

Uncertainty-Aware Variational Quantum Feature Learning for Reliable High-Dimensional Classification

Syed Basha Shaik

syba1139@gmail.com

KKR & KSR Institute of Technology and Sciences

Srihari Varma Mantena

SRKR Engineering College

Shaik Janbhasha

CVR College of Engineering

V. Malsoru

CMR Technical Campus

Ravikiran Reddy Kandadi

Vignana Bharathi Institute of Technology

RADHAKRISHNAN S

KKR & KSR Institute of Technology and Sciences

Santhi Tadikonda

KKR & KSR Institute of Technology and Sciences

Research Article

Keywords: Hybrid Quantum Machine Learning, Variational Quantum Feature Learning, Uncertainty-Aware Learning, Adaptive Feature Confidence, Variational Quantum Circuit, Hybrid Quantum-Classical Classification

Posted Date: August 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10455837/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

Uncertainty-Aware Variational Quantum Feature Learning for Reliable High-Dimensional Classification

Dr Syed Basha Shaik^{1*}, Dr Srihari Varma Mantena², Dr Shaik Janbhasha³, Dr V.Malsoru⁴,
Ravikiran Reddy Kandadi⁵, Dr RADHAKRISHNAN S⁶, Dr Santhi Tadikonda^{7*}

^{1,6,7} Department of Computer Science and Engineering (AI & ML),
KKR & KSR Institute of Technology and Sciences, Guntur, Andhra Pradesh, India

²Department of Computer Science and Engineering,
SRKR Engineering College, Bhimavaram, Andhra Pradesh, India

³Dept.of Computer Science and Engineering(Data Science),
CVR College of Engineering, Hyderabad, Telangana , India.

⁴Department of Computer Science and Engineering,
CMR Technical Campus, Hyderabad, Telangana, India

⁵Dept.of Computer Science and Engineering (Cybersecurity)
Vignana Bharathi Institute of Technology, Hyderabad, Telangana. India.

Corresponding Authors*: Dr Syed Basha Shaik*(syba1139@gmail.com) and
Dr Santhi Tadikonda*(shaikarchi@gmail.com)

Abstract

High-dimensional classification remains challenging for hybrid quantum–classical learning because conventional variational quantum feature learning methods often assign equal importance to all input features and do not explicitly account for feature uncertainty. This paper presents **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)**, a hybrid framework that combines adaptive feature weighting, uncertainty estimation, and variational quantum representation learning for reliable classification. The proposed framework incorporates an **Adaptive Feature Confidence Module (AFCM)** to estimate the contribution of individual input features before quantum encoding and an **Uncertainty Estimation Module (UEM)** to characterize feature reliability during the learning process. The confidence-weighted features are encoded using a variational quantum circuit and subsequently processed by an enhanced multi-layer classical classifier. Benchmark datasets, including MNIST, Fashion-MNIST, and PathMNIST, are employed during data preprocessing and quantum feature generation to establish the proposed framework. The detailed experimental evaluation, performance analysis, and comparative results presented in this study are conducted on the MNIST benchmark dataset to provide a comprehensive assessment of the proposed architecture. A systematic investigation using 4-, 6-, and 8-qubit hybrid models examines the influence of quantum feature

dimensionality on classification performance. The final UVQFL model achieves a **best validation accuracy of 84.62%**, with **84.60% precision**, **84.62% recall**, and an **84.59% F1-score**. The results indicate that integrating adaptive feature confidence with uncertainty-aware quantum feature learning improves feature representation quality and enhances the reliability of hybrid quantum–classical classification.

Keywords: Hybrid Quantum Machine Learning, Variational Quantum Feature Learning, Uncertainty-Aware Learning, Adaptive Feature Confidence, Variational Quantum Circuit, Hybrid Quantum–Classical Classification.

