



False Data Injection Attack Detection Based on Local Linear Embedding and Extreme Learning Machine

Idisire Japhet Hikon¹, Moshood A. Hambali², Siman Emmanuel^{3*}

^{1,2,3}Kwararafa University Wukari, Nigeria

^{2,3}Federal University Wukari, Nigeria.

*Corresponding Author

Siman Emmanuel

Kwararafa University
Wukari, Nigeria.

Article History

Received: 13.08.2025
Accepted: 18.09.2025
Published: 29.10.2025

Abstract: False Data Injection Attacks (FDIAs) pose significant threats to the reliability and stability of smart grid operations by manipulating measurement data to mislead monitoring and control systems. This research presents a lightweight framework for real-time detection of FDIAs that balances computational efficiency with predictive accuracy. The proposed pipeline integrates four core components: data preprocessing, dimensionality reduction using Locally Linear Embedding (LLE), classification through an Extreme Learning Machine (ELM), and deployment via a Gradio-based interactive interface. Data preprocessing ensures noise reduction and consistency in raw sensor inputs, while LLE captures essential nonlinear relationships and reduces feature complexity. The ELM classifier, chosen for its fast training speed and strong generalization capability, effectively discriminates between genuine and manipulated data. To enhance usability, the framework is deployed on a Gradio interface, providing an accessible platform for real-time monitoring and decision support. Experimental results on a synthetic Smart Grid FDIA dataset demonstrate the framework's ability to achieve high detection accuracy with minimal computational overhead, validating its suitability for practical, large-scale grid environments. By combining accuracy, speed, and user-friendly deployment, this study contributes a scalable solution to strengthen cyber-physical security in smart grids against evolving FDIA threats.

Keywords: False Data Injection Attack (FDIA), Smart Grid Security, Locally Linear Embedding (LLE), Extreme Learning Machine (ELM), Dimensionality reduction.

Cite this article:

Hikon, I. J., Hambali, M. A., Emmanuel, S., (2025). False Data Injection Attack Detection Based on Local Linear Embedding and Extreme Learning Machine. *ISAR Journal of Science and Technology*, 3(10), 96-102.

1. Introduction

The modernization of traditional power systems into smart grids represents a significant advancement in the energy sector, leveraging digital communication, advanced sensors, and intelligent devices to enhance efficiency, reliability, and sustainability (Wang & Lu, 2013). This transformation enables two-way communication between consumers and utility providers, facilitating real-time monitoring and automated control (He, Zhang, & Liu, 2017). However, this increased connectivity and digitization also expand the attack surface, introducing critical cybersecurity vulnerabilities (Bompard, Huang, Wu, & Cremenescu, 2017). Among the most formidable cyber threats to smart grids are False Data Injection Attacks (FDIAs) (Wang & Wu, 2020). In an FDIA, adversaries deliberately manipulate sensor measurements—such as voltage, current, frequency, or power factor—to corrupt the data integrity of the grid's monitoring

systems (Liu, Ning, & Reiter, 2011). This malicious activity can mislead system operators or automated control algorithms into making erroneous decisions, potentially triggering false alarms, disrupting load balancing, degrading energy efficiency, and, in extreme cases, causing large-scale blackouts (Zhu, Liu, Chang, & Guo, 2023; Yuan, He, Zhu, & Li, 2020). The primary challenge in detecting FDIAs stems from their design to mimic legitimate system noise and remain within plausible operational ranges, thereby evading conventional threshold-based or rule-driven detection mechanisms (Paudel, 2024). Existing approaches to FDIA detection can be broadly categorized into statistical methods and computational intelligence techniques. Traditional methods, such as residual analysis based on state estimation (Monticelli, 1999) and statistical process control (Hu et al., 2023), are often lightweight and interpretable but have proven inadequate against sophisticated, stealthy FDIAs that are designed to bypass consistency tests (Li, Wei, Li, Dong, & Shahidehpour, 2022).



