

Anomaly Detection in Aircraft Gas Turbine Engines

Devendra Tolani, Murat Yasar, Asok Ray* and Vigor Yang
*The Pennsylvania State University,
University Park, PA 16802*

I. Introduction

PERFORMANCE, reliability, and operational flexibility are critical requirements of complex mechanical systems, such as aircraft propulsion, and these requirements may vary with different mission perspectives during the service life. In spite of meticulous engineering design, complex systems do eventually degrade and often fail to yield the anticipated performance during the later phases of their operational life. Recently, much attention has been devoted to condition-based maintenance using the online information of sensor-based observables. For example, life extending control^{1,2} and fault-tolerant control^{3,4} of gas turbine engines may require monitoring of behavioral patterns, based on real-time tracking of the current status of system health and degradation rate, for generating appropriate feedback information.

Anomaly in a dynamical system is defined as deviation from its nominal behavior and can be associated with parametric or non-parametric changes that may gradually evolve in the system. Anomalous operation of aircraft gas turbine engines is undesirable from the perspectives of both engine operation and aircraft mission management. Early detection of anomalies and their characterization are essential for health management, which includes prognosis of impending failures in critical components and mitigation of their detrimental effects on the engine operation. In this context, anomaly detection and health monitoring are formulated as a two-time-scale problem, where:

- The fast time scale is in the order of the response time of engine dynamics.
- The slow time scale is related to the time range over which the engine dynamics may exhibit non-stationary behavioral changes.

Over the span of a given data sequence in the fast time scale, the engine operational behavior is assumed to be statistically stationary. That is, the engine dynamics are essentially invariant on the fast time scale. In general, a long time span in the fast time scale is a very tiny (i.e., several order of magnitude smaller) interval in the slow time scale. For example, evolution of thermo-mechanical fatigue damage in the hot sections of an aircraft engine and the associated anomaly growth (causing a detectable change in measured outputs) usually occurs on the slow time scale. Nevertheless, the notion of fast and slow time scales is dependent on specific applications, engine load conditions and the operating environment.

From the above perspectives, health management of aircraft gas turbine engines is cast as a two-time-scale problem that requires early detection of anomalies occurring at a time scale that is several orders of magnitude slower than the engine dynamics.⁵ For example, the health condition of an engine may vary in tens of hours, whereas its dynamic response is on the order of seconds. Identification of the current state of the engine health at slow-time epochs is very important for maintenance engineers because necessary repairs must be carried out much before the engine becomes seriously damaged or permanently non-operable. Thus, it is essential to monitor slow-time-scale anomalies for gas turbine engines from the time series data of the engine observables on a fast time scale. To this end, a generic gas turbine engine simulation test bed has been used to validate the anomaly detection method.² The features of the engine simulation model are similar to those of the engine model reported in.³

Received 25 January 2005; revision received 25 January 2005; accepted for publication 14 November 2005. Copyright © 2006 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 1542-9423/04 \$10.00 in correspondence with the CCC.

* Corresponding Author: Email axr2@psu.edu Tel: (814) 865-6377

For detection of the slow-time-scale deviations, it might be necessary to rely on the time series data of macroscopic observables, generated from the available sensors, as well as other relevant information regarding pilot's experience. Since sufficiently accurate and computationally tractable spatial-temporal modeling of thermo-fluid and structural system dynamics is often infeasible solely based on the fundamental principles of physics, small changes in the system behavior may be inferred from both time series analysis of the sensor data and model-based information for the purpose of anomaly prediction and control systems design.⁶ By taking advantage of the available model-based information, the time series data can be converted into a symbol sequence^{7,8} that, in turn, generates a finite-state machine representation of the dynamical system behavior.

Recently, Ray⁹ has proposed a concept of anomaly detection, based on symbolic time series analysis (STSA) of measured macroscopic observables. The anomaly detection algorithm is built upon two-time-scale analysis of the stationary behavior of dynamical systems using the principles of symbolic dynamics,^{10,11} information theory,¹² and pattern recognition.¹³ Symbolic dynamics captures the essential dynamical features of the physical process through space partitioning. Information theory allows modeling of incipient catastrophic failures and chaotic behavior that are analogous to thermodynamic phase transitions. This is achieved by constructing finite-state machine models of the dynamical system behavior under nominal and anomalous conditions. Pattern discovery methods infer anomalies through quantitative evaluation of the deviations in statistical patterns of the respective state machines from those under the nominal condition.

