Fleet Level Anomaly Detection of Aviation Safety Data

Santanu Das
UARC, UC Santa Cruz
Nasa Ames Research Center
Moffett Field, CA 94035
Email: Santanu.Das-1@nasa.gov

Bryan L. Matthews
SGT Inc.
Nasa Ames Research Center
Moffett Field, CA 94035
Email: Bryan.L.Matthews@nasa.gov

Robert Lawrence
Consultant/Contractor
Nasa Ames Research Center
Moffett Field, CA 94035
Email: captndaddy@gmail.com

***Abstract*—The National Airspace System (NAS) is a highly complex dynamical system that is regularly sensitive to faulty aeronautical decision-making and risk management skills. In order to ensure a healthy system with safe flights a systematic approach to anomaly detection is very important when evaluating a given set of circumstances and also simultaneously determine the best possible course of action. Given the fact that the NAS is a vast and loosely integrated network of systems, it requires improved safety assurance capabilities to maintain an extremely low accident rate under increasingly dense operating conditions. Data mining based tools and techniques are required to support and aid operators' (such as pilots, management, or policy makers) overall decision-making capacity. The ability to analyze fleetwide data autonomously is still considered a significantly challenging task. Here, in this effort, we aim at developing a system level analysis scheme. In this paper we address the capability for detection of fleetwide anomalies as they occur, which itself is an important initiative towards the safety of the real-world flight operations. The flight data recorders archive millions of data points with valuable information on flights everyday. The operational parameters consist of both continuous and discrete (binary & categorical) data from several critical sub-systems and numerous complex procedures. In this paper, we discuss a systematic system level anomaly detection approach based on the theory of kernel learning to detect potential safety anomalies in a very large data base of commercial aircraft. We also demonstrate that the proposed approach uncovers some operationally significant events due to environmental, mechanical, or human factors issues in high dimensional, multivariate Flight Operations Quality Assurance (FOQA) data. We present the results of our detection algorithms on real FOQA data from a regional carrier.**

I. Introduction

Suppose the entire data that is available for a given system, covering both inputs and outputs, is given by the set $(\mathcal{U}, \mathcal{Y})$. In real-world flight operations, the input $\mathcal{U}$ is provided by the pilot following some standard-operating procedure: for a given airport, aircraft, weather conditions, instructions from air traffic control, and other contextual elements of the flight, the flight procedures are well determined. The observed characteristics $\mathcal{Y}$ are representive of the flight behavior and hold a complex relationship with the input $\mathcal{U}$. Such a system that can be functionally described by the following equations:

$$\begin{aligned} \mathbf{h}_t &= \Gamma(\mathbf{h}^*_{t-1}) \\ \mathbf{x}_t &= \Psi(\mathbf{x}^*_{t-1}, \mathbf{h}^*_t, u_t) \\ y_t &= \Omega(\mathbf{x}_t) \end{aligned} \tag{1}$$

Equation 1 describes a system with u_t as the observed system input, and y_t is the observed system output which can take any form of discrete, categorical, and continuous features.Here we assume that the function Γ and Ψ are unknown. The function Γ and Ψ determines the evolution of the hidden system state $\mathbf{h}_t$ and governs the evolution of the continuous state vector respectively. We assume that the vector $\mathbf{x}$ is an N dimensional state vector, and $\mathbf{x}^*_{t-1}$ is its history for the last D time steps: $\mathbf{x}^*_{t-1} = [\mathbf{x}_{t-D}, \mathbf{x}_{t-D+1}, ..., \mathbf{x}_{t-1}]$. The hidden state $\mathbf{h}_t$ is assumed to correspond to different mode configurations within the system and each mode affects the output dynamics Ψ. Any malfunction in the system, $\mathbf{h}_t$ could move to an abnormal state, thus also changing the nature of the observed output. Not necessarily $\mathbf{h}_t$ always has to move to an abnormal state, it can very well reside in some unfamiliar state within the normal operational regime. In practice a similar situation will arise when all pilots are attempting to follow the same standard operating procedures, while some of them may deviate from these procedures which could lead to a different input sequence $\mathcal{U}'$, resulting in a different set of observed flight characteristics $\mathcal{Y}'$ which may be unusual but not necessarily wrong. The problem that we address in this paper is to develop a method to discover whether or not the current observed vector $\mathcal{Y}$ along with the discrete pilot inputs $\mathcal{U}$ represents one that is atypical or an anomaly based on the observed history of the system. We want to emphasis here that in this particular research we

don't intend to model causal behaviour or relationship between input and output variables of such systems. Rather the objective is to simultaneously identify unusual patterns in any combinations of input ($\mathcal{U}$) and output ($\mathcal{Y}$).

Over the last few decades with improved sensing capabilities we have seen tremendous increase in information flow, in terms of the volume and complexity, created at an unprecedented pace in several disciplines. Aviation industry is not an exception. A great example is the Distributed National FOQA (Flight Operations Quality Assurance) Archive (DNFA) data base established by Nasa. DNFA contains millions of flight data from most of the major carriers in the U.S. Typical FOQA parameters consist of both continuous and discrete data from the avionics, propulsion system, control surfaces, landing gear, the cockpit switch positions, and other critical systems. These data sets can have up to 500 parameters and are sampled at 1 Hz. For a moderate sized fleet that operates 1000 flights per day, these FOQA data sets become very large. Today, we are left with the challenge of dealing with such a vast amount of heterogeneous information resources in varieties of semantic structures. Knowledge discovery from these heterogeneous resources is still a challenging task. With this increased complexities of data sources there is a potential need of building intelligent refinement and integration frameworks, focusing on information content and semantics.

