Journal of Open Aviation Science (2023), Vol.**x**, 1–4
doi:10.74800/joas.x.xxxx

POSTER | *The 11th OpenSky Symposium* (*Pre-print submitted to JOAS*)

ADS-B anomaly detection in the surveillance of low-altitude aircrafts

Pirolley Melvyn, Couturier Raphaël, Salomon Michel, and Ambert Fabrice

Université de Franche-Comté, Belfort, France

Abstract

In the past few years, the fast increase in air traffic load has brought new challenges for air traffic controllers. The air surveillance task has become harder and as a consequence, the actual monitoring tools need to be improved. In this work, a method based on deep learning that automatically detects ADS-B spoofing attacks is proposed. As autonomous drone technologies will, in the near future, be more and more developed, this study focuses on low-altitude traffic. Our tool is based on a classifier model that raises anomalies between true aircraft trajectory shapes and supposed aircraft categories (e.g. planes, helicopters). The proposed approach can detect spoofing attacks with a success rate of 96.2%.

Keywords: Cybersecurity, ADS-B, Low-altitude air traffic, Machine learning

Abbreviations: **ADS-B**: Automatic Dependent Surveillance-Broadcast, **SAMU**: Service d'Aide Médicale Urgente, **AI**: Artificial Intelligence

1. Introduction

In the last few years, the increasing number of autonomous drone tests have shown that drones will, in the near future, be part of the air traffic landscape. As these technologies are also evolving in a military context, drones could be used in a malicious way. An attack scenario based on the spoofing concept was developed: in the context of the 2023 Rugby World Cup in Toulouse, an attacker could use a drone to target the stadium. While flying the drone emits false tracking messages, with the intention of making the air surveillance believe it is a SAMU (French emergency) helicopter.
Currently, some projects are developing defense systems to detect drones half a kilometer away from a sensible point, based on sound with high-sensitivity microphones [1]. However, some malicious attacks could be detected in earlier stages, directly by searching anomalies in aircraft tracking protocols such as ADS-B and FLARM[1]. Our project is the continuation of two works [2, 3].

2. Dataset

To find a large and open set of data for our experiments, the OpenSky database was used. It has been providing a history of ADS-B records all over the world, for more than 7 years.
In order to choose a relevant dataset for our scenario, we selected a bounding box including Toulouse and its surroundings. It starts from Saint-Gaudens to Carmaux, including an important variety of zones such as urban areas, villages, and forested areas. The box's precise coordinates are starting from lat. 0.72561, lon. 43.11581 to lat. 2.16344, lon. 44.07449.
Then our dataset focuses on recent ADS-B data, as the full year 2022 to train our model, and the beginning of 2023 to evaluate the model over unseen data. In the end, after filtering aberrant flights, we have 10,158 flights for training and 819 for evaluation.

[1] FLARM is a proprietary electronic system used to selectively alert pilots to potential collisions between aircraft

Each trajectory is saved in a ".csv" file including classical ADS-B and FLARM features such as Latitude, Longitude, Altitude, Velocity, Track... Finally, trajectories are labeled, by associating the icao24 code (registration code) to an aircraft model and then an aircraft type along: **Commercial planes**, only at take-off and landing when they interact with the low altitude traffic, **Tourism planes**, small airplanes between 2 and 6 seats owned by an individual and **Helicopter**, from SAMU the French emergency services.

3. Method

To detect spoofing attacks, a deep-learning classifier was developed to try and determine aircraft types based on their trajectory only. A well-trained model would detect spoofing attacks by raising a mismatch between the aircraft trajectory and the registration code it pretends to have.
Our solution is based on convolutional neural networks (CNN). Various classic deep-learning architectures were experimented such as Long and Short-Term Memory (LSTM), Transformers [4], and a few others but they lacked the precision of CNNs.
Because our CNN processes time series of variable length, sliding windows were chosen. A prediction is generated using a history of the previous 128 messages, which corresponds to approximately 2 minutes of flight data. In the end, as a single prediction is required, we only keep the most reliable prediction from the entire time series. Indeed a majority vote is less efficient because the most determining patterns for aircraft type recognition are sparse. Typically, during cruising altitudes, flight trajectories exhibit regular and straight patterns.

3.1 Take-off context

To improve our model's detection, we added a take-off contextual vector as a secondary input. As seen in the last paragraph, the most determining patterns for aircraft type recognition are sparse. However, take-off trajectories can contain a lot of relevant features to determine the aircraft's categories. It allows the AI model to always remember the condition of take-off, even at cruising altitude. From a technical point of view, the take-off context is represented by the 128 first timesteps of the whole trajectory.

3.2 Geographical context

Similarly, we conceived the idea to add geographical context to our model's input. It adds first some natural information as surrounding fields, forests, and rivers as aircraft behavior can vary according to the environment. For example, SAMU helicopters whenever possible, tend to follow rivers at take-off to avoid low-flying over dense urban areas.
Secondly, it allows replacing absolute with relative coordinates as the map provides absolute localization information. Using relative coordinates has the effect of reducing the numeric space needed to represent trajectories. With relative coordinates, the space needed to represent trajectories decreases from a rectangle with a diagonal of 155km to 50km. This not only streamlines the learning process but also encourages the model to prioritize the shape of the trajectory over its specific location.
To generate the geographical context, a large image, covering the bounding box was generated with an aggregation of OpenStreetMap tiles. When making predictions, the AI model extracts small 128x128-pixel squares centered on the aircraft's latitude and longitude as can be seen in Figure 1."

