

Behavioral Model Anomaly Detection in Automatic Identification Systems (AIS)

Jacob Coleman

Computer Science and Engineering
University of Tennessee at Chattanooga
Chattanooga, TN
jcoleman@imsa.global

Farah Kandah

Computer Science and Engineering
University of Tennessee at Chattanooga
Chattanooga, TN
farah-kandah@utc.edu

Brennan Huber

Computer Science and Engineering
University of Tennessee at Chattanooga
Chattanooga, TN
khs616@mocs.utc.edu

Abstract—Over 90% of all goods in the world, at some point in their life, are carried on a vessel at sea. Currently, the maritime industry relies on the Automatic Identification System (AIS) for collision avoidance, vessel tracking, and vessel awareness while operating at sea. AIS is a plaintext, unencrypted, unauthenticated protocol and, as such, is vulnerable to various types of attacks. Malicious actors can alter the AIS location of a vessel by spoofing a vessel or alter the channel the AIS receiver is using to send nefarious information to a vessel privately. With the advent of the Ocean of Things (OoT), vessels are sharing more information than vessel location alone at sea. As this information becomes critical for safe and efficient operation at sea, we in this work present a novel approach of applying machine learning to build behavior models for vessels at sea. These models allow vessels to detect anomalous communication from vessels nearby, thus enable vessels to determine the quality of the messages shared between each other and, more critically, identify malicious vessels' behaviors.

Keywords—Automatic Identification System, machine learning, behavioral model, anomaly detection, Ocean of Things

I. INTRODUCTION

The maritime industry is looking forward to the future Smart Ocean to provide reduced operating costs while simultaneously increasing crew safety. The Smart Ocean will consist of a large number of connected devices comprising the Ocean of Things (OoT), a subcategory of the Internet of Things (IoT) [1], [2] (Fig. 1). While IoT devices are always on and connected to the internet, OoT devices operate both online and offline on vessels at sea.

Vessels at sea operate in a peer to peer (p2p) manner and broadcast directly with each other, with no intermediary by which communication occurs. The primary maritime communication protocol is the Automatic Identification System (AIS), which allows vessels to share their location along with many other bits of information in real-time. AIS distinguishes between vessels, aids to navigation (ATON), base stations, and search and rescue transmitter (SART). Each type of transmitter has different privileges and priorities through a self-organizing time division multiple access (SO-TDMA) system [3].

The 2002 International Maritime Organization (IMO) Safety of Life at Sea (SOLAS) requires vessels over 300 gross tonnages to be equipped with AIS [3]. It is estimated that there are over 400,000 AIS installments and up to 1,000,000 once



Figure 1: Ocean of Things (OoT)

fully deployed globally. The SOLAS requirement accounts for the widespread adoption of AIS in the maritime industry.

AIS is a plaintext protocol where vessels communicate without encryption or message authentication. Any message can be broadcast as if it is from any vessel. Within this type of communication, one can not be certain of the validity of any message nor the identity of the sender of any message.

A secure version of AIS exists but has not seen widespread adoption [4], [5]. Some of the factors in the delay of adoption are the need for an international consensus on the implementation and distribution of encryption keys globally, therefore, there is a critical need to model the behavior of the vessels at sea to add a layer of awareness in the system in which vessels can be identified by their behavior, thus helps in identifying malicious actors and behaviors.

To address the aforementioned strategy, we present an anomaly detection scheme using machine learning to model the vessel's normal behavior. We design multiple use cases to test our method using different machine learning modeling techniques to select the highest performing method by analyzing the results of each use case.

The remainder of our paper is organized as follows: We present the related work in Section II, followed by our motivations and contributions in Section III. We present the behavioral model anomaly detection methodology in Section IV, followed by our performance evaluation in Section V, including the threat model, the experiment design, the analysis

methodology and the numerical results. Finally, we conclude the paper and discuss the future work in Section VI.

II. RELATED WORK

In this section, we present current work in the maritime domain using machine learning in AIS as well as the current state of AIS research as a protocol.

A. Machine Learning

Sidibe et al. [6] survey techniques to identify anomalous behavior in the maritime domain using AIS. They categorize the detection methods based on three categories: Statistical, machine learning, and data mining. Within these categories two types of data are examined; location based and data driven approaches.

1) *Location Based:* Vessel location based approaches examine the vessels' current location and movement or trajectory.

Liang et al. in [7] proposed a two-step Long Short-Term Memory (LSTM) supervised learning method to reconstruct a vessel's trajectory when AIS location data is lost. AIS allots 4,500-time slots per minute in a highly congested region. An AIS transceiver can become starved for resources due to a lack of available time slots to transmit on. When this occurs, missing AIS data creates a gap in information for the location of a vessel. Missing AIS data can also happen in inclement weather. As the signal drops, the information is lost after transmission. This allows those monitoring a vessel's movement to project more accurately the ship's prior location to better understand the ship's previous and future movements.

