Journal of the Brazilian Computer Society, 2025, 32:1, doi: 10.5753/jbcs.2026.5309

Multiclass Classification for Detection of GPS Spoofing and Jamming Attacks on UAVs

Gustavo Gualberto Rocha de Lemos [**Federal University of ABC** | ***gustavo.gualberto@aluno.ufabc.edu.br***]
Rodrigo Augusto Cardoso da Silva [**Federal University of ABC** | ***cardoso.rodrigo@ufabc.edu.br***]

Center for Mathematics, Computing and Cognition, Federal University of ABC, Av. dos Estados, 5001, Santo André, 09210-580, SP, Brazil.

Received: 09 December 2025 • **Accepted:** 29 September 2025 • **Published:** 17 March 2026

Abstract. Unmanned Aerial Vehicles (UAVs) are increasingly being employed across various domains, making them more vulnerable to a range of attacks, particularly cyber threats. These vehicles usually rely on a global navigation satellite system (GNSS), such as the Global Positioning System (GPS) satellites, for location and navigation data, which can be exploited by adversaries launching attacks using fake GPS signals. To safeguard UAVs from GPS Jamming and GPS Spoofing attacks, this paper proposes an Intrusion Detection System (IDS) that utilizes machine learning techniques for detecting and identifying such attacks. The IDS analyzes GPS signal samples representing normal operation, GPS Jamming, and three types of GPS Spoofing attacks. It relies on machine learning, with models trained and tested for binary class and multiclass classification. The binary class version aims to identify an occurrence of any attack, irrespective of type, as suggested by previous literature. However, the novelty of this work lies in the multiclass version, which enables the identification of attack types — an essential factor in determining the most effective protective measures and providing data for forensic investigations. Stacking, an ensemble machine learning method, yielded the best results, achieving an accuracy rate of 96.91%. Furthermore, the proposed multiclass IDS reduced false negatives to 0.71%, leading to an improved IDS that reduces the likelihood of overlooking attacks compared to the binary class version, which is crucial in real UAV deployments.

Keywords: UAV, GPS, Jamming, Spoofing, Intrusion Detection System, Machine Learning

1 Introduction

Unmanned aerial vehicles (UAVs), commonly known as drones, trace their origins back to 1917 when the military scientist Archibald Low designed a UAV for potential use in the First World War Bartsch *et al*. [2016]. While this initial attempt was unsuccessful, subsequent technological advancements have facilitated the extensive utilization of UAVs across various sectors including agriculture, surveillance, smart cities, search and rescue, and computer networks Weber *et al*. [2021]; da Silva *et al*. [2021]; Khan *et al*. [2022]; da Silva and da Fonseca [2023]; Araújo *et al*. [2023]. The versatility of UAVs, equipped not only with mechanical devices but also with advanced processing and networking capabilities, makes them potential targets for cyber attacks Ranyal and Jain [2021]; Prasad and Khilar [2024].

UAVs can be operated remotely by a ground-based pilot or autonomously. During their operations, particularly in autonomous flights, UAVs rely on Global Navigation Satellite Systems (GNSS) signals, which are transmitted by satellites orbiting the Earth. The GNSS consists of different satellite constellations controlled by various countries. The first positioning satellite constellation for public use was the Global Positioning System (GPS), controlled by the USA. GNSS signals are commonly referred to as GPS signals. Therefore, this article adopts that convention, with the term GPS referring to the entire GNSS system, not just the American constellation. The GPS signal comprises identifying information for each satellite, along with orbital data about the satellite and its constellation. Devices on Earth, such as UAVs, utilize these data to determine their position. There are two types of GPS signals: one designed for civilian use, which lacks mechanisms to ensure confidentiality and integrity of data, and another encrypted signal reserved exclusively for military purposes. The military-grade signal significantly minimizes the risk of attacks that manipulate GPS signal data Manulis *et al*. [2021]. However, the majority of devices, including non-military UAVs, rely on the civilian signal, making them more susceptible to GPS attacks.

Two common attacks involving GPS signals are GPS Jamming and GPS Spoofing Kim *et al*. [2012]; Derhab *et al*. [2023]; Kerns *et al*. [2014]. In GPS Jamming, the attacker emits radiofrequency energy with sufficient power to disrupt GPS signals received by the devices in the target area. In GPS Spoofing, the attacker employs on-ground antennas to generate GPS signals containing false location information. These attacks can result in significant losses. For instance, in 2018, a GPS Jamming attack against a drone swarm caused losses exceeding one hundred thousand dollars McCarthy *et al*. [2018]. In the event of a successful GPS Spoofing attack, the UAV can be directed towards the attacker's location, enabling them to hijack the aircraft, steal its payload, and potentially access sensitive information stored in the onboard disks. Therefore, the detection of GPS attacks has become a critical concern in autonomous drone missions.

One way to mitigate GPS attacks is by analyzing the re-

ceived signal to verify whether it was transmitted by an actual satellite. In this scenario, the UAV employs software to conduct signal analysis—an Intrusion Detection System (IDS) capable of autonomously identifying false GPS signals. Upon detection, the IDS can promptly alert the aircraft owner, or the UAV can rely on alternative instruments for navigation. Implementing an IDS for GPS signals often involves the utilization of machine learning techniques. The system undergoes a training phase using data obtained from normal and attacked operational scenarios to establish a baseline for identifying subsequent attacks Abdulganiyu *et al*. [2023]. Once properly trained, the model can be deployed on the UAV to detect potential attacks.

A simple approach for the IDS is to determine whether the aircraft is subject to a GPS attack or not. However, different attacks have diverse implications for UAVs. In a GPS jamming attack, for instance, the aim of the attacker may be just to disorient the aircraft, whereas a GPS Spoofing attack aims at gaining control of the UAV trajectory. Hence, this study examines three distinct classes of GPS Spoofing attacks and the GPS Jamming attack. Instead of solely indicating that the UAV is under attack (binary class IDS), the IDS identifies the specific type of attack in progress (multiclass IDS). This is key information for deciding the most appropriate counter response by the UAV and becomes valuable in forensic investigations of the attack. Previous investigations Aissou *et al*. [2021]; Talaei Khoei *et al*. [2022]; Khoei *et al*. [2023]; Gasimova *et al*. [2022]; Khoei *et al*. [2022]; Aissou *et al*. [2022]; Talaei Khoei *et al*. [2023] proposed IDSs based on machine learning to detect GPS attacks targeting UAVs. However, they were based only on binary class classifiers.

This paper proposes a multiclass IDS to detect GPS Jamming and GPS Spoofing attacks across three sophistication levels. The IDS was initially tested using machine learning base classifiers to yield preliminary results. Subsequently, ensemble classifiers were employed within the IDS, which utilize multiple base classifiers to enhance performance. The IDS was ultimately based on a multiclass Stack that achieved an accuracy of 96.91% with 0.71% false negatives in the evaluations. We assessed both binary class and multiclass solutions. Our results indicate that the multiclass IDS achieved a high accuracy, which, although slightly lower by less than 2% compared to the binary class version, effectively reduced false negatives. In this context, a false negative signifies an undetected attack, potentially resulting in severe consequences such as UAV destruction or hijacking. Thus, the IDS proposed in this paper emphasizes the benefits of multiclass classifiers and could serve as an effective countermeasure against cyber threats involving fake GPS signals.

The remainder of this paper is structured as follows. Section 2 reviews prior research related to the GPS attacks targeting UAVs. Section 3 introduces the background related to this project and the different GPS attacks. Section 4 presents the proposed solution. Section 5 describes the methodology applied in the experiments, while Section 6 discusses the numerical results obtained during the evaluation. Finally, Section 7 draws conclusions from this paper.

