Machine learning detects anomalies in OPS-SAT telemetry*

Bogdan Ruszczak[1,3][0000−0003−1089−1778], Krzysztof Kotowski[3][0000−0003−2596−6517], Jacek Andrzejewski[3], Alicja Musiał[3], David Evans[4], Vladimir Zelenevskiy[4,5], Sam Bammens[4], Rodrigo Laurinovics[4,6], and Jakub Nalepa[2,3][0000−0002−4026−1569]

[1] Faculty of Electrical Engineering, Automatic Control and Informatics, Department of Informatics, Opole University of Technology, Prószkowska 76, 45-758 Opole, Poland
b.ruszczak@po.edu.pl

[2] Faculty of Automatic Control, Electronics and Computer Science, Department of Algorithmics and Software, Silesian University of Technology, Akademicka 16, 44-100 Gliwice, Poland
jnalepa@ieee.org

[3] KP Labs, Konarskiego 18C, 44-100 Gliwice, Poland
{bruszczak, jandrzejewski, amusial, jnalepa, kkotowski}@kplabs.pl

[4] European Space Agency/ESOC, Robert-Bosch-Str. 5, 64293 Darmstadt, Germany
{David.Evans, Sam.Bammens}@esa.int, {Rodrigo.Laurinovics, Vladimir.Zelenevskiy}@ext.esa.int

[5] Telespazio Germany GmbH, Europaplatz 5 64293 Darmstadt, Germany

[6] IrbGS ltd, "Viraki", Ance parish, Irbene, Latvia, LV-3601

Abstract. Detecting anomalies in satellite telemetry data is pivotal in ensuring its safe operations. Although there exist various data-driven techniques for the task of determining abnormal parts of the signal, they are virtually never validated over real telemetries. Analyzing such data is challenging due to its intrinsic characteristics, as telemetry may be noisy and affected by incorrect acquisition, resulting in missing parts of the signal. In this paper, we tackle this issue and propose a machine learning approach for detecting anomalies in single-channel satellite telemetry. To validate its capabilities in a practical scenario, we build a dataset capturing the nominal and anomalous telemetry data captured on board OPS-SAT—a nanosatellite launched and operated by the European Space Agency. Our extensive experimental study showed that the proposed algorithm offers high-quality anomaly detection in real-life satellite telemetry, reaching 98.4% accuracy over the unseen test set.

Keywords: machine learning · anomaly detection · feature engineering · satellite telemetry.

* This work was partially supported by the: "On-board Anomaly detection from the OPS-SAT telemetry using deep learning" project funded by the European Space Agency under contract No. 4000137339/22/NL/GLC/ov. JN was supported by the Silesian University of Technology Rector's grant (02/080/RGJ22/0026).

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

1 Introduction

Many critical systems are being monitored using various sensors for detecting their malfunction [11,20]. This is extremely important for satellites to ensure their safe operations and to appropriately respond to any abnormal events that are reflected in the telemetry data. In general, there are three types of anomalous events which are commonly observed in satellite telemetry. In point anomalies, telemetry signal values fall outside the nominal operational range, and hence can be easily detected using out-of-limit checks. On the other hand, collective anomalies correspond to the overall sequences of consecutive telemetry values that are abnormal, whereas in contextual anomalies, the single telemetry values are anomalous within their local neighborhood [17]. There exist data-driven algorithms for detecting such types of abnormal events [18], spanning across classic techniques exploiting expert systems [25], unsupervised approaches [8,14] and deep learning models [10] (often benefiting from recurrent neural networks which are especially well-fitted to process time-series data [1,9]). Such algorithms, however, are virtually never validated over real-life telemetry data. Also, they often require long time-series data to build a model reflecting the nominal operation of the satellite—capturing it on board is tedious and time-consuming, thus data-level digital twins have been blooming to simulate the correct telemetry [2]. Benchmark datasets commonly exploited to validate detection algorithms contain time-series data, where each time series is split into its training and test parts, presenting similar characteristics. Such data is not affected by the practical challenges commonly observed in on-board telemetry, such as data noisiness or missing data, due to e.g., inappropriate signal acquisition. Therefore, the estimated anomaly detection capabilities of data-driven techniques may easily become over-optimistic, and the experimental scenarios are often flawed by methodological issues in the field [23].

