Real-time Detection of Spoofed AIS: Enhancing Maritime Surveillance Against Induced Noise*

Alexandru POHONTU[1]*, Vasile ERMOLAI[2], Constantin VERTAN[1], Andreea POHONTU[3]

[1] Department of Applied Electronics and Information Engineering, National University of Science and Technology Politehnica Bucharest, 1-3 Iuliu Maniu Blvd., Bucharest, 061071, Romania
alexandru.pohontu@stud.etti.upb.ro (**Corresponding author*), constantin.vertan@upb.ro

[2] Department of Machine Manufacturing Technology, "Gheorghe Asachi" Technical University, 59A, Prof. Dr. Doc. Dimitrie Mangeron St., Iasi, 700050, Romania
vasile.ermolai@academic.tuiasi.ro

[3] Department of Electrical Engineering Sciences, Maritime University, 104 Mircea cel Batran St., Constanta, 900663, Romania
anda.chelariu@yahoo.com

Abstract: The Automatic Identification System (AIS) is vital for maritime safety and surveillance but it remains vulnerable to cyberattacks such as spoofing. In a previous work, a methodology was proposed for detecting falsified AIS data by analysing the stochastic properties of vessel trajectories, which achieved a classification accuracy of over 98% for standard spoofed datasets. However, in anticipation of future spoofing techniques, tests using injected noise in order to mimic the randomness of genuine AIS signals revealed that the initial method's accuracy declined to nearly 50%. This study proposes an enhanced approach that addresses this limitation by training on datasets where Gaussian noise was added to spoofed vessel coordinates. The noise parameters were statistically derived from genuine maritime trajectories through parametric probability distribution fitting. This aspect allows the employed models to better capture the natural variability of authentic AIS signals. Additionally, unlike the previous offline approach, the improved methodology enables the real-time analysis of vessel tracks rather than relying solely on the post-processing of previously recorded vessel trajectories. To that, two of the proposed deep learning architectures, namely 1D Convolutional Neural Networks (1D CNNs) and Gated Recurrent Units (GRUs) maintained a detection accuracy of 99% even under noisy spoofing conditions, demonstrating a superior robustness and adaptability for real-time maritime surveillance.

Keywords: AIS spoofing, Maritime anomaly detection, Simulated trajectories, Gaussian noise.

1. Introduction

The Automatic Identification System (AIS) is a digital tracking and communication platform developed for improving maritime safety and navigational efficiency. It allows ships to autonomously transmit and receive key data, supporting the continuous, real-time surveillance of maritime activity. AIS is extensively employed by ships and maritime authorities to reduce the risk of collisions and combat illegal activities such as unregulated fishing or smuggling (Androjna et al., 2024).

The AIS operates through the exchange of data between transceivers over the maritime VHF band, utilizing a Time Division Multiple Access (TDMA) protocol (Androjna et al., 2021). The system supports both static data (e.g., ship dimensions, the MMSI identification code and vessel name) and dynamic data (e.g. a ship's position, course over ground and speed), much of which relies on GPS receivers. While AIS signals are generally reliable, they are subject to the typical propagation limits of VHF communications, which means the effective range is generally the line of sight, typically up to 40 nautical miles (Nm) under ideal conditions, but often less depending on antenna height, atmospheric conditions, and interference (Iphar et al., 2015).

The AIS features significant cybersecurity vulnerabilities due to its lack of built-in security features. Notably, AIS lacks authentication, validity, and timing checks (Kessler & Zorri, 2024). Since the messages are neither encrypted nor digitally signed, their integrity cannot be ensured, making them susceptible to intentional tampering or unintentional errors.

Spoofing, a sophisticated form of cyberattack, involves creating entirely fake AIS messages to mislead nearby receivers, and it is often harder to detect than jamming (Sciancalepore et al., 2022). Jamming, on the other hand, disrupts AIS or Global Navigation Satellite System (GNSS) signal reception by interfering with the radio frequency (RF) spectrum, typically through the

* This article represents an extension of the paper "Detection of spoofed AIS: Simulated tracks vs. real maritime data", published in the *Romanian Journal of Information Technology and Automatic Control*, vol. 35(1), 2025, pp. 37-50.

https://doi.org/10.24846/v34i3y202505

emission of strong, overlapping signals that overwhelm legitimate transmissions and prevent receivers from decoding them (Suratman & Mansor, 2024). Additionally, vessels can engage in "dark activities" by simply turning off their AIS transponders to conceal illegal or covert operations (Görkem et al., 2023).

Without geographic or temporal validation, AIS receivers have no means to verify if position data is accurate or current, allowing attackers to manipulate the location, identity, or status of vessels. These flaws make AIS a vulnerable open broadcast system, exploitable for malicious purposes including electronic warfare, smuggling, and maritime hijacking.

One of the most common uses of AIS spoofing is for concealing illegal, unreported, and unregulated fishing, such as in the July 2020 case when a large fishing fleet falsely reported certain positions near New Zealand while likely operating illegally in another zone. Smugglers and sanctions evaders also exploit AIS spoofing to obscure the movement of sanctioned goods, such as the case of an oil tanker that, in December 2021, broadcast a false route to hide its true journey, later verified by satellite imagery (Kessler & Zorri, 2024). Another common use is identity laundering, where vessels assume the AIS identity of scrapped or benign ships to avoid detection. AIS spoofing has also played a role in geopolitical manoeuvring; for instance, in 2021, fake tracks suggested that warships were provocatively close to Crimea. Additionally, spoofing has facilitated illegal ship-to-ship transfers of oil in the Black Sea to circumvent European sanctions, as observed in the extended anchoring of a tanker during 2022 (Androjna & Perkovič, 2024).

AIS spoofing is often performed with specialized Warship AIS (W-AIS) systems, while free tools like the NMEA Simulator enable virtual track generation.

