NOVA IMS Information Management School

MDSAA

Master's Degree Program in

Data Science and Advanced Analytics

Maritime Anomaly Detection:

The use of AIS data to identify Dark Activity

Luís Filipe Freire Fernandes

Master Thesis

presented as a partial requirement for obtaining a Master's Degree in Data Science and Advanced Analytics

NOVA Information Management School
Instituto Superior de Estatística e Gestão de Informação
Universidade Nova de Lisboa

NOVA Information Management School
Instituto Superior de Estatística e Gestão de Informação
Universidade Nova de Lisboa

Maritime Anomaly Detection:

The use of AIS data to identify Dark Activity

by

Luís Filipe Freire Fernandes

Master Thesis presented as a partial requirement for obtaining the Master's degree in Data Science and Advanced Analytics, with a specialization in Data Science.

Supervised by

Supervisor: PhD Leonardo Vanneschi, NOVA Information Management School

Co Supervisor: PhD Rui Gonçalves Deus, Escola Naval

July, 2024

STATEMENT OF INTEGRITY

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism or any form of undue use of information or falsification of results along the process leading to its elaboration. I further declare that I have fully acknowledged the Rules of Conduct and Code of Honor from the NOVA Information Management School.

[Lisboa, June 21, 2024]

ACKNOWLEDGEMENTS

I would like to begin by expressing my deepest gratitude to my family and closest friends for their unwavering support and patience throughout the development of this research. I extend my thanks to the Portuguese Navy and the DAGI department for allowing me to delve into the field of Data Science and Advanced Analytics and complete this master's program. I am especially grateful to my co-supervisor, Professor Rui Gonçalves Deus, for granting me access to the valuable data sources essential for this study. Additionally, I appreciate Eng. Tiago Giesta Martins for his insightful suggestions and guidance, which significantly contributed to the direction and success of this research. I would also like to sincerely thank my fellow colleagues in this master's program. Their support and guidance have been invaluable, particularly as I navigated through this new and initially unfamiliar area of study. Lastly, I wish to convey my heartfelt thanks to my supervisor, Professor Leonardo Vanneschi, for his exceptional guidance, support, and motivation. His exemplary teaching and inspirational approach have been a constant source of encouragement since the first class he lectured to me. His influence has profoundly shaped my academic journey and aspirations.

ABSTRACT

Maritime security is critical to global trade, transportation, and defence. It necessitates advanced methods for monitoring and detecting abnormal behaviours, especially "dark activity", where vessels deactivate their Automatic Identification System (AIS) transponders to evade detection. This research aims to develop robust machine learning methodologies to identify anomalies within AIS data, focusing on detecting dark activity. The study leverages a comprehensive dataset comprising real AIS data from 2020, focusing on the Mediterranean Sea, provided by the Navy Information Analysis and Management Department (DAGI) and the National Maritime Authority (AMN). This dataset, consisting of over 330,000 entries, presents a significant challenge due to its highly imbalanced nature, with instances of dark activity constituting a mere fraction of the total data. Overcoming this imbalance was crucial to the success of the research. Advanced preprocessing techniques such as oversampling and synthetic sampling were essential to prevent the models from being biased towards the majority class and to ensure effective learning of minority class patterns. This study employs supervised and unsupervised machine learning methods to tackle different aspects of anomaly detection. Supervised models were primarily used to classify dark activity instances, while unsupervised models were implemented to detect general anomalies without using predefined labels. Evaluation metrics focused on F1-Score and recall for supervised and Silhouette Score for unsupervised methods. These models were deployed using FastAPI, enabling real-time classification and anomaly detection from new AIS data. By addressing the significant challenge posed by the highly imbalanced dataset and integrating advanced machine learning techniques, this study's findings demonstrate the potential of machine learning in enhancing maritime surveillance, where advanced stacking methods were able to classify dark activity cases with an outstanding level of certainty, making a substantial contribution to naval security offering practical solutions for identifying and responding to dark activities, ultimately enhancing the safety and security of naval operations.

KEYWORDS

Automatic Identification System (AIS); Anomaly Detection; Dark Activity; Imbalanced Data; Machine Learning; Maritime Security; Supervised Learning; Unsupervised Learning.

Sustainable Development Goals (SDG):

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS AND ACRONYMS

ADASYN	Adaptative Synthetic Sampling Approach for Imbalanced Learning
AIS	Automatic Identification System
AMN	National Maritime Authority
AUC-ROC	Area Under the ROC Curve
BS	Base Station
CNN	Convolutional Neural Networks
CoG	Course over Ground
CRISP-DM	Cross Industry Standard Process for Data Mining
DAGI	Information Analysis and Management Department
DBSCAN	Density-based Spatial Clustering of Applications with Noise
ET	Extra Trees
GP	Genetic Programming
GS	Genetic Semantic
IMO	International Maritime Organization
KB	Knowledge Based
KB	Knowledge-Based
KNN	K-Nearest Neighbours
LR	Logistic Regression
ML	Machine Learning
MLPs	Multi-Layered Perceptron
MMHS	Message Handling System
MMSI	Maritime Mobile Service Identity
PML	Parallel meta-learning
PML	Parallel meta-learning

RF	Random Forest
RoT	Rate of Turn
RSSI	Received Signal Strength Information
SAR	Search and Rescue
SMOTE	Synthetic Minority Oversampling Technique
SO	Specific Objective
SoG	Speed over Ground
SOLAS	Safety of Life at Sea
SOM	Self-organizing map
SRQ	Specific Research Question
STGP	Standard Genetic Programming
SVM	Support Vector Machine
VHF	Very High Frequency
VTS	Vessel Traffic Services
XGB	XGBoost

1. INTRODUCTION

This section outlines the motivation, scope, and structure of this research. The primary motivation is to enhance maritime security and monitoring by classifying and differentiating intentional concealment of activities in the naval environment using machine learning techniques. The motivation is elaborated in Section 1.1, while the scope is discussed in Section 1.2. The impact of this research is briefly covered in Section 1.3. Finally, Section 1.4 provides an overview of the structure for the remainder of this work.

1.1. MOTIVATION

The proliferation of electronic equipment and systems, encompassing technologies such as radars, Automatic Identification Systems (AIS), video, and infrared cameras, has significantly enhanced maritime awareness (Castaldo et al., 2014). Unlike in the past, when maritime surveillance suffered from a lack of data, systems now produce unprecedented amounts of data. These technologies can fall into two distinct categories: one being cooperative self-reporting systems, where vessels voluntarily transmit information about their identity and location (these include AIS) and non-cooperative surveillance systems, which detect, track, and identify vessels without any cooperation from the boats themselves (e.g. radar, and video) (Kontopoulos et al., 2020).

This surge in data acquisition and surveillance capabilities is vital for numerous reasons. Firstly, it vastly improves navigational safety, allowing for real-time tracking of vessels and early detection of potential collision risks (Durlik et al., 2023). Additionally, it enhances security by providing authorities with the means to identify and respond to suspicious or anomalous vessel activities promptly (Byerly et al., 2022). Furthermore, the wealth of data these systems generate facilitates in-depth analysis for optimising maritime routes, minimising fuel consumption, and promoting more efficient and sustainable naval practices.

Despite these advancements, detecting "dark activity"[1] remains a significant concern. Identifying such activities is crucial for preventing illegal actions such as smuggling[2], piracy, and unauthorised fishing (Gwilliam, 2023). Therefore, this research is motivated by the need to develop robust methodologies using machine learning techniques to detect and classify these dark activities effectively.

Moreover, the continuous evolution of these electronic systems holds the promise of further refining maritime awareness and contributing to safer, more secure, and environmentally conscious maritime operations (Li & Yang, 2023).

[1] When vessels intentionally turn off their AIS transponders to evade detection (Görkem et al., 2023)

[2] To import or export (goods) secretly, in violation of the law; (of migrants) The illegal entry of a person into a State Party which the person is not a national or resident in order to obtain a financial or other material benefit (Dictionary.Com | Meanings & Definitions of English Words, 2024);

1.2. PROBLEM DEFINITION

This research addresses the primary question: **"How can AIS data be effectively analysed to detect and identify periods of dark activity in the maritime industry?"** This aligns with the broader objective of **developing a methodology built from machine learning algorithms to detect dark activity using AIS data**. Several specific research questions (SRQs) and objectives (SO) were defined to achieve this general objective.

Specific Research Questions (SRQs):

- **SRQ1:** What are the key features in AIS data that can be used to identify periods of dark activity?
- **SRQ2:** How can data preprocessing techniques be optimised to handle the imbalanced nature of AIS data?
- **SRQ3:** Which supervised machine learning algorithms most effectively detect dark activity in AIS data?
- **SRQ4:** How can unsupervised learning methods contribute to detecting anomalous maritime behaviours, including dark activity?
- **SRQ5:** What are the most effective evaluation metrics for accessing the performance of dark activity detection models?
- **SRQ6:** How can the developed models be employed for real-time detection of dark activity using new data?
- **SRQ7:** What are the main challenges and limitations encountered in detecting dark activity, and how can they be mitigated?

Specific Objectives (SOs):

- **SO1:** To collect and preprocess AIS data from 2020, ensuring high-quality input for analysis.
- **SO2:** To apply feature engineering techniques to extract and select relevant features indicative of dark activity.
- **SO3:** To implement and compare the effectiveness of various supervised machine learning algorithms in detecting dark activity.
- **SO4:** To integrate unsupervised learning approaches for a broader understanding of maritime anomalies.
- **SO5:** To establish robust evaluation criteria for model performance, emphasising the detection of rare events.
- **SO6:** To develop a user-friendly deployment framework using FastAPI, facilitating real-time anomaly detection.
- **SO7:** To document the challenges faced during the research and suggest improvements for future studies.

To tackle these objectives and answer these questions, the research focuses on leveraging real AIS data from 2020, filtered explicitly for the Mediterranean Sea. This dataset, provided by the Navy Information Analysis and Management Department (DAGI) and the National Maritime Authority (AMN), includes both the AIS data capturing vessel information and a complementary dataset of confirmed dark activity instances. The inherent imbalance in the data, where instances of dark activity are significantly outnumbered by regular activity, presents a significant challenge. Addressing this imbalance is crucial for training effective machine learning models.

The research adopts the Cross Industry Standard Process for Data Mining (CRISP-DM) methodology, which provides a structured framework encompassing business understanding, data understanding, data preparation, modelling, evaluation, and deployment. The business understanding phase clarifies the importance of detecting maritime anomalies to enhance security and optimise resource allocation. The data understanding and preparation phases involve exploring and processing the AIS data and dark activity labels to create a robust dataset suitable for machine learning applications.

This research employs various supervised and unsupervised machine learning models. Supervised models, including Random Forest, Gradient Boosting, advanced ensembles, and deep learning architectures, are used to classify dark activity instances. Unsupervised models, such as Self-Organizing Maps (SOM) and Density-Based Spatial Clustering of Applications with Noise (DBSCAN), are implemented to detect general anomalies and determine if dark activity patterns fall within broader anomalous behaviours.

Evaluation metrics such as accuracy, precision, recall, F1-score, Silhouette Score, and R-squared values are used to assess the performance of the models. Particular emphasis is given to the F1-Score and recall for the minority class within the supervised algorithms, as the primary objective is to identify dark activity cases, and to the Silhouette Score in unsupervised algorithms to ensure the effective identification of clusters and their outliers.

The deployment phase involves implementing the developed models using the FastAPI tool, which allows for detecting and classifying dark activity from new AIS data using pre-trained models and their associated preprocessing steps. This practical application underscores the potential for real-world impact, enhancing the capabilities of maritime surveillance systems.

1.3. Work Contribution

This research bridges the gap between practical maritime security challenges and scientific advancements by thoroughly evaluating the field's current state and addressing gaps that hinder practical applications. By following a systematic approach, leveraging available literature, and implementing machine learning techniques, this work contributes significantly to several critical areas within maritime anomaly detection.

The contributions of this research are multi-faceted. Firstly, it enhances methodologies for detecting dark activities using AIS data, providing robust models that can be deployed in real-world scenarios. This includes the development and optimisation of both supervised and unsupervised machine learning. Secondly, the research addresses the challenges of imbalanced datasets through advanced preprocessing techniques, ensuring that the models trained are accurate and reliable. This work contributes to the broader understanding of how to effectively handle and preprocess large-scale and imbalanced data for anomaly detection.

Moreover, deploying these models using FastAPI underscores the research's practical applicability. By providing a user-friendly interface for real-time anomaly detection, the research ensures that maritime authorities can quickly implement and benefit from the developed methodologies. Additionally, the study highlights the importance of evaluation metrics tailored to the specific needs of anomaly detection in maritime contexts.

Overall, this research makes substantial contributions to the field of maritime anomaly detection by developing and validating advanced machine learning models, addressing data preprocessing challenges, and ensuring practical deployment. These contributions are supported by a thorough review of the literature and expert validation, providing comprehensive recommendations for achieving practical viability in the field.

1.4. Report Structure

This report is structured to provide a comprehensive overview of the research, starting from the foundational theories and extending to practical implementations and findings.

Chapter 2 presents a detailed literature review, summarising existing research and theories relevant to maritime anomaly detection using AIS data. The following section, chapter 3 delves into the theoretical background and framework, outlining the key concepts, methodologies, and machine learning techniques employed in this research. This chapter is structured around the CRISP-DM methodology, which guided the research process from data understanding to deployment.

Chapter 4 focuses on the results of the study. This includes insights from the data analysis, performance evaluation of both supervised and unsupervised algorithms, and their deployment for real-time detection. This chapter also provides a comprehensive discussion of the findings, the challenges encountered, the limitations of the study, and recommendations for future research. The last section, chapter 5, encompasses all the research conclusions, including the models' efficacy, classification accuracy, and patterns identified. This chapter reflects on the overall contributions of the work and its implications for maritime anomaly detection.

By following this structure, the report aims to provide a clear and logical progression from theoretical foundations to practical applications, ensuring a thorough understanding of the research conducted and its contributions to the field of maritime anomaly detection.

2. LITERATURE REVIEW

This chapter comprehensively reviews the literature and legislation relevant to maritime anomaly detection using AIS data.

2.1. AUTOMATIC IDENTIFICATION SYSTEM (AIS)

2.1.1. Background and Legislation

AIS is a system primarily employed in maritime settings, originated in the early 1990s intending to complement radar and visual observation with ship-to-ship awareness (Watch, 2017). The International Maritime Organization (IMO) outlined three fundamental purposes for the AIS (International Maritime Organization (IMO), 2019). Firstly, it serves as a collision avoidance tool by facilitating the exchange of AIS messages among nearby vessels. It provides crucial information, such as vessel identities and precise locations, to enhance situational awareness and navigation. Secondly, AIS contributes to on-shore Vessel Traffic Services (VTS), aiding in the guidance of vessels through congested or hazardous areas like ports and sea routes. Additionally, AIS transceivers on survival crafts and life jackets play a pivotal role in Search and Rescue (SAR) operations. Lastly, coastal states utilise AIS to identify vessels and their cargo within their territorial waters.

The International Convention on Safety of Life at Sea (SOLAS), 1974, sets out the mandatory provisions for installing of navigational systems and equipment aboard ships tailored to specific vessel types. In 2000, IMO, as part of a new revised chapter V, adopted a new requirement for all boats to carry AIS capable of automatically providing information about the ship to other ships and coastal authorities.

According to this convention's regulations, all ships of 300 gross tonnage and upwards engaged on international voyages and cargo ships of 500 gross tonnage and upwards not engaged on international voyages and passenger ships, irrespective of size, shall be fitted with an AIS. It is essential to mention that this equipment shall provide automatically to appropriately equipped shore stations, other ships and aircraft details such as the ship’s identity, type, position, course, speed, navigational status, and other safety-related information. The motivation for adopting AIS was its autonomous ability to identify other AIS-fitted vessels and provide extra precise details about target ships that can be used in collision avoidance (Harati Mokhtari et al., 2008).

AIS data is available to the public domain via transmission from AIS transponders utilising two Very High Frequency (VHF) channels (frequencies 161.975 and 162.025 MHz), including static and dynamic information about vessels (Ribeiro et al., 2023). A total of 23 distinct AIS messages exist, with dynamic messages exhibiting a higher frequency—transmitted every two seconds to three minutes per message per transponder. In contrast, static messages follow a less frequent pattern with a periodicity of six minutes (Riveiro & Falkman, 2009).

Messages 1, 2, and 3 encompass essential details related to a vessel's position, including timestamp, Maritime Mobile Service Identity (MMSI), navigational status, rate of turn (RoT), Speed over Ground (SoG), position accuracy, latitude, longitude, Course over Ground (CoG), and heading. Conversely, static information in message 5 includes the vessel's MMSI, IMO number, type, size, name, and callsign, which are the most useful ones, among others (Dobrkovic et al., 2015). Additional information regarding AIS messages can be consulted in Table 1.

Table 1 - Information announced via AIS messages. Adapted from Wolsing et al. (2022)

Type	Data
Static	• MMSI number • Call sign & call name • Length & beam • Ship type • Antenna location (aft/bow; port/starboard)
Dynamic	• Ship position (PO), accuracy, and integrity • Time in UTC • Course over ground (COG) • Speed over ground (SOG) • Heading (HE) • Rate of turn (ROT) • Navigational status, e.g., at anchor
Voyage related	• Draught • Hazardous cargo (type) • Destination (DST) and estimated time of arrival
Safety-related	• Text messages

2.1.2. System Weaknesses

AIS data quality is susceptible to three primary categories of issues: errors, which involve the unintentional dissemination of inaccurate information; falsification, which is characterised by the deliberate broadcast of false data; and spoofing, which occurs when an external entity generates or alters data and transmits it through the AIS system (Ray et al., 2016).

The integrity of data within AIS messages is susceptible to various vulnerabilities: there is no robust transmission verification, the transmission is done using a non-secured channel, potential gaps in the crew awareness of certain information, or intentional efforts to conceal specific data. These operations modify and handicap the understanding of maritime traffic (Harati Mokhtari et al., 2008).

Unintentional errors in AIS data may stem from transponder malfunctions, incorrect manual data entry, poor-quality manual inputs, or inaccuracies from external sensors. These errors can impact vessel details such as the vessel's name, physical characteristics, position, or destination. Consequently, the information may become false, incomplete, or even violate norms and physical principles (e.g. latitude values exceeding 90°) (Ray et al., 2016). On the other hand, falsification involves intentionally degrading a message by altering a genuine value with a false one or ceasing message broadcasts, aiming to mislead external observers. Identity theft, disappearances, broadcasting false coordinates and providing inaccurate activity information are examples of falsification (European Commission. Joint Research Centre., 2016). Spoofing, orchestrated by external actors, entails creating entirely false messages and broadcasting them on AIS frequencies (Balduzzi et al., 2014). These activities aim to mislead external observers and sea crews by generating phantom vessels, false closest point-of-approach triggers and fabricated emergency messages (Ray et al., 2016).

2.2. MARITIME DOMAIN AWARENESS

Most of the time, the detection of illegal activities such as piracy, fishing in protected zones, intrusion into economic zones, transhipment of narcotics, and degassing at sea, rely solely on the visual observation of vessels, manual analysis of collected data, and coastal administration officers intuition based on their long-term expertise (Bernabé et al., 2023).

Rogue maritime entities are aware of the forms of surveillance and observation conducted by government and commercial agencies. To avoid being detected many bad actors at sea will seek to avoid detection by visual or AIS means by concealing themselves (Harati Mokhtari et al., 2008). For example, in visual spectrum image data, the bad actor might conduct their activity at night during inclement weather (Jones et al., 2023). Although vessels must carry an AIS transponder, it is not compulsory to switch it on. The most common method employed by individuals seeking to circumvent vessel tracking is simply switching off the AIS transponder. This is colloquially known as “going dark” (Gwilliam, 2023) to hide potential illegal activity (Kontopoulos et al., 2020). Nevertheless, not every vessel operating “dark” necessarily harbours malicious intentions; in certain instances, the decision to remain undetected may be inadvertent or completely benign (Jones et al., 2023). Cargo and tanker vessels may conceal their location to avoid piracy attacks in debatable areas; fishing vessels do so to prevent other fishing vessels from fishing in the same waters; adverse weather conditions swaying AIS transmission; lack of shore-based receiving stations or restricting satellite positions; equipment anomaly; are all examples of situations leading to benign dark activity (Kontopoulos et al., 2020).

2.3. MARITIME ANOMALY DETECTION

Anomalies within AIS tracks refer to behaviours that deviate from the norm or, more precisely, actions that are not anticipated during standard operations (Laxhammar, 2008). According to Lane et al. (2010), five general anomalous behaviours derived from AIS can be defined and expressed in Figure 1:

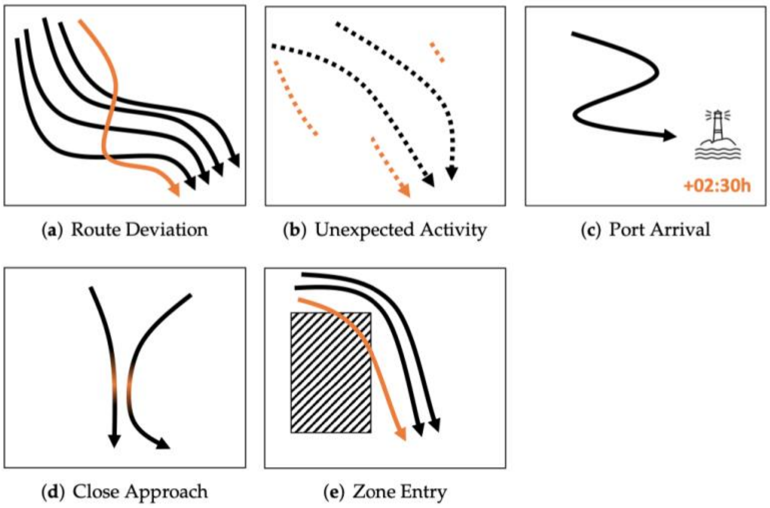

Figure 1 – General five AIS anomaly types derived by Wolsing et al. (2022).

Ships, especially cargo or passenger ships, generally take the most direct path possible between the origin and the destination, making repetitive routes. A deviation from a straight route may then indicate an anomaly (Figure 1 (a)). On the other hand - and the case that will be the focus of this research - unexpected AIS activity (Figure 1 (b)) encompasses two types: AIS signal loss in areas where reception is usually adequate (may indicate intentional on-off switching), and unexpected continuations of routes after an AIS outage (Wolsing et al., 2022). The unexpected port arrival (Figure 1 (c)) can be compared to voyage-related data of the AIS tracks and might represent an anomaly if it does not match. Extended periods of close approaches between vessels (Figure 1 (d)) should typically be rare unless, for instance, there is an emergency. Otherwise, such occurrences might indicate illicit activities like exchanging contraband or drugs (Wolsing et al., 2022). Lastly, anomalies related to zone entry (Figure 1 (e)) occur when vessels enter an area for a substantial duration where they are neither anticipated nor permitted to be, such as marine protected areas or other exclusion zones.

2.3.1. Machine Learning approaches

In their recent work, Bernabé et al. (2023) introduced a self-supervised method that leverages AIS data from geo-marine observation satellites provided by the Norwegian Coastal Administration. The approach involves examining a given ship's behaviour in the open sea, given a window of ω successive AIS messages and a time frame τ. The model is trained to utilise transformer models[3] to predict whether an AIS message is anticipated to be received in the subsequent timeframe τ. The author then compares its prediction with the actual observation performed in near real-time regarding the reception of the AIS message. For a visual exploration of the model architecture, refer to Figure 2.

[3] Self-supervised deep learning architectures designed for efficiently process and understand sequences (Vaswani et al., 2023)

From the maritime anomaly detection perspective, the model yields four distinct scenarios: 1) if the model predicts no AIS message, and indeed, no message is received during the timeframe, the trajectory is confidently classified as ordinary; 2) if the model predicts no AIS message, but a message is received within the timeframe, this indicates a prediction error in the model.; 3) if the model predicts the reception of an AIS message within the timeframe, and a message is indeed received, the trajectory is safely classified as ordinary.; 4) if the model predicts an expected AIS message within the timeframe, but no message is received, the trajectory is labelled as abnormal. This case interests coastal administrations, prompting them to thoroughly analyse the ship's trajectory to confirm or deny the abnormal dark activity (Bernabé et al., 2023).

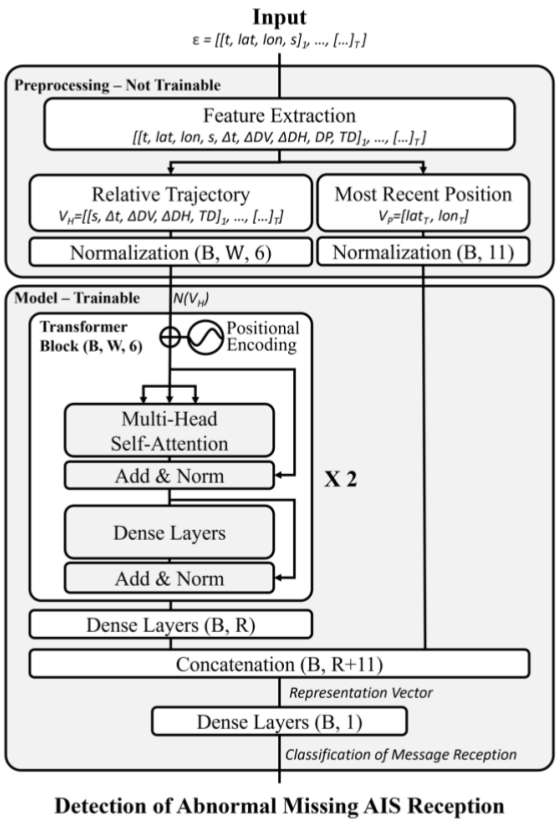

Figure 2 – Model Architecture: the relative vessel trajectory is encoded via a transformer network and forms a representation vector for classification with the most recent position (Bernabé et al., 2023).

