18th International Conference on Space Operations, Montreal, Canada, 26 - 30 May 2025.

SpaceOps-2025, ID # 334

Advancing spacecraft health monitoring and control with AI through continual learning

Gabriele De Canio[a]*, Rudy Semola[b], Valerio De Caro[b]

[a] *European Space Operations Centre, European Space Agency, Robert-Bosch-Str. 5, 64293 Darmstadt, Germany*
[b] *Department of Computer Science, University of Pisa, Largo B. Pontecorvo, 3, 56127 Pisa, Italy*
* Corresponding Author

Abstract

This paper presents what is thought to be the first development and validation of a Continual Learning paradigm for spacecraft health monitoring. It provides a case study demonstrating the feasibility and utility of Continual Learning as a Service in ground segment applications, enhancing computational efficiency and overall quality through effective and agile machine learning practices.

Keywords: Artificial Intelligence (AI), automation, satellite health monitoring, continual learning

1. Introduction

This paper introduces Continual Learning (CL) as a transversal paradigm that significantly enhances satellite health monitoring. Readers will gain an understanding of the fundamental concepts of CL and the challenges posed by non-stationary data, commonly found in multi-time series telemetry data. The document includes a case study on satellite health monitoring, focusing on state-of-the-art approaches for anomaly detection and forecasting. It evaluates CL for these tasks using European Space Agency's (ESA) public anomaly dataset, ESA-ADB [1-3]. Additionally, it discusses methodologies for effectively handling these challenges through agile machine learning (ML) models. Sec. 2 delves into the background of CL and the challenges posed by non-stationary environments, setting the stage for the subsequent analysis. In Sec. 3, we present a detailed case study on spacecraft health monitoring, focusing on adaptive approaches for time-series forecasting using the ESA-ADB Mission 1 Dataset [1-3]. Sec. 4 presents the results of our analysis. Finally, Sec. 5 and 6 summarize our findings and discuss potential future directions for further research, respectively.

2. Background

2.1 Overview and Motivation

The field of space operations is increasingly complex, demanding sophisticated tools to ensure the safe and continuous operation of spacecraft and ground stations. ESA's European Space Operations Centre (ESOC) is developing five AI and Data capabilities for ground segment engineering and operations to support the vision of operating twice as many missions by 2030 [4-5]. This involves spinning in innovation in technology and execution, solving real pain points in current processes, and developing user-centric solutions. One of the key capabilities is "Automated health monitoring and control" [5]. With the vision of going towards unified products that leverage AI to automate tedious and repetitive tasks while enhancing users' work, the Spacecraft Monitoring AI (SMAI) suite and the AI Investigation Assistant (AIIA) for Incident Investigation and Predictive Maintenance have been developed to support spacecraft and ground stations health monitoring, respectively [10]. SMAI comprises four key elements.

1. Anomaly Detection: identifying deviations from expected behavior to detect faults and anomalies.
2. Root Cause Investigation: analyzing anomalies to determine their underlying causes.
3. Time Series Forecasting: predicting future values of telemetry data to anticipate potential anomalies.
4. Trend analysis: identifying trends by looking at telemetry temporal evolution.

Monitoring spacecraft telemetry time series data for anomalies is a critical task for spacecraft operations engineers (SOEs) to ensure the safe and continuous operation of various satellites. While simple automatic anomaly detection systems exist, more complex anomalies require manual detection, which is both costly and prone to errors. Real-life spacecraft telemetry presents a particularly challenging example of multivariate time series, characterized by high dimensionality, large volume, and a complex network of dependencies between channels. The changing spacecraft modes and environments further complicate these relationships, causing visible shifts in data distribution and channel interactions. Additionally, the diverse types of channels, inherent noise, and measurement errors due to the harsh space environment add layers of complexity. These factors pose significant challenges for AI algorithms aimed at enhancing the automation of spacecraft operations, as demonstrated by recent efforts at the European Space Agency's European Space Operations Centre [4-10]. This paper focuses on the application of continual learning to improve the performance of all these four elements, demonstrating its potential to improve the efficiency, scalability, and effectiveness of spacecraft health monitoring systems.