Anneken et al. in [8] used Gaussian Mixture Model (GMM) and Kernel Density Estimator (KDE) to predict anomalies incurred a high rate of false alarms. Gaussian Process and Active Learning were used, but at the cost of high computational complexity in training models. Bayesian Networks have been trained to account for AIS data combined with real-world data such as weather and time with vessel interactions.

Another work by Pallotta et al. in [9] used Point-based anomalous behavior and Trajectory-Based anomalous behavior detection approaches. These two methods focus on the location of the vessel's travel: either where the vessel is currently located or the trajectory of the vessel's location.

Sidibe et al. in [6] noted that anomalous vessel behavior detection causes a high rate of false-positive anomalies detection. Data Mining methods seek to improve upon the high false-positive rates of trajectory and point-based methods.

2) *Data Driven:* Are created using a two-phase method. First, a vessel's normal behavior is modeled based on historical data. Second, the learned model is applied to current vessel movement data with any differentiation considered anomalous behavior.

Osekowska et al. proposed one such approach in [10] by developing and modeling traffic as a potential field for the geographic tracks that a vessel moves through at sea. The field is stronger with greater amounts of vessel traffic and weaker with less traffic. The field has three properties: strength, decay, and distribution. The field strength increases with greater traffic. As fewer vessels traverse a path, the path decays, and the strength value decreases. Distribution is the distance between two points and is described by a two-dimensional Gaussian smoothing, using Euclidean distance

between two points. In this system, a vessel whose current position is detected outside the local potential field is marked as anomalous.

Soleimani et al. in [11] proposed a geometrical method based on the vessel trajectory for the vessel's near-optimal path. A near-optimal path is generated using a graph search algorithm. If a vessel departs from the near-optimal path, then the movement generates an abnormality score.

Roy et al. in [12] generated alerts based on rules in ports of known port parameters, such as the maximum speed allowed in a port and marked restricted areas within a port. If the parameters are broken, then a vessel is marked as anomalous.

B. AIS Attacks

Balduzzi et al. in [13] detailed the various type of AIS attacks and categorized them into two categories: first, implementation-specific in software; second, protocol-specific in the AIS radio transponders. At the software layer, one could spoof another vessel's Maritime Mobile Service Identity (MMSI) and pose as another vessel. Spoofing would make the vessel broadcasting appear to be another vessel along with spoofing the location of the vessel one is broadcasting as. Spoofing as another vessel could also allow one to program a malicious route so that a vessel appears to have taken a false route. Software attacks occur when attacking the application layer based on the applications used by various systems that log AIS messages. An example of this is a port authority. If a port authority logs messages from AIS in a SQL database, one could craft a message to enter into the SQL database executing arbitrary code through AIS.

For radio attacks, one can alter the message broadcast by a physical vessel. This allows one to modify the location in real-time of vessels in transit. A type of attack is a man-in-water spoofing. This is an S.O.S that, once received by nearby vessels, compels them by regulation to attempt a rescue. Simulating an S.O.S would allow an attacker to lure a victim vessel to a hostile location. Closest Point of Approach (CPA) triggers a collision warning alert encouraging a vessel to alter course to avoid a collision. One can spoof a vessel's location so that it appears close to a vessel and that the direction indicates that a collision will occur. This will trigger an alarm on the victim's vessel that a collision is imminent.

Frequency Hopping (DoS++) can occur by an attacker spoofing as a port authority. This forces the vessel's transponder to a non-default frequency and masks the transponder to other vessels operating nearby. This would render a ship invisible to other vessels nearby on AIS.

Slot Starvation (DoS++) occurs when a base station, such as a port authority, exhausts all available slots for message broadcasting. A base station has a high priority compared to vessels. Spoofing, as a base station, one can book the next 100 milliseconds and then another 100 milliseconds continuously, so that all slots are continuously taken, barring any legitimate vessel's messages from being broadcast.

Timing Attack (DoS++) instructs an AIS transponder to delay its transmission for a period in time. One can broadcast continuously, causing an AIS transponder to delay transmission, essentially disabling the transponder continuously. One

can also change the transponder to transmit more often and flood all messages for a given region.

Hardware Panic (DoS) attacks saturate the channel's electromagnetic spectrum with copious quantities of noise. Based on the hardware, malfunctions can occur at the recipient's memory or processor, which could be overloaded.

