Federated Quantum Generative Adversarial Network for Intrusion Detection

Franco Cirillo
University of Salerno
Fisciano (SA), Italy
fracirillo@unisa.it

Christian Esposito
University of Salerno
Fisciano (SA), Italy
esposito@unisa.it

***Abstract*—In increasingly complex digital environments, Intrusion Detection Systems (IDS) are essential for maintaining network security. Traditional machine learning techniques face challenges, including difficulties with high-dimensional data and performance issues when dealing with unbalanced datasets. Although Generative Adversarial Networks (GANs) provide an acceptable alternative by improving data creation, their traditional implementations are computationally demanding and have limitations when it comes to examining intricate data distributions.**

Quantum GANs (QGANs) use the benefits of quantum computing to solve these problems. Therefore, a distributed approach improves load balancing and mitigate the impact of quantum noise for NISQ devices. The proposed model efficiently learns the distribution of benign data by employing a federated hybrid QGAN architecture that combines a quantum generators with a classical discriminators. This allows the model to produce samples that closely mimic patterns found in real data. The generator in the proposed QGAN is a variational quantum circuit (VQC), and the discriminator is a neural network. Tested on the NSL-KDD dataset, the federated QGAN offers a promising option for next-generation IDS.

***Index Terms*—Quantum Machine Learning (QML), Quantum Generative Adversarial Network (QGAN), Intrusion Detection System (IDS), Anomaly Detection, Federated Learning (FL)**

I. Introduction

Intrusion Detection Systems (IDS) are essential for safeguarding digital infrastructures by identifying malicious activities, unauthorized access, and policy violations [1]. Among IDS types, Network Intrusion Detection Systems (NIDS) play an important role in monitoring network traffic for anomalies that may indicate security threats. However, the growing sophistication of cyberattacks often surpasses the capabilities of traditional detection methods.

Anomaly detection is a robust technique for intrusion detection, relying on a baseline of normal behavior rather than predefined attack signatures [2]. It excels in identifying novel attack vectors but faces challenges like false positives and computational complexity.

Machine Learning (ML) has revolutionized intrusion detection, with supervised models classifying known threats and unsupervised models detecting new anomalies. Yet, ML's reliance on high-quality datasets and computational resources poses scalability issues [3]. A significant advancement in ML is the application of Generative Adversarial Networks (GANs) [4]. GANs enhance intrusion detection by generating synthetic data to augment training sets, addressing class imbalances, and detecting subtle anomalies. However, classical ML models, including GANs, face limitations in handling high-dimensional data complexity [5].

To overcome these constraints, Quantum Machine Learning (QML) offers a transformative solution [6]. Leveraging quantum computing, QML can address scalability and performance limitations of classical ML. Specifically, Quantum Generative Adversarial Networks (QGANs) extend GAN principles to the quantum domain, utilizing quantum generators and/or discriminators [7]. QGANs demonstrate promising efficiency and performance in anomaly detection, requiring fewer parameters with sparse or imbalanced datasets [8], making them ideal for cybersecurity applications [9].

In this scenario, distributed ML enhances scalability, efficiency, and privacy, making it ideal for large-scale, real-time intrusion detection [10]. The integration with QGANs can enhance the resiliency and efficiency of these models. By distributing the workload across multiple nodes, this approach improves load balancing, ensuring that computational demands are evenly managed. Additionally, it can mitigate the impact of quantum noise for NISQ devices by enabling quantum resources to work collaboratively.

This paper explores the integration of QGANs with a federated approach into NIDS for anomaly detection, making the following key contributions:

- Introduction of a federated QGAN-based model specifically designed for intrusion detection.
- Experimentation across various configurations, archiving an accuracy of 0.9125 and an F1-score of 0.9034, with a reduction of time complexity compared to the single version.

To the best of our knowledge, this is the first work that implement and evaluate a Federated QGAN for intrusion detection.

