

POLYTECHNIQUE MONTRÉAL

affiliée à l'Université de Montréal

**Quantum GAN Architectures With Successive Data Injection for Network
Anomaly Detection in Multivariate Time Series**

WAJDI EL HAMMAMI

Département de génie informatique et génie logiciel

Mémoire présenté en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
Génie informatique

Août 2025

POLYTECHNIQUE MONTRÉAL

affiliée à l'Université de Montréal

Ce mémoire intitulé :

**Quantum GAN Architectures With Successive Data Injection for Network
Anomaly Detection in Multivariate Time Series**

présenté par **Wajdi EL HAMMAMI**

en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
a été dûment accepté par le jury d'examen constitué de :

Alejandro QUINTERO, président

Soumaya CHERKAoui, membre et directrice de recherche

Heng LI, membre

DEDICATION

*I dedicate this work to my beloved family,
whose unwavering support, encouragement, and sacrifices
have been the foundation of all my achievements.
To my parents, for their unconditional love and constant guidance.
To my mentors, for inspiring me to reach higher.
This thesis is a reflection of your belief in me.
Thank you.*

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my supervisor, **Prof. Soumaya Cherkaoui**, for her guidance, unwavering support, and constructive feedback throughout this research. Her expertise and encouragement were instrumental in shaping both the direction and quality of this work.

I would also like to thank my co-authors and collaborators, **Prof. Wang Shengrui**, **Jean-Frédéric Laprade**, and **Ola Ahmad** for their valuable insights, technical discussions, and continuous contributions to our joint research efforts.

My sincere thanks go to my peers at the **Lincs Lab**, whose feedback, camaraderie, and collaborative spirit made the lab environment intellectually stimulating and enjoyable throughout this journey.

Finally, I would like to extend my heartfelt thanks to my family and friends for their constant encouragement, patience, and emotional support. This thesis would not have been possible without your belief in me.

RÉSUMÉ

La détection d'anomalies dans des séries temporelles multivariées représente un défi fondamental dans des domaines critiques tels que la cybersécurité réseau et les systèmes industriels. Face aux limites des approches classiques, ce mémoire propose une série d'architectures basées sur des réseaux antagonistes génératifs quantiques (Quantum GANs), intégrant des circuits quantiques variationnels (VQCs) et une technique d'injection successive des données (Successive Data Injection – SuDaI).

Deux contributions principales sont présentées : la première applique un QGAN conditionnel à la détection d'anomalies dans le trafic réseau à partir du jeu de données CIC-CIDDS-2018, tandis que la seconde exploite un modèle hybride Quantum-GRU avec un critère adversarial de Wasserstein pour identifier des anomalies dans les systèmes de contrôle industriels à trafic réseau (jeu de données HAI). Ces approches reposent sur des encodages quantiques compacts, ainsi que sur l'exploitation du générateur et du discriminateur pour la détection d'anomalies. Dans le second modèle, un mécanisme de score composite est introduit, combinant l'incertitude prédictive, les erreurs de reconstruction et le retour du critique Wasserstein.

Les résultats expérimentaux montrent que les modèles quantiques proposés surpassent leurs homologues classiques en termes de précision et de rappel, tout en utilisant un nombre réduit de paramètres. Le score TaF1 obtenu dépasse 89% dans le cas des systèmes industriels, tandis que le modèle réseau atteint une F1-score de 0.98 avec un MSE de 0.0013.

Ces travaux démontrent le potentiel des GANs quantiques pour la modélisation prédictive des séries temporelles complexes issues du trafic réseau et ouvrent la voie à de futures applications hybrides sur des dispositifs NISQ (Noisy Intermediate-Scale Quantum).

ABSTRACT

This thesis investigates the use of quantum generative adversarial networks (QGANs) for anomaly detection in multivariate time-series data, with applications in network traffic monitoring and industrial control systems. Motivated by the increasing complexity of time-series data and the potential of quantum computing to enhance learning models, we propose two quantum-enhanced architectures that leverage variational quantum circuits (VQCs) and successive data injection (SuDaI) for expressive and efficient learning.

The first model is a conditional QGAN applied to a network security dataset, where both the generator and discriminator are quantum circuits trained on time-windowed data. The second model introduces a hybrid Quantum-GRU architecture combined with adversarial training and probabilistic forecasting. This second approach integrates a composite anomaly scoring mechanism that combines predictive uncertainty, reconstruction error, and Wasserstein critic feedback.

Both models are evaluated on benchmark datasets under realistic settings. The first model achieves a high F1-score with low mean squared error, while the second model demonstrates strong time-series aware performance, exceeding 89% TaF1. These results illustrate the advantages of hybrid quantum-classical models in detecting subtle or sustained anomalies with fewer parameters and improved temporal awareness.

The contributions of this work pave the way for scalable anomaly detection frameworks designed to run on current and near-term quantum hardware.

TABLE OF CONTENTS

DEDICATION	iii
ACKNOWLEDGEMENTS	iv
RÉSUMÉ	v
ABSTRACT	vi
LIST OF TABLES	ix
LIST OF FIGURES	x
LIST OF SYMBOLS AND ACRONYMS	xi
CHAPTER 1 INTRODUCTION	1
1.1 Basic Concepts and Definitions	2
1.2 Problem Statement	2
1.3 Research Objectives	3
1.4 Approach Overview	3
1.5 Thesis Structure	4
CHAPTER 2 LITERATURE REVIEW	5
2.1 General Context	5
2.2 Network Anomaly Detection	5
2.3 The Role of Generative Models	6
2.4 Quantum Computing	7
2.5 Quantum Approaches to Time-Series Anomaly Detection	8
2.6 Hybrid Quantum-Sequential Modeling	9
2.7 Positioning of Our Work	10
CHAPTER 3 ARTICLE 1: ENHANCING NETWORK ANOMALY DETECTION WITH QUANTUM GANS AND SUCCESSIVE DATA INJECTION FOR MUL- TIVARIATE TIME SERIES	11
3.1 Introduction	11
3.2 Methodology	13
3.2.1 Data Preparation and Preliminary Analysis	13
3.2.2 Model Architecture	14

3.2.3	Training settings and process	15
3.2.4	Anomaly Detection	16
3.3	Results & Discussion	17
3.4	Conclusion & Future Work	20
CHAPTER 4 ARTICLE 2: QUANTUM GATED RECURRENT GAN WITH GAUS-		
SIAN UNCERTAINTY FOR NETWORK ANOMALY DETECTION		22
4.1	Introduction	22
4.2	Proposed Method	24
4.2.1	Problem Formulation	24
4.2.2	Data Preparation	24
4.2.3	Model Overview	26
4.2.4	Training Objectives	29
4.2.5	Anomaly Detection	31
4.3	Results	33
4.3.1	Experimental Setup	33
4.3.2	Evaluation Metrics	35
4.3.3	Results	36
4.3.4	Ablation Study	37
4.4	Conclusion	38
CHAPTER 5 CONCLUSION		40
5.1	Summary of Works	40
5.2	Limitations	41
5.3	Future Research	41
REFERENCES		43

LIST OF TABLES

Table 3.1	Comparison of different detection methods of the model	19
Table 3.2	Evaluation results of deep learning models for DDoS detection	20
Table 4.1	Hyperparameters Used in Our Model	35
Table 4.2	Time-Aware F1 Metrics on the HAI Dataset	37
Table 4.3	Ablation Study Results on the HAI Dataset	38

LIST OF FIGURES

Figure 3.1	Discriminator circuit	14
Figure 3.2	QGAN results for synthetic time series.	18
Figure 3.3	Classic GAN results for synthetic time series.	18
Figure 3.4	Showcasing the similar patterns between the time series of attacks and the time series of anomaly scores.	20
Figure 4.1	Top 16 Feature Importances ranked by Gini criterion.	26
Figure 4.2	Architecture of the Hybrid Quantum Layer (HQL).	28
Figure 4.3	Model overview	29
Figure 4.4	Scatter plot of scaled mean log-variance, with ground-truth anomalies highlighted by red square outlines.	38

LIST OF SYMBOLS AND ACRONYMS

AEs	Autoencoders
CIC	Canadian Institute for Cybersecurity
CNN	Convolutional Neural Network
CSE	Communications Security Establishment
DDoS	Distributed Denial-of-Service
DNN	Deep Neural Network
eTaP	enhanced Time-aware Precision
eTaR	enhanced Time-aware Recall
GANs	Generative Adversarial Networks
GRU	Gated Recurrent Unit (also referred to as GRUs)
HAI	HIL-based Augmented ICS (also referred to as HAI dataset)
HQLs	Hybrid Quantum Layers
ICS	Industrial Control System (also referred to as industrial control systems)
IDS	Intrusion Detection Systems
IETF	Internet Engineering Task Force
IT	Information Technology
LSTM	Long Short-Term Memory
MSE	Mean Squared Error
NISQ	Noisy Intermediate-Scale Quantum
OSI	Open Systems Interconnection
OT	Operational Technology
PLC	Programmable Logic Controller (also referred to as PLCs)
QGAN	Quantum GAN
QGANs	Quantum Generative Adversarial Networks
QGRUs	Quantum GRUs
QML	Quantum Machine Learning
QWGAN	Quantum Wasserstein Generative Adversarial Network
RNNs	Recurrent Neural Networks
SuDaI	Successive Data Injection
TaF1	Time-series aware F1 score
VAEs	Variational Autoencoders
VQCs	Variational Quantum Circuits

CHAPTER 1 INTRODUCTION

Quantum computing is emerging as a powerful paradigm that promises to redefine how we model, learn from, and interact with high-dimensional data. Unlike classical computation, which operates on binary states, quantum systems exploit principles such as superposition, entanglement, and interference, enabling the representation of complex patterns in exponentially larger spaces. This opens new opportunities for machine learning, particularly in scenarios where data complexity, dimensionality, and uncertainty present significant challenges—especially in network systems, where the increasing scale leads to a high number of variables that must be accounted for.

In recent years, the field of quantum machine learning (QML) has gained momentum, aiming to integrate quantum circuits into learning models to enhance their expressiveness and generalization capabilities. Among these models, quantum generative adversarial networks (QGANs) have shown promise in learning complex data distributions with fewer parameters by leveraging variational quantum circuits (VQCs). However, deploying such models on real-world data — especially multivariate network time-series — introduces unique challenges due to the sequential nature of the input and the limitations of near-term quantum hardware.

This thesis investigates how Quantum Generative Adversarial Network (QGAN) architectures can be adapted and extended to effectively detect anomalies in temporal network traffic data streams. To address the challenge of encoding sequential data within quantum circuits, we explore the use of an existing technique known as *Successive Data Injection* (SuDaI). Although not originally proposed in this work, SuDaI is evaluated and recontextualized in two distinct scenarios.

In the first study, SuDaI is integrated into a conditional QGAN to enhance quantum encoding for network traffic anomaly detection. In the second, SuDaI is integrated into a hybrid Quantum-GRU model, enabling the processing of multivariate industrial network time-series data under the constraints of NISQ (Noisy Intermediate-Scale Quantum) devices. This includes a tailored preprocessing pipeline and circuit design to preserve temporal dependencies across sliding windows.

By developing and evaluating these two architectures on benchmark datasets, this work aims to demonstrate the potential of quantum-enhanced generative models in achieving more efficient, robust, and temporally-aware network anomaly detection, paving the way for practical hybrid solutions in quantum machine learning.

1.1 Basic Concepts and Definitions

- **Multivariate Time Series:** A sequence of timestamped observations where multiple variables are recorded simultaneously over time.
- **Anomaly:** A data point or interval that significantly deviates from expected normal behavior.
- **Generative Adversarial Networks (GANs):** Generative models composed of a generator and a discriminator trained in a minimax setting.
- **Quantum Computing:** A computational paradigm that uses qubits and quantum logic to exploit superposition and entanglement.
- **Variational Quantum Circuit (VQC):** A quantum circuit with trainable parameters, optimized in a classical feedback loop.
- **Successive Data Injection (SuDaI):** A strategy originally designed to capture temporal context by progressively injecting time-dependent data into a quantum circuit across multiple layers or timesteps.

