Quantum Autoencoder-Based Anomaly Detection for Cyberattacks in Smart Power Systems

Franco Cirillo[1,*], Christian Esposito[1]

[1]University of Salerno, Via G. Paolo II 132, Fisciano, 84084, Italy

Abstract

The growing complexity and interconnectivity of modern power systems have increased their vulnerability to sophisticated cyberattacks, making real-time anomaly detection a critical challenge. Classical machine learning approaches, while effective, face limitations in scalability and representation power, especially when dealing with high-dimensional and temporally correlated data. Quantum autoencoders (QAEs) offer a promising alternative by leveraging the principles of quantum computing to encode and process information in more compact and expressive forms. In this work, we investigate the use of a stacked QAE architecture for cyberattack detection in smart grid environments, aiming to capture both spatial and temporal patterns through quantum-enhanced learning. Preliminary results show that even with minimal input dimensionality, the quantum model performs competitively with classical baselines, highlighting the potential of QAEs as a foundation for scalable, next-generation anomaly detection frameworks in critical infrastructure systems.

Keywords

Quantum Autoencoder (QAE), Quantum Machine Learning (QML), Power Systems, Attack Detection

1. Introduction

The digital transformation of power systems [1], driven by the integration of smart devices and automated control technologies, has significantly enhanced the efficiency and responsiveness of modern electrical grids [2]. However, this evolution has also introduced new challenges in terms of cybersecurity. The widespread deployment of Internet of Things (IoT) components in smart grids increases the attack surface, exposing critical infrastructures to malicious threats capable of disrupting system stability [3]. In such complex and interconnected environments, timely detection of anomalous behaviours becomes a critical requirement.

Machine Learning (ML) methods have shown promising results in the field of anomaly detection and cyber-intrusion recognition within smart grids. Classical approaches often rely on deep neural networks, which are effective but can be computationally intensive and face limitations when applied to high-dimensional, noisy, or partially labeled data [4]. In recent years, Quantum Machine Learning (QML) [5] has emerged as a novel paradigm that leverages quantum computing to address some of these challenges, offering new opportunities in terms of representational power and resource efficiency [6].

Among various QML architectures [7], Quantum Autoencoders (QAE) [8] stand out for their ability to learn compressed representations of input data using a limited number of qubits. Inspired by their classical counterparts, QAEs are particularly well-suited for feature extraction and anomaly detection tasks, making them a valuable tool for cybersecurity in quantum-aware environments. By learning to map high-dimensional input states into lower-dimensional quantum latent spaces, QAEs can facilitate efficient detection of outliers or abnormal patterns with reduced quantum resource requirements.

This work explores the quantum implementation of an anomaly detection pipeline for smart grids, by translating an existing classical neural network model [9] into a stacked Quantum Autoencoder architecture. The proposed model is designed to compress and reconstruct input data related to critical events in the power system, thereby learning a meaningful quantum representation that can differentiate

AIQxQIA 2025: International Workshop on AI for Quantum and Quantum for AI | co-located with ECAI 2025, Bologna, Italy

*Corresponding author.

fracirillo@unisa.it (F. Cirillo); esposito@unisa.it (C. Esposito)
0009-0006-9599-5996 (F. Cirillo); 0000-0002-0085-0748 (C. Esposito)

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-01-20

between normal operations and potential cyber threats. This integration of QML techniques into the cybersecurity domain aims to open a path toward scalable and quantum-native solutions for next-generation smart grid protection.

The paper is structured as follows: Section 2 introduces key background concepts, including classical and quantum autoencoders, and the characteristics of power system cyberattack datasets. Section 3 reviews the reference deep learning model that inspired our quantum approach and the literature review of QAE for anomaly detection. Section 4 presents the proposed methodology, describing each stage of the quantum learning pipeline. Finally, Section 5 summarizes the findings and outlines future directions.

