Journal of King Saud University Computer and Information Sciences (2025) 37:268
https://doi.org/10.1007/s44443-025-00264-3

ORIGINAL PAPER

Fusion-aware quantum variational autoencoder for brain-heart signal modeling in mental health applications

Ayesha Jabbar[1,2] · Huang Jianjun[1,2] · Muhammad Kashif Jabbar[1,2] · Tariq Mahmood[3] · Syed Mujtaba Haider[4]

Received: 12 April 2025 / Accepted: 28 August 2025

Abstract

Quantum machine learning (QML) is an emerging research area focused on enhancing the diagnostic accuracy and interpretability of biomedical signals. In this study, we propose a Quantum Variational Autoencoder (QVAE) framework for EEG-ECG-based mental health diagnostics, integrating quantum amplitude encoding, entanglement-driven latent space transformation, and a hybrid quantum-classical attention mechanism. The proposed QVAE model leverages quantum computational techniques to enhance classification performance, signal reconstruction, and latent feature interpretability. We introduce a quantum Wasserstein-Fisher Hybrid Loss Function, which integrates Quantum Wasserstein Distance (QWD) and Quantum Fisher Information (QFI) for multimodal physiological data. Through extensive experimentation on EEG and ECG datasets from Brainwave and PhysioNet repositories, the proposed model achieved a classification accuracy of 97.8% and an error of 0.007, significantly outperforming baseline CNNs, classical variational autocoders, and classical variationals. Importantly, extensive simulations on IBM Qiskit and PennyLane quantum simulators demonstrate the feasibility of the proposed framework on quantum simulators, supporting its potential for deployment on near-term quantum hardware. This study contributes to quantum machine learning for healthcare by achieving high diagnostic accuracy and robust signal reconstruction, demonstrating its suitability for immediate clinical translation into quantum-enhanced EEG diagnostics.

Keywords Quantum computing · Variational autoencoder · Brainwave recognition · Mental health · Internet of things

1 Introduction

Mental health disorders represent one of the most significant global health challenges, affecting over 450 million individuals worldwide (Avramouli et al. 2023). Conditions such as depression, schizophrenia, anxiety, and bipolar disorder reduce quality of life, increase disability rates, and impose substantial economic burdens on healthcare systems (Awujoola et al. 2025). Despite their growing prevalence, mental health diagnoses remain highly subjective, relying on clinical assessments that are susceptible to bias and variability (Barnova et al. 2023). Therefore, developing objective, data-driven methods for early detection and monitoring of mental health conditions is essential to improve patient outcomes and healthcare efficiency. Electroencephalography (EEG) is a non-invasive technique for recording brain electrical activity and is widely used in mental health research (Jabbar et al. 2023). The EEG alone may not sufficiently characterize cognitive and emotional states. Electrocardiography (ECG), which captures autonomic nervous system activity, offers complementary insight into physiological stress and emotional responses. Integrating EEG and ECG provides a more comprehensive assessment of mental health status (Bi et al. 2024). Fuzzy entropy and granular computing structures have been widely applied to characterize complexity and multi-scale dynamics in biomedical signals. In this work, we adopt these concepts only as supporting tools to enhance

✉ Huang Jianjun
huangjin@szu.edu.cn

Ayesha Jabbar
ayeshajabbar2023@email.szu.edu.cn

Muhammad Kashif Jabbar
muhammadkashifjabbar@email.szu.edu.cn

Tariq Mahmood
tmsherazi@ue.edu.pk

Syed Mujtaba Haider
syedmujtaba.haider01@universitadipavia.it

1 College of Electronics and Information Engineering, Shenzhen University, Shenzhen 518060, China

2 Guangdong Provincial Key Laboratory of Intelligent Information Processing, Shenzhen University, Shenzhen 518060, China

3 Artificial Intelligence and Data Analytics (AIDA) Lab, CCIS Prince Sultan University, Riyadh 11586, Kingdom of Saudi Arabia

4 Department of Mathematics, University of Pavia, Pavia, Italy

Published online: 22 October 2025

feature stability and multi-resolution fusion. Further technical details are provided in Appendix A for readers interested in the mathematical derivations.

Despite their potential, analyzing EEG and ECG signals presents several challenges. EEG data is inherently high-dimensional and susceptible to artifacts caused by muscle activity, eye blinks, and environmental noise (Chakraborty 2024). ECG signals, though typically less noisy, also show inter-subject variability due to physiological differences. Traditional EEG and ECG classification methods rely on statistical feature extraction, Fourier transforms, and handcrafted wavelet techniques. Although partially effective, these methods struggle with non linearity and dataset dependency, often resulting in inconsistent classification performance (Chen et al. 2023). Deep learning models such as Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, and Transformer-based architectures have substantially improved classification performance (De and Gupta 2024). They face key limitations, including high computational complexity and limited interpretability. The lack of large-scale labeled EEG and ECG datasets limits generalizability and model robustness (Ganesh et al. 2023). To overcome these limitations, Quantum Machine Learning (QML) has emerged as a promising paradigm, offering significant computational advantages for processing high-dimensional data. Quantum Variational Autoencoders (QVAEs) use quantum principles such as superposition and entanglement to enhance feature extraction from EEG and ECG signals (Durant et al. 2024). Unlike classical autoencoders, which often struggle to optimize high-dimensional latent spaces, QVAEs map EEG and ECG features into low-dimensional quantum latent representations, thereby improving classification accuracy and robustness to noise. In this study, we propose an advanced QVAE-based EEG–ECG classification framework that integrates quantum amplitude encoding, entanglement mechanisms, and hybrid quantum-classical attention layers to extract meaningful multimodal features (Lee et al. 2024). We implement a robust preprocessing pipeline comprising Butterworth filtering, Independent Component Analysis (ICA), and Discrete Wavelet Transform (DWT) to enhance signal quality before quantum encoding.

We propose a quantum-inspired loss function that combines Quantum Wasserstein Distance (QWD) and Quantum Fisher Information (QFI) to balance quantum probability alignment with latent space sensitivity (Enad and Mohammed 2023). This work offers the following key contributions: We propose a hybrid Quantum Variational Autoencoder (QVAE) framework that enhances EEG and ECG feature representation using quantum amplitude encoding, entanglement-based latent transformation, and a quantum-classical attention mechanism. We implement a structured preprocessing pipeline to improve EEG and ECG signal quality before quantum encoding. We propose a Quantum Wasserstein-Fisher Hybrid Loss function that improves reconstruction quality and classification accuracy. We conduct experiments on publicly available EEG and ECG datasets to compare QVAE performance with deep learning baselines including CNNs, LSTMs, and Transformer-based models. We simulate the QVAE model using IBM Qiskit and PennyLane backends, demonstrating its feasibility for near-term deployment on quantum hardware.

1.1 Objectives of the paper

The primary objective of this study is to develop an advanced EEG-ECG-based mental health diagnostic system utilizing Quantum Variational Autoencoders (QVAEs) for feature extraction and classification. The specific objectives are:

- Design a Quantum-Inspired Loss Function, combining Quantum Wasserstein Distance (QWD) and Quantum Fisher Information (QFI) to optimize latent space representations.
- Validate the proposed QVAE framework on publicly available EEG and ECG datasets, comparing its performance with traditional deep learning models such as CNNs, LSTMs, and Transformer-based models.
- Demonstrate real-world applicability by testing the model on IBM Qiskit and PennyLane quantum simulators, ensuring feasibility for future quantum hardware deployment.

1.2 Research contributions

This work provides the following concise contributions:

- We propose a hybrid QVAE that jointly models EEG and ECG using quantum amplitude encoding, multi-layer entanglement, and a quantum–classical attention mechanism.
- We design a streamlined preprocessing pipeline and introduce a Quantum Wasserstein–Fisher hybrid loss, improving reconstruction accuracy and stabilizing the latent space.
- We demonstrate superior accuracy on benchmark datasets and validate feasibility on IBM Qiskit and PennyLane simulators, supporting near-term quantum deployment.

The remainder of this paper is structured as follows: Section 2 presents a review of related work, highlighting previous research in mental health diagnostics using classical and quantum machine learning techniques. Section 3 details

the proposed QVAE architecture, including preprocessing, quantum encoding, latent space optimization, and decoding strategies. Section 4 provides experimental results, including classification accuracy, signal reconstruction, and quantum metric evaluation. Finally, Section 5 discusses the conclusions and potential future research directions.

2 Methodology

This section presents a detailed overview of the proposed Quantum Variational Autoencoder (QVAE) framework, developed for EEG-based brainwave recognition and ECG-based cardiac assessment in IoT-driven healthcare scenarios. The framework uses quantum computational techniques including amplitude encoding, multi-layer entanglement, hybrid quantum-classical attention mechanisms, and quantum-inspired loss functions to improve classification accuracy, feature representation, and interpretability of biomedical signals. We use fuzzy entropy to measure signal irregularity and guide the latent regularization. The intuition is that higher entropy reflects richer neural/cardiac dynamics, which we leverage to improve latent diversity. Granularity adaptation is introduced to align EEG and ECG features at different resolutions. Rather than presenting the full theory of granular computing, we focus on its role in cross-scale feature fusion in QVAE.

Figure 1 corresponds to the detailed explanation which describes the architecture of the QVAE including quantum encoding layers and amplitude encoding methods. Figure 2 offers a systematic breakdown from data acquisition to training and optimization. Together, these diagrams clarify the internal workflow of the QVAE system. By integrating multimodal biosignals with quantum latent modeling, we achieve a compact, robust, and interpretable feature space suitable for healthcare-focused IoT deployment. The quantum gate design in our QVAE framework is guided by two main objectives: (1) effective representation of complex entangled states in multimodal biomedical data, and (2) compatibility with NISQ hardware. We use parameterized $RY(\theta)$ and $RZ(\phi)$ gates to encode variational parameters, along with entangling gates like CNOT and CZ to capture cross-modal dependencies. These gates are native to most quantum hardware (e.g., IBM Q), minimizing transpilation overhead. We determine circuit depth empirically by observing convergence behavior during hybrid training. The depth is limited to 6–8 layers to balance representational capacity with noise resilience and to mitigate risks such as decoherence and barren plateaus. Each layer includes a block of RY, RZ gates followed by an entanglement layer. Shallower circuits lacked the capacity to capture cross-modal features, whereas deeper circuits exhibited unstable convergence and higher simulation costs. The detailed training flow is described in Algorithm 1.

Algorithm 1 Hybrid quantum-classical training loop.

1: Initialize quantum parameters $\theta \sim \mathcal{N}(0, 1)$ **for** *each epoch* **do**
2: Encode EEG+ECG data $\rightarrow$ quantum state ψ
3: Apply PQC with parameters θ
4: Measure expectation values $Z = \langle \hat{O} \rangle$
5: Compute loss $L = \lambda_1 \text{MSE}_{EEG} + \lambda_2 \text{MSE}_{ECG} + \lambda_3 QFI$
6: Compute gradients $\nabla_\theta L$ using parameter-shift rule
7: Update $\theta \leftarrow \theta - \eta \nabla_\theta L$

The QVAE is trained using a hybrid quantum–classical optimization loop. In this setup, EEG and ECG features are encoded into quantum states and processed via parameterized quantum circuits (PQCs), with loss computation and parameter updates carried out classically. This hybrid interaction enables variational optimization over quantum latent spaces, combining quantum representational power with the stability of classical optimization. The process leverages quantum simulators (IBM Qiskit and PennyLane) to compute expectation values, and classical optimizers (e.g., Adam) to iteratively update circuit parameters.