Wei et al. (2022) propose a maritime anomaly detection method based on a Support-Vector Machine (SVM) to address the problem of recognising abnormal vessel traffic behaviour. The approach involves three key steps: Firstly, a unique feature extraction method rooted in statistical theory is applied to each trajectory. This method captures spatiotemporal and motion characteristics, comprehensively considering trajectory features in a high-dimensional space. Secondly, the DBSCAN algorithm, complemented by the designed feature extraction for measuring similarity, recognises vessel traffic patterns based on spatiotemporal and motion features. Finally, an enhanced SVM, incorporating a hybrid kernel function, is designed to effectively identify abnormal behaviours in maritime traffic.

The authors' conclusion highlights the effectiveness of the DBSCAN algorithm in recognising vessel traffic patterns. Additionally, the SVM demonstrates high accuracy in detecting the majority of anomalies.

Yan et al. (2022) adopt a machine learning-based strategy for ship classification and anomaly detection in spaceborne AIS data. Conducting ship type analysis, they identify cargo ships, tanker ships, fishing ships, passenger ships, and tug ships as representatives for model training due to their substantial presence in the data. The authors categorise ship features into geometric and behavioural features, leveraging AIS messages, including dynamic data like longitude, latitude, SoG, and CoG. Employing SVM and Random Forest (RF) algorithms, they find that RF outperforms SVM, and integrating geometric and behavioural features enhances classification accuracy.

In the anomaly detection domain, Yan et al. (2022) uncover instances where classification labels deviate significantly from message types, suggesting potential evasion strategies involving the manipulation of AIS messages by certain ships.

Handling big data remains a significant hurdle many authors have yet to overcome. The present rate of AIS message reception is approximately 600 million per month from various sources. When applied under these circumstances, conventional machine learning algorithms encounter difficulties in building effective models and achieving satisfactory results (Wang et al., 2014). To address this challenge, Wang et al. (2014) implemented a two-level approach. In the first level, they employed an unsupervised technique to label vessels' normal and abnormal position points based on raw AIS data using DBSCAN. In the second level, they trained a supervised-learning Parallel Meta-Learning (PML)[4] algorithm, running on Hadoop[5] and leveraging the labelled data generated in the first level. This two-level methodology aimed to efficiently handle the large volume of AIS data and enhance the overall modelling process.

2.3.2. Rule-based approaches

Kontopoulos et al. (2020) propose an algorithm that utilises time-outs and network coverage, representing the surveillance area to address the challenge of AIS switch-off. This area is divided into cells, each receiving varying numbers of AIS messages – helping to identify whether an area has coverage or not. Lack of coverage in certain areas, referred to as “dark cells”, may result from weather conditions, absence of terrestrial base stations or intentional jamming of receivers. Among these scenarios, only the lack of terrestrial base stations is

[4] ML approach that trains multiple models simultaneously to quickly adapt and generalize to new tasks by leveraging shared experience and knowledge (Hospedales et al., 2020)

[5] Open-source Framework for distributed storage and processing of large datasets using clusters of computers(*Apache Hadoop*, n.d.)

known in advance; the other two can be estimated in real-time based on the number of AIS messages per cell.

The algorithm leverages network coverage to infer whether communication gaps are due to vessels intentionally switching off their transponders or genuine gaps in AIS messages within a cell. In the algorithm, positional information of vessels is distributed to individual actors within the Akka[6] system. Each actor monitors a single vessel, storing the last known position, SoG and Cog. Depending on their speed, vessels may transmit messages at varying rates, ranging from once every 2 seconds to once every 3 minutes. When an actor has not received a message for a defined time threshold (10, 30, 60 or 120 minutes in this case), it evaluates whether the vessel is in a dark cell or if its AIS transponder was switched off. Upon receiving this message, the algorithm calculates a possible position projection based on the last known position, speed and heading at the current timestamp. If the projection is not within a dark cell in the network coverage, the last known position is flagged as the start of an AIS switch-off. Otherwise, the projection inside a dark cell is flagged as the beginning of an AIS communication gap (Kontopoulos et al., 2020).

Mazzarella et al. (2016) employed another approach to analyse AIS data to identify whether the absence of measurements indicates an anomaly or is merely a result of communication channel dropout. As explained before, dropout instances can arise from various communication challenges, such as Line of Sight issues or intentional transmission termination by malicious users, posing a challenge due to the fluctuating nature of the AIS signal.

The methodology's flowchart, illustrated in Figure 3 for a single AIS-BS, outlines the process. Historical AIS data is employed to construct a Knowledge-Based (KB) model, capturing the base station (BS) coverage pattern by considering the Received Signal Strength Indicator (RSSI), vessel positions, and distances from the BS. This model plays a dual role: aiding in data reconstruction alongside observed test tracks and establishing a threshold to categorise the state space into regions representing normal (H_0) and anomalous (H_1) behaviour. Concurrently, the analysis periodically reviews observed test tracks to identify potential dropouts. When a dropout occurs in an initiated track, the AIS on/off detection procedure activates. Initially, an estimate of the vessel's position is obtained from the tracking system, allowing for the reconstruction of the missing RSSI value using the coverage pattern model and observed test track. A threshold test on the reconstructed sample triggers an alert if anomalies are detected, indicating that the dropout is likely not associated with the communication channel, necessitating further investigation (Mazzarella et al., 2016).

[6] Toolkit for building highly concurrent, distributed and resilient message-driven applications on Java Virtual Machine (*Akka Guide :: Akka Guide*, n.d.)

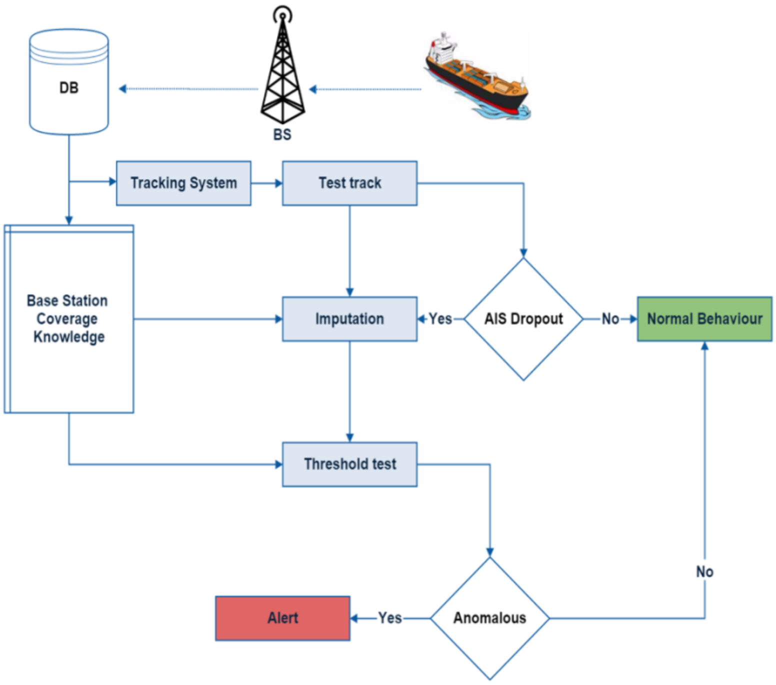

Figure 3 – General flowchart of the AIS on/off detection approach (Mazzarella et al., 2016).

2.3.3. Genetic Programming Approaches

Some researchers, such as Vanneschi et al. (2015) and Dobrkovic et al. (2018), employ genetic algorithms (GA) to predict vessel positions and pattern recognition using AIS data.

In the case of Vanneschi et al. (2015), the objective is to construct models predicting a vessel's position solely based on its previous positions within a given time interval. This approach proves valuable for emergency assistance and enhancing the efficiency of tracking ships engaged in illegal activities. The study utilises geometric semantic crossover and mutation, applying linear scaling to each individual in the population (GSGP-LIN). Comparisons with standard genetic programming (STGP), Geometric Semantic GP (GSGP), and three non-evolutionary machine learning techniques (Linear Regression, Multi-layered Feed-Forward Artificial Networks, and Support Vector Machines) were conducted. The authors concluded that GSGP-LIN outperformed all other methods, demonstrating high precision and accuracy and providing a reliable model describing vessel behaviour.

Dobrkovic et al. (2018) approach using Genetic Algorithms (GA) slightly differently, focusing on extracting sailing patterns through sequential waypoint discovery. As maritime routes tend to change over time due to the influence of other external factors such as weather and tides, the ability of the GAs to evolve with the new data gradually will make it the best candidate for providing real-time naval pattern awareness and accessing imperfect publicly available AIS data. In their methodology, GA is applied to discover pairs of waypoints, making extracting routes possible by checking the sequence of waypoints travelled by vessels for each MMSI and creating edges between them. The weight of an edge increases each time a unique MMSI passes through two nodes. Routes may appear and disappear based on external influences, causing nodes and edges to transition between active and inactive states. Notably, the algorithm is proficient in extracting potential routes even from missing data, leveraging available starting and ending points.

2.3.4. Hybrid Approaches

Venskus et al. (2017) present a novel self-learning adaptive classification algorithm combining a Self-Organizing Map (SOM) with a virtual pheromone to detect abnormal vessel movements in maritime traffic. The authors compared five neighbourhood functions: Gaussian, triangular, cut Gaussian, bubble, and Mexican hat, evaluating their impact on the classification results achieved by the modified SOM method.

This research uses SOM for data clustering and graphical result presentation. However, to classify observations, the authors introduced the biologically inspired concept of virtual pheromones in the last epoch of the model. The virtual pheromone intensity of a winning neuron increases with the number of observation vectors assigned to it, identifying the abnormal movement in cases outside the defined threshold. The process can be observed in Figure 4.

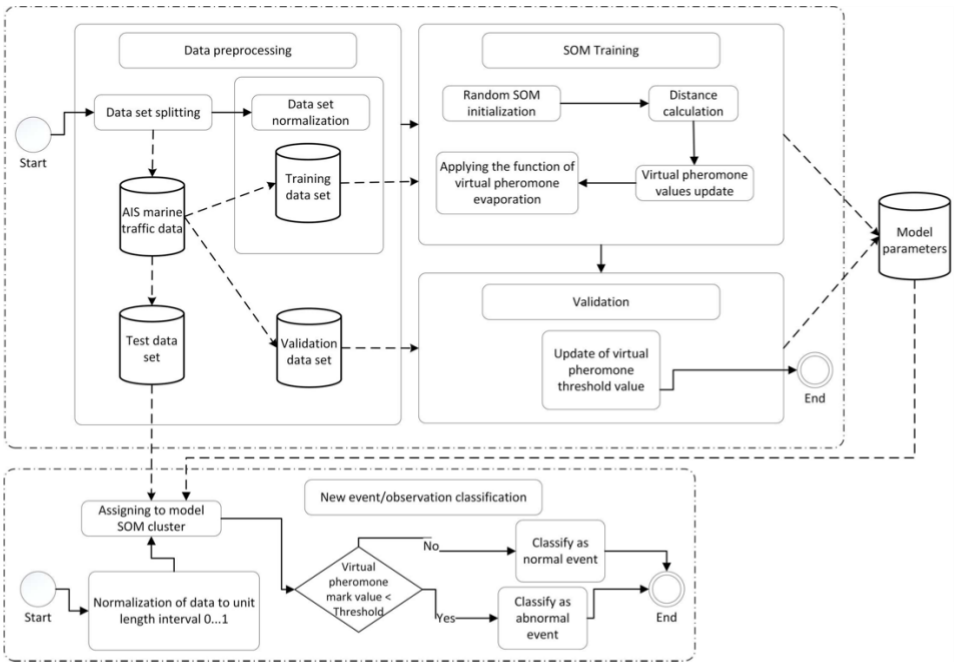

Figure 4 – Integration of a self-organizing map and virtual pheromone (Venskus et al., 2017)

Görkem et al. (2023) present an innovative fusion of machine learning and rule-based techniques to detect dark activity. Adopting a localised perspective, the method centres on discerning dark activity based on signals from nearby ships. When a vessel emits an AIS signal, neighbouring vessels equipped with AIS receivers in the coverage area receive it. The absence of a received signal, particularly after initial detection, may suggest that the transmitting vessel has either moved out of coverage or entered a dark activity phase. Vessels, designated as "detector vessels," possess the capability to validate the activities of their nearby counterparts.

The authors introduce two distinctive approaches, illustrated in Figure 5. The first, a rule-based dark activity detection (R-DAD) algorithm, involves a detector vessel continuously monitoring AIS signals from nearby ships, recording their transmitted parameters. Employing fundamental physics principles, the vessel calculates the expected positions of these nearby ships in the subsequent time step, storing both calculated positions and corresponding AIS parameters. In each time step, the detector vessel checks whether the estimated positions from the prior step fall within its reception range. If a vessel's estimated position is within range but no corresponding AIS signal is received, it triggers a potential alert for detecting dark activity. The second approach, machine learning-based dark activity detection (ML-DAD), focuses on enhancing accuracy. Each data instance is paired with a target value indicating dark activity presence, utilising details about the detector vessel's course, heading, speed, and those of the selected nearby vessel. This dataset is leveraged to train an ML model through a supervised learning algorithm, predicting the dark activity presence of nearby ships not transmitting AIS signals, emphasising decision-making over location estimation.

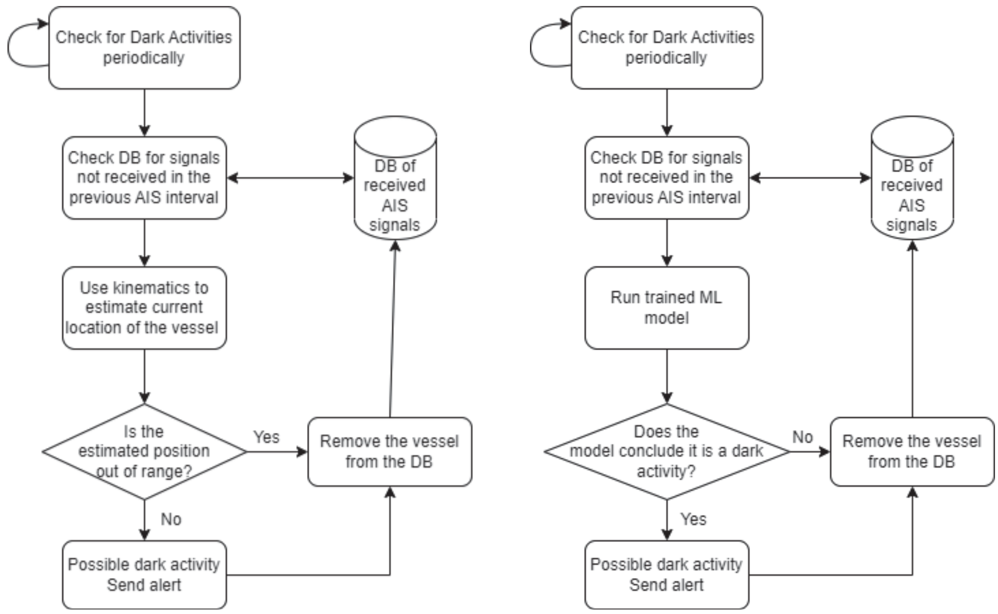

Figure 5 – Flowcharts of the Rule-based Dark Activity Detection (R-DAD) and Machine-learning-based Dark Activity Detection (ML-DAD), respectively (Görkem et al., 2023).

2.4. A COMPARATIVE ANALYSIS OF METHODS

Table 2 - Summarized survey results of 14 anomaly detection approaches for maritime AIS tracks. Adapted from Wolsing et al. (2022)

Method	Publication		Goal	Features								Scope			Dataset
	Authors	Year		t	PO	COG	SOG	HE	Type	Eng	Others	Region	Vessel	Time	
Machine-Learning	Bernabé et al.	2023	Dark Activity	●	●	●	●	●	○	○	○	●	○	●	Priv
Machine-Learning	Wei et al.	2022	Abnormal route	●	●	●	●	●	○	○	○	●	○	●	Pub
Machine-Learning	Yan et al.	2022	Ship Class.	○	●	●	●	●	●	○	○	○	●	○	Pub
Machine-Learning	Zhen et al.	2017	Abnormal route	●	●	●	○	○	○	○	○	●	○	●	Priv
Machine-Learning	Nguyen et al.	2022	Abnormal route	●	●	●	○	○	○	○	○	●	●	●	Priv
Clustering	Wang et al.	2014	Abnormal route	●	●	○	●	●	○	○	○	●	○	○	Priv
Clustering	Zissis et al.	2020	Abnormal route	●	○	○	○	○	○	○	●	●	●	○	Priv
Rule-Based	Kontopoulos et al.	2020	Dark Activity	●	●	●	●	●	○	○	○	●	○	●	Priv
Rule-Based	Mazzarella et al.	2016	Dark Activity	●	●	○	○	○	○	○	●	●	○	●	Pub
Rule-Based	Terroso-Saenz et al.	2016	Abnormal route	○	●	○	○	○	○	○	●	●	○	●	Pub
GP	Vanneschi et al.	2015	Vessel position	●	●	●	●	●	○	○	○	●	○	○	Priv
GP	Dobrkovic et al.	2018	Sailing patterns	●	●	○	○	○	○	○	●	●	○	○	Pub
Hybrid	Venskus et al.	2017	Abnormal route	●	●	●	●	●	○	○	○	●	○	○	Pub
Hybrid	Görkem et al.	2023	Dark Activity	○	●	●	●	●	○	●	○	●	○	○	Pub

Features: *Timestamp (t); Position (PO), Course over ground (COG), Speed over ground (SOG), Ship type (Type); Features obtained from feature engineering (Eng) and other features considered (Others)*
Dataset availability: *Public (Pub) or private (Priv)*

Table 2 presents a condensed overview of the outcomes from fourteen anomaly detection approaches designed for maritime AIS tracks. Despite the diverse methodologies, significant shared characteristics exist, particularly in identifying abnormal routes and dark activities, which emerge as predominant anomaly types. Additionally, a noteworthy observation is that many detectors are region-specific, necessitating re-training for application in different geographical areas. It is evident that recent studies predominantly fall within the realm of machine learning.

3. METHODOLOGY

The methodology chapter serves as the blueprint for navigating the intricate landscape of research, outlining the systematic approach undertaken to address the core objectives of the study.

3.1. Research Framework

The research framework adopted for this work is the Cross Industry Standard Process for Data Mining (CRISP-DM). The choice of CRISP-DM was deliberate, considering its widespread acceptance and usage within both industry and research domains (Ramirez, 2021). It offers a comprehensive framework comprising six distinct phases, facilitating the systematic exploration and extraction of insights from large datasets. Moreover, CRISP-DM's iterative nature aligns well with the data's dynamic and evolving nature, allowing for flexibility and adaptation as insights are uncovered and business requirements evolve. Figure 6 shows a suggestion for interacting these steps (Hendriksen, 2022).

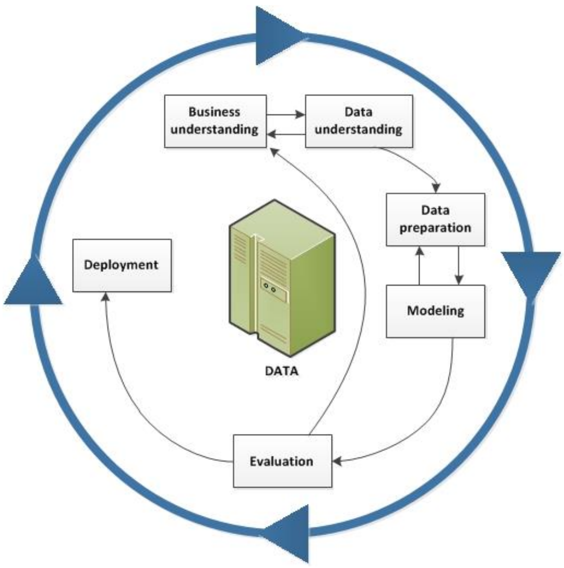

Figure 6 - Steps of the CRISP-DM Methodology figure from (*IBM Documentation*, 2021)

Let us outline each phase briefly:

1. *Business Understanding* – the primary aim of this starting phase is to grasp the project's goals and necessities from a business standpoint. Subsequently, this understanding is translated into a well-defined data mining problem statement and an initial strategy crafted to accomplish the objectives (Chapman et al., 2000);
2. *Data Understanding* - This phase of CRISP-DM involves taking a closer look at the available data for mining. This step is critical in avoiding unexpected problems during the data preparation phase. Data is accessed and thoroughly explored utilising tables and graphics, which can be conveniently organised to facilitate the assessment of the data quality (*IBM Documentation*, 2021);

3. *Data Preparation* – This crucial step involves addressing various issues, including handling missing values, verifying and converting data types, and normalising numeric data, amongst other procedures, to prepare data with the best characteristics for constructing the final dataset. Data preparation tasks will likely be performed multiple times and not in any prescribed order (Chapman et al., 2000). The preceding stage plays a vital role in guiding these decisions, as a lack of understanding of our data and objectives could adversely affect the data preparation process (Ramirez, 2021);
4. *Modelling* - During this phase, various modelling techniques are carefully chosen and implemented, focusing on fine-tuning their parameters to achieve optimal performance. Multiple methods are expected to address the same data mining problem type. Specific techniques may have unique prerequisites regarding the data format, necessitating a revisit to the data preparation phase. This iterative process ensures that the data is appropriately structured to accommodate the specific requirements of each modelling technique, thereby maximising their effectiveness in deriving meaningful insights (Schröer et al., 2021).
5. *Evaluating* – The evaluation stage holds significant importance as it allows us to assess the status and effectiveness of the models and presents an opportunity to uncover new insights. For instance, models may reveal unexpected outcomes or identify extraneous variables impacting the final results (*Most Popular Data Science Methodologies*, 2022). As illustrated by Figure 6, CRISP-DM operates as an iterative process, necessitating a return to the drawing board to adjust the data and models in light of newfound information – this may entail reassessing data quality, model structure, validation procedures or even overarching project objectives (Ramirez, 2021);
6. *Deployment*—Whether the model aims to enhance data understanding or facilitate decision-making, organising and presenting the acquired knowledge in a usable format becomes imperative. Depending on the project's specifications, deployment can range from generating a simple report to implementing a repeatable data mining process across the enterprise (Chapman et al., 2000).

3.2. BUSINESS UNDERSTANDING

As outlined in the Literature Review section, the maritime industry plays a critical role in global trade, transportation, and security, with vessel movements as a vital component of its operations. The Portuguese Navy's mission is to “promote and protect the nation's interests in and through the sea’ linked to defence, security and authority, and development, which allow Portugal the free, sustainable, and fair use at sea” (*A Missão*, n.d.). On the other hand, as the top entity, the National Maritime Authority (AMN) is responsible for coordinating the activities carried out by the Navy and the Coast Guard at a national level in the public and maritime domain spaces under national sovereignty and jurisdiction (*Missão e Competências*, n.d.).

3.2.1. Business Objectives

To align with the missions of the Portuguese Navy and the AMN and to further bolster security standards, this research endeavours to tackle the question: ***"How can AIS data be effectively analysed to detect and identify anomalous behaviours and periods of dark activity in the maritime industry?"*** This aligns seamlessly with the research objective of "***developing a model, constructed from machine learning algorithms, to detect anomalous behaviours, specifically focusing on dark activity, utilising AIS data"***.

3.2.2. Business Success Criteria

The criteria for a successful outcome to the project from a business point of view are:

- *Improved Maritime Security:* Identify anomalous behaviours and periods of dark activity, thus enabling timely intervention and response to potential threats or illegal activities.
- *Enhanced Regulatory Compliance:* The project outcomes have the potential to bolster adherence to legal requirements and standards, thus facilitating compliance with maritime legislation. A more sophisticated identification of 'dark activities' could serve as a strong incentive for intermediaries to comply with regulations to avoid detection.
- *Reduced Risk and Losses:* An associated reduction of maritime incidents, such as smuggling, piracy, and illegal fishing, ultimately safeguarding human and economic interests.
- *Optimised Resources Allocation:* Effective analysis of AIS data to detect and identify abnormal behaviour can enable more efficient allocation of resources, such as patrol vessels or surveillance equipment, by focusing efforts on areas or vessels exhibiting suspicious activities.
- *Improved Decision-Making:* The project's results should support informed decision-making processes within maritime organisations, enabling stakeholders to make proactive decisions to mitigate risks and enhance operational efficiency.

3.2.3. Assess Situation

As the author is a Portuguese Navy Officer, it is imperative to capitalise on the unparalleled resources and expertise. This includes access to specialised personnel and departments, facilitating entry to extensive data repositories and procedural norms, thus laying a robust foundation for analysis. The Portuguese Navy lacks a specific methodology or decision-making process for identifying potential targets for inspection or investigation. Maritime authorities and coastguards often rely on personal knowledge, experience, and expertise, potentially leading to biased decisions. Through meticulous risk assessment, this project ensures a well-informed and strategic approach to planning and decision-making, addressing this critical gap within the maritime domain.

3.3. DATA UNDERSTANDING AND PREPARATION

In pursuit of research efficiency and coherence, the author combined the chapters on understanding the data and preparing them into a unified section. This decision was driven by the recognition that conducting exploration following initial data filtering and feature engineering can enhance the research process. Although traditionally separated within the CRISP-DM framework, integrating these phases offers a more streamlined approach to data analysis.

The data acquisition process for this project involves obtaining AIS signals transmitted in near real-time by ships to antennas or satellites. These signals contain crucial vessel information encoded in Message Handling System (MMHS) format. Upon reception, the Portuguese Navy channels this AIS data to its Information Analysis and Management Department (DAGI), which is processed and compiled into MATLAB files daily.

While historical AIS data spanning at least a decade is available, the focus of this research centres on Satellite data from the year 2020, as it boasts the highest volume of data and offers the most comprehensive information for the research objectives.

In addition to AIS data, the project leverages information from the National Maritime Authority (AMN), providing access to labels of actual dark activity cases. This data offers invaluable insights into real-world occurrences of dark activity; however, it is worth noting that collecting this data is a time-intensive process, as it is outside a dataset format requiring collection day by day, necessitating significant time and resources.

By integrating AIS signals and supplementary data from the AMN, this project aims to construct a comprehensive dataset that forms the basis for subsequent analysis and model development. It will ultimately contribute to enhancing maritime security and decision-making processes within the Portuguese Navy.

