

# Online Learning of Spacecraft Simulation Models

**Justin R. Thomas\***

United Space Alliance  
600 Gemini  
Houston, TX 77058

**Christoph F. Eick**

University of Houston  
Department of Computer Science  
Houston, TX 77204

## Abstract

Spacecraft simulation is an integral part of NASA mission planning, real-time mission support, training, and systems engineering. Existing approaches that power these simulations cannot quickly react to the dynamic and complex behavior of the International Space Station (ISS). To address this problem, this paper introduces a unique and efficient method for continuously learning highly accurate models from real-time streaming sensor data, relying on an online learning approach. This approach revolutionizes NASA simulation techniques for space missions by providing models that quickly adapt to real-world feedback without human intervention. A novel *regional sliding-window* technique for online learning of simulation models is proposed that regionally maintains the most recent data. We also explore a knowledge fusion approach to reduce predictive error spikes when confronted with making predictions in situations that are quite different from training scenarios. We demonstrate substantial error reductions up to 74% in our experimental evaluation on the ISS Electrical Power System and discuss the early deployment of our software in the ISS Mission Control Center (MCC) for ground-based simulations.

## Introduction

Simulation plays a crucial role in NASA spaceflight. Software models recreate hardware behavior when using the real system is impossible due to costs, safety, or operational constraints. Spacecraft simulation is an integral part of mission planning, operations, training, and systems engineering.

Common practice at NASA in generating simulation models is to use hardware specifications, manufacturer test data, and physics to derive equation systems that describe system behavior. The resulting equations may range from simple algebraic expressions to complex integrals and differential equations that require advanced numerical analysis techniques. The final software implementations of these equation systems are known as *engineering models*. After initial model construction, engineers must tediously evaluate and adjust models in order to match true system behavior (Jannette et al. 2002). Engineering model adjustments range from modifying coefficients to changing equation forms.

\*Now at the Johns Hopkins University Applied Physics Laboratory  
Copyright © 2009, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

With the sheer size and evolving nature of the International Space Station, engineers cannot continuously adjust models to reflect current system. In this paper, we present a solution that uniquely solves this challenging problem by using machine learning to continuously and autonomously construct highly accurate models from real-time *telemetry* (streaming sensor data).

This solution advances the state of the art in NASA simulation techniques that traditionally require manual model construction and adaptation. The work in (Bay, Shapiro, and Langley 2002) recognizes the need for a machine learning solution to this problem, but does not address concerns such as model construction without the restrictions of specific equation forms, the use of supplemental knowledge during extrapolation, and efficient online learning. Our solution examines a drastically different approach that does not require any assumptions about the mathematical structure of the model, uses a model fusion approach to address extrapolation, and efficiently reacts to changes in system behavior within seconds.

Our approach pre-processes spacecraft telemetry, maintains a constant-size training dataset, and employs a regional sliding-window that preserves recent examples for localized regions within the input space, allowing for accurate predictions across all valid regions of the input space.

Using the resulting representative training dataset, we evaluate a set of candidate algorithms and learn a simulation model using the best training algorithm. Each resulting model is sent to the spacecraft simulator in real-time to match current system behavior. Since the simulator can request predictions for regions of the input space not contained in our training dataset, we employ a model fusion approach that uses the knowledge contained in existing engineering models to eliminate large spikes in error due to encountering situations that are quite different from training scenarios.

The unique contributions of the paper include:

- Online autonomous learning of highly accurate spacecraft simulation models from real-time telemetry
- Regional sliding-windows that outperform traditional sliding-windows for simulation model construction
- Model fusion that additionally uses traditional engineering models to make more reliable predictions in situations that are quite different from training scenarios

- A new system architecture that integrates machine learning into the model construction and adaptation process

We evaluate our technique using the ISS Electrical Power System (EPS) through NASA’s archives of spacecraft telemetry. Evaluation results demonstrate significant accuracy improvements over existing EPS engineering models.

