Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Feature-Adaptive Fusion in Hybrid Quantum–Classical Neural Networks for Robust Biomedical Image Classification

Yan-Yan Hou
Zaozhuang University
Jian Li
Beijing University of Posts and Telecommunications
Chongqiang Ye
Hangzhou City University
Hengji Li
Henan University
Zhuo Wang
China Information Technology Security Evaluation Center
Qinghui Liu
qinghuiliu810716@126.com

Zaozhuang University

Research Article

Keywords: hybrid quantum-classical neural networks, variational quantum circuits, quantum machine learning, adaptive fusion, robustness, biomedical image classification

Posted Date: July 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9777113/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at Quantum Machine Intelligence on November 4th, 2026. See the published version at https://doi.org/10.1007/s42484-026-00467-5.

Feature-Adaptive Fusion in Hybrid Quantum–Classical Neural Networks for Robust Biomedical Image Classification

Yan-Yan Hou[1], Jian Li[2†], Chongqiang Ye[3], Hengji Li[4], Zhuo Wang[5], Qinghui Liu[6*]

[1]College of Information Science and Engineering, Zaozhuang University, No. 1 Beian Road, Zaozhuang, 277160, Shandong, China.
[2]School of Cyberspace Security, Beijing University of Posts and Telecommunications, Beijing, 100876, China.
[3]Hangzhou City University, Hangzhou, 310015, China.
[4]School of Artificial Intelligence, Henan University, 85 Nanyang Road, Zhengzhou, 450000, Henan, China.
[5]China Information Technology Security Evaluation Center, Beijing, 100085, China.
[6]Network Center, Zaozhuang University, No. 1 Beian Road, Zaozhuang, 277160, Shandong, China.

*Corresponding author(s). E-mail(s): qinghuiliu810716@126.com;
Contributing authors: hyy@uzz.edu.cn; lijian@bupt.edu.cn;
chongqiangye@hzcu.edu.cn; lihengji@henu.edu.cn;
ZhuoWang@bupt.edu.cn;
†These authors contributed equally to this work.

Abstract

Hybrid quantum-classical neural networks provide a promising approach for incorporating quantum circuits into machine learning in the noisy intermediate-scale quantum regime. However, existing hybrid models often rely on fixed or globally shared fusion strategies, which may limit their ability to exploit complementary information carried by quantum branches, especially under distribution shifts. In this work, we propose a Feature-Adaptive Fusion Hybrid Quantum-Classical Neural Network (FAF-HQNN) for biomedical image classification. The model combines a classical deep feature encoder with a variational quantum circuit (VQC) and introduces a feature-adaptive fusion mechanism to dynamically weight classical and quantum predictions. We evaluate FAF-HQNN on two

MedMNIST benchmarks, PathMNIST and BloodMNIST, under clean and corrupted test conditions. FAF-HQNN achieves the strongest overall performance on clean data among the compared methods and shows improved robustness under Gaussian, salt-and-pepper, and Poisson corruptions. Further analysis of circuit layout, measurement basis, and depth shows that even shallow variational quantum circuits can provide useful complementary information. These results demonstrate that feature-adaptive fusion is an effective strategy for improving accuracy and robustness in hybrid quantum-classical models for biomedical image classification.

Keywords: hybrid quantum-classical neural networks, variational quantum circuits, quantum machine learning, adaptive fusion, robustness, biomedical image classification

1 Introduction

Hybrid quantum-classical neural networks have emerged as a practical paradigm for quantum machine learning in the noisy intermediate-scale quantum (NISQ) era, where current quantum devices remain constrained by limited qubit counts, shallow circuit depth, and imperfect quantum operations [1]. Under these hardware limitations, purely quantum models are difficult to scale to high-dimensional real-world tasks. Hybrid architectures provide a feasible alternative by combining the strong representation learning capabilities of classical neural networks with trainable quantum components, typically implemented as variational quantum circuits (VQCs), within end-to-end learning frameworks [2].

Biomedical image classification plays an important role in computer-aided diagnosis and medical data analysis, supporting tasks such as tissue recognition, blood cell identification, and disease screening. Deep learning methods, especially convolutional neural networks (CNNs), have achieved strong performance in these applications by learning hierarchical visual representations directly from biomedical images [3–5]. In recent years, there has been growing interest in hybrid quantum-classical CNN models for image analysis and biomedical data processing [6–10]. However, most studies have focused primarily on feasibility and clean-data performance [11–13].

In practice, biomedical images are frequently affected by common image corruptions, particularly Gaussian and Poisson noise, which are closely related to sensor noise and acquisition variability in imaging systems and can shift input distributions and degrade predictive reliability [14]. Robustness to such degradations is therefore an important requirement for practical biomedical image classification. Nevertheless, the behavior of hybrid quantum-classical neural networks under realistic input corruptions remains insufficiently understood. This issue is particularly important in biomedical applications, where robustness must be assessed under degradations commonly encountered in practical imaging workflows. In this work, our primary focus is on robustness to classical input corruptions applied to biomedical images.

A second limitation of existing hybrid quantum-classical neural networks lies in their typically fixed strategies for integrating classical and quantum components. Many existing architectures adopt fixed integration schemes, often taking sequential

or replacement-based forms, in which the quantum module is appended after a classical encoder or replaces part of a classical classifier. Although these designs are simple and compatible with near-term hardware constraints, they assign largely predetermined roles to the classical and quantum branches across different inputs. This design choice may be restrictive under corrupted input conditions, where the relative usefulness of classical and quantum representations may vary with image content, corruption type, and corruption severity. As a result, a fixed integration strategy may fail to fully exploit the complementary strengths of the two branches. These considerations suggest that hybrid learning may benefit from a more adaptive fusion mechanism that dynamically weights branch contributions according to the input.

To address the limitations of fixed integration strategies in existing hybrid quantum-classical models, we propose a Feature-Adaptive Fusion Hybrid Quantum-Classical Neural Network (FAF-HQNN) for biomedical image classification. The proposed architecture combines a ResNet-18-based deep feature encoder with a variational quantum circuit (VQC) branch in an end-to-end trainable framework. Its key innovation is a feature-adaptive class-wise fusion mechanism that predicts sample-dependent fusion weights from the learned feature representation and dynamically combines the outputs of the classical and quantum branches. By allowing branch contributions to vary with the input, the model provides a more flexible means of exploiting complementary classical and quantum information, which is particularly relevant under corrupted input conditions.

We evaluate the proposed framework on two representative biomedical image benchmarks, PathMNIST and BloodMNIST. Beyond standard evaluation on clean data, we systematically assess robustness under multiple types of classical input corruption. We further conduct supplementary analyses under simulated quantum noise to examine model behavior under NISQ-relevant conditions. To validate the effectiveness of the proposed adaptive fusion strategy, we compare FAF-HQNN with classical baselines, hybrid models without fusion, and alternative hybrid fusion baselines in terms of predictive accuracy and robustness.

The main contributions of this work are summarized as follows:

- We propose FAF-HQNN, a hybrid quantum-classical architecture for biomedical image classification that combines a deep classical feature encoder with a VQC-based quantum branch in an end-to-end trainable framework.
- We introduce a feature-adaptive class-wise fusion mechanism that predicts fusion weights from learned representations, enabling dynamic integration of classical and quantum outputs across inputs and classes.
- Experimental results on PathMNIST and BloodMNIST show that FAF-HQNN improves robustness to input perturbations while preserving competitive performance on clean data compared with classical baselines, hybrid models without fusion, and alternative hybrid fusion baselines.

The rest of this paper is organized as follows. Section 2 reviews related work on classical and hybrid quantum-classical approaches. Section 3 presents the proposed FAF-HQNN architecture. Section 4 describes the experimental setup and presents the

results. Section 5 discusses the findings, limitations, and directions for future research. Finally, Section 6 concludes the paper.

2 Related Work

This section reviews the prior research most relevant to the present study, focusing on robustness in classical biomedical image classification and hybrid quantum-classical learning for image-based tasks.

2.1 Classical Deep Learning and Robustness in Biomedical Image Classification

Deep learning has achieved strong performance in biomedical image classification, with applications in computer-aided diagnosis, digital pathology, and cellular image analysis [15]. Among existing approaches, convolutional neural networks (CNNs) have been widely adopted for tasks such as histopathology classification, radiographic diagnosis, and blood cell recognition, owing to their ability to learn hierarchical visual representations directly from image data [16–19]. Residual networks, in particular, remain popular CNN backbones due to their optimization stability, scalability, and strong empirical performance, including in limited-data settings [20]. Transformer-based architectures and other global-context modeling approaches have also been explored in medical imaging [21, 22]. However, despite their stronger capacity for modeling long-range dependencies, CNN-based encoders remain competitive in many biomedical classification tasks because of their computational efficiency and favorable inductive bias, particularly on compact benchmark datasets.

Despite strong performance on clean benchmark data, biomedical image classifiers remain vulnerable to image corruptions such as Gaussian and Poisson noise, which are relevant to sensor-related noise and acquisition variability in imaging systems [23]. These perturbations can shift input distributions and compromise predictive reliability in real-world deployment settings. To improve robustness, prior studies have explored data augmentation, corruption-based training, adversarial training, and regularization-based strategies [24, 25]. However, compared with standard clean-data evaluation, robustness evaluation across multiple corruption types and severity levels remains relatively limited in biomedical image classification.

Standardized benchmarks play an important role in robustness research by enabling controlled, reproducible evaluation under consistent data formats and protocols. MedMNIST provides a lightweight collection of biomedical image benchmarks designed for efficient experimentation and standardized evaluation, making it well suited for both classical and hybrid quantum-classical learning settings [26]. Among its subsets, PathMNIST serves as a representative benchmark for multi-class histopathology classification, whereas BloodMNIST offers a complementary task centered on blood cell recognition. Their biomedical relevance, standardized formulation, and compact scale make them practical testbeds for assessing predictive performance under controlled image corruptions.

Overall, prior work in classical biomedical image classification has established strong baselines and highlighted the practical importance of robustness under realistic perturbations. In contrast, systematic robustness evaluation has been incorporated only to a limited extent in hybrid quantum-classical models. As a result, a clear gap remains at the intersection of biomedical image robustness evaluation and hybrid quantum-classical learning.

2.2 Hybrid Quantum-Classical Learning for Image Classification

Quantum machine learning has attracted growing attention as a framework for integrating quantum information processing into classical learning pipelines [27]. However, current quantum hardware remains in the NISQ regime and is constrained by limited qubit counts, shallow circuit depth, and imperfect quantum operations [28]. These constraints make it challenging to apply purely quantum models directly to high-dimensional image classification problems, where input representations are typically large and structurally complex. As a result, hybrid quantum-classical architectures have become a common approach in image-based tasks, where classical neural networks are used to extract compact and informative features, while trainable quantum modules, typically variational quantum circuits (VQCs), are applied to low-dimensional latent representations within an end-to-end framework.

A common design in hybrid quantum-classical learning for image tasks is to use a classical encoder to compress high-dimensional image representations before passing them to a quantum circuit [29]. This strategy reduces the dimensionality of the quantum input, improves compatibility with near-term hardware, and preserves the representational strength of classical deep models. Because such architectures are primarily motivated by the need to combine expressive classical feature extraction with hardware-constrained quantum processing, much of the existing literature has focused on proof-of-concept validation and benchmark performance under standard evaluation settings. Recent studies have explored these hybrid architectures in image analysis and biomedical learning [30–33], providing initial evidence that hybrid quantum-classical models are feasible for small- to medium-scale problems under current hardware constraints.