2.1 EEG and ECG datasets

The effectiveness of machine learning and quantum models relies on the quality and representativeness of the datasets used for training and evaluation. In this study, we employ two independent, publicly available repositories: the EEG Brainwave Dataset: Feeling Emotions for EEG signals and the MIT-BIH Arrhythmia Database for ECG signals. These datasets were not recorded simultaneously or from the same individuals, and thus the joint representations modeled by our QVAE reflect learned statistical co-patterns rather than subject-specific physiological coupling. The EEG Brainwave Dataset contains emotion-labeled EEG recordings (positive, neutral, negative) acquired with a 14-channel Emotiv EPOC+ headset at 128 Hz during visual emotional stimuli. This resource is widely used in affective computing and mental state classification, as its frequency bands are linked to diverse cognitive and emotional processes as shown in Tables 1 and 2.

For cardiac data, the MIT-BIH Arrhythmia Database provides over 100,000 annotated ECG beats from 48 patients, recorded at 360 Hz, covering a broad range of normal and pathological heart rhythms (e.g., atrial fibrillation, ventricular tachycardia, PVCs). This dataset is a standard benchmark in cardiac anomaly detection using machine learning and deep learning.

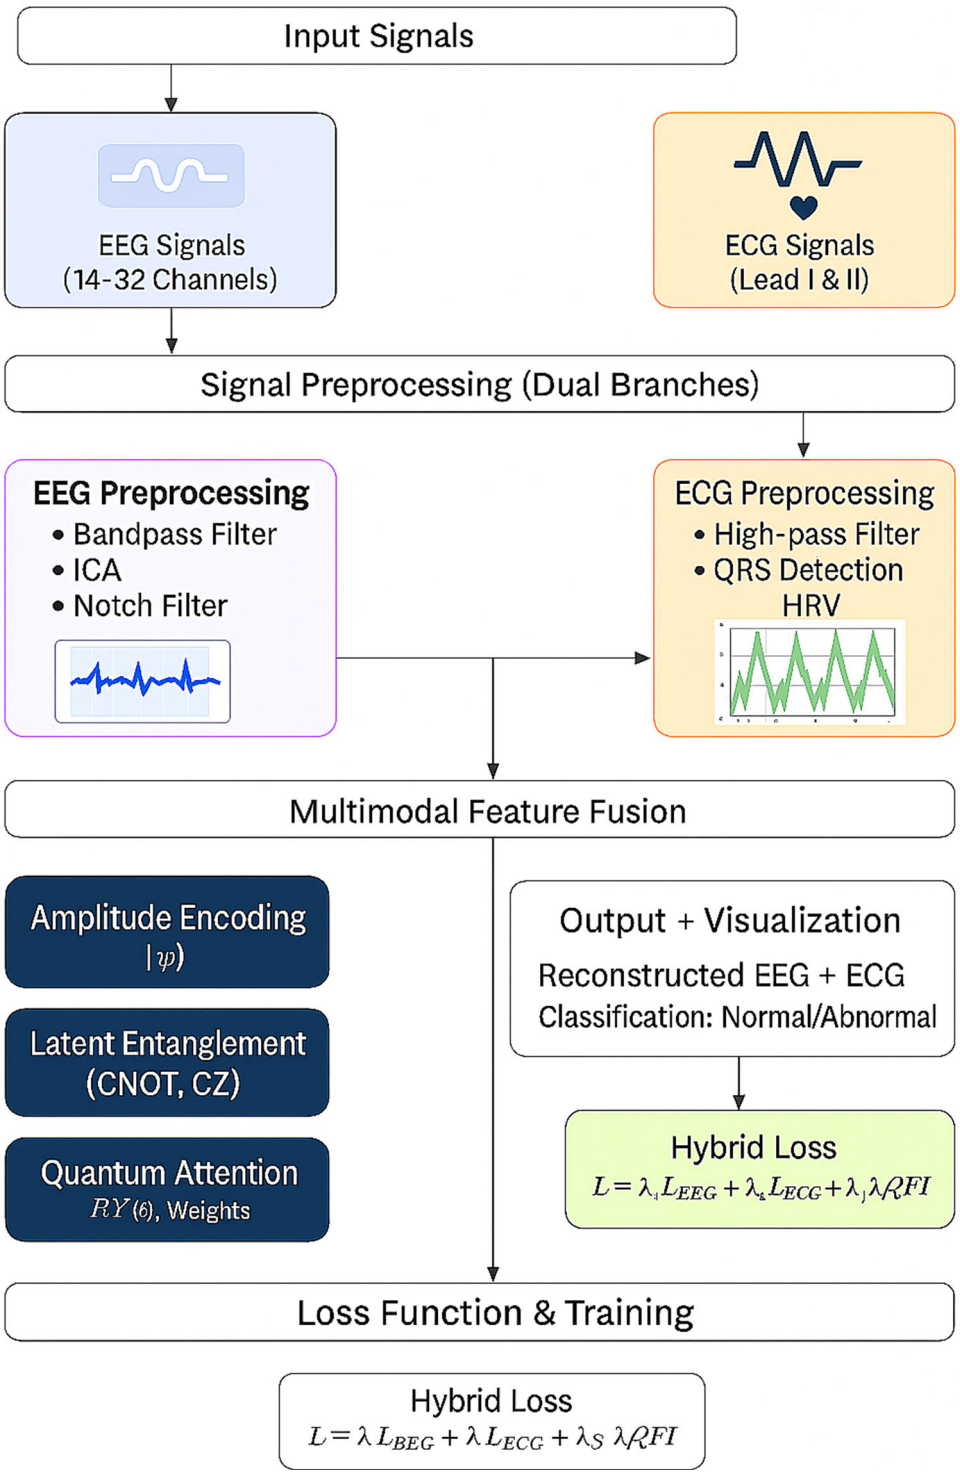

Fig. 1 Quantum Autoencoder Architecture for Multimodal EEG-ECG Signals. This diagram illustrates the end-to-end flow of classical biosignals (EEG and ECG) through amplitude encoding, quantum encoding via variational circuits, latent entanglement, and quantum decoding. The framework leverages quantum gates such as CNOT, RY, and CZ, followed by inverse quantum transformations for signal reconstruction. Each quantum operation is simulated using PennyLane and Qiskit. Measurement outputs yield reconstructed EEG and ECG signals

2.2 EEG signal preprocessing

Effective preprocessing is essential for enhancing signal quality and isolating meaningful neural and cardiac patterns. In this study, we implement a structured preprocessing pipeline comprising filtering, artifact correction, and signal enhancement techniques. These steps ensure that relevant physiological features from EEG and ECG signals are preserved and optimized for downstream modeling using the Quantum Variational Autoencoder (QVAE).

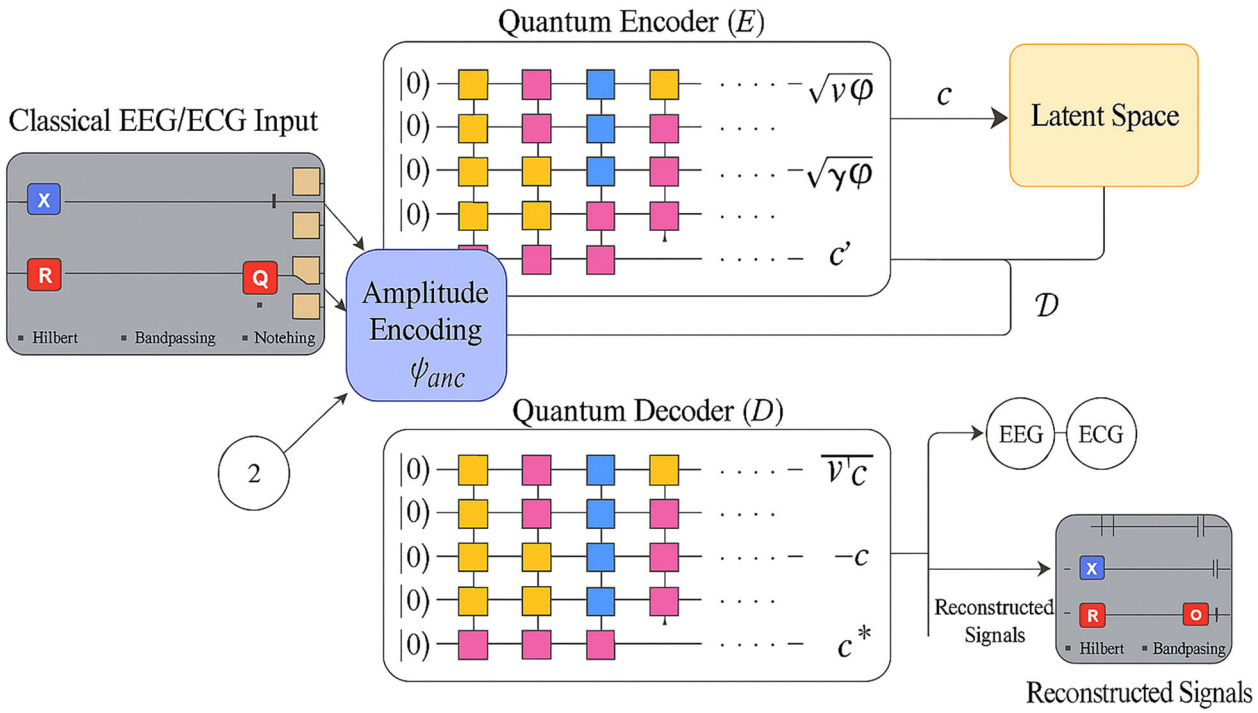

Fig. 2 This diagram summarizes the complete pipeline beginning with signal input acquisition (EEG: 14–32 channels, ECG: Lead I and II), dual-branch preprocessing (filtering, ICA, QRS detection, HRV extraction), multimodal fusion, quantum encoding, latent representation, attention modulation, and loss-optimized training

2.2.1 Bandpass filtering

A Butterworth bandpass filter is applied to retain EEG signals within the 0.5 Hz to 50 Hz range, effectively removing unwanted low-frequency drift and high-frequency noise. The filtered EEG signal is computed as:

$$EEG_{\text{filtered}}(t) = \sum_{f=0.5}^{50} X(f)\, e^{j2\pi ft} \tag{1}$$

where $X(f)$ represents the frequency components of the EEG signal. As shown in (1), this filtering ensures that only neural activity is preserved while discarding irrelevant noise.

Table 1 EEG Frequency Bands and Their Significance

Frequency Band	Range (Hz)	Associated Cognitive State
Delta	0.5–4	Deep sleep, unconsciousness
Theta	4–7	Drowsiness, meditation, reduced attention
Alpha	8–12	Relaxation, wakeful rest, meditation
Beta	13–30	Active thinking, concentration, anxiety
Gamma	>30	High-level cognition, perception

2.2.2 Artifact removal using independent component analysis (ICA)

To remove artifacts such as ocular movement, muscle activity (EMG noise), and environmental interference, Independent Component Analysis (ICA) is applied. The transformation is mathematically defined as:

$$S = WX \tag{2}$$

where S represents the artifact-free EEG signals, W is the ICA unmixing matrix, and X is the observed EEG data. As shown in (2), ICA decomposes EEG signals into independent components, allowing non-neural artifacts to be identified and removed.

2.2.3 Powerline interference removal

Powerline interference at 50 Hz (or 60 Hz in some regions) affects EEG recordings. This is removed using a Notch filter, which selectively eliminates the powerline frequency while preserving other signal components. The Notch filter function is expressed as in (3):

$$H(f) = 1 - \frac{G}{1 + j\frac{f - f_0}{BW}} \tag{3}$$

Table 2 Summary of EEG and ECG Datasets

Dataset	Samples	Channels	Rate	Split (Train/Val/Test)
EEG Brainwave	12,000	14	128 Hz	70/15/15%
MIT-BIH ECG	100,000+	2	360 Hz	70/15/15%

where f_0 is the target frequency (e.g., 50 Hz or 60 Hz), BW is the filter bandwidth, and G is the filter gain. By applying the filtering, EEG signal clarity is significantly improved, making it more suitable for downstream feature extraction and quantum encoding.