2 Related work

Security issues in UAVs, particularly those related to GPS attack, have received significant attention in the literature Omolara *et al*. [2023]; Yaacoub *et al*. [2020]. For instance, the survey in Omolara *et al*. [2023] reviewed 250 papers published between 2012 and 2022 and suggested different approaches to enhance UAV security, such as the development of more efficient and lightweight IDSs. In Yaacoub *et al*. [2020], an extensive analysis was carried out on various threats and vulnerabilities faced by UAVs, accompanied by countermeasures and recommendations for security improvements. The authors suggested the development of an efficient IDS tailored for drones with limited processing resources.

The remainder of this section reviews GPS Spoofing and Jamming attacks in Subsection 2.1, and papers that employed Machine Learning to counteract those attacks in Subsection 2.2.

2.1 GPS Spoofing and Jamming attacks

GPS Spoofing and GPS Jamming attacks have extensively been studied in the literature Derhab *et al*. [2023]; da Silva [2017]; Ranyal and Jain [2021]; Haider and Khalid [2016]; Noh *et al*. [2019]. The authors of Derhab *et al*. [2023] conducted an assessment of the impact and risk associated with various types of attacks on UAVs, classifying GPS Spoofing and GPS Jamming attacks as extremely high-risk. This classification stems from the ease with which attackers can access the GPS system, posing the potential for damage to both the drone and its operational environment. Consequently, these attacks require heightened attention and more robust countermeasures compared to other attack types. The authors also outlined future directions for enhancing drone security, with one of them involving research on IDSs specifically trained on datasets tailored for UAVs.

The study in da Silva [2017] investigated the feasibility of GPS jamming and GPS spoofing attacks by conducting tests on actual equipment. The authors confirmed that these attacks are achievable using Software Defined Radio, a technology that relies on readily available hardware and open-source software. This suggests that individuals do not need advanced expertise in radio technology to perform these attacks, highlighting the need for security mechanisms for UAVs.

GPS Spoofing attacks and their countermeasures have been covered in the literature Ranyal and Jain [2021]; Haider and Khalid [2016]; Noh *et al*. [2019]. The work in Ranyal and Jain [2021] reviewed various papers about GPS Spoofing attacks and the existing countermeasures. The authors emphasized that manufacturers of GPS-dependent devices should take this attack seriously and implement proper countermeasures. In Haider and Khalid [2016], the authors introduced a list of criteria that GPS Spoofing countermeasures should meet, and showed that no single method meets all defined criteria. GPS Spoofing can also be non-malicious, as indicated in Noh *et al*. [2019], which introduces a method for using this technique as a protective measure to divert potentially malicious drones away from sensitive areas.

A GPS Jamming attack was examined in the article Van den Bergh and Pollin [2019], revealing that UAVs are more sus-

ceptible to this attack compared to other GPS-dependent devices, given that jamming signals are more effective in high altitudes due to the lack of obstacles. The authors proposed to employ radio communication as a secondary location system in case of a GPS Jamming attack.

Detecting GPS Spoofing and GPS Jamming attacks on UAVs has been extensively explored Davidovich *et al*. [2022]; Liang *et al*. [2022]; Whelan *et al*. [2022]; Wei *et al*. [2022]; Aissou *et al*. [2021]; Talaei Khoei *et al*. [2022]; solutions employed various methods that involve GPS signal data and other sensors. In Davidovich *et al*. [2022], the approach evaluated images captured by the UAV camera to detect the presence of a GPS Spoofing attack. This detection is achieved by analyzing the correlation between video frames and the drone's location using computer vision algorithms like BF-Matcher and SURF. This method offers the advantage of not requiring additional hardware. However, it has limitations, such as relying on sufficient lighting conditions and being less effective in areas with minimal terrain variation. Differently, the authors of Liang *et al*. [2022] derived the current position of the aircraft based on signals received by nearby UAVs and a ground station, comparing it with received GPS data. The authors showed that the detection of a Spoofing attack can be made in a few seconds, but this approach relies on analyses of data other than the GPS received signal, which is not always possible.

2.2 Machine learning

Machine learning has become a common technique utilized to counteract GPS Spoofing and GPS Jamming attacks Whelan *et al*. [2022]; Wei *et al*. [2022]; Aissou *et al*. [2021]; Talaei Khoei *et al*. [2022]. The authors of Whelan *et al*. [2022] introduced the IDS named MAVIDS, which employs an anomaly-based detection method for identifying GPS Spoofing and GPS Jamming attacks. The authors trained the classifier using data collected from the sensors of the drone that would utilize the IDS, offering a viable solution in the face of limited availability of consistent, drone-specific datasets. The classifiers utilized include One-Class Support Vector Machine (OC-SVM), Local Outlier Factor (LOF), and Autoencoder, yielding F1 scores of 90.57% for GPS Spoofing and 94.3% for GPS Jamming. In the study by Wei *et al*. [2022], machine learning was employed to detect GPS Spoofing attacks using not only data from the GPS receiver but also information from the accelerometer, gyroscope, magnetometer, and barometer. A dataset with sensor parameters recorded during both normal operation and attacks was used to train various classifiers, including Support Vector Machine-linear (SVC-linear), Support Vector Machine-rbf (SVC-rbf), XGBoost, Random Forests (RF), K-Nearest Neighbors (KNN), and Gradient Boost. The most effective classifiers were Random Forests and XGBoost.

The paper by Aissou *et al*. [2021] introduces an IDS employing tree-based classifiers to detect Simple, Intermediate, and Sophisticated GPS Spoofing attacks. Training utilizes a dataset comprising GPS signal feature data obtained from an authentic test with a GPS receiver, along with data modified to reflect the signatures of the three types of GPS Spoofing attacks. The tested classifiers were XGBoost, Random Forests, Gradient Boost, and Light Gradient Boost Machine, with XGBoost achieving the highest accuracy (95.52%).

The dataset created by Aissou *et al*. [2021] has also been utilized in other studies: Talaei Khoei *et al*. [2022]; Khoei *et al*. [2023]; Gasimova *et al*. [2022]; Khoei *et al*. [2022]; Aissou *et al*. [2022]; Talaei Khoei *et al*. [2023]. In Talaei Khoei *et al*. [2022], two dynamic classifier selection methods, Metric-Optimized Dynamic Selection (MOD) and Weighted-Metric-Optimized Dynamic Selection (WMOD), were employed to optimize the performance of ten traditional classifiers, yielding an accuracy of 99.6% for both methods.

The authors of Khoei *et al*. [2023] proposed three Deep Learning techniques for GPS Spoofing attack detection: Deep Neural Network, U Neural Network, and Long Short Term Memory. The U Neural network technique yielded the best result, an accuracy of 98.80%. In Gasimova *et al*. [2022], ensemble models Stacking, Boosting, and Bagging were applied for detecting GPS Spoofing attacks, with the Stacking method achieving the highest accuracy at 95.43%.

The study by Khoei *et al*. [2022] compared supervised and unsupervised machine learning models for GPS attack detection, with the Decision Tree model demonstrating the highest accuracy at 99.87%. Aissou *et al*. [2022] evaluated instance-based classifiers, with the best accuracy (92.78%) achieved by the Nu-SVM model. Additionally, Talaei Khoei *et al*. [2023] investigated the impact of changes in classifier parameters and dataset characteristics on classifier performance, with the Decision Tree model achieving an accuracy of 99.99% on a balanced dataset with optimized parameters. The dataset employed by Aissou *et al*. [2021]; Talaei Khoei *et al*. [2022]; Khoei *et al*. [2023]; Gasimova *et al*. [2022]; Khoei *et al*. [2022]; Aissou *et al*. [2022]; Talaei Khoei *et al*. [2023] served as the basis for the present paper.

Table 1 provides a comparative analysis of the IDSs discussed in this section. The table categorizes each paper based on the classifiers used, the types of attacks addressed (GPS Jamming or Spoofing), the consideration of multiple classes of attacks, the number of features used for training, and the highest accuracy achieved. The accuracy of the present study might appear suboptimal compared to results obtained by other authors. However, prior research typically focused on binary class classifiers, overlooking multiple attack classes. In contrast, our study accounts for various classes of GPS attacks, which is crucial for effectively adopting countermeasures. Additionally, this paper addresses GPS Jamming attacks, which were often neglected in previous papers.

