An Application of Data Driven Anomaly Identification to Spacecraft Telemetry Data

Gautam Biswas[1], Hamed Khorasgani [2], Gerald Stanje[3], Abhishek Dubey[4], Somnath Deb[5], and Sudipto Ghoshal[6]

[1,2,3,4] *Inst. of Software-integrated Systems, Vanderbilt Univ., USA,*
gautam.biswas@vanderbilt.edu
hamed.g.khorasgani@vanderbilt.edu
gerald.stanje@vanderbilt.edu
abhishek.dubey@Vanderbilt.edu

[5,6] *Qualtech Systems, Inc. (QSI), USA,*
deb@teamqsi.com
sudipto@teamqsi.com

ABSTRACT

In this paper, we propose a mixed method for analyzing telemetry data from a robotic space mission. The idea is to first apply unsupervised learning methods to the telemetry data divided into temporal segments. The large clusters that ensue typically represent the nominal operations of the spacecraft and are not of interest from an anomaly detection viewpoint. However, the smaller clusters and outliers that result from this analysis may represent specialized modes of operation, e.g., conduct of a specialized experiment on board the spacecraft, or they may represent true anomalous or unexpected behaviors. To differentiate between specialized modes and anomalies, we employ a supervised method of consulting human mission experts in the approach presented in this paper. Our longer term goal is to develop more automated methods for detecting anomalies in time series data, and once anomalies are identified, use feature selection methods to build online detectors that can be used in future missions, thus contributing to making operations more effective and improving overall safety of the mission.

1. INTRODUCTION

As engineered systems have become more complex, and include a range of operations that vary widely, self-monitoring, self-diagnosis, and adaptability to maintain operability and safety have become focus areas for research and development. Typical goals of such self-diagnosis approaches are the detection and isolation of faults and anomalies, identifying and analyzing the effects of degradation and wear, and providing fault-tolerant and fault-adaptive control (Blanke & Schröder, 2006; Chen & Patton, 2012; Isermann, 2005; Ji, Zhang, Biswas, & Sarkar, 2003; Noura, Theilliol, Ponsart, & Chamseddine, 2009). The majority of projects dealing with monitoring and diagnosis applications rely on models created using physical principles or by human experts. However, these models are not always available, and are often incomplete, and sometimes erroneous. Moreover, it is hard to maintain the effectiveness of these models during a systems life-cycle.

More recently, promising data-driven alternatives that exploit the large amounts of operational data collected from these systems are being employed to better understand system behaviors and anomalies during system operations (Qin, 2012; Yin, Ding, Xie, & Luo, 2014). In data-driven approaches, monitoring and diagnosis knowledge can be learned by observing and analyzing system behavior (Mack, Biswas, Koutsoukos, & Mylaraswamy, 2016, in press). This large amount of data collected using new, more robust sensors and sensor networks, can be exploited in a reliable manner for the purpose of detecting and analyzing anomalous situations and faults in these large and complex systems. The vision is developing Cyber Physical Systems (CPSs) (Lee, 2008; Marwedel, 2010; Niggemann et al., 2015) that can observe their own behavior, recognize unusual situations during operations, and inform operators, who use this this information to modify system operations, or plan for repair and maintenance. Furthermore, system's experts and engineers can use the information gleaned from this data to update operations procedures and even redesign the system.

In this paper, we take on the challenges of developing an anomaly detection scheme for analyzing telemetry data gen-

