

NRC Publications Archive Archives des publications du CNRC

Improving Preciseness of Time to Failure Predictions: Application to APU starter

Létourneau, Sylvain; Yang, Chunsheng; Liu, Ziyang

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. /
La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version
acceptée du manuscrit ou la version de l'éditeur.

Publisher's version / Version de l'éditeur:

*Proceedings of the 1st International Conference on Prognostics and Health
Management (PHM 2008), October 6-9, 2008, Denver, Colorado, USA>, 2008*

NRC Publications Archive Record / Notice des Archives des publications du CNRC :
<https://nrc-publications.canada.ca/eng/view/object/?id=5f28f850-5b03-46ba-a475-05e6fc812626>
<https://publications-cnrc.canada.ca/fra/voir/objet/?id=5f28f850-5b03-46ba-a475-05e6fc812626>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at
<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site
<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at
PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the
first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la
première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez
pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.



National Research
Council Canada

Institute for
Information Technology

Conseil national
de recherches Canada

Institut de technologie
de l'information

NRC - CNRC

Improving Preciseness of Time to Failure Predictions: Application to APU starter *

Létourneau, S., Yang, C., Liu, Z.
October 2008

* published in Proceedings of the 1st International Conference on
Prognostics and Health Management (PHM 2008). Denver, Colorado,
USA. October 6-9, 2008. NRC 50409.

Copyright 2008 by
National Research Council of Canada

Permission is granted to quote short excerpts and to reproduce figures and tables
from this report, provided that the source of such material is fully acknowledged.

Improving Preciseness of Time to Failure Predictions: Application to APU starter

Sylvain Létourneau and Chunsheng Yang and Zhenkai Liu
Institute for Information Technology
National Research Council Canada
Ottawa, Ontario, Canada, K1A 0R6
Email: (Sylvain.Létourneau, Chunsheng.Yang)@nrc-cnrc.gc.ca

Abstract—Despite the availability of huge amounts of data and a variety of powerful data analysis methods, prognostic models are still often failing to provide accurate and precise time to failure estimations. This paper addresses this problem by integrating several machine learning algorithms. The approach proposed relies on a classification system to determine the likelihood of component failures and to provide rough indications of remaining life. It then introduces clustering and SVM-based local regression to refine the time to failure estimations provided by the classification system. The paper illustrates the applicability of the proposed approach through a real world aerospace application and discusses data pre-processing requirements. The preliminary results show that the proposed method can reduce uncertainty in time to failure estimates, which in turn helps augment the usefulness of prognostics.

I. INTRODUCTION

The need for higher equipment availability and lower maintenance cost is driving the development and integration of prognostic and health management (PHM) systems. Taking advantage of advances in sensor technologies, PHM systems favor a pro-active maintenance strategy by continuously monitoring data from the equipment and informing the maintenance staff whenever there is a risk for a component failure. A PHM system may also supplement the component failure predictions with an estimation of the time to failure (TTF), which is defined as the expected remaining time before the given component stops fulfilling its function. In order to avoid disruption and minimize maintenance costs, these time to failure estimates need to be as reliable and as precise as possible.

Traditional methods to estimate TTF include reliability analysis [1] and knowledge-based approaches from physics and material sciences [2]–[5]. These approaches help to understand the underlying physical mechanisms but they require enormous amounts of background information. They may also be difficult to apply as they tend to rely on difficult to obtain data on component damage or material properties. With the development and integration of data acquisition devices into complex equipment, data mining-based approaches are now starting to complement the traditional methods for building prognostic models [6]. Recent results show the potential of classification systems to identify the likelihood of component failures in a timely manner but none of the existing techniques can provide precise time to failure estimates required for the optimization

