Telemetry-mining: A Machine Learning Approach to Anomaly Detection and Fault Diagnosis for Space Systems

Takehisa Yairi, Yoshinobu Kawahara, Ryohei Fujimaki, Yuichi Sato, and Kazuo Machida
Research Center for Advanced Science and Technology
University of Tokyo
4-6-1 Komaba, Meguro-ku Tokyo, Japan
yairi@space.rcast.u-tokyo.ac.jp

Abstract

For any space mission, safety and reliability are the most important issues. To tackle this problem, we have studied anomaly detection and fault diagnosis methods for spacecraft systems based on machine learning (ML) and data mining (DM) technology. In these methods, the knowledge or model which is necessary for monitoring a spacecraft system is (semi-)automatically acquired from the spacecraft telemetry data.

In this paper, we first overview the anomaly detection / diagnosis problem in the spacecraft systems and conventional techniques such as limit-check, expert systems and model-based diagnosis. Then we explain the concept of ML/DM-based approach to this problem, and introduce several anomaly detection / diagnosis methods which have been developed by us.

1. Introduction

For any space mission, safety and reliability are the most important issues. Especially, development of advanced anomaly detection and fault diagnosis methods utilizing the latest information technologies such as artificial intelligence is essential for reliable operation of large scale complex space systems.

Conventionally, anomaly detection and diagnosis methods based on *apriori* expert knowledge and deductive reasoning process such as *expert systems* and *model-based reasoning* have been principally studied for this purpose. While these knowledge-intensive approaches have been proved to perform much better than the classical limit-checking method, it is often costly and time-consuming to prepare the knowledge-base or model-base which are required for them.

On the other hand, in recent years, inductive reasoning techniques so called data-mining (DM) or machine learning (ML) technologies have drawn much attention as alternative approaches to the anomaly detection problems in various application fields. The authors also have considered the potential value of the spacecraft telemetry data stored in ground stations, and studied DM/ML-based anomaly detection and fault diagnosis methods for spacecraft systems. These methods utilize large amount of past telemetry stored in the ground station as training data to appropriately update various parameter values contained in the diagnosis models which are given by experts, or automatically acquire the rules, patterns and models regarding the spacecraft systems. Then, they use the modified or acquired models or rules to detect and diagnose anomalies by comparing them with actual on-line telemetry. A significant advantage of this approach compared with the conventional expert systems and model-based diagnosis approach is that it does not require complete and accurate expert knowledge in the form of rule-base or model-base in advance.

In this paper, we explain the ML/DM-based approach to the anomaly detection / fault detection issues for spacecraft systems and introduce several specific methods based on this concept. Relationship with the conventional methods and other application fields is also discussed. The rest of this paper is organized as follows. In section 2, we consider how advanced anomaly detection and fault diagnosis technologies are significant for ensuring the highest level of safety and reliability which is requested by any spacecraft system, and review the three major conventional methods – limit checking, expert systems, and model-based diagnosis. In the latter part of this section, we survey the new approach to the anomaly detection problem based on machine learning / data-mining methodology, which has been actively studied in these years. In section 3, we introduce our studies on ML/DM-based anomaly detection / fault diagnosis methods for spacecraft systems. Finally, in section 4, we present some concluding remarks and future directions for further study.

2. Spacecraft and Anomaly Detection / Diagnosis Problem

2.1. Reliability of Space Systems in Operation Phase

Space systems such as artificial satellites, launchers, space stations, etc require the highest level of safety and reliability. Though this may be obvious, main reasons for it will include the followings:

- Generally, the development of a space system requires so much money and time that any fatal failure is not acceptable.
- Many space systems such as communication satellites, meteorological satellites, global positioning system are indispensable for our modern lives.
- As a space system is operating in a remote environment, it is practically impossible or very hard to repair it once a severe failure occurs.