This study draws upon two principal data sources. The first is the DAGI, which furnishes AIS information in near-real-time. The second source is the AMN, where records of confirmed instances of dark activity are maintained. The dataset labelled "SAT_20" was compiled from the former, encompassing global AIS data for 2020 obtained via satellite, which was subsequently refined to focus on the Mediterranean Sea. The latter source provided "DA_20," a catalogue of verified dark activity cases specific to the Mediterranean region for the same year. This process will be explained in detail in the next subchapters.

3.3.1. Describe Data

The AIS data is initially stored in MatLab (".mat") files, with one file generated daily. These files were processed using the h5py package in Python (*HDF5 for Python — H5py 3.10.0 Documentation*, n.d.) to compile a dataset spanning the entirety of the year 2020. The compiled dataset comprises 8 927 870 entries, each containing eight distinct features.

The features within the dataset are all of the float64 data types, all numerical data apart from the navigational status that, although float64 type might be considered categorical. It encompasses the following attributes:

- *MMSI—Maritime Mobile Service Identity—serves as the unique electronic identifier for each ship. It is important to note that much information is encoded within the MMSI*, such as the ship's flag, vessel type, and geographical registration area (*AIS Fundamentals*, n.d.).
- *GDH_dec* – Decimal DateTime represents the precise timestamp of each report in MAATLAB decimal time format.
- *Lat* – Indicates the latitude location of the vessel in degrees.
- *Long* – Indicates the longitude location of the vessel in degrees.
- *Nav_stat* – Navigational Status can take on a value from 0 to 15; each number represents a unique piece of metadata associated with the vessel's activity. The more common navigational status are the following (AIS Fundamentals, n.d.): 0 - underway using its engine; 1 – anchored; 2 – Not under command; 3 – Has restricted manoeuvrability; 4 – Ship draught is limiting its movement; 5 – Moored; 6 – Aground; 7 – Engaged in fishing; 8 – Under way sailing.
- *SOG* – Speed over ground, represents the vessel's speed over the ground, measured in knots.
- *COG* – Course over ground, indicates the direction of the vessel's movement over the ground, expressed in degrees from 0 to 360.
- *HDG* – The ship's heading denotes the direction in which the bow is pointed, measured in degrees from 0 to 360.

The second dataset utilised in this research identifies actual cases of 'dark activity' obtained by the author from records maintained by the AMN. It is pertinent to note that these records are confined to the Mediterranean Sea region. The original dataset containing labelled data of identified 'dark activity' consists of 44 244 entries specifically for 2020.

This dataset encompasses 18 columns, all of the type 'object' and 'float64', with the following features being relevant for the research:

- *MMSI* – As previously described, this serves as the unique identifier for each vessel.
- *Activity Type* – This feature characterises the activity observed, with the sole activity of interest labelled 'DARK_ACTIVITY'.
- *Start Date* – Denotes the timestamp marking the commencement of the identified 'dark activity'.
- *End Date* – Indicates the timestamp marking the identified 'dark activity' conclusion.

The utilisation of this dataset enhances the research's scope by providing real-world instances of 'dark activity', thereby facilitating the analysis and development of detection methodologies within the maritime domain.

3.3.2. Data Quality

Certain checks were conducted, albeit cursory, in terms of data quality assessment. This approach was adopted because some anomalous data values may correspond to the type of anomalous behaviour targeted by this research rather than being indicative errors (Zemicheal & Dietterich, 2019).

Among the essential quality checks conducted, the presence of missing values was examined, and it was determined that both datasets were devoid of any missing entries, obviating the need for further action in this regard. Additionally, the timestamps of both datasets were scrutinised to ensure conformity with the stipulated timeframe, revealing that all data entries fell within the year 2020. Furthermore, the MMSI number, unique 9-digit identifiers assigned to individual vessel vessels, were verified to confirm their adherence to the standard format. Moreover, the latitude and longitude values underwent validation to ascertain that they fell within the plausible ranges of -90 to 90 degrees for latitude and -180 to 180 degrees for longitude (GISGeography, 2017). The navigational status was also checked to verify if it was within the 15 possible numbers, concluding that it was coherent.

Despite the examination's superficial nature, data quality scrutiny remains pivotal in upholding the datasets' integrity and reliability. Notably, this process did not compromise the potential identification of abnormal results.

3.3.3. Data Construction

In terms of data processing, several approaches were considered, encompassing the manipulation of data types, the creation of new features, and the selection of relevant features for analysis. Specifically, in the case of the AIS dataset, the following procedures were undertaken:

- Conversion of the MMSI feature to integer type.
- Drop the HDG (Heading) feature, as it is highly related to COG (course over ground) and does not bring any extra insights for the project's purposes.
- The “DateTime” package was used to convert the “GDH_dec” feature into a new attribute named “Formatted_Datetime,” where timestamps were standardised to the “YYYY-MM-DD HH:MM: SS” format and represented as datetime64[ns].
- A novel feature named “Time_Diff” was generated, representing the time difference between consecutive reports for the same MMSI. This new attribute was stored in timedelta64[ns] format.
- A novel feature named “Country” was generated from the ship's MMSI number, representing the ship's flag. This new attribute was stored in string format.
- Filtering of dataset entries to encompass solely the Mediterranean Sea region, restricting latitude values to the range of 30°N to 46°N and longitude values to between 6°W and 36°E (Marine Regions · Mediterranean Sea (IHO Sea Area), n.d.)

Concerning the second dataset containing dark activity labels, the subsequent procedures were executed:

- Conversion of the MMSI feature to integer type.
- Extraction of entries solely featuring the 'DARK_ACTIVITY' label under the 'Activity Type' attribute
- Selection of pertinent features of interest, including 'MMSI,' 'Activity Type,' 'Start Date,' and 'End Date.

Following the data processing procedures outlined above, the two datasets underwent a merging process facilitated by an iterative approach. This merging operation involved aligning rows based on matching MMSI values, ensuring coherence between corresponding entries in both datasets.

As illustrated in Figure 7, by integrating these two datasets, we constructed the final dataset for this research, "SAT20_LAB," which encompasses AIS data paired with the identified instances of dark activity.

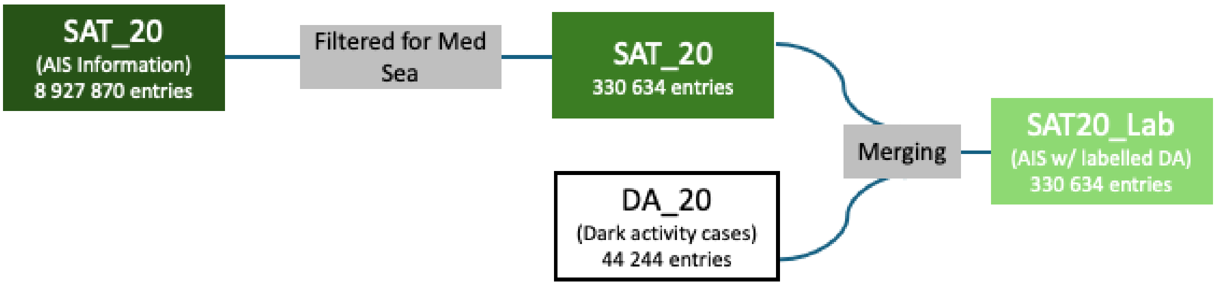

Figure 7 - Illustration of the Data flow stages.

For each MMSI value, the 'Formatted_Datetime' attribute from the AIS dataset was precisely synchronised with the 'Start Date' entries in the labels dataset. As the 'Start Date' signifies the commencement of a dark activity episode, it corresponds to the time when the dark activity begins. By aligning with the 'Start Date', we ensure that each AIS report corresponds to the initiation of a dark activity period, providing a clear temporal reference point for timely intervention and analysis.

Subsequently, the 'DARK_ACTIVITY' label was transformed into a binary target variable to facilitate subsequent analysis. Under this transformation, instances where dark activity was present were assigned a value of 1, while instances lacking dark activity were assigned a value of 0. This systematic merging and transformation process provided a unified dataset conducive to further analysis and modelling for identifying dark activity in maritime operations.

The final dataset for 2020, limited to the Mediterranean Sea, comprises 330,634 entries with a total of 10 columns. These columns include:

- *MMSI:* Maritime Mobile Service Identity (int64)
- *Country:* Flag of the ship
- *Formatted_Datetime:* Date and time in YYYY-MM-DD HH:MM:SS format (datetime64[ns])
- *Lat:* Latitude location in degrees (float64)
- *Long:* Longitude location in degrees (float64)
- *Nav_stat:* Navigational status, represented as a numerical value (float64)
- *SOG:* Speed over ground in knots (float64)
- *COG:* Course over ground in degrees (float64)
- *Time_Diff:* Time difference between consecutive reports for the same MMSI (timedelta64[ns])
- *Dark Activity:* Binary indicator of dark activity presence (int64)

3.3.4. Data Exploration

This subchapter aims to uncover insights, patterns, and relationships in the data that may inform subsequent analytical processes and model development.

Table 3 summarises the descriptive statistics for each variable within the labelled AIS dataset. All variables show a consistent count of 330,634 observations, underscoring the comprehensiveness of the data. The mean latitude and longitude are reported as 36.67 and -4.87, respectively, indicative of the central tendency of vessel positions within the Mediterranean Sea. The navigation status (Nav_stat) has a mean close to zero, suggesting an "undergoing using its engine" predominant status in the dataset.

The average Speed Over Ground (SOG) and Course Over Ground (COG) are 6.92 knots and 155.51 degrees, respectively, which provides a preliminary indication of vessel movement characteristics. The standard deviation for COG is notably high, indicating a wide variation in vessel courses, which is typical given the freedom of maritime navigation.

The 'Dark Activity' variable has a mean of 0.004, confirming the rarity of such events within the dataset. The distribution of this variable is significant, as the maximum value is 1, consistent with binary classification.

Table 3 - Descriptive Statistics of AIS numerical variables.

	Lat	Long	Nav_stat	SOG	COG	Dark Activity
count	330634.000000	330634.000000	330634.000000	330634.000000	330634.000000	330634.000000
mean	36.673102	-4.872181	0.629164	6.916737	155.515347	0.040622
std	2.089439	0.789259	1.752558	6.797694	103.239505	0.197413
min	32.187627	-6.000000	0.000000	0.000000	0.100000	0.000000
25%	35.975067	-5.434978	0.000000	0.100000	61.000000	0.000000
50%	36.128500	-5.181398	0.000000	7.100000	181.000000	0.000000
75%	36.198117	-4.330023	0.000000	12.000000	212.200000	0.000000
max	45.999955	-3.000000	7.000000	102.300000	360.000000	1.000000

Another focal point of analysis entails examining the distribution of the target variable. Notably, the distribution is highly imbalanced, with 13 431 instances of identified dark activity (4.1%) contrasted against 316 841 (95.9%) instances of non-dark activity entries within the merged dataset for 2020 in the Mediterranean Sea region, see Figure 8 (left).

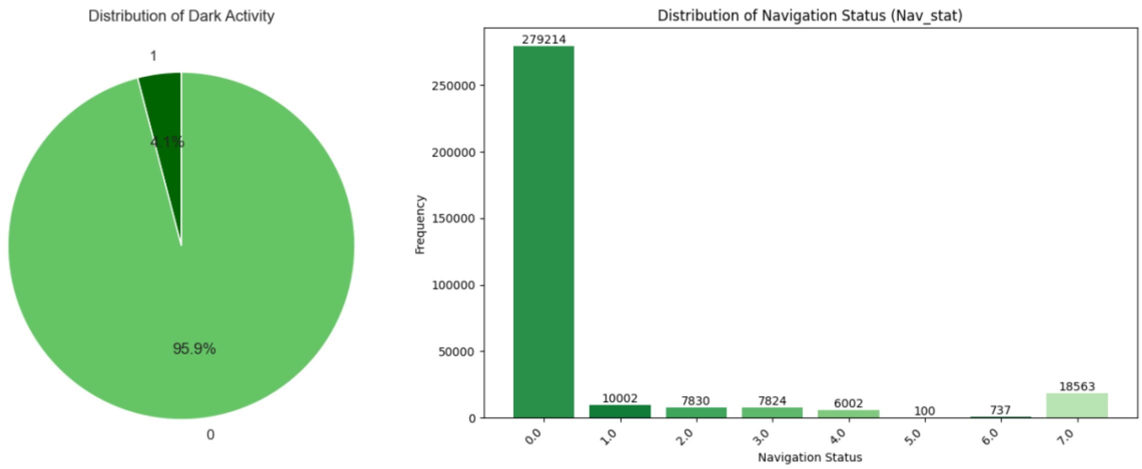

Figure 8 - Distribution of the variable ‘Dark Activity’ (left) and ‘Navigational Status’ (right)

In terms of navigational status, it is evident that the dataset only includes statuses ranging from 0 to 7, with the predominant status being "underway using its engine" (status 0), as observed in Figure 8 (right).

In analysing the distribution as a time series, it is evident that there is a significant peak in dark activity in May, reaching the highest counts. There are several fluctuations over time, with the lowest periods observed at the beginning and end of the year, as illustrated in Figure 9. Additionally, vessels tend to enter a 'dark' state predominantly during overnight hours or periods with reduced light, as shown in Figure 10.

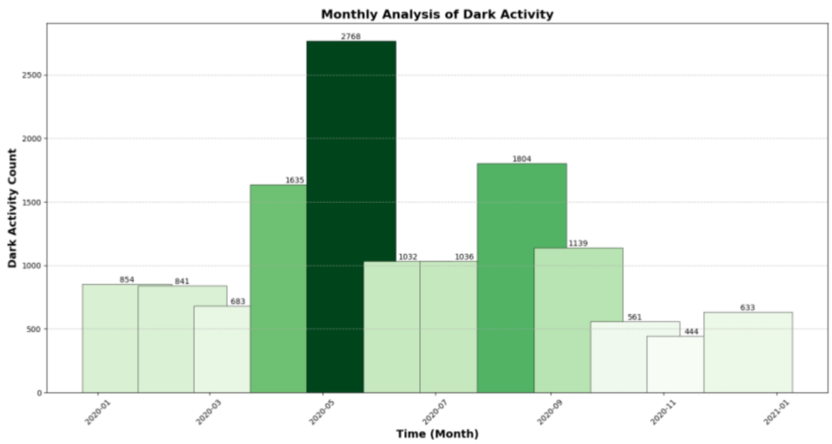

Figure 9 – Monthly Distribution of dark activity throughout 2020.

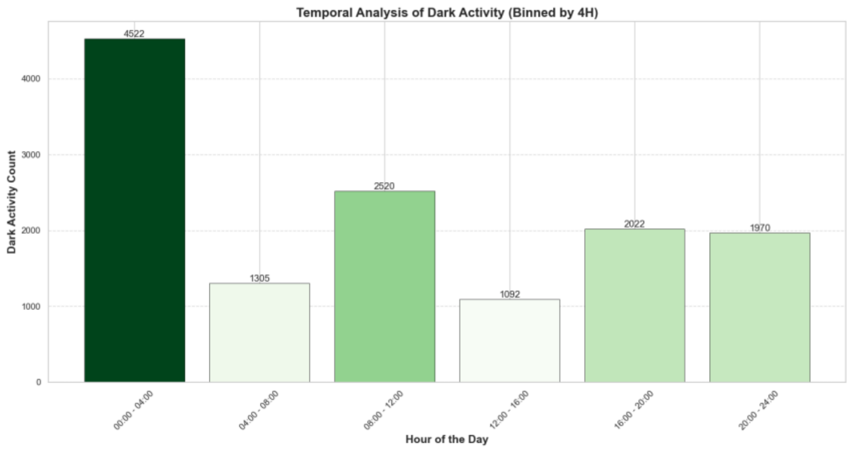

Figure 10 - Distribution of dark activity throughout the day, binned by 4-hour intervals.

Figure 11 ranks the top ten flags based on the number of dark activity cases associated with each. The x-axis quantifies the total count of dark activity incidents, while the y-axis categorises the flags accordingly. According to Figure 12, Malta has the highest recorded incidents, followed by Spain and Liberia, with significantly fewer cases. This graph underscores the disparities in dark activity incidence across different jurisdictions.

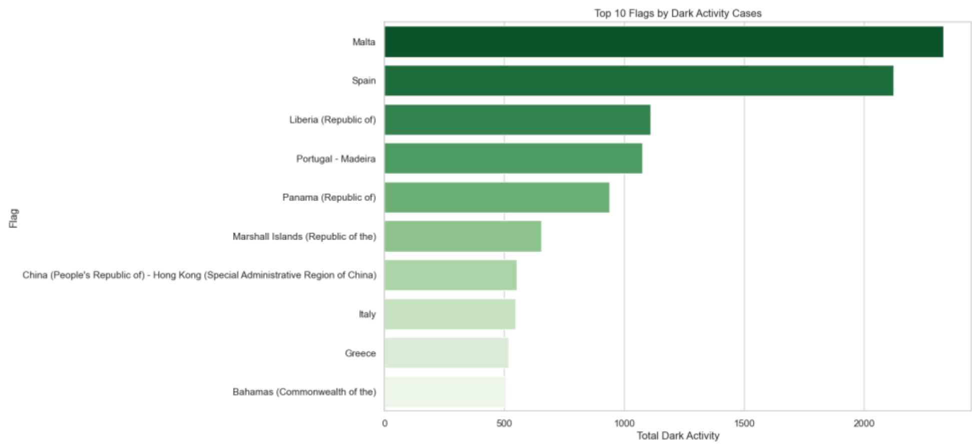

Figure 11 – Distribution of dark activity incidents by flag.

Another analytical approach involved examining and mapping the locations associated with confirmed instances of dark activity. The labelled instances were plotted on a map using the' folium' library (Story, n.d.). As depicted in Figure 12, these instances appeared to be concentrated within a specific zone of the Mediterranean Sea. The high concentration near the Strait of Gibraltar highlights this area as a critical point for maritime traffic, possibly due to its strategic position as an entry point to the Mediterranean. The map also indicates that vessels switch off their tracking systems more frequently in high-traffic areas.

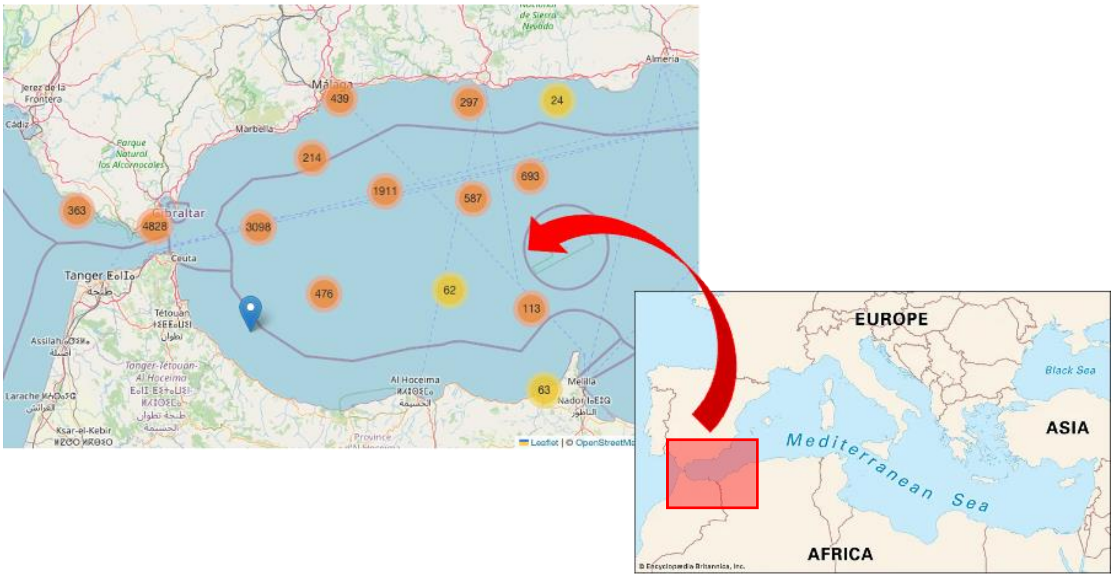

Figure 12 - Identification of the location of ‘Dark Activity’ cases (left) compared to the Mediterranean Sea (right).

3.4. Modelling

The modelling phase is critical in the research process and is central to addressing the research question (Hendriksen, 2022). This stage aligns seamlessly with the research objective of constructing a model using machine learning algorithms to identify anomalies in AIS data. This research aims to extract meaningful patterns and anomalies indicative of dark activity through advanced modelling techniques, thereby enhancing maritime security and safety standards. These models serve as essential stakeholder tools, offering valuable insights and aiding decision-making processes.

3.4.1. Handling Unbalanced Data

This section discusses various techniques for addressing class imbalance in the AIS and dark activity data, including a data level, algorithm level, and hybrid methods.

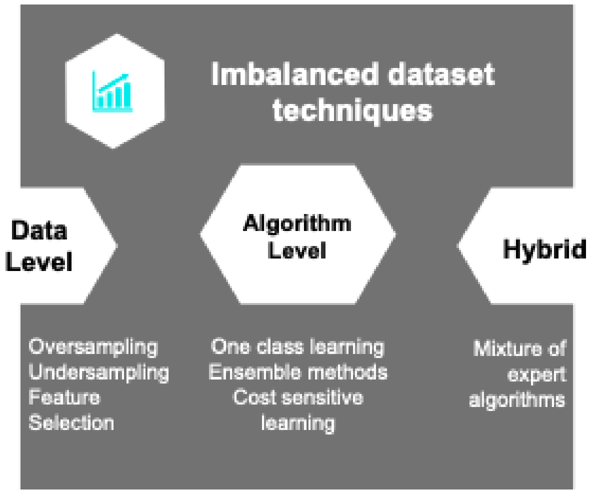

Figure 13 - Methods to handle the imbalanced data, adapted from Spelmen & Porkodi, 2018

The under-sampling method (refer to Figure 13) was excluded from this research due to the risk of losing crucial insights by reducing samples from the majority class and the limited breadth of data available for processing (Spelmen & Porkodi, 2018). Conversely, over-sampling, involving the addition of more examples from the minority class (refer to Figure 14), was employed to mitigate this issue and was validated to be effective in enhancing the results obtained.

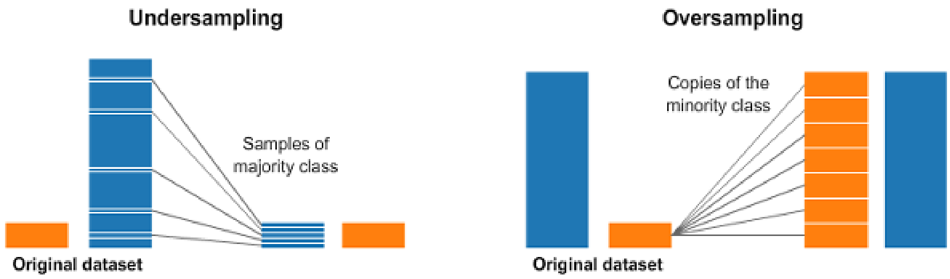

Figure 14 – Under sampling and oversampling methods adapted from guest_blog, 2020.

Another technique explored in this research, which has been proven to enhance the results obtained, is the Synthetic Minority Oversampling Technique (SMOTE). This method operates by randomly selecting a point from the minority class and identifying its k-nearest neighbours. Synthetic points are generated between the selected point and its neighbours, effectively augmenting the minority class data, see Figure 15.

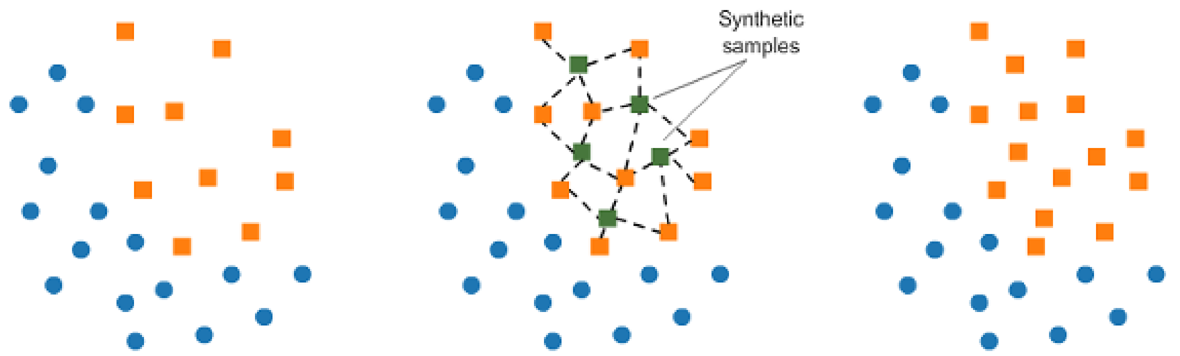

Figure 15 - SMOTE, Synthetic Minority Oversampling Technique, adapted from guest_blog, 2020.

The Adaptive Synthetic Sampling Approach (ADASYN) was also explored as it has been proven to enhance the results obtained. This method generates synthetic data points for the minority class, but unlike SMOTE, it does so in a more targeted manner. ADASYN focuses on harder-to-learn examples by adapting the decision boundary toward complex cases. It achieves this by calculating the density distribution of minority class samples and generating synthetic points inversely proportional to the density (He et al., 2008). This approach ensures that more synthetic data is created for minority samples that are harder to classify, thus improving the overall balance and robustness of the model.

ADASYN, SMOTE and oversampling techniques were applied using the 'imbalanced-learn' package (*Imbalanced-Learn*, n.d.).

3.4.2. Data Splitting

Splitting the dataset into training, validation, and testing is crucial for evaluating the performance of machine learning models (Muraina, 2022). In this research, the dataset was split into training and testing sets using the "train_test_split" function from the scikit-learn library (*Scikit-Learn: Machine Learning in Python — Scikit-Learn 1.4.1 Documentation*, n.d.).

The split ratio was 80% for training data and 20% for testing data. This ratio strikes a balance between having sufficient data for training the model and having a separate set for evaluating its performance. A more extensive training set allows the model to learn patterns from the data more effectively. In contrast, the testing set provides an unbiased estimate of the model's performance on unseen data. Choosing the stratifying version ensures that the distribution of the target variable (in this case, 'Dark Activity') is maintained in both the training and testing sets. This is particularly important when dealing with imbalanced datasets, where one class is significantly more prevalent than the other (Muraina, 2022).