of maintenance. For instance, the KDD methodology proposed in [6] can build classification models for prognostics. These models continuously assess the probabilities of a component failure within a pre-specified alert target window (e.g., between 1 and 20 days in advance of a functional failure), but often fail to provide precise TTF estimates. When a classifier detects patterns in the data that are characteristic of an incipient failure, it generates an alert indicating that the suspected component is likely to fail within the alert target window without being able to specify the exact number of days or hours of operation left. With this approach, the larger the alert target window, the larger the imprecision on the TTF estimates. In some specific applications, it is reasonable to try to increase precision by reducing the width of the target window. However, this is generally not suitable as it could prevent the end users from getting alerts as early as possible which, in turn, would reduce the opportunity for optimization and the benefits of prognostics. A too narrow target window may also have detrimental effects on the performance of the predictive models. For instance, when a component has various failure modes, each following their own time frame, there is a risk that a model specific to a narrow target window would only be able to detect a fraction of these failure modes.

Since predicting TTF can be seen as a regression problem, regression analysis and time-series forecasting methods could be used to build models that try to directly estimate TTF from the sensor data. To be successful, such models need to accurately map all the subtle changes in the data to specific life reduction estimates. These models also need to account for the fact that with complex components, we often observe significant variations in actual time to failure. Obviously, building such models is a very challenging task that requires ample amounts of high quality and relevant data. Since data from real world equipment is typically characterized by issues such as irregular sampling intervals, small signal/noise ratio, and sensor measurement errors, it is generally hopeless to try to develop a global regression model for TTF from sensor data. On the other hand, it is plausible that regression could be successfully applied locally on well chosen portions of the real world sensor data. This paper investigates this hypothesis by trying to demonstrate that regression analysis can help improve the preciseness of TTF estimates.

This paper proposes to improve the preciseness of TTF

estimates by combining classification and regression-based approaches. It relies on a comprehensive data mining methodology to develop the required classification system. The classifier developed is capable of identifying incipient component failures and providing rough TTF estimates. Clustering is used to partition the sensor data and a regression model is developed to estimate TTF within each cluster. When the classifier uncovers a potential component failure, a mapping function decides which regression model should be used to provide a TTF estimate. A final step produces the final TTF estimation based on the output from the classification and regression models. We name the proposed method "on-demand regression" since regression is only used once the classification system has identified the potential for a component failure. This paper extends a preliminary description of this work [7] by providing updated results and by discussing how the choice regarding the number of clusters affect performance.

Before detailing the approach, the paper explains the challenges with support from real world data from an aerospace application. The same application is also used to illustrate the applicability and the usefulness of the proposed approach.

II. CHALLENGES

Accurate and detailed health information on key systems and components is of utmost importance to help optimize the maintenance and management of complex equipment. Ideally, powerful prognostic models, well integrated into the organization's information system, would automatically combine sensor data, historical maintenance information, system's configuration, and other sources of information to continuously provide accurate and precise TTF information. Regression methods, which are specifically designed to predict numerical values such as TTF, appear well suited to develop these models. Unfortunately, many typical issues of real world data from complex equipment severely constrain the applicability and power of regression modeling. To illustrate, let us consider an aerospace application in which the objective is to build a prognostic model for the starter motor of the Auxiliary Power Unit engine (APU).

The data for this application has been produced by a fleet of 35 commercial aircraft over a period of 10 years. The specific dataset used consists of 18 attributes (5 symbolic, 11 numeric, and 2 for date and time of the event). More than 161000 observations are available for this task. Only a subset of these observations are relevant for learning the predictive models. These are the ones collected around each occurrence of component failures. In this particular task, we use engine operating hours as the time unit. We based our analysis on data generated between 250 operating hours prior to the failure and 30 hours after. A comprehensive search in the maintenance database revealed information on 83 occurrences of APU starter motor replacements. Since we do not have access to information from further testing of the components replaced, we assume that a replacement is equivalent to an actual failure. When an engine suffered consecutive failures in a short period of time, we constrained the above interval

to make sure that each observation is included only once. We used data from 61 failures for learning and kept data from the remaining 22 failures for testing.