On the other hand, it is difficult to eliminate the possibility of malfunction in space systems completely no matter how carefully they are designed and produced, because they are so-called large-scale, complex systems like airplanes, nuclear plants, and so on. In a sense, the history of space exploration and development has demonstrated indirectly that this is true.

Therefore, various efforts not only in the design, development, and test phases but also in the operation phase are required to ensure the highest level of safety and reliability. It is no doubt the development of advanced anomaly detection and fault diagnosis technologies is a major part of them. Actually, as will be mentioned below, several attempts to develop such advanced anomaly detection and fault diagnosis methods by applying various information technologies (IT) and artificial intelligence (AI) have been made.

2.2. Conventional Approaches

In this section, we review three representative conventional approaches to the anomaly detection and fault diagnosis problem in space systems. They are (i) (classical) limit checking, (ii) expert systems, and (iii) model-based diagnosis.

2.2.1 Limit Checking

Limit checking (also known as *limit check* and *limit sensing*) is the most fundamental and still the most widely used anomaly detection technique for spacecraft systems. Its basic function is to monitor whether various sensor values such as bus currency, voltage, angular velocity, temperature, and so on are within pre-determined ranges which are specified by upper and lower limits, and issue a warning if any of them is violated. Usually, limit checking is performed on a selected number of numeric telemetry series, rather than on all (up to several hundred) series.

While the limit values are initially decided by designers and engineers, they are often adjusted by operators' hands on orbit. Another popular extension to the basic limit checking is to use multiple sets of upper and lower limits on a variable depending upon the operational mode of the system, the level of urgency, and so on.

There seem to be two major reasons why the limit checking is still playing a central role in the space system operation. One is its *simplicity*, which means it is easy for human operators to implement a system, apply it to the spacecraft, and understand the detection results. The other is its *achievement* that a number of anomalies or malfunctions have been actually detected by this method.

However, the simplicity of limit checking means there are inevitable limitations at the same time. First, there exist a number of anomalies or their symptoms which cannot be detected just by monitoring whether sensor values are between upper and lower limits. In other words, some class of anomalies occur without violating the limits on the variables. This will be reasonable when considering the lack of representational power of limit checking. Another significant problem with limit checking is that it is laborious and costly for engineers and operators to pre-determine or adjust the limit values appropriately taking various conditions into account. If the limit values are inappropriate, it will either fail to detect any anomaly at all, or issue a great number of false alarms, which makes the operators insensitive to real anomalies.

To overcome these difficulties of the classical limit checking, a variety of machine learning methods (especially, non-linear regression methods) have been utilized recently. A famous example of them is ELMER which was developed by DeCoste[3], in which appropriate sets of limit values are automatically learned from past telemetry data. In section 3, we will describe another adaptive limit checking technique based on relevance vector regression method.

2.2.2 Expert Systems

The emergence and development of expert systems (a.k.a., rule-based systems, and knowledge-based systems) would be one of the most significant achievements of artificial intelligence (AI) in its early days. In the research of diagnosis systems for spacecraft, a number of expert systems were studied and developed mainly in 80s and early

90s[17, 2, 1, 15]. The principle common to all of them is to infer causes of anomalies using sets of rules that specify relationships between symptoms and failures. The set of rules is called knowledge-base and carefully prepared by domain experts.

While some of the above mentioned expert systems are based on nothing more than a table lookup[2], the other systems are employing more sophisticated mechanisms such as certainty factors (CF)[1, 15] and frame-based knowledge representation[17]. A noteworthy example of an expert system for spacecraft operation is ISACS-DOC (Intelligent Satellite Control Software DOCtor)[13]developed by ISAS and JAXA, which has been employed in a series of spacecraft for deep space exploration mission such as GEOTAIL, NOZOMI and HAYABUSA for more than ten years.

Obviously, expert systems outperform the limit checking in representational power and diagnosis capability. They have, however, their own problems. For example, they are unable to deal with "unknown" anomalies, because it is required that all possible failures and symptoms are enumerated and the relationships between them are described in advance. Besides, they have difficulties in keeping the consistency of the knowledge-base when some design changes take places in the middle of system development, and in determining appropriate CF values when the number of possible symptoms increases.

