

Article

SpectralNet-Enabled Root Cause Analysis of Frequency Anomalies in Solar Grids Using μ PMU

Arnabi Modak ¹, Maitreyee Dey ¹, Preeti Patel ^{1,*} and Soumya Prakash Rana ²

¹ GENESIS Research Lab, London Metropolitan University, London N7 8DB, UK;
arm0727@my.londonmet.ac.uk (A.M.); m.dey@londonmet.ac.uk (M.D.)

² School of Engineering, University of Greenwich, Chatham ME4 4TB, UK; s.rana@greenwich.ac.uk

* Correspondence: p.patel@londonmet.ac.uk

Abstract

The rapid integration of solar power into distribution grids has intensified challenges related to frequency instability caused by fluctuating renewable generation. These unexpected frequency variations are difficult to capture using traditional or supervised methods because they emerge from nonlinear, rapidly changing inverter grid interactions and often lack labelled examples. To address this, the present work introduces a unique, frequency-centric framework for unsupervised detection and root cause analysis of grid anomalies using high-resolution micro-Phasor Measurement Unit (μ PMU) data. Unlike previous studies that focus primarily on voltage phasors or rely on predefined event labels, this work employs SpectralNet, a deep spectral clustering approach, integrated with autoencoder-based feature learning to model the nonlinear interactions between frequency, ROCOF, voltage, and current. These methods are particularly effective for unexpected frequency variations because they learn intrinsic, hidden structures directly from the data and can group abnormal frequency behavior without prior knowledge of event types. The proposed model autonomously identifies distinct root causes such as unbalanced loads, phase-specific faults, and phase imbalances behind hazardous frequency deviations. Experimental validation on a real solar-integrated distribution feeder in the UK demonstrates that the framework achieves superior cluster compactness and interpretability compared to traditional methods like K-Means, GMM, and Fuzzy C-Means. The findings highlight SpectralNet's capability to uncover subtle, nonlinear patterns in μ PMU data, offering an adaptive, data-driven tool for enhancing grid stability and situational awareness in renewable-rich power systems.

Keywords: solar grid; micro-synchrophasor measurement; root cause analysis; anomaly detection; unsupervised learning



Academic Editors: Shilin Gao,
Zongsheng Zheng and Jianquan Liao

Received: 26 September 2025

Revised: 22 December 2025

Accepted: 26 December 2025

Published: 4 January 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The rapid integration of solar energy into distribution networks has introduced new challenges in maintaining grid stability, especially regarding frequency anomalies induced by variable generation and bidirectional power flows. These challenges necessitate advanced diagnostic tools for anomaly detection and fault localisation. Traditional monitoring systems often lack the spatiotemporal resolution required to capture subtle or transient disturbances, leading to the adoption of micro-synchrophasor units (μ PMUs), through which high-frequency, phase-synchronised measurements are enabled [1].

μ PMUs provide sub-second visibility into grid conditions through the synchronisation of measurements of voltage, current, and frequency across distributed nodes. Their deploy-

ment in solar-dense environments has revealed previously undetected instability patterns and frequency excursions [2]. As such, they are now pivotal in enhancing grid observability and supporting real-time analytics, especially when paired with scalable machine learning methods [3].

Unsupervised learning algorithms, including clustering and anomaly detection frameworks, offer promising solutions for parsing large, unlabelled μ PMU datasets. Prior studies have shown that techniques like K-Means, Isolation Forests, and spectral clustering can effectively detect voltage or frequency deviations without the need for predefined fault templates [4]. The classification of operational versus hazardous anomalies is enabled by these methods, thereby supporting early intervention and enhancing grid resilience strategies [5]. Unsupervised learning was chosen because μ PMU datasets are typically unlabelled, making supervised classification infeasible for detecting new or evolving anomalies [6]. Unlike supervised models, unsupervised clustering can reveal hidden patterns and emerging operational behaviors without requiring prior event annotation. This adaptability makes it particularly suited for frequency anomaly detection in complex, solar-integrated grids where system conditions vary dynamically. Recent work by Guato Burgos et al. provides a comprehensive review of AI-driven anomaly detection methods across smart grid domains, highlighting the strengths and limitations of current machine learning and deep learning approaches. Their findings emphasize the need for adaptive, data-efficient, and model-independent detection frameworks, a motivation that aligns with and supports the objectives of the present study [7]. Recent contributions in the field have underscored the importance of multimodal feature integration, combining frequency, ROCOF, voltage, and current patterns to infer the nature and origin of disturbances. For instance, multi-parameter clustering has revealed the linkage between unbalanced loads and localised frequency instability in solar-connected feeders [8,9]. Exploration into lightweight, interpretable models capable of running in real time on smart inverters and distribution-level controllers has been driven by the shift toward embedded learning and edge analytics [10].

This study builds on that foundation through the use of a layered approach in which μ PMU data is combined with advanced unsupervised learning models. By focusing on root cause analysis of frequency anomalies within a utility-scale solar grid, the demonstration is aimed at showing how clustering-based approaches can enhance situational awareness, distinguish between normal and anomalous conditions, and support operational decision-making in real-world scenarios where event data availability is limited.

2. Literature Review

Most prior work in the domain of μ PMU-based monitoring suffers from narrow modelling scopes or missing causal interpretations. Dey et al.'s [11] work is limited to univariate frequency forecasting and does not extend to root cause analysis. The methodology used is limited to prediction only despite high-resolution data, with no event categorisation or unsupervised approaches for anomaly interpretation. It lacks deeper root cause inference, and anomaly types are predefined rather than emerging from data patterns. Additionally, the temporal analysis is focused on individual events without multivariate interaction [12]. This study is limited to voltage phasors and lacks explicit frequency analysis. It does not pursue root cause inference or multivariate temporal correlation across event types, which restricts deeper situational understanding [4]. The methodology used in this study is effective for generic event clustering but omits detailed frequency anomaly analysis and lacks root cause attribution. The clustering process is applied uniformly across all signals without isolating frequency behaviour or exploring its distinct causal dynamics within the power system [13]. While the method integrates physical insights with data-driven models, it focuses primarily on event trajectory similarity and general anomaly patterns

without isolating frequency-specific behaviour. The clustering is tailored to global state transitions, limiting its application to fine-grained root cause analysis of frequency deviations [14]. The study is focused on supervised classification and relies on labelled datasets with specific event types, limiting its ability to discover unknown anomalies. Using conventional PMU data restricts temporal resolution and constraints the detection of fast-evolving frequency disturbances [15]. The framework depends on predefined anomaly labels and uses low-resolution PMU data, limiting its adaptability to emerging or unknown frequency disturbances. The study does not explore clustering-based methods or unsupervised root cause inference across multivariate signals [16]. The study is domain-agnostic despite its unsupervised design and lacks application to μ PMU-based frequency data. It does not explore clustering of anomaly types or root cause inference, limiting its diagnostic potential for power systems [17]. The study relies on supervised learning and predefined labels, limiting its ability to detect novel or evolving frequency anomalies. It also emphasises voltage and current metrics without isolating frequency-specific patterns or exploring unsupervised clustering for anomaly discovery [18]. The work is limited by its dependence on labelled events and focuses exclusively on voltage-related anomalies. It does not address frequency disturbances, nor does it employ unsupervised learning or investigate causal factors behind anomalies in a multivariate context [19]. Despite the strength of using real-time μ PMU measurements, the work is focused on labelled data and supervised classification, with limited exploration of unsupervised approaches or frequency anomaly-specific features. The analysis is restricted to predefined event types rather than general anomaly clusters [20]. The study emphasises temporal prediction for forecasting over anomaly detection and lacks clustering or root cause diagnostics. The model is also supervised and not tailored for frequency-specific anomaly interpretation [21]. This study is tailored to classification tasks and requires labelled data streams, limiting its generalisability to unlabelled anomaly detection. It also uses lower-resolution PMU data and does not specifically focus on frequency anomalies or causal diagnostics [22]. The analysis remains confined to a single feeder and does not generalise beyond solar-related contexts. The event detection is limited to voltage–current dynamics without deeper frequency anomaly or multi-parameter correlation, and clustering techniques are not explored [23]. The approach used in this study focuses on supervised classification and does not explore unsupervised clustering or causal interpretation of frequency anomalies. It also lacks analysis of spatial–temporal interactions across multiple sensors [24]. The study is limited by its singular site focus and reliance on a single clustering-free anomaly detection model, with no multivariate or root cause analysis applied to the detected events across spatially diverse μ PMU installations [5]. The study is based on low-resolution PMU data and transmission-level events, lacking focus on frequency-specific anomalies in distribution systems. Additionally, it does not incorporate root cause analysis or high-resolution μ PMU data from edge nodes [25]. The method, though effective, is primarily trained on labelled events and does not incorporate frequency anomalies or multivariate unsupervised clustering. It also lacks spatial diagnostic insights, which limits generalisation to unseen or composite grid disturbances [26]. Connelly et al. (2023) [27] proposed an Autoencoder and Incremental Clustering-Enabled Anomaly Detection framework that combines reconstruction-based unsupervised learning with adaptive cluster evolution, enabling real-time detection of behavioural deviations in electrical appliance power cycles. Their use of incremental clustering and ensemble autoencoders demonstrates how adaptive learning can manage unlabelled, evolving data streams a concept relevant to μ PMU-driven grid analytics. However, their work was limited to single-device time series and did not extend to frequency-domain interactions or causal diagnostics in power systems [27].

