Anomaly detection in approach operations at Sao Paulo/Guarulhos International Airport

Lucas Coelho e Silva*, Mayara Condé Rocha Murça
Aeronautics Institute of Technology (ITA), Civil Engineering Department (IEI)

*** Corresponding author e-mail adress:** coelho@ita.br

PAPER ID: SIT231

ABSTRACT

The pursuit of stringent targets on operational safety and efficiency in an increasingly complex aviation system has been driving the development of novel analytics capabilities for more proactive aviation performance management. Anomaly detection in flight operations data is a prominent approach to delivering actionable information, as anomalies are often related to critical safety events or inefficient operations. In this paper, we apply machine learning techniques to aircraft surveillance data for offline anomaly detection and explanation in approach operations at Sao Paulo/Guarulhos International Airport (SBGR). First, we build an autoencoder classifier for the automatic identification of anomalous approach performance from flight trajectory data. Then, we extend our analysis with runway configuration and weather information to develop models for anomaly explanation. We found that the autoencoder classifier was able to detect operationally relevant anomalies, while the explanatory models provided novel insights about contributing factors to the anomalies identified. We learned that anomalous flight trajectories are more likely to be associated with landing operations on runway 27 under wind scenarios, with an increase in the odds ratio of 63% and 48% for tailwinds and headwinds, respectively. In addition, we also observed a positive association between anomalies and wind gusts situations.

Keywords: Air traffic management, Aviation safety, Anomaly detection, Trajectory data analytics, Machine learning.

1 INTRODUCTION

Over the past years, anomaly detection initiatives with flight operations data, be it airplane sensor data or Air Traffic Management (ATM) system data, have become more prominent, driven mainly by the safety-oriented culture in aviation and the pursuit of improved operational efficiency in an increasingly complex and evolved airspace. Anomalies are patterns on data inconsistent with the expected behavior (Chandola et al., 2009). They arise in non-normal flight situations and operations, and they are often related to conditions that may lead to unsafe scenarios or generate inefficiencies in the airspace. In this sense, the modeling and discovery of anomalies shed light on hazardous or inefficient operations and substantiate the development of new safety and operational policies and practices for airlines and ATM.

Previous research has mostly focused on the development of anomaly detection models with various statistical learning techniques, without addressing explanatory aspects regarding anomalous situations. On the other hand, the explanation of machine learning models across different fields has been deemed of significant importance as a way to provide trust and improved usability in autonomous decision-making systems (Degas et al., 2022).

In this paper, we apply machine learning techniques to aircraft surveillance data for offline anomaly detection and explanation in approach operations at Sao Paulo/Guarulhos International Airport (SBGR). First, we build an autoencoder classifier for the automatic identification of anomalous approach performance from flight trajectory data. Then, we extend our analysis with runway configuration and weather information to develop models for anomaly explanation. The autoencoder classifier was able to detect operationally relevant anomalies, while the explanatory models provided novel insights about contributing factors to the anomalies identified.

This paper is organized in the following way: Section 2 reviews the related literature while discussing current gaps, and Section 3 presents the methodological approach for modeling and identifying anomalies in our particular case. Section 4 presents and discusses the results, and Section 5 details the conclusions.

2 BACKGROUND AND LITERATURE REVIEW

2.1. Anomaly detection

Anomaly detection refers to the process of identifying valid and practically significant data structures incompatible with a notion of expected normal behavior. Within flight operations, the aviation subsystem that deals with the day-to-day operations of flights, in which the principal stakeholders are Air Traffic Management (ATM) and airlines/aircraft operators, anomaly detection initiatives are often related to the investigation of safety or efficiency-related issues. In this sense, the objects of interest are often operation-wise qualities that interfere in the system behavior, such as aircraft trajectories, from the ATM standpoint; or the identification of safety-related events and the unveiling of new hazards in the case of airlines. Nevertheless, the analytical techniques are often shared with anomaly detection approaches for condition monitoring and proactive maintenance of the airframe and engine, thence we include some of those efforts in our literature assessment.