The paper is organized as follows: The Background section provides relevant background information, the Technical Approach section describes our solution in detail, the Experimental Results section offers results from our experimental evaluation, the Early Deployment section discusses the early deployment of our technique at NASA, and finally, the Conclusion section closes the paper.

## Background

### Spacecraft Simulation

For ISS mission operations, simulations allow engineers to answer critical questions such as the following: How long until the crew exhausts the oxygen supply during a cabin pressure leak? How can we orient the spacecraft to retain communication when mechanical antenna positioning fails? Is enough power available to run a science rack experiment for the required fourteen days?

Engineers generate initial models using hardware specifications, manufacturer test data, and physics to derive a set of equations that describe system behavior. After initial model construction, engineers must tediously evaluate and adjust models in order to match true system behavior. Simulation models need to reflect real-world system behavior, otherwise their value is limited. Engineers find it difficult and time-consuming to create and maintain accurate models for spacecraft systems for the following reasons:

- Human-rated spacecraft systems are complex
- True performance is only observable in space
- Abnormal scenarios are difficult to understand
- System behavior can evolve over time
- Extensive effort is required to refine models

We address these challenges by autonomously generating highly accurate models that reflect current system behavior.

### ISS Electrical Power System

We chose the ISS EPS (Gietl et al. 2000) to evaluate our technical approach due to our EPS domain knowledge, however, our approach is generically applicable to any system.

Our focus is on two core components: the battery charge/discharge unit (BCDU) and the battery itself. The hardware layout for a single power channel (eight total) is illustrated in Figure 1. The ISS absorbs sunlight through its massive solar arrays and converts the energy into a usable power source through a network of complex electrical equipment. An array of nickel-hydrogen batteries stores excess energy for use during orbital eclipse. The ISS orbit results in cyclic battery charge/discharge behavior due to the periodic transition from eclipse to insolation (periods of solar radiation) as demonstrated in Figure 2. The EPS model must forecast into the future for power availability planning.

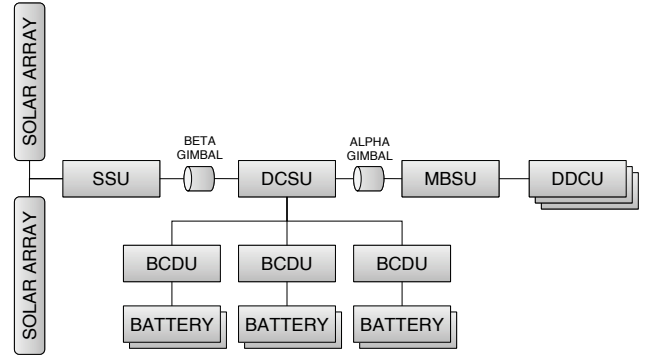


Figure 1: ISS Electrical Power System schematic




Figure 2: Battery charge level telemetry

## Technical Approach

In the Requirement Analysis section, we provide a requirement analysis for our solution. Then in the System Architecture section, we outline the system architecture and illustrate the flow of information from incoming telemetry to final simulation output. We then focus on the following solutions necessary to solve specific problems: regional sliding-windows (Regional Sliding-Window section), model learning and evaluation (Model Learning and Evaluation section), and model fusion (Model Fusion section).

### Requirement Analysis

We must solve a series of problems in order to effectively learn spacecraft simulation models in an online fashion:

- Process large continuous sensor data streams
- Create models that forecast into the future
- Account for frequent and long periods of missing data
- Create accurate models across valid input space regions
- Adjust models for system degradation and change
- Learn and evaluate models without human intervention

We discuss these requirements in detail throughout the discussion of our approach.

