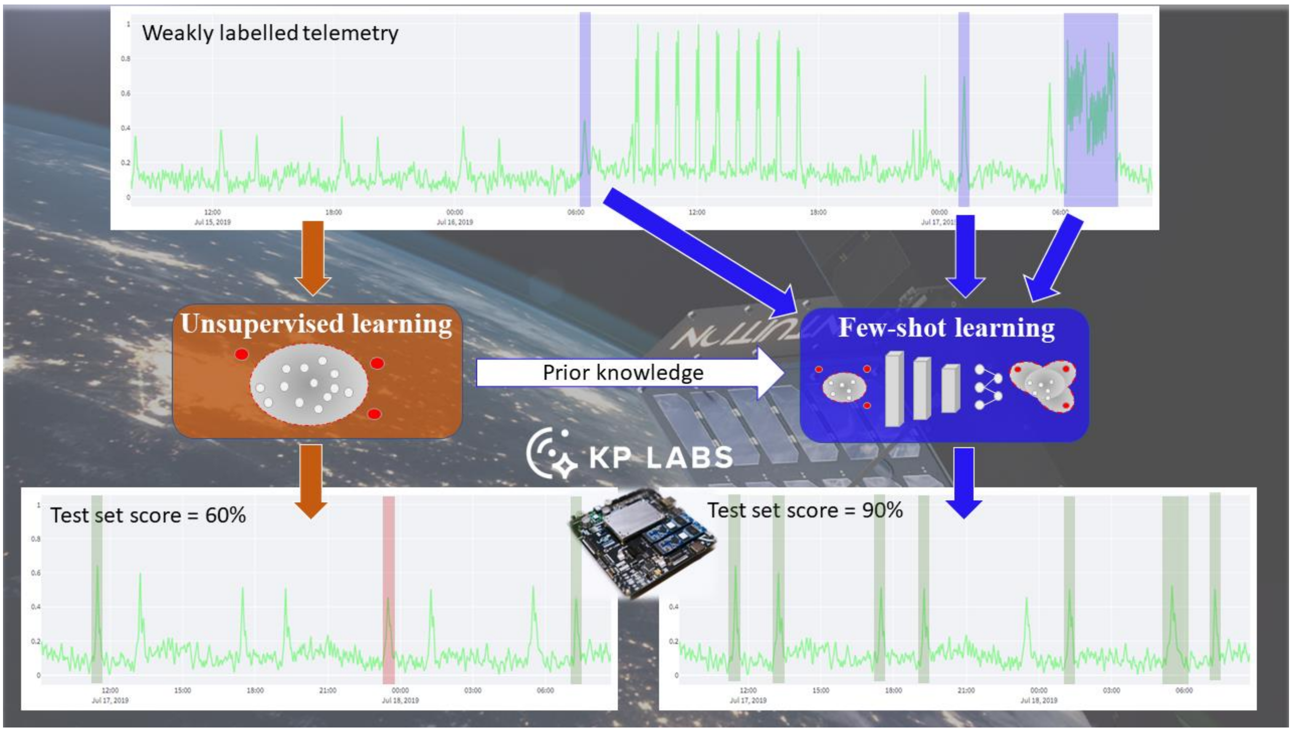

FEW-SHOT ANOMALY DETECTION IN SATELLITE TELEMETRY

Executive summary

Early technology development

OSIP Open Channel

Affiliation(s): KP Labs

Activity summary:

Anomaly detection in time series is an active area of research in which multiple classic and machine learning algorithms were proposed. However, satellite telemetry is a special case of time series characterized by high dimensionality, missing data, and – most importantly – limited information about known anomalies. The leading approach for time series anomaly detection (TSAD) in satellite telemetry is to establish a normality model based on available nominal telemetry periods and search for significant deviations from that model in the test data. Majority of the existing (deep and shallow) TSAD methods are not able to leverage prior knowledge (i.e., a few labelled anomalies) when such information is available. This technology gap can be addressed with a new approach called few-shot learning.