3.4.3. Model Selection

In this chapter, the first goal is to explore a variety of machine learning models for classifying dark activity in unlabelled maritime data. Given the imbalanced nature of the dataset, all the models were tested with the original dataset, with random oversampling and SMOTE. The chosen models are selected based on their potential to handle class imbalance, flexibility, and ability to capture complex patterns in data.

A **Random Forest** is a classifier consisting of a collection of tree-structured classifiers, $\{h(x, \theta_k), k = 1, \ldots\}$ where $\{\theta_k\}$ are independent identically distributed random vectors, and each tree casts a unit vote for the most popular class at input x. Random forests are an effective tool in prediction. Injecting the proper randomness and aggregating multiple decision trees mitigates overfitting and makes them accurate classifiers and regressors (Breiman, 2001).

The random forest classifier was selected as an initial baseline model due to its robustness, simplicity, and ability to handle class imbalance. It provided reasonable results, as expected, through the literature review analysis (refer to page 8).

Similarly to a random forest classifier, Gradient Boosting is an ML ensemble technique that sequentially combines the predictions of multiple weak learners, typically decision trees (Natekin & Knoll, 2013). Gradient Boosting demonstrated its effectiveness in handling complex datasets within this context. Initially, the model struggled with the imbalanced nature of the data. However, after applying advanced techniques such as oversampling and synthetic sampling to mitigate the effects of data imbalance, the performance of Gradient Boosting significantly improved.

Based on the comprehensive literature review conducted in this study (refer to subchapter page 8), the **Support Vector Machine (SVM)** algorithm has historically shown promise in similar classification tasks. The fundamental goal of the SVM algorithm is to delineate a hyperplane that effectively separates data points belonging to different classes. These hyperplanes are positioned to maximise the margin between classes, ensuring optimal classification ('All You Need to Know About Support Vector Machines', n.d.). SVMs are renowned for their ability to tackle intricate mathematical problems. However, for data classification purposes, it is often preferable to employ smooth SVMs, which utilise smoothing techniques to mitigate the influence of outliers and enhance pattern recognition (Pisner, 2020). In this research, despite initial expectations, the performance of the SVM algorithm was tested and found to fall short of achieving satisfactory results.

Stacking, also known as stacked generalisation, is an ensemble learning technique that involves training multiple base classifiers (level-1 models) and then combining their predictions using a meta-classifier (level-2 model) (see Figure 16). The goal is to leverage the strengths of each model to improve overall predictive performance (Wolpert, 1992). This research employed two stacking configurations to classify dark activity instances. This initial approach to the stacking model utilised Random Forest and Gradient Boosting as base classifiers, with one model employing a Random Forest as the meta-classifier and another using Gradient Boosting as the meta-classifier. The results demonstrated that stacking effectively improves classification performance by integrating the strengths of multiple algorithms.

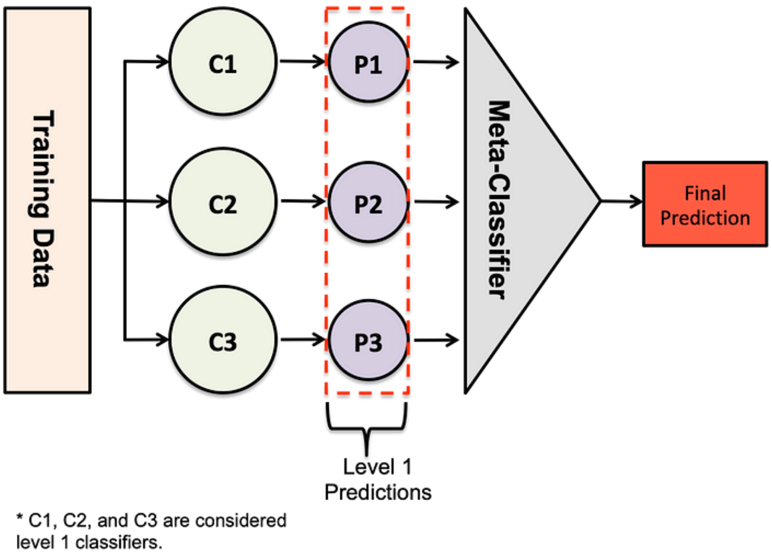

Figure 16 – Diagram of a Stacking Classifier Framework: This illustration depicts a stack comprising three classifiers, each trained independently. Their predictions are then combined and utilised to train a meta-classifier (Ceballos, 2019)

A second, more in-depth approach to the stacking classifier model was implemented, utilising a variety of base classifiers and a Logistic Regression meta-classifier. This method, referred to as "supermodels," involved stacking multiple models to enhance classification performance through a more comprehensive integration of algorithms. In one supermodel approach, a set of diverse base estimators was defined, including Random Forest (RF), Extra Trees (ET), XGBoost (XGB), K-Nearest Neighbors (KNN) and Logistic Regression (LR) with the meta classifier for this the LR. The second supermodel approach also employed the identical base learners but utilised the MLxtend library for stacking.

Both supermodels were designed with a single layer of stacking, meaning the meta-classifier was trained using the outputs of the base classifiers. No additional layers were added beyond this level. This approach combined the strengths of various algorithms, leveraging their capabilities to improve overall classification performance. The primary advantage of these supermodels was their ability to integrate predictions from different types of classifiers, enhancing robustness and accuracy. The use of diverse algorithms helped mitigate the weaknesses of any single model, resulting in a more reliable and effective classification system for detecting dark activities in maritime data.

The innovative **Deep Forest Cascade** structure emerged as a promising model, yielding commendable results in this study. Cascade Forest, renowned for its layered processing approach, adopts a hierarchical structure wherein each cascade level receives feature information processed by its preceding level (Zhou & Feng, 2019). This model leverages the strengths of ensemble learning by combining multiple types of classifiers in a deep forest architecture (Martins, 2023).

In the first cascade forest model, various classifiers were employed as base estimators. These base classifiers were set up on multiple levels within the Cascade Forest structure. For each level of the cascade, two random forests, two extra trees classifiers, two XGBoost classifiers, and two histogram-based gradient boosting classifiers were used.

The second cascade forest model employed Support Vector Machines (SVMs) as base estimators. The Cascade Forest classifier's unique feature lies in its ability to hierarchically process input data through multiple levels of diverse classifiers, effectively capturing complex patterns and improving classification performance. Each level in the cascade generates predictions concatenated with the original input features, creating a rich representation for subsequent levels to process.

For instance, as illustrated in Figure 17 if each level of the cascade comprises two random forests (depicted in black) and two completely random forests (depicted in blue) with three classes to predict, each forest will generate a 3D class vector. These vectors are then concatenated to re-represent the input data for subsequent processing.

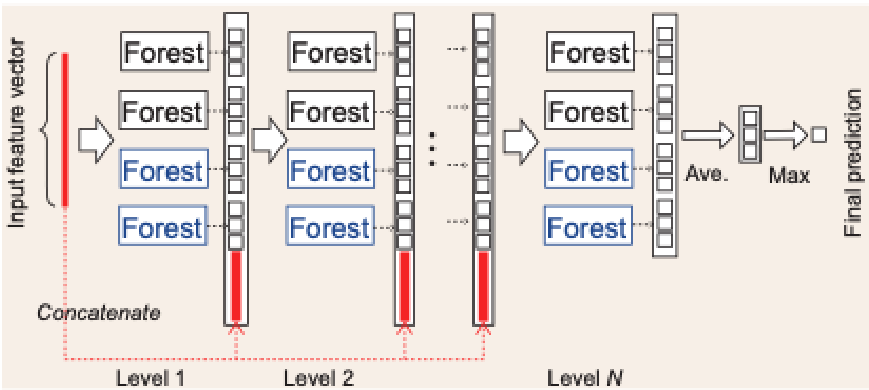

Figure 17 - Illustration of the Cascade Forest structure (Zhou & Feng, 2019)

With its capability to exploit ensemble learning, hierarchical structure, and adaptability to handle imbalanced data, the Deep Forest Cascade emerges as a strong contender for modelling purposes.

Additionally, Neural Networks were explored to assess their compatibility with solving this problem. **Convolutional Neural Networks (CNNs)** are investigated for their capacity to capture spatial and temporal patterns in the data. While they offer the potential for feature extraction and representation, their performance is hindered by the lack of diverse training samples for dark activity.

Multi-layer Perceptron (MLPs) are explored as a versatile neural network architecture capable of capturing non-linear relationships. However, their performance is limited by the imbalanced nature of the dataset and the risk of overfitting. Nonetheless, the results of both neural network methods were reasonably good, suggesting that with an increase in labelled data, there is potential for further improvement.

In addition to the supervised classification task aimed at identifying dark activity, this project encompasses an unsupervised approach to anomaly detection. The unsupervised task does not rely on labelled data and utilises **Self-Organizing Maps (SOM)** and **Density-Based Spatial Clustering of Applications with Noise (DBSCAN)** algorithms. SOMs are a type of artificial neural network trained using unsupervised learning to produce a low-dimensional, discretised representation of the input space of the training samples, called a map (*Self Organizing Map – Klaus Kayser*, n.d.). SOMs are particularly useful for clustering and visualising high-dimensional data in lower dimensions, making them suitable for identifying patterns and anomalies within complex datasets (Asan & Ercan, 2012).

DBSCAN is another unsupervised clustering algorithm commonly used for anomaly detection. It groups closely packed points while marking points in low-density regions as outliers. By defining clusters based on density rather than proximity to centroids, DBSCAN is robust to outliers and noise in the data (Cretulescu et al., 2019). By employing SOM and DBSCAN in the unsupervised part of the project, the objective is to identify anomalous patterns in the data and ascertain if the instances of dark activity identified through the supervised classification align with or are encompassed within the anomalies detected through unsupervised methods.

3.4.4. Hyperparameter Tuning

Hyperparameter tuning plays a crucial role in optimising the performance of machine learning models by fine-tuning the parameters that govern their behaviour. In this research project, **Grid Search** is the primary method for hyperparameter tuning across all models. Grid search systematically explores a predefined set of hyperparameters and their respective values, evaluating each combination using cross-validation to identify the optimal configuration (*Harnessing the Power of Grid Search for Optimized Machine Learning Models | by Nilimesh Halder, PhD | Analyst's Corner | Medium*, n.d.). Considering the computational resources available, this method is preferred for its simplicity, exhaustiveness, and ability to efficiently search through the hyperparameter space. By testing a range of hyperparameter values, grid search identifies the settings that yield the best model performance regarding accuracy, precision, and recall. Additionally, grid search provides transparency and reproducibility in the hyperparameter tuning process, allowing researchers to compare different configurations and select the most suitable ones systematically (Belete & D H, 2021).

3.4.5. Model Evaluation

Model evaluation is fundamental in assessing the performance of machine learning models and determining their effectiveness in solving the problem at hand. Given the classification nature of the problem addressed in this research, several evaluation metrics were considered to measure the performance of the models, with a particular emphasis on metrics suited for imbalanced datasets.

Although **accuracy** is a fundamental metric to evaluate the overall correctness of the model's predictions, in the case of imbalanced datasets, accuracy alone might not accurately depict the model's performance, as the majority class could skew it (Hancock et al., 2023).

Precision quantifies the proportion of correctly predicted positive instances among all instances predicted as positive. It is particularly relevant in scenarios with a high cost of false positives. On the other hand, **recall**, also known as sensitivity, measures the proportion of true positive instances correctly identified by the model (Hancock et al., 2023). In dark activity detection, high recall is crucial to ensure that the model captures a significant portion of actual dark activity cases. Furthermore, the **F1-score**, the harmonic mean of precision and recall, provides a balanced assessment of the model's performance, considering both false positives and false negatives. This metric is handy in scenarios where achieving a balance between precision and recall is essential.

Given the imbalanced nature of the dataset, particular emphasis was placed on these metrics for the minority class. This focus ensures that the evaluation accurately reflects the model's ability to detect dark activity, which constitutes the minority class in this study.

The **area under the ROC curve (AUC-ROC)** was also considered in the evaluation. The AUC-ROC measures the model's ability to distinguish between classes across different threshold settings, comprehensively evaluating the classifier's performance (Chicco & Jurman, 2020). A higher AUC indicates better performance in distinguishing between positive and negative classes, which is crucial for assessing the model's effectiveness in detecting dark activity.

For the unsupervised models, the evaluation focused on the percentage of dark activity cases that fell within the detected anomalies, the **Silhouette Score**, and the **R-squared**. The silhouette score measures the cohesion and separation of the clusters formed by the unsupervised model, indicating how well-separated the clusters are (Shahapure & Nicholas, 2020). A higher silhouette score suggests better-defined clusters, essential for identifying distinct anomalous behaviour patterns.

R-squared, or the coefficient of determination was used to evaluate the proportion of variance in the data captured by the model. In anomaly detection, a higher R-squared value indicates that the model effectively explains the variability in the dataset, providing insights into the underlying structure of the data (Frost, 2017).

By focusing on these metrics, this research provides a more accurate reflection of the model's ability to detect dark activity and anomalous situations, ensuring that both supervised and unsupervised approaches are thoroughly evaluated for their effectiveness in real-world applications.

3.5. DEPLOYMENT

For the deployment phase of the CRISP-DM implementation, the goal is to transition the developed machine learning models from the research environment to practical applications for real-time detection of dark activity in the maritime industry. To achieve this, leveraging modern tools and frameworks for model management, pipeline orchestration, and real-time data ingestion is essential (Chapman et al., 2000).

This research used FastAPI, a modern, fast (high-performance) web framework for building APIs with Python 3.6+ based on standard Python-type hints, for deployment. FastAPI allows for the seamless integration of machine learning models into a web service, providing an easy-to-use interface for real-time predictions and interactions with the models (*FastAPI*, n.d.). By leveraging FastAPI, the machine learning models developed in the research phase can be deployed as web applications, enabling users to interact with them through RESTful API endpoints (interfaces such as web or mobile applications, through which users interact with a web service using HTTP requests).

This approach ensures a seamless transition from research to deployment, allowing the developed machine learning models to be effectively utilised for real-time detection of dark activity in the maritime industry. The utilisation of FastAPI provides high performance and flexibility. It ensures scalability and ease of integration with other systems, making it ideal for deploying machine-learning models in practical applications (Mund, 2021).

4. RESULTS AND DISCUSSION

At the heart of this research lies the critical question: How can AIS data be effectively harnessed to detect and identify anomalous behaviours in the maritime industry, specifically focusing on periods of dark activity? This question points to the larger goal of enhancing maritime security by developing a robust model using ML algorithms, leveraging AIS data for predictive analysis.

The significance of this endeavour extends beyond academic pursuits; it is deeply rooted in addressing a pressing need within the Portuguese Navy and AMN for an objective, data-driven approach to maritime surveillance. Consequently, this research encapsulates the formulation of a sophisticated decision-making tool that promises to transform the current subjective methodologies into an unbiased, analytical process.

This chapter unfolds the outcomes of data analysis and modelling efforts, comprehensively examining the study's results. In doing so, it offers insights into the efficacy of various ML models in identifying maritime anomalies, their operational implications, and the overarching impact on maritime security and operations.

4.1. PRELIMINARY DISCUSSION

In the context of maritime security and the detection of dark activity, the implications of false positives and false negatives hold significant importance, though currently, the impact of both is relatively low. Presently, there is no dedicated decision-making tool for inspections or investigations of ships, and such processes rely heavily on shared knowledge and experience. Consequently, any technological aid, even if it produces some false positives and false negatives, is considered beneficial.

A false positive in this scenario means the system flags a compliant ship as engaging in dark activity. The primary negative impact of a false positive is the unnecessary allocation of resources. This could lead to unwarranted investigations or inspections, which, although maybe resource-intensive, ultimately reveal the ship to comply with regulations. While this does represent a cost in terms of time and effort, it does not compromise maritime security.

On the other hand, a false negative occurs when the system fails to detect a ship that is engaging in dark activity. This means a ship potentially violating regulations or conducting illegal activities goes unnoticed. While this could theoretically allow non-compliant or illegal activities to proceed undetected, it is essential to note that inspections and investigations are already somewhat random and opportunistic in the absence of a decision-support system. Therefore, introducing this technology, even with its false negatives, does not worsen the current situation but instead provides an additional layer of scrutiny that can enhance overall detection capabilities.

4.2. SUPERVISED MODEL RESULTS AND DISCUSSION

To address one of the primary objectives of this research, a suite of 22 distinct model configurations (see Table 5) was applied to the dataset to undertake the classification task of identifying dark activity instances. The underlying strategy for model selection was to commence with simpler models and incrementally progress toward more complex ones. This progressive approach allowed for a thorough evaluation of each model's performance, facilitating a comprehensive understanding of the intricacies of accurately classifying dark activity cases.

Given the imbalanced nature of the dataset, as previously detailed in this chapter and illustrated in **Error! Reference source not found.**, four distinct resampling methods from the ***imblearn.over_sampling*** package addressed the imbalance (*Imbalanced-Learn Documentation — Version 0.12.2*, n.d.). Careful consideration was given to the proportion of under or oversampling to prevent the model from developing a bias that might overemphasise or underrepresent the minority class. The ***sampling_strategy*** parameter was meticulously calibrated, with the optimal value set to 0.15. This adjustment allowed for the minority class to represent approximately 12% of the dataset, an increase from the initial 4.1%, thereby helping to mitigate the issues stemming from the initial class disproportion. The techniques considered were:

- ***Random Oversampling:*** Using RandomOverSampler from imblearn.over_sampling involves randomly duplicating examples in the minority class to achieve a balanced dataset (Ganganwar, 2012).
- ***Random Undersampling:*** Using RandomUnderSampler from imblearn.over_sampling, this technique reduces the number of examples in the majority class by randomly removing them until the dataset is balanced (Yap et al., 2014).
- ***SMOTE (Synthetic Minority Over-Sampling Technique):*** Using SMOTE from imblearn.over_sampling, this approach creates synthetic examples of the minority class by interpolating between existing examples (Chawla et al., 2002).
- ***ADASYN (Adaptative Synthetic Sampling Approach for Imbalanced Learning)***: Using ADASYN from imblearn.over_sampling similar to SMOTE, this method generates synthetic data points but focuses on generating more samples near the decision boundary where the classifier is most uncertain (He et al., 2008).

First, let us analyse the Random Forest and Gradient Boosting model variations. The literature review shows that Random Forest algorithms are feasible for maritime machine learning applications, typically yielding good results (Huang et al., 2024).

As shown in Table 4, various data sources were applied to determine the most effective sampling method. By examining the performance across different sampling methods in Table 5, it can be concluded that the methods that most improved the results were *Random Oversampling* and *SMOTE*. Given that the results were very similar, and *Random Oversampling* with the defined sampling strategy is a simpler and easier method to implement, it will be the chosen method for subsequent steps and models. Regarding the model parameters, it was observed that the *GridSearch* applied to the simple Random Forest model did not improve the results, as shown in Table 5, when comparing Model 1 and Model 2. However, when sampling strategies were applied, the model configured with the optimal parameters from *GridSearch* demonstrated superior performance. Consequently, the Random Forest model will be implemented using the best parameters identified through *GridSearch* for further analyses.

The literature also recognises gradient boosting as effective for classification problems and predictions within the maritime environment (Mondal et al., 2023). This research achieved commendable results, as Models 7 and 8 in Table 5. Furthermore, using *GridSearch* to identify the best parameters for the Gradient Boosting model significantly enhanced its performance. This improvement is evident in the same table, where the results of Models 7 and 8 are compared.

In this research, the stacking method combined Random Forest and Gradient Boosting models in two distinct configurations to enhance predictive performance. Stacking is an ensemble learning technique that combines multiple classification or regression models via a meta-classifier or meta-regressor (Wolpert, 1992). In this case, the base models were a Random Forest and a gradient-boosting model, with the final predictions made by a meta-classifier. For Model 9, the final estimator was a *RandomForestClassifier*, and for Model 10, it was a *GradientBoostingClassifier*, both configured with a random state of 42, a technique also used by (Najafzadeh et al., 2023). Both models were trained on the resampled training set and evaluated on the test set. While the results indicated an improvement over individual base models, they did not surpass the performance of the Gradient Boosting model fine-tuned with optimal parameters as identified through *GridSearch*.

Table 4 - Random Forest and Gradient Boosting models settings.

	Models		**Data**	**Parameters**
1	Random Forest		Original	Default
2	Random Forest		Original	Grid Search
3	Random Forest		Over Sampling	Grid Search
4	Random Forest		Under Sampling	Grid Search
5	Random Forest		ADASYN	Grid Search
6	Random Forest		SMOTE	Grid Search
7	Gradient Boosting		Over Sampling	Default
8	Gradient Boosting		Over Sampling	Grid Search
9	**Estimator**	**Meta Estimator**	Over Sampling	Grid Search
	Random Forest	Random Forest		
	Gradient Boosting			
10	**Estimator**	**Meta Estimator**	Over Sampling	Grid Search
	Random Forest	Gradient Boosting		
	Gradient Boosting			

In the second stage of modelling, more complex combinations were considered. The initial model in this phase was Model 11 (see Table 5), named Deep Forest Stack, which employed an ensemble learning strategy using a Cascade Forest Classifier (C*ascadeForestClassifier*). This approach was designed to harness the collective strengths of multiple well-regarded machine-learning algorithms to enhance prediction accuracy (Ali et al., 2024). The ensemble included two instances each of RandomForest (*RandomForestClassifier*, *random states 186 and 196*), ExtraTrees (*ExtraTreesClassifier, random states 286 and 296*), XGBoost (*XGBClassifier, random states 386 and 396*), and Histogram-based Gradient Boosting (*HistGradientBoostingClassifier, random states 486 and 496*). These models were integrated into the Cascade Forest Classifier and configured to perform internal cross-validation with two splits to ensure robust validation of the model's performance (Martins, 2023). The goal was to exploit the diverse learning mechanisms of the individual models to improve the system's generalisation ability on the scaled training data. This strategic combination into a cascading framework aimed to leverage the individual strengths of the models and mitigate their weaknesses, thereby enhancing predictive performance on classifying maritime dark activity.

In a subsequent modelling phase, the cascade ensemble strategy was extended to include Support Vector Machines (SVMs). Specifically, two instances of the SVM classifier with enabled probability estimation were deployed: svm1 with a *random_state* of 186 and svm2 with a *random_state* of 196. These classifiers were integrated into a Cascade Forest Classifier (*CascadeForestClassifier*), configured with a *random_state* of 86. The ensemble was set up with two splits for internal cross-validation, ensuring a thorough evaluation of the model's performance. This setup aimed to harness the robust classification capabilities of SVMs, known for their effectiveness in high-dimensional spaces, as already mentioned in the literature review (page 8). Although SVMs are known for their strong capabilities, in this instance, they did not perform as expected, achieving poor results, as detailed in Table 5.

For Model 13, termed 'Super Learner 1', a sophisticated stacking approach was utilised to integrate a diverse array of base estimators to enhance predictive performance. The base estimators included a *RandomForestClassifier*, *ExtraTreesClassifier*, *XGBClassifier* with a binary logistic objective, a *K-Nearest Neighbors Classifier* embedded within a pipeline featuring standard scaling, and a *Logistic Regression* also scaled and configured with a *liblinear* solver for optimisation. These models were strategically chosen to bring together a variety of learning algorithms, each with its strengths in handling different aspects of the data. The ensemble was orchestrated by a Logistic Regression meta-estimator, which consolidated the base estimators' predictions into a final decision.

In the development of Model 14, termed 'Super Model 2', the identical base learners used in 'Super Learner 1' were employed. However, a key differentiation in 'Super Model 2' lies in the configuration of the stacking mechanism, utilising the *MLxtend* library instead of *StackingClassifier*. This model was uniquely set to use the probability outputs from the base learners for the meta-model rather than just their predictions. This approach allows the meta-learner, a Logistic Regression, to make more informed decisions based on the nuanced probability landscape provided by the base models. Unlike 'Super Learner 1', 'Super Model 2' does not average these probability estimates but instead utilises the full spectrum of probabilistic data from each classifier. This method enhances the meta-learner's ability to discern more complex patterns indicative of maritime anomalies, potentially improving model robustness.

Deep learning methods were employed in the third modelling stage but did not achieve the expected results. The initial deep learning approach involved a *simple Multilayer Perceptron*, designated as Model 15 in Table 5. This model features a sequential architecture consisting of three layers: the first layer comprises *64 neurons* with *ReLU activation*, a second layer of *32 neurons* with *ReLU activation*, and a final output layer equipped with a *single neuron* using *sigmoid activation*. This configuration is tailored to process features from the scaled training data. Model 15 was compiled using the Adam optimiser, with binary cross-entropy as the loss function and accuracy as the evaluation metric. The model underwent training for 50 epochs with a batch size of 32, incorporating a validation split of 20% to monitor performance throughout the training phase. The final evaluation of the scaled test dataset involved translating probability outputs into binary classifications, using a 0.5 threshold to determine the model’s proficiency in identifying maritime dark activities.

Model 16, NN2, introduces an advanced approach to neural network configuration by employing hyperparameter tuning to optimise its architecture. Unlike the previous model, NN2 utilises a *RandomSearch* tuner to systematically explore a range of potential configurations. Unfortunately, despite the sophisticated hyperparameter tuning employed in Model 16, the anticipated improvement in performance was not realised. This outcome suggests that the initial tuning parameters and model configuration may not be optimally aligned with the specific characteristics of the maritime anomaly detection task. Further experiments with different tuning options and strategies will be conducted in response.

In pursuit of improved model performance, *Bayesian Optimization* was implemented for hyperparameter tuning in the development of Model 17 (NN3). This new tuning approach resulted in enhanced performance relative to the previous model. Despite these gains, the performance of NN3 still did not surpass that of the simpler *Multilayer Perceptron model*.

Continuing the exploration of hyperparameter tuning strategies, the *Hyperband* method was employed for Model 18 (NN4). Despite these efforts, the performance of NN4, as evaluated on the scaled test set, was not satisfactory, yielding results inferior to those of previous models.

Model 19, NN5, implemented a different architecture and configuration from previous models, yet it yielded the least effective results. This model was designed to prevent overfitting with a *deeper network* comprising *four dense layers* with *higher neuron counts* and *uniform dropout rates of 0.3*. Despite the more extensive and complex architecture aimed at capturing intricate patterns in the data, NN5 could not outperform simpler models, suggesting that increasing model complexity does not necessarily translate to better performance.