The paper is structured as follows: Section II introduces QGANs and the dataset used. Section III reviews relevant advancements in quantum machine learning and generative models. Section IV details the data preprocessing steps and the proposed QGAN architecture. Section V analyzes the results, focusing on statistical similarities, outlier detection, and evaluation metrics. Lastly, Section VI summarizes the contributions and suggests future research directions.

II. Background

QGANs extend classical GANs by integrating the principles of quantum computing, enabling them to tackle complex data modeling challenges more efficiently [11]. QGANs can be designed in two configurations: full quantum and hybrid.

In the full quantum version both the generator and discriminator are quantum systems [12]. The hybrid version combines a quantum generator with a classical discriminator, blending quantum computational advantages with the robustness of classical systems. This architecture is particularly well-suited for near-term quantum devices, as it balances practicality and performance [7].

A hybrid QGAN comprises two core components:

- Quantum Generator (G): A PQC that generates quantum states, encoding data samples drawn from a prior distribution such as uniform or Gaussian. The generator learns the data distribution by iteratively refining its parameters.
- Classical Discriminator (D): A neural network that evaluates the similarity between generated and real data, assigning a probability score to classify samples as real or fake.

The adversarial training process involves the interaction between them:

- Discriminator Objective: The discriminator maximizes its ability to correctly classify real samples as real and generated samples as fake. Its loss function, based on binary cross-entropy (BCE), is defined as:

$$\mathcal{L}_D = -\mathbb{E}_{x\sim p_{\text{real}}(x)}[\log D(x)] - \mathbb{E}_{z\sim p_z(z)}[\log(1-D(G(z)))] \tag{1}$$

  where $D(x)$ represents the probability that x is real, and $D(G(z))$ represents the probability that a generated sample $G(z)$ is real.
- Generator Objective: The generator minimizes the discriminator's ability to distinguish real from generated samples by maximizing $D(G(z))$. This is equivalent to minimizing:

$$\mathcal{L}_G = -\mathbb{E}_{z\sim p_z(z)}[\log D(G(z))] \tag{2}$$

In QGANs, the generator is realized using a PQC. This design begins with latent variables, z, which are typically drawn from simple prior distributions such as uniform or Gaussian distributions. These variables serve as compressed representations or "seeds" for data generation. The PQC encodes these variables into quantum states $|\psi(z)\rangle$, applies a sequence of quantum gates parameterized by θ, and ultimately measures the circuit to produce output samples. This process captures the stochastic nature of the data generation mechanism, enabling the generator to model complex data distributions effectively. The discriminator in a QGAN remains a classical neural network, tasked with evaluating the similarity between generated samples and real data. The hybrid architecture of QGANs, which integrates a quantum generator with a classical discriminator, provides both practical and performance advantages. One key benefit is enhanced scalability, as the quantum generator excels in handling high-dimensional data, while the classical discriminator leverages mature optimization techniques for efficient evaluation. Additionally, using a classical discriminator significantly reduces the quantum resource requirements, making QGANs more accessible and feasible for current and near-term quantum hardware. This blend of quantum and classical components positions hybrid QGANs as a promising and practical solution for leveraging quantum technologies in sophisticated data modeling tasks.

III. Related work

Recent advancements in machine learning, particularly GANs, have opened new avenues for enhancing NIDS. GANs are widely used in unsupervised learning tasks and are capable of generating synthetic data that closely resembles the original dataset. Their integration into intrusion detection has shown promise in improving feature extraction, augmenting datasets, and addressing limitations in both signature-based and anomaly-based detection systems.

One of the primary applications of GANs in NIDS is the generation of synthetic data to enhance training processes. The authors of [13] introduced an ANN-based GAN to generate synthetic samples for training an IDS on the NSL-KDD dataset. Their results demonstrated that an IDS incorporating GAN-generated data significantly outperformed standalone systems in attack detection. Similarly, in the work [14], GANs are employed to mitigate the class imbalance problem in NIDS datasets, finding GANs to be more effective than traditional oversampling techniques.

