

Satellite telemetry anomaly detection

Bc. Vojtěch Orava*

Abstract

Humankind has launched numerous satellites over the past decades, all operating in harsh space environments where anomalies in telemetry may indicate impending system failures. This paper addresses anomaly detection in satellite telemetry using deep learning methods applied to multivariate time-series data from onboard sensors. Several architectures were evaluated (LSTMs, autoencoders, convolutional and residual networks) employing both forecasting- and reconstruction-based approaches. Due to the problems of existing datasets, a synthetic telemetry generator was developed to simulate realistic sensor- and environment-driven anomalies. The proposed models were also assessed with respect to computational efficiency of edge platforms. Anomaly Explanation and Report System (AERS) extends the detection pipeline with a modular framework that enables domain experts to implement custom anomaly explanation methods and generate structured reports.

*xorava02@stud.fit.vut.cz, Faculty of Information Technology, Brno University of Technology

1. Introduction

In the context of a satellite, telemetry refers to the continuous stream of signals and time-series data gathered from its onboard sensors. These data enable engineers to remotely monitor the satellite's health and operational status. Much like a car that transmits telemetry such as throttle position, speed, or engine temperature, a satellite provides telemetry data including battery voltage, attitude (orientation), and communication signal strength.

Since launching a satellite into orbit is extremely costly, early detection of failures is essential. From a technological perspective, anomaly detection can be performed either on the ground, where telemetry is transmitted to Earth for processing, or directly onboard the spacecraft. Ground-based solutions allow the use of computationally intensive models, such as transformer architectures, without strict hardware limitations. However, onboard detection reduces latency, lowers communication demands, and increases system autonomy. This approach introduces significant constraints, including limited memory, restricted computational power, and strict inference time requirements.

Artificial intelligence-based anomaly detection methods typically follow two main paradigms: **forecasting-based** and **reconstruction-based** approaches. Forecasting models predict future telemetry values from

a window of past observations, while reconstruction-based models learn the distribution of nominal behavior and identify anomalies through elevated reconstruction error. Recurrent neural networks, particularly LSTM and GRU architectures, are commonly employed due to their suitability for sequential data. A well-known example is the **Teleanom model** introduced by NASA in 2018 [1], which applies LSTM-based forecasting for spacecraft telemetry anomaly detection.

Existing studies [2, 3, 4] focus primarily on detection accuracy, while onboard deployment constraints receive less attention. In addition, widely used benchmark datasets provided by space agencies (ESA and NASA) are significantly limited due to anonymization of channel values, which complicates interpretability of anomaly detection models. This work addresses these limitations by focusing on lightweight deep learning models suitable for onboard execution and by extending the detection framework with additional components that support realistic synthetic data generation and structured anomaly interpretation.

2. Spacecraft telemetry

Satellite telemetry consists of multivariate time series (where each time series is referred to as a channel). Compared to other domains, satellite telemetry data

typically contains very few anomalies relative to the overall dataset size. Other inherent challenges include high dimensionality and volume, complex inter-channel dependencies, and data errors caused by the harsh space environment [5].

In general, anomalies in satellite telemetry can be categorized into two main types. The first category includes **sensor-related anomalies**, such as hardware malfunctions, measurement drift, calibration errors, or permanent damage to onboard instruments. The second category consists of **environmental anomalies**, which arise from external conditions affecting the satellite, including radiation spikes, temperature extremes, space weather events, or unexpected interactions with the surrounding environment [5, 6, 7].

Due to the limited documentation and anonymization of channels in existing open datasets (such as NASA SMAP/MSL [6], ESA AD [5], and OPS-SAT [7]), a configurable synthetic telemetry dataset generator was developed [8]. This open-source tool allows users to generate datasets over a specified number of days and sampling frequencies. Anomalies can be generated (the number of them can be specified) or created based on the anomaly plan CSV file – the user can specify which, when and how long anomalies he wants to include. The user can also set how much environment-based anomalies should be generated.

3. Anomaly detection

To identify a suitable neural network architecture for onboard anomaly detection, multiple versions of dense autoencoders, LSTM autoencoders, ResNets, and Temporal Convolutional Neural Networks (TCNN) were evaluated. All models were tested on a synthetic dataset (86,400 timestamps across 20 channels) created using the above mentioned generator. An adapted version of Telemanom was utilized as a baseline.

