Using ADOPT Algorithm and Operational Data to Discover Precursors to Aviation Adverse Events

Vijay Manikandan Janakiraman
USRA/ NASA Ames Research Center, Moffett Field, CA, 94035, USA.

Bryan Matthews
SGT/ NASA Ames Research Center, Moffett Field, CA, 94035, USA.

Nikunj Oza
NASA Ames Research Center, Moffett Field, CA, 94035, USA.

In this paper, we describe a data mining method called ADOPT (Automatic Discovery of Precursors in Time series data) to identify precursors to aviation adverse events. An adverse event may refer to any unsafe event ranging from a negligible safety hazard to a catastrophic accident, depending on the scope of the analysis. A precursor is an early indicator of an increasing likelihood of the adverse event. Identifying precursors is important in the context of a proactive safety management because precursors detect the increasing severity of the underlying hazard much earlier, giving sufficient time to identify, analyze and implement corrective actions. ADOPT analyzes large volumes of historical data to find complex trends among several sensory variables simultaneously to find precursors. ADOPT's data mining approach captures real-world effects such as human factors, weather, geographic constraints, operating procedures, airline strategies etc that are difficult to capture using first-principle models. This paper describes the algorithm using two case studies including a take-off stall hazard and RNAV adherence. While the case studies are not intended to discuss the critical safety risks in aviation, they are used to demonstrate the various steps involved in ADOPT including data preparation, variable selection, parameter tuning, experiment setup and analyzing the results. The results show that ADOPT can be a powerful tool to identify and analyze performance and safety issues in Aviation.

I. Introduction

Many of the current aviation safety systems are reactive (i.e., an adverse event occurs first which is used to trigger a response/ recovery). For example, after an impending stall is sensed, an alert in the form of a stick shaker is given so that the pilot initiates a recovery maneuver. While reactive systems are abundant in current aviation systems, they may not always work and accidents do happen in rare cases. Often the time between the alert and the adverse event can be short in a reactive system. Further, during the recovery process, any secondary factors such as human fatigue, lack of situational awareness often results in a failed recovery.[1,2] Thus, to improve safety, a proactive approach to risk management is required. A proactive system identifies the latent risk factors early so that sufficient time is available to analyze the situation, make predictions about possible outcomes, and choose the best action to enforce a risk mitigation plan in order to prevent adverse event.[3] In this paper, we address the first aspect of identifying latent factors (which we call precursors) to aviation safety events. We propose to use operational data of past adverse events to identify the precursors in a retrospective way. We describe a recently developed data mining technique called ADOPT (Automatic Discovery of Precursors in Time series)[4] and show its application to aviation problems. We discuss how one can use historical aviation data and the ADOPT algorithm to setup a data mining work-flow, and automatically infer a knowledge base [a] of precursors to some adverse event of interest.

[a] A knowledge base could be a set of precursor rules or a model that infers precursor variables using sensory measurements.

A precursor is an early indicator of an increasing likelihood of the adverse event. For instance, a "hard touchdown" could be an adverse event for which a precursor could be a "high descent rate at the outer-marker" which gives an indication that a hard touchdown may occur. An adverse event in ADOPT may refer to any unsafe event ranging from a negligible safety hazard to a catastrophic accident that is of interest to the analyst. It may be an event that occurs at the subsystem level (as in engine failure), flight level (runway excursion), airport level (delays) or at the airspace level (traffic congestion). As precursors are early indicators, identifying them may give an increased response time to proactively determine subsequent risk. For example, knowing the high descent rate early in approach, the pilot may analyze the risk, evaluate corrective actions, and enforce the best action so that the hard landing is prevented.

Methods to identify precursors include incident/accident investigations, trend analysis, expert panel, simulation and testing, and using data mining. In this paper, we follow a data mining approach. The main advantage of basing the discovery process directly on operational data is that one can expect to capture the precursors arising from real-world effects such as human factors, weather, geographic constraints, operating procedures, airline strategies etc that are difficult to capture using models based on first-principles. Thus, ADOPT could be complementary to existing approaches. In addition, many of the modern aviation systems are equipped with sensors that collect rich data at multiple system and subsystem levels in the US National Airspace (NAS). Using such a rich source of information, it may now be possible to uncover some of the hidden risk factors in the NAS by mining operational data.

The ADOPT algorithm is a general methodology that does not make any special assumptions about the system and does not need specialized knowledge on the system states enabling it to operate directly on the observed time series data. Further, it can scale well to multivariate time series and can analyze large number of flights, which may enable a faster turnaround time for subsequent tasks such as hazard identification and safety risk analysis.[5] A subject matter expert may require a few hours to analyze a flight for precursors which is not scalable considering the thousands of flights operated on a daily basis. Also, human experts may not be able to visualize hundreds of time series variables to notice complex variations and trends in the data. ADOPT may be used to speed this process by analyzing the thousands of flights that operate every day to short-list only the significant precursors which may then be analyzed by a subject matter expert, reducing the turnaround time for safety analysis. Recently, ADOPT was published from a data mining perspective[4,6] with some preliminary demonstration using aviation safety problems. In this paper, we present further case studies using two scenarios including take-off stall hazard and RNAV procedure adherence. The case studies demonstrate the working of ADOPT under different problem scopes and data sources, and offer insights on how to efficiently apply data mining to discover safety precursors.

The rest of the paper is organized as follows. The precursor mining methodology and ADOPT algorithm are detailed in Section II which is followed by detailed case studies in Section III. A brief summary of relevant work and literature is presented in Section IV and is followed by concluding remarks in Section V.

II. Methodology

An adverse event such as a safety incident is a temporal phenomenon which is usually preceded by a sequence of events. Also, we usually have multiple sensors observing the system collecting time stamped data. Thus, our setup has a multivariate time series with d variables as shown in Figure 1. If the adverse event E_a occurs at time $L+1$, the data prior to the adverse event (corresponding to time $1, 2, .., L$) can be considered the search data where precursors may be present. Let $\overline{\mathcal{N}} = \{X_1, X_2, .., X_{\overline{N}}\}$ be a database containing $\overline{N}$ such time series records. Similarly let $\mathcal{N} = \{X_1, X_2, .., X_N\}$ be a database containing N time series records that are nominal; i.e., data where E_a does not occur. For example, flights that had a speed exceedance for a year could be considered the adverse database $\overline{\mathcal{N}}$. On the other hand, the flights where the exceedance did not occur during that year may be considered the nominal database $\mathcal{N}$.