III. MOTIVATION AND CONTRIBUTION

Recent records showed an increase in the capacity of vessels at sea between the years 2005 and 2015 of 40% within a single decade. This growth occurred, even during a global economic downturn [5]. This increases the demand on the AIS as the communication protocol, which has shown to be susceptible to various attacks. From the literature review we observed that:

- Machine learning has been applied for trajectory and location analysis, which can be expanded and used to study the behavior of vessels at sea through on vessel's real-time anomaly detection.
- Currently, AIS suffers from lack of encryption and authentication. A secure version of AIS exists but has not seen widespread adoption.

In this work we are proposing the use of machine learning to aid in identifying anomalous vessels behavior, which facilitates the identification of suspicious activities.

Our contribution in this work can be summarized as follows:

- Develop a machine learning anomaly detection method for vessels at sea.
- Analyze different machine learning methods to select the best method.
- Design multiple use cases that challenge the behavioral models built for vessels operating at sea.

From the literature review, vessels at sea are susceptible to various attacks via AIS that machine learning can help mitigate to provide a safer operating environment. Currently, ships do not share sensor data, but as autonomous shipping and the OoT increases in scope, vessels can and will be sharing more information [1], [2].

IV. BEHAVIORAL MODEL ANOMALY DETECTION METHODOLOGY

In this work, machine learning is applied to develop behavioral models aiming to aid AIS in identifying behavioral anomalies that could affect communications and decision making in such system. Our approach is carried through a set of consecutive phases as depicted in Fig. 3. AIS reading from vessels is collected to form a consensus about normal vessel behavior, which is fed to our machine learning methods for model training purposes. Once a model is trained, it is distributed to all vessels that request to use the model. In modeling normal ranges, anomalies or modifications are detected, which aids in identifying whether the vessel's sensor is sending false data.

Vessels commonly operate within shipping lanes along well-used vessel routes. Marine Traffic, a company that logs vessel movement via AIS, illustrates this fact with a heat map of vessel AIS locations (Fig. 2). Areas in red are well-used routes by vessels showing that ships share the same routes in the same region. In some cases, vessels will operate outside of conventional lanes for various reasons. This approach builds

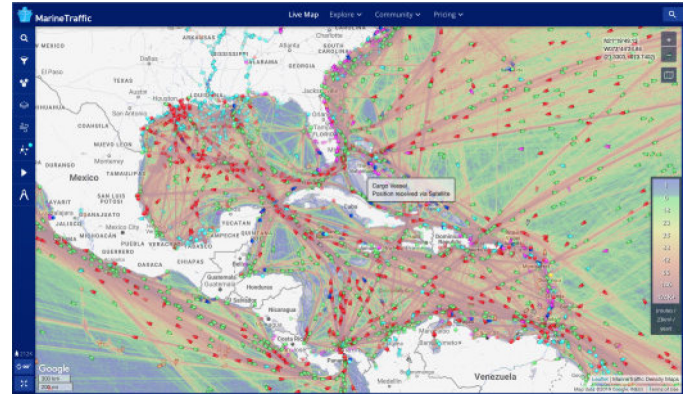


Figure 2: Marine Traffic Density Map adopted from [14]

off the cases in which vessels are within the same area within AIS range, and multiple vessels can communicate with each other. While vessels are within range of each other, AIS observations are recorded by each vessel. As each vessel arrives at a port, they offload the observations to port for consensus building. Figure 3 (top) is a microcosm of the larger case to demonstrate the principles of vessel communication.

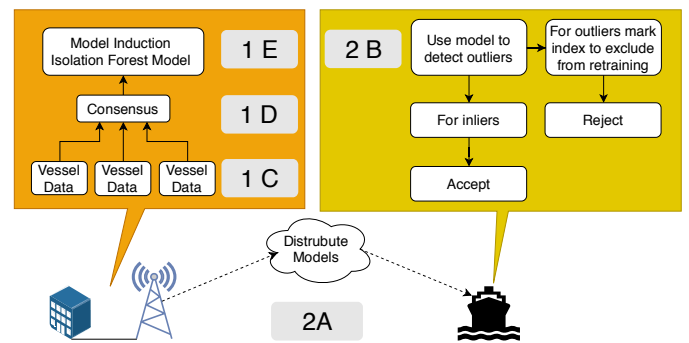



Figure 3: Behavioral Model Anomaly Detection

We will use the illustration presented in Fig. 3 to present our behavioral model anomaly detection approach. Our approach starts with the model induction, where vessel A broadcasts its reading using AIS to the surrounding vessels (1A). Upon reception of new information, each vessel calculates the difference in temperature to determine how different the vessel reporting is from itself (1B). Once a vessel arrives at shore side, the vessel uploads its observations from while at sea

to shore side(1C). Shoreside processing consists of forming a consensus data set based on multiple vessel observations to create a new dataset for model induction (1D). Using the new consensus data set, a model is fit with the consensus set, and a machine learning model is generated (1E).