2.2.3 Model-based Diagnosis

Model-based anomaly detection and diagnosis methods for spacecraft systems have been also actively studied[19, 8, 5, 4, 14], along with expert systems which were mentioned above. Their fundamental principle is to detect anomalies and reason about the causes by comparing simulation results obtained from some computational models with actual behavior of the target systems. A pioneering work in this approach is *Livingstone*[19], which was originally developed as a part of *ReactiveAgent* and tested on NASA's DeepSpaceOne (DS-1) mission.

Most of the early model-based diagnosis systems including Livingstone are based on *qualitative* models and reasoning[5, 4, 14]. In recent years, however, an alternative approach based on hybrid probabilistic models and reasoning has drawn much attention[8], in which a spacecraft system is represented by a stochastic model containing both discrete and continuous variables and its behavior is simulated by a Monte Carlo simulation based state estimation method called *particle filter*. A major factor in its popularity will be the dramatic advances in computer performance, as characterized by Moore's law.

Our anomaly detection and diagnosis method based on Dynamic Bayesian Networks (DBNs) which is explained in section 3 can also be classified into the latter category, though it has an advanced learning capability of estimating unknown parameters in the model from past data.

2.3. Anomaly Detection and Diagnosis Based on Machine Learning / Datamining Technologies

The conventional anomaly detection and fault diagnosis techniques reviewed above are driven by *apriori* expert knowledge such as sets of limit values, rule-bases, and model-bases. In other words, they are based on *deductive* reasoning processes.

On the other hand, an alternative approach to the anomaly detection problem based on machine learning (ML) or data mining (DM) has drawn much attention recently in a variety of areas. The basic idea of this approach is to acquire the system behavior models necessary for anomaly detection and diagnosis from past data (semi-)automatically, rather than from human experts. There are two main factors behind this trend:

- It is very laborious and costly to prepare and maintain by hand the complete and accurate expert knowledge such as rule-base and model-base, which are indispensable for the deductive approach.
- It becomes relatively easy to obtain, store and process a vast amount of sensor data, through the recent advances in technology.

In the research community of machine learning and datamining, a number of researchers are increasingly interested in applying the ML/DM technology to a variety of anomaly detection problems in the real world. For example, applications to detection of emerging disease outbreak[11], failure detection of computer networks[21], prediction of train wheel failures[22], fraud detection in stock market[12], and our work on spacecraft anomaly detection[7] have been reported in *The Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-05)*. In this conference, a special workshop named "Data Mining Methods for Anomaly Detection" was also held.

In these studies, not only classical machine learning methods such as supervised learning and clustering but also other closely related methods such as outlier detection, change-point detection, kernel theory and so on are utilized.

3. Proposed Methods

3.1. Overview

In this section, we introduce four ML/DM-based anomaly detection and diagnosis methods for spacecraft systems we have developed. Among them, the first two studies are considered as "data-driven" anomaly detection methods which require almost no apriori expert knowledge. The third one, on the other hand, is regarded as "hybrid" approach in which probabilistic inference (deductive process) is integrated with statistical estimation of unknown parameters using past data (inductive process). The last one is a visualization method of telemetry data (rather than an autonomous detection/diagnosis method) that is intended to support human operators in monitoring the large amount of telemetry data.

3.2. Data-driven Approach 1: Adaptive Limit Checking with Relevance Vector Regression

As was mentioned in section 2.2.1, limit checking is still utilized as the primary method of monitoring and fault detection in the operation of spacecraft systems because of its simplicity and reasonable performance.