In this paper, we tackle the issue of thorough validation of machine learning techniques for anomaly detection in satellite telemetry data. We not only propose an end-to-end pipeline for detecting such abnormal events in signal data. We also investigate its capabilities over a dataset capturing real-life telemetry data captured on board OPS-SAT—a nanosatellite launched and operated by the European Space Agency (ESA). OPS-SAT is a *flying laboratory*, providing a platform for on-board experiments (e.g., those which may be too risky to be executed on other operational satellites [7]), with one of them being anomaly detection from telemetry data using the pipeline discussed in this work. Other experiments performed on-board OPS-SAT include, among others:

– on-board unsupervised machine learning for spacecraft autonomy [13,22],
– the compression of housekeeping telemetry [5],
– assessing the stability of the attitude during inertial pointing mode using on-board images of the sky [21],
– deploying and maintaining the MO/MAL ground infrastructure with this aircraft [15],
– testing the on-board thermal vacuum [12],

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

– and addressing fail-safety and redundancy design challenges [24].

Here, we build upon the available OPS-SAT telemetry data to build a representative dataset that is used to train and validate our anomaly detection techniques in a practical setting, reflecting real on-board telemetry challenges, e.g., missing data and aperiodicity of the signal [6]. Our experimental validation performed over a carefully bundled OPS-SAT telemetry dataset revealed that our approaches exploiting hand-crafted feature extractors (which are independent of the length of the signal and can effectively operate with missing data) and classic supervised learners offer high-quality anomaly detection, reaching 98.4% classification accuracy over the unseen test set.

This paper is structured as follows. In Section 2, we discuss the main objective tackled in this paper and present the OPS-SAT telemetry channels of interest, together with our end-to-end machine learning pipeline for detecting abnormal events from time-series data. Our experimental study is reported and discussed in Section 3. Section 4 concludes the paper and highlights the most exciting future research pathways which may emerge from our work.

2 Materials and Methods

In this section, we discuss the investigated OPS-SAT telemetry (Sect. 2.1), together with our approach for detecting anomalies in such data (Sect. 2.2).

2.1 Dataset

In this work, we investigate the telemetry signals indicated by the ESA OPS-SAT team as the most interesting from the operational point of view, and we downloaded such data from the WebMUST client telemetry directory available for OPS-SAT [4]. Those telemetry channels include the magnetometer readouts, alongside the photo diode (PD) values—the following signals are collected:

– **Magnetometer telemetry channels:** I_B_FB_MM_0 (CADC0872) (see its fragment rendered in Fig. 1), I_B_FB_MM_1 (CADC0873), I_B_FB_MM_2 (CADC0874),
– **PD channels:** I_PD1_THETA (CADC0884), I_PD2_THETA (CADC0886), I_PD3_THETA (CADC0888), I_PD4_THETA(CADC0890), I_PD5_THETA (CADC0892), I_PD6_THETA (CADC0894).

It is of note that there are several practical challenges can be identified while exploring telemetry data captured by experimental satellites. They encompass—but are not limited to—the following issues:

– **High fragmentation of telemetry signals**. The data was registered mainly during the important stages of the OPS-SAT mission.
– **Missing data**, reflected as a high number of "gaps" in the telemetry readouts.
– Presence of the **recurring parts with noisy fragments** or with an unusual number of peaks (not necessarily abnormal) in the signal.

