

Cutoff Frequency Adjustment for FFT-Based Anomaly Detectors

Ellen Paixão Silva^{1,4}, Helga Balbi¹, Esther Pacitti²,
Fabio Porto³, Joel Santos¹, Eduardo Ogasawara¹

¹CEFET/RJ - Federal Center for Technological Education of Rio de Janeiro

²University of Montpellier / INRIA

³LNCC - National Laboratory for Scientific Computing

⁴ONS - National Electric System Operator

ellen.paixao@aluno.cefet-rj.br, helga.balbi@cefet-rj.br

Esther.Pacitti@lirmm.fr, fporto@lncc.br

joel.santos@cefet-rj.br, eogasawara@ieee.org

Abstract. *This article presents a time series anomaly detection method based on the Fast Fourier Transform (FFT) using a high-pass filter. The proposed method aims to remove low-frequency components, such as trends and seasonality, which represent the normal behavior of the series, while preserving high-frequency components associated with anomalies. The major challenge in constructing this method lies in determining the high-pass filter's cutoff frequency without prior knowledge of the intrinsic nature of the series. In addition to the traditional approach, four new distinct approaches were explored to determine the high-pass filter's cutoff frequency, making the method adaptable to various datasets. Experimental results show the effectiveness of the method in anomaly detection using high-pass FFT filters that have a cutoff frequency adjusted by change points, outperforming traditional techniques such as statistical and machine learning methods in terms of F1 score, precision, accuracy, and execution time.*

1. Introduction

The existence of anomalies in time series can compromise data analysis by disturbing their behavior and potentially causing biases in parameter estimation [Erkuş and Purutçuoğlu, 2021]. Therefore, detecting, correcting, and eliminating anomalies are important steps in analyzing various datasets in finance and economics, industry, geography, and medicine [Yu et al., 2014]. The literature contains various methods for anomaly detection that employ different approaches. Statistical methods, such as FBIAD and ARIMA [Lima et al., 2022], identify significant deviations from expected patterns [Chandola et al., 2009]. Machine learning methods, such as LSTM, ELM, and SVM, are employed to learn complex and non-linear patterns in time series [Pang et al., 2021]. Clustering methods, such as K-means and DBSCAN, identify anomalous data clustering similar data and highlighting distant outliers. There are also decomposition methods, such as Singular Value Decomposition (SVD) and Principal Component Analysis (PCA), that help reduce data dimensionality and highlight anomalous variations [Olteanu et al., 2023].

Additionally, methods in the frequency domain, such as Fast Fourier Transform (FFT) and Wavelet Transform, explore frequency-domain characteristics to differentiate anomalous patterns. However, current methodologies in the frequency domain have the underexplored potential for time series analysis [Zhou et al., 2023]. FFT is a mathematical tool used to convert time-domain data into frequency-domain representation, allowing the identification of periodicity and trends, as well as the creation of noise filters. It is widely used in signal processing, telecommunications systems, or electromagnetic fields in electrical engineering [Oppenheim et al., 1997]. In time series analysis, FFT is used to detect anomalies based on the idea that anomalies can introduce atypical frequency components in the series. These may manifest as unexpected peaks at specific frequencies or greater energy spread across multiple frequencies instead of being concentrated at specific frequencies associated with normal time series behavior [Zhou et al., 2023].

This work explores using FFT to create high-pass filters, which allow the passage of high frequencies and attenuate or eliminate low frequencies. Regarding time series, such filters suppress low-frequency components, such as trend and seasonality, that represent normal series behavior while preserving high-frequency components that may correspond to anomalies [Jiang et al., 2021]. The major challenge of this work is to construct a method that determines the cutoff frequency of the high-pass filter without prior knowledge of the intrinsic nature of the series to be analyzed.

Four new approaches were developed to determine the ideal cutoff frequency for anomaly detection using FFT, exploring possible solutions to this problem. To evaluate the performance in anomaly detection based on these approaches, an experiment was conducted on datasets with diverse properties to be tested, including volatility, trend, and the presence or absence of seasonality. FFT-based methods were compared against the cited statistical methods (FBIAD and ARIMA) and machine learning methods (LSTM, ELM, and SVM). These discussions reinforce the need to choose the appropriate anomaly detection method based on the specific requirements of the operational scenario, balancing response speed and accuracy.