3 GPS attacks

A UAV in autonomous flight must accurately determine its location to fulfill its mission. For instance, in aerial surveillance applications, an aircraft needs to know its precise location during its flight. Should an attacker launch a GPS Jamming attack against this UAV during its flight, it would lose access to its location data, preventing it from following its surveillance route. In the case of a GPS Spoofing attack, the attacker could alter the UAV route, redirecting it away from its surveillance area or causing a collision.

This paper considers two different attacks, namely GPS

Table 1. Comparison of papers proposing IDSs based on machine learning to detect GPS Spoofing and GPS Jamming attacks.

Paper	Tested Classifiers	Attacks Detected	Multiclass	Features in the Dataset	Best Accuracy (%)
Whelan *et al*. [2022]	OC-SVC LOF Autoencoder	GPS Spoofing GPS Jamming	No	14	94.3 (Autoencoder)
Wei *et al*. [2022]	SVC-linear SVC-rbf KNN RF Gradient Boost XGBoost	GPS Spoofing	No	21	99.22 (XGBoost)
Aissou *et al*. [2021]	XGBoost Gradient Boost Random Forests LightGBM	GPS Spoofing	No	13	95.52 (XGBoost)
Talaei Khoei *et al*. [2022]	MOD WMOD	GPS Spoofing	No	13	99.6 (WMOD)
Khoei *et al*. [2023]	DNN UNN LSTM	GPS Spoofing	No	13	98.80 (UNN)
Gasimova *et al*. [2022]	Stacking Boosting Bagging	GPS Spoofing	No	13	95.43 (Stacking)
Khoei *et al*. [2022]	GNB DT LR RF L-SVM ANN PCA K-MC Autoencoder	GPS Spoofing	No	13	99.87 (DT)
Aissou *et al*. [2022]	KNN RN L-SVM C-SVM Nu-SVM	GPS Spoofing	No	13	92.78 (Nu-SVM)
Talaei Khoei *et al*. [2023]	DT ANN RF LR GNB SVM	GPS Spoofing	No	13	99.99 (DT)
This work	AdaBoost Gradient Boost Random Forests Soft Voting Hard Voting Bagging Stacking	GPS Spoofing GPS Jamming	Yes	13	96.91 (Stacking)

Spoofing and GPS Jamming attacks. According to Haider and Khalid [2016], GPS Spoofing attacks can be classified into three types: Simple Spoofing, Intermediate Spoofing, and Sophisticated Spoofing. Table 2 displays the features of the GPS signal, which will be used to characterize these attacks. Figure 1 illustrates the GPS attacks considered in the present paper, described in the remainder of this subsection.

In the Simple Spoofing attack (Figure 1a), the attacker employs a single GPS antenna and does not know the position of the UAV. They generate a desynchronized false GPS signal, causing the Doppler shift measurements to deviate beyond the normal range of about 20 Hz. In this type of attack, the false GPS signals arrive with higher power compared to the authentic ones since they are generated from a nearby antenna, resulting in higher carrier to noise ratio ($C/N0$), the ratio between the received signal power and the noise. This makes this attack easy to detect Aissou *et al*. [2021].

In an Intermediate Spoofing attack, the attacker still employs a single antenna (Figure 1a) but they are aware of the UAV position and can manipulate the generated GPS signal to maintain normal Doppler Shift and pseudo-range values. Detecting such an attack requires meticulous monitoring of the following signal information: time of week, carrier phase shift, and correlator amplitude Aissou *et al*. [2021].

In a Sophisticated Spoofing attack, the attacker employs various synchronized GPS antennas to generate false signals across multiple channels to mimic the functionality of a legitimate GPS satellite constellation, gaining complete control over the GPS system used by the UAV, as displayed in Figure 1b. Synchronizing multiple antennas is challenging but can be achieved using advanced Radio Defined Software technologies. The signal features most affected by this type of attack are the correlators Aissou *et al*. [2021].

Lastly, in a Jamming attack, a high-power signal is sent to the UAV GPS receiver. Given the typically low power of the GPS signal, approximately -160dB Misra and Enge [2006], this interference directly impacts the $C/N0$, preventing the access to the genuine signal. Figure 1c illustrates this attack.

A successful GPS Jamming attack can disrupt the UAV flight, potentially causing collisions with obstacles or directing it towards inaccessible areas, resulting in aircraft loss. GPS Spoofing attacks can similarly lead to aircraft loss but carry the added risk of falling under attackers' control, enabling mission sabotage and sensitive data leak.

4 Proposed IDS

The solution proposed in this paper is an IDS designed to detect ongoing attacks and alert the UAV operating system. We assume that the drone will execute an IDS software to identify potential GPS Spoofing or Jamming attacks. This system will analyze data from received GPS signals, outputting whether an attack is occurring and, if so, specifying its type. The IDS will be lightweight to enable deployment across various UAV models.

The IDS can take necessary actions, such as notifying an alarm center of an ongoing attack or landing the aircraft. The IDS will undergo a training phase before the UAV operation using a multiclass machine learning model, trained using data

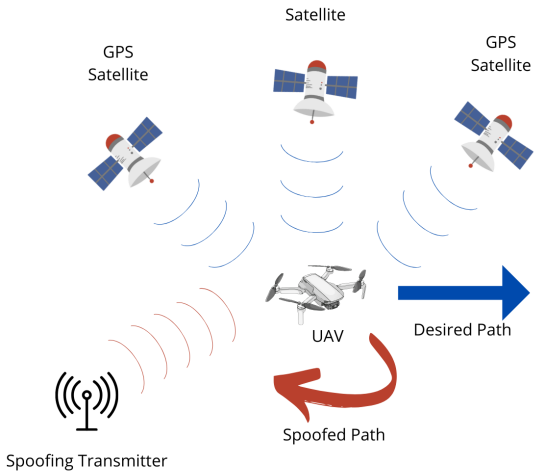

(a) Simple and Intermediate GPS Spoofing attack.

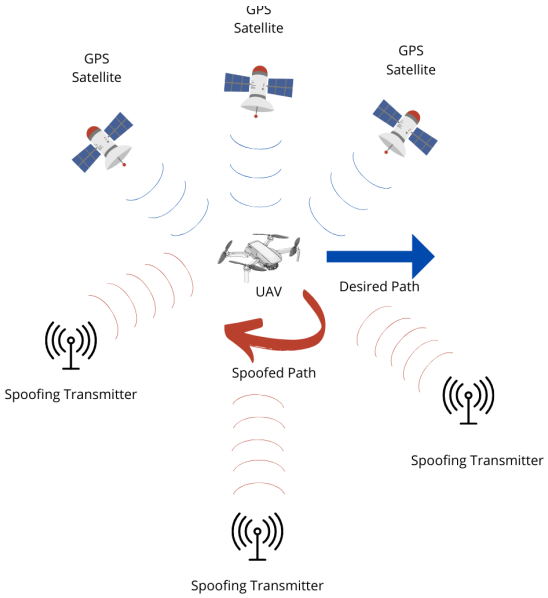

(b) Sophisticated GPS Spoofing attack.

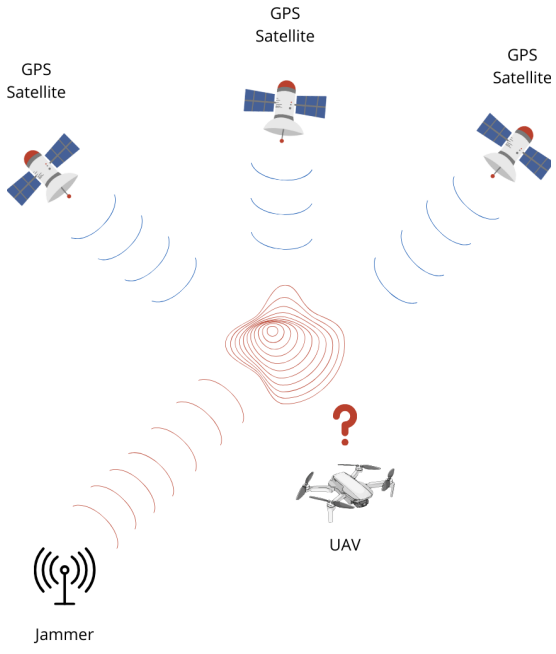

(c) GPS Jamming attack.