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

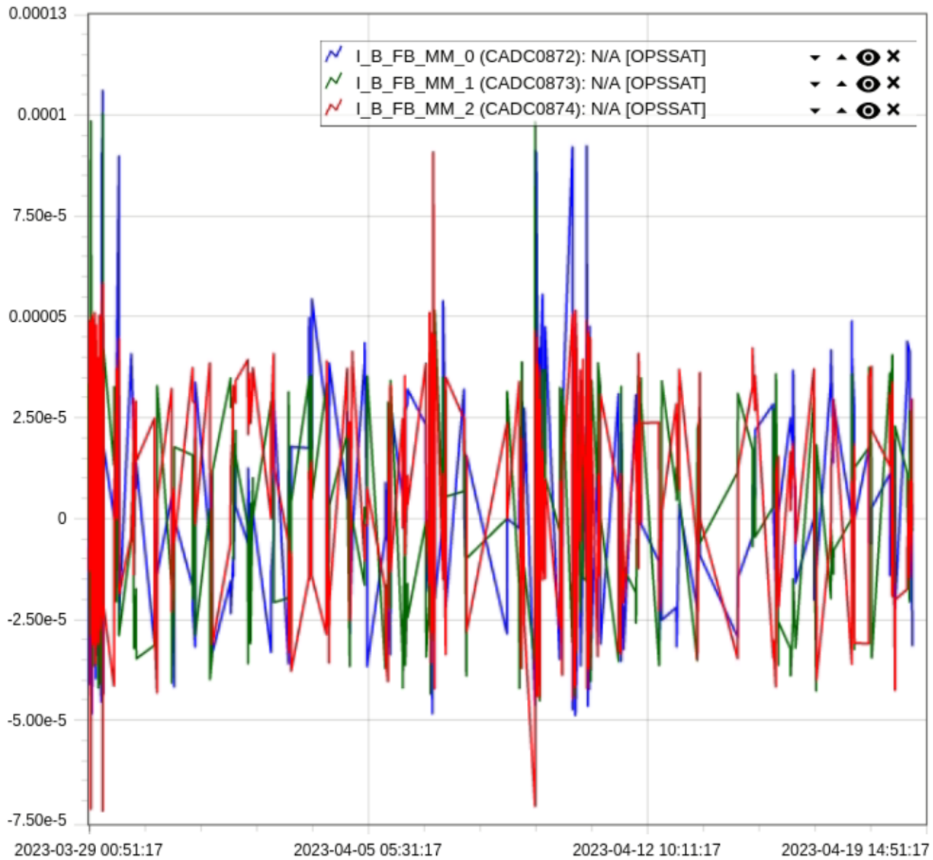

Fig. 1. A preview of three different OPS-SAT telemetries for a long (several weeks) timeframe. We can observe that the signal can easily become aperiodic and noisy, hence exploiting the approaches which focus on predicting the nominal signal and confronting it with the actual readouts [9] can easily become infeasible.

- **Non-uniform acquisition frequency rates** (the signal may be acquired hourly, daily, or even monthly), and there might be some radical changes in the signal characteristics along the analyzed months (due to e.g., various changes in satellite operations or hardware characteristics).
- **Lack of periodic and consistent nominal data** which would be captured over a long period of time.

The above-discussed telemetry channels were split into sub-parts, roughly corresponding to their periods, and they were labeled as nominal or anomalous by a human rater, and then the assigned labels were verified by the OPS-SAT

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

Operations Team using the OXI labeling system[7]. Overall, we collected **1273 nominal and 328 anomalous telemetry fragments for training**, and **416 nominal and 117 anomalous fragments for testing**, with the latter fragments, never seen during training (this training-test dataset split is stratified and it was generated randomly over all available telemetry fragments, to include approximately 25% of all telemetry fragments in the test set).

2.2 Detecting OPS-SAT Anomalies Using Machine Learning

In our machine learning anomaly detection pipeline, we **(1) extract an array of discriminative features from the telemetry fragments**[8], and such feature vectors are later **(2) fed to a supervised learner which performs prediction** (here, we tackle a binary classification task of determining abnormal vs. nominal telemetry fragments). For each telemetry segment, we extract the following 18 features (referred to as (a–p)), with some of them already proven to be discriminative for handling the classification tasks in telemetry data [16]:

– The features reflecting the **basic statistics** within the telemetry fragment (those include its *(a)* duration, *(b)* length, *(c)* mean value, *(d)* variance, and *(e)* standard deviation of a segment).
– The **number of peaks** within the fragment. More specifically, we sum the number of peaks with prominence of at least 10% of the segment value, and check this statistic for several segment representations, i.e., for *(f)* raw segment, alongside for its two variants after applying a smoothing transformation, using the "lighter" *(g)*, and "stronger" smoothing *(h)*, and for its *(i)* first and *(j)* second derivative.
– The **segment's variance** using its *(k)* first and *(l)* second derivative.
– The squared sum of *(m)* the **number of missing readouts**.
– The segment's *(n)* **weighted length** (relative to its sampling that also differs across different telemetry segments in the dataset), and *(o)* the **variance relative to segment's length** and *(p)* **to the segment's duration**.

Our extractors are independent of the length of the segment, and they can be effectively exploited to extract features from the segments with missing readouts—this may not be possible while operating over the raw telemetry data.

To tackle the issue of the relatively small amount and limited representativeness of OPS-SAT training samples, we introduce several data augmentation techniques specifically targeting the telemetry segments. Although they can be effectively used to synthesize both nominal and anomalous segments based on the actual samples, we augment the nominal examples only, in order to ensure that the abnormal signal segments were validated by the OPS-SAT Operations

[7] The OXI labeling system has been developed by KP Labs and it is available at `https://oxi.kplabs.pl`. OXI allows for not only investigating time-series data, together with the ground-truth information but also for generating ground truth.
[8] It is of note, however, that our approach can be applied to any time-series data, not only satellite telemetry.