In addition to this introduction, the article is organized into five more sections. Section 2 presents the literature review, and Section 3 details the new FFT-based method in its basic form, as well as the four new proposals used to define the cutoff frequency of the high-pass filter. Section 4 presents the experimental evaluation and its discussion. Finally, Section 5 provides the final considerations.

2. Literature Review

Collins Jackson and Lacey [2020] demonstrate how the discrete Fourier transform (DFT) can detect seasonality and anomalies in binary, rare data, introducing a new anomaly detection method based on the sum of distances. Bürger and Pauli [2013] present an unsupervised method for detecting and segmenting anomalies in sequential data, images, and volumetric data using a multiscale analysis based only on the phase of the Fourier transform. Herrera et al. [2021] propose a framework for anomaly detection in internet traffic in core and metro networks, using time series analysis of the Graph Fourier Transform to improve computational accuracy and efficiency.

Loyarte and Menenti [2008] investigate how rainfall anomalies impact the Fourier transform parameters of Normalized Difference Vegetation Index (NDVI) time series in

northwestern Argentina. Ye et al. [2023] presents a Fourier Time Series Transformer (FTST) model for anomaly detection in multivariate time series, combining features from the temporal and frequency domains to improve anomaly detection performance. Lindstrom et al. [2020] presents functional kernel density estimation (FKDE) methods for anomaly detection in time series, using point-based and Fourier-based approaches for aviation security.

Zhao et al. [2018] propose a Fourier series-based approach for extracting anomalies in power telecommunications network traffic to identify anomalous components. Bhattacharya et al. [2020] present a method based on the FFT for the detection and classification of thermoacoustic instability (TAI) and lean blowout (LBO) in turbulent combustors. Erkuş and Purutçuoğlu [2021] propose a frequency domain-based outlier detection (FOD) algorithm to identify quasi-periodic outliers in time series, demonstrating its effectiveness compared to traditional methods through simulations and real data applications.

As can be seen, FFT has many applications in anomaly detection. However, the cutoff frequency is not a subject of study, which reinforces the need to explore it, as proposed in this paper.

3. FFT-Based High-Pass Filter Anomaly Detection Method

Let X be a time series containing n observations, such that $X = \langle x_1, \dots, x_n \rangle$. Let Y be the frequency domain representation of the time series obtained from an FFT, such that $Y = FFT(X)$. Consider h to be a high-pass filter. In h , the power spectrum P of Y is computed, such that $P = Y^2$. Then, a cutoff frequency f is determined for the filter. When applying h to Y , it yields a frequency domain time series representation \dot{Y} , in which frequencies below the threshold f are removed. It should be noted that the challenge of this work is to study alternatives for choosing f .

Knowing the cutoff frequency, one can apply the Inverse Fourier Transform (IFFT) to the time series in the frequency domain, obtaining a residual time series $\omega = \langle \omega_1, \dots, \omega_n \rangle$ that are expected to exclude low frequencies associated with the trend and seasonality of X . Outliers in this series ω correspond to observations ω_i that are atypical, as defined by Equation 1, where $IQR(\omega)$ is the interquartile range of ω , and $Q_1(\omega)$ and $Q_3(\omega)$ are the first and third quartiles, respectively. Additionally, the observations characterized as outliers by Equation 1 can be mapped as anomalies into X because they occur in the same time instances, *i.e.*, $anomalies(X) = outliers(\omega)$.

$$outliers(\omega) = \{t, \omega_t \notin [Q_1(\omega) - 1.5 \cdot IQR(\omega), Q_3(\omega) + 1.5 \cdot IQR(\omega)]\} \quad (1)$$

Algorithm 1 summarizes the anomaly detection process using FFT. It begins by taking a time series X as input. First, FFT is applied to X , generating Y . Next, the high-pass filter h is applied to Y , resulting in \dot{Y} . Then, the inverse-transform FFT is applied to \dot{Y} leading to ω . Finally, outliers identified in ω are characterized as anomalies (A).