In image classification, existing hybrid quantum-classical models mainly adopt three types of integration patterns. A common design uses a classical backbone for feature extraction and places the quantum module at the classification stage, allowing the quantum circuit to operate on compressed representations rather than raw high-dimensional inputs [34–37]. A second line of work incorporates the quantum component into the prediction head, typically by replacing part of the final decision layers with a variational quantum circuit [38–40]. A third approach introduces a parallel quantum branch and fuses its outputs with those of the classical pathway through shared trainable decision layers [41–44]. Although these strategies are attractive because of their structural simplicity and hardware compatibility, they generally offer limited adaptivity in balancing classical and quantum contributions across

samples. This constraint may be particularly relevant in biomedical image classification, where image quality, structural complexity, and corruption severity can vary substantially across samples.

This limitation is further amplified by the fact that most prior work has focused primarily on feasibility demonstrations and clean-data performance, whereas robustness under realistic perturbations remains relatively underexplored. In particular, limited attention has been paid to how hybrid quantum-classical models behave under biomedical image corruptions and to whether more adaptive integration mechanisms can improve robustness beyond conventional fusion strategies. Motivated by these gaps, we investigate whether feature-adaptive fusion can provide a more robust integration mechanism for biomedical image classification under both standard and corruption-based evaluation settings.

3 Methodology

This section presents the proposed Feature-Adaptive Fusion Hybrid Quantum-Classical Neural Network (FAF-HQNN) for biomedical image classification. We first introduce the overall architecture and its design rationale, and then describe its main components in detail, including classical feature extraction, the feature compression module, hybrid prediction and feature-adaptive fusion, and the training objective.

3.1 Overall Architecture

Figure 1 presents the overall architecture of FAF-HQNN. As shown in Fig. 1(a), the model follows a preprocessing–feature extraction–feature compression–feature-adaptive fusion pipeline for biomedical image classification. In this pipeline, the feature compression module projects high-dimensional visual features into a low-dimensional shared embedding space, which serves as the interface between classical representation learning and downstream hybrid prediction.

Built upon this shared embedding, FAF-HQNN departs from conventional serial hybrid designs by adopting a parallel prediction strategy. Specifically, the shared embedding is fed into both the quantum and classical branches, while a lightweight fusion coefficient predictor generates sample-dependent, class-wise fusion coefficients to adaptively combine the logits produced by the two branches. In this way, the quantum branch serves as a complementary predictor rather than a direct replacement for the classical classifier, allowing the model to dynamically balance the relative contributions of the two branches across different input samples and categories. Overall, FAF-HQNN integrates classical deep feature extraction, quantum-enhanced prediction, and feature-adaptive fusion within a unified end-to-end trainable framework.

3.2 Feature Extraction

This module extracts high-level visual representations from biomedical images for downstream hybrid prediction. In FAF-HQNN, feature extraction is implemented

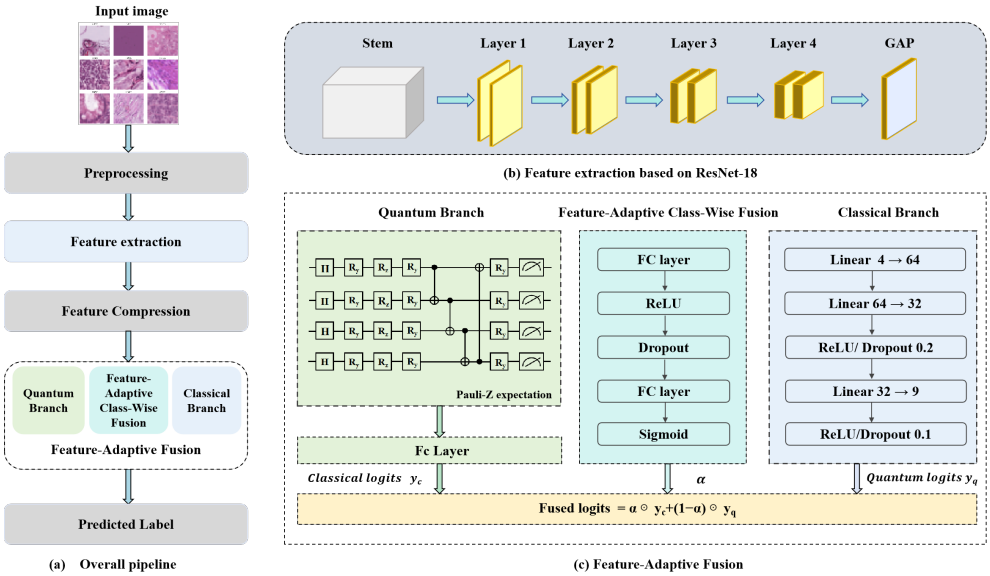


Fig. 1: Overall framework of the proposed FAF-HQNN model. (a) Pipeline of FAF-HQNN. Input biomedical images are first preprocessed and then fed into the feature extraction module, followed by a feature compression module for quantum-compatible embedding. The resulting shared embedding is subsequently passed to feature-adaptive fusion module, including a quantum branch, a feature-adaptive class-wise fusion mechanism, and a classical branch, whose outputs are combined to produce the final prediction. (b) Architecture of feature extraction based on ResNet-18, consisting of a stem block, four residual stages with two BasicBlocks in each stage, and a global average pooling (GAP) layer. (c) Structure of the feature-adaptive fusion. The embedding features are processed in parallel by a quantum branch implemented with a variational quantum circuit and a classical branch implemented as a multilayer perceptron. In parallel, a fusion coefficient predictor estimates class-wise adaptive fusion coefficients, which are used to combine the classical and quantum logits into the final output.

using ResNet-18 initialized with ImageNet-pretrained weights, which offers a suitable trade-off between representational capacity and computational complexity.

Specifically, the feature extractor consists of the convolutional stem and four residual stages of ResNet-18, while the original fully connected classification head is removed. Since the backbone is used solely for representation learning, it outputs deep visual features rather than dataset-specific classification logits. To obtain a fixed-length representation, the output of the final residual stage is further processed by a global adaptive average pooling layer. As a result, each input image is mapped to a 512-dimensional feature vector.

This extracted 512-dimensional feature vector is then passed to the subsequent feature compression module. By using an ImageNet-pretrained ResNet-18 as the classical backbone, the model benefits from transferable visual representations while

maintaining a relatively lightweight architecture suitable for the proposed hybrid quantum-classical framework.

3.3 Feature Compression for Quantum-Compatible Embedding

Although the ResNet-18 backbone provides informative deep features, the resulting 512-dimensional representation remains too high-dimensional for direct quantum processing under NISQ constraints. To address this issue, we introduce a feature compression module that maps the backbone output to a compact shared embedding suitable for both quantum and classical prediction.

Let $\mathbf{h}_i \in \mathbb{R}^{512}$ denote the feature vector extracted from the i-th input image. The compression module first applies a learnable nonlinear projection to reduce the dimensionality of $\mathbf{h}_i$. Specifically, it consists of two fully connected layers with batch normalization and ReLU activation, which transform the feature dimension from 512 to 64. This step suppresses redundancy in the backbone representation while preserving the information most relevant to classification. This compact feature is then passed through a learnable embedding layer to produce the final shared embedding for downstream hybrid prediction.

Let $\mathbf{h}'_i \in \mathbb{R}^{64}$ denote the compressed feature, and let $f_{\text{emb}} : \mathbb{R}^{64} \to \mathbb{R}^{n_q}$ denote the learnable embedding mapping. To ensure numerical compatibility with parameterized quantum gates, the embedding output is passed through a hyperbolic tangent activation and rescaled to the interval $[-\pi, \pi]$. The shared embedding is defined as $\mathbf{z}_i = \pi \cdot \tanh(f_{\text{emb}}(\mathbf{h}'_i))$, where $\mathbf{z}_i \in \mathbb{R}^{n_q}$ denotes the shared embedding.

This shared embedding $\mathbf{z}_i$ serves as the interface between the classical backbone and the downstream hybrid prediction module. In this way, the feature compression module alleviates the dimensional mismatch between high-dimensional visual features and quantum-compatible inputs.

3.4 Hybrid Prediction and Feature-Adaptive Fusion

Given the shared embedding $\mathbf{z}_i$, FAF-HQNN performs parallel hybrid prediction and adaptive fusion. Specifically, $\mathbf{z}_i$ is simultaneously fed into the quantum branch, the classical branch, and the feature-adaptive class-wise fusion mechanism. The quantum and classical branches produce complementary predictive outputs, while the feature-adaptive class-wise fusion mechanism assigns sample-specific fusion weights for each class to combine them adaptively. This design allows the model to dynamically balance the contributions of the two branches across different input samples and classes.

3.4.1 Quantum Branch

The quantum branch processes the extracted features through three stages: classical-to-quantum data encoding, variational quantum transformation, and measurement-based readout. Its overall structure is illustrated in Fig. 2.

Classical-to-quantum data encoding

Let $\mathbf{z}_i = \left[z_i^1, z_i^2, \ldots, z_i^{n_q}\right]^\top \in \mathbb{R}^{n_q}$ denote the shared embedding of the i-th sample used for classical-to-quantum data encoding, where z_i^k denotes the k-th element of

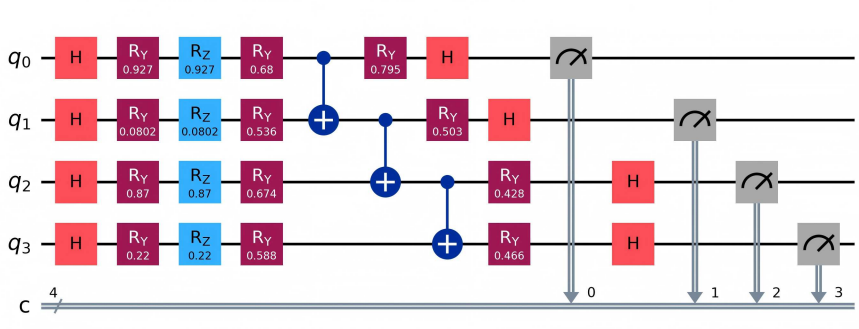


Fig. 2: Overall architecture of the quantum branch in FAF-HQNN. The branch consists of three stages: classical-to-quantum data encoding, variational quantum transformation, and measurement-based readout.

$\mathbf{z}_i$ for $k = 1, 2, \ldots, n_q$. The quantum register is first initialized in the computational basis state $|\psi_0\rangle = |0\rangle^{\otimes n_q}$. To enrich the accessible state space while maintaining a shallow circuit structure, a layer of Hadamard gates is applied to all qubits, resulting in a uniform superposition state

$$|\psi_H\rangle = H^{\otimes n_q}|\psi_0\rangle = \bigotimes_{k=0}^{n_q-1} \frac{|0\rangle_k + |1\rangle_k}{\sqrt{2}}. \tag{1}$$

The uniform superposition state $|\psi_H\rangle$ is then encoded through single-qubit R_y rotations,

$$|\psi_{R_y}\rangle = \left(\bigotimes_{k=0}^{n_q-1} R_y(z_i^k)\right)|\psi_H\rangle, \tag{2}$$

where each component z_i^k controls the rotation angle of qubit k. To further incorporate phase information, R_z rotations are additionally applied to the state $|\psi_{R_y}\rangle$. The encoded quantum state is thus written as

$$|\psi_i\rangle = \left[\left(\bigotimes_{l=0}^{n_q-1} R_z(z_i^l)\right)\right]|\psi_{R_y}\rangle. \tag{3}$$

This encoding scheme combines complementary angle and phase encoding operations to enrich the quantum representation while preserving a hardware-efficient and shallow circuit design.

