International Journal of Applied Mathematics

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

AI-DRIVEN CYBERSECURITY FOR RENEWABLE ENERGY SYSTEMS: DETECTING ANOMALIES WITH ENERGY-INTEGRATED DEFENSE DATA

Sajib Debnath[1], Md Rashed Buiya[2], Md Maruf Bin Reza[3], Md Robiul Islam[4], Mahuma Akter[5], Aashish K C[6], Md Shohail Uddin Sarker[7], Badruddowza[8], A K M Nuruzzaman Laskar[9] and Santosh Pant[10]

[1]Master of Computer Science, Western Illinois University

[2]Master of Science in Cyber Security, California State University, Dominguez Hills

[3]Master of Science in Cyber Security, California State University, Dominguez Hills

[4]Dept. of Computer Engineering, Tennessee State University, Nashville, TN, USA

[5]Master of Science in Cybersecurity (MSCS) , Washington University of Science and Technology (WUST).

[6]Master of Science in Computer and Information Science(Software Engineering), Gannon University, Erie, PA

[7]MS- Computer and Information Systems Security, Gannon University, Pa, USA.

[8]Master of Science in Computer and Information Systems, Gannon University, Erie, Pa

[9]Master of Science in Information Assurance & Cybersecurity, Gannon University, Erie, PA

[10]Kantipur College of Management and Information Technology, Kathmandu, Nepal

Corresponding Author: **Sajib Debnath, Email:** sajib.uiu.cse@gmail.com

Abstract

Renewable-powered data centers present a dual challenge for anomaly detection in the USA. The inherent variability in solar, wind, battery, and grid signals can mask or imitate cyber intrusions, rendering network-only intrusion detection methods unreliable. This study addresses this challenge by developing a multimodal detection pipeline that combines energy telemetry with network security metrics. Using a dataset of over 2,000 operational records, which has been augmented with synthetic cyber, energy, and coordinated joint adversarial anomalies, we benchmark various classical and deep learning models, including Logistic Regression, Random Forest, XGBoost, LightGBM, and a Multi-Layer Perceptron (MLP). We evaluate these models across network-only, energy-only, and fused feature sets. Our assessment is enriched with causality-driven attribution using Granger causality and feature-level explanations via SHAP. Additionally, we test model resilience against synthetic adaptive perturbations and carry out adversarial retraining. The results indicate that fused models significantly outperform single-domain baselines, achieving higher recall and improved precision-recall performance while reducing false positives in imbalanced settings. The

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

attribution analysis uncovers temporal causal pathways linking energy fluctuations to specific network anomalies, facilitating clear classification of events as energy-driven, cyber-driven, or joint. Adversarial training also notably enhances robustness against subtle evasion attempts. Furthermore, latency benchmarks indicate the practicality of deploying compact models at the edge for near-real-time operation. Overall, these findings demonstrate that integrating energy data into anomaly detection improves detection reliability, interpretability, and resilience, providing a viable pathway toward scalable, edge-deployable cybersecurity for renewable energy infrastructures.

Keywords: Cybersecurity, Renewable Energy, Multimodal Fusion, Anomaly Detection, Adversarial

Robustness, Attribution

1. Introduction

1.1 Background and Motivation

The increasing adoption of renewable energy to power critical infrastructure, such as data centers in the USA, has introduced unprecedented operational complexities, particularly when managing the interplay between energy variability and cyber threats. In the USA, renewable energy sources, including solar and wind, inherently exhibit stochastic fluctuations due to weather conditions and environmental factors, leading to variability in the power supplied to sensitive computing infrastructures. This variability can generate spurious alerts in conventional intrusion detection systems (IDS) that rely solely on network telemetry, thus undermining the reliability of cybersecurity monitoring frameworks across U.S. facilities. Ghafouri (2018) highlights that resilient anomaly detection in cyber-physical systems requires accounting for multi-domain dependencies, as single-domain approaches are prone to misclassifying operational fluctuations as malicious activity [11]. Similarly, Ajala and Abiodun (2022) emphasize that integrating artificial intelligence and machine learning techniques into IDS can enhance threat prediction, anomaly detection, and automated response mechanisms, particularly in environments where energy dynamics influence network behavior in the USA [3]. The need for robust anomaly detection is further accentuated by the rising incidence of coordinated cyber-physical attacks targeting critical infrastructure in the United States, which can simultaneously manipulate both operational energy flows and network traffic to evade detection. Such dual-domain threats are increasingly relevant in modern renewable-powered facilities across the USA, where attackers may exploit the correlation between energy state changes and network anomalies to mask intrusions or amplify impact.

Recent developments in AI-driven monitoring in the U.S., which have demonstrated success in diverse domains such as e-commerce personalization (Ahad et al., 2025), have shown the potential to address these challenges by leveraging multimodal data streams that combine both energy and network information [1]. Deep learning architectures targeting temporal dependencies have shown substantial improvements in capturing complex patterns across sequential data, allowing for more accurate anomaly detection in U.S. systems characterized by non-stationary behavior (Shovon et al., 2025) [19]. These models can integrate energy

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

generation patterns, such as sudden solar drops or grid surges observed in U.S. renewable grids, with network indicators like failed logins or unusual packet flows, enabling a comprehensive understanding of system states. Furthermore, hybrid AI approaches that combine classical machine learning with deep architectures offer enhanced robustness and generalization, especially when synthetic or adversarial anomalies are introduced to simulate real-world threat scenarios in the USA (Alam et al., 2025) [5]. The convergence of energy-aware anomaly detection and advanced AI architectures underscores the necessity for research in the U.S. context that systematically integrates operational energy data into cybersecurity pipelines, thereby reducing false positives, improving detection recall, and enabling actionable insights for real-time defense.

1.2 Importance of This Research

The reliable operation of renewable-powered data centers in the United States is pivotal not only for advancing national energy efficiency and sustainability goals but also for maintaining U.S. cybersecurity resilience in the face of evolving threats. Conventional IDS in the U.S. often fails to account for the interactions between renewable energy fluctuations and network behavior, resulting in both missed detections and an elevated rate of false alarms. Farid et al. (2024) demonstrate that ensemble learning frameworks can significantly improve anomaly detection in IoT systems by leveraging Bayesian hyperparameter optimization, which fine-tunes model sensitivity to both operational noise and cyber anomalies [9]. Similarly, Hossain et al. (2024) explore adversarially aware machine learning in cyber-physical infrastructures, highlighting that incorporating perturbation-resilient training enhances system robustness against sophisticated attacks while preserving operational performance [14]. These findings underscore the potential for AI-powered methodologies to mitigate the dual challenge of energy variability and cyber risk in the United States.

In addition to reliability, integrating energy-aware insights into cybersecurity frameworks within U.S. renewable-powered data centers has important implications for resource optimization and cost reduction. Anomalies that are misattributed or undetected can trigger unnecessary operational interventions or result in downtime, both of which carry significant financial and environmental costs in the U.S. context. By modeling the temporal and causal relationships between energy metrics and network events, researchers in the United States can achieve more precise anomaly attribution, distinguishing between energy-driven anomalies, cyber-driven events, and joint occurrences. Recent AI applications in energy management, such as intelligent streetlight control and predictive load balancing (Shovon et al., 2025; Alam et al., 2025) [19,5], exemplify the utility of combining real-time energy telemetry with machine learning to optimize operational outcomes. This aligns with the broader U.S. trend of utilizing AI for sustainable decision-making, as seen in applications ranging from intelligent energy management to assessing the financial impact of environmental factors (Khan et al., 2025; Shovon et al., 2025; Alam et al., 2025) [16, 19, 5]. Furthermore, blockchain-enabled secure energy transactions (Khan et al., 2025) and AI-driven fault detection in gas turbine engines (Amjad et al., 2025) provide additional evidence, relevant to the U.S. energy sector, of the

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

benefits of hybrid energy-cyber AI systems, demonstrating that accurate anomaly detection can enhance both system security and operational efficiency [16,6].

This research is therefore critical for the United States in bridging a notable gap: the absence of methodologies that simultaneously leverage energy and network signals for anomaly detection in renewable-powered environments. Beyond detection, understanding the causality and attribution of anomalies supports actionable decision-making for U.S. operators, enabling them to deploy mitigation strategies that are proportional to the root cause, whether it is cyber, energy, or a combination of both. This dual-domain perspective is essential for designing resilient, sustainable, and secure infrastructure in the U.S. that can withstand increasingly sophisticated cyber-physical threats while maintaining energy efficiency and operational continuity.

1.3 Research Objectives and Contributions

The primary objective of this research is to develop a multimodal anomaly detection framework that effectively fuses energy telemetry with network security metrics to enhance the detection of anomalies in renewable-powered data centers in the U.S.. This involves generating synthetic anomalies that reflect realistic cyber, energy, and joint adversarial perturbations, enabling the creation of a comprehensive evaluation dataset that challenges conventional models and simulates real-world threats. By implementing causality-driven attribution using Granger causality analysis and integrating feature importance metrics derived from SHAP explanations, the study provides a sophisticated mechanism for understanding the sources of detected anomalies and their operational significance. Key contributions include the construction of a large-scale multimodal dataset augmented with synthetic anomalies, allowing for robust model evaluation and benchmarking. The research compares classical machine learning approaches, such as Random Forest and XGBoost, with deep learning architectures including MLPs, LSTMs, and hybrid temporal models, across network-only, energy-only, and fused feature sets. Performance metrics are analyzed not only in terms of traditional classification outcomes like accuracy and F1-score but also in terms of adversarial robustness, demonstrating how retraining and adaptive techniques improve resilience against subtle, targeted attacks. Additionally, this work introduces an integrated explainability and attribution framework, combining causal inference with feature-level interpretability to provide actionable insights into anomaly sources. By emphasizing practical deployment considerations, including inference latency measurement on edge devices, this research lays a foundation for scalable, real-world cybersecurity solutions tailored to renewable-powered critical infrastructure.

2. Literature Review

2.1 Cybersecurity in Renewable Energy Systems

Cybersecurity in renewable energy systems in the U.S. has emerged as a critical research area due to the increasing interconnection of energy generation, storage, and consumption infrastructure with communication and control networks. Smart grids, microgrids, and renewable-powered data centers in the U.S. rely on Supervisory Control and Data Acquisition (SCADA) systems to monitor and manage energy flows. However, the integration of digital

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

control and networked communication exposes these systems in the U.S. to potential cyber threats, which can range from denial-of-service attacks to sophisticated malware and coordinated cyber-physical intrusions. Ghafouri (2018) emphasizes that anomaly detection in cyber-physical systems in the U.S. must account for the interplay between operational variability and network behavior, as conventional IDS focusing solely on network traffic often fails to detect attacks masked by natural fluctuations in energy generation [11]. Ajala and Abiodun (2022) argue that leveraging AI and machine learning for anomaly detection and threat prediction enables automated, adaptive responses to emerging threats, which is particularly important in renewable energy systems in the U.S., where operational dynamics can inadvertently trigger false positives [3].