The key aspect of any data analysis method depends on how the input data was measured within the process and transformed into information. In aviation domain the data sources are categorized as structured data and unstructured data. A typical example of unstructured data is free text data, for example reports or scripts from pilots describing some events, experts feedback on the process etc. However in this paper we will restrict our discussion to analysis using structured data which has two main categories, continuous and discrete data. Discrete attributes can be either binary, categorical or logical order i.e. sequential in nature.

In our earlier research we have demonstrated the potential of the multiple kernel based anomaly detection algorithm [9] in detecting anomalies. In this paper we attempt to expand on the analysis by reporting a variety of different interesting anomalies that we have observed in a commercial aircraft data. Here we conduct the analysis on a much large-scale data. In additional we provide some useful insights in some disgnosis information with a resonable explanations. We demonstrate the capability of the proposed methodology in terms of it's flexiblity to process a variety of heterogeneous data sources without reformulating the problem whenever there is a change in information content or data structure. The formulation we demonstrate in this paper is very simple and can easily be adopted for fleetwide analysis in various domains including medical applications, airspace safety, business analysis etc.

II. Anomaly Tracking Algorithms

Since the theme of this paper is anomaly detection, we will mostly discuss in the context to anomaly detection which is also known as outlier detection or surprise pattern detection. Outlier or anomaly detection refers to the task of identifying new or unknown patterns which, in many cases, are abnormal or inconsistent. The problem of outlier detection has been extensively studied using several approaches [14], [15], [16], [7]. Supervised and un-supervised are the two broader categories. In supervised approach a model is built for detection purpose and this model assumes known class labels. Typical classification based techniques such as Bayesian inference, decision trees, Support Vector Machines (SVMs)or neural networks models are built on previously labeled instances of both normal and abnormal data instances. However class labels are expensive and they are not easily available, especially for most of the historical data. Given the fact that it is impossible to always have prior knowledge about all possible classes or have known data labels or have data representing all possible scenarios or classes, unsupervised techniques hold an edge in many applications and play an important role in identifying "*what is desired and what is not*" in a dataset. Novelty detector is such a specialized tool that classifies the members of a given set of objects into two groups on the basis of whether the model have seen those objects before or not. Kernel based classification methods like single class SVMs, one-class kernel Fisher Discriminants etc fall under novelty detection category and are unsupervised in nature. In these techniques a model is built on the normal data and the idea is to come up with a threshold for determining abnormality and using a distance based score for evaluating the extent of abnormality. In nearest neighbor based approach, the aim is to infer the outliers based on the data itself e.g. by finding those points which are at a greater distance from most of the other data points or by finding those points which are in low density region. Since in most cases k-nn based solutions have quadratic time complexity, researchers like Angiulli and Pizzuti [3], Angiulli and Fassetti [2], Ramaswamy et al. [17], and Bay and Schwabacher [5] have proposed promising techniques with improved run time.

A. Production Level Implementations

In airline industry algorithms that are chosen for production level implementation must be heavily tested and developed to produce reliable and meaningful results before they are selected by the airlines for everyday use. The most widespread method for detecting operationally significant anomalies is with domain expert defined threshold exceedances. Methods for detecting unstable and high energy approaches are commonly exceedance based. Another very similar technique that is currently in use is just detecting the min and max for a given parameter. Any deviation from that range may be of interest. These methods have been in use for as long as the FOQA program has been in existence and has provided analysts with valuable results. This is in part due to the fact that the exceedance events identified are easily interpretable since the user has defined what the algorithm is looking for. Another attractive feature is the highly scalable implementation of the algorithm on large data sets. However, the drawback to exceedance based analysis is the fact that typically only the anomalies that are defined are reported to the analyst, leaving the undefined anomalous events undetected. In other words the method only answers the questions that someone thought to ask. Algorithms discussed in this paper address the issue of detecting ”unknown anomalies”. However, they are still in the research level stage, but with gaining interest by the airlines may one day be running in a production level environment.

B. Research Level Implementations

In this section we will describe some anomaly detection techniques that have been extensively used to analyze FOQA data. Some of these algorithms include Morning Report, Orca, IMS, SequenceMiner, and one-class Support Vector Machines. These research level algorithms help compliment the exceedance based methods by being able to identify the ”unknown anomalies”. Once the anomalies found are analyzed, new parameter exceedances can be developed and incorporated into the airlines’ daily analysis to track future or past events.

1) Morning Report: Morning Report [20] & [1] is an algorithm designed to detect atypical flights over a set of aircraft and identify the contributing anomalous parameters and phases of flight. The algorithm calculates statistical signatures across the parameters of a given flight and clusters the flights based on the multivariate signatures. Similar flights are grouped together and atypical flights are considered to be far away from a cluster and therefore have higher scores. The distribution of the anomaly scores are a function of the Mahalanobis distance from centroids of the multivariate cluster. The results provide both the degree of the anomalous flight along with the contributing parameters, which can be useful for the analysts. The algorithm was designed to be executed on a large set of flights overnight and return the results the next morning (hence the name Morning Report).

2) Orca: “Orca” [5] is a method used for detecting anomalies in both continuous and discrete (binary format) data in vector space, using a nearest neighbors based approach to detect anomalous points. For continuous data, Orca takes a nominal reference data set and calculates the nearest neighbors’ using euclidean distance to all test points in the original vector space. For binary data points the hamming distance [21] is used. “Orca” is a k-nearest neighbor based algorithm adopting nested loop structure in conjunction with randomization and simple pruning rule. Pruning used in this algorithm helps in achieving near linear time performance with high dimensional data. This makes the algorithm scalable for analyzing large data sets. The algorithm uses a distance-based metric for finding outlier by examining the distance of any test point to k existing examples those considered as nearest neighbors. If one looks at the local neighborhood and finds that the test points are relatively close, then the examples are considered normal or else unusual. In this algorithm, each data point is scored independently and therefore anomalies in the temporal domain are undetectable. The pseudo code of “Orca” is shown below.