Variational quantum transformation

After classical-to-quantum encoding, the encoded quantum state $|\psi_i\rangle$ is subsequently processed by an L-layer variational quantum circuit (VQC). In the l-th layer, each qubit j is first acted upon by one trainable single-qubit rotation gate, $R_y(\theta_{l,j})$. These

local operations are followed by a nearest-neighbor entangling chain composed of CNOT gates. After entanglement, another $R_y(\beta_{l,j})$ rotation is applied to each qubit. The overall unitary transformation implemented by the VQC is given by

$$U_{\mathrm{VQC}} = \prod_{l=0}^{L-1} \left[\left(\bigotimes_{j=0}^{n_q-1} R_y(\beta_{l,j}) \right) \left(\prod_{j=0}^{n_q-2} \mathrm{CNOT}(j, j+1) \right) \left(\bigotimes_{j=0}^{n_q-1} R_y(\theta_{l,j}) \right) \right], \quad (4)$$

where operators are applied from right to left.

This variational circuit combines trainable rotations with entangling operations, thereby enabling nonlinear state transformations in Hilbert space and allowing the quantum branch to model correlations among qubits.

Measurement-based readout

After applying U_{VQC}, the transformed quantum state is measured on each qubit to obtain a real-valued readout vector. In our implementation, the measurement outcome of the k-th qubit is defined as the expectation value of the Pauli-X observable,

$$q_i^k = \langle \psi_i | U_{\mathrm{VQC}}^\dagger X_k U_{\mathrm{VQC}} | \psi_i \rangle, \quad k = 0, \ldots, n_q - 1. \quad (5)$$

Collecting the measurement outcomes from all qubits yields the quantum feature vector

$$\mathbf{q}_i = [q_i^0, \ldots, q_i^{n_q-1}] \in \mathbb{R}^{n_q}. \quad (6)$$

This vector is then mapped to the class space through a learnable linear layer,

$$\mathbf{y}_i^q = f_q(\mathbf{q}_i), \quad (7)$$

where $f_q(\cdot)$ denotes the linear mapping of the quantum branch, and $\mathbf{y}_i^q \in \mathbb{R}^C$ represents the class logits produced by the quantum branch.

The quantum branch is constructed to remain compatible with NISQ devices. Each variational layer consists of parameterized single-qubit rotations and nearest-neighbor CNOT gates, resulting in a hardware-efficient circuit with limited depth. By combining local trainable operations with shallow entangling interactions, this design supports expressive quantum feature transformation while remaining practical for near-term quantum implementations.

3.4.2 Classical Branch

In parallel with the quantum branch, the shared embedding $\mathbf{z}_i \in \mathbb{R}^{n_q}$ is fed into a classical branch. This branch is implemented as a lightweight multilayer perceptron (MLP), which serves as an efficient classical predictor on the shared embedding.

Specifically, the classical branch consists of three fully connected layers with ReLU activations and dropout regularization. The output class logits are computed as

$$\mathbf{y}_i^c = f_c(\mathbf{z}_i), \quad (8)$$

where $f_c(\cdot)$ denotes the classical MLP, and $\mathbf{y}_i^c \in \mathbb{R}^C$ represents the class logits produced by the classical branch.

Since this branch operates entirely in the classical domain, it provides a stable predictive reference within the hybrid framework. Together with the quantum branch, it forms a complementary dual-branch prediction structure for subsequent adaptive fusion.

3.4.3 Feature-Adaptive Class-Wise Fusion Mechanism

Based on the shared embedding $\mathbf{z}_i$, the classical and quantum branches produce class logits $\mathbf{y}_i^c$ and $\mathbf{y}_i^q$, respectively. To adaptively combine these two outputs, a lightweight class-wise fusion predictor is applied to $\mathbf{z}_i$ to generate a fusion coefficient vector

$$\boldsymbol{\alpha}_i = f_\alpha(\mathbf{z}_i), \qquad \boldsymbol{\alpha}_i \in [0,1]^C, \tag{9}$$

where $f_\alpha(\cdot)$ is implemented as a lightweight two-layer MLP with a sigmoid output layer, ensuring that each coefficient lies in $[0,1]$.

For each class c, the coefficient α_i^c controls the contribution of the classical branch, while $1-\alpha_i^c$ controls the contribution of the quantum branch. The fused logits are defined as

$$\hat{\mathbf{y}}_i = \boldsymbol{\alpha}_i \odot \mathbf{y}_i^c + (\mathbf{1} - \boldsymbol{\alpha}_i) \odot \mathbf{y}_i^q, \tag{10}$$

where $\odot$ denotes element-wise multiplication and $\hat{\mathbf{y}}_i \in \mathbb{R}^C$ is the fused output vector. This formulation enables sample-dependent and class-specific fusion of the classical and quantum predictions while introducing only a small additional parameter overhead through the lightweight fusion coefficient predictor.

3.5 Training Objective

FAF-HQNN is trained end-to-end in a supervised manner for multi-class biomedical image classification. Given an input image $\mathbf{x}_i$ with ground-truth label $y_i \in \{1,\dots,C\}$, the model produces fused class logits $\hat{\mathbf{y}}_i = [\hat{y}_{i,1},\dots,\hat{y}_{i,C}] \in \mathbb{R}^C$ through the feature-adaptive class-wise fusion mechanism.

The model is optimized using the standard cross-entropy loss computed from the fused logits:

$$\mathcal{L}_{\mathrm{CE}} = -\frac{1}{N}\sum_{i=1}^{N} \log \frac{\exp(\hat{y}_{i,y_i})}{\sum_{c=1}^{C}\exp(\hat{y}_{i,c})}, \tag{11}$$

where N denotes the mini-batch size, and $\hat{y}_{i,c}$ denotes the c-th element of the fused logit vector $\hat{\mathbf{y}}_i$.

4 Experiments

This section presents the empirical evaluation of the proposed FAF-HQNN. The experiments are designed to assess four aspects of the proposed method: classification performance under standard conditions, training behavior and the effect of quantum module design, robustness under common image corruptions, and generalization across datasets. We first conduct the main analysis on PathMNIST as the

primary benchmark, and then examine whether the observed trends generalize to BloodMNIST.

4.1 Experimental Setup

This subsection describes the datasets and preprocessing procedures, compared methods, implementation details, and evaluation metrics and robustness protocol used throughout the experiments.

4.1.1 Dataset and Preprocessing

PathMNIST is used as the primary benchmark in this study. It is a subset of the MedMNIST collection and consists of 28×28 RGB histopathological images from nine tissue classes. Additional cross-dataset validation on BloodMNIST is presented in a separate experimental section.

For all experiments on PathMNIST, input images are resized to 64×64, converted to tensors, and normalized using the ImageNet mean and standard deviation. Because sample-wise quantum circuit simulation incurs substantial computational cost, we construct a reduced training subset by sampling 9000 images from the original training split. This reduced-data setting is adopted primarily for computational feasibility under quantum simulation; therefore, the reported results should be interpreted as controlled comparisons under matched data budgets rather than as performance estimates under full-data training. The original test split is retained unchanged for evaluation. To ensure fairness, all compared methods are trained and evaluated on exactly the same sampled training subset under each random seed.

4.1.2 Compared Methods

To evaluate the contribution of the quantum branch, the effect of hybrid prediction, and the benefit of the proposed feature-adaptive class-wise fusion mechanism, we compare FAF-HQNN with several model variants built on the same ResNet-18 backbone and feature compression pipeline. Unless otherwise specified, all compared methods share the same backbone, feature compression module, embedding dimension, training protocol, and data splits, so that the comparison isolates the effect of branch design and fusion strategy.

- **Classical:** a purely classical baseline in which the shared embedding vectors are fed only to the classical branch for final prediction.
- **HQNN:** a quantum-only baseline in which the shared embedding vectors are fed exclusively to the quantum branch for prediction.
- **H-Fix:** a fusion-based hybrid model in which the shared embedding vectors are processed by both the classical and quantum branches, and the resulting logits are fused by fixed equal-weight averaging.
- **H-Scalar:** a fusion-based hybrid model in which the shared embedding vectors are processed by both branches, and the final prediction is obtained through scalar-weight fusion using a single globally learnable coefficient shared across all samples and classes.

- **FAF-HQNN:** a fusion-based hybrid model, i.e., the proposed model, in which the shared embedding vectors are jointly processed by the classical and quantum branches and then adaptively fused through the proposed feature-adaptive class-wise fusion mechanism.

4.1.3 Implementation Details

All experiments are implemented in PyTorch, and the quantum components are realized using PennyLane with the `default.qubit` simulator. The quantum branch employs a 4-qubit variational quantum circuit. Unless otherwise specified, all models are optimized using Adam with a weight decay of 10^{-4}. The learning rate is set to 10^{-4} for the ResNet-18 backbone and 10^{-3} for the newly introduced modules, including the feature compression layers, embedding layer, classical branch, variational quantum circuit, and fusion coefficient predictor.

Each experiment is repeated across five random seeds, and the final performance is reported as the mean ± standard deviation. For experiments conducted under the reduced-data setting, all compared methods use the same sampled training subset under each random seed to ensure fair comparison.

4.1.4 Evaluation Metrics and Robustness Protocol

Classification performance is evaluated using accuracy (ACC), weighted precision, weighted recall, weighted F1-score, and one-vs-rest area under the ROC Curve (AUC). ACC measures the proportion of correctly classified test samples. Weighted Precision, Recall, and F1-score are reported to account for potential class-frequency differences, while one-vs-rest AUC measures the model's ability to distinguish each class from the remaining classes in the multi-class setting.

In addition to evaluation under clean conditions, robustness is assessed under three representative image corruption types: Gaussian noise, salt-and-pepper noise, and Poisson noise. In all robustness experiments, corruptions are applied only to the test images, whereas all models are trained exclusively on clean data. This train-clean/test-corrupted protocol enables a controlled assessment of robustness under progressively degraded input quality while isolating the effect of test-time perturbations from changes in the training distribution.

4.2 Main Results on PathMNIST

We first report the classification performance of all compared methods on PathMNIST under clean conditions and then analyze their training behavior. Table 1 summarizes the classification results of all compared methods on PathMNIST under clean conditions. Overall, FAF-HQNN attains the highest mean performance among the compared methods across all reported metrics under clean conditions, with an ACC of 0.8910 ± 0.0149, a weighted F1-score of 0.8880 ± 0.0151, and an AUC of 0.9816 ± 0.0030. Relative to the purely classical baseline, FAF-HQNN improves ACC by 2.63 percentage points, suggesting that feature-adaptive fusion of classical and quantum predictions is beneficial in this setting.

Table 1: Performance comparison of different model variants on PathMNIST under clean conditions. Values are reported as mean ± standard deviation over five runs.

