Detecting Spacecraft Anomalies Using LSTMs and Nonparametric Dynamic Thresholding

Kyle Hundman
NASA Jet Propulsion Laboratory
California Institute of Technology
kyle.a.hundman@jpl.nasa.gov

Valentino Constantinou
NASA Jet Propulsion Laboratory
California Institute of Technology
vconstan@jpl.nasa.gov

Christopher Laporte
NASA Jet Propulsion Laboratory
California Institute of Technology
christopher.d.laporte@jpl.nasa.gov

Ian Colwell
NASA Jet Propulsion Laboratory
California Institute of Technology
ian.colwell@jpl.nasa.gov

Tom Soderstrom
NASA Jet Propulsion Laboratory
California Institute of Technology
tom.soderstrom@jpl.nasa.gov

Jet Propulsion Laboratory
California Institute of Technology

https://github.com/khundman/telemanom

Why LSTMs?

High data complexity
Less people in flight operations, more data
Limitations of current methods (e.g. thresholds)

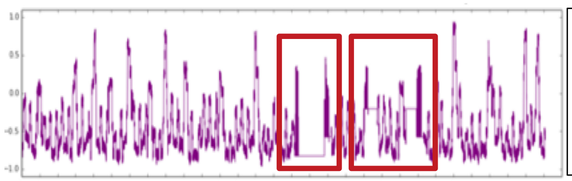

Example of telemetry anomalies not detectable using thresholds or distribution-based methods.

Predicting future values using recent data

Model inputs at step t

$$X = \left\{ \begin{bmatrix} x_1^{(t-l_s)} \\ x_2^{(t-l_s)} \\ \vdots \\ \boxed{x_m^{(t-l_s)}} \end{bmatrix}, \ldots, \begin{bmatrix} x_1^{(t-1)} \\ x_2^{(t-1)} \\ \vdots \\ \boxed{x_m^{(t-1)}} \end{bmatrix}, \begin{bmatrix} x_1^{(t)} \\ x_2^{(t)} \\ \vdots \\ \boxed{x_m^{(t)}} \end{bmatrix}, \begin{bmatrix} x_1^{(t+1)} \\ x_2^{(t+1)} \\ \vdots \\ \boxed{x_m^{(t+1)} = y^{(t)}} \end{bmatrix} \right\}$$

Telemetry Values

$l_p = 1$

$$e^{(t)} = \left[\hat{y}^{(t)} - y^{(t)}\right]$$

$$\mathbf{e} = \left[e^{(t-h)}, \ldots, e^{(t-l_s)}, \ldots, e^{(t)}\right]$$

h = length of historical window of errors (used to evaluate new errors)
ls = sequence length (# time steps to pass to model)

Prediction example

- High prediction error = abnormal behavior
- One model for each channel to maintain anomaly interpretability
- Command data same for all single-channel models

$$\mathbf{t} = \{[106], [107], [108], [109], [110], [111]\}$$

$$X = \left\{ \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 1.40 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 1.40 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 1.40 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 1.45 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ \vdots \\ 1.45 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 1.40 \end{bmatrix} \right\}$$

Cmd sent to Module A (T/F)
Cmd received by Module A (T/F)
Cmd sent to Module B (T/F)
Telemetry Value

One-hot encoded command information

$$\hat{\mathbf{y}} = \{[1.39], [1.39], [1.36], [1.48], [1.46], [1.41]\}$$

$$\mathbf{e} = \{[0.01], [0.01], [0.04], [0.03], [0.01], [0.01]\}$$

$$\mathbf{e}_s = \{[.016], [.014], [.015], [.017], [.015], [.012]\}$$

Smoothed prediction error

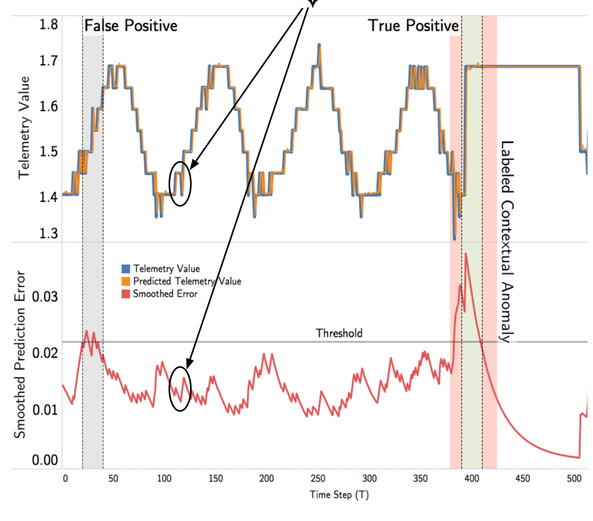

Actuals (blue) and prediction (orange)

Smoothed prediction error
High error = anomalies

Finding thresholds

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

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

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 $e_a \in \mathbf{e}_a$

Results

	Recall - *point*	Recall - *collective*
MSL	80.0%	42.1%
SMAP	97.7%	79.2%
Total	91.3%	62.8%

Collective anomalies are those that are not detectable by thresholds (0% recall)

Thresholding Approach	Precision	Recall	F_1 score
Non-Parametric w/ Pruning (p = 0.1)			
MSL	50.9%	63.6%	0.57
SMAP	62.6%	91.2%	0.74
Total	58.4%	80.4%	**0.68**

80% of all anomalies were identified (~115 in total)

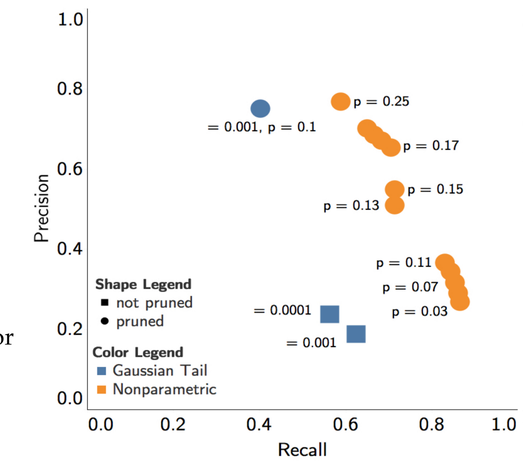

Thresholding performance compared to Gaussian tail method

p = pruning parameter (see paper)