A data record X_i can be represented in matrix form as follows

$$X_i = \begin{bmatrix} x^1(1) & x^1(2) & . & x^1(L_i) \\ x^2(1) & x^2(2) & . & x^2(L_i) \\ . & . & . & . \\ x^d(1) & x^d(2) & . & x^d(L_i) \end{bmatrix}$$

where L_i is the length of the multivariate time series X_i. Here, a data record X may be a flight with d sensory variables such as velocity, altitude, etc. that are measured at regular sampling intervals. The event

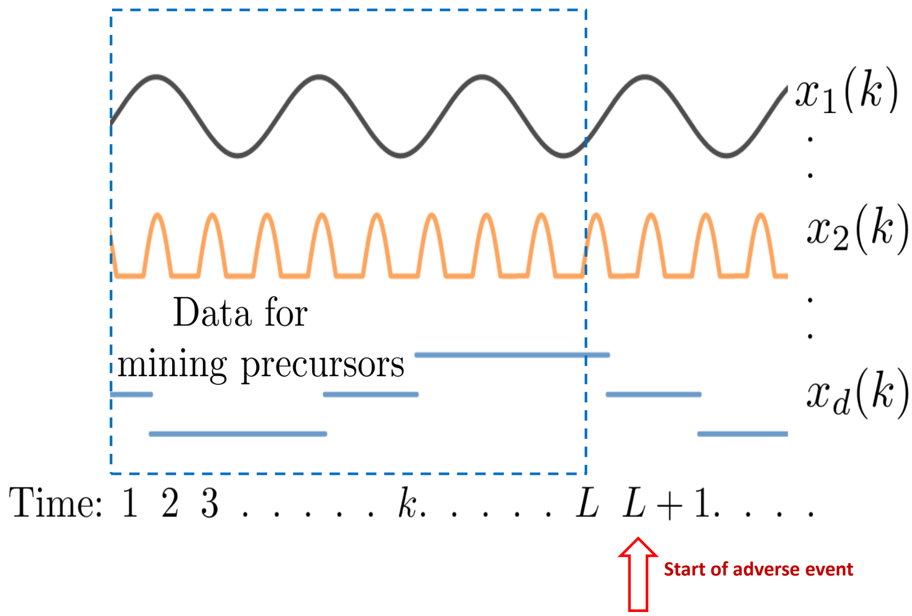

Figure 1: Schematic showing the data setup for precursor discovery in ADOPT. Each data record is a time series of d dimensions.

at time k is given by

$$\mathbf{x(k)} = [x_1(k), x_2(k), .., x_d(k)]\,. \tag{1}$$

The time series record X_i can now be represented in terms of events as $X_i = [\mathbf{x(1)}, \mathbf{x(2)}, .., \mathbf{x(L_i)}]$. With this setup, ADOPT defines a precursor as follows. Given a sequence of events $X = [\mathbf{x(1)}, \mathbf{x(2)}, .., \mathbf{x(L)}]$, an action is any transition $a_k : \mathbf{x(k)} \longrightarrow \mathbf{x(k+1)}$ where $1 \leq k \leq L$, then a_k is a precursor to E_a if

$$V(a_k) - V(a_k^*) > \delta, \tag{2}$$

where a_k^* is the expert action at time k and $V(.)$ is the expert's value function.[4]

A value function in this context, is a metric that evaluates a given transition (a_k in this case) for its long term consequence with respect to the adverse event. It was shown in[6] that the value function is equivalent to the conditional probability $P(E_a|a_k)$; i.e., a high value of a_k translates to a high probability of adverse event occurring in the future. Thus, the expert's value $V(a^*)$ is always less than or equal to $V(a)$. We assume that the nominal data is generated by an expert who manages the state transitions so that the adverse event is prevented. Consequently, the expert's action a_k^* is the best action that can be taken at time k to prevent the adverse event.

The algorithm begins by taking data corresponding to nominal $\mathcal{N}$ and adverse $\overline{\mathcal{N}}$ time series as trajectory demonstrations of the expert and non-expert respectively. The expert's reward model $R(\mathbf{x})$ is defined based on the knowledge of the adverse event in the data as

$$R(\mathbf{x}) = \begin{cases} 0, & \text{if} \quad \mathbf{x} = \mathbf{x(L)} \quad \text{and} \quad X \in \mathcal{N} \\ 1, & \text{if} \quad \mathbf{x} = \mathbf{x(L)} \quad \text{and} \quad X \in \overline{\mathcal{N}} \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

In other words, we assign a reward of 1 for the last state transition in the adverse time series data; i.e., when the adverse event happens. Everywhere else, the state transition gets a reward of 0 for the adverse time series. For the nominal time series, the reward is set to 0 for all state transitions. Using the reward model and the trajectory data, the expert's value model $\hat{V}^{\pi_E}(\mathbf{x})$ is estimated using reinforcement learning. The idea behind the value function is to capture the long-term consequence of the state. Note that the reward is an instantaneous consequence of the state. For our purposes, we need to evaluate the long-term consequence to identify if a given state/action is a precursor. A support vector machine (SVM) classification model is used to model the expert's value function. For a detailed description of the steps involved in ADOPT, we refer to.[4,6] Using the value model, every data transition (abstracted as action a_k) of a test flight is evaluated for its correlation to the adverse event. A high correlation is inferred if

$$PS(a_k) = V^{\pi_E}(a_k) - V^{\pi_E}(a_k^*) > \delta_P, \tag{4}$$

In other words, when a pilot takes an action that is similar to an "expert", then $PS(a_k)$ will be low while if the pilot takes a "poor" action that lowers safety, then $V^{\pi_E}(a_k)$ for that action will be high which makes $PS(a_k)$ to be higher than the threshold δ_P. Such an action will be detected as a precursor.

III. Case Studies

In this section, we detail two case studies to demonstrate the working of ADOPT. While the case studies are not intended to discuss the critical safety risks in aviation, they are used to demonstrate the various steps involved in ADOPT including data preparation, variable selection, parameter tuning, experiment setup and analyzing the results.

A. Take-off Stall Hazard

This case study is a demonstration of how monitoring data from an aircraft can be used to find safety precursors at a flight level. Here the goal is to find precursors to a stall related hazard during take-off. This is an important problem because airspeed management based accidents were prevalent in the last decade.[1,2] While aerodynamic stall is a well studied problem and a reactive safety system (stick shaker for example) exists, we aim to identify precursors that may inform proactive safety systems to prevent the hazard. In our prior work, we introduced this problem and developed an algorithmic framework using ADOPT.[6] Here, we aim to reuse this case study and discuss additional flight examples.