Method	ACC	Precision	Recall	F1-score	AUC
Classical	0.8647 ±0.0193	0.8716 ±0.0213	0.8647 ±0.0193	0.8619 ±0.0222	0.9770 ±0.0076
HQNN	0.7806 ±0.0713	0.7499 ±0.0996	0.7806 ±0.0713	0.7483 ±0.0932	0.9239 ±0.0132
H-Fix	0.8729 ±0.0191	0.8820 ±0.0142	0.8729 ±0.0191	0.8727 ±0.0188	0.9804 ±0.0053
H-Scalar	0.8792 ±0.0128	0.8830 ±0.0091	0.8792 ±0.0128	0.8769 ±0.0114	0.9786 ±0.0041
FAF-HQNN	**0.8910 ±0.0149**	**0.8900 ±0.0131**	**0.8910 ±0.0149**	**0.8880 ±0.0151**	**0.9816 ±0.0030**

Among the baseline variants, HQNN yields the lowest mean performance and the largest standard deviations across multiple seeds, indicating that the quantum-only branch is less stable and less effective than the classical or hybrid alternatives in the present setting. By contrast, both hybrid baselines, H-Fix and H-Scalar, clearly outperform HQNN and remain competitive with, or slightly stronger than, the Classical baseline, suggesting that combining classical and quantum predictors is more effective than relying solely on the quantum branch.

A further comparison between the two hybrid baselines shows that H-Scalar achieves slightly better overall classification performance than H-Fix, suggesting that introducing a learnable fusion weight is beneficial. However, a single globally shared coefficient remains limited in its ability to capture sample- and class-dependent branch contributions. FAF-HQNN further improves upon both hybrid baselines, with ACC gains of 1.81 percentage points over H-Fix and 1.18 percentage points over H-Scalar, providing empirical support for the effectiveness of the proposed feature-adaptive class-wise fusion mechanism on PathMNIST.

4.3 Analysis of the Quantum Module

Since the classical backbone is fixed throughout these experiments, the observed performance differences can be mainly attributed to changes in the quantum module. To better understand the role of the quantum component in FAF-HQNN, we further analyze how the variational quantum circuit (VQC) design affects the overall performance of the model. In particular, we examine three factors: circuit layout, measurement basis, and circuit depth.

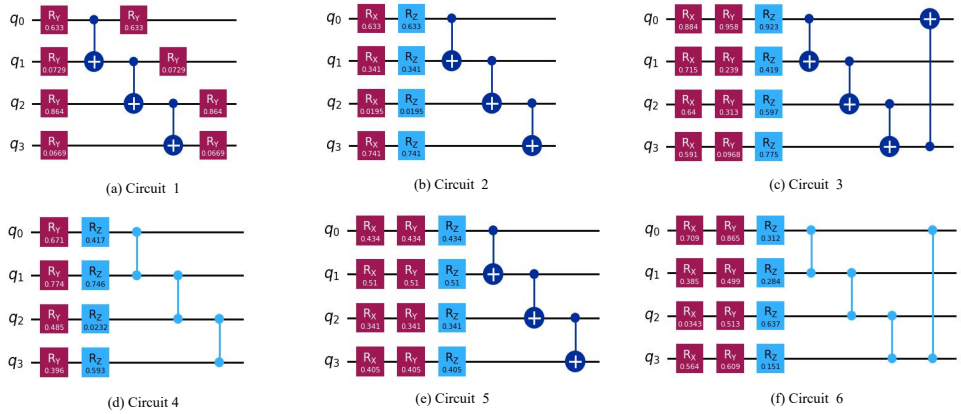


Fig. 3: Six variational quantum circuit layouts considered in this work.

4.3.1 Effect of Circuit Layout and Measurement Basis

We compare six VQC layouts that differ in their combinations of single-qubit rotation gates and entangling operations, as illustrated in Fig. 3. For each layout, we first consider three single-basis measurement settings, namely Pauli-X, Pauli-Y, and Pauli-Z. In each case, the quantum features are given by the expectation values of the corresponding Pauli observable for all qubits. We then consider a combined Pauli-XYZ setting, in which the expectation values of all three Pauli observables are concatenated for every qubit, yielding a quantum feature vector of dimension $3n_q$, where n_q denotes the number of qubits. Unless otherwise specified, all other model components and training settings are kept fixed so that the comparison isolates the effect of VQC design choices. The resulting classification accuracies on PathMNIST are summarized in Table 2.

Table 2 shows that the performance of FAF-HQNN depends on both the circuit layout and the measurement setting. Although the overall variation is moderate, the observed differences indicate that the design of the quantum branch can affect the final hybrid performance. No single measurement setting consistently performs best across all circuit layouts, suggesting that the usefulness of the quantum readout depends on its interaction with the underlying circuit structure. Although the Pauli-XYZ setting provides a richer representation than any single-basis measurement setting, it does not consistently yield the best performance, implying that the additional observables may introduce redundancy or increase the optimization burden of the downstream fusion module.

Among all configurations, the best result is obtained with Circuit 1 under Pauli-X measurement, achieving an ACC of 0.8910 ± 0.0149. Other competitive configurations include Circuit 2 with Pauli-Y and Circuit 6 with Pauli-Z, although neither surpasses the best setting. Overall, these results suggest that, within the present hybrid framework and simulation setting, relatively simple circuit designs are sufficient to provide useful complementary quantum features.

Table 2: Classification accuracy of FAF-HQNN on PathMNIST under different VQC layouts and measurement settings. Values are reported as mean ± standard deviation over five runs.

Measurement	Circuit 1	Circuit 2	Circuit 3	Circuit 4	Circuit 5	Circuit 6
Pauli-X	**0.8910 ±0.0149**	0.8826 ±0.0125	0.8796 ±0.0046	0.8739 ±0.0151	0.8754 ±0.0163	0.8808 ±0.0154
Pauli-Y	0.8769 ±0.0165	0.8843 ±0.0061	0.8794 ±0.0123	0.8705 ±0.0085	0.8770 ±0.0091	0.8810 ±0.0106
Pauli-Z	0.8813 ±0.0218	0.8748 ±0.0100	0.8763 ±0.0099	0.8785 ±0.0145	0.8780 ±0.0126	0.8870 ±0.0108
Pauli-XYZ	0.8717 ±0.0135	0.8674 ±0.0117	0.8752 ±0.0070	0.8835 ±0.0148	0.8704 ±0.0148	0.8812 ±0.0101

Table 3: Performance of FAF-HQNN on PathMNIST under different circuit depths. Values are reported as mean ± standard deviation over five runs.

Layers	ACC	Precision	Recall	F1-score	AUC
1	**0.8910±0.0149**	**0.8900±0.0131**	**0.8910±0.0149**	**0.8880±0.0151**	**0.9816±0.0030**
2	0.8772±0.0187	0.8834±0.0150	0.8772±0.0187	0.8740±0.0175	0.9774±0.0069
3	0.8692±0.0026	0.8778±0.0066	0.8692±0.0026	0.8691±0.0028	0.9742±0.0063

4.3.2 Effect of Circuit Depth

We further vary the number of VQC layers to examine the effect of circuit depth. The corresponding results on PathMNIST are summarized in Table 3. As shown in the table, increasing the circuit depth does not lead to improved performance in the present setting. Instead, the single-layer circuit achieves the best mean results across all reported metrics, including an ACC of 0.8910 ± 0.0149 and an AUC of 0.9816 ± 0.0030. When the depth is increased to two and three layers, all metrics show lower mean values. These findings suggest that, for the current task and simulation setting, a shallow quantum circuit is sufficient to provide useful complementary information within the hybrid framework. A possible explanation is that deeper circuits enlarge the parameter space and make optimization more challenging, while the task itself may not require more expressive quantum transformations to further improve the fused representation.

4.4 Robustness under Common Image Corruptions on PathMNIST

To assess robustness under common image corruptions, we evaluate all methods on three representative corruption types, namely Gaussian noise, salt-and-pepper noise,

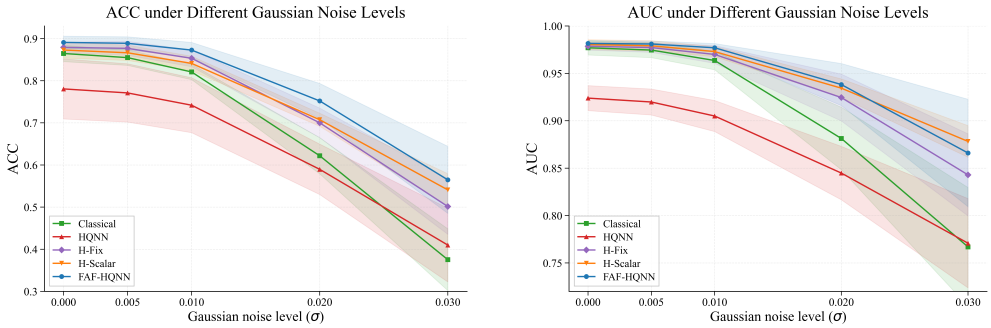
Table 4: Classification performance of different model variants on PathMNIST under Gaussian noise with varying corruption intensity. Results are reported as mean ± standard deviation over five runs.

σ	Metric	Classical	HQNN	H-Fix	H-Scalar	FAF-HQNN
0.00	ACC	0.8647±0.0193	0.7806±0.0713	0.8729±0.0191	0.8792±0.0128	**0.8910±0.0149**
	AUC	0.9770±0.0076	0.9239±0.0132	0.9804±0.0053	0.9786±0.0041	**0.9816±0.0030**
0.005	ACC	0.8549±0.0184	0.7711±0.0692	0.8664±0.0182	0.8769±0.0167	**0.8889±0.0154**
	AUC	0.9746±0.0080	0.9198±0.0137	0.9791±0.0054	0.9772±0.0050	**0.9810±0.0032**
0.01	ACC	0.8209±0.0183	0.7419±0.0652	0.8533±0.0199	0.8412±0.0171	**0.8728±0.0182**
	AUC	0.9636±0.0098	0.9050±0.0164	0.9700±0.0094	0.9730±0.0054	**0.9770±0.0045**
0.02	ACC	0.6219±0.0429	0.5897±0.0599	0.7074±0.0155	0.7001±0.0350	**0.7521±0.0415**
	AUC	0.8813±0.0330	0.8448±0.0282	0.9344±0.0098	0.9245±0.0247	**0.9382±0.0222**
0.03	ACC	0.3752±0.0733	0.4099±0.0875	0.5409±0.0385	0.5014±0.0660	**0.5647±0.0796**
	AUC	0.7669±0.0627	0.7708±0.0471	**0.8782±0.0167**	0.8429±0.0433	0.8660±0.0567

and Poisson noise. Following the protocol defined above, all models are trained on clean data and evaluated exclusively on corrupted test images.

4.4.1 Gaussian Noise

Table 4 and Fig. 4 summarize the classification results under Gaussian noise at five corruption levels, i.e., $\sigma \in \{0.00, 0.005, 0.01, 0.02, 0.03\}$. As expected, the performance of all methods gradually declines with increasing noise intensity in terms of both ACC and AUC. Nevertheless, FAF-HQNN achieves the best mean performance across most noise levels. In particular, it consistently attains the highest ACC and AUC from clean conditions up to moderate noise ($\sigma = 0.02$); for example, at $\sigma = 0.02$, it achieves an ACC of 0.7521 ± 0.0415 and an AUC of 0.9382 ± 0.0222. Under the strongest corruption level ($\sigma = 0.03$), FAF-HQNN still yields the highest ACC, while its AUC remains competitive and is only slightly lower than the best mean AUC achieved by H-Fix. By comparison, the Classical model degrades markedly as the noise level increases, and HQNN consistently underperforms the hybrid variants, especially under severe corruption. Overall, fusion-based hybrid models outperform the two single-branch baselines across corrupted settings, indicating that fusing quantum and classical information improves robustness against additive Gaussian perturbations in the present setting. Among them, FAF-HQNN provides the most favorable overall trade-off between clean-data performance and robustness to Gaussian noise.