1 Introduction

Recent advances in artificial intelligence (AI) and deep learning have significantly transformed image classification across diverse domains, including medical diagnosis, autonomous driving, remote sensing, industrial inspection, and intelligent surveillance [10,14,15]. Despite their remarkable success, conventional deep learning models often require substantial computational resources and struggle to efficiently process increasingly high-dimensional data [5,10,15]. Quantum computing has emerged as a promising paradigm that exploits the principles of superposition and entanglement to enable novel approaches for information processing, leading to growing interest in quantum machine learning (QML) [3,12,16,19]. Among various QML techniques, hybrid quantum–classical learning has gained considerable attention because it combines the expressive capability of variational quantum circuits with the optimisation efficiency of classical neural networks, making it suitable for current Noisy Intermediate-Scale Quantum (NISQ) devices [3,16,22,23]. Although existing hybrid models have demonstrated encouraging performance, they generally focus on improving quantum feature extraction while overlooking the reliability of the generated quantum features [4,9,17,18]. Consequently, uncertain or noisy feature representations may adversely affect classification performance and reduce model robustness [1,6,8,11]. This limitation is particularly critical in applications such as medical image analysis, where reliable and trustworthy predictions are essential [1,6]. Motivated by these challenges, this work proposes an **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** framework that integrates adaptive feature confidence estimation and uncertainty-aware feature refinement to improve the reliability of quantum feature representations and enhance hybrid quantum–classical image classification performance.

Despite the promising progress of hybrid quantum–classical learning, several challenges continue to limit its effectiveness in practical image classification tasks [3,16,17,18,21]. Most existing approaches primarily concentrate on designing more expressive variational quantum circuits, improving quantum feature encoding, or enhancing classical classifier architectures, while assuming that all extracted quantum features contribute equally to the final prediction [17,18,22,23,24]. In practice, however, quantum feature representations may contain uncertainty arising from quantum noise, limited qubit resources, variational optimization, and complex high-dimensional data distributions [12,16,22,25,26]. The absence of mechanisms to estimate feature confidence or refine uncertain representations can reduce classification robustness and limit generalization performance, particularly in safety-critical applications such as medical image

analysis [1,6,8,11]. Furthermore, many existing hybrid models rely on shallow classical classifiers that may not fully exploit the discriminative information contained in quantum features [4,9,17,18]. These limitations highlight the need for hybrid quantum learning frameworks that not only generate expressive quantum representations but also evaluate their reliability and adaptively refine them before the final classification stage.

To address the limitations of existing hybrid quantum–classical learning approaches, this paper proposes an **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** framework for reliable image classification. The proposed framework integrates three key components within a unified architecture: an **Adaptive Feature Confidence Module (AFCM)**, an **Uncertainty Estimation Module (UEM)**, and a **Variational Quantum Circuit (VQC)** followed by an enhanced classical classifier. Initially, input features are processed to estimate their confidence levels, allowing the model to identify the relative importance of different feature representations. These confidence-aware features are subsequently refined through the uncertainty estimation module to reduce the influence of unreliable information before quantum feature encoding. The refined features are then transformed into quantum representations using a variational quantum circuit [22,23,24,26] and finally classified by a multi-layer classical neural network [10,15]. By incorporating uncertainty-aware feature refinement prior to quantum feature learning, the proposed UVQFL framework aims to improve the reliability, robustness, and discriminative capability of hybrid quantum–classical models while remaining compatible with current Noisy Intermediate-Scale Quantum (NISQ) computing environments [3,12,16].

The main contributions of this work are as follows. First, an **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** framework is proposed to improve the reliability and robustness of hybrid quantum–classical image classification. Second, an **Adaptive Feature Confidence Module (AFCM)** is introduced to estimate the importance of extracted features before quantum processing, while an **Uncertainty Estimation Module (UEM)** is developed to refine feature representations by reducing the influence of uncertain information. Third, an enhanced hybrid architecture integrating variational quantum feature learning with a deeper classical classifier is designed to improve the utilization of refined quantum features and enhance predictive performance. Finally, experiments conducted using the proposed hybrid quantum–classical learning frameworks demonstrate the effectiveness of uncertainty-aware feature refinement, with the detailed experimental evaluation presented on the MNIST benchmark dataset.

2 Related work

2.1 Deep Learning-Based Image Classification

Deep learning has revolutionized image classification by enabling automatic learning of hierarchical feature representations from large-scale datasets [10,14,15]. Convolutional Neural Networks (CNNs), including architectures such as AlexNet, VGG, ResNet, DenseNet, and EfficientNet, have consistently demonstrated outstanding performance across a wide range of computer vision applications, including medical image analysis, object recognition, remote sensing, and intelligent surveillance [10,14,15]. More recently, Vision Transformers (ViTs) and hybrid CNN–Transformer architectures have further improved classification accuracy by

effectively modeling both local and global feature dependencies [5]. Despite these significant advances, deep learning models generally require large annotated datasets, extensive computational resources, and considerable training time, limiting their applicability in resource-constrained environments [10,15]. Furthermore, as data dimensionality and model complexity increase, conventional deep learning methods may experience challenges related to computational efficiency, scalability, and robustness [6,10]. These limitations have motivated researchers to investigate alternative computational paradigms, including quantum machine learning, to complement classical deep learning and improve the efficiency of high-dimensional image classification tasks [3,16,19].

