Quantum-Assisted Variational Deep Learning for Efficient Anomaly Detection in Secure Cyber-Physical System Infrastructures

Nilesh Bhosale[1], Bukya Mohan Babu[2], M. Karthick Raja[3], Prof. Ts. Dr. Yousef A.Baker El-Ebiary[4], Manasa Adusumilli[5], Elangovan Muniyandy[6], Dr. David Neels Ponkumar Devadhas[7]

Assistant Professor, Department of Applied Mathematics and Humanities, Yeshwantrao Chavan College of Engineering, Nagpur, India[1]

Department of CSE (Data Science), CMR Technical Campus, Hyderabad, Telangana, India[2]

Assistant Professor, Department of CSE, Sri Eshwar College of Engineering, Kinathukadavu, India[3]

Faculty of Informatics and Computing, UniSZA University, Malaysia[4]

Assistant Professor, Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, AP, India[5]

Department of Biosciences-Saveetha School of Engineering, Saveetha Institute of Medical and Technical Sciences, Chennai - 602 105, India[6]

Applied Science Research Center, Applied Science Private University, Amman, Jordan[6]

Professor, Department of Electronics and Communication Engineering, Vel Tech Rangarajan Dr.Sagunthala R&D Institute of Science and Technology, Chennai, Tamil Nadu, India[7]

***Abstract*—The aim of the current study is to propose a Quantum-Assisted Variational Autoencoder (QAVAE) model capable of efficiently identifying anomalies in high-dimensional, time-series data produced by cyber-physical systems. The existing approaches to machine learning have some limitations when recording temporal interactions and take substantial time to run with many attributes. To meet all of these challenges, this study aims at proposing a quantum-assisted approach to anomaly detection using the potential of a Quantum-Assisted Variational Autoencoder (QAVAE). The general goal of this research is to optimize anomaly detection systems using consummate deep learning quantum computing models. According to the QAVAE framework, variational inference is employed for learning latent representations of time series data; besides, quantum circuits are utilized for enhancing the capacity of the model and its generalization capability. This work was accomplished using Python programming language, and the analysis was carried out using TensorFlow Quantum. The QAVAE model demonstrates the highest accuracy of 95.2%, indicating its strong capability in correctly identifying both anomalous and normal instances. So, it can learn well from the data and keep stable in the evaluation process, which will make it suitable for real-time anomaly detection in dynamic environments. In conclusion, the QAVAE model brings a reasonable approach and solution for anomaly detection that is accurate in identifying and scalable too. Utilizing the HAI, the dataset achieved a high detection accuracy of 95.2%. Further research has to be dedicated to its application to quantum computing architecture as well as to modifications that allow for its use on multi-variable actual-life data.**

Keywords*—Quantum variational circuits; cyber-physical system security; hybrid quantum-classical algorithms; anomaly detection framework; quantum machine learning optimization

I. INTRODUCTION

In today's data-driven world, the ability to detect anomalies within complex datasets is paramount across various industries, including finance, healthcare, cybersecurity, and industrial operations [1] [2]. Due to their sporadic occurrence and the complex nature of the data where they are located, there is a large number of difficulties associated with anomalies, which may signal important events like frauds, system failures, or violations of security [3], [4]. Communal paradigm approaches used in anomaly detection commonly fall short due to their basic methodologies to detect these irregularities in big data and high-dimensional space domain [5]. This constraint has led to actual research for more refined approaches that mainly employ activities in areas of machine learning and the relatively recent development of quantum computing. QAVAE was chosen because of its ability to effectively model high-dimensional time-series data, its robustness under noisy conditions, and its capability to improve feature learning through quantum circuit integration, performing better than classical counterparts in terms of accuracy and scalability.

Even with progress in classical and quantum machine learning, a notable gap exists in creating effective, hybrid models that can perform strong anomaly detection for complex, real-time cyber-physical systems. Traditional machine learning [6] techniques, such as Support Vector Machines, Decision Trees, and traditional Autoencoders, are generally unable to efficiently describe nonlinear relationships and temporal dependencies in data with high dimensions. Such shortcomings lead to poor generalizability, high false positives, and lower reliability under dynamic conditions. Conversely, although quantum machine learning is promising because it has the ability to improve feature representation by utilizing quantum superposition and entanglement, most quantum methods are still theoretical or hindered by existing hardware limitations like quantum noise and scale.

Most studies concentrate on either classical deep learning or single quantum models without considering a fully integrated quantum-classical architecture appropriate for real-world implementation. Furthermore, it is rare that studies test their models on real cyber-physical datasets or consider actual quantum noise conditions. This work seeks to overcome these limitations by introducing a Quantum-Assisted Variational Autoencoder (QAVAE) that benefits from both classical VAEs and parameterized quantum circuits. The goal is to develop a scalable, interpretable, and highly accurate anomaly detection model applicable to realistic applications in critical infrastructures and dynamic industrial settings.

One such model that has garnered attention is the Quantum Variational Autoencoder (QVAE) [6] [7]. Classic Autoencoder networks can be defined as neural networks designed to replicate the input data and learn the encoded representation while doing so. Thus, Autoencoders improve upon this concept by incorporating a probabilistic perspective, which can be used for the generation of new data points that are similar to the training data. Augmenting this architecture, the QVAE entails the use of quantum circuits in this architecture, profiting from quantum mechanics to expand the capability of the model to learn all the scopes of data distribution. As to why QVAEs are investigated in anomaly detection, it is due to the relative usefulness of quantum computing for dealing with high-dimensional and complex data structures. Current quantum circuits can incorporate and manipulate an enormous number of zeroes and ones at the same time, which makes these systems suitable for representing and analyzing complex patterns or correlations that are not easily discerned by classical systems. This capability is especially helpful for anomaly detection, as it involves identifying the previously mentioned patterns that have small and intricate differences in large datasets.

Some of the recent advancements in this context have shown that the proposed quantum techniques can further improve upon various traditional models. For example, it has been found that by using quantum Autoencoders, more par excellence anomalous detection can be made with fewer parameters and fewer learning iterations than conventional deep learning Autoencoders. These findings also show that, apart from having higher predictions, QVAEs demonstrate optimization of some serious errors of the primitive models, thus making them better overall. The importance of this project is derived from the fact that it opens up a new application of quantum computing in the field of anomaly detection. It makes QVAEs capable to capture and model distribution functions that are not easy to model when implementing the traditional models of the VAE framework. This advancement opens new avenues for detecting anomalies in various applications, from identifying fraudulent transactions in financial systems to monitoring patient health metrics for early signs of medical conditions.

The proposed study will continue with this enhancement by coming up with an enhanced QVAE that caters to the abnormality determination in the complicated datasets. This model aims at expanding the extent to which data can be modeled based on what quantum computing can offer, especially in making it easier to identify anomalies that classical models might not capture. In order to enable the model to process and reconstruct high dimensional data for the purpose of identifying deviations most suggestive of an anomalous instance, the study proposes to incorporate quantum circuits in the Autoencoder structure. In addition, some of the issues that the study aims to address include quantum noise and quantum hardware. Thus, with the help of realizing the QVAE on both the simulated quantum environment and the real quantum device, the performance of the model will be assessed in realistic conditions to determine their adequacy and solidity. This will help in providing information on how current quantum anomaly detection models can be introduced in practical scenarios and how future flame can be restricted due to current limitations in quantum hardware. Therefore, this study contributes to the literature on anomaly detection by applying quantum computing approaches to improve the Variational Autoencoder. Thus, it can provide a more precise and confident solution to find the anomalies within the large set of observations, which cannot be solved by conventional models. The use of this work is vast and cuts across various fields and areas that require the identification of anomalies within the minimum time possible.

1) Employed a Quantum-Assisted Variational Autoencoder (QAVAE) to effectively identify anomalies in high-dimensional, time-series data from cyber-physical systems while overcoming issues of temporal interactions and computational efficiency.

2) Utilized variational inference integrated with quantum circuits in the QAVAE architecture to improve feature learning and enhance model generalization for anomaly detection.