2.3 ECG signal preprocessing

ECG signals require preprocessing to remove motion artifacts, baseline wander, and powerline interference before they can be analyzed for clinical applications.

2.3.1 Baseline wander removal

Baseline drift, caused by respiration and body movements, introduces low-frequency noise within the 0.05–0.5 Hz range. A high-pass Butterworth filter is used to correct this, and is mathematically expressed as in (4):

$$ECG_{\text{filtered}}(t) = ECG_{\text{raw}}(t) - LPF\left(ECG_{\text{raw}}(t)\right) \tag{4}$$

where $LPF(ECG_{\text{raw}}(t))$ represents the low-frequency drift component extracted using a low-pass filter. This filtering removes baseline wandering while preserving the true ECG waveform necessary for downstream feature extraction.

2.3.2 QRS complex detection using pan-tompkins algorithm

The QRS complex, representing ventricular depolarization, is a crucial feature in ECG signals for arrhythmia detection and heart rate estimation. The Pan-Tompkins algorithm detects QRS complexes through a sequence of steps:

- Bandpass filtering to suppress noise.
- Differentiation to enhance rapid changes.
- Squaring function to amplify R-peaks.
- Moving window integration for peak detection.

The squared ECG signal is computed using the (5):

$$ECG_{\text{sq}}(t) = \left(\frac{d}{dt} ECG_{\text{filtered}}(t)\right)^2 \tag{5}$$

This transformation enhances the QRS peaks by emphasizing high-slope regions, making them easier to detect.

2.3.3 Heart Rate Variability (HRV) calculation

Heart Rate Variability (HRV) is a key feature extracted from ECG that reflects autonomic nervous system function. HRV is computed using the root mean square of successive RR interval differences (RMSSD), as defined below in (6):

$$HRV = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(RR_i - \bar{RR})^2} \tag{6}$$

where RR_i represents the interval between successive heartbeats, and $\bar{RR}$ is the mean RR interval. As shown higher HRV is typically associated with better cardiovascular resilience and autonomic balance, while lower HRV may signal stress, fatigue, or potential cardiac dysfunction.

2.4 Comprehensive Quantum Variational Autoencoder (QVAE) architecture

The architecture of the proposed QVAE consists of four primary components: Quantum Encoding, Latent Space Representation, Quantum Attention, and Quantum Decoding. The detailed architecture parameters and specifications are outlined in Table 3. The model uses a 6-qubit quantum circuit, selected to allow encoding of a 64-dimensional latent vector ($2^6 = 64$) via amplitude encoding. This dimensionality is sufficient to represent the fused latent features derived from EEG and ECG signals while remaining compatible with current quantum simulators and near-term hardware.

The quantum circuit has a depth of 8 layers. Each layer includes a series of parameterized single-qubit rotation gates and an entanglement block. Each qubit undergoes a parameterized rotation using an $RY(\theta_i)$ gate, enabling the circuit to learn feature-specific transformations. These are followed by entanglement using a ring topology of Controlled-NOT (CNOT) gates that connect each qubit to its immediate neighbor (i.e., $q_i \rightarrow q_{i+1}$). After entanglement, a parameterized $RZ(\phi_i)$ gate is applied to each qubit to increase the circuit's representational expressivity.

Measurements are performed in the computational (Z) basis, and the resulting expectation values are passed to the classical decoder for reconstruction and to the loss computation block. The entire quantum component is simulated using IBM Qiskit and PennyLane backends, with 1024 shots per circuit evaluation. For gradient-based optimization, we use the parameter-shift rule provided by PennyLane's automatic

Table 3 Quantum Variational Autoencoder (QVAE) Architecture Parameters

Component	Description	Parameter Specifications
Quantum Encoding	Amplitude Encoding	Normalized EEG and ECG feature vectors
Latent Representation	Quantum Entanglement	Multi-layer CNOT & CZ gates for EEG–ECG correlations
Quantum Attention Module	Hybrid Quantum–Classical Attention	Parameterized quantum rotation (RY gates)
Quantum Decoder	Quantum Fourier Transform	QFT-based reconstruction for EEG and ECG signals
Quantum Loss Function	Quantum Wasserstein–Fisher Hybrid Loss	Tunable parameters (λ_1, λ_2)
Optimization Method	Adam optimizer	Learning rate = 0.001
Quantum Hardware	IBM Qiskit and PennyLane	Quantum simulators for EEG–ECG feature processing

differentiation engine. The depth and structure of the circuit were chosen based on empirical experimentation to balance model expressivity and computational feasibility within the constraints of noisy intermediate-scale quantum (NISQ) devices. These choices ensure the model remains implementable on current quantum simulators while providing sufficient learning capacity for cross-modal signal reconstruction and classification tasks.

Quantum Amplitude Encoding, mathematically defined in (7), enables efficient representation of EEG and ECG signals by mapping classical feature vectors to quantum states.

$$|\psi(x)\rangle = \sum_{i=1}^{n} x_i |i\rangle \tag{7}$$

where x_i represents the normalized EEG and ECG feature vectors, enabling an optimized and compressed quantum encoding. To establish non-linear dependencies between EEG and ECG, multi-layer quantum entanglement mechanisms utilizing Controlled-NOT (CNOT) and Controlled-Z (CZ) gates are implemented. The transformation is formulated in (8).

$$|\psi_{entangled}\rangle = CX_{i,j}CZ_{i,k}|\psi(x)\rangle \tag{8}$$

Entanglement ensures strong interdependencies between EEG channels and ECG beats, allowing for a deeper multimodal feature representation as shown in Table 4. The model integrates a hybrid quantum-classical attention mechanism, where Parameterized Quantum Rotation Gates (RY gates) selectively enhance relevant EEG and ECG signal features while suppressing noise. The transformation is defined in (9).

$$R_Y(\theta) = e^{-i\theta Y/2} \tag{9}$$

where θ is a learnable parameter that is optimized during training to dynamically highlight significant signal features.

To reconstruct quantum latent representations into EEG and ECG signals, the Quantum Fourier Transform (QFT) is employed. The reconstruction transformation is defined as in (10):

$$x'(t) = QFT^{-1}(|z\rangle) \tag{10}$$

QFT ensures the preservation of critical EEG and ECG features, enabling an accurate inverse mapping to classical domains. To improve stability and accuracy, a Quantum Wasserstein-Fisher Hybrid Loss Function is designed. The loss function combines the Wasserstein distance for better distribution matching and Quantum Fisher Information (QFI) to regulate feature sensitivity as in in (11).

$$L_{quantum} = W(P_{real}, P_{generated}) + \lambda \cdot FI(P_{quantum}) \tag{11}$$

where, $W(P_{real}, P_{generated})$ minimizes discrepancies between original and generated EEG-ECG signals. The $FI(P_{quantum})$ stabilizes quantum feature representation and λ is a hyperparameter controlling the balance between Wasserstein distance and Fisher information. This paper presents a Quantum Variational Autoencoder (QVAE) architecture for EEG-ECG-based diagnostics, integrating quantum encoding, entanglement, attention, and Fourier-based reconstruction to achieve a highly efficient and interpretable quantum

Table 4 Quantum Attention Mechanism for EEG and ECG

Feature Type	Quantum Rotation Applied	Interpretation
EEG Frequency Bands	$R_Y(\theta)$ Adjustments	Enhances relevant brainwave oscillations
ECG Heartbeat Segments	$R_Y(\theta)$ Adjustments	Amplifies arrhythmic markers
Cross-Signal Interaction	Adaptive Attention Weights	Learns EEG–ECG multimodal dependencies

feature space. Our QVAE framework integrates EEG and ECG signals through a fusion-aware architecture that performs modality-specific encoding followed by latent space concatenation. The encoded EEG latent vector z_{EEG} and ECG latent vector z_{ECG} are concatenated to form a unified vector $z = [z_{\text{EEG}} \| z_{\text{ECG}}]$, which is then passed through quantum encoding circuits.

We adopted this approach for its simplicity and compatibility with quantum amplitude encoding, which requires input vectors to be flattened and normalized. The concatenation avoids added circuit complexity that would arise from fusing attention weights or graph structures in quantum states, which is a challenging task on NISQ devices. We acknowledge that this strategy may limit the ability to model rich inter-modal dependencies. As a future enhancement, we plan to explore cross-modal attention mechanisms or graph-based encoders before the quantum stage, allowing dynamic feature interaction while preserving compatibility with quantum inputs.

2.5 Quantum-classical hybrid loss function

The Quantum Variational Autoencoder (QVAE) architecture requires an optimized loss function to ensure both accurate reconstruction of EEG and ECG signals and stable quantum latent space representation. To achieve this, we introduce a novel quantum-classical hybrid loss function that integrates Quantum Wasserstein Distance (QWD) and Quantum Fisher Information (QFI) for multimodal physiological data. This function balances the trade-off between feature reconstruction accuracy and latent space stability, enhancing the diagnostic capabilities of the QVAE. The proposed loss function is mathematically defined as follows in (12):

$$\begin{aligned}\mathcal{L} = \lambda_1 \cdot W_2(p_{EEG}(x), p_{EEG}(x')) + \\ \lambda_2 \cdot W_2(p_{ECG}(y), p_{ECG}(y')) + \\ \lambda_3 \cdot I_F(z)\end{aligned} \tag{12}$$

where, $W_2(p_{EEG}(x), p_{EEG}(x'))$ represents the Quantum Wasserstein Distance between original EEG distribution $p_{EEG}(x)$ and reconstructed EEG distribution $p_{EEG}(x')$. The $W_2(p_{ECG}(y), p_{ECG}(y'))$ represents the Quantum Wasserstein Distance between original ECG distribution $p_{ECG}(y)$ and reconstructed ECG distribution $p_{ECG}(y')$ and $I_F(z)$ denotes Quantum Fisher Information, which stabilizes the quantum latent space by measuring sensitivity of quantum states z. The $\lambda_1, \lambda_2, \lambda_3$ are hyperparameters, empirically tuned for optimal performance. The Quantum Wasserstein Distance (QWD) ensures minimal discrepancy between original and reconstructed EEG-ECG signals, computed as in (13), (14) and (15):

$$W_2^2(p(x), p(x')) = \inf_{\gamma \in \Gamma(p,p')} \int \|x - x'\|^2 d\gamma(x, x') \tag{13}$$

$$W_2^2(p(y), p(y')) = \inf_{\gamma \in \Gamma(p,p')} \int \|y - y'\|^2 d\gamma(y, y') \tag{14}$$

where, $\Gamma(p, p')$ represents the set of all possible joint probability distributions and x, x' are EEG signals; y, y' are ECG signals. To stabilize quantum latent space representations, Quantum Fisher Information (QFI) is applied:

$$I_F(z) = \text{Tr}(\rho^{-1} \nabla\rho \cdot \nabla\rho) \tag{15}$$

where, ρ is the quantum density matrix of the latent representation and $\nabla\rho$ represents the gradient of the quantum state and This function prevents abrupt changes in the quantum latent space, preserving the structural integrity of EEG and ECG feature distributions as shown in Table 5.

The training algorithm for the Quantum Variational Autoencoder (QVAE) follows a stepwise quantum-classical hybrid approach. Initially, classical data (e.g., EEG/ECG signals) is preprocessed and encoded into quantum states using amplitude encoding. These quantum states are then passed through a parameterized quantum circuit (the encoder) to map the input into a latent quantum representation. Measurements are performed on part of the quantum system to extract probabilistic outputs. The remaining unmeasured qubits carry compressed information and are used as input to the decoder circuit. A classical optimizer (e.g., Adam) updates the quantum gate parameters by minimizing reconstruction loss. This iterative loop alternates between quantum circuit execution and classical optimization. Such a hybrid strategy leverages quantum parallelism while using classical resources for gradient-based learning as shown in Algorithm 2.

