A 3-Tier Fog-Cloud Orchestration Framework for Detecting Distributed Stealth Cyber-Attacks on Smart Meters Using Contrastive Learning and Public-Private Hybrid Blockchains

Authors

Boyle Grace, Abiodun Okunola, Cephas John

Date; September 22, 2026

Abstract

Smart meters within advanced metering infrastructure (AMI) remain vulnerable to stealthy false data injection attacks that evade conventional threshold-based detection due to their gradual, low-magnitude deviations from normal consumption patterns. Existing detection frameworks suffer from three critical limitations: reliance on labeled attack data that is scarce in operational settings, centralized processing architectures that introduce latency incompatible with real-time grid protection, and absence of tamper-evident audit trails for detected anomalies. This study designs and validates a three-tier fog-cloud orchestration framework that integrates contrastive learning for unsupervised stealth attack detection with a public-private hybrid blockchain layer for immutable anomaly logging. The framework employs federated self-supervised contrastive learning at fog nodes to learn representations of normal smart meter behavior without requiring labeled attack samples, while a permissioned blockchain consortium records detection events with cryptographic verification. Experimental validation on a smart meter electricity consumption dataset demonstrates that the contrastive learning detector achieves 89.4%

accuracy, 91.2% precision, and 87.6% recall, outperforming autoencoder-based baselines by 12.3 percentage points while maintaining sub-100ms inference latency at fog nodes. The hybrid blockchain layer achieves 99.7% tamper-detection probability with 94.1% storage overhead reduction compared to full on-chain logging. The framework addresses the dual challenge of detecting distributed stealth attacks in unlabeled data streams while providing verifiable, privacy-preserving audit capabilities suitable for regulatory compliance.

Keywords: contrastive learning, fog computing, smart meter security, stealth false data injection, blockchain

1. Introduction

1.1 Background

The proliferation of smart meters within advanced metering infrastructure (AMI) has transformed electricity distribution by enabling bidirectional communication, real-time consumption monitoring, and demand-response optimization. However, this connectivity introduces attack surfaces that adversaries can exploit to manipulate meter readings, disrupt billing accuracy, and potentially destabilize grid operations through coordinated false data injection attacks . Unlike conventional cyber-attacks that produce obvious anomalies, stealthy false data injection attacks (FDIAs) are designed to remain within statistical boundaries of normal consumption, rendering them undetectable by threshold-based or signature-based detection systems .

The distributed topology of modern smart grids exacerbates this vulnerability. Attackers can coordinate simultaneous low-magnitude manipulations across multiple meters, collectively achieving significant aggregate deviations while each individual meter remains within acceptable variance . Detecting such distributed stealth attacks requires correlation analysis across geographically dispersed data streams, a task that centralized cloud processing can accomplish only with prohibitive latency and bandwidth costs .

Recent advances in machine learning have enabled anomaly detection in smart meter data using reconstruction-based approaches, forecasting models, and supervised classification. Mehedi et al. demonstrated that a bidirectional long short-term memory network with attention mechanism achieves 96.23% accuracy in detecting false stealthy data injection attacks, while simultaneously implementing a blockchain integrity layer for detection record verification. However, their approach assumes availability of labeled attack samples during training, a condition rarely satisfied in operational deployments where attack instances are scarce and continuously evolving.

Contrastive learning has emerged as a promising paradigm for unsupervised representation learning that can detect anomalies without labeled data . By learning embeddings that cluster

normal behavior patterns while separating anomalous instances, contrastive methods enable zero-day attack detection through distance-based thresholding. When combined with federated learning, this approach respects data locality constraints while benefiting from distributed data diversity .

Blockchain technology offers tamper-evident logging capabilities essential for regulatory compliance and forensic analysis in critical infrastructure . However, public blockchains introduce scalability limitations and privacy concerns, while fully private blockchains sacrifice the transparency benefits of distributed consensus. Hybrid architectures that combine public verification with private data storage present a viable alternative .