3) IMS: The Inductive Monitoring System (IMS) [11] is a distance base anomaly detection tool that uses an unsupervised clustering algorithm to build models of the expected operation of the system on a set of nominal data. The models are used to test new data to determine whether an anomaly is present or not. The underlying concept states that if the system behaves similar to the normal operating modes that the data was trained on, the distance scores will be lower than data that is generated from a system that is in an anomalous state. IMS evaluates each sample, which is a multivariate vector, by calculating the Euclidean distance to the cluster bounds of each cluster in the model, and reporting the distance to the closest cluster as the anomaly score. A 2D representation in figure II-B3 can be seen. The normal operating regions are defined by the two boxes, with the distances computed to the edge of the nearest box. In the context of flight data IMS will train on a set of nominal flights, either identified by domain experts or another algorithm, and test on the remaining flights. Each time point within the test flight will be evaluated,

Algorithm 1 Orca Algorithm

Input:
A_{pq} : Matrix with q dimensional dataset having p instances arranged in a random order
k : Number of nearest neighbors (default 5)
n : Number of outliers used (default: $n = p$)
Output:
$O(n)$: Set of outliers
S_{global} : Global scores
Argument:
a : Entries in A
B : Block of examples from A
b : Entries in B
C : Cut-off threshold
w_d : Weight of discrete parameters
Definitions:
$x \in (x_c, x_d)$
$D \in (D_c, D_d)$
$H_d(x_d, D_d) = w_{d_1}(x_{d_1}! = D_{d_1}) + w_{d_2}(x_{d_2}! = D_{d_2}) + \ldots + w_{d_n}(x_{d_n}! = D_{d_n})$
$d(x, D) = \sqrt{(x_{c_1} - D_{c_1})^2 + \ldots + (x_{c_n} - D_{c_n})^2} + H_d(x_d, D_d)$
$d(x, D)$: maximum distance between x and an example in D
$M^k_{x,D}$:k closest example in D to x
$S(D, x) = \frac{1}{m}\sum_{i=1}^{m} d(x, D)$ distance based score
Initialize:
Let p instances of A_{pq} be divided in N_B blocks and $K_{nn}(x)$ be the matrix that keeps track of the nearest neighbors/examples of x. And $C = 0$ and $O = \Phi$, where Φ is a null vector
For $block = 1 : N_B$ {
$B = A(:, block); K_{nn}(b) = \Phi$
For each a in A, {
For each b in B and $b \neq a$ {
If $Length(K_{nn}(b)) < k$ or $d(b, a) < dm(b, K_{nn}(b))$ {
$K_{nn}(b) \leftarrow M^k_{b,K_{nn}(b)\cup a}$
If $S(K_{nn}(b), b) < c$ {
Remove example b from set B
}}}}
$O = O \cup B$
$S_{global} = score(O)$
$C \leftarrow min(S_{global}(0))$
}

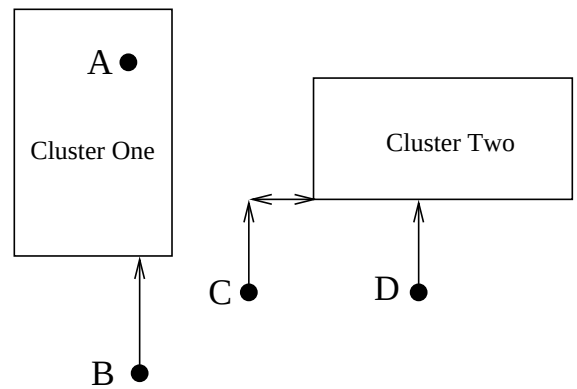

Fig. 1. A conceptual diagram to describe the working principle of IMS.

producing a profile of anomaly scores for each flight. The anomaly scores for each flight can be combined in many different ways, however, typically the scores are averaged for each flight and the flights with the highest average score are ranked most anomalous. As with Orca IMS evaluates each point independently and therefore suffers from the same drawback of not being able to detect anomalies in the temporal domain. The pseudo code for IMS is shown in figures **??** & **??**.

4) SequenceMiner: SequenceMiner [6] was developed to address the problem of detecting and describing anomalies in large sets of high dimensional symbol sequences such as recordings of switch sensors in the cockpits of commercial aircraft. SequenceMiner works by first using an unsupervised clustering algorithm to cluster the sequences using the normalized longest common subsequence (LCS) as a similarity metric. Once the clusters are defined anomalies can be detected using the LCS as the distance measure. In this context anomalies are determined to have low similarities between the clusters of other sequences and are defined to be far away from a cluster. Once anomalies are identified, SequenceMiner applies a genetic algorithm to modify the sequence to draw it closer to the cluster. Keeping track of the changes made to the sequence the algorithm reports back the missing and extra symbols giving the user some context of the anomaly. Since SequenceMiner focuses on the sequential nature of the anomalies it can find anomalies that other algorithms such as Orca and IMS are unable to detect, however it is ineffective at handling continuous parameters without somehow drastically changing the nature of the data.