This technical note presents the application of a novel method for real-time anomaly detection in aircraft gas turbine engines, which makes use of the time series data, generated from available sensors such as temperature, pressure, flow rate, and shaft speed. The algorithm of the anomaly detection software relies on symbolic time series analysis (STSA).^{8,9} The anomaly detection method has been tested on a simulation model of a generic gas turbine engine,² and its efficacy is evaluated relative to other existing pattern recognition techniques from the perspectives of early detection of turbine efficiency, which can be deemed as an indicator of the turbine's health state. The results are compared with those derived from traditional pattern recognition tools, such as Principal Component Analysis (PCA) and Artificial Neural Network (ANN); details are available in.¹⁴ The issues on real-time health monitoring of gas turbine engines are also discussed.

This technical note is organized in six sections including the present one. Section 2 reviews the concepts of STSA for finite state machine based anomaly detection. Section 3 summarizes the radial basis function approach to solve the anomaly detection problem. Section 4 discusses the principal component analysis technique. The simulation results are discussed in Section 5. The technical note is summarized and concluded in Section 6.

II. Symbolic Time Series Analysis (STSA)

As a consequence of its simplicity and ability to capture macroscopic dynamical behavior in a low dimensional state space, symbolic dynamics has been used as a tool for pattern recognition.⁸ In general, the state space of a physical process may have a very high dimension, which is difficult to handle numerically. Symbolic dynamics makes use of the data sequence for partitioning the continuous space with a predefined alphabet of symbols.^{10,15} The underlying concepts of symbolic dynamics and wavelet space partitioning are available in⁹ and,¹⁶ respectively.

The time series data of the combustor outlet temperature, employed herein for early detection of incipient faults, is converted to a symbol sequence by partitioning a finite region in the phase-space of engine operation dynamics, over which the time series data evolves, into finitely many discrete cells.^{7,17} Each cell is labeled as a symbol $\sigma \in \Sigma$, where the symbol set Σ is called *alphabet* consisting of m different symbols. In this way, a data sequence, obtained from a trajectory of the dynamical system, is converted to a symbol sequence $\{\sigma_{i1}, \sigma_{i2}, \sigma_{i3}, \dots\}$ that characterizes the system dynamics represented by the data sequence. Critical steps in the symbol generation process are: (i) partitioning of a finite region in the phase space; and (ii) construction of a mapping from the partitioning into the symbol alphabet, which becomes a representation of the system dynamics defined by the trajectories. There are potentially a number of ways to create a symbol sequence from the time series data. Two such methods to perform the partitioning are described below.

Kennel and Buhl¹⁸ have formulated a phase-space partitioning method based upon the concept of Symbolic False Nearest Neighbors (SFNN), where an algorithm is introduced to refine empirical partitions for symbolic state reconstruction. This method avoids topological degeneracy that is an essential feature of a generating partition.⁷ The major advantage of this method is that the partitioning is accomplished by an algorithm based on the time