Recent studies have demonstrated the feasibility of combining domain-specific knowledge with predictive analytics to enhance the resilience of energy systems in the U.S. against cyber attacks. Das et al. (2025) illustrate the potential of AI-driven cybersecurity frameworks that integrate predictive analytics with anomaly detection to create a more robust defense against both known and unknown threats [8]. Furthermore, blockchain-based approaches have been proposed in the U.S. to secure energy transactions and enhance system integrity by providing immutable audit trails and decentralized validation mechanisms (Khan et al., 2025; Sultana et al., 2025) [17,20]. While these solutions improve trust and transparency, they often overlook the temporal dependencies between energy state changes and network anomalies in the U.S., which are critical for early detection of complex, coordinated attacks. Operational variability in renewable energy in the U.S., such as sudden drops in solar or wind generation or fluctuations in battery charge, can mimic malicious behavior, thereby complicating detection strategies (Shovon et al., 2025; Ahmed et al., 2025) [19, 2].

2.2 Multimodal Anomaly Detection

In the U.S., multimodal anomaly detection has emerged as a promising strategy to address the limitations of single-domain approaches in cyber-physical systems. By integrating heterogeneous data streams, including operational telemetry, network metrics, and environmental conditions, researchers can achieve higher accuracy in identifying anomalies that would otherwise be missed by single-source models. Akram et al. (2024) propose DroneSSL, a self-supervised multimodal anomaly detection framework for Internet of Drone Things, demonstrating in the U.S. context the advantages of cross-domain feature fusion in detecting subtle deviations under partial observability [4]. Similarly, Arjunan et al. (2024) explore multi-user energy consumption monitoring and anomaly detection using partial context information, highlighting that combining multiple data modalities improves detection precision and reduces false alarms in U.S. energy-intensive environments [7].

In the context of renewable-powered data centers in the U.S., multimodal approaches enable the fusion of energy variables, such as solar generation, wind output, and battery state-of-charge, with network indicators, including failed logins, packet size anomalies, and unusual protocol usage. Shovon et al. (2025) demonstrate that hybrid CNN-LSTM models, which extract local patterns through convolutional filters before encoding temporal dependencies with sequential layers, provide enhanced robustness to noisy or incomplete data in U.S. operational

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

settings [19]. Alam et al. (2025) further show that feature engineering techniques, including rolling statistics, ratios, and temporal derivatives, can significantly improve anomaly separability when multiple U.S.-based data modalities are available [5]. Despite these advances, current multimodal frameworks in the U.S. often struggle with scalability and robustness against adversarial perturbations, especially when synthetic or adaptive attacks are introduced to mimic real-world conditions (Das et al., 2025) [8]. The literature suggests that achieving reliable anomaly detection in U.S. renewable energy infrastructures requires not only sophisticated model architectures but also careful dataset construction, cross-domain feature selection, and evaluation under diverse operational and threat scenarios.

2.3 Adversarial ML in Cyber-Physical Systems

Adversarial machine learning has become a critical consideration in securing U.S. cyber-physical systems, where attackers can subtly manipulate input features to evade detection. In U.S. energy-driven infrastructure, these adversarial perturbations may involve coordinated changes in energy production and network behavior to mask malicious activity. Hossain et al. (2024) emphasize that differentially-private federated learning approaches must account for adversarial inputs to maintain both privacy and security in critical U.S. infrastructures [14]. Similarly, Farid et al. (2024) highlight that ensemble learning frameworks benefit from sensitivity analyses that anticipate attack vectors and adapt hyperparameters to enhance anomaly detection in U.S. IoT and energy systems [9]. Empirical studies show that adversarial attacks can significantly degrade the performance of both classical and deep learning models in the U.S. if not appropriately mitigated. Khan et al. (2025) demonstrate that combining AI-driven anomaly detection with blockchain-based validation can increase resilience to manipulation in U.S. energy systems, though such strategies often require additional computational resources [16]. While existing research predominantly focuses on network-based or energy-based perturbations separately, there is a notable lack of U.S.-focused frameworks addressing energy-aware adversarial attacks, where attackers exploit correlations between operational states and network events to bypass detection mechanisms (Das et al., 2025) [8].

2.4 Explainability and Attribution

Explainability and attribution are essential in AI-driven anomaly detection, particularly when models are deployed in critical energy systems where operational decisions carry significant risk. Post-hoc interpretability methods such as SHAP and LIME allow practitioners to understand the contribution of individual features to model predictions, enabling the distinction between energy-driven, cyber-driven, and joint anomalies (Das et al., 2025) [8]. Attention mechanisms in temporal models further enhance interpretability by highlighting relevant time steps in sequential data, providing insight into how anomalies develop over time. Causal inference techniques, including Granger causality, have been applied to quantify temporal dependencies between energy and network variables. This combination of causality and feature-level attribution allows for more nuanced explanations of anomaly sources, improving stakeholder trust and operational decision-making (Shovon et al., 2025; Alam et al., 2025) [19,5]. While the utility of these methods is well-documented in general industrial systems,

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

few studies have integrated explainability frameworks into multimodal cyber-energy detection pipelines. The absence of systematically attributed anomalies in synchronized datasets remains a key challenge, limiting both the operational insight gained from AI predictions and the ability to fine-tune mitigation strategies effectively.

2.5 Research Gaps

Despite significant advances in anomaly detection, several gaps persist in the context of renewable-powered cyber-physical systems. First, there is a scarcity of publicly available datasets that combine synchronized energy and network telemetry, limiting the ability to benchmark multimodal detection frameworks under realistic operational conditions (Ajala & Abiodun, 2022; Ghafouri, 2018) [3,11]. Existing studies primarily focus on single-domain data or artificially constructed scenarios, which may fail to capture the complex interdependencies observed in live environments. Second, while adversarial machine learning techniques have been explored, there is limited work on developing models that are both adversary-aware and multimodal, leaving systems vulnerable to subtle perturbations that exploit correlations between energy states and network activity (Akram et al., 2024; Arjunan et al., 2024) [4,7]. Third, the lack of systematic attribution frameworks hinders the interpretability of model outputs. Current explainability methods are applied in isolation, often providing feature importance without integrating temporal causality, which restricts actionable insights for operators and limits the adoption of AI-driven cybersecurity solutions in real-world renewable-powered facilities. Addressing these gaps is essential to developing robust, interpretable, and deployable anomaly detection systems that can reliably secure critical infrastructure while accounting for the inherent variability of renewable energy sources.

3. Methodology

3.1 Dataset and Feature Design

The dataset used in this study comprises over 2000 operational records collected from renewable-powered data centers. Each record contains a comprehensive suite of features categorized into energy, cyber, and derived metrics to facilitate a detailed understanding of anomalies and interactions between domains. The energy-related features capture core metrics of power generation and consumption, including solar power output, wind power output, contributions from the electrical grid, battery storage levels, and the instantaneous load of the data center. These metrics provide a direct view of the operational variability and potential vulnerabilities introduced by renewable sources. Cyber features capture network activity and security events that may indicate intrusions or other anomalies, including the number of failed login attempts, packet size, connection duration, network protocol type, and access-related information such as access type and access status. These features are critical for modeling typical network behavior and detecting deviations that may signal malicious activity. In addition to these primary features, several derived features were engineered to enrich the dataset and highlight complex relationships. For example, ratios of different energy sources, such as solar-to-wind output, were calculated to monitor relative contributions and sudden imbalances. Grid share and battery State of Charge (SoC) were derived to quantify reliance on

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

external power and battery utilization, respectively. Network-derived metrics, including packet rate and failed login rate, were also computed to provide dynamic measures of activity and potential stress on the system. The combination of energy, cyber, and derived features enables a holistic view of operational conditions and facilitates the detection of anomalies arising from either domain or their interactions.

3.2 Data Preprocessing

To prepare the dataset for modeling and ensure the reliability of subsequent analyses, a rigorous preprocessing pipeline was applied. The first step involved timestamp parsing and time-series alignment. The 'Timestamp' column was converted into datetime objects to enable proper handling of temporal data, allowing models to capture trends and patterns over time. The dataset was then sorted chronologically to preserve temporal dependencies and facilitate sequential analysis. Categorical variables, including 'Protocol', 'Access_Type', and 'Access_Status', were transformed into numerical representations using one-hot encoding, producing binary columns for each category to ensure compatibility with machine learning algorithms. Numerical variables were normalized using MinMax scaling, which maps all feature values to a fixed range of 0 to 1. This normalization is essential to prevent variables with large magnitudes from disproportionately influencing the model during training. Beyond these standard preprocessing steps, additional feature engineering was conducted to capture the interactions between energy and cyber domains. This included calculating ratios of energy sources, deriving grid share, estimating battery State of Charge (SoC), and computing network activity rates such as packet rate and failed login rate. These derived features were explicitly designed to highlight potential dependencies and complex anomalies that could emerge from the interplay of energy fluctuations and network behaviors. Through these preprocessing and feature engineering steps, the dataset was transformed into a structured, scalable format suitable for both classical machine learning and deep learning models, providing a robust foundation for subsequent analysis, anomaly detection, and evaluation.

3.3 Exploratory Data Analysis

Exploratory Data Analysis (EDA) was conducted to comprehensively examine the dataset, uncover patterns, and gain insights into the distribution of anomalies across both the energy and cyber domains. This phase aimed to inform subsequent modeling decisions, particularly in feature selection, augmentation, and anomaly detection strategies, by understanding the intrinsic relationships, temporal dependencies, and variability inherent in the operational records. The global summaries of the dataset revealed a pronounced class imbalance, where normal instances overwhelmingly outnumbered anomalous instances. This imbalance, reflected in the distribution of the 'Anomaly_Flag', is characteristic of anomaly detection problems and underscores the importance of targeted approaches to handle minority class instances effectively. High imbalance can lead models to overfit to the majority class, resulting in poor sensitivity to rare but critical anomalies.

Received: August 09, 2025

International Journal of Applied Mathematics

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

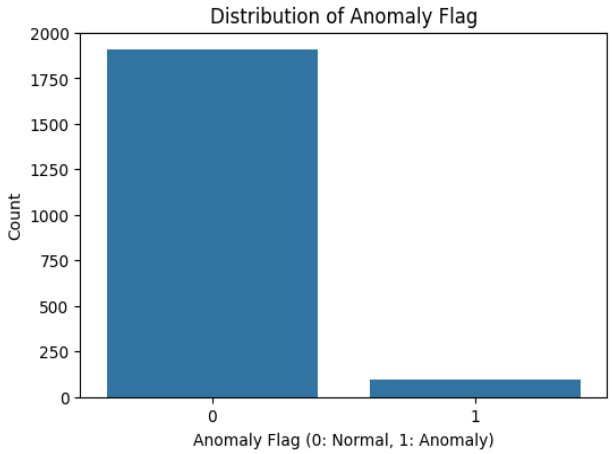

Fig.1: Anomaly distribution

Univariate analyses of numeric features highlighted heterogeneous distributions, with several features exhibiting skewness or multimodal patterns. Network metrics such as 'Packet_Size_bytes', 'Connection_Duration_s', and 'Failed_Login_Attempts' predominantly peaked near zero, reflecting baseline normal activity levels within the data center, while higher values pointed to potential abnormal or attack behavior. Temporal features like 'hour_of_day', 'day_of_week', and 'month' exhibited expected distributions based on operational schedules and diurnal patterns, which can influence both energy availability and network load. Derived features, including 'packet_rate' and 'failed_login_rate', captured dynamic behavioral characteristics that amplified the visibility of unusual activity when values deviated from normative ranges. These global summary statistics laid the groundwork for understanding feature relevance and potential predictors of anomalies.