In this case study, a stall hazard is characterized by a drop in airspeed during take-off with a severity of at least 20 knots in airspeed reduction. This severity level detects events that are operationally significant as well as gives sufficient data records to train our models. The reason for choosing airspeed to study precursors is as follows. The airspeed is an important component of the energy state of the aircraft. The airspeed profile during take-off is determined using the gross weight of the airplane, air temperature and using the performance characteristics of the airplane. Any human errors in the calculations or during take-off, lack of state awareness and external disturbances will be observed in the airspeed prior to the hazard. While the definition of the adverse event is based purely on the airspeed, the precursor analysis considers multiple time series variables in combination to find latent factors present in the explanatory variables.

Using the Flight Operational Quality Assurance (FOQA) data from a de-identified airlines, we obtained two sets of flights - the ones that had the drop in airspeed events and the ones that did not (the nominal). Let these two datasets be identified as $\overline{\mathcal{N}}$ and $\mathcal{N}$ respectively. The FOQA data has more than 350 time series variables in each flight and it is not trivial to decide which ones to choose for precursor analysis. We did feature selection based on Granger Causality[7] to get an initial set of variables from which shortlisted further using domain knowledge for further ADOPT modeling. More than 150 variables were eliminated using Granger causality method. From the rest, the top 10% of the variables that causally affects airspeed were short listed. We verified the short-listed variables with a team of domain experts who eliminated some variables that were unrelated to the airspeed drop events and added a few that were missed by Granger causality method.

The data is split as 90% for training (about 36000 flights) and the remaining (about 4000 flights) for testing. The class proportions in the data are balanced with equal number of nominal and adverse flights by randomly sub-sampling the nominal flights which were in excess. The Support Vector Machine (SVM) model for ADOPT's value function was trained using Gaussian kernel parameters $\gamma = 10, C = 50$, achieving an accuracy of over 85% on the unseen test data.[6] The values of γ and C are selected based on cross-validation. Using the trained models, we analyze the precursors to the drop in airspeed events in the following sections. We identified some commonly occurring scenarios and included an example flight for each. For each example, we refer the flight evaluated by ADOPT as the "test flight". The plots show nominal distribution of parameter values in green (10-90 percentile) and blue (25-75 percentile) as references. The blue area shows what a nominal flight (with no adverse event) values look like. The green area includes more variations in the nominal flights and thus shows possible values that may be slightly off-nominal but still ends up with no adverse event. Note that these distributions are not given to ADOPT during learning and are available only to serve as references. The flight progress (or time) is along the x-axis from left to right where the right most data point corresponds to the adverse event. Markers of type 'X' will be used to indicate precursor time instances in the flight.

1. *Flight Analysis 1: Reference speed set incorrectly*

In this example, the speed reference (PFD [b] selected speed) is incorrectly set during take-off which turns out to be the main precursor. The flight takes off nominally until about 30 seconds when the test flight's PFD was set to less than 150 knots. Note that during this time, most nominal flights had their PFD speed increased above 200 knots (see Figure 2). The PFD speed is a reference based on which the pitch commands are derived (a high pitch command is followed to match the low PFD speed reference). Correspondingly, between 30 and 45 seconds after take-off, the PFD Speed was ranked as a top precursor (see Table 1). Although the pitch is defined by the PFD, it is also identified as one of the top precursors by ADOPT. In this way, ADOPT ranks the time series variables at each time step that contributed to the high precursor score. This is important to understand and interpret the results of the model, particularly when the data is high dimensional. ADOPT breaks down the influence of each variable independently and ranks the top precursors at a given time instant. To get a visual understanding, ADOPT's scoring is shown in Figure 2. The "Prob" plot shows the probability of the speed drop occurring calculated every time step of the flight. The precursor score is shown in the "Precursor Index" plot along with the learned threshold. The threshold gives the cut-off at every time-step on the precursor score; i.e., a value greater than the threshold indicates a precursor. The precursors discovered for this flight and for others were qualitatively validated by subject matter experts who did an independent analysis of the flights for precursors.

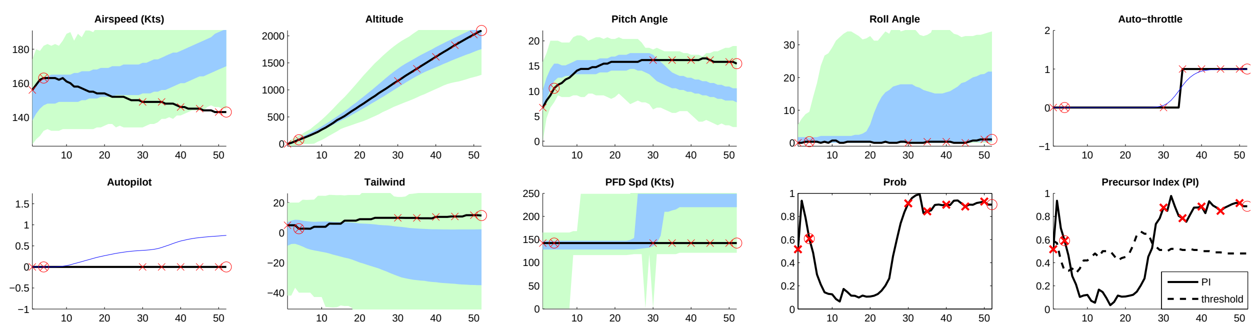

Figure 2: Flight time series variables plotted (thick black curve) against time (seconds) after lift-off. The distribution of nominal values are shown in blue (25-75 percentile) and green (10-90 percentile) as references. The X marker in red shows the precursors identified by ADOPT while the two O markers in red show the start of the speed drop and where the drop reaches 20 knots respectively.

2. *Flight analysis 2: High headwind*

In this example, the flight took-off in a strong headwind reaching about 25 knots. There was no significant roll and the correct PFD speed was set around 30 seconds (see Figure 3) unlike the previous example. Initially, the flight appears to have a high airspeed because of the headwind which became normal after the headwind dropped low which caused the overall drop in airspeed for this flight to exceed our defined threshold of 20 knots. While this is purely external, ADOPT correctly identifies headwind (note that a negative tailwind is a headwind) as the top precursor until about 20 seconds (see Table 2). The corresponding precursor scores and the probabilities are shown in Figure 3. Note that although pitch angle seems to be within nominal bounds, it is listed as a top precursor. One reason could be that, for the given values of high winds, the pitch

[b]Primary Flight Display

Table 1: Precursors ranked by ADOPT at various points during the flight after lift-off. The precursor at the top is contributing most to the high precursor score.

