Fractal-based Satellite Health Monitoring

Lucio Pinello[1], Lorenzo Brancato[1], Alessandro Lucchetti[1], Francesco Cadini[1], and Marco Giglio [1]

[1] *Department of Mechanical Engineering, Politecnico di Milano, Via La Masa 1, Milano, 20156, Italy*
lucio.pinello@polimi.it
lorenzo.brancato@polimi.it
alessandro.lucchetti@polimi.it
francesco.cadini@polimi.it
marco.giglio@polimi.it

ABSTRACT

Satellites and space systems are crucial to the success of modern space exploration, particularly as the scale and complexity of missions continue to increase. Given the high investment costs and the impossibility of physical intervention once in orbit, designing reliable and fault-tolerant platforms is crucial for success. Nevertheless, the extreme and unpredictable conditions of the space environment frequently lead to anomalies that threaten mission success. Telemetry data is therefore indispensable for real-time and predictive monitoring of system health. However, its complexity, multidimensionality, and the presence of noise pose significant challenges to traditional analytical techniques.

In this context, fractal analysis provides a robust set of tools for uncovering hidden patterns in telemetry signals, enabling the early detection of system degradation and anomalies. Unlike conventional threshold-based approaches, fractal methods are sensitive to changes in signal regularity and complexity, making them suitable for pre-failure diagnostics and trend forecasting.

This work investigates the use of fractal-based techniques for satellite health monitoring. The methods are applied to the Mission 1 dataset of the ESA Anomaly Detection Benchmark (ESA-ADB) database, enabling performance evaluation under realistic operational conditions. A comparative analysis is conducted to assess the diagnostic capability, robustness, and computational efficiency of each method, with a focus on identifying subtle anomalies and facilitating proactive decision-making.

The results highlight the potential of fractal techniques to enhance the interpretability and autonomy of satellite prognostics and health management (PHM) systems. By enabling more sensitive and timely diagnostics, this approach contributes to improving the operational resilience and life-cycle management of future space missions.

1. INTRODUCTION

As satellite missions become more complex and autonomous, real-time monitoring of spacecraft health is critical to ensure mission success and system longevity. Since physical intervention is infeasible in orbit, early detection of anomalies using telemetry data is essential. Traditional rule- or threshold-based approaches often fail to detect gradual or non-obvious faults, while machine learning (ML)-based methods may suffer from high training data requirements, lack of transparency, or overfitting (Hundman, Constantinou, Laporte, Colwell, & Soderstrom, 2018; Lakey & Schlippe, 2024; Zamanzadeh Darban, Webb, Pan, Aggarwal, & Salehi, 2024).

In this context, fractal-based signal analysis offers a promising and interpretable alternative for anomaly detection. Fractal descriptors measure the structural complexity and self-similarity of time series, properties that often change when systems begin to degrade. Unlike many ML-based algorithms, fractal measures are nonparametric, lightweight, and require no labelled data, making them especially well-suited to the low-resource, safety-critical environment of space missions (Yuan & Wu, 2022; Zamanzadeh Darban et al., 2024).

This study investigates the potential of four fractal techniques— Katz Fractal Dimension (FD), Kaguchi FD, Petrosian FD, and the Hurst Exponent — for satellite health monitoring. These techniques are applied to a real telemetry channel from Mission 1 of the ESA Anomaly Detection Benchmark (ESA-ADB) (Kotowski et al., 2024), a high-quality dataset annotated by spacecraft operations engineers to support benchmarking of anomaly detection systems in the space domain.

Each fractal method offers a distinct approach to quantifying signal irregularity. Higuchi FD uses a scale-space

method to estimate fractal behaviour in time series (Higuchi, 1988); Katz FD captures geometric complexity by considering the ratio of trajectory length to the furthest point (Katz, 1988); Petrosian FD simplifies computation by counting zero-crossings (Petrosian, 1995); and the Hurst exponent reflects long-range correlations and memory effects (Hurst, 1951; Feder, 1988). These methods can detect both abrupt and slowly evolving anomalies that may be invisible to simple statistical metrics or traditional classifiers.