GANs have also been directly applied as detection mechanisms. The authors of [15] proposed a bidirectional GAN-based framework for anomaly detection, evaluated using the KDDCUP-99 dataset [16], and demonstrated its superior performance compared to other deep learning models. Truong et al. [17] advanced this approach by utilizing two GAN models with custom neural network architectures for generators and discriminators. Their experiments on CIC-IDS 2017 and UNSW-NB15 datasets [18] showcased the efficacy of GAN-based systems against traditional unsupervised methods. Furthermore, GANs offer a unique advantage in generating diverse attack scenarios, addressing the inherent limitations of traditional datasets and boosting system resilience against emerging threats. Building upon GAN advancements, Quantum Generative Adversarial Networks (QGANs) [19], [20] have emerged as a promising extension, particularly for image generation. Recent works illustrate their potential to outperform classical counterparts even with fewer trainable parameters.

Despite the promising advances of QGANs in image generation, their application to anomaly detection remains relatively underexplored, with only a few studies addressing this intersection. In [21] it is demonstrated the use of QGANs for detecting specific physical particles. Their approach involved an initial phase where the quantum generator learns the data distribution, followed by training a quantum discriminator,

implemented as a parameterized quantum circuit, to identify anomalies. They tested their QGAN on up to eight features, achieving 0.88 accuracy on noiseless simulators for seven features, outperforming classical GANs. They also validated its feasibility on actual quantum hardware for three features. However, the results were dataset-specific, limiting applicability across domains.

The authors of [12] introduced a novel high-dimensional encoding approach called Successive Data Injection (SuDaI), which expands encoded data size without increasing qubits. They utilized a state-fidelity QGAN with both generator and discriminator implemented as quantum circuits. The design incorporated the SWAP test to compare generated and real data representations. Using the Numenta Anomaly Benchmark (NAB) database, they analyzed temporal anomaly detection but did not report specific performance metrics, leaving its practical effectiveness unclear. The work [22] proposed a QGAN-based IDS but highlighted challenges due to the quantum generator producing discrete output values. For each feature represented by n qubits, the results were limited to 2^n discrete values, potentially reducing fidelity in real-world applications. Their work lacked specific implementation details, reproducibility, extensive testing, and performance metrics such as accuracy.

In this work [8] the authors presented a hybrid quantum–classical anomaly detection model using Variational Quantum-Classical Wasserstein GANs (WGANs) with gradient penalty. Their method extended the AnoGAN framework [23] for anomaly detection, integrating recent GAN advancements with hybrid quantum–classical neural networks trained via variational algorithms. Testing on a credit card fraud dataset, they evaluated performance with an anomaly score combining generator and discriminator losses, achieving an F1 score of 0.85. While their results were promising, the dataset's simplicity limited the validity of the results in more complex and diverse real-world scenarios.

Regarding the distributed approach, there are some works like [24], that applied a federated learning algorithm to a classical GAN, proving the convergence of their model. However, in the state of the art, only two preprint articles have addressed Federated Quantum GAN. The work [25] implements a federated version of QGAN, but on a different problem, not intrusion detection, and does not report performance metrics such as accuracy or execution time. The study [26] proposes an unsupervised federated quantum generative adversarial network (UF-QGAN) to address optimization challenges in wireless communications. This study employs a distributed scheme where the generator is in the cloud and multiple discriminators operate at the edge. However, this approach has several limitations as it requires the transmission of quantum states, because it operates in a fully quantum setting. Therefore, it focuses on a different domain, not performing classification and not reporting accuracy.

At the present time, it is not useful to compare the performance of this work based on Quantum ML with other works based on federated classical ML. For example, work [27] already achieves 98-99% accuracy on the same dataset. The goal is to improve existing QML models so that they can be used one day when the hardware becomes sufficiently mature. This work is motivated by the need to thoroughly evaluate Federated QGANs in the intrusion detection domain, focusing on their scalability, robustness, and performance. Existing studies have not explored intrusion detection and are either based on different problem settings or lack key performance metrics such as accuracy. This study aims to address these gaps by evaluating Federated QGANs on high-dimensional intrusion detection datasets and providing a comprehensive analysis of their effectiveness through key metrics such as accuracy and F1-score.