Models 20 (NN6) and 21 (NN7) introduced variations in activation functions and network architecture to improve the performance of previous neural network models. However, these modifications did not result in enhanced predictive outcomes.

NN6 utilised *LeakyReLU* activation functions with a small *alpha value of 0.05* to address the issue of "dying neurons" seen with regular *ReLU activations* in deeper networks (Desiani et al., 2023). It also featured *multiple dropout layers* to reduce overfitting.

NN7 further increased model complexity by integrating *Batch Normalization* with *wider layers*, each having 256 units initially, then narrowing down to 128 and 64 units. This model aimed to stabilise learning by normalising the input layer by adjusting and scaling activations. It also employed an *RMSprop* optimiser instead of Adam, known for efficiently handling noisy gradients (Müller, 2023).

Model 22, NN8, is the most effective deep learning model tested in this research due to its advanced architecture and training techniques. It features multiple dense layers with high neuron counts and integrates *LeakyReLU* activations and Batch Normalization to improve stability and efficiency. Distinctively, NN8 incorporates *Dropout layers* to prevent overfitting and uses a *Learning Rate Scheduler* that reduces the learning rate exponentially after the initial ten epochs. This approach enhances fine-tuning during training, aiming to achieve better generalisation on unseen data. Although NN8 surpassed other deep learning models in performance, it did not outperform previous non-deep learning models.

The Loss and accuracy curves for Deep Learning methods are provided in Annexe 10 through Annexe 17.

Table 5 provides a comprehensive overview of the evaluation metrics for all the models previously discussed:

Table 5 - Supervised models evaluation table.

	Class 0		Class 1		Accuracy	Area under ROC
	F1 - Score	Recall	F1 - Score	Recall		
1. RF	0.99	1.00	0.66	0.49	0.98	0.75
2. RF (best parameters)	0.99	1.00	0.57	0.40	0.98	0.70
3. RF BP Oversampling	0.99	1.00	0.69	0.53	0.98	0.77
4. RF BP Undersampling	0.99	1.00	0.64	0.48	0.98	0.74
5. RF BP ADASYN	0.98	0.98	0.58	0.63	0.96	0.80
6. RF BP SMOTE	0.99	1.00	0.69	0.53	0.98	0.77
7. GB Over Sampling	0.98	1.00	0.42	0.27	0.97	0.64
8. GB BP Over Sampling	1.00	1.00	0.87	0.80	0.99	0.90
9. Stacking 1 RF GB + RF	0.99	1.00	0.85	0.77	0.99	0.88
10. Stacking 2 RF GB + GB	0.99	1.00	0.86	0.77	0.99	0.89
11. Deep Forest Stack	0.99	1.00	0.78	0.66	0.99	0.83
12. Deep Forest SVM	0.98	1.00	0.10	0.05	0.96	0.44
13. Super Learner 1	1.00	1.00	**0.92**	**0.87**	0.99	**0.94**
14. Super Learner 2	0.99	1.00	0.78	0.65	0.99	0.82
15. NN 1 - MLP	0.98	1.00	0.31	0.19	0.97	0.59
16. NN 2 BP random S	0.98	1.00	0.24	0.14	0.96	0.57
17. NN 3 BP Bayesian	0.98	1.00	0.26	0.16	0.96	0.58
18. NN 4 BP Hyperband	0.98	1.00	0.21	0.12	0.96	0.56
19. NN 5	0.98	1.00	0.15	0.08	0.96	0.54
20. NN 6 LeakyRelu	0.98	1.00	0.22	0.12	0.96	0.56
21. NN 7 Batch Norm RMSprop	0.98	1.00	0.15	0.08	0.96	0.54
22. NN 8 LR scheduler	0.98	1.00	0.33	0.20	0.97	0.60

4.3. BEST MODELS DISCUSSION

Regarding Table 5, it is crucial to focus on specific, particularly relevant evaluation metrics due to the highly imbalanced dataset used in this study. The metrics of primary importance are recall, F1-score, and the area under the ROC curve (AUC), each of which provides critical insights into the models' performance:

- **Recall (Sensitivity)**: This metric is essential because it measures the model's ability to correctly identify positive instances (in this case, actual maritime anomalies) among all actual positives. The high recall is vital in this scenario because a missing true positive could lead to safety concerns or risk.
- **F1-score**: Since the dataset is imbalanced, precision and recall are crucial. The F1 score harmonically combines these two metrics. In this case, it is advantageous to analyse the F1-score to identify the class 1 – dark activity.
- **Area Under the ROC Curve (AUC)**: The AUC provides an aggregate performance measure across all possible classification thresholds. This metric is beneficial for evaluating the performance in this case because it is threshold-independent and reflects the model's ability to discriminate between the classes at various threshold settings.

In this study, as observed in Table 5, all models demonstrated high performance in terms of accuracy, F1-score, and recall for class 0 (non-dark activity), which is anticipated given the unbalanced nature of the dataset where predicting the majority class is relatively more straightforward due to its larger volume. Consequently, this research will primarily focus on the F1-score and recall for class 1 (dark activity) and the area under the ROC curve to provide a more comprehensive evaluation of model effectiveness in identifying the less frequent but critical instances of dark activity.

In the initial phase of modelling, it is evident from Figure 18 that models 8, 9, and 10 are the most effective at predicting dark activity, achieving F1-scores for class 1 between 0.85 and 0.87, as detailed in Table 5. The recall for class 1 and the area under the ROC curve is also considered suitable for these models. However, the standout performer at this stage is Model 8, which employs tuned Gradient Boosting. The superior performance of Model 8 may be attributed to the precision of its hyperparameter optimisation, which effectively enhances its ability to manage the complexities and nuances of the imbalanced dataset, thus improving its predictive accuracy for the less-represented class of dark activity.

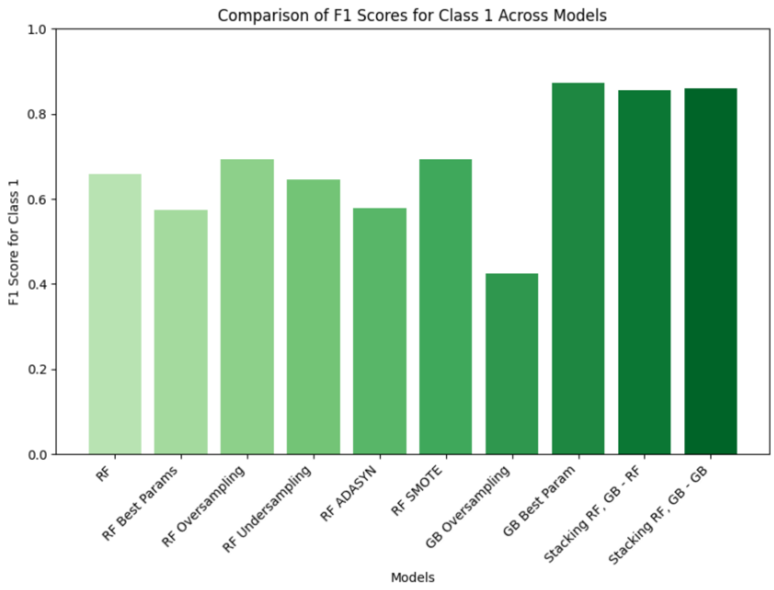

Figure 18 – F1 score for class 1 for 1st stage of modelling.

In the second modelling stage, the most favourable outcomes for this research were achieved, as evidenced by the superior performance of all models. As depicted in Figure 19, Super Learner 1 and Super Learner 2 emerged as the top-performing models, demonstrating an exceptional ability to discriminate between dark and non-dark activity classes (class 1 and class 0). Their ROC curves closely approach the left-hand and top borders of the ROC space, indicative of a high true positive rate and a minimal false positive rate, signifying their robustness in correctly identifying instances of dark activity.

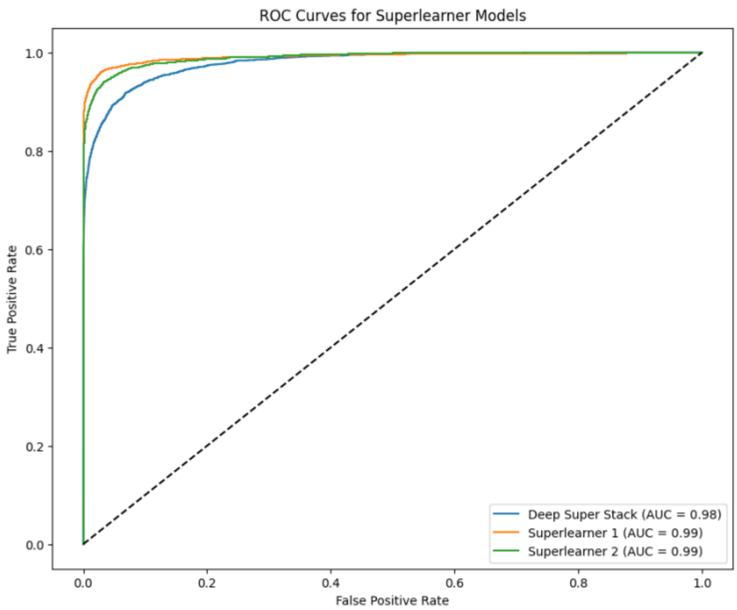

Figure 19 – ROC curves for Super Learner Models.

For a more comprehensive analysis, Figure 20, reveals that Superlearner 1 stands out with the highest recall for class 1 and a strong AUC, marking it as highly effective for identifying dark activity without many false negatives. Superlearner 2, while having the highest precision, slightly lags in recall and F1 for class 0, which may indicate a trade-off that prioritises correctly identifying dark activities at the expense of misclassifying some non-dark activities. The Deep Super Stack presents a balanced model with solid metrics across the board but does not excel to the same degree as the Superlearner models in class 1 metrics, which are critical for detecting dark activity in this research.

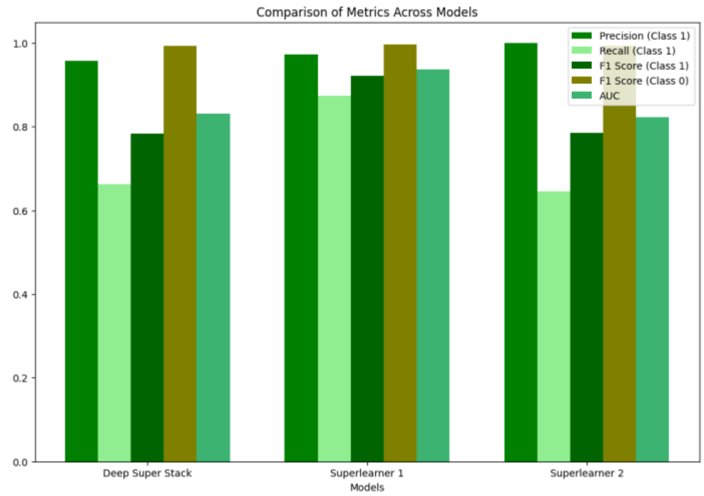

Figure 20 – Comparison of metrics across Super Learner Models.

Table 6 presents the confusion matrix for the top-performing models. Each model demonstrates consistent true positives, indicating a reliable detection rate for actual dark activities. Super Learner 1 has the highest number of true positives and the fewest false negatives, signifying its superior capability to identify dark activities and reduce the incidence of overlooked detections. Additionally, Super Learner 1 records many true negatives, underscoring its robustness in correctly recognising instances of non-dark activities.

Table 6 - Confusion matrix for best models.

	True Positives	False Negatives
GB BP Over Sampling	63 294	75
Stacking 1 RF GB + RF	63 290	79
Stacking 2 RF GB + GB	63 296	73
Deep Forest Stack	63 360	81
Super Learner 1	63 376	65
	False Positives	True Negatives
GB BP Over Sampling	546	2 140
Stacking 1 RF GB + RF	622	2 064
Stacking 2 RF GB + GB	608	2 078
Deep Forest Stack	907	1 779
Super Learner 1	340	2 346

4.4. UNSUPERVISED LEARNING

In this research, including an unsupervised approach to general anomaly detection was deemed essential. This approach enhances the ability to identify anomalies and helps clarify the extent to which "Dark Activity" can be classified as a specific instance of anomalous behaviour.

The first unsupervised method implemented was Isolation Forest, which was chosen for its simplicity and effectiveness in detecting anomalies. The isolation forest algorithm operates on the principle that anomalies are "few and different". It isolates observations through random selections of features and split values, making it particularly effective for identifying outliers in high-dimensional datasets (Liu et al., 2008). Initially, the model was configured with 100 estimators and a contamination rate of 0.15, corresponding to the expected proportion of anomalies. This configuration identified 49,585 outliers, with 1,509 (11.5%) correctly matching the confirmed dark activity instances.

To improve this, parameters such as the number of estimators, maximum samples, contamination rate, and maximum features were systematically varied to identify the best configuration – hyperparameter tuning using stability score. The optimal settings, determined through this tuning process, included 50 estimators, a maximum of 100 samples, a contamination rate of 0.2, and 50% of the available features. With these parameters, the refined Isolation Forest model identified 66,113 outliers, of which 2,418 were actual dark activities, increasing the detection proportion to 18.00%. Two metrics were used to assess the quality and reliability of the clustering and anomaly detection. One is the R-squared value with a result close to one (0.98), suggesting that the model has effectively captured the data variability, highlighting its potential in identifying patterns and anomalies. However, the silhouette score (second metric), which measures the separation distance between the resulting clusters, was 0.101. This relatively low value suggests the clustering structure is weak, with clusters overlapping significantly. Additionally, the low silhouette score and the high R-squared might indicate that while the model explains the variance in individual features well, it struggles to form well-defined clusters, potentially impacting anomaly detection accuracy.

The second unsupervised model applied in this research was DBSCAN, a method effective for identifying clusters of varying shapes and densities and distinguishing outliers (Cretulescu et al., 2019). To determine the optimal EPS parameter for DBSCAN, the Nearest Neighbours algorithm was employed, plotting the distances to the 500th nearest neighbour to identify the "elbow" point, which suggested an EPS value of 1.25 (see Annexe 18). This parameter selection process ensured that DBSCAN could effectively cluster the data and identify noise points. The implementation revealed eight clusters and identified 31,524 noise points, indicating dense and sparse data distribution areas (see Figure 21).

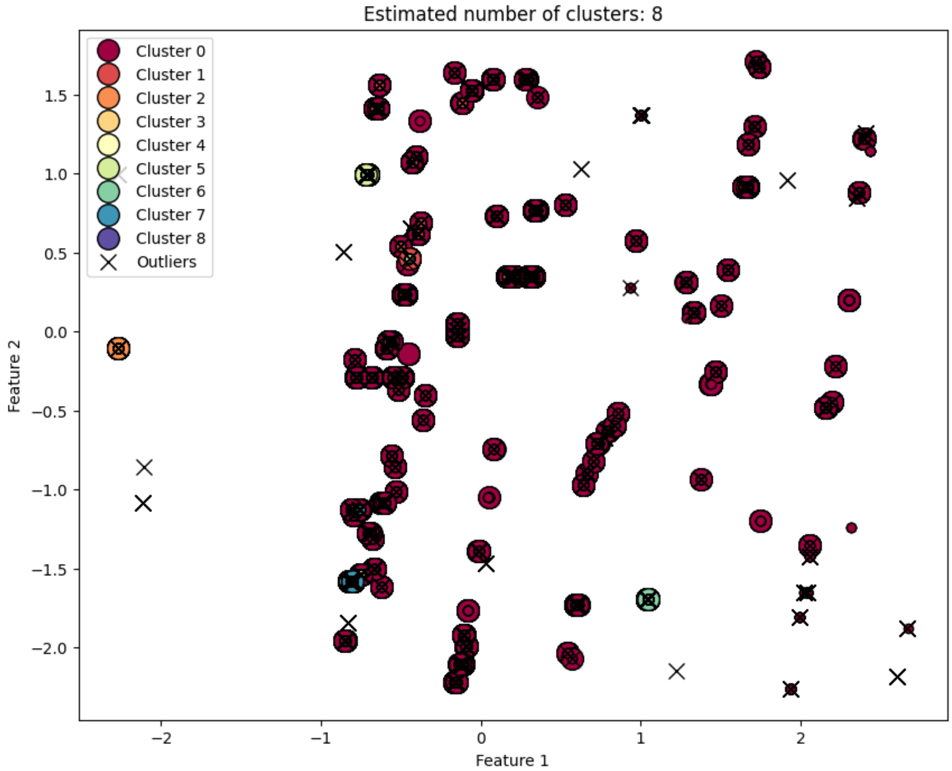

Figure 21 – Clustering of AIS data using DBSCAN.

Two approaches were considered in the identification of the anomalies. The first approach classified only the noise points identified by DBSCAN as anomalies, resulting in 31 524 anomalies detected (9% of the whole dataset), yielding a proportion of 11.74% of the dark activity cases in those anomalies. Based on the observation that Cluster 0 contained most data points, the second approach considered all points not in Cluster 0 as outliers. This revised approach increased the detection to 44 460 anomaly instances (13% of the data), achieving a proportion of 13% of the dark activity in those anomalies. Despite the improvement in the dark activity proportion, the silhouette score of 0.11 indicated a weak clustering structure, and the R-squared value of 0.26 suggested a limited variance explanation by the model. These findings highlight the complexities of using unsupervised methods for anomaly detection in maritime data.

The third and final unsupervised algorithm applied in this research was the Self-Organizing Map (SOM), which ended up being the most extensively explored and yielded the best results among the unsupervised learning methods. The initial implementation utilised the MiniSom library (Vettigli, 2013/2024), where the SOM was trained on the dataset to identify anomalies. Various thresholds for anomaly detection were evaluated, including the 85th percentile of quantisation errors, the mean plus one standard deviation of the quantisation error, and the mean plus half a standard deviation, with values above these thresholds considered anomalous. The 85th percentile threshold demonstrated the highest efficacy, identifying 49,585 anomalies out of 330,565 entries, representing 15.00% of the total data. Within the 'Dark Activity' instances, 1,828 were classified as anomalies, accounting for 18.8%. The evaluation metrics indicated a Silhouette Score of 0.45 and an R-squared value of 0.99, suggesting a reasonably good clustering performance. These metrics indicate that while most samples were accurately clustered, there was some overlap between clusters, and the model explained a high proportion of the variance in the data.

A parameter tuning process was conducted using a grid search approach to further enhance the SOM's performance. The parameters tuned included the dimensions of the SOM grid (x and y), the sigma (which controls the radius of the neighbourhood function), and the learning rate. Parameter tuning involves evaluating different combinations of these parameters to identify the best configuration. This configuration identified 59,502 anomalies, representing 18% of total data, and 19.6% of dark activity instances are included in those anomalies identified. This yielded an improved Silhouette Score of 0.64, indicating better-defined clusters, and maintained a high R-squared value of 0.999, reflecting a strong explanatory power of the variance in the data.

These two approaches are represented in Figure 22, where the light regions represent low distances between neurons, indicating high similarity among data points mapped to these neurons, and the dark regions represent the opposite, higher distances between neurons suggesting transitions between clusters. The highlighted anomalies (red circles) are data points with significantly higher quantisation errors, meaning they do not fit well into the learned clusters. They are located near the cluster boundaries and in regions with higher U-Matrix values (darker areas) in both cases, as anomalies are expected to be less similar to their neighbouring data points.

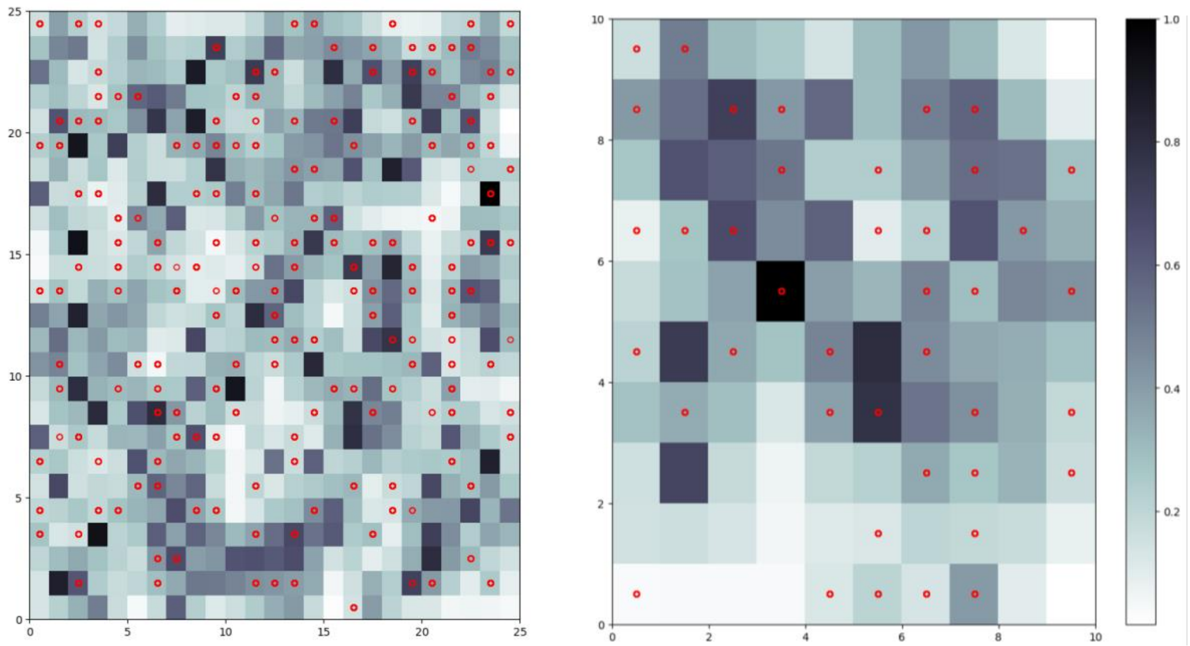

Figure 22 – Clustering of AIS data and respective identified anomalies using SOM. Without tunning (left) and with tunning (right), respectively.

In the first image of Figure 22, representing the results without parameter tuning, the clusters appear less distinct, with many areas showing overlapping regions of data points and a higher degree of dispersion. The anomalies are scattered more evenly across the map, indicating a broader spread and less clear differentiation between normal and anomalous data points. This dispersion suggests that the model's initial parameters may not have been optimal for clustering or detecting anomalies effectively.

In contrast, the second image, which depicts the results after parameter tuning, shows more clearly defined clusters with sharper boundaries between high and low U-Matrix values. The anomalies are more concentrated near the boundaries of clusters, particularly in regions with higher U-Matrix values (darker areas). This concentration suggests that the tuned parameters have improved the model's ability to identify anomalies more accurately by better defining the boundaries of standard behaviour clusters. The improved clustering and concentrated anomalies indicate that the parameter tuning has enhanced the model's effectiveness in distinguishing between normal and anomalous behaviours, providing a more precise identification of anomalous regions. Overall, the MiniSom approach with parameter tuning yielded the best results among the unsupervised methods in this research.

Another package, SOMPY (sevamoo, 2014/2024), was employed to refine the SOM approach further, allowing more flexibility and better parameter tuning. Using SOMPY, two different approaches for anomaly detection were tested. The first approach used the same 85th percentile of quantisation errors as the threshold. In contrast, the second approach used the dynamic threshold based on the mean plus one standard deviation of quantisation errors. The dynamic threshold approach yielded better results, detecting 51,422 anomalies out of 330 565 entries (15.5%). Among these, 3,060 'Dark Activity' instances were identified, representing 22.8% of the total confirmed instances of dark activities. The component planes (Picture in Annexe 19) revealed distinct patterns for different features, underscoring the effectiveness of the SOM in clustering the AIS data.

Another approach involved visualising the data using a hit-map visualisation (Figure 23). In this map, each hexagon represents a neuron, and the number within the hexagon indicates how many data points are mapped to that neuron. Areas with fewer hits can indicate outliers or rare patterns in the data (Concetti et al., 2023). Based on the rationale that these regions represent less common, potentially anomalous behaviours, neurons with less than 1,000 hits were considered anomalous. This method identified 68,722 anomalies, approximately 20.80% of the data, with 3,037 'Dark Activity' instances included, representing 22.61% of the identified dark activity cases.

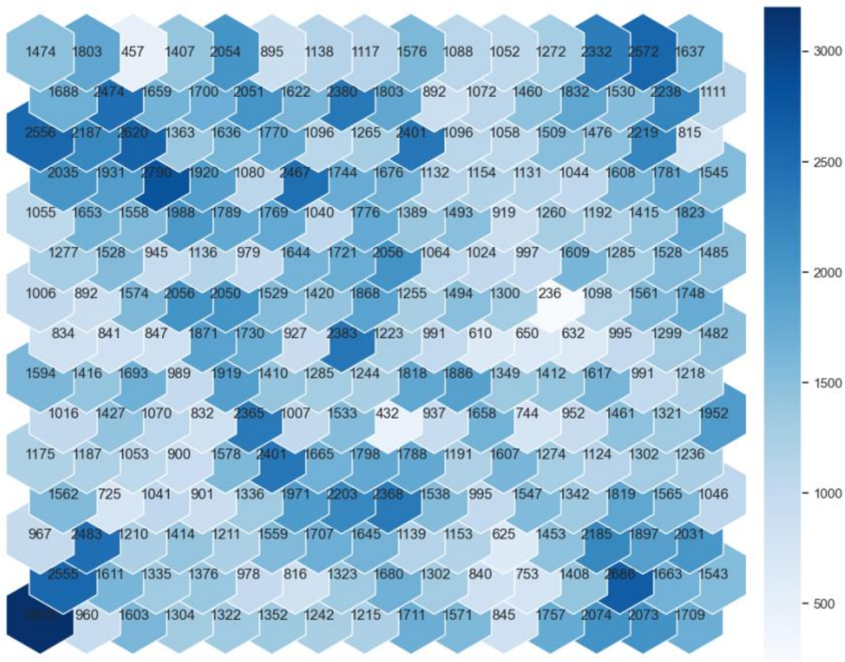

Figure 23 – Clustering of AIS data using data point density Map: Neuron Hit Frequency Representation.