Taking this fact into consideration, we developed an advanced limit checking method with the capability of learning appropriate limit values from past telemetry data[6]. The learning is performed by the *Relevance Vector Machine* (RVM). RVM (also known as *Sparse Bayesian Learning*) is a kernel-based regression and classification method which was originally introduced by Tipping[18]. In our method, RVM is used to learn a prediction model of upper and lower limit values for each target variable (i.e., telemetry series) from past normal data. Then, obtained models are used in the subsequent operations to predict appropriate ranges of the target variables and detect anomalies online. Although another kernel-based classification and regression method called *Support Vector Machine* (SVM) is more widely known than RVM at the moment, we consider RVM is more suitable for limit checking, because it is able to predict the probability density distribution or range of a target variable whereas SVM gives only a point estimation.

Figure1(upper) shows actual pitch rate of *ETS-VII* which was a testing satellite of JAXA (solid line) and upper and lower limits (break lines) predicted by the proposed method. The the change of likelihood is also shown in the bottom. The predicted range is much tighter than ordinary constant limit values, which means the false negative rate can be dramatically reduced. We reported the results of another experiment applying this method to a rendezvous simulation of orbital transfer vehicle with ISS. In this experiment, it was demonstrated that this method is able to detect different types of thruster failures quickly.

On the automated learning of limit values, several similar studies have been reported. ELMER[3] learns the upper and lower limits of a target variable by linear functions with respect to heuristic features. In our previous work[20], we

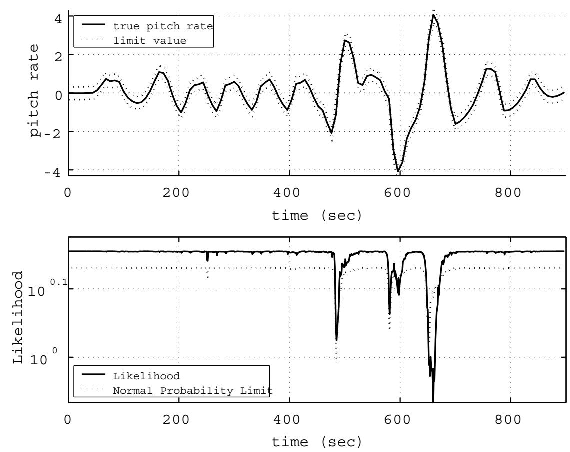

Figure 1. Pitch Rate of ETS-VII and Limit Values Predicted By Proposed Method (Top : Actual Pitch Rate, Bottom : Likelihood)

proposed another learning method of limit values using regression tree, in which an appropriate range of a target variable is predicted with respect to other symbolic variables (or status variables).

3.3. Data-driven Approach 2: Anomaly Detection from Telemetry with Kernel Principal Component Analysis

In this study, we developed another data-driven anomaly detection method which requires almost no *apriori* expert knowledge. This method is based on *Kernel Principal Component Analysis* (Kernel PCA)[16]. While the detail of this method is explained in [7], we present the basic idea here.

First, if a spacecraft system is made stationary by a proper feedback control law, we can assume that some static cause-effect relationships exist among the observation variables consisting the telemetry data. For example, when we consider only the orbit and attitude motion, the static cause-effect relationships should be represented by a set of equations of motion with fixed parameters. This implies that the behavior of a (controlled) spacecraft system is constrained on a intrinsic low-dimensional manifold, no matter how high the dimension of observation space is, or no matter how many series the telemetry contains. Now consider some significant anomaly occurs to the system, then we can expect the shape of manifold to be also changed, because the cause-effect relationships are affected by the anomaly. The key idea in the proposed method is to map non-linearly the telemetry (or observation data) of each time segment into a high dimensional feature space by the polynomial kernel, and detect the system anomaly or transformation of