Compared to machine learning methods like autoencoders or recurrent neural networks (Hundman et al., 2018; Zamanzadeh Darban et al., 2024), fractal techniques require neither training nor tuning and offer improved interpretability—a key factor in mission-critical systems where explainability and validation are essential (Guidotti et al., 2018; Adadi & Berrada, 2018; Li, Zhu, & Van Leeuwen, 2023). They are also resilient to the class imbalance and sparsity that often plague space telemetry datasets.

By comparing the performance of these fractal estimators under a unified health monitoring framework, this work contributes to the development of interpretable, low-complexity, and reliable anomaly detection tools for current and future space missions.

2. METHODOLOGY

2.1. Dataset and Preprocessing

The analysis was conducted on telemetry data from the `ESA-ADB` repository (Kotowski et al., 2024), specifically using `Mission1`. The dataset comprises 84 months of satellite operations. The data was resampled with a sampling frequency of 0.033 Hz to achieve uniform data on the mission's dominant sampling frequency.

2.2. Algorithm

Let $x = \{x_1, x_2, \ldots, x_N\}$ be a univariate time series of length N. To characterise its complexity, four fractal-based estimators were computed on each window: Higuchi's Fractal Dimension (HFD), Katz's Fractal Dimension (KFD), Petrosian's Fractal Dimension (PFD), and the Hurst exponent (H).

Higuchi's method estimates the fractal dimension HFD by computing the average curve length $L(k)$ over different scale factors $k \in \{1, 2, \ldots, k_{\max}\}$ (Higuchi, 1988). For each k, k new time series are constructed:

$$X_m^{(k)} = \{x_m, x_{m+k}, x_{m+2k}, \ldots\}, \quad \text{for } m = 1, 2, \ldots, k \tag{1}$$

Each has an associated length:

$$L_m(k) = \frac{1}{k} \left(\sum_{i=1}^{\lfloor \frac{N-m}{k} \rfloor} |x_{m+ik} - x_{m+(i-1)k}| \right) \cdot \frac{(N-1)}{\lfloor \frac{N-m}{k} \rfloor \cdot k} \tag{2}$$

The average $L(k)$ is obtained by averaging $L_m(k)$ over m, and the slope of $\log(L(k))$ versus $\log(1/k)$ gives an estimate of the fractal dimension HFD.

Katz's method quantifies complexity based on signal geometry (Katz, 1988):

$$KFD = \frac{\log_{10}(L/a)}{\log_{10}(d/a)} \tag{3}$$

where $L = \sum_{i=2}^{N} |x_i - x_{i-1}|$ is the total path length of the time series, $d = \max_i |x_i - x_1|$ is the diameter (maximum distance to the first point), $a = L/(N-1)$ is the average step length between successive points.

Petrosian's method uses the number of sign changes in the derivative to assess signal complexity (Petrosian, 1995):

$$PFD = \frac{\log_{10}(N)}{\log_{10}(N) + \log_{10}\left(\frac{N}{N+0.4N_\Delta}\right)} \tag{4}$$

where N is the length of the time series, while N_Δ is the number of sign changes in the first derivative of x, indicating the number of local extrema.

The Hurst exponent estimates long-term memory in time series (Hurst, 1951; Feder, 1988):

$$H = \frac{\log(R/S)}{\log(N)} \tag{5}$$

where R is the range of cumulative deviations from the mean and S is the standard deviation. The ratio R/S is known as the rescaled range. N is the length of the time series.

To detect deviations in complexity, a two-sided cumulative sum (CUSUM) algorithm (Page, 1954) is applied independently to each score stream. For a given method, scores are standardised using z-scores:

$$z_t = \frac{s_t - \mu_t}{\sigma_t} \tag{6}$$

where s_t is the complexity score at time t, μ_t and σ_t are the running mean and standard deviation of all previous scores (i.e., up to t).

