

Predicting Adverse Events and their Precursors in Aviation Using Multi-Class Multiple-Instance Learning

Marc-Henri Bleu Laine*, Tejas G. Puranik†, Dimitri N. Mavris‡
Georgia Institute of Technology, Atlanta, GA, 30332

Bryan Matthews§
KBR Inc., NASA Ames Research Center

In recent years, there has been a rapid growth in the application of machine learning techniques that leverage aviation data collected from commercial airline operations to improve safety. Anomaly detection and predictive maintenance have been the main targets for machine learning applications. However, this paper focuses on the identification of precursors, which is a relatively newer application. Precursors are events correlated with adverse events that happen prior to the adverse event itself. Therefore, precursor mining provides many benefits including understanding the reasons behind a safety incident and the ability to identify signatures, which can be tracked throughout a flight to alert the operators of the potential for an adverse event in the future. This work proposes using the multiple-instance learning (MIL) framework, a weakly supervised learning task, combined with carefully designed binary classifiers leveraging a Multi-Head Convolutional Neural Networks-Recurrent Neural Networks (MHCNN-RNN) architecture. The classifiers are then combined to perform a multi-class task, which enables the prediction of different adverse events for any given flight and the identification of their precursors with minimum post-processing. Results obtained show that the MHCNN-RNN is able to accurately forecast high speed and high path angle events during the approach, and that it is also capable of determining the aircraft's parameters that are correlated to these events. The identified parameters can be considered precursors to the events and may be studied/tracked further to prevent these events in the future.

<i>CNN</i>	=	Convolutional Neural Network
<i>DT – MIL</i>	=	Deep Temporal Multiple-Instance Learning
<i>FDR</i>	=	Flight Data Recorder
<i>FN</i>	=	False Negative
<i>FP</i>	=	False Positive
<i>FOQA</i>	=	Flight Operations Quality Assurance
<i>GRU</i>	=	Gated Recurrent Unit
<i>MHCNN</i>	=	Multi-Head Convolutional Neural Network
<i>MIL</i>	=	Multi-Instance Learning
<i>QAR</i>	=	Quick Access Recorder
<i>RNN</i>	=	Recurrent Neural Network
<i>TN</i>	=	True Negative
<i>TP</i>	=	True Positive

I. Introduction

The aviation industry brings tremendous amount of social and economic benefits. It has been observed the size of the air transportation industry has doubled every 15 years [1], and was expected to continue growing [2] prior to

*Graduate Research Assistant, Aerospace Systems Design Laboratory, School of Aerospace Engineering, AIAA Student Member

†Research Engineer II, Aerospace Systems Design Laboratory, School of Aerospace Engineering, AIAA Senior Member

‡S.P. Langley NIA Distinguished Regents Professor and Director of Aerospace Systems Design Laboratory, School of Aerospace Engineering, AIAA Fellow

§Data Science Research Engineer

COVID-19. In 2018 alone, airlines around the world carried a total of 4.3 billion passengers while the total global economic impact of the industry reached USD 2.7 Trillion in 2016 [1]. Furthermore, even though the industry continues to grow and more people are flying everyday, aviation safety has improved over the past decades as seen on Figure 1.



Fig. 1 Commercial Aviation Fatalities from FY96 to FY19 [3]

The number of fatalities for commercial air carriers decreased from 81 per 100 million persons on board in 1996 to 0.6 in 2019, which is below the target rate set by the Federal Aviation Administration (FAA). This reduction is the result of the efforts undertaken by agencies such as the National Aeronautics and Space Administration (NASA), the Federal Aviation Administration (FAA), the National Transportation Safety Board (NTSB), and others. In particular, these efforts led to better certification standards, better operating procedures, and decision-support systems [4]. However, it is important to keep reducing the accident rate even further so that we do not observe a rise in the number of accidents given the industry's expected continued growth [2, 5]. In order to significantly improve safety, the aviation industry has been moving towards a proactive approach to safety assessment and vulnerability identification, which consists of characterizing potential risks in terms of anomalies or deviations from nominal operations and the precursors to adverse events. This knowledge can be leveraged to increase awareness of emerging vulnerabilities amongst the operators and be incorporated into automated monitoring tools to flag the risks before they result in a near-miss, incident or accident [6].

Extensive data collection and advancement in data-mining methodologies are key enablers to proactive risk management in aviation [6, 7]. Airline programs such as the Flight Operations Quality Assurance (FOQA) have enabled the creation of large, and heterogeneous data sets. FOQA is a voluntary safety program that is designed to make commercial aviation safer. Data is collected using devices such as Quick Access Recorder (QAR) or directly from Flight Data Recorder (FDR) [8]. Traditional techniques of flight data analysis have focused on a continuous cycle of data collection from on-board recorders, retrospective analysis of flight-data records, identification of operational safety exceedances, design and implementation of corrective measures, and monitoring to assess their effectiveness. Airlines work with the FAA to reduce and eliminate safety risks, and flight safety divisions within airlines generally use FOQA data to perform exceedence and statistical analyses. Exceedence analysis consist of setting specific limits to the recorded airborne data so that particular parameters that fall outside of the normal operating conditions can be flagged [8]. The level of exceedence can be programmed for different severities of events. These profiles are used to create distributions of various criteria, which enables airlines to evaluate flight risk levels and trend known vulnerabilities over time [8]. A validation step at the end of each analysis is performed to determine the nature of corrective actions required and to store valid events in databases for analyzing trends [8].

Research efforts have been made towards advancing the development of data-driven methodologies, applying data science techniques, and using modern machine learning, including deep learning algorithms, in the context of

aviation safety [9–16]. Majority of the data mining effort in aviation is directed at detecting anomalies in aviation data [13, 15, 17–22] and leveraging unsupervised learning techniques due to lack of labeled data. Data mining has also been applied towards predictive maintenance in aviation [9] and predicting future trajectory states [23]. While identifying anomalies is important, it is also critical to investigate the causal factors or precursors to these anomalies or other safety events in order to understand them better and prevent them in the future. Recent literature has focused on identifying precursors to safety events and anomalies using FOQA data [10, 11, 24]. Precursors can be defined as any event that are correlated to a safety incident and occurs prior the incident itself [11]. They are useful for forecasting safety events, and provide insights to why the event happened. Knowing the precursors can then be used to initiate actions to avoid events from occurring [24]. Therefore, being able to identify and monitor precursors is an important step towards proactive safety enhancement in aviation operations. One of the main challenges in identification of precursors observed in literature is the lack of subject-matter-expert validated labels for the identified anomalies and safety events. To account for the lack of labels usually encountered, this paper proposes the application of a weakly-supervised learning technique called multi-class multiple-instance learning (MIL) [25] to detect different multiple adverse events and their precursors.

