IN11C-0798

AGU 100 FALL MEETING
Washington, D.C. | 10–14 Dec 2018

Detection and Interactive Exploration of Telemetry Anomalies

Valentino Constantinou, Christopher Laporte, Ian Colwell, Kyle Hundman (vconstan@jpl.nasa.gov)

NASA Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA

Jet Propulsion Laboratory
California Institute of Technology

OVERVIEW

As scientific instruments collect increasing amounts of data, improved systems are needed to analyze, condense, and present this data to engineers and scientists.

Automated anomaly detection for spacecraft telemetry is an area where modern data science and machine learning methods could provide substantial value.

We use Recurrent Neural Networks (LSTMs) to predict incoming telemetry values using recent telemetry, commands, and event records (EVRs) as inputs.

Where predictions are substantially different from actual telemetry values, these are identified as potentially anomalous events. We use a novel nonparametric method for determining “substantially different”.

MOTIVATION

Increasing data rates for missions necessitate automated anomaly detection methods. For example, NASA/JPL’s upcoming SWOT, NISAR missions will be generating 3-4 terabytes (TB) of data daily.

Current spacecraft monitoring systems only target a subset of anomaly types and often require costly expert knowledge to develop and maintain.

Thresholding systems, for example, are unable to detect “contextual” anomalies (anomalies which stay within the limits of their channel’s values, yet exhibit abnormal behavior in the context of their surrounding “normal” telemetry values).

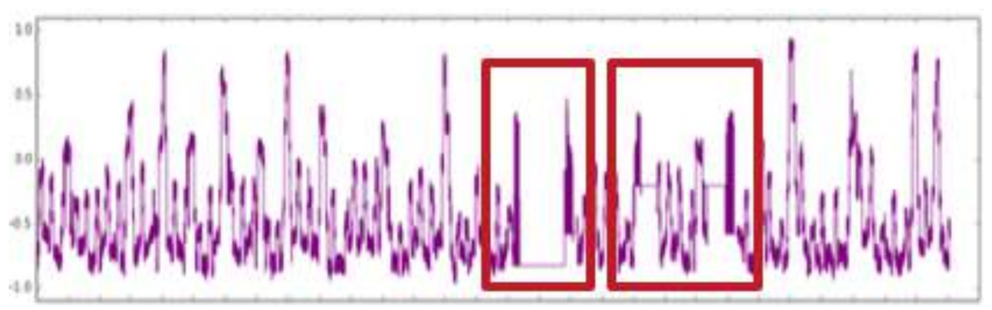

Example of anomalies not detectable using thresholds or distribution-based methods

Smaller missions (e.g. cubesats) will have fewer engineers dedicated to operations, and will have a greater need for automated monitoring methods.

APPROACH

Construct one model per telemetry channel. Modeling each telemetry channel independently allows traceability down to the channel level.

Inputs into the LSTM consist of prior telemetry values for a given channel and encoded command information sent to the spacecraft.

The combination of the module to which a command was issued and whether a command was sent or received are one-hot encoded and slotted into each time step.