Similarly, Kim et al. (2022) [28] introduced a Probabilistic Spectral Clustering Method for renewable-rich grids, applying spectral graph theory to capture variability from photovoltaic generation (fenrg-10-909611). This probabilistic spectral clustering approach effectively models renewable uncertainty through Monte Carlo power-flow simulations, revealing the value of spectral embedding in understanding nonlinear, stochastic grid behaviours. Nevertheless, their focus was on system-level voltage and topology partitioning rather than μ PMU-based frequency anomaly detection or unsupervised root cause analysis [28].

2.1. Research Gap

The μ PMU dataset has some limitations despite its rich detail, including power (real, reactive, apparent), phase-wise voltage and current, frequency, timestamps, and location. It lacks weather data, which is important for solar site analysis. The use of a proprietary device with built-in edge processing may affect transparency and limit access to raw signals. High-frequency sampling creates large volumes of data, making analysis resource-intensive. Cellular telemetry may lead to missing or delayed data, and GPS timing, while accurate, can still face occasional dropouts. The absence of labelled events limits the ability to rigorously validate models and hinders precise identification of specific grid anomalies.

Clustering unlabelled μ PMU data can be challenging due to high dimensionality, noise, and lack of ground truth for validation. Temporal dependencies, overlapping patterns, and varying scales across features may reduce cluster separability, while tuning parameters for methods like autoencoders or SpectralNet can be computationally intensive and sensitive.

Even with advanced grid monitoring, challenges remain. Data standards vary across systems, and high-quality, time-synchronised data is often limited to a few sites. Weather data is rarely integrated, reducing accuracy. Reliance on proprietary technologies also creates issues with scalability, transparency, and interoperability, slowing wider adoption of data-driven grid solutions [29].

This paper addresses key research gaps identified from the recent literature. First, existing studies focus primarily on voltage phasors, often overlooking detailed frequency-domain analysis. Second, most frameworks rely on labelled, low-resolution PMU data, limiting adaptability to new frequency disturbances and ignoring unsupervised clustering approaches. Third, studies using clustering methods treat all signals uniformly, lacking frequency-specific analysis or root cause attribution. These gaps frequency-focused analysis, unsupervised clustering, and multivariate root cause analysis are systematically addressed and resolved in this work through a high-resolution, μ PMU-driven approach.

2.2. Contribution

This work introduces a frequency-centric approach that integrates targeted clustering algorithms with spatial-temporal diagnostics across μ PMU measurements. This combination enables comprehensive root cause attribution and reveals the propagation pathways of system-level anomalies.

The primary objective is to explore unsupervised learning techniques for categorizing frequency anomalies using high-resolution μ PMU data. Since the dataset lacks labelled event information, clustering methods are employed to identify and group abnormal frequency patterns without prior information.

In addition, a multi-parameter root cause analysis is conducted to uncover the spatiotemporal relationships that drive these anomalies, thereby improving situational awareness in distribution grids. This analysis is validated against the UK National Grid's standard frequency limits to assess how frequency stability is influenced by other critical electrical parameters, such as voltage, power, and current.

3. Methodology

This section first outlines the end-to-end framework used to detect, cluster, and interpret frequency anomalies in a solar-integrated distribution feeder using high-resolution μ PMU data. To link the physical feeder environment to the data-driven steps used in this study, Figure 1 presents an end-to-end system overview. As shown, solar PV output interacts with substation and utility-grid dynamics, and these effects are captured by μ PMU/GDU measurements. The recorded three-phase FRVCP streams (frequency–ROCOF–voltage–current–power) are then processed and analysed through unsupervised clustering to identify anomaly families and support root-cause interpretation.

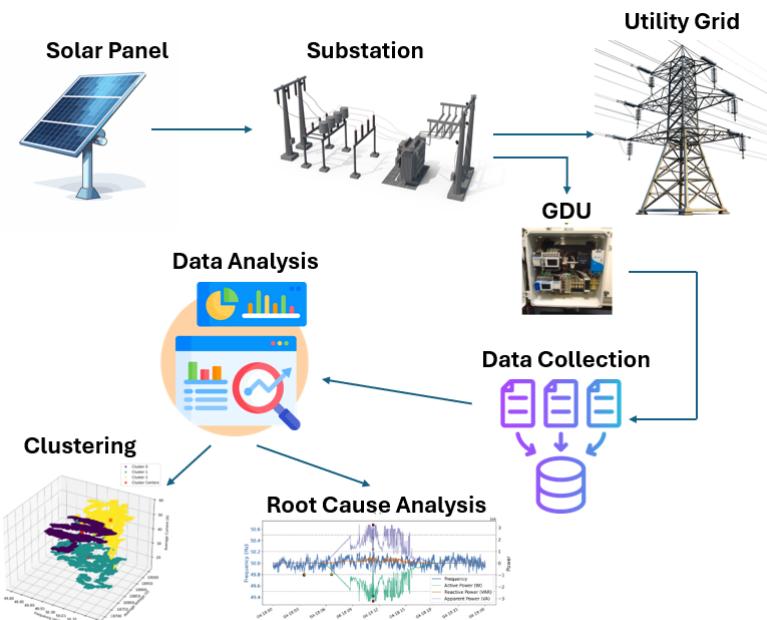


Figure 1. System overview for solar-integrated feeder monitoring.

The objective of this study is not to model or infer the behaviour of specific grid components. Instead, we focus on leveraging high-resolution μ PMU measurements for data-driven fault detection and prediction. Since no information about inverter controls, PLL designs, or other equipment is available, the proposed ML and DL-based clustering approach is intentionally hardware-agnostic and cannot be linked to particular inverter or PLL implementations.

3.1. Data Pre-Processing and Feature Construction

The analysis begins with preparation of the raw μ PMU streams. These data, sampled at 100 Hz and synchronised using GPS timestamps, contain frequency, angle, voltage, current, and power information for all three phases. The initial stage of processing involves cleaning the dataset by removing corrupted or missing entries resulting from cellular telemetry latency and creating a new column like ROCOF. Subsequently, the selected features such as frequency, ROCOF, voltage angle for three different lines, average voltage, and average current are standardised using z-score normalisation, which transforms each feature x as

$$x' = \frac{x - \mu}{\sigma} \quad (1)$$

where μ is the mean and σ is the standard deviation of the feature, computed over the training set to ensure equal contribution to the clustering algorithms and to prevent scale imbalances. The calculated rate of change of frequency uses the frequency data (scaled for

per second). ROCOF is a critical parameter in power systems that measures how quickly the system frequency changes over time. It is particularly important for detecting system disturbances, islanding conditions, and implementing protective relay schemes in electrical grids. ROCOF is calculated using the formula

$$\text{ROCOF} = \frac{df}{dt} = \frac{f_2 - f_1}{t_2 - t_1} \quad (2)$$

where df represents the change in frequency (in Hz), f_1 and f_2 are consecutive frequency measurements, and dt represents the change in time (converted from milliseconds to seconds), with typical units expressed as Hz/s. High ROCOF values (>0.2) indicate rapid frequency deviations that can trigger protective mechanisms to prevent equipment damage and maintain grid stability.

Following preprocessing, feature engineering is applied to extract relevant operational characteristics. Two types of readings are shown by the dataset. The two figures below represent a normal day (without any event) Figure 2a, and an event day in Figure 2b. The green dotted line represents the normal frequency range, whereas the red dotted line represents the hazardous frequency range. Eleven core features are retained, including a derived indicator representing compliance with the UK National Grid's operational frequency limits. In addition, average phase voltages and currents are computed to facilitate three-dimensional cluster visualisation. A sliding-window method, with a ten-second window and a one-second step, is used to identify transient disturbances such as voltage sags, current spikes, negative active power, and abnormal frequency deviations. This approach enables high-resolution localisation of event boundaries and preserves the temporal continuity essential for accurate anomaly interpretation.

