COMPARATIVE ANALYSIS OF DATA-DRIVEN ANOMALY DETECTION METHODS ON SOLID ROCKET MOTOR FAULTS

BRYAN MATTHEWS* AND ASHOK N. SRIVASTAVA, PH.D.

Abstract. This paper provides a review of three different advanced machine learning algorithms for anomaly detection in continuous data streams from a ground-test firing of a subscale Solid Rocket Motor (SRM). This study compares Orca, one-class support vector machines, and the Inductive Monitoring System (IMS) for anomaly detection on the data streams. We measure the performance of the algorithm with respect to the detection horizon for situations where fault information is available. These algorithms have been also studied by the present authors (and other co-authors) as applied to liquid propulsion systems. The trade space will be explored between these algorithms for both types of propulsion systems.

1. Introduction

Over the years as engineering systems have grown more complex there has been an increased interest in the analysis of integrated vehicle health monitoring (IVHM) systems. Applications of IVHM technologies range from commercial aircraft maintenance applications to the real time monitoring of the International Space Station subsystems. These systems will require the ability to analyze large, complex, multivariate time series in near-real time [10] [11] [3]. This requirement has spawned development of numerous algorithms for detecting anomalies in multivariate time-series with various applications.

For simplicity, we assume that we are observing the output of a dynamical system with d dimensions. This can be modeled as:

$$\mathbf{x}_t = \Psi(\mathbf{x}^*_{t-1}, \theta, u_t) \tag{1}$$

$$y_t = \Omega(\mathbf{x}_t) \tag{2}$$

where y_t is a d dimensional observed vector, $\mathbf{x}_t$ is the state variable, θ is a set of parameters, and u_t is a set of exogenous variables. The movement of the output vector y_t through the d dimensional space is under consideration. One can imagine that this vector stays within certain regions of the state space under normal operating conditions. For anomaly detection, it is useful to consider situations where this vector deviates with respect to a subset of components. In this situation, the deviation may actually not appear in any component by itself but can only be seen by examining multiple components. Consider a situation where we measure the temperature and pressure of a gas in an enclosed fixed volume. Normally, the temperature and pressure would be linearly related to each other. However, if we begin to notice that the temperature

*NASA Ames Research Center, Moffett Field, CA, 94035, bryan.l.matthews@nasa.gov.

increases in the volume but the pressure does not, that can indicate an anomalous behavior and perhaps a breach in the volume. There may be other reasons for such an anomaly to occur. Thus, the covariation of pressure and temperature is necessary to properly determine whether an anomalous condition occurs. In situations where the dimensionality is much larger than two, this problem becomes much more difficult to address. In systems where u_t is a driving function, we must keep track of the fact that this exogenous variable may actually be the source of the observed aberration, and thus the aberration is not a true anomaly but an expected event. Thus, for driven systems, a successful anomaly detection method must take into account the current values of the exogenous driving function. In the best case, we would have a physics-based model that explicitly states the functions Ψ and Ω. In this situation, we could use the models given above in order to estimate the correct value of the state variable and the observed variable and compare those estimates with the observed value. This generally leads to good results because all known dynamics are taken into account. This paper considers the situation where we do not know or cannot model the dynamics of the system under study. We resort to so-called data-driven approaches. These approaches do not require knowledge of the underlying physics of the system but instead build internal models of nominal system behavior using statistical and data mining approaches.

The development of the Ares I rocket is a present application and area of interest for vehicle health monitoring at NASA that has allowed various methods to employ a diverse set of techniques to monitor the liquid propulsion engines, the solid rocket booster engines as well as many other vital subsystems on the spacecraft. These techniques range from limits and rules on parameters derived from simulations [12] to integrated physics-based approaches [13] to advanced data driven methods [8]. With the recent advancements in the data mining field many data driven techniques have been proven to perform very well compared to some of the classical methods.

This paper addresses three data-driven methods Orca, one-class support vector machines (OCSVM), and Inductive Monitoring System (IMS), and their application on solid rocket motor data. Previous work was done using these methods in the liquid propulsion domain [7] where anomalies were successfully detected. The aim of this paper is to assess the performance of these three methods on solid rocket motor data and compare the performance of the algorithms. We also compare and contrast the performance of the algorithms with the results from the liquid propulsion study.