Progressive Analysis is the process of an offline vessel analyzing AIS readings in real-time. Once a model is trained at shoreside, the shoreside station sends the fit model to each vessel that requests a model (2A). Every vessel has a unique mode fit to its observations. Once a vessel leaves the range of the shore side station and is operating offline, vessels use the model to analyze received readings to classify readings as inliers or outliers (2B). If a message is an inlier then it is assumed to be generated from the vessel that is claiming to have generated that message and is accepted. If the message is classified as an outlier, it is assumed to be created by another device and rejected.

Model Re-Induction is the process of repeating the induction phase after more observations are recorded. After some time and more vessels encounter each vessel, those readings are reported to shore side to create a stronger consensus.

Vessels collect AIS data to facilitate learning while at sea and record the received values. Once alongshore or at the port, the values are transferred to a shoreside server for processing. With the addition of sharing sensor data, a model of the operation of a vessel can be created to determine each vessel's sensors' normal operating range. At shoreside, collections of all reports on a vessel are weighted together to provide a single truth of how a vessel historically operates. The more reports, the stronger the truth is. Using the weighted vessel reports at shoreside, a machine learning model is trained for each vessel. While a vessel is at the port, the models are transferred to the vessel for offline AIS operation.

V. PERFORMANCE EVALUATION

In analyzing the performance of our methodology against the threat model, we designed use cases representative of each type of threat for testing purposes.

A. Threat Model

Threat 1. Impersonation: AIS plain text messages are susceptible to various attacks [13]. Many of the attacks can be limited by identifying spoofing attacks where an attacker poses as another vessel; Slot Starvation by impersonating a base station when the attacker is not a base station, Frequency Hopping by impersonating as a port authority when the attacker is not a port authority, and Closest Point of Approach from a false collision being triggered by an attacker impersonating as another vessel.

Threat 2. Selective Transmission: When a vessel turns AIS off to conceal a vessel's location for a period of time and then turns it back on again when it is advantageous [6]. This can be dangerous as vessels are operating without broadcasting their location to other vessels nearby.

Threat 3. Model Manipulation: A model becomes susceptible to attack as nefarious actors attempt to manipulate the model used to classify observations onboard a vessel. Model manipulation attacks try to play the system and find locations where classification could be weak, allowing invalid

information not to be classified correctly. One such type of attack is breakout fraud, where an attacker maintains a good behavior for a period of time and then starts injecting invalid information [15].

B. Experiment Design

A python random weather generator was used to create weather samples. The weather generator randomly generates sample weather data for a given position by latitude and longitude for a date and time. By using historical weather measurements for these locations from Dark Sky API, a set of synthetic samples are generated for a location, date, and time. Weather samples of one week are generated, simulating interactions for vessels over six days which, are then used for training a model and one day used for model validation.

A range of dates, along with the number of requested samples, is given for sample generation. Five sets of samples are created for six days of 1000 samples for the two sets, simulating five vessels; the difference between each sample at each index is taken to create a single vessel behavior set. The single vessel behavior set simulates the interaction between two vessels with one set for the vessel receiving samples from another vessel within operating range. Using multiple vessel behavior data sets, three methods are used to determine a fit for behavior modeling. Consensus is performed using three methods; mean, median and max. Mean is the average of all readings for a given time period from all samples collected to create a new synthetic value from a mixture of all the readings. Median selects the middle reading from all the readings for a given time period. Max is taken from the absolute value in max either negative or positive for the largest difference recorded.

C. Model Analysis and Evaluation

Use cases are given chosen to demonstrate model fit and attempts to falsify information. Attempts to falsify AIS occur for many reasons, including spoofing a vessel, generating false readings by error or to degrade another vessel's ratings. Once a valid model is trained, an attacker might attempt to give false readings. The cases below demonstrate the case if an attacker attempts to send false readings at various time frames to demonstrate how the machine learning models would classify those readings.

We are considering the following set of factors as our preforming evaluation metrics:

- **True outliers:** are observations where the model detects behaviors as being an anomaly in the original data set.
- **False outliers:** are observations that the model identifies as an outlier but are in the original data set.
- **True inliers:** are observations the model identifies as inliers and that are in the original data set.
- **False inliers:** are observations the model identifies as inliers and they are not in the original data set.
- **Model accuracy:** is calculated as the total number of correct identifications over the total number of observations.

D. Numerical Analysis

In this subsection, we present performance evaluations for each model using a set of use cases. The results presented

in this section are based on designed use cases, where each use case is tested using 100 samples with various anomalies inserted at different locations during each test. The temperature readings are in Celsius and present the difference in readings between the broadcasting vessel and the local vessel.

Four models were used through our evaluation; Isolation Forest, Support Vector Machine, Local Outlier Factor, and Robust Covariance Elliptic Envelop.