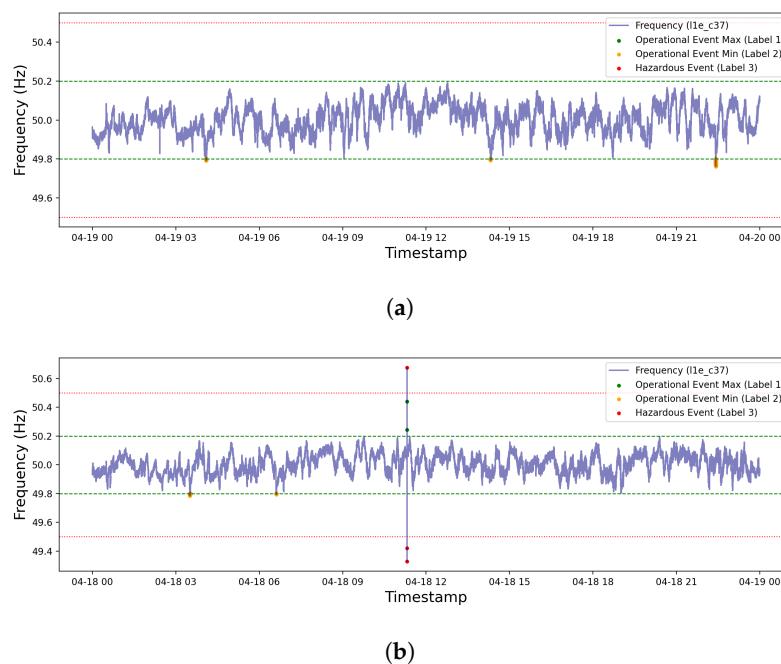


Figure 2. Frequency anomaly comparison: (a) normal day (19 April 2023), (b) event day (18 April 2023), timestamp format MM:DD:HH.

3.2. Clustering-Based Anomaly Detection

Following pre-processing and windowed feature construction, unsupervised clustering is performed to uncover hidden structures associated with normal and anomalous frequency behaviour. Five unsupervised clustering approaches are evaluated, K-Means, FCM, GMM, autoencoder-based clustering, and SpectralNet to identify hidden operational

and anomalous regimes. Conventional baseline models (K-Means, FCM, GMM) are compared with deep clustering methods (autoencoder, SpectralNet) to assess the benefit of nonlinear feature learning. To ensure methodological consistency, the optimal number of clusters is determined using the Elbow Method applied to the within-cluster sum of squares curve. Once the clustering is performed, the Davies–Bouldin Index is used to evaluate cluster compactness and separation.

K-Means is a widely used unsupervised learning algorithm by which a dataset is partitioned into k clusters through the minimisation of the within-cluster sum of squares (WCSS) [30]. The objective function for K-Means is defined as

$$J = \sum_{i=1}^k \sum_{x \in C_i} \|x - \mu_i\|^2 \quad (3)$$

where μ_i is the centroid of cluster C_i , and x represents the data points assigned to cluster i . The algorithm iteratively updates cluster centroids and reassigns points to the nearest centroid until convergence is achieved. This process aims to create clusters that are compact and well separated in the feature space.

The Fuzzy C-Means (FCM) algorithm is applied to allow partial membership of data points across multiple clusters, making it a soft clustering approach. Unlike K-Means, which uses hard assignments, a membership matrix U is computed by FCM, where each element u_{ij} is used to indicate the degree of belonging of x_j to cluster i [31]. The objective function of FCM is

$$J_m = \sum_{i=1}^c \sum_{j=1}^n u_{ij}^m \|x_j - v_i\|^2 \quad (4)$$

where v_i is the centroid of cluster i , and $m > 1$ is the fuzzifier parameter controlling the level of cluster fuzziness. The membership matrix U and cluster centroids are iteratively updated by FCM until convergence.

The data is modelled by the Gaussian Mixture Model (GMM) clustering algorithm as a mixture of multiple Gaussian distributions. A probabilistic approach is used, in which the probability that each data point belongs to each Gaussian component is estimated [32]. The likelihood function for GMM is given as

$$P(X|\Theta) = \sum_{i=1}^k \pi_i \mathcal{N}(x|\mu_i, \Sigma_i) \quad (5)$$

where π_i are the mixing coefficients, and $\mathcal{N}(x|\mu_i, \Sigma_i)$ is the Gaussian distribution with mean vector μ_i and covariance matrix Σ_i for component i . The Expectation-Maximization (EM) algorithm is used to optimise the parameters Θ .

An autoencoder-based deep learning model is employed to extract meaningful low-dimensional representations of the dataset before clustering. The autoencoder compresses the input $X \in \mathbb{R}^d$ into a latent representation $Z \in \mathbb{R}^p$ (where $p < d$) and reconstructs the data from Z [33]. The reconstruction loss is minimised as

$$\mathcal{L}(X, \hat{X}) = \frac{1}{n} \sum_{i=1}^n \|X_i - \hat{X}_i\|^2 \quad (6)$$

Spectral networks (SpectralNet) are used to improve nonlinear feature learning, a deeper autoencoder is used, reducing the input data to a 2-dimensional embedding space $Z \in \mathbb{R}^2$. This embedding captures complex relationships in the data while preserving structure. The same reconstruction loss as Equation (6) is applied during training.

For all clustering methods described above, the Davies–Bouldin Index (DBI) is computed to evaluate the cluster compactness and separation. A lower DBI score corresponds to better clustering quality. All clustering results are visualised in 3D plots, and the clustered data is stored for further analysis.

Elbow Method: This technique helps determine the optimal number of clusters k by plotting the within-cluster sum of squares (WCSS) against different values of k [34]. The WCSS is computed as

$$\text{WCSS} = \sum_{i=1}^k \sum_{x \in C_i} \|x - \mu_i\|^2 \quad (7)$$

The “elbow point”, where the rate of decrease sharply slows, indicates the optimal k .

Davies–Bouldin Index (DBI): The DBI measures the average similarity between each cluster and its most similar one. A lower DBI indicates better clustering [35]. It is defined as

$$\text{DBI} = \frac{1}{k} \sum_{i=1}^k \max_{j \neq i} \left(\frac{\sigma_i + \sigma_j}{d_{ij}} \right) \quad (8)$$

where σ_i is the average distance between points in cluster i and its centroid, and d_{ij} is the distance between centroids of clusters i and j .

Finally, clustering outputs are visualised in frequency–voltage–ROCOF space, and the resulting cluster patterns are analysed to support interpretation of the system’s operational and anomalous behaviours. The conventional clustering methods used in this study K-Means, FCM, and GMM exhibited very low runtime overhead. For the 10 s sliding-window datasets, K-Means and FCM converged within 10 iterations, while GMM required 20 iterations. These models therefore remain suitable for near real-time execution. In contrast, the deep clustering models introduced higher computational demand. The autoencoder was trained for 50 epochs with a batch size of 256. SpectralNet exhibited the highest training cost, requiring 15 epochs (batch size 2048) due to spectral regularisation and pairwise affinity learning. While these deep models improve cluster compactness and nonlinear feature representation, their training cost limits their direct deployment in low-power, real-time environments.

3.3. Experimental Design and Setup

The experimental objective was to assess each model’s ability to capture nonlinear relationships within frequency-centric μ PMU streams and to distinguish the underlying causes of abnormal events. This design demonstrates how data-driven clustering can reveal hidden operational patterns linked to inverter behaviour and load variation, providing a more adaptive and interpretable alternative to conventional threshold-based monitoring.

In this study, one month of data has been utilised; however, based on the analysis, a specific day (18th April 2023) was selected to show transitions among three distinct operational event conditions: normal operation, hazardous event (HE), and operational event (OE). A targeted sliding window technique was used to enable accurate identification and description of these transient states. In order to localise changes in signal behaviour, including voltage sags and swells, current spikes, negative active power, and abnormal frequency deviations, short-duration segments (10 s) were iteratively analysed by this method. All three operational states were present during a specific time window (9:00 to 12:00), which was isolated using the sliding-window technique. The sliding window was then used for the clustering and multiparameter root cause analysis to be performed.

A high-resolution comparison of frequency, ROCOF, voltage, current, and power (FRVCP) parameters across three lines during an event day and a normal day is presented in the below figures. All three line voltage graphs are represented by the colour red, all

three current lines are represented by the colour blue, all different powers (real, reactive, and apparent power) are shown in green, frequency is indicated in purple, and the ROCOF is represented by orange. On a normal day (19 April), stable operational characteristics are displayed by Figure 3. Voltage and frequency maintain tight bounds across all lines, with minimal fluctuation. The current remains around a steady 51 A, and the power waveforms retain their sinusoidal patterns with expected amplitude ranges, indicating a lack of abnormal activity. The timestamps (HH:MM:SS) in Figures 3 and 4 indicate the precise recording times of the high-resolution μ PMU measurements, captured at sub-second intervals. This timing enables accurate alignment of voltage, current, power, frequency, and ROCOF variations for interpreting events. In contrast, the event-day plots (18 April; Figure 4) show transient anomalies across all phases. In particular, Line 1 and Line 2 experience a sharp voltage sag near 11:18:52, falling below 17,000 V, alongside an immediate current surge above 50 A. Consistent with this disturbance, active power drops abruptly to below -3×10^6 W, while frequency dips to approximately 49.5 Hz. ROCOF also falls to approximately -1 at 11:18:52 and then within the next millisecond (11:18:53) increases rapidly, reaching nearly $+1$.