Conversely, deep learning models, including convolutional and recurrent neural networks (Energies, 2022; Peng et al., 2022), have demonstrated high detection accuracy by capturing complex, non-linear patterns in the data. Nevertheless, their practical deployment in real-time smart grid environments is often constrained by substantial computational demands, extensive hyperparameter tuning, and the need for large volumes of labeled training data (Kiani et al., 2024). This creates a critical research gap: the need for a detection mechanism that harmonizes high accuracy with computational efficiency and real-time operational capability. To bridge this gap, this study proposes a novel, lightweight hybrid framework that integrates Locally Linear Embedding (LLE) for non-linear dimensionality reduction and an Extreme Learning Machine (ELM) for rapid classification.

The contributions of this work are fourfold:

1. **A Hybrid FDIA Detection Model:** We propose and validate an integrated pipeline that leverages LLE to project high-dimensional smart grid data into a lower-dimensional manifold, preserving essential local structures and enhancing the separability between normal and attack instances. This transformed data is then classified using a single-hidden-layer ELM, chosen for its exceptional training speed and strong generalization performance.
2. **Demonstration of Efficacy on Imbalanced Data:** The model is rigorously evaluated on a synthetic smart grid dataset with a realistic class imbalance (16.67% attack instances), demonstrating robust performance without relying on synthetic oversampling techniques, thus reflecting real-world operational conditions.
3. **Comprehensive Performance Benchmarking:** The proposed LLE-ELM framework is shown to achieve high detection accuracy (99%), precision (100% for the attack class), and recall (96% for the attack class), significantly outperforming many traditional methods and rivaling the accuracy of more computationally expensive deep learning models.

2. Related Works

Traditional FDIA detection is often rooted in power system state estimation. These methods typically analyze the measurement residuals—the difference between observed values and values estimated from the system model—under the assumption that malicious injections will cause statistically significant anomalies in these residuals (Monticelli, 1999). Recent enhancements have incorporated adaptive thresholding and advanced statistical measures to improve sensitivity. For instance, Paudel (2024) conducted a comparative evaluation of statistical methods, including Median Absolute Deviation (MAD), Kullback–Leibler Divergence (KLD), and Cumulative Sum (CUSUM), noting their high interpretability and low computational cost. Similarly, Zhu, Liu, Chang, and Guo (2023) proposed the Measurement–Eigenvalue Residual Similarity (MERS) test, which achieved higher accuracy and lower false alarm rates compared to conventional residual tests. Despite these advancements, a fundamental limitation persists: stealthy FDIAs can be meticulously crafted to be consistent with the physical laws of the grid, thereby bypassing residual-based detection mechanisms (Liu,

Ning, & Reiter, 2011). These methods are often ineffective against coordinated attacks that subtly manipulate multiple measurements simultaneously (Paudel, 2024). To overcome the limitations of statistical methods, machine learning techniques have been widely adopted for their ability to learn complex patterns from data without explicit mathematical modeling of system physics. Algorithms such as Support Vector Machines (SVMs), Random Forests, and Extreme Learning Machines (ELMs) have been successfully applied to the FDIA detection problem (Hu et al., 2023; Kiani et al., 2024). Wei and Zhang (2023) demonstrated that such methods generally achieve higher accuracy and robustness compared to purely statistical approaches. The primary advantage of ML models lies in their generalization capability. However, their performance is highly contingent on the quality and representativeness of the training data. In real-world smart grid settings, obtaining a large, accurately labeled dataset that encompasses the vast spectrum of potential attack vectors is a significant challenge. Furthermore, the performance of many ML models can degrade severely in the presence of the inherent class imbalance typical of cybersecurity data, where attack instances are vastly outnumbered by normal operations (Peng et al., 2022). Deep learning models represent the forefront of detection accuracy, capable of identifying highly complex and non-linear attack patterns. Architectures such as Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, and autoencoders have been leveraged to model the temporal and spatial correlations in smart grid data (Yuan, He, Zhu, & Li, 2020). A CNN-LSTM autoencoder model applied to energy consumption forecasting achieved detection rates above 95% by using prediction errors as an anomaly signal (Energies, 2022). To address data privacy concerns in distributed grid systems, Li, Wei, Li, Dong, and Shahidehpour (2022) proposed a federated deep learning approach using Transformer models encrypted with the Paillier cryptosystem, enabling collaborative learning across grid nodes without sharing raw data. The main drawback of these deep learning approaches is their resource intensity. They require substantial computational power for training, extensive hyperparameter tuning, and large volumes of data, rendering them less suitable for real-time deployment on resource-constrained hardware like edge devices or legacy grid infrastructure (Wang & Wu, 2020).