Time = 1s	4s	30s	35s	40s	45s	50s
Tailwind	Pitch Angle	PFD Spd	PFD Spd	Pitch Angle	Pitch Angle	PFD Spd
Pitch Angle	Altitude	Pitch Angle	Tailwind	PFD Spd	PFD Spd	Pitch Angle
Roll Angle	Roll Angle	Auto-throttle	Pitch Angle	Tailwind	Tailwind	Tailwind
Altitude	Tailwind	Roll Angle	Roll Angle	Auto-throttle	Auto-throttle	Auto-throttle
Auto-throttle	Auto-throttle	Altitude	Altitude	Autopilot	Autopilot	Autopilot

could have been managed differently to reduce the speed drop. ADOPT works with all time series variables in combination and is efficient at finding precursors that involve multiple variables simultaneously.

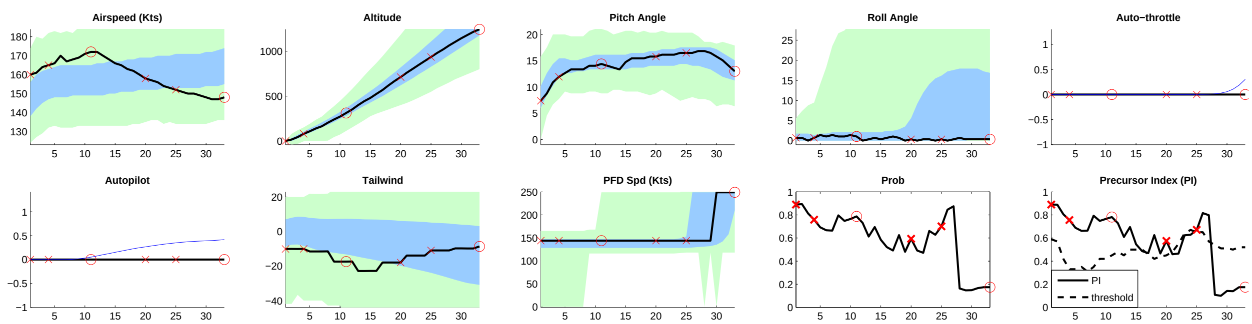

Figure 3: Flight time series variables plotted (thick black curve) against time (seconds) after lift-off. The distribution of nominal values are shown in blue (25-75 percentile) and green (10-90 percentile) as references. The X marker in red shows the precursors identified by ADOPT while the two O markers in red show the start of the speed drop and where the drop reaches 20 knots respectively.

Table 2: Precursors ranked by ADOPT at various points during the flight after lift-off. The precursor at the top is contributing most to the high precursor score.

Time = 1s	4s	20s	25s
Pitch Angle	Tailwind	Tailwind	Pitch Angle
Tailwind	Pitch Angle	Pitch Angle	PFD Spd
Altitude	Roll Angle	Roll Angle	Tailwind
Roll Angle	Altitude	PFD Spd	Roll Angle
Auto-throttle	Auto-throttle	Altitude	Altitude

3. Flight Analysis 3: Large roll at take-off

In this example, the flight takes off at moderate winds and the PFD speed is set correctly around 30 seconds (see Figure 4). However, the flight makes a roll and pitches up simultaneously possibly causing excessive drop in airspeed from about 10 seconds after lift-off. ADOPT lists the pitch angle and roll angle among the top precursors for this flight (see Table 3). The corresponding precursor scores and the prediction probabilities are shown in Figure 4.

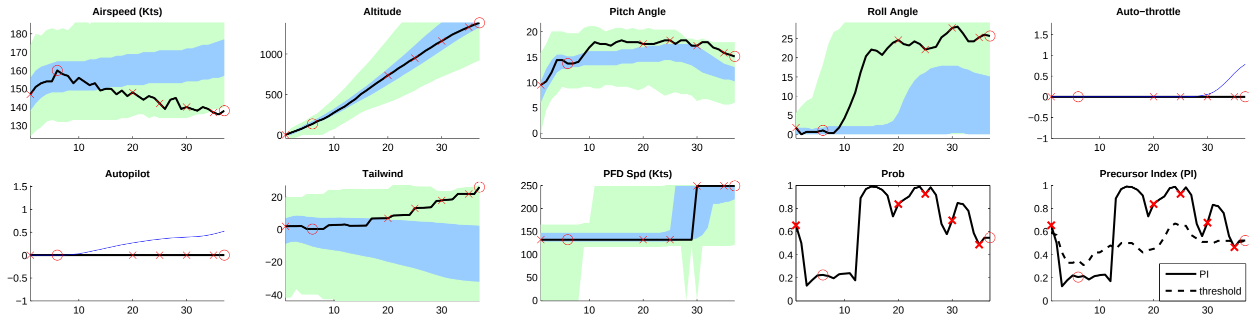

Figure 4: Flight time series variables plotted (thick black curve) against time (seconds) after lift-off. The X marker in red shows the precursors identified by ADOPT while the two O markers in red show the start of the speed drop and where the drop reaches 20 knots respectively.

4. Flight analysis 4: Nominal flight

In this example, most of the flight variables follow nominal values and thus, ADOPT does not find any precursors. Figure 5 shows the flight data as well as the probability and precursor scores. Note that the probability always stays under 0.2 and the precursor score always less than the threshold indicating no precursors.

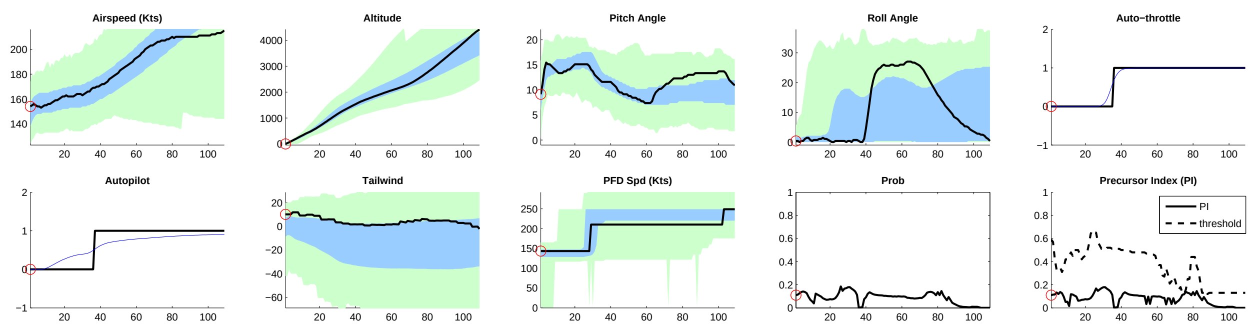

Figure 5: Flight time series variables plotted (thick black curve) against time (seconds). The distribution of nominal values are shown in blue (25-75 percentile) and green (10-90 percentile) as references. The X marker in red shows the precursors identified by ADOPT while the two O markers in red show the start of the speed drop and where the drop reaches 20 knots.