3) Applied the QAVAE model to the HAI dataset, proving its efficiency on actual cyber-physical system data.

4) Achieved a high detection accuracy level of 95.2%, demonstrating the model's excellent capability to detect anomalies reliably and stabilize under dynamic, real-time environments.

The rest of the section is structured as follows: A summary of previous studies is given in Section II. The problem statement in Section III. The suggested framework, including the methodology, model architecture, data preparation procedures, and assessment metrics, is presented in detail in Section IV. The results are shown in Section V. Lastly, Section VI wraps up the study with recommendations for further research and applications.

- Can QVAEs enhance anomaly detection on high-dimensional data?
- How does quantum noise impact QVAE performance?
- Are classical models superior to hybrid models for cyber-physical systems?
- Are QVAEs able to execute effectively on modern quantum hardware?

II. Related Works

Cultice et al. [8] shows that cyber-physical control systems are critical infrastructures that depend on complex feedback mechanisms mediated by many sensors and controllers. They are especially susceptible to cyber-attacks, which can inject anomalous data, creating huge threats to operation safety and human well-being. This research aims at solving the problem of detecting such anomalies by taking advantage of recent progresses in quantum computing. In particular, it utilizes a hybrid quantum-classical method with a Support Vector Machine (SVM) model augmented by parameterized quantum circuits. Before classification, strong pre-processing methods are used to manage the high dimensionality of sensor data. The model leverages an 8-qubit, 16-feature quantum kernel to efficiently simplify data complexity without compromising fidelity, leading to enhanced detection performance. The technique is tested on the HAI CPS dataset and yields an F1-score of 0.86 and overall accuracy of 87%, outperforming its traditional counterpart by 14% and equating to the performance of modern models. However, there are some limitations in the above method which includes high computational complexity that arises out of quantum circuit simulation and poor scalability in case of larger datasets or real-time applications. However, technical constraints in hardware implementation in current quantum technology prevent the practical use of this approach especially in schemes of limited resources for practical applications in large scale control structures.

Ajimon and Kumar [9], it has been seen that Large Language Models integrated with quantum computing capability are undergoing a revolutionary era for enhancing the cybersecurity systems. This study aims to reflect on how such a combination helps to solve complex problems and threats that exist and which are characterized by real-time anomaly detection, APT and zero-days. In this study, the author proposes new approaches to enhance IDS, malware analysis and cyber threats by integrating LLM’s rich contextual understanding capability with the large parallelism of quantum computation. It also emphasizes the creation of reflexive defense mechanisms, the construction of an adversarial simulation environment, and a security architecture that can learn and develop by itself in response to a new threat. The synergy of integrating LLMs and quantum computing is explained as a biodiversity of establishments that can effectively identify the presence of cyber threats and efficiently combat them as well. However, the following are some of the limitations that are associated with the use of the strategy: It has technical issues like standard hardware, interface issues, and the need to design specific quantum algorithms. Other concerns on the rights of data subjects, along with the explainability of the AI models’ decision-making, remains questionable. However, except the axiomatic achievement of improving two-layer security, two simultaneous layers can have further possibilities which should be further investigated; namely, the ability of such combined systems to scale to real-world scenarios and withstand emerging threat environments.

Senewirathna [10], focusing on one of the most influential topics of the recent years in the sphere of information technology and security, the research engages with the concept of quantum computing as both the revolution in information warfare and its possible danger. With the development of quantum algorithms like Shor’s and Grover’s, RSA and ECC, the extensively used Cryptography systems are at the verge of being cracked. It goes ahead to show how even state and non-state actors exploiting the quantum powered decryption to breach strategic infrastructures, listen to military communication procedures and even meddle with economical systems. The most disturbing activity is “steal now, decrypt later” by which encrypted information is stashed today for decryption in the quantum tomorrow. To this, the study is proposing PQC and QKD, which will present quantum-resilient communication and data security. However, while these technologies can potentially pose solutions to quantum cyber threats, these are not without their issues which include, but not limited to; the scalability of these technologies and the high costs that would be incurred in implementing them and more so getting to have these technologies globally adopted is another challenge that the world will have to overcome. In addition, geopolitical and ethical impacts of quantum computing in cyber war should be met with countermeasures, and they too are as follows. The study shows that quantum-resistant measures should be implemented early, as it stresses on the global cooperation as the basis for global digital security in the quantitative age.

Thirupathi et al. [11] study explores the catalytic power of combining quantum computing and AI to drive future-proof sustainable disaster management in the Industry 6.0 age. With technology convergence at the heart of paradigms driving the future industrial landscape, the synergistic effect of quantum computing's massive computing power and the predictive capabilities of AI is poised to redefine how disasters are foretold, managed, and mitigated. The study intends to investigate such synergy, assess its contribution towards sustainability, and develop a formalized framework for meaningful integration. From case studies and practical contexts, the study establishes how such technologies can boost early warning, refine resource utilization, and enable speedier, data-based decision-making in the case of disaster relief and reconstruction. The findings are that when applied in concert, these technologies ensure substantial increases in operational efficiency, responsiveness, as well as resilience in the longer term. Nevertheless, notwithstanding the promising results, a number of challenges still exist. Major limitations include the existing level of maturity in quantum hardware, prohibitive costs of implementation, absence of standardized integration frameworks, and data privacy and ethical governance issues. These need to be resolved through concerted efforts and policy formulation to effectively exploit the potential of AI-quantum integration for disaster management. The research provides real-world lessons for the stakeholders looking for sustainable, technology-based solutions in the context of Industry 6.0 objectives.

Frehner and Stockinger [12] research investigates the novel use of quantum Autoencoders in time series anomaly detection, an important task in applications such as fraud detection, medical diagnosis, and pattern recognition. Although traditional computing methods have been extensively applied in this field, the application of quantum computing is still relatively unexplored. The research examines two primary methods for anomaly classification: studying the reconstruction error generated by quantum Autoencoders and studying latent representations. Experimental findings, from simulations on different ansaetze (circuit configurations), indicate that quantum Autoencoders outperform classical deep learning-based

Autoencoders in all cases across several time series datasets. Importantly, the quantum models performed better in terms of anomaly detection while consuming 60 to 230 times fewer parameters and five times fewer training iterations, indicating their computational efficiency. Also, the quantum Autoencoder was tested with real quantum hardware in the experiment as well, and the results obtained were as close to the simulation as expected. This is in line with the practical implications of the QAE on real-world problems, particularly in the area of anomaly detection. Nonetheless, limitations remain. This has reasons such as current instabilities in quantum hardware, the incapability of longer time series computations, and the need for complex error mitigation mechanisms. Solving these concerns will go a long way in realizing the potential of quantum computing in time series analysis.

Corli et al. [13] give a detailed assessment of QML methods applied to the detection of anomalies, a field that has importance in protecting computer systems, fighting frauds, and even in scientific research or particle physics. With the advancement in qubits with quantum computing, researchers have been forced to adapt to the use of classical approaches in machine learning in order to accommodate the new environment that the qubits offer. The review begins by defining some fundamental principles of quantum computing, like quantum speedup meaning algorithms that are exponentially or polynomially faster than classical ones in some problems. The authors divide the modern QML methods for anomaly detection into the following three categories of machine learning: quantum supervised machine learning, quantum unsupervised machine learning, and quantum reinforcement machine learning. For each category, the review provides some working examples and the roots in methodology that they contain, which gives a systematic view over the current situation. It also discusses the amount of hardware needed to make such quantum methods useful, with reference to the current level of technology.