2. Data

The data examined in this paper consists of three ground firing tests of a 10-inch solid rocket motor TD-31. The motor was outfitted with pressure and temperature sensors along with strain gauges on both the forward and aft closures. Sampling rates range from 250Hz on the strain and temperature sensors to 2000 Hz on the pressure sensors. All three sensor arrays were recorded asynchronously.

The solid rocket casing was equipped with a total of 25 sensors covering pressure (in PSI), temperature (in degrees Fahrenheit), and strain (in Microstrain). The sensor arrays were comprised of 3 pressure sensors: 1 on the forward closure and 2 on the aft closure; 14 temperature sensors: 8 on the forward polar boss, 3 on the forward closure, and 3 on the case back from the forward polar boss; and 8 strain gauges: 5 on the forward polar boss, 2 on the boss-hoop, and 1 on the forward closure edge.

As mentioned previously the data was recorded from three ground firings. The first was a nominal solid rocket motor burn. The 2nd and 3rd ground firings contained seeded faults with 30 degree and 40 degree saw tooth cut out openings respectively on the forward end closure. This type of defect weakens and eventually ejects the O-ring shortly after startup, which produces a gas leakage around the forward end closure during rocket burn resulting in a large hole forming in the aft section of the rocket. This is considered a catastrophic failure, which can lead to a mission abort or even loss of vehicle. Critical failure time is within 8 to 10 seconds, however initial indication of the failure is present within the first second after ignition. The algorithms implemented in this paper will attempt to detect the initial indications of the O-ring fault within the first second.

3. Data Driven Methods

There are primarily two classes of models used in anomaly detection for engineering systems: physics based models and data driven models. Physics based approaches can provide very accurate forecasts, but require a high-fidelity model and understanding of the engineered system to perform effectively [5]. In data-driven methods the algorithms have no prior knowledge of the engineered system. The algorithms learn the operating boundaries of the system from an example of nominal recorded data and use this information to flag anomalies in an unseen data. The three data-driven methods discussed in this paper are considered unsupervised learning techniques, which means that only nominal data is used for reference [1].

ORCA

Orca [4] is an outlier detection algorithm which uses the Euclidean distance nearest neighbor based approach to determine outliers. For computational efficiency it employs a modified pruning technique which allows it to perform in near linear time. For each point in the test data set, where a point is a row in the data set consisting of measurements taken at a single point in time, Orca calculates the nearest neighbor points from the reference data set. In this study the reference set is a portion of the nominal ground firing and the test sets are the remainder of the nominal sets and the two fault firings. The output from Orca is a distance score which represents the average distance to its k nearest neighbors. The more anomalous the point is the higher the score, since the nearest neighbors are farther away.

ONE-CLASS SUPPORT VECTOR MACHINES

One-class support vector machine (OCSVM) [9, 1] is a modified support vector classification technique that is designed to find outliers by building a model based on nominal training data. Classical binary support vector machines map the data to a high dimensional space and determine the separating hyper plane that maximizes the distance between the two classes. Support vector machines can also take advantage of non-linear kernels to map the data to this higher or even infinite dimensional space where the data can be linearly separated. In the case of OCSVM a small percentage of the training data (called slacks) is separated from the rest of the data points and allowed to fall on the outlier side along with the origin. OCSVM attempts to optimize the hyperplane between the origin and the remaining nominal data. The output from OCSVM is a distance magnitude score, which represents the distance from the hyperplane in the outlier, or negative, direction. If the data falls on the nominal, or positive side of the hyperplane, the distance is ignored and considered zero.

INDUCTIVE MONITORING SYSTEM

Inductive monitoring system (IMS) [6] is an algorithm that determines outliers by calculating the distance from a learned set of nominal cluster bounds in a multi-dimensional space. During the training phase of IMS the algorithm builds a cluster base from a set of nominal data. In the monitoring phase IMS calculates the distances to the sides of the clusters hypercubes. If a point falls within the cluster bounds the IMS score is zero and is considered nominal. If the point falls outside of the cluster it will have a positive score. Higher scores are indicative of more anomalous behavior.