This work presents five (a baseline and four news) different approaches for adjusting the cutoff frequency of the high-pass FFT filter h : traditional (TF), AMOC (AF), BinSeg (BSF), CUSUM AMOC (CAF), and CUSUM BinSeg (CBSF). Besides the baseline TF, these new four approaches are inspired by the idea that cutoff frequencies can be characterized as a change point in the power spectrum. They are described as follows.

Algorithm 1 Anomaly detection using FFT

```

procedure AnomalyFFT( $X, h$ )
   $Y \leftarrow FFT(X)$ 
   $\dot{Y} \leftarrow h(Y)$ 
   $\omega \leftarrow IFFT(\dot{Y})$ 
   $A \leftarrow outliers(\omega)$ 
  return  $A$ 
end procedure

```

In the TF approach, the cutoff frequency f is initialized with the index of the maximum value of P , i.e., the frequency component that contributes the greatest power to the Y function. Suppose there is variation in the values of P , a threshold based on the average values of P plus 2.698 times the standard deviation of P using the Central limit theorem. The values of P are adjusted below this threshold to facilitate the identification of a significant cutoff point. The cutoff frequency f for the high-pass filter h corresponds to the new adjusted minimum index frequency in P .

In the AF approach, the AMOC (At Most One Change) method is used to identify a single point where a significant change in the data mean of P occurs. This method is well presented in [Killick and Eckley, 2014]. At each point, P is divided into two segments: before and after the change point. The mean for the data before and after each division point is calculated. The difference between the means before and after each point is evaluated using a test statistic. The change point is the one that maximizes the test statistic, indicating the largest difference in segment means. The significance of the change is verified to ensure it is not due to chance [Lykou et al., 2020]. The cutoff frequency f for the high-pass filter h corresponds to the change point detected in P by the AMOC.

In the BSF approach, the BinSeg (Binary Segmentation) method is used to identify multiple change points in the data mean. It adopts a recursive and greedy approach based on the AMOC. This method is well presented in [Lykou et al., 2020]. In this method, P is divided into segments and possible divisions are evaluated for the best segmentation. For each possible division, the mean of the resulting segments and the sum of the squared deviations within each segment are calculated. The method determines where significant changes occur by comparing the segments before and after the detected changes. The statistical test used in the BinSeg method tries to find the points that minimize the sum of errors between the segments. The method selects the points where the segment means differ significantly. This process continues iteratively until the desired number of change points is identified or an optimal solution is found. In this approach, the cutoff frequency f for the high-pass filter h corresponds to the last change point detected in P by BinSeg.

The CAF approach combines the Cumulative Sum Control Chart (CUSUM) and AMOC methods. CUSUM is a sequential analysis technique [Gama et al., 2014] based on accumulating deviations of observed data from a reference or target mean, allowing the detection of small variations. Applying CUSUM to the power spectrum allows detecting when the power distribution across different frequencies begins to change, which can be an early indicator of a change in the behavior of P . The CUSUM transformation on P leads to \hat{P} . In this approach, CUSUM highlights regions of interest where changes may occur,

and AMOC provides precision by locating the exact change point. The cutoff frequency \hat{f} for the high-pass filter h corresponds to the change point detected by the AMOC on the \hat{P} .

The CBSF approach is similar to CAF but differs using the BinSeg method following CUSUM. In this approach, the cutoff frequency \hat{f} for the high-pass filter h corresponds to the last change point detected in \hat{P} by BinSeg after the potential change points identified by CUSUM.

4. Results

This section evaluates the proposed anomaly detection methods presented in Section 3. For comparison purposes, other established methods employing different approaches in the anomaly detection process were chosen. Statistical methods (FBIAD and ARIMA) and machine learning methods (LSTM, ELM, and SVM) were considered. All methods are available in the Harbinger R package available at CRAN. The sliding window size parameter was set to 30 for the FBIAD method. For the LSTM method, the *epochs* parameter was set to 10000. The *actfun* parameter was set to *Purelin* for the ELM method, and for the SVM method, the *kernel* parameter was set to *Radial Basis*.

To compare the methods, the Yahoo Labs dataset was chosen, which consists of a collection of time series, including synthetic and real-time series related to data traffic on Yahoo services¹. The evaluations were conducted on an Intel Xeon w3-2423 processor with 512 GB of RAM, 12 cores, and an Ubuntu 22.04 LTS operating system.

