Hardware-in-the-Loop Evaluation of On-Chip Deep Learning and Smart Contract Execution for Securing Embedded AMI Devices Against FDI Attacks

Authors

Boyle Grace, Abiodun Okunola, Cephas John

Date; September 22, 2026

Abstract

Advanced Metering Infrastructure (AMI) devices face growing threats from false data injection (FDI) attacks that can manipulate energy consumption data and compromise grid operations. While machine learning and blockchain approaches have shown promise for detecting such attacks, their practical deployment on resource-constrained embedded smart meters remains unvalidated. This study presents a hardware-in-the-loop (HIL) evaluation framework that integrates an on-chip deep learning classifier with smart contract-based integrity verification for embedded AMI devices. The methodology employs a two-stage architecture: a quantized neural network deployed on an ARM Cortex-M7 microcontroller for real-time FDI detection, and an Ethereum-compatible smart contract executing on an FPGA-based hardware security module for tamper-proof logging. The HIL testbed reproduces realistic AMI communication constraints using MQTT protocol and evaluates detection performance under single and multi-feature FDI attack scenarios. Experimental results demonstrate that the hybrid architecture achieves 89.4% detection accuracy with 12.7 ms inference latency and 94.2% smart contract execution success

rate, while maintaining energy consumption below 45 mJ per inference. The framework outperforms centralized detection baselines by 8.3 percentage points in recall for stealthy attacks. These findings establish that on-chip deep learning combined with smart contract verification is feasible for embedded AMI security, providing a replicable HIL evaluation methodology for resource-constrained cyber-physical systems.

Keywords: Hardware-in-the-loop, Advanced Metering Infrastructure, False Data Injection, On-chip deep learning, Smart contract, Embedded security

1. Introduction

1.1 Background

The modernization of power grids through Advanced Metering Infrastructure (AMI) has transformed energy distribution by enabling bidirectional communication, real-time consumption monitoring, and demand-response optimization (Al-Turjman & Abujubbeh, 2019). Smart meters serve as the foundational sensing elements of this infrastructure, generating high-frequency measurement data that supports grid state estimation, billing verification, and load forecasting (Abir et al., 2021). However, the distributed and networked nature of AMI devices has significantly expanded the cyberattack surface, with false data injection (FDI) attacks emerging as a particularly insidious threat (Khalaf et al., 2024).

FDI attacks manipulate measurement streams by injecting crafted data that can bypass conventional bad-data detection mechanisms in power system monitoring (Musleh et al., 2020). In AMI contexts, successful FDI attacks can corrupt energy consumption records, mislead downstream scheduling decisions, and enable financial fraud through manipulated billing (Hou et al., 2019). Recent empirical evaluations have demonstrated that smart electric meters exhibit varying resilience to cyberattacks, with communication interfaces and data processing pipelines representing critical vulnerability points (Kumar et al., 2023). The stealthy nature of FDI attacks—particularly those exploiting temporal correlations in consumption data—renders signature-based detection inadequate, motivating the adoption of machine learning approaches capable of identifying anomalous patterns (Zhao et al., 2020).

1.2 Problem Statement

Existing research on FDI detection in smart grids has predominantly focused on centralized architectures where detection models execute at control centers or cloud servers (Mehedi et al., 2026; Al-Turjman & Abujubbeh, 2019). While these approaches demonstrate strong detection accuracy on benchmark datasets, they introduce fundamental limitations for practical AMI deployment. First, transmitting raw consumption data to centralized servers raises privacy concerns and incurs communication overhead that may violate latency requirements for real-time protection (Khalaf et al., 2024). Second, centralized detection creates a single point of failure—a

compromised control center disables protection across the entire AMI network (Abir et al., 2021). Third, the computational demands of deep learning models often exceed the capabilities of embedded smart meter processors, necessitating architectural compromises that degrade detection performance (Kumar et al., 2023).

Recent work has begun exploring edge-based detection using federated learning (Musleh et al., 2020) and lightweight neural networks optimized for embedded systems (Mehedi et al., 2026). However, these studies evaluate models primarily through simulation or software profiling, without validating performance under actual hardware constraints that govern embedded AMI devices (Khalaf et al., 2024). Furthermore, the integration of tamper-proof logging mechanisms—essential for forensic analysis and regulatory compliance—remains underexplored in embedded contexts. Blockchain-based integrity verification has been proposed for AMI security (Mehedi et al., 2026), but smart contract execution on resource-constrained hardware introduces latency and energy overheads that require systematic characterization (Zhao et al., 2020).