1.2 Problem Statement

Modeling long-term temporal dependencies, adapting to nonstationary data, and ensuring reliable performance under limited supervision are persistent challenges in time-series network anomaly detection. These difficulties are further compounded in quantum machine learning, where current devices are constrained by circuit depth, noise, and limited qubit availability.

In this context, the central research question is:

- How can we design a quantum model that captures complex temporal structures in multivariate time series?
- How can we leverage quantum representations while respecting the limitations of near-term quantum hardware?
- How can these models be used to detect meaningful and persistent anomalies in real-world data?

1.3 Research Objectives

The objectives of this thesis are as follows:

- Propose two quantum GAN-based architectures for anomaly detection in multivariate time-series data.
- Implement SuDaI both as a temporal encoding strategy to capture sequential dependencies and as a multidimensional injection mechanism to efficiently encode multiple features under quantum hardware constraints.
- Evaluate the proposed models on two benchmark datasets: **CIC-CIDDS-2018**, representing real-world network traffic, and **HAI**, a multivariate time-series dataset from industrial control systems with networked sensor and actuator data.
- Compare performance against classical baselines in terms of precision, recall, robustness, and efficiency, with improved metrics demonstrated in the second article.

1.4 Approach Overview

To address the challenges of detecting anomalies in multivariate time-series network data—particularly in network and industrial settings—we adopted a hybrid quantum-classical modeling approach grounded in recent advances in quantum machine learning. Our research explores the integration of **Generative Adversarial Networks (GANs)** with **Variational Quantum Circuits (VQCs)** to design robust, low-parameter models capable of capturing complex temporal dependencies and subtle deviations from normal traffic patterns.

In the first phase of our work, we developed a **Quantum GAN (QGAN)** framework that employs **Successive Data Injection (SuDaI)** and **data re-uploading** techniques. These methods allow for the effective encoding of multivariate time-series into limited quantum hardware by sequentially injecting input features into the quantum circuit across multiple layers. The model combines the outputs of the generator and the discriminator into a unified *anomaly score*, leveraging the strengths of both predictive reconstruction and adversarial discrimination.

In the second phase, we extended this work by incorporating temporal modeling and predictive uncertainty through a novel **Quantum Gated Recurrent Unit (QGRU)** architecture. Each gating mechanism of the GRU is implemented using **Hybrid Quantum Layers (HQLs)**, composed of a classical-quantum-classical pipeline. The generator outputs both

the *mean* and *variance* of the predicted next time-step, allowing the model to quantify its confidence in normal behavior. To detect anomalies, we design a gated scoring mechanism that first filters candidates based on **Gaussian interval violations** and then applies **critic feedback** and **reconstruction error** for refined assessment. This layered strategy improves robustness and reduces false positives.

This two-part approach enables the exploration of both architectural efficiency and temporal uncertainty in quantum-enhanced anomaly detection.

1.5 Thesis Structure

This thesis is structured around two core contributions:

- **Chapter 1** – Introduction, background concepts, problem formulation, and objectives.
- **Chapter 2** – Literature review: foundational methods in time-series network anomaly detection, GAN-based models, and recent advances in quantum GANs and hybrid architectures.
- **Chapter 3** – First article: Enhancing Network Anomaly Detection With Quantum GANs And Successive Data Injection For Multivariate Time Series.
- **Chapter 4** – Second article: Quantum Gated Recurrent GAN with Gaussian Uncertainty for Network Anomaly Detection.
- **Chapter 5** – General discussion and conclusion, including future research directions.

CHAPTER 2 LITERATURE REVIEW

2.1 General Context

Anomaly detection in multivariate time-series data plays a critical role in many application domains, with a particular emphasis on **network security**, where it is essential for identifying cyber threats and ensuring the resilience of digital infrastructures. These anomalies—unexpected deviations from normal patterns—can signal malicious activities, equipment failures, or systemic errors. Due to the sequential nature of time-series data, detecting anomalies is more challenging than in static datasets. Models must account not only for the values themselves but also for their temporal dependencies and context [1, 2].

Traditional rule-based or statistical approaches often fall short in high-dimensional, non-linear, and dynamic environments. In response, machine learning has emerged as a powerful alternative, capable of learning complex normal patterns from data. In particular, deep learning-based models have shown great promise in unsupervised anomaly detection, where only normal data is used during training [3, 4]. This setting is particularly relevant for real-world systems, where annotated anomalies are rare or unavailable.

Given the increasing volume and complexity of data in modern network systems, anomaly detection must strike a balance between *model expressiveness*, *computational efficiency*, and *adaptability* to evolving patterns. This challenge motivates the exploration of new modeling paradigms—such as quantum computing—that can offer enhanced representation capacity with fewer resources [5, 6].

2.2 Network Anomaly Detection

As digital infrastructures grow in complexity and scale, ensuring their security has become a major concern. Anomaly detection in network traffic plays a critical role in identifying malicious activities that do not conform to expected behavior, such as Distributed Denial-of-Service (DDoS) attacks, brute-force intrusions, or port scans. Traditional Intrusion Detection Systems (IDS) have primarily relied on rule-based or signature-based approaches, such as Snort and Bro, which are effective for known threats but fail to generalize to novel or evolving attack patterns.

To overcome these limitations, machine learning techniques have gained traction for their ability to detect anomalies in a data-driven, adaptive manner. These systems analyze various

traffic-related features, such as flow durations, packet inter-arrival times, byte volumes, and header fields, transforming raw traffic data into multivariate time series suitable for sequence modeling [7, 8]. However, challenges remain, including class imbalance, data scarcity for certain attack types, and the difficulty of setting appropriate detection thresholds in dynamic environments.

Several benchmark datasets have been proposed to evaluate the performance of anomaly detection models in network settings. The **CIC-CIDDS2018** dataset [9], for instance, includes labeled traffic with attack types such as brute-force, infiltration, and DDoS. It provides a realistic simulation of network behavior, capturing over 80 features per flow and enabling supervised and unsupervised detection tasks. Similarly, **NSL-KDD** [10] offers a cleaned and balanced version of the older KDD’99 dataset, frequently used for evaluating classical classifiers, although its features are relatively outdated for modern network topologies.

An interesting addition to the landscape of network-oriented datasets is the **HAI (HIL-based Augmented ICS)** [11] dataset, which exhibits key characteristics relevant to network anomaly detection. The HAI dataset contains time-series data from a simulated cyber-physical environment, including communication between networked sensors, actuators, and programmable logic controllers (PLCs). It captures normal and abnormal behaviors under staged attack scenarios, such as false command injection and control signal manipulation.

The inclusion of HAI in network anomaly detection studies is motivated by its hybrid nature. It features structured, protocol-based message exchanges between components over time, making it well-suited for sequence-based anomaly detection. Moreover, its realistic modeling of ICS operations allows researchers to study network anomalies in a setting that blends cyber and physical behaviors—an increasingly important consideration given the rise of Industry 4.0 and the convergence of IT and OT (Operational Technology) networks.

As anomaly detection models move toward more robust, generalizable, and explainable systems, datasets like CIC-CIDDS2018 and HAI offer rich, high-dimensional multivariate time series data that challenge conventional approaches. They also enable the development of hybrid models, such as those combining generative architectures with temporal modeling components, as explored in this thesis.

2.3 The Role of Generative Models

Generative models have become increasingly prominent in the domain of anomaly detection due to their ability to learn complex data distributions without relying on explicit labels. Among these, **Autoencoders (AEs)**, **Variational Autoencoders (VAEs)**, and **Gener-**

active Adversarial Networks (GANs) have shown substantial effectiveness, especially in time-series applications [12, 13].

Autoencoders attempt to reconstruct input data through a lower-dimensional latent space. When trained on normal samples, their reconstruction error on anomalous data typically rises, allowing for anomaly scoring. VAEs extend this principle by introducing probabilistic inference in the latent space, which adds robustness and enables uncertainty estimation.

GANs, first introduced by Goodfellow et al. [14], utilize a generator-discriminator architecture where the generator learns to produce synthetic data that fools the discriminator. In the context of anomaly detection, the discriminator can be trained to differentiate real (normal) data from either generated samples or anomalies [2, 3, 15].

Time-series specific GANs such as TadGAN and its variants have demonstrated that GANs can capture temporal patterns and generate realistic future sequences. However, these models often suffer from limitations such as unstable convergence, mode collapse, and difficulty in modeling long-term dependencies [16, 17]. Moreover, GANs traditionally lack an inherent measure of predictive uncertainty, which limits their ability to distinguish ambiguous deviations from true anomalies.

These limitations motivate the exploration of quantum-enhanced generative models, which can leverage the expressiveness of quantum circuits and potentially stabilize training in high-dimensional or low-sample scenarios.

2.4 Quantum Computing

Quantum computing has emerged as a disruptive paradigm in computational sciences, offering fundamentally new mechanisms for processing information based on the principles of quantum mechanics—namely, *superposition*, *entanglement*, and *quantum interference*. These properties enable quantum systems to encode and manipulate information in exponentially larger spaces compared to classical bits [5, 18].

In the context of machine learning, Quantum Machine Learning (QML) seeks to harness the advantages of quantum computation to enhance learning performance or resource efficiency. QML models can offer richer expressiveness and potentially faster training times, particularly in high-dimensional settings or when sample sizes are limited [6, 18].

Among the most prominent QML techniques are **Variational Quantum Circuits (VQCs)**, which combine classical optimizers with quantum circuits containing trainable parameters. These hybrid models can be used for classification, regression, and generative modeling. A key challenge in QML lies in encoding classical data into quantum states; techniques such

as angle encoding, basis encoding, amplitude encoding, and data re-uploading have been studied for this purpose [19, 20].

Recent work has also explored the use of quantum generative models such as **Quantum GANs (QGANs)**, which aim to leverage quantum circuits in either the generator, the discriminator, or both. These models have shown promise in early-stage benchmarks and image generation tasks [21, 22], and they represent a natural step toward enhancing generative learning in time-series data.

Nevertheless, current QML implementations are constrained by the limitations of Noisy Intermediate-Scale Quantum (NISQ) devices, including limited qubit counts, noise, and circuit depth constraints.

2.5 Quantum Approaches to Time-Series Anomaly Detection

While classical deep learning techniques have made significant strides in time-series anomaly detection, the application of quantum machine learning (QML) in this domain is still in its early stages. Quantum models such as Quantum Autoencoders, Quantum Support Vector Machines, and Quantum Classifiers have been explored for pattern recognition and dimensionality reduction, but only a limited number of works have extended these architectures to sequential or temporal settings [23, 24].

In the specific context of anomaly detection, recent efforts have investigated **Quantum GANs (QGANs)** for synthetic data generation and pattern modeling. Huang et al. [21] demonstrated an experimental implementation of a QGAN for image generation, showcasing the potential of quantum-enhanced generative models. However, these models generally focus on static data and do not explicitly incorporate temporal dependencies, which are crucial in time-series applications.

Furthermore, most QGAN frameworks do not integrate uncertainty modeling or probabilistic outputs, which are essential for distinguishing ambiguous fluctuations from true anomalies. The absence of uncertainty-aware mechanisms and the challenge of encoding multivariate sequences into quantum circuits without using a huge number of qubits remain major limitations.

To mitigate these issues, techniques such as **Successive Data Injection (SuDaI)** [25] and **data re-uploading** [20] have been proposed. **SuDaI** enables models to sequentially introduce features or time steps into shallow quantum circuits, effectively capturing temporal structure. Data re-uploading is a technique that enhances expressivity in quantum models, especially when operating under constraints on the number of available qubits.

Hybrid quantum-classical architectures have also been proposed to leverage the strengths of both paradigms. For example, Alhussein et al. [24] combined classical neural networks with quantum layers to enhance anomaly detection in industrial control systems. While promising, these works still face challenges related to noise robustness, lack of scalability, and limited benchmark evaluations.