The anomaly detection process is frequently classified in terms of the timeframes of the statistical learning process and model operation. First, the modeling process itself can happen offline, when the model training uses historical data on a batch processing schema, or online, often referred to as incremental learning, in which the statistical learning model gets updated as new data become available without the need to retrieve the previously used data. Second, the model operation itself can also be distinguished between the online and offline approaches. For the offline model operation approach, the anomaly detection process happens after the operations occur. For the online approach, on the other hand, the anomaly detection process happens in a real-time, streaming fashion as the operations develop. Finally, it is possible to have an offline learning model designed for online model operations, so as to construct an online/incremental learning approach that aids in offline model operations.

There are several approaches for formulating an anomaly detection problem as well as algorithms and techniques suitable for modeling it, and Chandola et al. (2009) presents a thorough general review. The anomaly detection problem does not have a universal solution, applicable to all cases. In practice, the existing techniques solve specific formulations of the general problem (Chandola et al., 2009). Over the literature, one finds several categories of anomaly detection methods, models, and tools. Here, we discuss the methods applicable to our methodological approach: reconstruction methods, and isolation methods.

Reconstruction methods rely on models that learn how to reconstruct normal data. When the model fails to reconstruct an observation - defined by comparing the reconstruction error to a previously tuned threshold -, that observation is considered an anomaly. Autoencoders, a particular type of neural network architecture, are commonly used within the reconstruction methods.

Isolation methods aim at separating an instance from the rest. It measures the susceptibility of each instance to be isolated, with the anomalies being those more easily isolated, based on the principle that anomalies are few and different than normal instances (Liu et al., 2008, 2012). The approach relies on a binary tree structure that isolates every instance via a recursive and random partitioning process that generates a path length metric for isolating each instance. By averaging the path lengths produced by several trees, one can identify anomalies corresponding to the shortest paths.

2.2. Anomaly detection in flight operations

One of the first initiatives for anomaly detection applied to the ATM domain is that of Matthews et al. (2013), in which the authors presented an approach for discovering operationally significant anomalies in flight track data generated from surveillance equipment. It is an extension of the studies in the Flight Operations Quality Assurance (FOQA) domain to ATM data conducted by Das et al. (2010), as it leverages the Multi-Kernel Anomaly Detection (MKAD) algorithm presented by the latter. The algorithm identified approximately 40 anomalous flights, which were then further analyzed. Domain experts confirmed operationally significant anomalies in 15 of these flights.

A different line of study is that of Murça (2018), which discusses anomaly detection via Conformal Prediction for identifying non-conforming trajectories within the proposed framework for characterization of air traffic flows based on flight trajectory data. The Conformal Prediction model presented better values for recall, precision, and F1-score when compared to K-Nearest Neighbors (KNN) and Gaussian Mixture Model (GMM).

On another approach, Deshmukh & Hwang (2019) presented TempAD, an unsupervised learning algorithm that uses temporal logic for anomaly detection for terminal airspace operations. It generates normal-flight parameter intervals that are easily interpreted and converted to natural language. Subsequently, Deshmukh et al. (2019) leverage the identified normal-flight parameter intervals and discuss an approach to identify precursors for the detected anomalies in surveillance data for terminal airspace operations. For that, the authors presented a supervised learning algorithm for precursor detection, Reactive TempAD.

Another effort is that of Olive & Basora (2019), which presented a methodology to analyze flight track data from Automatic Dependent Surveillance-Broadcast (ADS-B) and identify operationally significant anomalies. The authors obtained the principal flows in the airspace via trajectory clustering and used autoencoders for identifying anomalies.

While current literature discusses different scenarios in which anomaly detection is viable within the ATM context, it still lacks anomaly explanation efforts that could provide novel insights about contributing factors to the anomalies identified, in addition to providing trust and improved usability in autonomous systems. Furthermore, with the ongoing development of machine learning models and techniques, there are still opportunities for applying novel or previously unexplored modeling approaches with proper comparison to baseline methods or current practice. With the unprecedented application of isolation forests for anomaly detection in arrival trajectories, we contribute to the exploration of different modeling approaches for anomaly detection formulations with flight track data, while addressing their explanation to shed light on potential causal factors for the anomalies identified.

3 METHODOLOGICAL APPROACH

In this paper, we first conduct an offline anomaly detection learning process based on flight tracking ADS-B data regarding terminal airspace arrival operations at the Sao Paulo/Guarulhos International Airport (SBGR). Then, we augment our analysis with runway configuration and weather information to develop models for anomaly explanation. In this sense, our objective is to replicate a scenario in which the goal is the post-operation discovery of anomalous operations.