Previously, Pohontu et al. (2025) achieved a 98% detection accuracy on noise-free spoofed AIS tracks, but the accuracy dropped under noisy spoofing conditions. This paper proposes an enhanced model that overcomes these limitations and additionally supports real-time classification.

The remainder of this paper is organized as follows: Section 2 outlines direct strategies for securing AIS and explores indirect spoofing detection techniques discussed in prior research. It also revisits a previously developed method that identifies spoofed AIS tracks by analysing the stochastic patterns of kinematic errors in vessel trajectories. While this method demonstrated a high detection accuracy, additional tests evaluated its robustness when Gaussian noise was artificially added to the data, simulating a likely future tactic in which spoofers introduce noise to evade detection. Section 3 introduces a new, enhanced approach specifically designed for detecting AIS spoofing in real time, even under noisy conditions. Further on, Section 4 presents the experimental setup and results, assessing the performance of the proposed method. Finally, Section 5 concludes this study and outlines potential directions for future research.

2. Related Work

2.1. Measures for Securing AIS

The lack of authentication in AIS broadcasts exposes maritime communications to a variety of spoofing and security threats. To address this vulnerability, the Auth-AIS algorithm is proposed. It represents a secure, flexible, and backward-compatible authentication framework designed specifically for AIS. Auth-AIS relies on well-established cryptographic primitives, integrating the TESLA protocol for a delayed authentication and Bloom Filters for a bandwidth-efficient probabilistic verification. A key innovation of the Auth-AIS algorithm is its standard-compliant, software-only implementation, which makes it deployable on the existing AIS infrastructure without requiring any hardware modifications (Sciancalepore et al., 2022).

Another notable contribution to AIS security is an approach which focuses on authenticating AIS broadcasts by using Public Key Cryptography (PKC) while maintaining full backward compatibility. This method involves digitally signing AIS messages by using the sender's private key, with the signature transmitted in a separate follow-on AIS message, thus preserving the interoperability with existing AIS receivers (Wimpenny et al., 2022).

Although several secure versions of AIS have been proposed, they have not been widely adopted

due to the need for global consensus on how to distribute public encryption keys. Even with immediate approval, AIS spoofing would remain a threat for decades due to the slow global upgrading of vessel systems (Coleman et al., 2020).

2.2. Measures for Detecting AIS Spoofing

Until AIS is upgraded to a secure version, various indirect methods for detecting AIS spoofing, such as analysing vessel behaviour, identifying trajectory inconsistencies or detecting signal anomalies, can be implemented to help mitigate risks, including the disruption of maritime surveillance or the concealment of illegal vessel activities.

One particularly promising approach involves verifying the compliance of AIS messages with the TDMA protocol. This method leverages a real-time, low-cost strategy based on a Kalman filter for tracking vessel dynamics, comparing the reported position and velocity data against the predicted values. Inconsistencies trigger statistical tests that indicate potential falsification. Additionally, the method cross-checks whether AIS transmissions adhere to the expected reporting intervals and time-slot allocations as defined by the TDMA standard. Spoofed or fabricated AIS messages often fail to replicate the complex timing behaviours typical of the TDMA protocol, making this approach highly effective in identifying both falsification and spoofing with minimal false alarms (Louart et al., 2023).

Another effective method for detecting AIS spoofing involves cross-verifying the AIS-reported positions based on independent radar observations. By modelling the spoofing detection problem as a statistical hypothesis test, this approach evaluates whether the incoming AIS data aligns with radar-based estimates of the vessel's position. When discrepancies exceed a defined threshold, the method flags the AIS data as potentially spoofed (Katsilieris et al., 2013). However, a key drawback of this approach is its reliance on radar technology, which may not be available in all monitoring environments, particularly in remote maritime areas.

Complementing these approaches, another recent method focuses on distinguishing simulated AIS tracks from genuine maritime data by analysing stochastic kinematic errors. This technique assumes that authentic vessel trajectories naturally exhibit irregularities due to GPS inaccuracies and environmental influences such as wind and currents, factors which do not appear in mathematically generated spoofed tracks (Pohontu et al., 2025). The study applies a series of calculated parameters, including velocity, bearing, and distance prediction errors, alongside Kalman filter error distributions, to train various machine learning models. These models successfully detect the spoofed tracks based on their unnatural precision and consistency, achieving classification accuracies above 98%. This approach demonstrates the efficacy of using error-based motion analysis and statistical learning for identifying AIS spoofing, especially in environments lacking a high-cost radar infrastructure. Even though this method demonstrated a high performance, further research, presented in the following sections, focused on enhancing its capabilities. First, the robustness of the approach was challenged by a key observation: since spoofed AIS tracks are characterized by low variability with regard to prediction errors, future spoofing techniques are likely to evolve by injecting artificial noise to mimic the randomness of genuine signals. When such noise was added to the spoofed tracks, the performance of the initial model drastically dropped, its accuracy fell from 98% to almost 50%, and it was unable to detect spoofed vessels after noise injection. Second, the model was adapted to enable the real-time classification of AIS trajectories, rather than relying solely on post-processing after all parameters had been computed. Therefore, the subsequent section introduces an enhanced version of the original algorithm, designed to maintain a high detection performance even when facing more sophisticated spoofing attempts that incorporate induced noise. This improved solution also supports real-time operation, which enables on-the-fly spoofing detection and makes it more suitable for deployment in live monitoring systems.

3. Research Methodology

In a previous study, it was found that spoofed AIS tracks usually show limited variability because they are generated by computer models. In contrast, genuine vessel trajectories exhibit natural deviations due to GPS imprecision and external influences like wind, currents, or minor course adjustments. The main idea in that study

was that predicting a ship's future locations is affected by two types of errors: measurement errors from GPS and process errors from external influences like the weather and navigation behaviour. Since simulated tracks are often created using mathematical models, they usually don't include these natural variations. By focusing on these differences, the study demonstrated the possibility of differentiating between genuine and fake simulated AIS tracks. Several machine learning models were tested and they achieved an accuracy over 98% in detecting falsified AIS data (Pohontu et al., 2025). Building on that, this new study takes the next step by proposing a method that can handle more advanced spoofing techniques, especially those that might, in the future, add fake noise to make the analysed data appear more realistic.