4. Analysis

Due to the limited number of examples (1 nominal and 2 fault scenarios) it was decided to break the nominal data set into two randomly sampled subsets. Each algorithm was given one of the nominal subsets as a reference and evaluated on the remaining nominal subset as well as the two fault runs. The validation test on the remaining nominal subset was done to verify each algorithm's scores were within nominal ranges. The algorithms reported scores on the validation set that were within the range or lower than the scores found in the nominal periods of Tests 2 and 3. Each sensor was zscore normalized based on the same

[1] More detailed information on these algorithms and other applications can be found at the Dashlink website (http://dashlink.arc.nasa.gov/).

sensor's statistics calculated from the training data set before the algorithms were applied to the data. Due to the asynchronous sampling rates for each sensor array the algorithms were run separately on the parameters (pressure, temperature, and strain). This also allowed for comparison of the detection ability of the parameter sets as well as algorithms. Ground truth was considered to be when the O-ring was ejected from the rocket. This was approximately 0.25 seconds after ignition for both Test Fire 2 and 3. The O-ring ejection has been reported by the domain experts at a continuous process, therefore symptom identification and anomaly flagging rather than exact detection is emphasized in this analysis process.

The algorithms applied only provide a distance measure or score; therefore an additional detector must be used for flagging faults. This was achieved by employing a Bollinger band [2] detector. Bollinger bands have been used for tracking market trends in the financial sector for many years. The concept involves a running statistical threshold based on a sliding window. Once the score crosses the threshold more than N consecutive times a flag is raised and the running threshold update is paused until the score falls beneath that threshold again. All scores after a threshold crossing are considered anomalies until the scores fall beneath the threshold again. The dynamic threshold is determined by $\mu + 2\sigma$, where μ represents the mean and σ represents the standard deviation of the window. For this study the window length used was 100 milliseconds and N was 10 samples. Since the scores will always be positive we are only interested in the upper limit threshold.

PRESSURE

Figure 1 (a) and (b) show the Orca, IMS, and OCSVM scores for the pressure sensors from ignition time to 1 second for Tests 2 and 3 respectively. Note that OCSVM appears to detect some anomalous behavior at the beginning. This time frame is ignored by the detector due to startup transients. In all three algorithms the scores for the pressure sensors increase sharply after 0.8 seconds and result in a fault flag from the Bollinger band detector. This sharp increase is due to a rapid drop in pressure on the forward closure sensor following the O-ring ejection. Figure 1 (c) shows the forward closure sensor track the aft sensors up until the point where it begins a sharp decline at the time of the detection. All algorithms do a good job in highlighting this anomaly where the forward pressure sensor deviates from the others, however there does not seem to be a definite indication at the time of the O-ring ejection in any of the algorithm's scores around 0.25 seconds while using the pressure sensors array.

TEMPERATURE

In Figure 2 (a) and (b) the Orca, IMS, and OCSVM scores for the temperature sensors during the same duration can be seen. The scores yield a small upsurge around the 0.25 second range for both the Orca and IMS scores. The upward trend was sharp enough and long enough in duration for the Bollinger band detector to pick up the fault. Notice IMS has a slightly delayed response for both Test 2 and Test 3 when compared to Orca. For OCSVM the startup transient scores are ignored by the detector as in the pressure sensor analysis. The fault also seems to be picked up earlier than both Orca and IMS. All algorithms were run using all 14 temperature sensors. The top three scoring sensors (identified by both IMS and Orca), located on the forward closure near the O-ring, are displayed in Figure 2 (c). For Test 1 the three sensors can be seen to operate under nominal thermal conditions, however in Tests 2 and 3 the temperature sensors begin to exceed the normal operating conditions around the time of the O-ring ejection. It is important to note that the temperatures for Tests 2 and 3 swing back down within normal thermal conditions temporarily after 1 second, however all three rise sharply to extreme temperatures again within the next 5 seconds, while all algorithms continue to report anomalous scores throughout this period.