Then, the CUSUM statistics are updated as:

$$s_{pos}(t) = \max(0, s_{pos}(t-1) + z_t - \delta) \tag{7}$$

$$s_{neg}(t) = \min(0, s_{neg}(t-1) + z_t + \delta) \tag{8}$$

where δ is a drift factor (typically $\delta = 0.5$) that controls sensitivity.

An anomaly is triggered if either:

$$s_{pos}(t) > h \quad \text{or} \quad s_{neg}(t) < -h$$

where h is a threshold (e.g., $h = 5$), chosen as a multiple of the score standard deviation.

Anomaly predictions are evaluated using timestamp-aware binary classification metrics. Each prediction is compared against the ground truth anomaly labels within a tolerance window. The following standard metrics are computed:

- True Positives (TP): The number of correctly identified anomaly windows.
- False Positives (FP): The number of normal windows incorrectly marked as anomalies.
- False Negatives (FN): The number of anomaly windows that were not detected.

From these, the following metrics are computed separately for each fractal method by comparing the predicted anomaly labels (`window_preds`) against the ground truth (`window_labels`):

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{9}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{10}$$

$$\text{F1-score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{11}$$

Eventually, the Receiver Operating Characteristic (ROC) curves are computed for each method, along with the estimation of the respective Area Under the Curve (AUC), which will be used as an additional metric.

The threshold h is selected separately for each method with a grid search based on the maximisation of the F1-score. The choice of this metric rather than Precision or Recall is because the latter are used in the computation of the F1-score.

Algorithm 1: Fractal-based anomaly detection framework.

Input: Time series $x = \{x_1, \ldots, x_N\}$, window size `cur_size`, stride `slide`, threshold h
Output: Anomaly labels and evaluation metrics for each method

Initialize score lists $\mathcal{HFD}, \mathcal{KFD}, \mathcal{PFD}, \mathcal{H}$;
Initialize anomaly prediction lists $\mathcal{Y}_{\text{HFD}}, \mathcal{Y}_{\text{KFD}}, \mathcal{Y}_{\text{PFD}}, \mathcal{Y}_H$;
Initialize cumulative sums $s^m_{pos} = 0$, $s^m_{neg} = 0$ for each method m;
for $t = 0$ **to** `cur_size slide` **do**
  $x_{win} \leftarrow x[t : t + \texttt{cur\_size}]$;
  Compute scores: $s_{\text{HFD}}, s_{\text{KFD}}, s_{\text{PFD}}, s_H$;
  Append to $\mathcal{HFD}, \mathcal{KFD}, \mathcal{PFD}, \mathcal{H}$;
  foreach *method* $m \in \{\mathcal{HFD}, \mathcal{KFD}, \mathcal{PFD}, \mathcal{H}\}$ **do**
    Let s^m_t be the current score;
    Compute μ^m, σ^m from previous scores in m;
    $z^m_t \leftarrow \frac{s^m_t - \mu^m}{\sigma^m}$;
    $s^m_{pos} \leftarrow \max(0, s^m_{pos} + z^m_t - 0.5)$;
    $s^m_{neg} \leftarrow \min(0, s^m_{neg} + z^m_t + 0.5)$;
    if $s^m_{pos} > h$ ***or*** $s^m_{neg} < -h$ **then**
      Append 1 to $\mathcal{Y}_m$ (anomaly);
      $s^m_{pos} \leftarrow 0$, $s^m_{neg} \leftarrow 0$;
    else
      Append 0 to $\mathcal{Y}_m$ (normal);
    end
  end
end
foreach *method* $m \in \{HFD, KFD, PFD, H\}$ **do**
  Load `window_labels` (ground truth for method m);
  Load `window_preds` $= \mathcal{Y}_m$;
  Compute Precision, Recall, and F1-score
end

3. RESULTS AND DISCUSSION

The four fractal-based approaches presented in the previous section are tested on `Channel41` of `Mission1` data of the `ESA-ADB` database. The results presented focus on data from two three-month time windows: one spanning from March 01, 2000, to May 31, 2000 and the other does the same but for the year 2001.