Lastly, K-means clustering was applied on top of the SOM with 4 clusters considered (see picture in Annexe 20). After an analysis of the picture, clusters 2 and 3, as the minority, were considered as outliers. This approach identified 45,769 anomalous entries, with only 903 'Dark Activity' instances detected, making up only 6.72% of the identified cases. A similar hierarchical clustering approach yielded comparable results (see Annexe 21). The SOMPY model's silhouette Score of 0.201 and R-squared of 0.78 indicate moderate clustering performance and challenges in clearly differentiating between clusters.

These findings suggest that while SOM and its variations can effectively identify anomalies, the choice of threshold and additional clustering methods significantly impact the results.

A summary of the results of unsupervised learning methods is presented in Table 7.

Table 7 - Unsupervised algorithms evaluation table.

	Anomalies	% Anomalies	% Dark Activity	Silhouette	R-Squared
Isolation Forest	49 585	15%	11.5%	0.09	0.98
Isolation Forest w/ Tunning	66 113	18%	18%	0.11	0.98
DBSCAN (noise)	31 524	9%	11.74%	0.11	0.26
DBSCAN (Major cluster)	44 460	13%	13%		
MiniSOM	49 585	15%	18.8%	0.45	0.99
MiniSOM w/ Tunning	59 502	18%	19.6%	0.64	0.99
SOMPY	51 422	15.5%	22.8%	0.2	0.78
SOMPY (hit frequency)	68 722	20.8%	22.61%		
SOMPY w/ Knn	45 769	13.8%	6.72%		

Based on Table 7, it is evident that MiniSOM, both with and without tuning, stands out as the most reliable unsupervised model for detecting anomalous behaviours, as indicated by its high silhouette scores (0.45 and 0.64) and R-squared values (0.99). However, when considering the proportion of dark activity detected, SOMPY identifies the highest percentage at 22.8%, demonstrating its higher effectiveness in capturing these instances despite lower clustering performance metrics.

4.5. DEPLOYMENT

For the deployment phase of this research, we utilised FastAPI, a modern, fast (high-performance) web framework for building APIs with Python 3.6+ based on standard Python-type hints. FastAPI is designed to be easy to use and learn while also being highly performant and capable of handling production-grade applications (FastAPI, n.d.).

FastAPI was employed to deploy the supervised machine learning algorithms that were found to classify dark activity with high accuracy. The deployment allows users to interact with the models through a user-friendly web interface. Users can select a model from a dropdown menu and upload raw AIS data for processing and classification. The tool was designed to handle the preprocessing of the uploaded data according to the selected model's requirements. This includes necessary data transformations and scaling, ensuring that the data is in the correct format for the model. Once the data is pre-processed, the selected model is applied to classify the data, and the results are returned to the user. This is possible because the models were pre-trained and saved using the "JobLib" package (*Joblib: Running Python Functions as Pipeline Jobs — Joblib 1.4.2 Documentation*, n.d.), allowing them to be easily integrated into the FastAPI pipeline for efficient and seamless deployment.

In Figure 24, the user can select the desired model and upload the raw AIS data file. The backend, implemented using FastAPI, handles the data preprocessing and applies the pre-trained model to classify the data.

In Figure 25, the prediction response from the API is displayed, showing the classification results for the uploaded data. This output can be shaped to show different information types associated with the prediction response according to stakeholders' needs.

Implementing this deployment ensures that end-users can easily and effectively use the supervised models without requiring in-depth technical knowledge. This approach facilitates the application of the developed models in real-world scenarios, enabling timely and accurate detection of dark activities.

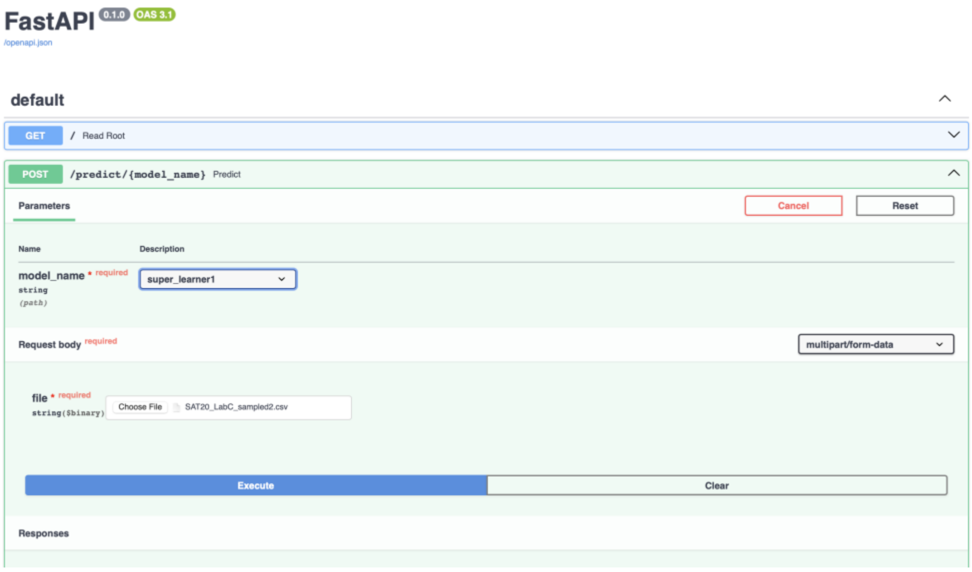

Figure 24 – FastAPI User Interface for Model Selection and Data Upload.

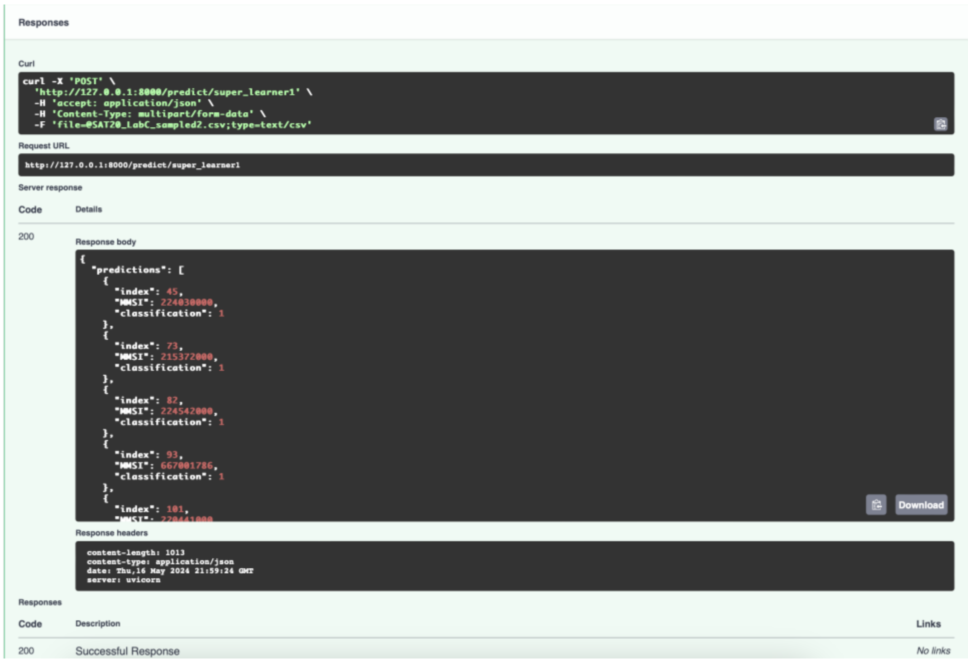

Figure 25 – FastAPI Prediction response.

4.6. CHALLENGES AND LIMITATIONS

Several challenges and limitations became apparent throughout this research, influencing the study's outcomes and providing valuable insights for future investigations.

One of the most significant challenges faced was the inherent imbalance within the dataset. The disproportionate number of instances of non-dark activity compared to dark activity presented difficulties in model training, as classifiers tend to be biased towards the majority class. The complexity of the models, particularly the deep learning architectures, introduced additional challenges. While more sophisticated models have the potential to capture complex patterns, they also require careful tuning and substantial computational resources. Moreover, the risk of overfitting increases with complexity, observed when the models performed well on training data but less on unseen test data.

The tuning of hyperparameters, especially in deep learning models and unsupervised learning, was time-consuming and computationally intensive. While techniques such as Grid Search, Random Search, and Bayesian Optimization were employed, finding the optimal set of hyperparameters proved to be a non-trivial task. The subtle nuances in hyperparameter values had significant impacts on model performance, requiring extensive experimentation that could not be exhaustively explored within the time constraints of this research.

Challenges also emerged with the unsupervised algorithms applied in this study. Obtaining good clustering results and effectively identifying anomalies was particularly difficult. The objective was to ensure that the detected dark activities were within the general anomaly detection framework. However, the difficulty in achieving this goal highlighted the limitations of unsupervised learning methods in distinguishing subtle patterns associated with dark activities. Nonetheless, the unsupervised models could form clusters and identify anomalous behaviour, differentiating it from normal behaviour, albeit not perfectly.

The application of semi-supervised learning techniques was also considered in this research. However, the imbalanced nature of the dataset posed a significant challenge for this approach. In semi-supervised learning, a portion of the data, typically around 10%, is labelled and used to guide the learning process (Chapelle et al., 2006). This necessitates careful selection of the labelled data to ensure meaningful representation. Random allocation of labelled data could result in none of the dark activity cases being included in the labelled subset, severely limiting the model's ability to learn from these critical instances. Conversely, if all confirmed dark activity cases were included in the labelled portion, it could introduce bias, skewing the model's understanding. Consequently, despite exploring semi-supervised methods, the inherent dataset imbalance prevented these techniques from yielding satisfactory results.

Moreover, the behaviour of ships that intentionally go dark or engage in illegal activities poses another significant challenge. Such ships often mimic standard behaviour patterns to avoid detection, which complicates the task for algorithms attempting to identify these activities. The sophistication of these evasion tactics means that unsupervised models, which rely on identifying deviations from typical patterns, may not always effectively detect these subtle anomalies. This issue underscores the need for more advanced techniques and more granular data to improve anomaly detection in maritime surveillance.

These challenges highlight the importance of continued research and development in this field.

5. CONCLUSIONS

This research set out to address the critical challenge of enhancing maritime security by detecting and identifying periods of dark activity using AIS data and machine learning techniques. The study aimed to bridge the gap between the abundant data generated by modern maritime systems and the actionable insights needed for effective monitoring and response. By leveraging a combination of supervised and unsupervised machine learning models, this research has made significant strides in advancing the capabilities of maritime anomaly detection.

To address the general objective of developing a model to detect dark activity using AIS data, the research was guided by several specific research questions:

SRQ1: What are the key features in AIS data that can be used to identify periods of dark activity?

The research identified key features such as vessel MMSI, position (Latitude and Longitude), timestamp, speed and course to be crucial in distinguishing periods of dark activity from normal behaviour. In contrast, the time difference between consecutive reports, the ship's country and the navigational status appeared less relevant for this classification (see feature importance graphs in Annexes).

SRQ2: How can data preprocessing techniques be optimised to handle the imbalanced nature of AIS data?

Various preprocessing techniques were employed, including oversampling, under-sampling, and synthetic sampling approaches. Among these methods, oversampling and synthetic strategies significantly improved the algorithms' performance. It is important to note that the sampling strategy is a crucial parameter that must be meticulously analysed to avoid excessively replicating the same samples. This could lead to model bias and incorrect predictions due to overfitting to repetitive situations. In this case, oversampling strategies were used not to balance the dataset completely (50/50) but to increase the minority class representation from 4.1% to 12%, allowing the model to learn those patterns better. Although these preprocessing strategies improved the results, the choice of robust models was the crucial step in handling the imbalanced nature of the data.

SRQ3: Which supervised machine learning algorithms most effectively detect dark activity in AIS data?

Among the supervised models, Super Learner 1 showed the best result in detecting dark activity. This was followed by Gradient Boosting, stacking 1 and 2 and deep forest, demonstrating promising results in achieving the research goal. These models effectively handled the imbalanced data and provided high F1 scores and recall for the minority class.

Additionally, they exhibited high true positive and true negative rates, along with low false positive and false negative rates.

SRQ4: How can unsupervised learning methods contribute to detecting anomalous maritime behaviours, including dark activity?

Unsupervised models such as Self-Organizing Maps (SOM) and Density-Based Spatial Clustering of Applications with Noise (DBSCAN) were implemented to detect general anomalies. Although these models effectively identified clusters of anomalous behaviour, they struggled to specifically identify instances of dark activity, with the best scenario identifying only 22% of the dark activity cases as anomalies. It should be noted that unsupervised learning is known for identifying deviations from patterns and clusters that are remarkably different from others. In this case, vessels tend to imitate standard patterns to avoid detection, which could explain the difficulty in identifying dark activity cases without specific knowledge of the behaviour's patterns.

SRQ5: What are the most effective evaluation metrics for assessing the performance of dark activity detection models?

The study used a variety of evaluation metrics, including accuracy, precision, recall, F1-score, confusion matrix, silhouette Score, and R-squared values. Emphasis was placed on the F1-Score and recall for the minority class within supervised algorithms to ensure effective detection of dark activities.

SRQ6: How can the developed models be employed for real-time detection of dark activity using new data?

The deployment phase involved implementing the developed models using the FastAPI tool. This practical application enabled real-time detection and classification of dark activity from new AIS data, providing a user-friendly interface for maritime authorities.

SRQ7: What are the main challenges and limitations encountered in detecting dark activity, and how can they be mitigated?

The research identified several challenges, including the complexity of models, high computational resource demands, and difficulties in achieving perfect clusters with unsupervised methods. Future research should focus on refining deep learning models with additional data and computational resources to maximise their potential. Additionally, new techniques such as active learning, autoencoders, and hybrid models could be explored to mitigate these challenges.

In addition to addressing the specific research questions, the study made several significant contributions. It highlighted the impact of geographical factors by analysing dark activity incidents by flag and monthly patterns, providing insights into geographical and temporal variations. The interdisciplinary approach, integrating data from multiple sources, including satellite data and additional maritime information, enhanced the robustness of dark activity detection models. The study advanced methodologies for detecting dark activities using AIS data, developed robust models for real-world applications, and addressed the challenges posed by imbalanced datasets through advanced preprocessing techniques. The practical deployment of these models using FastAPI demonstrated their potential for real-time anomaly detection.

Overall, the research successfully advanced the field of maritime anomaly detection by developing and validating advanced machine learning models, tackling data preprocessing challenges, and ensuring practical deployment. These contributions offer valuable insights and tools for maritime authorities. Future research should build on these findings, refining models and exploring new approaches to ensure even greater reliability. This study provides a comprehensive framework for enhancing maritime security and monitoring, contributing significantly to safer and more secure maritime operations. Ultimately, the methodologies and models developed here have the potential to revolutionise the way maritime security is managed, paving the way for more proactive and effective measures against illegal activities at sea.

6. FUTURE WORK

This research has made significant strides in developing and testing various machine-learning models to detect dark activity in maritime AIS data. Despite the promising results, several avenues remain for further investigation and improvement. Future research could benefit from acquiring more comprehensive and balanced datasets, involving collaborative efforts with maritime authorities and organisations to gather more instances of dark activity, ensuring a richer and more varied dataset. Integrating other sources, such as satellite imagery or radar data, could also provide complementary information to improve model accuracy.

While unsupervised models like Isolation Forest, DBSCAN, and SOM were explored, these models showed limitations in perfectly identifying dark activity. Future work could investigate more advanced anomaly detection techniques, such as hybrid models that combine supervised and unsupervised learning techniques to leverage the strengths of both approaches. Deep learning-based anomaly detection methods, such as autoencoders or Generative Adversarial Networks (GANs)[7], may capture more complex patterns in the data. Developing ensemble methods that combine multiple unsupervised algorithms could also enhance anomaly detection performance and reduce false positives.

Deploying the models in a real-time operational setting poses several challenges. Future work could address these issues by implementing real-time data processing pipelines using technologies like Apache Kafka[8] to handle continuous streams of AIS data. Another critical aspect is ensuring the models and preprocessing steps are scalable to handle large volumes of data efficiently. Exploring edge computing solutions to process data closer to the source, thereby reducing latency and bandwidth usage, is also worth investigating. For the models to be practically useful, they must be seamlessly integrated into existing maritime operations. It is essential to develop user-friendly interfaces and dashboards that allow maritime operators to interact with the models and interpret their outputs. Integrating the models into decision support systems that provide actionable insights and recommendations for maritime authorities could further enhance their utility. Another critical consideration is establishing feedback loops where the models' predictions can be validated by human experts, and the results can be used to refine and improve the models further.

The maritime environment is dynamic, and patterns of dark activity may evolve over time. Future research could investigate implementing online learning algorithms that allow the models to continuously update and adapt to new data. Developing adaptive models that can automatically adjust their parameters based on changing conditions and new information could further enhance the effectiveness and applicability of these models.

[7] ML models where two neural networks compete against each other to generate realistic synthetic data(Goodfellow et al., 2014)

[8] Distributed event streaming platform designed for managing and streaming large volumes of data in real-time across distributed systems (*Apache Kafka*, n.d.)

BIBLIOGRAPHICAL REFERENCES

A Missão. (n.d.). Marinha. Retrieved 19 March 2024, from http://www.marinha.pt/pt/a-marinha/Paginas/missao.aspx

AIS Fundamentals. (n.d.). Spire Maritime Documentation. Retrieved 20 March 2024, from https://documentation.spire.com/ais-fundamentals/

Akka Guide: Akka Guide. (n.d.). Retrieved 8 June 2024, from https://doc.akka.io/guide/index.html#

Ali, N., Chen, J., Fu, X., Ali, R., Hussain, M. A., Daud, H., Hussain, J., & Altalbe, A. (2024). Integrating Machine Learning Ensembles for Landslide Susceptibility Mapping in Northern Pakistan. *Remote Sensing*, *16*(6), Article 6. https://doi.org/10.3390/rs16060988

All You Need to Know About Support Vector Machines. (n.d.). *Spiceworks*. Retrieved 21 March 2024, from https://www.spiceworks.com/tech/big-data/articles/what-is-support-vector-machine/

Apache Hadoop. (n.d.). Retrieved 8 June 2024, from https://hadoop.apache.org/

Apache Kafka. (n.d.). Apache Kafka. Retrieved 9 June 2024, from https://kafka.apache.org/intro

Asan, U., & Ercan, S. (2012). *An Introduction to Self-Organizing Maps* (pp. 299–319). https://doi.org/10.2991/978-94-91216-77-0_14

Balduzzi, M., Pasta, A., & Wilhoit, K. (2014). *A security evaluation of AIS automated identification system* (p. 445). https://doi.org/10.1145/2664243.2664257

Belete, D., & D H, M. (2021). Grid search in hyperparameter optimization of machine learning models for prediction of HIV/AIDS test results. *International Journal of Computers and Applications*, *44*, 1–12. https://doi.org/10.1080/1206212X.2021.1974663

Bernabé, P., Gotlieb, A., Legeard, B., Marijan, D., Sem-Jacobsen, F. O., & Spieker, H. (2023). Detecting Intentional AIS Shutdown in Open Sea Maritime Surveillance Using Self-Supervised Deep Learning. *IEEE Transactions on Intelligent Transportation Systems*, 1–12. https://doi.org/10.1109/TITS.2023.3322690

Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Byerly, A. B., Zhang, W. C., Iwarere, S. A., Malik, W. A., Bish, S. F., Haque, M. A., & Sookoor, T. I. (2022). Neptune: An Automated System for Dark Ship Detection, Targeting, and Prioritization. *Johns Hopkins APL Technical Digest*, *36*(2).

Castaldo, F., Palmieri, F., Bastani, V., Marcenaro, L., & Regazzoni, C. (2014, August). *Online Bayesian learning and classification of ship-to-ship interactions for port safety | IEEE Conference Publication | IEEE Xplore*. 11th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), Seoul, Korea. https://ieeexplore.ieee.org/document/6918688

Ceballos, F. (2019, September 14). *Stacking Classifiers for Higher Predictive Performance*. Medium. https://towardsdatascience.com/stacking-classifiers-for-higher-predictive-performance-566f963e4840

Chapelle, O., Schölkopf, B., & Zien, A. (Eds.). (2006). *Semi-supervised learning*. MIT Press.

Chapman, P., Clinton, J., & Kerber, K. (2000). *CRISP-DM 1.0—Step-by-step data mining*.

Chawla, N., Bowyer, K., Hall, L., & Kegelmeyer, W. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *J. Artif. Intell. Res. (JAIR)*, *16*, 321–357. https://doi.org/10.1613/jair.953

Chicco, D., & Jurman, G. (2020). The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*, *21*(1), 6. https://doi.org/10.1186/s12864-019-6413-7

Concetti, L., Mazzuto, G., Ciarapica, F. E., & Bevilacqua, M. (2023). An Unsupervised Anomaly Detection Based on Self-Organizing Map for the Oil and Gas Sector. *Applied Sciences*, *13*(6), Article 6. https://doi.org/10.3390/app13063725

Cretulescu, R., Morariu, D., Breazu, M., & Volovici, D. (2019). DBSCAN Algorithm for Document Clustering. *International Journal of Advanced Statistics and IT&C for Economics and Life Sciences*, *9*, 58–66. https://doi.org/10.2478/ijasitels-2019-0007

Desiani, A., Adrezo, M., Kresnawati, E. S., Ermatita, Akbar, M., & Hasibuan, M. S. (2023). Back Translation-EDA and Transformer for Hate Speech Classification in Indonesian. *2023 International Conference on Informatics, Multimedia, Cyber and Informations System (ICIMCIS)*, 611–616. https://doi.org/10.1109/ICIMCIS60089.2023.10348979

Dictionary.com | Meanings & Definitions of English Words. (2024, June 8). Dictionary.Com. https://www.dictionary.com/browse/smuggle

Dobrkovic, A., Iacob, M.-E., & van Hillegersberg, J. (2015). Using machine learning for unsupervised maritime waypoint discovery from streaming AIS data. *Proceedings of the 15th International Conference on Knowledge Technologies and Data-Driven Business*, 1–8. https://doi.org/10.1145/2809563.2809573

Dobrkovic, A., Iacob, M.-E., & van Hillegersberg, J. (2018). Maritime pattern extraction and route reconstruction from incomplete AIS data. *International Journal of Data Science and Analytics*, *5*(2), 111–136. https://doi.org/10.1007/s41060-017-0092-8

Durlik, I., Miller, T., Cembrowska-Lech, D., Krzemińska, A., Złoczowska, E., & Nowak, A. (2023). Navigating the Sea of Data: A Comprehensive Review on Data Analysis in Maritime IoT

Applications. *Applied Sciences*, *13*(17), Article 17. https://doi.org/10.3390/app13179742

European Commission. Joint Research Centre. (2016). *Maritime knowledge discovery and anomaly detection: Workshop proceedings : Ispra, 5 6 July 2016.* Publications Office. https://data.europa.eu/doi/10.2788/025881

FastAPI. (n.d.). Retrieved 28 May 2024, from https://fastapi.tiangolo.com/

Frost, J. (2017, April 16). *How To Interpret R-squared in Regression Analysis*. Statistics By Jim. http://statisticsbyjim.com/regression/interpret-r-squared-regression/

Ganganwar, V. (2012). An overview of classification algorithms for imbalanced datasets. *International Journal of Emerging Technology and Advanced Engineering*, *2*, 42–47.

GISGeography. (2017, June 12). *Latitude, Longitude and Coordinate System Grids*. GIS Geography. https://gisgeography.com/latitude-longitude-coordinates/

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). *Generative Adversarial Networks* (arXiv:1406.2661). arXiv. http://arxiv.org/abs/1406.2661

Görkem, B. N., Çağlayan, B., Karaca, E., Karabulut, C., & Korçak, Ö. (2023). Dark Activity Detection in AIS-Based Maritime Networks. *2023 34th Conference of Open Innovations Association (FRUCT)*, 35–40. https://doi.org/10.23919/FRUCT60429.2023.10328171

guest_blog. (2020, July 23). 10 Techniques to Solve Imbalanced Classes in Machine Learning (Updated 2024). *Analytics Vidhya*. https://www.analyticsvidhya.com/blog/2020/07/10-techniques-to-deal-with-class-imbalance-in-machine-learning/

Gwilliam, B. (2023, January 10). *What Happens When a Vessel Goes 'AIS Dark'?* https://www.linkedin.com/pulse/what-happens-when-vessel-goes-ais-dark-ben-gwilliam

Hancock, J. T., Khoshgoftaar, T. M., & Johnson, J. M. (2023). Evaluating classifier performance with highly imbalanced Big Data. *Journal of Big Data*, *10*(1), 42. https://doi.org/10.1186/s40537-023-00724-5

Harati Mokhtari, A., Wall, A., Brooks, P., & Wang, J. (2008). *Automatic Identification System (AIS): A Human Factors Approach*.

Harnessing the power of Grid Search for optimized machine learning models | by Nilimesh Halder, PhD | Analyst's corner | Medium. (n.d.). Retrieved 21 March 2024, from https://medium.com/analysts-corner/harnessing-the-power-of-grid-search-for-optimized-machine-learning-models-5878e3abf2d6

HDF5 for Python—H5py 3.10.0 documentation. (n.d.). Retrieved 20 March 2024, from https://docs.h5py.org/en/stable/

He, H., Bai, Y., Garcia, E. A., & Li, S. (2008). ADASYN: Adaptive synthetic sampling approach for imbalanced learning. *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, 1322–1328. https://doi.org/10.1109/IJCNN.2008.4633969

Hendriksen, S. J. F. (2022). *Viability of anomaly detection aid for vessel traffic service operators*. University of Twente.