5) Multiple Kernel Anomaly Detection: The main concept of kernel-based theory can be illustrated using the conceptual diagram as shown in fig.**??**. In fig.**??**, the left side shows that the data is non-linearly separa-

Algorithm 2 IMS: Train

1: Input: X_t (nominal system data vectors)
i (initial tolerance percent)
e (expansion percent)
m (max distance from cluster centroid to input vector)
2: **for** each input vector $\mathbf{X}_t$ normalize the values of $\mathbf{X}_t$ and find the cluster with the closest centroid to $\mathbf{X}_t$.
3: **if** no clusters exist, create a new cluster centered on $\mathbf{X}_t$, adding initial tolerance percent i to each vector value to create upper and lower bounds.
4: **else if** a closest cluster is found and $\mathbf{X}_t$ is within distance m of the centroid of that cluster, expand the cluster parameter boundaries as necessary to include $\mathbf{X}_t$ adding the expansion percent e to each parameter bound that is changed.
5: **else if** a closest cluster is found and $\mathbf{X}_t$ is beyond distance m of the centroid of that cluster, create a new cluster centered on $\mathbf{X}_t$, as in step **3**.
6: **end if**
7: **end for**
8: **Output**: B_c (parameter boundaries for each cluster)

Algorithm 3 IMS: Test

1: **Input**: $\mathbf{x}_t$ (test input)
2: **for** each $\mathbf{x}_t$ normalize the values of $\mathbf{x}_t$ and find the closest nominal cluster in B_c to $\mathbf{x}_t$.
3: **if** all $\mathbf{x}_t$ parameter values fall within the bounds of a cluster, the distance from $\mathbf{x}_t$ to the cluster is zero.
4: **else if** no cluster contains $\mathbf{x}_t$ locate the cluster with the hyper-box boundary that is closest to $\mathbf{x}_t$. Calculate the distance between a vector of $\mathbf{x}_t$ and a cluster hyper-box by summing the squares of the differences between each $\mathbf{x}_t$ parameter and the nearest cluster boundary value for that dimension, then find the square root of that sum.
5: **end if**
6: **end for**
7: **Output**: D_t (distance of vector to nearest cluster)
d_{tp} (distance of each parameter to cluster bounds)

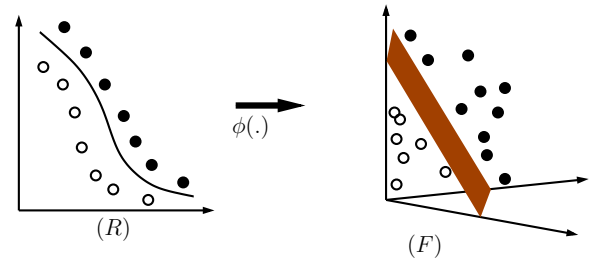

Fig. 2. In this figure we provide the illustration of higher dimensional mapping for linear separation fields. It shows that even if the patterns are nonlinearly separable in input space, it is possible to map them in higher dimensional feature space where they may be linearly separable. Here $\phi(.)$ is the mapping function.

ble (or sometimes non-seperable) in (2−D) data space whereas on the right hand side the same data is mapped into higher dimensional (3−D) feature space where the features are seperable. In the input space $\mathcal{R}$, suppose we are given the data $\mathcal{D} = \{(\vec{x}_i)\}_{i=1}^{n}$, where $\vec{x}_i \in \mathcal{R}^d$. We define a feature space $\mathcal{F}$, assuming that there exists a function ϕ that can be used to map any variable x from the input space to the feature space i.e. $\phi : \mathcal{R}^d \rightarrow \mathcal{F}$ and $\vec{x} \rightarrow \phi(\vec{x})$. Here $\vec{x}$ is an input vector ($\vec{x}_i \in \mathcal{R}^d$) which has been mapped into a much higher dimensional feature space. The mapping the data using $\phi(\vec{x})$ into $\mathcal{F}$ can be beneficial in defining a similarity measure using the dot product in $\mathcal{F}$ in terms of a function operating on the input data space and thus computing the inner products more efficiently which is commonly referred as the kernel trick in machine learning literatures. This results in a kernel matrix K which gives a relative similarity between objects ranging between maximum similar and maximum dissimilar. Unsupervised anomaly detection technique like one-class SVM builds on this kernel.

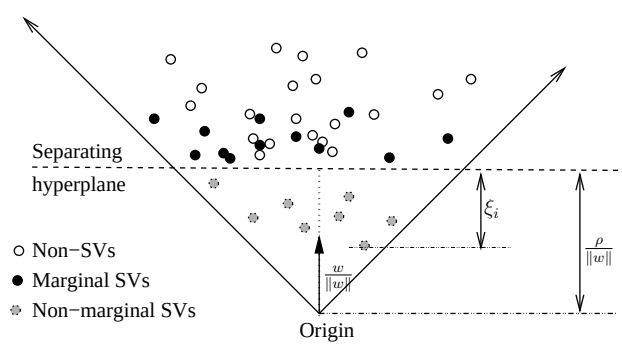

Fig. 3. This figure illustrates the geometric interpretation of optimal hyperplane for one class SVMs.

One-class SVMs builds a model on single (known)

class data and then finds a set of outliers using a decision boundary.The idea is to draw a separating hyperplane that can separate these outliers from the rest of the training examples. Schölkopf [18] showed that in the high dimensional feature space it is possible to construct an optimal hyperplane by maximizing the margin between the origin and the hyperplane in the feature space by solving the following dual problem,

$$\begin{aligned} \text{minimize} \quad & Q = \frac{1}{2}\sum_{i,j}\alpha_i\alpha_j(\sum_{\lambda}\beta_\lambda K^{\lambda}_{i,j}) \\ \text{subject to} \quad & 0 \leq \alpha_i \leq \frac{1}{\ell\nu}, \sum_i \alpha_i = 1 \\ & \sum \lambda = 1, \rho \geq 0, \quad \nu \in [0,1] \end{aligned} \tag{2}$$

where ν is an user specified parameter that defines the upper bound on the training error, and also the lower bound on the fraction of training points that are support vectors, β_λ is the weights to kernels and α_i is Lagrange multiplier. Once this problem is solved at least $\nu\ell$ training points with non-zero Lagrangian multipliers $(\vec{\alpha})$ are obtained and these points $\{x_i : i \in [\ell], \alpha_i > 0\}$ are called support vectors. The selected points can be marginal $\mathcal{I}_m = \{i : 0 < \alpha_i < 1\}$ and non-marginal $\mathcal{I}_{nm} = \{i : \alpha_i = 1\}$ support vectors. Once $\vec{\alpha}$ is obtained, SVMs compute the following decision function.

$$f(\vec{x}_z, \alpha, \beta, \rho) = sign(\sum_{i\in\mathcal{I}}\alpha_i(\sum_{\lambda}\beta_\lambda K^{\lambda}_{i,z}) - \rho) \tag{3}$$

where $\mathcal{I} = \mathcal{I}_m + \mathcal{I}_{nm}$. If the decision function predicts a negative label for a given test point z, then it is classified as an outlier. Test examples with positive labels are classified normal. The pseudo code of one class SVMs is shown below.