These findings underscore the effectiveness of the high-resolution FRVCP visualisations for distinguishing transient disturbances from steady-state behaviour. Such insights are critical for real-time anomaly detection in smart grid monitoring frameworks.

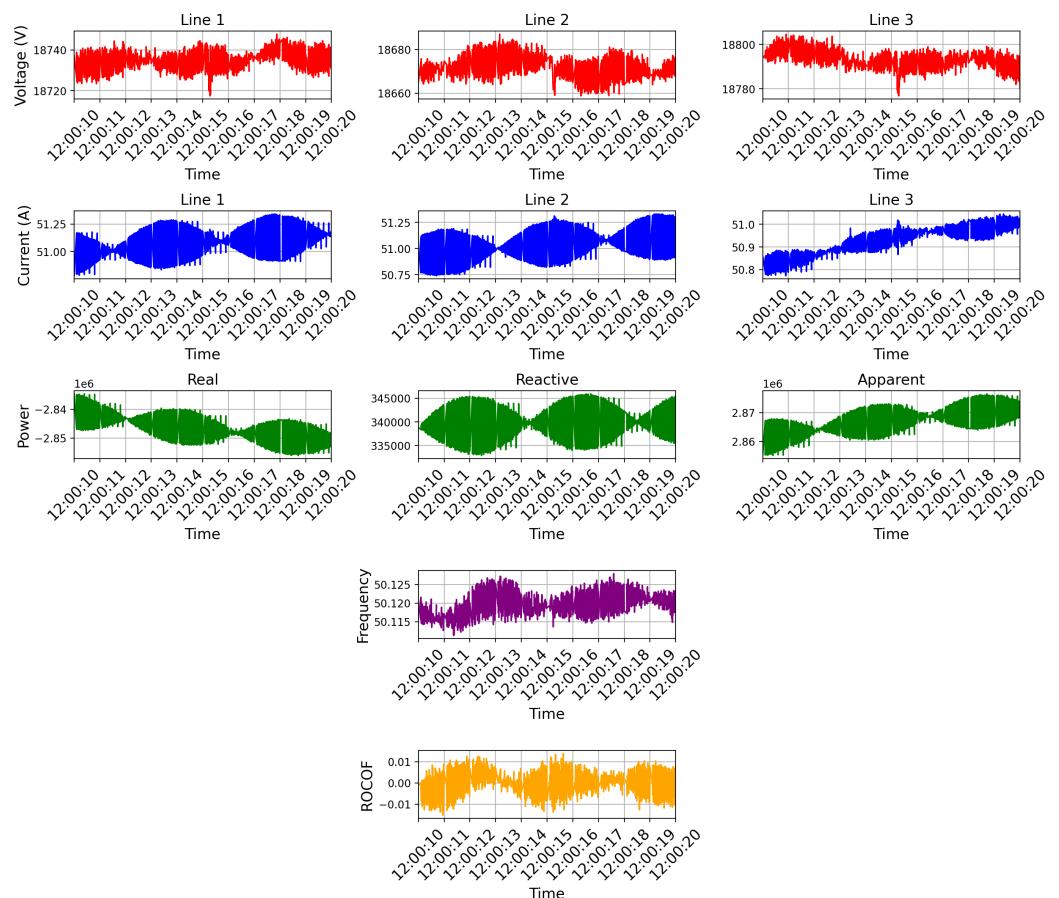


Figure 3. High-resolution FRVCP analysis showing voltage, current, power, frequency, and ROCOF across Lines 1–3 for normal day (19 April). Timestamp format HH:MM:SS.

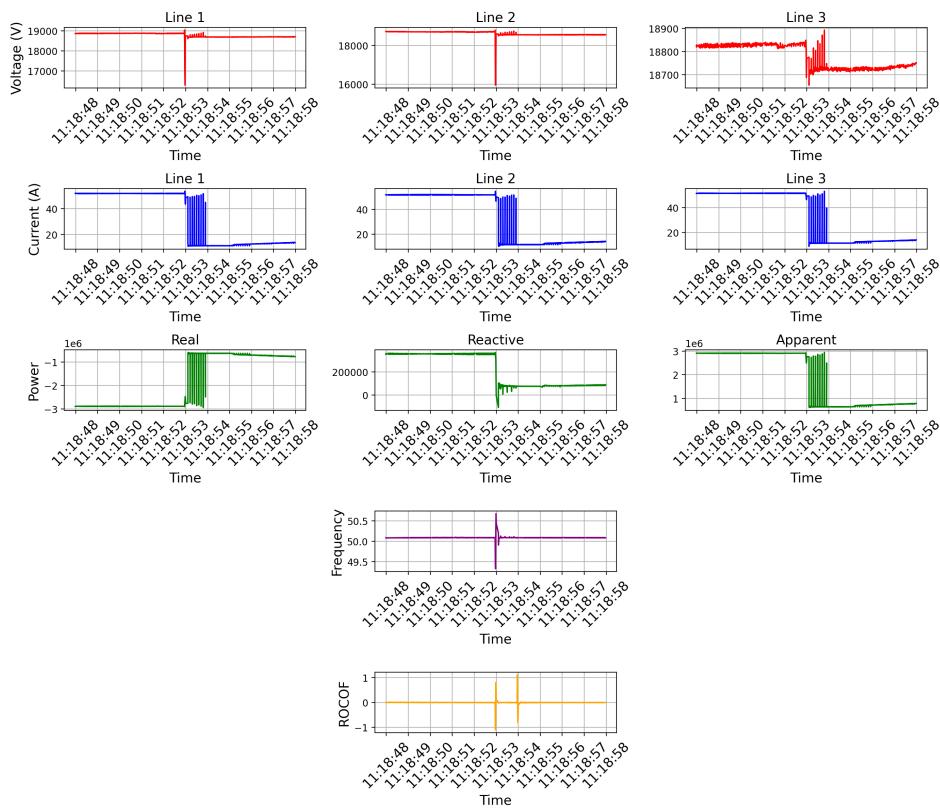


Figure 4. High-resolution FRVCP analysis showing voltage, current, power, frequency, and ROCOF across Lines 1–3 for event day (18 April). Timestamp format HH:MM:SS.

4. Results and Discussion

A detailed analysis of the study's findings is presented in this section, ensuring the conclusions are independently verifiable.

4.1. Clustering Model Comparison

A detailed discussion of the clustering of frequency, average voltage, and average current for both normal- and event-day conditions using five different models is contained in this section. Before applying clustering analyses, the optimal number of clusters was identified using the Elbow Method. As illustrated in Figure 5, the elbow point suggests that the optimal number of clusters for this study is $k = 3$.

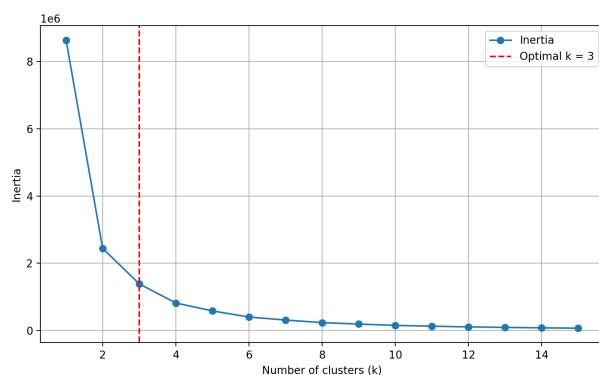


Figure 5. Elbow method graph with optimal k value.

The impact of grid disturbances on the quality of clustering is highlighted by the K-Means clustering results. Under stable grid conditions, as illustrated in Figures 6a and 7a,

under normal conditions, K-Means yields compact clusters with a DBI of 0.9188, demonstrating clear separation in the frequency voltage ROCOF space. The three clusters correspond to distinct operational states with ROCOF values remaining close to zero during stable operation. However, during the event day, performance degrades significantly with DBI increasing to 1.0245 (11.5% degradation). Sharp ROCOF spikes reaching ± 1 Hz/s create outliers that distort centroid positions, leading to broader distributions and increased overlap. This highlights the limitations of distance-based clustering when faced with rapid frequency changes and nonlinear transient dynamics.