From the confusion matrix, other metrics were considered: precision, recall, F1 score, accuracy, and elapsed time. Precision measures the accuracy of detections, showing the proportion of correct detections relative to the total detections. Recall measures completeness or the true positive rate, representing the correct proportion of detected anomalies relative to the total anomalies. The F1 score combines precision and recall in a harmonic mean, generating a balanced performance measure. Finally, accuracy evaluates the overall rate of correct predictions in the sample [Han et al., 2022]. The results presented here are the mean of these metrics for all explored time series.

Table 1 presents the results for the Yahoo dataset. The FBIAD method stands out for its high recall rate, indicating that it effectively identifies anomalies when they are present. However, its precision is relatively low, suggesting a high rate of false positives. This method also exhibits a low F1 score. Despite these limitations, FBIAD has a high accuracy and a relatively fast processing time compared to methods like ARIMA and LSTM.

Although with low precision and F1-score, the ARIMA method has a recall rate similar to FBIAD, indicating that it can also identify anomalies but with a high rate of false positives. ARIMA's processing time is significantly longer, which may limit its applicability in real-time scenarios. The LSTM method shows results similar to ARIMA regarding precision and recall but with a notably lower F1-score, indicating substantial issues in anomaly detection accuracy. LSTM has the longest processing time among all tested methods, which can be a major obstacle to its practical use.

¹<https://yaooresearch.tumblr.com/post/114590420346/>

The TF and AF methods stand out for their high precision and accuracy combination, with fast processing times. The CAF method presents the best overall combination of metrics, with the highest precision, good recall, the best F1 score, high accuracy, and a relatively fast processing time. Therefore, CAF achieved the best overall performance among the evaluated methods.

From these results, we can observe that methods based on deep learning models, like LSTM, despite being promising in other areas, face significant challenges in anomaly detection due to high processing time and low F1-score. Traditional methods, such as ARIMA and SVM, also show limitations in terms of precision and recall. On the other hand, methods based on FFT, like TF, AF, and especially CAF, prove to be more robust and efficient, combining good precision, high recall rates, and fast processing times. These results suggest that CAF can offer an ideal balance between performance and efficiency for practical applications, especially those requiring real-time processing.

Table 1. Results for Yahoo Dataset

Method	Precision	Recall	F1	Accuracy	Time (s)
FBIAD	0.14	0.69	0.18	0.94	9.01
ARIMA	0.06	<u>0.67</u>	0.10	0.93	130.94
LSTM	0.07	0.64	0.01	0.93	1280.58
ELM	0.06	0.64	0.10	0.93	4.55
SVM	0.04	0.66	0.07	0.91	73.80
TF	<u>0.49</u>	0.33	<u>0.28</u>	0.98	1.81
AF	0.41	0.35	<u>0.28</u>	0.98	2.09
BSF	0.22	0.44	0.24	0.98	2.86
CAF	0.54	0.42	0.39	0.98	<u>2.01</u>
CBSF	0.24	0.48	0.25	<u>0.97</u>	2.26

5. Conclusion

This paper presents innovative approaches to anomaly detection using high-pass filters with FFT. The main challenge is establishing the cutoff frequency for the high-pass filter so that it can be adaptable to any series without prior knowledge of its characteristics. Five alternative approaches presented in this paper were experimentally evaluated on the Yahoo Labs dataset. The approaches were also compared with other representative anomaly detection methods, including statistical approaches such as FBIAD and ARIMA and machine learning approaches such as LSTM, ELM, and SVM.

The approaches proved adaptable to distinct datasets and showed superior performance in anomaly detection compared to traditional methods such as FBIAD, ARIMA, LSTM, ELM, and SVM in terms of F1 score, precision, and accuracy. It also presented an advantage in terms of execution time. Among the five approaches presented using FFT, the one that stood out the most was the one that combined the CUSUM and AMOC methods (CAF) to determine the filter cutoff frequency.

Due to their computational efficiency, FFT-based methods are especially promising for real-time data analysis and can be exploited for anomaly detection in PMU (Phasor Measurement Unit) data, providing faster and more accurate identification of anomalous events in complex electrical systems.