For this problem, we use two model categories for anomaly detection. First, we apply an Isolation Forest model for the unsupervised discovery of potential anomalies. Then, we train an autoencoder based solely on the normal flights. By calibrating an allowable reconstruction error threshold for normal operations, the model is able to classify between normal and anomalous instances. Based on the identified anomalies, we then build the explanatory models based on support contextual data and obtain metrics of feature importance.

3.1. Datasets

The flight tracking ADS-B dataset features $10,209$ flights between December 31, 2019, and January 31, 2020, at the Sao Paulo/Guarulhos International Airport (SBGR). It presents the flight information in a time series from which we extracted the following parameters along with each flight:

- Record timestamp;
- Latitude;
- Longitude;
- Altitude;
- Speed;
- Whether the record takes place in the terminal manoeuvring area.

Figure 1 depicts the horizontal trajectories of the analyzed flights.

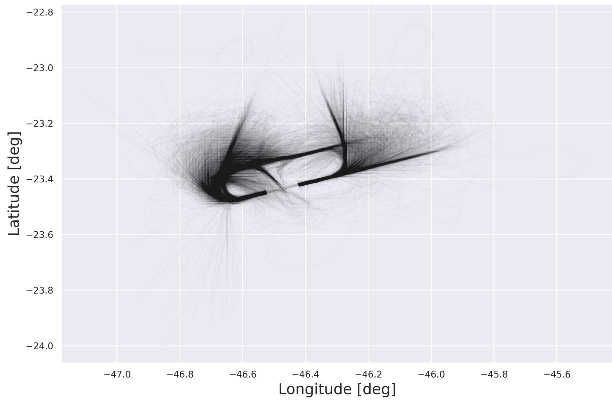

Figure 1 Horizontal trajectories of terminal airspace arrival operations at SBGR between December 31, 2019 and January 31, 2020.

The Meteorological Aerodrome Report (METAR) dataset contains information regarding winds, gusts, flight rules, presence of thunderstorms and so forth, for the same period.

To emulate the online model operation after training, we partition the data set temporally on the day which results in an approximately 80-20 data split. We use the first portion for training the anomaly detection models offline and the final 20% for replicating the model usage during real operations.

3.2. Anomaly detection

Model conception starts with the specification of a shared data-wrangling phase, common to both models. In this case study, we analyze approach procedures in SBGR terminal airspace, from 10,000 ft onwards. To ensure data consistency with the desired flight phase and scale, the wrangling processing subjects the flights through three filters. First, it narrows down the trajectories to the subsets within terminal airspace. Second, we discard all observations previous to the first instant a given flight reached 10,000 ft. To avoid climbing flights following a takeoff captured by the surveillance equipment, we also ensure a minimum value for the rate of descent. Finally, to avoid flights with few samples, we keep only the flights with more than five observations. The result of the wrangling phase is a set of flights in the desired scale and context for the development of the models.

Unsupervised learning: Isolation Forest For the preliminary anomaly discovery phase, we construct an unsupervised model using an Isolation Forest, or iForest. The goal is to obtain an initial distinction between normal and anomalous flights. Based on this categorization, the training process of the supervised model then considers only the flights labeled as normal.

For the application of the Isolation Forest model, we calculate summarizing metrics instead of directly using the wrangled time-series data for model building. There is thus an extra step comprising the transformation of the wrangled time-series data into the tabular metrics. For each flight, we obtain the following flight-wise summarizing metrics:

- Specific Total Energy (STE)'s total, average, and standard deviation;
- Specific Potential Energy Rate (SPER)'s total, average, and standard deviation;
- Specific Kinetic Energy (SKE)'s total, average, and standard deviation;
- Total time in terminal airspace;
- Latitude at 10000 ft;
- Longitude at 10000 ft;
- Last recorded latitude;
- Last recorded longitude;
- The total distance flown in terminal airspace;
- Hour of the day at the entrance in terminal airspace;
- Hour of the day at the exit of terminal airspace.

Based on the modeled data, we fit an Isolation Forest with 500 trees. Section 4.1. discusses the results.