Considering the above observations, the main aim of this paper is the development of a methodology that leverages highly dimensional aviation data to predict multiple adverse events and discover their precursors. This work will use a flight-level label called a bag-label in MIL terminology to predict the anomalies at the flight-level and identify their precursors along with their occurrence in-time during a flight. The current framework uses a deep learning model constructed as multi-headed convolutional neural network (CNN) where each flight sensor gets its own CNN and the multi-headed architecture is designed to work with the MIL framework. Successful implementation of the methodology uses the flight’s data to 1) predict adverse events, and 2) determine the precursors to the predicted adverse events. One major benefit from this framework over current approaches will be the extension of binary classifiers to perform multi-class predictions and the ability to retrieve the precursors with little to no post-processing. Additional benefits will include transferability of the model due its multi-head architecture, knowledge discovery to help analysts find root causes of safety events faster, and relative simplicity of the model to provide better transparency and explainability (compared to previous architectures).

II. Background

This section presents recent work conducted related to anomaly detection, precursor mining, and other applications of deep learning models in aviation safety. The data source utilized for this work is also presented along with the definition of multiple-instance learning and its accompanying assumptions.

A. Previous Work

Focusing on precursors to anomalies is important as it allows identifying potential causes to safety hazards. Recent work shows a growing interest in detecting precursors in various domains. Multiple techniques have been explored and the most relevant ones are summarized in this subsection.

Yue Ning et al. [26] presented an approach for precursor identification using nested Multi-Instance Learning (n-MIL). In their work, they forecast societal events in different cities. At the instance level, the probability of an article published on a given day is modeled using a simple logistic function. These probabilities are then aggregated over a day, and finally the probabilities for different days are aggregated together up to a certain number of days before the event, creating the nested structure. The authors explain that the probability of a news article on a given day can be used to estimate how related an article is to the target. Since the articles with higher probabilities influenced the classifier decision, they are likely to be precursors of the predicted societal event.

Janakiraman et al. [11] proposed to use a Deep Temporal Multiple-Instance Learning (DT-MIL) framework, which combines Multiple-Instance Learning and recurrent neural networks, in this case a gated recurrent unit (GRU), to mine precursors in FOQA data. In this approach, the individual time steps are considered low-level instances and the whole flight is considered a bag. The labels (occurrence of adverse event) are given at the bag level but the methodology takes advantage of MIL and uses the low-level instances to correctly predict the bag label, which allows to infer the instance-level label. The time-steps at which the probability of a safety event occurring is greater than a defined threshold are retrieved. The region of time for which the probability (called the precursor score) of a safety event occurring is high is then analyzed during post processing, where each feature is perturbed one at a time. The precursors are identified by finding the features whose perturbations had more significant impacts in reducing the precursor score. The DT-MIL model is also referred to as a newer version of the Automatic Discovery of Precursors in Time Series

(ADOPT) *, which is different from the architecture presented in [12].

Ackley et al. [10] have used a sequential backward selection technique along with Random Forest classification models for predicting Unstable Approach adverse events. They have identified the critical parameters using a cumulative feature importance score and grouped them into various categories of parameters (such as energy-related, configuration-related, etc.) that contribute significantly towards the identification of the adverse event. Their analysis is conducted at fixed altitudes above the event detection trigger altitude of 1000 feet above touchdown. Similarly, Lee et al. [27] have also used a Random Forest algorithm to identify precursors to two different aviation safety events using supervised learning. The normalized precursor score is obtained using a Gini importance for all parameters contained in the classification model.

Melynk et al. [28] have proposed a framework for detecting precursors to aviation safety incidents due to human factors based on Hidden Semi-Markov Models. They performed an empirical evaluation of their models against traditional anomaly detection algorithms and demonstrated better performance on synthetic and flight simulator data. Mangortey et al. [29] used a variety of clustering techniques to identify clusters of nominal operations and subsequently determine important parameters that differentiate outliers from those nominal clusters. Despite presenting interesting insights, their method is still unsupervised and lacks validation.

B. Dashlink Flight Data

The proposed framework is demonstrated using a publicly available data set obtained from NASA’s DASHlink website, which is a collaborative sharing network for researchers in the Data Mining and Systems Health Management field[†]. Flight data were recorded from a single type of regional jet operating in commercial service over a three-year period. The data contains detailed aircraft dynamics, system performance, and other engineering parameters but are de-identified such that it cannot be traced back to a particular manufacturer or airline. Since this data set is not part of any airline’s FOQA program, additional preprocessing is required to create FOQA-like flags and label safety events for individual flights, where the labeling was created by using domain-based rules. The definitions of resulting adverse events are presented in Table 1.

Table 1 Adverse Events Labeling

Adverse Event	Comments
High Speed in Approach	Flagged at 1,000 ft
Low Speed in Approach	Flagged at 1,000 ft
High Rate of Descent in Approach	Flagged between 1,000 - 500 ft
High Bank in Approach	Flagged between 1,000 - 400 ft
High Path Angle in Approach	Flagged at 1,000 ft
Low Path Angle in Approach	Flagged at 1,000 ft
Deviation from Localizer	Flagged between 1,000 - 500ft
Deviation below Glideslope	Flagged between 1000 - 500ft
Flaps Late Setting at Landing	None

Each of these events are characterized by a severity level ranging from 1 to 3, with 3 being the most severe. For this work, only flights without any safety events and the ones with safety events with a severity level 3 will be considered. This will allow limited overlaps between normal and abnormal operations. Moreover, to ensure a multi-class problem instead of a multi-label one, flights with multiple adverse events were omitted. The high speed in approach, and the high path angle in approach events were selected for this work because of their higher frequency in the data set. It is noted that the methodology developed in this work is applicable to any similar data set containing time series data of aircraft parameters and the known occurrence of an event.