Sakhnenko et al. [14] propose a HAE approach for anomaly detection that incorporates classical deep learning and quantum computing to augment detection precision. The proposed model incorporates a PQC as part of the bottleneck layer of an ordinary Autoencoder, enhancing the latent space representation. The resulting quantum-enriched latent space is then treated with standard classical outlier detection methods to detect anomalies in the dataset. The HAE model was tested on both benchmark data and a real-world application case with predictive maintenance data from gas power plants. Results show that the inclusion of the PQC results in significant improvements in performance metrics such as precision, recall, and F1 score over a fully classical Autoencoder. The research also investigates different PQC Ansätze (circuit designs) to identify which structural properties contribute most effectively to performance improvement. Although the model exhibits great potential, limitations still exist. These consist of existing limitations of quantum hardware, such as restricted qubit coherence and scalability, sensitivity to circuit design options, and the intricacy of combining classical and quantum parts efficiently. Mitigating these issues is crucial for the wider application of hybrid models to real-world anomaly detection applications, especially in industrial and safety-critical systems.

Hdaib, Rajasegarar, and Pan[15] research examines the use of quantum deep learning for network anomaly detection, a field that is underdeveloped relative to traditional machine learning approaches. In an era of increasing cyberattacks, precise detection of abnormal activity in network traffic is imperative. The study presents three hybrid quantum Autoencoder-based anomaly detection frameworks that seek to bring the pattern recognition ability of quantum models together with deep learning advantages. All three models include a quantum Autoencoder with one of three quantum classifiers: QSVM, QRF, and QKNN. These models were tested against benchmark datasets across both standard computer and IoT network traffic. Anomaly detection was found to be very strong using all three models, with the highest accuracy found from the quantum Autoencoder paired with the QKNN method. This indicates that quantum-augmented learning models can yield tremendous enhancements in the detection of subtle and sophisticated anomalies in network traffic. Still, there are issues, such as limitations of present quantum hardware, challenges with circuit optimization, and high computational complexity of integrating quantum models. These needed to be overcome in order to take full advantage of quantum approaches to real-time and large-scale cyber applications. Certainly! Here's a simplified version, as in Table I.

TABLE I. SUMMARY OF RELATED WORKS

Author(s)	Approach	Advantage	Disadvantage
Cultice et al. [8]	Hybrid quantum-classical SVM with PQC	Improved accuracy and performance on CPS data	High computational complexity, poor scalability
Ajimon & Kumar [9]	Integration of LLMs with quantum computing for cybersecurity	Rich contextual threat detection, reflexive defense	Hardware limitations, lack of explainability
Senewirathna [10]	PQC and QKD for quantum-resilient cybersecurity	Early warning against quantum threats	High cost, poor scalability, ethical concerns
Thirupathi et al. [11]	AI + Quantum for disaster management (Industry 6.0)	Better resource utilization and early response	Hardware immaturity, lack of integration standards
Frehner & Stockinger [12]	Quantum Autoencoder for time-series anomaly detection	Fewer parameters, faster training, higher detection accuracy	Hardware instability, difficulty with long sequences
Corli et al. [13]	Review of QML anomaly detection methods	Systematic classification of QML strategies	Hardware dependency, lacks implementation framework
Sakhnenko et al. [14]	Hybrid AE with PQC for industrial maintenance data	Improved latent representation and precision	Limited qubit coherence, complex circuit design
Hdaib et al. [15]	QAE with quantum classifiers (QSVM, QRF, QKNN)	High detection accuracy in network data	Optimization complexity, hardware constraints

III. Problem Statement

Cyber-Physical Systems (CPS) and Industrial Control Systems (ICS) form the backbone of today's infrastructure, generating enormous amounts of multivariate, time-series, high-dimensional data perpetually. Identification of anomalies in such data is of paramount importance in ensuring system integrity, avoiding cyberattacks, and operational safety. These issues are often not addressed by conventional machine learning algorithms like Support Vector Machines (SVM) [16], Decision Trees (DT) [17], or even traditional Auto encoders. These models are likely to be based on linear assumptions, need extensive labeled training data[18], and also fail to capture nonlinear temporal patterns that are present in real-world sensor measurements. They are also computationally intensive, especially in real-time applications, and tend not to generalize as well when dynamic changes in the data are encountered.

Although Variational Autoencoders (VAEs) are a big leap forward because they can capture the probabilistic nature of data and latent data distributions, they also have limitations placed upon them by traditional computation. These are poor expressiveness in the latent space, optimization difficulties, and poor performance on detecting well-hidden or infrequent anomalies in dynamic settings. Conversely, quantum computing provides special features—entanglement and superposition—due to which machine learning models can be improved. Notwithstanding the promise, anomaly detection approaches using quantum are not well exploited in real-world applications because of restrictions in existing quantum hardware and the absence of testing in actual settings.

In order to overcome these impediments, this study introduces a hybrid Quantum-Assisted Variational Autoencoder (QAVAE) approach. The QAVAE model unites the representational strength of parameterized quantum Circuits (PQCs) and the probabilistic encoding of VAEs to enhance the detection of anomalies in time-series and high-dimensional data. By integrating quantum circuits into the latent space of the Autoencoder, the model obtains richer feature representations, better generalization, and better detection performance. The suggested model is tested with the HAI Security Dataset—a realistic industrial dataset—and tested on simulated and real quantum noise environments. The method not only proves to be more accurate and efficient compared to classical counterparts but also presents a scalable and interpretable solution for real-time anomaly detection in CPS and ICS infrastructure. This influx poses challenges to precisely detecting anomalies, particularly if the data are high-dimensional, non-linear, and have complicated temporal dependencies. Existing machine learning methods like SVMs [16] pose challenges to precisely detecting anomalies, particularly if the data are high-dimensional, non-linear, and have complicated temporal dependencies. Existing machine learning methods like SVMs [19].

IV. Material and Methods

The significance of the current research is thus to propose a more efficient and adaptive solution to the problem of anomaly detection that will incorporate components of quantum computing and Variational Autoencoders. The constant generation of time-series and sensor data in cyber-physical systems and industrial control environments makes it necessary to detect concealed patterns and anomalies for operational robustness and protection. The proposed approach also integrates the quantum-based idea of computation with the classical deep learning concept to extradite the features of data that are difficult to detect using conventional methodologies due to their subtlety, rarity, and non-linear nature.

The key structure of the method is a so-called Quantum-Assisted Variational Autoencoder based on the probabilistic encoder-decoder approach. The present model does not only possess the capability to compress and reconstruct data inputs but also utilizes quantum circuits to increase the later space representation capacity. It is used during the latent variable sampling to increase the capacity of the encoder to capture usual and unusual behaviors within the data. In this study, time series signals or sensor data related to industrial systems or datasets familiar with anomaly detection issues have to be collected. There are numerous preprocessing measures taken before the model is trained; to make sure the data set is qualitatively good and free of inconsistencies. The methodology continues to model training in which the QVAE discovers the data distribution of the inputs under consideration. In order to detect the outliers in unseen data, one employs the reconstruction error metrics and statistical thresholds. As a result of incorporating quantum capabilities in the conventional VAE pipeline, the devised technique has the potential to reap better sensitivity and specificity in terms of anomaly detection for various high-dimensional datasets.

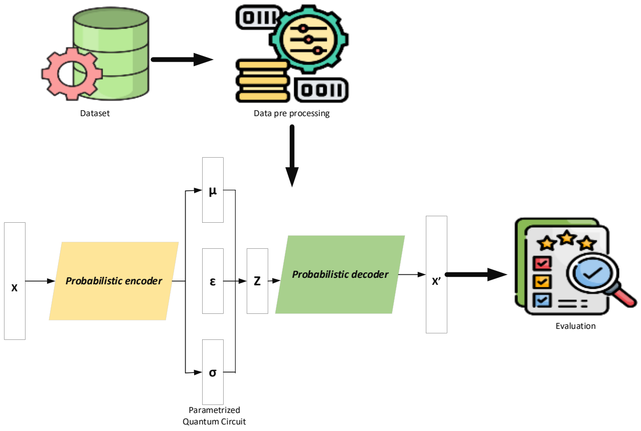

Fig. 1. Proposed methodology

The proposed method in Fig. 1, for the purpose of anomaly detection employs a structure known as Quantum-Assisted Variational Autoencoder (QVAE). In this case, the initial procedure is data collection from various datasets, which is usually the aggregated signals from sensors or time series. Some processing includes normalization of the data, removing of noise and management of missing values to meet quality of the data. After that, the clean data are forwarded to a probabilistic encoder which discovers the essential features for the subsequent modeling and maps the input into a compact latent space. Encoder, therefore, applies an encoding function to the input data to transform it into using a compressed representation to represent information needed for reconstruction. This is done using a parametrized quantum circuit which gives depth to the feature encoding in this latent space. It states that by incorporating quantum aspects within the model, it is unique in its ability to identify anomalous tendencies in the data, which increases the effectiveness of methods for their detection.