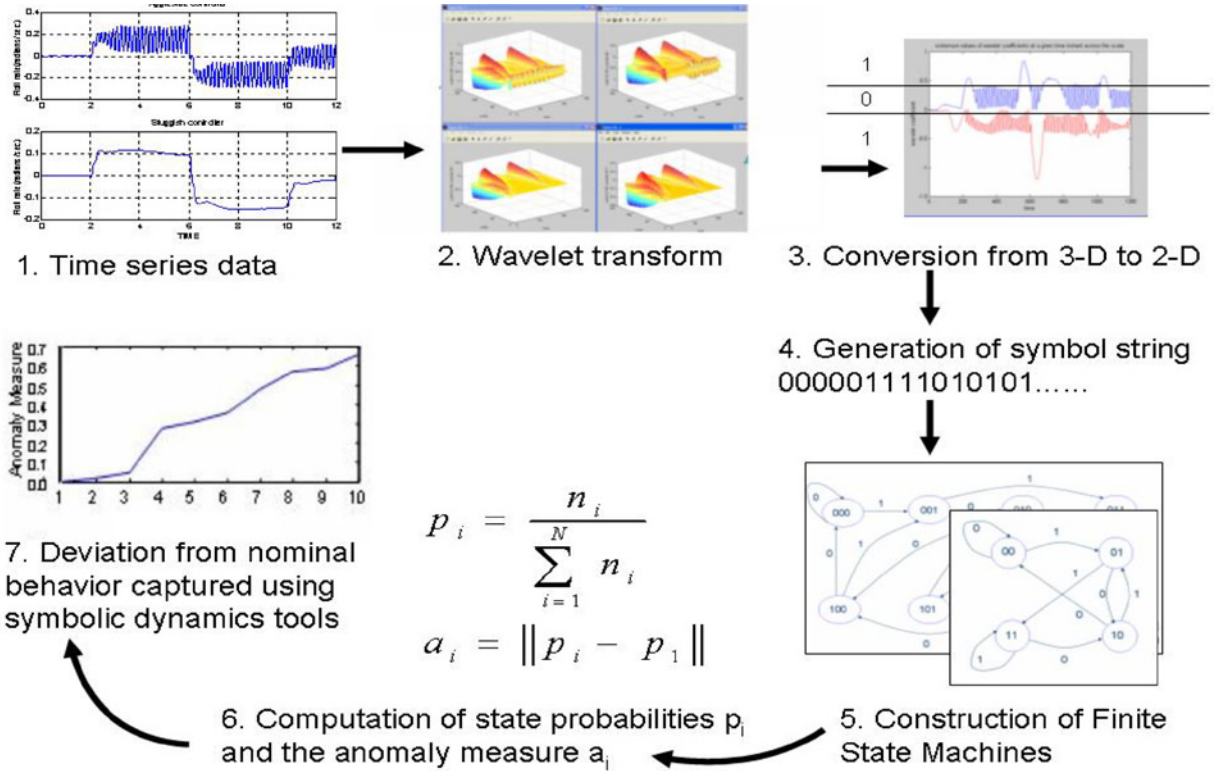


Fig. 1 Anomaly Measure Computation via Wavelet-based Symbolic Time Series Analysis

series data. For example, the partitions are defined with respect to a set of radial-basis influence functions, $f_k(x) = \alpha_k / (\|x - z_k\|^2)$, each associated with a symbol s_k with the center z_k and weight α_k . For each element x of the time series data set, one function $f_m(x)$ is generally expected to be greater than the other $f_k(x)$ with $k \neq m$. Then, the data point x in the phase space is transformed to a symbol s in the symbol space. The parameters z_k and α_k are the free optimization variables, with the constraint $\alpha_k \geq 0 \forall k$. There may be one or more influence functions assigned to each of the symbols in the alphabet. The partitions remain invariant at all epochs of the slow time scale. Partitioning of the phase space by the SFNN approach may often be difficult if the time series data are noise-corrupted.^{17,18} This problem is circumvented by if denoising is possible without any significant loss of dynamical information in the time series data.

An alternative partitioning scheme, which facilitates signal denoising, relies on preprocessing the time series data by wavelet transformation. This method has been recently proposed by Ray⁹ and subsequently expanded;¹⁶ this method is referred to as wavelet space (WS) partitioning in the sequel. The proposed scheme produces a graph of coefficients versus scale at each time shift. Since the wavelet transform is a function of two-variables: scale and time, these bivariate graphs are converted to univariate by stacking from end to end, starting with the smallest value of scale and ending with the largest value. For example, the wavelet coefficients versus scale at time shift t_k are stacked after those at time shift t_{k-1} to obtain the so-called *scale series data* in the wavelet space, which is analogous to the time series data in the phase space. In the present work, the WS is partitioned into horizontal slabs. Optimized partitioning of the WS will be addressed in subsequent work. The number of blocks in a partition is equal to the size of the alphabet and each block of the partition is associated with a symbol in the alphabet. For a given stimulus, the partitioning of wavelet space must remain invariant at all epochs of the slow time scale.

After the partitioning is finalized, the sequence of symbols is generated from the time series (or scale series) data at different epochs of the slow time at which anomalies may take place. Then, a probabilistic finite state automaton is constructed from the symbol sequence at each time epoch. The anomaly measure at a given epoch is obtained as a metric or distance function between the state probability vector of the finite state machine at that epoch and the state

probability vector of the finite state machine at the nominal condition. Thus, the above measure quantifies the growth of anomaly relative to the nominal condition in the slow time scale. The state probability vector under the nominal case serves as a benchmark. Obviously, the deviation measure at the benchmark condition is zero.