*<https://github.com/nasa/ADOPT>

[†]<https://c3.nasa.gov/dashlink/resources/?page=3&sort=-created&type=28>

C. Multiple-Instance Learning

The multiple-instance learning framework is used when labels are only available in sets called bags [30], and each bag contains many instances. The learning task is therefore supervised, but since the labels are not provided for each instance, we say that the task is weakly supervised. This framework is of interest because it alleviates the burden of weak supervision, which is common to many fields as labelling data is costly. In particular, it is useful in the context of aviation because of the lack of labeled data. In this context, we can consider the bag to be a flight and the instances to be the time-steps of the features. Furthermore, positive bags are flights that had a safety event while negative bags are the ones with no events recorded. The standard MIL assumption, which is used in this work, states that all negative bags contain only negative instances, and that all positive bags contain at least one positive instance [30]. This means a flight experiencing an event have at least one abnormal time-step that caused the whole flight to be considered a positive bag. Flights can later be divided into training, validation, and testing sets, and the performances of a chosen model can then be evaluated on those two subsets.

More specially, the learning task for this work is a classification problem. In a classification application, the task is to use given data and assign it to one of multiple predefined classes. In particular, given time-series from sensor data, we can assign the series a label such as a safety event [13]. MIL classification task can be performed at the bag and instance level. However, the methodology developed in this work will use the bag label to infer the instance label of each feature at each time-step. In other words for example if we classify a bag as positive, the algorithm will determine (infer) the time instances that enabled the classification of the bag.

III. Methodology

The Intelligent Methodology for the Discovery of Precursors of Adverse Events (IM-DoPE) will be described in this section. The four steps of the methodology as seen on fig. 2 are the data processing, the development of the model, the extraction of precursors and the derivation of precursor scores. These steps allow for the discovery of precursors related to adverse events of interest, which can be used to provide insights into the cause of these events.




Fig. 2 Intelligent Methodology for the Discovery of Precursors of Adverse Events

A. Data Processing

The DASHlink sample flight data is recorded at 1 Hz frequency, therefore no effort was placed in re-sampling the data set at the same frequency, as this is a sufficiently high frequency. The data set had already been cleaned prior

to accessing it as no missing data was observed. Subsequent tasks mainly focused on selecting events of interest, dimensionality reduction, data interpolation, and data formatting.

1. Event Definition and Selection

As previously stated, 2 adverse events were selected for this work due to their higher frequency in the available dataset. These events are pre-defined and characterized by the deviation of certain aircraft parameters from accepted nominal behaviors. Fig 3, shows the counts of nominal flights, and adverse flights that experienced one of the two events. As seen on the figure, the high speed event was the most common.




Fig. 3 Flight Counts for Nominal and Adverse Flights

2. Feature Selection

The dimensionality of the data was reduced by leveraging the correlations among the features. Certain parameters (e.g. computed airspeed and true airspeed) are different but convey similar information. This is due to the redundancy of parameters, the aircraft's physics, and the derivation of parameters [10]. Using the data, a correlation matrix was created. For each pair of features, the Pearson's correlation coefficient is computed. The coefficient is a measure of the strength of the association between the features in the pair[‡]. The Pearson correlation coefficient for two random variables X and Y is given by:

$$\rho = \frac{\sigma_{XY}^2}{\sqrt{\sigma_{XX}^2 \sigma_{YY}^2}} \quad (1)$$

The correlation coefficient ρ is bounded between $-1 < \rho < 1$, and values closer to 1 signify strong positive correlation while values closer to -1 suggest strong negative correlation. In eq. 1, σ_{XX}^2 is the sample variance for the variable X , and σ_{YY}^2 the one for the variable Y , and σ_{XY}^2 the co-variance between X and Y . Since features are considered highly correlated if they have a correlation coefficient higher than 0.9 [31], a threshold of $\rho = 0.90$ was set so that for a given pair of features, with the absolute value of their correlation coefficient greater or equal to the threshold, one of the correlated features was removed. Note that the absolute value of the correlated coefficient between feature pairs was taken since feature can be highly negatively correlated. For this work, only non highly correlated continuous variables were used as the input feature space. In addition to the correlation-based feature selection, trivial precursors were removed. This step was performed to avoid observing one of the speed parameters as precursors to the high speed event.

[‡]Data Analysis - Pearson's Correlation Coefficient: <http://learntech.uwe.ac.uk/da/Default.aspx?pageid=1442>

3. Data Re-sampling

All the events of interest for this work are flagged at a 1,000 ft above touch-down. Given the large number of flights and the differences between each of them, re-sampling the data was necessary to ensure uniformity across all flights. Thus, starting from a distance of 20 nautical miles away from the 1,000 ft mark, flights were re-sampled at every quarter nautical miles. Each flight therefore contains 81 data points or time steps. A Python code was developed to interpolate each feature to retrieve their values at every quarter mile distance, yielding a data set similar to table 2 for each flight. Fig. 4 shows two examples of interpolated data. The original flight data contained 281 points for each feature while the interpolated data only contain 30 points.




Fig. 4 Example of Altitude and Ground Speed Interpolations

Table 2 Example Results of Interpolating Parameters For a Given Flight

Distance away from 1,000 ft	Feature 1	Feature 2	...	Feature d
20	X	X	...	X
19.75	X	X	...	X
$\vdots$	X	X	...	X
0	X	X	...	X

4. Data Formatting

In practice, modern algorithms such as Convolutional Neural Networks (CNNs), and Recurrent Neural Networks (RNNs) usually take as inputs 3 dimensional tensors. In order to account for this, data from each flight, as specified in table 2, is concatenated together into a tensor as shown on fig. 5, with $X_{i,j}^f$ referring to feature j for flight f at time step i .

The data takes a 3 dimensional form where the first dimension (N) corresponds to the batch size or the number of flights that the deep learning algorithm will process at once. The second dimension (L) is the number of time-steps in each flight equal to 81 for this work. Finally the last dimension (D) refers to the number of features/aircraft parameters which corresponds to 58 continuous variables.

B. Model Architecture

The model architecture was chosen to take advantage of the Multiple-Instance Learning (MIL) framework, the feature extraction capabilities of Convolutional Neural Networks (CNNs) [32, 33], and the temporal pattern recognition ability of Recurrent Neural Networks (RNNs) yielding a MHCNN-RNN architecture. This architecture choice allows the aircraft's parameters to be initially processed individually by the MHCNNs, which yields feature maps. The complex




Fig. 5 Flight Data Reshaping

(time-dependent) correlations between the feature maps are then learned by the RNN. The model was developed using PyTorch [34], a high performance Python deep learning library.