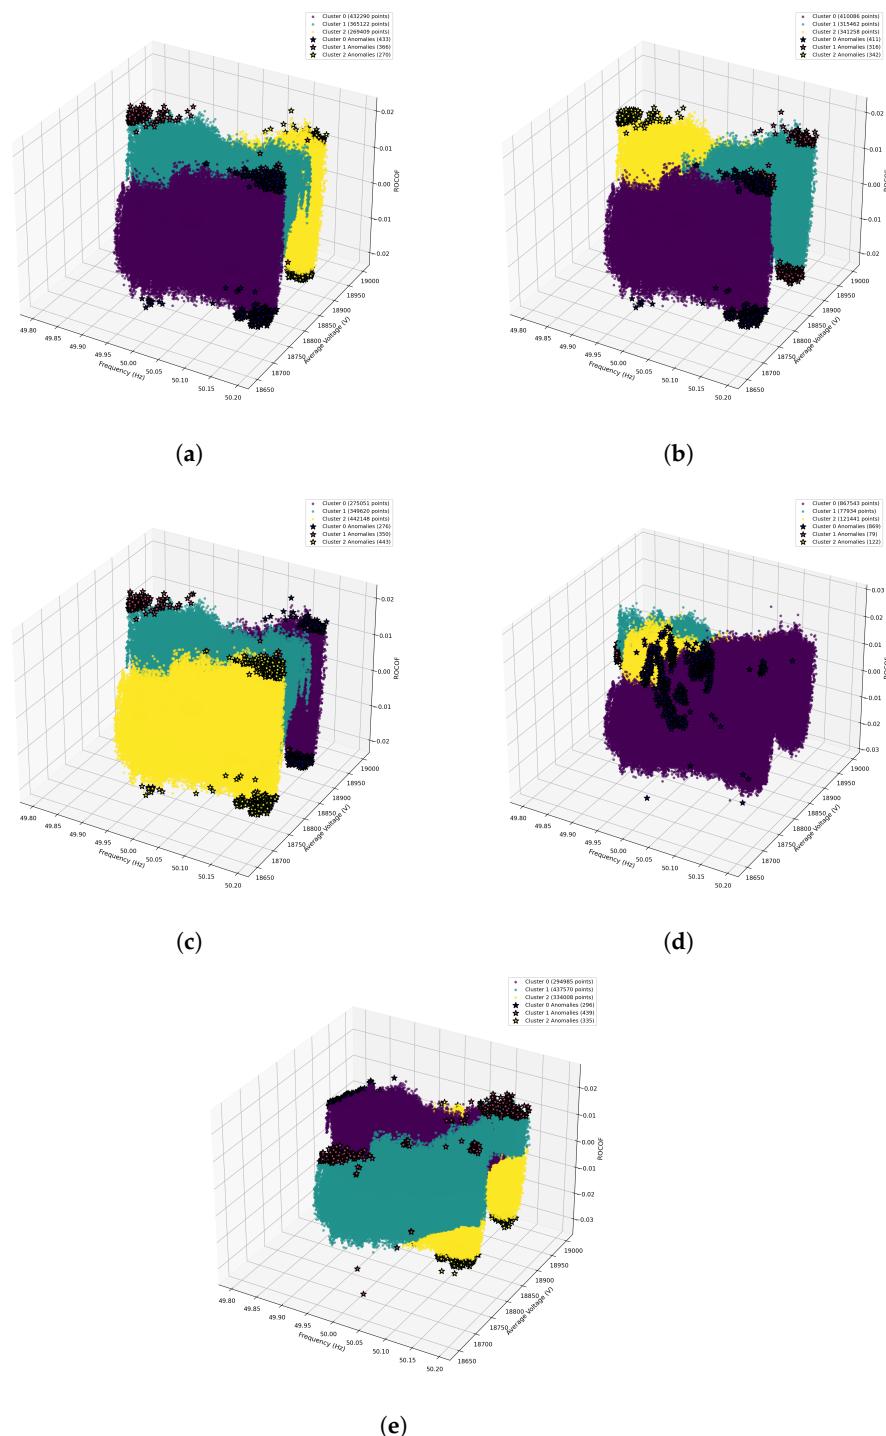


Figure 6. Comparison of clustering techniques on frequency, ROCOF and voltage data under normal conditions: (a) K-Means, (b) Fuzzy C-Means, (c) GMM, (d) autoencoder, (e) SpectralNet.

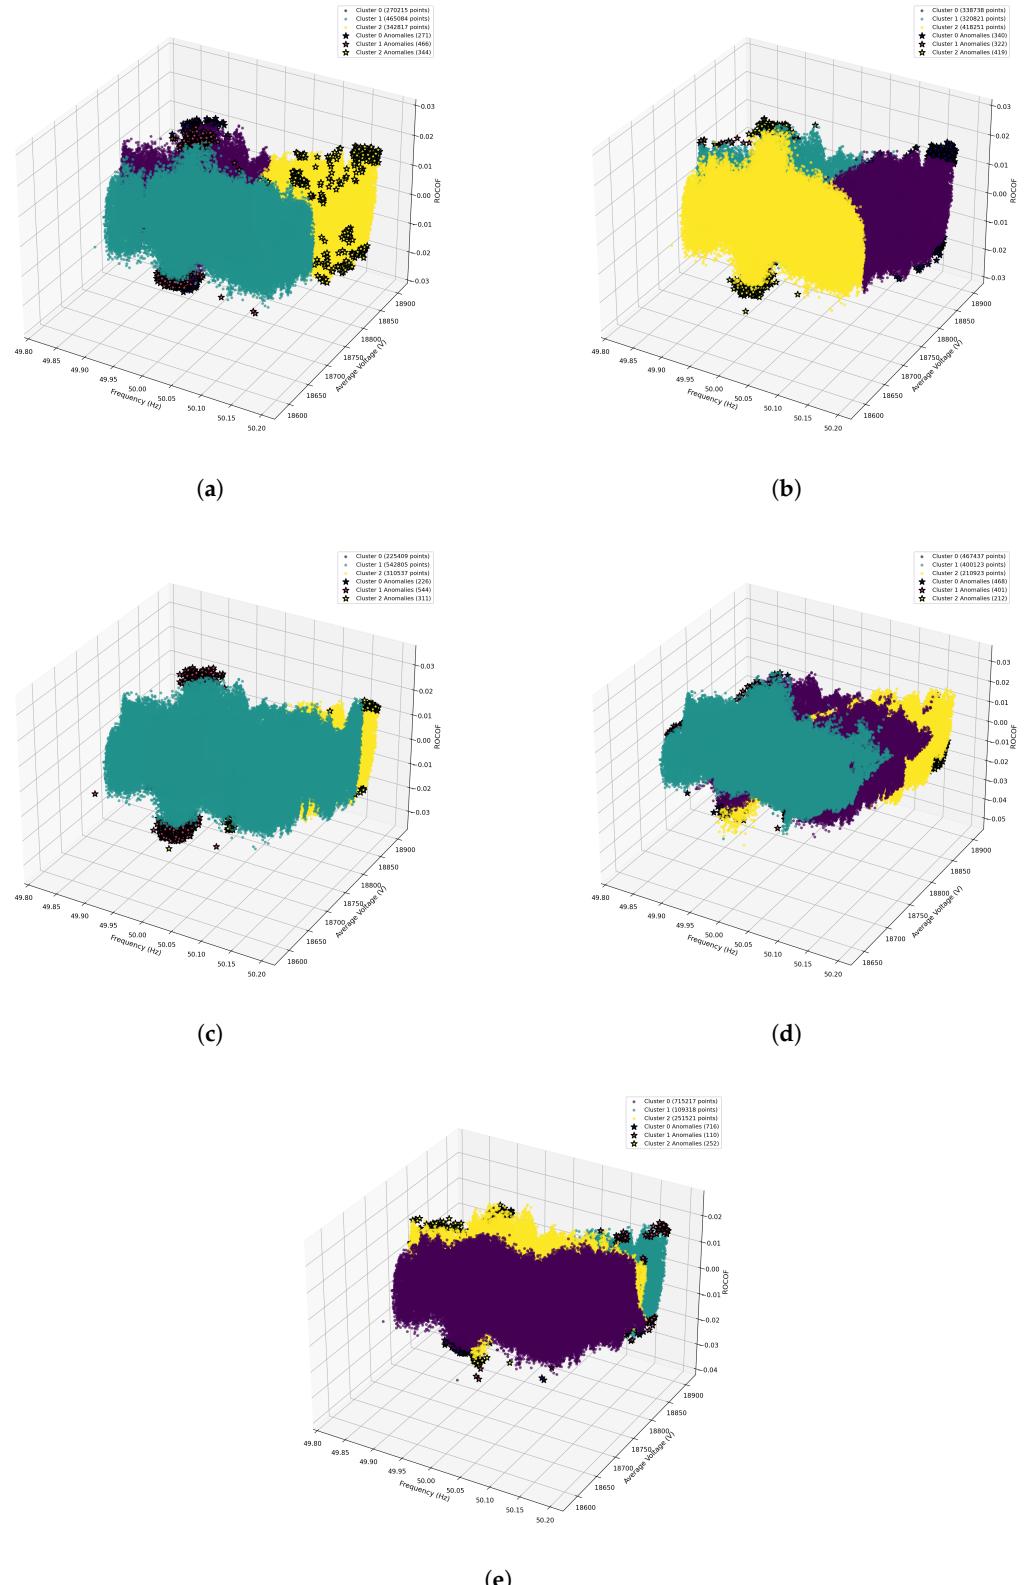


Figure 7. Clustering comparison during event scenario: (a) K-Means, (b) Fuzzy C-Means, (c) GMM, (d) autoencoder, (e) SpectralNet.