3.1. Data Collection and Pre-processing

The initial stage of the study focused on collecting real-time tracking data for an interval of 10 days by using several AIS transponders deployed along the Romanian coastline. In parallel, live AIS data feeds from different regions were collected using several Maritime Situational Awareness (MSA) web platforms. Tools like AIS Hub, MarineTraffic, and VesselFinder provided access to raw AIS-NMEA messages via configured APIs. In addition to live feeds, publicly available databases were utilized to supplement the analysed dataset with historical and structured AIS records. For example, the AIS Exploratorium platform contains over 1.2 million raw AIS entries suitable for a large-scale maritime analysis (Anon., n.d.). A processed AIS dataset is also accessible on Kaggle (Jose, n.d.).

To ensure that trajectories were extracted only from genuine vessels, a filtering process was applied for removing all the tracks associated with Maritime Mobile Service Identity (MMSI) codes specific to Aids to Navigation (AtoNs) and other non-navigational entities.

For each recorded MMSI code, as part of the pre-processing phase, several parameters were computed for each segment of k consecutive positions on the recorded vessel trajectories. This analysis focused on evaluating successive pairs of geographic coordinates, namely latitude (ϕ) and longitude (λ), corresponding to the vessel's previous ($k-1$), current (k), and next ($k+1$) positions (Figure 1).

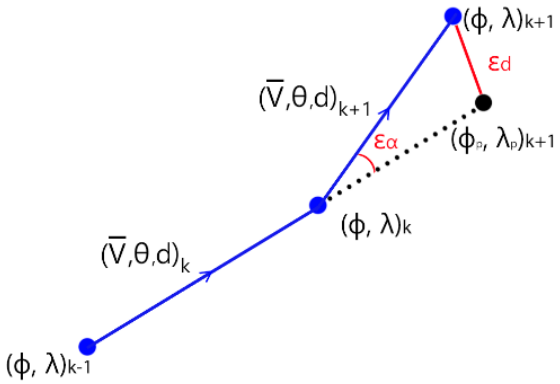

Figure 1. Sequential Plot of Vessel Location Data (Pohontu et al., 2025)

For every k-sampling interval, multiple parameters were calculated, such as velocity magnitudes $|\bar{V}|$, the travel distances denoted by d which is also expressed in equation (1), and the bearing angles Θ between adjacent positions. The values of d and Θ were derived using spherical geometry to compute the distance and direction taking into account two successive geographic points. Using the current position, velocity, and bearing information, the upcoming coordinates $(\varphi_p, \lambda_p)_{k+1}$ were predicted to quantify the discrepancies in latitude denoted by ε_φ in equation (2), the discrepancies in longitude denoted by ε_λ in equation (3), the travelled distance ε_d expressed in (4) and the bearing ε_α as expressed in equation (5) between the predicted and actual future locations:

$$d = 2r \cdot \arcsin\sqrt{\sin^2\left(\frac{\varphi_{k-1}-\varphi_k}{2}\right) + \cos(\varphi_{k-1}) \cdot \cos(\varphi_k) \cdot \sin^2\left(\frac{\lambda_{k-1}-\lambda_k}{2}\right)} \quad (1)$$

$$\varepsilon_\varphi = (\varphi)_{k+1} - (\varphi_p)_{k+1} \quad (2)$$

$$\varepsilon_\lambda = (\lambda)_{k+1} - (\lambda_p)_{k+1} \quad (3)$$

$$\varepsilon_d = d((\varphi,\lambda)_{k+1}, (\varphi_p,\lambda_p)_{k+1}) \quad (4)$$

$$\varepsilon_\alpha = \theta((\varphi,\lambda)_{k-1}, (\varphi,\lambda)_k) - \theta((\varphi,\lambda)_k, (\varphi,\lambda)_{k+1}) \quad (5)$$

where (φ, λ) denotes the latitude and longitude in radians, and r refers to the Earth's radius.

In addition to the above, several differential parameters were computed for each pair of consecutive locations. These included the change in elapsed time between successive AIS reports which is denoted by Δt_{k+1} in equation (6), the change in vessel velocity Δv_{k+1} as expressed in equation (7), and the variations in course represented by $\Delta\theta_{k+1}$ in equation (8). Furthermore,

spatial vector differences were calculated, including $\Delta\vec{x}$ in equation (9) for changes in longitudinal position and $\Delta\vec{y}$ in equation (10) for changes in latitudinal position:

$$\Delta t_{k+1} = t_{(k+1,k)} - t_{(k,k-1)} \quad (6)$$

$$\Delta v_{k+1} = v_{(k+1,k)} - v_{(k,k-1)} \quad (7)$$

$$\Delta \theta_{k+1} = \theta_{(k+1,k)} - \theta_{(k,k-1)} \quad (8)$$

$$\Delta \overrightarrow{x_{k+1}} = v_{k+1} \cdot \cos(\theta_{k+1}) - v_k \cdot \cos(\theta_k) \quad (9)$$

$$\Delta \overrightarrow{y_{k+1}} = v_{k+1} \cdot \sin(\theta_{k+1}) - v_k \cdot \sin(\theta_k) \quad (10)$$

3.2. Data Analysis and Parametric Probability Distribution Fitting

Following the pre-processing phase, multiple vessel trajectories were selected and analysed based on both the ship type (e.g. fishing, cargo, towing ship) and navigational status (e.g. at anchor, docked, drifting, navigating). The analysis covered a wide range of operational conditions, with vessel speeds ranging from 0 up to 20 knots (kt). For each selected vessel, the mean values (μ) and standard deviations (σ) were calculated for all sequential parameters defined in subsection 3.1, including latitude ε_φ and longitude ε_λ prediction errors, the variations in course denoted by $\Delta\theta$ and in speed denoted by Δv, and the changes in AIS reporting intervals denoted by Δt.