Publishing Date: 24.06.2024
Contract Number: 4000141301
Implemented as ESA Standard Procurement

Anomaly detection in time series is an active area of research in which multiple classic and machine learning algorithms were proposed. However, satellite telemetry is a special case of time series characterized by high dimensionality (up to thousands of parameters per satellite), missing data (idle states and varying sampling rates during a mission), and – most importantly – limited information about known anomalies. The leading approach for time series anomaly detection (TSAD) in satellite telemetry is to establish a normality model based on available nominal telemetry periods and search for significant deviations from that model in the test data. Majority of the existing (deep and shallow) TSAD methods are not able to leverage prior knowledge (i.e., a few labelled anomalies) when such information is available. This technology gap can be addressed with a new approach called few-shot learning.

In this activity, 4 distinct approaches for few-anomaly detection (FSAD) have been developed and validated on two different datasets: a set of real-life satellite telemetry fragments from the OPS-SAT mission and simulated telemetry from the CATS dataset (zenodo.org/records/8338435). The datasets were adjusted to the few-shot setting by defining new training splits based on the distribution of distinct anomaly classes. DeepSAD and DevNet approaches were validated on the tabular dataset containing handcrafted features calculated for segments of OPS-SAT telemetry. DevNet turned out to be more promising for this use-case, showing clear effects of few-shot learning even with very limited numbers of anomalous samples as presented in Figure 1. DeepSAD showed results better than the majority of standard supervised and unsupervised anomaly detection algorithms, but the effect of few-shot learning was not significant.

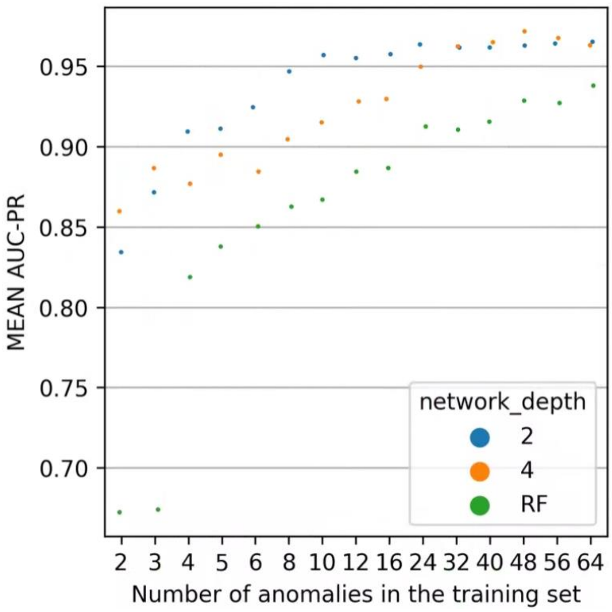

Figure 1. The impact of the number of anomalies on the performance of the few-shot DevNet (depths 2 and 4) and the supervised Random Forest (RF).

For purposes of time series data in CATS, the DC-VAE model was modified in two ways, by adding additional loss terms (DC-VAE-FPS) and by novel adaptation of memory modules for FSAD. Despite many efforts, DC-VAE-FPS did not provide satisfactory results and the conclusion is that it is not the best choice for this task. The effect of few-shot learning was somehow noticeable in the qualitative comparisons, but it was not reflected in the quantitative metrics. However, there is a small chance that using proper postprocessing or pruning routines for false alarms could improve the results in the future. Much more promising results were obtained in the proof of concept for the model-based FSAD using memory modules (MM). Qualitatively, it gives much better results for known anomalies while doing it faster and keeping the same or better performance for unknown ones. There was no time left to incorporate this idea into the actual DC-VAE pipeline, but it is a very promising future direction.

Table 1. summarizes the 4 implemented approaches in terms of 9 requirements defined for the target FSAD algorithm at the beginning of the project (using “**shall**” for 4 obligatory requirements and “**should**” for 5 desired requirements):

R1. The FSAD algorithm **shall** learn from as few as a 1 anomalous telemetry fragment per anomaly class (one-shot or few-shot anomaly detection) and orders of magnitude larger dataset of nominal satellite telemetry data.

R2. The FSAD algorithm **shall** achieve statistically significantly better performance for known anomaly classes than the equivalent algorithm without using FSL techniques (measured using metrics defined in the Methods section).

R3. The FSAD algorithm **shall** achieve at least as good performance for novel anomaly classes (i.e., unseen during training) as the equivalent algorithm without using FSL techniques (measured using metrics defined in the Methods section).