### System Architecture

Our architecture consists of the offline analysis components and online learning components as shown in Figure 3. We

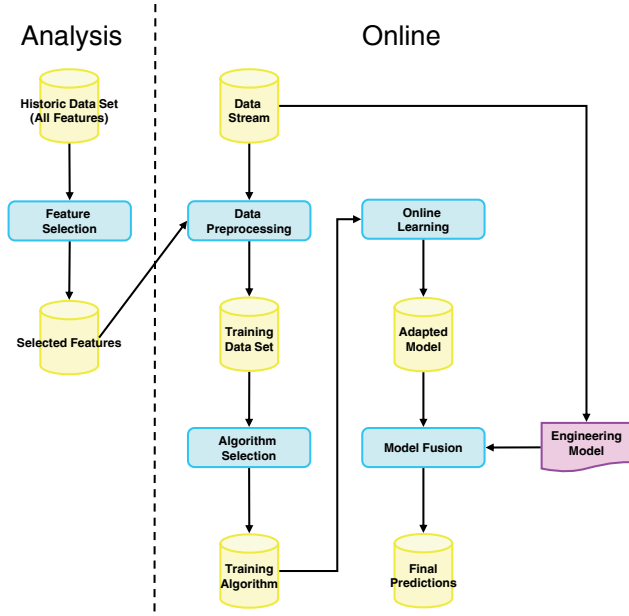


Figure 3: System architecture

---

#### Algorithm 1 Regional Sliding-Window Algorithm

---

**Input:** trainingSet, dataStream  $\vec{x}[]$ , size  $n$ , size  $k$   
Initialize  $trainingSet \leftarrow \emptyset$   
**loop**  
     $newPoints \leftarrow$  next  $n$  samples from  $x$   
     $trainingSet \leftarrow trainingSet \cup newPoints$   
     $trainingSet \leftarrow k\text{-means from } trainingSet$   
**end loop**

---

now describe the flow of data through these system components. Before running our system, the Feature Selection component determines the optimal set of sensors that generate the best possible model. Once our system is running, all incoming telemetry runs through the Data Preprocessing component to condition the data appropriately for model learning and evaluation. At a fixed frequency, the system will revise the training dataset using the regional sliding-window. The system sends this training dataset to the Algorithm Selection component where the best algorithm is selected from a set of candidate training algorithms. The Online Learning component then learns a system model using the current training set. The system then sends the resulting *learned models* to the simulator where the Model Fusion component combines the knowledge contained in existing engineering models to prevent highly inaccurate predictions.

Detailed descriptions of the Feature Selection and Data Preprocessing components are provided in (Thomas 2007) and are not the focus of this paper.

### Regional Sliding-Window

When learning from data streams (Babcock et al. 2002; Gaber, Zaslavsky, and Krishnaswamy 2005), it is not realistic to use the entire data stream for model learning. In our

case, a year of operational data sampled every 10 seconds results in a training set with over 3 million data points (1 GB). A popular technique to handle data streams is to only retain the  $n$  most recent data points. This is known as a sliding-window (Widmer and Kubat 1996). This simplistic model works well in many situations, but suffers from the drawback that knowledge is lost when it leaves the sliding-window (Figure 4). For a spacecraft simulation model, we desire the most recent examples covering the entire input space as last known behavior. We propose a regional sliding-window to perform this task.

We use  $k$ -means clustering to create a representative dataset across all exercised operational regions. Algorithm 1 takes an initial training dataset and appends  $n$  samples from data stream  $\vec{x}[]$ . Then  $k$ -means is used to generate a large number of representatives, keeping the training set a constant size  $k$ . Examples in the training set that are nearby the new samples from  $\vec{x}[]$  will be shifted towards the new data as  $k$ -means determines the new representatives. Examples from the training set that are not nearby the new data will not move. We desire this behavior since we do not want to discount or remove elder data that are most recent for a particular region in the input space. By selecting a sufficiently large value for  $k$ , we can maintain data points across the entire input space as shown in Figure 4.

The selection of the  $k$  and  $n$  parameters are important in the performance of this regional sliding window approach. Our experimental evaluations found that  $n$  must be much smaller than  $k$  (typically around 90-95% smaller). The value of  $k$  is related to the number of data points necessary to preserve valid operating values across the input space. We experimentally determined  $k$  by producing a series of plots on a data set with increasing values of  $k$ , and then increasing the value of  $k$  by 20% to account for future data not included in our input dataset. If  $k$  is too large, our runtime performance will suffer due to increased dataset size during training, and if  $k$  is too small, then we might not be able to preserve enough data for quality training. Also  $n$  controls the training frequency, since after  $n$  input points are buffered, we then execute  $k$ -means to perform a reclustering. As the value of  $n$  decreases, system performance decreases.