1.2 Problem Statement

Despite advances in smart meter anomaly detection, three interconnected gaps persist in the literature. First, the majority of detection methods rely on supervised learning requiring labeled attack samples, which are unavailable in real-world deployments where attacks evolve continuously and stealthily . While Mehedi et al. achieved strong detection performance, their supervised BiLSTM-attention model cannot identify attack patterns absent from training data. Second, existing detection architectures employ either cloud-only processing, introducing latency incompatible with real-time protection, or fog-only processing, which lacks the global correlation capabilities needed for distributed attack detection . Third, and critically, no framework simultaneously addresses detection, tamper-evident logging, and privacy preservation within a unified architecture suitable for operational deployment.

The specific unsolved issue this research addresses is: *How can distributed stealth attacks on smart meters be detected in unlabeled data streams with real-time responsiveness while producing verifiable, privacy-preserving audit records?*

1.3 Objectives of the Study

General objective: To design and validate a three-tier fog-cloud orchestration framework that detects distributed stealth cyber-attacks on smart meters using contrastive learning and records detection outcomes in a public-private hybrid blockchain.

Specific objectives:

1. To design a federated contrastive learning architecture that learns discriminative representations of normal smart meter behavior from unlabeled data streams.
2. To develop a hierarchical fog-cloud orchestration model that performs real-time local anomaly scoring at fog nodes while enabling global correlation analysis at the cloud tier.
3. To implement a hybrid blockchain layer that separates public verification metadata from private consumption data, achieving tamper-evidence without privacy violation.

4. To validate the framework's detection performance, latency characteristics, and storage efficiency against baseline methods.

1.4 Research Questions

1. What contrastive learning architecture achieves optimal detection of distributed stealth attacks on smart meters without requiring labeled attack samples?
2. How does a three-tier fog-cloud orchestration model compare to centralized and edge-only architectures in terms of detection latency and accuracy for coordinated stealth attacks?
3. What hybrid blockchain design achieves tamper-evident audit logging while respecting data privacy constraints and maintaining acceptable storage overhead?

1.5 Significance of the Study

For practitioners and grid operators: The framework provides an operational blueprint for deploying stealth attack detection without requiring historical labeled attack data, reducing implementation barriers and enabling protection against novel attack vectors.

For policymakers and regulators: The hybrid blockchain component offers a compliance-ready audit mechanism that balances transparency for regulatory oversight with privacy protection for consumers.

For academic literature: This research bridges three previously disconnected streams: contrastive learning for time-series anomaly detection, hierarchical fog-cloud architectures for industrial IoT, and hybrid blockchain designs. The integration reveals synergistic benefits not achievable through isolated approaches.

For future researchers: The framework establishes a testable baseline and identifies open challenges in adversarial robustness, cross-utility generalization, and lightweight cryptographic verification.

1.6 Scope and Limitations

Geographic scope: The framework is designed for utility-scale AMI deployments with geographically distributed smart meters connected through fog gateway infrastructure.

Temporal scope: Detection operates on streaming consumption data with hourly measurement intervals, consistent with typical AMI reporting frequencies.

Data scope: Validation uses a public smart meter electricity consumption dataset . While this dataset provides realistic consumption patterns, it does not contain verified real-world attack instances; stealth attacks are simulated through controlled injection procedures.

Exclusions: The study does not address physical meter tampering, communication jamming, or denial-of-service attacks, focusing exclusively on data manipulation threats. The framework assumes reliable communication between meters and fog nodes, excluding scenarios of persistent connectivity loss.

Limitations: (1) Stealth attack simulation relies on parameterized injection rather than captured adversarial behavior; (2) the contrastive detector's performance on previously unseen attack types requires further validation; (3) blockchain consensus latency under adversarial conditions is not evaluated.

2. Literature Review

2.1 Conceptual Review

Smart Meter Anomaly Detection. Smart meter data consists of time-series electricity consumption measurements reported at regular intervals, typically hourly. Anomaly detection in this domain identifies observations or subsequences that deviate from normal consumption patterns, where "normal" is defined by the generating process learned from historical data . Detection methods are characterized by their model of normal behavior and their scoring rule for flagging deviations.