Subsequently, a statistical evaluation was performed to determine the most appropriate Probability Density Functions (PDFs) that accurately model the empirical distributions of these parameters. The recorded data was compared against a suite of theoretical distributions, including Gaussian, Gamma, Rayleigh, Log-normal, and Weibull distribution models. The objective was to identify the "*best fit*" distributions that most accurately captured the statistical behaviour of each variable. To achieve this, the study employed the Fitter library developed in Python, which enabled the quantitative comparison of empirical data with data related to theoretical models. The best-fitting distributions were identified based on the Sum of Squared Errors (SSE) index in equation (11):

$$SSE = \sum_{i=1}^{n} (e_{x_i} - p_{x_i})^2 \quad (11)$$

where n represents the number of records, e_{x_i} is the predicted density value based on empirical probability for all x_i and p_{x_i} is the predicted density value based on the parametrical distribution fit.

3.3. Spoofed Data Generation and Noise Addition

After analysing the genuine vessel trajectories, multiple simulated tracks were algorithmically generated to replicate typical navigation patterns. Each spoofed trajectory consisted of a predefined number of waypoints, with associated parameters such as speed and course, to mimic realistic movement patterns. These tracks were designed to continuously report positional data over intervals ranging from 3 minutes to up to 24 hours.

To emulate diverse navigational scenarios, the simulated tracks were configured with different levels of movement complexity. Some vessels maintained a constant speed and bearing over a single waypoint segment, while others featured up to 5 distinct transitions in speed and direction.

The time intervals between successive AIS reports were dynamically determined based on the vessel's velocity, in accordance with the AIS TDMA reporting protocol. For instance, the vessels marked as "underway" in the simulation transmitted position updates every 2 to 10 seconds, depending on their speed. Conversely, stationary vessels were configured to report their positions at a fixed interval of 3 minutes, mirroring the standard AIS behaviour for anchored or moored ships.

To replicate real-world signal propagation issues, a subset of messages, ranging randomly from 0% to 10% of the total number of waypoints per track, was deliberately discarded. The aim was to simulate the effect of meteorological or environmental phenomena that can occasionally disrupt VHF communication and lead to an intermittent AIS message loss.

Furthermore, based on a Gaussian distribution function F as expressed in equation (12), noise $N(\mu, \sigma^2)$ was added to the latitude and longitude coordinates of each reported position. The induced noise was characterized by a mean value of $\mu = 0°$ and a standard deviation $\mu = 1$. The choice of Gaussian noise to perturb the simulated spoofed tracks was motivated by the empirical analysis of

genuine AIS trajectories presented in the previous subsection. When statistically characterizing these deviations, several parameters exhibited distributions closely aligned with the Gaussian distribution, particularly the latitude and longitude prediction errors (subsection 4.3).

$$F(\mathrm{x}) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(\mathrm{x}-\mu)^2}{2\sigma^2}} \tag{12}$$

3.4. Machine Learning Implementations

Following the extraction of sequential kinematic parameters described in subsection 3.1, each vessel trajectory was transformed into a multivariate time-series signal representing its dynamic movement behaviour. To ensure consistency across different samples, only trajectories containing at least 32 AIS recordings were retained. These sequences were pre-processed using a custom mirror padding strategy, which symmetrically reflects edge values to extend each signal to a fixed length of 4096 timesteps.

The input to the model consisted of normalized time-series vectors comprising features such as latitude and longitude prediction errors, speed and course variations, and changes in AIS reporting intervals. Each sample was scaled using a MinMax normalization technique to ensure consistent feature ranges. The output was a binary label: 1 for genuine AIS trajectories and 0 for spoofed ones.

The dataset used for the training and evaluation phases comprised approximately 5,000 vessel trajectories, evenly balanced between genuine AIS tracks and simulated spoofed tracks. Genuine data was collected from genuine maritime traffic using AIS transponders and online platforms, while spoofed data was programmatically generated to replicate plausible movement patterns.

The primary deep learning architecture implemented was a 1D Convolutional Neural Network (1D CNN) (Kiranyaz et al., 2021). The model architecture included three convolutional blocks with increasing filter sizes (32, 64 and 128), each followed by batch normalization and max-pooling layers. These were followed by a flattening layer and three fully connected dense layers (1024, 254 and 128 neurons) with ReLU activations and dropout regularization (with a rate of 0.2). The final layer included a sigmoid activation for binary classification.

The model was compiled based on the Adam optimizer (with a learning rate of 0.001) and trained using binary cross-entropy as the loss function. Training was performed for 250 epochs with a batch size of 32. A *ReduceLROnPlateau* callback was used for dynamically adjusting the learning rate based on validation loss, with a minimum threshold of 0.0001 and a patience of 5 epochs. The dataset was split into 65% training and 35% testing using stratified sampling to preserve class balance.

In addition to 1D-CNN, several other machine learning architectures were employed for benchmarking purposes, including Long Short-Term Memory networks (LSTM), Gated Recurrent Units (GRU), Multi-Layer Perceptrons (MLP), K-Nearest Neighbors (KNN), and Support Vector Machines (SVM).

4. Results

4.1. General Aspects

Each day, around 1000 unique vessel trajectories were retrieved from the AIS data, encompassing both stationary vessels and those in active motion within the observed area. To exclude AtoNs and ensure data quality, only MMSI values within the range of 201000000 to 775999999 were retained for analysis.

As observed in a previous study, stationary moored vessels exhibit slight positional fluctuations over time, primarily due to inherent GPS inaccuracies and minor environmental influences (Pohontu et al., 2025). This natural variation is visible in Figure 2(a). In contrast, spoofed stationary vessels typically report fixed coordinates with little to no variation, as they lack the stochastic errors present in genuine AIS data (Figure 2(b)). To overcome this limitation and generate more realistic spoofed scenarios, Gaussian noise was added to the simulated tracks. As illustrated in Figure 2(c), the resulting noise-injected trajectories more closely mimic the behaviour of genuine stationary vessels, thereby increasing the difficulty of spoofing detection.