Finite state machines, generated from derived symbol sequences, identify behavioral patterns of the dynamical system. As the system trajectory evolves, different states are visited with different frequencies. While sliding the window from a state leading to another state under triggering of individual symbols from a sequence, counting of the number of times state visits generates the state probability vector under different anomalous conditions. After the state probability vectors have been obtained, the next step is to calculate the anomaly measure that signifies the change in stationary behavior of the dynamical system as the fault progresses. The step-by-step process is depicted in the flow chart shown in Figure 1.

III. Radial Basis Function Neural Network

Neural networks are capable of learning complex input-output relationships, using sequential training procedures, and adapting to the data.¹⁹ The most commonly used family of neural networks for pattern classification is the feed-forward network. The learning process involves updating the network architecture and its connection weights, so that the network can efficiently perform a specific classification/clustering task. Although neural networks provide a suite of nonlinear algorithms for feature extraction and classification, they are implicitly equivalent or similar to classical statistical pattern recognition methods.¹³ Despite these similarities, neural networks do offer several advantages such as, unified approaches for feature extraction and classification and flexible procedures for finding moderately nonlinear solutions. Among the most commonly used family of feed-forward neural networks for pattern classification, the Radial Basis Function Neural Network (RBFNN) was found to be the most suitable for anomaly detection based on time series data of observables(s). RBFNN is essentially a nearest neighbor type of classifier, where the activation of a hidden unit is determined by the distance between the input vector and the prototype vector.²⁰

The RBFNN technique used herein for the anomaly detection is an extension of the standard RBFNN to form a statistical model of nominal data. As new data enters into the anomaly detection system, it is compared with the RBFNN model. If it falls within the boundaries defined by the model, then it is considered as a nominal data; otherwise, the data is considered as anomalous. The approach is generic and has been applied to a variety of problems, including advanced military aircraft subsystems.²¹ A key requirement for RBFNN is appropriate selection of the radial basis function and the order of the statistics of the model. From this perspective, a radial basis function for anomaly detection is chosen as:

$$f(x) = \exp\left(-\frac{1}{\theta_\alpha} \sum_k |x_k - \mu|^\alpha\right) \quad (1)$$

where the parameter $\alpha \in (0, \infty)$; and μ and θ_α are the center and α th central moment of the data set, respectively. For $\alpha = 2$, $f(\bullet)$ becomes Gaussian, which is a typical radial basis function used in the neural network literature.

From a sampled time series data under the nominal condition, the mean μ and the central moment θ_α are calculated as:

$$\mu = \frac{1}{N} \sum_{k=1}^N x_k \quad \text{and} \quad \theta_\alpha = \sum_{k=1}^N |x_k - \mu|^\alpha \quad (2)$$

The distance between any vector \mathbf{x} and the center μ is obtained as:

$$\|\mathbf{x} - \mu\|_{\ell_\alpha} = \left(\sum_k |x_k - \mu|^\alpha\right)^{1/\alpha} \quad (3)$$

Hence, at the nominal condition, the radial basis function $f_{nom} = f(x)$. For different anomalous conditions, the parameters, μ and θ , are kept fixed; and the radial basis function f_k is evaluated from the data set under the (possibly anomalous) condition at the k^{th} epoch is defined as a distance function

$$M_k \equiv d(f_{nom}, f_k) \quad (4)$$

IV. Principal Component Analysis

In the statistical approach, each pattern is represented in terms of d features or measurements and is viewed as a point in a d -dimensional space. The goal is to choose those features that allow pattern vectors belonging to different categories to occupy compact and disjoint regions in the d -dimensional feature space. The effectiveness of the representation space, generated from the feature set, is determined by how well patterns belonging to different classes can be separated. Given a set of training patterns from each class, the objective is to establish decision boundaries in the feature space to separate patterns belonging to different classes. In the statistical decision-theoretic approach, the decision boundaries are determined by the specified or learnt probability distributions of the patterns belonging to each class.

Feature extraction methods determine an appropriate subspace of dimension m in the original feature space of dimension d with $m \leq d$. The best known linear feature extraction technique is the Principal Component Analysis (PCA)¹³ that makes use of Karhunen-Loève expansion to compute the m largest eigenvectors of the $d \times d$ covariance matrix of the N patterns, each of which is d -dimensional. Since PCA uses the most expressive features (eigenvectors with the largest eigenvalues), it effectively approximates the data on a linear subspace using the mean squared error criterion.