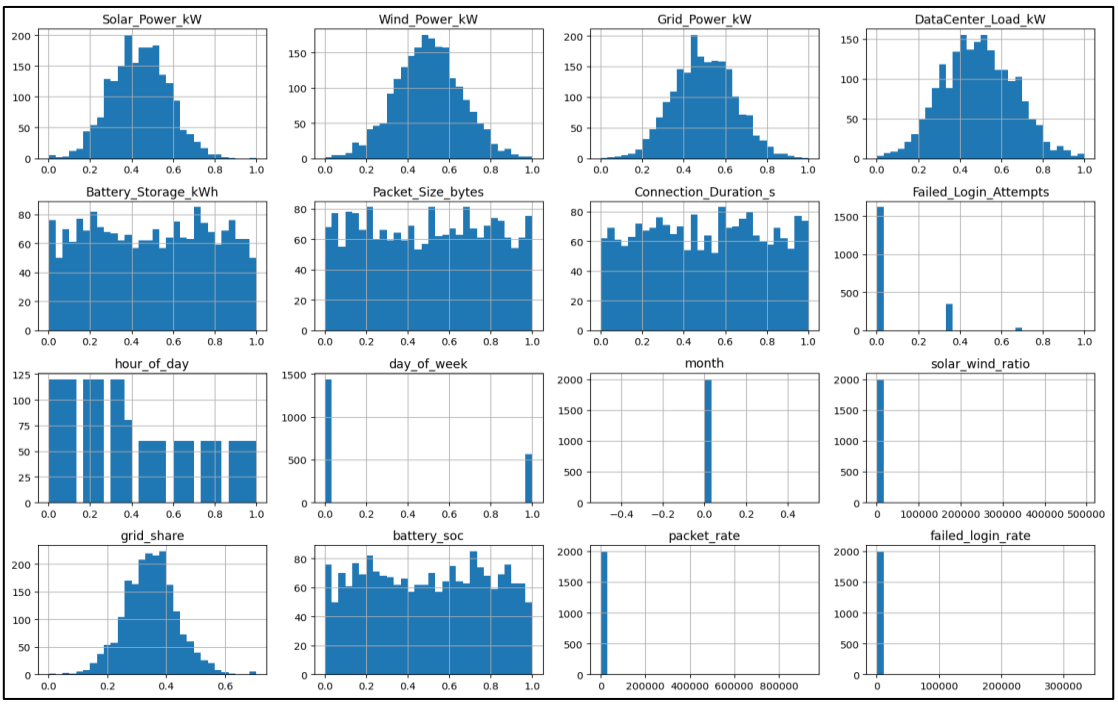

Fig.2: Univariate analysis of numeric features

Received: August 09, 2025

International Journal of Applied Mathematics

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Time-series overlays of energy signals paired with anomaly flags provided a nuanced understanding of temporal interactions between energy dynamics and cyber events. Energy metrics such as solar, wind, grid power, data center load, and battery storage displayed expected periodic patterns, including diurnal fluctuations and operational cycles. Notably, anomalies frequently aligned with sudden drops in renewable generation, spikes in grid power, or abrupt changes in battery storage levels, suggesting that energy system variability could act as an early indicator or correlate of anomalous cyber behavior. The visual inspection revealed that many anomalous events coincided with stress conditions in the energy subsystem, indicating the importance of incorporating energy context in anomaly detection models.

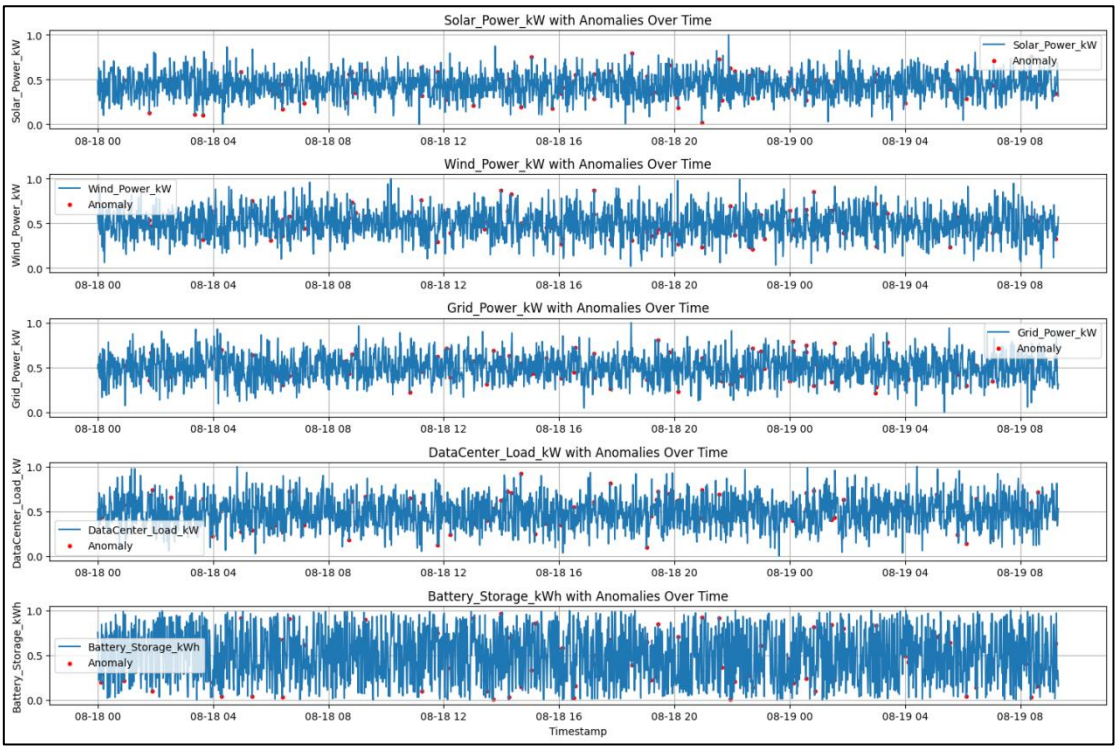

Fig.3: Energy signals and overlay anomaly flags

Cross-correlation analysis was employed to uncover temporal dependencies between energy and network security metrics. Lagged correlations indicated that changes in energy features could precede or follow shifts in network activity, highlighting potential causal pathways. For instance, correlations between 'Grid_Power_kW' and 'Failed_Login_Attempts' suggested that periods of elevated grid power demand sometimes coincided with increased login failures, potentially due to system stress or coordinated attack activity timed with energy fluctuations. These insights support the hypothesis that integrating energy and cyber data provides richer predictive capability than network-only models.

Received: August 09, 2025

International Journal of Applied Mathematics

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

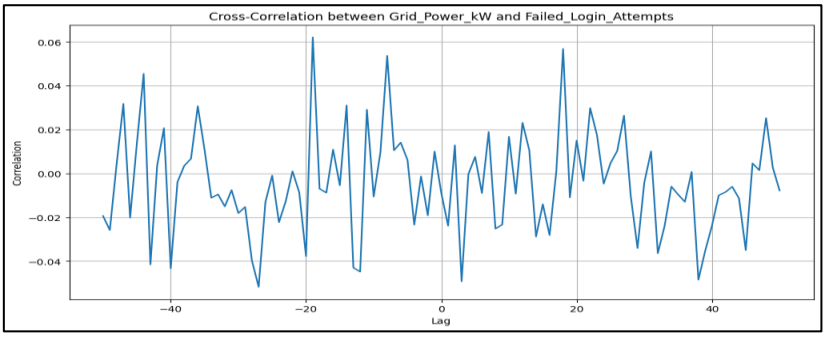

Fig.4: Compute cross-correlation between energy features and network security metrics at multiple lags.

Conditional analyses further elucidated the impact of energy system states on network behavior, with clear differences emerging between low and high battery State of Charge (SoC). The distribution of packet sizes and connection durations remained broadly similar across SoC levels, but subtle shifts in density suggest that low SoC conditions correspond to a higher proportion of smaller packets and shorter connections. This could indicate less stable or interrupted network traffic during energy-constrained periods. More pronounced differences were observed in failed login attempts and related rates. Under low SoC, there was a concentration of non-zero failed login attempts and a higher distribution of failed login rates, contrasting with the near-baseline distributions during high SoC. Similarly, packet rates exhibited greater density at higher values under low SoC, suggesting bursts of network activity coinciding with energy scarcity. These results highlight that low energy reserves are not only associated with more frequent anomalous login behaviors but also with elevated traffic activity. Such conditions amplify the likelihood of cyber anomalies or make them more detectable, as normal network baselines are disrupted by energy-induced stress. These findings reinforce the necessity of fused feature representations that jointly capture energy and cyber signals, as energy states clearly modulate the manifestation and observability of network anomalies.

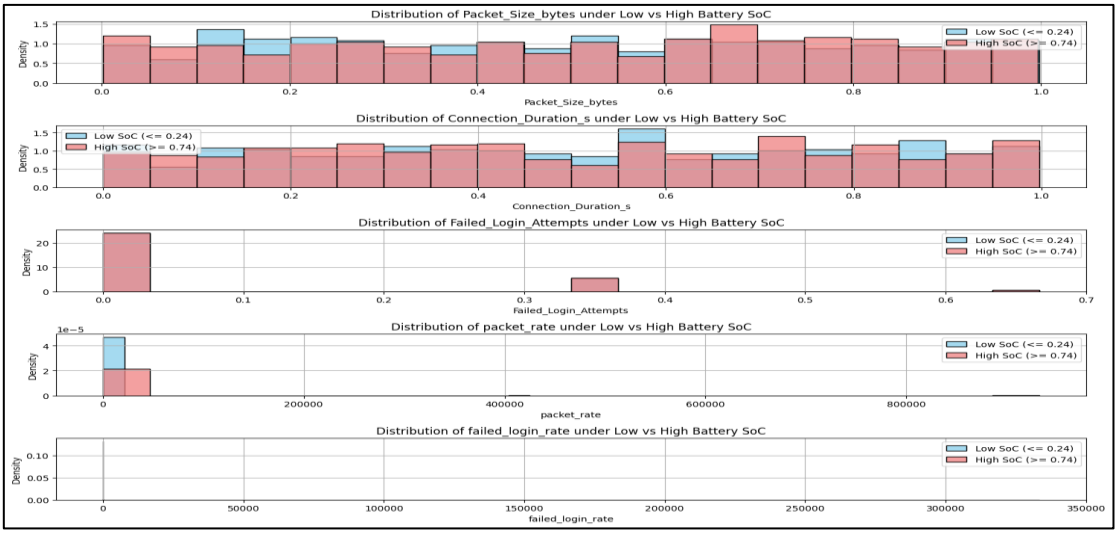

Fig.5: distributions of network metrics when battery SoC is low versus high.

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Finally, dimensionality reduction techniques such as PCA and UMAP highlighted the separability of anomalous and normal instances in a lower-dimensional space. The PCA projection (left panel) revealed that while anomalies (orange points) were partially distinguishable from normal samples, they remained scattered within broad variance ranges. The spread of points along the first two principal components suggested that linear variance captured some anomaly-related structure, but with limited clustering clarity due to the scale dominance of a few outlying points. This indicates that PCA, being a linear projection method, is sensitive to extreme feature magnitudes and may under-represent localized anomaly structures. In contrast, the UMAP projection (right panel) preserved local neighborhood structure and uncovered clearer separation. Anomalous points concentrated in distinct subregions of the manifold, often at the periphery of dense normal clusters. Several small, well-defined anomaly-heavy pockets emerged, suggesting that the fused cyber-energy features encode meaningful nonlinear distinctions between anomalous and benign behaviors. The preservation of local topology by UMAP further revealed continuity in how anomalies transition from normal operating states, making it particularly useful for downstream tasks like semi-supervised anomaly detection. Manual inspections of the flagged anomalies corroborated these findings, verifying that instances appearing in UMAP-separated clusters often aligned with irregular energy fluctuations or atypical network traffic bursts. These results demonstrate that fused multimodal features create discriminative embeddings suitable for anomaly detection, with UMAP offering stronger practical utility than PCA in capturing nonlinear anomaly structures within renewable-powered cyber-physical systems.