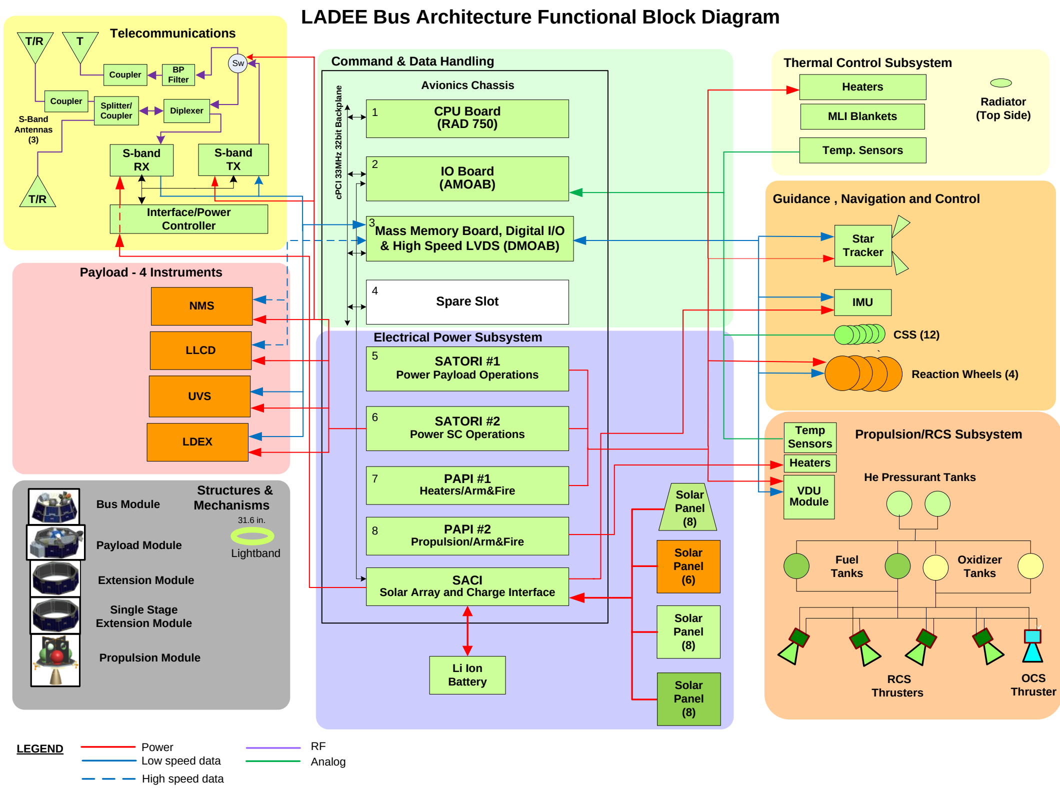

Figure 1. Functional block diagram of the LADEE spacecraft (image credit: NASA/ARC) `https://directory.eoportal.org/web/eoportal/satellite-missions/l/ladee`

erated by NASA's Lunar Atmosphere and Dust Environment Explorer (LADEE) spacecraft[1]. LADEE was a robotic mission that orbited the moon to gather detailed information about the structure and composition of the thin lunar atmosphere, and determine whether dust is lofted into the lunar sky. The LADEE spacecraft's modular common spacecraft bus, or body, designed, developed, and operated by NASA's Ames Research Center, innovated away from custom designs and transitioned toward multi-use designs and assembly-line production, which could drastically reduce the cost of spacecraft development (Hine, Spremo, Turner, & Caffrey, 2010). The LADEE system block diagram, shown in Figure 1, consists of five primary subsystems: (1) the Integrated Avionics system, (2) the Propulsion system, (3) the Attitude Control system (ACS)), and the Electrical Power Subsystem (EPS). Our focus is to develop a general data-driven monitoring approach for telemetry (i.e., streaming time series) data for purposes of health monitoring, which includes fault and anomaly detection, prognosis, and performance analysis of the monitored system.

Our specific focus in this paper is on developing a general unsupervised method for data-driven anomaly detection in complex systems. The rest of this paper is organized as follows. Section 2 defines anomaly detection problem for spacecraft longterm missions. Section 3 formally describes our approach to anomaly detection, and lays out the description of our methodology. Section 4 shows an application of our methodology to telemetry data from the EPS of the LADEE spacecraft. Using examples, we illustrate the intertwining of the mode and anomaly detection problem. Finally, Section 5 presents a discussion and conclusion based on the results or our case study, and briefly discusses how we will extend this approach in future work.

2. PROBLEM FORMULATION

A spacecraft is a mixed discrete-continuous system (hybrid system). For example, a reaction wheel which is a continuous system with an electric motor to rotate the spacecraft around its center of mass is controlled by a discrete processor. The discrete states of the spacecraft are called modes and the switches or external events define the discrete dynamics of the system. The behavior of the spacecraft depends on the system mode: each mode change can be: (1) *controlled,* i.e., they are initiated by a set of discrete switches by the system controller or a human operator; and (2) *autonomous,* i.e., they are induced by external events (e.g., when the spacecraft moves from region where it receives sunlight to one where it

[1] see `https://www.nasa.gov/mission_pages/ladee/main/index.html`

is dark) or they may be caused by a change internal to the system (e.g., a battery may be completely drained, or a fuel tank overflows). In both cases, the mode change changes the system model, and, therefore, the system dynamics. In this work, we assume the spacecraft is a hybrid system that is modeled as a hybrid automaton (Henzinger, 2000) with the following definition.

Definition 1 (Hybrid Automaton) *A hybrid automaton H is defined by a 5-tuple, $H = \{Q, R^n, f, \varphi, \rho\}$, where Q represents a set of discrete states; R^n represents the space of continuous behaviors; $f : Q \times R^n \to R^n$ represents the vector field that defines continuous behaviors in a mode; $\varphi : Q \times R^n \to Q$ represents the discrete transition function, and $\rho : Q \times R^n \to R^n$ represents the reset map.*

Generally, before a spacecraft launch, scientists and system developers plan the entire mission. This is called space mission planning. Some components of the plan may be uploaded to controllers on the spacecraft, others are manually commanded or uploaded as a mission progresses by the mission controllers. Considering the spacecraft as a hybrid automata, we can define the expected mission trace as a sequences of time transitions that the spacecraft is designed to follow during the mission.

Definition 2 (Expected Mission Trace) *The expected mission trace, MT consists of an initial mode $q_{start} \in Q$, a finite set of mode transitions, $T_r : Q \to Q$, which may be controlled or autonomous; and a sequence of intermediate modes $Q_m = \{q_i, q_j, \cdots q_k\}$, where each intermediate mode $q_i \in Q_m$ has a start time t_{si} and an end time t_{ei}. The final mode that ends a mission is q_{end}.*

A mode trace is considered normal if and only if every mode the system enters is expected according the hybrid automata model, H.