To simulate real-time operations, where data is analysed as soon as it arrives, a sliding window approach was adopted. A window size of 720 samples (approximately 3 hours) was used as the time series length N for the computation of fractal-based estimators for anomaly detection, with a stride of 60 samples (about 30 minutes). This allowed for temporally dense estimation of signal complexity while maintaining reasonable processing intervals.

Figures 1 and 2 show the value of `Channel41` data and the four fractal-based estimators adopted, on the upper and lower subplots, respectively, over time. Figure 1 refers to the time window going from 2000-03-01 to 2000-05-31, while Figure 2 refers to the time window from 2001-03-01 to 2001-05-31. From the figure, it can be observed that not all the metrics are

sensitive to anomalies.
The Higuchi FD shows pronounced drops or shifts within anomaly windows, and its response appears well-aligned temporally with the anomalies. It maintains relatively stable behaviour outside anomalies (low false positives). It is susceptible to structural or complexity changes in the time series. Therefore, it can be an effective anomaly detector, showing a good balance of sensitivity and specificity, particularly for abrupt changes in frequency or self-similarity.
The Katz FD exhibits a strong, distinct response to anomalies, which is larger than that of the other metrics, as it spikes sharply and consistently. The amplitude and sharpness of response suggest high signal-to-anomaly contrast. However, this metric is highly sensitive to changes in signal regularity, as observed during the second anomaly of Figure 1, for example, where the consistency of the signal during the anomaly results in a KFD value close to that of the healthy signal. In addition, it may be prone to overreacting to minor fluctuations if not smoothed or normalised, meaning that this metric is excellent for detecting clear structural changes, but caution is advised for noise-sensitive data.
The Petrosian FD exhibits minimal variation throughout the entire time series, and there is no consistent reaction to anomalies, as shown in Figures 1 and 2. This may be due to the type of complexity changes induced by the anomaly in the dataset, to which this metric may not be significantly sensitive. The Petrosian FD is based on path length and maximum distance, so it may not capture fine-grained irregularities. Thus, this metric shows poor anomaly detection capability in this context and is not recommended as a primary metric.
Finally, the Hurst Exponent responds to anomalies with relatively modest yet consistent changes. Its behaviour is similar to what is described for the Higuchi FD, as it clearly detects the beginning and end of the anomaly. During the anomaly, its value becomes similar to that associated with the healthy signal if the analysed signal is consistent during the anomaly.

After evaluating the metrics, the optimal anomaly detection threshold for each method was determined based on maximising the F1-score using a grid search within the threshold range $[1, 20]$. Table 1 shows the threshold, Precision, Recall, and F1-score for each fractal-based metric. The Katz FD achieves the best balance between detecting true anomalies and minimising false positives. Likely, it benefits from a narrow value range that exaggerates subtle changes in its statistical behaviour.
The Hurst Exponent exhibits good performance, nearly as good as the Katz FD. This suggests that, despite more subtle and smooth temporal changes in the curve, the underlying memory properties captured by the Hurst Exponent are strong indicators of anomalies. Its high precision and recall indicate low noise sensitivity and consistent true detections.
The Higuchi FD exhibits high recall but low precision, as it tends to detect most anomalies while also yielding many false positives (it is prone to over-detection).
Eventually, the Petrosian FD exhibits moderate recall and low precision, suggesting somewhat unreliable performance.

Table 1. Optimal threshold and respective performance metrics for the four methods implemented.

