Enabling Safety from Data: Machine Learning/Artificial Intelligence for In-Time Aviation Safety

Nikunj C. Oza, Ph.D., Chad Stephens
NASA System-Wide Safety Project

Modern jet airliners record nearly one gigabyte of raw data per flight, nearly double that recorded by the previous jet airliners brought into service less than ten years ago. Given this treasure trove of data, data analysis is an ever-important capability to convert these data into knowledge that permits understanding and achieving safe operations. The practice of Data Analytics involves applying Artificial Intelligence (AI) and Machine Learning (ML), among other approaches, to derive insights and identify meaningful relationships in the data.

AI is the field of study focused on developing simulated human intelligence in computer-based agents. ML, a subdiscipline of AI, involves development of prediction or decision algorithms not explicitly programmed to predict or decide but rather that learn from data representing past predictions or decisions. You may have experienced ML-enabled capabilities such as customized recommendations in Netflix or Amazon. Virtual assistants, such as Apple's Siri or Amazon's Alexa, and partially or fully-autonomous vehicles are possible due to the ability of ML algorithms to learn from past operations.

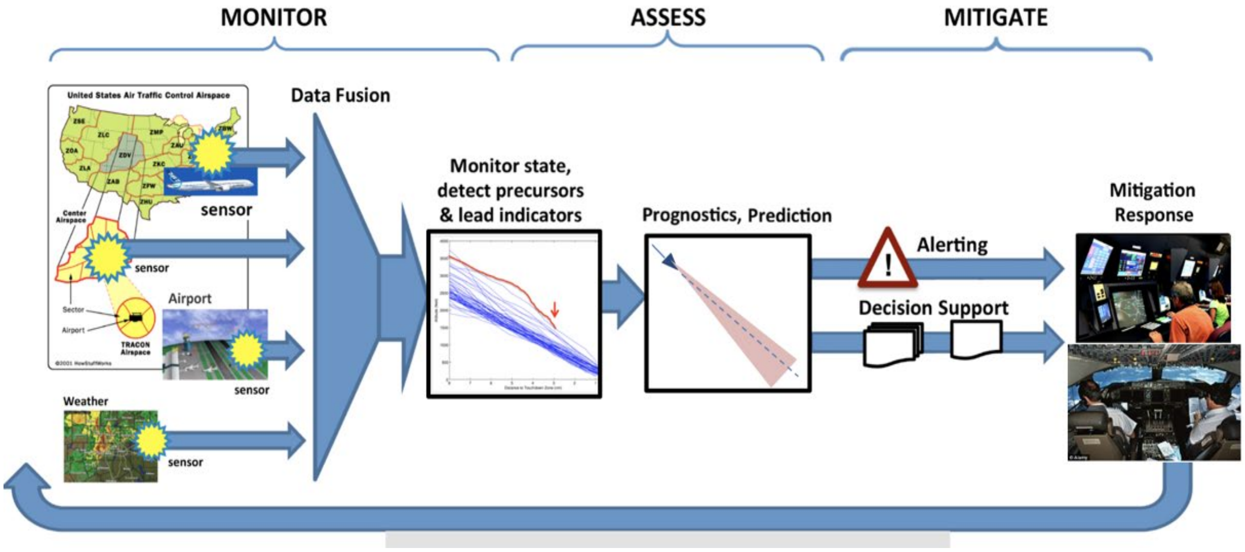

Figure 1

Our team of NASA Data Scientists in the System-Wide Safety Project have been developing ML algorithms to discover patterns and relationships that were previously undetectable by exceedance-based methods or traditional simple statistical models. Figure 1 shows a high-level view of the process of identifying and mitigating safety issues. The steps involve using data sources to monitor operations, performing analyses that allow for assessing the nature of operations and determining if safety issues are present or likely to arise, and then taking actions to mitigate any safety issues. Currently, flight safety monitoring is mostly done using exceedances, which are rules, typically over a few variables that describe conditions that are best avoided, such as a drop in airspeed during takeoff or excessive speed on approach at 1,000 feet altitude. Such exceedances are clearly effective at finding known safety issues, but are not designed to look for vulnerabilities, which are previously unknown safety issues. Additionally, exceedance thresholds may have been ideal when they were set but are not necessarily ideal in current operating conditions.