(a) ACC (b) AUC

Fig. 4: ACC and AUC of different model variants on PathMNIST under Gaussian noise with increasing corruption intensity.

4.4.2 Salt-and-Pepper Noise

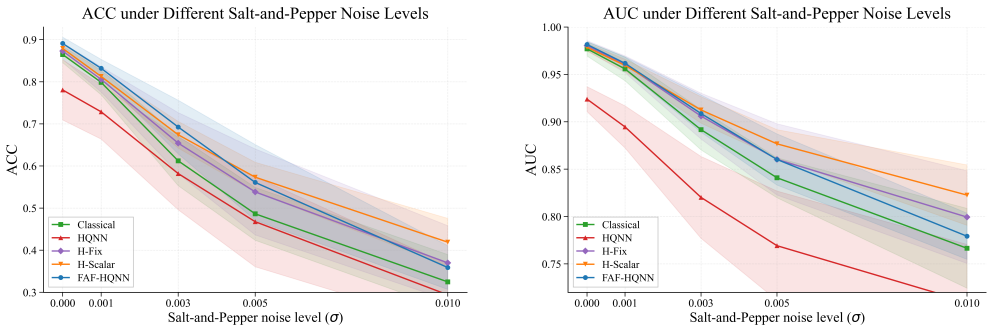
We further evaluate robustness under salt-and-pepper noise at five corruption levels, i.e., $\sigma \in \{0.000, 0.001, 0.003, 0.005, 0.010\}$, with the corresponding results summarized in Table 5 and Fig. 5. Compared with Gaussian noise, salt-and-pepper corruption leads to a much sharper performance degradation for all methods, indicating that impulse noise is more disruptive to PathMNIST classification in the present setting. FAF-HQNN achieves the best mean performance under clean condition and low corruption levels, attaining the highest ACC and AUC under clean conditions and at $\sigma = 0.001$, and still yielding the highest ACC at $\sigma = 0.003$. As the corruption severity increases further, however, H-Scalar becomes the best-performing variant, achieving the highest ACC and AUC at both $\sigma = 0.005$ and $\sigma = 0.010$. Overall, the fusion-based hybrid models consistently outperform the two single-branch baselines across corrupted settings, suggesting that the fusion of quantum and classical information improves robustness to salt-and-pepper corruption, while the most effective fusion strategy appears to depend on the noise severity.

4.4.3 Poisson Noise

We further evaluate robustness under Poisson noise using one clean setting and four noisy settings with peak values of 3200, 1600, 800, and 400, where smaller peak values indicate stronger corruption. The corresponding results are summarized in Table 6 and Fig. 6. As expected, all methods exhibit progressive performance degradation as the corruption becomes stronger. Nevertheless, FAF-HQNN consistently achieves the highest mean ACC across all evaluated settings, with its advantage being particularly evident under moderate and strong corruption. For example, it attains an ACC of 0.7396 ± 0.0430 at a peak value of 1600 and remains the best-performing method in terms of ACC at a peak value of 400. In terms of AUC, FAF-HQNN performs best under the clean and relatively mild corruption settings, whereas H-Fix achieves the highest mean AUC under stronger corruption (at peak values of 800 and 400).

Table 5: Classification performance of different model variants on PathMNIST under salt-and-pepper noise with varying corruption intensity. Results are reported as mean ± standard deviation over five runs.

σ	Metric	Classical	HQNN	H-Fix	H-Scalar	FAF-HQNN
0.000	ACC	0.8647±0.0193	0.7806±0.0713	0.8729±0.0191	0.8792±0.0128	**0.8910±0.0149**
	AUC	0.9770±0.0076	0.9239±0.0132	0.9804±0.0053	0.9786±0.0041	**0.9816±0.0030**
0.001	ACC	0.7986±0.0297	0.7286±0.0647	0.8056±0.0300	0.8126±0.0125	**0.8320±0.0212**
	AUC	0.9558±0.0125	0.8946±0.0222	0.9607±0.0089	0.9595±0.0065	**0.9617±0.0069**
0.003	ACC	0.6122±0.0590	0.5820±0.0860	0.6543±0.0725	0.6742±0.0317	**0.6925±0.0637**
	AUC	0.8917±0.0214	0.8202±0.0432	0.9060±0.0242	**0.9125±0.0140**	0.9087±0.0197
0.005	ACC	0.4864±0.0628	0.4677±0.1070	0.5384±0.1021	**0.5734±0.0355**	0.5610±0.0899
	AUC	0.8409±0.0211	0.7692±0.0577	0.8603±0.0376	**0.8767±0.0150**	0.8602±0.0270
0.010	ACC	0.3248±0.0649	0.2940±0.0728	0.3700±0.0883	**0.4191±0.0566**	0.3589±0.0519
	AUC	0.7665±0.0421	0.7098±0.0608	0.7994±0.0489	**0.8226±0.0319**	0.7792±0.0244



(a) ACC (b) AUC

Fig. 5: ACC and AUC of different model variants on PathMNIST under salt-and-pepper noise with increasing corruption intensity.

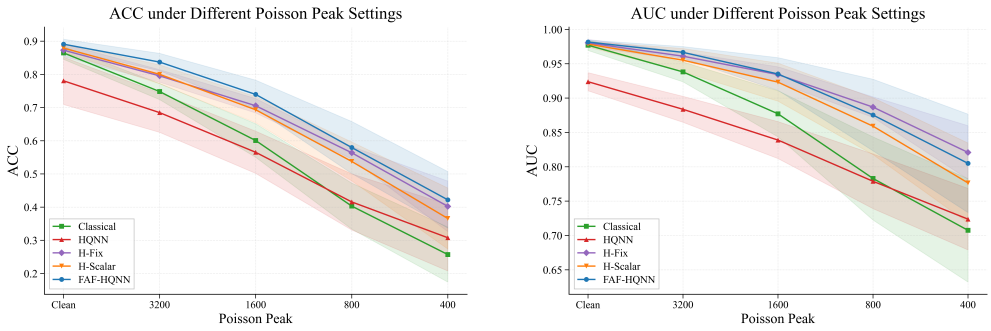
Overall, the fusion-based hybrid models generally show stronger robustness than the two single-branch baselines under Poisson noise, while FAF-HQNN offers the most favorable overall trade-off between classification accuracy and robustness.

4.4.4 Overall Robustness Trends

Across all three corruption types, the performance of all methods degrades as the corruption severity increases. Nevertheless, the fusion-based hybrid models generally

Table 6: Classification performance of different model variants on PathMNIST under Poisson noise with different peak settings. Results are reported as mean ± standard deviation over five runs.

Peak	Metric	Classical	HQNN	H-Fix	H-Scalar	FAF-HQNN
Clean	ACC	0.8647±0.0193	0.7806±0.0713	0.8729±0.0191	0.8792±0.0128	**0.8910±0.0149**
	AUC	0.9770±0.0076	0.9239±0.0132	0.9804±0.0053	0.9786±0.0041	**0.9816±0.0030**
3200	ACC	0.7483±0.0243	0.6848±0.0594	0.7955±0.0188	0.8002±0.0285	**0.8371±0.0263**
	AUC	0.9381±0.0144	0.8837±0.0189	0.9610±0.0060	0.9552±0.0170	**0.9666±0.0084**
1600	ACC	0.6007±0.0506	0.5653±0.0636	0.7055±0.0217	0.6934±0.0393	**0.7396±0.0430**
	AUC	0.8771±0.0336	0.8389±0.0269	0.9342±0.0114	0.9232±0.0277	**0.9351±0.0238**
800	ACC	0.4031±0.0694	0.4157±0.0846	0.5638±0.0187	0.5377±0.0586	**0.5796±0.0784**
	AUC	0.7830±0.0605	0.7789±0.0398	**0.8869±0.0150**	0.8591±0.0414	0.8754±0.0521
400	ACC	0.2571±0.0829	0.3078±0.1001	0.4023±0.0762	0.3658±0.0919	**0.4220±0.0852**
	AUC	0.7075±0.0753	0.7238±0.0448	**0.8209±0.0390**	0.7765±0.0581	0.8050±0.0717



(a) ACC (b) AUC

Fig. 6: ACC and AUC of different model variants on PathMNIST under Poisson noise with increasing corruption severity.

show stronger robustness than the single-branch baselines in most settings, indicating that combining quantum and classical information is beneficial under common image corruptions in the present setting. Among the hybrid variants, FAF-HQNN achieves the strongest overall performance under Gaussian, Poisson, salt-and-pepper corruptions, whereas H-Scalar is more robust under severe salt-and-pepper corruption. These

results suggest that adaptive fusion is effective for robust biomedical image classification, while the most suitable fusion strategy may vary with the corruption type and severity.

4.5 Cross-Dataset Validation on BloodMNIST

To assess whether the observations on PathMNIST generalize to another dataset, we further evaluate all compared methods on BloodMNIST, a MedMNIST benchmark for multi-class peripheral blood cell image classification. BloodMNIST contains 17,092 RGB images of size 28×28 and defines an 8-class single-label classification task, with 11,959 training samples, 1,712 validation samples, and 3,421 test samples in the official split.

To maintain consistency with the primary experiments, we apply the same preprocessing pipeline: input images are resized to 64×64, converted to tensors, and normalized using the ImageNet mean and standard deviation. In addition, for each random seed, all compared methods are trained and evaluated under the same protocol, so that the BloodMNIST experiments provide a controlled cross-dataset validation of the proposed hybrid framework.

4.5.1 Performance under Clean Conditions

Table 7 summarizes the classification results on BloodMNIST under clean conditions. Among all compared methods, FAF-HQNN achieves the best overall performance, reaching an ACC of 0.9503 ± 0.0052 and an AUC of 0.9943 ± 0.0008. It consistently outperforms both the classical baseline and HQNN in terms of mean performance, indicating that the proposed feature-adaptive fusion strategy generalizes effectively to blood cell image classification.

Similar to the observations on PathMNIST, HQNN shows the weakest overall performance and substantially larger standard deviations than the other methods, suggesting lower stability of the quantum-only branch in the present setting. In contrast, all fusion-based hybrid models outperform the two single-branch baselines, further supporting the benefit of fusing quantum and classical information.

4.5.2 Robustness under Image Corruptions

We further evaluate all compared methods on BloodMNIST under Gaussian, salt-and-pepper, and Poisson corruptions. The corresponding results are summarized in Tables 8–10 and Figs. 7–9. Overall, the trends on BloodMNIST are broadly consistent with those observed on PathMNIST: performance degrades as the corruption severity increases, while all fusion-based hybrid models remain markedly more robust than the quantum-only baseline and generally outperform the purely classical model under corrupted conditions.

Under Gaussian noise, FAF-HQNN shows the strongest overall robustness, achieving the best mean ACC and AUC from the clean setting to moderate corruption levels. Under the most severe Gaussian setting, H-Scalar attains the highest mean ACC, whereas FAF-HQNN remains highly competitive and achieves the best mean

Table 7: Performance comparison of different model variants on BloodMNIST under clean conditions. Values are reported as mean and standard deviation.