5. Discussion

In this section, ADOPT algorithm was applied to find precursors to the drop in airspeed events during take-off. The precursors identified important latent factors to the airspeed drop events and helped assist the domain experts in our group to identify possible root causes. For majority of the cases, the main precursor was an incorrect setting of the PFD speed reference. For some of the flights we analyzed, human errors prior to takeoff were speculated including possible errors in performance calculations and overcompensation to a high headwind during take-off. While additional investigations are required to confirm the true causes, the case study shows how ADOPT can help distill the large volumes of flight data into a set of focus points in the flight where further human analysis can be performed.

B. STAR Procedure Adherence

The second case study is an example of how ADOPT can be used to quantify performance of flights following standard operating procedures such as the standard terminal arrival routes (STAR) Area Navigation

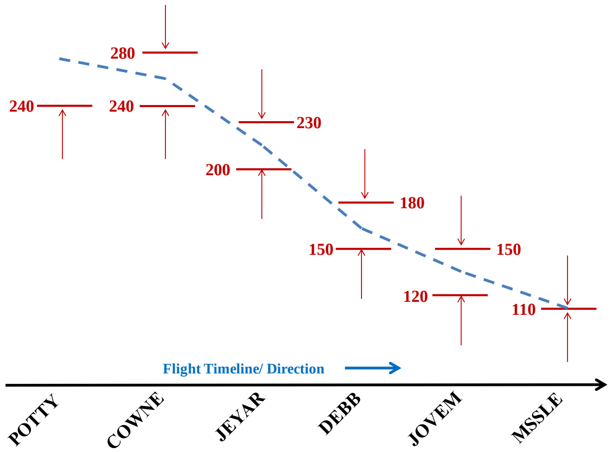

Figure 6: Schematic of the JOVEM ONE procedure with altitude restrictions. The altitudes are shown as hundreds of feet. Note that POTTY has only lower restrictions while MSSLE has an "at" restriction which means the window size is zero.

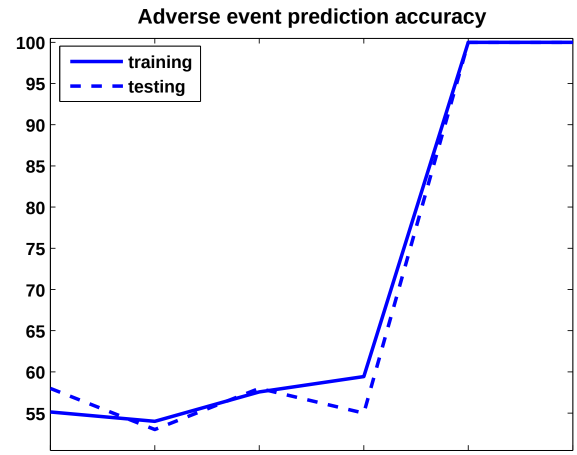

Figure 7: ADOPT's performance on the STAR RNAV adherence case study.

(RNAV). RNAV aims to achieve an optimized flight path, reduce traffic congestions and a streamlined controlled flow of flights to the airport. The performance of STAR RNAV procedures were quantified using a combination of lateral and vertical adherence rates of flights following instrument procedures in the literature.[8] As a case study to demonstrate ADOPT, we consider the vertical adherence as defined in[8] as the adverse event and aim to find precursors. The adverse event is flagged if a flight is above the defined altitude limits of a waypoint. Although intentional non-adherence may not involve any safety risk, any non-intentional ones could be potentially dangerous and create inefficiencies. For this case study, we looked at a subroute (POTTY.JOVEM1) from the JOVEM1 procedure at Dallas-Fort Worth Intl Airport, and applied ADOPT to find precursors to altitude excursion (from above) at the MSSLE waypoint. The sequence of waypoints in the subroute is shown in Figure 6.

The MSSLE waypoint has an "at" restriction where the crossing altitude should be at 11,000 ft within a tolerance of +/-300 ft. The data for this study comes from the Center-TRACON Automation System (CTAS) at NASA, Aviation System Performance Metrics (ASPM), procedure definitions from the Coded Instrument Flight Procedures (CIFP) data defined by the FAA for every 56 day chart cycle release. The time-stamped data is grouped based on flight identifier to get a time series of variables for each fligh. Then, we divided them into nominal and adverse data sets as required by ADOPT; the ones that had an altitude excursion (altitude above the procedure specified value at MSSLE waypoint) constitute the adverse flights. The data had about 35 variables from which a subset were chosen based on domain knowledge. The variables that we chosen include altitude, ground speed, tail wind, descent rate at the waypoint and required slope to reach the next waypoint. In addition, we also included some artificially made binary variables to characterize the flight path including lateral skip feature (to flag if the flight skipped a waypoint laterally), altitude skip feature (to flag if the flight had an excursion above the limits), slope-possible feature (flags if the required slope is higher than the slope defined by the current and the next waypoint), late-entry feature (flags if the flight entered the route from a waypoint other than a defined entry point) and early-exit feature (flag if the flight left the route from a waypoint other than a defined exit point). The data was preprocessed by ADOPT to split the data into a set each for nominal and adverse time series, normalization, cleanup, splitting into training and validation sets etc. The training set had about 400 examples each for nominal and adverse flights while the test set had about 50 examples each.

The SVM model for ADOPT's value function was learned with parameters Gaussian kernel, $\gamma = 5$, $C = 1$ with an overall accuracy of 61%. The accuracy at specific waypoints varies (see Figure 7) as data early in the flight may not have enough information to predict the excursion far in the future. The model was applied on the training data to find the precursor scores for both the nominal and the adverse time series. A simple threshold was learned from this data so that when applied on the precursor score, it rightly classifies the data as being nominal or adverse. The threshold was learned independently at every time step. The corresponding accuracies are shown in Figure 7. It can be seen that the training and testing accuracies follow the same trend. Also, the accuracy is 100% at JOVEM waypoints indicating that the state transitions at JOVEM can predict the altitude excursions at MSSLE with a high accuracy. ADOPT was able to automatically identify this from data. We also analyzed several flights which is discussed below. Similar to the previous case studies, the results were qualitatively evaluated by subject matter experts. In the analysis, the waypoints POTTY, COWNE, JEYAR, DEBB, JOVEM and MSSLE are labeled as P, C, Je, D, Jo, M respectively to make easy appearance in the plots. The adverse event is checked at waypoint M (for MSSLE). The abnormalities at the previous waypoints (P through J) are detected as precursors.