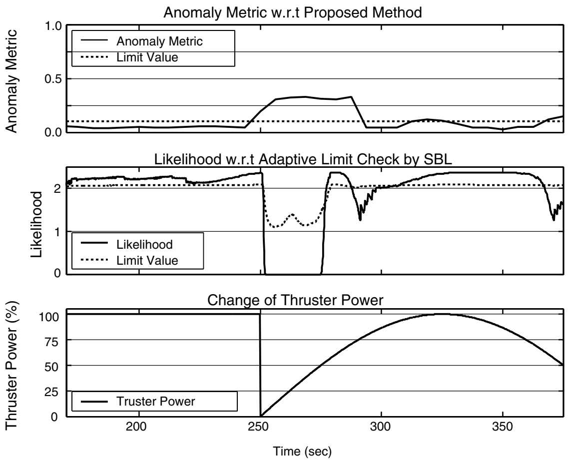

Figure 2. Detection Results of Thruster Anomaly in Rendezvous Simulation (Top : Result with Kernel PCA, Middle : Result with Adaptive Limit Checking, Bottom : Ratio of Actual Thruster Output to Expected Output)

the manifold as a change in *principal directions* of the data.

We tested this method on the data set of rendezvous simulation of orbital transfer vehicle mentioned above, and confirmed that it was able to detect several different kinds of anomalies in the thrusters. Figure 2 shows an example of the results. The bottom figure shows the ratio of actual output of *Thruster 9* to the expected output. (Note that the actual outputs of thrusters are *unobservable*.) As can be seen from this figure, a fault occurs at time 250[sec] and the thruster output goes down to zero sharply, then repeats rises and falls. The top figure shows the transition of *anomaly metric* used in the proposed method. This anomaly metric is defined as the change of directions of principal axes in the kernel feature space. We can see that this method detects the fault when the anomaly metric exceeds a threshold. The threshold is semi-automatically determined by the statistical criterion, assuming that the principal directions in normal operation follow the *von Mises-Fisher distribution*. For comparison, detection result with the adaptive limit checking by RVM mentioned previously is also shown in the middle figure. Though both methods successfully detected the anomaly in this case, there were several anomaly patterns which the kernel PCA method was successful with but the adaptive limit checking was not.

3.4. Model-Data Hybrid Approach : Anomaly Detection and Diagnosis with Dynamic Bayesian Networks

As discussed in section 2.2.3, a common major problem in the conventional model-based diagnosis methods is that the system models, whether qualitative or quantitative, are required to be complete and accurate in advance.

To cope with this problem, we have developed a *hybrid* approach which combines a deductive anomaly detection / diagnosis process based on probabilistic inference and an inductive model estimation process based on learning from past telemetry data[9]. More specifically, the behavior of spacecraft is modeled by a kind of probabilistic graphical model called *Dynamic Bayesian Network*(DBN), whereas unknown parameters and structures contained in the model are learned from the telemetry in normal operation.

The Dynamic Bayesian Network (DBN)[10] is an extension of the Bayesian Network (BN) intended to model a variety of dynamical systems. The DBN can be regarded as a generalized state space model that contains the Kalman filter (KF) and Hidden Markov Model (HMM) as special cases. It is especially suitable for modeling "hybrid systems" which contain both continuous variables (e.g.,velocity, attitude angle, etc.) and discrete variables (e.g.,operation modes, statuses of instruments, etc.) Besides, efficient approximation algorithms for the inference on DBNs such as junction-tree algorithm, Rao-Blackwellized Particle Filter have been developed so far. In our method, uncertain parameters and structures in the DBN are learned from past telemetry data in normal operation by a modified version of EM (Expectation-Maximization) algorithm.

We applied this method to the data of rendezvous simulation mentioned before, and examined its validity. Figure 3 shows a DBN which models the orbit and attitude motions of the spacecraft. Shaded nodes in this network represent *unobservable* variables. The DBN performs the anomaly detection and diagnosis by sequentially estimating the unobservable variables from observable variables (represented as white nodes) In this experiment, too, we assumed several different patterns of anomalies in the thrusters of vehicle. For example, in *Case 1*(Figure 4) which is a simplest pattern, the ratio of actual *Thruster 4* output to expected output begins to go down at time 250[sec] then falls to 0 after 60 seconds. In *Case 2*(Figure 5), a more complicated pattern is assumed in which the output ratio of *Truster 9* abruptly falls to zero at time 250[sec] and then repeatedly rises and falls with an interval of 60 seconds. Figures 6 and 7 show the outputs of all thrusters estimated by the DBN in the two cases. Again, please note that they are not observable. As can be seen from these figures, anomalies of the right thrusters are successfully detected, though estimated