IV. Methodology

A. Data preprocessing

The NSL-KDD dataset is a well-established benchmark in the field of intrusion detection. It consists of a combination of numerical and categorical features, along with target labels that classify each instance as either normal or anomalous network activity. To prepare the data for the QGAN framework, which requires normalized and compact input, the feature values are scaled to a [0, 1] range using Min-Max normalization. Additionally, to further enhance the dataset's suitability for QGAN processing, Principal Component Analysis (PCA) is employed. PCA reduces the dimensionality of the feature space by transforming it into a lower-dimensional latent representation, preserving only the components that account for the majority of the data's variance. The number of principal components is set to 4, as this configuration achieves strong performance while maintaining manageable computational complexity.

B. Federated QGAN configuration

The single QGAN model proposed is a specialized adaptation tailored for intrusion detection, featuring a custom interpret function and a carefully designed quantum circuit. As illustrated in Figure 1, the model comprises two core components: a quantum generator implemented as a VQC and a classical discriminator built as a neural network.

The process begins with a random noise vector fed into the generator, introducing variability to produce diverse synthetic samples. The generator operates on n-qubits, starting from the initial state $|0\rangle$. It applies parameterized quantum gates, defined by weights θ_G, to generate a quantum state:

$$G(\theta_G)|0\rangle^{\otimes n},$$

which encodes the generated data in a Hilbert space of dimension 2^n. Since the quantum state cannot be directly evaluated, an interpret function maps it to classical data.

The discriminator, a classical neural network, receives input data, either from the generator or the real dataset, and outputs a probability score indicating whether the input is benign or an attack. The generator and discriminator are trained adversarially: the discriminator $D(\theta_D)$ learns to distinguish between real and synthetic data, while the generator $G(\theta_G)$ improves its ability to produce realistic data, minimizing

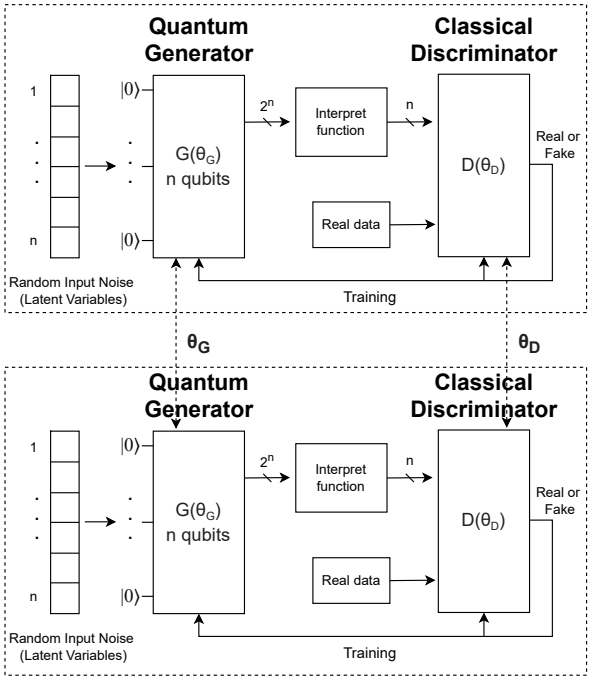

Fig. 1. Federated QGAN training process

Fig. 2. Quantum circuit for the generator

the discrepancy between real and synthetic samples. This iterative process enhances the generator's capability to create increasingly authentic data.