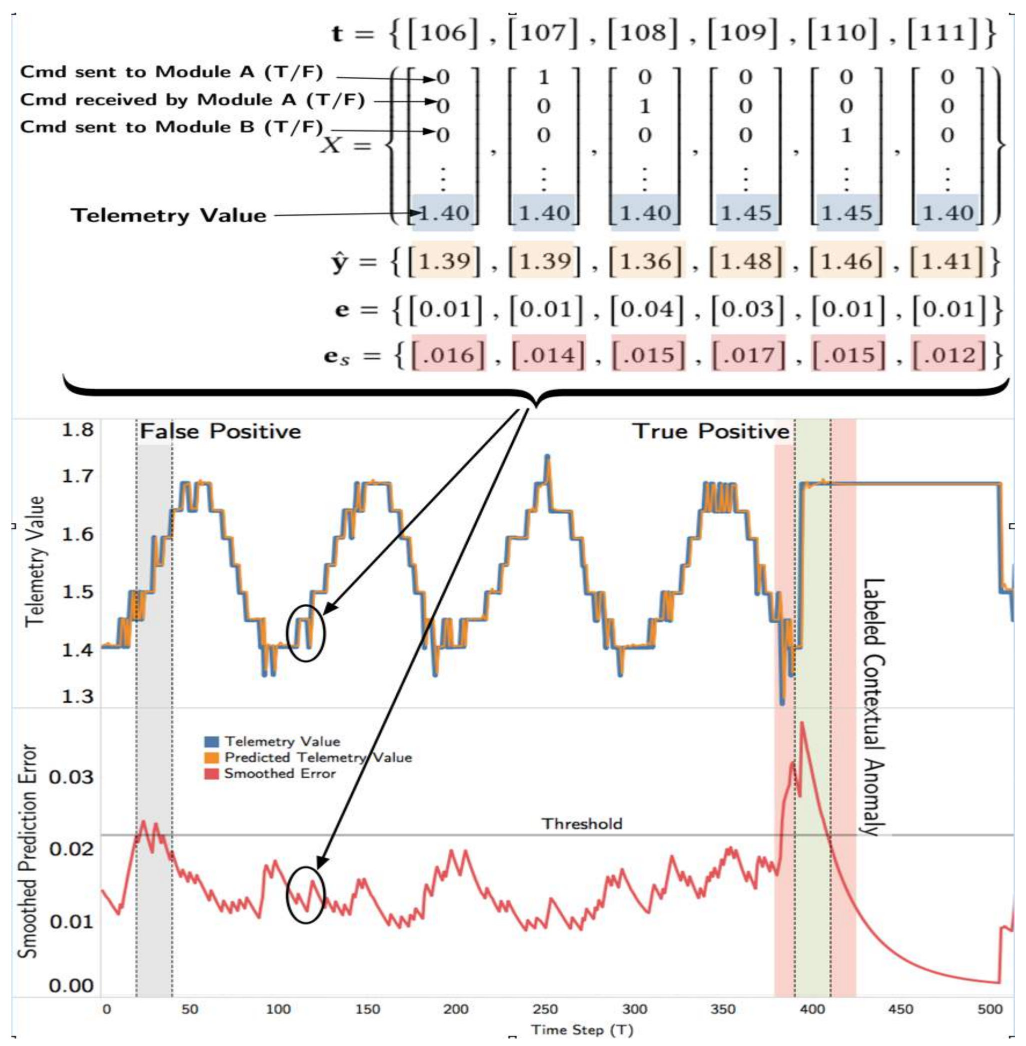

Smoothed errors $\mathbf{e}_s = [e_s^{(t-h)}, \ldots, e_s^{(t-l_s)}, \ldots, e_s^{(t-1)}, e_s^{(t)}]$

Candidate thresholds $\boldsymbol{\epsilon} = \mu(\mathbf{e}_s) + \mathbf{z}\sigma(\mathbf{e}_s)$

Threshold $\epsilon = argmax(\boldsymbol{\epsilon}) = \dfrac{\Delta\mu(\mathbf{e}_s)/\mu(\mathbf{e}_s)) + (\Delta\sigma(\mathbf{e}_s)/\sigma(\mathbf{e}_s)}{n(\mathbf{e}_a) + n(\mathbf{E}_{seq})^2}$

Definitions
$\Delta\mu(\mathbf{e}_s) = \mu(\mathbf{e}_s) - \mu(\{e_s \in \mathbf{e}_s | e_s < \epsilon\})$
$\Delta\sigma(\mathbf{e}_s) = \sigma(\mathbf{e}_s) - \sigma(\{e_s \in \mathbf{e}_s | e_s < \epsilon\})$
$\mathbf{e}_a = \{e_s \in \mathbf{e}_s | e_s > \epsilon\}$
$\mathbf{E}_{seq}$ = continuous sequences of $e_a \in \mathbf{e}_a$