Use Case: Errors–Large Uniform: This checks whether a model correctly classifies errors outside of the set and the model does not determine those readings to be inlier observations, even if those errors appear on a consistent regular basis.

Table I: Errors: Large Uniform Summary

| Model | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|-------------------|--------------|---------------|-------------|--------------|---------------|
| Isolation Forest | 10 | 8 | 82 | 0 | 0.92 |
| SVM | 10 | 23 | 67 | 0 | 0.77 |
| LOF | 10 | 18 | 72 | 0 | 0.82 |
| Elliptic Envelope | 10 | 24 | 66 | 0 | 0.76 |

(a) Mean


| Model | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|-------------------|--------------|---------------|-------------|--------------|---------------|
| Isolation Forest | 10 | 3 | 87 | 0 | 0.97 |
| SVM | 10 | 22 | 68 | 0 | 0.78 |
| LOF | 10 | 10 | 80 | 0 | 0.9 |
| Elliptic Envelope | 10 | 23 | 67 | 0 | 0.77 |

(b) Median

| Model | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|-------------------|--------------|---------------|-------------|--------------|---------------|
| Isolation Forest | 10 | 0 | 90 | 0 | 1 |
| SVM | 10 | 35 | 55 | 0 | 0.65 |
| LOF | 0 | 1 | 89 | 10 | 0.89 |
| Elliptic Envelope | 10 | 3 | 87 | 0 | 0.97 |

(c) Max


For large uniform anomalies, Table I shows the mean, median, and max consensus fit test results. It can be seen that the highest attained true accuracy is from an isolation forest using a max consensus at 100%.



| Name | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|--------|--------------|---------------|-------------|--------------|---------------|
| Mean | 10 | 8 | 82 | 0 | 0.92 |
| Median | 10 | 3 | 87 | 0 | 0.97 |
| Max | 10 | 0 | 90 | 0 | 1 |

Figure 4: Isolation Forest – Large Uniform

Figure 4 plots the large uniform case using an isolation forest for mean, median, and max, along with the numerical analysis of each case with the highest true accuracy attained using the max consensus at 100%.




| Name | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|--------|--------------|---------------|-------------|--------------|---------------|
| Mean | 10 | 23 | 67 | 0 | 0.77 |
| Median | 10 | 22 | 68 | 0 | 0.78 |
| Max | 10 | 35 | 55 | 0 | 0.65 |

Figure 5: Support Vector Machine – Large Uniform

Figure 5 plots the large uniform case for a support vector

machine using mean, median, and max, along with the numerical analysis of each case with the highest true accuracy attained using the median consensus at 78%.


Figure 6 plots the large uniform case using the local outlier factor using the mean, median, and max, along with the numerical analysis of each case with the highest true accuracy attained using the median consensus at 90%.



| Name | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|--------|--------------|---------------|-------------|--------------|---------------|
| Mean | 10 | 18 | 72 | 0 | 0.82 |
| Median | 10 | 10 | 80 | 0 | 0.9 |
| Max | 0 | 1 | 89 | 10 | 0.89 |

Figure 6: Local Outlier Factor – Large Uniform

Figure 7 plots the large uniform case using an robust covariance elliptic envelope for mean, median, and max, along with the numerical analysis of each case with the highest true accuracy attained using the median consensus at 97%.



| Name | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|--------|--------------|---------------|-------------|--------------|---------------|
| Mean | 10 | 24 | 66 | 0 | 0.76 |
| Median | 10 | 23 | 67 | 0 | 0.77 |
| Max | 10 | 3 | 87 | 0 | 0.97 |

Figure 7: Robust Covariance Elliptic Envelope – Large Uniform

Use Case–Significant Errors: This case demonstrates the model fit for large values outside of the training set to see whether the model can accurately classify them as outliers. This might be the case in an on-off attack where a vessel shuts off its AIS. If a vessel does not receive a reading from another vessel, then the difference would be significant compared to previous readings. This would indicate that the vessel is not sending accurate readings or potentially no readings at all.

Table II: Errors: Significant Errors Summary

| Model | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|-------------------|--------------|---------------|-------------|--------------|---------------|
| Isolation Forest | 11 | 5 | 82 | 2 | 0.93 |
| SVM | 11 | 19 | 68 | 2 | 0.79 |
| LOF | 11 | 16 | 71 | 2 | 0.82 |
| Elliptic Envelope | 11 | 21 | 66 | 2 | 0.77 |

(a) Mean

| Model | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|-------------------|--------------|---------------|-------------|--------------|---------------|
| Isolation Forest | 10 | 2 | 85 | 3 | 0.95 |
| SVM | 11 | 19 | 68 | 2 | 0.79 |
| LOF | 11 | 7 | 80 | 2 | 0.91 |
| Elliptic Envelope | 11 | 19 | 68 | 2 | 0.79 |