Models were primarily evaluated using the **Area Under the Precision-Recall Curve (PR AUC)**. The developed LSTM-based autoencoder achieved the highest PR AUC (area under the precision–recall curve), reaching a value of **0.946**. Another well performing model was the classical (dense) autoencoder, which achieved a PR AUC of 0.870. Adapted version of Telemanom achieved 0.768. Forecasting-based models generally achieved better result in these tests.

However, experiments conducted on real mission data from the ESA AD [5] demonstrated that as the dimensionality increases (58 channels), forecasting approaches reach practical limitations. In these complex, high-dimensional scenarios, reconstruction-based anomaly

detection significantly outperformed the models that excelled on the simpler synthetic dataset.

Another important metric for space-related detection is inference time. Models were exported to ONNX and tested on three edge-computing platforms: Raspberry Pi 4B, AMD ZCU104, and NVIDIA Jetson Orin Nano. The fastest model (a dense autoencoder) achieved inference times under 0.01 seconds per 1000 input sequences (length 250) on the Jetson Nano. Inference power consumption was 13.04 W for AMD ZCU104, 8.02 W for Jetson Nano and around 3 – 5 W for Raspberry Pi 4.

4. Anomaly explanation

The last output of this work is Anomaly Explanation Report System (AERS) which is a Python-based modular wrapper for the explanation of anomalies and the generation of reports. It is available at GitHub [9]. AERS ingests anomaly predictions and raw telemetry data to generate standardized, interpretable reports (including JSON outputs). Its modular architecture allows domain experts to write custom analysis modules.

5. Conclusions

This work considers satellite anomaly detection for spaceflight. Real satellite data is often restricted or hard to understand, an open-source generator was built to create realistic test data. It was shown that while forecasting-based models work well for smaller systems (less channels), different approaches are needed to handle the complex realities of a real mission. Created detection models are fast and lightweight enough to run directly on a satellite's limited hardware. Finally, the development of the AERS framework ensures that these rapid onboard detections can be effectively interpreted by ground-based domain experts, fostering greater trust and autonomy in satellite operations.

Acknowledgements

I would like to thank my supervisor Ing. Tomáš Kašpárek, Ph.D. for his help and guidance. I also thank Zaitra s.r.o. (Bc. Martin Javorka and Ing. Daniel Kyselica) for consultations and for the opportunity to access specific hardware. Computational resources were provided by the e-INFRA CZ project (ID:90254), supported by the Ministry of Education, Youth and Sports of the Czech Republic.

References

- [1] Kyle Hundman, Valentino Constantinou, Christopher Laporte, Ian Colwell, and Tom Söderström. Detecting spacecraft anomalies using lstms and

nonparametric dynamic thresholding. *CoRR*, abs/1802.04431, 2018.

- [2] James Murphy, Maria Buckley, Léonie Buckley, Adam Taylor, Jake O'Brien, and Brian Mac Namee. Deploying machine learning anomaly detection models to flight ready ai boards. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 6828–6836, 2024.
- [3] James Murphy, John Ward, and Brian Mac Namee. An overview of machine learning techniques for onboard anomaly detection in satellite telemetry*. pages 1–6, 10 2023.
- [4] Asma Fejjari, Alexis Delavault, Robert Camilleri, and Gianluca Valentino. A review of anomaly detection in spacecraft telemetry data. *Applied Sciences*, 15(10), 2025.
- [5] Krzysztof Kotowski, Christoph Haskamp, Jacek Andrzejewski, Bogdan Ruszczak, Jakub Nalepa, Daniel Lakey, Peter Collins, Aybike Kolmas, Mauro Bartesaghi, Jose Martinez-Heras, and Gabriele De Canio. European space agency benchmark for anomaly detection in satellite telemetry, 2025.
- [6] astropat. Nasa anomaly detection dataset smap & msl, Oct 2022.
- [7] Bogdan Ruszczak, Krzysztof Kotowski, David Evans, and Jakub Nalepa. The ops-sat benchmark for detecting anomalies in satellite telemetry. *Scientific Data*, 12(1):710, Apr 2025.
- [8] Vojtěch Orava. Satellite telemetry generator with annotated anomalies, 2026. https://github.com/vorava/satellite_telemetry_generator.
- [9] Vojtěch Orava. Anomaly explanation report system, 2026. <https://github.com/vorava/aers>.