The impact of event-based data is highlighted by applying FCM clustering to both normal and event days. In the normal day (Figure 6b), FCM produces slightly dispersed groupings under normal conditions with a DBI of 0.9875, revealing transitional zones

during load ramping periods where ROCOF briefly deviates from steady-state. During the event day (Figure 7b), performance worsens with DBI increasing to 1.0707 (8.4% degradation). Data points with high ROCOF values receive distributed membership across multiple clusters, creating ambiguous classifications. While this captures the continuous spectrum of ROCOF severity, it complicates real-time decision-making and reduces diagnostic clarity for definitive event categorisation.

We compared the performance of the Gaussian Mixture Model (GMM) clustering on two different days. The GMM achieves a DBI of 0.9453 under normal conditions (Figure 6c), effectively capturing an ellipsoidal structure with ROCOF values concentrated near zero. However, the GMM shows the most significant degradation during events (Figure 7c), with DBI rising to 1.2333 (30.5% increase). The Gaussian assumption becomes severely violated as ROCOF exhibits non-Gaussian, heavy-tailed distributions during transients. Extreme ROCOF outliers exceeding ± 0.5 Hz/s force excessive covariance expansion, resulting in elongated, poorly separated clusters that merge during extreme events.

The autoencoder demonstrates superior performance under normal conditions (Figure 6d) with a DBI of 0.6137 (33% improvement over K-Means), learning nonlinear representations that capture complex frequency voltage ROCOF interactions. Critically, it maintains robust performance during events (Figure 7d) with a DBI of 0.9389, only a 53% increase and the smallest degradation among all baseline methods. The latent space representation captures ROCOF patterns as continuous trajectories rather than discrete outliers, successfully distinguishing between normal operational fluctuations (low ROCOF) and genuine hazardous events (high ROCOF), demonstrating superior adaptability to dynamic grid conditions.

Once the frequency, voltage, and ROCOF feature vector was formed using the selected time window, SpectralNet was applied to identify operating patterns and detect frequency anomalies from PMU data under normal and event-day conditions. Figures 6e and 7e show the clustering results for the normal and event-day scenarios, respectively. Three dominant clusters are observed in the frequency, voltage, and ROCOF feature space, representing distinct operating regimes. The inclusion of ROCOF enables effective separation between steady-state conditions and transient events characterised by rapid frequency changes. Anomaly points, marked by star symbols in Figures 6e and 7e, are primarily associated with elevated ROCOF values, even when frequency deviations are relatively small. The event-day scenario exhibits a higher concentration of anomaly points at the boundaries of dense clusters, indicating transient disturbances, while the normal day scenario shows fewer anomalies corresponding to short duration frequency ramps. These results confirm that incorporating ROCOF significantly improves sensitivity to frequency anomalies in unsupervised clustering-based power system monitoring.

4.2. Model Performance Comparison

Table 1 presents the Davies–Bouldin index (DBI) values for five different clustering models under normal and event-day conditions.

Table 1. DBI values of clustering models for normal and event-day data.

Category	Model	DBI (Normal Day)	DBI (Event Day)
Baseline	K-Means	1.2885	1.2985
	GMM	1.2948	1.6152
	Fuzzy C-Means	1.3564	1.3689
Deep Learning	Autoencoder	0.4208	0.6416
	SpectralNet	0.8011	0.5755

From the Table 1 and Figure 8, it is observed that the autoencoder and SpectralNet outperform traditional clustering methods in both normal and event-day scenarios, achieving significantly lower DBI values. The autoencoder achieves the lowest DBI on normal-day data (0.4208), while SpectralNet performs comparably well on event-day data (0.5755) and maintains slightly better consistency across both conditions. In contrast, conventional baseline models like K-Means, GMM, and Fuzzy C-Means show higher DBI values, indicating less compact and well-separated clusters. These results suggest that deep learning-based models, particularly SpectralNet and Autoencoder, offer more robust and reliable clustering performance for detecting anomalies in electrical systems under varying operational conditions.

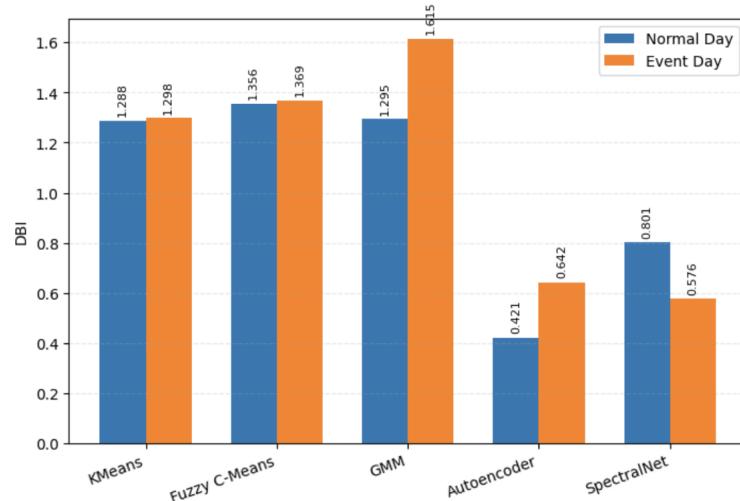


Figure 8. Comparison of DBI value of all experimental clustering models.

For operational use, real-time feasibility could be improved by training the deep models offline and deploying only the inference modules, or by adopting model-compression strategies such as pruning, quantisation, or incremental/online learning. These steps would help ensure that the improved accuracy of SpectralNet and the autoencoder can be achieved without compromising real-time responsiveness.

4.3. Root Cause Analysis: Hazardous Frequency Disturbance

To investigate the underlying causes of frequency disturbances in power grid systems, this section presents a detailed analysis of synchronised frequency data alongside all key electrical parameters, like phase currents, phase voltages, and power. By comparing anomalous events, we aim to identify patterns and correlations that reveal the origins of hazardous frequency deviations. Below is a detailed discussion one by one on all the abovementioned parameters' relationships with frequency anomalies to isolate contributing factors that highlight potential fault sources.

Phase current variations are essential for detecting frequency anomalies and preserving grid stability because they reflect shifts in load dynamics and power quality. Here, the frequency and individual phase currents are shown during an anomalous event day. During the event day (Figure 9a), there are significant spikes in the Line1 (L1), Line2 (L2), and Line3 (L3) phase currents, with values sharply increasing and fluctuating over time. The timestamps shown in these figures follow an MM:DD:HH format, indicating the month, day, and hour of each high-resolution μ PMU recording. This format allows precise temporal alignment of current, frequency, and power variations during the event. These abnormal fluctuations are accompanied by anomalies in frequency, dropping below the normal

operational range of 49.8–50.2 Hz, as indicated by the red crosses. These observations highlight the effectiveness of anomaly detection in distinguishing in power grid systems.

