

A Data Driven Health Monitoring Approach to Extending Small Sats Mission

Fangzhou Sun¹, Abhishek Dubey², Chetan S. Kulkarni³, Nagbhusan Mahadevan⁴ and Ali Guarneros Luna⁵

^{1,2,4} *Institute for Software Integrated Systems, Vanderbilt University, Nashville, TN, 37212, USA*
fzsun316@gmail.com abhishek.dubey@vanderbilt.edu nag@isis.vanderbilt.edu

³ *SGT, Inc., NASA Ames Research Center, Moffett Field, CA, 94035, USA*
chetan.s.kulkarni@nasa.gov

⁵ *NASA Ames Research Center, Moffett Field, CA, 94035, USA*
ali.guarerosluna@nasa.gov

ABSTRACT

In the next coming years, the International Space Station (ISS) plans to launch several small-sat missions powered by lithium-ion battery packs. An extended version of such mission requires dependable, energy dense, and durable power sources as well as system health monitoring. Hence a good health estimation framework to increase mission success is absolutely necessary as the devices are subjected to high demand operating conditions. This paper describes a hierarchical architecture which combines data-driven anomaly detection methods with a fine-grained model-based diagnosis and prognostics architecture. At the core of the architecture is a distributed stack of deep neural network that detects and classifies the data traces from nearby satellites based on prior observations. Any identified anomaly is transmitted to the ground, which then uses model-based diagnosis and prognosis framework to make health state estimation. In parallel, periodically the data traces from the satellites are transported to the ground and analyzed using model-based techniques. This data is then used to train the neural networks, which are run from ground systems and periodically updated. The collaborative architecture enables quick data-driven inference on the satellite and more intensive analysis on the ground where often time and power consumption are not constrained. The current work demonstrates implementation of this architecture through an initial battery data set. In the future we propose to apply this framework to other electric and electronic components on-board the small satellites.

Fangzhou Sun et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

1. INTRODUCTION

In the near future several organizations will be entering the era of small satellite technology for validating science missions around the earth as well as for deep space studies. Such missions will require the satellites to be healthy during the course of the mission to accomplish mission goals. Majority of the small sat missions at present are powered by lithium-ion battery packs.

A battery pack consisting of lithium-ion cells has been chosen to fly based on previous flight heritage and compliance with NASA battery safety requirements. Before batteries can be used for small satellite missions on ISS, both the individual cells as well as packs of multiple cells must be certified for safe operations. Certification tests for the cells and packs include electrical cycling characterization, over charging/discharging, external shorting, vibrational excitation, and exposure to vacuum. Only after each cell and pack has passed these certification tests can the batteries be installed in small satellites and delivered to the ISS.

Dependable, energy dense, and durable power sources are critical components for small satellite missions. They are subjected to the same strenuous operating conditions that the satellite is subjected to during transit to the ISS, deployment into space, and for the duration of the mission after launch from the ISS. Since the batteries come into close proximity to the astronauts on the ISS, it is critical to establish rigorous testing procedures to certify their safety.

Normal operation requires repeated charging and discharging of the batteries that age the packs and can lead to activation of internal fault protection systems. These critical internal fault protection systems prevent the batteries from destructively

failing in adverse scenarios, and protect electronic equipment from becoming damaged.

Through regular cycling at varying loads during mission operation the battery ages, losing the ability to hold full charge and to recharge the same amount. Through proper use and high quality construction, lithium-ion batteries can survive hundreds to thousands of these cycles depending on their operating conditions. In order to properly charge and discharge efficiently, it is important to understand the batteries charging characteristics. Cycling the batteries from maximum charge to minimum charge provides valuable information on the batteries, health and ability to recharge. This cycling also ensures that the batteries perform as specified and can be expected to perform appropriately when in operation.

One of the key conditions leading to degradation leading to faults is due to over charging/discharging of the cells. The ability to prevent such damaging cycling protects sensitive electronic systems powered by the batteries from being exposed to high or low voltages that could lead to cascading failures. Another adverse effects due to degradation is due to unintentional battery shorting leading to failure.

In long duration missions, small satellites may face such challenges where the batteries used may fail leading to unfinished mission goals. To overcome this issue, we propose a health monitoring framework for batteries in this work and probably extended to other systems in our future work. Earlier work on battery prognostics (Daigle & Kulkarni, 2013; Saha, Quach, & Goebel, 2012; Hogge et al., 2015) investigated and implemented physics and lumped parameter models for different systems. In this work deep-learning methodology for health monitoring of small satellites is proposed to develop a framework for low on-board computation on such systems limited by resources.

Simulated battery data sets for small satellites (Kulkarni & Guarneros Luna, 2018) are used from the NASA PCOE data repository. There are two data sets from different battery packs which run a set mission. Our deep learning framework is implemented to detect any abnormalities in the data sets.