Supervised learning: autoencoder For the supervised learning process, we train an autoencoder classifier. Autoencoders are a particular type of neural network architecture used for learning representation of data via reconstruction. It does so by encoding the data into a different feature space - also referred to as a compressed representation - before decoding it back to the original space. The underlying assumption in this approach is that if one successfully trains a model to reconstruct normal flights only, the execution of this model in an anomalous flight would result in a higher reconstruction error, enabling the classification of the flight as an anomalous instance.

As in the iForest training, the starting point is the wrangled time-series data. Instead of constructing summarizing metrics from the data, however, as performed in the Isolation Forest model training, the goal here is to use a representation as close as possible to the sequential data. Among the parameters available, we select those representatives of the aircraft trajectory and positioning: latitude, longitude, altitude, heading, and speed.

Nevertheless, one of the challenges of anomaly detection in the aviation domain is the varying number of samples between each flight. This property of the data often requires preprocessing steps that resample and reshape the data, which is the case for the application of the autoencoder. Therefore, to unify data dimensions, we resample the time series of each individual flight with a fixed number of observations. The next step is to transform the time series of the four parameters into a single high-dimensional vector. To achieve this, we rearrange the data interspersedly. Each flight can be represented in the high dimensional space as a vector of the following shape:

$$x = [p0_{t_0}, p1_{t_0}, \ldots, pn_{t_0}, \ldots, p0_{t_m}, p1_{t_m}, \ldots, pn_{t_m}]$$

where n is the number of parameters p and m is the number of time instances t.

Finally, the data goes through a scaling process. The scaling process contributes to the proper assessment of feature importance by the autoencoder neural network structure. The data is scaled via the Robust Scaler. It centers the data by removing the median and scales the result based on a specified quantile range, making the scaling process robust to outliers. Since the autoencoder models the behavior of normal flights, the contamination of the training data set with anomalous flights becomes relevant. The Robust Scaler is used precisely to tackle these effects of data set contamination.

The training process of the autoencoder classifier relies on two fundamental steps. The first step concerns the development of a suitable neural network architecture capable of reconstructing the data with a low reconstruction error - in this case, the squared error between the norm of the actual flight vector and the one predicted by the autoencoder. The second training aspect refers to the definition of a reconstruction error threshold that enables classification between normal and anomalous instances.

For training a suitable neural network architecture, we further subdivide the normal data set - as flagged by the Isolation Forest model - into training and validation subsets at an 80:20 ratio.

The autoencoder neural network architecture consists of five hidden layers of 500, 300, 2, 300, and 500 neurons, respectively. For each hidden layer, we use the Rectified Linear Unit (ReLU) activation function. The neural network model was built on top of the open-source machine library scikit-learn, written in Python (Pedregosa et al., 2011). Hyperparameter tuning was not required as a simple neural network configuration sufficed for reconstructing the normal data.

For defining the reconstruction error threshold, the model performance metrics are evaluated under the expected operational context. As a reference, we assume the expected number of daily flights as 500 and a capacity constraint of being able to investigate ten flagged anomalies in the same period. Section 4 discusses the results regarding the anomaly detectioin model.

3.3. Anomaly explanation

Based on the class labels - anomalous or normal - obtained for each flight during model operation, our next goal is to investigate the relationship between the operational context to the identified class. For that, we build two explainable classifiers (logistic regression and random forests) that predict the anomalous quality of a flight, as per identified by the autoencoder, in accordance to operational information extracted from METAR data and from the surveillance data itself.

For each flight assessed during the model operation phase - hence 20% of the original data set - we process the corresponding METAR information to obtain data on visibility, thunderstorms, flight rules, wind gusts, and cross and aligned wind components referenced to the runway, the latter inferred via the heading parameter on the surveillance data set. Based on the wrangled support data, we explore two supervised learning classifiers: a random forests classifier, and a logistic regression classifier. For both modeling approaches, we divide the data set into train and test subsets for validating the modeling process in terms of the model capability of predicting the anomaly class of a flight given the operational context. Because the anomaly classification problem is an imbalanced one, we weigh the models referenced to the class given the heuristics proposed by King & Zeng (2001). After validating the modeling process regarding the accuracy and recall on the test subset, we fit the models on the complete data set and explore the explanatory metrics.