(b) Median

| Model | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|-------------------|--------------|---------------|-------------|--------------|---------------|
| Isolation Forest | 10 | 0 | 87 | 3 | 0.97 |
| SVM | 10 | 31 | 56 | 3 | 0.66 |
| LOF | 2 | 1 | 86 | 11 | 0.88 |
| Elliptic Envelope | 10 | 2 | 85 | 3 | 0.95 |

(c) Max

Table II presents the results of each machine learning model fit to a mean, median, and max consensus set for data with

significant errors, both positive and negative. In this case, the highest performing model with the greatest true accuracy is an isolation forest using the max consensus set at 97%.

Figure 8 shows model classification for significant positive and negative anomalies using an isolation forest, along with the numerical results of each test. The highest attained true accuracy is through the max consensus fit at 97%.

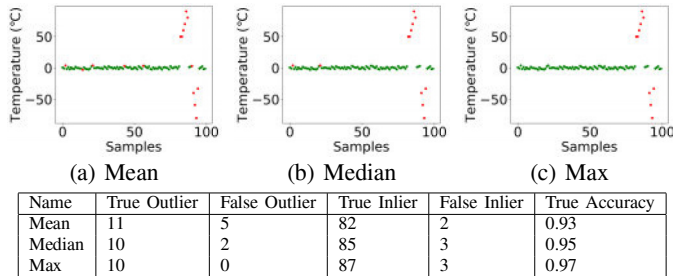


Figure 8: Isolation Forest – Significant Errors

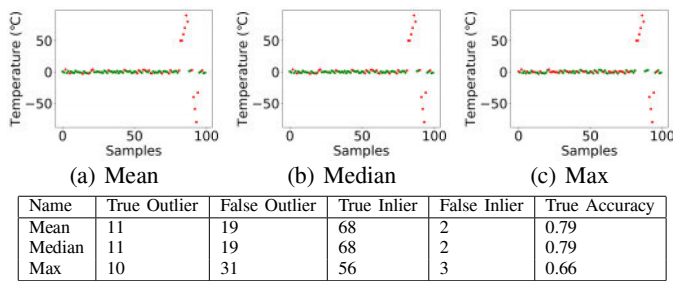


Figure 9: Support Vector Machine – Significant Errors

Figure 9 shows classification for significant positive and negative anomalies using a support vector machine, along with the results of each test. The highest attained true accuracy is through the mean and median consensus fit at 79%.

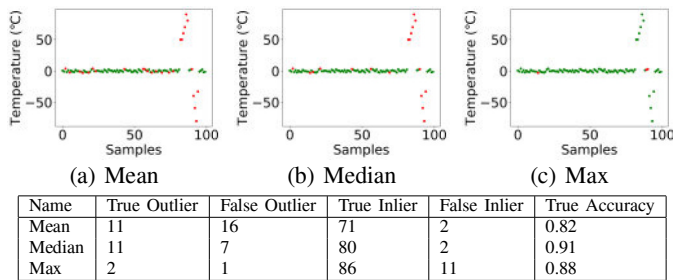


Figure 10: Local Outlier Factor – Significant Errors

Figure 10 shows model classification for significant positive and negative anomalies using the local outlier factor, along with the numerical results of each test. The highest attained true accuracy is through the median consensus fit at 91%.

Figure 11 shows model classification for significant positive and negative anomalies using a robust covariance elliptic envelope, along with the numerical results of each test. The highest attained true accuracy is through the max consensus fit at 95%.

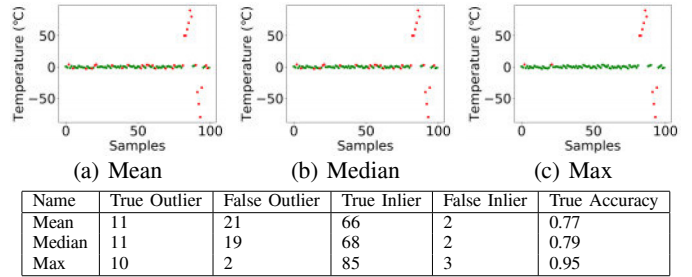


Figure 11: Robust Covariance Elliptic Envelope – Significant Errors

Use Case: Errors–Breakout Fraud: This case tests the model against the case where one might begin spoofing as a user within the range of the original vessel reading but tries to push the readings to a new normal outside of the vessel model.