2.1. Hybrid and Lightweight Frameworks

A promising research direction involves hybrid frameworks that combine feature extraction or dimensionality reduction with fast classifiers. These models aim to strike a balance between the high accuracy of deep learning and the efficiency of traditional ML. The typical pipeline involves transforming the raw feature space to remove redundancy and enhance separability before applying a lightweight classifier. For example, Dou, Wu, and Yue (2022) used Variational Mode Decomposition (VMD) to preprocess power signals and an Online Sequential ELM (OS-ELM) for classification, reporting strong performance with low latency. Wang et al. (2022) explored deep variants of Locally Linear Embedding (LLE-Net) to learn robust low-dimensional embeddings for downstream classifiers. The Extreme Learning Machine (ELM) is particularly favored in such pipelines due to its unique architecture: it randomly initializes input weights and biases and computes output weights analytically via a least-squares solution, leading to extremely fast training times without iterative

tuning (Huang, Zhu, & Siew, 2006; Peng et al., 2022). This makes ELM exceptionally well-suited for real-time applications. However, a gap remains in the exploration of non-linear manifold learning techniques for this specific purpose. While Principal Component Analysis (PCA) and autoencoders are commonly used, the application of Locally Linear Embedding (LLE), a technique specifically designed to preserve local neighbourhood structures in non-linear data, is less explored. LLE is particularly advantageous for anomaly detection as it can unfold the intrinsic geometry of data, making subtle, attack-induced outliers more separable in a lower-dimensional space (Wang et al., 2022). Its unsupervised nature is also beneficial when labelled attack data is scarce. Therefore, this work proposes a novel hybrid framework that integrates LLE for non-linear dimensionality reduction with ELM for rapid classification. This combination leverages the complementary strengths of both algorithms: LLE's ability to reveal the underlying manifold structure of smart grid data for improved separability, and ELM's capability to provide fast, accurate, and efficient classification, making the entire pipeline suitable for real-time FDIA detection.

3. The Proposed Hybrid LLE-ELM Framework

This section details the architecture and implementation of the proposed lightweight framework for detecting False Data Injection Attacks (FDIAs) in smart grids. The overall pipeline, illustrated in Figure 1, is designed for accuracy and real-time efficiency, integrating four core stages: (1) Data Preprocessing, (2) Dimensionality Reduction via Locally Linear Embedding (LLE), (3) Classification using an Extreme Learning Machine (ELM), and (4) Real-time Deployment via an interactive interface.




Figure 1: The integrated pipeline showing the flow from data input through preprocessing, feature transformation via LLE, classification by ELM, to final prediction and deployment.

3.1. Dataset Description and Preprocessing

The model was developed and evaluated using a synthetic Smart Grid FDIA dataset, meticulously designed to emulate realistic grid operations and sophisticated cyber-attack scenarios (OpenAI, 2025). The dataset consists of 6,000 instances, each characterized

by four critical sensor measurements that are common targets for adversarial manipulation:

- Voltage (V): Electric potential difference.
- Current (A): Flow of electric charge.
- Power Factor: Measure of power utilization efficiency.
- Frequency (Hz): Grid frequency stability.

The target variable is a binary label where 0 indicates normal operation and 1 signifies an FDIA attack. The dataset exhibits a realistic class imbalance, with 5,000 normal samples (83.33%) and 1,000 attack samples (16.67%), accurately reflecting the infrequency of security incidents compared to normal operations in a real grid.