The generator, implemented as a variational quantum circuit, is designed to generate synthetic data from a noise vector. Figure 2 illustrates one of the several configurations tested. Each qubit in the circuit represents a dataset feature. The circuit begins by applying Hadamard gates to all qubits, placing them in a superposition state to ensure a uniform distribution over possible states. A feature map, such as the ZZFeatureMap, is then applied to encode latent variables into the quantum state, introducing entanglement and interdependence among qubits. The trainable portion of the circuit, or ansatz, employs the EfficientSU2 architecture, which combines single-qubit rotation gates and two-qubit entangling gates. By repeating these layers, the circuit gains expressive power, enabling it to approximate complex probability distributions. The number of repetitions is a hyperparameter that balances circuit depth and representational capacity.

Trainable parameters associated with the rotation gates in the ansatz, along with noise injected via rotation gates, allow the generator to sample from a broad range of data distributions. The generator's gradients are computed using the Adam optimizer. The interpret function transforms the generator's output, measurements of all possible n-qubit combinations, into a lower-dimensional representation of size n. Each of the 2^n values represents the probability of observing a specific qubit state combination. For each qubit i, the interpret function calculates its marginal probability by summing over all measurement outcomes where qubit i is in state 1. Let $P(x)$ denote the probability of observing an n-qubit state $x = (x_1, x_2, \ldots, x_n)$, where $x_i \in \{0, 1\}$. The marginal probability p_i for qubit i is computed as:

$$p_i = \sum_{x\,:\,x_i=1} P(x),$$

where the summation includes all 2^n possible states x with the i-th qubit in state 1. The resulting vector $(p_1, p_2, \ldots, p_n)$, of dimension n, represents the marginal probabilities of each qubit being in state 1. Each p_i lies in the range $[0, 1]$, making it directly comparable to the features in the real dataset.

The discriminator is a classical feedforward neural network with dense hidden layers and LeakyReLU activations, which introduce non-linearity while mitigating vanishing gradients. The output layer uses a sigmoid activation to produce a probability score indicating whether the input is real. Trained with Binary Cross-Entropy loss, the discriminator iteratively improves through adversarial feedback with the generator.

To further enhance the scalability and robustness of the QGAN framework, we implemented a federated version of the model, Figure 1. This approach involves running multiple independent instances of the QGAN across distributed nodes, each processing its own subset of the data and using different seeds for each quantum generator as QML is shown to be very sensitive to seeding choices. After each training epoch, the weights of both the generators and discriminators from all instances are aggregated separately. Specifically, each QGAN instance begins the next epoch with the average of all the weights from the previous iteration, ensuring a collaborative and synchronized learning process across the network.

This federated setup offers several advantages. First, it enables decentralized training, allowing multiple nodes to contribute to the model's learning without sharing raw data, thus preserving privacy and security. Second, by averaging the weights, the model benefits from the collective knowledge of all instances, leading to more generalized and robust performance. This is particularly important in intrusion detection, where the diversity of network behaviors across different environments can be vast.

The following section discusses the various configurations tested and their corresponding results. The experiments are replicable, and the source code is available on [28].

V. Results and Discussion

To assess the effectiveness of the proposed Federated QGAN model, a comprehensive analysis of various hyper-

parameters has been conducted. The evaluation process aims to determine the optimal settings that balance performance, stability, and computational efficiency, ensuring that the model achieves reliable results while remaining feasible for real-world applications.

All experiments were executed over a span of 80 epochs to guarantee sufficient convergence of the quantum generator and discriminator. This choice was made based on empirical observations that indicated a stabilization of loss values and performance metrics around this number of iterations. To encode noise within the quantum generator, a ZZFeatureMap was employed. This mapping technique plays a crucial role in allowing the quantum generator to effectively learn and represent complex data distributions, thereby enhancing the overall expressiveness of the model.

The selection of an appropriate quantum ansatz significantly impacts the expressiveness and efficiency of the quantum generator. Two widely used ansatzes, EfficientSU2 and RealAmplitudes, were compared across multiple configurations. While both ansatzes demonstrated competence in capturing the underlying data distribution, EfficientSU2 exhibited marginally superior performance in most experimental settings. This slight advantage likely stems from its enhanced ability to encode correlations within the quantum state space, making it a preferable choice for the Federated QGAN model.