1. Architecture Selection

The MHCNN-RNN architecture combines a Multi-Head CNN (MHCNN) to a Gated Recurrent Unit (GRU), a type of RNN, as shown on fig. 6. The MHCNN is effective at processing each feature independently [32]. This allows the model to find relevant patterns in each aircraft parameter, which subsequently leads to the identification of the parameters that are highly correlated to the adverse event. Indeed, the most important aircraft parameters are the ones that helped the algorithm correctly predict an anomaly, and therefore can be considered precursors to the anomaly. A window-based approach was used to process the data with the MHCNN, allowing for information extraction of each feature at different time regions. The window is defined by the kernel size of the CNN and the networks all use a stride of 1, such that overlapping windows are created. Some instances of time are more useful than others to predict an anomaly and this approach allows the algorithm to have a more granular extraction of features. Similar to the architecture proposed in [32], four CNNs are used in each head to process the time-series and expand the unique channel of each time series as seen in fig 6 and fig 7, resulting in the extraction of relevant information from each window of the time-series. In Figure 6, k_1 , k_2 and k_3 represent the number of channels, and W_{N_1} , W_{N_2} , and W_{N_3} represent the length of the time series after each convolution operation. The length of the time series after each operation depends on CNN parameters such as the kernel size, the stride, and whether or not padding was used. The first three convolutional layers use batch normalization to reduce the internal co-variance shift, and bring a regularization effect [32], and each normalization is followed by a ReLU activation function which has become a default function to use when developing neural networks due to its many advantages [35]. The fourth convolutional layer reduces the number of channels back to 1, and applies a sigmoid activation function, allowing the output (or feature map) to be interpreted as the probability of a given feature to be a precursor. The chronological order of the feature maps was kept, which helped maintain the temporal information of the data. A concatenation step was then performed to combine the feature maps of each aircraft parameters into a 3 dimensional precursor score tensor.

The concatenated tensor is the input to the GRU, which processes the features all together unlike the MHCNN that processed the features independently. This layer learns the temporal patterns present in the extracted information by the previous convolutional layers. The GRU is appropriate for handling temporal data due to its structure characterized by two gates: the update and reset gates [33]. Both of these gates are used to select which information to keep and which to ignore, and allow the GRU to handle long-term dependencies better than simpler architectures [33]. An hyperbolic tangent activation is then applied to the output of the GRU and followed by a time distributed dense layer that squeezes all the features into one unique dimension, allowing for additional approximation capability [11]. The

output is passed through a sigmoid function, and can be interpreted as the probability that a precursor occurs at a given instances of time. Finally to classify a given flight as positive (event occurred), it is assumed that if any instance of time is positive then the whole flight would be labeled as positive (MIL assumption). Thus, a max pooling layer is used to get the maximum probability across time to classify the flight. A threshold is set such that a flight is positive only if the maximum probability across time is greater or equal to the threshold. The architecture is thus expected to achieve two critical goals:

- 1) Label the bag for identifying adverse events
- 2) Extract precursors and the time instance at which they occur

Although the two approaches are similar because they take advantage of the MIL framework for precursor mining, it is important to note the differences between this architecture and the previously developed algorithm ADOPT [11]. The architecture proposed in this paper is designed such that the precursor probabilities for each feature can be extracted directly from the neural network layers, this approach aims at creating a more interpretable model. ADOPT can only identify temporal precursors directly, and needs to perform a sensitivity analysis during a post-processing step, which consist of perturbing one individual feature at a time and measuring its effect on the time instance probabilities to determine precursors in the feature space. This method potentially misses interactions between features. Thus being able to retrieve precursors by leveraging the combination of the MHCNNs (extraction of individual information for each parameter) and the GRU (extraction of temporal interactions of parameters) mitigates this drawback while eliminating the overhead required to setup the post-processing analysis.

2. Multi-Class Extension

The architecture described in III.B.1 corresponds to a binary classifier. Both IM-DoPE and ADOPT make use of binary classifiers as their task is not only to correctly predict the occurrence of an adverse event but to also identify precursors to the event. Both models could be extended to include multiple outputs allowing for the training of multiple events at once [11]. However, for this work the multi-class model is obtained by combining multiple binary classifiers since it was observed that they performed well at completing their individual tasks. Two popular approaches of combining the classifiers are the "one-vs-all" and the "one-vs-one" strategies [36]. The first strategy learns to classify one class at a time, where the class is distinguished from the other ones. The second strategy requires creating multiple classifiers for different pairs of classes. In this work, the approach taken is a mix between the two strategies. Indeed, 2 classifiers are trained for each event, similarly to "one-vs-all" approach, but each of them are trained to either recognize the event or the default nominal operation. This approach was chosen to allow for the interpretation of nominal flights when performing predictions and to avoid training extensive number of models. When using the models to perform an inference, the output class is chosen to be the output of the model with the highest probability (i.e. the highest confidence) that is greater than a specified threshold. If none of the probabilities are greater than the threshold, then the output is defaulted to the class corresponding to normal operations since each model either predicts an event or normal operation. When using sigmoid functions, the decision threshold is problem dependent and could be treated as a hyperparameter to tune. However, a default decision threshold of 0.5 was used for this work since it yielded good performances for the model.

3. Hyperparameter Search

Given the novelty of the chosen architecture for the precursor mining task, there is a lack of knowledge on what parameters to choose for the model, and on how to train it. To mitigate this problem, a hyperparameter search was performed to determine the optimal parameters of the architecture for each event. A grid-search is the most straight-forward strategy to perform a hyperparameter search, as it entails searching through the space created by all possible combinations of hyperparameters. This means that each parameter is given an array of values to try, and the model is evaluated using metrics described in III.B.4 for each of the parameters combinations. Table 3 defines the hyperparameter space containing 36 possible combinations.