III. Fleetwide Analysis of Heterogeneous Data

Besides Orca, rest of the algorithms described above have been explored either in context to continuous real-valued data attributes or on discrete domain. Having said that, with increasing number of data sources, there is a need of developing intelligent knowledge refinement and integration techniques, focusing on the descriptions of underlying heterogeneous data sources. In aviation safety data heterogeneity may result due to the presence of multiple attributes as compared to single attribute at each data observation. In multivariate data, the attributes may or may not belong to the same data type. For example the attributes can be either continuous or discrete or it may be a mixture of both. The other source heterogeneity can be the behavioral or functional properties of these attributes. Not only is aviation data extremely large in size, it also has many aspects that create natural sources of heterogeneity. Some examples include flights that have common: origin or destination airports, city pair routes, tail numbers, aircraft models, as well as seasonal aspects such as flights within a month. Even within a flight there exists several phases like take offs, landings and cruise. It is important to note that in order to enable knowledge discovery, algorithms in general require an integrated and merged view of the data available across various resources.

Once an "appropriate" kernel function is chosen and the kernel matrix K is formed, it can be incorporated into any kernel based classification, regression, or anomaly detection methods where the kernel matrix is the sufficient representation of the input data. In another word, by "kernelizing" any method we encode knowledge about the data, expressed in terms of pairwise similarities. This provides us with the opportunity to incorporate vast amount of knowledge from heterogeneous sources using particular kernel functions. This field of research is know as Multiple Kernel Learning (MKL) [**?**], [4], [13]. MKL takes advantage of the mathematics of kernels allowing us to derive new kernels from kernels, provided each kernel satifies Mercer condition which states that the kernel function must be continuous, symmetric, and positive definite. There are several class of kernels those coincide with Mercer kernel. There are several existing literatures describing many different types of data with various types of features such as graphs and multiple feature types in computer vision such as color, shape, texture, and graphs based on image segmentations. Interested readers can explore literature [8], [10], [12], [19], [22] that look into various other classes of kernels like rbf, polynomical, bag-of-word, sigmoid, spline, graph based, tree based, mismatch based functions etc. A common practice is to use a convex combination (i.e. $\sum_\lambda \beta_\lambda = 1$) of various kernels which may be constructed on very different feature sets.

When analyzing FOQA data the concept of a system level analysis is paramount. The flight data consists of many parameters that monitor the various subsystems within a plane. Given a single flight this is not a simple task, and when considering additional flights and multiple aircraft the task can quickly grow beyond the classical timeseries analysis problem. With this increased complexity it is important to understand the hierarchical system structure and design algorithms to address this paradigm. Another challenging aspect is the sheer size

of the data that must be consumed by the algorithm. Typically flights are recorded at 1 Hz and may last anywhere from a couple hours for some regional flights and up to 10 to 19 hours for some international flights. This is compounded by the number of parameters that are recorded, which is typically hundreds of measurements. Some airlines are flying thousands of flights a day and millions of flights a year, which soon can add up to terabytes of data very quickly. To address this in the system level approach the algorithm must treat the flights as a fleet of aircraft with some sort of intelligent way of compressing the important features of each flight, and identifying the anomalies at the flight level.

IV. FOQA Data Analysis

The real world data set chosen for analysis is from a U.S. regional carrier. All aircraft analyzed were of the same fleet and type (narrow body jet), over a one year period resulting in over 176,000 flights. Each flight consists of 160 parameters sampled at 1 Hz with the average flight length between 1.5 to 2 hours. Due to privacy reasons, each pilot's identity and the exact date of the flight is kept confidential by the airline industry.

A. Data Preparation

Data analysis was focused on the approach portion of the flight below 10,000 ft. Mean Sea Level (MSL) to landing, using the deployment of the thrust reversers as a means to determine touchdown. Flights that were not found to reach 10,000 ft. or did not have thrust reversers deloped were removed from the data sets. For parameter selection a domain expert provided a list of 26 relevant continuous parameters that were extracted for analysis. Using information from the domain expert in conjuction with the statistics from the data, the flap parameter, which is categorical in nature, was decomposed into 3 binary state variables and then combined with landing gear and ground spoilers for sequence analysis.

The working data set consists of approximately 174,000 flights with varying lengths with each of these flights multidimensional heterogeneous time series. A random set of 2048 flights were chosen for training and the remainig used for testing. For continuous data, the mean and standard deviation are calculated for each parameter across all training flights. These statistics are then used in both training and testing to z-score normalize each parameter and flight to maintain consistency. Once the continuous parameters are normalized they are converted into the SAX representation (details can be found in [9][1]). The discrete parameters are handled by marking the on and off transitions between switch states with unique symbols and concatinating the symbols, while preserving the time ordering, into a sequence vector.

[1]The source code of SAX can be obtained from the authors' website at http://www.cs.ucr.edu/ eamonn/SAX.htm.