2. Background

2.1. Classical Autoencoders

Autoencoders [10] are a class of unsupervised neural networks designed to learn efficient representations of input data by compressing and reconstructing it. They consist of two main components: an encoder, which maps the input into a lower-dimensional latent space, and a decoder, which reconstructs the original input from this compressed representation. The model is trained to minimize the reconstruction error, typically measured using a loss function such as mean squared error (MSE).

In anomaly detection, autoencoders are trained on normal data only. When presented with anomalous or outlier data during inference, the model generally exhibits higher reconstruction errors, thus enabling the detection of deviations from expected patterns. This makes autoencoders particularly suitable for intrusion detection tasks, where malicious behavior often manifests as statistically rare or unseen input patterns.

Variants such as denoising autoencoders, variational autoencoders, and stacked autoencoders have been proposed to improve robustness, generalization, and expressive power. However, as the dimensionality of the data and the size of the neural network increase, the training and inference costs become prohibitive, especially in resource-constrained or real-time environments such as smart grids.

2.2. Convolutional and Recurrent Neural Networks in Autoencoder Architectures

Convolutional Neural Networks (CNNs) [11] and Recurrent Neural Networks (RNNs) [12] are widely used in modern machine learning due to their ability to extract structured features from high-dimensional data. When combined with autoencoders, they enhance the model's capacity to learn meaningful representations by exploiting spatial and temporal patterns in the input data.

CNNs are particularly effective in identifying local spatial correlations. Originally developed for image recognition tasks, CNNs apply convolutional filters that move across input data to detect edges, textures, and other hierarchical features. In autoencoder architectures, a CNN can be used in the encoder to compress input data into a latent representation that retains the most relevant spatial information. The decoder then reconstructs the original input by applying transposed convolutional layers. This approach is especially suitable for grid-like data structures, such as sensor readings organized as matrices or time-frequency maps.

RNNs, on the other hand, are designed to process sequential data by maintaining a hidden state that evolves over time. This makes them well-suited for learning temporal dependencies, such as trends or recurrent patterns in time-series data. In the context of autoencoders, RNNs are typically used in the encoder and decoder to compress and reconstruct sequences of data, respectively. This allows the model to capture long-term dependencies that may not be apparent in individual observations.

The integration of CNNs and RNNs into autoencoder frameworks enables the extraction of both spatial and temporal features from complex datasets [13]. This is particularly useful in domains like power system monitoring, where anomalies may manifest through subtle variations across multiple sensors and over time. These capabilities make CNN and RNN based autoencoders valuable building blocks in anomaly detection architectures, and serve as the foundation for the quantum variants explored in this work.

2.3. Quantum Autoencoders

In the context of QML, Quantum Autoencoders (QAEs) have been proposed as the quantum analog of classical autoencoders. QAEs are designed to compress quantum states into lower-dimensional subspaces using parameterized quantum circuits. By learning an encoding circuit that maps input quantum states to a smaller number of qubits, and a decoding circuit that reconstructs the original state, QAEs aim to preserve essential information while reducing quantum resource usage.

From a practical standpoint, QAEs operate within variational quantum circuits, which are trained using classical optimization methods in a hybrid quantum-classical loop. They are particularly well-suited for near-term quantum hardware (NISQ devices), as they can be implemented using shallow circuits and a small number of qubits. In anomaly detection scenarios, a QAE trained on "normal" quantum states can be used to identify anomalous states by measuring deviations in the reconstruction fidelity.

The potential of QAEs lies not only in their ability to reduce the dimensionality of quantum data but also in offering fundamentally different representational capabilities compared to classical autoencoders. When applied to cybersecurity tasks, QAEs may enable efficient, real-time detection of cyber threats in quantum-aware infrastructures, such as future quantum-enhanced smart grids.

In analogy with their classical counterparts, various architectural enhancements can be envisioned for quantum autoencoders to improve their robustness, scalability, and expressive power. Several classical variants of autoencoders have proven effective in different tasks, each introducing specific inductive biases tailored to particular data characteristics or learning goals. Extending these ideas to the quantum domain may further enrich the functionality of QAE models.