Similarly to [38], the data set was divided into 3 subsets: the training set, the validation set, and the testing set as seen on fig. 8. Stratified subsets are used to ensure that the proportion of nominal flights to abnormal flights remain consistent in each subset. Additionally, stratified mini-batch were used during the training process. The training set was used to train the model using a given set of hyperparameters (i.e. one hyperparameter combination). Google Colab's Tesla V100 16 GB of RAM GPU was leveraged to train the models. The performances of the models on the validation set was used to determine the best model.




Fig. 6 Layer configuration with output size of layer

4. Classification Metrics

Several metrics are used to quantitatively evaluate the model's performance. The evaluation allows for the selection of the best model obtained from the hyperparameter search and its assessment on new unseen data (i.e. on the test set). The best model correctly forecast the occurrence of an adverse event and therefore is best suited to determine the




Fig. 7 Layer configuration showing where the precursor score is extracted

Table 3 Hyperparameters and Search Space

Hyperparameter	Search Range
Convolutional Layers 1,2,3 kernel sizes	[8, 5, 3], [6, 3, 2]
Convolutional Layers 1,2,3 output channels	[10, 15, 20], [16, 32, 64], [32, 64, 128]
Learning Rate	0.001, 0.0001
Weight Decay (L2 regularization)	0.01, 0.001, 0.0001
Mini-batch size percent	1% of training set size
Optimizer	ADAM [37]
Number of epochs	30




Fig. 8 Data Split for Model Training, Hyperparameter Tuning, and Model Evaluation

aircraft's parameters that are the most correlated to the event. The following metrics were used for the model selection and evaluation:

- 1) **Confusion Matrix:** Contains the counts for true positives, true negatives, false positives and false negatives [39], as shown by:

Table 4 Confusion Matrix

	Predicted: No Event	Predicted: Event
Actual: No Event	True Negative (TN)	False Positive (FP)
Actual: Event	False Negative (FN)	True Positive (TP)

- 2) **Precision:** Measures the proportion of positively labeled examples that are actually positive [40]. The closer to one the better, the following equation defines the precision:

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

- 3) **Recall:** Measures the the fraction of positives labels that were actually detected [40]. The closer to one the better, the following equation defines the recall:

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

- 4) **F1 score:** Harmonic mean of the the precision and recall [40]. The closer to one the better, the following equation defines the recall:

$$\text{F1 score} = \frac{2 \times TP}{2 \times TP + FN + FP} \quad (4)$$

- 5) **Distance from ADOPT (DFA):** Metric created to measure the resemblance in the precursor score ranking of IM-DoPE and ADOPT [12]. Assuming precursor scores p_i for feature i ranked between 0 and 1, N flights, and d feature, the DFA for each combination of hyperparameters is given by:

$$sse_j = \frac{1}{d} \sum_i^d (p_{i,IMDoPE} - p_{i,ADOPT})^2 \quad (5)$$

$$\text{DFA} = \frac{1}{N} \sum_j^N sse_j \quad (6)$$

DFA for a combination of hyperparameters is the mean sum square error between the precursor scores of ADOPT and IM-DoPE, therefore the closer to zero, the more alike the rankings of the two algorithms are. The precursor

score sum square errors are averaged across all features 5, and across all flights 6. ADOPT has been validated by expert for few events such as an high speed exceedence and can thus be used as a second-hand validation if no human expert is available.

C. Precursor Identification

1. Raw Precursor Scores

Once acceptable performances are achieved by the model, it can be used to identify precursors. The precursor score $p_{i,t}$ of each feature i at time step t is extracted directly from the architecture. As seen on fig. 9, the raw precursor score for each feature can be extracted from the last convolution layer of each head. The raw scores are bounded between 0 and 1 due to the sigmoid function used in the fourth convolutional layer. From empirical results, it was observed that the model learns to set non-important parameters to 0.5. As an example on Fig. 9, the rudder position, the rudder pedal, and the spoiler position are deemed not important by the model since their individual scores remain at 0.5 for last 20 nautical miles before a 1,000 ft. On the other hand, the radio altitude was flagged as a precursor. As the less important features go through the different convolutional layers, they get reduced to zero. This zeroed out feature is then passed through the sigmoid activation resulting in a raw score of 0.5.




Fig. 9 Extraction of Precursors for a High Speed Event

The identified precursors have instance of time (precursor score over time) at which they are more correlated to the event of interest (grayed area in Fig. 9 and 10). The combination of GRU and dense layer allows for the extraction of these time instances, as seen in Fig. 10. Similar to ADOPT, abnormal flights will have the precursor score increase towards 1 and nominal flights will have the score fall towards 0. Additionally, since the precursors are more correlated to an event at time windows identified by the GRU, an overall score can be given to the precursor by averaging its $p_{i,t}$ values within an identified window T .

2. Adjusted Precursor Scores

A more natural way to express the precursor score is to have a score of zero for parameters that are not important. Therefore, the raw precursor scores were adjusted. The adjusted precursor score p_i of feature i is therefore defined by the following equation:

$$p_i = \frac{1}{|T|} \sum_t (|p_{i,t} - 0.5|), \forall t \in T \text{ where } T \subseteq L \quad (7)$$

where $|T|$ is the number of time steps within the time window T . The adjusted precursor score is thus bounded between 0 and 0.5. From this definition, it can be seen that features with higher scores are flagged as precursors.




Fig. 10 Extraction of Precursors Score Over Time from Dense Layer for a High Speed Event

IV. Results and Discussion

A. Quantitative Results

For each event, 36 combination of hyperparameters were evaluated and the best classifiers were selected. Table 5 summarises the quantitative results obtained when evaluating the models on the test set.

Table 5 Model Evaluation Results (Binary Classifications)

Event	Algorithm	F1 Score	Precision	Recall	DFA
High Speed	IM-DoPE	0.90	0.84	0.97	0.0392
High Path Angle	IM-DoPE	0.83	0.81	0.87	0.0152
High Speed	ADOPT	0.88	0.90	0.86	N/A
High Path Angle	ADOPT	0.70	0.56	0.90	N/A

While the best MHCNN-RNN model achieves high performances, better results for the classical classification metrics are obtained for the high speed event due to the higher number of training examples and the the smaller imbalance in the classes. However, the larger DFA metric for the high speed event suggest a greater difference with ADOPT feature ranking for this event. Note that the DFA values obtained are still relatively good as it is among the lowest DFA values obtained as seen on fig. 11. Additionally, best model of IM-DoPE performs better than ADOPT when using the same training set, and testing set. It is however important to note that while a grid-search was performed for IM-DoPE for each event, the default setting of ADOPT were used. Some models obtained during the grid-search have lower scores that are worse than the ones obtained from using ADOPT default settings. For example, on fig 11, several models reached scores lower than 0.88.