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

Team (and they are *not* the result of synthetic data augmentation). We introduce the following data transformations to augment such telemetry samples:

- **Mirroring a segment over the OX axis** (ω_1)—a segment is flipped along the horizontal axis.
- **Mirroring a segment over the OY axis** (ω_2)—in a single telemetry channel, two consecutive nominal segments are flipped along their median value, starting from the half length of the first segment to the half of the second one.
- **Shifting a segment** (ω_3)—in a single telemetry channel, two consecutive nominal segments are shifted for a certain number of steps (steps of 15% and 25% of the segments' length were applied). The resulting segment starts at the shifted (by a given number of steps) starting position, and it finishes at the shifted ending position of the first segment. Hence, the last point of the first (shifted) segment will overlap the second segment—to avoid information leaks across the training and test sets, both neighboring segments that undergo this augmentation step must be originally included in the training set.

Finally, once the features are extracted and the training set has been potentially augmented, such nominal and anomalous examples are fed to the classifier to perform training. We may easily exploit any supervised learner in the proposed approach—this flexibility of our processing chain is shown in Section 3.

3 Experimental Validation

The main objective of our experimental study is to investigate the classification performance of the proposed classification engine for detecting anomalies in real-life OPS-SAT telemetry data. On top of that, we are aimed at verifying the flexibility of the processing chain, alongside the impact of improving the classification part of the pipeline working on a stable set of extracted features. We focus on classic classification models, including widely-adopted random forests, multi-layer perceptrons, adaptive boosting algorithms, k-nearest neighbors and support vector machines (SVMs) with a linear kernel function. To quantify the performance of the models, we calculate their accuracy, precision, recall and F_1 score, as well as the Matthews correlation coefficient (r_ϕ) over the unseen test set, with r_ϕ being the measure commonly used in imbalanced classification tasks [3]. All metrics should be maximized, with one indicating their perfect score. We split the experimental study into two experiments: in Experiment 1 (Section 3.1), we train the classification models over the original data only, whereas in Experiment 2 (Section 3.2) we benefit from data augmentation of nominal training samples.

3.1 Experiment 1: Exploiting Original Training Dataset

In the preliminary experiment, we optimized the most important hyperparameters of the classification models using a five-fold cross-validation procedure over the training set [19]. Thus, we ultimately had the maximum depth of random

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

forests of 20 and a maximum allowed number of features of 50% and 50 estimators; the size of the hidden layer in the multi-layer descriptor was 400; for adaptive boosting, we used 500 estimators; for SVMs, the regularization (C) parameter was set to 8; and $k = 5$ for k-nearest neighbors. The results obtained for such optimized models are gathered in Table 1 (they are sorted according to r_ϕ). Although precision exceeds 0.92 for the k-nearest neighbor classifier, we can observe rather low r_ϕ values across all investigated algorithms, indicating that the number of false positives (i.e., abnormal telemetry segments incorrectly classified as nominal) is large. This kind of error is, however, unacceptable in practice, as it could lead to missing anomalous events, e.g., in the on-board fault detection, isolation, and recovery system. On the other hand, a too large number of false negatives can deteriorate the usability of the system, if such incorrectly classified nominal segments are to be reviewed by the operations team before taking action in response to potential on-board anomalies. In practice, the latter issue is commonly a more severe problem, especially if the number of false alarms gets (very) large. In this work, we hypothesize that deploying data augmentation routines to synthesize nominal training examples may help deal with those issues by generating a larger and more representative training set that will allow us to elaborate well-generalizing models.

Table 1. The classification results obtained over the test set using machine learning models with default parameterization (sorted by r_ϕ). The best results are **boldfaced**.

Model	Accuracy	Precision	Recall	F_1 score	r_ϕ
Random forest	**0.9294**	**0.8826**	0.8050	**0.8233**	**0.7830**
Multilayer perceptron	0.8958	0.7814	**0.8077**	0.7688	0.7215
Adaptive boosting	0.9101	0.8138	0.7774	0.7646	0.7151
k-nearest neighbors	0.9051	0.9204	0.7495	0.7643	0.7133
SVM with linear kernel	0.8957	0.8249	0.7132	0.7369	0.6765

3.2 Experiment 2: Augmenting Training Datasets

Table 2. The summary of the segment data set for OPS-SAT telemetry anomalies.