A key aspect influencing the model's performance is the number of ansatz repetitions. A clear trend shows that increasing the number of reps generally leads to incremental improvements in performance metrics, but this improvement comes at a significant computational cost, as the training time increases non-linearly with additional repetitions. A balance between performance and efficiency was observed at three ansatz repetitions.

One of the most critical hyperparameters in training the QGAN model is the learning rate, which manage how quickly the generator and discriminator adapt their weights. Initial trials with a learning rate of 0.01 led to instability, as the steep learning curve caused erratic updates, preventing smooth convergence. Refining the learning rate to 0.001 yielded more stable training dynamics, although the discriminator adapted faster than the generator, resulting in an imbalance in adversarial competition.

Through further fine-tuning, an optimal configuration was established at 0.001 and 0.0006, ensuring that both the generator and discriminator evolved at comparable rates. This balance was crucial for maintaining the adversarial nature of the training process, where neither component dominated the other prematurely, ultimately leading to more robust and realistic generated data.

To evaluate the effectiveness of the federated learning approach, a network of four QGANs was implemented. Each QGAN was trained independently on different subsets of the data, with their insights aggregated to improve generalization and robustness. The federated setup facilitated parallelization, distributing the workload among multiple quantum generators while maintaining comparable performance levels to a single QGAN trained on the entire dataset. This approach not only accelerated the training process but also ensured greater model adaptability across diverse data distributions.

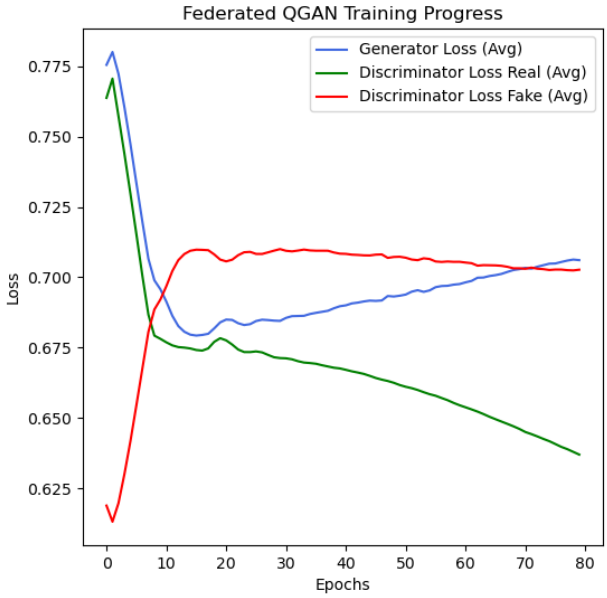

Fig. 3. Generator and Discriminator loss during the training process

The training process of the Federated QGAN is illustrated in Figure 3, which depicts the average loss curves of the generator and discriminator across all federated instances. Initially, the generator's loss starts at a high value, reflecting its poor initial capability to produce realistic samples. However, as training progresses, the generator's loss gradually decreases, indicating improved data synthesis. The discriminator's loss for fake data follows an inverse trajectory, signifying its increasing difficulty in distinguishing generated samples from real data.

Additionally, the loss associated with real data classification exhibits a steady decline, showcasing the discriminator's growing proficiency in correctly classifying genuine samples. By the end of training, both the generator and discriminator achieve a stable adversarial equilibrium, with neither model overpowering the other, an essential condition for generating high-quality synthetic data.

The final evaluation of the model's predictive performance revealed an accuracy of 0.9125 and an F1-score of 0.9034, confirming the efficacy of the Federated QGAN in learning and generating realistic data representations.

One of the most significant advantages of the federated approach is its impact on computational efficiency. While training a single QGAN for the full dataset required approximately 24 hours, the federated setup achieved equivalent performance in just 6 hours, reducing execution time to a quarter of the standalone model's duration. This acceleration is attributed to the parallelization of QGAN training, which enables multiple

instances to simultaneously process and refine their respective datasets before aggregation. Importantly, despite this drastic reduction in training time, the model maintained comparable predictive performance, highlighting the practical viability of the federated learning paradigm in quantum machine learning applications.