To detect anomaly growth from time-series data, the PCA is performed for dimensionality reduction and thus serves as a feature selector in the pattern analysis. For a time series data with length L , an $M \times N$ data matrix is to be created by dividing the time-series data string of length L into substrings of length $M = L/N$, where each row of the data matrix is a substring of length N . Then, the $N \times N$ covariance matrix, obtained from the data matrix, generates the orthonormal eigenvectors $v_1 \dots v_N$ and the corresponding non-negative real eigenvalues $\lambda_1 \geq \dots \geq \lambda_N \geq 0$, where the eigenvalues are arranged in the increasing order of magnitude. The m largest eigenvalues and associated eigenvectors are selected such that $\sum_{i=1}^m \lambda_i > \eta \sum_{i=1}^N \lambda_i$, where $\eta < 1$ is a real positive number close to 1 (e.g., $\eta = 0.95$). The feature matrix F is defined as:

$$F = \left[\sqrt{\frac{\lambda_1}{\sum_{k=1}^N \lambda_k}} v_1 \quad \dots \quad \sqrt{\frac{\lambda_m}{\sum_{k=1}^N \lambda_k}} v_m \right] \quad (5)$$

The feature matrix F_{nom} represents the engine health status derived from the time series data at the nominal condition. Similarly, feature matrices $F_1, F_2 \dots$ are obtained from time series data at different epochs in the slow time scale. The anomaly measure for at a slow time k is obtained as:

$$M_k = d(F_{nom}, F_k) \quad (6)$$

where d is a metric signifying the distance between the nominal and the k^{th} anomalous conditions. For the results shown in the next section, metric was chosen to be the Euclidian norm.

V. Results of Simulation Experiments

This section presents the results of the application of various anomaly detection methods on the simulation test bed of a generic gas turbine engine, similar to that reported in.³ The model contains steady state performance maps for all the components and has control volumes where continuity and energy balances are maintained. Rotor dynamics and the duct momentum dynamics are also included in this model.

In the health monitoring of the engine, the degradation of engine efficiency from the nominal condition (e.g., a new or an overhauled engine) can be treated as anomaly. For implementation purposes, the high and low pressure turbine efficiencies, fan and compressor tip velocity ratios are reduced to observe the effects on the time series data. Here, the degradation in efficiencies of different engine components is assumed to be uniform. This assumption is reasonable for the comparative study presented herein.

The health condition (e.g., efficiencies) of engine components changes at a very slow time scale. Thus, for a short span of time, the efficiency values are assumed to remain approximately constant for all practical purposes. Replication of these conditions may require hundreds of hours of engine simulation. Since this is not a feasible solution, engine component efficiencies are reduced for each run of the simulation and a certain period of operation is observed for each health condition. This is conceptually similar to experimental methods, where an engine is tested

under extreme conditions (accelerated degradation) to simulate many hours of engine operation within small time duration. Only steady state data was used for this analysis.

For anomaly detection, the time series data from a number different sensor sources was analyzed. Based on these simulation experiments it was found that the combustor outlet temperature (T_c) and the main burner fuel flow rate (F_f) variables, best captured the degradation of the engine. From thermodynamic perspective, it was expected that these two variables would be affected by the change in engine efficiencies. The results of the simulation corroborate this fact. Upon further investigation, it was found that the combustor outlet temperature (T_c) essentially captured the dynamics of degradation and using the fuel flow rate (F_f) did not add anything significant to the analysis. Therefore, the results on anomaly detection are based on the time series analysis of the combustor outlet temperature. The sets of time-series data were collected after the dynamic response attained the stationary behavior. These data sets were used to compare the anomaly detection capability of the symbolic dynamics approach relative to that of two existing pattern recognition techniques: PCA and RBFNN. Since symbol generation from time series data is the crucial step in symbolic-dynamics-based anomaly detection, this paper investigates two alternative approaches: Symbolic False Nearest Neighbors (SFNN) partitioning and Wavelet Space (WS) partitioning.

Time series data belonging to different health conditions (up to 2.5% efficiency drop) have been used for the analysis purpose. It is noticed that the output data are almost impossible to distinguish with a simple threshold check or visual inspection, or a peak pass test in the frequency domain. However, when the saturation occurs, high fluctuations in temperature are observed and visual inspection might be sufficient for anomaly detection. However, the goal here is to capture the anomalous behavior much before saturation sets in, and a permanent damage occurs to the engine.