2.2 Non-stationary Environments

The dynamic nature of human learning allows for continuous and robust knowledge acquisition and retention. In contrast, traditional machine learning models optimize their parameters based on historical data, assuming that the data distribution remains constant over time. This assumption limits the model's generalization to new, unseen data. In practical applications, raw data often follows a non-stationary distribution, leading to changes or drifts in data patterns. These drifts pose significant challenges for models trained on static datasets, resulting in performance degradation. For example, in satellite health monitoring, telemetry data such as temperature, voltage, and orientation angles can fluctuate due to varying environmental conditions, satellite maneuvers, and equipment degradation over time. Similarly, in predictive maintenance scenarios, drifts can arise from internal factors (e.g. sensor changes) and external factors (e.g. temperature variations). Traditional machine learning models struggle to adapt to these evolving patterns, underscoring the need for adaptable techniques that can effectively analyze and interpret data streams in real-time production environments.

2.3 Continual Learning

The constant drift of input data necessitates the continuous adaptation of trained models to maintain their performance against new distributions. Ideally, similar to human learning, we want a model capable of accumulating additional knowledge by reusing and complementing what has been learned in the past with new data. The concept of a continual, or lifelong, learning agent has been proposed in several works over the past few years [11-20]. However, there is no definitive solution to the problem, as proposals often encounter the plasticity-stability dilemma [15, 19-20]. This dilemma arises when a model's structure is modified to learn new data (plasticity), but performance drops dramatically without access to previously collected data, leading to catastrophic forgetting. The challenge lies in balancing the model's ability to learn new distributions without forgetting previous experiences [15, 19-20].

3. Case Study

3.1 Problem Domain

Spacecraft monitoring applications have traditionally relied on offline-trained supervised ML algorithms, where all training data is compiled before the training process. However, the dynamic space environment introduces distinct challenges. Factors like thermal noise, atmospheric conditions, and onboard noise can substantially alter data characteristics, leading to difficulties or failures in offline-trained models when confronted with novel and unfamiliar data distributions. A crucial drawback is that, since the training and deployment phases in offline learning are separate, the system does not learn during deployment. Consequently, every time there are significant changes in the input data, the system must be taken offline, re-trained, and adjusted to the new normality. This often leads to "catastrophic forgetting," where the model's performance on the original data degrades as it attempts to learn from the new data. As a result, most AI systems, especially on the ground, cannot automatically generate accurate predictions and adapt models to changing environments. Previous studies tackled these challenges and proposed effective solutions [7-8, 21]. Here, we demonstrate that CL can be a valid and successful paradigm, complementing shortcomings of previous approaches [7-8,21] using the ESA-ADB [1-3].

3.2 Methodology

Our analysis uses solely on the ESA-ADB Mission 1 (anonymized) dataset. This dataset presents several challenges, including:

- hard-to-spot anomalies;
- several huge outliers;
- low signal-to-noise ratio;
- monotonically non-decreasing signals;
- severe concept drift in channels from groups 4, 7, and 13 and Subsystem 6;
- visible seasonality with a long period length;
- overabundance of telecommands.

The Telemanom Anomaly Detector, a semi-supervised algorithm proposed by NASA engineers, is widely regarded as the most popular algorithm for anomaly detection in satellite telemetry [1-3]. Its core component is a Long Short-Term Memory (LSTM) [22] based Recurrent Neural network (RNN) [23] that forecasts a small number of time points (10 by default) for a single channel based on hundreds of preceding samples (250 by default) from multiple input channels. The mean absolute difference between the forecasted samples and the actual signal is used as an anomaly score, which is then thresholded using the non-parametric dynamic algorithm (NDT) to identify anomalies.

Time series forecasting is a critical component of SMAI, providing the foundation for anomaly detection by predicting expected values of telemetry data. Here, we focus on this problem.

3.2.1 Model Architecture

The models used for time series forecasting are Long Short-Term Memory (LSTM) networks [22]. The architecture varies based on the input data:

- For groups of channels: LSTM with 2 layers and 80 units.
- For subsystem or target channels: LSTM with 2 layers and 134 units.