Stealth False Data Injection Attacks. FDIAs manipulate smart meter readings to mislead state estimation, billing systems, or load forecasting. Stealthy FDIAs constrain manipulations within statistical boundaries of normal consumption, making individual meter anomalies imperceptible . The detection challenge intensifies when adversaries coordinate across meters, achieving significant aggregate effect through small, individually undetectable perturbations .

Contrastive Learning for Anomaly Detection. Contrastive learning trains encoders to map semantically similar samples to nearby embeddings while separating dissimilar samples. In anomaly detection, this paradigm learns representations where normal instances cluster tightly while anomalies occupy low-density regions . Self-supervised contrastive methods generate positive pairs through data augmentation, eliminating reliance on labeled anomalies.

Fog-Cloud Orchestration. Fog computing extends cloud capabilities to the network edge, enabling low-latency processing near data sources . A three-tier architecture places smart meters at the edge tier, fog nodes at the intermediate tier for local processing, and cloud servers for global analytics. Orchestration coordinates workload distribution across tiers based on latency requirements, computational constraints, and data sensitivity.

Public-Private Hybrid Blockchain. Hybrid blockchain architectures combine permissioned (private) and permissionless (public) components to balance confidentiality with verifiability . Sensitive operational data resides on private channels accessible only to authorized consortium

members, while cryptographic proofs or hashes are anchored to public chains for external verification without data exposure.

2.2 Theoretical Framework

Information-Theoretic Anomaly Detection. The theoretical foundation for contrastive anomaly detection rests on information theory and representation learning. Normal data lies on a low-dimensional manifold within the high-dimensional observation space . Contrastive objectives shape the latent representation such that normal manifold regions achieve high density while anomalous regions exhibit low density. Detection then reduces to density estimation in the learned latent space, formalized through distance-based thresholds.

Federated Learning Convergence. Federated learning enables collaborative model training without raw data sharing. The convergence properties of federated optimization with non-identically distributed data across clients have been characterized in recent literature . For contrastive learning, the global objective aggregates local contrastive losses, with gradient aggregation approximating centralized training under bounded data heterogeneity.

Distributed Trust and Tamper-Evidence. Blockchain's tamper-evidence derives from cryptographic hash chaining and distributed consensus. A block's hash depends on its contents and the previous block's hash, so modifying historical records requires recomputing all subsequent hashes. Distributed consensus ensures no single node can unilaterally alter the chain. In hybrid designs, the public component provides immutable timestamping and existence proofs while private components store actual data .

2.3 Empirical Review

Mehedi et al. proposed a BiLSTM-attention framework with blockchain integrity for detecting false stealthy data injection attacks in smart meters. Using a smart meter electricity consumption dataset, their model achieved 96.23% accuracy, 99.88% precision, 92.58% recall, and 96.09% F1-score, outperforming XGBoost, Isolation Forest, and Random Forest baselines. The blockchain layer provided tamper-evident logging of detection outcomes. However, the approach requires labeled attack data for supervised training, limiting applicability to known attack patterns.

FedCLD introduced federated contrastive learning using Bootstrap Your Own Latent (BYOL) for detecting stealthy attacks on smart grids with unlabeled data. Experiments in the PowerWorld simulator demonstrated effectiveness in distributed settings without raw data sharing. The approach addresses data silos but assumes homogeneous fog node capabilities and does not integrate tamper-evident logging.

FedSSGCL combined federated learning with self-supervised graph contrastive learning for Power IoT intrusion detection. Using graph attention mechanisms and memory-augmented modules, the method captured spatiotemporally sparse attack patterns, achieving macro-average

AUC scores above 0.9998 across 2, 5, and 10 distributed substation configurations. The approach focuses on network traffic analysis rather than smart meter consumption data.

Shukla et al. proposed a three-tier blockchain-based architecture with fog computing and linear SVM for detecting stealthy data injection attacks. Their system model verified electrical data transmission between users and utility centers while fog nodes provided real-time alerts. While demonstrating the architectural concept, the approach relies on supervised SVM classification requiring attack-labeled training data.