Table III: Errors: Breakout Fraud Summary

| Model | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|-------------------|--------------|---------------|-------------|--------------|---------------|
| Isolation Forest | 12 | 5 | 74 | 9 | 0.86 |
| SVM | 17 | 17 | 62 | 4 | 0.79 |
| LOF | 14 | 16 | 63 | 7 | 0.77 |
| Elliptic Envelope | 17 | 18 | 61 | 4 | 0.78 |

(a) Mean

| Model | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|-------------------|--------------|---------------|-------------|--------------|---------------|
| Isolation Forest | 10 | 2 | 77 | 11 | 0.87 |
| SVM | 16 | 17 | 62 | 5 | 0.78 |
| LOF | 13 | 7 | 72 | 8 | 0.85 |
| Elliptic Envelope | 17 | 17 | 62 | 4 | 0.79 |

(b) Median

| Model | True Outlier | False Outlier | True Inlier | False Inlier | True Accuracy |
|-------------------|--------------|---------------|-------------|--------------|---------------|
| Isolation Forest | 7 | 0 | 79 | 14 | 0.86 |
| SVM | 10 | 30 | 49 | 11 | 0.59 |
| LOF | 1 | 1 | 78 | 20 | 0.79 |
| Elliptic Envelope | 11 | 2 | 77 | 10 | 0.88 |

(c) Max

Table III shows the highest attained true accuracy is through an elliptic envelope using a max consensus model at 88%.

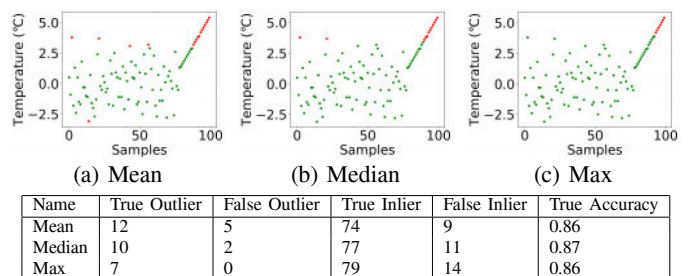


Figure 12: Isolation Forest – Breakout Fraud

Figure 12 illustrates breakout fraud model testing for the mean, median, and max consensus methods using an isolation forest, along with the numerical results of model analysis. The highest attained is by using a median consensus at 87%.

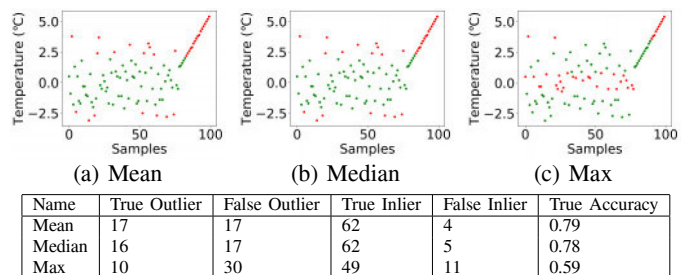


Figure 13: Support Vector Machine – Breakout Fraud

Figure 13 illustrates breakout fraud model testing for the mean, median, and max consensus methods of a support vector machine, along with the numerical results of model analysis. The highest attained is by using a mean consensus at 79%.

Figure 14 illustrates breakout fraud model testing for the mean, median, and max consensus methods using the local outlier factor, along with the numerical results of model analysis. The highest attained is by using a median consensus at 85%.

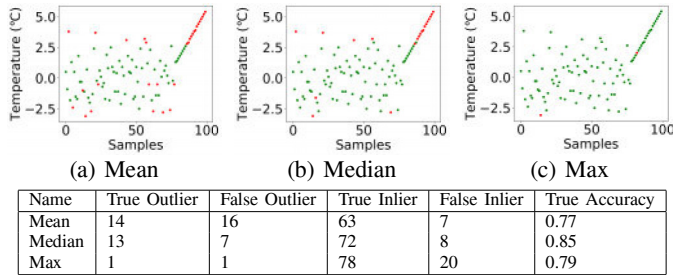


Figure 14: Local Outlier Factor – Breakout Fraud

Figure 15 illustrates breakout fraud model testing for the mean, median, and max consensus methods using the elliptic envelope, along with the numerical results of model analysis. The highest attained is by using a max consensus at 88%.

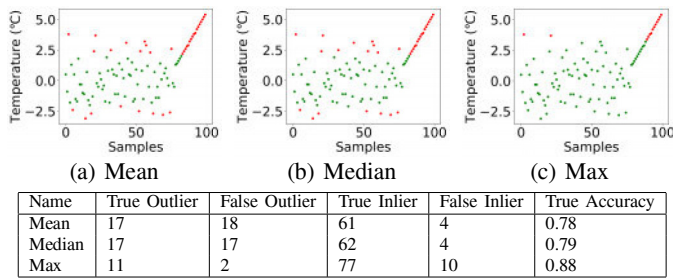


Figure 15: Robust Covariance Elliptic Envelope – Breakout Fraud

VI. CONCLUSION AND FUTURE WORK

The current maritime vessel communication protocol (AIS) lacks encryption and authentication leaving vessel communication susceptible to manipulation. Through the usage of different machine learning models it is possible to model a vessel's normal behavior, thus when abnormal behavior is detected the appropriate mitigation strategy can be applied.