STRAIN

In Figure 3 (a) and (b) the scores for the strain sensors for IMS, Orca, and OCSVM are displayed and are successfully able to pick up the fault around the 0.20 second mark for Test 2 and also flag what appears to be a precursor in Test 3. Even though the beginning of the precursor for Test 3 is very close to the end of

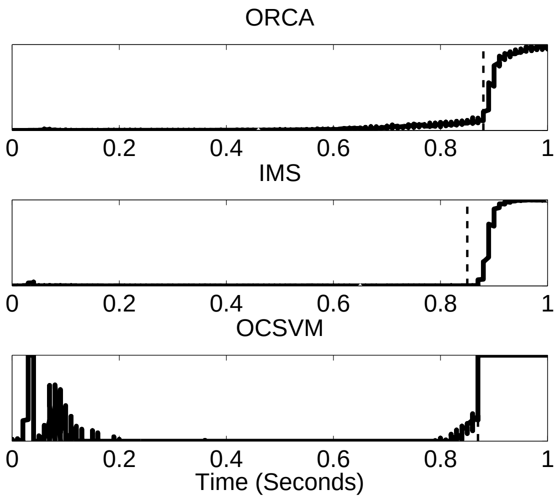

(a) Algorithm Scores for Test 2.

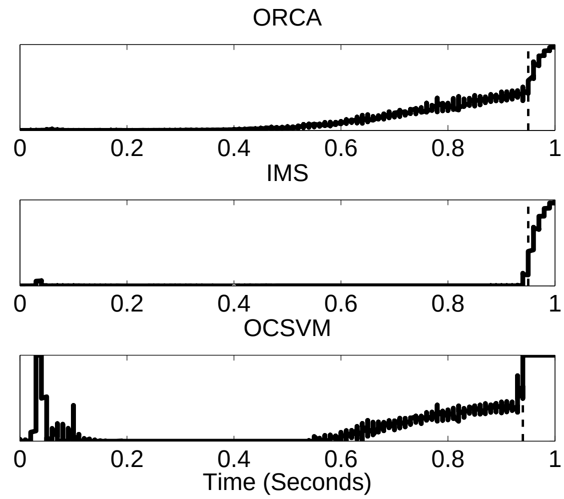

(b) Algorithm Scores for Test 3.

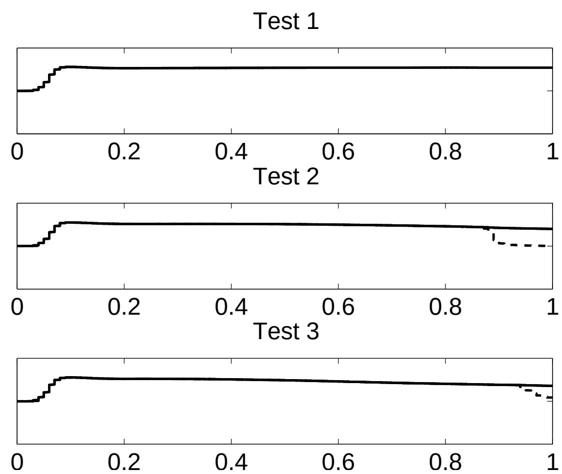

(c) Pressure Sensors for Test 1, 2, and 3.

FIGURE 1. The solid lines in (a) and (b) are the respective algorithms' scores. The dotted vertical lines represents the point at which the Bollinger band signaled a fault. In (c) the solid lines are the pressure sensors on the aft closure. The dotted line is the pressure sensor on the forward closure

the startup transient phase it can still be seen in the sensor data and therefore considered a valid precursor. In Figure 3 (c) the strain gauge sensors mounted near the forward closure are plotted. In the nominal run all 3 sensors track with each other and are within normal operating ranges. In Tests 2 and 3 the sensor mounted closest to the O-ring begins to deviate sharply from the other 2 sensors around the 0.20 second mark. In Test 3 the precursor can be seen in the same sensors when it fails to track with the other sensors slightly before 0.20 second.