No validated framework exists that jointly evaluates on-chip deep learning inference and smart contract execution for FDI detection on embedded AMI hardware under realistic communication and computational constraints. This gap leaves critical questions unanswered: Can quantized neural networks achieve acceptable detection accuracy on microcontroller-class processors? What are the latency and energy costs of smart contract verification at the network edge? How do these components interact under hardware-in-the-loop conditions that reproduce actual AMI deployment scenarios?

1.3 Objectives of the Study

General Objective: To develop and evaluate a hardware-in-the-loop framework integrating on-chip deep learning and smart contract execution for securing embedded AMI devices against FDI attacks.

Specific Objectives:

1. To design a quantized neural network classifier optimized for execution on ARM Cortex-M-class microcontrollers that detects FDI attacks in smart meter measurement streams.
2. To implement a lightweight smart contract verification mechanism on FPGA-based hardware security module for tamper-proof logging of detection outcomes.
3. To validate the integrated framework using a HIL testbed that reproduces realistic AMI communication constraints and attack scenarios.

1.4 Research Questions

1. What detection accuracy can a quantized neural network achieve on embedded AMI hardware when processing realistic FDI attack patterns?

2. How do inference latency and energy consumption of on-chip deep learning compare with centralized detection baselines under hardware-in-the-loop conditions?
3. What are the smart contract execution success rates and latency characteristics when deployed on FPGA-based hardware security modules within resource-constrained AMI networks?

1.5 Significance of the Study

For practitioners and grid operators: The framework provides empirical evidence on the feasibility of edge-based FDI detection, enabling informed decisions about deploying on-device intelligence versus centralized alternatives. The quantified latency and energy metrics support infrastructure planning for AMI security upgrades.

For policymakers and regulators: The study demonstrates technical pathways for achieving regulatory requirements around data integrity and auditability in smart metering, offering a replicable architecture for compliance verification.

For academic literature: This research extends the body of knowledge on embedded AI security by providing HIL-validated performance characterization, addressing a methodological gap where simulation-only evaluations have predominated.

For future researchers: The testbed design and evaluation methodology provide a foundation for investigating additional attack vectors, alternative neural architectures, and emerging blockchain consensus mechanisms in cyber-physical security contexts.

1.6 Scope and Limitations

This study focuses on FDI detection in residential AMI devices using a two-stage on-chip architecture. The scope encompasses: (a) synthetic and benchmark FDI attack patterns representative of single and multi-feature manipulation; (b) ARM Cortex-M7 microcontroller execution of quantized neural networks; (c) Ethereum-compatible smart contract execution on FPGA-based hardware; (d) MQTT-based AMI communication protocol.

The study excludes: (a) evaluation of large-scale deployment scenarios exceeding 100 devices; (b) real-world utility network integration; (c) comparison of alternative consensus mechanisms beyond proof-of-authority; (d) economic analysis of deployment costs. Limitations include reliance on simulated attack data for certain scenarios and the assumption that benchmark consumption patterns adequately represent diverse geographic and demographic contexts.

2. Literature Review

2.1 Conceptual Review

Advanced Metering Infrastructure (AMI) constitutes an integrated system of smart meters, communication networks, and data management systems that enables bidirectional information flow between utilities and customers (Al-Turjman & Abujubbeh, 2019). AMI devices perform three core functions: measurement of electrical parameters, local computation for analytics, and communication for data reporting and control (Abir et al., 2021).

False Data Injection (FDI) Attacks represent a class of cyberattacks in which adversaries deliberately manipulate measurement data to mislead system state estimation and decision-making (Musleh et al., 2020). In AMI contexts, FDI attacks can target consumption readings, voltage measurements, or power quality indicators. Stealthy FDI attacks are specifically crafted to remain undetected by conventional statistical bad-data detection by exploiting knowledge of system topology and measurement correlations (Hou et al., 2019).

On-Chip Deep Learning refers to the deployment of trained neural network models directly on embedded processors, enabling local inference without data transmission to external servers (Mehedi et al., 2026). Quantization—reducing numerical precision from floating-point to integer representations—is a critical technique for fitting deep learning models within the memory and computational constraints of microcontrollers (Zhao et al., 2020).

Smart Contract Execution involves the automated, self-executing code that enforces rules and records state changes on a blockchain (Mehedi et al., 2026). For AMI security, smart contracts provide tamper-evident logging of detection events, creating immutable audit trails while enabling automated responses to identified threats.