In summary, while foundational components for quantum time-series modeling have emerged, the field still lacks cohesive architectures that combine sequential encoding, probabilistic reasoning, and robust anomaly scoring—gaps that our work directly addresses.

2.6 Hybrid Quantum-Sequential Modeling

A key limitation in early Quantum GAN (QGAN) designs for anomaly detection is their lack of explicit temporal modeling. Since time-series data involves complex dependencies over time, models that capture sequential patterns—such as Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, or Gated Recurrent Units (GRUs)—have proven essential in classical settings [2, 13].

To bridge this gap, hybrid architectures have been proposed that integrate classical recurrent models with quantum components. In these designs, classical recurrent layers handle temporal structure, while quantum circuits model nonlinear or high-dimensional relationships in the latent space. However, these early hybrid models often treat the quantum layers as post-processing units or feature augmenters, rather than embedding them directly into the core recurrence mechanism [26].

Recent advances have proposed embedding quantum logic within recurrent architectures. For instance, Quantum GRUs (QGRUs) replace the classical gates (update, reset, and candidate) with variational quantum circuits, allowing temporal memory to evolve under quantum transformations [27]. These models can exploit quantum entanglement and interference to encode richer dependencies between time steps.

To further enhance interpretability and robustness, some hybrid models incorporate probabilistic outputs—such as Gaussian parameters (mean and variance)—which allow for the estimation of prediction uncertainty. For example, Bayesian quantum-enhanced architectures output trainable Gaussian means and variances to quantify uncertainty in anomaly detection. This is particularly important in safety-critical applications where false positives or undetected anomalies can be costly [28].

In our work, we build on these developments by proposing a QGRU-based architecture integrated with variational quantum circuits and a Wasserstein adversarial framework. Our

model generates uncertainty-aware predictions and employs a gated anomaly detection mechanism to robustly identify abnormal behavior in multivariate network time-series data.

2.7 Positioning of Our Work

Despite recent advances in quantum approaches to time-series anomaly detection, several key limitations remain unresolved. Quantum models, while promising in theory, are generally limited to simplified datasets, static scenarios, or shallow circuits, due to hardware constraints. Furthermore, most existing QGAN implementations either ignore temporal dynamics or lack the ability to produce interpretable, uncertainty-aware outputs.

Our research addresses these gaps through a two-phase contribution:

- In **Chapter 3**, we introduce a novel **Quantum GAN (QGAN)** framework based on **Variational Quantum Circuits (VQCs)**, enhanced with **Successive Data Injection (SuDaI)** and **data re-uploading**. The model efficiently maps multivariate time-series data into limited quantum hardware by encoding features as rotation angles and alternating encoding with variational layers. Anomaly detection is performed using a combined score derived from both the generator (reconstruction error) and the discriminator, achieving superior performance over classical GAN baselines on synthetic and real datasets.
- In **Chapter 4**, we build on this foundation by proposing a **Quantum Wasserstein GAN** with a **Quantum Gated Recurrent Unit (QGRU)** architecture. Here, each gating mechanism in the GRU is implemented via Hybrid Quantum Layers (HQLs). The generator outputs a *mean* and *variance* for the next time-step, allowing for predictive uncertainty. A **gated anomaly detection mechanism** is introduced, which activates further anomaly scoring only when confidence bounds are violated. This strategy combines interval violation, top- k reconstruction error, and Wasserstein critic feedback to yield high recall and precision.

These two contributions reflect complementary directions: the first focuses on efficient quantum encoding and generation in shallow circuits; the second addresses temporal modeling, uncertainty, and robust scoring strategies. Together, they demonstrate that hybrid quantum models can not only match but sometimes outperform their classical counterparts with significantly fewer parameters, while offering interpretability and adaptability for real-world anomaly detection.

CHAPTER 3 ARTICLE 1: ENHANCING NETWORK ANOMALY DETECTION WITH QUANTUM GANS AND SUCCESSIVE DATA INJECTION FOR MULTIVARIATE TIME SERIES

Authors: Wajdi Hammami, Soumaya Cherkaoui, Shengrui Wang

Publication: Proceedings of the IEEE International Wireless Communications and Mobile
Computing Conference (IWCMC 2025)

Publication date: May 16, 2025

Contribution

I contributed to the model design, data processing (preprocessing, feature selection), implementation of quantum circuits, result analysis, as well as the full writing of the paper and the literature review.

3.1 Introduction

Anomaly detection plays a vital role across numerous domains by enabling continuous monitoring of system behavior and providing early warnings of potential issues, allowing for proactive intervention to address detected irregularities. This capability is particularly crucial in areas such as cybersecurity, where growing system complexity and surface attack has led to an increased risk of cybersecurity vulnerabilities in recent years. 6G networks will face an even broader attack surface—from massive IoT deployments and AI-orchestrated services to potential quantum-powered adversaries—demanding quantum-resistant encryption, zero-trust architectures, secure network slicing, and AI-driven threat analytics to safeguard next-generation connectivity [29–31].

In this work, we address the challenge of network anomaly detection by leveraging multivariate time series to effectively capture the temporal dynamics present in network data. To develop an efficient anomaly detection system, the approach presented in this work draws on two important fields that have shown considerable promise. The first is machine learning, which excels at identifying patterns and extracting insights. This has made it a key tool of anomaly detection, with models such as Generative Adversarial Networks (GANs) showing particular effectiveness for this task due to their powerful generative capabilities [32–37].

The second field is quantum computing, which has the potential to tackle computational

problems that are intractable for classical computers [38–42]. Quantum computing uses two fundamental principles: superposition and entanglement. Superposition allows qubits—the basic units of quantum information—to exist in multiple states simultaneously, increasing computational parallelism. Entanglement enables qubits to share correlations that can be leveraged to perform certain computations more efficiently than classical systems.

Integrating quantum computing with machine learning can potentially address challenges related to high computational overhead. Notably, some quantum machine learning approaches, especially those that employ generative models, have in some cases demonstrated advantages over classical methods, such as improved stability of the loss function during training [43–47].

Much work has been done on network anomaly detection using machine learning [33, 48]. For instance, the authors in [34] examine the opportunities and challenges associated with deep learning techniques for anomaly detection, while recent studies have explored the use of self-attention mechanisms [35, 49, 50].

Quantum Machine Learning (QML) has emerged as a promising field, aiming to harness quantum computing for machine learning tasks [42]. However, despite its potential, QML faces several challenges, including hardware limitations, noise sensitivity, and the difficulty of encoding high-dimensional classical data into quantum representations [39]. In the context of network anomaly detection, the use of QML remains relatively unexplored [40, 41, 51], particularly with quantum generative adversarial networks (QGANs), as only a few works examining their potential [40]. This highlights the gap in current research and underscores the novelty of this approach which applies a QGAN to detect anomalies in multivariate network time-series data.

In this work, we build upon existing approaches by designing both the generator and discriminator as quantum models. By leveraging successive data injection [40] jointly with data re-uploading techniques [52], we extend the framework to multivariate time series. The model is trained on a real-world network traffic dataset [9, 53], and to mitigate noise, we optimize the training process by separately training the generator and discriminator, thereby reducing the number of qubits required in a single circuit. By the end of the study, we demonstrate that the model performs on par with state-of-the-art models used on the same dataset and attack types. Moreover, the model achieves this performance with only 80 parameters, demonstrating competitive results with a compact architecture.

3.2 Methodology

3.2.1 Data Preparation and Preliminary Analysis

In this study, we use the CIC-CIDDs-2018 dataset [9], a collaborative project between the Communications Security Establishment (CSE) and the Canadian Institute for Cybersecurity (CIC). Designed to simulate a realistic network environment, the dataset captures over 80 features, providing detailed metrics such as flow duration, packet counts, and byte volumes. It includes six types of attacks: brute-force, DoS, web attacks, infiltration, botnet, and DDoS+PortScan. The dataset is publicly available through the CIC and AWS Open Data Registry. This accessibility ensures that the experimental results can be reproduced and compared with state-of-the-art methods [53].

Before conducting advanced modeling, we performed a preliminary analysis to select the most informative features, a critical necessity as qubit resources are scarce on NISQ platforms. Even with successive data injection (detailed later), quantum circuits remain highly sensitive to noise, constraining the volume of tolerable data.

To meet the objective, we applied the Granger causality test a method for determining whether past values of one time series can predict future values of another [36]. We used it to identify significant dependencies between various network metrics and their labels, with the goal of selecting a minimal set of features that best captures anomalous behavior. The analysis is confined to the DDoS attack day (Thursday, February 20, 2018), using one-second aggregated traffic intervals.

Taking into account both domain expertise and the constraints of today’s quantum devices—where each added input dimension deepens and noisifies the circuit—we narrowed the model to two core features:

- **Forward Inter-Arrival Time:** During a port scan aimed at evading detection, an increase in inter-arrival time is often observed. Attackers may deliberately slow down packet transmission to avoid triggering detection systems based on high-frequency traffic.
- **Forward Packet Length:** This measures the size of transmitted packets. DDoS attacks frequently involve large payloads, which increase the average packet size.

3.2.2 Model Architecture

Given the constraints of quantum circuits —particularly the limited number of qubits and the sensitivity to circuit depth— we adopted a sliding window approach to manage temporal data efficiently [37]. Instead of processing long sequences, we analyzed fixed-size windows of $\tau = 3$ at each iteration. For the multivariate time series, this involved extracting τ consecutive values from each variable and encoding them as rotation angles within the quantum circuit. This method balances the need to capture temporal dependencies with the practical limitations of current quantum hardware.

In the model, the time window serves as the conditional input for the GAN, both at the generator and the discriminator. Rather than using pure noise as input, we control the generator’s output by providing contextual data. Similarly, the discriminator receives the same input as the generator, along with the next data point to determine whether it is real or fake. By leveraging a conditional GAN [43], we reduce the temporal dimensionality of the data while ensuring controlled generation of the next data point.

The proposed QGAN employs Variational Quantum Circuits (VQCs), which are fundamental to hybrid quantum-classical algorithms, especially in QML. VQCs consist of parameterized quantum gates optimized through classical methods, enabling learning on quantum devices. VQCs are composed of three parts:

a) Encoding layer: This is the part of the circuit where classical input data is encoded. There are several techniques for input encoding, including basis encoding, amplitude encoding, and angle (rotation) encoding [45] [46]. In the model, we use angle encoding, i.e. we take each input value (i.e., each timestamp in the time window), transform it into a rotation angle using the $\arccos$ function, and then encode it into the circuit via a rotation gate about the Y-axis of the Bloch sphere. We chose this technique because it pairs well with SuDaI [40]. Since




Figure 3.1 Discriminator circuit

SuDaI trades off the number of qubits used with the circuit depth, selecting an encoding technique that uses a minimal number of gates is crucial for limiting the amount of noise.

b) Variational Layer: This is the core of the variational circuit [54]. It consists of parameterized quantum gates followed by an entangling layer. The parameterized gates transform the encoded inputs into expressive quantum states, while the entanglement layer interconnects the qubits, effectively integrating different pieces of information from the input. In the implementation, we use two rotation gates—one for the X-axis and another for the Z-axis—which allow us to explore the entire surface of the Bloch sphere, as two rotations are sufficient. For the entanglement layer, we employ CX gates to entangle every pair of adjacent qubits, where the last qubit is entangled with the first.

c) Measurement Layer: This is the final stage of the circuit, where quantum states collapse into classical values. We use the expectation value of the Pauli Z operator as the output of the model.

Usually, in VQCs, the variational layer (Ansatz) is repeated multiple times within a circuit to increase the model’s expressive power. In this work, we employ a technique called SuDaI [40], where the authors not only repeat the Ansatz but also the input layer. In each new input layer, the successive timestamps are injected allowing the model to process temporal dependencies more effectively. In [40], the successive data injection technique is applied to a univariate time series. In the work, we extend this approach by adapting it to a multivariate time series.