Existing anomaly detection techniques can be Classification based, State based, Statistical/Consensus based, Clustering and Nearest Neighborhood based, or Information Theoretic (Chandola, Banerjee, & Kumar, 2009). *Classification* methods typically include Support Vector Machines (SVM), Neural Networks, Markovian and Bayesian Models. *State* based techniques often use Extended Kalman filtering or linear quadratic estimation techniques to predict normal behavior. However, this is not feasible in a data-driven system where we do not have a model of the system. While clustering techniques are useful (Mack, Biswas, Koutsoukos, & Mylaraswamy, 2015; Biswas et al., 2016), the success of *clustering* and *information theoretic* approaches is limited due to

unavailability of good and bad labels. Therefore, people prefer *Statistical* approaches that achieve statistical invariance under no fault scenarios. However, a major challenge is to find stable invariants, which can be made more difficult due to lack of data. Second, inherent uncertainties or high non-linearity, owing to behavioral randomness, makes statistical invariance hard to achieve without a large error residual. Further, in addition to the lack of the data the dimensionality of the problem also becomes a challenge in large systems. Therefore, our approach is to use deep-learning techniques and use them for both data-augmentation and challenges related to large dimensionality of a big system.

However, applying such techniques to cyber-physical systems like small satellites requires to (a) solve the challenges of developing mechanisms for learning the spatio-temporal patterns of the power system networks and (b) developing mechanisms for online-learning.

The two key problems with online learning mechanisms specifically in system like small satellites are : (a) they are limited in computation resources, (b) there is not enough bandwidth to stream all the real-time data to the central ground station for data processing. Therefore, in the work we propose on a novel-hybrid structure that uses pre-trained models and then re-learn the weights in the middle layers as new data is received during operation. The approach along with implemented framework is discussed in the later sections of this work.

The next section 2 discuss the current state of the art and challenges in developing a methodology. In section 3, our generalized approach to solve the problem is discussed. Section 4 discusses the datasets and setup for conducting experiments along with the results. The paper concludes with section 5 discussing the results and future work.

2. RESEARCH CHALLENGES

This section explores the key research challenges associated with anomaly detection for the battery systems on small satellites using deep learning techniques.

2.1. Capturing Spatial and Temporal Dynamics of Battery Sets

The performance of lithium-ion batteries degrades as they are cycled during mission operation. A prediction-based methodology to estimate health state of the batteries requires monitoring and a good understanding of the battery operations. This implementation of the methodology can either be data driven (Saha, Goebel, Poll, & Christophersen, 2007; Chen & Pecht, 2012) or model driven (Daigle & Kulkarni, 2013). In this work a data driven approach is presented to detect anomaly in battery operations and use this information for future health state estimation. Anomalous behaviors can be

identified when the differences between actual and predicted values of some battery characteristics (like capacity and voltage) cross a pre-defined operating threshold. Several model-based methods have been developed such as Kalman filter and particle filter (Saha et al., 2007; Daigle & Kulkarni, 2013; Bole, Kulkarni, & Daigle, 2014).

Fortunately, heterogeneous sensors have been developed to monitor key battery features and produce data traces, where data-driven analytics approaches can be employed. An approach can be developed based on long-short term memory (LSTM) models (Section 3.1.1) to learn the dynamics of a battery when it's being charged/discharged, make continuous one-step predictions, and compare with the ground truth to identify anomalies.

However, monitoring data traces from individual batteries is not enough, since a small satellite usually has more than one battery set and there are heterogeneous sensors (current, voltage, temperature, energy, etc.) tracking the status of the battery sets. There are still collective anomalies possible when sensor signals of individual batteries are normal. This motivates us to develop auto-encoder models (Sakurada & Yairi, 2014) to detect anomalies in the behaviors of battery set groups as discussed in Section 3.1.2. Auto-encoders are unsupervised machine learning models that reconstruct the original data with the low dimension representations. The reconstruction error between the original data point and its low dimensional reconstruction can be then used as a anomaly score.

2.2. Quick Inference on Small Satellites with Time, Resource and Energy Constraints

The dynamics of batteries are complex and vary widely due to many factors such as battery manufacturing processes, temperature, cycling profiles and rates, etc. A traditional way for training a data-driven model for battery anomaly detection is to (1) build models using data of normal behaviors that are observed in the past, and (2) trigger an alarm if the actual behaviors divert from the pre-trained models. However, there is no single model that could fit the dynamics of all batteries and models trained for batteries (that are monitored and already died) are usually specific to individuals, operation profiles and environment conditions. Because of the lack of knowledge (like degradation evolution and operation behaviors) of the specific batteries to be deployed, the models pre-trained using historical datasets on the ground can result in over-fitting problems.

Training predictive model at run-time could be a solution to learn the dynamics of new batteries in an adaptive way. But the limitations on computation resources on satellites make it very challenging. Not only the traditional training methods are time and power-consuming, but also a lot of data is needed

to train a complex model like deep neural networks. Thus it remains an open problem.

2.3. Application Aware Thresholds for Anomaly Detection

While the deep networks can learn the underlying patterns in the data series sequence and this prediction can be improved by using recurrent networks like LSTM and dimensionality reducer like auto-encoders (Section 3.1), we still need to develop mechanisms for identifying the thresholds that will finally become indicators of an anomaly. The deep learning models are updated in an online manner at runtime and the dynamics when batteries are charged/discharged are expected to be captured better over time. However, when anomalous charging characteristics occur, the previously learned patterns may change. Therefore we try to identify the anomalies by analyzing the overall trend of the residuals between the predicted and actual features.