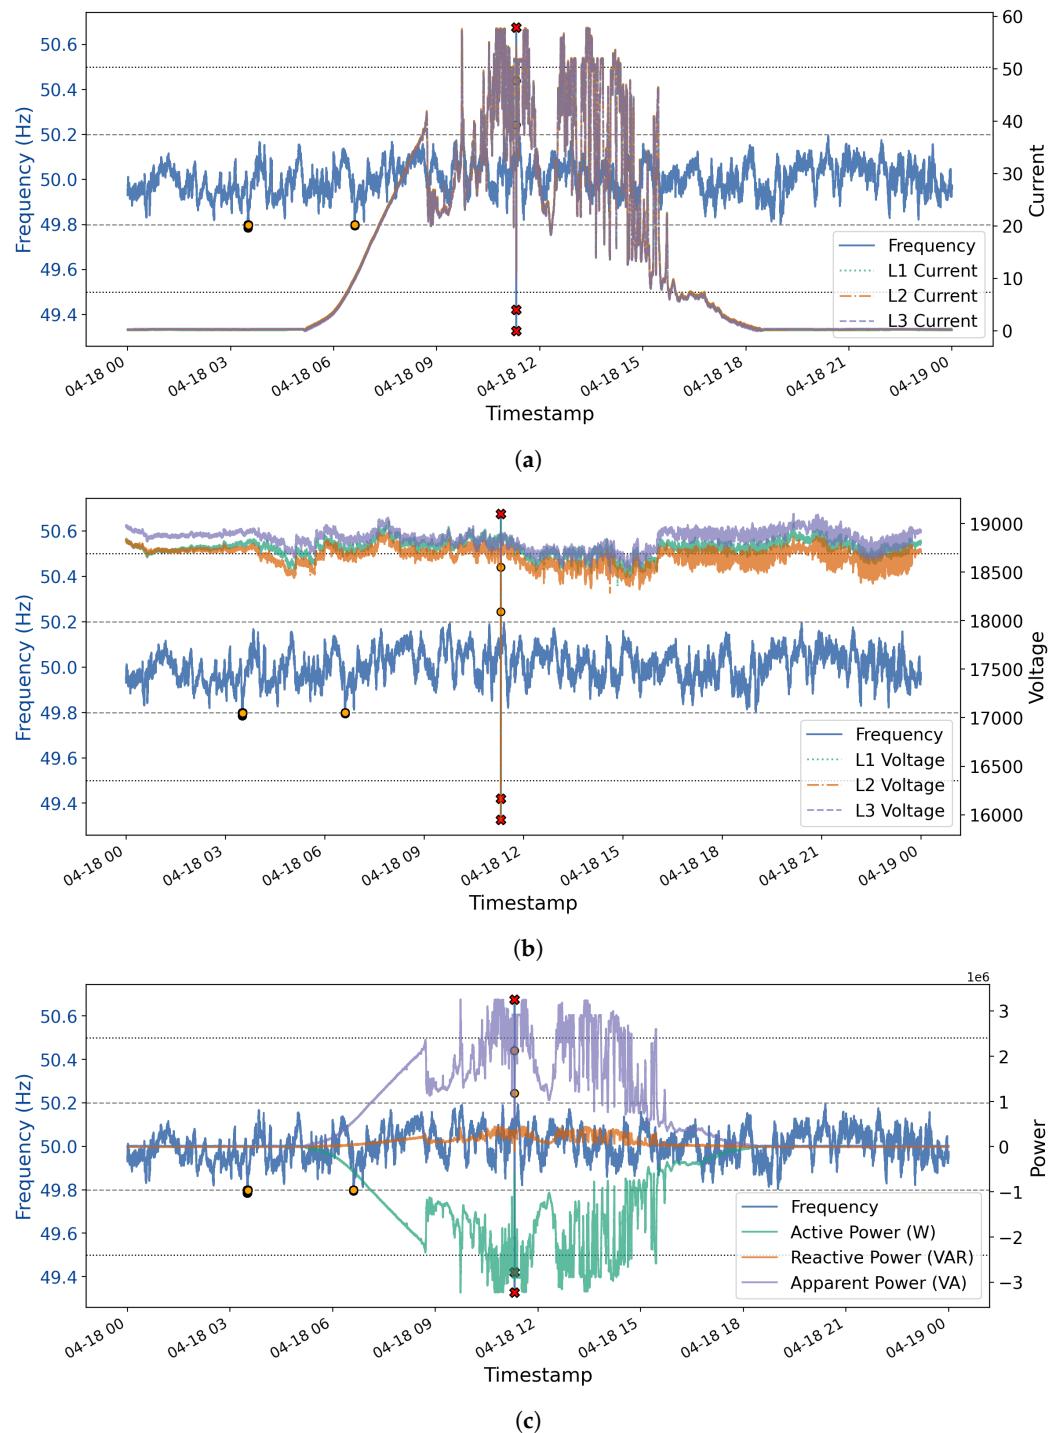


Figure 9. Comparison on event day (timestamp MM:DD:HH) with anomalies highlighted: (a) frequency and individual phase currents, (b) frequency and individual phase voltages, (c) frequency and three different powers.

Phase voltage variations reflect changes in power flow and grid operating conditions, making them critical for detecting and understanding frequency-related disturbances. We discuss frequency along with individual phase voltages (L1, L2, and L3) on the event day. In the anomalous case (Figure 9b), the same MM:DD:HH timestamp format is used to indicate when specific voltage and frequency fluctuations occur. During the anomalous

period, significant deviations in voltage profiles are observed across all three phases (green, orange, and purple, respectively, for L1, L2, and L3), particularly around the timestamp where the frequency sharply drops below the nominal threshold of 49.8 Hz. The voltage in L2 shows a sharper dip (orange line) compared to L1 and L3, highlighting an unbalanced load condition or potential fault. Anomalous instances are also marked with a red 'x' and brown circular indicators, showing operational instances. This comparison emphasises the effectiveness of voltage–frequency joint analysis in identifying power quality anomalies and underlines the temporal alignment of voltage anomalies with frequency disturbances.

Further analysis was conducted using comparative plots of frequency alongside the individual power parameters active, reactive, and apparent during an anomalous day. As shown in Figure 9c, notable deviations emerge within the MM:DD:HH timestamped anomaly period: active power exhibits sharp negative spikes, while apparent power rises significantly beyond normal operating thresholds. These shifts align with abnormal fluctuations in both frequency and reactive power, suggesting the presence of a grid disturbance or atypical load behavior. This correlation highlights the model's effectiveness in identifying and characterizing abnormal operational conditions through integrated analysis of power and frequency signals.

Our investigation identified the root cause as an unbalanced load or phase-specific fault on L2, which generates excessive current draw and voltage instability. These conditions are likely to subsequently trigger ROCOF variations and induce frequency disturbances across the entire system. The findings emphasize the vital role of comprehensive multi-parameter monitoring in achieving timely anomaly detection and precise fault localisation within power distribution systems.

5. Conclusions and Future Work

A robust clustering framework for identifying and analysing frequency anomalies in the power grid system using high-resolution μ PMU data was presented in this study. FRVCP features were extracted and analysed using unsupervised clustering techniques. Among the evaluated methods, the SpectralNet demonstrated superior performance, as measured by the Davies–Bouldin Index (DBI), particularly under event-day operating conditions. The proposed approach successfully identified transient grid disturbances and associated operational patterns, such as unbalanced loading and phase-specific faults. This study adopts a data-driven approach using real operational measurements, aiming to detect anomalous behaviour and provide actionable insights for system operators. However, some limitations should be acknowledged. First, the study was constrained by a single-site, one-month deployment, which limits seasonal and geographic generalisability. The absence of exogenous contextual data, such as weather conditions and grid event labels, restricted the rigor of validation and prevented causal attribution of detected anomalies. In addition, the computational overhead associated with deep learning models poses challenges for real-time deployment on resource-constrained edge devices. Density-based clustering algorithms, including DBSCAN and HDBSCAN, were excluded due to their quadratic time complexity ($O(n^2)$), rendering them impractical for high-volume, high-resolution time-series data streams. Furthermore, the use of UK-specific grid frequency standards (49.8–50.2 Hz) necessitates adaptation for deployment in regions with differing operational thresholds. Despite these constraints, the proposed framework demonstrates strong potential for global scalability. This can be achieved through (1) configurable parameters adapting to different frequency standards; (2) model compression (pruning, quantisation) reducing computational overhead for edge deployment; (3) federated learning helping utilities build models together whilst protecting data privacy; (4) standardised

protocols ensuring vendor-agnostic interoperability; and (5) multi-site validation across diverse climates and renewable penetration levels.

Future work will focus on integrating weather and grid event APIs to enhance contextual awareness and validation robustness; extending the framework to hybrid renewable energy systems; incorporating explainable AI techniques to improve operator trust and interpretability; expanding evaluations to multi-site and multi-regional datasets to improve generalisability; and conducting large-scale pilot deployments across different grid infrastructures. Collectively, these efforts aim to establish the proposed approach as an adaptive and globally deployable tool for real-time grid stability monitoring.

Author Contributions: Conceptualisation, A.M. and M.D.; Data curation, A.M. and M.D.; Formal analysis, A.M.; Funding acquisition, P.P.; Investigation, A.M. and M.D.; Methodology, A.M. and M.D.; Project administration, P.P.; Resources, P.P.; Software, A.M. and S.P.R.; Supervision, M.D.; Validation, M.D. and S.P.R.; Visualisation, S.P.R.; Writing—original draft, A.M. and M.D.; Writing—review and editing, P.P. and S.P.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. Data availability is restricted by the policies of the external agency for this research.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

PMU	Phasor Measurement Unit
μ PMU	Micro-Phasor Measurement Unit
GDU	Grid Data Unit
DBI	Davies–Bouldin Index
FCM	Fuzzy C-Means
GMM	Gaussian Mixture Model
OE	Operational Event
HE	Hazardous Event
OR	Operational Range
HR	Hazardous Range
FRVCP	Frequency, ROCOF, Voltage, Current, and Power
ROCOF	Rate of Change of Frequency

References

1. Arghandeh, R.; Brady, K. Micro-synchrophasors for power distribution systems. *Eng. Technol. Ref.* **2016**, *2016*, 1–16. [\[CrossRef\]](#)
2. Meydani, A.; Shahinzadeh, H.; Nafisi, H.; Gharehpetian, G.B. Synchrophasor Technology Applications and Optimal Placement of Micro-Phasor Measurement Unit (μ PMU): Part II. In Proceedings of the 2024 28th International Electrical Power Distribution Conference (EPDC), Zanjan, Iran, 23–25 April 2024; IEEE: Piscataway, NJ, USA, 2024; pp. 1–18.
3. Lal, M.D.; Varadarajan, R. A review of machine learning approaches in synchrophasor technology. *IEEE Access* **2023**, *11*, 33520–33541. [\[CrossRef\]](#)
4. Dey, M.; Rana, S.P.; Simmons, C.V.; Dudley, S. Solar farm voltage anomaly detection using high-resolution μ PMU data-driven unsupervised machine learning. *Appl. Energy* **2021**, *303*, 117656. [\[CrossRef\]](#)
5. Singh, U.; Dey, M.; Patel, P. Enabling Grid Stability: Harnessing μ PMU Data for Data-Driven Analysis of Grid Frequency Events. In *Proceedings of the International Conference on Frontiers of Intelligent Computing: Theory and Applications*; Springer: Singapore, 2024; pp. 99–112.
6. Lan, T.; Lin, Y.; Wang, J.; Leao, B.; Fradkin, D. Unsupervised power system event detection and classification using unlabeled PMU data. In Proceedings of the 2021 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe), Espoo, Finland, 18–21 October 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 1–5.