Method	h	Precision	Recall	F1-score
HFD	19	0.41	0.73	0.53
KFD	14	0.71	0.82	0.76
PFD	9	0.44	0.64	0.52
H	8	0.71	0.78	0.74

4. CONCLUSION

This study evaluated the effectiveness of four fractal-based complexity estimators—Higuchi Fractal Dimension (HFD), Katz Fractal Dimension (KFD), Petrosian Fractal Dimension (PFD), and the Hurst Exponent (H)—as unsupervised anomaly detectors in spacecraft telemetry data. The methods were tested on a subset of the ESA-ADB dataset, using a sliding window framework to simulate real-time monitoring conditions.

The experimental results demonstrate that Katz FD and the Hurst Exponent outperform the other metrics in terms of both detection accuracy and robustness. Specifically, the Katz FD achieved the highest F1-score (0.76), indicating its strong ability to detect true anomalies while maintaining a low false positive rate. The Hurst Exponent also performed well (F1-score of 0.74), suggesting that metrics capturing long-term memory properties are particularly useful in identifying subtle temporal deviations in complex signals.

The Higuchi FD, although sensitive to structural changes, exhibited lower precision (0.41), resulting in a higher rate of false positives. Its high recall, however, suggests it could be helpful in scenarios where missing anomalies are more critical than over-detection. In contrast, the Petrosian FD showed the weakest performance across all metrics, likely due to its insensitivity to the specific type of signal irregularities present in the data.

Overall, the study shows that not all fractal-based measures are equally effective for anomaly detection, and that careful selection and tuning of the metric and its threshold are crucial for reliable deployment in real-world monitoring systems. Among the tested approaches, the Katz FD and Hurst Exponent stand out as promising candidates for anomaly detection in space mission telemetry, with potential applicability to other domains involving complex, high-dimensional time series.