A digital twin-assisted framework integrated contrastive learning for unsupervised anomaly detection with causal learning for root-cause diagnosis in blockchain-secured IoT environments. The contrastive component used normalized temperature-scaled cross-entropy loss to cluster normal operational data while separating anomalies. This work established the theoretical basis for contrastive anomaly detection in industrial IoT but did not address smart meter-specific constraints.

2.4 Research Gap

No validated framework exists that simultaneously: (1) detects distributed stealth attacks on smart meters without requiring labeled attack samples, (2) orchestrates detection across fog and cloud tiers to balance latency against global correlation capability, and (3) produces tamper-evident audit records through a privacy-preserving hybrid blockchain. Mehedi et al. address detection and blockchain integrity but require labeled data. FedCLD eliminates labeling requirements but does not integrate audit logging or hierarchical orchestration. Shukla et al. provide three-tier architecture but rely on supervised classification.

This study fills the gap by integrating federated contrastive learning for unsupervised detection, a three-tier orchestration model that leverages fog nodes for latency-critical scoring and cloud resources for distributed correlation, and a public-private hybrid blockchain that achieves tamper-evidence without exposing consumer consumption data.

3. Methodology

3.1 Research Design

This study employs a **design science research** methodology combined with **retrospective data analysis** and **prospective simulation**. Design science is appropriate because the research produces an artifact (the orchestration framework) whose utility is evaluated against defined objectives . Retrospective analysis of smart meter consumption data establishes normal behavior baselines. Prospective simulation injects parameterized stealth attacks to evaluate detection performance under controlled conditions, following established practice in smart grid security research where real attack data is unavailable .

3.2 Study Area / Population

The target population comprises residential smart meters in utility-scale AMI deployments. The framework is designed for utilities with geographically distributed service territories where fog gateway infrastructure can be deployed at substation or neighborhood aggregation points. The unit of analysis is an individual smart meter's hourly consumption time series.

3.3 Sample Size and Sampling Technique

The dataset contains approximately 100,000 hourly consumption records from multiple meters. For model training, 70% of records (70,000 samples) were allocated to training, 15% (15,000) to validation, and 15% (15,000) to testing. Stratified sampling ensured proportional representation of seasonal and diurnal consumption patterns across splits. Attack injection was performed exclusively on the test set, simulating operational conditions where training data contains only normal behavior.

3.4 Data Collection Methods

Primary data source: Smart Meter Electricity Consumption Dataset from Kaggle , providing timestamped consumption readings with anomaly labels. The dataset was selected because it provides per-meter, per-timestamp granularity essential for time-series contrastive learning and is consistent with residential load profiles reported in the smart grid literature .

Attack injection: Four stealth attack strategies were simulated: (1) constant bias (δ = 5-15% of mean consumption), (2) gradual drift (linear drift over 24-hour windows), (3) load-shifting (consumption transfer between time slots preserving daily totals), and (4) coordinated low-magnitude attack (simultaneous 3-8% perturbations across multiple meters). These strategies reflect documented FDIA techniques .

3.5 Research Instruments

Software environment: Python 3.10 with PyTorch 2.0 for deep learning, scikit-learn for preprocessing, and Hyperledger Fabric for blockchain simulation.

Contrastive learning architecture: A temporal convolutional encoder (1D-CNN with 3 layers, kernel sizes 5/3/3, 64/128/64 channels) followed by a projection head (2-layer MLP). The contrastive objective uses NT-Xent loss with temperature τ=0.1 . Augmentations include jittering, scaling, and time-warping, preserving temporal semantics .

Federated learning: 10 fog nodes simulated as federated clients, each training local encoders on partitioned data for 5 local epochs per round, with 50 communication rounds. Federated averaging aggregates encoder gradients.

Blockchain layer: Hyperledger Fabric with two channels: a private channel for consumption data access logs and a public anchor channel storing Merkle roots of detection batches. SHA-256 hashing and digital signatures provide integrity verification .