Hardware-in-the-Loop (HIL) Evaluation is a testing methodology that couples simulated environments with physical hardware components, enabling realistic assessment of embedded system performance under controlled conditions (Kumar et al., 2023).

2.2 Theoretical Framework

Prospect Theory (Kahneman & Tversky, 1979) provides a behavioral foundation for understanding attack and defense dynamics in AMI security. Prospect theory posits that decision-makers evaluate outcomes relative to reference points and exhibit loss aversion—preferring to avoid losses over acquiring equivalent gains. In the FDI context, attackers weigh potential financial gains from data manipulation against detection risks, while defenders assess the perceived severity of undetected intrusions against the costs of implementing countermeasures. This framework explains why stealthy attacks—those minimizing detection probability—are particularly attractive to adversaries (Hou et al., 2019).

Distributed Systems Theory informs the architectural design by addressing fundamental trade-offs between latency, consistency, and availability in edge computing deployments (Khalaf et al.,

2024). The decision to process detection locally on AMI devices versus centralized servers represents a specific instantiation of this trade-off, with on-chip processing favoring latency and availability while accepting reduced model complexity relative to centralized alternatives (Abir et al., 2021).

2.3 Empirical Review

Musleh et al. (2020) conducted a comprehensive survey of FDI attack vulnerabilities in smart grids, identifying measurement manipulation as the most critical threat vector. Their analysis demonstrated that conventional bad-data detection methods fail against stealthy attacks that respect system constraints. However, the study remained theoretical, proposing detection architectures without experimental validation.

Zhao et al. (2020) developed an LSTM-based detection model for power grid FDI attacks, achieving 94.2% detection accuracy on simulated data. Their work established the efficacy of temporal deep learning for detecting stealthy manipulations but did not address deployment constraints on embedded hardware, limiting practical applicability to AMI contexts.

Kumar et al. (2023) performed experimental evaluation of smart electric meters under cyberattacks, identifying communication interfaces as primary vulnerability points. Their study provided valuable hardware-specific threat characterization but lacked integrated detection mechanisms, focusing exclusively on vulnerability identification rather than mitigation.

Mehedi et al. (2026) proposed a BiLSTM-attention framework combined with blockchain integrity verification, achieving 96.23% detection accuracy. This work represents the closest prior effort to the present study, combining deep learning detection with cryptographic verification. However, their evaluation employed software simulation exclusively, without characterizing performance on resource-constrained embedded hardware or evaluating smart contract execution overhead.

Al-Turjman and Abujubbeh (2019) reviewed IoT-enabled smart grid architectures, noting that edge intelligence represents a critical requirement for scalable AMI security. Their analysis identified on-device processing as a design imperative but did not provide experimental validation of specific detection frameworks.

2.4 Research Gap

Existing literature demonstrates strong detection accuracy for centralized machine learning approaches to FDI detection (Mehedi et al., 2026; Zhao et al., 2020) and establishes the theoretical benefits of edge processing (Al-Turjman & Abujubbeh, 2019). However, **no validated framework exists that jointly evaluates on-chip deep learning inference and smart contract execution for FDI detection on embedded AMI hardware under hardware-in-the-loop conditions**. The gap encompasses three dimensions: (1) absence of quantized neural network performance characterization on microcontroller-class processors for FDI detection; (2)

lack of smart contract execution profiling on embedded hardware security modules; and (3) no integrated HIL methodology that reproduces realistic AMI communication constraints during security evaluation.

This study fills this gap by developing and validating an integrated framework that addresses all three dimensions, providing empirical evidence on detection accuracy, latency, energy consumption, and execution success rates for the proposed architecture.

3. Methodology

3.1 Research Design

This study employs a design-based research methodology combined with quantitative performance evaluation. The design-based approach is appropriate for developing and iteratively refining the integrated hardware-software architecture, while quantitative evaluation enables statistical characterization of detection performance and resource consumption. The research proceeds in three phases: (1) system design and implementation; (2) HIL testbed configuration; (3) experimental evaluation under controlled attack scenarios.

3.2 Study Area / Population

The study targets residential AMI deployments representative of North American distribution networks. The reference topology is based on the IEEE 13-node test feeder, which captures characteristic features of residential distribution including unbalanced loads, mixed overhead/underground lines, and voltage regulation equipment (Khalaf et al., 2024). The embedded device population consists of ARM Cortex-M7 class microcontrollers representative of modern smart meter processing units.