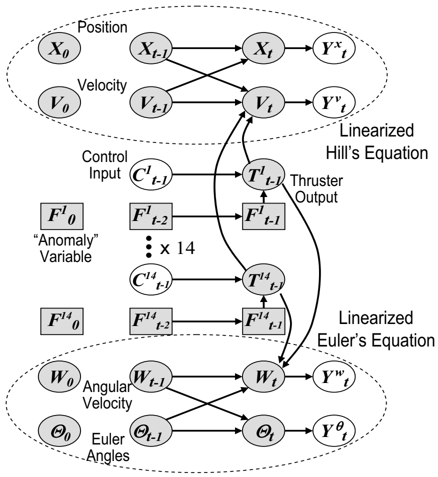

Figure 3. Dynamic Bayesian Network for Orbit and Attitude Motion of Transfer Vehicle

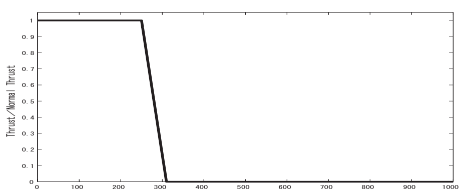

Figure 4. Case 1 : Ratio of *Thruster 4* **output to expected output**

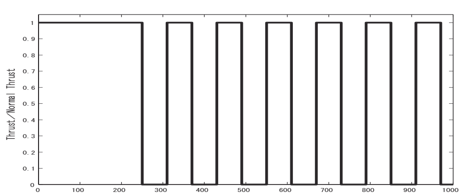

Figure 5. Case 2 : Ratio of *Thruster 9* **output to expected output**

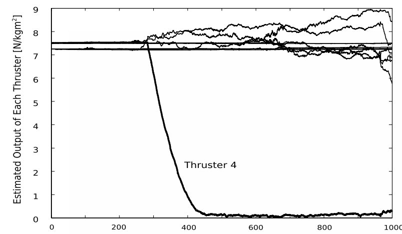

Figure 6. Estimated Outputs of All Thrusters By DBN (Case 1)

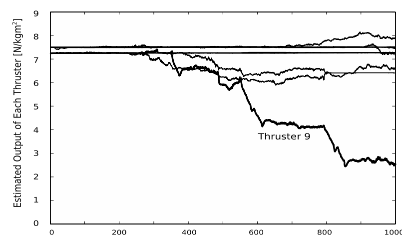

Figure 7. Estimated Outputs of All Thrusters By DBN (Case 2)

outputs of other thrusters are slightly affected in the latter case.

3.5. Telemetry Visualization Based on Change-point Detection

In the three studies above, our purpose of using various ML/DM techniques was to develop telemetry monitoring systems that directly detect anomalies in the spacecraft systems. Another promising application of the ML/DM technologies is to provide the operators with useful information by summarizing and visualizing the large amount of data, and to assist them in understanding the health status of the spacecraft and discovering symptoms of anomalies. In the last case study, we introduce a telemetry visualization method which extracts only important patterns in the high-dimensional time-series data. Of course, a specific definition of "important patterns" depends on a number of factors including spacecraft systems, operation modes, skill and knowledge of operators, and so on. In general, however, it is reasonable to assume "change points" where the characteristics of the time-series change drastically contain important information, because they are likely to correspond to some actual changes or events occurred in the system. Based on this idea, the proposed visualization method first extracts all change-points in each series of telemetry, then displays the results for all series after sorting them so that highly associated series are placed in neighborhood. We utilized the locally stationary autoregressive model and spectral ordering method for the detection of change-points and ordering of the series, respectively. We applied this method to the telemetry of a commercial communication satellite owned by *Space Communication Corporation* (SCC). Figure 8 shows the telemetry for a day *as it is*. On the other hand, Figure 9 shows the result of change-points detection and ordering of series. In addition, Figure 10 shows the strengths of associations among the series in a matrix form. As a result of ordering, there appear several partially overlapping blocks in this figure. We consider that this means there are several process groups which are loosely coupled. The block pattern like Figure 10 varies day by day. We are now investigating how effective this visualization method is.