3. DATA ACQUISITION AND ANOMALY DETECTION

This section describes proposed data-driven approach for on-board battery health monitoring and anomaly detection on small satellites. The overall architecture of the system is illustrated in Figure 1, where the cloud layer runs on powerful GPU-enabled servers on the ground and the edge layer is deployed on small satellites to detect faults in battery packs at run time. Two types of deep learning models are involved: (1) long short-term memory networks continuously make one-step predictions by looking back for multiple time steps of data from a battery set's sensors; (2) auto-encoder networks that focus on reconstructing individual time step's data from multiple sensors and multiple battery sets. These deep learning models are pre-trained to capture the overall spatio-temporal dynamics using historical and simulated battery operation datasets, and then adapted to the specific batteries deployed on-board through transfer and online learning. The models detect anomalies by evaluating the prediction error (i.e., the difference between predicted and actual sensor values) of LSTM models and the overall reconstruction error of auto-encoder models. Time series decomposition is conducted on the prediction error and reconstruction error to get three components: trend, seasonal and residual. The overall trends of residuals are then analyzed for each charging/discharging period to identify anomalous batteries. The term *cycle* used in the paper is defined as the complete process of either charging or discharging a battery.

3.1. Capturing Spatio-temporal Dynamics by Deep Learning Models

Accurately modeling the battery dynamics is critical for a prediction-based anomaly detection system. Our approach

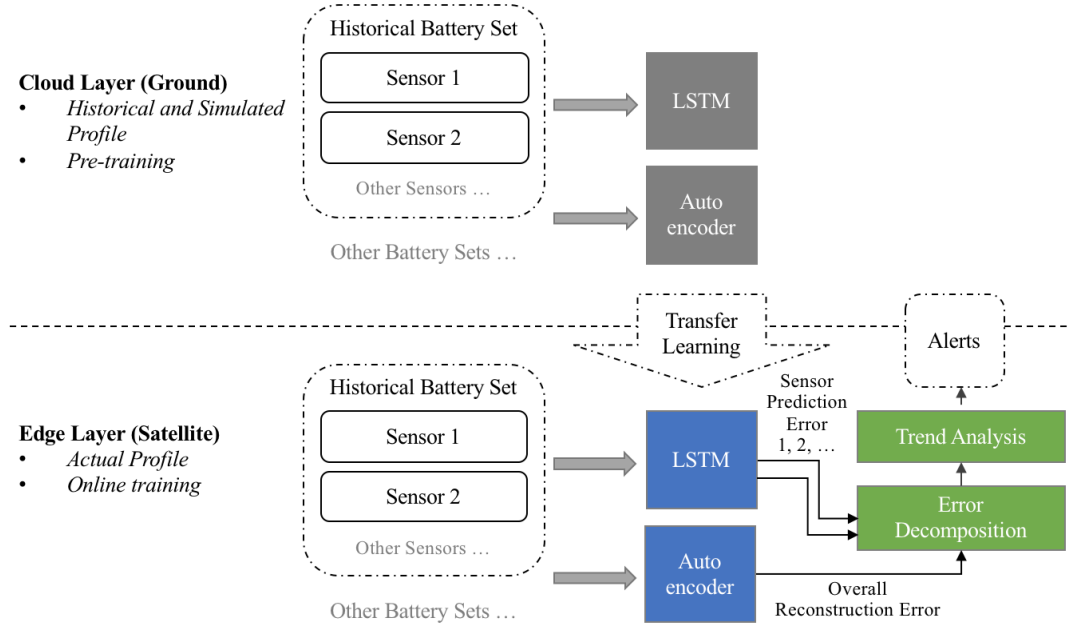


Figure 1. The overall architecture of the system contains a cloud layer and an edge layer. The cloud layer runs on powerful GPU-enabled servers on the ground to pre-train the deep learning models, while the edge layer is on small satellites for transfer and online training.

uses an architecture with deep learning models to learn the dynamics of on-board battery sets in two dimensions:

- **Temporal:** The data traces from sensors monitoring a battery set are time series, where deep learning techniques such as recurrent neural networks and their variants like LSTM models can be applied to learn the temporal dynamics.
- **Spatial:** For a satellite with one or more battery sets cooperating together, the data traces at a time step from multiple battery sets and sensors show spatial dynamics between them. We utilize auto-encoder models to learn the normal spatial patterns and identify the anomalous ones.

3.1.1. Long-Short Term Memory Networks

LSTM networks have been successful in modeling, classifying and predicting time series in many domains because of their ability to remember the short-term memory for a long period of time (Hochreiter & Schmidhuber, 1997). LSTM is a variant of recurrent neural networks that solves vanishing and exploding gradient issues by utilizing a gating mechanism - an input gate, an output gate and a forget gate. In order to address the temporal part of Challenge 1 in Section 2.1 (i.e., capturing the structure of time series from single battery sets), we develop stacked LSTM networks.

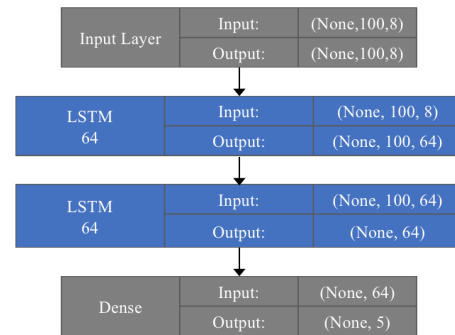


Figure 2. Stacked LSTM for sequence prediction.