Definition 3 (System normal operating mode) *A mode (q_i, t_{si}, t_{ei}) is a normal operating mode if and only if $(q_i, t_{si}, t_{ei}) \in MT$.*

Because of unpredicted events and possible faults and degradation in the system, it is possible that the system starts behaving in an unexpected way for an interval of time during the mission. Since this behavior cannot be explained or justified using the mission plan or by a mode transition in the definition of H, then the system is considered to be in an abnormal mode.

Definition 4 (System abnormal operating mode) *A mode (q_i, t_{si}, t_{se}) is an abnormal operating mode if and only if $(q_i, t_{si}, t_{se}) \notin MT$, or there is no defined transition from a mode $q_k \in Q$ to mode q_i.*

The objective here is to develop a method to detect the abnormal behavior modes during spacecraft operation. This helps the system developers to study the abnormal modes and analyzes their root causes to prevent them from occurring in future missions. A common approach to detect normal and abnormal operating modes in hybrid systems is using the state estimation approaches (Hanlon & Maybeck, 2000; Blom & Bar-Shalom, 1988). Estimation approaches typically use multiple-model-estimation schemes to track state estimates over time, and therefore, require at least as many filters as there are modes in the system. In the LADEE spacecraft, we have at least 67 switches in the power system, which means we have at least 2^{67} modes. A model-based approach to designing 2^{67} filters is unrealistic and impractical – most of these modes are not likely to occur in any spacecraft mission. A more realistic approach would be to build detectors corresponding to the system Mission Trace, but that would miss unanticipated anomalous behaviors in the system.

3. DATA DRIVEN ANOMALY IDENTIFICATION

The objective here is to develop a method to detect abnormal behaviors that may have occurred during the spacecraft operations for a long-term mission. Since the possible discrepancies, faults, or errors that may occur are unknown before the mission takes place, we have to develop approaches that can discover them by analyzing mission telemetry data. If this can be accomplished, system designers and mission specialists can perform detailed studies in the time intervals when the abnormal behaviors occur. This will help them identify the root causes, which in turn will influence the design of future spacecraft to avoid such anomalies. Alternately, monitors that can detect such anomalies in an automated fashion can be designed that allows mission specialists to come up with corrective actions or change the mission plan and avoid adverse incidents.

In our work, we have applied unsupervised learning techniques to find groupings in a large database of time series data. Our approach is to divide the time series representing the entire mission trajectory into segments, and each segment represents an object of interest on the mission time line. The selected size of each time interval (time window) is a trade-off between the number of objects created from the time series, and having enough data per window to detect and characterize abnormal versus nominal behavior. Typically, most of the objects will represent nominal operations of the spacecraft during the mission, but a small subset of the time segments (objects) may represent anomalous or faulty behaviors. Since the anomalous behaviors or faults are not known beforehand, we apply a clustering algorithm to group the objects. Our hypotheses is that the larger groups of clusters will represent nominal operations, whereas outliers and smaller groups may represent anomalous situations. In general, researchers have developed classifier or supervised methods for characterizing known faults and semi-supervised and unsupervised methods for discovering and characterizing unknown faults

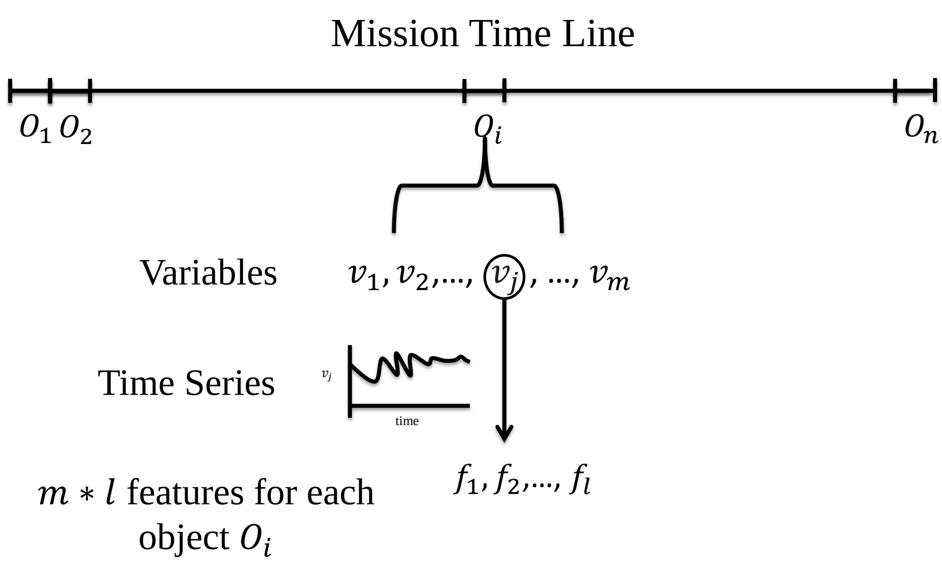

Figure 2. Data preprocessing and feature extraction.

and anomalies. In this work, we propose a mixed method for anomaly detection in a robotic space mission.