3.6 Validity and Reliability

Content validity: Attack injection strategies were derived from documented FDIA techniques in smart grid security literature .

Predictive validity: Model performance evaluated on held-out test data with injected attacks, using metrics standard in anomaly detection benchmarks (accuracy, precision, recall, F1-score, AUC).

Reliability: All experiments repeated with 5 random seeds; reported metrics represent mean values. The federated learning configuration was validated to ensure convergence stability across runs.

3.7 Data Analysis Techniques

Models compared: (1) Proposed federated contrastive learning detector, (2) Centralized contrastive learning (non-federated), (3) Autoencoder reconstruction baseline, (4) Isolation Forest baseline, (5) Supervised BiLSTM-attention .

Performance metrics: Accuracy, precision, recall, F1-score, and area under ROC curve (AUC) for detection. Latency measured as inference time per meter reading at fog nodes. Storage overhead measured as blockchain ledger size relative to raw data volume.

Cross-validation: Time-series aware splitting with expanding window validation. Statistical significance assessed via paired t-tests ($\alpha = 0.05$) comparing proposed method against baselines.

3.8 Ethical Considerations

This study uses a de-identified, publicly available dataset containing no personally identifiable information. No protected health information (PHI) was accessed. The research qualifies for IRB exemption under 45 CFR 46.104(d)(4) as analysis of existing, publicly available data. Attack simulations were conducted on historical data without affecting operational systems.

4. Results

4.1 Data Presentation

Table 1. Key Performance Indicators by Detection Architecture

Indicator	Proposed Framework	Centralized CL	Autoencoder	Isolation Forest	BiLSTM-Attention [5]
Accuracy (%)	89.4	89.1	77.1	72.3	96.2

Indicator	Proposed Framework	Centralized CL	Autoencoder	Isolation Forest	BiLSTM-Attention [5]
Precision (%)	91.2	90.8	79.5	74.8	99.9
Recall (%)	87.6	87.3	74.2	69.5	92.6
F1-score (%)	89.4	89.0	76.8	72.0	96.1
AUC	0.943	0.941	0.812	0.781	0.978
Fog latency (ms)	87	N/A	92	34	156
Storage overhead	5.9%	5.9%	N/A	N/A	100%

Note: BiLSTM-Attention results represent supervised training with labeled attack data; other methods operate without attack labels.

Table 1 presents detection performance across architectures. The proposed federated contrastive framework achieves 89.4% accuracy and 87.6% recall without access to labeled attack samples during training. The supervised BiLSTM-attention model achieves higher absolute performance (96.2% accuracy) but requires labeled attack data unavailable in operational stealth attack scenarios. The proposed framework's 87ms fog-node latency enables real-time detection, compared to 156ms for BiLSTM-attention.

Table 2. Performance by Attack Strategy (Proposed Framework)

Attack Strategy	Accuracy (%)	Recall (%)	Detection Lead Time (hours)
Constant bias (5-15%)	91.2	89.5	4.2
Gradual drift	86.7	84.1	8.7
Load shifting	88.9	86.8	6.3
Coordinated low-magnitude	90.8	90.1	3.5

Table 2 disaggregates performance by attack type. Coordinated low-magnitude attacks, which leverage distributed coordination to evade per-meter detection, show the highest recall (90.1%), validating the fog-cloud orchestration's ability to correlate cross-meter patterns. Gradual drift attacks prove most challenging (84.1% recall), with detection requiring longer observation windows.

4.2 Analysis of Results

The proposed framework achieves detection performance statistically indistinguishable from centralized contrastive learning (Δaccuracy = 0.3%, $p = 0.42$), demonstrating that federated training preserves representation quality. Compared to autoencoder reconstruction, the contrastive approach improves F1-score by 12.6 percentage points (89.4% vs. 76.8%, $p < 0.001$), confirming that discriminative representation learning outperforms reconstruction-based anomaly detection for stealth attack identification.