In this work, since we use a QGAN, we use two VQC, one for the generator and one for the discriminator. The generator circuit uses four qubits, while the discriminator uses six qubits. For the former, the first two qubits represent the first feature, a_1 , of the time series. The last two qubits encode the second feature, b_1 , following the same pattern. A series of variational circuits follows after encoding. We alternate between an input layer and a variational layer three times. In each iteration i , we encode a_i , the next timestamp, (respectively b_i for 3rd and 4th qubit). For the discriminator, we follow the same approach but introduce a fifth and sixth qubits, where we encode the predicted/real timestamp. Instead of using SuDaI, we employ data re-uploading [52] for these additional qubits. Figure 1 illustrates the structure of the discriminator.

3.2.3 Training settings and process

The time series consists of a total of 32,401 entries. To set up the training process, we first divide the time series into two parts: one for training and one for testing. The training-to-

testing ratio is 0.5. For training phase, we use data from the second half of the day which contain no attacks. The generator and discriminator are trained exclusively on benign data to enable them to detect anomalies.

Prior to training, we begin by scaling the input. Since we use angle encoding, we first scale the data to the interval $[-1, 1]$ and then apply the arccos function to obtain the corresponding angles. We follow the standard procedure used for training GANs to train the proposed QGAN. We iterate through the time series, and at each step, we take a time window of size 3, (x_{t-2}, x_{t-1}, x_t) , with $x_i = [a_i, b_i]$ and generate a prediction for x_{t+1} . Next, we run the discriminator separately on both the real value of and the predicted value. We then compute the loss using the Mean Squared Error (MSE). The objective is for the discriminator to assign a score as close as possible to 1 for real data and 0 for fake data. Finally, we optimize the discriminator's parameters using the parameter-shift rule [47].

For the generator, we generate a prediction of x_{t+1} and pass it to the discriminator for scoring. We then compute the loss using the MSE, aiming for the discriminator's score to be as close as possible to 1.

Regarding the hyperparameters, we set the learning rate to $\alpha = 0.001$. Both the generator and discriminator were equally trained. The generator was configured with 30 parameters, while the discriminator had 50 parameters.

3.2.4 Anomaly Detection

Algorithm 1: Anomaly Detection Using QGAN

Input : Trained Generator G , Trained Discriminator D , Time series window

$X_t = \{x_{t-2}, x_{t-1}, x_t\}$, Actual target x_{t+1}

Output: Anomaly Score S_t

- 1 $\hat{x}_{t+1} \leftarrow G(X_t)$
 - 2 $MSE_t \leftarrow \frac{1}{n} \sum (x_{t+1} - \hat{x}_{t+1})^2$
 - 3 $D_t \leftarrow D(X_t, x_{t+1})$
 - 4 $w_G \leftarrow \frac{L_D}{L_G + L_D}, w_D \leftarrow \frac{L_G}{L_G + L_D} \quad S_t \leftarrow w_G \cdot MSE_t + w_D \cdot (1 - D_t)$
 - 5 Compute anomaly scores for benign traffic S_{benign}
 - 6 Set threshold T as 99.99th percentile of S_{benign}
 - 7 **if** $S_t > T$ **then**
 - 8 Flag x_{t+1} as anomaly
 - 9 **else**
 - 10 Mark x_{t+1} as normal
-

We leverage both the generator and discriminator of the trained model for anomaly detection, which is the main objective. The core idea behind our approach is to use the generator for

predictive modeling and the discriminator for distinguishing between normal and anomalous behavior. By combining their outputs, we create a robust anomaly detection mechanism.

To achieve this, we compute an anomaly score based on the model’s outputs. Specifically, the generator takes the time window preceding the target data point to generate a prediction. We then compute MSE between the predicted and actual values.

For the discriminator, we provide it with the time window and the actual target value to determine whether it aligns with the distribution of benign traffic. Finally, we combine the MSE and discriminator output through a weighted summation. The weights are determined based on the generator and discriminator losses from the last training iteration, assigning higher confidence to the component with the lower loss.

To determine the threshold for anomaly detection, we first calculate the anomaly scores from the benign traffic (training data). We then compute the 99.99th percentile, meaning only the top 0.01% of benign values exceeds this threshold, making them potential outliers or anomalies. It should be noted that the threshold value of 99.99% here is user-fixed. Automatic determination of this threshold is not addressed in this work. Algorithm 1 provides a detailed illustration of this process.

3.3 Results & Discussion

To evaluate the proposed model, we conducted a series of tests to assess each component’s contribution. We first trained a classical GAN using the same number of parameters as the QGAN, which resulted in poor performance. Competitive results were only achieved after increasing the classical GAN’s parameters to 51 for the generator and 55 for the discriminator.

To test the generative power of the QGAN, we first trained it with a synthetic multivariate time series combining a random walk with a modulated sinusoidal wave. The first variable introduces randomness, while the second adds periodic complexity through amplitude and phase modulation. The results of the QGAN are illustrated in Figure 3, while the results of the classical GAN are illustrated in Figure 4 where both models were trained for the same number of epochs.

The QGAN achieves an MSE of 0.015, significantly outperforming the classical GAN’s MSE of 0.078. The QGAN better captures fine details for this time series, particularly noise fluctuations, highlighting its superior generative capability compared to its classical counterpart.

Following this, we assessed the model’s ability to detect anomalies for the real dataset by evaluating three approaches: the generator alone (Generator-Only Detection), the discriminator alone (Discriminator-Only Detection), and a combined approach (Combined Detection).




Figure 3.2 QGAN results for synthetic time series.




Figure 3.3 Classic GAN results for synthetic time series.

Moreover, we also assessed the combined approach of the QGAN with a noisy simulator using the Fake Vigo backend in order to evaluate the model's robustness under realistic quantum noise conditions. This setup allows us to simulate the behavior of a quantum device affected by noise, providing preliminary insights into the performance of the QGAN when deployed on actual quantum hardware proving a good foundation for future hardware deployment. The results in terms of accuracy, recall, precision, F1-score and MSE for all approaches are shown in Table I

For each test, we compute standard metrics such as accuracy, recall, and F1-score, along with MSE for the training set. The MSE measures how well the generator predicts benign time-series data. The results indicate that the combined approach (Combined Detection), which integrates both the generator and the discriminator, achieves superior performance. The approach attains an accuracy of 0.99 and an F1-score of 0.98, along with the lowest MSE of

0.0013. In contrast, the discriminator-only approach (Discriminator-Only Detection) shows high precision (0.98) but suffers from low recall (0.35) which indicates that it misses many anomalies. The generator-only approach (Generator-Only Detection) achieves perfect recall (1), but its low precision (0.26) suggest that it produces numerous false positives. The classical GAN however, even after increasing the number of parameters, achieves competitive performance (accuracy of 0.96, recall of 0.86, precision of 0.99, and F1-score of 0.92) but exhibits a significantly higher MSE of 0.0157. Moreover, testing the quantum model on a noisy simulator (Noisy anomaly) yields strong performance (accuracy 0.97, F1-score 0.94) with an MSE of 0.0042, which demonstrate its robustness to noise.

Table 3.1 Comparison of different detection methods of the model

Detection Method	Accuracy	Recall	Precision	F1-Score	MSE
QGAN Discriminator-Only Detection	0.83	0.35	0.98	0.51	0.0013
QGAN Generator-Only Detection	0.26	1.00	0.26	0.41	0.0013
QGAN Combined Detection	0.99	0.97	0.99	0.98	0.0013
Classic GAN Combined Detection	0.96	0.86	0.99	0.92	0.0157
Noisy QGAN Combined Detection	0.97	0.89	0.99	0.94	0.0042

The individual performances of Discriminator-Only Detection and Generator-Only Detection highlight the effectiveness of combining both components in the QGAN model. For instance, Discriminator-Only Detection exhibits high precision, accurately identifying true positives. On the other hand, Generator-Only Detection demonstrates high recall, effectively capturing a wide range of anomalies. By integrating both, we leverage the strengths of each—Discriminator-Only Detection’s precision and Generator-Only Detection’s recall—resulting in a more robust and comprehensive anomaly detection system.

Table II summarizes the performance of several classical machine learning models for anomaly detection, evaluated on the same dataset employed in this work [55]. While deep learning models such as CNNs and RNNs achieve slightly higher accuracy, they require thousands of parameters. In contrast, the QGAN uses only 80 parameters while still delivering competitive performance. This indicates the potential of quantum models to achieve good results with significantly fewer resources in terms of number of parameters.

Finally, in Figure 2, we present the results of the anomaly scores in the test data set for the QGAN using the combined approach. Yellow highlights indicate moments where the model failed to detect attacks. However, we observe that these errors still occur within time periods where attacks have already been detected. In real-world scenarios, missing these anomalies




Figure 3.4 Showcasing the similar patterns between the time series of attacks and the time series of anomaly scores.

Table 3.2 Evaluation results of deep learning models for DDoS detection

Model	Accuracy	Precision	Recall	F1-Score	Trainable Parameters
DNN	98.80%	98.80%	98.80%	98.80%	2,977
CNN	99.33%	99.33%	98.34%	99.33%	8,065
RNN	99.78%	99.78%	99.78%	99.78%	12,705

would have a minimal impact.

3.4 Conclusion & Future Work

In this work, we introduced a QGAN for network anomaly detection in a real-world multi-variate dataset. To address qubit limitations, we optimized qubit usage by leveraging data re-uploading, and successive data injection. To enhance feature selection, we applied the Granger causality test, identifying the most relevant network features. Experimental results show that the QGAN achieves high accuracy, recall, and F1-scores in anomaly detection, while also yielding a lower MSE compared to the classical model. Remarkably, the QGAN delivers this performance using just 80 parameters, highlighting its competitive capability with a compact architecture.

Despite the promising results, several areas offer opportunities for future enhancement. First, the model was evaluated on a single type of attack (DDoS) and a limited set of features, which may not generalize to more complex or diverse network threats. Second, we tested the model on a noisy simulator to assess the QGAN’s resilience to quantum noise. However, testing on an actual quantum computer would provide a more definitive evaluation of the model’s performance under real hardware constraints, including decoherence, gate errors, and measurement noise. Such testing is essential to fully validate the model’s practical viability and guide future improvements. Lastly, the current thresholding method for anomaly detection is manually set, which may not adapt well to dynamic environments. Addressing these aspects presents valuable directions for future research and development.

CHAPTER 4 ARTICLE 2: QUANTUM GATED RECURRENT GAN WITH GAUSSIAN UNCERTAINTY FOR NETWORK ANOMALY DETECTION

Authors: Wajdi Hammami, Soumaya Cherkaoui, Jean-Frederic Laprade, Ola Ahmad, Shengrui Wang

Publication: Submitted to IEEE Journal on Selected Areas in Communications

Submission date: July 6, 2025

Contribution

I contributed to the model design, data processing (preprocessing, feature selection), implementation of quantum circuits, result analysis, as well as the full writing of the paper and the literature review.

4.1 Introduction

Anomaly detection in time-series data plays a vital role in monitoring the behavior of complex and dynamic communication systems, where temporal dependencies heavily influence traffic patterns [56]. These anomalies—subtle and often transient deviations from normal network behavior—can indicate serious issues such as cyberattacks or system failures. Unlike static datasets, time-series network data presents unique challenges, since irregularities must not only be flagged, but also explained within the evolving temporal context of network activity.

Traditional statistical approaches often fall short in capturing non-linear dependencies or adapting to evolving patterns over time, necessitating more sophisticated learning-based methods to model the temporal and structural intricacies of modern networks.

In recent years, machine learning has become a cornerstone in solving network anomaly detection problems across various domains [57]. By learning patterns directly from data, machine learning models can adapt to complex, high-dimensional environments where manual rule-setting is impractical. Techniques such as autoencoders [58], recurrent neural networks [2], and generative models [59] have shown strong capabilities in modeling normal behavior and flagging deviations. These models are particularly valuable in settings where data distributions shift over time or labeled anomalies are scarce. As a result, machine learning based anomaly detection techniques has become integral component for modern network security systems [60].