The preprocessing pipeline involved two critical steps to ensure data quality and model performance:

- Stratified Train-Test Split:** The dataset was partitioned using a stratified random split, allocating 70% (4,200 samples) for training the model and 30% (1,800 samples) for testing. Stratification ensured that the original class imbalance was preserved in both subsets, providing a realistic basis for evaluating model performance.
- Feature Standardization:** All four numerical features were standardized using Scikit-learn's StandardScaler. This process transformed the features to have a mean of zero and a standard deviation of one. Standardization is crucial as it ensures that all features contribute equally to the distance calculations in the subsequent LLE algorithm, preventing variables with inherently larger scales from dominating the model.

3.2. Dimensionality Reduction with Locally Linear Embedding (LLE)

Although the raw feature space is only four-dimensional, the complex, non-linear relationships between parameters like Voltage, Current, and Power Factor can obscure the discriminative boundary between normal and attack states. To address this, we employed Locally Linear Embedding (LLE), a powerful non-linear manifold learning technique (Roweis & Saul, 2000). LLE works by preserving the local neighborhood relationships between data points, effectively unfolding the intrinsic geometric structure of the data into a lower-dimensional space where patterns become more apparent. The algorithm operates in two main phases: first, it identifies the local neighborhood for each data point and computes weights that best reconstruct each point from its neighbors. Second, it finds a low-dimensional embedding that preserves these local reconstruction weights. In this study, the number of neighbors was set to 10 through empirical tuning to optimally capture local structure without overfitting to noise. The output dimensionality was set to 2, which significantly reduces computational load for the classifier and provides a clear visual representation of the separation between classes.

3.3. Classification with Extreme Learning Machine (ELM)

The 2D feature vectors produced by the LLE transformation serve as the input to an Extreme Learning Machine (ELM) classifier. ELM is a single-hidden-layer feedforward neural network chosen for its exceptional training speed and strong generalization performance (Huang, Zhu, & Siew, 2006). The key innovation of ELM lies in its simplified training mechanism. Unlike traditional

neural networks that use iterative gradient-based methods like backpropagation, ELM randomly initializes the weights and biases connecting the input layer to the hidden layer. The output weights connecting the hidden layer to the output layer are then computed analytically using a least-squares solution. This approach eliminates the need for computationally expensive iterative tuning, reducing training time from minutes to seconds. In this implementation, the hidden layer consisted of 50 neurons with a sigmoid activation function, a configuration determined empirically to provide a good balance between model capacity and computational efficiency. The model was trained on the LLE-transformed features from the training set.

3.4. Real-Time Deployment via Gradio Interface

To transition from a theoretical model to a practical tool, the trained LLE-ELM pipeline was deployed as an interactive web application using the Gradio library (Abid et al., 2019). This deployment demonstrates the framework's suitability for real-time monitoring.

The deployment process involved:

- i. **Model Serialization:** The trained Standard Scaler, LLE model, and ELM parameters were saved to disk using the joblib library for persistence.
- ii. **Interface Development:** A user-friendly Gradio interface was built with four numerical input fields corresponding to the four sensor measurements.

- iii. **Backend Inference Engine:** A prediction function was created to (a) scale any new input data using the saved scaler, (b) apply the LLE transformation to project the data into 2D, and (c) perform a fast forward pass through the ELM network to generate a prediction.
- iv. **Instantaneous Feedback:** The interface provides immediate output, displaying either "Normal (No Attack Detected)" or "FDIA Attack Detected," thereby acting as a proof-of-concept for a grid operator's real-time diagnostic tool.