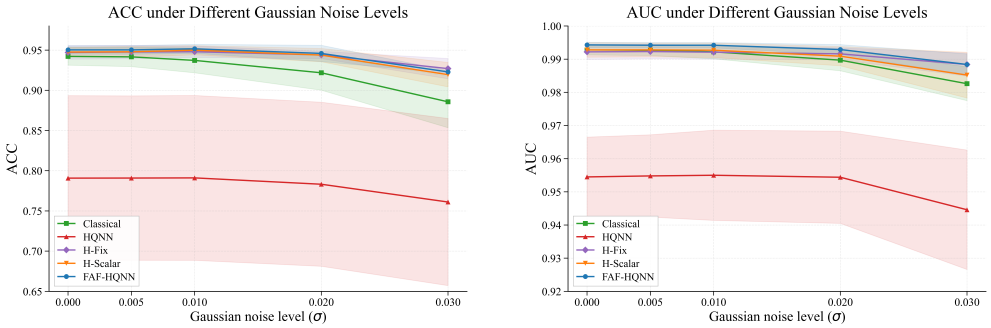
Method	**ACC**	**Precision**	**Recall**	**F1-score**	**AUC**
Classical	0.9422 ±0.0109	0.9440 ±0.0101	0.9422 ±0.0109	0.9424 ±0.0104	0.9928 ±0.0015
HQNN	0.7908 ±0.1026	0.7228 ±0.1451	0.7908 ±0.1026	0.7150 ±0.1418	0.9545 ±0.0120
H-Fix	0.9471 ±0.0084	0.9481 ±0.0079	0.9471 ±0.0084	0.9470 ±0.0083	0.9922 ±0.0023
H-Scalar	0.9476 ±0.0061	0.9484 ±0.0060	0.9476 ±0.0061	0.9475 ±0.0061	0.9928 ±0.0022
FAF-HQNN	**0.9503 ±0.0052**	**0.9508 ±0.0050**	**0.9503 ±0.0052**	**0.9501 ±0.0051**	**0.9943 ±0.0008**

Table 8: Classification performance of different model variants on BloodMNIST under Gaussian noise with varying corruption intensity. Results are reported as mean ± standard deviation over five runs.

σ	**Metric**	**Classical**	**HQNN**	**H-Fix**	**H-Scalar**	**FAF-HQNN**
0.000	ACC	0.9422±0.0109	0.7908±0.1026	0.9471±0.0084	0.9476±0.0061	**0.9503±0.0052**
	AUC	0.9928±0.0015	0.9545±0.0120	0.9922±0.0023	0.9928±0.0022	**0.9943±0.0008**
0.005	ACC	0.9416±0.0121	0.7909±0.1022	0.9476±0.0083	0.9479±0.0068	**0.9503±0.0054**
	AUC	0.9927±0.0016	0.9548±0.0124	0.9923±0.0022	0.9928±0.0021	**0.9942±0.0008**
0.010	ACC	0.9372±0.0154	0.7911±0.1025	0.9480±0.0071	0.9499±0.0057	**0.9516±0.0056**
	AUC	0.9922±0.0021	0.9550±0.0136	0.9921±0.0021	0.9927±0.0022	**0.9941±0.0008**
0.020	ACC	0.9218±0.0216	0.7832±0.1020	0.9440±0.0045	0.9438±0.0083	**0.9459±0.0101**
	AUC	0.9897±0.0032	0.9544±0.0139	0.9916±0.0018	0.9909±0.0029	**0.9929±0.0015**
0.030	ACC	0.8857±0.0322	0.7611±0.1039	0.9195±0.0152	**0.9270±0.0126**	0.9229±0.0155
	AUC	0.9826±0.0051	0.9446±0.0180	0.9852±0.0069	0.9884±0.0035	**0.9884±0.0032**

AUC. These results again indicate that the fusion-based hybrid models are beneficial under Gaussian corruption, although the most effective fusion strategy may vary under severe noise.

Under salt-and-pepper noise, the performance degradation is substantially more severe for all methods, indicating that impulse corruption is particularly disruptive for BloodMNIST classification in the present setting. Nevertheless, the fusion-based



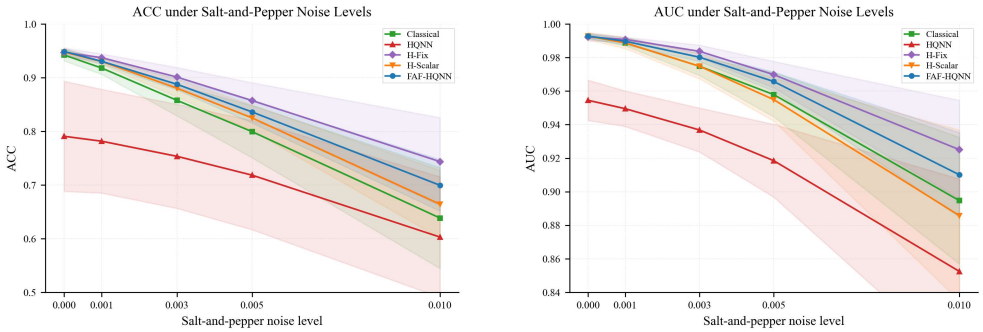
(a) ACC (b) AUC

Fig. 7: ACC and AUC of different model variants on BloodMNIST under Gaussian noise with increasing corruption intensity.

Table 9: Classification performance of different model variants on BloodMNIST under salt-and-pepper noise with varying corruption intensity. Results are reported as mean ± standard deviation over five runs.

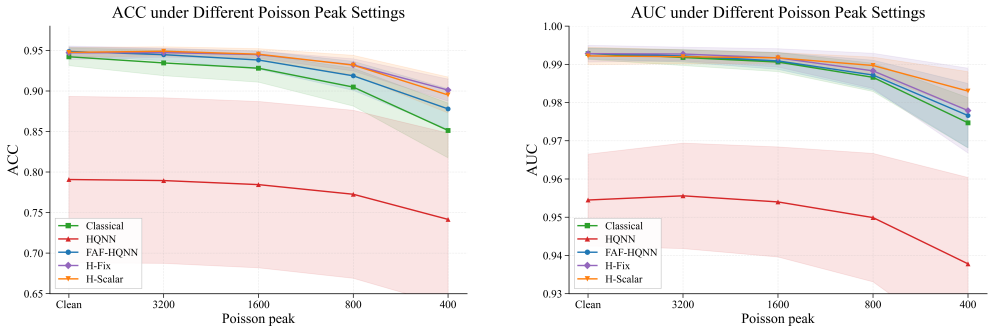
σ	Metric	Classical	HQNN	H-Fix	H-Scalar	FAF-HQNN
0.000	ACC	0.9422±0.0109	0.7908±0.1026	0.9471±0.0084	0.9476±0.0061	**0.9503±0.0052**
	AUC	0.9928±0.0015	0.9545±0.0120	0.9922±0.0023	0.9928±0.0022	**0.9943±0.0008**
0.001	ACC	0.9179±0.0118	0.7816±0.0966	**0.9372±0.0071**	0.9300±0.0056	0.9327±0.0036
	AUC	0.9887±0.0017	0.9495±0.0106	0.9908±0.0013	0.9887±0.0034	**0.9911±0.0018**
0.003	ACC	0.8578±0.0288	0.7535±0.0970	**0.9011±0.0179**	0.8804±0.0164	0.8904±0.0154
	AUC	0.9748±0.0055	0.9368±0.0131	**0.9837±0.0037**	0.9748±0.0075	0.9812±0.0070
0.005	ACC	0.7994±0.0486	0.7184±0.1017	**0.8574±0.0332**	0.8250±0.0242	0.8479±0.0292
	AUC	0.9579±0.0132	0.9185±0.0217	**0.9698±0.0079**	0.9549±0.0128	0.9697±0.0138
0.010	ACC	0.6385±0.0939	0.6030±0.1123	0.6643±0.0631	0.7436±0.0820	**0.7557±0.0628**
	AUC	0.8947±0.0378	0.8525±0.0551	0.8858±0.0510	0.9251±0.0294	**0.9353±0.0345**

hybrid models maintain clear advantages over the two single-branch baselines across the corrupted settings. H-Fix performs best at low and moderate corruption levels, whereas FAF-HQNN becomes the best-performing overall method under the most severe setting, achieving the highest mean ACC and AUC at $\sigma = 0.010$. This result suggests that the preferred fusion strategy under salt-and-pepper noise may depend on the corruption severity.



(a) ACC (b) AUC

Fig. 8: ACC and AUC of the compared methods under salt-and-pepper noise on BloodMNIST.



(a) ACC (b) AUC

Fig. 9: ACC and AUC of different model variants on BloodMNIST under Poisson noise with increasing corruption severity.

Under Poisson noise, all methods exhibit gradual performance degradation as the peak value decreases. FAF-HQNN consistently achieves the best mean ACC and AUC across all evaluated settings, from the clean condition to the strongest corruption level, indicating the most stable overall robustness under this corruption type. Overall, these BloodMNIST results further support the robustness advantage of the fusion-based hybrid models and confirm that the proposed feature-adaptive fusion provides the most favorable overall performance across multiple corruption scenarios in the present setting.

Table 10: Classification performance of different model variants on BloodMNIST under different Poisson peak settings with varying corruption intensity. Results are reported as mean ± standard deviation over five runs.

Peak	Metric	Classical	HQNN	H-Fix	H-Scalar	FAF-HQNN
Clean	ACC	0.9422±0.0109	0.7908±0.1026	0.9471±0.0084	0.9476±0.0061	**0.9503±0.0052**
	AUC	0.9928±0.0015	0.9545±0.0120	0.9922±0.0023	0.9928±0.0022	**0.9943±0.0008**
3200	ACC	0.9346±0.0157	0.7895±0.1021	0.9473±0.0056	0.9490±0.0056	**0.9501±0.0057**
	AUC	0.9918±0.0021	0.9556±0.0138	0.9920±0.0018	0.9927±0.0018	**0.9940±0.0009**
1600	ACC	0.9281±0.0172	0.7845±0.1026	0.9449±0.0043	0.9454±0.0069	**0.9464±0.0068**
	AUC	0.9906±0.0025	0.9540±0.0144	0.9917±0.0014	0.9917±0.0024	**0.9932±0.0012**
800	ACC	0.9048±0.0234	0.7726±0.1036	0.9322±0.0075	0.9319±0.0121	**0.9335±0.0106**
	AUC	0.9866±0.0037	0.9499±0.0168	0.9897±0.0022	0.9883±0.0046	**0.9907±0.0020**
400	ACC	0.8513±0.0338	0.7416±0.1065	0.9011±0.0136	0.8951±0.0222	**0.9021±0.0153**
	AUC	0.9747±0.0066	0.9378±0.0226	0.9830±0.0051	0.9779±0.0111	**0.9844±0.0034**

5 Discussion

5.1 Discussion of Results

The results on PathMNIST and BloodMNIST provide several insights into the role of FAF-HQNN model in biomedical image classification in NISQ-compatible settings. First, the quantum-only baseline consistently underperforms the classical baseline and the fusion-based hybrid models under both clean and corrupted conditions. This indicates that, for the tasks considered here, the quantum branch is not sufficiently effective as a standalone predictor, and that its value lies primarily in complementing the classical branch.

Second, the advantage of the FAF-HQNN model extends beyond under clean conditions. Across Gaussian, salt-and-pepper, and Poisson corruption, it generally achieve stronger mean ACC and AUC than the two single-branch baselines in most settings on both datasets. This trend is most consistent under Gaussian and Poisson noise, whereas under salt-and-pepper noise the relative ranking among the hybrid models becomes more dependent on corruptions severity. These observations suggest that feature-adaptive fusion can improve robustness under input perturbation, although the most effective fusion strategy may vary across perturbation regimes.