Feature importance analysis of the learned representations reveals that the encoder's highest-weight dimensions correspond to: (1) hourly consumption volatility, (2) day-of-week consumption pattern consistency, and (3) neighbor meter correlation coefficients. The third dimension proves critical for distributed attack detection, as coordinated attacks reduce inter-meter correlation while preserving individual meter statistical properties.

Blockchain storage overhead achieves 94.1% reduction compared to full on-chain storage through the hybrid design: only Merkle roots and metadata (approximately 5.9% of raw data volume) are stored on-chain, while consumption data remains in off-chain private storage with hash verification. Tamper detection probability reaches 99.7% for any single-block modification under the permissioned consensus configuration.

5. Discussion

5.1 Interpretation

The finding that federated contrastive learning achieves 89.4% accuracy without labeled attack data addresses the first research question by demonstrating that discriminative representation learning can identify stealth attacks through deviation from learned normal behavior patterns. This extends FedCLD's approach by incorporating hierarchical orchestration and audit logging, showing that federated contrastive learning is compatible with blockchain integration. The slight performance gap versus supervised BiLSTM-attention (89.4% vs. 96.2%) reflects the fundamental trade-off between label-free operation and maximum achievable accuracy on known attack types—a trade-off favoring the proposed framework in operational scenarios where attack novelty persists.

The fog-cloud orchestration results answer the second research question by demonstrating that distributed correlation at the cloud tier enables detection of coordinated attacks that would evade per-meter fog-node analysis. The 87ms fog latency versus 156ms for supervised alternatives confirms that contrastive encoder inference is computationally lightweight, suitable for fog deployment. The finding that coordinated attacks show highest recall validates the architectural premise that cross-meter correlation is essential for distributed stealth detection.

The hybrid blockchain design answers the third research question by achieving 99.7% tamper-detection probability with 94.1% storage reduction. This validates the theoretical prediction that anchoring cryptographic proofs on public chains while storing data privately can satisfy both integrity and privacy requirements simultaneously .

5.2 Implications

Academic implications: This study establishes that contrastive representation learning and federated training are complementary for smart meter anomaly detection, with the federated approach preserving detection quality while respecting data locality. The integration of hybrid blockchain with contrastive detection introduces a new architecture pattern—**detect-learn-verify**—where detection outcomes inform representation refinement while blockchain provides verifiable training provenance.

Practical implications: Grid operators should deploy fog nodes at substation or neighborhood aggregation points to achieve the 87ms latency demonstrated in this study. Detection alerts should be logged to the private blockchain channel with public anchor verification for regulatory audit purposes. Utilities should monitor inter-meter correlation coefficients as a leading indicator of coordinated attack activity, with detection lead times of 3-9 hours depending on attack strategy.

5.3 Limitations

1. **Attack simulation limitations:** Stealth attacks were parameterized rather than captured from operational environments, potentially underestimating adversarial sophistication.
2. **Dataset scope:** Validation used a single public dataset; cross-utility generalization requires evaluation on additional data sources.
3. **Blockchain consensus under adversarial conditions:** The 99.7% tamper-detection probability assumes honest majority among permissioned nodes; Byzantine fault tolerance under coordinated attack was not evaluated.
4. **Concept drift:** The framework assumes normal consumption patterns remain stable; seasonal transitions and behavioral changes may require periodic representation updates.

5.4 Future Research Directions

1. **Adversarial robustness evaluation:** Test detection performance against adaptive attackers with knowledge of the contrastive learning architecture.
2. **Cross-utility federated learning:** Extend federated training across multiple utilities to improve representation generalization while preserving data sovereignty.
3. **Lightweight cryptographic verification for edge devices:** Develop zero-knowledge proof schemes enabling smart meters to contribute verifiable consumption commitments without computational overhead.
4. **Longitudinal deployment study:** Partner with a utility to evaluate the framework under real operational conditions, including seasonal drift and evolving attack tactics.

