A Federated Physics-Informed Neural Network (f-PINN) and Zero-Knowledge Blockchain Architecture for Real-Time Stealthy FDIA Detection in Edge-Computing Smart Meters

Authors

Adaan Ashun, Abbas Ashun, Ade John, Monica Bright

Date; September 22, 2026

Abstract

False Data Injection Attacks (FDIAs) targeting edge-computing smart meters represent a critical and evolving threat to smart grid stability, with stealthy variants specifically engineered to evade conventional bad data detection mechanisms while maintaining physical consistency with power system state estimation. Existing detection approaches predominantly rely on centralized data aggregation, which introduces privacy vulnerabilities, communication bottlenecks, and single points of failure, while physics-agnostic machine learning models remain vulnerable to adversarially crafted stealthy injections that respect power flow constraints. This study proposes a federated physics-informed neural network (f-PINN) integrated with a zero-knowledge blockchain architecture for real-time, privacy-preserving detection of stealthy FDIAs at the edge. The methodology combines physics-informed neural network training constrained by power flow equations with federated learning to enable collaborative model development without raw data exchange, while zero-knowledge proofs on a permissioned blockchain ensure detection integrity and tamper-evident auditability. Experimental validation on IEEE 14-bus and 118-bus systems demonstrates that the f-PINN framework achieves 89.4% detection accuracy for stealthy FDIAs, outperforming conventional physics-agnostic baselines by 12.7 percentage points. The architecture achieves sub-second inference latency suitable for real-time SCADA deployment while maintaining data sovereignty across distributed grid operators. The framework provides a

replicable blueprint for deploying privacy-preserving, physics-aware cybersecurity in critical energy infrastructure.

Keywords: False Data Injection Attack, Physics-Informed Neural Network, Federated Learning, Zero-Knowledge Blockchain, Smart Meter Security, Edge Computing

1. Introduction

1.1 Background

The modernization of electrical power systems into smart grids has introduced unprecedented operational efficiency through advanced metering infrastructure, real-time state estimation, and distributed energy resource integration. Central to this transformation is the smart meter, which functions as both a measurement device and a communication node within the grid's cyber-physical architecture. These devices continuously transmit voltage, current, and power measurements to utility control centers, enabling dynamic state estimation, demand forecasting, and automated control decisions. However, this increased connectivity has expanded the attack surface for malicious actors seeking to compromise grid integrity.

Among the most sophisticated threats facing smart grid infrastructure are False Data Injection Attacks (FDIAs), in which adversaries manipulate measurement data to induce incorrect state estimation, potentially triggering erroneous control actions or masking physical system disturbances. Mehedi et al. (2026) recently demonstrated that machine learning and blockchain technologies can be combined to detect false stealthy data injection attacks in smart meters, achieving significant improvements over conventional detection approaches . Their work establishes the foundational principle that intelligent detection mechanisms can identify anomalous patterns in metering data, while blockchain technology provides tamper-evident logging of detected events.

The challenge becomes substantially more complex when considering *stealthy* FDIAs, which are explicitly designed to circumvent bad data detection mechanisms by ensuring that injected measurements remain consistent with power flow equations and network topology constraints. Unlike random or naive attacks, stealthy FDIAs leverage knowledge of the system's physical model to construct attack vectors that satisfy the residual test used in conventional state estimation. This class of attacks exploits a fundamental limitation in existing detection paradigms: purely statistical or data-driven approaches lack the physics awareness necessary to distinguish between legitimate state variations and maliciously constructed measurement patterns that maintain physical consistency.

1.2 Problem Statement

Despite significant advances in FDIA detection research, three interconnected gaps persist in the literature. First, the majority of existing detection frameworks assume centralized data collection

and processing, requiring smart meter measurements to be transmitted to a central utility server for analysis. This architecture introduces privacy concerns regarding consumer energy usage patterns, creates communication bottlenecks in large-scale deployments, and establishes single points of failure that are attractive targets for adversaries. Second, physics-agnostic machine learning models, while effective against naive FDIAs, demonstrate degraded performance against stealthy attacks that respect physical constraints. Third, existing blockchain-based integrity solutions often lack the computational efficiency required for real-time edge deployment and the privacy-preserving properties necessary to protect sensitive metering data.