For instance, Denoising Autoencoders (DAEs) are designed to learn representations that are robust to noise by artificially corrupting the input and training the network to reconstruct the original, clean version. This is particularly relevant in quantum systems, where noise is intrinsic due to decoherence and imperfect gate operations. A quantum analog of DAE could enable the QAE to learn stable encodings even in the presence of hardware-level noise.

Another important variant is the Sparse Autoencoder (SAE), which imposes a sparsity constraint by penalizing activations that exceed a predefined threshold. This forces the network to activate only a limited number of neurons (or qubits in a quantum context), helping to prevent overfitting and promoting the learning of essential features. A quantum sparse autoencoder could implement similar constraints on the activation probabilities of qubits or on the usage of specific quantum gates, potentially improving generalization while keeping circuit depth low—a critical factor for near-term quantum devices.

Additionally, Variational Autoencoders (VAEs) introduce a probabilistic framework that learns a distribution over latent variables, enabling generative capabilities such as sampling new data similar to the input distribution. While the direct implementation of VAEs in quantum computing is non-trivial due to differences in probability modeling, research on Quantum Variational Autoencoders is ongoing and may lead to novel generative models with applications in anomaly synthesis and attack simulation.

Finally, Stacked Autoencoders represent a hierarchical architecture in which multiple autoencoders are layered, with each level learning increasingly abstract features based on the output of the previous one. Training typically involves two phases: a pre-training stage where each layer is trained individually to reconstruct its input, followed by a fine-tuning phase where the entire deep model is optimized jointly. The quantum analog would involve stacking multiple quantum encoding and decoding layers, potentially implemented as a sequence of variational circuits, each trained to refine the latent representation. This design could improve the QAE's ability to capture complex, high-level patterns associated with sophisticated attack signatures or temporal correlations in smart grid data.

Integrating these classical strategies into the quantum realm remains an open research direction. Nevertheless, exploring such variants could lead to more powerful and resilient quantum anomaly detection models capable of operating under practical constraints and noisy conditions. The flexibility of the QAE framework makes it a promising candidate for future quantum-native cybersecurity solutions.

2.4. Power System Cyberattack Datasets

In the context of machine learning-based anomaly detection for power systems, the availability of realistic and diverse datasets is essential for both training and evaluation. A widely adopted dataset in this domain was developed through a collaboration between Mississippi State University and Oak Ridge National Laboratory. This dataset [14] provides a rich collection of simulated events designed to reflect a broad spectrum of real-world scenarios, including both natural system behaviors and malicious cyberattacks.

The dataset models a simplified yet representative power grid infrastructure, comprising key components such as generators, transmission lines, circuit breakers, and Intelligent Electronic Devices (IEDs). These IEDs are responsible for controlling breakers through distance protection schemes and are susceptible to both valid operational commands and malicious interference.

The dataset encompasses 37 distinct scenarios, grouped into three main categories: natural disturbances, normal system operation, and cyber-induced attack events. Attack types include remote tripping command injection, relay setting manipulation, and false data injection—each simulating adversarial strategies aimed at compromising the stability of the power system.

To support various machine learning tasks, the original data was processed into multiple formats and classification schemes, including binary, three-class, and multiclass versions. This structure facilitates the development and benchmarking of detection algorithms under different levels of granularity and complexity, making it a valuable resource for evaluating quantum and classical models in the smart grid cybersecurity context.

3. State of the Art

3.1. Stacked Autoencoder-Based Deep Learning for Cyberattack Detection in Power Control Systems

A work that inspired the quantum adaptation presented in this study is the framework proposed in [9]. This work introduces a deep learning architecture specifically designed to detect cyberattacks in electric power control systems by distinguishing malicious interventions from naturally occurring system faults.