Class	Training set	ω_1	ω_2	ω_3	Validation set
Nominal	1273	2677	703	1406	416
Anomalous	328	—	—	—	117

The augmentation methods can substantially increase the size of the training set (Table 2). Several examples of augmented training samples are visualized in Fig. 2, where each column represents one telemetry segment (marked with a solid line) and its augmented version (rendered as a dashed line). In this experiment,

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

we initially trained a random forest classifier (as it achieved the largest F_1 score in the previous experiment across all supervised models) using all versions of the augmented training sets. To better understand the impact of specific data augmentation routines on the overall performance of the classifier, we employed them separately (and then collectively), hence obtaining four augmented training sets. In Table 3, we gather the performance gain for each augmented dataset separately, and for all of them collectively (see the last row in this table). We can appreciate that the biggest improvement in the random forest's accuracy is obtained while exploiting ω_1 (mirroring training segments over the OX axis), as well as while utilizing all augmentation techniques ($\omega_1 \wedge \omega_2 \wedge \omega_3$). Here, the results are reported for the very same test set to ensure fair comparison.

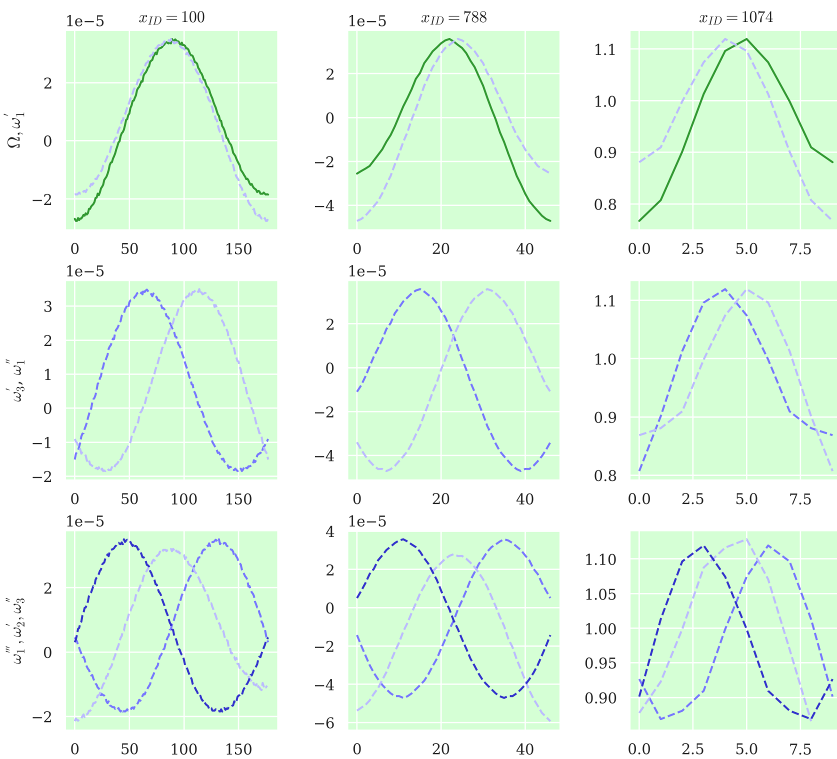

Fig. 2. The examples of the nominal telemetry segments (Ω, green solid lines), with their augmented versions with different parameterizations ($\omega_1, \omega_2, \omega_3$, blue and violet dashed lines). Each column provides one selected telemetry segment and its augmented versions. Some augmentations have been applied more than once, thus we have synthesized segments denoted as: $\omega_1^{'}, \omega_1^{''}, \omega_1^{'''}$. The presented segments are extracted from the following OPS-SAT telemetry channels: $x_{ID} = 100$ from *CADC0872*, $x_{ID} = 788$ from *CADC0874*, and $x_{ID} = 1074$ from *CADC0888*.

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

Table 3. The benefit of exploiting data augmentation, quantified as the gain in a metric when compared to the classifier trained over the original training set. The positive and negative values indicate an increase or decrease in the model's performance.