3.2.2 Update Approaches

To evaluate the impact of continual learning, three update approaches are compared:

- Cumulative: This traditional approach involves retraining the model using all historical data plus the current month's data. It offers excellent long-term generalization but is computationally expensive.
- Naïve: This continual learning approach uses only the current month's training data, discarding all historical information. It is computationally efficient but has limited generalization capacity.
- Replay: This continual learning strategy maintains a fixed buffer of historical data alongside the current month's data. It balances computational efficiency and predictive accuracy. In this case, we arbitrarily chose 10% for the buffer in which data was randomly selected.

3.2.3 Evaluation Metrics

The performance of the time series forecasting models is evaluated using the Mean Squared Error (MSE). Additionally, memory usage and training time are measured to assess the computational efficiency of each approach.

3.2.4 Experimental Setup

In our experiments, the dataset is divided into monthly batches to simulate a real-world operational scenario, with models updated monthly. The training set comprises the first three weeks of each month's data, while the test set uses the last week. This setup allows for an iterative evaluation of the models' adaptability and generalization.

4. Results

4.1 Efficiency of Update Approaches

The efficiency of the update approaches is evaluated based on memory usage. The results, shown in Fig. 1, 2, and 3, demonstrate that continual learning strategies significantly reduce memory overhead compared to the traditional cumulative approach.

- Naïve: This approach shows the highest memory savings, averaging 90.45% reduction, as it only retains the current training batch.
- Replay: Both fixed and variable memory replay strategies also demonstrate substantial memory savings, averaging around 72% reduction.

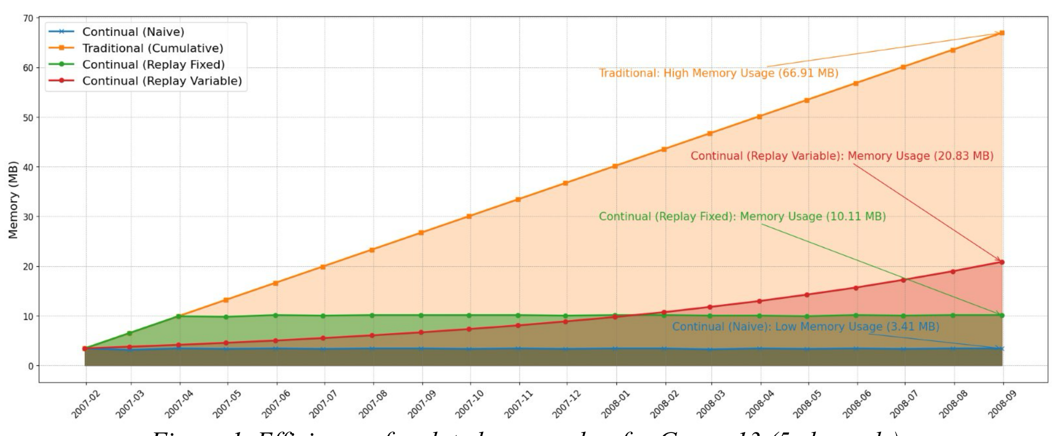

Figure 1. Efficiency of updated approaches for Group 13 (5 channels).

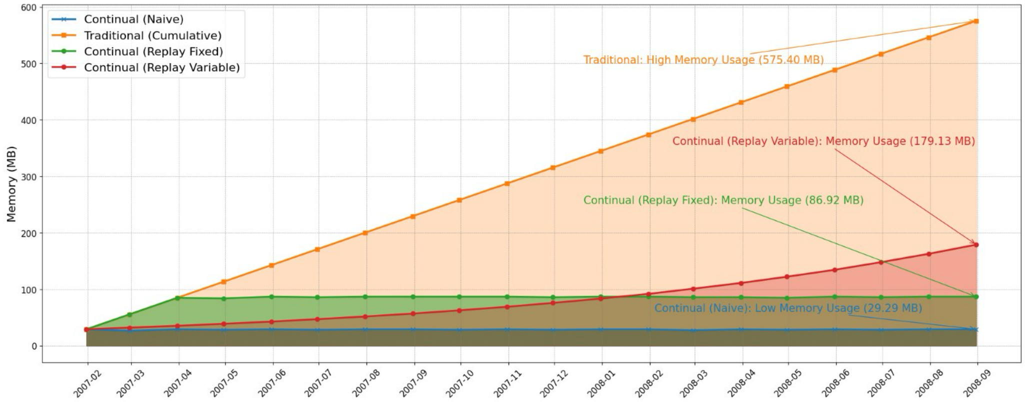

Figure 2. Efficiency of updated approaches for Subsystem 6 (42 channels).

