations of vessels that are not traveling in a straight line as previously done.

Since real data for detecting intentional AIS switch-off is not publicly available, we carefully design experiments to artificially create such situations and test the baseline algorithm and our modifications. The results indicate that our modifications indeed improve the detection of on purpose AIS switch-off AIS. Extensive testing has been done, and results have been shown to improve the baseline significantly. Furthermore, we have been able to find specific areas in which the original implementation has failed and have been able to determine reasons as to why that is the case. While there are some potential limitations to this approach, it is better in almost every way and to solve some of the major deficiencies of the original approach.

6.1 Limitations

While we have been able to see several improvements in the previous work in this area, there are some limitations to this approach. Firstly, as opposed to the original extrapolation algorithm, these approaches require some previous points from the vessel to always be present. While the number of points needed is low and should usually be able to be obtained, it still provides an additional requirement on the model, which can cause complications in edge scenarios. However, with the improvements provided by these newer extrapolation strategies, this seems to be a fair trade. Another limitation in these extrapolation algorithms compared to the original baseline is that these extrapolation algorithms add an extra step of complexity to models that are meant to be run in real-time.

Also, most of the limitations mentioned in Section 3.2.6 still exist here. This work intended to fix what was viewed as the most significant problem. There still exist, however, some minor issues that could cause problems. An example of this is that there could be signals lost due to vessel congestion in an area. This work will look at that area and see plenty of signals coming from it. Even though the vessel is still attempting to broadcast, its signal cannot get through and will be flagged as an intentional switch-off.

The assumption that is being made by the original authors is that when a single vessel cannot send a signal, all ships in that region cannot send signals as well. This assumption has not been proven and seems as though it could be a dangerous one to make.

The reason that this assumption cannot be tested leads us to the final limitation. That is the fact that while we are able to validate these models using synthetic data, no dataset currently exists that can genuinely validate this work. We do not have datasets of either of the vessels where we know that they have intentionally switched off their AIS dataset, and we do not have datasets where we know that there are gaps that exist because of lack of network coverage that we can compare against.

6.2 Future Work

The problem of intentional AIS switch-off is a large one. Illegal fishing is ubiquitous in many parts of the world, and without regulation, current rates of fishing cannot be sustained. Solutions need to be put into place to curb this problem before severe repercussions to our maritime ecosystems occur. Several future works could be done to improve these models and better detect AIS switch-offs. The first and most important work that could be done would be to compile datasets that could validate these models correctly. There are ways to go about doing this, and the effort should be made to ensure that it happens. Various organizations, both governmental and non-governmental, in the world attempt to monitor smuggling and illegal fishing.

Researchers could partner with either to develop a dataset that includes in it known labelled gaps that can validate such models. In terms of extrapolation techniques, several other strategies could be tested such as kinematic interpolation or neural networks (Long short-term memory (LSTM)) to check if they could properly estimate the ship position while a signal is lost.

Bibliography

[1] How often do the positions of the vessels get updated on MarineTraffic? `help.marinetraffic.com`, (accessed on August 4 2020).

[2] Fernando Henrique Oliveira Abreu, Amilcar Soares, Fernando V Paulovich, and Stan Matwin. Local anomaly detection in maritime traffic using visual analytics. In *In EDBT/ICDT Workshops*, 2021.

[3] Fernando HO Abreu, Amilcar Soares, Fernando V Paulovich, and Stan Matwin. A trajectory scoring tool for local anomaly detection in maritime traffic using visual analytics. *ISPRS International Journal of Geo-Information*, 10(6):412, 2021.

[4] Marco Balduzzi, Alessandro Pasta, and Kyle Wilhoit. A security evaluation of AIS automated identification system. In *Proceedings of the 30th Annual Computer Security Applications Conference on - ACSAC '14*, pages 436–445. ACM Press.

[5] T.N. Bondaroff. The illegal fishing and organized crime nexus: Illegal fishing as transnational organized crime. In *The Global Initiative Against Transnational Organized Crime and The Black Fish*, volume 24.

[6] Emanuele Carlini, Vinicius Monteiro de Lira, Mohammad Etemad, Bruno Brandoli Machado, and Stan Matwin. Uncovering vessel movement patterns from ais data with graph evolution analysis. In *In EDBT/ICDT Workshops*, 2020.