4. Experimental Results and Discussion

This section presents a comprehensive evaluation of the proposed LLE-ELM framework. The analysis begins with an exploratory data analysis (EDA) to uncover inherent data characteristics, followed by an assessment of the dimensionality reduction performance using LLE. The core classification results of the ELM model are then detailed using standard performance metrics. Finally, the practical deployment of the model via an interactive Gradio interface is demonstrated. A thorough EDA was conducted to understand the underlying structure of the synthetic smart grid dataset and to identify preliminary patterns that distinguish FDIA instances from normal operations. Descriptive Statistics: Key statistical measures for the four features are summarized in Table 1. The wide ranges observed in Voltage (211.44V to 279.30V) and Current (4.12A to 22.21A), along with minimum Power Factor values of 0.5, strongly suggest the presence of anomalous data points, which are confirmed to be FDIA instances.

Table 1: Descriptive Statistics of the Smart Grid FDIA Dataset.

Feature	Mean	Std Dev	Min	25%	50%	75%	Max
Voltage (V)	233.41	9.72	211.44	227.35	231.27	235.95	279.30
Current (A)	10.80	2.44	4.12	9.20	10.33	11.83	22.21
Power Factor	0.879	0.116	0.500	0.864	0.911	0.956	0.999
Frequency (Hz)	49.66	0.88	44.34	49.76	49.95	50.11	50.93

Class Distribution: The dataset is inherently imbalanced, comprising 5,000 normal instances (83.33%) and 1,000 FDIA attack instances (16.67%). This distribution was intentionally preserved to mirror the real-world scenario where cyber-attacks are rare events compared to normal system operations, providing a realistic challenge for the detection model. **Feature Distribution and Correlation:** Visual analysis through histograms and boxplots revealed distinct patterns. FDIA attacks typically manifest as abnormally high Voltage and Current readings, and significantly lower Power Factor values. Frequency deviations during attacks were more subtle but present. A correlation matrix heatmap (Figure 2) quantified these relationships, revealing a strong negative correlation between the attack label (1) and both Power Factor ($r = -0.888$) and Frequency ($r = -0.866$), and a strong positive correlation with Voltage ($r = 0.779$) and Current ($r = 0.751$). This confirmed that the selected features are highly relevant for distinguishing FDIAs.




Figure 2: Heatmap showing strong negative correlations between FDIA attacks and Power Factor/Frequency, and strong positive correlations with Voltage and Current.

4.2. Dimensionality Reduction with LLE

The application of Locally Linear Embedding (LLE) successfully transformed the original 4-dimensional feature space into a 2-dimensional manifold while preserving the critical local structures of the data. The result of this transformation is visualized in the scatter plot in Figure 3.

The LLE projection yielded a clear and compelling separation between the two classes:

- i. Normal Operations (Label 0): These instances form a single, dense, and well-defined cluster in the central region of the transformed space. This tight clustering indicates stable and consistent system behavior.
- ii. FDIA Attacks (Label 1): Attack instances are more dispersed and form several distinct sub-clusters on the periphery of the main cluster, particularly along the horizontal axis. This dispersion reflects the diverse strategies and magnitudes of false data injection.

The visible margin of separation between the primary cluster of normal data and the clusters of attack data demonstrates that LLE effectively enhanced the separability of the classes. It successfully unfolded the complex, non-linear relationships in the original data, making subtle attack patterns more distinguishable and thus simplifying the subsequent classification task.




Figure 3: The LLE projection showing clear separability between the dense cluster of normal operations (blue) and the dispersed FDIA attack instances (red).

4.3. Model Performance Evaluation

The transformed 2D data from LLE was used to train and evaluate the Extreme Learning Machine (ELM) classifier. The model was tested on a held-out set of 1,800 samples. The confusion matrix (Figure 4) and the classification report (Table 2) provide a detailed account of the model's performance.

The results are exceptional:

Overall Accuracy: The model achieved an accuracy of 99%.

Precision: For the FDIA attack class, precision is 100%. This means every instance predicted as an attack was a true attack, resulting in zero false positives. This is a critical performance

metric for operational environments, as false alarms can lead to unnecessary and costly operational interventions.

Table 2: Classification report of the LLE-ELM model on the test set.