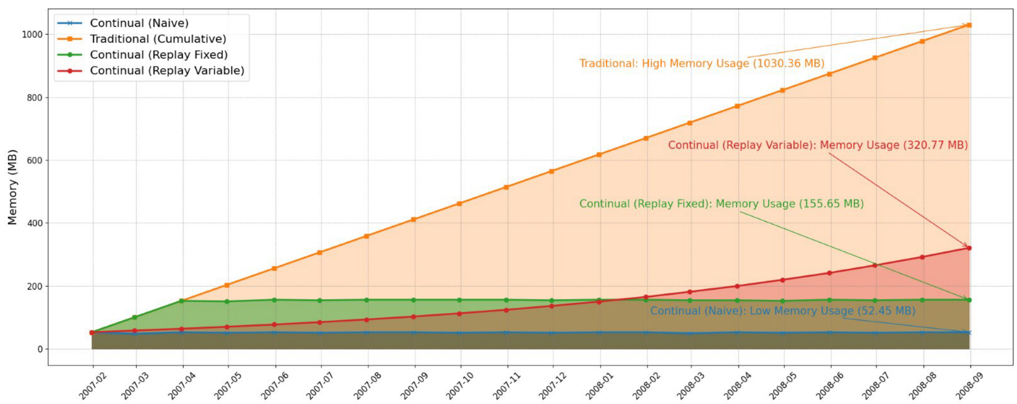

Figure 3. Efficiency of updated approaches for Target channels (56 channels).

4.2 Performance of Update Approaches

The performance of the update approaches is measured using the Mean Squared Error (MSE) for time series forecasting. The results, presented in Fig. 4, 5, and 6 and in Table 1, show that the Continual Replay (Variable) approach achieves the best balance between efficiency and efficacy.

- Cumulative: This approach generally has the lowest MSE, indicating high accuracy, but it is computationally expensive.
- Naïve: This approach has the highest MSE, reflecting its limited generalization capacity.
- Replay (Variable): This approach achieves a MSE close to the Cumulative approach while maintaining the efficiency of the Naïve approach.

In general, we can correlate large deviations in the various approaches with changes in the telemetry, as shown in sample channels depicted in Fig. 7, 8, and 9.

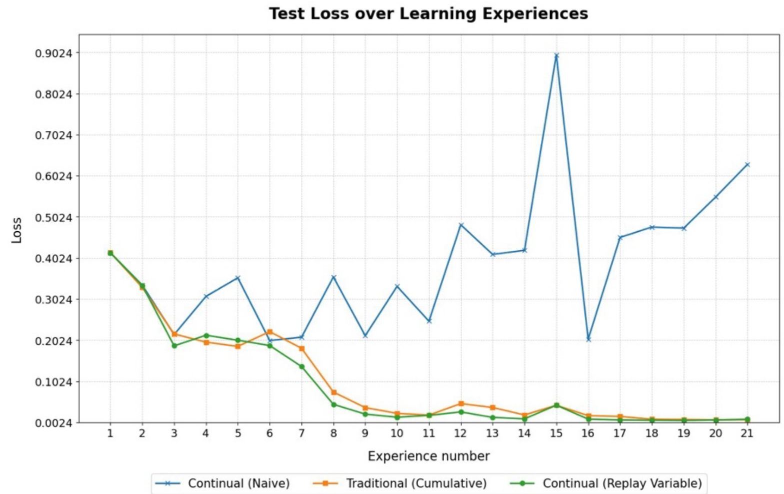

Figure 4. Performance of updated approaches for Group 13 (5 channels).