Many system models must also account for system degradation. In the field of machine learning, this is known as the concept drift problem (Widmer and Kubat 1996). Concept drift is best described as dynamic change to the underlying data-generating distribution. In our domain, the EPS components tend to degrade over time, resulting in changes to the system behavior and corresponding output. Our regional sliding-window technique also indirectly accounts for the concept drift phenomenon since incoming data points will influence the elder training examples in that local region.

### Model Learning and Evaluation

When learning from data streams, our system takes in the regional sliding-window of telemetry and generates a model at a fixed frequency. During the learning process, the system selects the best algorithm from a candidate set based on mean absolute error on independent test sets. We selected the candidate algorithms by examining our problem charac-




Figure 4: Regional sliding-window example – A full day of battery data (left) shows a loop in the middle of the input space. This loop is lost the next day using a traditional sliding-window (middle). However, a regional sliding-window maintains this important knowledge (right).

teristics (real-valued inputs/outputs, efficient prediction performance, etc...) and experimenting on sample datasets. We used the following algorithms in our experiments:

- Artificial Neural Networks (single and boosted)
- Linear Regression
- Regression Trees (single and bagged)
- Model Trees (single and boosted) (Quinlan 1992)

The EPS telemetry shown in Figure 2 contains significant gaps which represent data unavailability due to communication outages between the ISS and the MCC. The Tracking and Data Relay Satellite (TDRS) communications system is a shared resource and we expect such communication outages, even for up to 30 minutes.

For time-independent models (such as the BCDU model), we construct training sets and independent model evaluation sets by simply removing all data points with missing data from the full dataset. This basic approach is valid since the occurrence of missing data is completely independent of EPS system operation and we have ample training data.

For time-dependent forecasting models (e.g. battery model), we create a *lagged variable* (Chatfield and Weigend 1994) for the time-delayed output. To create a lagged variable, the system creates a new feature for the data point at time  $t$  with the output from the data point at time  $t + 1$ .

To evaluate forecasting models, we cannot use traditional  $n$ -fold cross validation since we must have a contiguous time-series of data. Instead we construct time-series test sets from independent data. We desire datasets without missing data for at least 90 minutes (one ISS orbit), but these datasets are extremely scarce. Instead we employ linear interpolation to complete missing values. To avoid overly biasing our evaluation, we experimentally determine the maximum duration of missing data that meets our quality threshold. For our evaluations, we generated ample test sets using a 40-second threshold (under the 100-second maximum threshold).

## Model Fusion

Simulation models provide engineers a means of experimenting and trying out many potential scenarios. This means system models must be able to make predictions in any valid region of the input space at any time. We must




Figure 5: Battery model error comparison (engineering model vs. learned model)

account for cases when our system model must predict in a region in the input space where no training examples exist; a common problem for inductive algorithms.

Model fusion is a knowledge fusion approach that uses traditional engineering models to supplement the learned models. Our software uses the output from the learned model when the incoming data point is very similar to the training dataset, otherwise we use the output from the engineering model. We observed that engineering models provide useful background knowledge to reduce predictive error spikes when confronted with making predictions in situations that are quite different from the training scenarios used when learning the machine learning model. We refer to the resulting model as the *fused model*.

As we show in Figure 5, battery models created using our online learning technique outperform the existing engineering model with the exception of a few large spikes. Visual inspection of the telemetry data in Figure 6 shows that during one of these spikes, there were novel examples reflected by the loop in the middle of the normal operating region. This scenario demonstrates the need for model fusion since the learned model did not accurately predict these inputs.