To evaluate the feasibility of regression as a direct way to predict TTF, we augmented the initial representation with a TTF attribute. This attribute is simply defined as the difference between the engine operating hours in the current observation and the operating hours of this engine at the next starter failure. To make the regression task easier, we removed the instances observed after the failures. We built an SVM-based regression model using the training data set and then applied it on the testing set. Figure 1 shows results from one of the best SVM (Support Vector Machine) models developed. The scatter clearly illustrates the lack of fit of the model. The expected error from this model is 58.7 with a standard deviation of 43.8. Other regression methods such as NN (Neural Network) and linear regression lead to similar performance. The following paragraphs discuss some of the reasons that explain the lack of success of global regression models.




Fig. 1. TTF versus actual TTF from a global SVM-based regression model.

Irregular sampling rate Most traditional forecasting techniques require a fixed interval between measurements (e.g., every hour, every day, or every month). In this particular application, the data is systematically collected at every start of the APU. On the other hand, the data is not sent to the central system unless the on-board system decides to do so. Due to different configurations across the fleet and over time, we ended up having APU data at every start for some of the aircraft, at every other start for other aircraft and so on up to every 8 starts. As a result, some forecasting methods can simply not be applied. Unequal sampling rate also means that the time-series corresponding to the different failures have different lengths. This may lead to an unbalanced representation of the failure cases. A possible solution is to re-sample through

interpolation or smoothing but given the high variability in the original sampling rate and the small signal/noise ratio, such processing is likely to hurt the modeling for TTF estimation.

Lack of relevant information In order to accurately predict remaining life, the model needs to be able to estimate the current life consumption. The information required to evaluate life consumption could come from highly informative core measurements that adequately account for the internal state of the component. Such high value information is typical from laboratory testing equipment but rarely available from sensors deployed on today's complex equipment. Information about life consumption can also be captured directly by means of a counter. For instance, in the APU starter motor application, we use the engine operating hours to approximate life consumption but it is far from perfect since the starter motor and the engine may consume their life differently. Moreover, engine overhauls which happen at regular intervals result in a reset of the engine operating hours counter. Since we do not have access to detailed information on work performed during the overhaul, it is impossible to determine if the reset of the engine operating hours counter also corresponds to the repair of the starter motor. Consequently, these resets can possibly introduce cuts in the evaluation of life consumption and cause great difficulty for the regression models.

Large variance due to contextual effects Equipment such as aircraft operates in a very dynamic environment. Changes in this environment affect the behavior of the system. In some cases, these changes also affect the measurements taken. For instance, all measurements related to temperature, flow, and pressure are likely to be affected by the altitude of the aircraft. The mode of operation and the status of internal sub-systems and components are also likely to affect the behavior of the performance parameters. All of these contextual effects need to be accounted for in order to understand the behavior of the key parameters and properly use them to infer reliable TTF estimates.

Notwithstanding all of the difficulties mentioned above, this paper argues that regression can still play a role in helping to improve the TTF estimates in prognostic applications. As explained in the following section, the main idea is to partition the data space into relatively homogeneous data subsets and then use different regression models for these subsets.

III. ON-DEMAND REGRESSION

Figure 2 illustrates the proposed hybrid approach to improve the preciseness of TTF estimates. The tree steps are described below.

A. Classification-based prognostic

The first step involves a binary classifier that can identify incipient component failures and provide a rough estimate of the remaining useful time. We build this classifier using the KDD methodology documented in [6], [8]. This methodology consists of several steps which we now succinctly describe.