Alongside the advancements in classical machine learning, quantum computing has emerged as a promising frontier for solving complex computational problems. Quantum machine learning (QML), a fusion of quantum computing principles with data-driven learning, is attracting increasing attention for its potential to outperform classical models in specific tasks [61], thanks to its leveraging quantum phenomena such as superposition and entanglement to encode rich feature spaces and capture complex correlations that are challenging for classical models. While empirical evidence on large-scale real-world datasets remains limited, recent studies suggest that hybrid quantum-classical models may offer advantages in scenarios involving high-dimensional data, limited samples, or probabilistic modeling [62]—challenges commonly encountered in network anomaly detection tasks. This growing synergy between quantum computing and machine learning opens new avenues for addressing the complexity of network anomaly detection in ways that classical methods may struggle to achieve.

In this work, we focus on the problem of anomaly detection within network time-series data—a setting where temporal structure plays a crucial role in identifying irregularities. Given the sequential nature of the data as well as the need to capture contextual patterns and handle uncertainty, we propose to leverage the strengths of both classical machine learning and quantum computing. Specifically, we explore the integration of recurrent neural architectures, known for their temporal modeling capabilities, with quantum-enhanced learning mechanisms that offer expressive generative potential [63]. This convergence allows us to build a hybrid model tailored for detecting anomalies in network time-series, where subtle and context-dependent deviations must be identified with high precision.

We address several key limitations that arise in time-series network anomaly detection. While traditional GAN architectures can model data distributions, they often struggle to capture complex temporal dependencies inherent in sequential data. To address this limitation, we integrate Gated Recurrent Units (GRUs), which are well-suited for modeling sequential patterns due to their efficient parameterization. To enrich the model’s expressive capacity, we introduce Variational Quantum Circuits (VQCs) that embed the latent representations into richer Hilbert spaces, enabling the capture of complex feature interactions [64]. Since directly encoding large input dimensions into quantum circuits remains constrained by qubit availability, we employ Successive Data Injection [65] (SuDaI), which allows multiple segments of temporal data to be sequentially injected into the quantum circuit, thereby efficiently utilizing limited quantum resources. Finally, we combine these components within a hybrid GAN architecture, and introduce a multi-metric anomaly detection pipeline that leverages predictive uncertainty, adversarial feedback, and reconstruction-based criteria to robustly detect anomalies with greater contextual sensitivity.

We have evaluated our method using the HAI (HIL-based Augmented Industrial Control Systems) dataset [11], which provides multivariate time-series data collected from a network-integrated industrial control system under both normal operations and staged network-based cyber-attacks. Performance is assessed using the extended time-series aware F1 score (TaF1) metric [66], which accounts for anomaly detection accuracy over temporal intervals. Our approach achieves an TaF1 score of 88.63%, substantially surpassing existing methods.

The chapter proceeds as follows: Section 4.2 outlines our proposed model’s architecture. Section 4.3 details the experimental design and results. Section 4.4 concludes the chapter.

4.2 Proposed Method

4.2.1 Problem Formulation

Let $X = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$ with $\mathbf{x}_t \in \mathbb{R}^d$ denote a multivariate time-series consisting of T time steps and d features. The goal is to detect whether a given point x_t at time t is anomalous based on the context provided by a preceding window of size w . In our unsupervised setting, the training data is assumed to contain predominantly normal patterns, and no anomaly label is used during model training.

We define the input to the model to be a sliding window $W_t = [x_{t-w}, \dots, x_{t-1}]$, which conditions the model to predict the next point $\hat{x}_t$. Deviations between the predicted distribution and the actual observation are often used to compute anomaly scores. Since real-world time series often exhibit temporal drift, subtle deviations, or bursty behaviors, anomalies may arise under changing conditions and be difficult to detect consistently. As a result, the model must capture not only the expected value, but also the associated uncertainty and the contextual consistency of each prediction. Therefore, we formulate anomaly detection as a conditional sequence modeling task: the model learns a conditional distribution $p(x_t|W_t)$, and decides whether x_t is an anomaly depending on whether it significantly deviates from this distribution according to a composite scoring function.

4.2.2 Data Preparation

Our experiments are conducted on version 21.03 of the HAI dataset [11], which consists of multivariate time-series data collected from a network-integrated industrial control system (ICS) operating under both normal and attack conditions. Each sample corresponds to sensor and actuator readings comprising 86 features that reflect the physical and network states of the system. The dataset spans approximately 25.5 days of recorded time-series data, including

around 20 days of normal operation and 5.5 days under attack scenarios. Data are recorded at a frequency of 59 measurements per second. The anomaly events in the dataset correspond to 38 cyberattacks generated using 14 distinct attack primitives –elementary malicious actions such as Set Point (SP) manipulation, Control Output (CO) override, or Process Variable (PV) spoofing–. These attacks vary in duration from 154 to 2880 seconds and range from simple one-primitive intrusions to more complex, stealthy, or combined attacks. Notably, the dataset includes 15 stealth attacks designed to conceal abnormal behaviors by replaying pre-recorded PV values, as well as gradual attacks exhibiting linear transitions (temporal drift) and discrete bursts of activity. This diversity makes the HAI dataset particularly challenging and realistic for network anomaly detection in ICS environments.

The dataset is divided into a training set and a test set. One-third of the test set is used as a validation set for feature selection and hyperparameter tuning, while the remaining two-thirds are held out and kept completely unseen for final performance evaluation.

To reduce dimensionality and improve learning efficiency, we applied feature selection using the Gini importance criterion, following the same methodology as proposed in [67], which was conducted on the same dataset.

This method evaluates the relevance of each feature by measuring its ability to reduce uncertainty across classification splits. The top 16 ranked features are retained, capturing the most informative aspects.

Figure 4.1 illustrates the results of our feature selection. We begin by preprocessing the training set using MinMax normalization to scale all features to the range $[0, 1]$. Subsequently, the normalized data is segmented into fixed-length windows of size w with no overlap, enabling the model to capture local temporal patterns.

For a time-series of length T , our dataset is therefore composed of $\lfloor T/w \rfloor$ inputs, a number that can be excessive given the high computational cost of simulating and training quantum circuits. We employ a clustering-based downsampling strategy to reduce the size of the training set while preserving the diversity of temporal patterns. Each time-series window is flattened into a feature vector and the ensemble of the feature vectors is clustered using K-Means. Each window is therefore assigned to a cluster and the centroid of each cluster is used as a representative sample. Then, the centroid of each cluster is used to represent all windows assigned to that cluster. This results in a more compact yet representative training set, well-suited for efficient quantum model training. The number of clusters is chosen based on validation performance to strike a balance between reduction and pattern preservation.




Figure 4.1 Top 16 Feature Importances ranked by Gini criterion.

4.2.3 Model Overview

In this section, we first introduce the Hybrid Quantum Layers (HQLs), the fundamental building blocks of our model. We then describe the Quantum-Enhanced GRU (QGRU) architecture, followed by the overall design of our generative adversarial network (GAN).

Hybrid Quantum Layers (HQLs)

HQLs have been designed to serve as the core quantum computational units in our model, integrating a classical preprocessing stage, a quantum variational circuit, and a classical post-processing stage. Each HQL consists of a classical linear transformation, a Variational Quantum Circuit (VQC), and a final classical projection layer. The initial dense layer projects the GRU input into a lower-dimensional latent space, which is then injected into the VQC using a strategy known as *Successive Data Injection* (SuDaI) that we have developed recently [65]. In this scheme, the input vector is not injected all at once but introduced progressively across the depth of the circuit. Specifically, the vector components are encoded as rotation angles of single-qubit gates (typically along the R_y or R_z axes) and applied sequentially at multiple circuit layers. In our model, we adopt a SuDaI-inspired scheme not to encode temporal information, but to progressively inject different components of a high-dimensional classical representation (e.g., the output of a dense layer) into a limited number of qubits. This

approach enables us to incrementally construct a rich quantum state while staying within hardware constraints.

The quantum circuit’s output is measured by computing the expectation value of the Pauli-Z observable on each qubit individually, resulting in one scalar output per qubit. These expectation values are used as the quantum layer’s output and are optimized jointly with the rest of the model.

and passed through a final classical linear layer to produce the output. This hybrid design enables the model to exploit both classical and quantum representations, thereby enhancing its ability to capture complex, non-linear temporal dependencies while remaining compatible with NISQ devices. Figure 4.2 illustrates the architecture of the HQL. In the figure, A small set of classical input neurons is first mapped through a fully connected classical layer. The resulting activations are then successively injected into the rotation angles of single qubit gates. Each qubit undergoes a sequence of parameterized rotations (e.g., R_y , R_z , R_x), followed by an entangling layer composed of CNOT gates. This structure is repeated across multiple layers of the variational quantum circuit (VQC). The quantum state is measured at the end of the circuit, and the expectation value of the Pauli-Z observable is computed for each qubit. These values are then passed through a lightweight classical post-processing layer to produce the final output. Gray arrows indicate the direction of data flow.

Quantum-Enhanced GRU (QGRU)

The QGRU is a recurrent unit in which the traditional gating mechanisms—update, reset, and candidate gates—are implemented by HQLs.

At each time step, the QGRU takes the current input vector and the previous hidden state and processes them through HQLs to compute the update and reset gate values. The update gate determines how much of the past hidden state should be retained, while the reset gate controls the influence of the previous state in generating the new candidate activation [68]. The quantum component consists of Variational Quantum Circuits (VQCs), where classical input is encoded via angle encoding and each qubit’s output is computed as the expectation value of the Pauli-Z observable.

Our architecture draws inspiration from the hybrid QGRU proposed by De Falco et al. [69], which also interfaces classical dense layers with quantum circuits to implement gating mechanisms. However, unlike their design that includes both an input and an output fully connected layer surrounding each VQC ($FC_{in} \rightarrow VQC \rightarrow FC_{out}$), we employ only a single classical-to-quantum interface layer before the VQC. This simplification reduces the num-




Figure 4.2 Architecture of the Hybrid Quantum Layer (HQL).

ber of trainable parameters while maintaining the expressive power of the quantum gates. Furthermore, instead of encoding the entire input at once, we inject different components of a high-dimensional vector across successive quantum layers, allowing us to build a richer quantum state using a limited number of qubits.

Quantum Wasserstein Generative Adversarial Network (QWGAN)

Generative Adversarial Networks (GANs) are composed of two competing neural networks: a generator and a discriminator (or critic). The generator learns to produce data that mimics the real distribution, while the critic attempts to distinguish between real and generated samples. Through this adversarial training, the generator improves by minimizing the discrepancy between generated and real samples. In our setting, the generator aims to model the conditional distribution of the next time-step given a sequence of previous observations, making it suitable for time-series forecasting and anomaly detection tasks. The adversarial nature of GANs enables learning complex distributions in an unsupervised manner, a key advantage in settings where labeled anomalies are scarce or unavailable which is the case here as the dataset training set contain only benign data.

The generator works in such a way that at each time step, it receives a fixed-length sliding window of historical multivariate data $W_t = [x_{t-w}, \dots, x_{t-1}]$ and is tasked with forecasting the next time step. This input is first processed by a QGRU to extract sequential features, which are then passed through two parallel HQLs. Each HQL produces one of the parameters of a Gaussian distribution: the mean μ_t and the log-variance $\log(\sigma_t^2)$. This design enables the generator to learn a full predictive distribution over the next time step, capturing both the expected value and the associated uncertainty.

Following that, the model generates a sample from the predicted Gaussian distribution using the reparameterization trick [70]:

$$\hat{x}_t = \mu_t + \sigma_t \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I)$$

This formulation expresses the sampling operation as a differentiable transformation of a noise term, which enables gradient-based optimization despite the presence of stochasticity. In our context, where the training dataset is downsampled using K-Means clustering, this added randomness helps compensate for the reduced data diversity by encouraging the model to generalize beyond the discrete cluster centroids.