Figure 1. Types of GPS attacks targeting UAVs.

Table 2. List of features, based on the Table 2 of Talaei Khoei *et al*. [2022]

Abbreviation	Feature	Description
PRN	Pseudo Random Number	Identification of the satellite
DO	Doppler Shift Measurement	Difference in the signal frequency related to the Doppler Effect
PR	Pseudo Range	Difference between the transmission and the reception time
RX	Receiver Time	Receiver reference time
TOW	Time of Week	Seconds elapsed since the start of each week (from 0s to 604799s)
CP	Carrier Phase Cycles	Beat frequency difference between the received carrier and a receiver-generated carrier replica
EC	Early Correlators	0.5 chip spacing before Prompt Correlator
LC	Late Correlators	0.5 chip spacing after Prompt Correlator
PC	Prompt Correlator	Operation used in synchronizing with the incoming GPS signal, expressed as $\sqrt{PIP^2 + PQP^2}$
PIP	Prompt In-Phase Component	Component of the Prompt Correlator amplitude
PQP	Prompt Quadrature Component	Quadrature component of the Prompt Correlator amplitude
TCD	Tracking Carrier Doppler	Doppler shift that is measured during the correlation stage
CN0	Carrier to Noise Ratio (C/N0)	Ratio of the signal power to the noise

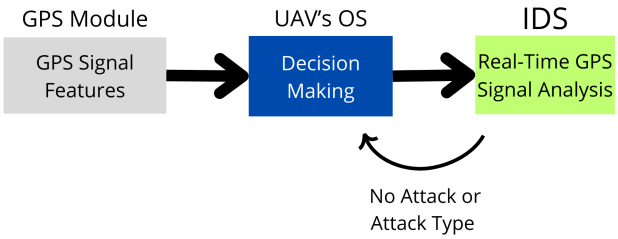

Figure 2. IDS operation and data pipeline.

of GPS signal data from usual drone operations as well as data representing various Spoofing and Jamming attacks. The IDS recognizes potential attacks based on features of the received GPS signals.

Figure 2 illustrates the operation of the IDS in a real-world environment and the data pipeline between it and the UAV's operating system (OS). The IDS would be pre-installed in UAV's disk and loaded into main memory during operation. The UAV's OS would collect signal feature values from the sensors and promptly forward them to the IDS. The IDS would then analyze the data and return a classification result to the OS. If an attack is detected, the OS could apply specific countermeasures based on the type of attack. For example, if the IDS identifies a GPS jamming attack, the most appropriate action would be to switch to manual and remote control, since the UAV would lose access to its positioning. In the case of Simple Spoofing, the OS could attempt to change the drone's route to move away from the attack source, since this type of spoofing is typically performed using only one RF antenna. If the attack is classified as Intermediate or Sophisticated, countermeasures may vary — such as returning to base, switching to manual control, or initiating a landing. In any attack scenario, the OS should also log the time and location of the incident and send an alert to the control center to assist in identifying and investigating the attackers.

Other countermeasures include, for instance, suspending the mission as soon as IDS detects a GPS Jamming attack, and scan for areas where the $C/N0$ parameter restores to normal values. If this fails to happen within a predefined distance limit, it may engage secondary sensors like a camera, gyroscope, or barometer to attempt orientation while awaiting further instructions. Different responses can be pre-programmed based on the sophistication level of the identified GPS Spoofing attack. As the IDS might produce false positives, the system might incorporate specific tolerance thresholds, activating protective measures only after a certain number of positive detections during monitoring. This minimum threshold may differ depending on the sophistication level. Furthermore, the protective measures themselves could vary, including different search areas for the true signal depending on the identified attack. The array of strategies to address GPS attacks can be extensive and is beyond the scope of the present work. Nonetheless, the multiclass IDS proposed in this paper furnishes additional data to enable more efficient strategies compared to IDSs based on binary classification.

5 Experimental Methodology

This section introduces the methodology employed in the experiments to evaluate the proposed IDS. Subsection 5.1 details the data utilized, with the covered attack types, and its application in training the classifiers. Subsection 5.2 describes the classifiers used and is divided into two subsections: Subsection 5.2.1 describes the individual classifiers, while Subsection 5.2.2 details the ensemble methods employed. Lastly, Subsection 5.3 presents the metrics employed to numerically assess the proposed IDS.

5.1 Dataset

The data utilized for the training phase and subsequent tests was derived from the dataset introduced in the article Aissou *et al*. [2021], available at Aissou [2022]. The original dataset comprises real-time data extracted from an actual 8-channel GPS receiver, yielding the 13 distinctive features depicted in Table 2. Furthermore, following the categorization of GPS Spoofing attacks into the three categories (Section 3), the authors of Aissou *et al*. [2021] modified the dataset to

simulate GPS Spoofing attacks, altering authentic signals within the dataset to reflect the attack signatures.

The original dataset Aissou [2022] contains entries for normal operation and the three categories of Spoofing attacks. We extended this dataset to include entries simulating the signature of the Jamming attack described in Section 3. The simulation of this attack was based on the work in da Silva [2017], which analyzed the impact of GPS Jamming on the $C/N0$ parameter, revealing that increased noise levels mainly lead to a reduction in $C/N0$ values. When they drop below 25 dB-Hz, noise is sufficient to disrupt the UAV localization function. The authors of Misra and Enge [2006] demonstrated that environmental obstacles such as enclosed spaces and buildings also degrade the carrier-to-noise ratio. However, considering that drones predominantly operate in open spaces and at significant altitudes, it is reasonable to assume that high noise levels are due to Jamming attacks rather than environmental obstructions. Therefore, the entries simulating this attack exhibit $C/N0$ parameter values ranging between 21 and 24 dB-Hz. This range was selected since Jamming attacks that drop the $C/N0$ to these levels are already sufficient to impair the operation of the GPS receiver Bauernfeind *et al*. [2011]. Consequently, including values below 21 dB/Hz is unnecessary.

The final dataset, available at Lemos [2023], consists of 397,646 samples (73,46%) representing normal operation, 36,438 samples (6,73%) of Simple Spoofing attacks, 44,212 samples (8,17%) of Intermediate Spoofing, 31,995 samples (5,91%) of Sophisticated Spoofing, and 31,022 samples (5,73%) of Jamming. The complete dataset contains synthetic data for all attack types: the synthetic Spoofing attack samples were provided in the original dataset, and the Jamming samples were appended specifically for the present paper. These data are based on real measurements of normal UAV operation and are effective for use in training; however, they may not capture the full environmental variability present in real-world attack scenarios.

For preprocessing data, we performed a correlation analysis between the features and made a balancing of both dataset versions. The correlation analysis helped eliminate redundant features, reducing dataset size and improving classifier performance. Through the calculation of the Spearman's Correlation, two pairs of highly correlated features in the dataset were detected: DO and TCD (95% correlation), and TOW and RX (94% correlation). This correlation is expected, as DO and TCD are both based on Doppler Shift measurements, while TOW and RX are based on time measurements. We subsequently removed the features with the lowest Information Gain from each pair, namely TCD and RX. This operation does not compromise the results of the classifiers. The dataset balancing was carried out through undersampling. For the binary version, the dataset was balanced to contain 143,667 attack entries and an equal number of No Attack entries. For the multiclass version, each class (No Attack, Simple Spoofing, Intermediate Spoofing, Sophisticated Spoofing, and Jamming) was adjusted to contain 31,022 samples. A balanced dataset helps prevent classifier bias and increases the reliability of the accuracy metric.