Third, the comparison among the fusion-based hybrid models indicates that robustness gains depend not only on the presence of a quantum component but also on the fusion mechanism. The strong overall performance of FAF-HQNN supports the value of feature-adaptive fusion, whereas the competitive behavior of H-Fix and H-Scalar in some severe corruption settings suggests that different fusion rules induce different

robustness profiles. A plausible interpretation is that the quantum branch provides complementary nonlinear transformations and may respond differently to corrupted inputs than the classical branch. In this view, the quantum module is most useful as a complementary representation learner rather than as an independent classifier, which is consistent with the large gap between FAF-HQNN and HQNN.

The analysis of the quantum module further shows that performance depends on circuit layout, measurement basis, and circuit depth, although the variation across configurations remains moderate. The best results are obtained with a shallow single-layer circuit, and increasing depth does not improve performance in the present simulation setting. In addition, the Pauli-XYZ measurement scheme does not consistently outperform single-basis measurement, suggesting that richer quantum readout does not necessarily yield better downstream classification.

Overall, these findings support the FAF-HQNN model as a promising approach for robust biomedical image classification in NISQ-compatible settings. These conclusions should nevertheless be interpreted within the scope of simulation-based experiments on MedMNIST datasets.

5.2 Limitations

While the results on PathMNIST and BloodMNIST provide encouraging empirical evidence for the proposed FAF-HQNN model, several limitations should be acknowledged when interpreting the present findings.

1. **Simulation-based evaluation.** All experiments were conducted in a classical simulation setting rather than on real quantum hardware. As a result, the reported results do not reflect hardware-specific effects such as gate noise, decoherence, restricted qubit connectivity, or measurement error. The practical behavior of the proposed framework on NISQ devices remains to be validated through hardware experiments.
2. **Limited dataset scale and modality diversity.** PathMNIST and BloodMNIST are widely used biomedical image benchmarks, but they remain relatively small-scale, low-resolution, and standardized compared with real clinical imaging scenarios. Accordingly, it remains unclear whether the observed advantages of the FAF-HQNN model extend to larger datasets, higher-resolution inputs, or imaging modalities with greater visual and statistical complexity.
3. **Restricted robustness protocol.** The robustness analysis is limited to a train-clean/test-corrupted setting with three synthetic corruption types: Gaussian noise, salt-and-pepper noise, and Poisson noise. Although this setup enables a controlled evaluation of robustness under test-time perturbations, it does not cover other practically important forms of distribution shift, such as staining variation, acquisition artifacts, or annotation noise.
4. **Configuration dependence and limited mechanistic understanding.** The observed gains depend not only on the inclusion of a quantum branch, but also on the fusion mechanism. Moreover, although the empirical results suggest that the quantum branch can provide complementary representations after fusion, the present study does not yet offer a direct mechanistic account of why these gains

emerge. Further investigation of representation structure, optimization behavior, and corruption sensitivity will be important for clarifying the role of the quantum component and the conditions under which adaptive fusion is most effective.

These limitations suggest that the present results should be interpreted primarily as controlled empirical evidence for the promise of feature-adaptive hybrid fusion, rather than as definitive evidence of superiority across all medical imaging and quantum computing settings.

5.3 Future Work

Several directions may help further establish the practical relevance and scientific understanding of the FAF-HQNN model for biomedical image analysis.

First, an important next step is to move beyond idealized classical simulation toward hardware-aware evaluation. Assessing the proposed framework on real NISQ devices would make it possible to examine how gate noise, decoherence, limited qubit connectivity, calibration drift, and measurement error influence both predictive performance and robustness. Such studies will be essential for determining whether the trends observed here remain valid under realistic quantum hardware constraints.

Second, broader validation on more demanding biomedical imaging datasets is needed. While MedMNIST provides a useful controlled benchmark, future studies should investigate larger-scale, higher-resolution, and more heterogeneous clinical datasets that better reflect real-world visual complexity, anatomical variability, and acquisition diversity. This will be important for assessing how far the present observations generalize beyond standardized benchmark conditions.

Third, robustness evaluation should be extended beyond synthetic corruption toward more realistic sources of distribution shift. In particular, future work should consider staining variation, scanner heterogeneity, acquisition artifacts, or annotation noise. Testing the FAF-HQNN model under such conditions would provide a more practically meaningful characterization of its robustness and potential clinical relevance.

Finally, the design principles and underlying mechanisms of fusion strategies warrant deeper investigation. The present results suggest that the effectiveness of the FAF-HQNN model depends strongly on how quantum and classical information are integrated, and that the preferred fusion strategy may vary across datasets and perturbation regimes. Future research should therefore pursue more principled adaptive fusion methods, while also examining representation structure, branch complementarity, optimization dynamics, and corruption sensitivity to better understand when and why feature-adaptive fusion yields practical advantages.

6 Conclusion

This study investigated fusion-based hybrid models for robust biomedical image classification by evaluating several representative architectures on PathMNIST and BloodMNIST under both clean and corrupted conditions. Across the considered tasks,

hybrid models consistently outperform the quantum-only baseline and often surpass the purely classical baseline, particularly under corruption-induced distribution shifts.

Among the evaluated hybrid variants, FAF-HQNN achieves the strongest overall performance in many key settings, especially under clean conditions and across most Gaussian and Poisson corruption levels. At the same time, the results indicate that no single fusion strategy is uniformly optimal across all datasets, corruption types, and corruption severities. In particular, H-Fix and H-Scalar remain competitive in certain scenarios, including severe salt-and-pepper corruption on BloodMNIST. These findings suggest that the effectiveness of hybrid quantum-classical learning depends not merely on the inclusion of a quantum branch, but more fundamentally on how classical and quantum representations are integrated.

Overall, the present study provides empirical evidence that quantum feature transformations can contribute useful complementary information when combined with expressive classical representations through appropriate fusion mechanisms. From this perspective, the practical value of hybrid quantum-classical learning may lie less in replacing classical deep learning than in enriching representation diversity and improving robustness under challenging conditions. Although these conclusions are currently limited to simulation-based experiments on relatively controlled MedMNIST benchmarks, they support adaptive hybrid fusion as a promising direction for robust biomedical image classification in NISQ-compatible settings.

Acknowledgements. The authors would like to thank Chaoyang Li from Zhengzhou University of Light Industry for valuable discussions and technical support.

Declarations

- Funding
  This work was supported by the Shandong Provincial Natural Science Foundation (Grant No. ZR2025MS1063), the Scientific Research Fund of Zaozhuang University (Grant No. 102061901), the National Natural Science Foundation of China under Grant No. 62501523, and the Zhejiang Province Selected Funding for Postdoctoral Research Projects (No. ZJ2025023).
- Conflict of interest/Competing interests (check journal-specific guidelines for which heading to use)
  Not applicable
- Ethics approval and consent to participate
  Not applicable
- Consent for publication
  Not applicable
- Data availability
  The data that support the findings of this study are available from the corresponding author upon reasonable request.
- Materials availability
  The materials that support the findings of this study are available from the corresponding author upon reasonable request.
- Code availability

The code that supports the findings of this study is available from the corresponding author upon reasonable request.

- Author contribution
  Yan-Yan Hou wrote the manuscript and conducted the experiments, ensuring the methodology was rigorously followed and that the data collected were reliable and valid. Jian Li conceived the experimental design, providing critical insights into the research framework and guiding the overall direction of the study. Chongqiang Ye analyzed the results, employing statistical methods to interpret the data accurately and draw meaningful conclusions. Hengji Li reviewed the manuscript, offering valuable feedback on the clarity and coherence of the writing, as well as ensuring that the findings were presented in a logical manner. Zhuo Wang validated the experiments, confirming the reproducibility of the results and providing additional verification of the experimental protocols. Qinghui Liu supervised the project, contributed to the interpretation of the results, and provided overall guidance on the revision of the manuscript.

Appendix A Additional Analysis under Simulated Quantum Noise

This appendix provides a supplementary analysis of the behavior of FAF-HQNN under simulated quantum hardware noise. Whereas the main text focuses on robustness to classical image corruptions, the present analysis examines the sensitivity of the quantum branch to several representative circuit-level noise channels introduced during quantum circuit execution. These results are intended as complementary evidence rather than as the primary basis for the conclusions of this study.

A.1 Experimental Setting

All experiments in this appendix are conducted on PathMNIST and BloodMNIST, utilizing the same preprocessing pipeline, model configuration, and in the main text, unless otherwise specified. The quantum module is implemented in PennyLane, and noisy quantum circuit simulations are performed by injecting quantum noise channels into the variational quantum circuit during execution.

We consider four representative quantum noise types of quantum noise that commonly encountered in noisy intermediate-scale quantum (NISQ) devices.

- **Bit-flip noise**, flips the qubit state with a probability of p;
- **Phase-flip noise**, alters the relative phase of the qubit state with a probability of p;
- **Amplitude damping**, models energy dissipation from $|1\rangle$ to $|0\rangle$ with a damping parameter γ;
- **Depolarizing noise**, replaces the original qubit state with a maximally mixed state with a probability of p.

To reflect the fact that two-qubit operations are generally more error-prone than single-qubit gates, we adopt a gate-dependent noise setting. Specifically, the noise coefficient is set to $p_1 = 0.001$ for single-qubit gates and $p_2 = 0.01$ for two-qubit gates.

For amplitude damping, the same values are used for the corresponding damping parameters of the two gate categories.

All reported results are averaged over five random seeds and presented as mean ± standard deviation.

A.2 Quantum Noise Models

For completeness, we briefly summarize the quantum noise channels considered in this analysis.

Bit-flip Noise.

Bit-flip noise serves as a quantum analogue to classical binary inversion. For a density matrix ρ, the bit-flip channel is defined as

$$\rho' = (1-p)\rho + pX\rho X, \tag{A1}$$

where p represents the bit-flip probability and X is the Pauli-X operator.

Phase-flip Noise.

Phase-flip noise alters the relative phase between the computational basis states without affecting their amplitudes. Its action on the density matrix ρ is given by

$$\rho' = (1-p)\rho + pZ\rho Z, \tag{A2}$$

where Z denotes the Pauli-Z operator.

Amplitude Damping.

Amplitude damping models energy relaxation processes, such as spontaneous emission, where the excited state $|1\rangle$ decays toward the ground state $|0\rangle$. It is described by the Kraus operators

$$E_0 = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{bmatrix}, \quad E_1 = \begin{bmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{bmatrix}, \tag{A3}$$

and the corresponding state evolution is described by

$$\rho' = E_0\rho E_0^\dagger + E_1\rho E_1^\dagger. \tag{A4}$$

Depolarizing Noise.

Depolarizing noise simulates the uniform loss of quantum information due to interactions with the environment. In the single-qubit case, it can be expressed as

$$\rho' = (1-p)\rho + \frac{p}{2}I, \tag{A5}$$

where I is the identity operator.

Table A1: Performance of FAF-HQNN on PathMNIST under simulated quantum noise (mean ± std). The noise coefficients are set to $p_1 = 0.001$ for single-qubit gates and $p_2 = 0.01$ for two-qubit gates.

Setting	ACC	Precision	Recall	F1-score	AUC
Noiseless	0.8910±0.0149	0.8900±0.0131	0.8910±0.0149	0.8880±0.0151	0.9816±0.0030
Bit-flip	0.8688±0.0118	0.8746±0.0094	0.8688±0.0118	0.8668±0.0125	0.9760±0.0068
Phase-flip	0.8725±0.0246	0.8815±0.0214	0.8725±0.0246	0.8723±0.0237	0.9747±0.0066
Amplitude damping	0.8865±0.0076	0.8870±0.0057	0.8865±0.0076	0.8837±0.0076	0.9795±0.0034
Depolarizing noise	0.8674±0.0082	0.8710±0.0084	0.8674±0.0082	0.8657±0.0082	0.9769±0.0021

A.3 Results under Simulated Quantum Noise

Table A1 and Table A2 report the performance of FAF-HQNN on PathMNIST and BloodMNIST under different simulated quantum noise channels.