In order to use classification learning, we need to add a class attribute to the sensor data. We proceed with an

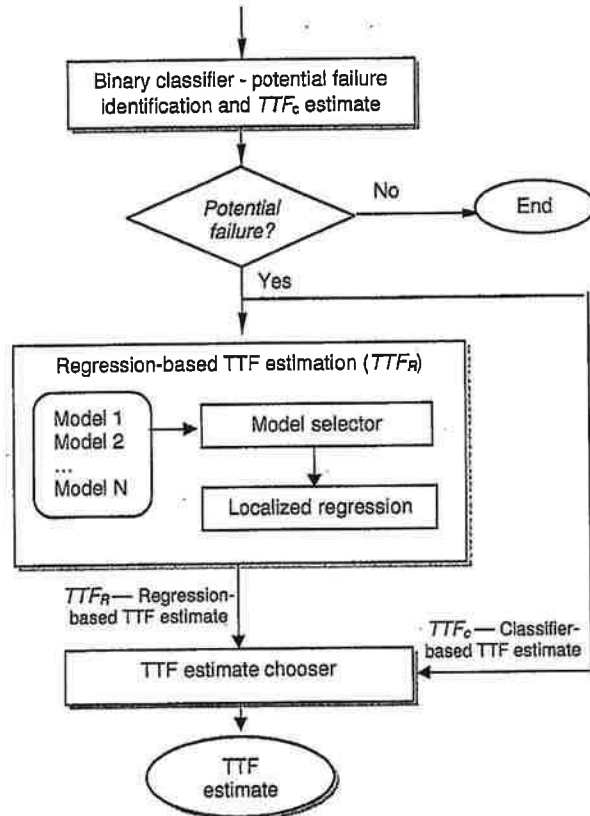


Fig. 2. Overview of the hybrid approach proposed to improve TTF estimates.

automated approach. This approach labels as positive ("1") all instances that fall in a pre-determined target window before the occurrence of a starter motor failure and as negative ("0") all other instances. This labeling scheme allows us to build a classifier that generates an alert whenever the patterns in the data are similar to the ones observed near a failure. In practice, we define the length of the target window by taking into account the optimal period for the end users to receive the alerts and the balance between positive and negative instances. As a rule of thumb, we try to keep a minimum of 15% as positive instances to simplify the learning.

Since data representation is often a key factor, we systematically try to improve the initial representation by augmenting it with new informative features. We construct these features using methods from signal processing, time-series analysis, and constructive induction. Feature selection is also applied on the augmented data representation to automatically remove correlated or irrelevant features [9], [10].

After updating the initial dataset with the class attribute and incorporating data representation enhancements, we build the required classifier. We use data from a subset of all failures for learning the models and keep the remaining data for testing. Any classifier learning algorithm can be used. In early experiments, we tended to prefer simple algorithms such as decision trees and naive-Bayes over more complex ones

because of their efficiency and because they produce models that we can easily explain to the end users. We apply the same algorithm several times with varying attribute subsets and cost information. To compare the classifiers obtained, we apply a score-based approach that we have developed to evaluate classifiers for prognostic systems [6]. The one with the maximal score on testing data is selected as the best classifier.

As illustrated in Figure 2, the classifier needs to provide a rough TTF estimate (TTF_C) whenever it predicts a potential component failure. We define this TTF estimate based on the expected number of operating hours left between a positive prediction and the actual failure. We use only the training data to compute this expected value. Precisely, $TTF_C = 1/N \sum_{i=1}^N \text{RemainingOPH}_i$ where RemainingOPH_i is for the difference between the engine operating hours in the i^{th} positive prediction from the training set and the operating hours of the corresponding engine at the next starter failure, and N is the number of positive predictions made by the classifier on the training data set. This value is constant for all positive predictions made by the classifier.

B. Regression-based TTF estimation

The objective of the second step is to try to improve the preciseness of the TTF estimates provided by the classifier described above. This is done through localized regression models. Each model accounts for a specific area of the data space. Every time the classifier makes a positive prediction, one of the local regression models is selected and applied to compute a new TTF estimate. The construction of the models required goes as follows.