Mehedi et al. (2026) addressed several of these challenges by integrating machine learning with blockchain technology for stealthy FDIA detection, demonstrating the viability of intelligent detection at the metering layer . However, their approach employs centralized model training and does not explicitly incorporate power system physics into the detection model, leaving a residual vulnerability to adversarially constructed stealthy attacks. Furthermore, the computational overhead of blockchain operations in resource-constrained edge environments remains insufficiently characterized.

The specific problem addressed by this research is the absence of a validated, privacy-preserving, physics-aware detection framework that operates at the edge computing layer of smart meter infrastructure, achieves real-time inference latency, and provides cryptographically verifiable detection integrity without relying on a trusted central authority. No existing study has demonstrated that federated physics-informed neural networks can be effectively combined with zero-knowledge blockchain mechanisms to detect stealthy FDIAs while satisfying the stringent latency and privacy requirements of edge-deployed smart metering systems.

1.3 Objectives of the Study

General objective: To design, implement, and validate a federated physics-informed neural network architecture integrated with zero-knowledge blockchain mechanisms for real-time detection of stealthy false data injection attacks in edge-computing smart meter environments.

Specific objectives:

1. To develop a physics-informed neural network model that incorporates power flow equation constraints for distinguishing stealthy FDIAs from legitimate state variations in smart meter measurements.
2. To design a federated learning framework that enables collaborative model training across distributed smart meter edge nodes without requiring raw measurement data exchange.
3. To integrate zero-knowledge proof mechanisms on a permissioned blockchain for tamper-evident recording of detection events and model updates while preserving data confidentiality.

4. To evaluate the detection accuracy, inference latency, and communication overhead of the proposed f-PINN and zero-knowledge blockchain architecture against baseline methods on standard IEEE test systems.

1.4 Research Questions

1. To what extent does the incorporation of physics-informed constraints improve detection accuracy for stealthy FDIAs compared to physics-agnostic neural network models in edge-computing smart meter environments?
2. How does the federated learning architecture affect detection performance, communication efficiency, and privacy preservation relative to centralized training approaches?
3. What is the computational overhead and latency impact of zero-knowledge blockchain integration on real-time FDIA detection at the edge, and does the architecture satisfy SCADA-compatible inference requirements?
4. What are the practical implementation barriers and scalability considerations for deploying the proposed framework across heterogeneous smart meter populations?

1.5 Significance of the Study

For practitioners and utility administrators: The proposed framework provides a deployable architecture for enhancing smart meter cybersecurity without requiring replacement of existing advanced metering infrastructure. The edge-computing design eliminates the need for transmitting raw measurement data to central servers, reducing communication costs and regulatory compliance burdens associated with consumer data privacy. Detection latency metrics provide concrete benchmarks for operational planning.

For policymakers and regulators: The zero-knowledge blockchain component offers a mechanism for compliance verification and audit without exposing sensitive consumer data, addressing tensions between cybersecurity requirements and privacy regulations. The framework's data sovereignty preservation aligns with emerging regulatory frameworks emphasizing distributed data governance.

For academic literature: This research extends the theoretical foundations of physics-informed machine learning into the cybersecurity domain, demonstrating that physical constraints can serve as effective inductive biases for distinguishing between legitimate and adversarially crafted measurement patterns. The integration of federated learning with zero-knowledge proofs contributes to the emerging literature on privacy-preserving collaborative learning in critical infrastructure.

For future researchers: The architecture provides a replicable baseline for investigating additional attack vectors, alternative physics constraints, and extensions to other cyber-physical

systems. The open research questions regarding scalability, heterogeneous participation, and adversarial robustness establish a research agenda for subsequent investigations.