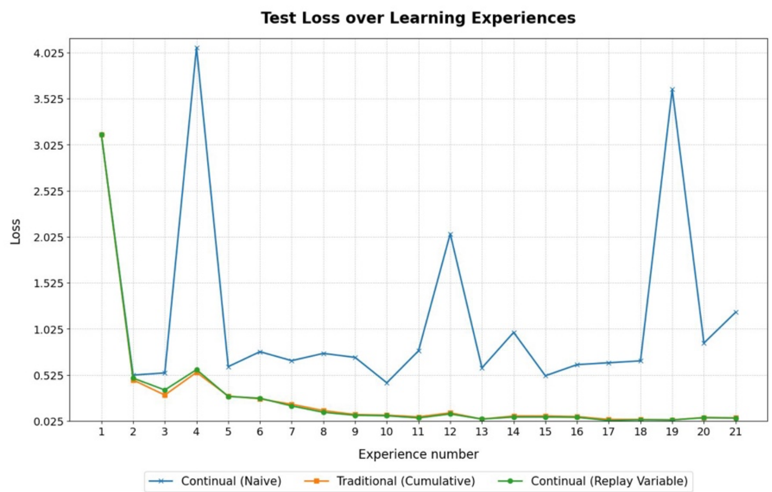

Figure 5. Performance of updated approaches for Subsystem 6 (42 channels).

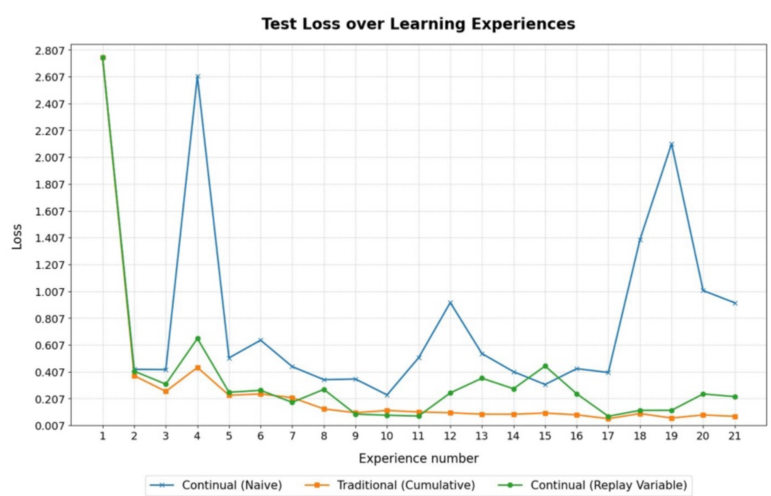

Figure 6. Performance of updated approaches for Target channels (56 channels).

Method	MSE Group 13	MSE Subsystem 6	MSE Target Channels
Cumulative	0.1018	0.2981	0.2766
CL Naïve	0.3901	1.1852	0.8427
CL Replay	0.0920	0.2972	0.3276

Table 1. Mean Squared Error (MSE) for the various updated approached.