Network: The LSTM model uses a stacked architecture that contains two LSTM layers. The architecture of the LSTM is illustrated in Figure 2. The first LSTM returns its full output sequences and the second one only returns its output sequences. The dense layer uses sigmoid activation functions to output final predicted values. Dropout is a regularization method that probabilistically exclude some inputs and neurons during training phase. We add dropout to the input and recurrent signals on the LSTM units to reduce over-fitting to the pre-trained dataset and improve the model's generalization performance. The training details are presented in Section 3.2.

Feature vector: The behaviors of batteries have many char-

acteristics, such as voltage, current, temperature, capacity, energy, step, cycle, etc. We combine data traces of such characteristics and contextual information (e.g., charging mode) to construct feature vectors, where the first five features are numerical but the charging mode feature is categorical - charging, discharging and rest. So the mode feature is encoded using one-hot encoding and appended to the end of the vectors (Figure 4).

Data Preparation: LSTM is sensitive to the scale of the input data, especially when sigmoid activation functions are used. So the raw data for each feature is normalized and re-scaled to the range of 0-to-1. We observed that there are short term patterns in the data, so a relatively long look back (i.e., the number of previous time steps to use as input variables to predict the next time period) is defaulted to 100 steps (the interval between two steps is 1 second).

3.1.2. Auto-encoder Models

In order to deal with the spatial aspect of Challenge 1 in Section 2.1 (i.e., learning the behaviors of multiple battery sets from heterogeneous sensors), we develop deep auto-encoder neural network models to learn from normal battery samples (Sakurada & Yairi, 2014). Auto-encoder is a machine learning model for non-linear dimension reduction that tries to learn a function that maps output to the same input via hidden layers. The lower dimension of hidden layers along with the goal that the difference between inputs and outputs are as small as possible forces the underlying structure of data is learned and noise is abandoned. Another advantage of auto-encoders is that it's an unsupervised learning technique which is suitable in context to the battery data set (Kulkarni & Guarneros Luna, 2018) for small satellites.

Network: The overall architecture of the auto-encoder network is illustrated in Figure 3. We use mean squared error as the loss function for training, which measures the similarity between input x and reconstructed output z :

$$J(x, z) = \|x - z\|^2 \quad (1)$$

The training details are presented in Section 3.2.

Feature Vector: Similar to the LSTM model in Section 3.1.1, the feature vectors of the auto-encoder are composed of both observed and contextual features. The difference is that they belong to the same time-step and the input and output vectors for training are exactly the same (Figure 5).

3.2. Online Learning

Data driven approaches, especially deep learning algorithms, typically rely on large amounts of data to be statistically sufficient to train models. Challenge 2 from Section 2.2 described the problem with battery anomaly detection on small satellites that different batteries vary in charging/discharging char-

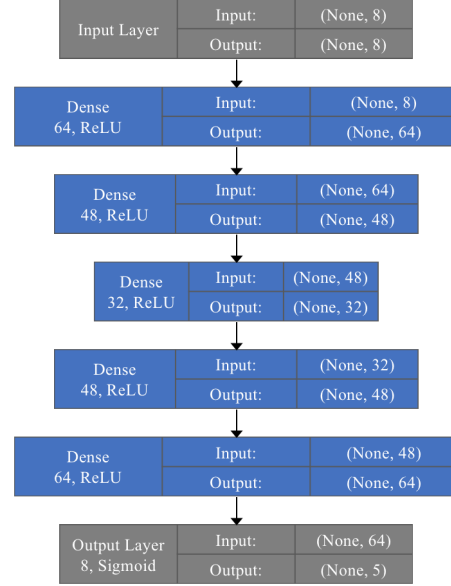


Figure 3. The deep auto-encoder architecture for learning the behaviors of multiple battery sets from heterogeneous sensors.

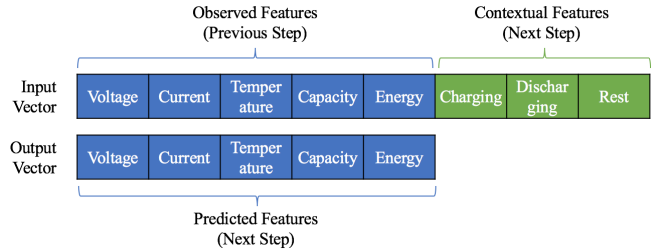


Figure 4. The observable features and contextual features are combined to construct feature vectors for LSTM networks.

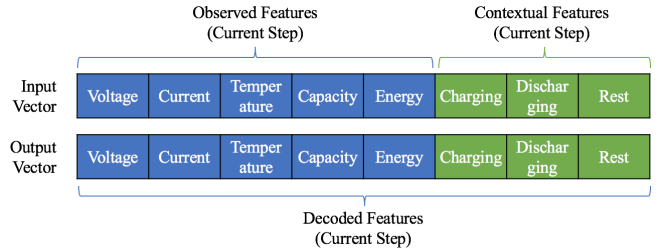


Figure 5. The input and output of the auto-encoder model are consisted of the same observed and contextual feature vectors.

acteristics but training deep neural networks on-board is energy and time-consuming. We address the challenge by transfer and online learning. The entire training process is divided into two phases:

- **Pre-Training:** The pre-training step learns initial weights for the LSTM and auto-encoder models. The computing centers on the ground usually have GPU-

enabled machines that are powerful enough for efficient deep learning training.

- **Online Learning:** The pre-trained models are reused and deployed on the small satellites. Online learning enables very low computational cost but accelerated on-board tuning based on the pre-trained models. The architecture of the pre-trained models as well as the weights in neural layers are reused and initialized.