First, clustering is used to partition the time-series associated to the various failures. The intent is to obtain clusters of time-series as homogeneous as possible with respect to the performance of the core measurements. This is done by clustering based on the attributes that represent meaningful contexts for the component of interest. In other words, we use clustering to obtain subsets into which the potentially negative effect of contextual conditions is minimized. In the case of the APU starter motor application, the predominant contextual attribute is the age of the starter motor at the time of the failure. As explained above, we approximate this age using the engine operating hours at the time of the failure. As it is often the case in clustering task, we need to pre-determine the number of clusters required for the given application. This number must be sufficiently high to obtain acceptable homogeneity within each cluster but not too high as to avoid over partitioning the data. Additionally, we also need to ensure that each cluster contains at least one test time-series for the evaluation purpose. Since we have 22 occurrences of APU starter failure in the test data, we can have at most 22 clusters. As we will explain in the discussion section, it turns out that there was no benefit in using more than 16 clusters for this specific application.

Second, a model selector is developed in order to assign each positive prediction to a given data subset (Fig 2). The clustering model built for partitioning the data cannot be

deployed for this task as it relies on the operating hours at the failure time, which is unknown for yet to fail components. We resolve this issue with an N -class classifier, where N is the number of clusters. Once the clustering scheme has been established, we tag each instance with its cluster ID and learn a classifier that can tell apart the instances as accurately as possible using the measurements available. Based on our experiments, simple decision trees and naive-Bayes classifiers perform very well for this task with a typical accuracy of 90% on test data.

Finally, a regression model is built for each cluster. SVM, NN, and many statistical regression techniques can be used to build these models. A key aspect of this step is that it uses only a subset of the training data available in each cluster to learn the models. Precisely, it uses use only the instances where TTF (i.e., remaining engine operating hours) is less than the expected TTF from the classifier. This allows us to further limit the scope of the regression models to the areas with the greatest potential for enhancing the precision of TTF estimates. The evaluation procedure starts by running the classifier on the test data to identify positive predictions. For each of these positive predictions, the model selector chooses an adequate regression model. The chosen model computes the regression-based TTF estimate noted TTF_R .

C. Selecting which TTF estimate to use

Two TTF estimates are produced for each positive prediction: one from the classifier (TTF_C) and one from a local regression model (TTF_R). We now need to decide how to combine them into a single TTF estimate. Our approach is very simple; it returns TTF_R if $TTF_R < TTF_C$. This is to avoid potentially significant errors that could come from an extrapolation of a regression model. TTF_C corresponds to the value used to limit the range of the output attribute while learning the regression models. If at deployment (or testing) time, a regression model outputs a prediction that is higher than the maximal value observed during training (i.e., TTF_C) then the model is extrapolating. Since extrapolation from local models is risky, we prefer to disregard such predictions and rely on the default classifier-based estimate.

IV. EXPERIMENTS AND RESULTS

This section reports experimental results on the application of the proposed method to try to improve the preciseness of TTF estimates for prognostic of APU starter failures on commercial aircraft. Detailed information about the data have been discussed in the *Challenges* section. All models have been built using the WEKA package.

To evaluate the performance, we conducted a 4-fold cross validation experiment. Table I summarizes the training and testing datasets for each fold. The experimental results are shown in Table II. The last line in the table is the average performance over the cross validation runs. The classifier for identifying potential failures was a multiple classifier systems (MCS) [6], [8] which combined two binary classifiers which are built using the J48.PART and J48 algorithms with default options. We configured the automatic

Fold #	Test data		Train data	
	failure cases	instances	failure cases	instances
1	22	2307	61	8814
2	20	2704	63	8560
3	20	2735	63	8243
4	21	3375	62	7746

TABLE I
DATA SUMMARY FOR EACH OF THE FOLD IN THE CROSS-VALIDATION EXPERIMENT.

labeling step so that it tags all observations with remaining engine operating hours less than 50 hours as positive and all others as negative. This provides sufficient time for the maintenance staff to plan the repair of the starter prior to the actual failure. In this experiment, we used only the raw measurements without any data representation enhancement. We also used default cost information. The expected TTF estimate from this classifier is an average 22.5 hours. As reported in the last line of Table II, the average error of the TTF estimates from this model alone on the test data is 14.5 with a standard deviation of 14.4. Figure 3 shows the graph of the TTF estimates versus actual TTF when using only the binary classification system. We notice that all the points are around 22.7 which is the estimate that this model returns for all positive instances from the test data set.