In more detail, we have developed a multi-step unsupervised learning method to distinguish the outliers which include specialized modes of operation and abnormal behaviors from the normal operation. Then, we employ a supervised method of consulting human mission experts to differentiate between specialized modes and anomalies. First, we select a set of objects from a curated data set, where each object is represented by a set of features, typically segments of a time series signal. The next step is to convert each feature signal into a set of feature values that make the data amenable to traditional numeric clustering methods. To do this we apply a wavelet transform to each time series segment, and represent that waveform in terms of coefficients that define the wavelet transform. In addition to generating numeric feature values, this approach also serves as an approach for compressing and smoothing a signal. Then a hierarchical clustering approach is used to generate clusters from the extracted features. The outputs of the clustering algorithm are preliminary groups of time intervals for further consideration.

The input to the process is operational data, which is the telemetry data transferred to the earth from the spacecraft during the mission. This dataset contains many variables from different subsystems of the spacecraft which are measured or computed with different sample rates. In general, this dataset is extremely large. To reduce computational complexity, we use the following approaches; 1) data reduction, 2) feature reduction. The data reduction here means selection of relevant time series waveforms for object definition to describe anomalies. $V = \{v_1, v_2, ..., v_m\}$ represents the selected variables. Figure 2 illustrates the data preprocessing procedure. As discussed earlier, we divide the mission into segments, and each segment represents an object of interest on the mission time line. k is the number of samples in each time interval (time window). The time windows are the objects, $O = \{O_1, O_2, ..., O_n\}$. In fact, each object $O_i \in O$ is a time series which contains k samples of each select variable.

In this work, we use the discrete Haar wavelet transform to extract the scaling coefficients which correlate to the low frequency of the time series signals. We consider the first l coefficients of the wavelet transform for each variable (See Figure 2). The set of coefficients for object O_i, is presented with $f_i \in R^{m*l}$ in Figure 3, where $m * l$ represents the number of wavelet coefficients derived from the set of time series signal segments defining the object. The algorithm then uses the generated features and a weighted Euclidean distance measure to build the dissimilarity matrix between every pair of objects, i.e., $D_{nn} = dist(O_i, O_j), 1 \leq i, j \leq n$. We then run a UPGMA (Unweighted Pair Group Method with Arithmetic Mean), agglomerative (bottom-up) hierarchical clustering algorithm to generate a dendrogram that represents the order in which the objects group into clusters. The advantage of this approach is that the number of clusters does not have to be pre-determined. We then use heuristic methods to cut the dendrogram at a level (i.e., a distance measure) that provides a distinct grouping of clusters. In other words moving the distance level at which the dendrogram is cut by small levels will not change the number of clusters that are generated.

The next step is the interpretation process and the identification of anomalous groups. We assume the results of the clustering produce one or more large clusters where most of time intervals reside, with the implication that these large clusters represent mostly nominal behaviors because overall the mis-

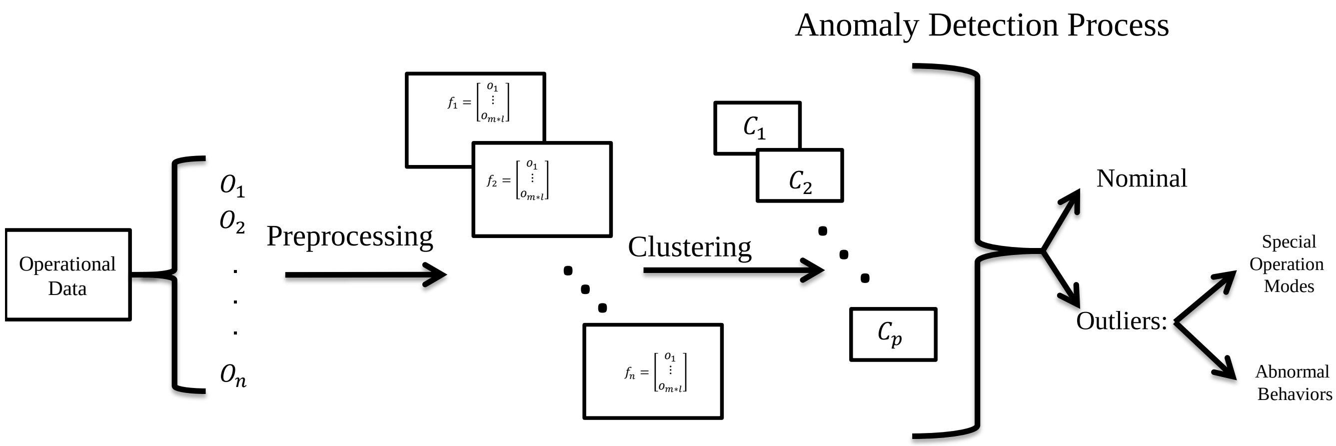

Figure 3. Our proposed approach for anomaly detection.

sion was a success. We label these clusters as corresponding to nominal operations of the spacecraft. Since a spacecraft is on a complex mission that involves multiple maneuvers and also a number of scientific experiments, we then time focus on the smaller clusters, and start with the hypotheses that each one of the cluster represent: (i) special operating modes , or (ii) anomalous or faulty behavior. We identify special operating modes by looking at the discrete waveforms that are available, such as switches that indicate a particular load representing a science experiment or a spacecraft maneuver was turned on. This helps us assign non anomalous labels to a number of the smaller clusters. For the rest, we seek expert input.