To determine if the incoming data point is similar to the training dataset, we turn to unsupervised density estimation techniques as proposed by Bishop (Bishop Aug 1994). We use a Gaussian Mixture Model (GMM) density estimator to determine if the new data point is novel when compared to the training dataset. Our software achieves this by measuring the distance from the clusters generated by the standard




Figure 6: Previously unobserved battery data points

---

**Algorithm 2** Model Fusion Algorithm

---

**Input:**  $\text{dataPoint}$ ,  $\text{means}[]$ ,  $\text{covars}[]$ ,  $\text{threshold}$   
 $c \leftarrow \text{NearestCluster}(\text{dataPoint}, \text{means}, \text{covars})$   
 $\text{distance} \leftarrow \text{MahalanobisDistance}(\text{dataPoint}, c)$   
**if**  $\text{distance} > \text{threshold}$  **then**  
     $\text{output} \leftarrow \text{EngineeringModelOutput}(\text{dataPoint})$   
**else**  
     $\text{output} \leftarrow \text{LearnedModelOutput}(\text{dataPoint})$   
**end if**

---

Expectation-Maximization training procedure as shown in Algorithm 2. If the distance is greater than our global novelty threshold, the incoming data point is deemed novel and we use the output from the engineering model. The global threshold is set so that 99% of the training set is deemed nominal to allow for some noise. We use the Mahalanobis distance function rather than the normal probability density function to increase simulator runtime performance since the probability density function uses several exponential terms that impact performance.

Figure 7 shows the large error spikes by the learned model in Figure 5 are reduced in the final fused model. In some instances the mean error in the final fused model is larger than the learned model, but these small sacrifices are necessary to bound error spikes caused by extreme extrapolation.

Model fusion provides a safety net to ensure that adap-




Figure 7: Battery model error comparison (learned model vs. fused model)




Figure 8: Battery model error comparison (engineering model vs. fused model)

tive model extrapolation does not introduce bizarre results. Unusual output will hurt the user's confidence and decrease their trust in this intelligent system. This safeguard is necessary for real-world deployment in spacecraft simulation.

## Experimental Results

For our experimental evaluation, we used 90 days of ISS historical telemetry data from year 2006. First, we performed feature selection for the battery and BCDU devices. Next, we started a regional sliding-window with  $k=2000$  and  $n=100$  to generate training datasets. Then, our software learned new models every 24 hours. Finally, all models were evaluated on independent test sets one week into the future.

A summary of the results is provided in Table 1. The regional sliding-window approach improves over traditional sliding-windows by reducing mean absolute error (MAE) by 53% (battery), 25% (BCDU current), and 21% (BCDU voltage). Model fusion slightly increases MAE for all learned models, but removes drastic spikes in error by reducing maximum errors by 70% (battery), 2% (BCDU current), and 73% (BCDU voltage). Our final models improve over the existing engineering models with 25% (battery), 73% (BCDU current), and 74% (BCDU voltage) reductions in MAE. In all cases, the system selected boosted model trees as the learning algorithm due to EPS characteristics, but our approach remains generic for future systems.

## Early Deployment

ISS power system flight controllers approved our proposal to integrate this method into a ground-based EPS model used for ISS power planning in the Mission Control Center. This is the only method in use at the NASA Johnson Space Center (JSC) that applies machine learning to spacecraft system simulation. We split this project into two phases in order to evaluate this method in stages and build user confidence in machine learning techniques for critical mission support.

The first phase integrated learned models generated offline from archived data. Models are not continuously learned and updated within the simulator. Without full on-line adaptation, we must manually trigger the learning process to generate new models from recent data. This allows human validation of new models, but does not meet our goal of fully adaptive and automated model learning.



Approach	Battery	BCDU Current	BCDU Voltage
Existing engineering model	0.00831	1.1305	4.8653
Learned model with traditional sliding-window	0.01096	0.3037	0.8237
Learned model with regional sliding-window	0.00513	0.2273	0.6538
Final model (model fusion and regional sliding-window)	0.00623	0.3079	1.2574

Table 1: Results summary (mean absolute errors)

The first phase completed in April 2008 with deployment into the Solar power planning software application to increase accuracy for a critical constraints-based solar array planning problem (Thomas and Downing 2008). We received positive feedback from the flight controllers who use Solar, but the exact accuracy increases are hard to quantify beyond results provided in this paper because the model predicts the amount of ISS available power; a number that cannot be measured under normal circumstances because the ISS must not operate under maximum power consumption.