Similar to the previous case study, we refer the flight evaluated by ADOPT as the "test flight". The plots show nominal distribution of parameter values in green (10-90 percentile) and blue (25-75 percentile) as references. The blue area shows what a nominal flight (with no adverse event) values look like. The green area includes more variations in the nominal flights and thus shows possible values that may be slightly off-nominal but still ends up with no adverse event. Note that these distributions are not given to ADOPT during learning and are available only to serve as references. The flight progress (or time) is along the x-axis from left to right where the right most data point corresponds to the adverse event. Markers of type 'X' will be used to indicate precursor time instances in the flight.

1. *Flight Analysis 1: Consistent high altitude profile*

This is an example of a Boeing 747-8 descending with an altitude higher than nominal at the waypoints C, Je, D and Jo (see Figure 8 for the black curve outside distribution). Consequently, the required slope to reach the next waypoint is also higher than nominal (more negative indicates steeper slope down). The ground speed and tailwind are marginally high during these times because of which further descent may have been difficult. The flight not being able to descend faster, has an excursion at waypoint M. ADOPT identifies the flight to have precursors throughout its path in the route and ranks the slope possible feature, descent rate and tail winds as the top precursors (Figure 8 gives the probability of excursion at M while the Table 4 for precursor ranking).

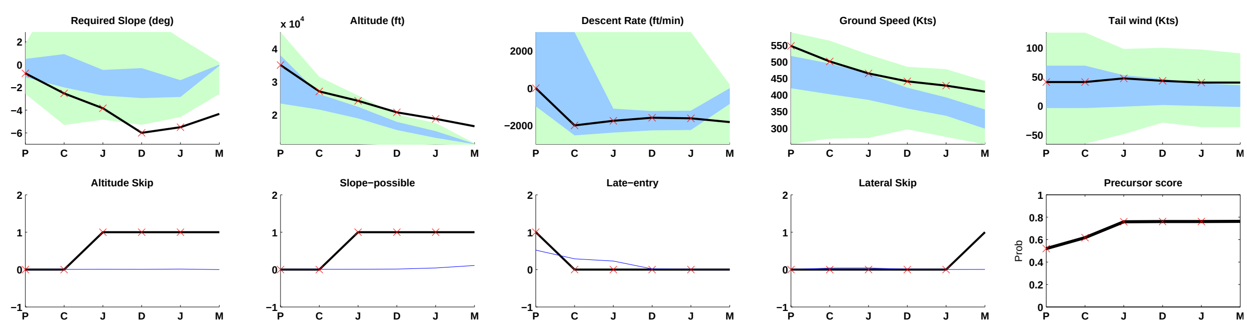

Figure 8: Precursor analysis of a Boeing 747-8 showing flight data as a black curve in each subplot. The nominal data distribution are shown in blue (25-75 percentile) and green (10-90 percentile) as references. The X marker in red shows the precursors identified by ADOPT.

2. *Flight Analysis 2: Marginally high altitude*

In this example, the flight was an Embraer E170 aircraft having nominal tail wind and ground speed with altitudes marginally higher (see Figure 9). However, the flight does not descend fast enough at waypoints D and Jo causing an excursion at waypoint M. At Jo, the slope-possible feature changes to 1 (flagged ON which means the slope required to reach the next waypoint is higher than slope defined by the waypoints). This indicates that the flight was managed sub-optimally from waypoint D to Jo increasing the precursor score. At these waypoints, ADOPT identifies this feature along with improper altitude and descent rates to be the main precursors (see Figure 9 for precursor scores and Table 5 for the precursor ranking). This example shows a case of marginal drift in energy management that is left uncorrected just before the waypoint at MSSLE which ADOPT identified.

Table 3: Precursors ranked by ADOPT at various points during the flight after lift-off. The precursor at the top is contributing most to the high precursor score.

Time = 1s	20s	25s	30s	35s
Tailwind	Pitch Angle	Pitch Angle	Autopilot	Autopilot
Roll Angle	Autopilot	Roll Angle	Roll Angle	Auto-throttle
Pitch Angle	Roll Angle	Tailwind	Auto-throttle	Roll Angle
Altitude	PFD Spd	PFD Spd	Pitch Angle	Tailwind
Auto-throttle	Tailwind	Auto-throttle	Tailwind	Pitch Angle

Table 4: Precursors ranked by ADOPT at various waypoints for the Boeing 747-8 flight example.

P	C	Je	D	Jo
Altitude	Slope-possible	Slope-possible	Slope-possible	Slope-possible
Ground Spd	Descent Rate	Ground Spd	Descent Rate	Altitude Skip
Descent Rate	Tailwind	Descent Rate	Tailwind	Descent Rate
Tailwind	Ground Spd	Tailwind	Ground Spd	Ground Spd
Altitude Skip	Altitude	Altitude	Altitude	Tailwind
Slope-possible	Altitude Skip	Altitude Skip	Altitude Skip	Altitude

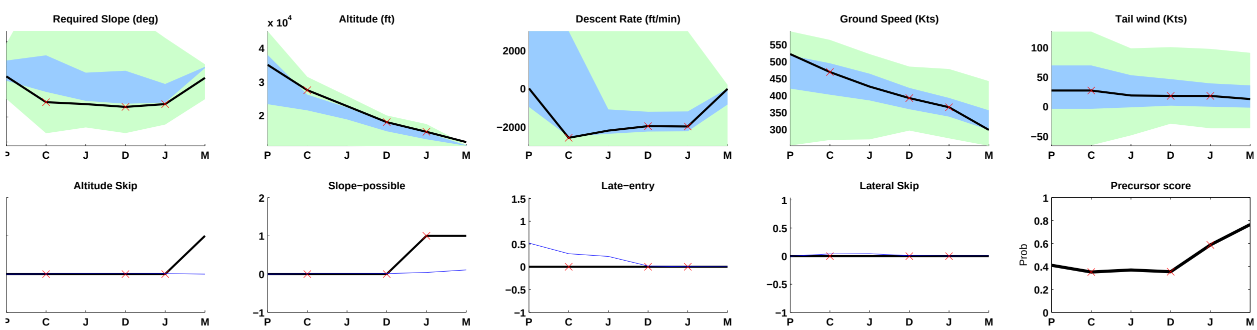

Figure 9: Precursor analysis of an Embraer E170 showing flight data as a black curve in each subplot. The nominal data distribution are shown in blue (25-75 percentile) and green (10-90 percentile) as references. The X marker in red shows the precursors identified by ADOPT.

Table 5: Precursors ranked by ADOPT at various waypoints for the Embraer E170 flight example.