For the random forests classifier, the goal is to assess feature importance to understand the operational factors contributing to the model predictions. We construct a random forests classifier with 500 trees using a split rule based on the Gini impurity, maximum depth of 15 and 3 as the number of predictors randomly sampled as split candidates. Then, we compute the importance metrics - in terms of mean decrease in impurity for each feature - and their standard deviations based on the values of each tree within the forest.

For the logistic regression classifier, we first augment the wrangled support data with features regarding the interaction terms between each operational factor. After fitting the model, we then analyze the coefficients of the independent variables to assess the changes in the odds ratio for each operational term.

4 RESULTS AND DISCUSSION

4.1. Anomaly detection model training

The Isolation Forest trained with 500 trees identified 669 anomalies within the 8073 flights selected for model training - approximately 8.3%

of the data. Figure 2 presents the two principal components of the evaluated data, obtained via Principal Component Analysis (PCA), hued by the operation category as identified by the Isolation Forest.

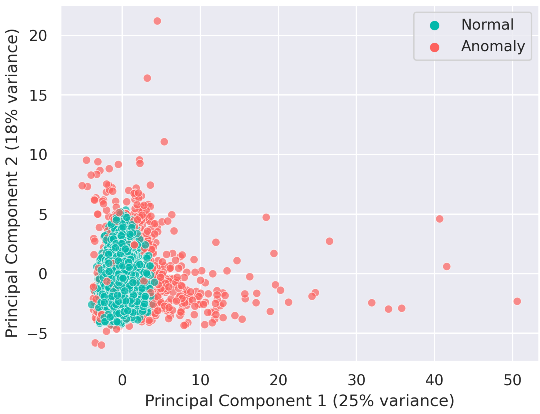

Figure 2 Principal components, hued by the operation category (normal/anomalous).

Based on the flights labeled as normal by the forest, the autoencoder anomaly detector was trained. We assess its reconstruction error for the training and validation subsets, as well as for the flights flagged as anomalous by the Isolation Forest, as shown in Figure 3. The figure also shows that the reconstruction error distribution of the flights in the training and validation subsets are similar, indicating the model did not overfit the data. In addition, there is a statistically significant difference between the error distribution for the normal data and the ones labeled as anomalous, with higher reconstruction error values in the latter. This indicates that the model did not underfit the data and learned a distinction between normal and anomalous instances.

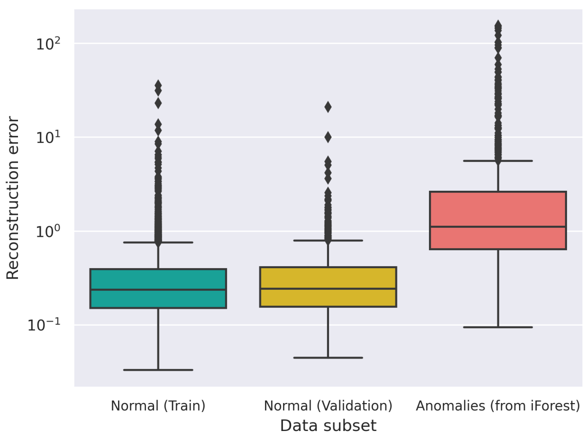

Figure 3 Boxplot of the autoencoder reconstruction errors.

Performance metric evaluation happens via bootstrapping. We first calculate the reconstruction errors for observations in the training, validation, and iForest-flagged anomalies subsets. Next, we define a range of candidate thresholds. For each potential threshold value, we sample, with replacement, 500 flights and evaluate the number of flagged anomalies and the value for each performance metric. Figure 4 shows the results when experimenting 500 times for each potential threshold. It displays the value for recall, precision, F1-Score, and false positive rate (FPR) for each threshold value. The solid lines represent the mean value for the performance metrics, while each shaded region covers the average ± one standard deviation. Additionally, the green line presents the average number of flagged anomalies per 500 flights, given a threshold.