Time series data were gathered by executing the simulation model under different anomalous conditions, by altering the efficiency values and flow constants to generate simulated faults. The nominal relative parameter for each efficiency was set to 1.0 and, under anomalous conditions; the parameters were reduced down to 0.97. The effects of the decreasing efficiency on the output data were observed under persistent excitation. For example, in these simulation experiments, output data was observed by perturbing the booster vane angle. The perturbation used for this purpose was a square wave with amplitude 10% of the nominal value. The four plots in Figure 2 compare the anomaly measures obtained by using the afore-described anomaly detection approaches, SFNN, WS, PCA, and RBFNN. The nominal condition is chosen at the time epoch when all transients have decayed.

For the WS partitioning method, the alphabet size used in partitioning the wavelet coefficient space is 8. For the SFNN method, the alphabet size to generate the finite state machine is 4. Comparison of the two Symbolic Dynamics based partitioning techniques shows that the SFNN method performs marginally better than the WS method in

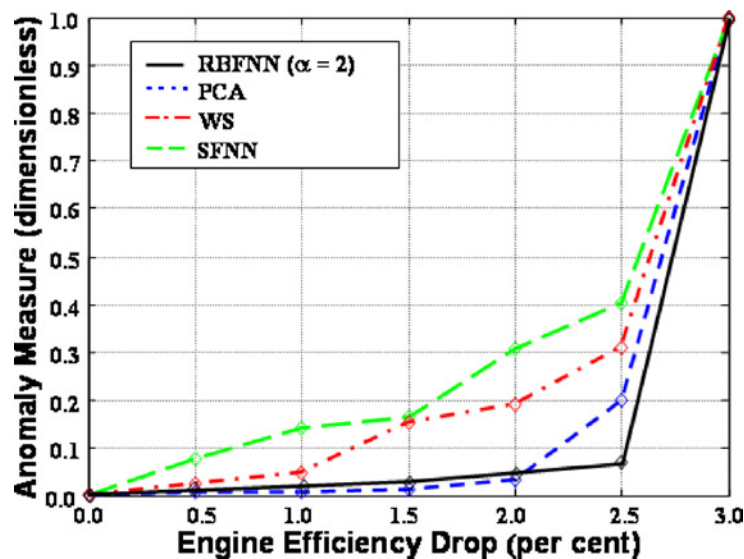


Fig. 2 Comparison of Different Methods of Anomaly Detection

predicting the onset of anomaly. For the RBFNN method exponent value $\alpha = 2$ (Gaussian RBF) is used. Results are normalized with respect to maximum anomaly condition. For the PCA method, time series data is divided into 10 subsections to form the data matrix. PCA technique also agrees with the other ones predicting the health condition of the gas turbine engine before saturation. Figure 2 clearly shows that the results for RBFNN and PCA based methods are comparatively inferior to the symbolic dynamics based methods in terms of early detection of anomalies.

VI. Summary and Conclusions

This technical note presents the application of a novel method for early detection of anomalies in complex mechanical systems, where the physical process is approximately stationary in the fast time scale of the process dynamics and any non-stationarity due to gradually evolving anomalies occur in the slow time scale. The anomaly detection method is based on symbolic time series analysis and its efficacy has been examined on the simulation test bed of a twin-engine propulsion system. Time series data of observed macroscopic variables, generated on the fast time scale from the simulation model, are analyzed at slow time scale epochs for early detection of anomalies. The results are compared with those derived from traditional pattern recognition tools, such as Principal Component Analysis (PCA) and Artificial Neural Network (ANN).

Four methods for anomaly detection are compared quantitatively. Although all techniques were able to predict the saturation point correctly, the best results are obtained using phase space partitioning based on the Symbolic False Nearest Neighbors (SFNN) and Wavelet Space (WS) methods. Computational complexity is an important issue for online implementation. The symbolic time series analysis (STSA) method based on both SFNN and WS partitioning are compared with more traditional techniques using Radial Basis Function Neural Network (RBFNN) and Principal Component Analysis (PCA). The STSA method performed significantly better compared to PCA and RBFNN techniques in terms of early detection of incipient anomalies. The WS partitioning technique has been investigated in detail. Future research will focus on health management and control of aircraft propulsion systems based on anomaly detection.