4. Conclusion

In this paper, we discussed the anomaly detection and diagnosis problem in the space system operation, reviewed conventional approaches to it, and introduced a new approach based on machine learning and datamining technologies. The key idea of this approach is to maximally utilize past telemetry data for acquiring system behavior models that can be used for anomaly detection, compensating for the lack of *apriori* expert knowledge, and providing useful information for operators' decision making.

While the methods described in section 3 have shown promising performance so far, there are still problems for practical use. For example, a significant issue is how we can guarantee the reliability and generality of the acquired information from data. Effective and intuitive ways of presenting outputs from the detection / diagnosis systems must be also considered, because a ML/DM technique is often used as a "black-box". Though there is no absolute answer to these questions, we believe that development of systems promoting interaction between experts (or their knowledge) and information derived from data is a right direction.

Acknowledgment

Telemetry data was provided by JAXA (Japan Aerospace Exploration Agency) and SCC (Space Communication Corporation). We really appreciate their cooperation. This research was partially supported by the Ministry of Education, Science, Sports and Culture, Grant-in-Aid for Young Scientists(B), 14750717, 16760641.

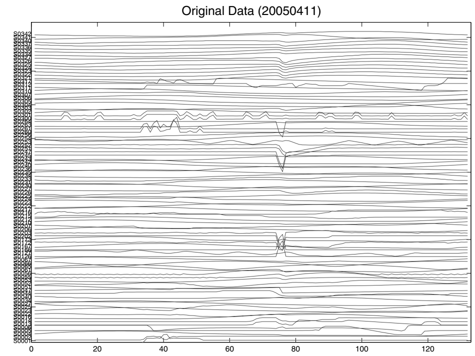

Figure 8. Raw Telemetry Data for One Day

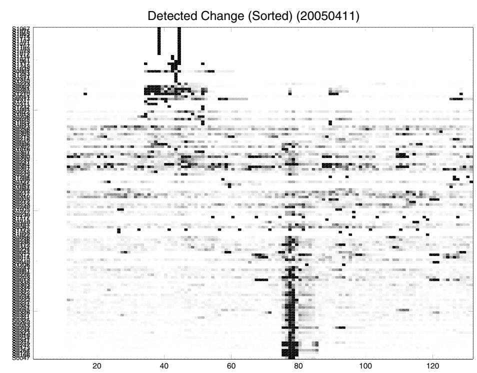

Figure 9. Visualization Result Based on Change-Point Detection for The Data of Figure 8

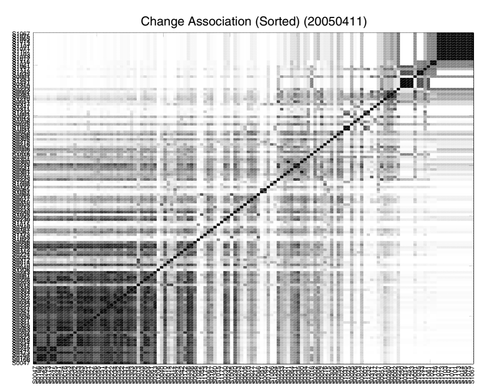

Figure 10. Change Association Among Telemetry Series

References

[1] C. Chang, W.Nallo, R.Rastogi, D.Beugless, F.Mickey, and A.Shoop. Satellite diagnostic system: An expert system for intelsat satellite operations. In *Proc. IVth European Aerospace Conference (EAC)*, pages 321–327, 1992.