5.2 Classifiers

The training phase for all classifiers was conducted using 5-fold cross-validation, with 80% of the samples used for training and 20% for testing in each fold, selected randomly. All classifiers employed are from scikit-learn, an open-source Python library for machine learning Pedregosa *et al*. [2011], using the default configuration and the hyperparameters in Table 3. Initially, base classifiers were tested, followed by ensemble classifiers, which are methods of utilizing multiple classifiers to achieve improved results. All classifiers are described in Scikit-learn [2025].

The Bayesian Optimization algorithm determined the best hyperparameters for each classifier. Bayesian Optimization is an efficient technique for fine-tuning machine learning model parameters, as it quickly converges to good values. Table 3 displays the hyperparameter values used by each classifier in both versions of the dataset.

5.2.1 Base Classifiers

This subsection briefly introduces the base classifiers employed in this work, selected based on their adaptability to both binary class and multiclass classification. First, Naive Bayes classifiers are based on Bayes' theorem and assume that all features are independent. We employed the Gaussian Naive Bayes (GNB) version, which assumes a Gaussian distribution of data features.

The K-Nearest Neighbors (KNN) classifier is based on the proximity of a new sample to known examples within the training dataset, conducting a majority vote among the neighbors. The value of K determines the number of neighbors taken into account, and various distance measures can be utilized. Differently, the Decision Tree (DT) classifier builds a decision tree by splitting the data into smaller subsets based on the features. It starts with all the data set as a set and splits it into two subsets using a specific feature chosen. The selected feature is the one that yields the greatest information gain or the most significant decrease in subset impurity. The splitting process continues recursively until all subsets are pure or nearly pure concerning a class or output value.

The Multilayer Perceptron Neural Network (NN) employs the Backpropagation algorithm for its training process. It comprises various connected layers of neurons, with an output layer that transforms values from the last hidden layer into an output value. Linear Discriminant Analysis (LDA) generates a between-class scatter matrix and a within-class scatter matrix. These matrices help identify discriminant directions that effectively separate the classes, and data points are projected along these directions, resulting in a new set of features that optimally distinguish between classes.

Finally, Logistic Regression (LR) models the probability of an instance belonging to a particular class using the sigmoid logistic function, transforming linear output into probability values. Subsequently, this probability is assigned to a particular class based on a decision threshold.

5.2.2 Ensemble Methods

Four ensemble methods were utilized: Boosting, Bagging, Voting, and Stacking. Whereas the Boosting method utilizes

Table 3. Hyperparameter Tuning

Classifier	Hyperparameter Setting	Best Hyperparameters (Binary Class)	Best Hyperparameters (Multiclass)
GNB	var_smoothing = range(1e-9, 1e+9)	var_smoothing = 5.749221387201998e-07	var_smoothing = 8.482673276130252e-08
KNN	n_neighbors = range(5, 100), p = range(1, 3)	n_neighbors = 17, p = 1	n_neighbors = 16, p = 1
DT	criterion = ['gini', 'entropy'] max_depth = range(1, 35)	criterion = 'entropy' max_depth = 17	criterion = 'entropy' max_depth = 19
NN	activation = ['identity', 'logistic', 'tanh', 'relu']	activation = 'tanh'	activation = 'tanh'
LDA	solver = ['svd', 'lsqr', 'eigen']	solver = 'svd'	solver = 'eigen'
LR	solver = ['lbfgs', 'liblinear', 'newton-cg', 'newton-cholesky', 'sag', 'saga']	solver = 'newton-cholesky'	solver = 'liblinear'
AB	n_estimators = range(10, 1000)	n_estimators = 994	n_estimators = 957
RF	criterion = ['gini', 'entropy', 'log_loss']	criterion = 'log-loss'	criterion = 'entropy'
GB	learning_rate = range(0.1, 1)	learning_rate = 0.20225246845483644	learning_rate = 0.13320966081361318

its internal classifiers, the other ones require specifying the classifiers to be utilized. The remainder of this subsection introduces those ensemble classifiers.

Boosting ensembles combine multiple weak base classifiers to make more accurate predictions. For instance, Adaptive Boosting (AdaBoost) combines weak classifiers in a weighted fashion. It starts by fitting a classifier to the original dataset and then fits additional copies of the classifier to the same dataset, giving more weight to misclassified samples in each iteration. Thus, AdaBoost focuses on the most challenging samples to classify, improving its overall performance. Differently, the Random Forests classifier creates multiple independent decision trees and combines their predictions, determining the final decision by majority voting among the trees. Gradient Boosting also combines diverse decision trees but differs in their construction and combination methodology. It sequentially trains multiple trees, adjusting each tree to correct errors made by its predecessors. This progressive technique improves the classifier's performance by leveraging the strengths of multiple weak decision trees.

The Bagging method (Bootstrap Aggregating) trains multiple copies of the same classifier on random subsets and their predictions are combined through majority voting. For the Bagging method, the KNN and Decision Tree classifiers were used. Note that the previously mentioned Random Forests method is also considered a special type of Bagging. For the Voting and Stacking methods, the 5 classifiers that performed best in previous tests were selected: KNN, Decision Tree, Neural Network, Random Forests, and Gradient Boosting. The Voting method combines conceptually different classifiers and uses either majority voting (hard) or average predicted probabilities (soft) to predict result labels. The Stacking method trains a set of classifiers on a dataset, and the predictions from these base classifiers are used to train a final classifier, also known as a meta-classifier. Gradient Boost was chosen as the meta-classifier for the Stacking method due to its performance.

5.3 Evaluated metrics

The metrics employed to evaluate the classifiers in this paper were Accuracy, Confusion Matrix, Precision, Recall, and F1-Score. Accuracy is the ratio of correct predictions to the total number of samples and is particularly useful when errors in predicting all labels carry equal weight Burkov [2019]. This metric served as the basis for comparing the performance of the classifiers during the initial evaluation, since the objective is to identify a classifier with the best overall results. Equation (1) illustrates how it is calculated, where $\hat{y}_i$ represents the predicted value of the i-th sample, y_i is the corresponding true value, and $1(x)$ denotes the Indicator Function.

The confusion matrix is a table that demonstrates the distribution of the classifier's predictions across the dataset labels. The horizontal axis denotes the predicted labels by the classifier, while the vertical axis corresponds to the actual labels. This matrix was employed to assess the distribution of errors of the classifier with the highest accuracy.

After the initial evaluation, the Precision, Recall, and F1-Score metrics were employed for a deeper evaluation of the classifier with the highest accuracy. Precision is the ratio of accurately predicted positive samples relative to the total samples predicted as positive. A higher Precision value implies a stronger capability of the classifier to reduce the number of false positive samples. Recall represents the ratio of accurately predicted positive instances in relation to the total instances that are genuinely positive. Higher Recall values imply a better ability of the classifier to identify all positive instances. The F1-score is the harmonic mean between Precision and Recall. Equations (2)–(4) display the calculation formulas for Precision, Recall, and F1-Score, respectively, where TP denotes the number of true positives, FP the number of false positives, and FN the number of false negatives.

$$Accuracy(y, \hat{y}) = \frac{1}{n_{\text{samples}}} \sum_{i=0}^{n_{\text{samples}}-1} 1(\hat{y}_i = y_i) \quad (1)$$

$$Precision = \frac{TP}{TP + FP} \quad (2)$$

$$Recall = \frac{TP}{TP + FN} \quad (3)$$

$$F1\text{-}score = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (4)$$

6 Numerical results

This section presents the numerical evaluation of the IDS proposed in this paper by discussing the results of all evaluated classifiers. Table 4 displays the accuracy results obtained from the Base Classifiers in both their binary class and multiclass versions. The top-performing classifiers were k-Nearest Neighbours, Decision Tree, Neural Network, Random Forests, and Gradient Boost. Subsequently, these five

Table 4. Accuracy obtained by Base Classifiers.