C	D	Jo
Altitude	Slope-possible	Altitude Skip
Tailwind	Altitude	Slope-possible
Descent Rate	Descent Rate	Descent Rate
Ground Spd	Tailwind	Tailwind
Altitude Skip	Ground Spd	Ground Spd
Slope-possible	Altitude Skip	Altitude

3. Flight Analysis 3: Nominal flight having precursors but corrected

This is an example of flight Embraer E170 aircraft that has some mild factors but are corrected just before the excursion checkpoint at M. The tailwind and ground speed are moderate but the transition from waypoint D to Jo does not have sufficiently descent rate which makes it high in altitude at waypoint Jo. This is very similar to the previous example. However, this was corrected before it reached waypoint M making sure the altitude was within acceptable window. This is an example where a marginal suboptimal state transition is seen at waypoint D which was corrected between Jo and M. ADOPT identifies altitude, slope-possible and descent rates as the top precursors for waypoint D shown in Table 6. The flight data distribution and prediction probabilities are shown in Figure 10.

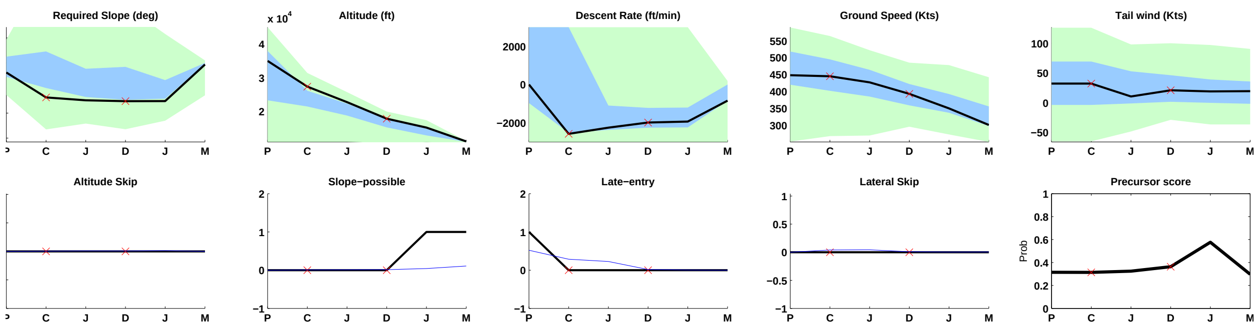

Figure 10: Precursor analysis of an Embraer E170 showing flight data as a black curve in each subplot. The nominal data distribution are shown in blue (25-75 percentile) and green (10-90 percentile) as references. The X marker in red shows the precursors identified by ADOPT.

Table 6: Precursors ranked by ADOPT at various waypoints for the Embraer E170 flight example.

C	D
Altitude	Slope-possible
Ground Spd	Altitude
Descent Rate	Descent Rate
Tailwind	Tailwind
Altitude Skip	Ground Spd
Slope-possible	Altitude Skip

4. Flight Analysis 4: Nominal flight with no precursors

This is an example of a Boeing 737-8 flight that was nominal in all variables (see Figure 11) and thus ADOPT in this case identifies no precursors.

5. Discussion

While the take-off stall hazard is an example of high frequency data (measured every second) , the RNAV case study is an example of a coarsely sampled data where the data is recorded only at waypoints (measured every tens of seconds to minutes). This is reflected in ADOPT's low accuracy for the RNAV study as the level of uncertainty increases with coarse sampling. For the RNAV case study with altitude excursions from above at MSSLE waypoint, it appears that majority of the flights failed to descend fast enough at previous waypoints could be the main reason for not able to descent enough at MSSLE. Sometimes this was corrected while sometimes not depending on many factors including ground speed, aircraft type, controller communication etc. For some cases, a combination of altitude and ground speed is higher than nominal which doesn't give any room for the crew to descend at the required rate, causing an excursion. The recovery in

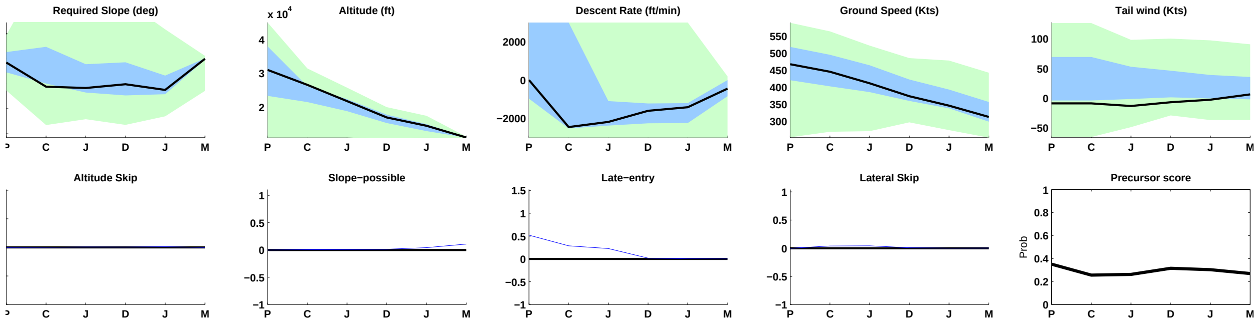

Figure 11: Precursor analysis of a Boeing 737-8 showing flight data as a black curve in each subplot. The nominal data distribution are shown in blue (25-75 percentile) and green (10-90 percentile) as references. The X marker in red shows the precursors identified by ADOPT.

this case would be difficult. It would be interesting to know the operational factors involved such as if there was restriction from the controllers on the ground or if the crew was unaware etc. Unfortunately, it is not easy to make a conclusion without having the voice communication or other data to confirm. However, for the case study, ADOPT was able to take time series flight data and just minimal knowledge of the adverse event's occurrence, and able to identify precursors to such incidences.

IV. Literature Review

Precursor analysis is a recent area of research and very few algorithms exist that can automatically mine the data and identify precursors.[4,9] Other data mining algorithms such as rule mining,[10,11] causal mining[12] and motif mining[13] may be adapted to find precursor rules. However, such methods are often limited because of the exponential growth in the number of rules[14,15] for systems involving multiple sensory variables such as Aviation. This problem is prevented by directly operating in a continuous vector space which avoids the need for discretization of continuous data. Relevant to Aviation, precursors are usually analyzed using anomaly detection methods.[16] For example, issues in flight deck human-automation interactions may be detected using anomaly detection.[17,18] Our prior work involved finding precursors to go-arounds,[4] high-energy approach and landing issues[19] and stall hazard[6] using the ADOPT algorithm discussed in this paper. This paper is an extension of the previous works that includes several scenarios of safety precursors under different problem setups.