Through the application of anomaly detection it is possible to identify when a maritime vessel is not operating in a manner that is consistent with previous observations. This allows for additional threats to be identified in real-time. Due to the widespread adoption of AIS, any vessel will benefit from implementing behavior modeling even when communicating with other vessels which lack the necessary equipment.

Temperature sensor data transmitted via AIS was used to illustrate the principle of machine learning in order to model vessel behavior in real-time. Future work would consist of also adding additional sensors as features to add degrees of information for the machine learning model. An example of this type of data would be AIS reported position along with a vessel's local radar reading to determine the accuracy of each

vessel's location information. Additional sensors could also increase the accuracy and confidence of a vessel's model.

Future work could also investigate using these machine learning models to influence trust networks built to operate in the maritime domain for vessel to vessel communication. For an example of trust-building, a vessel which broadcasts normally would have a maintained or increased trust rating, while a vessel broadcasting abnormally would have its trust rating decreased.

ACKNOWLEDGEMENT

Research reported in this publication was supported by the 2020 Center of Excellence for Applied Computational Science and Engineering grant competition (CEACSE).

REFERENCES

- [1] Defense Advanced Research Projects Agency. Ocean of things aims to expand maritime awareness across open seas. <https://www.darpa.mil/news-events/2017-12-06>, December 2017. (Accessed on 11/08/2019).
- [2] Yujie Li, Shinya Takahashi, and Seichi Serikawa. Cognitive ocean of things: a comprehensive review and future trends. *Wireless Networks*, Jan 2019.
- [3] International Maritime Organization. Automatic Identification System. <http://www.imo.org/en/OurWork/Safety/Navigation/Pages/AIS.aspx>, 2019. [Online; accessed 21-March-2019].
- [4] Athanasios Goudosis and Sokratis Katsikas. Towards a secure automatic identification system (ais). *Journal of Marine Science and Technology*, 05 2018.
- [5] Pedro Luis Sanchez Gonzalez, David Díaz-Gutiérrez, T.J. Leo, and Luis Núñez. Toward digitalization of maritime transport? *Sensors*, 19(4):926, 02 2019.
- [6] Abdoulaye Sidibé and Gao Shu. Study of automatic anomalous behaviour detection techniques for maritime vessels. In *THE JOURNAL OF NAVIGATION*, 2017.
- [7] M. Liang, R. W. Liu, Q. Zhong, J. Liu, and J. Zhang. Neural network-based automatic reconstruction of missing vessel trajectory data. In *2019 IEEE 4th International Conference on Big Data Analytics (ICBDA)*, pages 426–430, March 2019.
- [8] M. Anneken, Y. Fischer, and J. Beyerer. Evaluation and comparison of anomaly detection algorithms in annotated datasets from the maritime domain. In *2015 SAI Intelligent Systems Conference (IntelliSys)*, pages 169–178, Nov 2015.
- [9] G. Pallotta and A. Joussetme. Data-driven detection and context-based classification of maritime anomalies. In *2015 18th International Conference on Information Fusion (Fusion)*, pages 1152–1159, July 2015.
- [10] Ewa Osekowska, Henric Johnson, and Bengt Carlsson. Grid size optimization for potential field based maritime anomaly detection. *Transportation Research Procedia*, 3:720 – 729, 2014. 17th Meeting of the EURO Working Group on Transportation, EWGT2014, 2-4 July 2014, Sevilla, Spain.
- [11] B. H. Soleimani, E. N. De Souza, C. Hilliard, and S. Matwin. Anomaly detection in maritime data based on geometrical analysis of trajectories. In *2015 18th International Conference on Information Fusion (Fusion)*, pages 1100–1105, July 2015.
- [12] Jean Roy. Anomaly detection in the maritime domain. In Craig S. Halvorson, Daniel Lehtfeld, and Theodore T. Saito, editors, *Optics and Photonics in Global Homeland Security IV*, volume 6945, pages 180 – 193. International Society for Optics and Photonics, SPIE, 2008.
- [13] D. M. Balduzzi. Ais exposed understanding vulnerabilities and attacks. Marine Traffic. Density Map. <https://www.marinetraffic.com/>, 2019. [Online; accessed 19-October-2019].
- [14] F. Kandah, J. Cancellieri, D. Reising, A. Altarawneh, and A. Skjellum. A hardware-software codesign approach to identity, trust, and resilience for iot/cps at scale. In *2019 International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, pages 1125–1134, July 2019.