5. Comparison with Tests on Liquid Propulsion Systems and Discussion

Martin et. al. [7] performed an extensive comparison between several anomaly detection algorithms for detecting faults on the Space Shuttle Main Engine. The algorithms used in that study include Orca, Gritbot, One Class SVM, IMS, Linear Dynamical Systems, and Gaussian Mixture Models (GMM). In this study, the authors concluded that OCSVM has the highest accuracy, while Orca and GMM are tied for

Table 1

Sensor Type	Detection Time (Seconds)					
	Orca		IMS		OCSVM	
	Test 2	Test 3	Test 2	Test 3	Test 2	Test 3
Pressure	0.88	0.95	0.85	0.95	0.87	0.94
Temperature	0.25	0.11	0.35	0.19	0.21	0.12
Strain	0.19	0.12	0.19	0.13	0.21	0.07

having the earliest detection times. In the present study (see Table 1), we observed that OCSVM has the earliest detection time with the strain gauge sensor array.

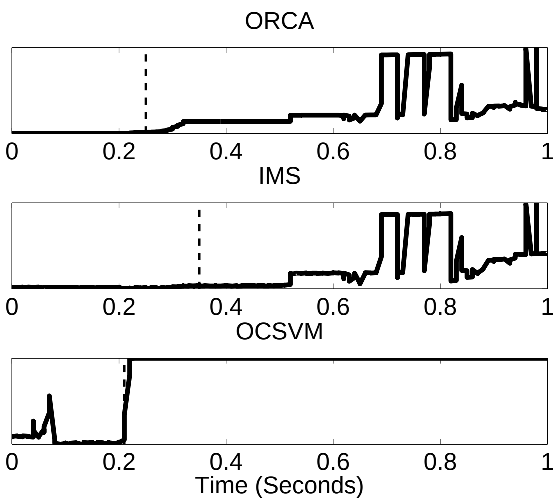

(a) Algorithm Scores for Test 2.

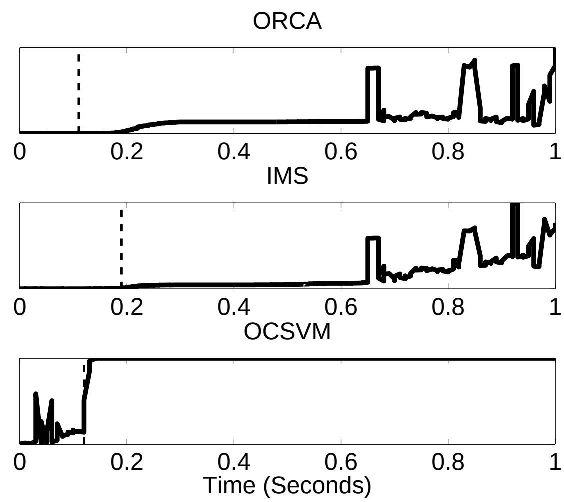

(b) Algorithm Scores for Test 3.

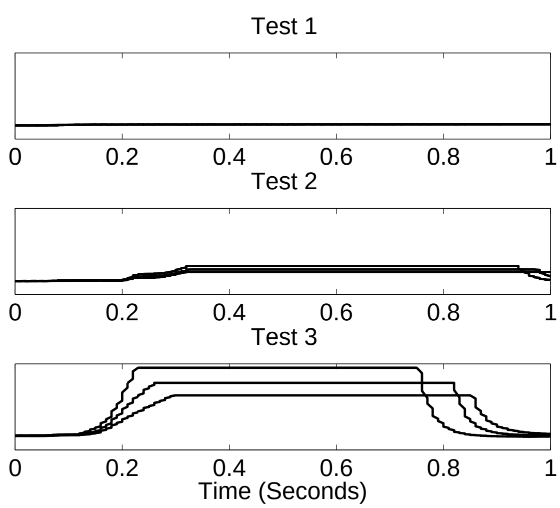

(c) Temperature Sensors for Test 1, 2, and 3.

Figure 2. The solid lines in (a) and (b) are the respective algorithms' scores. The dotted vertical lines represents the point at which the bollinger band signaled a fault. In (c) the 3 temperature sensors on the forward closure are shown for Tests 1, 2, and 3.