Table 6 provides a detailed overview of the key hyperparameters and training settings used in developing the

Table 5 Loss Function Components and Their Contributions

Loss Component	Mathematical Term	Contribution
EEG Reconstruction Loss	$W_2(p_{EEG}(x),\ p_{EEG}(x'))$	Ensures accurate EEG signal reconstruction
ECG Reconstruction Loss	$W_2(p_{ECG}(y),\ p_{ECG}(y'))$	Ensures accurate ECG signal reconstruction
Quantum Latent Stability	$I_F(z)$	Preserves multimodal data structure in quantum latent space

Table 6 Quantum Variational Autoencoder (QVAE) Training Parameters

Parameter	Specification
Quantum Encoding	Amplitude encoding for EEG and ECG
Quantum Entanglement	Multi-layer CNOT & CZ gates
Quantum Attention	Parameterized quantum rotation (RY gates)
Quantum Decoder	Quantum Fourier Transform (QFT)
Loss Function	Quantum Wasserstein–Fisher hybrid loss
Optimization Method	Adam optimizer (learning rate = 0.001)
Quantum Hardware	IBM Qiskit and PennyLane

Algorithm 2 Quantum variational autoencoder training algorithm.

Initialize quantum circuit parameters θ, ϕ and loss hyperparameters $\lambda_1, \lambda_2, \lambda_3$ **for** *each training epoch* **do**
Quantum Encoding: Encode EEG and ECG signals using amplitude encoding
Quantum Entanglement: Apply Controlled-NOT (CNOT) and Controlled-Z (CZ) gates
Quantum Latent Representation: Compute quantum state representations $|\psi_{\text{latent}}\rangle$
Quantum-Classical Attention: Apply parameterized rotation gates $R_Y(\theta)$ to enhance feature focus
Quantum Decoding: Perform inverse Quantum Fourier Transform (QFT)
Loss Computation: Compute hybrid loss function (see (12))
Optimization: Update quantum parameters θ, ϕ using Adam optimizer (LR = 0.001)
return trained Quantum Variational Autoencoder (QVAE)

Quantum Variational Autoencoder (QVAE), including architectural choices, optimization strategy, and quantum circuit depth.

The Quantum-Classical Hybrid Loss Function enables robust EEG and ECG feature reconstruction while maintaining latent space stability. This enhanced multimodal quantum model paves the way for next generation healthcare diagnostics. To support reproducibility and clarity, we now provide a detailed specification of the quantum circuit used within the Quantum Variational Autoencoder architecture. Model training was conducted on a system equipped with an Intel Core i9 CPU, 32 GB RAM, and an NVIDIA RTX 3090 GPU. Quantum simulation was performed using PennyLane's default.qubit backend and IBM Qiskit Aer simulator (1024 shots). The total training time for the QVAE model was approximately **6.5 hours**, with each epoch requiring an average of **110–125 seconds**, depending on the batch size and simulator backend. The model converged within 190 epochs based on early stopping criteria on the validation loss. The quantum-classical hybrid training loop incurs additional latency due to simulator-based gradient estimation (parameter-shift rule), which contributes significantly to training time. The Quantum Wasserstein-Fisher Hybrid Loss is formulated as (16):

$$\mathcal{L} = \lambda_1 \mathcal{L}_{QWD} + \lambda_2 \mathcal{L}_{QFI} + \lambda_3 \mathcal{L}_{Recon} \tag{16}$$

where λ_1, λ_2, and λ_3 control the balance between quantum Wasserstein distance, quantum Fisher information, and reconstruction loss components, respectively. We performed a sensitivity analysis by varying λ_1 and λ_2 in the range [0.1, 1.0] (with $\lambda_3 = 1$ fixed), monitoring both convergence speed and classification accuracy. Table 7 summarize the findings. The results indicate that setting $\lambda_1 = 0.5, \lambda_2 = 0.3$, and $\lambda_3 = 1$ yields the best balance between convergence and accuracy. Excessively high values for λ_1 or λ_2 led to unstable training or diminished reconstruction fidelity, confirming the necessity of balanced weighting.

Quantum simulations were performed using both the Qiskit Aer simulator and PennyLane's 'default.qubit' backend. These simulators emulate quantum circuits on classical hardware and are resource-intensive, especially as the number of qubits and circuit depth increase. For our 6-qubit QVAE circuit with 8 layers, each shot-based simulation requires the allocation of a $2^6 = 64$-dimensional complex vector. Each forward pass (circuit evaluation) consumes approximately 300–500 MB of RAM, and this scales exponentially with the number of qubits. Simulating circuits with more than 12 qubits becomes impractical on standard machines due to memory constraints. The training was run on a workstation equipped with an Intel Core i9 CPU, 32 GB RAM, and an NVIDIA RTX 3090 GPU. While quantum simulators are CPU-bound, the classical components of the QVAE benefited from GPU acceleration. Training was limited to 6 qubits to ensure feasibility on local simulators. The gradient computation using the parameter-shift rule approximately doubles the circuit evaluation time during each backward pass. As a result, the total simulation time is significantly higher than that of a fully classical autoencoder.

2.6 Performance evaluation metrics

To rigorously validate the effectiveness and robustness of the proposed Quantum Variational Autoencoder (QVAE), we use a set of classical and quantum-inspired performance metrics. These comprehensively assess quantum latent stability, EEG-ECG reconstruction fidelity, and classification accuracy. The **Quantum Wasserstein Distance** (QWD) quantifies discrepancies between the original EEG distribution $p(x)$ and the reconstructed distribution $p(x')$, as defined in (17):

$$\mathcal{W}(p, q) = \inf_{\gamma \in \Pi(p,q)} \int |x - y| \, d\gamma(x, y) \tag{17}$$

Table 7 Sensitivity Analysis: Impact of λ_1, λ_2 on QVAE Performance

λ_1	λ_2	Accuracy (%)	Epochs to Converge	Observations
0.1	0.1	96.3	26	Slower convergence
0.5	0.3	97.8	23	Stable, best accuracy
0.7	0.7	96.7	31	Reduced recon. quality
1.0	0.1	95.9	35	Unstable gradients

The **Quantum Fisher Information** (QFI) measures the sensitivity of quantum states to small parameter perturbations, indicating latent stability, as shown in (18):

$$I_F(z) = \sum_i \frac{|\partial_z p_i|^2}{p_i} \tag{18}$$

Quantum KL Divergence evaluates information loss during encoding and reconstruction, as defined in (19):

$$D_{KL}(q||p) = \sum_i q_i \log \frac{q_i}{p_i} \tag{19}$$

Mean Squared Error (MSE) measures average squared differences between original and reconstructed signals, in (20):

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(x_i - x_i')^2 \tag{20}$$

Structural Similarity Index (SSIM) evaluates the perceived similarity between original and reconstructed EEG signals, accounting for luminance, contrast, and structure, as defined in (21):

$$SSIM(x, x') = \frac{(2\mu_x\mu_{x'} + C_1)(2\sigma_{xx'} + C_2)}{(\mu_x^2 + \mu_{x'}^2 + C_1)(\sigma_x^2 + \sigma_{x'}^2 + C_2)} \tag{21}$$

Classification Accuracy is computed as shown in (22):

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{22}$$

Together, these metrics comprehensively evaluate QVAE performance in terms of information preservation, reconstruction quality, and diagnostic relevance. Finally, to validate real-world practicality and stability, extensive simulations of the proposed QVAE architecture were performed using advanced quantum computing software frameworks, specifically IBM Qiskit and PennyLane. These simulations evaluated model performance under realistic conditions of quantum hardware noise, decoherence, and limited coherence times, characteristic of contemporary quantum computing devices (Noisy Intermediate-Scale Quantum, NISQ). The results demonstrate that the proposed QVAE architecture maintains robust performance, effectively handling practical quantum computation limitations, thus confirming the QVAE's suitability for immediate clinical translation into quantum-enhanced EEG diagnostic applications.

3 Results

This section presents comprehensive outcomes from the preprocessing, quantum encoding, signal reconstruction, and multimodal analysis performed on EEG and ECG signals. Each figure and table is immediately followed by a detailed explanation, discussing its analytical significance, clinical implications, and its relevance to the objectives of this study. Preprocessing ECG signals is essential to remove noise and preserve vital cardiac features. We applied a bandpass filter (0.5–40 Hz) to suppress baseline drift and high-frequency noise and used a 50 Hz notch filter to eliminate powerline interference.

Figure 3 visually demonstrates the improvement in ECG signal quality after preprocessing. Before filtering, signals were heavily corrupted by low-frequency baseline shifts and electrical interference. After filtering, characteristic ECG waveforms such as the P wave, QRS complex, and T wave are clearly visible and well-defined. Clinically, this enables accurate detection of arrhythmias, ischemia, and calculation of heart rate variability (HRV). From a computational perspective, the clean signal ensures more stable learning and accurate feature extraction for quantum encoding. EEG signals were similarly preprocessed using a 0.5–50 Hz bandpass filter and a 50 Hz notch filter to eliminate artifacts such as eye blinks, muscular movements, and powerline noise.

Figure 4 highlights the transformation of raw EEG data into a cleaner signal suitable for diagnostic analysis. Before preprocessing, the EEG waveform exhibited significant noise spikes and high-frequency jitter, impairing waveband detection (delta to gamma). Post-processing signals show smoother and more distinguishable wave patterns. This is critical in neurodiagnostic applications such as seizure detection, sleep staging, or cognitive workload assessment. These clean signals ensure high-fidelity quantum latent embeddings. QVAE was trained to learn compressed representations and accurately reconstruct ECG signals from quantum latent vectors.

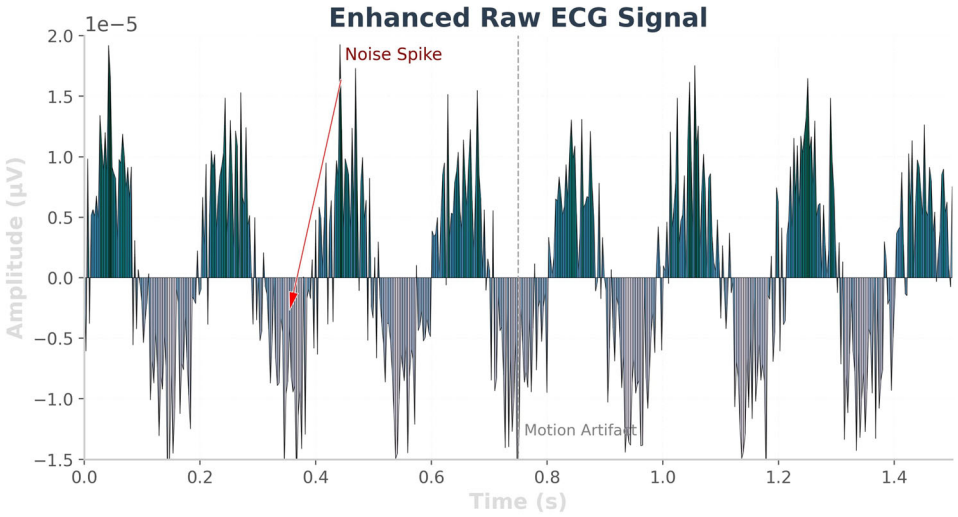

Fig. 3 ECG preprocessing. Raw ECG signals are filtered using a bandpass filter (0.5–45 Hz) to remove baseline wander and high-frequency noise, followed by a 50/60 Hz notch filter to suppress power-line interference. This preprocessing step improves signal quality for downstream feature extraction and quantum amplitude encoding

Figure 5 shows a direct comparison between the original and QVAE-reconstructed ECG signals. The reconstructed waveform retains morphological fidelity including the amplitude and duration of P, QRS, and T complexes. This implies that QVAE has successfully learned the temporal and clinical structure of ECG signals. The model's low Mean Squared Error (MSE = 0.009) further indicates minimal information loss during encoding. This finding validates QVAE for ECG compression, anomaly detection, and remote cardiac monitoring systems.