3.3 Sample Size and Sampling Technique

The experimental design comprises 10 virtual smart meter instances communicating via MQTT protocol, consistent with the HIL co-simulation architecture validated in prior AMI security research (Kumar et al., 2023). Attack scenarios are stratified into two categories: single-feature FDI attacks (n=500 instances) and multi-feature FDI attacks (n=500 instances). This stratification enables comparative evaluation of detection performance across attack complexity levels.

3.4 Data Collection Methods

Data sources: The study utilizes the publicly available smart meter electricity consumption dataset employed by Mehedi et al. (2026), which contains residential consumption patterns with labeled normal and fraudulent records. For the present study, this dataset serves as the foundation for generating synthetic FDI attack instances that respect temporal and statistical properties of legitimate consumption.

Data types: Measurement features include active power, reactive power, voltage magnitude, and current magnitude sampled at 15-minute intervals. Labeled attack instances are generated through controlled manipulation of these features, following methodologies established in prior FDI attack simulation literature (Hou et al., 2019; Zhao et al., 2020).

Time periods: The dataset spans 12 months of hourly consumption records, with preprocessing to extract 15-minute interval features for model training and evaluation.

3.5 Research Instruments

Hardware components:

- **Target microcontroller:** STM32H747XI (ARM Cortex-M7 at 480 MHz, 2 MB Flash, 1 MB SRAM) representing smart meter processing capabilities.
- **FPGA-based hardware security module:** Xilinx Artix-7 FPGA (Nexys 4 DDR platform) configured for smart contract execution.
- **Network interface:** Ethernet connectivity for MQTT communication emulation.

Software tools:

- **Neural network training:** TensorFlow Lite for Microcontrollers and TensorFlow Model Optimization Toolkit for quantization-aware training.
- **Smart contract development:** Solidity for Ethereum-compatible contract implementation, Ganache for local blockchain simulation.
- **HIL orchestration:** Custom Python scripts for testbed synchronization and data collection.

Preprocessing: Consumption data is normalized using min-max scaling to [0,1], with temporal windowing creating 24-step sequences for model input. Attack labels are binary (0=normal, 1=FDI attack).

3.6 Validity and Reliability

Content validity: Attack scenarios are designed based on documented FDI attack methodologies from peer-reviewed literature (Musleh et al., 2020; Zhao et al., 2020), ensuring representative coverage of threat space.

Predictive validity: Detection accuracy is evaluated against held-out test data not used during model training, with stratified sampling preserving attack type distribution.

Reliability: Each experimental configuration is executed across 10 independent trials, with mean and standard deviation reported for all performance metrics. The HIL testbed is calibrated before each experimental session to ensure consistent timing and energy measurement.

3.7 Data Analysis Techniques

Detection model: A quantized 1D convolutional neural network (1D-CNN) architecture is evaluated, comprising three convolutional layers followed by two dense layers. The model is trained using quantization-aware training with 8-bit integer weights and activations.

Smart contract logic: The contract implements a simple logging function that records detection events (timestamp, device ID, classification result, confidence score) to the blockchain. Gas consumption and execution success rates are measured for each transaction.

Performance metrics: Detection accuracy, precision, recall, F1-score, inference latency (milliseconds), energy consumption (millijoules per inference), and smart contract execution success rate are recorded.

Statistical analysis: Comparisons between on-chip and centralized baselines employ paired t-tests with significance threshold $\alpha=0.05$.

3.8 Ethical Considerations

The study utilizes de-identified, publicly available consumption data with no personally identifiable information. No protected health information is accessed. The research protocol was reviewed and determined to be exempt from institutional review board oversight under 45 CFR 46.104(d)(4) as analysis of publicly available, de-identified data. Attack simulation is conducted in a controlled laboratory environment with no access to operational grid infrastructure.

4. Results

4.1 Data Presentation

Table 1 presents key performance indicators for the on-chip detection framework across attack categories.

Table 1. Detection Performance by Attack Type

Metric	Single-Feature FDI	Multi-Feature FDI	Combined
Accuracy (%)	92.1	86.7	89.4
Precision (%)	94.3	88.2	91.2
Recall (%)	90.8	84.3	87.6

Metric	Single-Feature FDI	Multi-Feature FDI	Combined
F1-Score (%)	92.5	86.2	89.4
Latency (ms)	11.8 (1.2)	13.6 (1.4)	12.7 (1.3)
Energy (mJ)	41.2 (3.8)	48.6 (4.2)	44.9 (4.0)

Values in parentheses represent standard deviation across 10 trials.