The Adam optimization algorithm (Kingma & Ba, 2014) is an extension to stochastic gradient descent (SGD) that update neural network weights iterative based on training data in an adaptive way. In the online training phase, because the new training data collected at runtime contains noise and there are training constraints like time and power on board, we keep the training light weighted with a low learning rate, and a mini-batch gradient descent optimizer. Mini-batch size is a hyperparameter that affects the progress speed and the variance of the stochastic gradient updates. 50 is chose as the batch size.

3.3. Anomaly Detection

The previous sections described a stacked deep learning architecture of LSTM and auto-encoder models and a two-phase training process for transfer and online learning. The next step is to identify the anomalous battery charging/discharging behaviors by examining time series of prediction and reconstruction errors between predicted and actual values.

In this work we demonstrate the implementation of the framework to anomaly detection which can be extended to estimate future health state and predict any degradations in the system.

Error Decomposition: Time series often exhibit a variety of overall trend and seasonal patterns. In order to identify the underlying patterns separately, we utilize a technique for time series decomposition with moving averages to split a time series into three components (trend, seasonal, and residual):

$$Y[t] = T[t] + S[t] + e[t] \quad (2)$$

where $Y[t]$ denotes the final output that consists of $T[t]$ (trend), $S[t]$ (seasonal) and $e[t]$ (residual).

Residual Trend Analysis: The residuals decomposed from the prediction and reconstruction errors indicate the difference between the predicted and actual dynamics of battery sets. Since the deep learning models are updated in an online manner at runtime, the models will capture the dynamics better and better and the residuals will probably decrease over time. However, as battery ages through cycling, the charging characteristics change over time and may not follow the previous learned patterns. Therefore we analyze the overall trend of the residuals and try to identify the periods when the battery dynamics vary from the past.

To analyze the trend of residuals, the first step to divide the data traces in separate periods according to charging/discharging mode, and then quantify the benign range and find the outliers that deviate from the majority data points. The inter-quartile range (IRQ) is employed to identify the outliers where upper outliers are $1.5 \cdot \text{IQR}$ above the third quartile or $1.5 \cdot \text{IQR}$ below the first quartile. An example of the IQR results is illustrated in Figure 6.

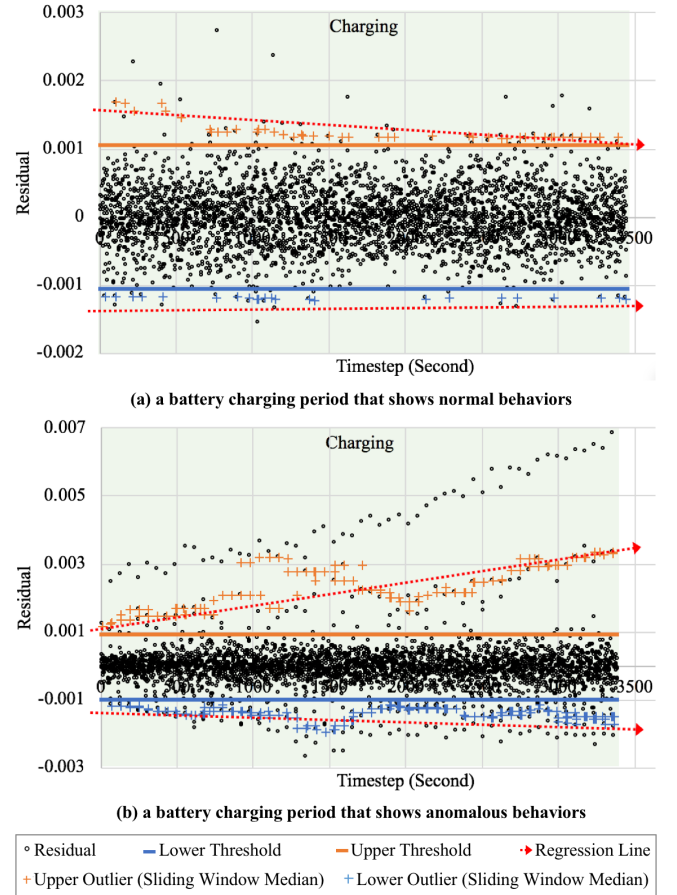


Figure 6. Normal and anomalous behaviors detected when a battery is being charged: (a) normal period – the residuals beyond the upper and lower thresholds become closer to zero over time; (b) anomalous period – the residuals beyond the upper and lower thresholds become more anomalous over time.

We use sliding window to re-sample the original sequence of data.

$$R(i) = \text{median}(r[i - 5], r[i - 4], \dots, r[i + 5]) \quad (3)$$

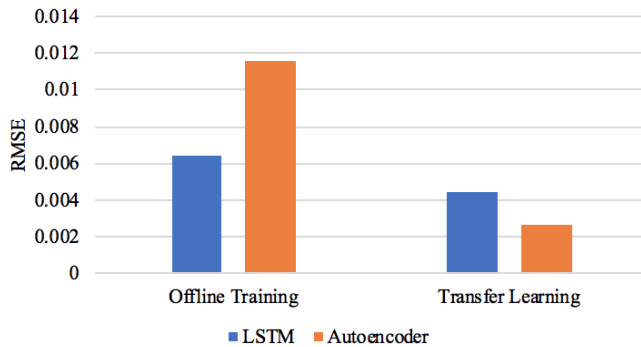
where $r[i]$ is the data point to be re-sampled and its value is replaced with the median of the 11 points around it. We then try to find the overall trend of the residual outliers. Two re-

gression lines are created to best-fit the outlier points above and below. Thresholds are decided from the historical battery datasets. Analysis results of two charging periods are illustrated in Figure 6.