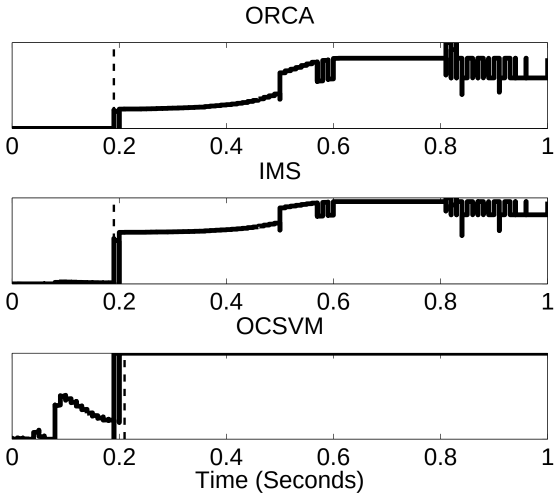

(a) Algorithm Scores for Test 2.

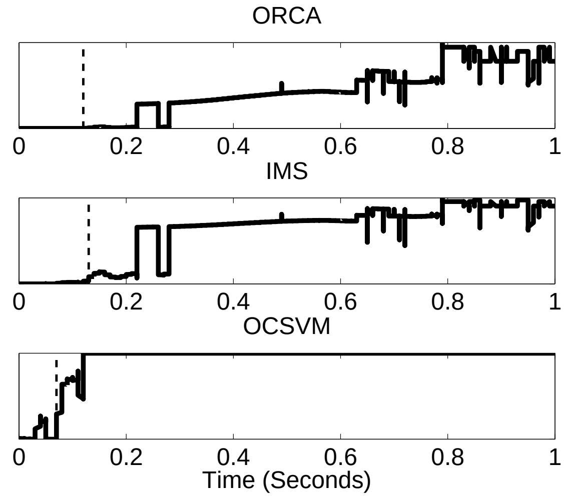

(b) Algorithm Scores for Test 3.

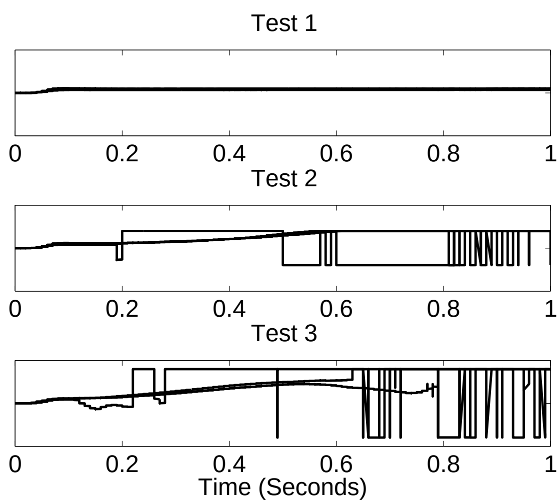

(c) Temperature Sensors for Test 1, 2, and 3.

FIGURE 3. The solid lines in (a) and (b) are the respective algorithms' scores. The dotted vertical lines represents the point at which the bollinger band signaled a fault. In (c) the 3 strain gauge sensors on the forward closure are shown for Tests 1, 2, and 3.

It is useful to note that the Orca and IMS algorithms are both distance based outlier detection methods. As such, their behavior may be well-suited for comparison across a broad spectrum of tests where the variables measured are continuous. However, in situations where multiple data streams arise from sensors where the streams represent both continuous and discrete measurements, the behavior of distance-based algorithms may be unexpected. This is due to the fact that for example Euclidean distance is not necessarily a good representation of anomalies that are occurring in discrete sequences [3]. In hybrid situations, i.e., situations where both discrete and continuous signals exist, the anomaly detection algorithm must take into account that anomalies in discrete sequences are usually of the form of misplaced subsequences or incorrectly inserted or deleted sequences. In this situation, it is critical that the anomaly detection method be flexible enough that a combination of multivariate continuous anomalies in state space are detected along with the type of discrete anomalies described here. This is the subject of future work in the field.