2.2 Hybrid Quantum–Classical Learning

Hybrid quantum–classical learning has emerged as one of the most practical paradigms in quantum machine learning by combining the computational advantages of variational quantum circuits (VQCs) with the optimization capabilities of classical neural networks [3,16,19,22]. Unlike fully quantum models, hybrid approaches leverage classical preprocessing and optimization while employing parameterized quantum circuits to encode and transform feature representations into higher-dimensional quantum spaces [22,23,26]. Recent studies have demonstrated the effectiveness of hybrid architectures for image classification by integrating quantum feature encoding with classical classifiers, achieving encouraging performance on benchmark datasets such as MNIST, Fashion-MNIST, and MedMNIST [4,9,17,18,20]. Various techniques, including angle encoding, amplitude encoding, strongly entangling layers, data re-uploading, and hybrid convolutional quantum networks, have been explored to enhance feature representation and improve classification accuracy [22,23,24,26]. However, most existing hybrid quantum–classical models primarily focus on optimizing quantum circuit architectures, encoding strategies, or classical classifiers, while directly forwarding the extracted quantum features to the classification stage [17,18,22,23]. Consequently, the reliability and uncertainty of quantum feature representations are seldom considered during the learning process, potentially limiting the robustness and generalization capability of hybrid quantum learning systems, particularly for complex and high-dimensional image classification tasks [1,6,16,19]. These limitations motivate the development of more reliable hybrid frameworks that incorporate uncertainty-aware feature refinement before the final classification stage.

2.3 Uncertainty-Aware Learning

Uncertainty-aware learning has become an important research area in machine learning because it enables predictive models to estimate the reliability of their decisions in addition to generating classification outputs [1,6,8,11]. Several techniques, including Bayesian neural networks, Monte Carlo dropout, ensemble learning, evidential deep learning, and confidence estimation methods, have been developed to quantify predictive uncertainty and improve model robustness, particularly in safety-critical applications such as medical image analysis and autonomous systems [1,6,8,11]. By identifying uncertain or unreliable feature representations, these approaches help reduce overconfident predictions and enhance the interpretability and generalization capability of learning models [1,6,8]. Although uncertainty estimation has been extensively investigated in classical deep learning, its integration into hybrid quantum–classical learning remains relatively limited [3,16,19]. Most existing quantum machine learning

approaches primarily emphasize quantum feature extraction and classification accuracy without explicitly evaluating the reliability of the learned quantum features [17,18,21,22]. Consequently, uncertainty-aware feature refinement represents a promising direction for improving the robustness and trustworthiness of hybrid quantum–classical image classification systems [1,6,16]. The proposed **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** framework addresses this research gap by incorporating adaptive feature confidence estimation and uncertainty-aware feature refinement prior to the final classification stage.

2.4 Research Gap

Although significant progress has been achieved in deep learning and hybrid quantum–classical image classification, several important challenges remain unresolved [3,16,17,18,20]. Existing hybrid quantum learning approaches predominantly focus on improving variational quantum circuit design, quantum feature encoding strategies, or classical classifier architectures, while assuming that all quantum features contribute equally to the final prediction [22,23,24,26]. As a result, the reliability and uncertainty of learned quantum feature representations are rarely considered during the feature learning process [1,6,8,11]. This limitation may reduce classification robustness, particularly when handling noisy, ambiguous, or high-dimensional data [6,11]. Moreover, few studies have investigated adaptive mechanisms for evaluating feature confidence before quantum classification, leaving an important gap in the development of reliable hybrid quantum learning systems [16,19,21]. To address these limitations, this work proposes an **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** framework that integrates an **Adaptive Feature Confidence Module (AFCM)** and an **Uncertainty Estimation Module (UEM)** to assess and refine quantum feature representations prior to classification. By incorporating uncertainty-aware feature refinement into the hybrid quantum–classical learning pipeline, the proposed framework aims to improve the reliability, robustness, and overall classification performance of image classification systems.