V. Conclusions

ADOPT is a recently developed data mining algorithm that can help discover latent factors indicative of adverse events using data. A characteristic feature of ADOPT is in its ability to observe the temporal nature of the problem and retrospectively identify the time instances corresponding to precursors. This paper detailed the application of the ADOPT algorithm to find aviation safety precursors using real operational data. The ideas behind precursor discovery is described followed by analysis using two adverse events including take-off stall hazard and RNAV procedure adherence. For each of the cases, the motivation, problem setup, data and model building are discussed followed by flight analyses that demonstrate the use of ADOPT algorithm in finding precursors. The results are validated qualitatively by subject matter experts and precursors are identified for the two adverse events early in the flight which gives significant lead time before the incidents occur. A list of features and benefits of ADOPT are summarized below

- The data input to ADOPT include a set of time series data leading up to the adverse event and a set of nominal time series data. The data may be multivariate and may contain continuous and categorical variables. ADOPT operates in a continuous vector space and needs no prior discretization of variables.

- ADOPT's results include a set of time instances in the time series data that correspond to precursors, a ranked list of variables that are precursors at each of the time instances, recommended actions (increase, decrease or stay constant) for each of the precursor variables, probability of the adverse event at each time-step useful for further risk analysis and forecasting.

- Any domain knowledge may be used for feature selection or select time periods for ADOPT to analyze which speeds up modeling.
- The underlying model in ADOPT may be any classifier of choice. In this paper, we used support vector machines (SVM) but any other classification model such as decision trees or neural networks may be used.
- ADOPT algorithm can be modified to find precursors to multiple adverse events simultaneously.
- ADOPT can be easily parallelized by performing precursor discovery on independent batches of flights.
- A Python based open source version of ADOPT algorithm will be made available in the future for further development and industrial applications.

ADOPT is a general algorithm that can be applied to a wide variety of aviation adverse events. It has to be noted that ADOPT only identifies correlation and not causation. While the precursors may indicate an impending safety event by showing an abnormality in certain measured quantities, further investigations are usually required to identify root causes. However, ADOPT identifies the time instances in the flight as well as the variables in ranked order as precursors, which helps speed up this process. In other words, ADOPT "simplifies" the data from a flight or an airspace (often in the order of hundreds of variables sampled at a high frequency) and outputs a condensed information set (the precursors) so that safety analysts can start from the precursors instead of the original data. Future work will focus on implementing the knowledge base of ADOPT in an online system that proactively detects safety precursors and recommends corrective actions to pilots or controllers to aid in proactive safety management.

Acknowledgments

This research is supported by the NASA Airspace Operation and Safety Program. The authors would also like to thank the subject matter experts Dr. Michael Feary, Mike Stewart and John Kaneshige for their insightful comments and perspective on the identified precursors.We would also like to thank the NASA ATM Data Warehouse for access to some of the data used for the case studies.

References

[1]"The Bureau of d'Enquetes et d'Analyses (BEA) pour la securite de l'avation civile. Final Report: On the accident on 1 June 2009 to the Airbus A330-203 registered F-GZCP operated by Air France flight 447 Rio de Janeiro-Paris." July 2012.

[2]"National Transportation Safety Board Accident Report. Loss of Control on Approach Colgan Air, Inc. Operating as Continental Connection Flight 3407 Bombardier DHC-8-400," *Report Number NTSB/AAR-10/01, PB2010-910401*, N200WQ Clarence Center, NY, USA, February 2009.

[3]Maurino, D., Reason, J., Johnston, N., and Lee, R., *Beyond Aviation Human Factors: Safety in High Technology Systems*, Taylor & Francis, 2017.

[4]Janakiraman, V. M., Matthews, B., and Oza, N., "Discovery of Precursors to Adverse Events using Time Series Data," *Proceedings of the 2016 SIAM International Conference on Data Mining*, Miami, FL, USA, 2016.

[5]"System Safety Process Steps," .

[6]Janakiraman, V. M., Matthews, B., and Oza, N., "Finding Precursors to Anomalous Drop in Airspeed During a Flight's Takeoff," *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2017, pp. 1843–1852.

[7]Granger, C. W., "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica: Journal of the Econometric Society*, 1969, pp. 424–438.

[8]Stewart, M. and Matthews, B., "Objective Assessment Method for RNAV STAR Adherence," *Proceedings of the IEEE/AIAA 36th Digital Avionics Systems Conference (DASC)*, AIAA, 2017.

[9]Ning, Y., Muthiah, S., Rangwala, H., and Ramakrishnan, N., "Modeling Precursors for Event Forecasting via Nested Multi-Instance Learning," *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, ACM, New York, NY, USA, 2016, pp. 1095–1104.

[10]Laxman, S. and Sastry, P., "A survey of temporal data mining," *Sadhana*, Vol. 31, No. 2, 2006, pp. 173–198.

[11]Agrawal, R. and Srikant, R., "Mining Sequential Patterns," *Proceedings of the Eleventh International Conference on Data Engineering*, ICDE '95, IEEE Computer Society, Washington, DC, USA, 1995, pp. 3–14.

[12]Karimi, K. and Hamilton, H., "Discovering Temporal/Causal Rules: A Comparison of Methods," *Advances in Artificial Intelligence*, edited by Y. Xiang and B. Chaib-draa, Vol. 2671 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2003, pp. 175–189.

[13]McGovern, A., Rosendahl, D., Brown, R., and Droegemeier, K., "Identifying predictive multi-dimensional time series motifs: an application to severe weather prediction," *Data Mining and Knowledge Discovery*, Vol. 22, No. 1-2, 2011, pp. 232–258.

[14]Das, G., ip Lin, K., Mannila, H., Renganathan, G., and Smyth, P., "Rule Discovery From Time Series," AAAI Press, 1998, pp. 16–22.

[15]Chiu, B., Keogh, E., and Lonardi, S., "Probabilistic Discovery of Time Series Motifs," *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '03, ACM, New York, NY, USA, 2003, pp. 493–498.

[16]Das, S., Matthews, B. L., Srivastava, A. N., and Oza, N. C., "Multiple kernel learning for heterogeneous anomaly detection: algorithm and aviation safety case study," *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, ACM, 2010, pp. 47–56.

[17]Lee, S., Hwang, I., and Leiden, K., "Flight deck human-automation issue detection via intent inference," *Proceedings of the International Conference on Human-Computer Interaction in Aerospace*, ACM, 2014, p. 24.

[18]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*, Dec 2013, pp. 407–412.

[19]Janakiraman, V. M., "Explaining Aviation Safety Incidents Using Deep Learned Precursors," *CoRR*, Vol. abs/1710.04749, 2017.