The binary classifiers are then combined as explained in III.B.2 in order to obtain a multi-class classifier. Similar to binary classification, each metric can be computed for each class. For instance, the precision of the high speed class is the number of flights that actually faced a high speed event over the total number of flight that were labeled as a high speed event by the classifier. Table 6 presents similar scores as in table 5 showing that even when combined, each model learns to identify its positive class correctly. The confusion matrix in 7 confirms this behavior, as the numbers of FP and TN for each class are relatively low.




Fig. 11 F1 Score and DFA for All Combinations (High Speed Event)

Table 6 Model Evaluation Results (Multi-Class Classifications)

Event	F1 Score	Precision	Recall
Nominal	0.91	0.97	0.85
High Speed	0.90	0.84	0.97
High Path Angle	0.82	0.80	0.83

Table 7 Confusion Matrix (Multi-Class Classifications)

	Predicted Nominal	Predicted High Speed	Predicted High Path Angle
Actual Nominal	1717 (47%)	263 (7.2%)	33 (0.9%)
Actual High Speed	38 (1.0%)	1422 (39%)	1 ($2.7 \times 10^{-4}\%$)
Actual High Path Angle	16 (0.44%)	12 (0.33%)	138 (3.8%)

B. Identified Precursors

1. Average Adjusted Precursor Scores

As previously mentioned, the final model is expected to predict adverse events and identify their precursors. The adjusted precursor score can be obtained for each feature and each flight. Considering true positive flights only, the average adjusted precursor score can be determined and the precursors can be identified on a fleet level, as seen in table 8 for the top 5 precursors of each event. Noticeable differences are observed in the precursor rankings for the two events. For example the glideslope deviation and the pitch angle are characteristics of a high path angle event while the N1 target relates to engine power which can be related to speed. Some resemblances are also observed, for instance the

altitude is seen to be the precursor for both events. This is expected since the altitude above touch-down is used to define both events. The trained model can also be used to analyze individual flights. Two flights experiencing a high speed event, and a high path angle event respectively, are analyzed IV.B.2 and in IV.B.3.

Table 8 Average Adjusted Precursor Scores for High Speed and High Path Angle Events

Precursor Rank	Average Adjusted Precursor Score	
	High Speed Event	High Path Angle Event
#1	Altitude: 0.31	Radio Altitude:0.31
#2	Radio Altitude : 0.28	Glideslope Deviation: 0.24
#3	N1 Target: 0.25	Pitch Angle: 0.21
#4	Body Longitudinal Acceleration: 0.23	Altitude: 0.18
#5	Lateral Acceleration: 0.14	Flight Path Acceleration: 0.12

2. High Speed Event Flight Analysis

When performing an inference using the trained model, the precursor scores can be extracted from the model's MHCNN outputs. Following previously highlighted steps in III.C.2, the adjusted precursor score can be computed and used to identify the precursors for a flight. Fig. 12 shows the precursor ranking obtain for a flight that experienced a high speed event. Similar precursors were identified by ADOPT for this flight. The top 5 precursors were the altitude, the radio altitude, the flight path acceleration and the N1 target. Additionally, the the total pressure was also identified by ADOPT as a top precursor. Once the precursors are discovered, the aircraft's parameters can then be plotted to assess the them.




Fig. 12 Precursor Ranking for a High Speed Event

In particular, the time series is of interest when the precursor score over time is greater than the 0.5 threshold, which occurs at around 1.5 nautical miles away from a 1,000 ft above touch-down on fig. 13 (grey shaded area). On the figure, the red and blue curves in the first cell represent the precursor at different time steps obtained from the high speed and high path angle classifiers, respectively. Since the flying conditions do not correspond to a high path angle

event, the classifier for the event cannot find precursors, while the classifier for the high speed event detects abnormal pattern leading to the identification of precursors. Note that if only the high path angle binary classifier was used, the model would have confidently classified the flight as normal. In the remaining cells, the dotted blue line represents the feature value. For instance this flight reached a pitch angle of 5 degrees 7.5 nautical miles away from a 1,000 ft. The dotted black line represents the mean nominal values for each feature at every time steps and the shaded purple region is defined as ± 2 standard deviations away from the mean. From the figure and knowledge of the precursors it can be inferred that a potential cause to the excessive speed experienced by the flight at a 1,000 ft is correlated to the combination of negative accelerations (body longitudinal acceleration and flight path acceleration) and low pitch angle in the last 2.5 nautical miles. The oil temperature of the flight is also plotted to show that even though it has a lower values than the mean of nominal flights, the relatively low precursor score obtained is likely due to the fact that the temperature is relatively constant and is not highly correlated to the event.




Fig. 13 Precursor Score and Aircraft's Parameters during a High Speed Event

3. High Path Angle Event Flight Analysis

Similar to the high speed event case, precursor to the high path angle event can be probed from the model. Again, the identified precursors provide a list of aircraft's parameters that can be assessed through visualization. For the selected flight the presence of a dominating precursor is observed. Indeed, there is a larger difference between the top two precursor than there was for the high speed event, as seen on fig. 14. This larger difference suggest a highly abnormal glideslope deviation. Moreover, the glideslope was also identified by ADOPT as the top parameter. Other important parameters were the pitch angle, the radio altitude, the airbrake position, and the flight path acceleration.

The abnormal behavior is confirmed by fig. 14 since the glideslope deviation is much greater than 2 times the standard deviation. The potential cause of the event is likely mainly related to the high deviation in the glideslope. However, other precursors such as the pitch angle and the flight path acceleration are also identified and observed to have abnormal patterns. It is important to note that these two parameters can also be precursors to the high speed event as previously observed. This flight in particular was not classified as a high speed event but some of the parameters behaved similarly to how they would behave during such event. This led the high speed event classifier to increase the precursor score over time probability to be closer to 1. Ultimately, the multi-class classifier correctly labeled this flight as a high path angle event because the high path angle classifier had a stronger confidence.