The model is structured around two sparse autoencoders (SAEs), each integrated with a specific type of neural network to capture different aspects of the data: a CNN and a RNN. This design enables the system to learn both spatial dependencies arising from sensor correlations, and temporal dependencies arising from events occurring over time.

The first component, referred to as CNN-SAE, uses convolutional layers within the encoder-decoder structure to learn spatial patterns from input data arranged as a two-dimensional matrix, emulating the layout of an image. This approach facilitates the detection of spatial anomalies across multiple sensor readings or measurement types.

The second component, RNN-SAE, operates on the latent representation generated by the CNN-SAE. It leverages a recurrent structure in both encoding and decoding phases to capture temporal correlations across sequential observations. This layered, stacked autoencoder architecture allows the model to progressively refine its internal feature representation, enabling the detection of subtle and sophisticated attack patterns.

The final output of the stacked model is passed through a Softmax classification layer, which produces a prediction indicating the presence or absence of a cyberattack. One of the key advantages of this framework is its ability to learn directly from raw monitoring data, without the need for manual feature engineering or domain-specific preprocessing. This makes it applicable to a wide range of industrial control systems, provided that a sufficient volume of time-series data is available.

The success of this approach in detecting complex attacks in critical infrastructure motivated the exploration of a quantum-native counterpart. By adopting a similar architectural principle, stacked autoencoders enhanced with convolutional and recurrent mechanisms, this work aims to investigate the feasibility and potential advantages of implementing such a model using quantum computing components.

3.2. Related Work on Quantum Autoencoders for Anomaly Detection

Recent years have seen increasing interest in the application of quantum autoencoders for anomaly detection across a variety of domains. These models leverage the unique properties of quantum computing, such as high-dimensional feature representation and entanglement, to improve the efficiency and performance of machine learning tasks.

In [15], the authors explore the use of variational quantum circuits for anomaly detection for physics purposes. Their quantum autoencoder demonstrates superior performance compared to classical models in identifying anomalous physics events, while also showing that such results are reproducible on current quantum hardware. This work highlights the potential of QAEs in high-energy physics for learning complex background distributions in a fully unsupervised manner.

In [8], quantum autoencoders are applied to time series anomaly detection. The authors evaluate performance across multiple ansätze and show that quantum models outperform classical autoencoders while using significantly fewer parameters and training steps. This reinforces the hypothesis that QAEs can be more efficient than classical models, even on sequential data.

A hybrid approach is presented in [16], where a classical autoencoder is augmented with a parametrized quantum circuit (PQC) inserted into the bottleneck layer. Applied to standard benchmark datasets and a predictive maintenance use case in gas power plants, the hybrid model improves precision, recall, and F1-score. This demonstrates how embedding quantum modules into classical pipelines can enhance anomaly detection performance in industrial settings.

In the domain of network security, [17] proposes three quantum anomaly detection frameworks by integrating QAEs with quantum classifiers such as quantum SVM, random forest, and k-nearest neighbor. These hybrid models are tested on benchmark IoT and computer network datasets and demonstrate promising detection capabilities, especially the QAE-kNN combination, which achieves the highest accuracy.

Another notable application is credit card fraud detection, as presented in [18]. The proposed QAE-based fraud detection model (QAE-FD) achieves high performance metrics, including a G-mean of 0.946 and an AUC of 0.947. The study shows that quantum autoencoders can effectively address the challenge of imbalanced datasets and offer efficient, scalable solutions for real-time detection.

While all these works underline the strong potential of quantum autoencoders in anomaly detection, it is important to note that none of them focus on the domain of power systems or smart grids. To the best of our knowledge, no existing work has applied QAEs to cyberattack detection in interconnected power control environments, despite the critical importance of securing such infrastructure. The present study fills this gap by proposing and evaluating a quantum autoencoder-based architecture specifically tailored for anomaly detection in smart grid systems. Inspired by a classical deep learning framework based on stacked convolutional and recurrent autoencoders, this work investigates the advantages of its quantum counterpart using real-world power system attack datasets.