A typical approach here would be to compare the features of the possible anomalous groups against the labeled nominal groups. We can do this systematically by running a classifier algorithm and determining the features that best separate these groups. These features then define the anomalous patterns that do not confirm to expected behavior. In this work, to identify anomalies we perform the following analysis on the small clusters (outliers): 1) we study the features that sufficiently distinguish the outliers from the nominal data, 2) we study the switches that went on or off during the outlier intervals, 3) we talk with the experts to confirm our hypothesis. With possible anomalous groups and a nominal base of clusters, a feature selection algorithm is applied to identify the relevant features that differentiate each anomalous group. A switch can explain special operating conditions due to the mode transitions in the system. An expert will use the distinguishing features and binary transitions during the outliers to characterize the anomalies and their level of failure. Coupled with the nominal sets, these groups can be used to produce new models of anomaly detection.

4. CASE STUDY: LADEE EPS

The data set was a collection of time series data recorded over a 1 year long lunar mission. We extracted only voltage and current features (total 34) from this dataset. We believe this provided sufficient information for a mode and anomaly identification. This paper explores a data-driven mode and anomaly detection approach where the algorithm has access to the electric power subsystem data of a robotic satellite. The dataset has the following characteristics:

1. There are 265 time series variables in the electrical subsystem dataset.
2. There are 17 time series binary variables in zero or one format among the variables.
3. There are 50 time series binary variables in on or off format among the variables.
4. There are 7 time series variables with voltage index among the variables. These variables include:
   - Battery voltage
   - Solar array voltages
   - Load voltages
5. There are 27 time series variables with current index among the variables. These variables include:
   - Battery current
   - Solar array currents
   - Load currents
6. The mission is 223 days long and there are 574687 samples for each time series variable.
7. The sampling rate is not constant and the time between two samples can vary from $0.4s$ to $10195s$ in the dataset.

4.1. Data preprocessing and feature extraction

We have decided to break each time series into 1512 time windows where each time window includes 380 samples. The sampling rate of the recorded data was not constant, therefore a time window may represent from 5 minutes to 10 hours of operation. In average each time window represents 3 hours and 31 minutes. The selected size of the time window was a tradeoff between detection accuracy and having enough data per window to detect/identify an abnormal behavior.

The task of the data processing is to break the recorded time series data into distinct features, which were used as inputs of the clustering algorithm. Wavelet transform is an important mathematical tool to analyze time series data because it contains both time and frequency information of a signal. A wavelet is an oscillation function which increases from zero, and then returns to zero. The wavelet transform of a signal is a convolution of a wavelet function with the signal. Those wavelet functions have different shapes and sizes and are implemented as band-pass filters. The output of the wavelet transform is a set of coefficients which capture the time (position in time) and frequency characteristics of the signal.

In signal processing wavelets are typically used in peak detection, noise reduction, and data compression. The discrete wavelet transform was applied to reduce the computational complexity and to filter the noise of all variables by extracting the scaling coefficients which correlate to the low frequency subbands. We used the package of functions for computing wavelet filters, wavelet transforms and multiresolution analyses (Aldrich, 2010) in R software environment[2] to extract the wavelet coefficients. The wavelet transform was decomposed up to the 8th coefficient for each variable. This resulted in 272 features per time window.

4.2. Clustering

In the next step the euclidean distance was employed to compute the pairwise distance between the temporally segmented objects. The time series waveforms were converted to a set of features represented by the wavelet transforms as discussed earlier. We used the R function *hclust* to generate the dedrogram shown in Figure 4. We selected UPGMA as the agglomerative clustering in this case study. The dendrogram represents the order in which the objects group into clusters. We used a heuristic, which suggested cutting the dendrogram at a height, where small changes in the location of the cut would not cause changes in the number of groups extracted from the dendrogram. The small clusters generated were potential anomalies. We will investigate them in more detail using the binary variables (switches) as context to understand the spacecraft's mode of operation, and experts recommendations to further characterize the condition. Cutting the tree at lower heights will generate more clusters and anomalies. Therefore, the level at which the dendrogram is cut represents a trade-off between increasing the precision of the anomaly definitions, while also reducing the number of anomalies an expert would have to analyze.

[2]see `http://www.R-project.org/`

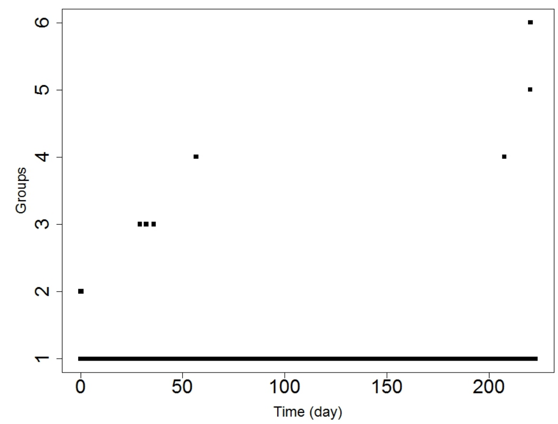

Figure 5. Anomalies and special modes during the mission

4.3. Anomaly Detection

Figure 5 shows the most relevant groups plotted over time (mission days) detected by the clustering algorithm. Each group represents a mode or an anomaly. Table 1 shows an abstract description of the most important modes and anomalies that were discovered from the groups generated.

4.3.1. Group 1: Normal operation

As we expected, the clustering algorithm generates a large cluster where significant number of time intervals reside. This cluster represents the normal operation of the spacecraft.