Figure 4.3 Model overview

Finally, the sampled value $\hat{x}_t$ is then forwarded to the critic, which estimates the Wasserstein distance between real and generated samples. The generator and critic are trained adversarially to improve the fidelity of synthetic predictions under normal operating conditions. This structure allows the generator to perform probabilistic, context-aware forecasting of the next time step, forming the core of our anomaly detection mechanism. Figure 4.3 provides an overview the QWGAN work flow: The generator (left) uses a QGRU backbone followed by two Hybrid Quantum Layers (HQLs) to produce Gaussian parameters μ_t and $\log \sigma_t^2$, applies the reparameterization trick, and yields a sample $\hat{x}_t$. The critic (right) processes both the real next point x_t and generated sample through parallel QGRU+HQL pipelines to compute Wasserstein scores. Colored blocks distinguish classical (orange) and quantum-augmented (blue) components.

4.2.4 Training Objectives

Training Objective To train our quantum-enhanced generative model, we design a composite loss function comprising three complementary terms:

- **Critic Score:** A Wasserstein loss encourages the generator to produce samples $\hat{x}_{t+1}$ that closely resemble real data x_{t+1} by minimizing the critic’s ability to distinguish between them:

$$\mathcal{L}_{\text{critic}} = \mathbb{E}[D(\hat{x})] - \mathbb{E}[D(x)].$$

This adversarial objective improves the fidelity of generated predictions.

- **KL Divergence:**

$$\mathcal{L}_{\text{KL}} = \frac{1}{2} \mathbb{E} \left[\exp(\log \sigma^2) + \mu^2 - 1 - \log \sigma^2 \right]$$

This regularization serves as a statistical reference baseline and encourages the learned distribution to remain close to a known, unimodal prior. In our anomaly detection context, it helps shape a stable representation of normal behavior — deviations from which can later signal anomalies.

Empirically, we observed that including the KL term significantly improved training stability, especially during the early stages of optimization. This effect is particularly important in our setting, where training is performed on quantum circuit simulators with high computational cost. In such environments, where each epoch is expensive, stabilizing the optimization trajectory is critical for efficient convergence.

- **Variance Penalty:** While the KL term promotes a variance near 1, this can lead to overly diffuse predictions for data normalized in $[0, 1]$. We thus include a regularization term:

$$\mathcal{L}_{\text{var}} = \mathbb{E} \left[\exp(\log \sigma^2) \right],$$

which encourages the generator to remain confident (i.e., to keep uncertainty low) in non-anomalous contexts. The dynamic between the KL divergence and the variance penalty is tuned by the hyperparameter λ_{kl} , which allows us to control the modeled variance of normal behavior and modulate the trade-off between flexibility and certainty.

The full generator loss is given by:

$$\mathcal{L}_{\text{gen}} = -\mathbb{E}[D(\hat{x})] + \mathcal{L}_{\text{var}} + \lambda_{\text{kl}} \mathcal{L}_{\text{KL}},$$

This composite loss enables the model to generate realistic, low-uncertainty predictions during normal operation, while establishing a structured probabilistic baseline suitable for detecting anomalies.

Critic Architecture and Objective The critic network is designed to estimate the Wasserstein distance between real and generated samples, guiding the generator to produce more realistic outputs. Both the real future observation x_{t+1} and the generated prediction $\hat{x}_{t+1}$ are passed through identical parallel pipelines composed of a QGRU followed by a HQL. These pipelines produce latent representations which are evaluated to yield scalar scores. The

critic is trained to maximize the difference between these scores, following the Wasserstein GAN framework:

$$\mathcal{L}_{\text{critic}} = \mathbb{E}[D(\hat{x})] - \mathbb{E}[D(x)].$$

To enforce the Lipschitz continuity required by the Wasserstein objective, we incorporate a gradient penalty term as proposed in WGAN-GP [71]. Specifically, we penalize the squared deviation of the gradient norm from 1 over random interpolations $\tilde{x}$ between real and fake data:

$$\mathcal{L}_{\text{gp}} = \lambda_{\text{gp}} \mathbb{E}_{\tilde{x}} \left[(\|\nabla_{\tilde{x}} D(\tilde{x})\|_2 - 1)^2 \right].$$

The total critic loss thus becomes:

$$\mathcal{L}_{\text{total_critic}} = \mathcal{L}_{\text{critic}} + \mathcal{L}_{\text{gp}}.$$

This gradient penalty stabilizes training by ensuring the critic satisfies the Lipschitz condition.

4.2.5 Anomaly Detection

To identify anomalies in time-series sequences, our approach employs a two-stage detection strategy. At each time step, the model outputs a predicted mean and variance, which are then used to derive multiple diagnostic signals: (i) a reconstruction-based top- k error, (ii) a Wasserstein critic score, and (iii) an interval-based uncertainty score.

In the first stage (Step 1), we apply an interval-based gating mechanism to determine whether further anomaly evaluation is warranted. Only when the observed value significantly deviates from the model’s confidence interval do we proceed to compute additional metrics.

The second stage (Step 2) involves calculating and combining the remaining diagnostic components to produce a final anomaly score. This gated scoring process helps reduce false positives and avoids unnecessary computation in regions where the model remains confident.

The following subsections detail each part of this process: (a) the interval violation gate, (b) the multi-metric scoring signals used for anomaly assessment, and (c) the final anomaly score derived from their combination.

Step 1 – Interval Violation Gating Let the model’s predicted mean and standard deviation at time step t be $\boldsymbol{\mu}_t \in \mathbb{R}^d$ and $\boldsymbol{\sigma}_t \in \mathbb{R}^d$, and let the ground-truth observation be $\mathbf{x}_t \in \mathbb{R}^d$. For each feature $j \in \{1, \dots, d\}$, we define the per-feature violation score as:

$$v_t^{(j)} = \max\left(0, |x_t^{(j)} - \mu_t^{(j)}| - \kappa \cdot \sigma_t^{(j)}\right)$$

where κ is a tunable hyperparameter controlling the size of the confidence interval. For instance, setting $\kappa = 2$ corresponds to a 95% confidence interval under the Gaussian assumption.

To obtain a robust interval-based uncertainty score, we compute the average of the top- k largest per-feature violations at time t :

$$S_{\text{iv}}(t) = \frac{1}{k} \sum_{i=1}^k \text{TopK}_i \left(\left\{ v_t^{(j)} \right\}_{j=1}^d \right) \quad (4.1)$$

This formulation prioritizes the most significant deviations while suppressing minor noise.

The interval violation score quantifies how far the observed input deviates from the model's predicted confidence bounds. Intuitively, it reflects how "surprised" the model is when confronted with the actual input, given what it expected under normal behavior. A low score indicates that the input lies within plausible variation, while a high score signals statistical improbability. Thus, this score serves as a lightweight, uncertainty-aware gating mechanism that filters out benign fluctuations before invoking more computationally intensive diagnostics.

Step 2 – Multi-Metric Investigation When the interval gate is activated, we compute two complementary signals:

- **Wasserstein Critic Score:**

$$\mathcal{S}_{\text{critic}}(t) = -D(\hat{x}_t) \quad (4.2)$$

where a lower critic score implies a sample appears more fake and thus more anomalous.

- **Reconstruction Top- k Error Score:**

We compute the squared error between the predicted mean $\boldsymbol{\mu}_t$ and the true value $\mathbf{x}_t$ at time t across all features. To focus on the most significant deviations, we retain only the top- k largest squared errors and average them:

$$S_{\text{top-}k}(t) = \frac{1}{k} \sum_{i=1}^k \text{TopK}_i \left(\left\{ \left(x_t^{(j)} - \mu_t^{(j)} \right)^2 \right\}_{j=1}^d \right) \quad (4.3)$$

This score emphasizes dominant deviations while reducing the influence of small, potentially noisy fluctuations.

Final Anomaly Score The final anomaly score is computed by summing the reconstruction-based and critic-based components:

$$A(t) = S_{\text{top-}k}(t) + S_{\text{critic}}(t)$$

Here, $S_{\text{top-}k}(t) \in [0, \infty)$ is always non-negative, while $S_{\text{critic}}(t)$ may be positive or negative depending on the critic’s output.

To ensure comparability between the reconstruction error and the critic score, we normalize both $S_{\text{top-}k}(t)$ and $S_{\text{critic}}(t)$ to the range $[0, 1]$ using min-max normalization computed over the validation set. The final anomaly score is then given by:

$$A(t) = \tilde{S}_{\text{top-}k}(t) + \tilde{S}_{\text{critic}}(t) \tag{4.4}$$

where $\tilde{S}_{\text{top-}k}$ and $\tilde{S}_{\text{critic}}$ denote the normalized forms of the respective metrics. This normalization ensures that both signals contribute proportionally to the final score, preventing any single component from dominating due to scale differences.

In practice, we observed that combining these two signals enhances the detection of anomalous patterns by capturing both structural deviation and adversarial detectability.

This scoring is invoked only when the interval violation score $S_{\text{iv}}(t)$ exceeds a threshold. This conditional mechanism avoids unnecessary computation during normal behavior and focuses diagnostics only when the model expresses high uncertainty about its prediction.

4.3 Results

4.3.1 Experimental Setup

Quantum Simulation Configuration. All quantum components in our model, including the QGRU gates and Hybrid Quantum Layers (HQLs), are simulated using **PennyLane** with the `lightning.qubit` simulator for accelerated training and efficient quantum circuit evaluation. Each Variational Quantum Circuit (VQC) operates on **5 qubits**, balancing between expressive power and manageable simulation cost.

Time Window Size Estimation. To determine an appropriate sliding window size w , we analyze statistics from the validation set. Since the attack periods in the validation data are generally short and infrequent, naively averaging durations would lead to overly small window sizes that fail to capture stable temporal patterns. To address this, we compute a **probability-weighted geometric mean** that balances short attack intervals against longer normal periods:

$$w = \mu_d^p \cdot g^{1-p}$$

Here, μ_d represents the average duration of labeled attack intervals, computed from ground-truth annotations. The value of g corresponds to the average time between successive attacks (i.e., normal intervals). The probability p reflects the proportion of time the system is under attack, estimated by dividing the total number of attack-labeled time steps by the overall sequence length. These statistics are derived empirically from the validation set.

Thresholding Mechanism. To improve robustness and temporal adaptability, we apply a local, adaptive thresholding strategy at each time step t . A sliding window W_t is centered around the current index, and we compute the mean and standard deviation of anomaly-related scores within that window. The adaptive threshold T_t is then defined as:

$$T_t = \mu(W_t) + k \cdot \sigma(W_t)$$

where k is a tunable sensitivity hyperparameter. This threshold is used in two places within our detection pipeline:

- First, it serves as a gate: we invoke deeper diagnostics (i.e., compute $A(t)$) only if the interval violation score $S_{iv}(t)$ exceeds T_t .
- Second, it is used for final decision-making: an anomaly is flagged if the computed final score also satisfies $A(t) > T_t$.

This dual usage enables both efficient screening and adaptive alerting, ensuring that anomalies are flagged relative to the local behavior of the system.

Hyperparameter Selection To ensure reproducibility and interpretability, we report in Table 4.1 the main hyperparameters used in our model, along with their description.

Table 4.1 Hyperparameters Used in Our Model

Hyperparameter	Value	Description
Learning rate (α)	0.001	Used for training both generator and discriminator via gradient descent
Window size (τ)	3	Number of past timesteps used as conditional input for QGAN
Number of clusters (n)	300	KMeans clusters used for successive data injection (SuDaI)
Rotation encoding	$\arccos(x)$	Maps scaled input x to quantum gate rotations
Ansatz depth	12	Total number of alternating (encoding + variational) layers used in the Variational Quantum Circuit (VQC), controlling overall model expressiveness
Ansatz layers (SuDaI)	6	Number of ansatz blocks where new timestamps are injected using the SuDaI technique; among the 12 total layers, data is injected into 6 of them
Number of qubits	5 (G), 5 (D)	Generator uses 5 qubits; discriminator uses 5 qubits
Input scaling range	$[0, 1]$	Range to which raw features are scaled before encoding
k (Thresholding mechanism)	1.5	Number of windows used to compute adaptive threshold score
p (attack probability)	$\frac{A \cdot \mu_d}{N}$	Estimated proportion of time the system is under attack; N is total number of timesteps; A is the total number of attacks; μ_d Average duration of labeled attack intervals
g (normal interval)	$\frac{N}{A} - \mu_d$	Average time between successive attacks (i.e., average benign interval)
Top- k (Interval Violation Gating)	3	Number of top features used in Interval Violation Gating
Top- k (Reconstruction error)	3	Number of most significant dimensions used to compute final score
Feature selection count	16	Number of selected features from Gini method

4.3.2 Evaluation Metrics

Evaluating time-series anomaly detection models requires metrics that accurately reflect both detection correctness and the temporal consistency of detected anomaly intervals. Traditional point-wise metrics such as precision, recall, and F1-score fail to capture the sequential structure of anomalies, which are typically characterized by sustained abnormal behavior over continuous intervals rather than isolated points.