C. Discussion of Results

The MHCNN-RNN architecture yields satisfying results. The scores obtained from the conventional classification metrics show the model's ability to extract information from the data in order to make accurate predictions. In fact, the model accurately forecast the events since it predicts them before they actually occur, though variability in the prediction is observed. The model also identifies precursors, and it was observed that different precursors are discovered for different events. which is expected. The high speed top precursors relate more to poor energy management while the precursors for the high path angle include trajectory related parameters, as seen in table 8. The identified precursors are partially validated given their resemblances with ADOPT identified precursors, which is characterized by the low DFA




Fig. 14 Precursor Ranking for a High Path Angle Event




Fig. 15 Precursor Score and Aircraft's Parameters during a High Path Angle Event

values. Additional support towards the validation of the discovered precursors is obtained through visualization and observing that the precursors exhibited abnormal behaviors, outside of nominal operations.

V. Conclusion

The work presented in this paper tackles the precursor mining task. A public flight data set was leveraged to create FOQA-like labels. After performing the required preprocessing steps, a novel architecture for the task was developed to take advantage of the Multi-Instance Learning framework, the feature extraction capabilities of Convolutional Neural Networks and the temporal pattern recognition capabilities of Recurrent Neural Networks. Binary classifiers were trained to predict both high speed and high path angle events. A grid-search was implemented to determine the best parameters for the neural networks, and thus the best model for each of the events. The best two models were then combined to form an unique multi-class classifier. The final model was evaluated on a test set and high scores were observed for classification metrics such as F1 score, precision, and recall. The model was also able to identify precursors and provide the average precursor score across all flights that experienced each of the two events. Finally, visualizations were used to observe the behaviors of the identified precursors, which exhibited patterns different from normal flight

operations. Future work will include enhancing the interpretability of the precursor score tensor. While empirically the parameters that deviate from 0.5 are understood to be correlated to the event of interest, the meaning of the sign of the deviation (greater than or lower than 0.5) still needs to be investigated. Further improvements to the model should also be made towards increasing the prediction window, and allowing the classification of unknown precursors instead of defaulting unknown behaviors to the nominal class.

Acknowledgments

The authors would like to thank Dr. Nikunj Oza, Dr. Milad Memarzadeh, and Dr. Hamed Valizadegan for the feedback and insights they provided.

References

- [1] ICAO, “Aviation Benefits Report,” 2019. URL <https://www.icao.int/sustainability/Documents/AVIATION-BENEFITS-2019-web.pdf>.
- [2] “Commercial Market Outlook 2019–2038,” 2019. URL <https://www.boeing.com/resources/boeingdotcom/commercial/market/commercial-market-outlook/assets/downloads/cmo-sept-2019-report-final.pdf>.
- [3] Allmond, R., “Aviation Safety 2019 Year in Review,” 2019. URL https://www.faa.gov/about/office_org/headquarters_offices/avs/media/2019_Year_In_Review.pdf.
- [4] Menon, P. K., Dutta, P., Chen, O., Iyer, H., and Yang, B.-J., “A Modeling Environment for Assessing Aviation Safety,” *AIAA Aviation 2019 Forum*, 2019, p. 2937.
- [5] Matthews, S., “The Potential for Improving Aviation Safety and Reducing the Accident Rate,” *AIAA International Air and Space Symposium and Exposition*, 2003. doi:10.2514/6.2003-2796.
- [6] Howell, C., “How to Do Predictive Risk Management Data Mining in Aviation SMS,” 2019. URL aviationsafetyblog.asms-pro.com/blog/best-data-mining-methods-predictive-aviation-safety-risk-management.
- [7] Logan, T. J., “Error Prevention as Developed in Airlines,” *International Journal of Radiation Oncology Biology Physics*, Vol. 71, No. 1, 2008, pp. S178–S181. doi:10.1016/j.ijrobp.2007.09.040.
- [8] “Federal Aviation Administration Advisory Circular, 120-82 - Flight Operational Quality Assurance,” , April 2004. Url: https://www.faa.gov/regulations_policies/advisory_circulars/index.cfm/go/document.information/documentID/23227.
- [9] Basora, L., Olive, X., and Dubot, T., “Recent Advances in Anomaly Detection Methods Applied to Aviation,” *Aerospace*, Vol. 6(11), 2019, p. 117. doi:10.3390/aerospace6110117.
- [10] Ackley, J. L., Puranik, T. G., and Mavris, D., “A Supervised Learning Approach for Safety Event Precursor Identification in Commercial Aviation,” *AIAA AVIATION 2020 FORUM*, 2020, p. 2880. doi:10.2514/6.2020-2880.
- [11] Janakiraman, V. M., “Explaining Aviation Safety Incidents Using Deep Temporal Multiple Instance Learning,” 2018, pp. 406–415. doi:10.1145/3219819.3219871.
- [12] Janakiraman, V. M., Matthews, B., and Oza, N., “Using ADOPT Algorithm and Operational Data to Discover Precursors to Aviation Adverse Events,” 2018. doi:10.2514/6.2018-1638.
- [13] Gavrilovski, A., Jimenez, H., Mavris, D., Rao, A. H., Marais, K., Shin, S., and Hwang, I., “Challenges and Opportunities in Flight Data Mining: A Review of the State of the Art,” 2016. doi:10.2514/6.2016-0923.
- [14] Matthews, B., Das, S., Bhaduri, K., Das, K., Martin, R., and Oza, N., “Discovering Anomalous Aviation Safety Events Using Scalable Data Mining Algorithms,” Vol. 10, No. 10, 2013, pp. 467–475. doi:10.2514/1.1010080.
- [15] Puranik, T. G., and Mavris, D. N., “Anomaly Detection in General-Aviation Operations Using Energy Metrics and Flight-Data Records,” *Journal of Aerospace Information Systems*, Vol. 15, No. 1, 2018, pp. 22–35. doi:10.2514/1.1010582.
- [16] Puranik, T. G., and Mavris, D. N., “Identification of Instantaneous Anomalies in General Aviation Operations using Energy Metrics,” *Journal of Aerospace Information Systems*, Vol. Article in Advance, 2019. doi:10.2514/1.1010772.