[2] F. Ciceri and L. Marradi. Event diagnosis and recovery in real-time on-board autonomous mission control. In *Proc. of the 1st Ada-Europe Symposium*, 1994.

[3] D. DeCoste. Automated learning and monitoring of limit functions. In *Proc. of International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, pages 287–292, 1997.

[4] A. Fijany, F. Vatan, A. Barrett, M. James, and R. Mackey. An advanced model-based diagnosis engine. In *Proc. of International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*, 2003.

[5] A. Finzi, M. Lavagna, and G.Sangiovanni. Fuzzy inductive reasoning and possibilistic logic for space systems failure smart detection and identification. In *Proc. of International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*, 2003.

[6] R. Fujimaki, T. Yairi, and K. Machida. An anomaly detection method for spacecraft using relevance vector learning. In *Proc. of The 9th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, pages 785–790, 2005.

[7] R. Fujimaki, T. Yairi, and K. Machida. An approach to spacecraft anomaly detection problem using kernel feature space. In *Proc. of The Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 401–410, 2005.

[8] F. Hutter and R. Dearden. Efficient on-line fault diagnosis for non-linear systems. In *Proc. of Seventh International Symposium on Artificial Intelligence and Robotics in Space (i-SAIRAS)*, 2003.

[9] Y. Kawahara, T. Yairi, and K. Machida. Diagnosis method for spacecraft using dynamic bayesian networks. In *Proc. of 8th International symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2005.

[10] K. P. Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, University of California, 2002.

[11] D. Neill, A. Moore, M. Sabhnani, and K. Daniel. Detection of emerging space-time clusters. In *Proc. of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 218–227, 2005.

[12] J. Neville, O. Simsek, D. Jensen, J. Komoroske, K. Palmer, and H. Goldberg. Using relational knowledge discovery to prevent securities fraud. In *Proc. of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 449–458, 2005.

[13] N. Nishigori, M. Hashimoto, A. Choki, and M. Mizutani. Fully automatic and operator-less anomaly detecting ground support system for mars probe 'nozomi'. In *Proc. of the 6th International Symposium on Artificial Intelligence and Robotics and Automation in Space (I-SAIRAS)*, 2001.

[14] P. I. Robinson, M. Shirley, D. Fletcher, R. Alena, D. Duncavage, and C. Lee. Applying model-based reasoning to the fdir of the command and data handling subsystem of the international space station. In *Proc. of International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*, 2003.

[15] M. Rolincik, M. Lauriente, H. C. Koons, and D. Gorney. An expert system for diagnosing environmentally induced spacecraft anomalies. In *Proc. of 5th Annual Space Operations and Applications Research Symposium*, pages 36–44, 1992.

[16] B. Schölkopf, A. Smola, and K.-R. Muller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10:1299–1319, 1998.

[17] D. P. Tallo, J. Durkin, and E. J. Petrik. Intelligent fault isolation and diagnosis for communication satellite systems. *Telematics and Informatics*, 9(3-4):173–190, 1992.

[18] M. Tipping. Sparse bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1:211–244, 2001.

[19] B. C. Williams and P. P. Nayak. A model-based approach to reactive self-configuring systems. In *Proc. of the Thirteenth National Conference on Artificial Intelligence*, pages 971–978, 1996.

[20] T. Yairi, M. Nakatsugawa, K. Hori, S. Nakasuka, and K. Machida. Adaptive limit checking for spacecraft telemetry data using regression tree learning. In *Proc. of International Conference on Systems, Man and Cybernetics (IEEE SMC)*, pages 5130–5135, 2004.

[21] K. Yamanishi and Y. Maruyama. Dynamic syslog mining for network failure monitoring. In *Proc. of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 499–508, 2005.

[22] C. Yang and S. Letourneau. Learning to predict train wheel failures. In *Proc. of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 516–525, 2005.