VI. Conclusion and future work

This study explored the effectiveness of a Federated QGAN model for intrusion detection. The evaluation focused on balancing performance, stability, and computational efficiency to ensure the feasibility of deploying the model in real-world applications. The final model achieved promising predictive performance, with an accuracy of 0.9125 and an F1-score of 0.9034, while significantly reducing training time compared to a single QGAN setup.

Despite these positive outcomes, further investigations are necessary to explore more complex configurations and optimize the model's architecture. Future work will focus on evaluating alternative quantum feature maps, ansatz structures, and aggregation strategies to enhance both expressiveness and efficiency. Additionally, testing different numbers of federated clients and data partitioning strategies could provide insights into the scalability and robustness of the approach. Exploring hybrid quantum-classical models and incorporating reinforcement learning techniques may further improve training stability and adversarial dynamics.

Acknowledgements

This research was funded by the NGIsargasso project (Europe Horizon Grant No. 101092887), Open Call 4 FRQ-GAN4AD project.

References

[1] O. H. Abdulganiyu, T. Ait Tchakoucht, and Y. K. Saheed, "A systematic literature review for network intrusion detection system (ids)," *International journal of information security*, vol. 22, no. 5, pp. 1125–1162, 2023.

[2] S. H. Rafique, A. Abdallah, N. S. Musa, and T. Murugan, "Machine learning and deep learning techniques for internet of things network anomaly detection—current research trends," *Sensors*, vol. 24, no. 6, p. 1968, 2024.

[3] S. Muneer, U. Farooq, A. Athar, M. Ahsan Raza, T. M. Ghazal, and S. Sakib, "A critical review of artificial intelligence based approaches in intrusion detection: A comprehensive analysis," *Journal of Engineering*, vol. 2024, no. 1, p. 3909173, 2024.

[4] J. Gui, Z. Sun, Y. Wen, D. Tao, and J. Ye, "A review on generative adversarial networks: Algorithms, theory, and applications," *IEEE transactions on knowledge and data engineering*, vol. 35, no. 4, pp. 3313–3332, 2021.

[5] H. Yamasaki, N. Isogai, and M. Murao, "Advantage of quantum machine learning from general computational advantages," 2023. [Online]. Available: https://arxiv.org/abs/2312.03057

[6] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, "Challenges and opportunities in quantum machine learning," *Nature Computational Science*, vol. 2, no. 9, pp. 567–576, 2022.

[7] T. A. Ngo, T. Nguyen, and T. C. Thang, "A survey of recent advances in quantum generative adversarial networks," *Electronics*, vol. 12, no. 4, p. 856, 2023.

[8] D. Herr, B. Obert, and M. Rosenkranz, "Anomaly detection with variational quantum generative adversarial networks," *Quantum Science and Technology*, vol. 6, no. 4, p. 045004, Oct. 2021. [Online]. Available: https://iopscience.iop.org/article/10.1088/2058-9565/ac0d4d

[9] F. Cirillo and 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.

[10] J. Verbraeken, M. Wolting, J. Katzy, J. Kloppenburg, T. Verbelen, and J. S. Rellermeyer, "A survey on distributed machine learning," *Acm computing surveys (csur)*, vol. 53, no. 2, pp. 1–33, 2020.

[11] C. Zoufal, A. Lucchi, and S. Woerner, "Quantum generative adversarial networks for learning and loading random distributions," *npj Quantum Information*, vol. 5, no. 1, p. 103, 2019.

[12] B. Kalfon, S. Cherkaoui, J. Laprade, O. Ahmad, and S. Wang, "Successive data injection in conditional quantum GAN applied to time series anomaly detection," *IET Quantum Communication*, vol. 5, no. 3, pp. 269–281, Sep. 2024. [Online]. Available: https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/qtc2.12088

[13] M. H. Shahriar, N. I. Haque, M. A. Rahman, and M. Alonso, "G-ids: Generative adversarial networks assisted intrusion detection system," in *2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC)*, 2020, pp. 376–385.