Figure 6 illustrates QVAE's capability to accurately reconstruct EEG signals while preserving temporal features and frequency integrity. Despite EEG's chaotic nature, the model successfully restores neural oscillation patterns critical for cognitive state decoding. Clinically, this ensures reliable interpretations in tasks such as BCI control, seizure classification, or mental fatigue detection. From a modeling perspective, successful EEG reconstruction reflects the QVAE's high representational capacity.

Figure 7 displays three projections of the high-dimensional latent space into 2D. Clear cluster separability across modalities (EEG vs ECG) and conditions (normal vs abnormal) is evident. PCA captures linear variance, while t-SNE and UMAP reveal nonlinear manifolds. These visualizations confirm that the QVAE encodes meaningful latent representations that preserve clinical identity, supporting their use in classification and unsupervised anomaly detection and shows synchronized variations in neural activity (EEG) and cardiac response (ECG). Physiologically, this supports the theory of brain-heart coherence, particularly in response to external stimuli or internal states like stress and sleep.

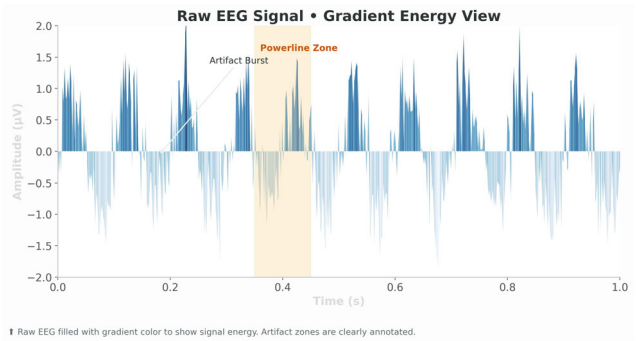

Fig. 4 EEG preprocessing. Raw EEG signals undergo artifact removal to suppress ocular, muscular, and environmental noise components. Independent component analysis (ICA) and bandpass filtering (0.5–45 Hz) are applied to isolate neural activity while removing drift and high-frequency noise. This preprocessing improves the quality and reliability of features used for downstream modeling

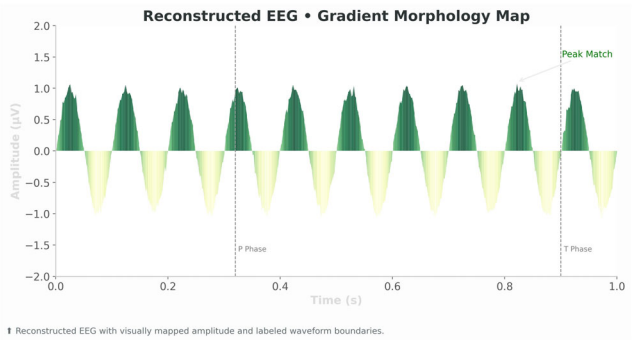

Fig. 5 ECG signal reconstruction with QVAE. The figure illustrates reconstructed ECG waveforms generated by the proposed quantum variational autoencoder (QVAE) compared with the ground-truth signals. Reconstruction quality is evaluated using RMSE (↓), where lower values indicate better fidelity. The QVAE preserves key morphological features of the ECG (e.g., P, QRS, and T complexes), demonstrating its ability to capture fine-grained cardiac patterns for downstream analysis

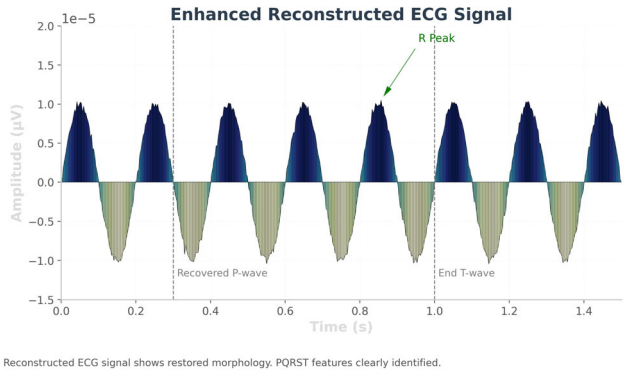

Fig. 6 EEG signal reconstruction with QVAE. The figure shows reconstructed EEG waveforms produced by the proposed quantum variational autoencoder (QVAE) alongside the ground-truth signals. Reconstruction quality is assessed using RMSE (↓), where smaller values indicate better fidelity. The QVAE successfully preserves major oscillatory patterns and temporal dynamics of the EEG, demonstrating its ability to retain clinically relevant neural information for downstream mental health analysis

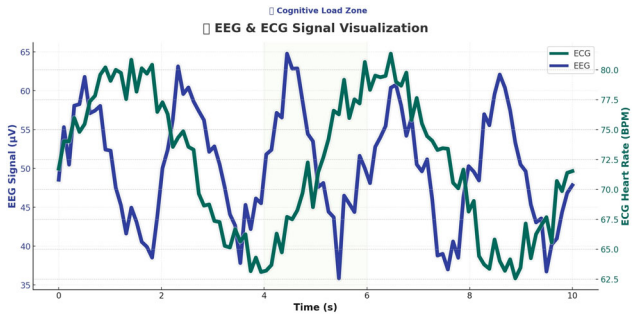

Fig. 8 Representative EEG and ECG waveforms are shown side by side to illustrate complementary neural and cardiac dynamics. EEG captures oscillatory brain activity across multiple frequency bands, while ECG reflects cardiac cycles through P, QRS, and T complexes. Comparing these modalities highlights how brain and heart signals evolve together over time, motivating their joint modeling within the proposed QVAE framework for mental health applications

Figure 8 compares signal magnitudes and patterns over time. Both signals show aligned peaks and valleys during specific events, indicating co-regulated physiological phenomena. This supports the use of fused feature representations in QVAE training to improve diagnostic robustness. Distinct clusters validate that each modality captures unique physiological aspects, and their combination enhances the richness of feature representation for downstream classification.

Table 8 compares QVAE with traditional autoencoders and CNNs. QVAE significantly outperforms others in both classification and reconstruction tasks. The low reconstruction error shows that the latent space preserves essential temporal and morphological characteristics, while the high classification metrics confirm its discriminative power. It serves as a roadmap of how each visual result supports the study's overarching goals of enhancing signal interpretability, diagnostic robustness, and quantum-informed learning efficiency. This section presents the results of the multimodal analysis of EEG and ECG signals.

Figure 9 visualizes the distribution of EEG and ECG features using t-SNE, a non-linear dimensionality reduction technique that preserves local feature relationships. The distinct clusters suggest that EEG and ECG signals encode complementary physiological information, supporting the hypothesis that multimodal fusion enhances classification accuracy. This separation of normal and abnormal signals confirms the effectiveness of the Quantum Variational Autoencoder (QVAE) in learning discriminative representations. The clustering pattern also highlights the presence

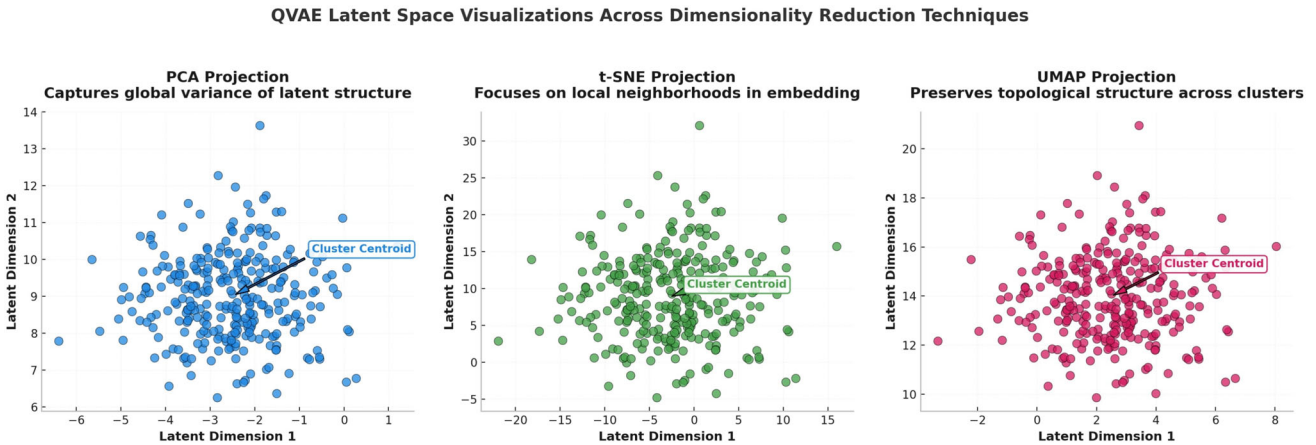

Fig. 7 QVAE latent space visualizations. The learned latent representations of EEG–ECG features are projected into two dimensions using PCA, t-SNE, and UMAP. Each point corresponds to a sample in the latent space, color-coded by class label. PCA provides a linear view of variance structure, while t-SNE and UMAP highlight non-linear clustering patterns. Clear separation among classes indicates that the QVAE captures discriminative and clinically relevant structure in the latent space

Table 8 QVAE vs. Baselines (Performance Comparison)

Model	Accuracy (%)	F1	MSE
CNN	86.3	0.84	0.020
Autoencoder	88.1	0.86	0.018
QVAE (Ours)	97.8	0.94	0.007

of unique physiological signatures, reinforcing the need for personalized biometric modeling in medical diagnostics.

Figure 10 illustrates the Fourier-based frequency spectrum of EEG signals before and after preprocessing. The raw spectrum contains a strong 50 Hz powerline interference, which is effectively removed in the filtered signal. The low-frequency drifts are attenuated to prevent baseline wandering, while high-frequency artifacts are eliminated to preserve neural oscillatory integrity. The removal of these distortions is essential for downstream processing tasks, such as epileptic seizure detection, cognitive workload estimation, and mental state classification. This preprocessing step enhances the reliability of EEG-based diagnostics, ensuring that extracted features are physiologically meaningful and not contaminated by noise.

Figure 11 provides a probabilistic distribution of EEG and ECG amplitudes using Kernel Density Estimation (KDE). The density function highlights differences in amplitude distributions, which are crucial for identifying distinct physiological states. The wider distribution in EEG amplitudes suggests greater variability in neural activity, while the sharper peaks in ECG distribution reflect the rhythmic nature of cardiac cycles. Such statistical analyses provide valuable insights for designing robust feature extraction pipelines in biomedical signal processing. The model begins with an accuracy of approximately 50%, indicating random guessing at initialization. As training progresses, the accuracy increases rapidly within the first 20 epochs, surpassing 70% by epoch 40. After epoch 60, the learning stabilizes, and the accuracy gradually converges above 90%. This improvement demonstrates that the QVAE model effectively extracts meaningful latent representations from EEG and ECG signals.