3. Proposed Uncertainty-Aware Variational Quantum Feature Learning (UVQFL) Framework

The proposed **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** framework is developed to improve the reliability and discriminative capability of hybrid quantum–classical image classification by progressively enhancing feature representations through uncertainty-aware learning and variational quantum computing [3,16,19,22]. As illustrated in **Figure 1**, the overall framework consists of five models, including three baseline Hybrid Variational Quantum Feature Learning (Hybrid VQFL) models with 4-, 6-, and 8-qubit quantum circuits, the proposed UVQFL model, and the Enhanced UVQFL model. Initially, the input images are preprocessed and transformed into low-dimensional feature vectors suitable for quantum encoding. The baseline Hybrid VQFL models investigate the influence of different quantum circuit sizes on feature representation and classification performance. Building upon the best-performing baseline configuration, the proposed UVQFL model introduces an **Adaptive Feature Confidence Module (AFCM)** to estimate the importance of individual features and an **Uncertainty Estimation Module (UEM)** to identify and refine uncertain feature representations before quantum processing [1,6,8,11]. The refined features are then encoded into a parameterized

Variational Quantum Circuit (VQC), where trainable quantum operations extract expressive quantum features that are subsequently measured and transferred to the classical learning stage [22,23,24,26]. Finally, the Enhanced UVQFL model employs an improved multi-layer classical classifier to learn more discriminative decision boundaries from the refined quantum features, resulting in superior classification performance. By integrating adaptive feature confidence estimation, uncertainty-aware feature refinement, variational quantum feature extraction, and enhanced classical learning into a unified framework, the proposed architecture provides a systematic progression from conventional hybrid quantum learning to a more robust and reliable uncertainty-aware classification framework [16,19,22]. Figure 1 illustrates the overall architecture and workflow of the proposed UVQFL framework, highlighting the relationship among the five evaluated models.




Figure 1. Overall architecture of the proposed Uncertainty-Aware Variational Quantum Feature Learning (UVQFL) framework

3.2 Baseline Hybrid Variational Quantum Feature Learning (Hybrid VQFL)

To establish a strong baseline for evaluating the proposed uncertainty-aware framework, three Hybrid Variational Quantum Feature Learning (Hybrid VQFL) models with different quantum circuit sizes were developed and evaluated. These baseline models employ a conventional hybrid quantum–classical architecture in which classical preprocessing is followed by variational quantum feature extraction and classical classification [3,17,18,22]. The primary objective is to investigate the influence of quantum circuit dimensionality on feature representation and classification performance. Accordingly, three configurations consisting of **4-, 6-, and 8-qubit Variational Quantum Circuits (VQCs)** are implemented while maintaining the same classical preprocessing and training strategy. In each model, the input features are first normalized and encoded into quantum states using an appropriate data encoding scheme [23]. The encoded

quantum states are then processed through parameterized variational quantum circuits composed of trainable rotation gates and entanglement operations to generate expressive quantum feature representations [22,24,26]. After quantum measurement, the extracted features are supplied to a classical neural network classifier for final prediction. Increasing the number of qubits enables the quantum circuit to represent more complex feature interactions and higher-dimensional latent representations, allowing the effect of quantum feature dimensionality on classification performance to be systematically analyzed [25,26]. These baseline Hybrid VQFL models provide the reference performance against which the proposed **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** and Enhanced UVQFL models are evaluated in the subsequent sections.

3.3 Proposed Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)

The proposed **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** model extends the baseline Hybrid VQFL architecture by incorporating uncertainty-aware feature refinement prior to quantum feature extraction [1,6,8,11]. Unlike conventional hybrid quantum learning methods that treat all input features with equal importance, the proposed model evaluates both the confidence and uncertainty associated with each feature to generate more reliable feature representations for quantum processing.

3.3.1 Adaptive Feature Confidence Module (AFCM)

The Adaptive Feature Confidence Module (AFCM) estimates the relative importance of each input feature before quantum encoding. Instead of treating all extracted features equally, AFCM learns adaptive confidence weights that emphasize discriminative features while suppressing less informative components. This adaptive weighting mechanism improves the quality of the feature representation before quantum processing.

The confidence weights are computed as

$$\mathbf{C} = \sigma(\mathbf{W}_c \mathbf{X} + \mathbf{b}_c)$$

Where $\mathbf{X} \in \mathbb{R}^d$ denotes the input feature vector, $\mathbf{W}_c$ represents the trainable weight matrix, and $\mathbf{b}_c$ denotes the bias vector, $\sigma(\cdot)$ is the sigmoid activation function, $\mathbf{C} \in [0,1]^d$ represents the confidence weights.