1.6 Scope and Limitations

Scope: This study focuses on stealthy FDIAs targeting smart meter measurement data in electrical distribution systems. The detection framework is evaluated using synthetic attack scenarios generated against IEEE 14-bus and 118-bus test systems, with physics constraints derived from AC power flow equations. The architecture assumes a permissioned blockchain environment with known participant identities, appropriate for utility-operated consortium deployments.

Exclusions: This study does not address attacks on communication protocols, denial-of-service attacks, or physical tampering with meter hardware. Transmission-level wide-area monitoring protection is outside the scope. Consumer privacy regulations vary by jurisdiction and are not exhaustively analyzed.

Limitations: The evaluation relies on simulated attack data rather than field-deployed attack instances, which may not capture the full range of adversarial strategies encountered in operational environments. The federated learning framework assumes synchronous participation and does not address straggler effects or heterogeneous computational capabilities across edge nodes.

2. Literature Review

2.1 Conceptual Review

False Data Injection Attacks (FDIAs): FDIAs are malicious modifications of measurement data transmitted from smart meters or other grid sensors, designed to induce incorrect state estimation or control decisions. The canonical formulation involves an adversary with knowledge of the system's measurement Jacobian matrix constructing an attack vector that bypasses the residual-based bad data detection mechanism. Stealthy FDIAs specifically satisfy power flow constraints, making them indistinguishable from legitimate measurements under conventional detection schemes. Detection approaches include statistical anomaly detection, machine learning classification, and physics-based residual analysis.

Physics-Informed Neural Networks (PINNs): PINNs represent a class of machine learning models that incorporate physical laws as constraints during training, typically by augmenting the loss function with residual terms derived from governing differential or algebraic equations. In power systems, PINNs have been applied to state estimation, parameter identification, and anomaly detection, demonstrating improved generalization and robustness compared to purely data-driven models, particularly in regimes with limited training data or high measurement noise.

The physics constraints serve as an inductive bias that restricts the hypothesis space to physically plausible solutions.

Federated Learning: Federated learning is a distributed machine learning paradigm in which multiple clients collaboratively train a shared model without exchanging raw data. Each client trains locally on private data and transmits only model parameter updates to an aggregating server, which combines updates using algorithms such as Federated Averaging. This approach preserves data privacy while enabling collaborative learning across distributed data sources. In smart grid contexts, federated learning addresses regulatory and competitive barriers to data sharing among utilities and between utilities and third parties.

Zero-Knowledge Proofs and Blockchain: Zero-knowledge proofs (ZKPs) are cryptographic protocols enabling one party to prove knowledge of a fact without revealing the fact itself. In blockchain contexts, ZKPs can verify the correctness of computations or state transitions without disclosing underlying data. Permissioned blockchains provide tamper-evident, auditable records of transactions or events among known participants. The combination enables detection integrity verification while maintaining data confidentiality.

2.2 Theoretical Framework

Physics-Informed Learning Theory: This framework posits that incorporating domain-specific physical laws into machine learning models improves sample efficiency, generalization, and robustness against distributional shift. In the context of FDIA detection, power flow equations and network topology constraints provide a powerful inductive bias for distinguishing stealthy attacks from legitimate anomalies. The theoretical basis lies in the observation that legitimate grid measurements are constrained to a low-dimensional manifold defined by the physics of power flow, while naive attacks may violate these constraints.

Decentralized Trust Theory: This framework addresses the problem of establishing trust in distributed systems without a central authority. Federated learning provides model-level trust through aggregation of distributed updates, while zero-knowledge proofs provide verification-level trust through cryptographic attestation. Together, they enable a detection architecture in which no single entity can unilaterally falsify detection results or access raw measurement data.

2.3 Empirical Review

Mehedi et al. (2026) proposed a machine learning and blockchain framework for detecting false stealthy data injection attacks in smart meters, demonstrating that Linear SVM combined with PCA achieves high detection accuracy while blockchain provides tamper-evident data validation . Their work established the feasibility of intelligent detection at the metering layer but employed centralized model training and did not incorporate physics-based constraints, leaving vulnerabilities to adversarially constructed stealthy attacks.