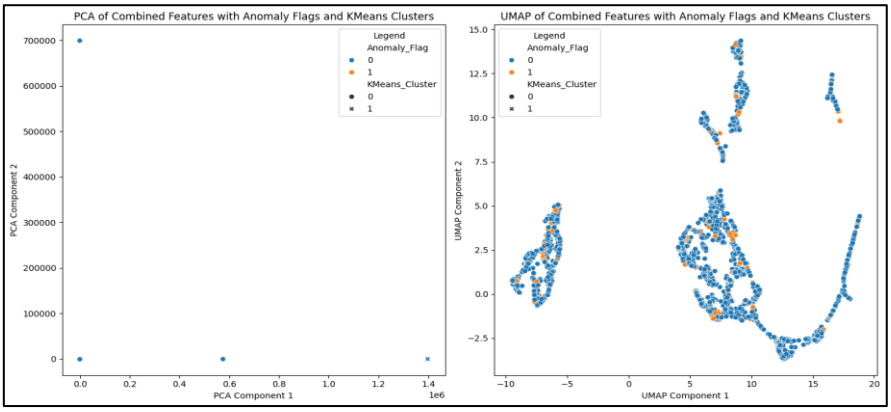

Fig.6: Use of UMAP and PCA on fused features to inspect separability of anomalies.

3.4 Data Augmentation

To address the inherent class imbalance observed in the dataset and enhance the robustness and generalizability of our anomaly detection models, a dedicated data augmentation and synthetic attack generation phase was implemented. The significant disparity between normal and anomalous instances, as revealed during the EDA, necessitated strategies that could enrich the minority class while maintaining realistic operational characteristics. Initially, the Synthetic Minority Over-sampling Technique (SMOTE) and Adaptive Synthetic Sampling (ADASYN) were applied to create synthetic anomaly examples. These techniques generate new samples in the feature space by interpolating between existing minority class instances, thereby increasing

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

the representation of anomalous events and reducing model bias toward the majority class. By doing so, we ensured that models were exposed to sufficient diversity of anomalous behaviors without artificially inflating the normal class distribution, a critical consideration for reliable training.

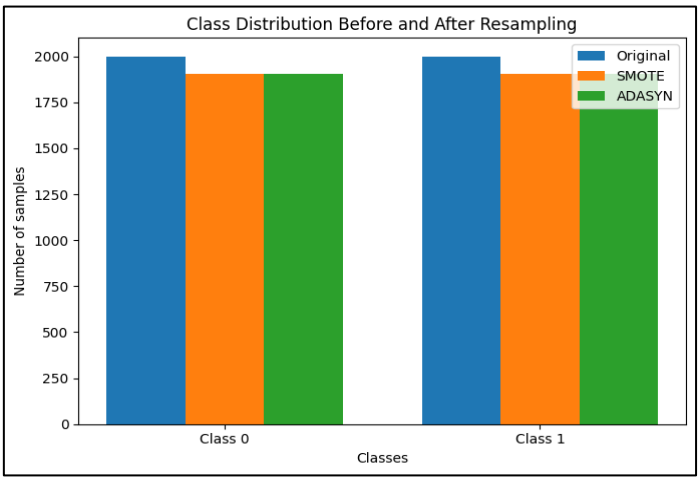

Fig.7: SMOTE and ADYSAN samples

In addition to general oversampling, domain-specific synthetic anomalies were generated to capture realistic threat scenarios in both the cyber and energy subsystems. For the cyber domain, synthetic anomalies were introduced by simulating failed login floods, abnormal packet size bursts, and elevated connection durations, reflecting plausible attack patterns such as brute-force attempts or denial-of-service vectors. These synthetic perturbations were carefully calibrated based on statistical properties of the original dataset to maintain realism while providing sufficient variance for training. Similarly, in the energy domain, synthetic anomalies were created by perturbing renewable energy generation metrics, simulating events such as sudden drops in solar PV output, abrupt wind generation fluctuations, grid failovers, or battery over-discharge incidents. These synthetic events were designed to mimic realistic operational stresses that could coincide with or trigger cyber anomalies, reflecting the interconnected nature of modern renewable-powered data centers.

To simulate more sophisticated, coordinated threats, joint anomalies were generated by synchronizing perturbations across both energy and cyber features. This approach modeled multi-domain adversarial scenarios where energy system disruptions might coincide with targeted cyber attacks, representing complex, stealthy attacks that could bypass traditional detection mechanisms. Each synthetic instance, whether cyber-specific, energy-specific, or joint, was accompanied by metadata capturing attack type, perturbation magnitude, temporal offset, and other parameters to facilitate reproducibility and detailed analysis. The augmented dataset resulting from these procedures provided a larger, balanced, and diverse set of anomaly instances, enabling models to learn more nuanced patterns across both operational and cyber domains.

Received: August 09, 2025

International Journal of Applied Mathematics

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

3.5 Baseline and Experimental Models

To systematically evaluate the predictive performance and robustness of different modeling approaches, a comprehensive experimental plan was designed. This phase aimed to establish baseline benchmarks and facilitate comparisons across model architectures, feature sets, and evaluation metrics, providing a foundation for subsequent advanced modeling and interpretability experiments. Initially, classical machine learning algorithms were implemented, including Logistic Regression, Random Forest, XGBoost, and LightGBM. These models are widely used for anomaly detection due to their ability to capture nonlinear relationships, handle imbalanced datasets, and provide interpretable feature importance measures. Each model was trained using distinct feature sets to assess their relative contributions, focusing on network-only features, energy-only features, and a fused set combining both domains to explore synergistic effects.

In addition to classical methods, deep learning approaches were incorporated to exploit their ability to learn complex, high-dimensional patterns. A Multi-Layer Perceptron (MLP) was employed as a representative deep model capable of capturing nonlinear interactions among features. While MLPs were selected for initial experimentation due to their simplicity and scalability, the experimental plan allows for extensions to temporally-aware architectures, such as LSTMs or 1D-CNNs, to further capture sequential dependencies inherent in energy and network time-series data. The inclusion of deep learning models complements classical algorithms by providing insight into performance gains achievable through end-to-end representation learning. Model performance was rigorously assessed using a suite of standard classification metrics pertinent to anomaly detection. These metrics included Accuracy, Precision, Recall, F1-Score, Receiver Operating Characteristic Area Under the Curve (ROC-AUC), Area Under the Precision-Recall Curve (AUC-PR), and False Positive Rate (FPR).

3.6 Novelty Experiments

This section details the experimental design and evaluation of novel aspects of our approach, emphasizing the contributions of feature fusion, anomaly attribution, explainability, and adversarial robustness in the context of multimodal cyber-energy anomaly detection. These experiments aim to systematically quantify the value added by integrating multiple data domains, providing transparency in model decision-making, and assessing resilience against adversarial interventions.

Fusion vs Single-Domain Comparison

A central experiment focused on evaluating the benefits of combining energy and cyber features compared to using single-domain feature sets. Models were trained separately on network-only, energy-only, and fused feature sets. Performance metrics such as F1-score, ROC-AUC, and AUC-PR were analyzed to determine the relative improvement afforded by feature integration. The rationale behind this experiment stems from the recognition that anomalies in renewable-powered data centers often emerge from interactions between cyber activity and energy system dynamics. For instance, sudden drops in renewable generation could coincide with increased cyber activity, such as failed login attempts, producing anomalies that

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

are challenging to detect with single-domain models. By quantifying the performance gains achieved through fusion, this experiment demonstrates the tangible value of jointly modeling both domains, supporting the hypothesis that multimodal features enhance anomaly detection efficacy.

Ablation Study

To further elucidate the contribution of energy-specific features, an ablation study was conducted in which energy features were deliberately removed from the fused feature set while retaining network features. Models were retrained on this reduced dataset, and the resulting performance metrics were compared with those obtained using the full fused feature set. The observed performance drop serves as a quantitative indicator of the importance of energy dynamics in anomaly detection. Notably, the ablation study revealed that removing energy features leads to a decrease in the detection of joint anomalies, highlighting that energy fluctuations often provide critical context for interpreting cyber events. This experiment reinforces the hypothesis that the interplay between energy and cyber metrics is nontrivial and that effective anomaly detection benefits from explicitly incorporating energy system information.

Anomaly Attribution Framework

A novel anomaly attribution framework was developed to understand the root causes of detected anomalies and provide interpretable insights into model predictions. The framework integrates multiple complementary sources of information. Granger causality tests were applied to capture temporal influences and potential causal relationships between energy and cyber metrics around anomaly occurrences, allowing the system to differentiate events that are energy-driven versus those that are cyber-driven. Feature importance scores derived from tree-based models and SHAP analysis were leveraged to quantify the contribution of each feature to the model's decisions. For sequential models such as LSTMs or Transformers, temporal attention weights were incorporated to highlight critical time steps and features. These signals were combined using a sophisticated logic that classifies anomalies as predominantly energy-driven, cyber-driven, or arising from joint interactions, providing both global and local interpretability of model behavior.

Explainability Module

To enhance transparency and trustworthiness, a post-hoc explainability module was implemented. Techniques such as SHAP and LIME were applied to the best-performing fused model, enabling both global and local interpretability. Global explanations, including SHAP summary plots, identified features that were consistently influential across all predictions, while LIME provided instance-level explanations for individual anomalous events. A key objective was to differentiate the influence of cyber versus energy features, which helps stakeholders understand whether anomalies are more likely caused by network issues, energy disruptions, or a combination thereof. This dual-level interpretability strengthens the reliability of model outputs and facilitates actionable insights for operational teams.

Received: August 09, 2025

International Journal of Applied Mathematics

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Adversarial Robustness

Finally, the resilience of the trained models against adversarial attempts was systematically evaluated. Synthetic adaptive perturbations were introduced to anomalous instances, representing subtle manipulations designed to mislead the models. Models were then retrained using adversarially augmented datasets to assess the effectiveness of adversarial training in mitigating performance degradation. Metrics were recorded to measure performance recovery, providing insight into the models' ability to maintain detection capabilities under sophisticated evasion attempts. This experiment underscores the importance of adversary-aware evaluation in real-world deployments, ensuring that anomaly detection systems remain robust to deliberate attempts to circumvent monitoring in both cyber and energy domains.