4. Methodology

The proposed system is designed as a structured pipeline composed of multiple stages, each fulfilling a specific role within the learning and classification process. The goal is to replicate and evaluate, in a quantum context, the structure presented in prior reference implementations [9] and to analyze the results comparatively to a classic version implementation. The pipeline consists of the following phases:

- Data preprocessing
- A quantum autoencoder (QAE) with a quantum convolutional neural network (QCNN) at its core
- A second QAE using a quantum recurrent neural network (QRNN)
- A classifier for supervised learning and output generation

Throughout the experimentation, a fine-tuning process was performed on the main hyperparameters, including:

- learning rate (`lr`)
- batch size (`batch_size`)
- number of qubits used for encoding (`n_qubits`)
- number of quantum layers (`n_layers`)

The following metrics were recorded to evaluate performance:

- final training and test accuracy (`final_train_acc`, `final_val_acc`)
- final training and test loss (`final_train_loss`, `final_val_loss`)

These metrics were stored for analysis and comparison between quantum and classical classifiers under optimized conditions.

4.1. Data Preprocessing

The preprocessing stage prepares the data for training and testing within the quantum learning pipeline. A random sample of 2,000 examples was selected from the dataset to balance computational efficiency and data variability. Columns exhibiting highly skewed distributions or extremely large values were log-transformed using the `log1p` function to stabilize the distributions. Outliers were clipped at an upper threshold, and missing values were imputed using the mean of each column.

The dataset was split into training and test sets using an 80/20 ratio. A `StandardScaler` was applied to center and normalize the data, ensuring zero mean and unit variance across all features. This transformation improves convergence during training and enables compatibility with dimensionality reduction techniques such as Principal Component Analysis (PCA).

PCA was used to reduce dataset dimensionality while preserving at least 95% of total variance. After PCA, a `MinMaxScaler` normalized the resulting features into the $[0, 1]$ interval, preparing them for the quantum circuits. The transformed data was divided into fixed-size blocks of 4 features.

4.2. Quantum Autoencoder with QCNN

The first QAE utilizes a quantum convolutional neural network (QCNN) as its encoder. The QCNN transforms each input block into a quantum feature representation. A quantum simulator was instantiated with 4 qubits, and a 6-layer variational quantum circuit was constructed. Each layer applied parameterized rotations to each qubit, followed by entangling operations modeled through random layers composed of rotation gates and CNOT gates.

The encoding process involved mapping classical values to quantum states using `RY` rotations, followed by entangling layers. At the output, each qubit was measured via the Pauli-Z operator, yielding four numerical values representing the quantum-transformed features.

A decoder stage was implemented symbolically by applying a transformation that inverts the encoded outputs (i.e., $1-$measurement), completing the QAE reconstruction phase. The entire dataset was passed through the encoder-decoder pair, generating a set of quantum feature matrices used in subsequent stages.

4.3. Quantum Autoencoder with QRNN

The second QAE adopts a quantum recurrent neural network (QRNN) architecture to capture dependencies between adjacent qubits, analogously to how classical RNNs handle sequential data. Each input block was encoded using `AngleEmbedding` with `RY` rotations, initializing the quantum state for processing.

The encoder then applied three types of parameterized single-qubit rotations (`RX`, `RY`, `RZ`) on each qubit, followed by entangling only adjacent qubits via CNOT gates. This localized entanglement mimics the recurrence in sequential learning. The decoder, mirroring the encoder, applied another set of `RX`, `RY`, and `RZ` rotations to complete the reconstruction of the data.