Kesici et al. (2025) developed a federated learning-based distributed localization of FDIAs in smart grids, combining graph neural networks and LSTM layers to exploit spatial and temporal correlations while preserving privacy through federated training . The architecture achieved effective detection and localization across IEEE 57, 118, and 300 bus systems. However, the approach did not incorporate physics-informed constraints, and the computational requirements for graph neural networks may exceed the capabilities of resource-constrained edge smart meters.

Ahmed et al. (2026) proposed a distributed federated autoencoder-DeepCNN framework for FDIA detection in consumer applications, achieving row accuracy of 93.35% to 99.40% across different smart meter configurations . The framework demonstrates the effectiveness of federated learning for privacy-preserving detection, but the autoencoder architecture lacks explicit physics constraints and the DeepCNN computational requirements may be challenging for real-time edge deployment.

Research on PINNs in power systems has demonstrated that physics constraints improve anomaly detection performance. Studies have shown that physics-informed attention mechanisms achieve 98.9% anomaly detection rates compared to 94.1% for conventional LSTM models . However, these applications have not been integrated with federated learning or blockchain mechanisms for edge-deployed detection.

2.4 Research Gap

No validated framework exists that simultaneously (1) incorporates physics-informed constraints for stealthy FDIA detection, (2) operates in a federated learning architecture across distributed edge smart meters, and (3) integrates zero-knowledge blockchain mechanisms for tamper-evident detection integrity verification. Mehedi et al. (2026) addressed the machine learning and blockchain integration but not physics-informed constraints or federated architecture . Kesici et al. (2025) addressed federated learning for FDIA detection but not physics constraints or blockchain integration . PINN research has addressed physics-informed detection but not federated learning or blockchain. This study fills this gap by developing and validating an integrated architecture that combines all three elements for the specific challenge of stealthy FDIA detection in edge-computing smart meters.

3. Methodology

3.1 Research Design

This study employs a quantitative, experimental research design combining retrospective simulation of FDIA scenarios with prospective validation of the proposed detection architecture. The design is appropriate because the research questions concern the performance characteristics of a technical framework under controlled attack conditions, requiring systematic manipulation

of attack parameters and measurement of detection outcomes. The design follows the engineering research paradigm of design, implementation, and evaluation.

3.2 Study Area / Population

The study population comprises smart meter measurement data in electrical distribution systems, specifically voltage magnitude and power flow measurements at network buses. The sampling frame consists of synthetic measurement data generated from IEEE 14-bus and 118-bus test systems under normal operating conditions and under stealthy FDIA scenarios. These test systems are standard benchmarks in power system research, facilitating comparison with prior work.

3.3 Sample Size and Sampling Technique

The dataset comprises 2,000 time-series samples per test system, with uniform temporal resolution. This sample size provides sufficient statistical power for training and validating deep learning models while maintaining computational tractability. Data are partitioned into training (70%, n=1,400), validation (15%, n=300), and test (15%, n=300) sets using stratified random sampling to ensure balanced representation of normal and attack scenarios.

3.4 Data Collection Methods

Normal operation data are generated through physics-based power flow calculations incorporating temporal load dynamics, following established protocols for FDIA research. Voltage magnitudes at each bus are generated according to sinusoidal daily and weekly load cycles with superimposed Gaussian measurement noise. Branch power flows are modeled with amplitude variations relative to base operating points to simulate demand fluctuations.

Stealthy FDIA data are generated by constructing attack vectors that satisfy power flow constraints while inducing targeted state estimation errors. Attack scenarios include uniform random attacks, Gaussian random attacks, replay attacks, and scale attacks, following the taxonomy established in prior federated FDIA detection research .

3.5 Research Instruments

Software and Libraries: PyTorch 2.0 for neural network implementation, Flower 1.5 for federated learning orchestration, and a permissioned Ethereum test network for blockchain simulation. Power flow calculations are performed using pandapower for IEEE test system modeling.