To address this limitation, we adopt the enhanced time-series aware evaluation framework

introduced in [72], which extends conventional precision and recall to better account for partial and fragmented detections in time-series contexts. Specifically, the *enhanced time-aware precision* (**eTaP**) rewards consistent detection of anomalous intervals, while the *enhanced time-aware recall* (**eTaR**) penalizes fragmented or incomplete coverage of anomaly ranges.

We combine these two components into a harmonic mean to compute the overall time-series aware F1 score, which we refer to as **TaF1**:

$$\text{TaF1} = \frac{2 \cdot \text{eTaP} \cdot \text{eTaR}}{\text{eTaP} + \text{eTaR}}.$$

This combined metric offers a more realistic assessment of anomaly detection performance in scenarios where anomalies unfold over extended periods, such as in industrial control systems.

4.3.3 Results

Table 4.2 presents the performance of our proposed model compared to several baselines on the HAI dataset, evaluated using time-aware metrics. The baseline results (RNNv1, RNNv2) are directly reported from prior work [72], and were not retrained in our setting.

Our method achieves the highest TaF1 score of 0.89, outperforming both RNN-based baselines [72]—with scores of 0.83 for RNNv1 and 0.78 for RNNv2—as well as the naive last-timestep predictor, which yields a score of only 0.04. The latter serves as a non-learned baseline that simply reuses the previous value as the prediction, illustrating the difficulty of the forecasting task.

To further assess the contribution of our proposed model, we introduce an additional baseline: the *Cluster-Likelihood Baseline*. This method relies solely on statistical modeling of the target distribution within each cluster obtained from KMeans applied to training windows. At test time, anomaly detection is based on the log-likelihood of the predicted point under the corresponding cluster’s Gaussian distribution. Despite its principled design, this baseline achieves a TaF1 of only 0.23, with a recall (eTaR) of 0.15—highlighting its limited ability to capture temporally extended anomalies.

In addition to its superior TaF1 score, our model attains an eTaP of 0.93, reflecting its prompt detection capability, and an eTaR of 0.85, indicating strong coverage of complete anomaly intervals. These results collectively demonstrate that our quantum-enhanced GAN achieves robust and temporally-aware anomaly detection, significantly outperforming both deep learning and statistical alternatives.

Table 4.2 Time-Aware F1 Metrics on the HAI Dataset

Model	TaF1	eTaP	eTaR
RNNv1	0.83	0.90	0.84
RNNv2	0.78	0.75	0.69
Last Timestep	0.04	0.03	0.09
Cluster-Likelihood Baseline	0.23	0.48	0.15
QWGAN (Ours)	0.89	0.93	0.85

4.3.4 Ablation Study

The ablation results in Table 4.3 highlight the contribution of each component of our composite anomaly scoring strategy (as defined in Equations 4.1, 4.3, and 4.2). Specifically, we assess the impact of interval violation gating, reconstruction-based error, and Wasserstein critic feedback.

When all three components are combined into the final score (Equation 4.4), the model achieves the highest TaF1 (0.89) and the best early-time precision (eTaP = 0.93), confirming that the fusion of diverse signals leads to more timely and precise anomaly detection. Additionally, we observe an eTaR of 0.85, indicating strong recall across complete anomaly intervals.

These findings support the effectiveness of the full scoring strategy and justify the use of multi-signal integration in our final anomaly detection framework.

Removing the interval-violation gate (“Critic Only” and “Reconstruction Only” rows) lowers recall markedly (eTaR drops to 0.63 and 0.72, respectively), confirming that the gate is essential for sustaining detection across the full duration of an attack. Conversely, using *only* the interval-violation score achieves an eTaR of 0.85—on par with the full model—yet loses precision (eTaP falls to 0.91) and slightly trails in TaF1. This suggests that variance-driven gating is a strong standalone detector for sustained anomalies but benefits from the critic’s realism check to reduce false positives.

Finally, the “Critic Only” configuration achieves high precision (0.85) but the lowest recall, illustrating that critic feedback excels at flagging obvious outliers yet misses subtle deviations. Taken together, these findings show that interval violation provides broad coverage, the critic sharpens precision, and their combination —augmented by reconstruction— offers the most balanced and discriminative anomaly signal.

To better understand the role of predictive variance in our model, we visualize in Figure 4.4 the evolution of $\log(\sigma^2)$ across time. The plot reveals a clear correlation between elevated log-

Table 4.3 Ablation Study Results on the HAI Dataset

Configuration	eTaP	eTaR	TaF1
Full Model	0.93	0.85	0.89
Critic Only	0.85	0.63	0.77
Reconstruction error Only	0.85	0.72	0.78
Interval Violation Only	0.91	0.85	0.88

variance values and anomalous events. Specifically, during anomalous intervals, the generator consistently outputs higher uncertainty, reflected by spikes in $\log(\sigma^2)$. This suggests that the model becomes less confident in its predictions when exposed to inputs that diverge from the learned distribution of normal behavior.




Figure 4.4 Scatter plot of scaled mean log-variance, with ground-truth anomalies highlighted by red square outlines.

This behavior validates the intuition behind penalizing the variance in our generator loss: by discouraging excessive uncertainty on normal data, we encourage sharper, more confident predictions in the expected regime. As a result, sudden increases in variance can be interpreted as early indicators of distributional shift, thus improving anomaly localization. These findings support the use of variance as an informative signal for anomaly detection at inference time.

4.4 Conclusion

In this work, we introduced a novel quantum-enhanced anomaly detection framework that integrates a Quantum Gated Recurrent Unit (QGRU) with a hybrid adversarial training scheme. Our architecture leverages the expressiveness of Variational Quantum Circuits

(VQCs) and SuDaI for sequence modeling trick to generate uncertainty-aware predictions. To improve training stability and model reliability, we incorporated a KL-divergence regularizer and a log-variance penalty alongside a Wasserstein-based critic with gradient penalty. Furthermore, we proposed an interpretable multi-stage anomaly scoring mechanism that begins with interval violation analysis leading to better scoring based on critic evaluation and reconstruction error. Experimental results demonstrated the model’s strong performance across extended time-aware metrics (TaF1: 0.89, eTaR: 0.85), validating the utility of our quantum-inspired design and scoring strategy.

This work marks an interesting step toward scalable and explainable quantum machine learning systems for network anomaly detection, combining quantum expressiveness with principled uncertainty modeling.

CHAPTER 5 CONCLUSION

5.1 Summary of Works

This thesis presents two contributions to the field of anomaly detection in multivariate time-series data, with a focus on network security applications. Both contributions leverage quantum-inspired techniques and generative modeling to improve detection performance while operating within the constraints of current quantum hardware.

Article 1: Enhancing Network Anomaly Detection with Quantum GANs and Successive Data Injection

In this work, we propose a conditional Quantum Generative Adversarial Network (QGAN) architecture integrated with the Successive Data Injection (SuDaI) mechanism. This technique enables the sequential injection of input features into shallow quantum circuits, enhancing the model's ability to learn temporal dependencies without increasing quantum depth. The proposed model is evaluated on the CIC-CIDDs2018 and HAI datasets and demonstrates superior performance over classical baselines in terms of precision, recall, and robustness.

Article 2: Quantum Gated Recurrent GAN with Gaussian Uncertainty for Network Anomaly Detection

The second contribution introduces a hybrid architecture combining Gated Recurrent Units (GRUs) with parameterized quantum circuits. The model generates Gaussian parameters (mean and variance) at each timestep, enabling uncertainty-aware anomaly scoring. This approach captures both temporal structure and predictive uncertainty, which is particularly valuable in safety-critical applications. Evaluation results on the HAI dataset highlight the effectiveness of the proposed hybrid model in capturing both temporal dynamics and predictive uncertainty, leading to improved robustness in anomaly detection for industrial network environments.

Together, these works highlight the potential of quantum and hybrid quantum-classical models in advancing time-series anomaly detection, particularly in network security contexts.

5.2 Limitations

While the proposed quantum-enhanced architectures offer promising results in time-series anomaly detection, several limitations remain:

- **Scalability:** The current quantum models operate under the constraints of Noisy Intermediate-Scale Quantum (NISQ) hardware, which limits the number of qubits and circuit depth. This restricts the ability to scale to high-dimensional inputs or long temporal sequences.
- **Simulation Overhead:** Due to limited access to real quantum hardware, our experiments were conducted using classical quantum simulators, which are computationally expensive and do not fully capture hardware-specific noise or decoherence effects.
- **Architecture Generalization:** The hybrid architectures were designed and tuned for specific datasets (e.g., CIC-CIDS2018 and HAI). Their generalizability to other network environments or anomaly types has not been fully explored.

These limitations open avenues for future work, including the design of more expressive quantum circuits, efficient training algorithms, and domain adaptation techniques for broader deployment.

5.3 Future Research

This work lays the foundation for several promising research directions in quantum-enhanced time-series anomaly detection:

- **Hardware Execution:** Future efforts will focus on deploying the proposed models on real quantum hardware to assess performance under realistic noise and decoherence constraints, going beyond classical simulation environments.
- **End-to-End Quantum Recurrence:** Building on our preliminary integration of quantum components, future models may embed quantum circuits directly into the recurrence mechanism (e.g., within GRU gates), enabling fully quantum temporal modeling.
- **Model Compression and Efficiency:** Research into optimizing quantum circuit depth, gate complexity, and hybrid-classical resource sharing can further enhance scalability and reduce inference latency for real-time deployment.

Pursuing these directions will not only advance the applicability of quantum machine learning but also bridge the gap between theoretical innovation and practical network security needs.

REFERENCES

- [1] Z. Zamanzadeh Darban *et al.*, “Deep learning for time series anomaly detection: A survey,” *ACM Computing Surveys*, vol. 57, no. 1, pp. 1–43, 2024.
- [2] P. Malhotra *et al.*, “Long short term memory networks for anomaly detection in time series,” in *ESANN*, 2015.
- [3] B. Luer and K. Böhm, “Anomaly detection using generative adversarial networks: Reviewing methodological progress and challenges,” in *KDD*, 2022.
- [4] A. Dawoud, S. Shahrstani, and C. Raun, “Deep learning for network anomalies detection,” in *iCMLDE*. IEEE, 2018.
- [5] V. Dunjko and H. J. Briegel, “Machine learning & artificial intelligence in the quantum domain: a review of recent progress,” *Reports on Progress in Physics*, vol. 81, no. 7, p. 074001, 2018.
- [6] M. Hdaib, S. Rajasegarar, and L. Pan, “Quantum deep learning-based anomaly detection for enhanced network security,” *Quantum Machine Intelligence*, vol. 6, no. 1, p. 26, 2024.
- [7] S. Roy, “A comprehensive survey on network traffic anomaly detection using deep learning,” *arXiv preprint*, vol. arXiv:2406.12345, 2024.
- [8] J. Koumar, K. Hynek, T. Čejka, and P. Šiška, “Cesnet-timeseries24: Time series dataset for network traffic anomaly detection and forecasting,” *Scientific Data*, 2024.
- [9] Canadian Institute for Cybersecurity, “CSE-CIC-IDS2018 Dataset,” <https://www.unb.ca/cic/datasets/ids-2018.html>, 2018.
- [10] M. Tavallaei, E. Bagheri, W. Lu, and A. A. Ghorbani, “A detailed analysis of the kdd cup 99 data set,” in *Proceedings of the 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA)*. IEEE, 2009, pp. 1–6. [Online]. Available: <https://doi.org/10.1109/CISDA.2009.5356528>
- [11] H.-K. Shin, W. Lee, J.-H. Yun, and H. Kim, “Hai 1.0: Hil-based augmented ics security dataset,” in *13th USENIX Workshop on Cyber Security Experimentation and Test (CSET 20)*. USENIX Association, 2020. [Online]. Available: <https://www.usenix.org/conference/cset20/presentation/shin>