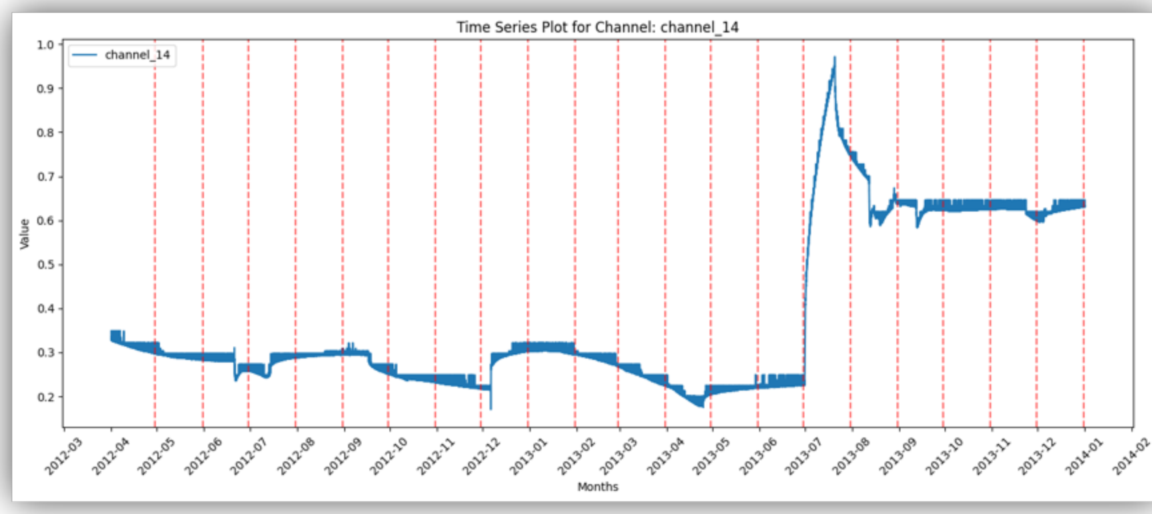

Figure 7. Time evolution of channel 14 (blue) with overlaid the times of the experiments (red).

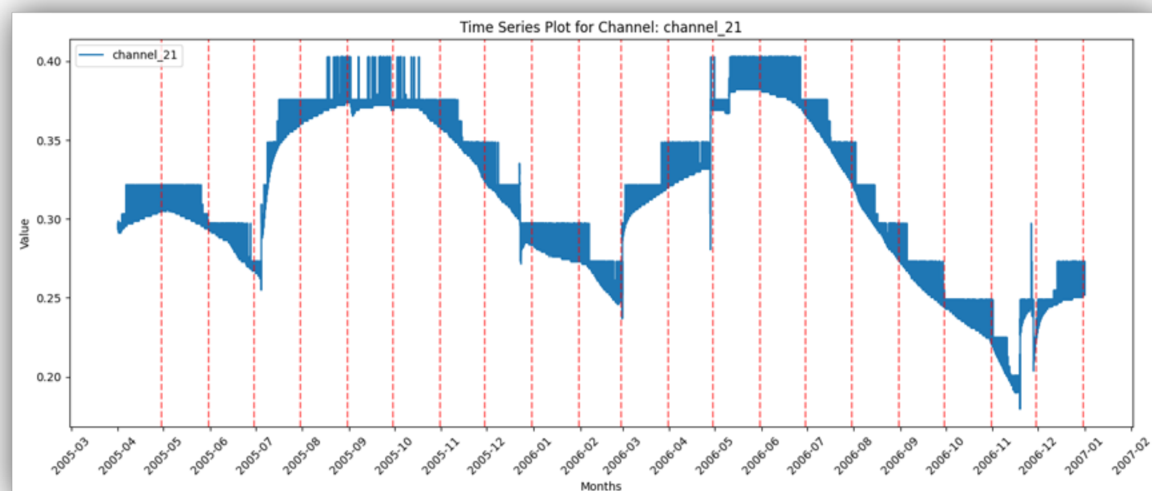

Figure 8. Time evolution of channel 21 (blue) with overlaid the times of the experiments (red).