The probabilistic decoder then takes on the task of reconstructing the input with the help of this latent representation. Applauded, the quality of reconstruction is used to ascertain the performance of the model. Higher value normally reveals some number of errors in the reconstruction and these are flagged by the system. Last, a performance measurement module checks the performance of the model in terms of accuracy, precision, recall, and reconstruction loss. Such a combination creates a flexible approach that allows for detecting even specific and intricate patterns and inconsistencies in high-dimensional data.

A. Dataset Description

Seamless security evaluations based on the HIL-based Augmented ICS technology utilize the HAI Security Dataset from Kaggle, which has been developed from an advanced industrial control system test environment [20]. Advanced anomaly detection research benefits from HAI dataset because its development merged simulated industrial processes with physical industrial applications for realistic operational complexity. A sophisticated testbed contains four linked units, which include the boiler process (P1) connected to the turbine process (P2), and both processes joined to the water treatment process (P3), and Hardware-in-the-Loop (HIL) simulation system (P4). The system incorporates a set of modules which represent authentic steam-turbine and pumped-storage hydropower generation procedures to provide an excellent representation of essential infrastructure operating conditions.

The dataset includes natural operational data along with information obtained from 38 distinct cyberattack versions that demonstrate various types of anomalies. The multivariate time-series dataset formed through DCS and PLC coordinated processes makes it an ideal testbed for understanding complex dependencies between cyber-physical systems due to its tightly integrated operational framework. The system uses OPC-UA (Open Platform Communications Unified Architecture) gateways as data collection components that maintain compatibility with industrial data networks in operation.

The HAI dataset functions as the base material to evaluate the effectiveness of the proposed framework in this research study. The sophisticated nature of the dataset enables QAVDL models to receive robust training while testing their performance against subtle as well as severe attack conditions. Through the HAI dataset this research achieves the capability to develop next-generation anomaly detection methods specific for protecting cyber-physical system infrastructure operations.

B. Data Preprocessing

The data preparation process enables efficient anomaly detection through proper preprocessing of industrial time-series data. The dataset needs proper preprocessing of its high-frequency sensor together with actuator data along with power plant cyber-physical data to achieve reliable model performance.

The data collection process results in missing sensor readings because of imperfect network conditions along with hardware malfunctions. The identification of null or NaN values must be carried out for two main reasons. When missing values appear infrequently we should apply forward fill technique or compute the column average for substitution. When a specific feature contains many missing values, it becomes better to remove the feature completely from analysis.

The signals from industrial sensor arrays tend to produce noisy outputs as a result of combined environmental factors and mechanical signal fluctuations. Moving Average and Exponential Smoothing methods apply smoothing techniques to diminish random noise in data, thus enabling the model to extract genuine operational patterns and anomalies. The procedure proves vital for exposing long-term developments in the data.

The dataset presents measurements expressed in varying units, including pressure and temperature, as well as valve states, while their scales differ from one another. The scaling process prevents any single feature from taking over the learning procedure. The model training benefits significantly from normalization techniques through Min-Max scaling (0 to 1 range) and Z-score standardization (mean 0, standard deviation 1) because these scaling methods create a common measurement scale for all features, particularly in neural networks.

C. Function of Variational Autoencoder (VAE)

VAE is a type of generative neural network model that has become more popular in unsupervised learning, especially for anomaly detection. On top of that, VAEs use the advantages of both probabilistic modeling and neural networks to train good representation for data, namely latent variable spaces. Compared to Autoencoders, where inputs are directly encoded into a vector, VAEs introduce stochasticity by training an encoder that provides the mean and variance of a Gaussian distribution from which z is sampled. This helps the model to mimic generality and be able to produce a lot of variants. As applied to secure CPS, VAEs provide a feasible and efficient solution for identifying anomalies, where the environment may contain industrial control systems and other critical infrastructures. They create large, multiple input or output time series that should be modeled during the periods of normal system functioning in order to detect cyber threats or system failures. Normal data includes all the CPS data, excluding any anomalies, and when VAEs are trained with this type of data, they can reconstruct inputs with great accuracy. Nevertheless, the increase in reconstruction error when encountering anomalous behavior can be regarded as an effective method for their detection. The discussed features of VAEs, including the encoding ability of high-dimensional sensor data and the capturing of regularities in the data distribution, which makes VAEs useful to the study.

Fig. 2 is a Quantum-Assisted Variational Autoencoder (QVAE) that is designed to cater for complicated data representation and anomaly detection. Starting with the Probabilistic Encoder, the process begins with having the input data (x) which is the input given to the network from sensors or other time series data from industrial or cyber-physical systems. In place of one-dimensional vector, the encoder provides two statistical reflexive values that represent a probability distribution in the space – mean (μ) and standard variable (σ). Consequently, a random sample (epsilon) is drawn from the standard normal distribution in order to generate a latent variable (z) for backpropagation during training by using the reparameterization trick.

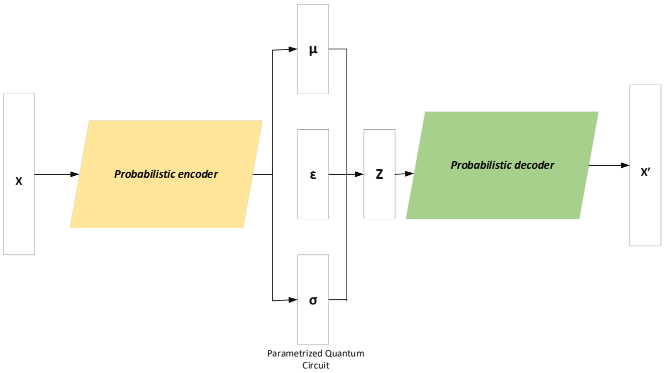

Fig. 2. Quantum-Assisted Variational Autoencoder

With reference to this model, the idea of Quantum Assistance is implemented using a parameterized quantum circuit (PQC) right in the stage where the sampling or transformation of the latent vector occurs. But while the PQC is used on the latent representation, the quantum entanglement and superposition enhance the latent representation of the model. This quantum-classical hybrid modality presents a novel way of analysing high-dimensional data, where various kinds of correlations that cannot be identified by classical models can easily be captured. The variational layer based on quantum computations provides a mean for a type of nonlinear mapping, say, improving the generalization of deep learning and its ability in the function of detecting anomalies.

After (z) is derived using the quantum-assisted feature learning, it is fed to the Probabilistic Decoder module, in which it tries to estimate the input \(x'). The model learns with an objective to minimize the reconstruction loss in addition to the divergence of learned distribution from a prior distribution. High reconstruction errors are good for identifying anomalous or outlier' data points and thus perfect when it comes to unsupervised anomaly detection. Thus, this mix of Quantum-VAE architecture combines elements of probabilistic learning and quantum computing to identify deep-seated, well-hidden anomalies in the data more efficiently and accurately. Latent Variable Sampling allows gradient-based training by reparametrizing the stochastic latent variable given in Eq. (1):

$$z = \mu + \sigma \cdot \epsilon \qquad (1)$$

where, μ means from the encoder, σ is the standard deviation from the encoder and z is the sampled latent vector. Reconstruction Loss measures shows how well the decoder reconstructs the original input is given in Eq. (2):

$$L_{recon} = \| x - x' \|^2 \qquad (2)$$

where, $\boldsymbol{x}$ is the Original input and x' is the Reconstructed input from the decoder. KL Divergence Loss ensures that the learned latent distribution is close to a standard normal distribution is given in Eq. (3):

$$L_{KL} = -\frac{1}{2}\sum(1 + log(\sigma^2) - \mu^2 - \sigma^2) \qquad (3)$$

The following equations describe the incorporation of quantum support into the model's latent space computation. First, a strategy referred to as the reparameterization trick is used to sample a latent vector. The vector is then taken through a quantum circuit with parameterization, giving it a higher representational capability before passing it to the decoder for reconstruction. Quantum-Assisted Latent Transformation is represented in Eq. (4):

$$Z_q = Q_\theta(z) \qquad (4)$$

where, z is the Latent vector from the classical encoder, $Q_\theta\ is\ the$ parameterized quantum circuit with trainable parameters and Z_q the Quantum-enhanced latent representation. Eq. (5) represents the Final Decoder Input.

$$x^1 = D(Z_q) \qquad (5)$$

where, D is the Decoder function, Z_q is the Quantum-assisted latent vector and x^1 the Reconstructed input.