Table 2 presents smart contract execution characteristics.

Table 2. Smart Contract Execution Metrics

Metric	Value	Standard Deviation
Success Rate (%)	94.2	2.1
Gas Consumed	43,847	3,421
Execution Time (ms)	238	31
Throughput (tx/sec)	4.2	0.6

Table 3 compares the proposed on-chip framework with centralized detection baseline.

Table 3. Comparison with Centralized Baseline

Metric	On-Chip HIL	Centralized (Simulated)	Δ
Accuracy (%)	89.4	94.8	-5.4
Recall (%)	87.6	79.3	+8.3
Latency (ms)	12.7	156.4	-143.7

Metric	On-Chip HIL	Centralized (Simulated)	Δ
Energy (mJ)	44.9	N/A (server)	—

4.2 Analysis of Results

The on-chip detection framework achieved 89.4% combined accuracy, meeting the predetermined threshold of 85% for acceptable performance in embedded AMI contexts. Detection performance varied significantly by attack type, with single-feature attacks detected more reliably (92.1% accuracy) than multi-feature attacks (86.7% accuracy). This differential reflects the increased complexity of multi-feature manipulations, which more closely resemble legitimate consumption variations.

The recall advantage over centralized baseline (+8.3 percentage points) indicates that on-chip processing enables detection of stealthy attacks that centralized models miss, potentially due to temporal feature preservation at the edge. The latency reduction (143.7 ms improvement) supports real-time detection requirements for AMI protection.

Feature importance analysis identified active power variance (weight=0.31), consumption pattern deviation (weight=0.27), and voltage-current phase relationship (weight=0.19) as the most discriminative features for FDI detection.

Smart contract execution achieved 94.2% success rate with mean execution time of 238 ms—within acceptable bounds for non-critical logging operations. The 5.8% failure rate was attributed to network congestion during simultaneous multi-device transactions, suggesting optimization opportunities for high-density AMI deployments.

5. Discussion

5.1 Interpretation

The detection accuracy of 89.4% demonstrates that quantized neural networks can achieve operationally meaningful FDI detection on microcontroller-class hardware. The gap relative to centralized performance (-5.4 percentage points) represents the quantization and architectural constraints inherent to embedded deployment. However, the recall improvement (+8.3 percentage points) suggests that on-chip processing captures temporal attack signatures that centralized models with coarser feature representations fail to detect. This finding aligns with the theoretical expectation that edge processing preserves information fidelity (Al-Turjman & Abujubbeh, 2019) and extends prior work by quantifying this advantage in the FDI context.

The latency characteristics (12.7 ms mean inference) satisfy real-time detection requirements, enabling protection decisions within a single measurement cycle. The energy consumption (44.9

mJ per inference) remains within the power budget of typical smart meter processors (Mehedi et al., 2026), though continuous operation at 15-minute intervals yields approximately 4.3 J per day—negligible relative to meter power consumption.

Smart contract execution at 94.2% success rate validates the feasibility of blockchain-based integrity logging on embedded hardware. The 238 ms execution time introduces latency that precludes synchronous verification for every detection event; however, batch verification or selective logging for high-confidence detections offers viable operational strategies.

5.2 Implications

Academic implications: This study provides the first HIL-validated performance characterization of integrated on-chip deep learning and smart contract execution for AMI security. The methodology establishes a replicable framework for evaluating embedded AI security systems under realistic hardware constraints, addressing the simulation-only limitation identified in prior work (Mehedi et al., 2026). The differential detection performance across attack types contributes new knowledge about the boundary conditions of quantized neural network capabilities.

Practical implications: For grid operators, the results demonstrate that on-chip FDI detection is technically feasible with commercially available microcontrollers, requiring no specialized AI accelerators. Recommended monitoring metrics include: detection accuracy by attack category (target ≥85%), inference latency (target ≤20 ms), and smart contract success rate (target ≥90%). The recall advantage suggests that edge deployment should be prioritized for critical infrastructure nodes where stealthy attack detection is paramount.

5.3 Limitations

1. **Sample size and generalizability:** The evaluation employs 1,000 synthetic attack instances derived from a single geographic dataset. Detection performance may vary across diverse consumption patterns and attack methodologies not represented in the training data.
2. **Simulated attack data:** While attack generation follows established methodologies (Hou et al., 2019), the absence of real-world FDI attack traces limits ecological validity.
3. **Single hardware configuration:** Evaluation on a single microcontroller and FPGA platform may not capture performance variations across the broader embedded device ecosystem.
4. **Assumption of stable consumption patterns:** The framework assumes that legitimate consumption patterns remain stable over the detection window; seasonal or behavioral shifts may introduce false positives.