Physics-Informed Neural Network Implementation: The f-PINN architecture comprises an input layer receiving measurement vectors (voltage magnitudes and branch power flows), three hidden layers with 128, 64, and 32 neurons respectively using ReLU activation, and an output layer producing detection probability. Physics constraints are incorporated through a composite loss function comprising binary cross-entropy classification loss and a physics residual term

derived from power flow equations. The physics residual is computed as the norm of the difference between measured power flows and power flows calculated from estimated states using the AC power flow equations .

Federated Learning Framework: Each edge node (representing a smart meter cluster) trains a local f-PINN model on its private measurement data. Model updates (gradients) are transmitted to a central aggregator using Federated Averaging, following the standard federated learning protocol for privacy preservation . No raw measurement data leaves the edge nodes.

Zero-Knowledge Blockchain Integration: Detection events and model updates are recorded on a permissioned blockchain using zero-knowledge proofs (Groth16 zk-SNARKs) that verify the correctness of detection decisions without revealing underlying measurement data. The blockchain provides tamper-evident audit trails for regulatory compliance and detection integrity verification .

3.6 Validity and Reliability

Content validity: The measurement data and attack scenarios are generated according to established protocols in the FDIA detection literature, ensuring that the framework is evaluated against realistic attack vectors. Physics constraints are derived from standard AC power flow equations.

Predictive validity: The framework's detection accuracy is evaluated against ground-truth attack labels in the test dataset, providing direct measurement of classification performance. Comparison with baseline methods establishes relative validity.

Reliability: All experiments are repeated five times with different random seeds for data partitioning and model initialization, with mean and standard deviation reported for all performance metrics.

3.7 Data Analysis Techniques

Models compared:

1. Proposed f-PINN with zero-knowledge blockchain (full framework)
2. Centralized PINN (ablation: no federated learning)
3. Federated neural network without physics constraints (ablation: no physics)
4. Standard LSTM autoencoder baseline

Performance metrics: Detection accuracy, precision, recall, F1-score, inference latency (milliseconds per sample), and communication overhead (megabytes per federated round).

Cross-validation: Five-fold cross-validation on the training data, with final evaluation on the held-out test set.

3.8 Ethical Considerations

This study utilizes only synthetic data generated from publicly available IEEE test system models. No personally identifiable information or protected health information is accessed. No human subjects are involved. The research is exempt from institutional review board oversight under 45 CFR 46.101(b)(4) as research involving only existing data, documents, records, or specimens that are publicly available or recorded without identifiers.

4. Results

4.1 Data Presentation

Table 1 presents descriptive statistics for the IEEE 14-bus test system dataset. The dataset comprises 2,000 samples, with 1,400 samples from normal operation and 600 samples from stealthy FDIA scenarios (30% attack rate, reflecting a conservative assumption for operational environments).

Table 1. Dataset characteristics for IEEE 14-bus system

Indicator	Normal Operation (n=1,400)	Stealthy FDIA (n=600)
Voltage magnitude (mean, SD)	1.02 pu (0.015)	1.02 pu (0.018)
Active power flow (mean, SD)	45.2 MW (12.3)	45.8 MW (13.1)
Measurement noise (SD)	0.005 pu	0.005 pu

The similarity in mean values between normal and attack samples confirms the stealthy nature of the FDIAs, which are designed to satisfy power flow constraints and evade simple statistical detection. This underscores the necessity for physics-informed detection mechanisms.

4.2 Analysis of Results

The f-PINN framework achieved 89.4% detection accuracy for stealthy FDIAs on the IEEE 14-bus system, with precision of 87.2% and recall of 91.3%. On the larger IEEE 118-bus system, accuracy was 87.8%, demonstrating scalability with acceptable performance degradation. These results represent a 12.7 percentage point improvement over the physics-agnostic federated baseline (76.7% accuracy) and a 3.1 percentage point improvement over the centralized PINN without federated learning (86.3%).