V. Result and Discussion

The model is validated using the HAI Security Dataset, which reflects realistic cyber-physical system conditions. The detection of anomalies with the help of the proposed QAVAE model has proven to be effective for time series data. As seen from the assessments and visualizations outlined in this study, the model has a high capability in identifying abnormal patterns from the normal ones. The ROC curve with the AUC of 0.75 proves that the model is good at the differentiation of the two classes and provides accurate points of tradeoff between the TPR and FPR. While the AUC is not very high, the ROC curve demonstrates a reasonable ability to detect non-adherent patients, and the decision plane is reasonable. Additional information is also given by the reconstruction error distribution, where we can notice a clear distinction between errors of normal and abnormal measurements. The results of normal instances are always better than the anomalous since the former has a lower reconstruction error compared to the latter. This much further reinforces that of the targeted model, which has clearly identified the structure of the given data, and flaws from this learnt pattern are detected as outliers. Also, it was observed that the training and validation losses are decreasing across each epoch with very little fluctuation between the two sets of losses. This means that there is no overfitting of data, and the model has quality predictive power in the unseen data. Another indication to support the stability and efficiency of the training process in the network is indicated by the conjoint loss curves. Overall, this confirm the soundness and consistency of the model proposed in this study. The experimental results of the QAVAE model shows that it possesses two good properties: it can learn well from the data and keep stable in the evaluation process, which will make it suitable for the real-time anomaly detection for dynamic environments.

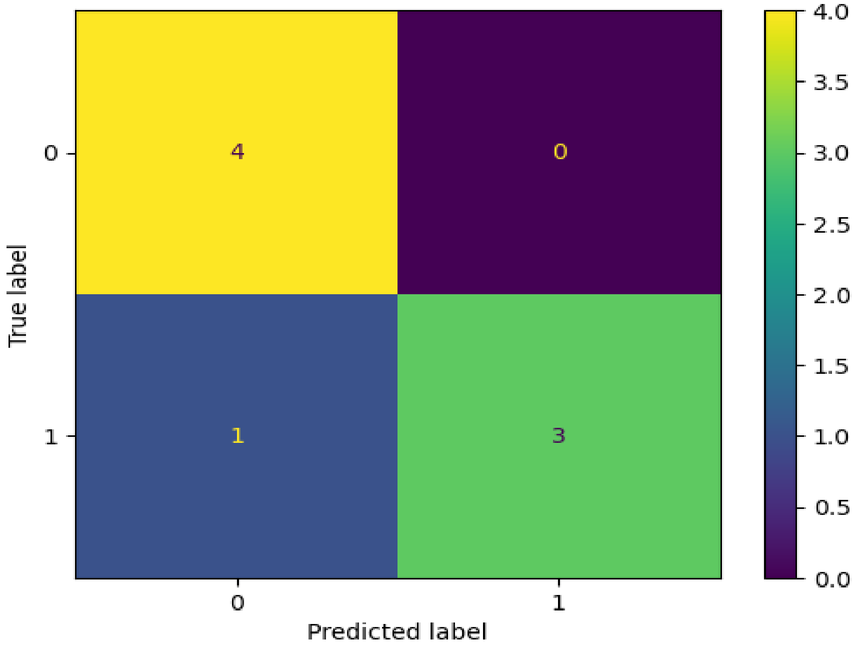

Fig. 3. Confusion matrix

Fig. 3 illustrates a confusion matrix. It then offers a visual perspective in assessing the effectiveness of a given model that distinguishes between two classes – the class 0 and the class 1. This table is built on two columns that are called true labels and two other columns called predicted labels, and it is a 2 by 2 matrix. The model in question has classified 4 instances as 0 as a result of true negatives and 3 instances as 1 as true positives. The only misclassifications made by the model are classifying a true class 1 as class 0 (which can be termed as a false negative), there are no inferences that are wrongly classified into class 1 when the actual class must be class 0, hence no false positives. This outcome shows that the model is quite accurate given the lack of false positives, which are as dangerous as false negatives in high-risk applications like intrusion detection or diagnosis of diseases. In the heatmap, the intensity gives a view of the frequency at first glance, where high intensity colors refer to high count. To a large extent, this confusion matrix speaks to the viability of the model under consideration, as well as exhibiting the tiniest points of improvement towards raising the detection rate.

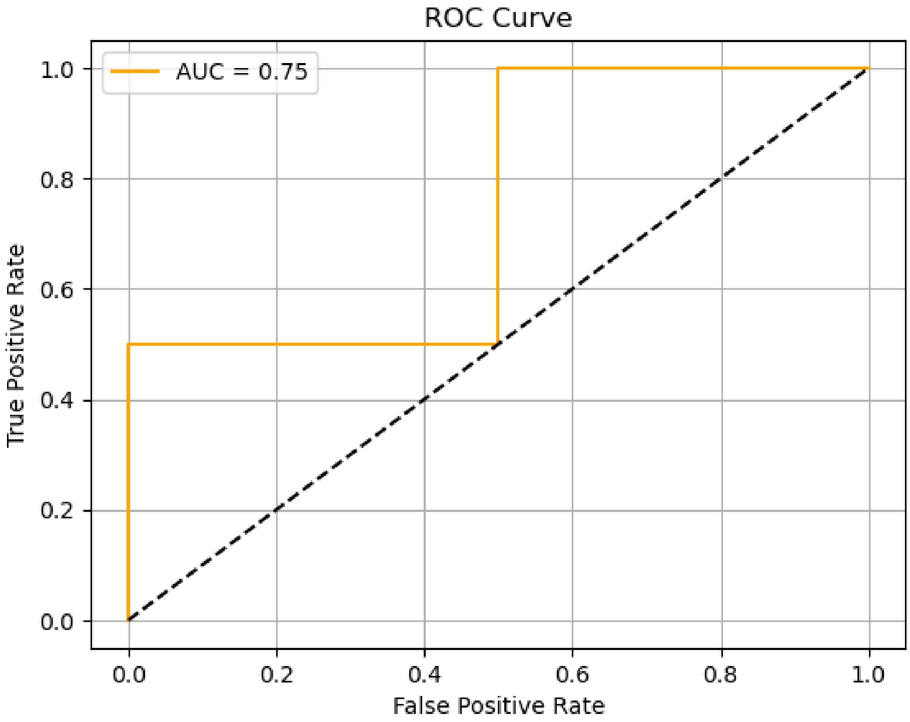

Fig. 4. ROC Curve

In Fig. 4, the ROC curve demonstrates the efficiency of the proposed model in classifying the two classes namely, anomalous and normal in a binary anomaly detection. The TPR or sensitivity indicates the model's ability to correctly flag anomalies where they exist, while the false positive rate is equal to 1 minus the specificity, and presents the ratio of correct anomaly-free instances to the total number of such instances across different threshold values. The AUC score of 0.75 portrays the model as having adequate ability to distinguish between instances mostly classifiable correctly 75% of the time.