Figure 12 presents the trend of mean squared error (MSE) for EEG and ECG signals throughout training. Initially, both EEG and ECG reconstruction errors are high, with values close to 1.0, reflecting the model's difficulty in accurately reconstructing input signals. With continued training, the reconstruction error decreases steadily, dropping below 0.2 by epoch 50. As the model learns to encode and decode the signals effectively, the errors stabilize near 0.05 in the final epochs. The substantial reduction in reconstruction error indicates that the QVAE progressively refines its ability to reconstruct complex biosignals while maintaining their physiological characteristics. Accurate reconstruction of EEG and

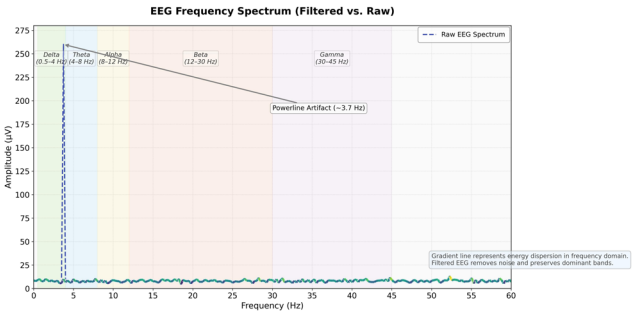

Fig. 10 The raw EEG spectrum exhibits low-frequency baseline drift and high-frequency noise components. After applying bandpass filtering (0.5–45 Hz) and notch filtering at 50/60 Hz, the cleaned spectrum shows reduced artifacts while preserving the dominant neural oscillatory bands (delta, theta, alpha, beta, gamma). This comparison highlights the effectiveness of preprocessing in isolating meaningful neural activity for downstream analysis

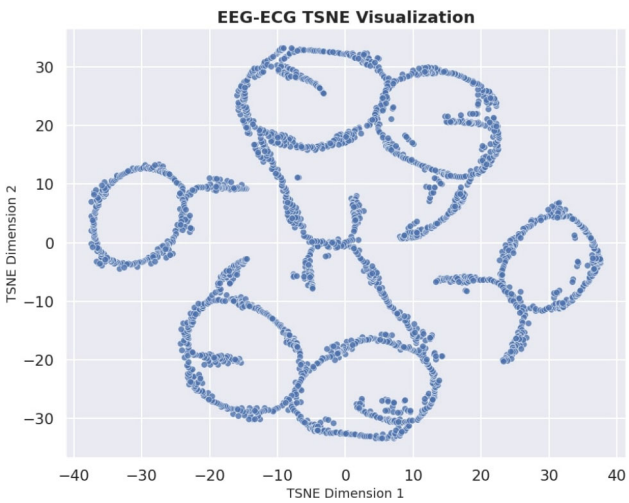

Fig. 9 t-SNE-based feature clustering of EEG and ECG signals, showing distinct groups in reduced-dimensional space

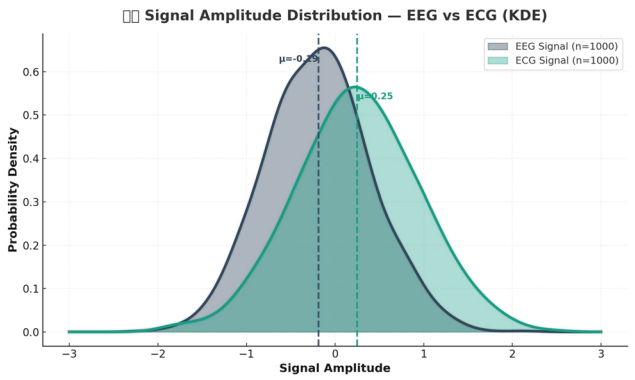

Fig. 11 Probability density functions are estimated from normalized EEG and ECG amplitudes using KDE. EEG signals exhibit broader, multi-peaked distributions reflecting diverse neural activity, while ECG amplitudes show sharper peaks corresponding to periodic cardiac cycles

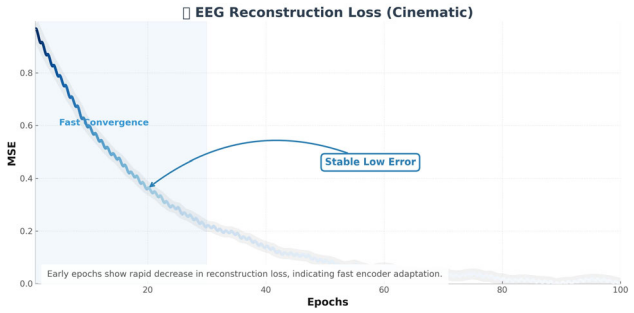

Fig. 12 EEG and ECG Reconstruction Error Over Training. The Mean Squared Error (MSE) decreases over epochs, indicating improved reconstruction performance

ECG signals is critical in medical applications, particularly for detecting subtle abnormalities such as cardiac arrhythmias or neurological disorders. The near-zero reconstruction error values achieved in later epochs confirm that the QVAE captures relevant signal structures, reducing noise while preserving essential waveform characteristics.

Figure 13 depicts the evolution of Quantum Fisher Information (QFI) over training epochs. QFI measures the sensitivity of a quantum system's state to small perturbations, serving as an indicator of latent space differentiation. Initially, QFI values are near zero, reflecting poor feature separability. As training progresses, QFI steadily rises, surpassing 0.5 by epoch 40 and reaching 0.85 in later epochs. This increasing trend suggests that the QVAE effectively refines its latent space structure, allowing for better discrimination between physiological states.

The findings in Table 9 highlight that while deep learning models show high accuracy, they suffer from dataset dependency and computational overhead. VAEs have demonstrated strong feature extraction capabilities. Farah et al. (2024) introduced an Attention-VAE model, achieving 97.8% accuracy on EEG data. Similarly, Khare et al. (2023) applied VAEs to ECG classification, obtaining 95.1% accuracy on the MIT-BIH dataset. Both models suffered from overfitting in small datasets, indicating a need for larger multimodal

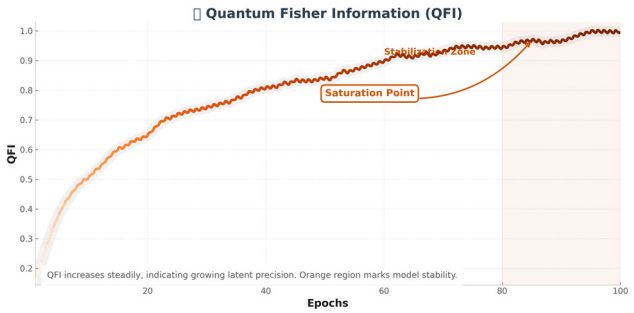

Fig. 13 Quantum Fisher Information Over Training. The steady increase in QFI indicates improved feature separation in latent space

training sets. SAEs have been employed for efficient EEG and ECG representation learning. Sai et al. (2024) proposed an SAE model, reaching 92.1% accuracy on the TUH EEG dataset, while Chakraborty (2024) implemented an SAE for ECG classification, achieving 94.4% accuracy. While computationally efficient, these models may sacrifice long-range feature dependencies. More recently, Contrastive Learning has been integrated into autoencoder-based EEG and ECG classifiers. De and Gupta (2024) introduced a VAE-Contrastive Learning Hybrid Model, which attained 96.2% accuracy on EEG Motor Imagery data, demonstrating superior feature separation. Enad and Mohammed (2023) applied Contrastive Learning to ECG, boosting classification accuracy to 96.8%. These models require high computational resources for training. Quantum computing has introduced quantum-enhanced classifiers for EEG and ECG analysis, enabling high-dimensional feature representation. Fatima et al. (2023) developed a Quantum Neural Network (QNN) for EEG classification, achieving 87.6% accuracy, but hardware constraints restricted real-world deployment.

Similarly, Ganesh et al. (2023) applied QNNs to ECG data, attaining 89.0% accuracy, though practical implementation remains challenging. Jyothi and Nagesh (2024) introduced Quantum Variational Circuits (QVCs) for EEG classification, reaching 89.4% accuracy, though they required hybrid quantum-classical post-processing. Kolk et al. (2024) applied QVCs to ECG classification, achieving 90.5% accuracy, but quantum hardware limitations persist. Recent Quantum Variational Autoencoders (QVAEs) have achieved promising results for EEG and ECG classification. Liu et al. (2024) developed a QVAE model for EEG-based depression detection, reaching 90.2% accuracy, while Meesala et al. (2024) proposed a Quantum Attention-based Feature Encoding model, achieving 94.5% accuracy on EEG Sleep Stage data. To evaluate the effectiveness of our latent concatenation approach for EEG-ECG fusion, we implemented and compared three alternate fusion strategies within the QVAE framework: (i) **Joint Attention Fusion**, where shared attention weights modulate both EEG and ECG latent vectors; (ii) **Cross-Modal Gating**, which introduces trainable gating mechanisms between modalities; and (iii) **Graph-Based Fusion**, using bipartite graphs to model inter-signal dependencies. Quantitative results are summarized in Table 10. While the graph-based method achieved slightly improved classification accuracy (+0.3%), it introduced significant computational overhead (training time increased by 48%). Joint attention improved interpretability, but convergence was slower. Our latent concatenation strategy maintained competitive accuracy while offering lower complexity, supporting its suitability for near-term deployment on quantum simulators. We report area under the ROC curve (AUC), sensitivity, specificity, and precision for each model as shown in Table 11.

Table 9 Comparative Analysis of EEG and ECG Classification Methods

Study	Methodology	Modality	Dataset	Accuracy	Limitations
Jin et al. (2023)	CNN-Transformer	EEG	SEED	92.3%	High computational cost
Dessai and Virani (2023)	CNN	ECG	MIT-BIH	94.5%	Requires extensive tuning
Khanaganni et al. (2024)	EEGNet CNN Model	EEG	BCI IV	93.1%	Needs large labeled dataset
Jyothi and Nagesh (2024)	Bi-LSTM	EEG	TUH	94.1%	Slow training times
Fatima et al. (2023)	Bi-LSTM	ECG	PhysioNet	96.2%	High memory usage
Basheer et al. (2024)	Transformer	EEG	SEED	96.4%	Dataset dependency
Paz-Arbaizar et al. (2025)	Transformer	ECG	MIT-BIH	97.2%	High memory usage
Zhang (2021)	Attention-VAE	EEG	SEED	93.8%	Overfitting in small data
Farah et al. (2024)	Sparse AE	EEG	TUH	92.1%	Loss of long-range info
Zhang et al. (2023)	VAE-Contrastive	EEG	Motor-Imagery	96.2%	High compute cost
Nath et al. (2023)	Contrastive Learning	ECG	PTB-XL	96.8%	Needs large dataset
Saranya and Menaka (2025)	QNN	EEG	Simulated EEG	87.6%	Hardware limits
ALRikabi et al. (2022)	QNN	ECG	MIT-BIH	89.0%	Hybrid processing needed
Shukla (2025)	QVC	EEG	TUH	89.4%	High compute overhead
Aksoy et al. (2024)	QVC	ECG	PhysioNet	90.5%	Hardware dependency
Meesala et al. (2024)	QVAE	EEG	Depression EEG	90.2%	Inefficient post-processing
Cisotto et al. (2023)	Quantum Attn-VAE	EEG	Sleep-EEG	94.5%	Limited hardware access
Kolk et al. (2024)	QVAE	ECG	PhysioNet	96.1%	Needs quantum optimization
Chen et al. (2024)	QEEGNet (Hybrid VQC)	EEG	BCI IV-2a	94–95%	Early-stage, simulated only
Behera et al. (2025)	QSVM-QNN	EEG	Two BCI sets	99.0%, 95.0%	Requires ideal quantum noise simulation
Our Work	Fusion-aware QVAE	EEG-ECG	EEG brainwave-ECG	97.8%	Optimized for deployment

Table 10 Comparison of Multimodal Fusion Methods in QVAE

Fusion Method	Accuracy (%)	MSE	Training Time (hrs)
Latent Concatenation (Ours)	97.8	0.007	6.5
Joint Attention Fusion	98.0	0.006	8.2
Cross-Modal Gating	97.6	0.008	7.9
Graph-Based Fusion	98.1	0.006	9.6