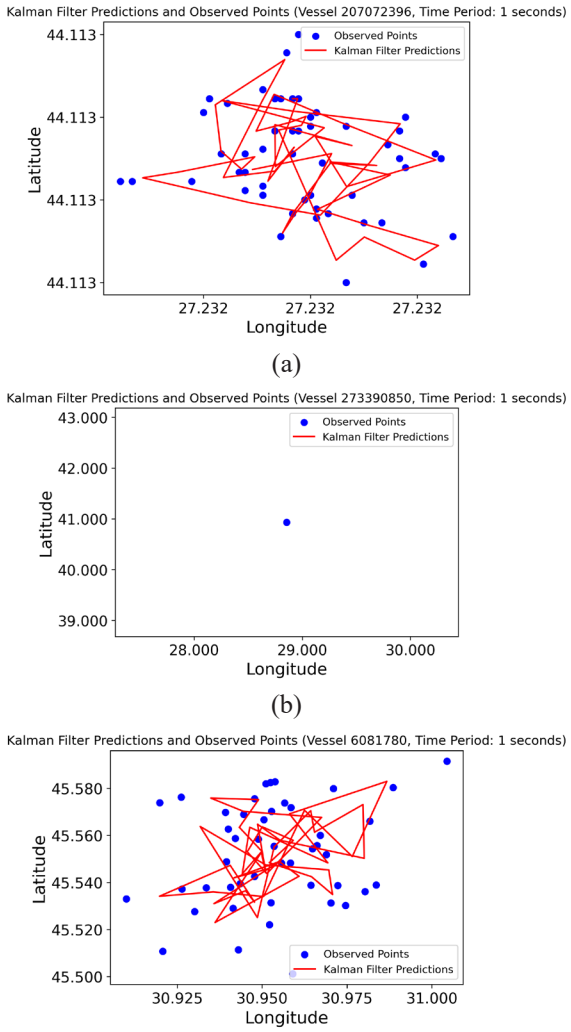

Figure 2. Stationary vessels: (a) Genuine; (b) Simulated and (c) Simulated with Noise

Further on, Figure 3(a) illustrates the trajectory of a genuine underway vessel, which is naturally affected by both measurement and process errors. Measurement errors mainly stem from imprecisions in Global Navigation Satellite Systems (GNSSs), whereas process errors are caused by dynamic influences like wind, ocean currents, weather variations, and slight deviations in a vessel's steering. In contrast, Figure 3(b) depicts a simulated vessel trajectory generated using mathematical models. These tracks lack the stochastic variability typical of genuine AIS data, resulting in overly smooth and predictable motion profiles when forecasting future positions, velocities, or courses. To address this limitation, Figure 3(c) shows a simulated vessel with Gaussian noise injected into its positional data to emulate the natural randomness observed in genuine vessel behaviour. Although the injected noise was sufficient to deceive previously state-of-the-art spoof detection algorithms (Pohontu et al., 2025), further refinements can potentially be implemented.

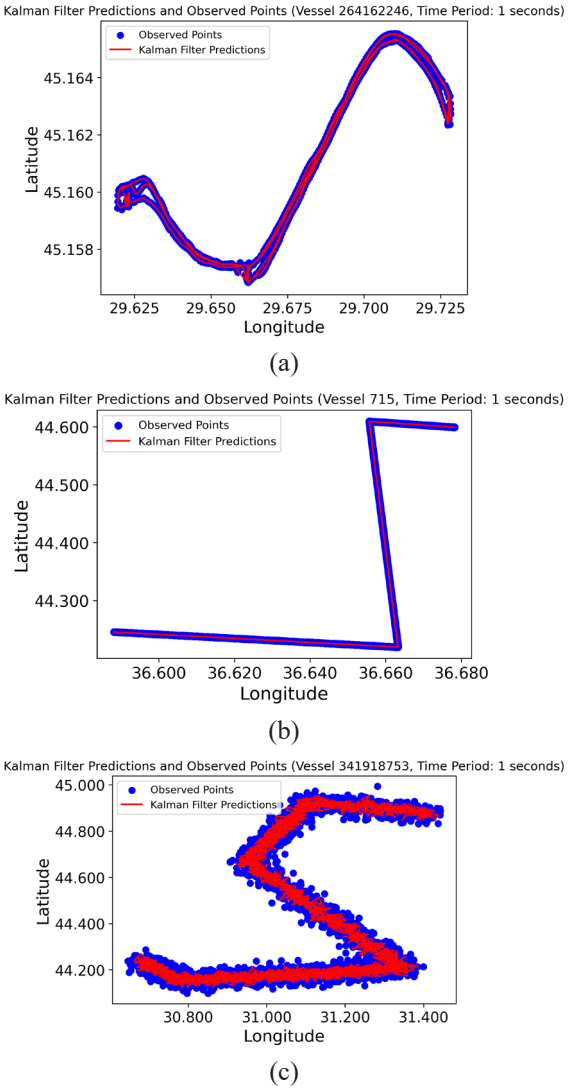

Figure 3. Underway vessels: (a) Genuine; (b) Simulated and (c) Simulated with Noise

To sum up, despite the added variability, these tracks still display artificially linear patterns and lack the irregular curves and turning behaviour typical of authentic vessel movements. A more realistic perspective could be reached by introducing variability in AIS reporting intervals, mimicking the TDMA protocol, and by simulating the effects of meteorological conditions on AIS signal propagation.