The confidence-weighted feature vector is then obtained as $\mathbf{X}_c = \mathbf{C} \odot \mathbf{X}$

Where $\odot$ denotes element-wise multiplication, $\mathbf{X}_c$ represents the confidence-enhanced feature vector.

3.3.2 Uncertainty Estimation Module (UEM)

The confidence-enhanced features are subsequently processed by the Uncertainty Estimation Module (UEM), where a deterministic confidence margin score is computed to estimate the relative reliability of the feature representations before quantum encoding. Rather than performing probabilistic uncertainty estimation such as Bayesian inference or Monte Carlo

sampling, the proposed module utilizes the prediction confidence obtained from the current model to attenuate less reliable feature representations prior to quantum processing. This confidence-guided refinement was deliberately adopted as a lightweight alternative to probabilistic uncertainty estimation, improving feature reliability while maintaining computational efficiency suitable for hybrid quantum–classical learning.

The deterministic confidence margin score is computed as $U = 1 - \max(\mathbf{P})$

Where $\mathbf{P}$ denotes the predicted probability distribution obtained from the confidence-enhanced feature representation and $U \in [0,1]$ represents the estimated uncertainty score. Higher values of the confidence margin score indicate lower prediction confidence and therefore greater attenuation of the corresponding feature representation before quantum encoding.

The refined feature vector is then obtained as $\mathbf{X}_r = (1-U) \odot \mathbf{X}_c$

Where $\mathbf{X}_c$ denotes the confidence-enhanced feature vector obtained from the AFCM, $\mathbf{X}_r$ represents the uncertainty-refined feature vector, and $\odot$ denotes element-wise multiplication. This refinement reduces the influence of noisy or ambiguous features while preserving highly reliable feature information. Although the proposed score is used as a confidence-guided feature attenuation mechanism, it should not be interpreted as a formal estimator of epistemic or aleatoric uncertainty. Instead, it provides a computationally efficient deterministic confidence measure suitable for adaptive feature refinement within the proposed hybrid quantum–classical framework.

3.3.3 Variational Quantum Feature Encoding

The uncertainty-refined feature vector $\mathbf{X}_r$ is subsequently encoded into a parameterized **Variational Quantum Circuit (VQC)** to generate expressive quantum feature representations. Prior to quantum processing, the refined classical features are transformed into quantum states using angle encoding, where each feature is mapped to the rotation angle of a corresponding qubit. This encoding strategy preserves the numerical characteristics of the refined features while enabling quantum superposition for enhanced feature representation.

The VQC consists of multiple trainable quantum layers composed of parameterized single-qubit rotation gates and entanglement operations. The rotation gates learn nonlinear feature transformations, whereas the entanglement gates capture complex correlations among the encoded features that are difficult to model using classical feature transformations alone. The trainable quantum circuit is optimized jointly with the classical classifier through gradient-based optimization, allowing end-to-end learning of discriminative quantum representations [22–24,26].

Following quantum evolution, measurements are performed on each qubit in the computational basis to obtain expectation values corresponding to the quantum feature representation. These measured quantum features are subsequently transferred to the classical neural network classifier, where they are processed through fully connected layers to produce the final class prediction.

By integrating adaptive feature confidence estimation, uncertainty-aware feature refinement, and variational quantum feature encoding into a unified hybrid learning framework, the proposed UVQFL model generates robust and informative quantum feature representations, resulting in improved classification accuracy, reliability, and generalization compared with conventional Hybrid VQFL models [17,18,20,21].

3.4 Enhanced UVQFL

To further improve the performance of the proposed framework, an **Enhanced UVQFL** model is developed by strengthening the classical learning component while preserving the uncertainty-aware quantum feature extraction process. The enhanced model retains the **Adaptive Feature Confidence Module (AFCM)**, **Uncertainty Estimation Module (UEM)**, and **Variational Quantum Circuit (VQC)** introduced in the UVQFL framework but replaces the conventional classifier with a deeper multi-layer neural network capable of learning more discriminative decision boundaries from the refined quantum features [10,15]. The enhanced classifier incorporates multiple fully connected layers with nonlinear activation functions and dropout regularization to improve feature learning, reduce overfitting, and enhance model generalization [10,15]. During end-to-end training, the parameters of the AFCM, UEM, variational quantum circuit, and classical classifier are jointly optimized through backpropagation, enabling effective interaction between the classical and quantum components [22,23]. This integrated optimization strategy allows the enhanced architecture to exploit high-quality uncertainty-refined quantum features more effectively than the standard UVQFL model. Consequently, the Enhanced UVQFL achieves superior classification performance by combining uncertainty-aware feature refinement, expressive variational quantum feature learning, and a more powerful classical classifier, demonstrating the effectiveness of the proposed hybrid quantum–classical learning framework for reliable high-dimensional image classification [16,17,18,20,21].