4.3.2. Group 2: The reaction wheels control problem

This time interval belongs to the beginning of the mission and is distinguished from the rest of the data set because of relatively high current in the SATORI #2 subsystem. As it is shown in Figure 6, the reaction wheels go off during this time interval. Note that Figure 6 shows that the reaction wheels go off two times during the mission. However, the experts in NASA confirmed that the reaction wheels only went off once during the mission and the second zero in the figure is a result of bad data. Considering the reaction wheels and the high current in the SATORI #2 subsystem, we found it very likely that this incident related to the guidance navigation and control unit. The experts from NASA confirmed that this time interval represents an anomaly and the incident is as follows.

In the first few orbits around the earth, the spacecraft began to spin at a faster rate than was expected, and the reaction wheels were turned off by the control software to avoid a high current load on the battery by the guidance, navigation and control (GNC) system. When the reaction wheels were turned off, the

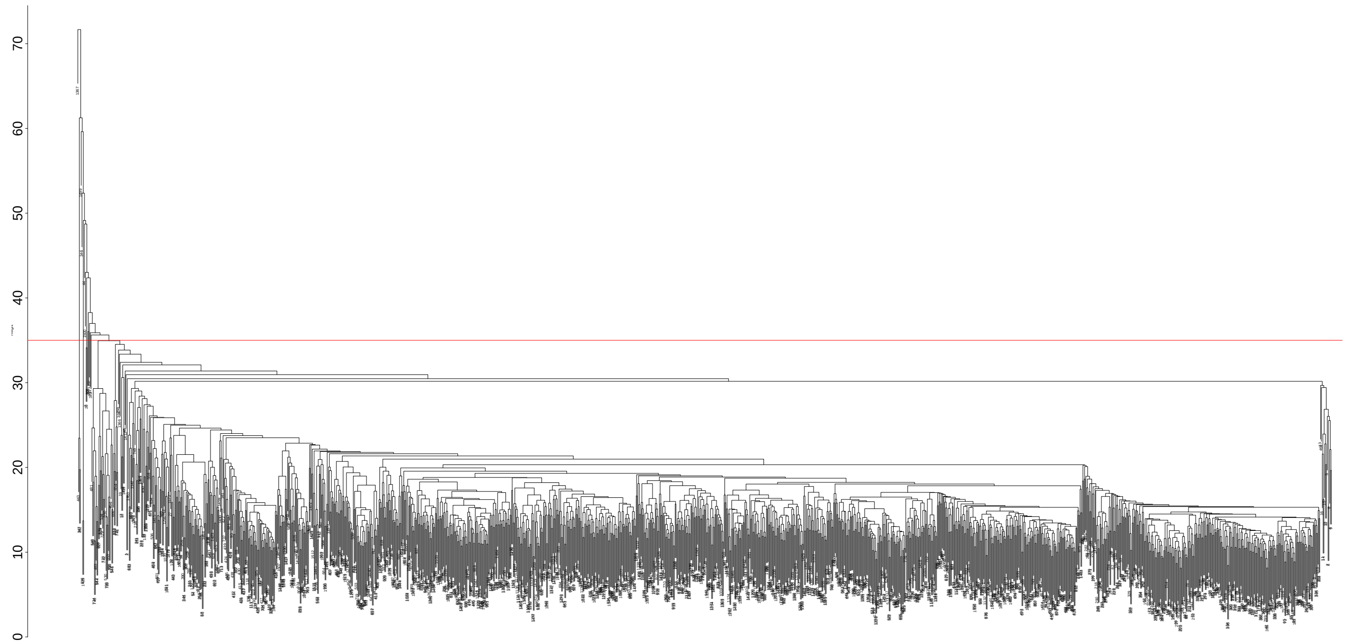

Figure 4. Hierarchical clustering

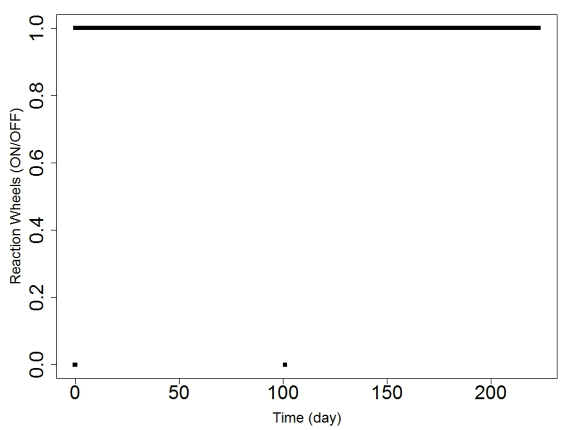

Figure 6. Reaction wheels (OFF=0, ON=1)

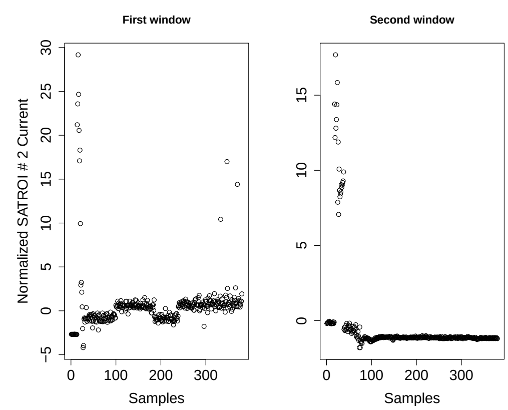

Figure 7. Normalized SATORI # 2 current during group 2

spacecraft stopped rotating, and, therefore, one side became too hot, and the other side too cold. To keep the temperature balanced several heaters went on which led to high current in the SATORI #2 subsystem. Figure 7 shows normalized SATORI # 2 current during group 2.