Figures 4 and 5 illustrate the temporal variation of latitude prediction errors and Course Over Ground (COG) changes for consecutive locations for a set of vessels. As shown in Figure 4(a), the genuine vessel displays significant and persistent

error variations throughout its movement phase. Notably, approximately after the index 2000, the magnitude of these variations decreases substantially, indicating that the vessel likely became stationary (e.g. it was moored or anchored) and continued transmitting its position without the same dynamic influences. In contrast, the simulated track without noise (Figure 4(b)) exhibits almost no variation in the latitude prediction error, except during abrupt changes in the vessel's course, further emphasizing the deterministic nature of mathematically generated trajectories. Figure 4(c) shows the latitude error profile for a spoofed vessel with Gaussian noise. Further on, in terms of COG changes between successive locations, both genuine vessels and spoofed vessels, in the context of a simulation with added Gaussian noise, exhibited similar variation patterns. A comparable convergence was also observed with regard to the variations in AIS reporting intervals, longitude prediction errors, travelled distance errors, or speed variations, suggesting that noise injection significantly enhances the realism and deception potential of spoofed AIS behaviour. The observed fluctuations closely resemble those related to the genuine trajectory, making such tracks more challenging to distinguish when using conventional detection methods.

4.2. Parametric Results

To enable the injection of Gaussian noise during the spoofed vessel data generation stage (subsection 3.3), it was first necessary to fit the empirical distributions associated with genuine AIS parameters, in order to statistically characterize authentic maritime movement patterns. This was achieved by analysing the sequential parameters computed based on the genuine AIS data (e.g., prediction errors regarding latitude and longitude, variations in speed and

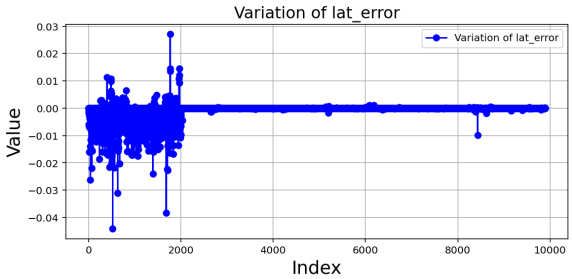

(a)

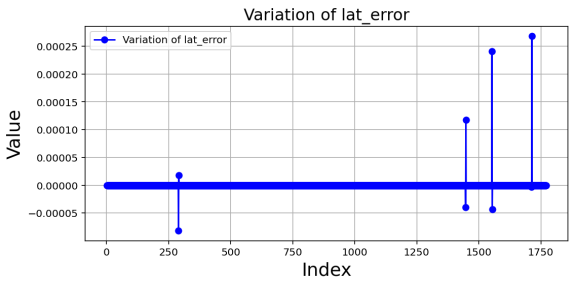

(b)

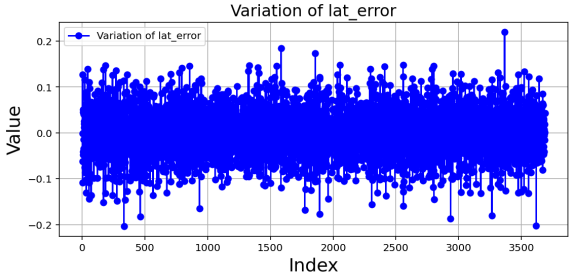

(c)

Figure 4. Latitude prediction error variation in predicting future locations: (a) Genuine; (b) Simulated and (c) Simulated with Noise

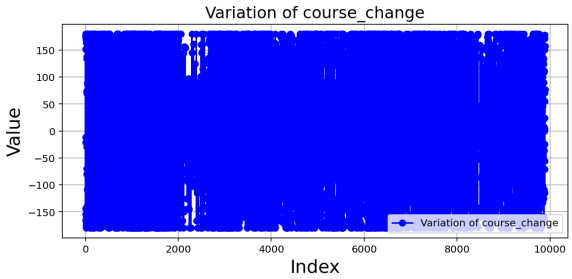

(a)

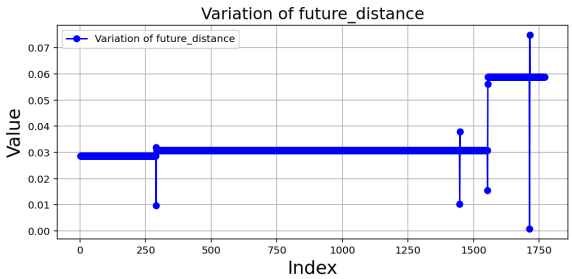

(b)

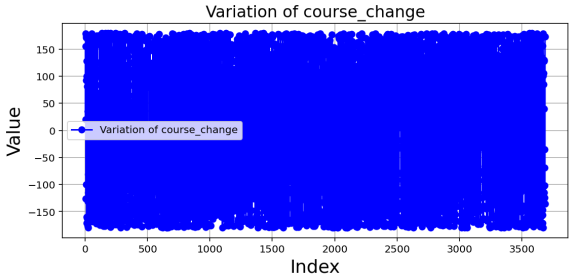

(c)

Figure 5. COG variation in predicting future locations: (a) Genuine (A); (b) Simulated and (c) Simulated with Noise

course, and changes in AIS reporting intervals). These parameters were then fitted to standard parametric probability distributions in order to to capture their underlying statistical behaviour.

As illustrated in Figure 6, histograms with Kernel Density Estimation (KDE) plots were used to visualize the empirical distribution of key features, revealing a clear variability in parameters like latitude prediction error, vessel speed change, and AIS reporting interval variation.

The results for parametric distribution fitting, summarized in Figure 7 and Table 1, show that each parameter aligns with a parametric distribution. For example, speed change follows a Wrapped Cauchy distribution, the COG change fits a Normal-Inverse Gaussian distribution, and both latitude and longitude prediction errors conform to Exponential Normal and Gaussian distributions.

