Federated Quantum Generative Adversarial Network for Intrusion Detection

Franco Cirillo & Christian Esposito

University of Salerno, Department of Computer Science

{fracirillo, esposito}@unisa.it

Introduction

As cyberattacks grow more complex, traditional methods struggle to detect new threats. Anomaly-based detection offers a promising alternative but **classical ML** approaches face **limitations** in scalability and performance, particularly with **high-dimensional or imbalanced data**.

Quantum Generative Adversarial Networks (QGANs) provide a novel solution by combining quantum computing with adversarial learning. In a QGAN, a quantum circuit acts as the generator, producing synthetic data samples, while a classical neural network serves as the discriminator, evaluating their authenticity. Through iterative training, the generator learns to mimic the underlying data distribution.

Federated learning can be integrated with quantum machine learning (QML) approaches to improve scalability, enhance privacy, reduce computational load on individual devices, and increase robustness against quantum noise through distributed collaboration.

Objective

While QGANs hold promise, their application to network intrusion detection remains largely unexplored, particularly in federated contexts. Most existing studies focus on different problem domains or lack critical evaluation metrics.

This work addresses these limitations by:

- Proposing a **Federated QGAN architecture** specifically designed **for intrusion detection**.
- We conducted **experiments across various configurations**, achieving good performance while reducing time complexity compared to the single-instance setup.

Methodology

We use the **NSL-KDD dataset** and apply **PCA** to reduce dimensionality, retaining 4 components to ensure compatibility with quantum circuits while preserving key variance. Each QGAN instance combines a Quantum Generator and a Classical Discriminator:

- The **generator** is a Variational Quantum Circuit (VQC) (see Figure 1), initialized with Hadamard gates, followed by a Feature Map and an ansatz. The circuit outputs a quantum state that is measured, and using our designed interpret function, marginal probabilities are extracted to form synthetic feature vectors.
- The **discriminator** is a neural network with dense layers, LeakyReLU activations, and sigmoid output, trained to distinguish real from generated data.

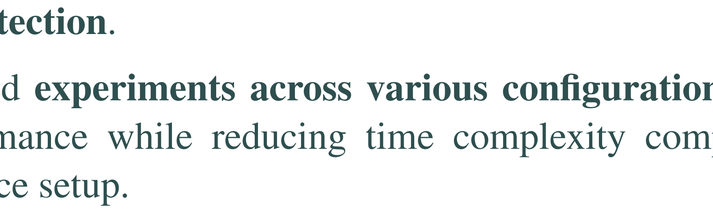

Figure 1: Quantum circuit for the generator

Both components are trained adversarially using binary cross-entropy loss and the Adam optimizer.

The **federated version** of the QGAN is shown in Figure 2. Multiple QGAN instances are trained in parallel across distributed nodes, each using different random seeds and local data. After each epoch, generator and discriminator weights are aggregated and averaged, then synchronized across nodes. This setup preserves data privacy, reduces the impact of quantum noise, and enables robust, decentralized intrusion detection.

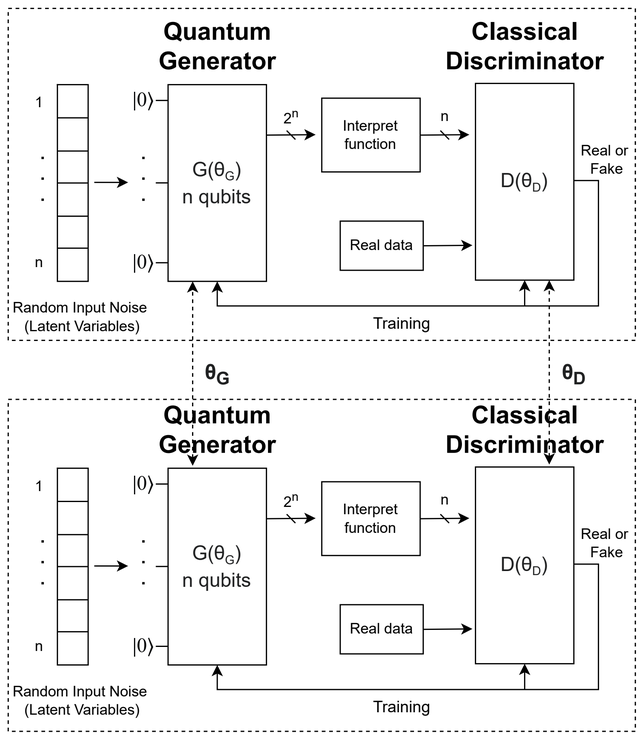

Figure 2: Federated QGAN training process

Results

We conducted extensive experiments focusing on generator fidelity, accuracy, and training efficiency.

Training was carried out over 80 epochs for each settings, and the best performance was obtained using the **EfficientSU2 ansatz** and a **ZZFeatureMap** to enhance the expressiveness of the quantum generator. Testing multiple configurations, we found that **three ansatz repetitions** provided the best trade-off between performance and computational cost. Fine-tuning the **learning rates** (0.0006 for the generator and 0.001 for the discriminator) ensured stable adversarial training and convergence.

As shown in Figure 3, the loss curves of the generator and discriminator demonstrate a well-balanced training process, converging to a stable equilibrium. The model achieved an accuracy of 0.9125 and an F1-score of 0.9034, confirming its ability to detect anomalies effectively.

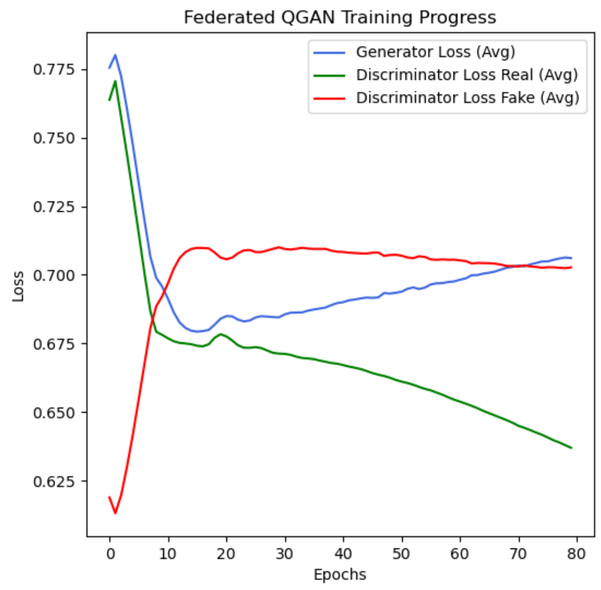

Figure 3: Generator and Discriminator loss during the training process

The federated setup reduced training time by n times, matching the number of devices used, without compromising performance, enabling scalable and privacy-preserving intrusion detection.

Conclusions

This work shows that Federated QGANs are effective and efficient for intrusion detection. Future research will focus on scaling to more clients, exploring new quantum ansatzes and feature maps, and integrating hybrid or reinforcement learning to improve training dynamics and adaptability.

Acknowledgements

This research was funded by the NGIsargasso project (Europe Horizon), Open Call 4 FRQGAN4AD project.

Code: https://github.com/francocirill/fqgan