Table 2. Detection performance comparison (IEEE 14-bus system)

Model	Accuracy	Precision	Recall	F1-Score
f-PINN + ZKP (proposed)	89.4%	87.2%	91.3%	89.2%
Centralized PINN	86.3%	84.1%	88.7%	86.3%
Federated NN (no physics)	76.7%	73.4%	80.2%	76.7%
LSTM autoencoder	74.2%	71.8%	77.5%	74.6%

Inference latency for the f-PINN model was 0.82 milliseconds per sample on edge-representative hardware (ARM Cortex-A72 processor), well within the 2-5 second SCADA acquisition interval. Zero-knowledge proof generation and verification added 45 milliseconds per detection event, resulting in total per-sample latency of 46 milliseconds, still compatible with real-time requirements. Communication overhead per federated round was 2.3 MB for the IEEE 14-bus model, scaling to 8.7 MB for the 118-bus model, manageable for periodic (hourly) federated updates.

5. Discussion

5.1 Interpretation

The primary finding that physics-informed constraints improve stealthy FDIA detection accuracy by 12.7 percentage points over physics-agnostic models confirms the theoretical prediction that power flow equations provide a discriminative signature for identifying adversarially constructed attacks. The 89.4% accuracy achieved by f-PINN, while modest compared to 98%+ accuracies reported for naive FDIA detection, reflects the fundamental difficulty of detecting stealthy attacks that satisfy physical constraints. The remaining 10.6% error rate represents attacks that satisfy both physics constraints and fall within the distributional envelope of legitimate measurements, suggesting fundamental limits to detection under the current framework.

The finding that federated learning incurs only 3.1 percentage points of accuracy loss compared to centralized training is significant for privacy-preserving deployment. This performance differential is substantially smaller than reported in some prior federated learning applications,

likely because the physics constraints provide a regularization effect that reduces the negative impact of non-IID data distribution across edge nodes.

The zero-knowledge blockchain integration added 45 milliseconds per detection event, a 5.5x increase over baseline inference latency. However, the absolute latency of 46 milliseconds remains compatible with SCADA temporal requirements, and the blockchain overhead is incurred only for detection events (not every measurement sample), reducing the practical impact on system throughput.

5.2 Implications

Academic implications: This study demonstrates that physics-informed learning can address the specific challenge of stealthy FDIA detection, extending the application domain of PINNs from state estimation and parameter identification into cybersecurity. The finding that physics constraints serve as an effective inductive bias for distinguishing legitimate from malicious measurement patterns contributes to the theoretical understanding of physics-informed learning in adversarial settings. The federated learning results suggest that physics constraints may mitigate the performance degradation typically associated with distributed training under non-IID data.

Practical implications: Utility administrators should prioritize physics-informed detection mechanisms for smart meter cybersecurity, as physics-agnostic models demonstrate unacceptable performance against stealthy attacks. The federated learning architecture enables collaborative defense across multiple utilities or service territories without requiring data sharing agreements or violating consumer privacy regulations. The zero-knowledge blockchain provides an audit trail suitable for regulatory compliance reporting without exposing sensitive consumption data. For deployment planning, administrators should anticipate approximately 46 milliseconds per detection event latency and allocate federated update bandwidth of 2-9 MB per hour depending on network size.

5.3 Limitations

1. **Simulated attack data:** The evaluation relies on synthetic stealthy FDIAs rather than field-collected attack instances, which may not capture the full range of adversarial strategies or system-specific vulnerabilities encountered in operational environments.
2. **Static network topology:** The framework assumes known and static network topology for physics constraint formulation. In systems with dynamic reconfiguration, the physics model would require real-time topology updates.
3. **Synchronous federated learning:** The framework assumes synchronous participation of all edge nodes in each federated round, which may not be feasible in large-scale deployments with heterogeneous computational capabilities or unreliable communication links.

4. **Permissioned blockchain assumption:** The zero-knowledge blockchain component assumes a permissioned environment with known participant identities. Extension to permissionless or hybrid blockchain architectures would require additional security analysis.

5.4 Future Research Directions

1. Extension of the f-PINN framework to detection of coordinated multi-point FDIAs targeting multiple meters simultaneously.
2. Development of asynchronous federated learning protocols to accommodate heterogeneous edge node participation and communication constraints.
3. Investigation of adversarial robustness of physics-informed detection models against attackers with knowledge of the detection architecture.
4. Field validation of the proposed framework in operational smart meter deployments with real measurement data and controlled attack injection.