[7] Shwu-Jing Chang. Vessel identification and monitoring systems for maritime security. In *IEEE 37th Annual 2003 International Carnahan Conference onSecurity Technology, 2003. Proceedings.*, pages 66–70.

[8] Renata Dividino, Amilcar Soares, Stan Matwin, Anthony W Isenor, Sean Webb, and Matthew Brousseau. Semantic integration of real-time heterogeneous data streams for ocean-related decision making. In *Big Data and Artificial Intelligence for Military Decision Making*. STO, 2018.

[9] Enrica d'Afflisio, Paolo Braca, Leonardo M. Millefiori, and Peter Willett. Detecting anomalous deviations from standard maritime routes using the ornstein–uhlenbeck process. In *IEEE Transactions on Signal Processing*, volume 66, pages 6474–6487. Conference Name: IEEE Transactions on Signal Processing.

[10] Torkild Eriksen, Gudrun Høye, Bjørn Narheim, and Bente Jensløkken Meland. Maritime traffic monitoring using a space-based AIS receiver. 58(10):537–549.

[11] EU-Commission. Communication from the commission to the european parliament, the council, the european economic and social and the committee of the regions.

[12] Food and Agriculture Organization of the United Nations. Contributing to food security and nutrition for all. (2016). OCLC: 957090850.

[13] Jessica H. Ford, David Peel, David Kroodsma, Britta Denise Hardesty, Uwe Rosebrock, and Chris Wilcox. Detecting suspicious activities at sea based on anomalies in automatic identification systems transmissions. 13(8):e0201640.

[14] Fisheries Governement of Canada and Oceans Canada. Illegal, unreported and unregulated (IUU) fishing - international fisheries. Last Modified: 2019-12-09 Library Catalog: www.dfo-mpo.gc.ca.

[15] Shaoqing Guo, Junmin Mou, Linying Chen, and Pengfei Chen. Improved kinematic interpolation for ais trajectory reconstruction. *Ocean Engineering*, 234:109256, 2021.

[16] Salman Haidri, Yaksh J Haranwala, Vania Bogorny, Chiara Renso, Vinicius Prado da Fonseca, and Amilcar Soares. Ptrail–a python package for parallel trajectory data preprocessing. *arXiv preprint arXiv:2108.13202*, 2021.

[17] Amílcar Soares Júnior, Chiara Renso, and Stan Matwin. Analytic: An active learning system for trajectory classification. *IEEE Computer Graphics and Applications*, 37(5):28–39, 2017.

[18] Ioannis Kontopoulos, Konstantinos Chatzikokolakis, Dimitris Zissis, Konstantinos Tserpes, and Giannis Spiliopoulos. Real-time maritime anomaly detection: detecting intentional ais switch-off. *International Journal of Big Data Intelligence*, 7(2):85–96, 2020.

[19] Ioannis Kontopoulos, Iraklis Varlamis, and Konstantinos Tserpes. A distributed framework for extracting maritime traffic patterns. *International Journal of Geographical Information Science*, pages 1–26, 2020.

[20] Sofia Maria Karadimitriou, Ellen Marshall Reviewer, and Chris Knox Sheffield. Mann-Whitney U Test. In *Sheffield Hallam University*, Community Project, Sheffield, UK.

[21] MarineCadastre.gov. Vessel traffic data. `https://www.marinecadastre.gov/ais/`.

[22] Lucas May Petry, Amilcar Soares, Vania Bogorny, Bruno Brandoli, and Stan Matwin. Challenges in vessel behavior and anomaly detection: From classical machine learning to deep learning. In *Advances in Artificial Intelligence*, pages 401–407, Cham, 2020. Springer International Publishing.

[23] Fabio Mazzarella, Alfredo Alessandrini, Harm Greidanus, Marlene Alvarez, Pietro Argentieri, Domenico Nappo, and Lukasz Ziemba. Data fusion for wide-area maritime surveillance. page 5.

[24] Fabio Mazzarella, Michele Vespe, Alfredo Alessandrini, Dario Tarchi, Giuseppe Aulicino, and Antonio Vollero. A novel anomaly detection approach to identify intentional AIS on-off switching. 78:110–123.