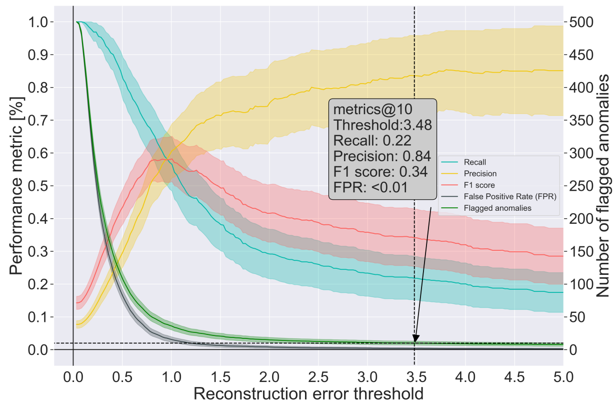

Figure 4 Autoencoder performance metrics obtained via bootstrapping in 500 rounds, considering 500 landings per day and processing capacity of k=10 flagged anomalies.

The selected threshold value was the one that resulted in, on average, ten flagged anomalies per 500 flights - i.e., the analysis capacity in this proposed scenario. Finally, given this threshold of 3.48, we calculate the metrics of 0.22 for recall, 0.84 for precision, 0.34 for F1-score, and a false positive rate of less than 0.01.

4.2. Anomaly detection model operation

For replicating the model usage during real operations, we use the 20% of the flights originally set apart and hence did not go through any processing step during model conception, as discussed in Section 3.1.

Out of the 2136 flights, the autoencoder classifier flagged 47 - or 2.20% - as potentially anomalous, throughout the seven days of operation. Figure 5 displays the number of flagged anomalies per day.

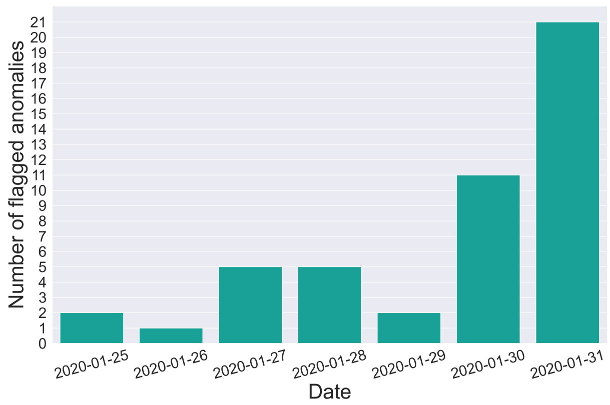

Figure 5 Number of flagged anomalies by the autoencoder during simulated model operation.

Figure 6 highlights the horizontal profiles of three flights flagged as potentially anomalous by the autoencoder. We see that the flights are indeed associated with operationally relevant scenarios that can impact safety/efficiency, presenting either a holding pattern or multiple landing attempts. This is reinforced by Figure 7, which displays the vertical profiles of the same flights along the 95th percentile region shaded in gray.

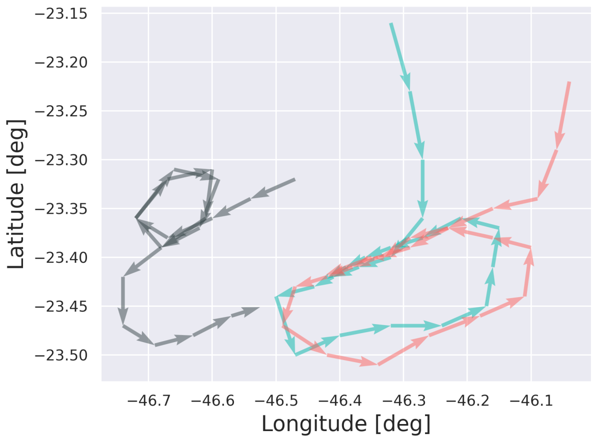

Figure 6 Horizontal profile of three flights flagged as anomalous by the autoencoder.

4.3. Anomaly explanation

To assess the performance of the explanatory models, we first consider the confusion matrices as well as the metrics of accuracy and recall. Table 1 displays the confusion matrix for the random forests classifier, with an accuracy of 0.86 and a recall of 0.75. Table 2 displays the confusion matrix for the logistic regression classifier, with an accuracy of 0.71 and a recall of 0.62. Both results indicate that the explanatory models were able to link anomalous situations to the environmental and contextual conditions.