7. Guato Burgos, M.F.; Morato, J.; Vizcaino Imacaña, F.P. A review of smart grid anomaly detection approaches pertaining to artificial intelligence. *Appl. Sci.* **2024**, *14*, 1194. [[CrossRef](#)]
8. Kamal, M.B. Analysis of Synchrophasor Measurements for Cybersecurity and Situational Awareness in Power Distribution Systems. Ph.D. Thesis, University of California, Riverside, CA, USA, 2022.
9. Baba, M.; Nor, N.B.M.; Shiekh, M.A.; Alharthi, Y.Z.; Shutari, H.; Majeed, M.F. A Review on Microgrid Protection Challenges and Approaches to Address Protection Issues. *IEEE Access* **2024**, *12*, 175278–175303. [[CrossRef](#)]
10. Resende, É.C.; Simoes, M.G.; Freitas, L.C.G. Anti-islanding techniques for integration of inverter-based distributed energy resources to the electric power system. *IEEE Access* **2024**, *12*, 17195–17230. [[CrossRef](#)]
11. Dey, M.; Wickramarachchi, D.; Rana, S.P.; Simmons, C.V.; Dudley, S. Power grid frequency forecasting from μ PMU data using hybrid vector-output LSTM network. In Proceedings of the 2023 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE), Grenoble, France, 23–26 October 2023; IEEE: Piscataway, NJ, USA, 2023; pp. 1–5.
12. Jamei, M.; Scaglione, A.; Roberts, C.; Stewart, E.; Peisert, S.; McParland, C.; McEachern, A. Anomaly Detection Using Optimally Placed PMU Sensors in Distribution Grids. *IEEE Trans. Power Syst.* **2017**, *33*, 3611–3623. [[CrossRef](#)]
13. Aligholian, A.; Shahsavari, A.; Stewart, E.M.; Cortez, E.; Mohsenian-Rad, H. Unsupervised event detection, clustering, and use case exposition in micro-pmu measurements. *IEEE Trans. Smart Grid* **2021**, *12*, 3624–3636. [[CrossRef](#)]
14. Dwivedi, D.; Yemula, P.K.; Pal, M. DynamoPMU: A physics informed anomaly detection, clustering, and prediction method using nonlinear dynamics on μ PMU measurements. *IEEE Trans. Instrum. Meas.* **2023**, *72*, 3536309. [[CrossRef](#)]
15. Liu, Y.; Yang, L.; Ghasemkhani, A.; Livani, H.; Centeno, V.A.; Chen, P.Y.; Zhang, J. Robust event classification using imperfect real-world PMU data. *IEEE Internet Things J.* **2022**, *10*, 7429–7438. [[CrossRef](#)]
16. Shi, X.; Qiu, R. Dimensionality increment of PMU data for anomaly detection in low observability power systems. *arXiv* **2019**, arXiv:1910.08696. [[CrossRef](#)]
17. Geiger, A.; Liu, D.; Alnegheimish, S.; Cuesta-Infante, A.; Veeramachaneni, K. Tadgan: Time series anomaly detection using generative adversarial networks. In Proceedings of the 2020 IEEE International Conference on Big Data (Big Data), Atlanta, GA, USA, 10–13 December 2020; IEEE: Piscataway, NJ, USA, 2020; pp. 33–43.
18. Shahsavari, A.; Dubey, A.; Stewart, E.M. Situational Awareness in Distribution Grid Using Micro-PMU Data: A Machine Learning Approach. In Proceedings of the 2019 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT), Washington, DC, USA, 17–20 February 2019; IEEE: Piscataway, NJ, USA, 2019; pp. 1–5. [[CrossRef](#)]
19. Rodrigues, N.M.; Janeiro, F.M.; Ramos, P.M. Deep learning for power quality event detection and classification based on measured grid data. *IEEE Trans. Instrum. Meas.* **2023**, *72*, 9003311. [[CrossRef](#)]
20. Chandrakar, R.; Dubey, R.K.; Panigrahi, B.K. Deep-Learning based Multiple Class Events Detection and Classification using Micro-PMU Data. In Proceedings of the 2024 8th International Conference on Green Energy and Applications (ICGEA), Singapore, 14–16 March 2024; IEEE: Piscataway, NJ, USA, 2024; pp. 137–142.
21. Moazzen, F.; Hossain, M. Multivariate deep learning long short-term memory-based forecasting for microgrid energy management systems. *Energies* **2024**, *17*, 4360. [[CrossRef](#)]
22. Miyata, Y.; Ishikawa, H. Handling concept drift in data-oriented power grid operations. *Meas. Energy* **2025**, *7*, 100052. [[CrossRef](#)]
23. Khaledian, P.; Aligholian, A.; Mohsenian-Rad, H. Event-based analysis of solar power distribution feeder using micro-PMU measurements. In Proceedings of the 2021 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT), Washington, DC, USA, 16–18 February 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 1–5.
24. Dey, M.; Rana, S.P.; Wylie, J.; Simmons, C.V.; Dudley, S. Detecting power grid frequency events from μ PMU voltage phasor data using machine learning. In *Proceedings of the IET Conference Proceedings CP811*; IET: London, UK, 2022; Volume 2022, pp. 125–129.
25. Aligholian, A.; Mohsenian-Rad, H. GraphPMU: Event clustering via graph representation learning using locationally-scarce distribution-level fundamental and harmonic PMU measurements. *IEEE Trans. Smart Grid* **2022**, *14*, 2960–2972. [[CrossRef](#)]
26. Ehsani, N.; Aminifar, F.; Mohsenian-Rad, H. Convolutional autoencoder anomaly detection and classification based on distribution PMU measurements. *IET Gener. Transm. Distrib.* **2022**, *16*, 2816–2828. [[CrossRef](#)]
27. Connelly, A.C.; Zaidi, S.A.R.; McLernon, D. Autoencoder and incremental clustering-enabled anomaly detection. *Electronics* **2023**, *12*, 1970. [[CrossRef](#)]
28. Kim, J.; Lee, J.; Kang, S.; Hwang, S.; Yoon, M.; Jang, G. Probabilistic optimal power flow-based spectral clustering method considering variable renewable energy sources. *Front. Energy Res.* **2022**, *10*, 909611. [[CrossRef](#)]
29. Energy Systems Catapult. *Neuville Grid Data: Network Monitoring*; Energy Systems Catapult: Birmingham, UK, 2020.
30. Sharma, P. Comprehensive Guide to K-Means Clustering. 2019. Available online: <https://www.analyticsvidhya.com/blog/2019/08/comprehensive-guide-k-means-clustering/> (accessed on 24 September 2025).
31. Cebeci, Z.; Yildiz, F. Comparison of K-Means and Fuzzy C-Means Algorithms on Different Cluster Structures. *J. Agric. Inform.* **2015**, *6*, 13–23. [[CrossRef](#)]
32. GeeksforGeeks. Gaussian Mixture Model. GeeksforGeeks. Available online: <https://www.geeksforgeeks.org/machine-learning/gaussian-mixture-model/> (accessed on 24 September 2025).

33. Autoencoder, 2025. Available online: <https://en.wikipedia.org/wiki/Autoencoder> (accessed on 24 September 2025).
34. Tomar, A. Elbow Method in K-Means Clustering: Definition, Drawbacks, vs. Silhouette Score. *GeeksforGeeks*, 5 July 2025. Available online: <https://builtin.com/data-science/elbow-method> (accessed on 24 September 2025).
35. GeeksforGeeks. Davies–Bouldin Index. *GeeksforGeeks*. Available online: <https://www.geeksforgeeks.org/machine-learning/davies-bouldin-index/> (accessed on 24 September 2025).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.