Hospedales, T., Antoniou, A., Micaelli, P., & Storkey, A. (2020). *Meta-Learning in Neural Networks: A Survey* (arXiv:2004.05439). arXiv. http://arxiv.org/abs/2004.05439

Huang, I.-L., Lee, M.-C., Nieh, C.-Y., & Huang, J.-C. (2024). Ship Classification Based on AIS Data and Machine Learning Methods. *Electronics*, *13*(1), Article 1. https://doi.org/10.3390/electronics13010098

IBM Documentation. (2021, August 17). https://www.ibm.com/docs/en/spss-modeler/saas?topic=dm-crisp-help-overview

imbalanced-learn documentation—Version 0.12.2. (n.d.). Retrieved 26 April 2024, from https://imbalanced-learn.org/stable/index.html

imbalanced-learn: Toolbox for imbalanced dataset in machine learning. (0.12.0). (n.d.). [C, Python; MacOS, Microsoft :: Windows, POSIX, Unix]. Retrieved 21 March 2024, from https://github.com/scikit-learn-contrib/imbalanced-learn

International Convention on Safety of life at Sea (SOLAS). (1974). International Maritime Organization (IMO). https://www.liscr.com/sites/default/files/SOLAS%20V_Reg19.pdf

International Maritime Organization (IMO). (2019). *AIS transponders*. International Maritime Organization (IMO) Website. https://www.imo.org/en/OurWork/Safety/Pages/AIS.aspx

Joblib: Running Python functions as pipeline jobs—Joblib 1.4.2 documentation. (n.d.). Retrieved 5 June 2024, from https://joblib.readthedocs.io/en/stable/

Jones, A., Koehler, S., Jerge, M., Graves, M., King, B., Dalrymple, R., Freese, C., & Von Albade, J. (2023). BATMAN: A Brain-like Approach for Tracking Maritime Activity and Nuance. *Sensors (14248220)*, *23*(5), 2424. https://doi.org/10.3390/s23052424

Kontopoulos, I., Chatzikokolakis, K., Zissis, D., Tserpes, K., & Spiliopoulos, G. (2020). Real-time maritime anomaly detection: Detecting intentional AIS switch-off. *Int. J. Big Data Intell.*, *7*(2), 85–96. https://doi.org/10.1504/IJBDI.2020.107375

Lane, R. O., Nevell, David. A., Hayward, S. D., & Beaney, T. W. (2010). Maritime anomaly detection and threat assessment. *2010 13th International Conference on Information Fusion*, 1–8. https://doi.org/10.1109/ICIF.2010.5711998

Laxhammar, R. (2008). Anomaly detection for sea surveillance. *2008 11th International Conference on Information Fusion*, 1–8. https://ieeexplore.ieee.org/document/4632192

Li, H., & Yang, Z. (2023). Incorporation of AIS data-based machine learning into unsupervised route planning for maritime autonomous surface ships. *Transportation Research Part E: Logistics and Transportation Review*, *176*, 103171. https://doi.org/10.1016/j.tre.2023.103171

Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation Forest. *2008 Eighth IEEE International Conference on Data Mining*, 413–422. https://doi.org/10.1109/ICDM.2008.17

Marine Regions · Mediterranean Sea (IHO Sea Area). (n.d.). Retrieved 20 March 2024, from https://www.marineregions.org/gazetteer.php?p=details&id=1905

Martins, T. P. G. (2023). *Radar Emitter Classification based on Deep Ensemble*.

Mazzarella, F., Vespe, M., Tarchi, D., Aulicino, G., & Vollero, A. (2016). AIS reception characterisation for AIS on/off anomaly detection. *2016 19th International Conference on Information Fusion (FUSION)*, 1867–1873. https://ieeexplore.ieee.org/abstract/document/7528110

Missão e Competências. (n.d.). Retrieved 19 March 2024, from https://www.amn.pt/AMN/Paginas/Missao.aspx

Mondal, S., Roy, S., R, J., E, M. J., Ghosh, M., & Nag, A. (2023). Random Forest-based Underwater Temperature Prediction Model for IoT Devices. *2023 IEEE 3rd Mysore Sub*

Section International Conference (MysuruCon), 1–6. https://doi.org/10.1109/MysuruCon59703.2023.10396911

Most Popular Data Science Methodologies. (2022, July 28). https://www.knowledgehut.com/blog/data-science/data-science-methodology

Müller, L.-M. (2023). *3D Generative Adversarial Networks to Autonomously Generate Building Geometry*. https://repository.tudelft.nl/islandora/object/uuid%3Ab4a44f69-e3a4-4a29-85c4-77a7e899b81a

Mund, S. K. (2021, July 26). FastAPI: The Modern Age API Framework For Pythonista. *Analytics Vidhya*. https://medium.com/analytics-vidhya/fastapi-the-modern-age-api-framework-for-pythonista-4b2cd1e6652

Muraina, I. (2022). *IDEAL DATASET SPLITTING RATIOS IN MACHINE LEARNING ALGORITHMS: GENERAL CONCERNS FOR DATA SCIENTISTS AND DATA ANALYSTS*. https://www.researchgate.net/publication/358284895_IDEAL_DATASET_SPLITTING_RATIOS_IN_MACHINE_LEARNING_ALGORITHMS_GENERAL_CONCERNS_FOR_DATA_SCIENTISTS_AND_DATA_ANALYSTS

Najafzadeh, M., Basirian, S., & Li, Z. (2023). Vulnerability of the rip current phenomenon in marine environments using machine learning models. *Results in Engineering*, *21*, 101704. https://doi.org/10.1016/j.rineng.2023.101704

Natekin, A., & Knoll, A. (2013). Gradient Boosting Machines, A Tutorial. *Frontiers in Neurorobotics*, *7*, 21. https://doi.org/10.3389/fnbot.2013.00021

Nguyen, D., Vadaine, R., Hajduch, G., Garello, R., & Fablet, R. (2022). GeoTrackNet—A Maritime Anomaly Detector Using Probabilistic Neural Network Representation of AIS Tracks and A Contrario Detection. *IEEE Transactions on Intelligent Transportation Systems*, *23*(6), 5655–5667. https://doi.org/10.1109/TITS.2021.3055614

Pisner, D. (2020). Chapter 6—Support vector machine. In *Machine Learning Methods and Applications to Brain Disorders*. https://www.sciencedirect.com/science/article/abs/pii/B9780128157398000067

Ramirez, A. (2021, September 5). 3 Most Popular Data Science Metho. *Medium*. https://medium.com/@aj.ramirez23/3-most-popular-data-science-metho-e61f6600b83f

Ray, C., Iphar, C., & Napoli, A. (2016). *Methodology for Real-Time Detection of AIS Falsification*.

Ribeiro, C. V., Paes, A., & Oliveira, D. de. (2023). AIS-based maritime anomaly traffic detection: A review. *Expert Systems with Applications*, *231*, 120561. https://doi.org/10.1016/j.eswa.2023.120561

Riveiro, M., & Falkman, G. (2009). Interactive Visualization of Normal Behavioral Models and Expert Rules for Maritime Anomaly Detection. *Imaging and Visualization 2009 Sixth International Conference on Computer Graphics*, 459–466. https://doi.org/10.1109/CGIV.2009.54

Schröer, C., Kruse, F., & Gómez, J. M. (2021). A Systematic Literature Review on Applying CRISP-DM Process Model. *Procedia Computer Science*, *181*, 526–534. https://doi.org/10.1016/j.procs.2021.01.199

Scikit-learn: Machine learning in Python—Scikit-learn 1.4.1 documentation. (n.d.). Retrieved 21 March 2024, from https://scikit-learn.org/stable/

Self organizing map – Klaus Kayser. (n.d.). Retrieved 30 May 2024, from https://www.ki-universum.de/wordpress/index.php/2021/03/22/self-organizing-map/

sevamoo. (2024). *Sevamoo/SOMPY* [Jupyter Notebook]. https://github.com/sevamoo/SOMPY (Original work published 2014)

Shahapure, K. R., & Nicholas, C. (2020). Cluster Quality Analysis Using Silhouette Score. *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, 747–748. https://doi.org/10.1109/DSAA49011.2020.00096

Spelmen, V. S., & Porkodi, R. (2018). A Review on Handling Imbalanced Data. *2018 International Conference on Current Trends towards Converging Technologies (ICCTCT)*, 1–11. https://doi.org/10.1109/ICCTCT.2018.8551020

Story, R. (n.d.). *folium: Make beautiful maps with Leaflet.js & Python* (0.16.0) [Python]. Retrieved 21 March 2024, from https://github.com/python-visualization/folium

Terroso-Saenz, F., Valdes-Vela, M., & Skarmeta-Gomez, A. F. (2016). A complex event processing approach to detect abnormal behaviours in the marine environment. *Information Systems Frontiers*, *18*(4), 765–780. https://doi.org/10.1007/s10796-015-9560-7

Vanneschi, L., Castelli, M., Costa, E., Re, A., Vaz, H., Lobo, V., & Urbano, P. (2015). Improving Maritime Awareness with Semantic Genetic Programming and Linear Scaling: Prediction of Vessels Position Based on AIS Data. In A. M. Mora & G. Squillero (Eds.), *Applications of Evolutionary Computation* (pp. 732–744). Springer International Publishing. https://doi.org/10.1007/978-3-319-16549-3_59

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2023). *Attention Is All You Need* (arXiv:1706.03762). arXiv. http://arxiv.org/abs/1706.03762

Venskus, J., Treigys, P., Bernatavičienė, J., Medvedev, V., Voznak, M., Kurmis, M., & Bulbenkienė, V. (2017). Integration of a Self-Organizing Map and a Virtual Pheromone for Real-Time Abnormal Movement Detection in Marine Traffic. *Informatica*, *28*(2), 359–374. https://doi.org/10.15388/Informatica.2017.133

Vettigli, G. (2024). *JustGlowing/minisom* [Python]. https://github.com/JustGlowing/minisom (Original work published 2013)

Wang, X., Liu, X., Liu, B., de Souza, E. N., & Matwin, S. (2014). Vessel route anomaly detection with Hadoop MapReduce. *2014 IEEE International Conference on Big Data (Big Data)*, 25–30. https://doi.org/10.1109/BigData.2014.7004464

Watch, G. F. (2017, March 31). AIS for Safety and Tracking: A Brief History. *Global Fishing Watch*. https://globalfishingwatch.org/data/ais-brief-history/

Wei, Z., Xie, X., & Zhang, X. (2022). Maritime anomaly detection based on a support vector machine. *Soft Computing*, *26*(21), 11553–11566. https://doi.org/10.1007/s00500-022-07409-w

Wolpert, D. H. (1992). Stacked generalization. *Neural Networks*, *5*(2), 241–259. https://doi.org/10.1016/S0893-6080(05)80023-1

Wolsing, K., Roepert, L., Bauer, J., & Wehrle, K. (2022). Anomaly Detection in Maritime AIS Tracks: A Review of Recent Approaches. *Journal of Marine Science and Engineering*, *10*. https://doi.org/10.3390/jmse10010112

Yan, Z., Song, X., Zhong, H., Yang, L., & Wang, Y. (2022). Ship Classification and Anomaly Detection Based on Spaceborne AIS Data Considering Behavior Characteristics. *Sensors*, *22*(20), Article 20. https://doi.org/10.3390/s22207713

Yap, B. W., Rani, K. A., Rahman, H. A. A., Fong, S., Khairudin, Z., & Abdullah, N. N. (2014). An Application of Oversampling, Undersampling, Bagging and Boosting in Handling Imbalanced Datasets. In T. Herawan, M. M. Deris, & J. Abawajy (Eds.), *Proceedings of the First International Conference on Advanced Data and Information Engineering (DaEng-2013)* (pp. 13–22). Springer. https://doi.org/10.1007/978-981-4585-18-7_2

Zemicheal, T., & Dietterich, T. G. (2019). Anomaly detection in the presence of missing values for weather data quality control. *Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies*, 65–73. https://doi.org/10.1145/3314344.3332490

Zhen, R., Jin, Y., Hu, Q., Shao, Z., & Nikitakos, N. (2017). Maritime Anomaly Detection within Coastal Waters Based on Vessel Trajectory Clustering and Naïve Bayes Classifier. *Journal of Navigation*, *70*(3), 648–670. https://doi.org/10.1017/S0373463316000850

Zhou, Z.-H., & Feng, J. (2019). Deep forest. *National Science Review*, *6*, 74–86. https://doi.org/10.1093/nsr/nwy108

Zissis, D., Chatzikokolakis, K., Spiliopoulos, G., & Vodas, M. (2020). A Distributed Spatial Method for Modeling Maritime Routes. *IEEE Access*, *8*, 47556–47568. https://doi.org/10.1109/ACCESS.2020.2979612