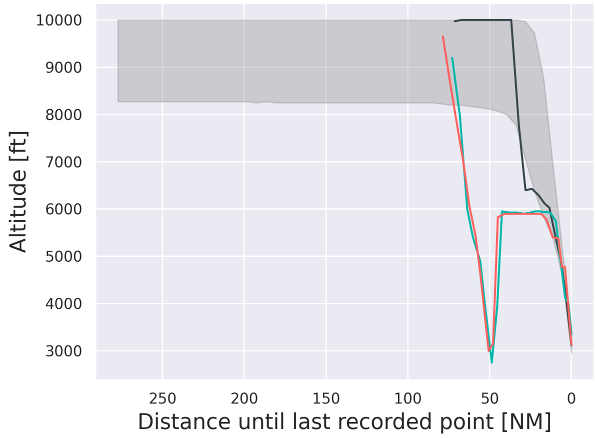

Figure 7 Vertical profile of three flights flagged as anomalous by the autoencoder.

Table 1 Random forests classifier confusion matrix.

		Prediction		
		Normal	**Anomalous**	**Total**
Actual	**Normal**	1800	289	2089
	Anomalous	12	35	47
	Total	1812	324	

Table 2 Logistic regression classifier confusion matrix.

		Prediction		
		Normal	**Anomalous**	**Total**
Actual	**Normal**	1481	608	2089
	Anomalous	18	29	47
	Total	1499	637	

In terms of assessing the contributing factors for the anomalous scenarios, we first evaluate the feature importance metrics regarding the random forests classifier, as shown in Figure 8. The cross-component of the wind, the headwind, and the tailwind values are the top three features contributing to the predictions of the model beyond the reference dashed line of equally contributing features.

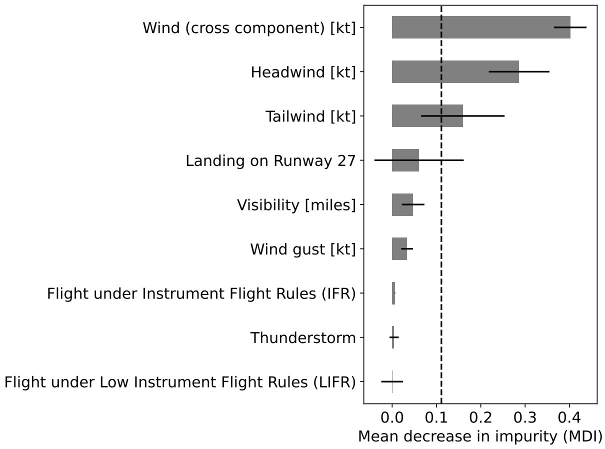

Figure 8 Random forest classifier feature importances.

After evaluating the feature importance of the random forests classifier, we analyze the logistic regression model coefficients in terms of percent change in the odds ratio, as shown in Figure 9. According to the values, we learn that the anomalous situations are associated with landing operations on runway 27 under wind scenarios, with an increase in the odds ratio of 63% and 48% for tailwinds and headwinds, respectively. In addition, we also see a positive association between anomalous scenarios and wind gusts and also thunderstorms accompanied by wind gusts. On the other hand, flights under Instrument Flight Rules (IFR) and increased visibility with flights under IFR showed a reduction in the odds ratio.

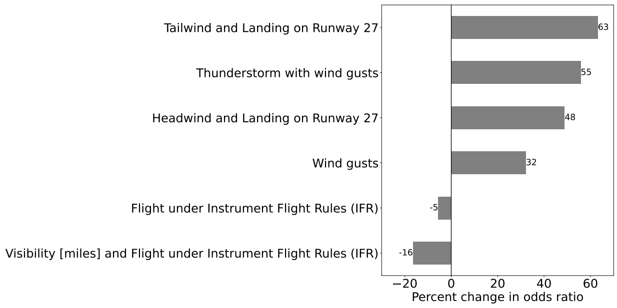

Figure 9 Percent change in the odds ratio for selected parameters of the logistic regression model.

The constructed pipeline demonstrates its applicability regarding the objective of post-operation detection of anomalous approach performance as well as discovery of contributing factors to anomalies. From the operational perspective, the results contribute to the development of novel capabilities that can support performance analysis and monitoring processes, which in turn may be used in training programs or aid in subsequent developments of operational procedures. The pipeline enables efficient highlighting of flights that deviate from normality. In addition, it does so without resorting to a previous mapping of what to look for: operationally relevant scenarios surfaced without the need to resort to heuristic rules that specify every possible case.