B. Experimental Details

In Multiple Kernel Anomaly Detection (MKAD) algorithm, since we want to model switching sequences for a given process and where the order of the switching is important, normalized Longest Common Subsequence (nLCS) based kernel was chossen as a potential candidate. Given two sequences $\vec{x}_i$ and $\vec{x}_j$, if z denotes a subsequence of them it means that removing some symbols from $\vec{x}_i$ produces $\vec{x}_j$ or vice versa. The longest such subsequence of $\vec{x}_i$ and $\vec{x}_j$ is called the longest common subsequence (LCS) and is denoted by $LCS(\vec{x}_i, \vec{x}_j)$ and $|LCS(\vec{x}_i, \vec{x}_j)|$ is its length. Such a kernel over discrete sequences, when normalized, takes the form of,

$$k(\vec{x}_i, \vec{x}_j) = nLCS(\vec{x}_i, \vec{x}_j) = \frac{|LCS(\vec{x}_i, \vec{x}_j)|}{\sqrt{l_{\vec{x}_i} l_{\vec{x}_j}}}, \quad (4)$$

where $l_{\vec{x}}$ is the number of symbols in sequence $\vec{x}$. Each sequence of switches is compared against other sequences by using the longest common subsequence (LCS) as the metric for comparison. Sequences that are similar are bound to hold high nLCS values, while dissimilar sequences will hold very low nLCS values.For the MKAD algorithm, once the sequences are generated the discrete kernel is computed pairwise across all possible flight combinations in the training set. For the continuous data, each time series was SAX transformed. In the original version of SAX, the z-score normalization is an integral part of the algorithm. However in this research, we normalized each time series (only once) before it is SAX transformed. We are able to maintain consistency in choosing the alphabet size for both reference and test sets. The window size was also kept fixed throughout the analysis. The window size and alphabet size were both set to 10. Once the SAX representations are obtained, another kernel is computed pairwise across for all possible flight combinations. Each element of this kernel is the average of the pairwise comparison across the parameters of any two flights. In the optimization, we have set the ν parameter of one-class SVMs to 0.1. For testing, the support vectors are used to calculate the pairwise similarity between all testing flights. The discrete and continuous kernels for test data were generated in a similar fashion as the training.

C. Flight Analysis

Out of 174,000 flights the algorithm identified over 4,700 flights as anomalous. This paper will present two flights identified by a domain expert to be operationally significant.

The first flight can be categorized as a high energy approach. There are basically two conditions which result in a high energy approach - the aircraft may be too high or too fast or both. Each of these conditions can be converted into the other in a process called ”trading altitude for airspeed” or vice versa. But ultimately, drag devices such as flight spoilers, gear or flaps must be deployed to permit a return to the ideal flight profile. However, complicating this effort, most of the drag devices have a maximum speed for deployment, limiting corrective options in the ”too fast” scenario, and airlines and aircraft have rate of descent limits which affect the ”too high” situation. Finally, Air Traffic Control (ATC) is responsible for traffic separation, which may limit the maneuvering needed to dissipate excess energy. The flight shown in figures IV-C & IV-C managed to encounter every one of these issues. Techniques were employed, rather competently, which enabled the aircraft to ”go down and slow down”, a supposed dichotomy. It was both too high and too fast to make a stabilized approach, so drag devices; in this case flaps and gear were deployed. The landing gear was extended just below the maximum allowable airspeed for that operation. This helped some, but it was still too fast to deploy flaps beyond an initial setting, so the pilot did something entirely counterintuitive when too high - he climbed for a few moments, bleeding off sufficient airspeed to permit him to extend the flaps further, which in turn enabled him to descend more rapidly. But even all this was insufficient to permit a return to an ideal flight profile, so the pilot presumably obtained an ATC clearance to make a 360 degree turn which enabled him to lose sufficient altitude to return to the approach in a much better position for a stabilized approach. Due to considerations of conflicting traffic this is always something which must be coordinated with ATC, and sometimes a simple 360 is not an option, resulting in the performance of a go around maneuver.

The second flight falls under the category of a turbulent approach. The amount of lift supplied by an aircraft wing is a function of the speed of the air flowing across it. In turbulent conditions, it is possible to experience enough loss of lift that one or both wings stall, possibly causing roll or pitch changes, which in extreme cases could be dangerous. All of this has been well known since the earliest days of flight, but recent

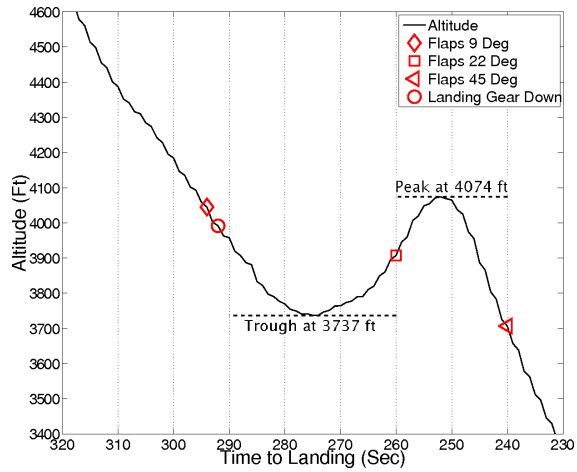

Fig. 4. Altitude plot showing when the drag devices were deployed.

crashes in the South Atlantic and in upper New York State remind us that despite advanced aircraft design and current operational procedures, the issue of turbulence, especially clear air turbulence, remains a threat. Airline flight data monitoring programs watch for excessive aircraft attitudes and speeds, but generally downplay environmental factors like wind speed and direction, as well as throttle position and engine speed (N1). This flight descended from about 5,000 ft to landing on a fairly straight path, but the pilot had to work hard to accomplish this (see figure 6(a)).