4. EXPERIMENTAL RESULTS

This section describes the results of the tests on the proposed deep learning based health monitoring approach through battery data sets collected from simulated small satellite operations (Cameron, Kulkarni, Luna, Goebel, & Poll, 2015; Kulkarni & Guarneros Luna, 2018). Particularly, two BP930 Lithium-ion batteries packs (identified as PK31 and PK35) were operated continuously using a simulated satellite operation profile. The two battery packs were operated under similar operating conditions and the same loading profiles. The simulation data consists of the satellite traveling in and out of the sun which affects its charging and discharging cycles. While in the sun the batteries are charged at specific rates depending on the on-board solar panels. When in the dark the batteries are in discharge mode and the loading conditions change depending on mission requirements. The battery set of PK31 operated in overall good condition and finally failed at 39th cycles (total running time: 95496.77 seconds), while the battery set of PK35 died much faster at 5th cycles (total running time: 7783.95 seconds). Keras (Chollet et al., 2015) with TensorFlow (Abadi et al., 2015) back-end is used to implement the presented deep learning models.

4.1. Evaluating Transfer and Online Learning



Hypothesis \Rightarrow Transfer and online learning would enable quick training and accurate inference on small satellite with limited data and computation resources. As described in Section 3.2, detecting battery operation anomalies on small satellites can be challenging because of the variance in battery characters and the complexity of the deep learning models. For off-line models, even though they have been trained extensively using large-scale historical datasets, their actual performance may downgrade since the dynamics of the specific battery sets deployed on-board can differ greatly from the normal ones. On the other hand, using online train-

ing alone is not enough for real-time detection on small satellites. Although the total number of neural layers of our LSTM and auto-encoder models is relatively small, the weights are still in high dimension and optimizing them on weak computers without powerful GPUs could cost hours of training time and valuable power on satellites. Based on these considerations, the first experiment evaluates how the transfer learning approach accelerates the training process with light-weighted online updating compared with traditional off-line learning mechanisms.

Simulation Setup: To demonstrate this approach, we train the proposed LSTM and auto-encoder models in two ways: (1) off-line training only using the first half of PK31 dataset and then inferring the feature values in PK35 step by step, (2) pre-training using the first half of PK31 dataset in advance, and then using PK35 dataset to do single epoch light-weighted online training while inferring the values for the next step at the same time. For online training, Mini-batch gradient descent updating is utilized and the batch size is 50.

Simulation Results: Root-mean-square errors (RMSE) between the actual and the predicted feature vectors (capacity, current, energy, temperature, and voltage) are calculated for individual time steps. Figure 9 illustrates the average RMSE across all measured variables of LSTM and auto-encoder models between the two training mechanisms. The RMSE decreases for LSTM and % for auto-encoder. The results validate our assumption that transfer and online training mechanism is efficient and accurate to capture the battery charging dynamics compared with off-line training mechanism alone.

4.2. Evaluating Anomaly Detection

Hypothesis \Rightarrow The spatial and temporal dynamics of batteries could be captured by the LSTM and auto-encoder models. The anomalous battery charging or discharging behaviors would trigger anomaly detection. As described in Section 2.1, capturing the dynamics of battery sets is challenging since there are usually more than one battery sets and several data sensor traces available. We develop LSTM models for the temporal dynamics of an individual battery set and auto-encoder models for spatial dynamics of multiple battery sets and features.

Simulation Setup: We pre-train an LSTM model and an auto-encoder model using the first half of PK31 dataset and then transfer the parameters and conduct online training and inference using PK35 dataset. The battery dynamics that we analyzed include (1) Capacity, (2) Current, (3) Energy, (4) Temperature, (5) Voltage.

Simulation Results: Since the PK35 battery set failed much earlier than the PK31 battery set (7783.95 seconds vs. 95496.77 seconds), the PK35 is determined as an abnormal battery set. We expect that the anomalous spatial and tem-

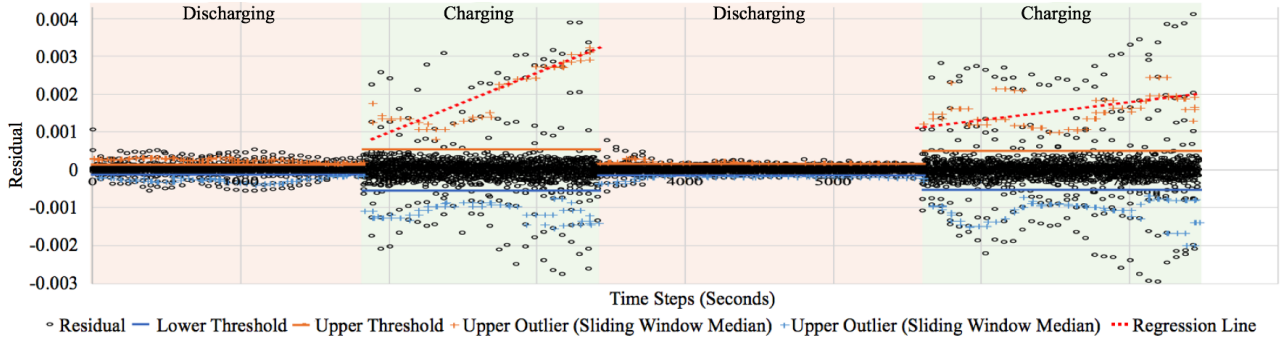


Figure 7. The residuals are decomposed from the reconstruction errors between auto-encoder inputs and outputs.