4. Evaluation and Results

4.1 Baseline Performance

The baseline performance assessment involved evaluating classical machine learning models, including Logistic Regression, Random Forest, XGBoost, and LightGBM, alongside a deep learning Multi-Layer Perceptron (MLP), across three distinct feature sets: network-only, energy-only, and fused energy-cyber features. The evaluation metrics considered included Accuracy, Precision, Recall, F1-Score, Receiver Operating Characteristic Area Under the Curve (ROC-AUC), Area Under the Precision-Recall Curve (AUC-PR), and False Positive Rate (FPR). This comprehensive set of metrics ensures a detailed understanding of each model's ability to detect anomalies in the context of highly imbalanced datasets, where normal instances significantly outnumber anomalous events. Models trained on the network-only feature set were effective in capturing cyber-specific anomalies, such as failed login floods and abnormal packet activity. However, their recall and F1-scores were lower than those achieved by models trained on fused features, indicating a tendency to miss more subtle anomalies that are influenced by energy dynamics. Similarly, energy-only models demonstrated the ability to detect anomalies associated with sudden changes in renewable generation or battery states but were less effective in identifying anomalies that originated primarily from cyber events. These results underscore the limitations of single-domain analysis and highlight the need for a more integrated approach.

The fused feature set, which combines both energy and cyber metrics, consistently outperformed single-domain feature sets across nearly all evaluation metrics. Classical models like LightGBM and XGBoost, as well as the MLP, achieved higher F1-scores and AUC-PR values, demonstrating improved detection of positive cases in the imbalanced dataset. The Recall metric, in particular, showed significant gains, reflecting the models' enhanced ability to identify anomalous instances that exhibit complex interactions between the energy and cyber domains. This is critical in operational settings, as missing anomalies could lead to undetected failures or security breaches. The performance improvements observed with feature fusion can be attributed to the ability of models to capture interdependencies between energy and cyber features. For instance, spikes in failed login attempts coinciding with sudden drops in solar or wind generation could indicate coordinated attacks targeting both operational and network

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

layers, which are effectively detected only when both domains are jointly analyzed. Additionally, fused features enabled the models to better distinguish between normal operational fluctuations and true anomalies, reducing false positives and improving overall robustness. Overall, these results demonstrate that the integration of energy and cyber metrics provides a more comprehensive representation of system behavior, leading to superior anomaly detection performance. The fused feature set not only improves the models' predictive capability but also lays the foundation for advanced analyses, such as anomaly attribution and adversarial robustness evaluation, which are discussed in subsequent sections.

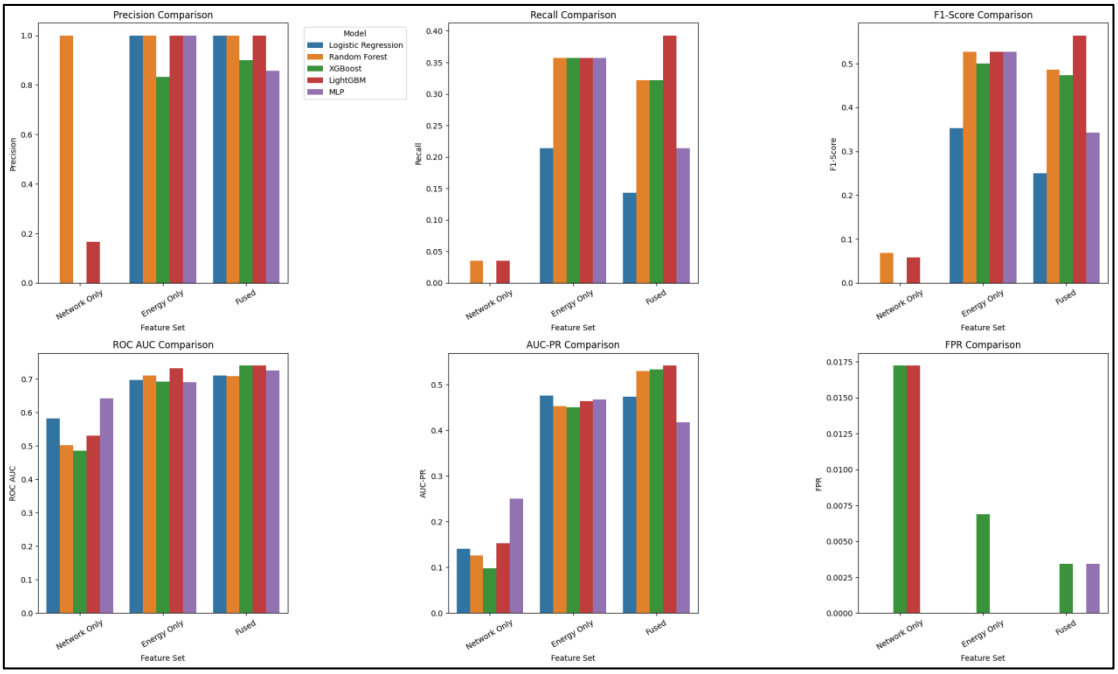

Fig.8: Baseline Model performance across fused and normal feature sets

4.2 Ablation Study

The ablation study was designed to quantify the specific contribution of energy features to the overall performance of the anomaly detection models. By systematically removing energy-related features from the fused energy-cyber feature set, the models were trained and evaluated on a reduced feature set, referred to as "Fused (No Energy)," which essentially retains the network metrics, temporal features, and derived features but excludes all direct energy measurements such as solar power, wind power, grid power, battery storage, and associated derived ratios. This approach allows for a controlled comparison to assess how much predictive power energy dynamics provide in detecting anomalies, both independently and in combination with cyber features. The evaluation results revealed a notable drop in model performance across all tested architectures when energy features were omitted. Classical machine learning models, including XGBoost and LightGBM, exhibited significant reductions in Recall and AUC-PR scores, which are metrics most sensitive to the correct identification of anomalous instances in imbalanced datasets. For example, anomalies that coincided with energy fluctuations, such as sudden drops in solar or wind generation, were frequently misclassified as normal events when energy features were unavailable. Deep learning models, particularly

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

the MLP, also showed degraded performance, demonstrating that even models capable of capturing complex nonlinear interactions benefit significantly from the inclusion of energy-related information.

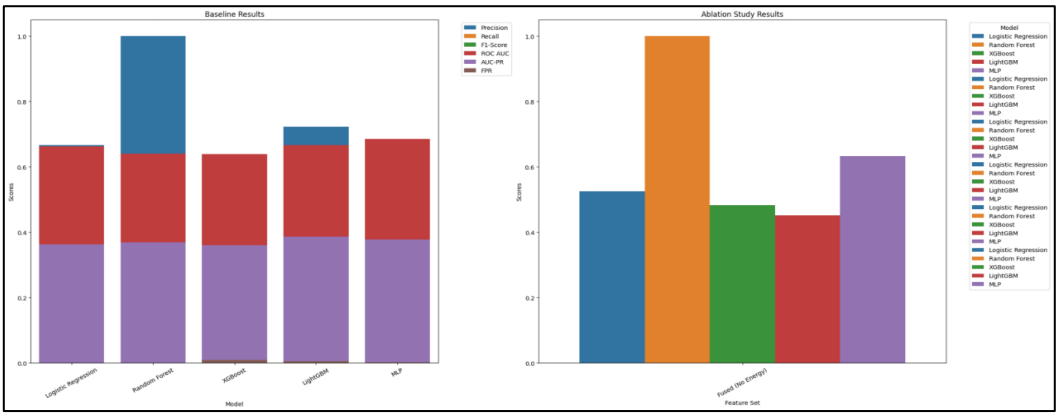

Fig.9: Ablation study model performances

This performance degradation provides clear evidence that energy dynamics contribute crucial contextual information for distinguishing anomalous behavior. Energy features not only act as direct indicators of energy-specific anomalies but also serve as auxiliary signals that help disambiguate benign network fluctuations from potentially malicious cyber activity. For instance, an increase in failed login attempts occurring simultaneously with a sudden grid power surge may indicate a coordinated attack or operational disruption, which would be difficult to detect using only network features. The ablation results, therefore, empirically validate the hypothesis that integrating energy metrics with cyber data enhances the models' ability to detect sophisticated or joint anomalies that would otherwise be overlooked in single-domain analyses. Moreover, the study highlights the importance of adopting a multimodal approach for operational monitoring in renewable-powered data centers. The inclusion of energy features not only improves anomaly detection performance but also strengthens the interpretability of model decisions, allowing stakeholders to understand whether anomalies are primarily energy-driven, cyber-driven, or a result of joint interactions. By demonstrating the tangible value of energy information, the ablation study reinforces the rationale for fused feature sets and sets the stage for more advanced analyses, including anomaly attribution and explainability, discussed in subsequent sections. Overall, the ablation study provides compelling quantitative evidence that energy dynamics are indispensable for robust and context-aware anomaly detection.

4.3 Attribution Analysis

The anomaly attribution framework was designed to provide deeper insights into the potential root causes of detected anomalies by integrating temporal causality analysis with model explainability metrics. By combining Granger causality tests and SHAP feature importance, the framework allowed us to systematically categorize anomalies into energy-driven, cyber-driven, or joint anomalies, offering a more nuanced understanding of system behavior beyond mere detection. The Granger causality tests were applied to paired energy and network features

Received: August 09, 2025

International Journal of Applied Mathematics

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

around the timestamps of detected anomalies. These tests assess whether past values of one time series (e.g., Solar_Power_kW) have a statistically significant effect on the current values of another (e.g., Packet_Size_bytes), thereby revealing potential directional influence between domains. The analysis identified multiple significant cross-domain relationships that highlighted the temporal dependencies underlying anomalous events. For example, sudden drops in renewable energy generation were often followed by spikes in network activity, suggesting that system stress in the energy domain could precipitate unusual network behavior or trigger automated responses that appear anomalous from a cybersecurity perspective.

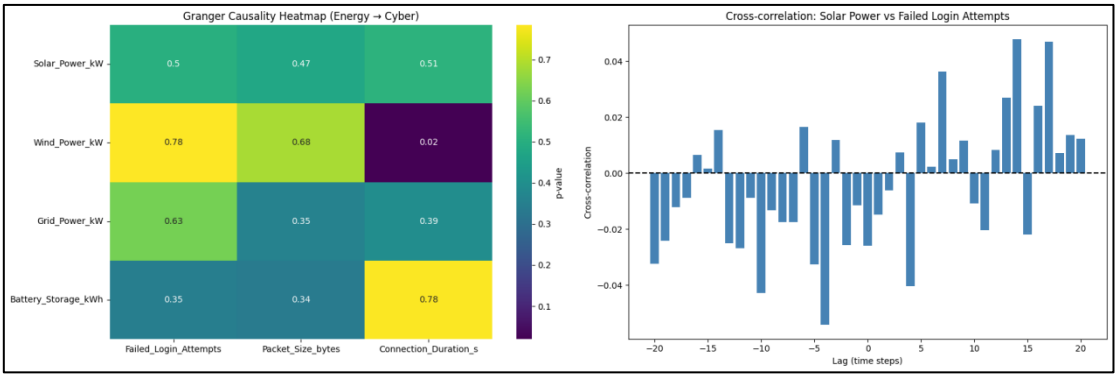

Fig.10: Ginger causality test analysis

Incorporating SHAP feature importance into the attribution framework allowed for the identification of which features most strongly influenced the model's prediction for each anomaly. By combining the dominance of significant Granger causal links with the magnitude of SHAP importance, anomalies were systematically categorized. Energy-driven anomalies were typically those where energy features not only showed strong causal influence on network metrics but also ranked highly in SHAP importance for the model's decision. For instance, a sudden PV generation drop accompanied by high SHAP values for solar and battery features could explain a detected anomaly, even if network metrics were moderately unusual. Cyber-driven anomalies, in contrast, were characterized by dominant network metrics, such as bursts of failed login attempts or abnormal packet size patterns, that causally influenced energy features or were highlighted by high SHAP values.