4.3.3. Group 3: The lunar orbit insertion

As it is shown in Figure 5 group 3 consists of three time intervals which occur in three different days. A high variation in the PAPI # 2 (propulsion, see Figure 1) subsystem's current distinguishes this group from the rest of the dataset. Figure 8 shows normalized PAPI # 2 current during this group.

Moreover, the valve driver unit which controls the propulsion subsystem, and the pressurant tanks heaters which also belongs to the propulsion subsystem (see Figure 1) goes ON in all the three time intervals. This group represents a unique behavior in the dataset, however, the experts confirmed that this does not represent an anomaly behavior and we should classify this group as a special operation mode. The experts informed us that this group represents the lunar orbit insertion process. In fact, there were three firing process in the propulsion subsystem to get into lunar orbit and our algorithm was able to classify them in a single group.

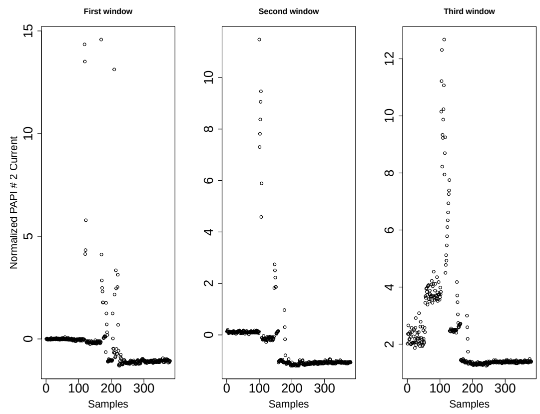

Figure 8. Normalized PAPI # 2 current during group 3

4.3.4. Group 4: The laser communication test

This mode of operation consists of two time windows where each of them is almost 20 minutes long. To explain this mode, we should notice that a laser communication test occurs during each time window. These tests are part of the mission plan and occur multiple times during the mission. However, because of the new moon, the solar array current is almost zero during this mode. The high current demand due to the laser communication test in the absence of solar energy puts too much pressure on the battery and leads to a battery voltage drop. Figure 9 shows the normalized battery voltage during each time window.

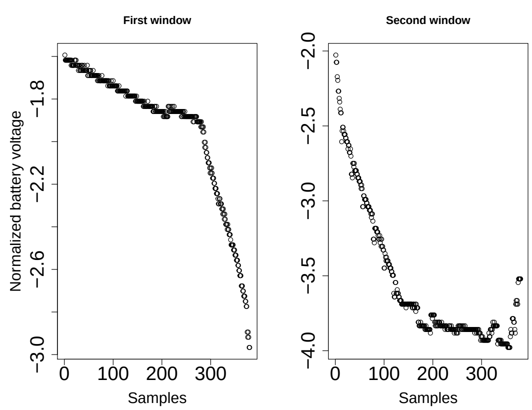

Figure 9. Normalized battery voltage during group 4

This group represents an abnormal behavior because the laser test in the absence of solar energy leads to unintended consequences.

4.3.5. Group 5: The eclipse

This group is related to the eclipse. The time interval is almost 5 hours long, because the sampling rate gets very low

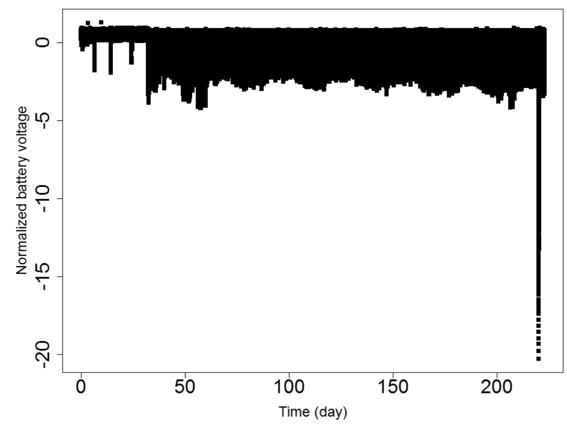

Figure 10. Normalized battery voltage during the mission

at the end of the mission. The solar array current is zero during the first 2 hours. Then it starts turning on, and switches between zero and none zero after that. Several heaters went on to keep different subsystems temperatures in normal range in the absence of sun light. This increases the load current. A simultaneous increase in the load current and decrease in the solar array current put an unprecedented pressure on the battery, which leads to battery voltage drop. We can classify the eclipse as an abnormal behavior because the battery is not designed to operate in this situation. Figure 10 shows the normalized battery voltage during the mission. We can see the eclipse at the end of the mission (day 219).

4.3.6. Group 6: The safe mode

The system goes to the safe mode right after the eclipse. During this mode, several loads were switched off to limit energy consumption and provide enough current for the battery recharge. This group represents a unique behavior in the dataset, however, it does not represent an anomaly behavior and we should classify this group as a special operation mode.