MKAD discovered atypical fluctuations in engine speed, which is represented by an RPM parameter called N1. Engine speed is normally fairly consistent during a stabilized approach, increasing a bit each time flaps are extended, to make up for the increased drag. On this flight, N1 varied between 35% and 80% (maxium being 100%) in more than 14 cycles see figure 7(a). This is a considerable amount of engine speed variation.

After some analysis, the reason for all this variation was that the pilot had to cope with gusty wind conditions which caused the airspeed to vary between 140 - 180 knots. Part of the reason for this variation, of course, was the extension of drag devices (gear and flaps), but an examination of the airspeed trace seen in figure 6(b) shows about 9 cycles of acceleration and deceleration. The headwind over this same time segment varied between 23 - 28 knots in about 8 cycles (see figure 7(b)), not dangerous, but certainly not ideal.

It is important to note that the anomalies identified by MKAD were determined to have contributions from both the discrete and continuous parameters. In the high energy approach the landing gear was statistically out of order from the normal approaches in the data set and

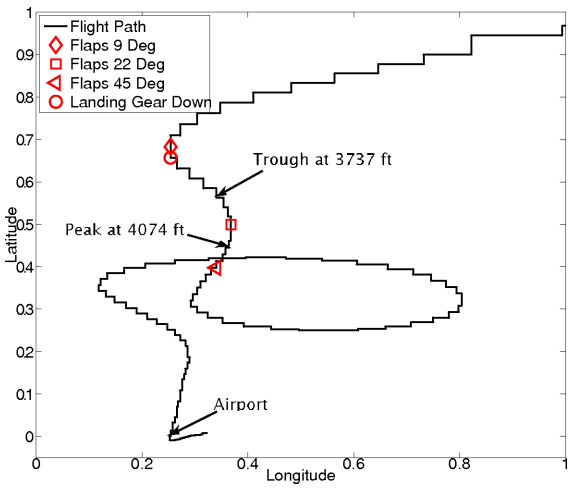

Fig. 5. Flight path plot showing when flaps and gear were deployed. The arrows indicate the peak and trough occured.

when combined with the unusual altitude profile signaled an anomaly. Additionally the preprocessing steps used to discretize the continuous parameters preserved the sequential nature of the data allowing the kernel function to highlight an unusual parameter profile such as the throttle changes in the turbulent approach flight. This is different from other algorithms that treat each time sample independantly and cannot detect events in this way.

V. Conclusion

Parameter anomalies discovered by MKAD are frequently like little windows into a larger reality. It will frequently require analysis by a domain expert to ferret out whether an anomaly stemmed from a hazardous issue, or whether the condition, while atypical, was safe and explainable. One thing is certain - while a large amount of experience and expertise underlie the current aviation data analysis programs, they only answer questions that someone thought to ask. There will always be value in the search for the unexpected, which is the purpose of MKAD.

An interesting observation is that detection techniques which involve a metric or distance related function, in mathematical sense, to define the pairwise distances or similarities between elements of the set somewhere in the algorithm stage, are more versatile to address different data structures. Choosing the right *metric function* and integrating knowledge from multiple sources must be done judiciously.

Acknowledgment

The authors would like to thank...

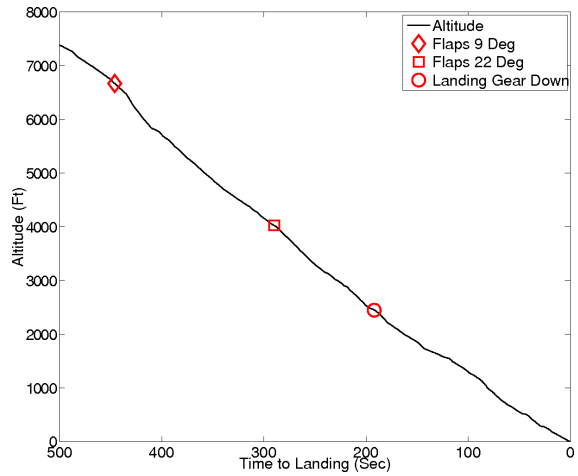

(a) Altitude

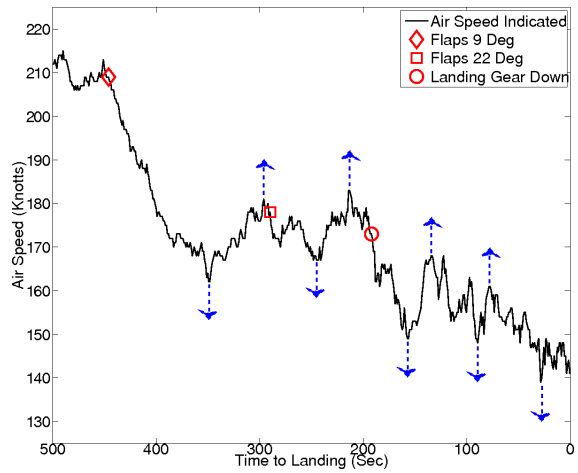

(b) Air Speed

Fig. 6. Figure 6(a) shows straight altitude approach profile. The arrows in figure 6(b) are shown to help visualize the frequent increase and decrease in air speed.

References

[1] Amidan, B.G., Swickard, A.R., Allen, R.E. and Ferryman T. A. Identifying In-Close-Approach-Changes in Air Traffic Control (ATC) Data. 2002. PNWD-3210. Battelle Pacific Northwest National Laboratory: Richland, WA.

[2] F. Angiulli and F. Fassetti. Very Efficient Mining of Distance-based Outliers. In *Proceedings of CIKM '07*, pages 791–800, 2007.

[3] F. Angiulli and C. Pizzuti. Outlier Mining in Large High-Dimensional Data Sets. *IEEE Transactions on Knowledge and Data Engineering*, 17(2):203–215, 2005.