Classifier	Binary Class Accuracy (%)	Multiclass Accuracy (%)
GNB	62.99	53.80
KNN	90.18	87.56
DT	96.62	95.33
NN	95.36	94.08
LDA	62.95	56.10
LR	62.93	55.23
AB	90.23	78.90
RF	95.57	93.89
GB	96.70	95.50

classifiers were employed in conjunction with Ensemble methods, such as Bagging, Hard Voting, Soft Voting, and Stacking, in an attempt to improve the results, whose results are displayed in Table 5. As described in Subsection 5.2.2, the Bagging method is applied individually to the KNN and DT classifiers, while the other Ensemble methods utilize the classifiers collectively. The Stacking method yielded the highest accuracy, achieving 96.91% accuracy with the multiclass version and 98.03% with the binary class version.

Figure 3 compares the highest accuracies obtained by all classifiers, both in multiclass and binary class versions. For all of them the binary class version yielded the best results. The Accuracy in the binary version, as expected, is higher than in the multiclass version, since adding more classes increases the complexity of the prediction task. Additionally, in the multiclass setting, misclassifying the type of attack is also counted as an error by this metric. GNB, LDA, LR, and AdaBoost showed the largest performance differences between the binary and multiclass versions, while the other classifiers had differences within a 2% range between both versions.

The discrepancies in results highlight that certain classifiers perform better than others for the same data. The No Free Lunch Theorem Wolpert [1996] asserts that it is impossible to determine the best classifier for a dataset beforehand, which emphasizes the importance of testing various classifiers to achieve optimal results Géron [2023].

Since the Stacking method yielded the highest accuracy, it underwent evaluation using additional metrics to gain further insights into the proposed IDS. Figure 4 displays the confusion matrix with the predicted labels against the actual labels in the multiclass dataset test. Label 0 denotes normal operation, while labels 1, 2, and 3 represent Simple, Intermediate, and Sophisticated Spoofing attacks, respectively. Label 4 indicates a Jamming attack. Stacking led to 224 instances of false negatives, where an attack was present but misclassified as normal operation; this number is the sum of three types of misclassification: 164 instances with label 1, 42 with label 2 and 18 with label 3. This constitutes 0.71% of the total 29,255 negatives predicted in the test, and 1,991 instances of false positives, where no attack was present, but the classifier indicated otherwise, representing 6.41% of the total 31,022 non-attack samples in the test.

Additionally, there were 1,136 instances of Simple Spoofing attacks misclassified as Intermediate Spoofing, 1,445 instances of Intermediate Spoofing attacks misclassified as Simple Spoofing, and a single instance of Simple Spoofing misclassified as Sophisticated Spoofing. The Jamming attack remained distinct and was not misclassified. This distinction arises from the consistent $C/N0$ values below 25 dB-Hz, a characteristic that classifiers can readily learn.

Figure 5 displays the confusion matrix generated by the Stacking method when applied to the binary class data. It revealed 1406 false negatives, accounting for 0.97% of the total 140,806 predicted attack entries, and 4,266 false positives, representing 2.97% of the total 146,526 non-attack predictions.

The multiclass Stacking exhibits a slightly lower accuracy compared to its binary class version. However, a close inspection of their respective confusion matrices reveals that the multiclass version produces fewer false negatives, which represent the most critical errors in an IDS since they indicate undetected attacks. The binary class version had 0.97% cases of false negatives, while the multiclass had 0.71%, implying that in a real scenario, the binary class IDS would overlook more attacks than the multiclass. The reduced accuracy of the multiclass version primarily arises from challenges in distinguishing between types of attacks, which, within an IDS context, is much less severe than a false negative. This comparison is outlined in Table 6.

When the IDS accurately identifies the type of ongoing attack, the UAV control system can take necessary actions to prevent the attack from succeeding, thus avoiding potential damages. A false negative indicates that the drone is vulnerable to an attack; for instance, in a Jamming attack, the drone loses its location information, potentially resulting in collisions or navigation into restricted areas, which leads to the aircraft loss and financial implications. Spoofing attacks may also result in drone loss but pose an additional risk of falling under attackers' control, potentially sabotaging missions and causing more extensive damages. Accurately identifying the type of attack also supports forensic investigation, offering insights into the attackers' behavior and helping to identify geographic areas that are more frequently targeted by each attack.

Table 6 displays the Precision, Recall, and F1-score metrics of the Stacking classifier. For the multiclass version, it shows that the No Attack class had the lowest Recall, meaning this class was most often confused with others. In the context of an IDS, this implies that the system incorrectly indicated attacks where there were none (false positives). When analyzing the Simple and Intermediate classes together with the confusion matrix, results make evident that the model struggles to distinguish between these two types of attacks. Their lower precision values, compared to the other classes, stem from this confusion. This can be justified in part by the similar attack signatures of Simple and Intermediate spoofing compared to that of the Sophisticated attacks, as both would typically rely on a single antenna, unlike Sophisticated attacks that employ multiple antennas. The Sophisticated class achieved the best performance, which might seem counterintuitive since it would presumably be better at deceiving detection systems due to the more sophisticated approach with multiple antennas that should better mimic real GPS samples. However, Sophisticated class presents more distinct features, such as changes in correlator values, which could be more accurately detected by the Stacking classifier.

Table 5. Accuracy obtained by Ensemble Classifiers.

Ensemble	Classifiers	Binary Class Accuracy (%)	Multiclass Accuracy (%)
Bagging	KNN	90.20	87.19
	DT	96.40	95.23
Hard Voting	KNN, DT, NN, RF, GB	96.61	95.27
Soft Voting	KNN, DT, NN, RF, GB	96.48	94.98
Stacking	KNN, DT, NN, RF, GB	98.03	96.91

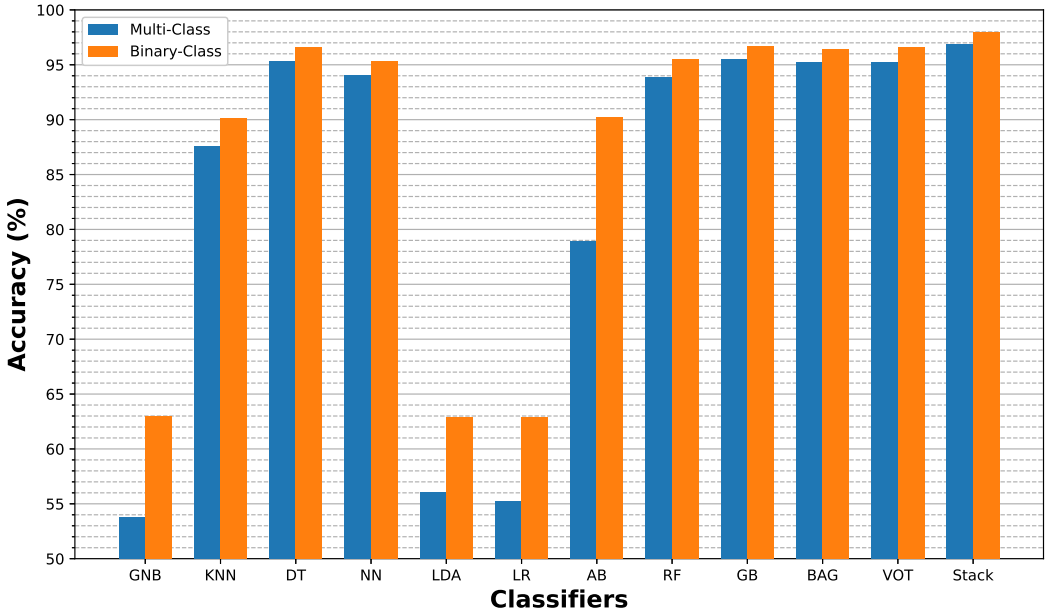

Figure 3. The accuracy obtained by the classifiers in this work.