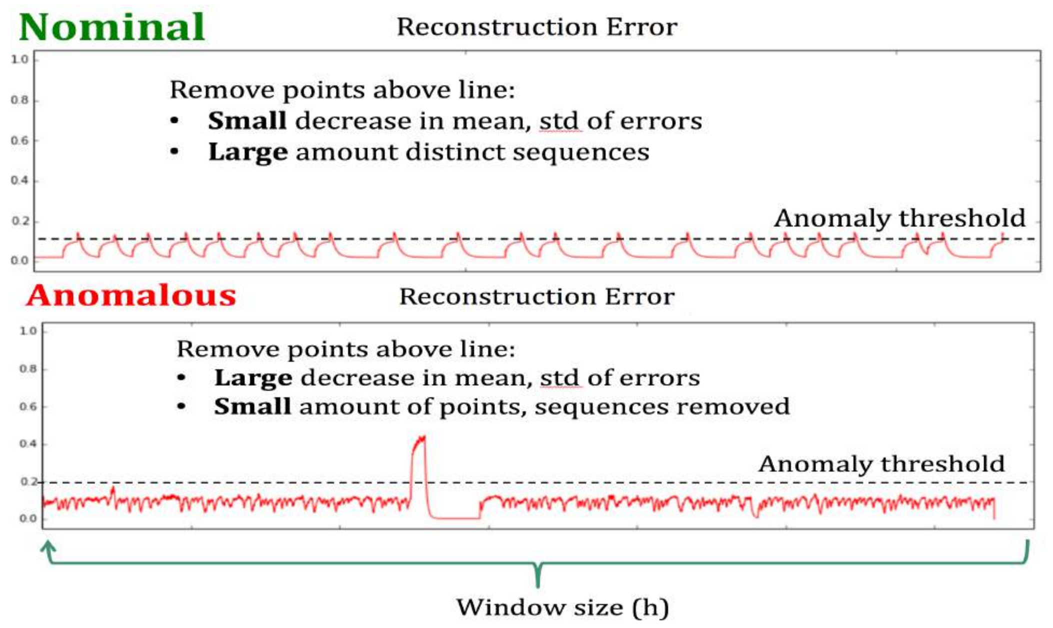

A threshold is found that, if all values above are removed, would cause the greatest percent decrease in the mean and standard deviation of the smoothed error.

The function penalizes for having larger numbers of anomalous values and sequences to prevent overly greedy behavior.

The highest smoothed error in each sequence of anomalous errors is given a normalized score based on its distance from the chosen threshold.

Pruning procedure (outlined in [1]) helps mitigate false positives, limit memory and compute cost, and ensure anomalous sequences are not the result of regular noise within a stream.

[1] Detecting Spacecraft Anomalies Using LSTMs and Nonparametric Dynamic Thresholding: https://arxiv.org/pdf/1802.04431