Augmentation	Accuracy	Precision	Recall	F_1 score	r_ϕ
ω_1	0.0491	0.1055	−0.0764	0.0109	0.0535
ω_2	0.0064	−0.0363	−0.0615	−0.0592	−0.0535
ω_3	0.0347	0.0870	−0.0521	−0.0054	0.0243
$\omega_1 \wedge \omega_2 \wedge \omega_3$	0.0548	0.1174	−0.0978	−0.0079	0.0393

Knowing that the latter approach results in the largest accuracy and precision gains (hence, we minimize the number of false positive alarms), we selected this augmented version of the training set for further experiments with all other supervised models. Here, we optimized their hyperparameters following the five-fold cross-validation procedure over the augmented training set, resulting in 400 neurons in the hidden layer of the multi-layer perceptron, 500 estimators in adaptive boosting, $C = 8$ for an SVM, and $k = 5$ for k-nearest neighbors. In Table 4, we summarize the classification measures obtained using all optimized classifiers, with the random forest model (max. depth of 25) exceeding the accuracy of 0.98, and the corresponding r_ϕ amounting to 0.826, significantly outperforming the random forest model trained over the original training set (Table 1).

To provide a comparison for those metrics, a deep learning algorithm [9] obtained precision: 0.855 and recall: 0.855 (it was evaluated for the SMAP Spacecraft data), and precision: 0.926 and recall: 0.694 (for the Curiosity telemetry dataset). Although we are aware that comparing the algorithms over different sets may easily become biased, we can indeed observe that our classification models were able to exceed the reported metrics for very challenging OPS-SAT telemetry data which indicates their significant generalization capabilities and robustness against noisy and difficult time-series data.

Table 4. The results (obtained over the test set) elaborated using all investigated classification models with optimized hyperparameters and trained over the augmented training set ($\omega_1 \wedge \omega_2 \wedge \omega_3$), sorted by r_ϕ. The best results are **boldfaced**.

Model	Accuracy	Precision	Recall	F_1	r_ϕ
Random forest	**0.9843**	**1.0000**	0.7317	**0.8239**	**0.826**1
Multi-layer perceptron	0.9792	0.8716	**0.7437**	0.7795	0.7815
Adaptive boosting	0.9812	0.9422	0.6858	0.7817	0.7893
SVM with linear kernel	0.9729	0.9097	0.6154	0.6683	0.6802
k-nearest neighbors	0.9491	0.8662	0.6637	0.7186	0.7008

Finally, for the best model (random forest), we present the confusion matrix in Fig. 3, which—besides presenting all measures (the number of true positives, true negatives, false positives, and false negatives: N_{TP}, N_{TN}, N_{FP}, and N_{FN}, respectively) brings the examples of correctly and incorrectly classified test seg-

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

ments. The classifier correctly identified 82% abnormal segments (96/117), and 99% (413/416) nominal segments, showing its potential practical utility for on-board anomaly detection. Additionally, we can observe a significant heterogeneity of the test segments, highlighting the difficulty of the classification task.