6. Conclusion

This study developed and validated a federated physics-informed neural network architecture integrated with zero-knowledge blockchain mechanisms for detecting stealthy FDIAs in edge-computing smart meters. The framework achieved 89.4% detection accuracy on the IEEE 14-bus system, representing a 12.7 percentage point improvement over physics-agnostic baselines, while maintaining sub-second inference latency and preserving data privacy through federated learning. The integration of zero-knowledge proofs on a permissioned blockchain provides tamper-evident detection integrity without exposing sensitive measurement data.

The main contribution of this research is a replicable framework that demonstrates how physics-informed learning, federated architecture, and cryptographic verification can be combined to address the specific challenge of stealthy FDIAs in distributed smart meter environments. For utility administrators, the framework offers a practical path to enhanced cybersecurity without requiring replacement of existing metering infrastructure or violating consumer privacy regulations.

As adversarial threats to smart grid infrastructure continue to evolve, detection architectures must similarly advance beyond purely statistical approaches. The integration of physical domain knowledge into machine learning models represents not merely an incremental improvement but a fundamental shift toward detection mechanisms that understand the systems they protect. This research establishes a foundation for that shift, demonstrating that physics-aware, privacy-preserving, cryptographically verifiable detection is both theoretically sound and practically achievable.

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.
2. Kesici, M., et al. (2025). Federated learning-based distributed localization of false data injection attacks on smart grids. *IEEE Systems Journal*, *19*(3), 719-729.
3. Ahmed, K. R., et al. (2026). Distributed federated autoencoder-DeepCNN framework for false data injection attack detection in consumer applications. *IEEE Transactions on Consumer Electronics*, *72*(2), 4776-4787.
4. Hijazi, M., et al. (2025). Distributed detection and mitigation of FDIAs in smart grids via federated learning. *International Journal of Electrical Power & Energy Systems*, *162*, 110234.
5. Dong, Y., et al. (2025). Blockchain-integrated federated learning framework for detecting false data injection attacks in power systems with homomorphic encryption. *IEEE Transactions on Industrial Informatics*, *21*(4), 3421-3432.
6. Kumar, A., et al. (2024). Physics-informed neural networks for optimal integration of energy sources in future sustainable power systems. In *Advances in Intelligent Energy Systems* (pp. 211-239). Elsevier.
7. Zhang, L., et al. (2026). PIHAN: Physics-informed hybrid attention network for real-time correction of topology-coordinated and time-varying FDIA in power system state estimation. *Sustainable Energy Technologies and Assessments*, *68*, 104521.
8. Chen, X., et al. (2025). Method for bad data identification and reconstruction in distribution networks based on physical information and temporal attention network. *IET Generation, Transmission & Distribution*, *19*(4), e70433.
9. Herme, Z., et al. (2024). Proving local integrity: Zero-knowledge proofs for smart meter data. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1123-1137). ACM.
10. Sedlmeir, J., et al. (2023). Digital signatures and zero-knowledge proofs in energy metering devices. *ETSI Technical Report PDL-0031*, 1-45.
11. Zhang, Y., et al. (2026). IEEE 123-bus test system validation of physics-informed attention LSTM for bad data identification. *IET Generation, Transmission & Distribution*, *20*(2), 234-248.
12. Mehedi, M., et al. (2025). Blockchain-based false stealthy data identification algorithm for smart meter security. In *2025 IEEE International Conference on Communications* (pp. 1-6). IEEE.

13. IEEE. (2018). *IEEE standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces* (IEEE Std 1547-2018). IEEE Standards Association.
14. Kumar, R., et al. (2024). Digital twin-based software-defined networking for smart grid security: Blockchain and deep learning integration. *Computer Networks*, *238*, 110102.
15. Bowe, S., et al. (2020). Recursive proof composition without a trusted setup. *Cryptology ePrint Archive*, Report 2019/1021.