The existing EPS model was implemented in Java, and therefore we also implemented the software supporting on-line learning of spacecraft simulation models in Java. Our software currently runs within a Windows environment on Intel x86-based laptops, although the Java implementation allows for portability across a variety of platforms. Total project effort was estimated at 500 hours from problem description to deployed software (first phase). We leveraged the WEKA data mining software library (Witten and Frank 2005) for the training algorithm (ANN, Model Tree, etc...) and  $k$ -means implementation to reduce costs, however we implemented our own data preprocessing, results evaluation, and Gaussian Mixture Model algorithm due to custom project requirements.

The second phase consists of deploying the full on-line learning and evaluation approach. Software will continuously and autonomously generate models from real-time telemetry. Our model evaluation technique will collect statistics and validation test sets to perform trustworthy automated model validation. The second phase is under consideration for 2009.

## Conclusion

This paper introduced a unique and efficient method for continuously learning highly accurate models from real-time streaming sensor data, relying on an online learning approach. This approach revolutionizes NASA simulation techniques for space missions by providing models that quickly adapt to real-world feedback without human intervention. A novel regional sliding-window technique for on-line learning of simulation models was proposed that regionally maintains the most recent data. We also explored a knowledge fusion approach to reduce predictive error spikes when confronted with making predictions in situations that are quite different from training scenarios. We demonstrated substantial error reductions up to 74% in our experimental evaluation on the ISS Electrical Power System which resulted in our approach being the only application of machine learning in spacecraft simulation in use at NASA JSC.

## Acknowledgments

We would like to thank the reviewers, Susan Ahrens, Natalie Ducote, Ray Halyard, Mike Genest, Jon Yemin, Patrick Walter, and the PRO and PHALCON flight control groups for their participation and contributions. This work was funded by NASA JSC under contract NAS9-20000.

## References

- Babcock, B.; Babu, S.; Datar, M.; Motwani, R.; and Widom, J. 2002. Models and issues in data stream systems. In *PODS '02*, 1–16. New York, NY, USA: ACM.
- Bay, S. D.; Shapiro, D. G.; and Langley, P. 2002. Revising engineering models: Combining computational discovery with knowledge. In *ECML '02*, 10–22.
- Bishop, C. Aug 1994. Novelty detection and neural network validation. *Vision, Image and Signal Processing, IEE Proceedings* 141(4):217–222.
- Chatfield, C., and Weigend, A. S. 1994. Time series prediction: Forecasting the future and understanding the past. *International Journal of Forecasting* 10(1):161–163.
- Gaber, M. M.; Zaslavsky, A.; and Krishnaswamy, S. 2005. Mining data streams: a review. *SIGMOD* 34(2):18–26.
- Gietl, E. B.; Gholdson, E. W.; Manners, B. A.; and Delventhal, R. A. 2000. The electric power system of the international space station - a platform for power technology development. Technical Report TM-2000-210209, NASA.
- Jannette, A. G.; Hojnicky, J. S.; McKissock, D. B.; Fincannon, J.; Kerslake, T. W.; and Rodriguez, C. D. 2002. Validation of international space station electrical performance model via on-orbit telemetry. In *ECEC '02*, 45–50.
- Quinlan, J. R. 1992. Learning with continuous classes. In *5th Australian Joint Conference on AI*, 343–348.
- Thomas, J., and Downing, N. 2008. A flexible spacecraft operations strategy for complex constraint management. In *AIAA SpaceOps 2008*.
- Thomas, J. 2007. Adaptive modeling of the international space station electrical power system. Master's thesis, University of Houston, Department of Computer Science.
- Widmer, G., and Kubat, M. 1996. Learning in the presence of concept drift and hidden contexts. *Machine Learning* 23(1):69–101.
- Witten, I. H., and Frank, E. 2005. *Data Mining: Practical machine learning tools and techniques*. San Francisco, CA: Morgan Kaufmann, 2nd edition.