- [12] M. Sakurada and T. Yairi, “Anomaly detection using autoencoders with nonlinear dimensionality reduction,” in *MLSDA*, 2014, pp. 4–11.
- [13] M. Wang *et al.*, “Anomaly detection in multidimensional time series for water injection pump operations based on lstma-ae and mechanism constraints,” *Scientific Reports*, vol. 15, 2025.
- [14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in Neural Information Processing Systems*, vol. 27, 2014, pp. 2672–2680.
- [15] A. Di Mattia *et al.*, “A survey on gans for anomaly detection,” in *arXiv preprint arXiv:1906.11632*, 2019.
- [16] A. Geiger *et al.*, “Tadgan: Time series anomaly detection using generative adversarial networks,” in *arXiv preprint arXiv:2009.07769*, 2020.
- [17] K. Saravanan, X. Zhang, and P. Zhang, “Tsi-gan: Unsupervised time series anomaly detection using convolutional cycle-consistent generative adversarial networks,” in *arXiv preprint arXiv:2303.12952*, 2023.
- [18] M. Schuld, I. Sinayskiy, and F. Petruccione, “An introduction to quantum machine learning,” *Contemporary Physics*, vol. 56, no. 2, pp. 172–185, 2015.
- [19] M. Rath and H. Date, “Quantum data encoding: a comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy,” *EPJ Quantum Technology*, vol. 11, no. 72, 2024.
- [20] A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster, and J. I. Latorre, “Data re-uploading for a universal quantum classifier,” *Quantum*, vol. 4, p. 226, 2020.
- [21] H.-L. Huang, Y. Du, M. Gong *et al.*, “Experimental quantum generative adversarial networks for image generation,” *Nature Communications*, vol. 12, p. 2631, 2021.
- [22] 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.
- [23] Z. Chai *et al.*, “Quantum anomaly detection with a spin processor in diamond,” *Advanced Quantum Technologies*, vol. 7, no. 5, p. 2300385, 2024.
- [24] M. Alhussein, F. Yu *et al.*, “A quantum-classical hybrid solution for deep anomaly detection,” *Entropy*, vol. 25, no. 3, p. 427, 2023. [Online]. Available: <https://www.mdpi.com/1099-4300/25/3/427>

- [25] B. Kalfon, S. Cherkaoui, J.-F. Laprade, O. Ahmad, and S. Wang, “Successive data injection in conditional quantum gan applied to time series anomaly detection,” *arXiv preprint*, 2023.
- [26] P. K. Choudhary, N. Innan, M. Shafique, and R. Singh, “A hybrid classical–quantum neural network for time series regression,” *arXiv preprint*, vol. arXiv:2503.15403, 2025.
- [27] F. De Falco *et al.*, “Evolving hybrid quantum-classical gru architectures for multivariate time series,” in *IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2024, pp. 73–78.
- [28] Anonymous, “Deep learning advancements in anomaly detection,” *arXiv preprint*, vol. arXiv:2503.13195, 2025.
- [29] A. Filali *et al.*, “Dynamic sdn-based radio access network slicing with deep reinforcement learning for urllc and embb services,” *IEEE Transactions on Network Science and Engineering*, vol. 9, no. 4, pp. 2174–2187, Dec. 2022.
- [30] A. Abouaomar *et al.*, “Federated deep reinforcement learning for open ran slicing in 6g networks,” *IEEE Communications Magazine*, vol. 61, no. 2, pp. 126–132, Feb. 2023.
- [31] S. Cherkaoui *et al.*, “Research landscape – 6g networks research in europe,” *IEEE Network*, vol. 35, no. 6, pp. 6–13, 2021.
- [32] B. Luer *et al.*, “Anomaly detection using generative adversarial networks: Reviewing methodological progress and challenges,” in *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022.
- [33] S. Wang *et al.*, “Machine learning in network anomaly detection: A survey,” *IEEE Access*, vol. 9, pp. 152 379–152 396, 2021.
- [34] A. Dawoud *et al.*, “Deep learning for network anomalies detection,” in *2018 International Conference on Machine Learning and Data Engineering (iCMLDE)*, 2018.
- [35] W. Hu *et al.*, “Research on anomaly network detection based on self-attention mechanism,” *Sensors*, vol. 23, no. 11, p. 5059, 2023.
- [36] A. R. Al-Ghuwairi *et al.*, “Intrusion detection in cloud computing based on time series anomalies utilizing machine learning,” *Journal of Cloud Computing*, vol. 12, no. 1, p. 127, 2023.

- [37] A. Blázquez-García *et al.*, “A review on outlier/anomaly detection in time series data,” *ACM Computing Surveys*, vol. 54, no. 3, 2021.
- [38] S. Cherkaoui, “Quantum leap: Exploring the potential of quantum machine learning for communication networks,” in *Proceedings of the Int’l ACM Conference on Modeling Analysis and Simulation of Wireless and Mobile Systems*, 2023.
- [39] T. A. Ngo *et al.*, “A survey of recent advances in quantum generative adversarial networks,” *Electronics*, vol. 12, no. 4, p. 856, 2023.
- [40] B. Kalfon *et al.*, “Successive data injection in conditional quantum gan applied to time series anomaly detection,” arXiv preprint arXiv:2310.05307, 2023.
- [41] A. Vieloszynski *et al.*, “Latentqgan: A hybrid qgan with classical convolutional autoencoder,” in *Proceedings of the 2024 IEEE 10th World Forum on Internet of Things (WF-IoT)*, Nov. 2024, pp. 1–7.
- [42] M. Schuld *et al.*, “An introduction to quantum machine learning,” *Contemporary Physics*, vol. 56, no. 2, pp. 172–185, Oct. 2014.
- [43] M. Mirza *et al.*, “Conditional generative adversarial nets,” arXiv preprint arXiv:1411.1784, Nov. 2014. [Online]. Available: <https://arxiv.org/abs/1411.1784>
- [44] H.-L. Huang *et al.*, “Experimental quantum generative adversarial networks for image generation,” *Nature Communications*, 2021.
- [45] M. Rath *et al.*, “Quantum data encoding: A comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy,” *EPJ Quantum Technology*, 2024.
- [46] D. Ranga *et al.*, “Quantum machine learning: Exploring the role of data encoding techniques, challenges, and future directions,” *Mathematics*, vol. 12, no. 21, 2024.
- [47] M. Schuld *et al.*, “Evaluating analytic gradients on quantum hardware,” *Physical Review A*, vol. 99, no. 3, 2019.
- [48] J. Rezgui *et al.*, “Detecting faulty and malicious vehicles using rule-based communications data mining,” in *Proceedings of the 2011 IEEE 36th Conference on Local Computer Networks (LCN)*, Oct. 2011, pp. 827–834.
- [49] W. Hamhoum *et al.*, “Mistralbsm: Leveraging mistral-7b for vehicular networks misbehavior detection,” arXiv preprint arXiv:2407.18462, 2024.

- [50] —, “Fortifying open ran security with zero trust architecture and transformers,” in *ICC 2024 – IEEE International Conference on Communications*, 2024.
- [51] A. Aaraba *et al.*, “Quack-tsf: Quantum-classical kernelized time series forecasting,” in *Proceedings of the 2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1, Sep. 2024, pp. 1628–1638.
- [52] A. Pérez-Salinas *et al.*, “Data re-uploading for a universal quantum classifier,” *Quantum*, vol. 4, p. 226, 2020.
- [53] J. L. Leevy *et al.*, “A survey and analysis of intrusion detection models based on cse-cic-ids2018 big data,” *Journal of Big Data*, vol. 7, no. 1, p. 104, 2020.
- [54] M. Cerezo *et al.*, “Variational quantum algorithms,” *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, Sep. 2021.
- [55] D. Agostinello *et al.*, “Anomaly-based intrusion detection system for ddos attack with deep learning techniques,” in *Proceedings of the 20th International Conference on Security and Cryptography (SECRYPT 2023)*, 2023.
- [56] M. Ahmed, A. Mahmood, and J. Hu, “A survey of network anomaly detection techniques,” *Journal of Network and Computer Applications*, vol. 60, pp. 19–31, 2016.
- [57] H. Zhang, Y. Wang, J. Wang, K. Li, M. Liu, and S. McKeever, “Network anomaly detection: A survey and comparative analysis of stochastic and machine learning methods,” *Computer & Security*, vol. 86, pp. 147–169, 2019.
- [58] J. An and S. Cho, “Variational autoencoder based anomaly detection using reconstruction probability,” *Special Lecture on IE*, vol. 2, no. 1, pp. 1–18, 2015.
- [59] W. Hammami, S. Cherkaoui, and S. Wang, “Enhancing network anomaly detection with quantum gans and successive data injection for multivariate time series,” *arXiv preprint arXiv:2505.11631*, 2025. [Online]. Available: <https://doi.org/10.48550/arXiv.2505.11631>
- [60] M. Landauer, F. Skopik, B. Stojanović, W. Kastner, and T. Zseby, “A review of time-series analysis for cyber security analytics: From intrusion detection to attack prediction,” *International Journal of Information Security*, vol. 24, no. 3, 2025. [Online]. Available: <https://doi.org/10.1007/s10207-024-00921-0>
- [61] M. Hdaib, S. Rajasegarar, and L. Pan, “Quantum deep learning-based anomaly detection for enhanced network security,” *Quantum Machine Intelligence*, vol. 6, no. 1, p. 26, 2024. [Online]. Available: <https://link.springer.com/article/10.1007/s42484-024-00163-2>

- [62] V. Dunjko and H. J. Briegel, “Machine learning & artificial intelligence in the quantum domain: a review of recent progress,” *Reports on Progress in Physics*, vol. 81, no. 7, p. 074001, 2018.
- [63] S. Jerbi, M. F. Huber, V. Dunjko, and H. J. Briegel, “Quantum machine learning for sequence modeling: a review and outlook,” *Machine Learning: Science and Technology*, vol. 4, no. 1, p. 013002, 2023.
- [64] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, “Parameterized quantum circuits as machine learning models,” *Quantum Science and Technology*, vol. 4, no. 4, p. 043001, 2019.
- [65] B. Kalfon, S. Cherkaoui, J.-F. 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, 2024. [Online]. Available: <https://doi.org/10.1049/qtc2.12088>
- [66] W.-S. Hwang, J.-H. Yun, J. Kim, and B. G. Min, “Do you know existing accuracy metrics overrate time-series anomaly detections?” in *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing*. ACM, 2022, pp. 403–412.
- [67] T. Cultice, M. S. Hassan Onim, A. Giani, and H. Thapliyal, “Anomaly detection for real-world cyber-physical security using quantum hybrid support vector machines,” in *2024 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, 2024, pp. 619–624.
- [68] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, “Empirical evaluation of gated recurrent neural networks on sequence modeling,” *arXiv preprint arXiv:1412.3555*, 2014.
- [69] F. De Falco, L. Lavagna, A. Ceschini, A. Rosato, and M. Panella, “Evolving hybrid quantum-classical gru architectures for multivariate time series,” in *2024 IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP)*, 2024, pp. 73–78. [Online]. Available: <https://ieeexplore.ieee.org/document/10734792>
- [70] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2014.
- [71] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, “Improved training of wasserstein gans,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017. [Online]. Available: <https://arxiv.org/abs/1704.00028>

- [72] W.-S. Hwang, J.-H. Yun, J. Kim, and B. G. Min, “Do you know existing accuracy metrics overrate time-series anomaly detections?” in *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing*, 2022, pp. 1234–1241. [Online]. Available: <https://dl.acm.org/doi/10.1145/3477314.3507024>