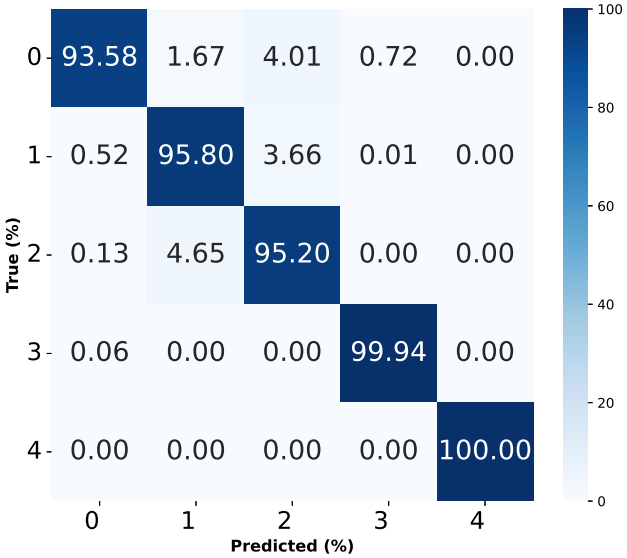

Figure 4. Multiclass Stacking Classifier Confusion Matrix.

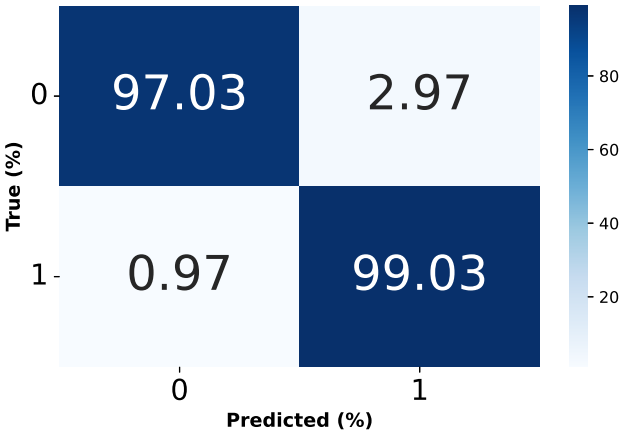

Figure 5. Binary Class Stacking Classifier Confusion Matrix.

Table 6. Precision, Recall, F1-score, accuracy, and false negatives for the Stacking classifier with multiclass and binary classifications.

Multiclass classification			
Class	**Precision (%)**	**Recall (%)**	**F1-Score (%)**
No Attack	99	94	96
Simple Spoofing	94	96	95
Intermediate Spoofing	93	95	94
Sophisticated Spoofing	99	100	100
Jamming	100	100	100
Accuracy		96.91%	
False negatives		0.71%	
Binary classification			
Class	**Precision (%)**	**Recall (%)**	**F1-Score (%)**
No Attack	99	97	98
Attack	97	99	98
Accuracy		98.03%	
False negatives		0.97%	

Table 6 also presents the metric results of the binary class version of the Stacking classifier. Although both classes yielded the same F1-score, the No Attack class has lower recall than precision, while the Attack class has lower precision than recall. This indicates that, despite being well-balanced, the classifier is slightly better at detecting Attack cases than No Attack cases.

This study aimed to identify the most effective classifier for multiple classes of GPS attacks, aiming to use it in a reliable IDS for GPS signals received by UAVs. The top five classifiers in terms of high precision were employed within a Stacking ensemble method to achieve the best result, a multiclass classifier with 96.91% accuracy and only 0.71% false negatives. This method not only detects attacks but also accurately diagnoses the specific attack types, facilitating the implementation of countermeasures after the identification of attacks, providing records for future forensic investigations. Furthermore, the multiclass version notably demonstrated a lower incidence of false negatives compared to its binary class version, highlighting another significant advantage. For an IDS, the false negative error is the most severe since, in a real-world scenario, it signifies an undetected attack. Lastly, it is noteworthy that the trained IDS is lightweight and readily deployable on an actual UAV: the pre-trained Stack binary class classifier measures 132.1MB in size, while its multiclass version occupies 186.9MB.

7 Conclusions

Autonomous UAVs serve numerous critical applications, and their usage is on the rise as technology becomes more accessible. There are various security threats to drones, and some of the most severe attacks exploit the GPS signal.

This work proposed an IDS to detect GPS Jamming and three types of GPS Spoofing attacks. Tests were carried out using both binary class and multiclass versions of the classifiers. Results indicate Stacking as the most effective classifier. The multiclass version of the Stack method achieved an accuracy of 96.91% with only 0.71% false negatives.

Previous work on this topic relied on binary class classifiers. However, the results of this paper showed that the multiclass approach holds an advantage in the number of false negatives, the most critical error in an IDS, since it implies an undetected attack. Moreover, the multiclass approach enables the drone to implement more targeted and efficient countermeasures against a specific attack.

GPS attacks are feasible, but setting up a real environment to replicate them is a task that requires various equipment as well as proper care to avoid damage. For future research, it would be beneficial to test the IDS using an actual UAV exposed to different types of GPS attacks, as well as generating new datasets by collecting data from GPS sensors during actual Spoofing and Jamming attacks. Additionally, continuing the investigation with machine learning models based on alternative techniques might potentially enhance performance compared to the results obtained in this study.

Declarations

Authors' Contributions

All authors contributed to the study conception, analysis of solution, and writing the manuscript. G.G.R.L. reviewed the literature, processed the dataset, coded the IDS, carried out the experiments, and designed the figures. R.A.C.S. supervised the study and reviewed the manuscript. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Funding

This study was financed in part by the São Paulo Research Foundation (FAPESP), grant 2015/24494-8. This research was partially sponsored by CNPq grant 405940/2022-0 and CAPES grant 88887.954253/2024-00.

Availability of data and materials

The datasets generated and analysed during the current study are available in `https://docs.google.com/spreadsheets/d/1srN7w4dO2NU8XKeyeLlwjNZbaz9Sx4Wj/edit?usp=sharing&ouid=107994669384648426370&rtpof=true&sd=true`.

References

Abdulganiyu, O. H., Tchakoucht, T. A., and Saheed, Y. K. (2023). A systematic literature review for network intrusion detection system (IDS). *International Journal of Information Security*, 22:1125–1162. DOI: 10.1007/s10207-023-00682-2.

Aissou, G. (2022). Dynamic selection techniques for detecting GPS spoofing attacks on UAVs. Available at:https://data.mendeley.com/datasets/z7dj3yyzt8/3.

Aissou, G., Benouadah, S., El Alami, H., and Kaabouch, N. (2022). Instance-based supervised machine learning models for detecting GPS spoofing attacks on UAS. In *2022 IEEE 12th Annual Computing and Communication Workshop and Conference (CCWC)*, pages 0208–0214. DOI: 10.1109/CCWC54503.2022.9720888.

Aissou, G., Slimane, H. O., Benouadah, S., and Kaabouch, N. (2021). Tree-based supervised machine learning models for detecting GPS spoofing attacks on UAS. In *2021 IEEE 12th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*, pages 0649–0653. DOI: 10.1109/UEMCON53757.2021.9666744.

Araújo, F., Araújo, F., Alencar, D., Pinheiro, S., Oliveira, H., and Rosário, D. (2023). Dynamic video service migration in flying edge computing networks. *Journal of the Brazilian Computer Society*, 29(1):63–72. DOI: 10.5753/jbcs.2023.2228.

Bartsch, R., Coyne, J., and Gray, K. (2016). *Drones in Society*. Routledge, London, England. DOI: 10.4324/9781315409658.

Bauernfeind, R., Kraus, T., Dötterböck, D., Eissfeller, B., Loehnert, E., and Wittmann, E. (2011). Car jammers: Interference analysis. *GPS World*, 22:28–35. Available at:https://www.researchgate.net/publication/296740298_Car_Jammers_Interference_Analysis.