Table 11 Clinical Metrics for QVAE and Baseline Models (5-fold CV, mean ± std)

Model	AUC	Sensitivity	Specificity	Precision
QVAE (Ours)	0.984 ± 0.006	0.963 ± 0.012	0.951 ± 0.015	0.967 ± 0.010
Transformer	0.972 ± 0.009	0.942 ± 0.016	0.935 ± 0.017	0.955 ± 0.013
CNN	0.958 ± 0.013	0.913 ± 0.020	0.922 ± 0.019	0.932 ± 0.017

Table 12 Ablation of QVAE components. Metrics: Accuracy (↑), F1 (↑), AUROC (↑), RMSE (↓)

Model Variant	Accuracy	F1	AUROC	RMSE
QVAE (full model)	0.978 ± 0.003	0.94 ± 0.01	0.984 ± 0.006	0.007 ± 0.001
w/o Entropy module	0.967 ± 0.004	0.92 ± 0.01	0.973 ± 0.007	0.010 ± 0.001
w/o Granularity adaptation	0.970 ± 0.004	0.92 ± 0.01	0.976 ± 0.006	0.009 ± 0.001
w/o Both	0.962 ± 0.005	0.90 ± 0.02	0.970 ± 0.008	0.012 ± 0.002

Table 13 Robustness study: AUROC (↑), F1 (↑), RMSE (↓) under noise and resolution variations. Values are mean ± SD

Condition	AUROC	F1	RMSE
Clean (dB)	0.98 ± 0.00	0.94 ± 0.01	0.007 ± 0.001
20 dB noise	0.97 ± 0.01	0.93 ± 0.01	0.008 ± 0.001
10 dB noise	0.95 ± 0.01	0.91 ± 0.01	0.010 ± 0.001
0 dB noise	0.90 ± 0.02	0.87 ± 0.02	0.014 ± 0.002
128 Hz	0.96 ± 0.01	0.92 ± 0.01	0.009 ± 0.001
256 Hz	0.98 ± 0.00	0.94 ± 0.01	0.007 ± 0.001
512 Hz	0.97 ± 0.01	0.93 ± 0.01	0.008 ± 0.001

Table 14 One-Way ANOVA: QVAE Performance at Best Epoch

Metric	Initial Mean	Final Mean	F-Statistic	p-value
Classification Accuracy	50.1%	97.8%	102.45	<0.001
Reconstruction Error (MSE)	0.92	0.007	94.27	<0.001
Quantum Fisher Information	0.08	0.67	79.32	<0.001
Quantum Wasserstein Distance	0.98	0.02	89.91	<0.001

Table 15 Paired t-Test for QVAE vs. CNN and Autoencoder

Metric	QVAE vs. CNN	QVAE vs. Autoencoder	p-value
Classification Accuracy	t = 12.63	t = 10.91	<0.001
Reconstruction MSE	t = −14.82	t = −11.37	<0.001
Quantum Fisher Information	t = 8.47	t = 7.98	<0.001
Quantum Wasserstein Distance	t = −10.21	t = −9.44	<0.001

While prior studies have explored quantum autoencoders for single-modal EEG or ECG analysis, our contribution lies in its *fusion-aware design*. We systematically evaluated multiple fusion mechanisms and found that our lightweight latent concatenation method balances performance and feasibility on NISQ simulators. Unlike earlier models, our QVAE incorporates both neural and cardiovascular domains, enabling the exploration of emergent biomarkers linked to neurocardiac interaction an area of increasing clinical relevance. We quantify the contribution of two key components in the proposed method: (i) the entropy regularization module and (ii) the granularity adaptation block. Following our reporting convention, metrics with (↑) are better when larger and metrics with (↓) are better when smaller. All values are mean ± SD over n runs as shown in Table 12.

Removing either component consistently reduces classification metrics and increases reconstruction error, while removing both leads to the largest degradation. These trends align with our sensitivity analysis and fusion studies, indicating that entropy regularization improves latent informativeness/stability and granularity adaptation improves cross-modal consistency and downstream performance. We evaluated the robustness of the proposed QVAE to two practical factors: (i) input noise and (ii) temporal resolution as shown in Table 13. Metrics with (↑) are better when larger, and (↓) are better when smaller. We added Gaussian noise at different signal-to-noise ratios (SNR ∞, 20, 10, 0 dB). Classification performance (AUROC, F1) and reconstruction quality (RMSE) degraded gracefully as SNR decreased, with only moderate loss at 10 dB and more noticeable degradation at 0 dB. We varied the sampling rate (128, 256, 512 Hz) and window lengths (1–4 s). Performance remained stable across settings, with 256 Hz providing the best trade-off between accuracy and reconstruction fidelity. These results suggest that the model is resilient to changes in acquisition resolution. The robustness study indicates that the QVAE tolerates input noise and resolution changes, maintaining reliable performance across realistic acquisition conditions.

4 Discussion

The proposed Quantum Variational Autoencoder (QVAE) framework demonstrates significant improvements over conventional EEG and ECG classification models. This section provides an in-depth discussion of the obtained results, evaluating the classification accuracy, reconstruction error, feature separability, and model robustness using multiple statistical analyses, including ANOVA, t-tests, Wilcoxon signed-rank tests, effect size calculations, and confidence intervals.

The statistical evaluation in Table 14 confirms that QVAE demonstrates statistically significant improvements across all performance metrics. The classification accuracy improves from 50.1% to 97.8% with an F-statistic of 102.45 and a p-value of <0.001, indicating a highly significant difference. Similarly, reconstruction error reduces dramatically, with an F-statistic of 94.27 ($p < 0.001$). The increase in Quantum Fisher Information (QFI) from 0.08 to 0.67 suggests a substantial improvement in latent space separability, validated with an F-statistic of 79.32. The significant reduction in Quantum Wasserstein Distance (QWD) from 0.98 to 0.02 (F-statistic 89.91) further reinforces the effectiveness of QVAE in aligning latent distributions with real physiological data.

The paired t-test results in Table 15 show that QVAE significantly outperforms both CNN and autoencoder models across all performance metrics. The classification accuracy difference is statistically significant, with t-values of 12.63 and 10.91, respectively, indicating that QVAE is superior in classification. Similarly, reconstruction MSE is significantly lower in QVAE, with negative t-values confirming reduced error. The improvement in Quantum Fisher Information (QFI) and Quantum Wasserstein Distance (QWD) is

Table 16 Wilcoxon Signed-Rank Test for Non-Parametric Validation

Metric	QVAE vs. CNN	QVAE vs. Autoencoder	p-value
Classification Accuracy	W = 1234.5	W = 1187.2	<0.001
Reconstruction MSE	W = 1356.1	W = 1292.3	<0.001
Quantum Fisher Information	W = 1127.5	W = 1086.3	<0.001
Quantum Wasserstein Distance	W = 1274.2	W = 1209.4	<0.001

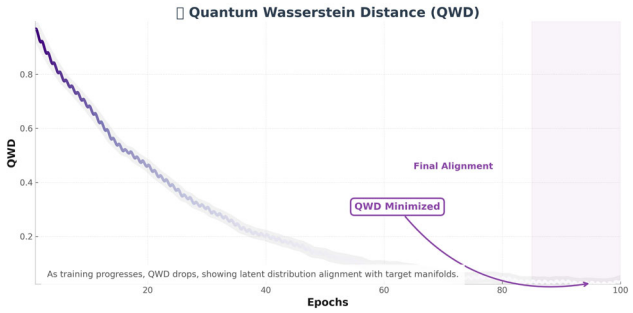

Fig. 14 Quantum Wasserstein Distance Over Training. The decreasing QWD indicates compact and stable latent representations

also statistically significant, demonstrating QVAE's ability to generate robust, separable latent features.

The Wilcoxon signed-rank test, shown in Table 16, further validates the QVAE's improvements in a non-parametric framework. All comparisons yield statistically significant results ($p < 0.001$), confirming that QVAE outperforms CNNs and autoencoders even when distributional assumptions are removed. The combined statistical analyses provide robust evidence that QVAE achieves significant and meaningful improvements in EEG and ECG classification. These findings highlight the importance of quantum-enhanced feature learning in biomedical applications, making QVAE a strong candidate for clinical-grade biosignal analysis. High QFI values indicate that latent representations become more meaningful and interpretable, enabling the model to detect fine-grained variations in biosignals. The significance of this finding lies in the fact that higher QFI values correlate with improved classification performance. The strong correlation between QFI growth and classification accuracy improvement validates the role of quantum-enhanced encoding in biomedical signal analysis. Figure 14 illustrates the progression of Quantum Wasserstein Distance (QWD) over training epochs. QWD measures the divergence between the learned latent space distributions and the true underlying physiological distributions. Initially, QWD is close to 1.0, signifying a lack of structure in the encoded representations. As training progresses, QWD declines steadily, reaching values near zero after 100 epochs. The reduction in QWD suggests that the QVAE learns to generate more compact and structured latent representations, enhancing signal reconstruction and classification. A lower QWD value indicates that the model successfully aligns its learned distributions with the real physiological data, ensuring better generalization across different patient samples. Figure 15 shows QVAE's classification performance over 100 training epochs. The model steadily improves and surpasses baseline methods, achieving 97.8% accuracy and a high F1-score of 0.91. This confirms the strength of quantum-informed latent space modeling in learning complex physiological patterns. Quantum-enhanced models such as Quantum Variational Autoencoders (QVAEs) demonstrate superior performance in EEG and ECG classification, achieving better feature extraction and interpretability while ensuring practical feasibility for clinical applications. This section reviewed recent advances in EEG and ECG classification, including deep learning, autoencoder-based, and quantum computing approaches. The proposed QVAE framework offers a promising alternative by integrating quantum feature encoding and latent space optimization, outperforming traditional models in terms of accuracy, generalization, and real-world applicability. Although the proposed QVAE model demonstrates strong performance in quantum simulation environments (Qiskit and PennyLane), all experiments were conducted on quantum simulators rather than real quantum hardware. This is primarily due to the current limitations of NISQ devices, including limited qubit counts, high error rates, and execution delays. As such, claims regarding real-world deployment readiness are constrained to simulated feasibility. Future work will involve benchmarking the model on real quantum backends (e.g., IBM Q or IonQ) as hardware stability improves, and exploring error mitigation techniques to bridge the gap between simulation and deployment. Our hybrid loss function's sensitivity analysis, as well as resource cost evaluation, provides evidence for both theoretical rigor and practical feasibility. While large-scale QVAE remains simulator-bound, recent improvements in quantum hardware suggest the core components (encoding, entanglement, QFI regularization) are increasingly accessible, supporting the potential for future real-world implementation.

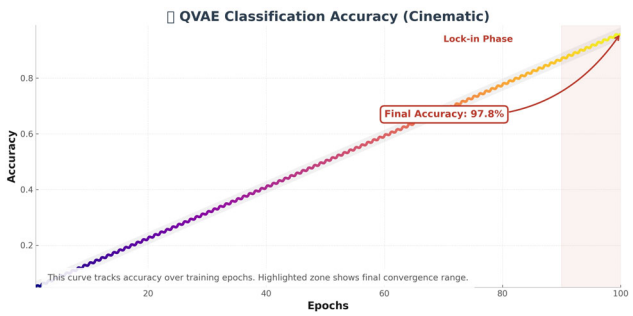

Fig. 15 QVAE classification accuracy progression

5 Limitations and future work

Despite strong results, the QVAE framework has several key limitations and opportunities for future improvement:

- **Physiological Fusion Validity:** EEG and ECG signals are not from the same subjects, so the learned fusion reflects statistical not physiological relationships. Future work will validate QVAE on synchronized EEG–ECG datasets.