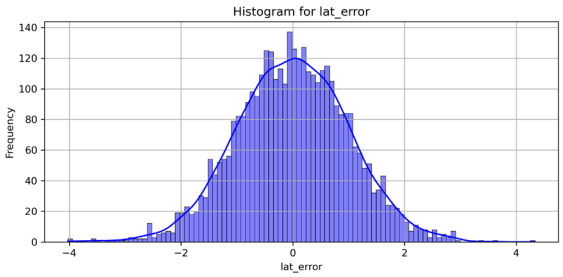

(a)

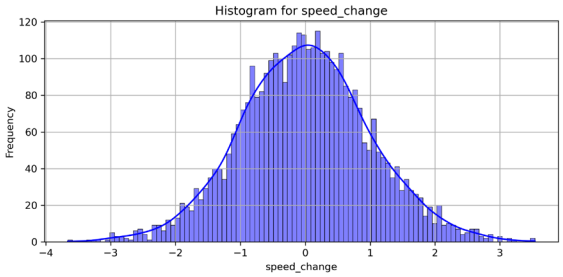

(b)

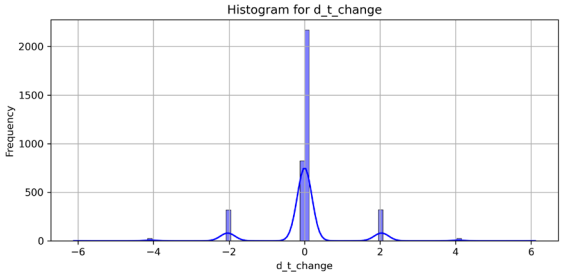

(c)

Figure 6. Histograms with KDE plots for sequential parameters: (a) Latitude prediction error variation; (b) Speed change and (c) AIS reporting interval change

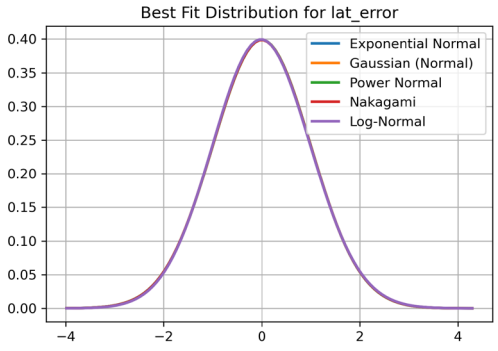

(a)

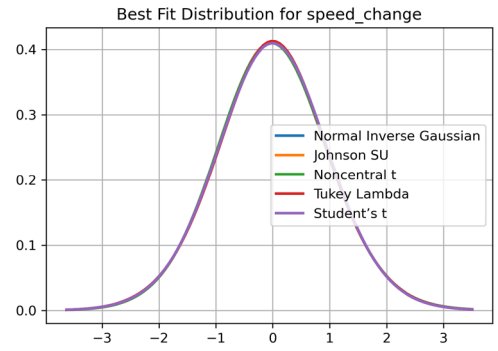

(b)

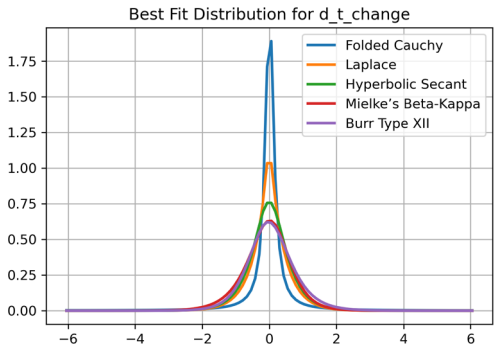

(c)

Figure 7. Parametric distributions fitting: (a) Latitude prediction error variation; (b) Speed change and (c) AIS reporting interval change

Table 1. Best fitted distribution for the sequential parameters

Parameter	Best fitted distribution
Speed change (speed_change)	Normal Inverse Gaussian
COG change	Normal-Inverse Gaussian
AIS reporting interval change (d_t_change)	Folded Cauchy
Error in position	Exponentiated Weibull
Error in total travelled distance	Crystal Ball
Error in predicting latitude (lat_error)	Exponential Normal and Gaussian (Normal)
Error in predicting longitude	Exponential Normal and Gaussian
Vector x change	Generalized Normal
Vector y change	Generalized Normal

4.3. Analysing Machine Learning Models' Effectiveness

To assess the effectiveness of the implemented machine learning models, two categories of datasets were employed during training and evaluation. The former consisted of AIS trajectories collected from genuine maritime traffic, reflecting natural movement patterns and inherent GPS inaccuracies. The latter included artificially generated spoofed trajectories, which were further divided into two subsets: noise-free tracks and tracks enhanced with Gaussian noise to replicate realistic stochastic variability. This comprehensive dataset design allowed the testing of the employed models not only under standard spoofing scenarios but also against more sophisticated attacks that attempt to mimic the randomness of authentic signals.

As shown in Table 2, the best overall classification performances were achieved by the 1D Convolutional Neural Network (1D CNN) and the Gated Recurrent Unit (GRU) models, both reaching a global accuracy of 99%. These models demonstrated a superior ability to learn and generalize based on the subtle spatial and temporal dynamics typical of the AIS time-series data. The LSTM and MLP models also performed strongly, each reaching an accuracy of 97% accuracy, while the K-Nearest Neighbors (KNN) and the Support Vector Machine (SVM) algorithms achieved an accuracy of 98% and 92% respectively. These results underline the advantage of deep learning approaches, particularly the convolutional and recurrent architectures, in identifying spoofed maritime tracks, even under adversarial conditions involving noise injection.

Table 2. Accuracy of the implemented ML models

Model	Accuracy
1D CNN	0.99
LSTM	0.97
GRU	0.99
MLP	0.97
SVM	0.92
KNN	0.98

5. Conclusions

The AIS remains a cornerstone of maritime safety, navigation, and surveillance, providing real-time vessel tracking and facilitating secure maritime operations. However, despite its critical role, AIS continues to exhibit significant vulnerabilities, most notably, its susceptibility to spoofing attacks due to the lack of built-in authentication and validation mechanisms. These weaknesses can be exploited by malicious actors to falsify vessel locations, conceal illicit activities, and disrupt maritime situational awareness.