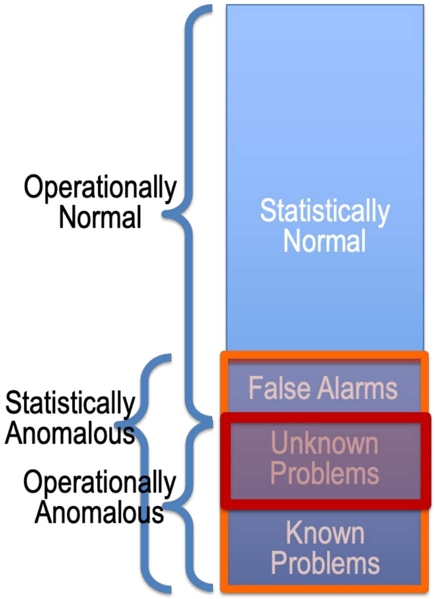

Figure 2 (not to scale)

Data analytics aims to “let the data speak” by using techniques from ML, data mining, statistics, and related fields to transform data into knowledge of the system that generated the data. Our team is developing ML algorithms for anomaly detection, which involves identifying those few data points that are unusual or “stick out” compared to most of the data that represent normal operations. These methods are relevant for vulnerability discovery because both known and unknown problems, which are operationally anomalous (see Figure 2), are relatively rare and therefore, are among the statistically anomalous data points that anomaly detection methods would find. The process of vulnerability discovery involves using an anomaly detection algorithm to find the statistical anomalies, removing those data points representing known problems, and examining the remaining statistical anomalies to separate the vulnerabilities (unknown problems) from the false alarms. The problem of having too many false alarms is a well-known problem in anomaly detection. One can choose to restrict the algorithm’s definition of statistically anomalous to reduce the number of false alarms, but this raises the chance of missed detections---operationally anomalous data points that are marked statistically normal---which is the problem with exceedance-based methods.

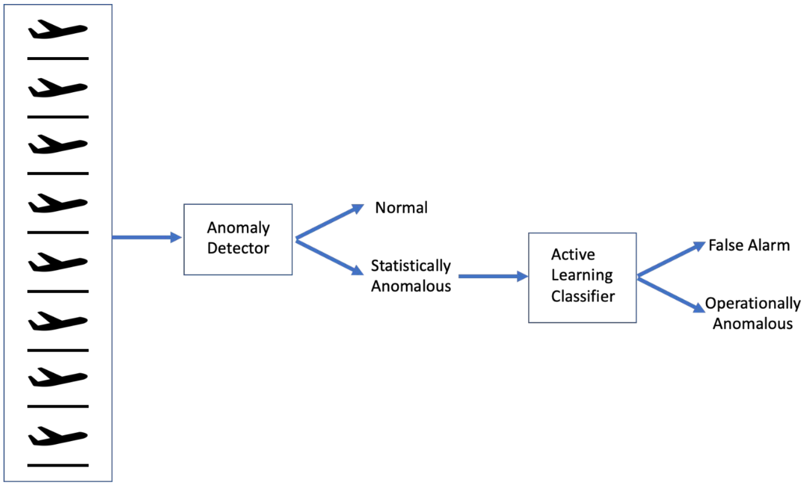

Figure 3

To reduce the problem of false alarms, we are working on the use of active learning methods to learn from domain experts how to distinguish operationally anomalous problems from false alarms. Active learning is an area of ML that accounts for the cost of getting labels for data---in this case, having domain experts spend time to label data points as operationally anomalous or false alarms---and only asks domain experts to label those data points that are most helpful in learning how to distinguish between the labels. After learning and during operations, new data points are processed as described in figure 3---each data point is passed to the anomaly detector. If the detector decides that the data point represents normal operations, then nothing more is done. If the data point is deemed statistically anomalous then it is passed to the classifier that results from active learning. The classifier determines whether the statistically anomalous point is also operationally anomalous or a false alarm.