5.4 Future Research Directions

1. Extension of the HIL framework to evaluate alternative neural architectures (e.g., temporal convolutional networks, spiking neural networks) for improved multi-feature attack detection.
2. Longitudinal deployment study examining detection performance drift and model adaptation requirements over extended operational periods.
3. Investigation of lightweight consensus mechanisms optimized for embedded smart contract execution, reducing the 238 ms latency observed with proof-of-authority.
4. Cross-dataset validation using consumption patterns from diverse geographic and demographic contexts to assess generalizability.

6. Conclusion

This study developed and validated a hardware-in-the-loop framework integrating on-chip deep learning and smart contract execution for securing embedded AMI devices against FDI attacks. The framework achieved 89.4% detection accuracy with 12.7 ms inference latency on ARM Cortex-M7 hardware, demonstrating that quantized neural networks can provide operationally meaningful protection within embedded computational constraints. The smart contract verification layer achieved 94.2% execution success rate on FPGA-based hardware, enabling tamper-evident logging of detection events. The on-chip approach outperformed centralized baselines in recall for stealthy attacks by 8.3 percentage points, establishing a replicable framework for HIL evaluation of embedded AI security systems. For practitioners, these results indicate that edge-based FDI detection is technically feasible with commercially available microcontrollers, providing a pathway toward more resilient AMI deployments that preserve privacy, reduce latency, and maintain detection efficacy for evolving threat landscapes.

References

1. Abir, S. M. A., Anwar, A., Choi, J., & Kayes, A. S. M. (2021). IoT-enabled smart energy grid: Applications and challenges. *IEEE Access*, *9*, 50961–50981.
2. Al-Turjman, F., & Abujubbeh, M. (2019). IoT-enabled smart grid via SM: An overview. *Future Generation Computer Systems*, *96*, 579–590.
3. Hou, L., Luo, X., & Wang, X. (2019). An adaptive control defense scheme of false data injection attacks in smart grids. *2019 3rd International Symposium on Autonomous Systems (ISAS)*, 522–527.
4. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.
5. Khalaf, M., Ayad, A., & Tushar, M. H. K. (2024). A survey on cyber-physical security of active distribution networks in smart grids. *IEEE Access*, *12*, 29414–29444.
6. Kumar, H., Alvarez, O. A., & Kumar, S. (2023). Experimental evaluation of smart electric meters' resilience under cyber security attacks. *IEEE Access*, *11*, 55349–55360.
7. McKenna, K., Gotseff, P., & Chee, M. (2022). Advanced metering infrastructure for distribution planning and operation: Closing the loop on grid-edge visibility. *IEEE Electrification Magazine*, *10*(4), 58–65.
8. 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.
9. Musleh, A. S., Chen, G., & Dong, Z. Y. (2020). Vulnerabilities, threats, and impacts of false data injection attacks in smart grids: An overview. *2020 International Conference on Smart Grids and Energy Systems (SGES)*, 77–82.
10. Zhao, Y., Jia, X., & An, D. (2020). LSTM-based false data injection attack detection in smart grids. *2020 35th Youth Academic Annual Conference of Chinese Association of Automation (YAC)*, 638–644.
11. Alimi, O. A., Ouahada, K., & Abu-Mahfouz, A. M. (2020). A review of machine learning approaches to power system security and stability. *IEEE Access*, *8*, 113512–113531.
12. Avancini, D. B., Rodrigues, J. J., Martins, S. G., Rabêlo, R. A., Al-Muhtadi, J., & Solic, P. (2019). Energy meters evolution in smart grids: A review. *Journal of Cleaner Production*, *217*, 702–715.
13. Ben Fredj, O., Mihoub, A., Krichen, M., Cheikhrouhou, O., & Derhab, A. (2020). Cybersecurity attack prediction: A deep learning approach. *13th International Conference on Security of Information and Networks*, 1–6.

14. Breitenbach, J., et al. (2022). A systematic literature review of deep learning approaches in smart meter data analytics. *2022 IEEE 46th Annual Computers, Software, and Applications Conference (COMPSAC)*, 1337–1342.
15. Dhaou, I. B. (2023). Design and implementation of an internet-of-things-enabled smart meter and smart plug for home-energy-management system. *Electronics*, *12*(19), 4041.