- [17] Das, S., Matthews, B., Srivastava, A., and Oza, N., "Multiple Kernel Learning for Heterogeneous Anomaly Detection: Algorithm and Aviation Safety Case Study," 2010, pp. 47–56. doi:10.1145/1835804.1835813.
- [18] Sheridan, K., Puranik, T., ey, E., Pinon, O., Kirby, M., and Mavris, D., "An Application of DBSCAN Clustering for Flight Anomaly Detection During the Approach Phase," *AIAA SciTech Forum*, 2020. doi:10.2514/6.2020-1851.
- [19] Li, L., Hansman, R. J., Palacios, R., and Welsch, R., "Anomaly detection via a Gaussian Mixture Model for flight operation and safety monitoring," *Transportation Research Part C: Emerging Technologies*, Vol. 64, 2016, pp. 45 – 57. doi:10.1016/j.trc.2016.01.007.
- [20] Lee, H., Li, G., Rai, A., and Chattopadhyay, A., "Real-time anomaly detection framework using a support vector regression for the safety monitoring of commercial aircraft," *Advanced Engineering Informatics*, Vol. 44, 2020, p. 101071. doi:10.1016/j.aei.2020.101071.
- [21] Memarzadeh, M., Matthews, B., and Avrekh, I., "Unsupervised Anomaly Detection in Flight Data Using Convolutional Variational Auto-Encoder," *Aerospace*, Vol. 7, 2020, p. 115. doi:10.3390/aerospace7080115.
- [22] Corrado, S. J., Puranik, T. G., Pinon, O. J., and Mavris, D. N., "Trajectory Clustering within the Terminal Airspace Utilizing a Weighted Distance Function," *Proceedings*, Vol. 59, No. 7, 2020. doi:10.3390/proceedings2020059007.
- [23] Puranik, T. G., Rodriguez, N., and Mavris, D. N., "Towards online prediction of safety-critical landing metrics in aviation using supervised machine learning," *Transportation Research Part C: Emerging Technologies*, Vol. 120, 2020, p. 102819. doi:10.1016/j.trc.2020.102819.
- [24] Deshmukh, R., Sun, D., and Hwang, I., "Data-Driven Precursor Detection Algorithm for Terminal Airspace Operations," *Proceedings of the 13th USA/Europe Air Traffic Management Research and Development Seminar, Vienna, Austria*, 2019, pp. 17–21. URL http://www.atmseminar.org/seminarContent/seminar13/papers/ATM_Seminar_2019_paper_46.pdf.
- [25] Li, X.-L., "A Multiclass Multiple Instance Learning Method with Exact Likelihood," , 2018.
- [26] Ning, Y., Muthiah, S., Rangwala, H., and Ramakrishnan, N., "Modeling Precursors for Event Forecasting via Nested Multi-Instance Learning," 2016. doi:10.1145/2939672.2939802.
- [27] Lee, H., Madar, S., Santusht, S., Puranik, T. G., Payan, A. P., Kirby, M., Pinon, O. J., and Mavris, N., Dimitri, "Critical Parameter Identification for Safety Events in Commercial Aviation Using Machine Learning," *Aerospace*, Vol. 7(6), 2020, p. 73. doi:10.3390/aerospace7060073.
- [28] Melnyk, I., Yadav, P., Steinbach, M., Srivastava, J., Kumar, V., and Banerjee, A., "Detection of precursors to aviation safety incidents due to human factors," *2013 IEEE 13th International Conference on Data Mining Workshops*, IEEE, 2013, pp. 407–412. doi:10.1109/ICDMW.2013.55.
- [29] Mangortey, E., Monteiro, D., Ackley, J., Gao, Z., Puranik, T. G., Pinon, O. J., Kirby, M., and Mavris, D., "Application of Machine Learning to Parameter Selection for Flight Risk Identification," *AIAA SciTech Forum*, 2020. doi:10.2514/6.2020-1850.
- [30] Carbonneau, M.-A., Cheplygina, V., Granger, E., and Gagnon, G., "Multiple instance learning: A survey of problem characteristics and applications," *Pattern Recognition*, Vol. 77, 2018, p. 329–353. doi:10.1016/j.patcog.2017.10.009, URL <http://dx.doi.org/10.1016/j.patcog.2017.10.009>.
- [31] Calkins, K. G., "Applied Statistics-Lesson 5: Correlation Coefficients," , July 2005.
- [32] Canizo, M., Triguero, I., Conde, A., and Onieva, E., "Multi-head CNN-RNN for multi-time series anomaly detection: An industrial case study," *Neurocomputing*, Vol. 363, 2019, pp. 246 – 260. doi:10.1016/j.neucom.2019.07.034.
- [33] Goodfellow, I., Bengio, Y., and Courville, A., *Deep Learning*, MIT Press, 2016. <http://www.deeplearningbook.org>.
- [34] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S., "PyTorch: An Imperative Style, High-Performance Deep Learning Library," *Advances in Neural Information Processing Systems 32*, Curran Associates, Inc., 2019, pp. 8024–8035. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [35] Brownlee, J., "A Gentle Introduction to the Rectified Linear Unit (ReLU)," , Aug 2019. URL <https://machinelearningmastery.com/rectified-linear-activation-function-for-deep-learning-neural-networks/>.

- [36] Galar, M., Fernández, A., Barrenechea, E., Bustince, H., and Herrera, F., “An overview of ensemble methods for binary classifiers in multi-class problems: Experimental study on one-vs-one and one-vs-all schemes,” *Pattern Recognition*, Vol. 44, No. 8, 2011, pp. 1761 – 1776. doi:10.1016/j.patcog.2011.01.017.
- [37] Kingma, D. P., and Ba, J., “Adam: A Method for Stochastic Optimization,” *CoRR*, Vol. abs/1412.6980, 2015. doi:<https://doi.org/10.1007/s11704-019-8457-x>.
- [38] Mangortey, E., Bleu-Laine, M.-H., Puranik, T. G., Pinon Fischer, O., and Mavris, D. N., “Machine Learning Approach to the Analysis of Traffic Management Initiatives,” *Journal of Air Transportation*, Vol. Article in Advance, 2020, pp. 1–14. doi:10.2514/1.D0212.
- [39] Lantz, Brett, *Machine Learning with R: Discover How to Build Machine Learning Algorithms, Prepare Data, and Dig Deep into Data Prediction Techniques with R*, Packt Publishing, 2015.
- [40] Murphy, K., *Machine Learning: a Probabilistic Perspective*, the MIT press, 2013.