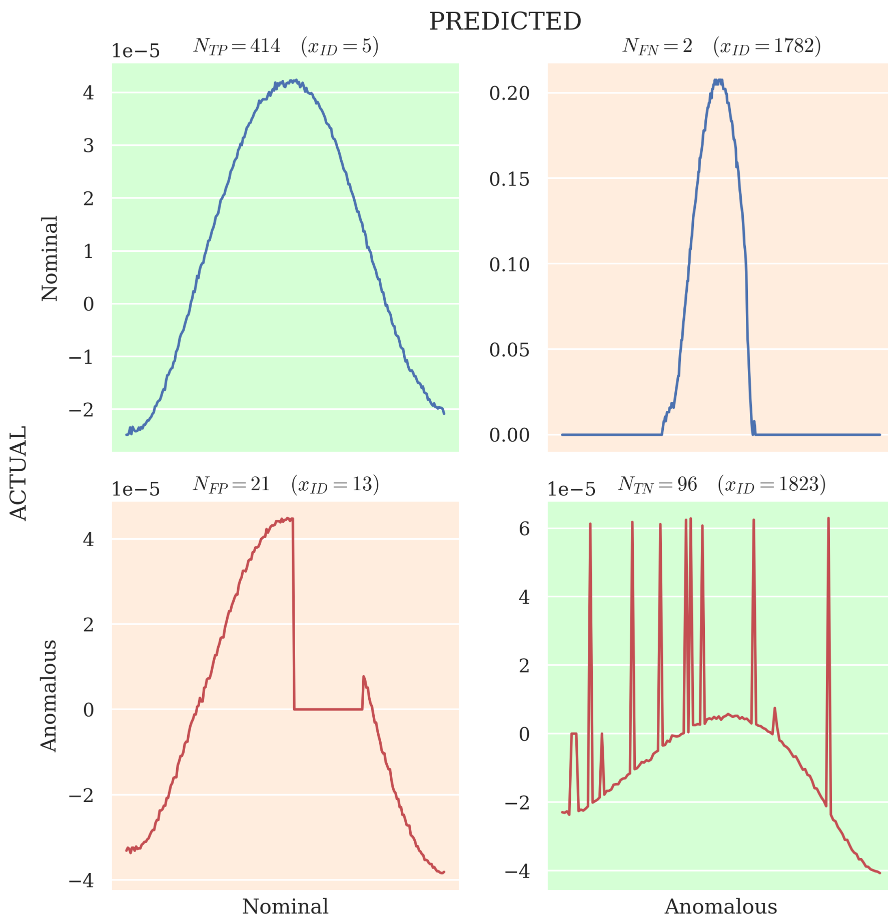

Fig. 3. The confusion matrix for the best classifier (random forest trained over the augmented training set with optimized hyperparameters), and the examples of correctly and incorrectly classified test segments (with the green and orange background, respectively). These example segments were extracted from the following OPS-SAT telemetries: $x_{ID} = 1782$ from *CADC0894*, $x_{ID} \in \{5, 13, 1823\}$ from *CADC0872*.

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

4 Conclusions and Future Work

Detecting anomalies in spacecraft telemetry data is of paramount practical importance to appropriately respond to various unexpected events that may happen on-board an operational satellite. Albeit there exist data-driven approaches toward this task, they are virtually never validated over telemetry channels captured by a real spacecraft. In this paper, we tackled this problem and proposed an end-to-end machine learning pipeline for detecting such abnormalities. Our approach benefits from hand-crafted feature extractors which are independent of the length of the telemetry segments which are later classified by a supervised learner as nominal or anomalous. To understand its generalization capabilities, we validated it over a curated (and validated by the ESA Operations Team) set of nominal and abnormal telemetry channels acquired on-board OPS-SAT—a nanosatellite operated by ESA. The experiments indicated that our technique offers high-quality detection of anomalies from OPS-SAT telemetry. Also, we showed that exploiting the suggested data augmentation routines allows for significantly improving the generalization capabilities of the classification models.

Our current research efforts are focused on deploying the proposed anomaly detection system on-board OPS-SAT. It is of note that the model has been already trained on the ground, and we are in the process of uplinking it to the satellite. Also, we are working on utilizing our method for other missions, to further prove its generalizability. Finally, it would be interesting to exploit unsupervised clustering over large-scale telemetry datasets while utilizing our feature extractors. This may help accelerate the process of generating ground truth for supervised models (as pre-classified telemetry segments would have to be verified by humans), and possibly uncover intrinsic characteristics of such time-series data that might have not been observed by the Operations Teams.

References

1. Benecki, P., Piechaczek, S., Kostrzewa, D., Nalepa, J.: Detecting anomalies in spacecraft telemetry using evolutionary thresholding and lstms. In: Proc. GECCO. p. 143–144. GECCO '21, ACM, New York, NY, USA (2021)
2. Castellani, A., Schmitt, S., Squartini, S.: Real-world anomaly detection by using digital twin systems and weakly-supervised learning. CoRR **abs/2011.06296** (2020), https://arxiv.org/abs/2011.06296
3. Chicco, D., Jurman, G.: The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. BMC Genomics **21**(1), 6 (Jan 2020)
4. ESA: WebMUST - web client for OPS-SAT directory (2021), https://opssat1.esoc.esa.int/webclient-must
5. Evans, D., Labrèche, G., Marszk, D., Bammens, S., Hernández-Cabronero, M., Zelenevskiy, V., Shiradhonkar, V., Starcik, M., Henkel, M.: Implementing the New CCSDS Housekeeping Data Compression Standard 124.0-B-1 (Based on POCKET+) on OPS-SAT-1. Small Satellite Conference (8 2022), https://digitalcommons.usu.edu/smallsat/2022/all2022/133