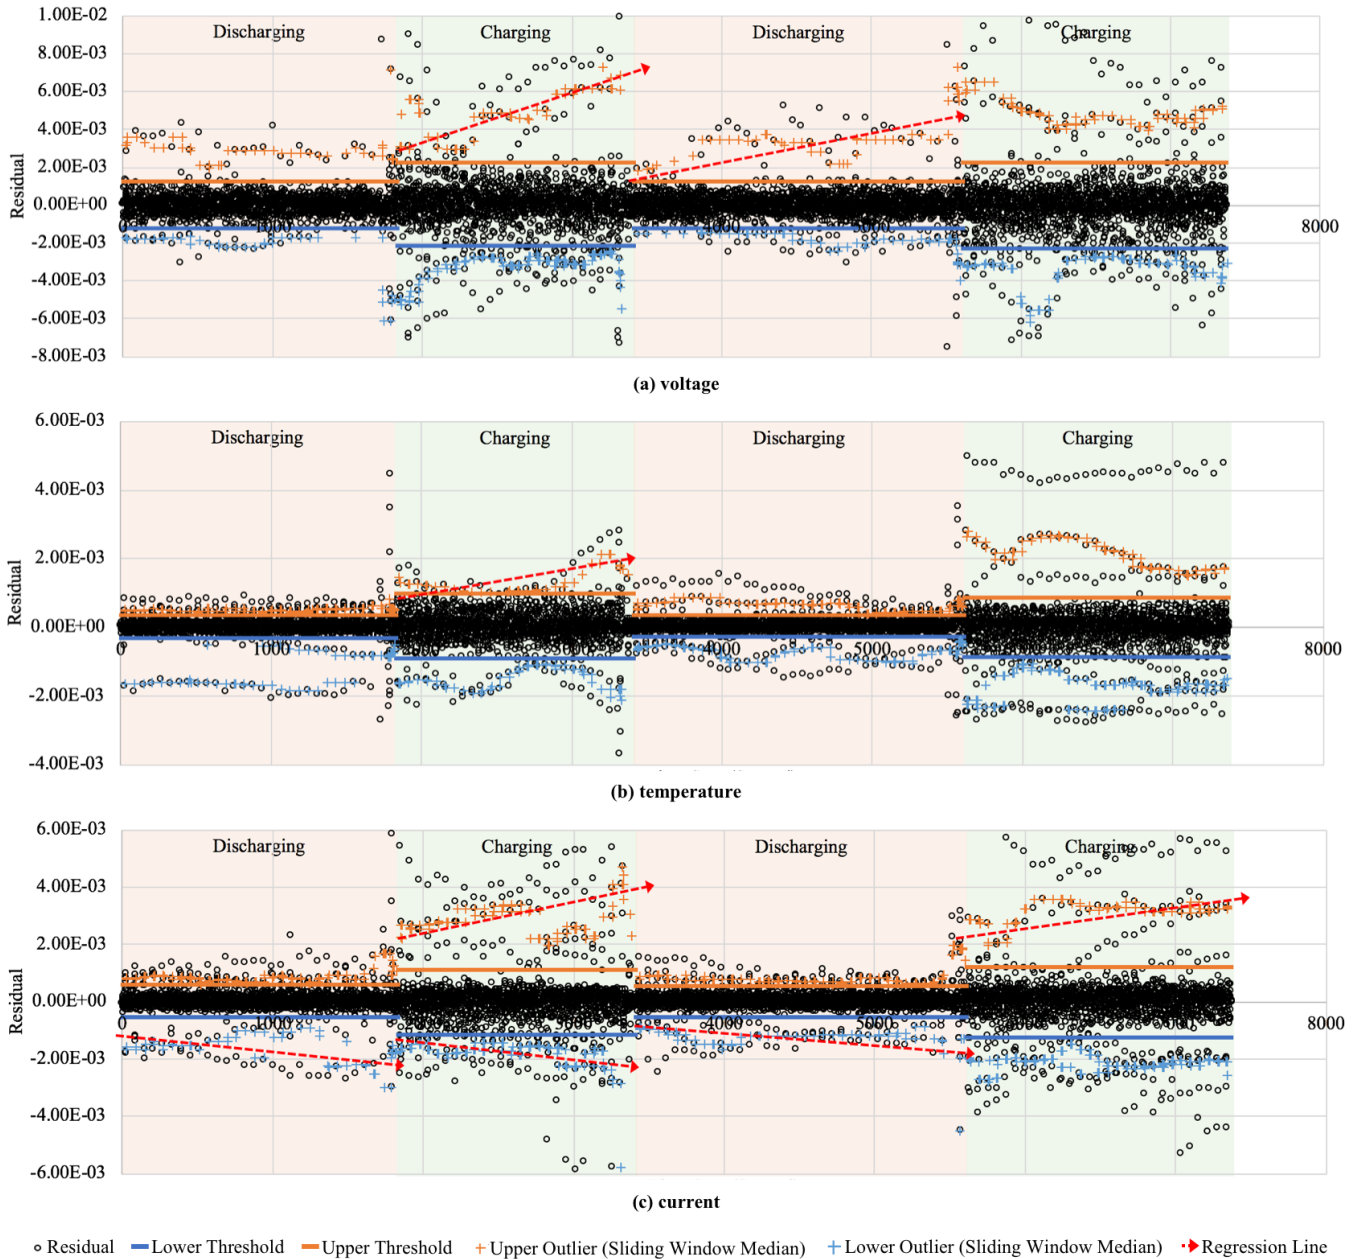


Figure 8. The residuals are decomposed from the time series of differences between actual and predicted voltage, temperature and current values (LSTM) outputs

Figure 9. The RSME between the normalized actual and predicted values of all measured variables in two training modes – off-line training and transfer learning.

poral behaviors on PK35 can be identified by the proposed anomaly detection mechanism using auto-encoder and LSTM models.

- **Auto-encoder:** Figure 7 illustrates the residuals decomposed from the reconstruction errors that are calculated using the auto-encoder’s inputs and outputs. The upper regression lines for the residual outliers in second and fourth phases are larger than thresholds learned from PK31 datasets, which validates our assumption that the auto-encoder model captures the anomalous behaviors from multiple sensors.
- **LSTM:** The residuals decomposed from the time series of differences between actual voltage, temperature, and current values and LSTM’s predicted values are shown in Figure 8. The overall trend of the anomalies from some phases are higher than thresholds learned from PK31 datasets. Generally, the battery began to fail in the first charging phase, and the phenomenon is captured by the LSTM model.

5. CONCLUSIONS

The paper demonstrates implementation of a novel hybrid framework which uses pre-trained models on off-board systems and then re-learn the weights in the middle layers on low computational edge processing to cyber-physical systems such as small satellites.

As more and more of such small satellites are launched processing health monitoring parameters off-board on ground stations and updating the model parameters on respective small satellites will improve capability of the system to estimate its health state and contribute to missions success.

Anomaly detection is the first step to analyze health state of the system. In the future we propose an updated framework to incorporate prognostics estimates for battery health. Implementing prognostics framework would enable small satellites perform certain mission profiles more efficiently based on predicted health state estimate. This enables the system to take correct decisions and perform required task efficiently for mission success and minimum computational and power requirements.

REFERENCES

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... Zheng, X. (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems*. Retrieved from <https://www.tensorflow.org/> (Software available from tensorflow.org)