In other application domains, we have found that anomaly detection methods such as OCSVM, Orca, and IMS can detect anomalies even in situations where poor signal-to-noise ratios exist. We are in the process of testing and documenting these findings. However, it is important to note that such findings can be expected because we are using multiple signals for the anomaly detection task. Consider a situation where we have N signals, each of which have a signal-to-noise ratio associated with them. When we study a combination of these N signals, it may be possible to detect anomalies because of the added information from the movement of the signals together. This can be monitored by synchronizing multiple sensor arrays and allowing the algorithms to track the multivariate interaction over the new higher dimensional space. We will study this effect further and report on it both from a theoretical and empirical point of view.

6. Conclusions

This paper has addressed three data driven methods (Orca, IMS, and OCSVM) and discussed their performance on the three sensor arrays for two ground fire tests of a solid rocket motor. In conclusion all sensor arrays and algorithms do a good job at picking up the fault either by directly detecting the O-ring failure or indirectly picking up on the drop in pressure as the hole begins growing. Table 1 shows the detection times for both tests and each algorithm versus the corresponding sensor arrays. It appears that the IMS and Orca tie for the earliest detection times for Test 2 at 0.19 seconds using the strain gauge sensor arrays. In Test 3 the earliest detection time is OCSVM at 0.07 seconds with the strain sensors array. For these particular fault scenarios the strain sensor array appears to have the best insight for early detection when using all three algorithms. However all sensor arrays appear to provide comprehensive indication of a fault well before the critical failure point of 8 to 10 seconds.

7. Acknowledgements

This work was supported through funding from the NASA Aeronautics Research Mission Directorate, Aviation Safety Program, Integrated Vehicle Health Management project. The authors thank Vadim Smelyanskiy, Dmitry Luchinsky, and Slava Osipov for their insight and suggestions.

References

1. Das S. Srivastava A. and Chattopadhyay A., *Classification of damage signatures in composite plates using one-class svm's*, In Proceedings of the IEEE Aerospace Conference, Big Sky, MO. (2007).
2. John. A. Bollinger, *Bollinger on bollinger bands.*, McGraw-Hill 1st edition (2001).
3. Srivastava A. N. Budalakoti, S. and M. Otey, *Anomaly detection and diagnosis algorithms for discrete symbol sequences with applications to airline safety*, IEEE Transactions on Systems, Man, and Cybernetics, Part C (2008).
4. Bay S. D. and Schwabacher M., *Mining distance-based outliers in near linear time with randomization and a simple pruning rule.*, Proceedings of The Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. (2003).
5. Luchinsky D.G. Osipov V.V. Smelyanskiy V.N. Timucin D.A. and Uckun S., *Model based ivhm system for the solid rocket booster*, IEEE Aerospace Conference (2008).
6. 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).
7. R. A. Martin M. Schwabacher N. Oza and A. Srivastava, *Comparison of unsupervised anomaly detection methods for systems health management using space shuttle main engine data*, Proceedings of the Joint Army Navy NASA Air Force Conference (2007).
8. Mark A. Schwabacher, *A survey of data-driven prognostics*, American Institute of Aeronautics and Astronautics Inc. (2005).
9. B. Schlkopf J. Platt J. Shawe-Taylor A. J. Smola and R. C. Williamson, *Estimating the support of a high-dimensional distribution.*, Neural Computation (2001), 1443–1471.
10. A. N. Srivastava, *Discovering system health anomalies using data mining techniques*, Proceedings of the 2005 Joint Army Navy NASA Airforce Conference on Propulsion (2005).
11. A. N. Srivastava and S. Das, *Detection and prognostics for low-dimensional systems*, IEEE Transactions on Systems, Man, and Cybernetics, Part C (2008).
12. James R. Beaty Brett R. Starr and John W. Gowan Jr., *Ares-i-x vehicle preliminary range safety malfunction turn analysis*, American Institute of Aeronautics and Astronautics Inc. (2007).
13. Dimitry Gorinevsky Robert W. Mah Dogan Timucin, *Early detection of solid rocket motor case-breach failure*, American Institue of Aeronautics and Astronautics Inc. (2007).