We would also like to identify precursors to known adverse events. These precursors are conditions after which the known adverse event is more likely to occur. For example, for excessive speed on approach at 1,000 feet, we may be able to identify a threshold on speed at top of descent beyond which the probability of violating the high-speed exceedance is unacceptably high. In figure 4, we would prefer to find the precursor depicted by p_1, which is the earliest point at which one action has an unacceptably high probability of reaching the hazard state while the other action has a high probability of reaching the safe state. We are also interested in finding corrective actions, which correct for an action that could lead to a hazard state. We are developing algorithms that allow users to specify the hazard states vs. safe states and automatically find the precursors like p_1.

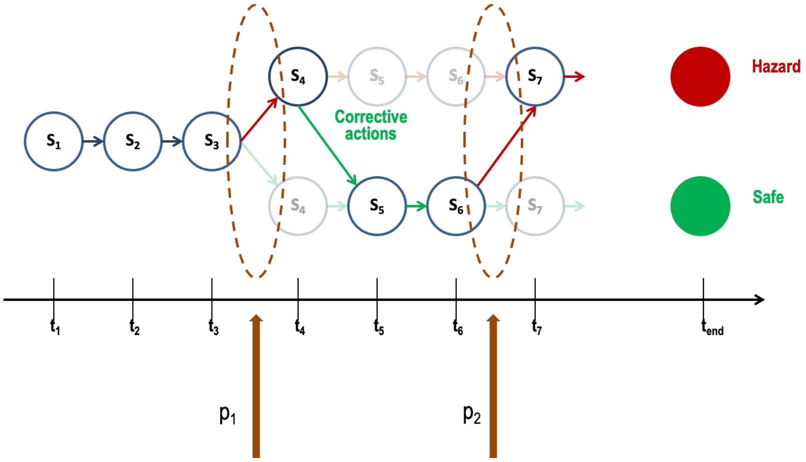

Figure 4

Airlines invest significantly in collecting and storing data that describe many aspects of operations, including flight operations, airspace, and maintenance. These data can suggest how an airline is really operating versus how the airline thinks it is operating, as well as how the operations are changing over time. Data analytics methods, such as those described above, can contribute to this understanding. The examples given above are largely applied to numeric data such as FOQA and radar track data. We are collaborating with airlines on data analytics to extract knowledge from the substantial data that they collect, including not only FOQA, but also ASAP, LOSA, and others. The goal of this work is to:

- Assess how beneficial different data sources can be at identifying vulnerabilities and precursors by themselves and jointly,
- Determine what ML techniques can be used on these data sources,
- Identify what visualization techniques we can use to reveal these vulnerabilities to domain experts (e.g., pilots, IOC personnel) in the most intuitive way possible, and
- See how best to incorporate their feedback on the operational significance of the identified vulnerabilities and which subsets of the data are most relevant.

The greater the number of data sources and amount of data examined, the greater the knowledge that can be gained and the more the different parts of an airline's operations can coordinate to mitigate vulnerabilities.

The data analytics work described so far is being done offline, on data representing operations that are not currently in progress, and will generate knowledge based on those data. A second aspect of our data analytics work involves using the generated knowledge, in the form of models generated by the ML techniques, to monitor operations in progress to identify potential vulnerabilities and precursors when or before they happen. This will require us to develop software that operates in the Integrated Operations Center (IOC), ingests data from the data feeds that the IOC receives, and runs the ML models on these data to indicate whether the data represent operationally significant anomalies or precursors. We will also develop software to allow IOC personnel to visualize the ML results and related data so that they can take appropriate actions. We will develop the tools needed to allow the results obtained during operations to inform and update the ML models developed offline. This will always be necessary as the nature of operations continually changes.

Our Data Analytics collaboration will allow airlines to use the data already being collected to better understand how operations are proceeding, leading to more effective actions and mitigations of safety issues as well as a more accurate assessment of how beneficial these actions are and when these actions need to change. These analyses will be done while keeping the data resident on the airlines' systems and maintaining confidentiality and anonymity of the data. We expect the collaboration to allow us to develop ML and Visualization tools and related software that are more helpful to U.S. commercial aviation than what they currently use and thereby yield greater value in terms of safety, efficiency, and passenger comfort.