References

- Bhattacharya, C., De, S., Mukhopadhyay, A., Sen, S., and Ray, A. (2020). Detection and classification of lean blow-out and thermoacoustic instability in turbulent combustors. *Applied Thermal Engineering*, 180.
- Bürger, F. and Pauli, J. (2013). Unsupervised segmentation of anomalies in sequential data, images and volumetric data using multiscale fourier phase-only analysis. In *LNCS*, volume 7944, pages 44 – 53.
- Chandola, V., Banerjee, A., and Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3).
- Collins Jackson, A. and Lacey, S. (2020). The discrete Fourier transformation for seasonality and anomaly detection of an application to rare data. *Data Technologies and Applications*, 54(2):121 – 132.
- Erkuş, E. C. and Purutçuoğlu, V. (2021). Outlier detection and quasi-periodicity optimization algorithm: Frequency domain based outlier detection (FOD). *European Journal of Operational Research*, 291(2):560 – 574.
- Gama, J., Zliobaite, I., Bifet, A., Pechenizkiy, M., and Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4).
- Han, J., Pei, J., and Tong, H. (2022). *Data Mining: Concepts and Techniques*. Morgan Kaufmann, Cambridge, MA, 4th edition edition.
- Herrera, M., Proselkov, Y., Perez-Hernandez, M., and Parlikad, A. K. (2021). Mining Graph-Fourier Transform Time Series for Anomaly Detection of Internet Traffic at Core and Metro Networks. *IEEE Access*, 9:8997 – 9011.
- Jiang, J.-R., Kao, J.-B., and Li, Y.-L. (2021). Semi-supervised time series anomaly detection based on statistics and deep learning. *Applied Sciences (Switzerland)*, 11(15).
- Killick, R. and Eckley, I. A. (2014). Changepoint: An R package for changepoint analysis. *Journal of Statistical Software*, 58(3):1 – 19.
- Lima, J., Salles, R., Porto, F., Coutinho, R., Alpis, P., Escobar, L., Pacitti, E., and Ogasawara, E. (2022). Forward and Backward Inertial Anomaly Detector: A Novel Time Series Event Detection Method. In *Proceedings of the IJCNN*, volume 2022-July, pages 1–8.
- Lindstrom, M. R., Jung, H., and Larocque, D. (2020). Functional kernel density estimation: Point and fourier approaches to time series anomaly detection. *Entropy*, 22(12):1 – 15.
- Loyarte, M. G. and Menenti, M. (2008). Impact of rainfall anomalies on Fourier parameters of NDVI time series of northwestern Argentina. *International Journal of Remote Sensing*, 29(4):1125 – 1152.
- Lykou, R., Tsaklidis, G., and Papadimitriou, E. (2020). Change point analysis on the Corinth Gulf (Greece) seismicity. *Physica A: Statistical Mechanics and its Applications*, 541.
- Olteanu, M., Rossi, F., and Yger, F. (2023). Meta-survey on outlier and anomaly detection. *Neurocomputing*, 555.
- Oppenheim, A. V., Willsky, A. S., and Nawab, S. H. (1997). *Signals & Systems*. Prentice Hall.
- Pang, G., Shen, C., Cao, L., and Van Den Hengel, A. (2021). Deep Learning for Anomaly Detection: A Review. *ACM Computing Surveys*, 54(2).
- Ye, Y., He, Q., Zhang, P., Xiao, J., and Li, Z. (2023). Multivariate Time Series Anomaly Detection with Fourier Time Series Transformer. In *2023 IEEE 12th CloudNet 2023*, pages 381 – 388.
- Yu, Y., Zhu, Y., Li, S., and Wan, D. (2014). Time series outlier detection based on sliding window prediction. *Mathematical Problems in Engineering*, 2014.
- Zhao, H., Lu, B., Yu, L., Zhao, S., Zeng, L., Zhang, Z., and You, P. (2018). A fourier series-based anomaly extraction approach to access network traffic in power telecommunications. In *2017 ICCSEC*, pages 550 – 553.
- Zhou, L., Guo, W., Cao, J., Zhang, X., and Wang, Y. (2023). Wavelet-SVDD: Anomaly Detection and Segmentation with Frequency Domain Attention. In *LNAI*, volume 14177, pages 230 – 243.