Burkov, A. (2019). *The Hundred-Page Machine Learning Book*. Andriy Burkov, Quebec. Book.

da Silva, D. A. M. (2017). GPS jamming and spoofing using software defined radio. Master's thesis, University Institute of Lisbon. Available at:https://repositorio.iscte-iul.pt/bitstream/10071/15244/1/GPS%20Jamming%20and%20Spoofing%20Using%20Software%20Defined%20Radio.docx.

da Silva, R. A. and da Fonseca, N. L. (2023). Location of fog nodes mounted on fixed-wing UAVs. *Vehicular Communications*, 41:100600. DOI: 10.1016/j.vehcom.2023.100600.

da Silva, R. A. C., da Fonseca, N. L. S., and Boutaba, R. (2021). Evaluation of the employment of UAVs as fog nodes. *IEEE Wireless Communications*, 28(5):20–27. DOI: 10.1109/MWC.101.2100018.

Davidovich, B., Nassi, B., and Elovici, Y. (2022). Towards the detection of GPS spoofing attacks against drones by analyzing camera's video stream. *Sensors*, 22(7). DOI: 10.3390/s22072608.

Derhab, A., Cheikhrouhou, O., Allouch, A., Koubaa, A., Qureshi, B., Ferrag, M. A., Maglaras, L., and Khan, F. A. (2023). Internet of drones security: Taxonomies, open issues, and future directions. *Vehicular Communications*, 39:100552. DOI: 10.1016/j.vehcom.2022.100552.

Gasimova, A., Khoei, T. T., and Kaabouch, N. (2022). A comparative analysis of the ensemble models for detecting GPS spoofing attacks on UAVs. In *2022 IEEE 12th Annual Computing and Communication Workshop and Conference (CCWC)*, pages 0310–0315. DOI: 10.1109/CCWC54503.2022.9720738.

Géron, A. (2023). *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*. O'Reilly Media, Sebastopol. Book.

Haider, Z. and Khalid, S. (2016). Survey on effective GPS spoofing countermeasures. In *2016 Sixth International Conference on Innovative Computing Technology (INTECH)*, pages 573–577. DOI: 10.1109/INTECH.2016.7845038.

Kerns, A. J., Shepard, D. P., Bhatti, J. A., and Humphreys, T. E. (2014). Unmanned aircraft capture and control via GPS spoofing. *Journal of Field Robotics*, 31(4):617–636. DOI: 10.1002/rob.21513.

Khan, A., Gupta, S., and Gupta, S. K. (2022). Emerging UAV technology for disaster detection, mitigation, response, and preparedness. *Journal of Field Robotics*, 39(6):905–955. DOI: 10.1002/rob.22075.

Khoei, T. T., Aissou, G., Al Shamaileh, K., Devabhaktuni, V. K., and Kaabouch, N. (2023). Supervised deep learning models for detecting GPS spoofing attacks on unmanned aerial vehicles. In *2023 IEEE International Conference on Electro Information Technology (eIT)*, pages 340–346. DOI: 10.1109/eIT57321.2023.10187274.

Khoei, T. T., Gasimova, A., Ahajjam, M. A., Shamaileh, K. A., Devabhaktuni, V., and Kaabouch, N. (2022). A comparative analysis of supervised and unsupervised models for detecting GPS spoofing attack on UAVs. In *2022 IEEE International Conference on Electro Information Technology (eIT)*, pages 279–284. DOI: 10.1109/eIT53891.2022.9813826.

Kim, A., Wampler, B., Goppert, J., Hwang, I., and Aldridge, H. (2012). Cyber attack vulnerabilities analysis for unmanned aerial vehicles. In *Infotech@Aerospace 2012*, pages 1–30. DOI: 10.2514/6.2012-2438.

Lemos, G. (2023). Original dataset. https://docs.google.com/spreadsheets/d/1srN7w4dO2NU8XKeyeLlwjNZbaz9Sx4Wj/edit?usp=sharing&ouid=107994669384648426370&rtpof=true&sd=true.

Liang, C., Miao, M., Ma, J., Yan, H., Zhang, Q., and Li, X. (2022). Detection of global positioning system spoofing attack on unmanned aerial vehicle system. *Concurrency and Computation: Practice and Experience*, 34(7):e5925. DOI: 10.1002/cpe.5925.

Manulis, M., Bridges, C. P., Harrison, R., Sekar, V., and Davis, A. (2021). Cyber security in new space. *International Journal of Information Security*, 20:287–311. DOI: 10.1007/s10207-020-00503-w.

McCarthy, S., Zheng, W., and Tsang, D. (2018). Hk$1 million in damage caused by GPS jamming that caused 46 drones to plummet during hong kong show. Available at:https://www.scmp.com/news/hong-kong/law-and-crime/article/2170669/hk13-million-damage-caused-gps-jamming-caused-46-drones accessed on June 2, 2023.

Misra, P. and Enge, P. (2006). *Global Position System Signals, Measurement and Performance*. Ganga Jamuna Press, Lincoln. Book.

Noh, J., Kwon, Y., Son, Y., Shin, H., Kim, D., Choi, J., and Kim, Y. (2019). Tractor beam: Safe-hijacking of consumer drones with adaptive gps spoofing. *ACM Trans.*

Priv. Secur., 22(2). DOI: 10.1145/3309735.

Omolara, A. E., Alawida, M., and Abiodun, O. I. (2023). Drone cybersecurity issues, solutions, trend insights and future perspectives: a survey. *Neural Computing and Applications*, 35:23063–23101. DOI: 10.1007/s00521-023-08857-7.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Édouard Duchesnay (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12(85):2825–2830. DOI: 10.48550/arxiv.1201.0490.

Prasad, S. V. R. V. and Khilar, P. M. (2024). SVM-SFL based malicious UAV detection in wireless sensor networks. *Concurrency and Computation: Practice and Experience*, 36(13):e8049. DOI: 10.1002/cpe.8049.

Ranyal, E. and Jain, K. (2021). Unmanned aerial vehicle's vulnerability to GPS spoofing a review. *Indian Society of Remote Sensing*, 49:585–591. DOI: 10.1007/s12524-020-01225-1.

Scikit-learn (2007-2025). 1. Supervised learning. `https://https://scikit-learn.org/stable/supervised_learning.html` (Accessed on May 4, 2025).

Talaei Khoei, T., Ismail, S., and Kaabouch, N. (2022). Dynamic selection techniques for detecting GPS spoofing attacks on UAVs. *Sensors*, 22(2). DOI: 10.3390/s22020662.

Talaei Khoei, T., Ismail, S., Shamaileh, K. A., Devabhaktuni, V. K., and Kaabouch, N. (2023). Impact of dataset and model parameters on machine learning performance for the detection of GPS spoofing attacks on unmanned aerial vehicles. *Applied Sciences*, 13(1). DOI: 10.3390/app13010383.

Van den Bergh, B. and Pollin, S. (2019). Keeping UAVs under control during GPS jamming. *IEEE Systems Journal*, 13(2):2010–2021. DOI: 10.1109/JSYST.2018.2882769.

Weber, J. S., Neves, M., and Ferreto, T. (2021). VANET simulators: an updated review. *Journal of the Brazilian Computer Society*, 27(1):8. DOI: 10.1186/s13173-021-00113-x.

Wei, X., Wang, Y., and Sun, C. (2022). Perdet: Machine-learning-based UAV GPS spoofing detection using perception data. *Remote Sensing*, 14(19). DOI: 10.3390/rs14194925.

Whelan, J., Almehmadi, A., and El-Khatib, K. (2022). Artificial intelligence for intrusion detection systems in unmanned aerial vehicles. *Computers and Electrical Engineering*, 99:107784. DOI: 10.1016/j.compeleceng.2022.107784.

Wolpert, D. H. (1996). The Lack of A Priori Distinctions Between Learning Algorithms. *Neural Computation*, 8(7):1341–1390. DOI: 10.1162/neco.1996.8.7.1341.

Yaacoub, J.-P., Noura, H., Salman, O., and Chehab, A. (2020). Security analysis of drones systems: Attacks, limitations, and recommendations. *Internet of Things*, 11:100218. DOI: 10.1016/j.iot.2020.100218.