- Biswas, G., Khorasgani, H., Stanje, G., Dubey, A., Deb, S., & Ghoshal, S. (2016). An application of data driven anomaly identification to spacecraft telemetry data. In *Prognostics and health management conference*.
- Bole, B., Kulkarni, C., & Daigle, M. (2014, September). Adaptation of an electrochemistry-based li-ion battery model to account for deterioration observed under randomized use. In *Annual conference of the prognostics and health management society 2014* (p. 502-510).
- Cameron, Z., Kulkarni, C. S., Luna, A. G., Goebel, K., & Poll, S. (2015). A battery certification testbed for small satellite missions. In *Ieee autotestcon, 2015* (pp. 162–168).
- Chandola, V., Banerjee, A., & Kumar, V. (2009, July). Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3), 15:1–15:58. Retrieved from <http://doi.acm.org/10.1145/1541880.1541882> doi: 10.1145/1541880.1541882
- Chen, C., & Pecht, M. (2012, May). Prognostics of lithium-ion batteries using model-based and data-driven methods. In *Proceedings of the IEEE 2012 prognostics and system health management conference (phm-2012 beijing)* (p. 1-6). doi: 10.1109/PHM.2012.6228850
- Chollet, F., et al. (2015). *Keras*. <https://keras.io>.
- Daigle, M., & Kulkarni, C. (2013, October). Electrochemistry-based battery modeling for prognostics. In *Annual conference of the prognostics and health management society 2013* (p. 249-261).
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780.
- Hogge, E., Bole, B., Vazquez, S., Celaya, J., Strom, T., Hill, B., ... Quach, C. (2015). Verification of a remaining flying time prediction system for small electric aircraft. In *Annual conference of the prognostics and health management society 2015. annual conference of the prognostics and health management society 2015*.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kulkarni, C. S., & Guarneros Luna, A. (2018). Description of simulated small satellite operation data sets.
- Mack, D. L., Biswas, G., Koutsoukos, X., & Mylaraswamy, D. (2015). Learning bayesian network structures to augment aircraft diagnostic reference models. *IEEE Transactions on Automation Science and Engineering, in Press*.
- Saha, B., Goebel, K., Poll, S., & Christophersen, J. (2007). An integrated approach to battery health monitoring using Bayesian regression and state estimation. In *2007 IEEE Autotestcon* (pp. 646–653).
- Saha, B., Quach, C., & Goebel, K. (2012). Optimizing battery life for electric uavs using a bayesian framework. In *2012 IEEE aerospace conference*.

Sakurada, M., & Yairi, T. (2014). Anomaly detection using autoencoders with nonlinear dimensionality reduction.

In Proceedings of the mlsda 2014 2nd workshop on machine learning for sensory data analysis (p. 4).