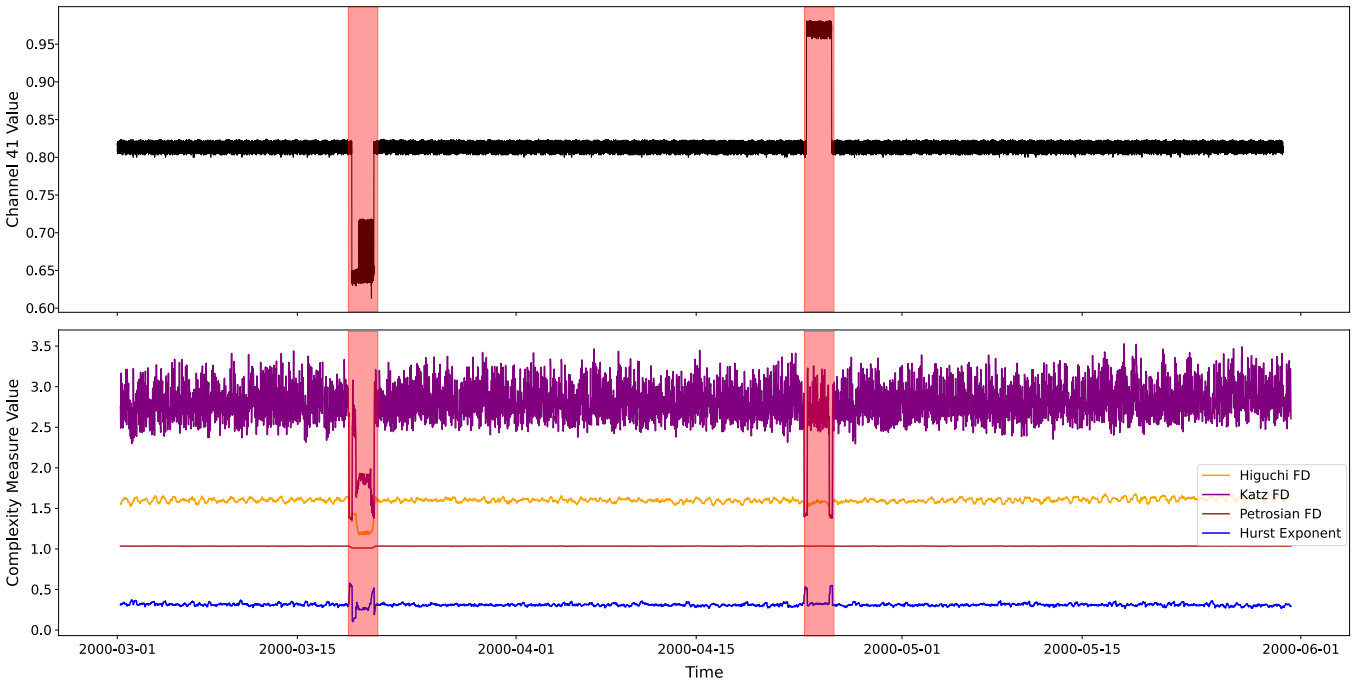

Figure 1. `Channel41` time series from 2000-03-01 to 2000-05-31 (top), and corresponding fractal-based metrics (bottom): Higuchi FD (orange), Katz FD (purple), Petrosian FD (brown), and Hurst exponent (blue). Anomalies are shaded to illustrate deviations and evaluate the responsiveness of the metric.