[25] Fabio Mazzarella, Michele Vespe, Dario Tarchi, Giuseppe Aulicino, and Antonio Vollero. AIS reception characterisation for AIS on/off anomaly detection. page 7.

[26] Fabrizio Natale, Maurizio Gibin, Alfredo Alessandrini, Michele Vespe, and Anton Paulrud. Mapping fishing effort through AIS data. 10(6):e0130746. Publisher: Public Library of Science.

[27] International Maritime Organization. Solas convention treaty. `https://www.imo.org/en/About/Conventions/Pages/International-Convention-for-the-Safety-of-Life-at-Sea-(SOLAS),-1974.aspx`, 2000 (accessed August 16 2020.

[28] Ziqiang Ou. AIS database powered by GIS technology for maritime safety and security - ProQuest. In *The Journal of Navigation*, volume 61 of *The Royal Institute of Navigation*, pages 655–665. Library Catalog: search.proquest.com.

[29] Giuliana Pallotta, Steven Horn, Paolo Braca, and Karna Bryan. Context-enhanced vessel prediction based on ornstein-uhlenbeck processes using historical AIS traffic patterns: Real-world experimental results. page 8.

[30] Giuliana Pallotta, Michele Vespe, and Karna Bryan. Vessel pattern knowledge discovery from AIS data: A framework for anomaly detection and route prediction. 15(6):2218–2245. Number: 6 Publisher: Multidisciplinary Digital Publishing Institute.

[31] KUCB Radio. Dutch Harbor Remains Nation's Top Port For 21st Consecutive Year.

[32] Kurt D. Schwehr and Philip A. McGillivary. Marine ship automatic identification system (AIS) for enhanced coastal security capabilities: An oil spill tracking application. In *OCEANS 2007*, pages 1–9. ISSN: 0197-7385.

[33] Amir Yaghoubi Shahir, Mohammad A. Tayebi, Uwe Glässer, Tilemachos Charalampous, Zahra Zohrevand, and Hans Wehn. Mining vessel trajectories for illegal fishing detection. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 1917–1927.

[34] Amílcar Soares, Renata Dividino, Fernando Abreu, Matthew Brousseau, Anthony W Isenor, Sean Webb, and Stan Matwin. Crisis: Integrating ais and ocean data streams using semantic web standards for event detection. In *2019*

International conference on military communications and information systems (ICMCIS), pages 1–7. IEEE, 2019.

[35] B.J. Tetreault. Use of the automatic identification system (AIS) for maritime domain awareness (MDA). In *Proceedings of OCEANS 2005 MTS/IEEE*, pages 1590–1594 Vol. 2. ISSN: 0197-7385.

[36] Iraklis Varlamis, Ioannis Kontopoulos, Konstantinos Tserpes, Mohammad Etemad, Amilcar Soares, and Stan Matwin. Building navigation networks from multi-vessel trajectory data. *GeoInformatica*, 25(1):69–97, 2021.

[37] Iraklis Varlamis, Konstantinos Tserpes, Mohammad Etemad, Amílcar Soares Júnior, and Stan Matwin. A network abstraction of multi-vessel trajectory data for detecting anomalies. In *Proceedings of the Workshops of the EDBT/ICDT 2019 Joint Conference, EDBT/ICDT 2019, Lisbon, Portugal, March 26, 2019.*, 2019.

[38] Wikipedia. Universal transverse mercator coordinate system. `https://en.wikipedia.org/w/index.php?title=Universal_Transverse_Mercator_coordinate_system`, accessed on 2021-02-16. Page Version ID: 991135085.

[39] Yimin Ye and Nicolas L. Gutierrez. Ending fishery overexploitation by expanding from local successes to globalized solutions. 1(7):1–5. Number: 7 Publisher: Nature Publishing Group.

[40] Daiyong Zhang, Jia Li, Qing Wu, Xinglong Liu, Xiumin Chu, and Wei He. Enhance the ais data availability by screening and interpolation. In *2017 4th International Conference on Transportation Information and Safety (ICTIS)*, pages 981–986, Aug 2017.