VII. Acknowledgments

This work has been supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under Grant No. DAAD19-01-1-0646 and NASA Glenn Research Center under Grant No. NNC04GA49G.

References

- ¹Dai, X., and Ray, A., "Life Prediction of the Thrust Chamber Wall of a Reusable Rocket Engine," *AIAA Journal of Propulsion and Power*, Vol. 11, No. 6, November–December 1995, pp. 1279–1287.
- ²Fu, J., Yasar, M., and Ray, A., "Optimal Discrete Event Supervisory Control of Gas Turbine Engines," *American Control Conference*, Boston, MA, 2004.
- ³Diao, Y., and Passino, K. M., "Stable Fault-Tolerant Adaptive Fuzzy/Neural Control for a Turbine Engine," *IEEE Trans. Control Systems Technology*, Vol. 9, No. 3, May 2001.
- ⁴Litt, J. S., Parker, K. I., and Chatterjee, S., "Adaptive Gas Turbine Engine Control for Deterioration Compensation Due to Aging," *International Symposium on Air Breathing Engines*, ISABE 2003-1056.
- ⁵Volponi, A., DePold, H., Ganguli, R., and Daguang, C., "The Use of Kalman Filter and Neural Network Methodologies in Gas Turbine Performance Diagnostics: A Comparative Study," *ASME Turbo Expo 2000*, 2000-GT-547, Munich, 2000.
- ⁶Khatkhate, A., Ray, A., Chin, S., Rajagopalan, V., and Keller, E., "Early Detection of Fatigue Crack Anomaly: A Symbolic Dynamic Approach," *American Control Conference*, Boston, MA, June–July 2004.
- ⁷Badii, R., and Politi, A., *Complexity, Hierarchical Structures and Scaling in Physics*, Cambridge University Press, Cambridge, U.K. 1997.
- ⁸Daw, C. S., Finney, C. E. A., and Tracy, E. R., "A review of symbolic analysis of experimental data," *Review of Scientific Instruments*, Vol. 74, No. 2, 2003, pp. 915–930.
- ⁹Ray, A., "Symbolic Dynamic Analysis of Complex Systems for Anomaly Detection," *Signal Processing*, Vol 84, No. 7, 2004, pp. 1115–1130.
- ¹⁰Lind, D., and Marcus, M., *An Introduction to Symbolic Dynamics and Coding*, Cambridge University Press, United Kingdom, 1995.
- ¹¹Shalizi, C. R., Shalizi, K. L., and Crutchfield, J. P., "An algorithm for pattern discovery in time series," *SFI Working Paper* 02-10-060, 2002.

- ¹²Cover, T. M., and Thomas, J. A., *Elements of Information Theory*, John Wiley, 1991.
- ¹³Duda, R., Hart, P., and Stork, D., *Pattern Classification*, John Wiley & Sons Inc., 2001.
- ¹⁴Chin, S., Ray, A., and Rajagopalan, V., "Symbolic Time Series Analysis for Anomaly Detection: A Comparative Evaluation," *Signal Processing*, Vol. 85, No. 9, 2005, pp. 1859–1868.
- ¹⁵Beck, C., and Schlogl, F., *Thermodynamics of Chaotic Systems: an Introduction*, Cambridge University Press, Cambridge, U.K., 1993.
- ¹⁶Rajagopalan, V., and Ray, A., "Wavelet-based Space Partitioning for Symbolic Time Series Analysis," *IEEE Conference on Decision and Control (CDC) and European Control Conference (ECC)*, Seville, Spain, December 2005.
- ¹⁷Abarbanel, H. D. I., *The Analysis of Observed Chaotic Data*, 1996, Springer-Verlag, New York, 1996.
- ¹⁸Kennel, M. B., and Buhl, M., "Estimating good discrete partitions from observed data: symbolic false nearest neighbors," *arXiv:nlin.CD/0304054 v1*, April 26, 2003.
- ¹⁹Haykin, S., *Neural Networks: A Comprehensive Foundation*, Prentice Hall, Upper Saddle River, NJ, 1999.
- ²⁰Bishop, C. M., *Neural networks for pattern recognition*, Oxford University Press Inc., New York, 1995.
- ²¹Brotherton, T., and Johnson, T., "Anomaly Detection for Advanced Military Aircraft Using Neural Networks," *Proceedings of 2001 IEEE Aerospace Conference*, Big Sky Montana, March 2001.