Class	Precision	Recall	F1-Score	Support
Normal (0)	0.99	1.00	1.00	1500
FDIA Attack (1)	1.00	0.96	0.98	300
Accuracy			0.99	1800
Macro Avg	1.00	0.98	0.99	1800
Weighted Avg	0.99	0.99	0.99	1800




Figure 4: Confusion matrix showing the model's predictions against the true labels.

Recall (Sensitivity): The recall for the attack class is 96%, indicating that the model correctly identified 96% of all actual attacks. The 13 missed attacks (false negatives) represent the primary area for potential improvement. F1-Score: The harmonic mean of precision and recall for the attack class is 0.98, indicating a very strong balance between the two metrics despite the class imbalance.

4.4 Discussion of Results

The LLE-ELM framework demonstrates a powerful balance between high accuracy and computational efficiency. The 100% precision is a standout feature, making the model highly reliable. The 96% recall on a minority class within an imbalanced dataset is also highly commendable and confirms that the LLE transformation effectively prepared the data for the ELM classifier. The entire training process for the ELM was completed in a fraction of a second, underscoring its suitability for real-time systems.

5. Conclusion

The escalating sophistication of cyber-physical attacks on critical infrastructure necessitates the development of intelligent, efficient, and deployable defense mechanisms. This research addressed the critical challenge of detecting False Data Injection Attacks (FDIAs) in smart grids by proposing a novel hybrid framework that

integrates Locally Linear Embedding (LLE) for non-linear dimensionality reduction and an Extreme Learning Machine (ELM) for rapid classification. The experimental results demonstrate the unequivocal success of this approach. The LLE algorithm proved highly effective in unraveling the complex, non-linear relationships inherent in smart grid sensor data, projecting it into a two-dimensional manifold where the separation between normal operations and malicious attacks became markedly clear. This transformation provided an optimal input for the classifier. The ELM, in turn, leveraged this prepared data to achieve an exceptional overall accuracy of 99% on an imbalanced test set. Crucially, the model attained a precision of 100% for the attack class, eliminating false alarms, and a high recall of 96%, ensuring that the vast majority of attacks were detected. Furthermore, the entire training process was completed in seconds, highlighting the framework's computational efficiency. Beyond algorithmic performance, this study underscored the practical applicability of the model through its deployment as an interactive web application via Gradio. This interface serves as a proof-of-concept, demonstrating how the proposed framework can be integrated into a real-time monitoring system, providing grid operators with immediate, actionable intelligence. In conclusion, the LLE-ELM framework presents a lightweight, accurate, and computationally efficient solution to the FDIA detection problem. It successfully balances the high accuracy typically associated with complex deep learning models with the low computational cost required for real-time deployment in resource-constrained smart grid environments, thereby offering a viable and practical tool for enhancing grid cybersecurity.

5.3. Future Work

The findings of this study open several compelling avenues for future research to enhance the robustness and applicability of the proposed framework:

- i. **Validation on Real-World Data:** The most critical next step is to validate the model using large-scale, real-world data streams from operational smart grids or high-fidelity testbeds. This would provide a more rigorous assessment of its performance under practical conditions.
- ii. **Incorporation of Adaptive Learning:** Future iterations could integrate an Online Sequential ELM (OS-ELM) or other incremental learning algorithms. This would enable the model to learn continuously from new data, adapting to new normal behaviors and emerging attack patterns without requiring full retraining.
- iii. **Expanded Feature Engineering:** Research should explore the inclusion of additional features such as harmonic distortions, synchrophasor data (PMU), and system-level variables to create a more comprehensive detection system.
- iv. **Multi-Attack Detection:** The framework could be extended to form a multi-classification system capable of not only detecting but also identifying and categorizing different types of cyber-attacks (e.g., FDIAs, DoS, malware).
- v. **Enhanced Deployment Architecture:** The Gradio prototype can be evolved into a full-fledged dashboard

integrated with alerting systems (e.g., SMS, email notifications) and visualization tools for grid operators, and deployed within a cloud or edge-computing architecture for scalable performance.