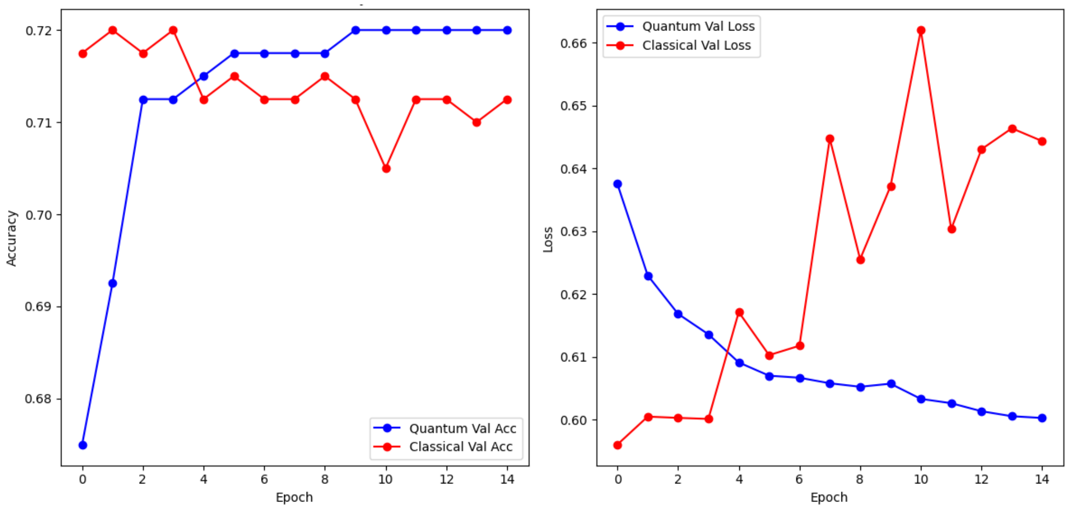

Figure 1: Test accuracy and loss plot comparing the proposed hybrid model to a fully classical classifier

4.4. Classifier and Training

After the autoencoding stages, the reconstructed outputs were used as input to a classifier. Two distinct models were introduced:

- A hybrid quantum-classical classifier, incorporating the stacked QAEs
- A fully classical classifier used as a baseline for comparison

A unified `train_model` function was implemented to encapsulate training, test, and performance tracking. The models were trained over multiple epochs using mini-batch gradient descent. For quantum models, the learning rate was set to 1×10^{-3}, while for classical models it was set to 1×10^{-2}.

The training loss included `CrossEntropyLoss` for classification, and additionally, `Mean Squared Error (MSELoss)` for quantum models to measure reconstruction error within the QAE. Each epoch consisted of random shuffling of training data, forward pass, loss calculation, backpropagation, and parameter update using the Adam optimizer.

4.5. Performance Evaluation and Comparison

Upon completion of training, both models were evaluated based on test accuracy and loss. The results indicate that both approaches converged to a similar accuracy of approximately 72%. Despite the simplicity of the input—only four features and a relatively small sample size—the quantum-enhanced model demonstrated encouraging behaviour, showing greater stability in the later epochs and a steady improvement trend.

Visual analysis was performed through two comparative plots in Figure 1:

- A test accuracy plot showing performance trends across epochs
- A test loss plot reflecting optimization behaviour and robustness

Regarding test loss, the quantum model exhibited a consistently decreasing profile, while the classical model, although initially lower, showed signs of instability and degradation over time. These observations suggest that the quantum model may be more robust under the current experimental conditions. Nevertheless, further experimentation with larger datasets, more complex feature sets, and extended training is required to draw definitive conclusions about its generalization capability.

5. Conclusion

This paper presents a novel approach for anomaly detection in smart grid systems based on stacked quantum autoencoders, integrating quantum convolutional and recurrent neural network components. Our results demonstrate that quantum-enhanced models can achieve performance comparable to or better than classical baselines, despite the use of reduced feature sets and limited training data. These findings support the feasibility of leveraging quantum representations to capture meaningful spatial-temporal dependencies in cyber-physical systems.

While the results are promising, this work serves primarily as a proof of concept. Additional experimentation on larger and more complex datasets, exploration of alternative quantum circuit designs, and real-device testing are required to further assess the advantages of quantum approaches in operational settings. Nevertheless, this study highlights the potential of quantum machine learning as a valuable tool for enhancing the robustness and adaptability of cybersecurity mechanisms in modern power infrastructures.