5 CONCLUSION

In this paper, we explored the development of machine learning models for anomaly detection and explanation in flight operations trajectory data based on aircraft surveillance data regarding approach operations at Sao Paulo/Guarulhos International Airport (SBGR). We considered first an autoencoder classifier for the identification of anomalous approach performance from flight trajectory data, trained according to normal flights labeled by an Isolation Forest unsupervised learner. Subsequently, the analysis was extended with runway configuration and weather information to develop models for anomaly explanation.

The autoencoder anomaly detector was able to detect operationally relevant anomalies, such as holding patterns and go-around maneuvers. Additionally, the construction of explanatory logistic regression and random forests classification models enabled the association of anomalous situations with operational factors. For instance, we learned that the anomalous situations are more likely to be associated with landing operations on runway 27 under wind scenarios, with an increase in the odds ratio of 63% and 48% for tailwinds and headwinds, respectively.

For future work, there is a need for extending tools that support the evaluation of potentially anomalous flight, such as the explanatory approach provided in this paper, with the assessment of the indication correctness while aiding the discovery process of unknown hazards in the data. Finally, another research direction is the development of online models for the real-time identification of anomalies from flight track data.

References

Chandola, V., Banerjee, A. & Kumar, V. (2009). Anomaly detection: A survey, ACM Computing Surveys 41(3), 58.

URL: *http://doi.acm.org/10.1145/1541880.1541882*

Das, S., Matthews, B. L., Srivastava, A. N. & Oza, N. C. (2010). Multiple kernel learning for heterogeneous anomaly detection: algorithm and aviation safety case study, KDD '10: Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data miningpp. 47–56.

Degas, A., Islam, M. R., Hurter, C., Barua, S., Rahman, H., Poudel, M., Ruscio, D., Ahmed, M. U., Begum, S., Rahman, M. A., Bonelli, S., Cartocci, G., Di Flumeri, G., Borghini, G., Babiloni, F. & Aricó, P. (2022). A Survey on Artificial Intelligence (AI) and eXplainable AI in Air Traffic Management: Current Trends and Development with Future Research Trajectory, Applied Sciences 12(3), 1295.
URL: *https://www.mdpi.com/2076-3417/12/3/1295*

Deshmukh, R. & Hwang, I. (2019). Anomaly Detection Using Temporal Logic Based Learning for Terminal Airspace Operations . AIAA Scitech 2019 Forum, American Institute of Aeronautics and Astronautics San Diego, California.
URL: *https://arc.aiaa.org/doi/10.2514/6.2019-0682*

Deshmukh, R., Sun, D. & Hwang, I. (2019). Data-Driven Precursor Detection Algorithm for Terminal Airspace Operations, Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019)p. 7.

King, G. & Zeng, L. (2001). Logistic regression in rare events data, Political Analysis 9, 137–163.

Liu, F. T., Ting, K. M. & Zhou, Z.-H. (2008). Isolation Forest . 2008 Eighth IEEE International Conference on Data Mining, IEEE Pisa, Italy pp. 413–422.
URL: *http://ieeexplore.ieee.org/document/4781136/*

Liu, F. T., Ting, K. M. & Zhou, Z.-H. (2012). Isolation-Based Anomaly Detection, ACM Transactions on Knowledge Discovery from Data 6(1), 1–39.
URL: *https://dl.acm.org/doi/10.1145/2133360.2133363*

Matthews, B., Srivsatava, A. N., Schade, J., Schleicher, D. R., Chan, K., Gutterud, R. & Kiniry, M. (2013). Discovery of Abnormal Flight Patterns in Flight Track Data . 2013 Aviation Technology, Integration, and Operations Conference, American Institute of Aeronautics and Astronautics Los Angeles, CA.
URL: *http://arc.aiaa.org/doi/10.2514/6.2013-4386*

Murça, M. C. R. (2018). Data-driven modeling of air traffic flows for advanced Air Traffic Management Thesis Massachusetts Institute of Technology, Accepted: 2019-02-14T15:22:40Z.
URL: *https://dspace.mit.edu/handle/1721.1/120378*

Olive, X. & Basora, L. (2019). Identifying Anomalies in past en-route Trajectories with Clustering and Anomaly Detection Methods, Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019)p. 11.

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. & Duchesnay, E. (2011). Scikit-learn: Machine learning in Python, Journal of Machine Learning Research 12, 2825–2830.