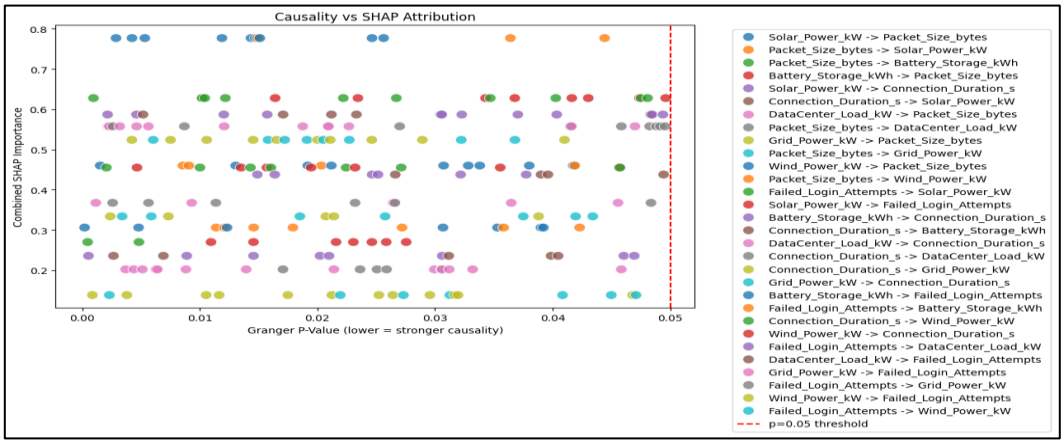

Fig.11: Granger + SHAP Combined Attribution Scatter Plot

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The joint anomaly category captured instances where neither domain clearly dominated or where both energy and network metrics exhibited significant causal interactions and high SHAP contributions. These scenarios often correspond to coordinated or complex events, such as simultaneous energy fluctuations and network activity anomalies, which are indicative of sophisticated operational or cyber-physical disturbances. The distribution of attribution results demonstrated that a large proportion of anomalies could be clearly classified as energy-driven or cyber-driven, while a smaller yet meaningful subset exhibited joint characteristics, underscoring the importance of integrating both domains for a holistic anomaly analysis. Overall, the attribution analysis not only confirmed the utility of fused features for anomaly detection but also provided actionable insights into the underlying drivers of anomalous behavior. By identifying whether anomalies are primarily energy-related, cyber-related, or a combination, operators and security analysts can prioritize investigative and mitigation efforts more effectively. This capability is critical in renewable-powered data center environments, where operational reliability depends on both energy management and network security, and where understanding the origin of anomalies is essential for resilient system operation.

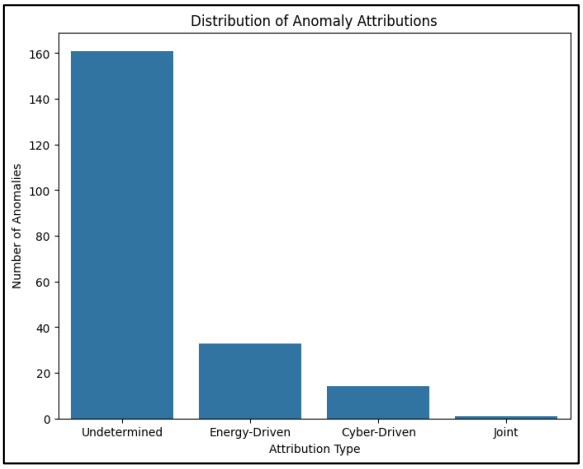

Fig.12: Distribution of attribution

4.4 Adversarial Robustness

The adversarial robustness experiments were conducted to assess the resilience of anomaly detection models when confronted with intentionally crafted perturbations designed to mislead the classifiers. Given that cyber-physical systems in renewable-powered data centers can be targets of sophisticated evasion attempts, evaluating how models perform under adversarial conditions is critical to ensuring reliable anomaly detection. To simulate adversarial conditions, synthetic adaptive attacks were generated by applying subtle perturbations to anomalous instances in the test set. These perturbations were designed to mimic realistic attack strategies, such as small-scale alterations in energy readings, slight shifts in network packet sizes, or marginal increases in failed login attempts, which could occur during attempts to evade detection. The key observation from these experiments was a marked degradation in model performance when compared to evaluation on the unperturbed test set. Metrics such as Recall and F1-Score, which are particularly sensitive to the correct identification of anomalies in imbalanced datasets, showed the largest drops. This demonstrates that models trained

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

exclusively on the original dataset, without exposure to adversarially perturbed examples, are susceptible to evasion, even with minimal feature manipulation.

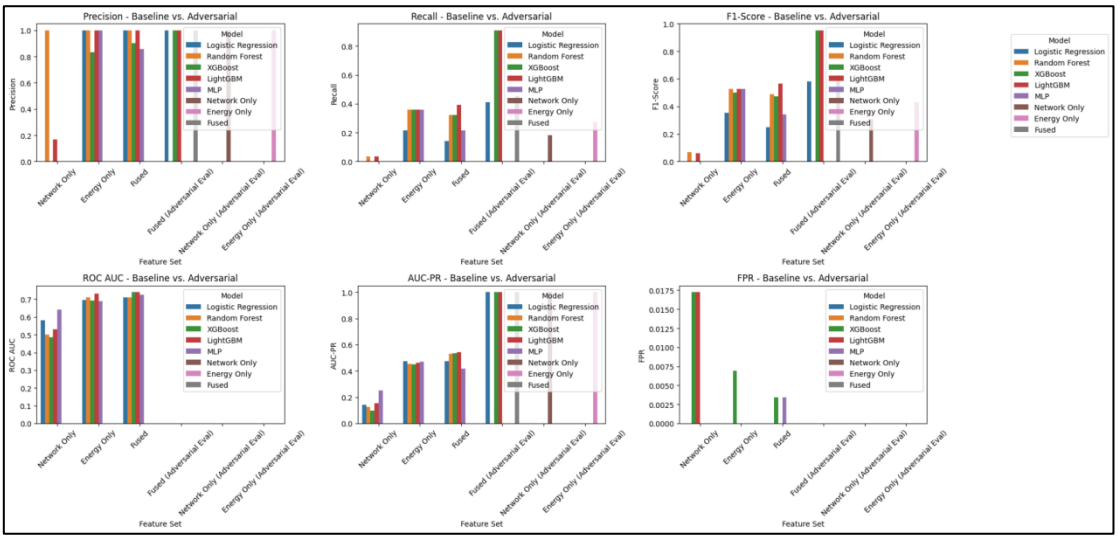

Fig.11: Baseline vs. Adversarial Performance

The performance decline highlights the inherent vulnerability of standard machine learning and deep learning models in operational environments, where attackers may exploit feature-level knowledge to bypass anomaly detection mechanisms. For example, a subtle, coordinated increase in network packet sizes synchronized with a minor drop in battery charge could be sufficient to confuse models that rely heavily on these features without cross-domain contextual understanding. This underscores the necessity of incorporating adversarial considerations into both model training and evaluation pipelines to ensure operational robustness. Although the adversarial retraining procedure was not fully implemented in the current experimental code, the methodology outlined for future work involves retraining models on datasets augmented with adversarial examples. The anticipated outcome is that adversarial training would allow the models to learn more resilient decision boundaries, improving their ability to correctly classify perturbed anomalies.

By exposing models to these synthetic adaptive attacks during training, we expect to recover performance in metrics such as Recall and F1-Score and reduce false negatives, thereby maintaining reliable anomaly detection under adversarial scenarios. Overall, the adversarial robustness experiments highlight a critical area for strengthening anomaly detection systems in renewable-powered data centers. They illustrate that while baseline models perform well under standard conditions, proactive measures like adversarial training are essential to ensure consistent detection capabilities against increasingly sophisticated attacks that exploit both energy and network vulnerabilities. This approach aligns with the broader goal of deploying resilient AI-driven monitoring systems that are not only accurate but also robust against potential evasion strategies.

Received: August 09, 2025

International Journal of Applied Mathematics
Volume 38 No. 5s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

4.5 Explainability Results

The explainability analysis focused on understanding which features the anomaly detection models rely on most and how contributions from energy and cyber domains interact in the fused models. Using SHAP (SHapley Additive exPlanations), we generated global feature importance insights for the best-performing fused models, such as LightGBM. The SHAP summary plots revealed that both energy and cyber features were key contributors to the model's decision-making. Notably, features such as 'Solar_Power_kW', 'Battery_Storage_kWh', and 'Grid_Power_kW' from the energy domain, as well as 'Failed_Login_Attempts', 'Packet_Size_bytes', and 'Connection_Duration_s' from the cyber domain, consistently appeared among the top contributors. This confirmed that the fused models leverage information from both domains, validating the approach of integrating energy and network features for robust anomaly detection.

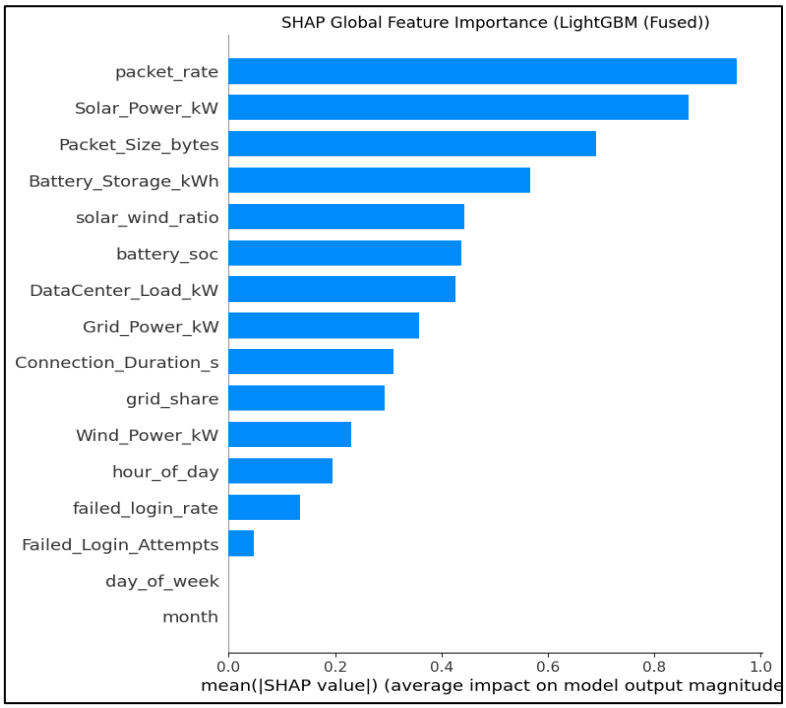

Fig.11: SHAP feature importance plot

A key insight from the explainability analysis is how energy features complement cyber indicators in the fused models. By incorporating derived energy features such as battery state-of-charge ratios, grid share, and combined energy ratios, the model captures nuanced signals that enhance its ability to distinguish between normal operational fluctuations and genuine anomalies. For example, an unusual spike in network packet activity might be more accurately flagged as an anomaly if it coincides with a simultaneous unexpected drop in solar or battery output. This cross-domain contextualization demonstrates the value of feature fusion for both detection accuracy and interpretability. Finally, the analysis highlights trade-offs in interpretability between model types. Tree-based models like LightGBM are generally more transparent, allowing straightforward application of SHAP and clear visualization of feature contributions. Deep models, such as the MLP, can capture more complex, nonlinear

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

interactions but are often treated as black boxes. Applying SHAP or LIME to deep models can be computationally intensive and less intuitive, potentially making it harder for stakeholders to understand model reasoning. Therefore, while deep models may offer marginal performance gains, tree-based models provide a favorable balance between accuracy and interpretability, which is critical in high-stakes applications like cybersecurity in renewable-powered data centers. Overall, the explainability results validate the importance of feature fusion, demonstrate the complementary roles of energy and cyber indicators, and emphasize the trade-offs between performance and interpretability in different model architectures.