Acknowledgments

This research was funded by the NGIsargasso project (Europe Horizon Grant No. 101092887), Open Call 4 FRQGAN4AD project.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to: Grammar and spelling check, Paraphrase and reword. After using this tool, the authors reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] U. Bhadani, Pillars of power system and security of smart grid, International journal of innovative research in science engineering and technology 13 (2024) 10–15680.

[2] M. Cavus, Advancing power systems with renewable energy and intelligent technologies: A comprehensive review on grid transformation and integration, Electronics 14 (2025) 1159.

[3] S. Abdelkader, J. Amissah, S. Kinga, G. Mugerwa, E. Emmanuel, D.-E. A. Mansour, M. Bajaj, V. Blazek, L. Prokop, Securing modern power systems: Implementing comprehensive strategies to enhance resilience and reliability against cyber-attacks, Results in engineering (2024) 102647.

[4] H. Yamasaki, N. Isogai, M. Murao, Advantage of quantum machine learning from general computational advantages, 2023. arXiv:2312.03057.

[5] F. Cirillo, C. Esposito, Quantum machine learning for intrusion detection on noisy quantum computers, in: 2025 IEEE International Conference on Quantum Computing and Engineering (QCE), 2025.

[6] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, P. J. Coles, Challenges and opportunities in quantum machine learning, Nature Computational Science 2 (2022) 567–576.

[7] F. Cirillo, C. Esposito, Intrusion detection system based on quantum generative adversarial network, in: Proceedings of the 17th International Conference on Agents and Artificial Intelligence - Volume 1: QAIO, INSTICC, SciTePress, 2025, pp. 830–838. doi:10.5220/0013397800003890.

[8] R. Frehner, K. Stockinger, Applying quantum autoencoders for time series anomaly detection, Quantum Machine Intelligence 7 (2025) 1–21.

[9] G. D'Angelo, F. Palmieri, A stacked autoencoder-based convolutional and recurrent deep neural network for detecting cyberattacks in interconnected power control systems, International Journal of Intelligent Systems 36 (2021) 7080–7102.

[10] D. Bank, N. Koenigstein, R. Giryes, Autoencoders, Machine learning for data science handbook: data mining and knowledge discovery handbook (2023) 353–374.

[11] L. Mohammadpour, T. C. Ling, C. S. Liew, A. Aryanfar, A survey of cnn-based network intrusion detection, Applied Sciences 12 (2022) 8162.

[12] A. Sherstinsky, Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network, Physica D: Nonlinear Phenomena 404 (2020) 132306.

[13] M. S. Ibrahim, S. M. Gharghory, H. A. Kamal, A hybrid model of cnn and lstm autoencoder-based short-term pv power generation forecasting, Electrical Engineering 106 (2024) 4239–4255.

[14] B. Barika, Power system, https://www.kaggle.com/datasets/bachirbarika/power-system, 2019. Accessed: June 30, 2025.

[15] V. S. Ngairangbam, M. Spannowsky, M. Takeuchi, Anomaly detection in high-energy physics using a quantum autoencoder, Physical Review D 105 (2022) 095004.

[16] A. Sakhnenko, C. O'Meara, K. J. Ghosh, C. B. Mendl, G. Cortiana, J. Bernabé-Moreno, Hybrid classical-quantum autoencoder for anomaly detection, Quantum Machine Intelligence 4 (2022) 27.

[17] M. Hdaib, S. Rajasegarar, L. Pan, Quantum deep learning-based anomaly detection for enhanced network security, Quantum Machine Intelligence 6 (2024) 26.

[18] C. Huot, S. Heng, T.-K. Kim, Y. Han, Quantum autoencoder for enhanced fraud detection in imbalanced credit card dataset, IEEE Access (2024).