Prior research introduced a novel detection method based on analysing the stochastic properties of vessel trajectories and the capabilities of currently available online spoofing tools, which achieved an excellent performance against spoofed tracks that closely resembled recent real-life scenarios. However, as spoofing techniques are likely to evolve, relying solely on this method may no longer be sufficient. This paper proposed a more advanced and adaptive approach designed to withstand potential enhancements in spoofing tactics, particularly the injection of synthetic noise into fabricated tracks, which can significantly lower the performance of traditional detection models.

Until now, no previous studies have addressed AIS spoofing detection under adversarial conditions involving artificially injected noise. By leveraging real-time deep learning architectures such as 1D CNNs and GRUs, the new methodology maintained a high detection accuracy even when exposed to noise-injected spoofed data. Nevertheless, this line of defence must continue to evolve. In order to increase the robustness of the obtained results and further optimize the models, similar experiments should be conducted in other maritime regions. Future research should also focus on detecting increasingly sophisticated spoofing attempts, including those capable of generating highly realistic virtual trajectories that mimic not only the motion patterns of genuine vessels but also the dynamic AIS TDMA behaviour under variable signal propagation conditions which can be influenced by weather and geographical conditions.

REFERENCES

Androjna, A., Perkovic, M., Pavic, I. et al. (2021) AIS Data Vulnerability Indicated by a Spoofing Case-Study. *Applied Sciences.* 11(11), art. no. 5015. https://doi.org/10.3390/app11115015.

Androjna, A., Pavic, I., Gucma, L. et al. (2024) AIS Data Manipulation in the Illicit Global Oil Trade. *Journal of Marine Science and Engineering.* 12(1), art. no. 6. https://doi.org/10.3390/jmse12010006.

Androjna, A. & Perkovič, M. (2024) AIS Data Falsification - How Long Will It Be Before We Can No Longer TrustAIS? In: *IEEE International Workshop on Metrology for the Sea; Learning to Measure Sea Health Parameters (MetroSea)*, *14-16 October 2024, Portorose, Slovenia.* New York, USA, IEEE. pp. 301-306.

Anon. (n.d.) *San Francisco Bay Area AIS feed.* http://ais.exploratorium.edu/ [Accessed 1st June 2025].

Coleman, J., Kandah, F. & Huber, B. (2020) Behavioral Model Anomaly Detection in Automatic Identification Systems (AIS). In: *10th Annual Computing and Communication Workshop and Conference (CCWC)*, *6-8 January 2020, Las Vegas, USA*. New York, USA, IEEE. pp. 481-487.

Görkem, B. N., Çağlayan, B., Karaca, E. et al. (2023) Dark Activity Detection in AIS-Based Maritime Networks. In: *34th Conference of Open Innovations Association (FRUCT)*, *15-17 November 2023, Riga, Latvia*. New York, USA, IEEE. pp. 35-40.

Iphar, C., Napoli, A. & Ray, C. (2015) Detection of false AIS messages for the improvement of maritime situational awareness. In: *OCEANS 2015 - MTS/IEEE*, *19-22 October 2015, Washington, USA*. https://doi.org/10.23919/OCEANS.2015.7401841.

Jose, A. (n.d.) *AIS maritime data*. https://www.kaggle.com/datasets/aswinjose/ais-maritime-data [Accessed 1st June 2025].

Katsilieris, F., Braca, P. & Coraluppi, S. (2013) Detection of malicious AIS position spoofing by exploiting radar information. In: *Proceedings of the 16th International Conference on Information Fusion*, *9-12 July 2013, Istanbul, Turkey*. New York, USA, IEEE. pp. 1196-1203.

Kessler, G. C. & Zorri, D. M. (2024) AIS Spoofing: A Tutorial for Researchers. In: *IEEE 49th Conference on Local Computer Networks (LCN), 8-10 October 2024, Normandy*, *France*. New York, USA, IEEE. https://doi.org/10.1109/LCN60385.2024.10639747.

Kiranyaz, S., Avci, O., Abdeljaber, O. et al. (2021) 1D convolutional neural networks and applications: A survey. *Mechanical Systems and Signal Processing*. 151, art. ID 107398. https://doi.org/10.1016/j.ymssp.2020.107398.

Louart, M., Szkolnik, J. J., Boudraa, A.-O. et al. (2023) Detection of AIS messages falsifications and spoofing by checking messages compliance with TDMA protocol. *Digital Signal Processing.* 136, art. ID 103983. https://doi.org/10.1016/j.dsp.2023.103983.

Pohontu, A., Vertan, C., Ciocioi, I. et al. (2025) Detection of spoofed AIS: Simulated tracks vs. real maritime data. *Revista Română de Informatică și Automatică* [*Romanian Journal of Information Technology and Automatic Control*]. 35(1), 37-50. https://doi.org/10.33436/v35i1y202503.

Sciancalepore, S., Tedeschi, P., Aziz, A. et al. (2022) Auth-AIS: Secure, Flexible, and Backward-Compatible Authentication of Vessels AIS Broadcasts. *IEEE Transactions on Dependable and Secure Computing.* 19(4), 2709-2726. https://doi.org/10.1109/TDSC.2021.3069428.

Suratman, M. F. & Mansor, M. F. (2024) Development of GPS Receive Antenna with Anti-Jamming Capabilities. In: *IEEE 7th International Symposium on Telecommunication Technologies (ISTT)*, *21-22 October 2024, Langkawi Island, Malaysia*. New York, USA, IEEE. 162-165. https://doi.org/10.1109/ISTT63363.2024.10750676.

Wimpenny, G., Grant, A., Safar, J. et al. (2022) Securing the Automatic Identification System (AIS): Using public key cryptography to prevent spoofing whilst retaining backwards compatibility. *Journal of Navigation.* 75(2), 333-345. https://doi.org/10.1017/S0373463321000837.