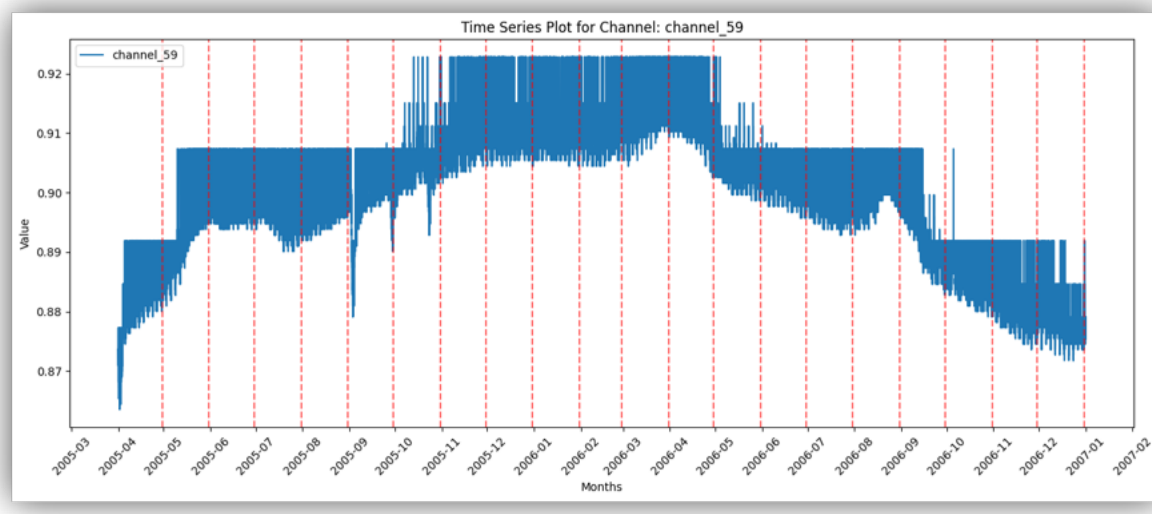

Figure 9. Time evolution of channel 59 (blue) with overlaid the times of the experiments (red).

5. Discussion

The results demonstrate the effectiveness of continual learning for spacecraft health monitoring. Continual learning strategies, particularly the Replay approach, offer a significant improvement in efficiency without compromising performance. The Naïve approach, while highly efficient, suffers from reduced accuracy due to its lack of historical context. The Cumulative approach provides high accuracy but is computationally expensive and not scalable for long-term missions with increasing data. The Continual Replay (Variable) approach strikes a balance by using a buffer of the most informative historical data, enhancing both short-term adaptability and long-term generalization while maintaining efficiency. This makes it a promising solution for on-ground satellite health monitoring, where adaptability and scalability are crucial.

6. Conclusion and Future Directions

This paper has explored the application of Continual Learning paradigms to satellite health monitoring, aiming to enhance computational efficiency and overall model quality. By leveraging the ESA-ADB dataset, we conducted a case study to validate the scalability and effectiveness of CL strategies in dynamic environments. The findings underscore the potential of CL to improve anomaly detection and forecasting, aligning with the broader goal of advancing time-series analysis for space operations. Future work will focus on extending these findings to additional spacecraft to analyze generalizability as well as to on-board settings. Key areas of development include advanced replay techniques to improve forecasting capabilities, onboard implementation for real-time adaptation, and deployment of CL as a service for ground segment applications, such as on SMAI, which will offer a more adaptable, efficient, and scalable solution for monitoring spacecraft telemetry data.

References

[1] K. Kotowski, C. Haskamp, J. Andrzejewski, B. Ruszczak, J. Nalepa, D. Lakey, P. Collins, A. Kolmas, M. Bartesaghi, J. Martínez, G. De Canio, European Space Agency Benchmark for Anomaly Detection in Satellite Telemetry, arXiv, 2024

[2] G. De Canio, K. Kotowski, C. Haskamp, ESA Anomaly Dataset (1.0), European Space Agency. https://doi.org/10.5281/zenodo.12528696, 2024

[3] K. Kotowski, C. Haskamp, J. Andrzejewski, B. Ruszczak, J. Nalepa, D. Lakey, P. Collins, A. Kolmas, M. Bartesaghi, J. Martínez, G. De Canio, The Making of the European Space Agency Benchmark for Anomaly Detection in Satellite Telemetry, SpaceOps ID # 152, 18th International Conference on Space Operations, Montreal, Canada, 26 - 30 May 2025

[4] G. De Canio, J. Eggleston, J. Fauste, A. M. Palowski, M. Spada, Development of an actionable AI roadmap for automating mission operations, SpaceOps ID # 303, 17th International Conference on Space Operations, Dubai, United Arab Emirates, 6 - 10 March 2023

[5] G. De Canio, E. Ntagiou, F. Antonello, J. Eggleston, Artificial Intelligence for mission operations automation roadmap: the European Space Operations Centre updates and vision, SpaceOps 2025, 18th International Conference on Space Operations, Montreal, Canada, 26 - 30 May 2025

[6] N. Salor Moral, P. Pilgerstorfer, F. Marino, V. Sicking, M. G. F. Kirsch, A. McDonald, E. Ntagiou, G. De Canio, Automation of flight dynamics planning for ESA's XMM-Newton, IAC-24,B6,2,6,x88657, 75th International Astronautical Congress, Milan, Italy, 2024, 14 – 18 October

[7] A. Krstova, P. Fleith, J. Martinez-Heras, N. Salor Moral, A. Buccione, P. Pilgerstorfer, S. Foley, D. Mesples, J. Eggleston, G. De Canio, Intelligent root-cause investigation and AI-assisted handling tool for flight control teams, IAC-24,B6,IP,52,x88661, 75th International Astronautical Congress, Milan, Italy, 2024, 14 – 18 October.