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

6. Evans, D.J., Ortega, M., Zeif, R., Sergert, T.: Ops-sat: Fdir design on a mission that expects bugs - and lots of them. In: SpaceOps 2016 Conference. American Institute of Aeronautics and Astronautics (2016). https://doi.org/10.2514/6.2016-2481
7. Fratini, S., Policella, N., Silva, R., Guerreiro, J.: On-board autonomy operations for OPS-SAT experiment. Applied Intelligence **52**(6), 6970–6987 (4 2022)
8. Gao, Y., et al.: An unsupervised anomaly detection approach for spacecraft based on normal behavior clustering. In: Proc. ICICTA. pp. 478–481 (2012)
9. Hundman, K., Constantinou, V., Laporte, C., Colwell, I., Söderström, T.: Detecting spacecraft anomalies using LSTMs and nonparametric dynamic thresholding. In: Proc. SIGKDD. pp. 387–395. ACM (2018)
10. Kacker, S., Meredith, A., Cahoy, K., Labrèche, G.: Machine Learning Image Processing Algorithms Onboard OPS-SAT. Small Satellite Conference (8 2022), https://digitalcommons.usu.edu/smallsat/2022/all2022/65
11. Kopp, M., Pevný, T., Holeňa, M.: Anomaly explanation with random forests. Expert Systems with Applications **149**, 113187 (2020)
12. Kubicka, M., Zeif, R., Henkel, M., Hörmer, A.J.: Thermal vacuum tests for the esa's OPS-SAT mission. Elektrotech. Informationstechnik **139**(1), 16–24 (2022)
13. Labreche, G., Evans, D., Marszk, D., Mladenov, T., Shiradhonkar, V., Soto, T., Zelenevskiy, V.: OPS-SAT Spacecraft Autonomy with TensorFlow Lite, Unsupervised Learning, and Online Machine Learning. In: IEEE Aerospace Conference Proceedings. vol. 2022-3. IEEE Computer Society (2022)
14. Li, K., et al.: A novel method for spacecraft electrical fault detection based on FCM clustering and WPSVM classification with PCA feature extraction. Proc. of the Institution of Mechanical Engineers **231**(1), 98–108 (2017)
15. Marszk, D., Evans, D., Mladenov, T., Labrèche, G., Zelenevskiy, V., Shiradhonkar, V.: MO Services and CFDP in Action on OPS-SAT. Small Satellite Conference (8 2022), https://digitalcommons.usu.edu/smallsat/2022/all2022/67
16. Nalepa, J., Andrzejewski, J., Ruszczak, B., Kotowski, K., Musiał, A., Evans, D., Zelenevskiy, V., Bammens, S., Laurinovics, R.: Toward On-Board Detection Of Anomalous Events From Ops-Sat Telemetry Using Deep Learning (2022), https://www.researchgate.net/publication/367091759
17. Nalepa, J., Myller, M., Andrzejewski, J., Benecki, P., Piechaczek, S., Kostrzewa, D.: Evaluating algorithms for anomaly detection in satellite telemetry data. Acta Astronautica **198**, 689–701 (2022)
18. Nassif, A.B., Talib, M.A., Nasir, Q., Dakalbab, F.M.: Machine learning for anomaly detection: A systematic review. IEEE Access **9**, 78658–78700 (2021)
19. Pedregosa, F., et al.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011)
20. Primartha, R., Tama, B.A.: Anomaly detection using random forest: A performance revisited. In: Proc. ICoDSE. pp. 1–6 (2017)
21. Segret, B., Bammens, S., Bras, S., Marszk, D., Shiradhonkar, V., Zelenevskiy, V., Evans, D.: On-board Images to Characterize a CubeSat's ADCS. Small Satellite Conference (8 2022), https://digitalcommons.usu.edu/smallsat/2022/all2022/174
22. Toussaint, F., Thomassin, J., Laurens, S.: ASTERIA In-Orbit Testing on OPSSAT: An On-Board Autonomous Orbit Control Solution Including Collision Risks Avoidance. Small Satellite Conference (8 2022), https://digitalcommons.usu.edu/smallsat/2022/all2022/103
23. Wu, R., Keogh, E.J.: Current time series anomaly detection benchmarks are flawed and are creating the illusion of progress (extended abstract). In: 2022 IEEE 38th International Conference on Data Engineering (ICDE). pp. 1479–1480 (2022)

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21

24. Zeif, R., Henkel, M., Hörmer, A., Kubicka, M., Wenger, M., Koudelka, O.: The redundancy and fail-safe concept of the OPS-SAT payload processing platform (3 2018), 69th International Astronautical Congress : IAC 2018
25. Zeng, Z., Jin, G., Xu, C., Chen, S., Zhang, L.: Spacecraft telemetry anomaly detection based on parametric causality and double-criteria drift streaming peaks over threshold. Applied Sciences **12**(4), 1–24 (2022)

ICCS Camera Ready Version 2023
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-35995-8_21