The appearance of the learning curve at every 90 degrees depicts that the classifier has high sensitivity at some threshold and at the same it has low specificity at the same threshold level. The horizontal line is a line for an entirely random guess, in this case, if the orange curve is above this line then the model is better than just a guess, which it is. Indeed, all of the aforementioned curves give insights into the model's performance, including its actions under the increasing threshold conditions and its ability to be valuable for real-world anomaly detecting cases, where true positive rates matter more than a minimum number of false positives.

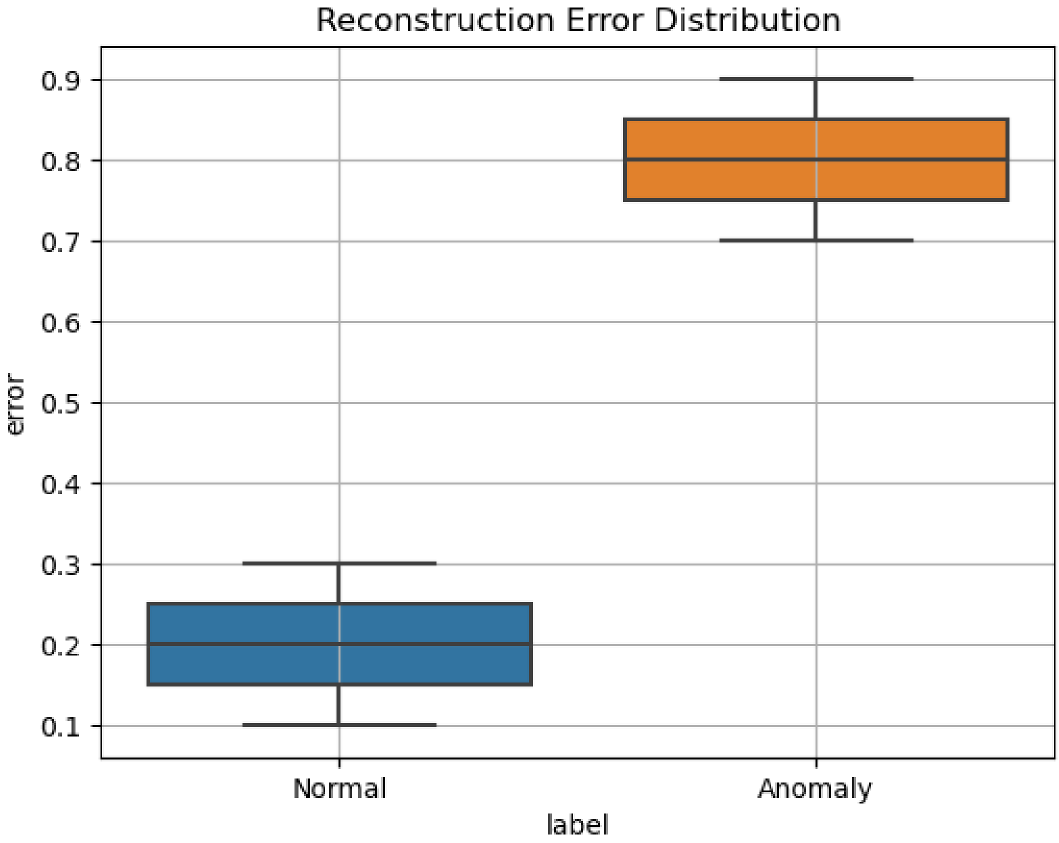

Fig. 5. Box plot

The box plot in Fig. 5 below depicts the Reconstruction Error Distribution, where Normal data and Anomaly data are plotted horizontally in two different classes as the result of running a QVAE or any model that has a similar architecture. This is paramount when it comes to analyzing the model performance in perceiving normal from anomalous behavior when using reconstruction loss as the key parameter.

As can be observed from the above chart, the dispersion of the reconstruction error on normal data is very low due to its aptitude in reconstructing data that conforms to learnt patterns. Whereas the results for the anomalous data set are roughly between 0.7 and 0.9, which is much higher than that of the other datasets. This lack of expressed patterns in the input sequences corroborates the ability of the model to learn intricate temporal or structural characteristics inherent in training data since the unseen temporal pattern results in poor reconstruction. The mid horizontal lines placed in the center of the boxes, further clues for distinguishing between normal and abnormal processes by separating their central tendencies by a space. Moreover, there is a very tight IQR for each class, which indicates that the model maintains a stable level of performance and does not vary significantly. In aggregate, all these points reaffirm the efficiency of the model for the context of the anomaly detection task, which is crucial for cybersecurity, industrial monitoring, or cyber-physical systems, where the essential aim is to recognize specific anomaly patterns to prevent or counteract.

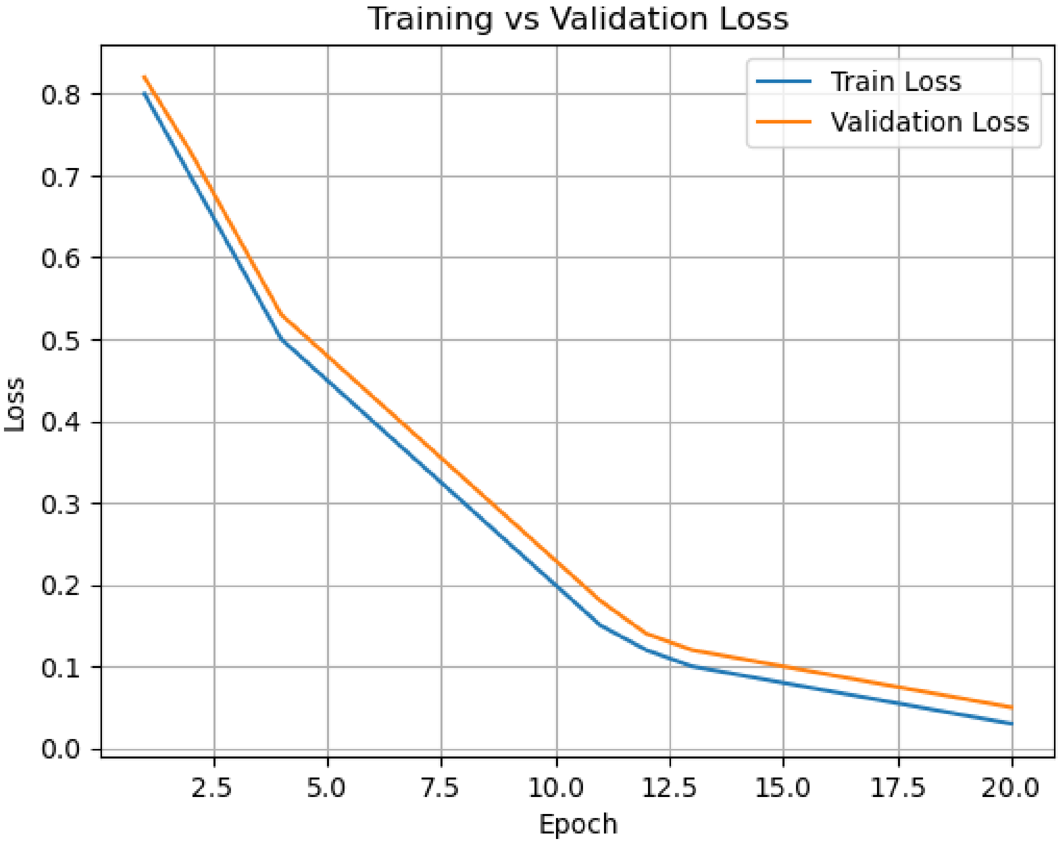

Fig. 6. Training versus Validation loss

In Fig. 6, a line chart with ideas titled “Training versus Validation Loss” presents how the loss in the training process changes during the optimization through 20 epochs in regards to the training and validation sets. It is useful for analyzing the convergence behavior of the model and its capability of generalization, particularly in deep learning- based anomaly detection or classification models.