5. CONCLUSIONS, AND FUTURE WORK

In this paper, we presented a data-driven anomaly detection method. We defined the problem and presented a new approach for anomaly detection. Our approach uses an unsupervised learning method to detect the outliers that may represent special modes of operation, or they may be anomalies. To differentiate between special modes and anomalies, the method uses a supervised approach of consulting human mission experts. We applied our approach to detect anomalies during the LADEE mission. To detect anomalies, we looked for the features that were significant actors in differentiating the outliers from the nominal set. We also considered the binary switches during each time interval as the indicators of mode changes. Finally, we used domain experts to validate the anomalies.

Table 1. Abstract overview of detected modes and anomalies

Group	Detected Mode or Anomaly	Voltage or Current	Switches
1	Normal operation mode		
2	Reaction wheels anomaly	• SATORI # 2 current had a higher variance	• Propulsion heater turned on • Star tracker went off
3	Lunar orbit insertion mode	• PAPI # 2 current had large variations • SATORI # 2 current had large variations	• Pressurant tank heater went on • Valve driver unit went on
4	Laser communication test anomaly	• Solar array current was zero because of new moon • Battery voltage dropped because of high current demand of the laser communication test	• Laser communications switch went on
5	Eclipse mode	• Solar array current was zero during the first 2 hours, and had high fluctuation afterward.	• Several heaters went on (e.g. Propulsion heater)
6	Safe mode	• Battery current shows a high fluctuation	• Several loads (e.g. star tracker) turned off

We found that a number of detected anomalies were quite interesting to the experts including a fault in reaction wheels control system, and a laser test that caused a drop in the battery voltage. In both cases the experts confirmed the anomalies and flagged them for further investigation. The proposed approach performs well when the number of outliers is small enough for examining them one by one. In future work, we will investigate different algorithms to further automate the process.

ACKNOWLEDGMENT

This work was partially supported by funding from NASA STTR grant # NNX15CA11C. The authors gratefully acknowledge the support provided by Scott Poll, Mark Shirley, Peter Berg, and other members of the LADEE Mission Ops team from NASA Ames in acquiring, analyzing, and interpreting the LADEE Telemetry data.

REFERENCES

Aldrich, E. (2010). Wavelets: a package of functions for computing wavelet filters, wavelet transforms and multiresolution analyses. *R package version 0.2-60.*.

Blanke, M., & Schröder, J. (2006). *Diagnosis and fault-tolerant control* (Vol. 691). Springer.

Blom, H., & Bar-Shalom, Y. (1988). The interacting multiple model algorithm for systems with markovian switching coefficients. *Automat. Contr., vol. 33, pp. 780-783, Aug. 1988.*

Chen, J., & Patton, R. J. (2012). *Robust model-based fault diagnosis for dynamic systems* (Vol. 3). Springer Science & Business Media.

Hanlon, P., & Maybeck, P. (2000). Multiple-model adaptive estimation using a residual correlation kalman filter bank. *IEEE Trans. Aerosp. Electron. Syst., vol. 36, pp. 393-406, Apr. 2000.*

Henzinger, T. A. (Ed.). (2000). *The theory of hybrid automata.* Springer Berlin Heidelberg.

Hine, B., Spremo, S., Turner, M., & Caffrey, R. (2010). The lunar atmosphere and dust environment explorer (ladee) mission. In *Ieee aerospace conference.*

Isermann, R. (2005). Model-based fault-detection and diagnosis–status and applications. *Annual Reviews in Control*, *29*(1), 71–85.

Ji, M., Zhang, Z., Biswas, G., & Sarkar, N. (2003).

Hybrid fault adaptive control of a wheeled mobile robot. *Mechatronics, IEEE/ASME Transactions on*, *8*(2), 226–233.

Lee, E. A. (2008). Cyber physical systems: Design challenges. In *Object oriented real-time distributed computing (isorc), 2008 11th ieee international symposium on* (pp. 363–369).

Mack, D. L., Biswas, G., Koutsoukos, X. D., & Mylaraswamy, D. (2016, in press). Learning bayesian network structures to augment aircraft diagnostic reference models. *IEEE Transactions on Automation Science and Engineering*.

Marwedel, P. (2010). *Embedded system design: Embedded systems foundations of cyber-physical systems*. Springer Science & Business Media.

Niggemann, O., Biswas, G., Kinnebrew, J. S., Khorasgani, H., Volgmann, S., & Bunte, A. (2015). Data-driven monitoring of cyber-physical systems leveraging on big data and the internet-of-things for diagnosis and control. *26th International Workshop on Principles of Diagnosis, Paris, France*.

Noura, H., Theilliol, D., Ponsart, J.-C., & Chamseddine, A. (2009). *Fault-tolerant control systems: Design and practical applications*. Springer Science & Business Media.

Qin, S. J. (2012). Survey on data-driven industrial process monitoring and diagnosis. *Annual Reviews in Control*, *36*(2), 220–234.

Yin, S., Ding, S. X., Xie, X., & Luo, H. (2014). A review on basic data-driven approaches for industrial process monitoring. *Industrial Electronics, IEEE Transactions on*, *61*(11), 6418–6428.