3.3 Model architecture

The model architecture (Figure 2) consists of three input layers: the ADS-B messages window, the take-off context, and the geographical context. The ADS-B messages and, the take-off context are processed in the same one-dimensional CNN. On the other side, the geographical context is pro-

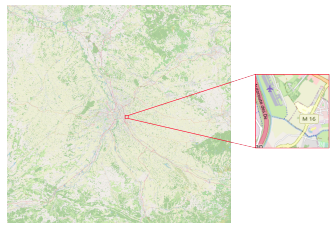

Figure 1. OpenStreetMap Tile extraction from aircraft's latitude and longitude

cessed by a two-dimensional CNN. Finally, the last part of the model is a fully connected network. It uses both flattened CNN outputs to produce probabilities for the three classes.

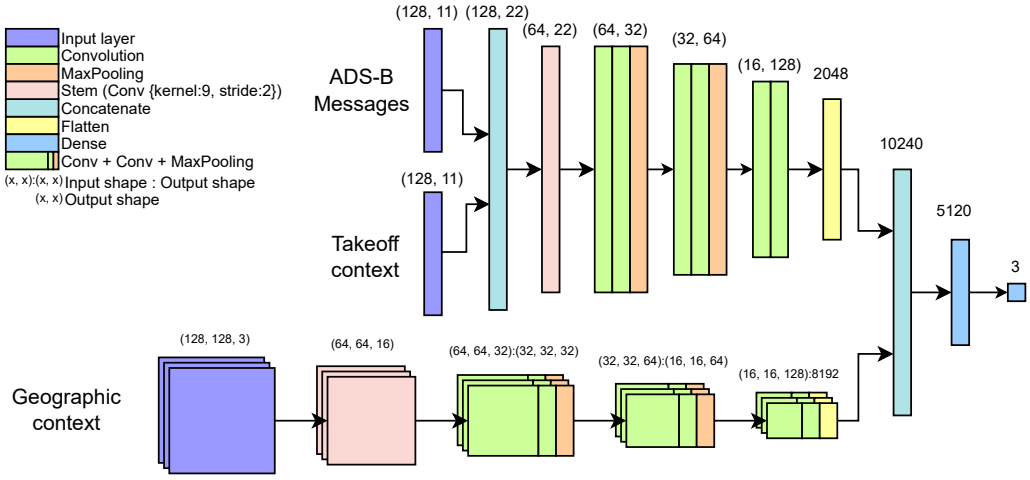

Figure 2. Model architecture

4. Results

In order to check the efficiency of each of our previously presented model's components, we started back with a simple CNN and added each of our components one by one. The following Table 1 shows the accuracy results of this ablation study.

Table 1. Ablation study

Model	Accuracy	Improvment
Bare-CNN	91.2%	
CNN with take-off context	93.5%	+2.3%
CNN with take-off and geographic context	95.1%	+1.6%
Add relative coordinates	96.2%	+1.1%

The results show that each of our components slightly improves the model's accuracy. We can see that the impact of the take-off context on the final prediction is important with an improvement of 2.3%. The use of relative coordinates shows an improvement of 1.1%. However, relative coordinates

are only relevant in combination with the map context. The absolute aircraft's position is still very important to know if the aircraft flies in a suspicious area. In the same way, the take-off context is not relative, in order to let the model know where the aircraft has taken-off.

5. Conclusion and future work

This article proposes a new deep-learning architecture designed in the field of detecting spoofing attacks in low-altitude ADS-B data. Our results are quite promising. We provided a tool that detects up to 96.2% of spoofing attacks around Toulouse. Moreover, the model could easily be re-trained in other cities by using the corresponding ADS-B and Geographic data. The reliability of our experiment shows that deep-learning methods can be used to detect spoofing attacks in the ADS-B protocol.

However, given the absence of readily available data on autonomous drones, our current model lacks the capacity to directly detect drones. In our future work, we plan to conduct experiments using simulated drone trajectories based on known drone specifications. This will allow us to assess the model's reliability in extreme cases, such as scenarios where a drone closely mimics the flight path of a SAMU helicopter. Then to improve our results we plan to increase our dataset quality by including take-off data for flights coming from the outside of the bounding box where, for instance, the take-off context is not available.

Author contributions

- Melvyn Pirolley: model conception, experiments, paper writing
- Other Authors: paper writing

Acknowledgement

This work was partially supported by the EIPHI Graduate School (contract "ANR-17-EURE-0002"). We also thank the mesocentre of Franche-Comté for the computing facilities.

Reproducibility statement

To reproduce the results presented in this paper, you can follow the instructions of the GitHub: https://github.com/DApIA-Project/Anomaly-Detection/tree/OpenSky

The dataset can be downloaded at: https://zenodo.org/doi/10.5281/zenodo.10050766

References

[1] Cătălin Dumitrescu, Marius Minea, Ilona Mădălina Costea, Ionut Cosmin Chiva, and Augustin Semenescu. "Development of an Acoustic System for UAV Detection". In: *Sensors* 20.17 (2020). ISSN: 1424-8220.

[2] Ralph Karam, Michel Salomon, and Raphaël Couturier. "A comparative study of deep learning architectures for detection of anomalous ADS-B messages". In: *2020 7th International Conference on Control, Decision and Information Technologies (CoDIT)*. Vol. 1. IEEE. 2020, pp. 241–246.

[3] Antoine Chevrot, Alexandre Vernotte, and Bruno Legeard. "CAE: Contextual auto-encoder for multivariate time-series anomaly detection in air transportation". In: *Computers & Security* 116 (2022), p. 102652.

[4] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention is all you need". In: *Advances in Neural Information Processing Systems*. 2017, pp. 5998–6008.