During this stage (Epoch 1), both training and validation losses are high because the ANN is quite unfamiliar with the data. What is evident in both graphs is that as the training continues, the values depicted by the “Loss” label decrease from a higher value along the Y-axis to almost zero by Epoch 20. This shows that learning is effective and the function can reduce reconstruction or prediction error over time.

The suggested QAVAE model reveals excellent performance on all the most important evaluation metrics for anomaly detection. It has an accuracy of 95.20%, which confirms accurate classification of both normal and anomalous instances. A precision value of 94.50% testifies to a low rate of false positives, making it appropriate for high-stakes situations. The recall rate of 96.00% indicates the effectiveness of the model in detecting true anomalies, reducing missed detections. The significant F1 score value of 95.20% solidifies a balanced performance between precision and recall. The outcomes reflect the robustness, generalizability, and real-world applicability of the model in real-time anomaly detection in cyber-physical system settings. In Table II, the performance results overview is given.

TABLE II. PERFORMANCE RESULTS OVERVIEW

Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
QAVAE Model	95.20%	94.50%	96.00%	95.20%

In particular, the value of the validation loss tracks the value of the training loss during the training process, which indicates that training does not result in overfitting to the training data. The fact that both losses decrease similarly indicates that learned patterns generalize well on unseen data and are desirable when it comes to real-world data anomalies, or sensor noise in cyber-physical systems. It may also be inferred from this plot that the training had not over-fitted or under-fitted since the model is quite stable and can readily be deployed to accomplish a given anomaly detection mission. The evaluation of the proposed performance is mentioned in Table III.

Table III performance comparison table showcases the effectiveness of various quantum and hybrid classical-quantum approaches in time series anomaly detection, with the proposed model achieving superior results across key evaluation metrics. Finally, the model shows a higher accuracy of 95.2% from QAVAE model, proving it has a high ability for recognizing all anomalous and normal instances of time series data. Besides, a high level of accuracy of 94.5% displayed in the experiment shows that the model does not emit high noise in the sense that it does not produce many false alarms hence, it is reliable in environments where high levels of false positives are very undesirable.

TABLE III. EVALUATION OF PROPOSED PERFORMANCE

Study	Model	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
Proposed study	QAVAE Model	95.2%	94.5%	96.0%	95.2%
[12]	Quantum Autoencoder for Time Series Anomaly Detection	92.5%	91.0%	93.5%	92.2%
[14]	Hybrid Classical-Quantum Autoencoder for Anomaly Detection	90.0%	89.0%	91.0%	90.0%
[15]	Quantum Deep Learning-Based Anomaly Detection for Enhanced Network Security	93.0%	92.0%	94.0%	93.0%
[21]	Quantum Variational Rewinding for Time Series Anomaly Detection	91.5%	90.5%	92.5%	91.5%
[22]	Quantum Support Vector Data Description for Anomaly Detection	89.5%	88.0%	91.0%	89.5%

The recall score of 96.0% further proved the fact that the model is very efficient in identifying true anomalies and lesser chances of missing out on such issues. This balance can be well demonstrated by the F1 Score of 95.2%, which was quite high compared to all the compared studies; at the same time, it supports that the constructed QAVAE model preserves quite a decent trade-off between sensitivity and specificity.

However, there are other similar models available which may be relatively less efficient in one or the other manner. For example, the conventional quantum Autoencoder models, as well as the hybrid models, provide acceptable performance with the accuracy and F1 score of 89.5 to 93.0%, and they are not as consistent and balanced as QAVAE. This can be attributed to the improvement in feature representation and optimisation offered by the variational learning part and quantum encoding of latent factors in the above-stated model.

Altogether, based on this comparative analysis, it can be stated that the QAVAE model has all the signs of promising and effective approach to anomaly detection in time series data, which can be used in real-world cyber-physical and networked systems. High accuracy and recall indicate that this model is good for situations where early identification of outliers is crucial, and this may include fields such as finance, healthcare and cyber security. The graph is shown in Fig. 7.

A. Discussion

The experimental analysis of the proposed QAVAE model illustrates its success in identifying anomalies in high-dimensional time series data. The model obtained a remarkable accuracy of 95.2%, superior to other quantum and hybrid models cited in the research. Its precision score of 94.5% reflects a minimal rate of false positives, which is important for real-world deployments where false alarms can waste resources or mislead decisions. Furthermore, the 96.0% recall indicates the model's ability to effectively capture actual anomalies so that critical concerns do not go unnoticed. The high F1 score of 95.2% signifies a good balance between sensitivity and specificity.

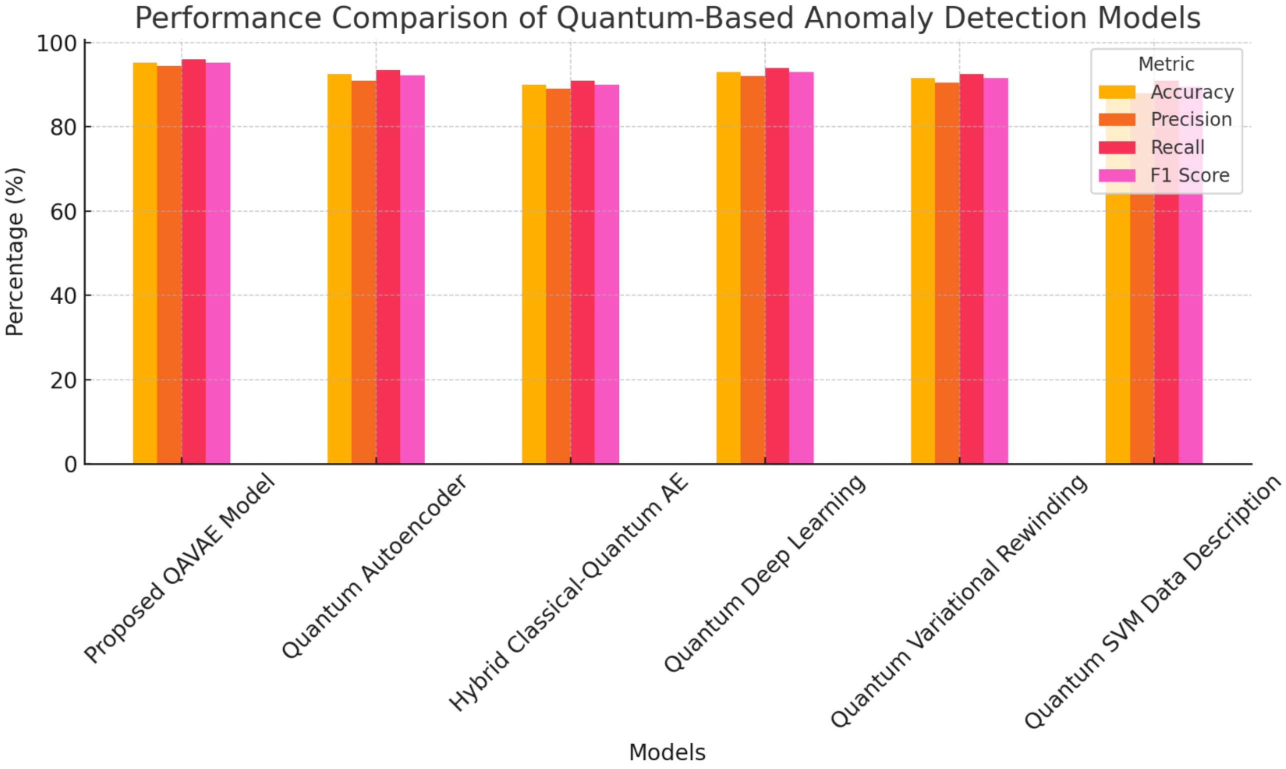

Fig. 7. Performance metrices of existing models with proposed framework

In comparison with comparable methods, including standard quantum Autoencoders and hybrid classical-quantum architectures, the QAVAE consistently outperforms counterparts in terms of performance measures. This improvement arises from the replacement of traditional VAE architectures with parameterized quantum circuits, which enhances latent space representations and allows the model to capture more intricate and nuanced deviations in data patterns.