On PathMNIST, all considered noise types lead to some degree of performance degradation relative to the noiseless setting, although the magnitude of the decline remains limited under the adopted noise levels. Among the evaluated channels, amplitude damping produces the smallest reduction across most metrics, whereas bit-flip, phase-flip, and depolarizing noise lead to somewhat larger but still moderate decreases in performance. These results suggest that the hybrid model retains a reasonable level of stability under mild circuit-level perturbations, while also indicating that the impact of noise depends on the specific channel type.

On BloodMNIST, the effect of simulated quantum noise is even smaller. Across all four noise channels, the differences relative to the noiseless baseline remain marginal, and some metrics vary only within a narrow range across random seeds. This pattern suggests that, for the present architecture and task setting, the model is comparatively insensitive to the simulated noise levels considered here.

Overall, these supplementary results indicate that FAF-HQNN exhibits a degree of tolerance to moderate simulated quantum hardware noise in the current NISQ-compatible setting. At the same time, these observations should be interpreted with caution, since they are derived from hardware-noise simulation under fixed noise coefficients and a relatively lightweight circuit configuration rather than from deployment on real quantum devices. Nevertheless, the results support the view that lightweight circuit design and effective hybrid fusion may help mitigate the practical impact of circuit-level noise in hybrid quantum-classical models.

References

[1] Alchieri, L., Badalotti, D., Bonardi, P., Bianco, S.: An introduction to quantum machine learning: from quantum logic to quantum deep learning. Quantum Machine Intelligence **3**(2), 28 (2021)

Table A2: Performance of FAF-HQNN on BloodMNIST under simulated quantum noise (mean ± std). The noise coefficients are set to $p_1 = 0.001$ for single-qubit gates and $p_2 = 0.01$ for two-qubit gates.

Setting	**ACC**	**Precision**	**Recall**	**F1-score**	**AUC**
Noiseless	0.9503±0.0052	0.9520±0.0042	0.9503±0.0052	0.9503±0.0052	0.9943±0.0008
Bit-flip	0.9493±0.0045	0.9521±0.0028	0.9493±0.0045	0.9495±0.0044	0.9914±0.0036
Phase-flip	0.9520±0.0045	0.9538±0.0044	0.9520±0.0045	0.9522±0.0046	0.9936±0.0012
Amplitude damping	0.9522±0.0037	0.9531±0.0035	0.9522±0.0037	0.9521±0.0038	0.9943±0.0013
Depolarizing noise	0.9502±0.0046	0.9519±0.0034	0.9502±0.0046	0.9501±0.0044	0.9929±0.0021

[2] Benedetti, M., Lloyd, E., Sack, S., Fiorentini, M.: Parameterized quantum circuits as machine learning models. Quantum Science and Technology **4**(4), 043001 (2019)

[3] Litjens, G., Kooi, T., Bejnordi, B.E., *et al.*: A survey on deep learning in medical image analysis. Medical Image Analysis **42**, 60–88 (2017)

[4] Esteva, A., Kuprel, B., Novoa, R.A., *et al.*: Dermatologist-level classification of skin cancer with deep neural networks. Nature **542**(7639), 115–118 (2017)

[5] Shen, D., Wu, G., Suk, H.-I.: Deep learning in medical image analysis. Annual Review of Biomedical Engineering **19**, 221–248 (2017)

[6] Ajlouni, N., Özyavaş, A., Takaoğlu, M., Takaoğlu, F., Ajlouni, F.: Medical image diagnosis based on adaptive hybrid quantum cnn. BMC Medical Imaging **23**(1), 126 (2023)

[7] Hdaib, M., Rajasegarar, S., Pan, L.: Quantum deep learning-based anomaly detection for enhanced network security. Quantum Machine Intelligence **6**(1), 26 (2024)

[8] Gujju, Y., Matsuo, A., Raymond, R.: Quantum machine learning on near-term quantum devices: Current state of supervised and unsupervised techniques for real-world applications. Physical Review Applied **21**(6), 067001 (2024)

[9] Chow, J.C.: Quantum computing in medicine. Medical Sciences **12**(4), 67 (2024)

[10] Suneel, S., Krishnamoorthy, R., Gopatoti, A., Maguluri, L.P., Kuncha, P., Sunil, G.: Retracted article: Enhanced image diagnosing approach in medicine using quantum adaptive machine learning techniques. Optical and Quantum Electronics **56**(4), 534 (2024)

[11] Huang, S.-Y., An, W.-J., Zhang, D.-S., Zhou, N.-R.: Image classification and adversarial robustness analysis based on hybrid quantum–classical convolutional neural network. Optics Communications **533**, 129287 (2023)

[12] Liu, J., Lim, K.H., Wood, K.L., Huang, W., Guo, C., Huang, H.-L.: Hybrid quantum-classical convolutional neural networks. Science China Physics, Mechanics & Astronomy **64**(9), 290311 (2021)

[13] Wei, L., Liu, H., Xu, J., Shi, L., Shan, Z., Zhao, B., Gao, Y.: Quantum machine learning in medical image analysis: A survey. Neurocomputing **525**, 42–53 (2023)

[14] Yang, J., Soltan, A.A., Clifton, D.A.: Machine learning generalizability across healthcare settings: insights from multi-site covid-19 screening. NPJ digital medicine **5**(1), 69 (2022)

[15] Zhou, S., Greenspan, H., Shen, D., al.: A review of deep learning in medical image analysis: Applications and challenges. Journal of Imaging **7**(12), 225 (2021) https://doi.org/10.3390/jimaging7120225

[16] Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4700–4708 (2017)

[17] Rajpurkar, P., Irvin, J., Ball, R., *et al.*: Deep learning for chest radiograph diagnosis: A retrospective comparison of the chexnext algorithm to practicing radiologists. PLoS Medicine **15**(11), 1002686 (2018) https://doi.org/10.1371/journal.pmed.1002686

[18] Shao, W., Wang, X., Wu, B., *et al.*: Histopathology image classification: A review with deep learning solutions. Computers in Biology and Medicine **144**, 105350 (2022) https://doi.org/10.1016/j.compbiomed.2022.105350

[19] Graham, S., Valcarcel Macua, J., Jansen, M.: Cells classification with deep features and random forests. In: IEEE International Symposium on Biomedical Imaging (ISBI), pp. 565–568 (2019). https://doi.org/10.1109/ISBI.2019.8759276

[20] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778 (2016). Introduces residual learning to enable very deep networks with identity shortcut connections, addressing degradation problem in training deep nets

[21] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)

[22] Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 10012–10022 (2021)

[23] Zech, J.R., Badgeley, M.A., *et al.*: Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. PLoS Medicine **15**(11), 1002683 (2018)

[24] Shorten, C., Khoshgoftaar, T.M.: A survey on image data augmentation for deep learning. Journal of big data **6**(1), 1–48 (2019)

[25] Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations (ICLR) (2018)

[26] Yang, J., Shi, R., Wei, D., Liu, Z., Zhao, L., Ke, B., Pfister, H., Ni, B.: Medmnist v2—a large-scale lightweight benchmark for 2d and 3d biomedical image classification. Scientific Data **10**(1), 41 (2023)

[27] Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., Lloyd, S.: Quantum machine learning. Nature **549**(7671), 195–202 (2017)

[28] Chen, F., Jiang, L., Müller, H., Richerme, P., Chu, C., Fu, Z., Yang, M.: Nisq quantum computing: A security-centric tutorial and survey [feature]. IEEE Circuits and Systems Magazine **24**(1), 14–32 (2024)

[29] Mari, A., Bromley, T.R., Izaac, J., Schuld, M., Killoran, N.: Transfer learning in hybrid classical-quantum neural networks. Quantum **4**, 340 (2020)

[30] Obayya, M., Saeed, M.K., Alruwais, N., Alotaibi, S.S., Assiri, M., Salama, A.S.: Hybrid metaheuristics with deep learning-based fusion model for biomedical image analysis. IEEE Access **11**, 117149–117158 (2023)

[31] Singh, A.K., Saxena, D., Kumar, J., Gupta, V.: A quantum approach towards the adaptive prediction of cloud workloads. IEEE Transactions on Parallel and Distributed Systems **32**(12), 2893–2905 (2021)

[32] Fan, F., Shi, Y., Guggemos, T., Zhu, X.X.: Hybrid quantum-classical convolutional neural network model for image classification. IEEE transactions on neural networks and learning systems **35**(12), 18145–18159 (2023)

[33] Li, W., Chu, P.-C., Liu, G.-Z., Tian, Y.-B., Qiu, T.-H., Wang, S.-M.: An image classification algorithm based on hybrid quantum classical convolutional neural network. Quantum Engineering **2022**(1), 5701479 (2022)

[34] Henderson, M., Shakya, S., Pradhan, S., Cook, T.: Quanvolutional neural networks: powering image recognition with quantum circuits. Quantum Machine Intelligence **2**(1), 2 (2020)

[35] Chen, S.Y.-C., Huang, C.-M., Hsing, C.-W., Kao, Y.-J.: An end-to-end trainable hybrid classical-quantum classifier. Machine Learning: Science and Technology **2**(4), 045021 (2021)

[36] Senokosov, A., Sedykh, A., Sagingalieva, A., Kyriacou, B., Melnikov, A.: Quantum machine learning for image classification. Machine Learning: Science and Technology **5**(1), 015040 (2024)

[37] Mahmud, J., Mashtura, R., Fattah, S.A., Saquib, M.: Quantum convolutional neural networks with interaction layers for classification of classical data. Quantum Machine Intelligence **6**(1), 11 (2024)

[38] Skolik, A., McClean, J.R., Mohseni, M., Van Der Smagt, P., Leib, M.: Layerwise learning for quantum neural networks. Quantum Machine Intelligence **3**(1), 5 (2021)

[39] Arthur, D., Date, P.: Hybrid quantum-classical neural networks. In: 2022 IEEE International Conference on Quantum Computing and Engineering (QCE), pp. 49–55 (2022). IEEE

[40] Wang, L., Liu, Y., Ji, J., Meng, F., Zhang, Z., Yu, X.: Hybrid quantum-classical inception neural network for image classification. Advanced Quantum Technologies **8**(10), 2400700 (2025)

[41] Gong, L.-H., Pei, J.-J., Zhang, T.-F., Zhou, N.-R.: Quantum convolutional neural network based on variational quantum circuits. Optics Communications **550**, 129993 (2024)

[42] Bartolucci, S., Birchall, P., Bombin, H., Cable, H., Dawson, C., Gimeno-Segovia, M., Johnston, E., Kieling, K., Nickerson, N., Pant, M., *et al.*: Fusion-based quantum computation. Nature Communications **14**(1), 912 (2023)

[43] Qu, Z., Li, Y., Tiwari, P.: Qnmf: A quantum neural network based multimodal fusion system for intelligent diagnosis. Information Fusion **100**, 101913 (2023)

[44] Gordienko, Y., Trochun, Y., Stirenko, S.: Multimodal quanvolutional and convolutional neural networks for multi-class image classification. Big Data and Cognitive Computing **8**(7), 75 (2024)