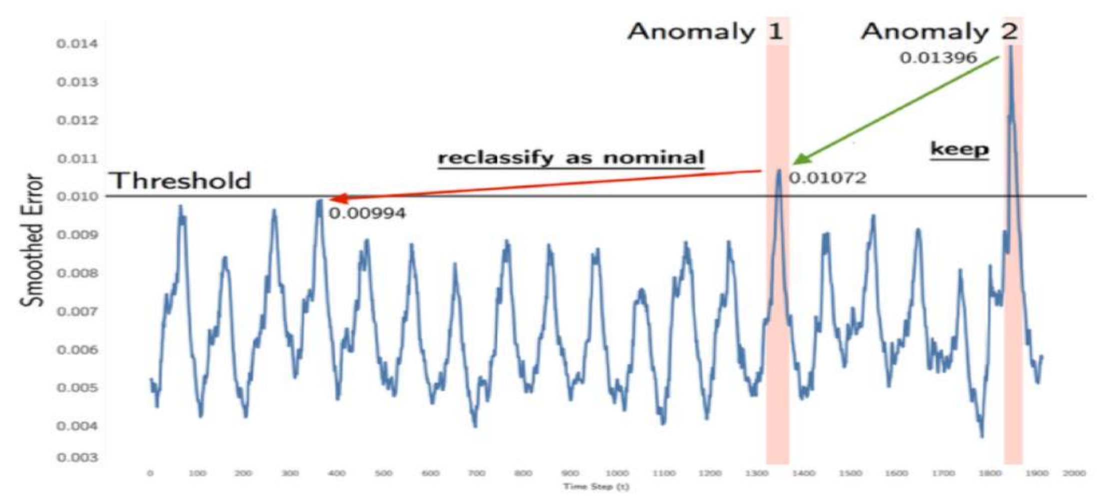

KEY PRELIMINARY RESULTS

	SMAP	MSL	Total
Total anomaly sequences	69	36	**105**
Point anomalies (% tot.)	43 (62%)	19 (53%)	**62 (59%)**
Contextual anomalies (% tot.)	26 (38%)	17 (47%)	**43 (41%)**
Unique telemetry channels	55	27	**82**
Unique ISAs	28	19	**47**
Telemetry values evaluated	429,735	66,709	**496,444**

Model Parameters

hidden layers	2
units in hidden layers	80
sequence length (l_s)	250
training iterations	35
dropout	0.3
batch size	64
optimizer	Adam
input dimensions	25 (SMAP), 55 (MSL)

Experiment to evaluate system looked at historical telemetry values for SMAP and MSL missions, and tried to detect ~115 expert-labeled anomalies (ISAs).

Performance was evaluated for different anomaly types.

Detected 69% of “contextual” anomalies (those that are not detectable by thresholds)

	Recall - *point*	Recall - *contextual*
MSL	78.9%	58.8%
SMAP	95.3%	76.0%
Total	90.3%	69.0%

Thresholding Approach	Precision	Recall	$F_{0.5}$ score
Non-Parametric w/ Pruning (p = 0.13)			
MSL	92.6%	69.4%	**0.69**
SMAP	85.5%	85.5%	**0.71**
Total	87.5%	80.0%	**0.71**

80% of all anomalies were detected.

Our nonparametric dynamic thresholding method outperformed gaussian tail method for thresholding.

Pruning decreased overall recall by 4.8%, but increased overall precision by 38.6%

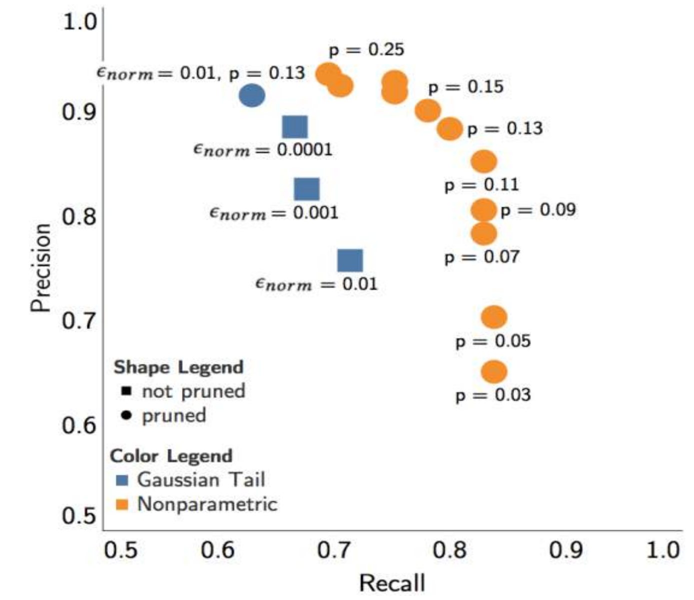

FUTURE WORK

This pilot was a key step in establishing that a large-scale telemetry monitoring system is feasible. Future work will focus on extending this system to other missions, and improving the telemetry predictions primarily through improved feature engineering.

ACKNOWLEDGEMENTS

This effort was supported by the Office of the Chief Information Officer (OCIO) at JPL, managed by the California Institute of Technology on behalf of NASA.