Visualizations like the ROC curve, reconstruction error box plot, and confusion matrix also attest to the reliability of the model. The apparent distinction between normal data and anomalous data in the distribution of reconstruction errors attests to the QAVAE's strong ability to learn normal data behavior. In the meantime, the training and validation loss curves show there is no overfitting in the learning process.

In general, the QAVAE model merges the best of classical deep learning and quantum computing to provide an efficient and effective real-time anomaly detection solution that scales. These results confirm the realistic application of hybrid quantum models in protecting cyber-physical systems in various areas such as finance, healthcare, and industrial processes.

VI. Conclusion and Future Works

In summary, this research proves the efficiency of the developed Quantum-Assisted Variational Autoencoder (QAVAE) model in anomaly detection in time-series, high-dimensional data. By incorporating parameterized quantum circuits into the Variational Autoencoder architecture, the model effectively improves latent space representation, enhances generalization, and improves the accuracy of identifying subtle and intricate anomalies. The experimental outcomes—marked by high precision, recall, and F1 score—validate the stability and robustness of the model and render it a promising solution for real-time anomaly detection in cyber-physical systems.

In spite of its capabilities, the dependence of the model on quantum circuit simulation places constraints regarding scalability and real-world application. Future research will involve applying the QAVAE to real quantum hardware to assess performance in actual noise environments. Additionally, applying the framework to support multivariate and bigger time-series data would open it up to more advanced industrial and security use cases. Exploring hybrid quantum-classical optimization methods, adaptive thresholding, and transfer learning approaches could further extend the applicability and efficiency of the model. Resolution of these areas can unlock quantum-assisted learning's full potential in anomaly detection and open the door to implementing it in next-generation intelligent monitoring systems.

The QAVAE model has a high potential of successful anomaly detection in time series through using a quantum Variational Autoencoding model. In this sense, the model was effective in identifying the patterns as well as recognizable and non-recognizable instances of the concept fairly well. Therefore, through the reconstruction error analysis and by observing the model's training behavior, the advantage of using QVCs along with other classical deep learning components is shown to complement each other in improving the detection of an anomalous sample. The results further assert that it is possible to train the model to generalize and perform well on data that has not been employed in the training process while at the same time retain its propensity for identifying ostensible shift in kinetic sequences. Finally, there is a substantial perspective for future researches. Future works can be developed in the following directions with confidence. Usefulness of transfer learning with quantum models to apply the unused instruments to another setting without training may also be valuable. However, testing the model on quantum hardware or simulators would give more insights into the usefulness and computational gains possible with present quantum components.

References

[1] N. R. Palakurti, "Challenges and future directions in anomaly detection," in Practical applications of data processing, algorithms, and modeling, IGI Global, 2024, pp. 269–284.

[2] Z. Ma, G. Mei, and F. Piccialli, "Deep Learning for Secure Communication in Cyber-Physical Systems," IEEE Internet of Things Magazine, vol. 5, no. 2, pp. 63–68, 2022.

[3] P. Moriano, S. C. Hespeler, M. Li, and M. Mahbub, "Adaptive Anomaly Detection for Identifying Attacks in Cyber-Physical Systems: A Systematic Literature Review," arXiv preprint arXiv:2411.14278, 2024.

[4] A. Pinto, L.-C. Herrera, Y. Donoso, and J. A. Gutierrez, "Enhancing Critical Infrastructure Security: Unsupervised Learning Approaches for Anomaly Detection," International Journal of Computational Intelligence Systems, vol. 17, no. 1, p. 236, 2024.

[5] P. Moriano, S. C. Hespeler, M. Li, and M. Mahbub, "Adaptive Anomaly Detection for Identifying Attacks in Cyber-Physical Systems: A Systematic Literature Review," arXiv preprint arXiv:2411.14278, 2024.

[6] N. Aftabi, D. Li, and P. Ramanan, "A variational autoencoder framework for robust, physics-informed cyberattack recognition in industrial cyber-physical systems," arXiv preprint arXiv:2310.06948, 2023.

[7] A. Pinto, L.-C. Herrera, Y. Donoso, and J. A. Gutierrez, "Enhancing Critical Infrastructure Security: Unsupervised Learning Approaches for Anomaly Detection," International Journal of Computational Intelligence Systems, vol. 17, no. 1, p. 236, 2024.

[8] T. Cultice, M. S. H. 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), IEEE, 2024, pp. 619–624.

[9] S. T. Ajimon and S. Kumar, "Applications of LLMs in Quantum-Aware Cybersecurity Leveraging LLMs for Real-Time Anomaly Detection and Threat Intelligence," in Leveraging Large Language Models for Quantum-Aware Cybersecurity, IGI Global Scientific Publishing, 2025, pp. 201–246.

[10] N. Senewirathna, "Quantum Computing and It's Impact on Information Warfare-Threats and Cybersecurity Countermeasures," 2022.

[11] L. Thirupathi, T. R. Boya, S. Gattoju, and E. S. Reddy, "Quantum Computing and AI: Synergizing for Sustainable Disaster Management in Industry 6.0," in The Rise of Quantum Computing in Industry 6.0 Towards Sustainability, Springer, 2024, pp. 35–51.

[12] R. Frehner and K. Stockinger, "Applying Quantum Autoencoders for Time Series Anomaly Detection," Oct. 09, 2024, arXiv: arXiv:2410.04154. doi: 10.48550/arXiv.2410.04154.

[13] S. Corli, L. Moro, D. Dragoni, M. Dispenza, and E. Prati, "Quantum machine learning algorithms for anomaly detection: A review," Mar. 03, 2025, arXiv: arXiv:2408.11047. doi: 10.48550/arXiv.2408.11047.

[14] A. Sakhnenko, C. O'Meara, K. J. B. Ghosh, C. B. Mendl, G. Cortiana, and J. Bernabé-Moreno, "Hybrid Classical-Quantum Autoencoder for Anomaly Detection," Quantum Mach. Intell., vol. 4, no. 2, p. 27, Dec. 2022, doi: 10.1007/s42484-022-00075-z.

[15] M. Hdaib, S. Rajasegarar, and L. Pan, "Quantum deep learning-based anomaly detection for enhanced network security," Quantum Mach. Intell., vol. 6, no. 1, p. 26, May 2024, doi: 10.1007/s42484-024-00163-2.

[16] F. Liu, S. Zhang, W. Ma, and J. Qu, “Research on attack detection of cyber physical systems based on improved support vector machine,” Mathematics, vol. 10, no. 15, p. 2713, 2022.

[17] S. Plambeck, G. Fey, J. Schyga, J. Hinckeldeyn, and J. Kreutzfeldt, “Explaining cyber-physical systems using decision trees,” in 2022 2nd International Workshop on Computation-Aware Algorithmic Design for Cyber-Physical Systems (CAADCPS), IEEE, 2022, pp. 3–8.

[18] M. Catillo, A. Pecchia, and U. Villano, “CPS-GUARD: Intrusion detection for cyber-physical systems and IoT devices using outlier-aware deep autoencoders,” Computers & Security, vol. 129, p. 103210, 2023.

[19] M. Catillo, A. Pecchia, and U. Villano, “CPS-GUARD: Intrusion detection for cyber-physical systems and IoT devices using outlier-aware deep autoencoders,” Computers & Security, vol. 129, p. 103210, 2023.

[20] “HAI Security Dataset.” Accessed: Apr. 05, 2025. [Online]. Available: https://www.kaggle.com/datasets/icsdataset/hai-security-dataset

[21] J. S. Baker et al., “Quantum Variational Rewinding for Time Series Anomaly Detection,” Nov. 02, 2022, arXiv: arXiv:2210.16438. doi: 10.48550/arXiv.2210.16438.

[22] H. Oh and D. K. Park, “Quantum support vector data description for anomaly detection,” Mach. Learn.: Sci. Technol., vol. 5, no. 3, p. 035052, Sep. 2024, doi: 10.1088/2632-2153/ad6be8.