R4. The FSAD algorithm **shall** handle anomaly detection for at least 5 telemetry channels simultaneously.

R5. It **should** be possible to run the FSAD algorithm inference on-board satellite using Antelope or Leopard DPUs by KP Labs (see Section **Error! Reference source not found.**), or similar hardware.

R6. It **should** be possible to adapt the FSAD algorithm on-board satellite when new anomalies are encountered.

R7. The FSAD algorithm **should** be able to learn from individual outlying nominal samples (e.g., commanded manoeuvres or special events such as eclipses) to avoid detecting them in the future (one-shot learning).

R8. The FSAD algorithm **should** detect anomalies no later than the half of their duration.

R9. The FSAD algorithm **should** learn from unlabelled data if available.

Table 1. Analysis of 4 implemented FSAD algorithms according to the 9 functional requirements. The boldfaced IDs mark obligatory requirements. Brackets mark values based on the theoretical analysis of an algorithm that were not empirically verified.

Requirement		Approach			
ID	Short description	**DeepSAD**	**DevNet**	**DC-VAE-FPS**	**MM***
R1	Few samples	1	1	1	1
R2	Better for known anom.	0	1	0.5	1
R3	Not worse for unknown	1	1	0	1
R4	>5 channels	(1)	(1)	1	1
R5	On-board inference	(1)	(1)	(0.5)	(0.5)
R6	On-board adaptation	(0)	(0)	(0)	(1)
R7	Rare nominal events	0	0	0	0.5
R8	Early detection	Not applicable		1	1
R9	Learn from unlabelled	(1)	(1)	(0.5)	(0.5)

* - hypothetical solution for time series data based on our initial proof of concept

The approaches based on the feature representation learning with additional loss terms, namely DeepSAD and DC-VAE-FPS, failed to meet important requirements related to the performance of FSAD in comparison to baselines (R2 and R3). In both cases, there was no significant performance improvement for known anomalies. There were some visible positive qualitative effects for many samples for DC-VAE-FPS (hence, the value of 0.5 in the Table 1) but they were not reflected in the most important quantitative performance metric for space operations, i.e., the event-wise F0.5-score. DeepSAD and DevNet are designed for tabular data, so they can theoretically handle also multiple channels of telemetry, but it was not verified empirically. We have not been explicitly focusing on learning rare nominal events or from unlabelled data, but based on theoretical analysis (and a few experiments for selected algorithms) all the proposed algorithms should be able to learn from unlabelled data assuming a small proportion of anomalies in them. Learning from rare nominal events is potentially possible using memory modules and it was proved experimentally, but we did not assess it quantitatively in this study.

All the proposed algorithms are good candidates to be implemented on-board satellites using Antelope or Leopard DPUs by KP Labs, i.e., in the future ESA OPS-SAT Volt flying laboratory. DeepSAD and DevNet use simple fully connected networks widely supported by TensorFlow Lite (TFLite) and Vitis AI. DeepSAD is implemented in PyTorch, but the model itself can be converted to desired format using the ONNX framework. DC-VAE and its few-shot versions, despite using specific 1D dilated causal convolutions, are compatible with version 2 of TFLite and were successfully compiles for on-board applications. The proposed prototype of the model-based few-shot anomaly detection with memory modules seems to be very promising approach in practical satellite applications, especially on-board. It is worth exploring in the future. KP Labs plans to further explore this approach in the ongoing ESA-funded project "On-board continual learning in SatCom systems" realized in consortium with OHB Hellas and Eutelsat OneWeb.

Proposed solutions have also a wide range of potential applications outside the space domain. They can be easily adapted to any anomaly detection use-case in any domain in which the scarcity of annotated anomalies or outliers is the key obstacle to achieve acceptable performance of the machine learning model. Examples include quality control in manufacturing, fraud detection, network security, and medical screening tests.

We assess the current Technology Readiness Level of the proposed solutions to 3. To increase the TRL to 5/6, future works should implement the proposed approaches in a relevant environment, validate them on a wider range of real-life use-cases, and establish test-time adaptation procedures when new knowledge is available. It would be also crucial to collect and address a feedback from end users, i.e., spacecraft operations engineers, on the usability of the system.