References

1. Abid, A., Abdalla, A., Abid, A., Khan, D., Alfozan, A., & Zou, J. (2019). Gradio: Hassle-Free Sharing and Testing of ML Models in the Wild. *arXiv preprint arXiv:1906.02569*.
2. Bompard, E., Huang, T., Wu, Y., & Cremenescu, M. (2017). Classification and trend analysis of threats origins to the security of power systems. *International Journal of Electrical Power & Energy Systems*, 86, 13–22.
3. Dou, C., Wu, D., & Yue, D. (2022). A hybrid method for false data injection attack detection in smart grid based on variational mode decomposition and OS-ELM. *[Conference/Journal]*.
4. Energies. (2022). False data injection attack detection in smart grid using energy consumption forecasting. *Energies*, 15(13), 4877.
5. He, H., Zhang, J., & Liu, X. (2017). A parallel detection scheme for false data injection attacks in smart grid. *IEEE Transactions on Smart Grid*, 8(4), 1809–1820.
6. Hu, P., Li, X., Zhang, Y., Wang, Q., & Chen, Z. (2023). Detection of false data injection attacks in smart grids based on statistical learning. *Sensors*, 23(3), 1683.
7. Huang, G. B., Zhu, Q. Y., & Siew, C. K. (2006). Extreme learning machine: theory and applications. *Neurocomputing*, 70(1-3), 489-501.
8. Kiani, R., et al. (2024). A survey on extreme learning machines for outlier detection. *[Survey Article]*.
9. Li, Y., Wei, X., Li, Y., Dong, Z., & Shahidehpour, M. (2022). Detection of false data injection attacks in smart grid: A secure federated deep learning approach. *arXiv preprint arXiv:2209.00778*.
10. Liu, Y., Ning, P., & Reiter, M. K. (2011). False data injection attacks against state estimation in electric power grids. *ACM Transactions on Information and System Security (TISSEC)*, 14(1), 1–33.
11. Monticelli, A. (1999). State Estimation in Electric Power Systems: A Generalized Approach. Springer Science & Business Media.
12. Paudel, S. (2024). An evaluation of methods for detecting false data injection attacks in the smart grid. *Frontiers in Computer Science*.
13. Peng, Xinggan; Li, Hanhui; Feng, Yong; Razul, Sirajudeen Gulam; Chen, Zhebin; Lin, Zhiping. *An extreme learning machine for unsupervised online anomaly detection in multivariate time series*. *Neurocomputing*, Vol. 501, August 2022, pp. 596-608. <https://doi.org/10.1016/j.neucom.2022.06.042>
14. Roweis, S. T., & Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500), 2323-2326.

15. Wang, W., & Lu, Z. (2013). Cyber security in the Smart Grid: Survey and challenges. *Computer Networks*, 57(5), 1344–1371.
16. Wang, Z., & Wu, L. (2020). An overview of false data injection attacks and their detection methods in smart grids. *IEEE Transactions on Information Forensics and Security*, 15, 2496–2513.
17. Wang, J., Shao, Z., Huang, X., Lu, T., Zhang, R., & Chen, X. (2022). *Deep locally linear embedding network*. Information Sciences, 614, 416–431. <https://doi.org/10.1016/j.ins.2022.10.036>
18. Wei, L., & Zhang, Q. (2023). Detection of false data injection attack in smart grid based on improved UKF. *Journal of System Simulation*, 35(7), 1508–1516.
19. Yuan, X., He, P., Zhu, Q., & Li, X. (2020). Adversarial examples: Attacks and defenses for deep learning. *IEEE Transactions on Neural Networks and Learning Systems*, 31(7), 2435–2451.
20. Zhu, Y., Liu, R., Chang, D., & Guo, H. (2023). Detection of false data injection attacks on power systems based on measurement–eigenvalue residual similarity test. *Frontiers in Energy Research*, 11, 1285317.
21. OpenAI. (2025). *ChatGPT (GPT-5) [Large language model]*. <https://chat.openai.com>