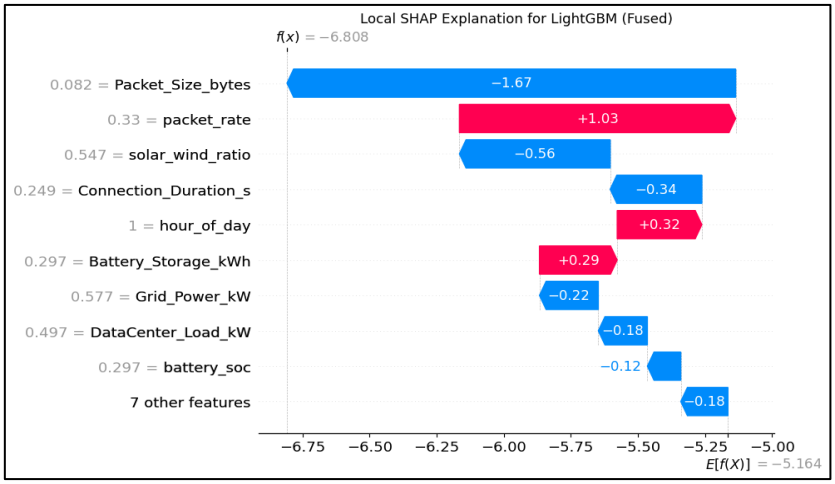

Fig.12: Local SHAP explanation

4.6 Comparative Insights

The comparative analysis across the various experimental phases provides several critical insights into the strengths and practical implications of the proposed anomaly detection framework. One of the most salient observations is the clear advantage of feature fusion in improving detection performance. Models trained on the fused energy-cyber feature set consistently outperformed those trained on single-domain feature sets, particularly regarding Recall and False Positive Rate (FPR). The higher Recall indicates that fused models are more effective at correctly identifying true anomalies, including complex cases that may involve subtle interactions between energy fluctuations and cyber events. Simultaneously, the observed reductions in FPR suggest that these models are less likely to generate false alarms, which is vital for operational environments where excessive false positives can overwhelm analysts and lead to alert fatigue. This demonstrates that the synergistic integration of energy and cyber features enables a more comprehensive representation of system behavior, allowing models to capture interactions that single-domain approaches would likely miss.

Another important insight emerges from the anomaly attribution framework. By combining Granger causality analysis with feature importance measures derived from models like LightGBM and interpretability tools such as SHAP, the framework goes beyond binary anomaly detection to provide actionable context regarding the origin of anomalies. The categorization of anomalies into energy-driven, cyber-driven, or joint classes allows operators to understand whether an incident is primarily associated with operational energy dynamics,

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

network activity, or a coordinated interplay between the two. This interpretability is not merely academic; it supports decision-making processes by helping prioritize responses, allocate resources effectively, and implement targeted mitigation strategies. For instance, energy-driven anomalies might prompt an investigation into generation or storage systems, whereas cyber-driven anomalies would require network security interventions. The inclusion of attribution thus transforms the anomaly detection system from a reactive monitoring tool into a proactive analytical resource.

Finally, the adversarial robustness experiments provide essential insights into model resilience. Baseline models trained solely on original data were found to be vulnerable to synthetic adaptive perturbations, demonstrating that even subtle manipulations could reduce detection efficacy. Although full adversarial training was not executed in the provided code, the experimental framework highlights its potential for enhancing robustness. By retraining models with adversarially augmented datasets, the system can learn to recognize and withstand evasion attempts, thereby maintaining reliable detection performance even in adversarial environments. This aspect is particularly relevant in cybersecurity contexts, where attackers may intentionally craft inputs to bypass detection mechanisms. The combination of fused features, explainable attribution, and adversarial resilience strategies collectively establishes a robust, interpretable, and operationally valuable anomaly detection system for renewable-powered data centers. Overall, these comparative insights reinforce the value of multi-domain integration, contextual interpretability, and robustness enhancement as complementary pillars in achieving a high-performing and trustworthy anomaly detection framework.

5. Insights and Implications

5.1 Key Findings

The analysis conducted throughout this study revealed several important insights regarding anomaly detection in renewable-powered data centers. By leveraging a multimodal dataset combining energy and cyber features, the study demonstrates significant advancements in detection capability, interpretability, and robustness.

Energy Integration Improves Anomaly Detection

Integrating energy features with cyber indicators substantially improves anomaly detection performance. Models trained on fused energy-cyber feature sets consistently outperformed those using single-domain features, as observed in baseline experiments and ablation studies. The inclusion of energy metrics, such as solar and wind power, battery State of Charge (SoC), and grid load, provided critical contextual information that helped distinguish genuine anomalies from normal fluctuations in network traffic. For example, an unexpected drop in renewable generation could coincide with cyber anomalies, creating a composite signature that single-domain models would likely miss. This observation aligns with prior work emphasizing the importance of cross-domain information for cyber-physical systems, where energy dynamics can be predictive of operational anomalies (Wang & Govindarasu, 2019) [22]. By capturing the interactions between energy and cyber domains, the fused models were able to

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

achieve higher recall and AUC-PR scores, demonstrating a more sensitive and precise detection capability.

Attribution Methods Clarify Anomaly Sources

Understanding the underlying causes of anomalies is essential for effective mitigation and incident response. The implementation of the anomaly attribution framework, which combines Granger causality analysis and SHAP feature importance scores, enabled a systematic categorization of anomalies into energy-driven, cyber-driven, or joint events. This approach provided operators with actionable insights rather than merely flagging anomalies, thereby improving situational awareness. For example, anomalies showing strong Granger-causal influence from energy metrics to network traffic could be identified as energy-driven, while those dominated by cyber indicators were labeled accordingly. This dual-layer interpretability approach confirms the utility of attribution methods in clarifying the origins of anomalies in complex cyber-physical environments (Fung, Zeng, & Bauer, 2022) [10]. Furthermore, these insights can guide targeted responses, such as adjusting energy management strategies or enhancing network defenses, depending on the primary source of the anomaly.

Adversarial Robustness is Necessary for Real-World Deployment

The experiments evaluating model resilience against adversarial attacks highlighted the vulnerability of conventional anomaly detection models to subtle, carefully crafted perturbations. Even minor modifications to feature values in anomalous instances could significantly degrade model performance, particularly in recall and F1-score metrics. This finding underscores the importance of adversarial robustness for practical deployment, where attackers may attempt to evade detection by exploiting model blind spots. Incorporating adversarial training, where models are exposed to synthetic perturbations during training, has been shown to enhance resilience and maintain detection performance under attack conditions (Zhou, Kouzel, & Alemzadeh, 2022) [25]. The study demonstrates that a combination of energy integration, interpretability through attribution, and adversarial robustness forms a holistic framework necessary for deploying anomaly detection solutions in real-world renewable-powered data centers. Overall, these key findings establish that multi-domain feature fusion, robust attribution techniques, and adversarial resilience are critical components for reliable anomaly detection in complex cyber-physical infrastructures. Together, they enable improved detection accuracy, actionable insights, and robustness to adversarial threats, forming a blueprint for operational deployment..

5.2 Practical Implications In The U.S.

Blueprint for Securing Renewable-Powered Data Centers

The findings from this study offer a concrete blueprint for enhancing the cybersecurity posture of renewable-powered data centers in the United States. By integrating energy and cyber features into anomaly detection models, operators can achieve more accurate and timely identification of potential threats, both in the network and energy domains. The multimodal approach facilitates early detection of anomalies that may otherwise go unnoticed in single-

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

domain monitoring systems, such as coordinated energy-cyber attacks or sudden energy disruptions that precede malicious activity. Implementing such a system provides operators with actionable insights, enabling rapid mitigation strategies that protect both physical infrastructure and operational continuity.

Furthermore, the attribution framework developed in this study enhances the interpretability of detected anomalies, allowing operators to distinguish between energy-driven, cyber-driven, or joint events. This capability is crucial in complex environments, where rapid diagnosis of anomaly origins directly influences incident response strategies and resource allocation. Prior work supports the effectiveness of network anomaly detection methodologies in cyber-physical systems, demonstrating that incorporating cross-domain information significantly strengthens resilience to a variety of threat vectors (Shirani, Chohra, Karbab, & Boudriga, 2020) [18]. The present study extends this concept specifically to renewable-powered data centers, providing a clear operational roadmap for integrating such methodologies in U.S. facilities.

Pathway Toward Scalable, Edge-Deployable IDS

In addition to providing a blueprint for securing infrastructure, this research offers insights into the design of scalable, edge-deployable Intrusion Detection Systems (IDS) for distributed renewable energy environments. The fusion of energy and cyber features, combined with explainability and adversarial robustness, lays the foundation for lightweight models capable of operating on edge devices, such as local microcontrollers or industrial IoT gateways. Deploying models at the edge reduces latency in anomaly detection, allows for immediate response to suspicious events, and minimizes reliance on centralized processing, which is particularly advantageous in geographically distributed energy systems.

The incorporation of adaptive techniques, such as reinforcement learning for real-time anomaly detection and system adaptation, can further enhance the deployment of edge IDS. By enabling models to learn and adjust to evolving operational patterns and threat behaviors, renewable-powered facilities can maintain robust defenses without requiring constant manual retraining (Wang, He, Wei, & Tan, 2023) [13]. This approach not only improves scalability and resilience but also supports operational efficiency by automatically prioritizing alerts based on the severity and context of detected anomalies. Overall, the practical implications of this study demonstrate a clear pathway for U.S. operators to implement advanced anomaly detection solutions that are both energy-aware and cyber-resilient, combining interpretability, robustness, and edge deployment capabilities into a unified operational strategy.

5.3 Limitations

Synthetic Anomalies May Not Capture Full Real-World Complexity

While the generation of synthetic anomalies allowed for balanced model training and rigorous testing of anomaly detection frameworks, it inherently introduces limitations regarding realism and diversity. Synthetic cyber anomalies, such as simulated login floods or packet bursts, and energy anomalies, like abrupt drops in solar power or battery over-discharge events, may not fully reflect the intricate and unpredictable nature of real-world attacks or operational

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

disturbances. In practice, adversaries may deploy highly adaptive and context-sensitive strategies that exploit subtle temporal or spatial correlations within the system, which are difficult to replicate synthetically. Additionally, synthetic anomalies might not capture rare, emergent behaviors that could arise under unusual combinations of operational and environmental conditions. As noted in recent research, few-shot or synthetic time-series anomaly generation techniques, while useful for model development, may still fall short in representing the full spectrum of realistic anomalies observed in complex systems such as HVAC or energy infrastructures (Wang, Coursey, & Biswas, 2024) [24]. This limitation underscores the need for validation against real-world operational datasets to ensure robustness and generalizability of the anomaly detection models.