Fig. 3. Predicted TTF versus actual TTF using only the binary classification system.

As mentioned earlier, the operating hours at the failure time was used to partition the 83 time-series (one for each failure case) into 16 clusters. Results reported are based on K-means clustering. Experiments with EM-based clustering produced similar results. The model selector was built using J48. Its accuracy on test data is slightly above 85%. We used SMOReg with a linear polynomial kernel to construct the 16 local regression models. If we assume that these models would

Fold	Regression		TTF_C	Binary		On-demand	
	Err±std	MSE		Err±std	MSE	Err±std	MSE
1	20.7±30.8	1373	21.7	14±14.1	393	5.9±16.7	311
2	20.5±36.7	1762	21.3	12.9±8.5	238	4.2±7.4	72
3	22.9±30.1	1433	23.3	14.9±17.2	517	7.9±12.4	216
4	26.9±35.3	1964	24.6	16.2±17.9	581	9.5±25.1	714
Avg.	22.7±33.2	1639	22.5	14.5±14.4	432	6.9±15.4	328

TABLE II
ERROR FOR TTF ESTIMATES ON TEST DATA FOR USING REGRESSION-ONLY, CLASSIFICATION-ONLY, AND THE ON-DEMAND REGRESSION APPROACHES.

be used to generate all TTF estimates, then the results would be as illustrated in Figure 4. The scatter in the graph shows the lack of fit between many of the estimates and the actual TTF values. This is also confirmed by the second column on the last line in Table II, which reports an average error of 22.7 ± 33.2 hours.




Fig. 4. Predicted TTF versus actual TTF using estimates from the regression-based models only.

The results from the on-demand regression approach are presented in Figure 5 as well as in the last two columns in Table II. With this approach, we observed a much better fit between the estimates and the actual TTF values. With an average error of 6.9 ± 15.4 hours and a MSE of 328, the proposed approach clearly outperforms the initial classification-based approach. The strong reductions in the average error and in the standard deviation suggest an improvement in the preciseness of TTF estimates by a factor of 3.

V. DISCUSSION

The APU starter engine application has demonstrated the applicability of the proposed on-demand regression approach to improve the precision of TTF estimations. On the other hand, a number of aspects deserve further attention.




Fig. 5. Predicted TTF versus actual TTF from the proposed on-demand regression approach.

One of the key parameter of the proposed approach is the number of clusters. In this application, we could have selected as many as 22 clusters since we had 22 occurrence of failures in the test set and that we needed a minimum of one test set in each cluster to evaluate the corresponding regression model. To decide on the optimal number of cluster for the given application and to better understand the impact of this choice, we experimented with 1, 4, 7, 10, 13, 16, 19, and 22 clusters. In each experiment, we built a regression model for each cluster, a model selector, and then ran the test data to evaluate the expected performance. From these experiments, we observed that increasing the number of clusters consistently improves the accuracy of the local regression models but also decreases the performance of the model selector. This later observation is explained by the fact that the probability of selecting the right local regression model decreases as the number of cluster increases. The challenge is to find a tradeoff that would minimize the overall average error on TTF predictions. As shown in Figure 6, 16 clusters appears to be the optimal choice in the case of the APU starter failure application. More than 16 clusters would not produce any significant benefit while less could reduce the expected precision of the estimates.

There are many ways to combine the results from a binary classification system and from one or more regression models. In this paper, we presented a simple combination rule. We only keep the TTF_R that are less than TTF_C . Although this simple rule appears useful in improving the TTF estimations, it could negatively affect the fault detection rate. In a real-world setting, the combination rule could be adapted by taking into account, the intended usage of the alerts from the on-demand regression model, the operational constraints, and the various costs involved. For instance, if one would prefer to avoid false negative, a decision rule that replaces TTF_R by TTF_C when




Fig. 6. Average error on TTF predictions for different numbers of cluster.