[4] F.R. Bach, G.R.G. Lanckriet, and M.I. Jordan. Multiple kernel learning, conic duality, and the smo algorithm. In *International Conference on Machine Learning*, 2004.

[5] S. D. Bay and M. Schwabacher. Mining Distance-based Outliers in Near Linear Time with Randomization and a Simple Pruning Rule. In *Proceedings of KDD'03*, pages 29–38, 2003.

[6] Suratna Budalakoti, Ashok N. Srivastava, and Matthew E. Otey. Anomaly detection and diagnosis algorithms for discrete symbol

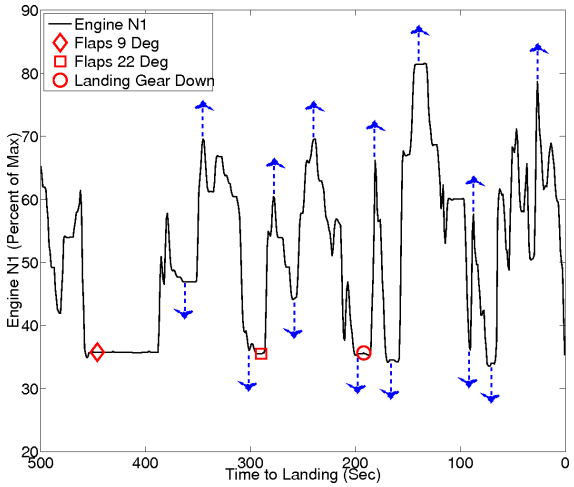

(a) Engines N1 Speed

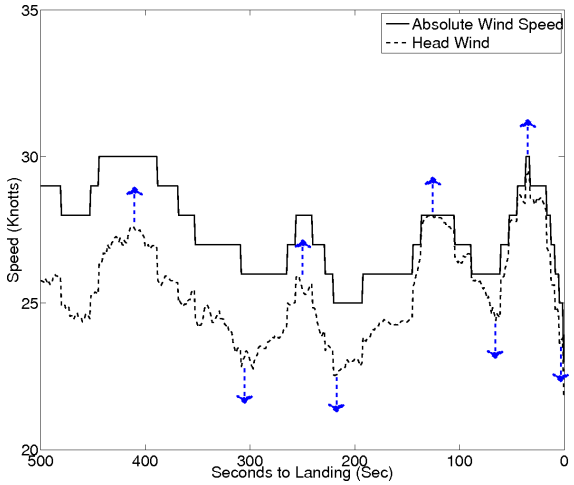

(b) Wind Speed

Fig. 7. The arrows in figure 7(a) & 7(b) are shown to help visualize the frequent increase and decrease in Engines and Wind speed.

sequences with applications to airline safety. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 39(1):101–113, 2008.

[7] Banerjee A. Chandola V. and Kumar V. Anomaly detection: A survey. *ACM Comput. Surv.*, 41:1–58, 2009.

[8] O. Chapelle and P. Haffner. Support vector machines for histogram-based classification. *IEEE Transactions on Neural Networks*, 10(5):1055–1064, 1999.

[9] Santanu Das, Bryan L. Matthews, Ashok N. Srivastava, and Nikunj C. Oza. Multiple kernel learning for heterogeneous anomaly detection: algorithm and aviation safety case study. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '10, pages 47–56, New York, NY, USA, 2010. ACM.

[10] Z. Harchaoui and F.R. Bach. Image classification with segmentation graph kernels. In *Computer Vision and Pattern Recognition*, pages 1–8, 2007.

[11] David L. Iverson. Inductive system health monitoring. *Proceedings of the 2004 International Conference on Artificial Intelligence (IC-AI'04), CSREA Press, Las Vegas, NV*, 2004.

[12] Hisashi Kashima, Koji Tsuda, and Akihiro Inokuchi. Kernels for graphs. *Kernel Methods in Computational Biology*, 39(1):101–113, 2004.

[13] G.R.G. Lanckriet, N. Cristianini, L. El Ghaoui, P. Bartlett, and M.I. Jordan. Learning the kernel matrix with semidefinite programming. *Journal of Machine Learning Research*, 5:27–72, 2004.

[14] Markos Markou and Sameer Singh. Novelty detection: A review - part 1: Statistical approaches. *Signal Processing*, 83:2003, 2003.

[15] Markos Markou and Sameer Singh. Novelty detection: A review - part 2: Neural network based approaches. *Signal Processing*, 83:2499–2521, 2003.

[16] John A. Quinn and Christopher K. I. Williams. Known unknowns: Novelty detection in condition monitoring. In *Proceedings of 3 rd Iberian Conference on Pattern Recognition and Analysis, Lecture Notes in Computer Science 4477. Springer-Verlag*, pages 1–6. Springer LNCS, 2007.

[17] S. Ramaswamy, R. Rastogi, and K. Shim. Efficient Algorithms for Mining Outliers from Large Data Sets. *SIGMOD Rec.*, 29(2):427–438, 2000.

[18] B. Schölkopf, J. C. Platt, J. C. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the Support of a High-Dimensional Distribution. *Neural Comput.*, 13(7):1443–1471, 2001.

[19] B. Schölkopf and A. Smola. *Learning with Kernels*. MIT Press, 2002.

[20] Thomas R. Chidester. Understanding Normal and Atypical Operations through Analysis of Flight Data. *In Proceedings of the 12th International Symposium on Aviation Psychology*, 2003. Dayton, Ohio.

[21] Wikipedia. Hamming distance, the free encyclopedia, 2010. [Online; accessed 04-December-2010].

[22] H. Zhang, A.C. Berg, M. Maire, and J. Malik. Svm-knn: Discriminative nearest neighbor classification for visual category recognition. In *Computer Vision and Pattern Recognition*, pages 2126–2136, 2006.