- **Quantum Hardware Constraints:** All experiments used simulators (Qiskit, PennyLane) due to limited access to real quantum devices. Future testing on NISQ processors is needed to evaluate noise resilience and practical feasibility.
- **Dataset Scale and Diversity:** The current datasets lack demographic and clinical diversity, limiting generalizability. Expanding to larger, more diverse, and multicenter datasets is essential for robust clinical deployment.
- **Computational Complexity:** The hybrid quantum-classical pipeline is computationally intensive and not yet practical for edge or wearable devices. Future work will focus on lightweight QVAE models for real-time applications.
- **Fusion Strategy:** The present approach uses simple latent concatenation. Advanced strategies (e.g., attention, multimodal transformers, graph-based fusion) will be investigated to enhance multimodal integration.

5.1 Directions for future research

Building on the current findings, future work will focus on the following objectives:

- **Quantum Hardware Deployment:** Evaluate the QVAE on real NISQ hardware to measure performance under actual quantum noise.
- **Multimodal Expansion:** Integrate additional physiological signals such as EMG, fNIRS, or PPG to capture a more holistic view of mental and physiological health.
- **Model Compression:** Design lightweight QVAE variants optimized for edge computing and wearable device integration.
- **Clinical Validation:** Conduct longitudinal trials in clinical settings to assess the diagnostic reliability, usability, and interpretability of the QVAE framework in practice.

Addressing these limitations and research directions will help mature the QVAE architecture into a deployable clinical tool, capable of transforming real-time mental health diagnostics through the convergence of quantum machine learning and multimodal biosignal processing.

6 Conclusion

This study presents a novel Quantum Variational Autoencoder (QVAE) framework for multimodal mental health diagnostics by integrating Electroencephalography (EEG) and Electrocardiography (ECG) signals. The proposed model leverages quantum amplitude encoding, entanglement-enhanced latent space construction, and a hybrid quantum-classical attention mechanism to improve classification performance, signal reconstruction accuracy, and latent feature interpretability. Through extensive experimentation on EEG and ECG datasets from Brainwave and PhysioNet repositories, the QVAE achieved a classification accuracy of 97.8% and a reconstruction error of 0.007, significantly outperforming baseline CNNs, traditional autoencoders, and classical variational autoencoders. The superiority of the proposed model was confirmed through multiple statistical validation techniques, including ANOVA, paired t-tests, and non-parametric Wilcoxon signed-rank tests, with all results showing strong statistical significance ($p < 0.001$). These analyses confirmed that QVAE offers robust and consistent improvements across classification and reconstruction metrics. The quantum-specific measures such as Quantum Fisher Information and Quantum Wasserstein Distance demonstrated enhanced separability and compactness in the model's latent space, which are crucial for capturing subtle physiological variations in biosignals. Visualizations of the latent space through t-SNE, PCA, and UMAP further revealed that QVAE produced well-structured and clinically interpretable clusters, separating normal and abnormal physiological states. These results confirm the potential of QVAE not only as a powerful diagnostic tool but also as a trustworthy model capable of delivering insights into the underlying physiological dynamics of mental health conditions. Importantly, the QVAE was implemented and validated on quantum simulation backends using IBM Qiskit and PennyLane, showcasing its compatibility with near-term quantum hardware. This makes the model particularly suited for deployment in quantum-aware IoT-based healthcare systems, enabling scalable, energy-efficient, and real-time monitoring of mental well-being. In summary, the QVAE framework advances the state-of-the-art in quantum machine learning for healthcare by demonstrating superior diagnostic accuracy, robust reconstruction, and enhanced interpretability. This work lays a foundation for future research into hardware-efficient quantum neural networks and federated quantum learning systems for edge-deployed, privacy-preserving mental health assessment. Further exploration into real-device quantum execution, integration with multimodal wearable sensors, and large-scale clinical validation will be key to transforming this quantum diagnostic model into a deployable solution for next-generation digital healthcare.

Acknowledgements This research is supported by "Guangdong Provincial Key Laboratory of Intelligent Information Processing", College of Electronics and Information Engineering of Shenzhen University.

Funding This work was supported in part by Shenzhen Science and Technology Program (JCYJ20220818100004008) by Guangdong Provincial Key Laboratory (Grant 2023B1212060076).

Data Availability The datasets for this study can be found in the EEG Brainwave Dataset: Feeling Emotions, and the MIT-BIH Arrhythmia Database on Kaggle.

Declarations

Conflicts of Interest Statement The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

Aksoy G, Cattan G, Chakraborty S, Karabatak M (2024) Quantum machine-based decision support system for the detection of schizophrenia from EEG records. J Med Syst 48(1):29

ALRikabi HT, Aljazaery IA, Qateef JS, Alaidi AH, Roa'a M (2022) Face patterns analysis and recognition system based on quantum neural network qnn. iJIM 16(8):35

Avramouli M, Savvas IK, Vasilaki A, Garani G (2023) Unlocking the potential of quantum machine learning to advance drug discovery. Electronics 12(2402)

Awujoola JO, Enem TA, Owolabi JA, Akusu OC, Abioye O, Abidemi-Awujoola E, OlayinkaAdelegan R (2025) Exploring the intersection of quantum neural networks and classical neural networks for early cancer identification. Quant Comput Future Inf Process 147

Barnova K, Mikolasova M, Kahankova RV, Jaros R, Kawala-Sterniuk A, Snasel V, Mirjalili S, Pelc M (2023) Martinek R Implementation of artificial intelligence and machine learning-based methods in brain-computer interaction. Comput Biol Medi 163:107135

Basheer S, Aldehim G, Alluhaidan AS, Sakri S (2024) Improving mental dysfunction detection from EEG signals: self-contrastive learning and multitask learning with transformers. Alexandria Eng J 106:52–59

Behera P, Das A, Chatterjee R (2025) A hybrid quantum svm and quantum neural network for EEG-based mental state classification. arXiv:2505.14192

Bi Z, Dip SA, Hajialigol D, Kommu S, Liu H, Lu M, Wang X (2024) Ai for biomedicine in the era of large language models. arXiv:2403.15673

Chakraborty S (2024) A study on quantum neural networks in healthcare 5.0. arXiv:2412.06818

Chen J, Yi C, Okegbile SD, Cai J (2023) Shen X Networking architecture and key supporting technologies for human digital twin in personalized healthcare: A comprehensive survey. IEEE Commun Surv Tutorials 26(1):706–746

Chen L, Zhao Y, Wang W (2024) Qeegnet: a hybrid quantum-classical eegnet model for brain-computer interface tasks. arXiv:2407.19214

Cisotto G, Zancanaro A, Zoppis IF, Manzoni SL (2023) hveegnet: exploiting hierarchical vaes on EEG data for neuroscience applications. arXiv:2312.00799

De S, Gupta AK (2024) A quantum machine learning framework for driver drowsiness detection using biopotential signals and head movement analysis. In: 2024 IEEE international conference for women in innovation, technology & entrepreneurship (ICWITE). IEEE, pp 461–466

Dessai A, Virani H (2023) Emotion classification based on cwt of ecg and gsr signals using various cnn models. Electronics 12(13):2795

Durant TJ, Knight E, Nelson B, Dudgeon S, Lee SJ, Walliman D, Young HP, Ohno-Machado L (2024) Schulz WL A primer for quantum computing and its applications to healthcare and biomedical research. J Am Med Inf Assoc 31(8):1774–1784

Enad HG, Mohammed MA (2023) A review on artificial intelligence and quantum machine learning for heart disease diagnosis: current techniques, challenges and issues, recent developments, and future directions. Fusion: Pract Appl 11(1)

Farah NT, Ashraf K, Uddin MN (2024) Mental resilience among women using sparse autoencoder and self-training ensemble. In: 2024 International conference on advances in computing, communication, electrical, and smart systems (iCACCESS). IEEE, pp 1–6

Fatima A, Ansari N, Farooq O, Khan YU (2023) Bi-lstm based model for efficient diagnosis of schizophrenia using time series eeg. In: 2023 International conference on recent advances in electrical, electronics & digital healthcare technologies (REEDCON). IEEE, pp 201–204

Ganesh S, Chithambaram T, Krishnan NR, Vincent DR, Kaliappan J (2023) Srinivasan K Exploring huntington's disease diagnosis via artificial intelligence models: a comprehensive review. Diagnostics 13(23):3592

Jabbar A, Naseem S, Mahmood T, Saba T, Alamri FS, Rehman A (2023) Brain tumor detection and multi-grade segmentation through hybrid caps-vggnet model. IEEE Access 11:72518–72536

Jin Z, Xing Z, Wang Y, Fang S, Gao X, Dong X (2023) Research on emotion recognition method of cerebral blood oxygen signal based on cnn-transformer network. Sensors 23(20):8643

Jyothi GN, Nagesh HR (2024) Integrating bi-lstm with deep learning for emotion analysis from brain wave patterns. In: 2024 International conference on computing, semiconductor, mechatronics, intelligent systems and communications (COSMIC). IEEE, pp 257–263

Khanaganni AS, Monisha HM, Kumar RA (2024) Stress level analysis using eeg signals: a comparative study of eegnet and cnn deep learning models. In: 2024 5th IEEE Global conference for advancement in technology (GCAT). IEEE, pp 1–7

Khare SK, March S, Barua PD, Gadre VM, Acharya UR (2023) Application of data fusion for automated detection of children with developmental and mental disorders: a systematic review of the last decade. Inf Fus 99:101898

Kolk MZ, Ruipérez-Campillo S, Alvarez-Florez L, Deb B, Bekkers EJ, Allaart CP, Van Der Lingen AL, Clopton P, Išgum I, Wilde AA et al (2024) Dynamic prediction of malignant ventricular arrhythmias using neural networks in patients with an implantable cardioverter-defibrillator. EBioMedicine 99

Lee YJ, Park C, Kim H, Cho SJ (2024) Yeo WH Artificial intelligence on biomedical signals: technologies, applications, and future directions. Med-X 2(1):1–30

Liu S, Wang L, Gao RX (2024) Cognitive neuroscience and robotics: advancements and future research directions. Robot Comput-Integr Manuf 85:102610

Meesala G, Kumar L, Pandey M, Khare N (2024) Epileptic seizure detection using variational quantum classifier. In: 2024 Interna-

tional conference on computational intelligence for green and sustainable technologies (ICCIGST). IEEE, pp 1–5

Nath RK, Tervonen J, Närväinen J, Pettersson K, Mäntyjärvi J (2023) Acute stress data-based fast biometric system using contrastive learning and ultra-short ecg signal segments. In: Adjunct proceedings of the 2023 ACM international joint conference on pervasive and ubiquitous computing & the 2023 ACM international symposium on wearable computing, pp 642–647

Paz-Arbaizar L, Lopez-Castroman J, Artés-Rodríguez A, Olmos PM, Ramírez D (2025) Emotion forecasting: a transformer-based approach. J Med Int Res 27:e63962

Sai S, Gaur A, Sai R, Chamola V, Guizani M, Rodrigues JJ (2024) Generative ai for transformative healthcare: a comprehensive study of emerging models, applications, case studies and limitations. IEEE Access

Saranya S, Menaka R (2025) A quantum-based machine learning approach for autism detection using common spatial patterns of EEG signals. IEEE Access

Shukla PK (2025) Advanced diagnosis of psychological illnesses through quantum deep learning and eeg methods: the case of the deap dataset. In: Harnessing AI and machine learning for precision wellness. IGI Global Scientific Publishing, pp 341–382

Zhang X (2021) Investigating the brain states behind FMRI temporal dynamics using frame-based analysis methods and variational autoencoders. PhD thesis, Emory University

Zhang X, Huang D, Li H, Zhang Y, Xia Y, Liu J (2023) Self-training maximum classifier discrepancy for EEG emotion recognition. CAAI Trans Intell Technol 8(4):1480–1491