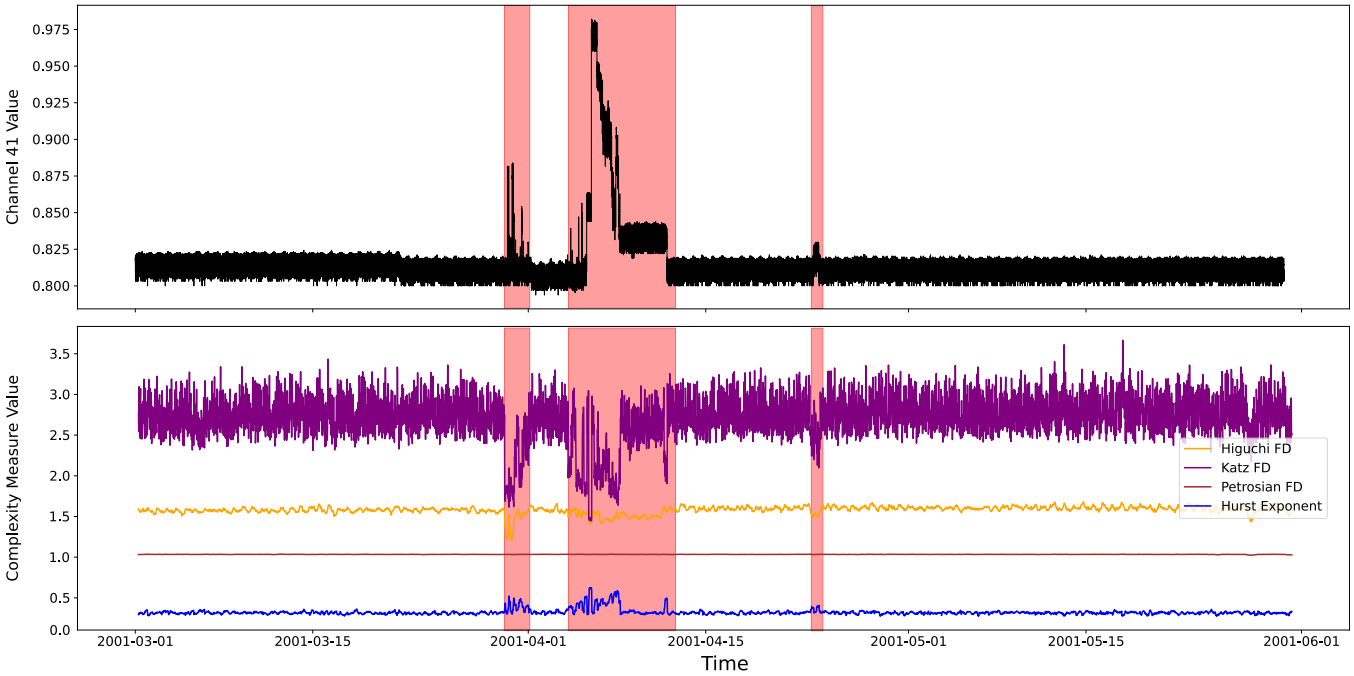

Figure 2. `Channel41` time series from 2001-03-01 to 2001-05-31 (top), and corresponding fractal-based metrics (bottom): Higuchi FD (orange), Katz FD (purple), Petrosian FD (brown), and Hurst exponent (blue). Anomalies are shaded to illustrate deviations and evaluate the responsiveness of the metric.

REFERENCES

Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on explainable artificial intelligence (xai). *IEEE Access*, *6*, 52138-52160. doi: 10.1109/ACCESS.2018.2870052

Feder, J. (1988). *Fractals*. Springer Science & Business Media.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018, August). A survey of methods for explaining black box models. *ACM Comput. Surv.*, *51*(5). Retrieved from https://doi.org/10.1145/3236009 doi: 10.1145/3236009

Higuchi, T. (1988). Approach to an irregular time series on the basis of the fractal theory. *Physica D: Nonlinear Phenomena*, *31*(2), 277-283. Retrieved from https://www.sciencedirect.com/science/article/pii/0167278988900814 doi: https://doi.org/10.1016/0167-2789(88)90081-4

Hundman, K., Constantinou, V., Laporte, C., Colwell, I., & Soderstrom, T. (2018). Detecting spacecraft anomalies using lstms and nonparametric dynamic thresholding. In *Proceedings of the 24th acm sigkdd international conference on knowledge discovery & data mining* (p. 387–395). New York, NY, USA: Association for Computing Machinery. Retrieved from https://doi.org/10.1145/3219819.3219845 doi: 10.1145/3219819.3219845

Hurst, H. E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, *116*(1), 770-799. Retrieved from https://ascelibrary.org/doi/abs/10.1061/TACEAT.0006518 doi: 10.1061/TACEAT.0006518

Katz, M. J. (1988). Fractals and the analysis of waveforms. *Computers in Biology and Medicine*, *18*(3), 145-156. Retrieved from https://www.sciencedirect.com/science/article/pii/0010482588900418 doi: https://doi.org/10.1016/0010-4825(88)90041-8

Kotowski, K., Haskamp, C., Andrzejewski, J., Ruszczak, B., Nalepa, J., Lakey, D., ... Canio, G. D. (2024). *European space agency benchmark for anomaly detection in satellite telemetry.* Retrieved from https://arxiv.org/abs/2406.17826

Lakey, D., & Schlippe, T. (2024). A comparison of deep learning architectures for spacecraft anomaly detection. In *2024 ieee aerospace conference* (p. 1-11). doi: 10.1109/AERO58975.2024.10521015

Li, Z., Zhu, Y., & Van Leeuwen, M. (2023, September). A survey on explainable anomaly detection. *ACM Trans. Knowl. Discov. Data*, *18*(1). Retrieved from https://doi.org/10.1145/3609333 doi: 10.1145/3609333

Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, *41*(1/2), 100–115. doi: 10.2307/2333009

Petrosian, A. A. (1995). Kolmogorov complexity of finite sequences and recognition of different preictal eeg patterns. *Proceedings Eighth IEEE Symposium on Computer-Based Medical Systems*, 212-217. Retrieved from https://api.semanticscholar.org/CorpusID:10910382

Yuan, S., & Wu, X. (2022). *Trustworthy anomaly detection: A survey.* Retrieved from https://arxiv.org/abs/2202.07787

Zamanzadeh Darban, Z., Webb, G. I., Pan, S., Aggarwal, C., & Salehi, M. (2024, October). Deep learning for time series anomaly detection: A survey. *ACM Comput. Surv.*, *57*(1). Retrieved from https://doi.org/10.1145/3691338 doi: 10.1145/3691338