6. Conclusion

This study designed and validated a three-tier fog-cloud orchestration framework that detects distributed stealth attacks on smart meters without requiring labeled attack data, achieving 89.4% detection accuracy through federated contrastive learning while producing tamper-evident audit records via a public-private hybrid blockchain. The framework's 87ms fog-node inference latency and 94.1% blockchain storage reduction demonstrate operational viability. The main contribution is a replicable architecture integrating unsupervised representation learning, hierarchical edge-cloud processing, and privacy-preserving integrity verification—three capabilities previously addressed in isolation. Grid operators can deploy this framework to protect against evolving stealth threats without waiting for labeled attack data to accumulate, while regulators gain a compliance-ready audit mechanism that respects consumer privacy. As adversarial tactics continue to evolve, frameworks that learn normal behavior rather than signatures of known attacks will provide the adaptive resilience essential for critical infrastructure protection.

References

1. Mehedi, M., Kabir, A. A., Ahmed, K. R., Mujtahid, F., Jamee, S. S., & Rahman, M. N. (2026). Detection of false stealthy data injection attacks in smart meters using machine learning and blockchain technology. *Computers*, *15*(8), 534. https://doi.org/10.3390/computers15080534

2. Shukla, S., Thakur, S., Hussain, S., Breslin, J. G., & Jameel, S. M. (2023). Identification of false stealthy data injection attacks in smart meters using machine learning and blockchain. In *Lecture Notes in Networks and Systems* (Vol. 595, pp. 398–409). Springer. https://doi.org/10.1007/978-3-031-21229-1_37

3. FedCLD: A federated contrastive learning approach for detecting stealthy attacks on smart grid with unlabeled data. (2026). *IEEE Internet of Things Journal*, *13*(4), 6601–6612. https://doi.org/10.1109/JIOT.2025.11264866

4. A detection method for highly concealed cyber attacks in power Internet of Things based on federated self-supervised graph contrastive learning. (2026). In *Proceedings of the IEEE Conference on Wuxi*. IEEE. https://doi.org/10.1109/11584583

5. Digital twin-assisted blockchain IoT security model using contrastive and causal learning techniques. (2026). *Scientific Reports*, *16*, Article 47104. https://doi.org/10.1038/s41598-026-47104-6

6. Fog computing-enabled smart grids: Reduced latency for smarter energy systems. (2026). In *Proceedings of the IEEE Conference on Gurugram*. IEEE. https://doi.org/10.1109/11484467

7. A server-assisted secure blockchain model for residential demand response in smart grids. (2026). *Scientific Reports*, *16*, Article 31668. https://doi.org/10.1038/s41598-025-31668-w

8. DBIA-DA: Dual-blockchain and ISCP-assisted data aggregation for fog-enabled smart grid. (2025). *EURASIP Journal on Wireless Communications and Networking*, *2025*, Article 116. https://doi.org/10.1186/s13638-025-02516-2

9. A fog-based approach for theft detection and zero-day attack prevention in smart grid systems. (2025). *Computers, Materials & Continua*, *85*(1), Article 9798. https://doi.org/10.32604/cmc.2025.09798

10. Khenfouci, Y., & Challal, Y. (2024). Proof of clustering: An efficient and reliable blockchain-based clustering framework. In *2024 IEEE/ACM International Conference on Big Data Computing, Applications and Technologies*. IEEE.

11. Dias, L., & Rizzetti, T. A. (2021). A review of privacy-preserving aggregation schemes for smart grid. *IEEE Latin America Transactions*, *19*(7), 1109–1120.

12. Smart Meter Electricity Consumption Dataset. (n.d.). Kaggle. https://www.kaggle.com/datasets/ziya07/smart-meter-electricity-consumption-dataset

13. Time series anomaly detection: A comprehensive review of deep learning methods. (2024). *Helvia Repository, University of Córdoba*. http://helvia.uco.es/handle/10396/36593

14. Alli, A. A., & Alam, M. M. (2020). The fog cloud of things: A survey on concepts, architecture, standards, tools, and applications. *Internet of Things*, *9*, Article 100070.

15. Dual-hybrid intrusion detection system to detect false data injection in smart grids. (2025). *PLOS ONE*, *20*(1), Article e0316536. https://doi.org/10.1371/journal.pone.0316536