[8] A. Krstova, F. Hegwein, J. Hansen, P. Fleith, J. Martinez-Heras, J. Lerch, G. De Canio, Artificial Intelligence-based short-term satellite health forecasting, IAC-24,B6,IP,36,x88664, 75th International Astronautical Congress, Milan, Italy, 2024, 14 – 18 October

[9] G. De Canio, J. Eggleston, E. Ntagiou, M. Unal, H. Dreihahn, J. Lerch, M. G.F. Kirsch, S. Foley, N. Salor Moral, P. Pilgerstorfer, M. Szybowski, S. Napoletano, J. Martinez-Heras, A. Krstova, S. Krikorian, Artificial intelligence-based automation of mission post-launch operations processes, IAC-24,B6,1,1,x88670, 75th International Astronautical Congress, Milan, Italy, 2024, 14 – 18 October

[10] N. Policella, J. Martinez, N. Salor Moral, S. Kotisis, K. Helmsauer, T. Goettfert, M. Cristoforetti, A. Gobbi, F. Hegweine, A. Riise, S Mejri, V. D. Tran, I. Tanco, G. De Canio, Innovating Space Operations with AI: The AISHGO Project, SpaceOps ID # 49, 18th International Conference on Space Operations, Montreal, Canada, 26 - 30 May 2025

[11] R. Semola, V. Lomonaco, D. Bacciu, Continual-Learning-as-a-Service (CLaaS): On-Demand Efficient Adaptation of Predictive Models, 1st International Workshop on Pervasive Artificial Intelligence, Hosted by the 2022 IEEE World Congress on Computational Intelligence

[12] J. Hurtado, D. Salvati, R. Semola, M. Bosio, V. Lomonaco, Continual learning for predictive maintenance: Overview and challenges, Intelligent Systems with Applications - Elsevier

[13] R. Semola, J. Hurtado, V. Lomonaco, D. Bacciu, Adaptive Hyperparameter Optimization for Continual Learning Scenarios, 1st ContinualAI Unconference - Proceedings of Machine Learning Research (PMLR)

[14] C. Lippi, Fault Diagnosis Systems, Springer International Publishing, Cham. pp. 249–270, 2014

[15] G. Ditzler, M. Roveri, C. Alippi, R. Polikar, Learning in nonstationary environments: A survey. IEEE Computational Intelligence Magazine 10, 12–25, 2015

[16] H. Hu, M. Kantardzic, T. S. Sethi, No free lunch theorem for concept drift detection in streaming data classification: A review, Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery 10, e1327, 2020

[17] F. Hutter, L. Kotthoff, J. Vanschoren, Automated machine learning: methods, systems, challenges, Springer Nature, 2019

[18] T. Lesort, M. Caccia, I. Rish, Understanding continual learning settings with data distribution drift analysis, arXiv:2104.01678, 2021

[19] T. Lesort, V. Lomonaco, A. Stoian, D. Maltoni, D. Filliat, N. Díaz-Rodríguez, Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges, Information fusion 58, 52–68, 2020

[20] E. Verwimp, S. Ben-David, M. Bethge, A. Cossu, A. Gepperth, T. L. Hayes, E. Hüllermeier, C. Kanan, D. Kudithipudi, C. H. Lampert, et al., Continual learning: Applications and the road forward, arXiv:2311.11908, 2023

[21] G. De Canio, K. Annus, Using machine learning for thermal gradients modelling of XMM-Newton propellant tanks, SpaceOps-2021,13,x1552, 16th International Conference on Space Operations, Virtual Edition, 2021, 3 – 5 May

[22] S. Hochreiter, J. Schmidhuber, Long Short-Term Memory, Neural Computation, 9(8):1735–1780, 1997

[23] I. Goodfellow, Y. Bengio, A. Courville. Deep Learning, MIT Press, 2016