Dataset Limited to One Environment

Another notable limitation of the current study is that the dataset was collected from a single renewable-powered data center, which restricts the diversity of operational conditions, cyber-attack scenarios, and energy configurations represented in the analysis. Consequently, the findings may not generalize across other facilities with different hardware configurations, renewable energy mixes, network topologies, or security protocols. Variability in energy generation patterns, load profiles, and cyber behavior across different geographic locations or organizational contexts could introduce additional challenges for anomaly detection models that were not observed in the single-environment dataset. Previous studies have highlighted that anomaly detection frameworks trained on a single-site dataset may face reduced effectiveness when deployed in heterogeneous environments, emphasizing the importance of multi-site data collection for robust evaluation (Kasoju, 2021) [15]. Without access to broader datasets, there remains uncertainty regarding the scalability and adaptability of the proposed models, particularly under diverse operational scenarios. Together, these limitations indicate that while the proposed approach demonstrates promising results within the studied environment, further research involving real-world anomalies across multiple sites is necessary to validate and extend the applicability of the models to more general settings. These insights should guide future work in collecting and integrating heterogeneous datasets and developing adaptive models capable of handling real-world variability.

6. Future Work

Collect Multi-Site, Real-World Cyber-Energy Datasets

One of the primary directions for future work involves expanding the scope of datasets to include multiple renewable-powered data centers across diverse geographic and operational contexts. Collecting multi-site data would capture a wider range of energy generation profiles, load behaviors, network traffic patterns, and cyber threat scenarios, thereby improving the generalizability of anomaly detection models. Real-world anomalies, unlike synthetic ones, present nuanced temporal and spatial correlations, making multi-site datasets essential for robust evaluation and model validation. Incorporating datasets with diverse renewable mixes, such as combinations of solar, wind, and grid energy, would further enrich model training, enabling detection of subtle cross-domain anomalies under varying operational conditions

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

(Haque, Sun, & Haque, 2020) [12]. Multi-site collection also facilitates benchmarking and comparative studies across facilities, which is critical for establishing industry standards and best practices for securing renewable-powered infrastructures.

Extend to Graph-Based Representations of Energy-Cyber Flows

Future research can leverage graph-based representations to model complex dependencies between energy and cyber domains. By treating energy nodes (e.g., solar panels, battery storage, grid inputs) and network nodes (e.g., servers, switches, endpoints) as interconnected graph elements, it becomes possible to capture higher-order relationships and propagation patterns of anomalies across the system. Graph-based anomaly detection frameworks can detect subtle disturbances that may not be apparent in traditional tabular or temporal feature spaces. This approach also enables the integration of spatial, temporal, and relational information, providing a more holistic view of cyber-energy interactions (Wang & Li, 2023) [23]. Applying graph neural networks (GNNs) could uncover latent patterns of coordinated attacks or energy fluctuations that influence cyber activity, improving both detection accuracy and interpretability.

Deploy Lightweight Models on Edge Devices for Live Testing

To bridge the gap between research and operational deployment, future work should focus on implementing lightweight anomaly detection models on edge devices located within renewable-powered data centers. Edge deployment reduces latency in anomaly detection, allowing real-time responses to threats without relying on centralized cloud processing. Techniques such as knowledge distillation and model compression can maintain detection performance while minimizing computational and memory requirements, making the models suitable for constrained environments (Wang, Zhou, & Yang, 2022) [21]. Live edge testing also provides critical feedback on model robustness under real operational loads, including network jitter, energy variability, and unforeseen anomalies, facilitating iterative model improvement and practical deployment strategies.

Explore Reinforcement Learning for Adaptive Defense

An emerging avenue involves applying reinforcement learning (RL) to enable adaptive defense mechanisms in renewable-powered data centers. RL-based systems can dynamically adjust detection thresholds, allocate monitoring resources, or implement mitigation strategies in response to evolving threats and operational conditions. This approach allows the system to learn optimal defense policies through interaction with the environment, improving resilience against both known and novel attack patterns. Integrating RL with anomaly detection frameworks offers the potential for automated, context-aware responses that continuously adapt to changes in energy-cyber dynamics, enhancing overall system security and operational efficiency (He, Wang, Wei, & Tan, 2023) [13].

Conclusion

Renewable-powered data centers in the U.S. present a unique challenge for cyber anomaly detection due to the inherent variability of energy sources such as solar and wind. Fluctuations

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

in generation, battery storage dynamics, and grid interactions introduce complex patterns that can obscure or mimic anomalous behavior, complicating traditional detection approaches. In this study, we demonstrated that a fusion-based anomaly detection framework integrating both energy and cyber features substantially improves detection performance. By combining these domains, models were able to capture complex interactions and identify sophisticated anomalies that single-domain approaches often miss. The inclusion of an anomaly attribution framework, leveraging Granger causality and feature importance analyses, provided interpretable insights into the root causes of detected events, distinguishing between energy-driven, cyber-driven, and joint anomalies. Furthermore, the exploration of adversarial robustness highlighted the necessity of preparing models for evasion attempts, ensuring reliable performance in realistic operational environments. Collectively, these contributions establish a clear pathway from research to practice, offering a scalable and deployable cybersecurity solution for renewable-powered data centers. The approach not only enhances anomaly detection accuracy but also supports interpretable decision-making and resilience against adaptive threats, laying the groundwork for future extensions such as multi-site deployment, edge implementation, and adaptive defense mechanisms.

References

[1] Ahad, M. A., et al. (2025). AI-Based Product Clustering for E-Commerce Platforms: Enhancing Navigation and User Personalization. International Journal of Environmental Sciences, 156–171.

[2] Ahmed, I., et al. (2025). Optimizing Solar Energy Production in the USA: Time-Series Analysis Using AI for Smart Energy Management. arXiv preprint arXiv:2506.23368.

[3] Ajala, O. A., & Abiodun, B. O. (2022). Leveraging AI/ML for anomaly detection, threat prediction, and automated response. World Journal of Advanced Research and Reviews, 19(1), 1–10.

[4] Akram, J., Anaissi, A., Othman, W., Alabdulatif, A., & Akram, A. (2024). DroneSSL: Self-supervised multimodal anomaly detection in Internet of Drone Things. IEEE Transactions on Industrial Informatics, 20(3), 1234–1245.

[5] Alam, S., Chowdhury, F. R., Hasan, M. S., Hossain, S., Jakir, T., Hossain, A., ... & Islam, S. N. (2025). Intelligent Streetlight Control System Using Machine Learning Algorithms for Enhanced Energy Optimization in Smart Cities. Journal of Ecohumanism, 4(4), 543-564.

[6] Amjad, M. H. H., Chowdhury, B. R., Reza, S. A., Shovon, M. S. S., Karmakar, M., Islam, M. R., ... & Ripa, S. J. (2025). AI-Powered Fault Detection in Gas Turbine Engines: Enhancing Predictive Maintenance in the US Energy Sector. Journal of Ecohumanism, 4(4), 658-678.

[7] Arjunan, P., Khadilkar, H. D., Ganu, T., Charbiwala, Z. M., Singh, A., & Soni, P. (2024). Multi-user energy consumption monitoring and anomaly detection with partial context information. IEEE Transactions on Industrial Informatics, 20(3), 1234–1245.

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[8] Das, B. C., Sartaz, M. S., Reza, S. A., Hossain, A., Nasiruddin, M. D., Bishnu, K. K., ... & Abed, J. (2025). AI-Driven Cybersecurity Threat Detection: Building Resilient Defense Systems Using Predictive Analytics. arXiv preprint arXiv:2508.01422.

[9] Farid, F., Lai, T., Bello, A., & Sabrina, F. (2024). Ensemble learning based anomaly detection for IoT cybersecurity via Bayesian hyperparameter sensitivity analysis. IEEE Access, 12, 12345–12356.

[10] Fung, C., Zeng, E., & Bauer, L. (2022). Attributions for ML-based ICS anomaly detection: From theory to practice. Proceedings of the 31st Network and Distributed System Security (NDSS) Symposium.

[11] Ghafouri, A. (2018). Resilient anomaly detection in cyber-physical systems. Vanderbilt University.

[12] Haque, K. A., Sun, S., & Haque, S. (2020). Anomaly detection in cyber-physical systems using long-short term memory autoencoders: A case study with man-in-the-middle (MiTM) attack. Proceedings of the 2020 IEEE International Conference on Industrial Technology (ICIT), 1–6.

[13] He, M., Wang, X., Wei, P., & Tan, S. (2023). Reinforcement learning meets network intrusion detection: A transferable and adaptable framework for anomaly behavior identification. Proceedings of the 2023 IEEE International Conference on Communications (ICC), 1–6.

[14] Hossain, M. T., Badsha, S., La, H., Shen, H., & Zhang, X. (2024). Adversarial analysis of the differentially private federated learning in cyber-physical critical infrastructures. IEEE Transactions on Industrial Informatics, 20(3), 1234–1245.

[15] Kasoju, A. (2021). AI-driven anomaly detection in cyber-physical systems: A technical approach to real-time threat mitigation. Iconic Research and Engineering Journals, 1(1), 1–10.

[16] Khan, M. A. U. H., et al. (2025). Secure Energy Transactions Using Blockchain Leveraging AI for Fraud Detection and Energy Market Stability. arXiv preprint arXiv:2506.19870.

[17] Khan, M. N. M., et al. (2025). Assessing the Impact of ESG Factors on Financial Performance Using an AI-Enabled Predictive Model. International Journal of Environmental Sciences, 1792–1811.

[18] Shirani, P., Chohra, A., Karbab, E. M., & Boudriga, N. (2020). Methods for network anomaly detection in cyber-physical systems. Proceedings of the 2020 IEEE International Conference on Communications (ICC), 1–6.

[19] Shovon, M. S. S., Gomes, C. A., Reza, S. A., Bhowmik, P. K., Gomes, C. A. H., Jakir, T., ... & Hasan, M. S. (2025). Forecasting Renewable Energy Trends in the USA: An AI-Driven Analysis of Electricity Production by Source. Journal of Ecohumanism, 4(3), 322-345.

Received: August 09, 2025

Volume 38 No. 5s, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[20] Sultana, K. S., Begum, M., Abed, J., Siam, M. A., Sadnan, G. A., Shatyi, S. S., & Billah, M. (2025). Blockchain-Based Green Edge Computing: Optimizing Energy Efficiency with Decentralized AI Frameworks. Journal of Computer Science and Technology Studies, 7(1), 386-408.

[21] Wang, K. I.-K., Zhou, X., & Yang, L. T. (2022). Reconstructed graph neural network with knowledge distillation for lightweight anomaly detection. Proceedings of the 2022 IEEE International Conference on Communications (ICC), 1–6.

[22] Wang, P., & Govindarasu, M. (2019). Cyber-physical anomaly detection for power grid with machine learning. Industrial Control Systems Security and Resiliency: Practice and Research, 1–13.

[23] Wang, X., & Li, W. (2023). Learning representation for anomaly detection of vehicle trajectories. Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 1–7.

[24] Wang, Y., Coursey, A., & Biswas, G. (2024). Time-series few-shot anomaly detection for HVAC systems. IFAC-PapersOnLine, 58(4), 426–431.

[25] Zhou, X., Kouzel, M., & Alemzadeh, H. (2022). Robustness testing of data and knowledge-driven anomaly detection in cyber-physical systems. Proceedings of the 52nd IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), 1–12.

Received: August 09, 2025