ANNEXES

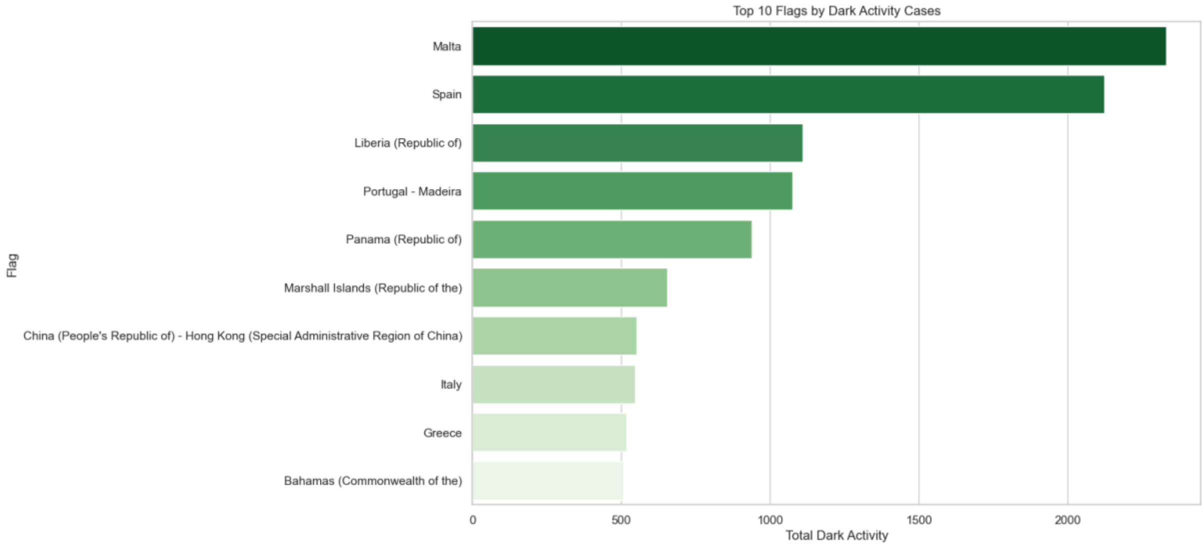

Annexe 1 - Top 10 Flags by Dark Activity

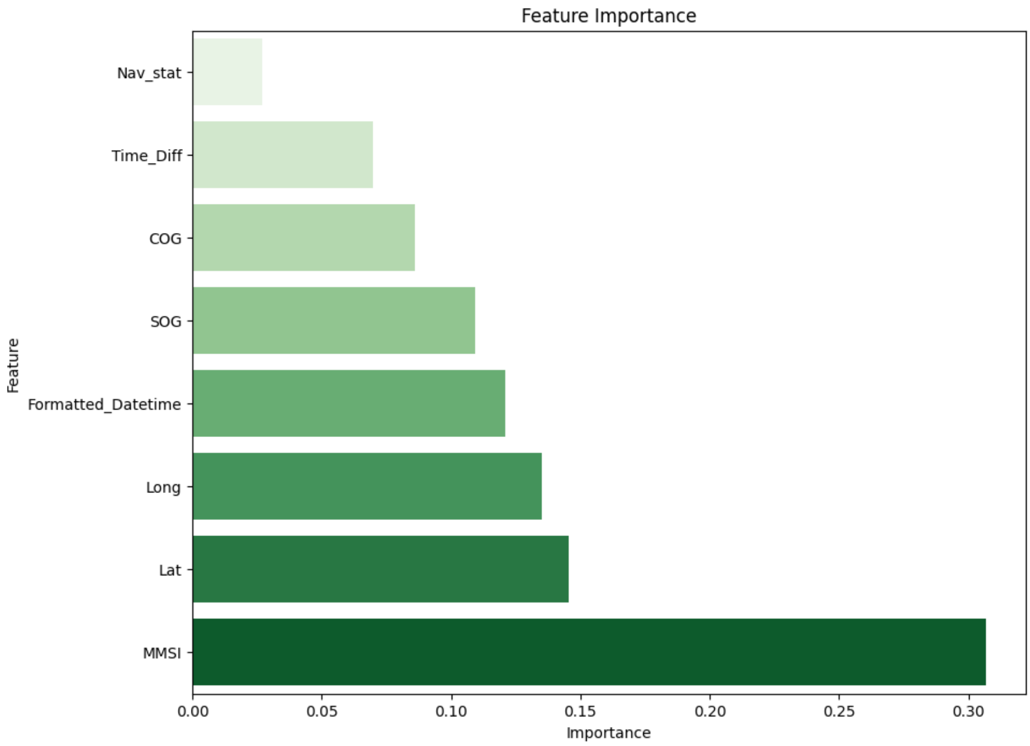

Annexe 2 - Feature importance for RF

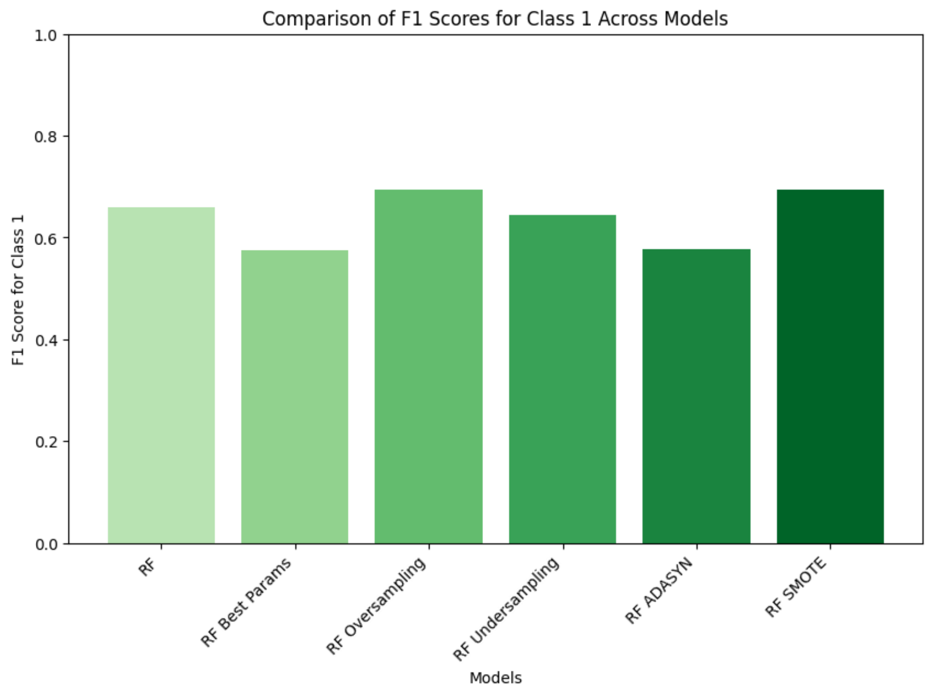

Annexe 3 - F1 Score for minority class across different implementations of RF

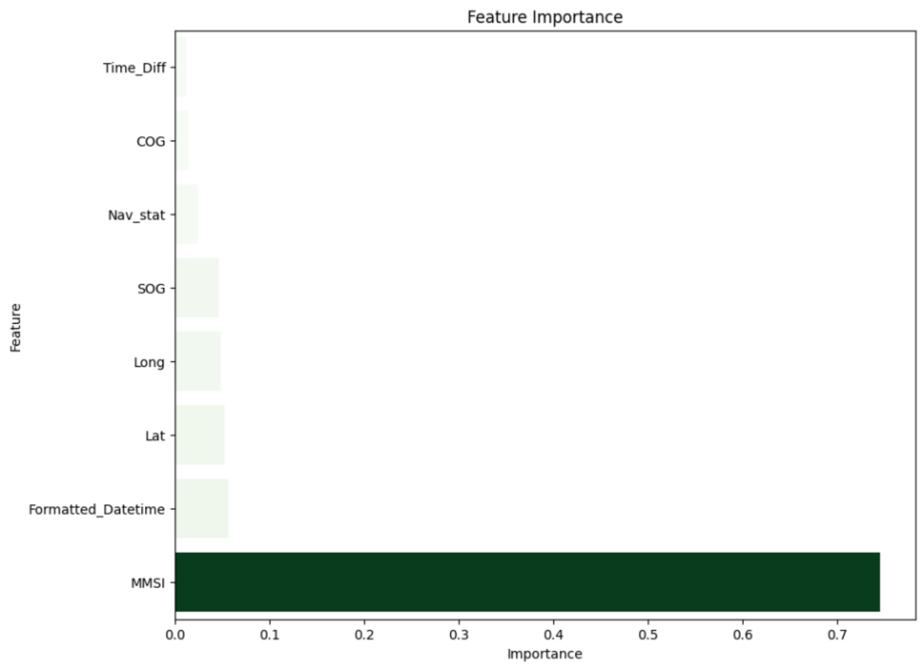

Annexe 4 - Feature Importance for GB

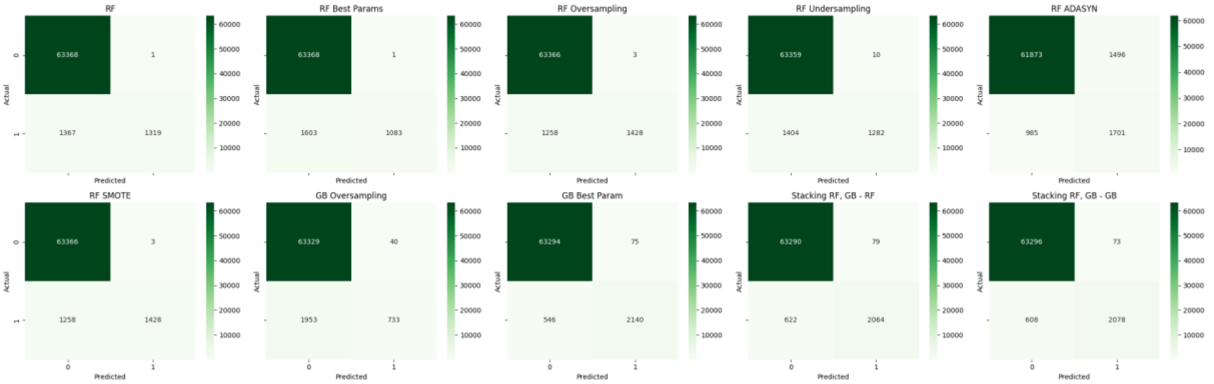

Annexe 5 - Confusion Matrixes for Stage 1 Modelling

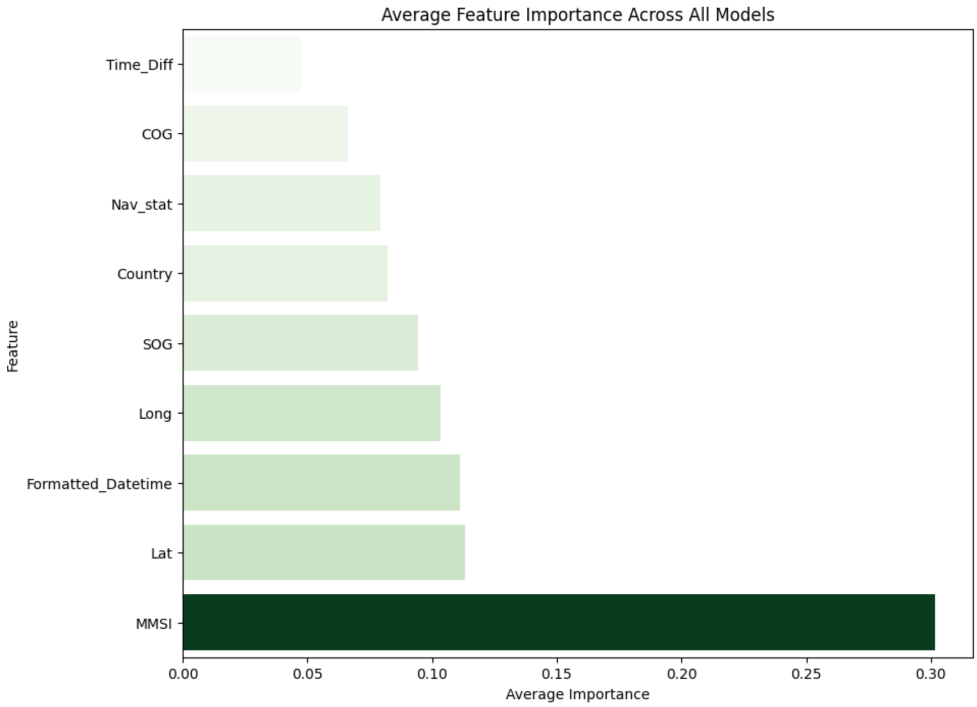

Annexe 6 - Feature Importance for Deep Forest Stack

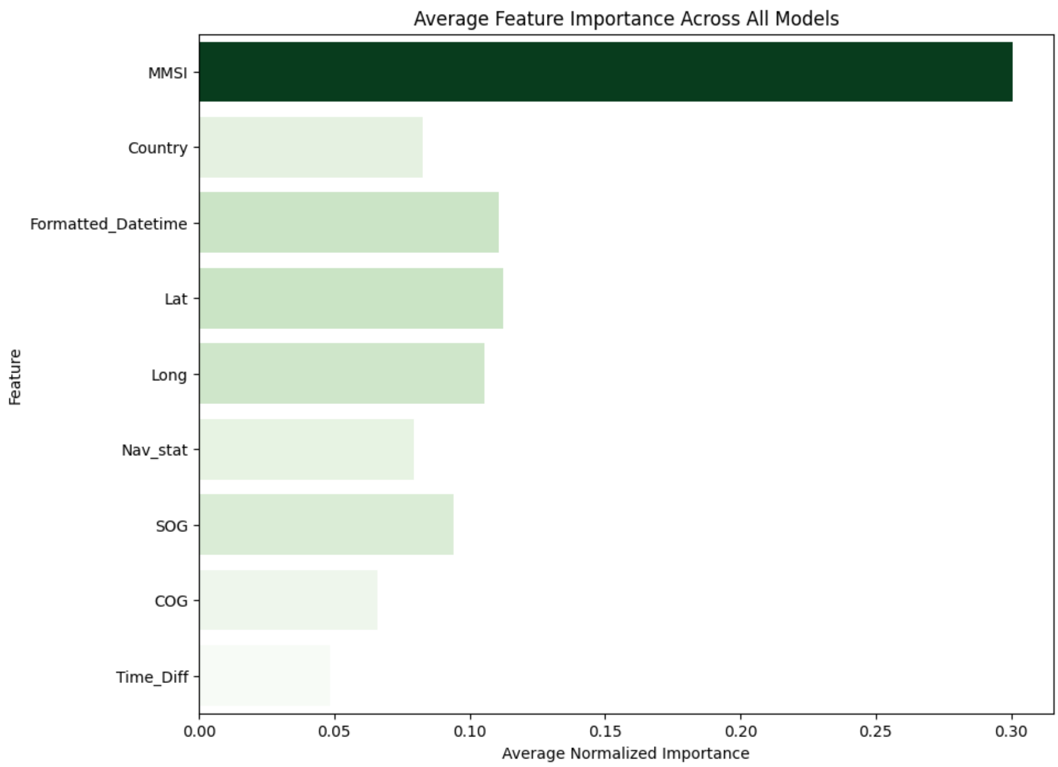

Annexe 7 - Feature Importance for Super Learner 1

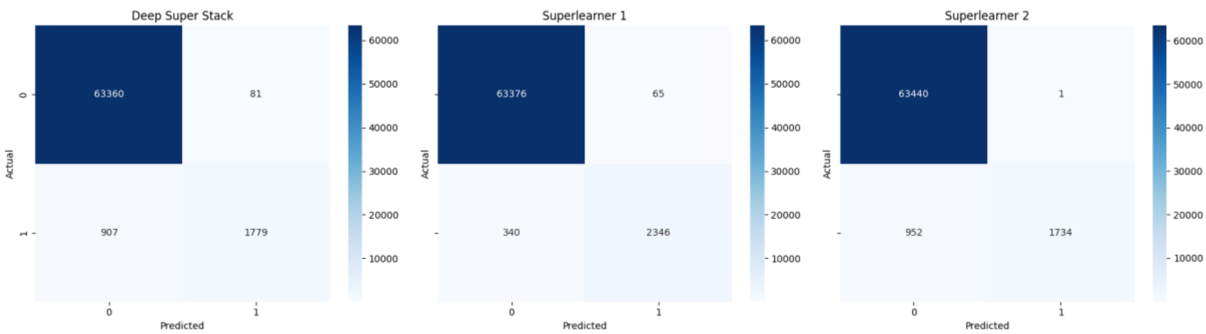

Annexe 8 - Confusion Matrixes for Stage 2 Modelling

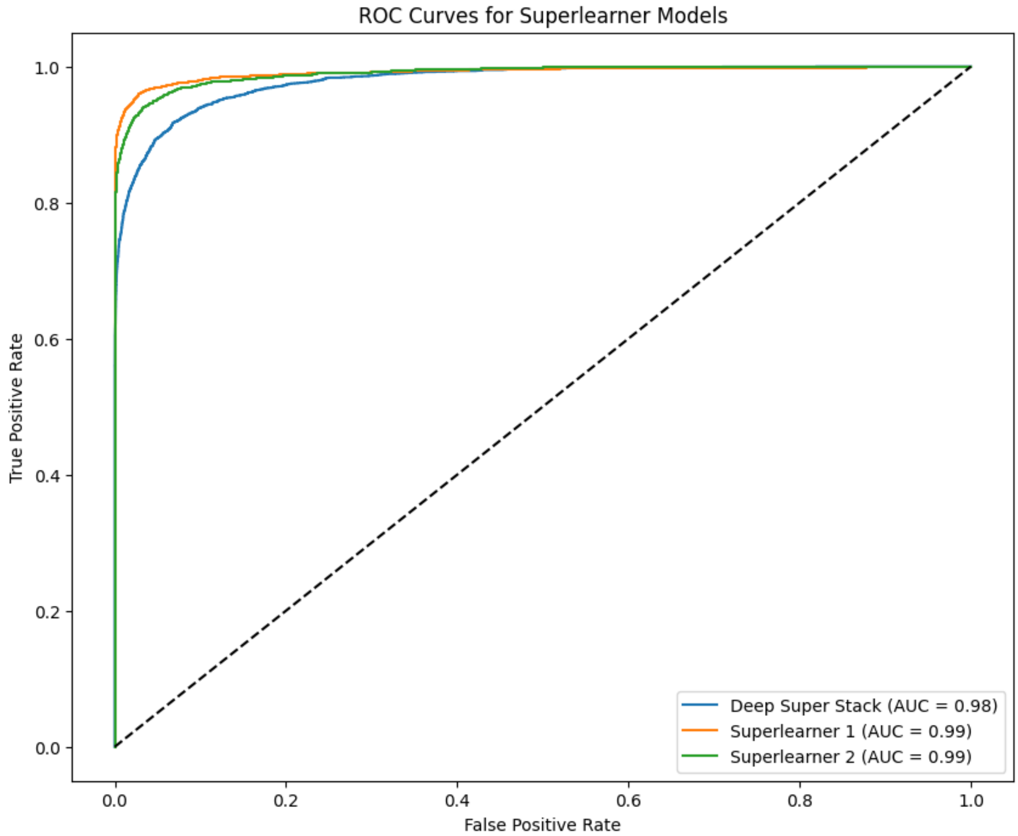

Annexe 9 - ROC Curves for Stage Modelling 2

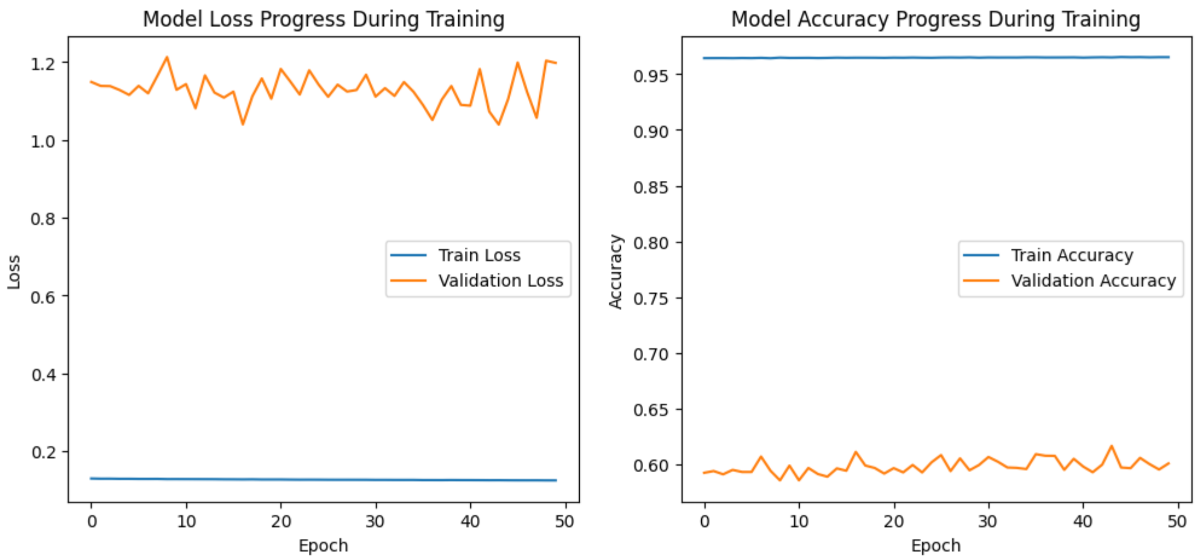

Annexe 10 - Loss and Accuracy Progress During Training for NN1

Annexe 11 - Loss and Accuracy Progress During Training for NN2

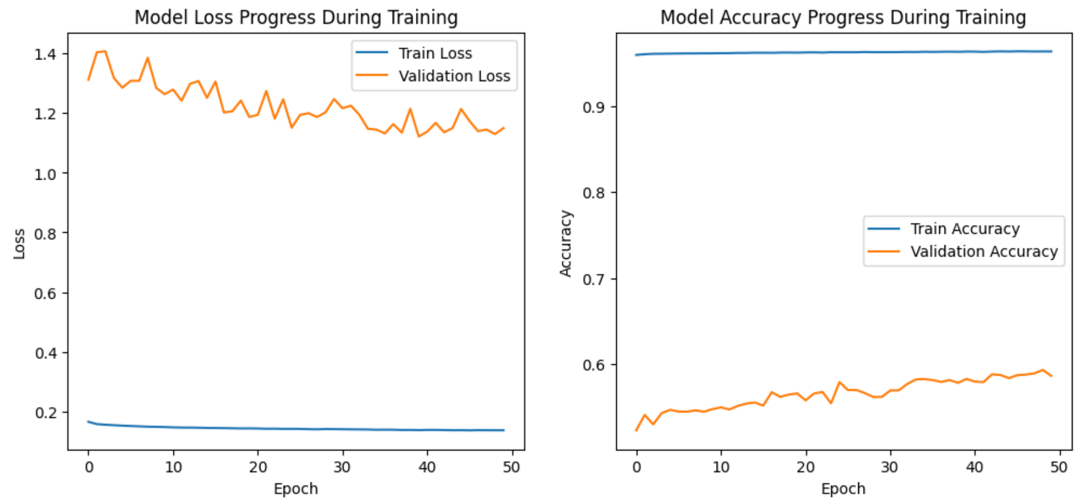

Annexe 12 - Loss and Accuracy Progress During Training for NN3

Annexe 13 - Loss and Accuracy Progress During Training for NN4

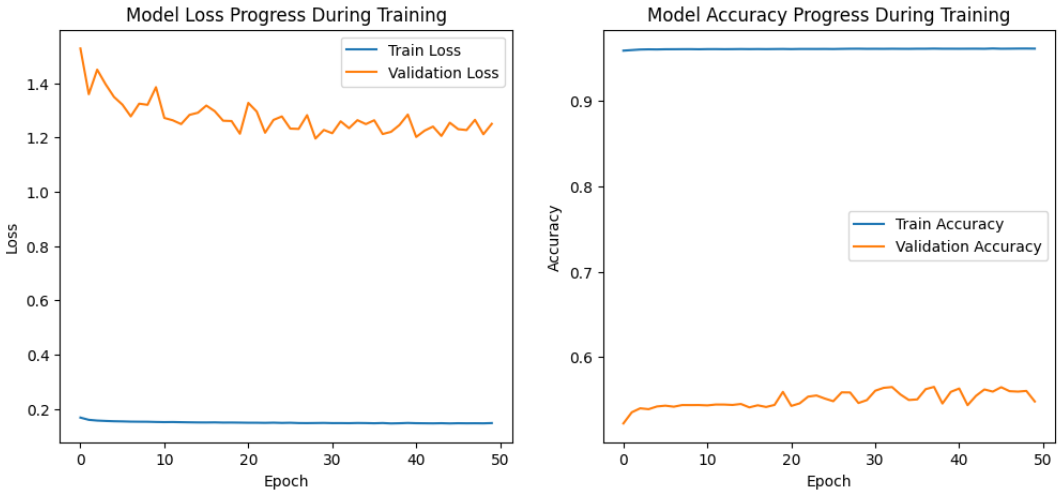

Annexe 14 - Loss and Accuracy Progress During Training for NN5

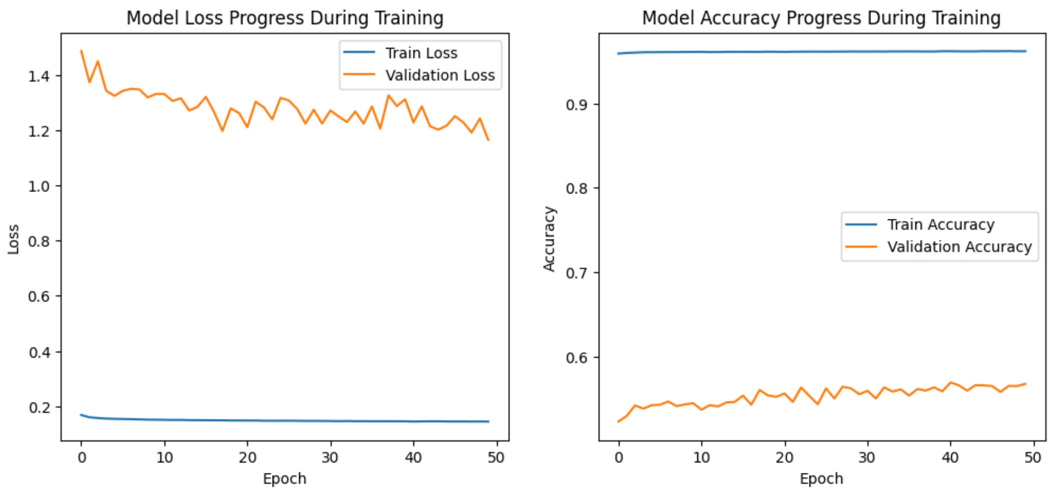

Annexe 15 - Loss and Accuracy Progress During Training for NN6

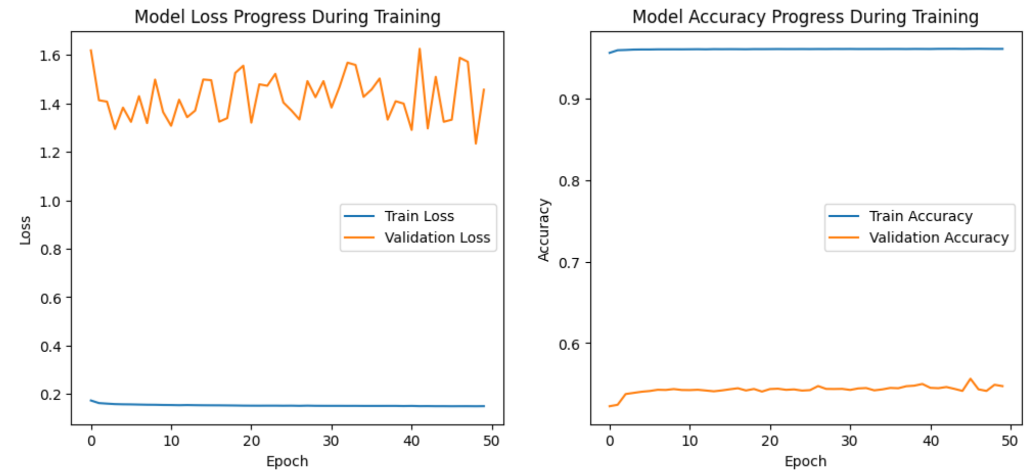

Annexe 16 - Loss and Accuracy Progress During Training for NN7

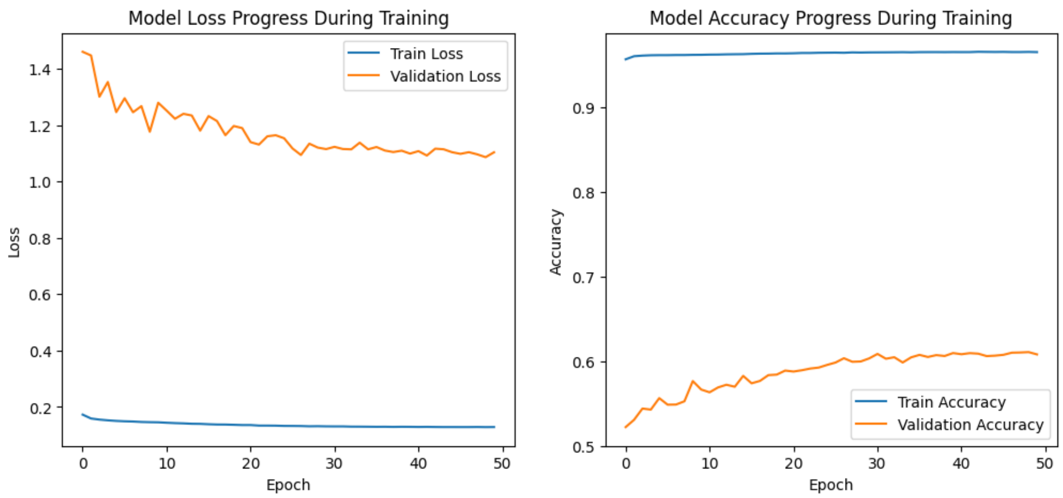

Annexe 17 - Loss and Accuracy Progress During Training for NN8

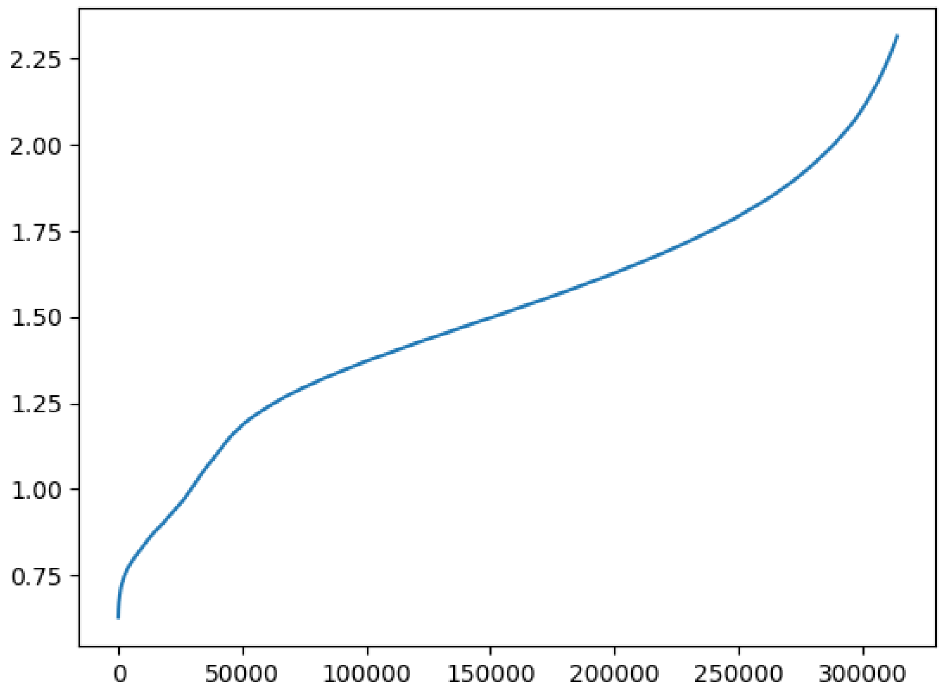

Annexe 18 - Plot of distances for DBSCAN algorithm in order to find the best EPS factor from the “elbow”

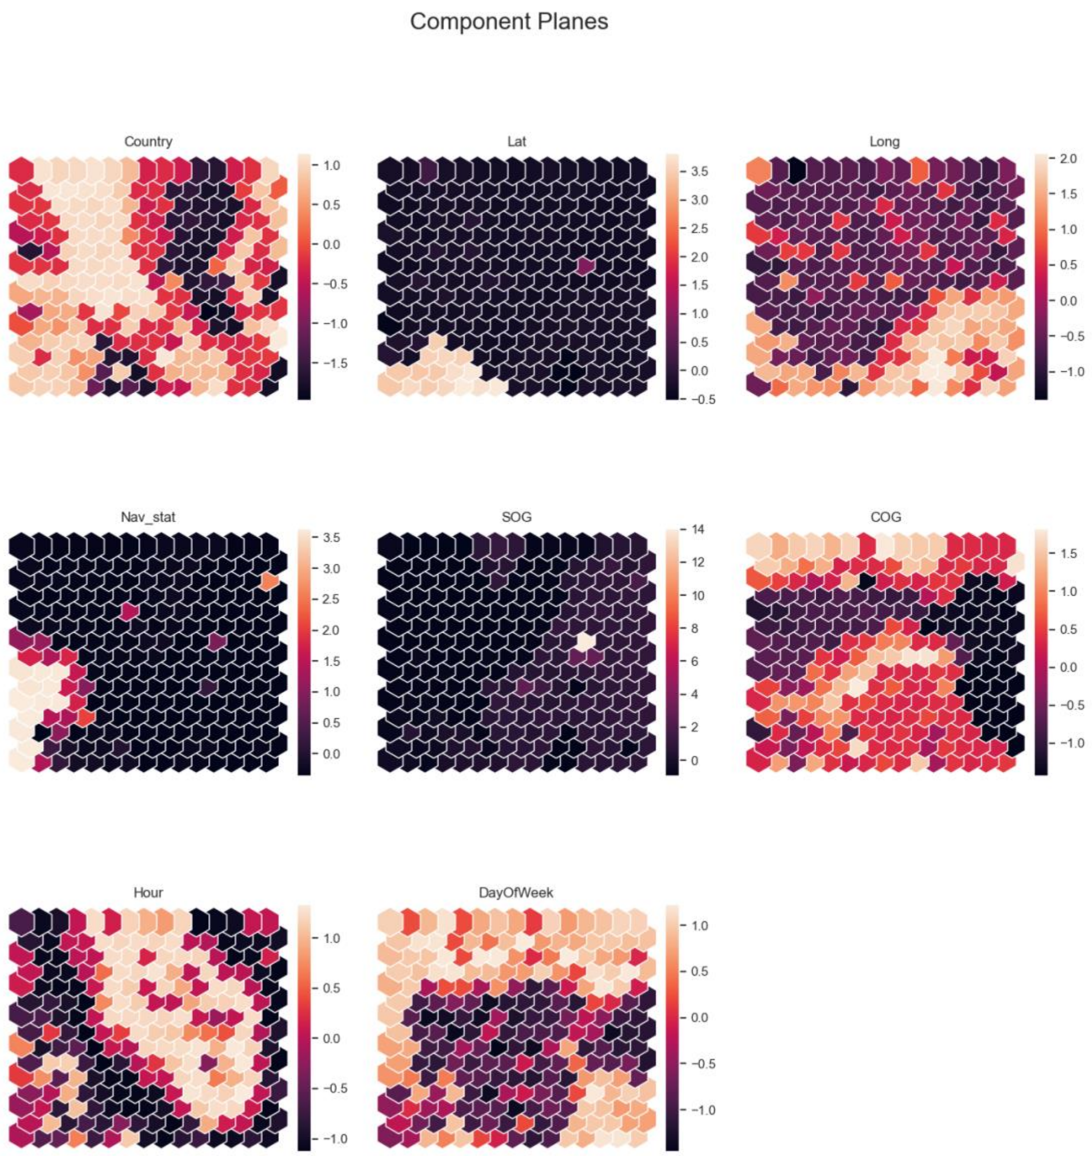

Annexe 19 - Component Planes of the SOM for AIS Data

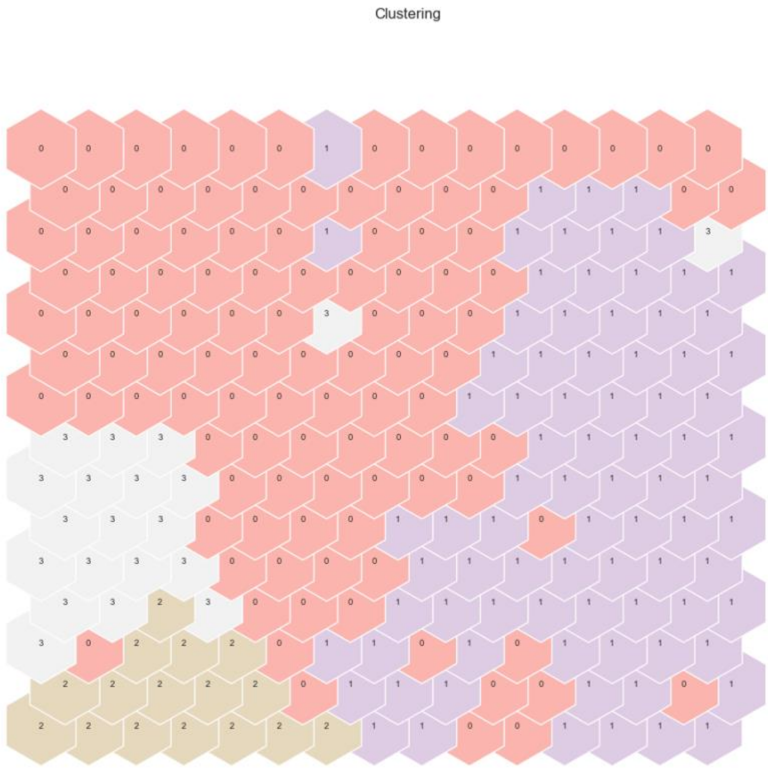

Annexe 20 - K-means on top of SOM, number of clusters defined as 4

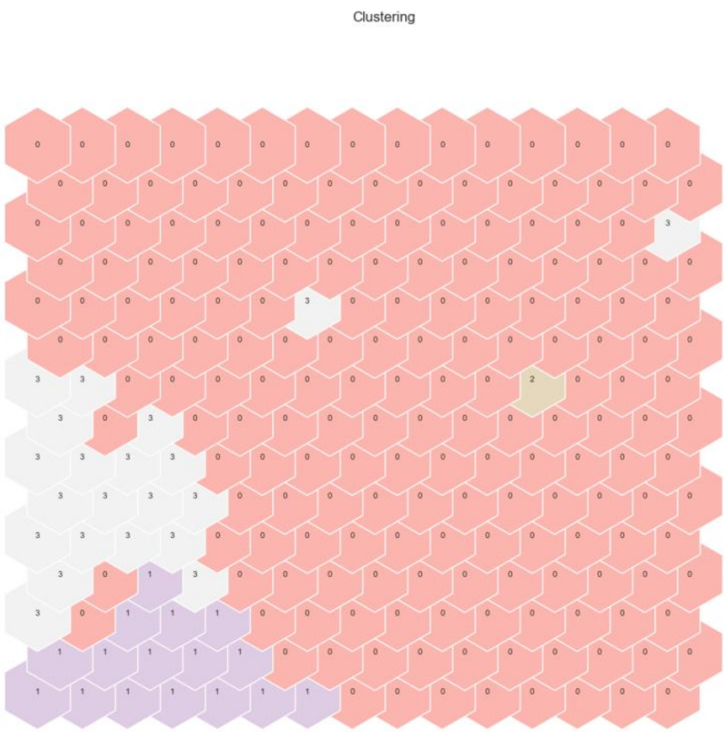

Annexe 21 – Hierarchical clustering on top of SOM, number of clusters defined as 4

NOVA IMS
Information Management School