[14] J. Lee and K. Park, "Gan-based imbalanced data intrusion detection system," *Personal and Ubiquitous Computing*, vol. 25, no. 1, pp. 121–128, 2021.

[15] R. Patil, R. Biradar, V. Ravi, P. Biradar, and U. Ghosh, "Network traffic anomaly detection using pca and bigan," *Internet Technology Letters*, vol. 5, no. 1, p. e235, 2022.

[16] M. Tavallaee, E. Bagheri, W. Lu, and A. A. Ghorbani, "A detailed analysis of the kdd cup 99 data set," in *2009 IEEE symposium on computational intelligence for security and defense applications*. Ieee, 2009, pp. 1–6.

[17] T. Truong-Huu, N. Dheenadhayalan, P. Pratim Kundu, V. Ramnath, J. Liao, S. G. Teo, and S. Praveen Kadiyala, "An empirical study on unsupervised network anomaly detection using generative adversarial networks," in *Proceedings of the 1st ACM Workshop on Security and Privacy on Artificial Intelligence*, 2020, pp. 20–29.

[18] N. Moustafa and J. Slay, "Unsw-nb15: a comprehensive data set for network intrusion detection systems (unsw-nb15 network data set)," in *2015 military communications and information systems conference (MilCIS)*. IEEE, 2015, pp. 1–6.

[19] S. Lloyd and C. Weedbrook, "Quantum generative adversarial learning," *Physical review letters*, vol. 121, no. 4, p. 040502, 2018.

[20] P.-L. Dallaire-Demers and N. Killoran, "Quantum generative adversarial networks," *Physical Review A*, vol. 98, no. 1, p. 012324, 2018.

[21] E. Bermot, C. Zoufal, M. Grossi, J. Schuhmacher, F. Tacchino, S. Vallecorsa, and I. Tavernelli, "Quantum Generative Adversarial Networks For Anomaly Detection In High Energy Physics," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 01, Sep. 2023, pp. 331–341. [Online]. Available: https://ieeexplore.ieee.org/document/10313887/?arnumber=10313887

[22] M. A. Rahman, H. Shahriar, V. Clincy, M. F. Hossain, and M. Rahman, "A Quantum Generative Adversarial Network-based Intrusion Detection System," in *2023 IEEE 47th Annual Computers, Software, and Applications Conference (COMPSAC)*. Torino, Italy: IEEE, Jun. 2023, pp. 1810–1815. [Online]. Available: https://ieeexplore.ieee.org/document/10197070/

[23] T. Schlegl, P. Seeböck, S. M. Waldstein, G. Langs, and U. Schmidt-Erfurth, "f-anogan: Fast unsupervised anomaly detection with generative adversarial networks," *Medical image analysis*, vol. 54, pp. 30–44, 2019.

[24] M. Rasouli, T. Sun, and R. Rajagopal, "Fedgan: Federated generative adversarial networks for distributed data," *arXiv preprint arXiv:2006.07228*, 2020.

[25] R. B. Gómez, C. O'Meara, G. Cortiana, C. B. Mendl, and J. Bernabé-Moreno, "Towards autoqml: A cloud-based automated circuit architecture search framework," in *2022 IEEE 19th International Conference on Software Architecture Companion (ICSA-C)*. IEEE, 2022, pp. 129–136.

[26] T. Jamaluddin, B. Narottama, and S. Y. Shin, "Unsupervised federated quantum gan for optimizing wireless communications," *Authorea Preprints*, 2023.

[27] P. K. Sarkar, H.-H. Nguyen, and D. M. Farid, "Flgan-ids: Intrusion detection using gans with federated learning," in *International Conference on Intelligent Systems and Data Science*. Springer, 2024, pp. 216–230.

[28] Franco Cirillo, "Fqgan code," https://github.com/francocirill/fqgan, 2025.