TTF_R is greater than TTF_C could be more appropriate. Such a rule would actually ensure that the overall approach has the same failure detection ability as a binary classification system.

Finally, it is also worth noting that the performance of the on-demand regression method could be improved by enhancing the fault detection algorithm. In this paper, we use an N-class classifier to perform this task but alternative methods could also be investigated. Actually, this should lead further research as a simple reduction in the rate of false alerts at the detection stage could have great impact on the the preciseness of final TTF predictions.

VI. CONCLUSION

This paper presented a method that we are developing to improve the preciseness of time to failure estimates for prognostic. Relying on real world data from an aerospace application, the paper describes the difficulties limiting the usefulness of regression for TTF estimation. In spite of these difficulties, the paper argues that regression can help improve TTF estimates. The method introduced to demonstrate this is named "on-demand regression". It carefully integrates classification-based prognostic, clustering, and local regression. The paper fully describes the process followed to build the various models involved and report the experimental results for the APU starter application. These results show the great potential of the approach for improving the preciseness of TTF estimates in prognostic applications. Future work includes additional validation through the application of the proposed approach to CF18 engine data and to the data provided for the PHM 2008 challenge problem.

Acknowledgments

Many people at the National Research Council of Canada have contributed to this work. Special thanks go to Bob Orchard for valuable suggestions and to Marvin Zaluski for

his help in automating the various experiments performed. We would also like to thank Air Canada for providing data used in this research.

REFERENCES

- [1] F. A. Schouten and P. Wartenhorst, "Time to failure, time to repair and availability of a two unit standby system with markovian degrading units," *Technical report BS-R9007, Center for Mathematics and Computer Science, Amsterdam, The Netherlands*, 1980.
- [2] N. Gebraeel, M. Lawley, R. Liu, and V. Parmeshwaran, "Residual life prediction from vibration-based degradation signals: A neural network approach," *IEEE Transactions on Industrial Electronics*, vol. 51, no. 3, pp. 692-700, 2004.
- [3] M. Grotke and K. Trivedi, "On a method for mending time to failure distributions," in *Proceedings of International Conference on Dependable Systems and Networks*, Los Alamitos, USA, 2005, pp. 560-569.
- [4] B. Promellec, I. Riant, and C. Tessier-Lescourret, "Precise life-time prediction using demarcation energy approximation for distributed activation energy reaction," *Journal of Physics: Condensed Matter*, vol. 18, no. 2006, pp. 2199-2216, 2006.
- [5] D. Sornette and J. V. Andersen, "Optimal prediction of time-to-failure from information revealed by damage," *Europhysics Letters*, vol. 74, no. 5, pp. 778-784, 2006.
- [6] S. Létourneau, C. Yang, C. Drummond, E. Scarlett, J. Valds, and M. Zaluski, "A domain independent data mining methodology for prognostics," in *Proceedings of the 59th Meeting of the Society for Machine Failure Prevention Technology, MFPT '05*, Virginia Beach, Virginia, USA, April 2005.
- [7] S. Létourneau, C. Yang, and Z. Liu, "On-demand regression to improve preciseness of time to failure prediction," in *Technical Report for 2007 AAAI Fall Symposium on AI for Prognostics*, Arlington, Virginia, USA, November 2007.
- [8] C. Yang and S. Létourneau, "Learning to predict train wheel failures," in *Proceedings of the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD2005)*, Chicago, USA, August 2005, pp. 516-525.
- [9] M. Hall, "Correlation-based feature selection for discrete and numeric class machine learning," in *Proceedings of the 17th International Conference on Machine Learning*, A 2000, pp. 359-366.
- [10] K. Kira and L. Rendell, "A practical approach to feature selection," in *Proceedings of the 9th International Conference on Machine Learning*, 1992, pp. 249-256.