4 Methods

4.1 Datasets

To comprehensively evaluate the effectiveness and generalization capability of the proposed **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** framework, experiments were conducted on three widely used benchmark image classification datasets: **MNIST** [27], **Fashion-MNIST** [28], and **PathMNIST** [29]. These datasets were selected because they represent increasing levels of visual complexity, enabling a systematic assessment of the proposed framework across handwritten digit recognition, object classification, and medical image analysis.

MNIST consists of grayscale handwritten digit images belonging to 10 classes (digits 0–9), with each image having a resolution of 28×28 pixels [27]. The dataset contains **60,000 training images** and **10,000 testing images** and serves as a standard benchmark for evaluating image classification algorithms.

Fashion-MNIST contains grayscale images of 10 categories of fashion products, also with an image resolution of 28×28 pixels [28]. The dataset comprises **60,000 training images** and

10,000 testing images. Compared with MNIST, Fashion-MNIST presents greater intra-class similarity and inter-class variability, making it a more challenging benchmark for feature learning.

PathMNIST, a subset of the **MedMNIST v2** benchmark collection [29], consists of **nine classes** of histopathological tissue images with a resolution of **28×28** pixels. The dataset includes **89,996 training images**, **10,004 validation images**, and **7,180 testing images**, providing a representative benchmark for medical image classification.

The diversity of these datasets enables the proposed framework to be evaluated under varying levels of classification difficulty, ranging from simple handwritten digit recognition to complex medical image analysis. Using these benchmark datasets ensures a comprehensive evaluation of the proposed UVQFL framework and demonstrates its ability to learn robust and reliable quantum feature representations across different application domains.

Table 1 summarizes the benchmark datasets used in this study, including their characteristics and reasons for selection.

Table 1. Summary of the datasets used in this study.

Dataset	Domain	Classes	Image Size	Training Images	Validation Images	Testing Images	Reason for Selection
MNIST	Handwritten digit recognition	10	28×28	60,000	–	10,000	Standard benchmark for evaluating image classification models with relatively simple visual patterns.
Fashion-MNIST	Fashion object classification	10	28×28	60,000	–	10,000	More challenging than MNIST due to higher inter-class similarity and intra-class variability.
PathMNIST	Histopathological tissue image classification	9	28×28	89,996	10,004	7,180	Medical imaging benchmark used to evaluate robustness and generalization on complex real-world datasets.

Benchmark datasets, including MNIST, Fashion-MNIST, and PathMNIST, are employed during data preprocessing and quantum feature generation to establish the proposed framework. The detailed experimental evaluation, performance analysis, and comparative results presented in this study are conducted on the MNIST benchmark dataset to provide a comprehensive assessment of the proposed architecture.

4.2 Data Preprocessing

Prior to model training, all input images underwent a standardized preprocessing pipeline to generate compact and quantum-compatible feature representations suitable for variational quantum learning. Initially, the grayscale images from the **MNIST** [27], **Fashion-MNIST** [28], and **PathMNIST** [29] datasets were normalized by scaling the pixel intensity values from the range $[0, 255]$ to $[0, 1]$, thereby reducing numerical instability and facilitating faster model convergence during training. Each normalized image of size 28×28 pixels was subsequently flattened into a one-dimensional feature vector containing **784 features**, enabling efficient processing by the classical preprocessing modules.

Since current **Noisy Intermediate-Scale Quantum (NISQ)** devices have limited qubit resources and cannot directly encode high-dimensional image vectors [12,16], **Principal Component Analysis (PCA)** was employed to reduce the feature dimensionality while preserving the most informative variance within the data. The dimensionality reduction process generated low-dimensional feature vectors corresponding to the number of qubits used in the respective models, namely **4, 6, and 8 principal components** for the **Hybrid VQFL (4Q)**, **Hybrid VQFL (6Q)**, and **Hybrid VQFL (8Q)** architectures, respectively. Similarly, the proposed **UVQFL** and **Enhanced UVQFL** models utilized the reduced feature vectors as inputs to the uncertainty-aware feature refinement modules.

Following PCA, the reduced feature vectors were further standardized using feature scaling to produce normalized numerical values with consistent distributions, ensuring stable optimization during both classical and quantum learning. Finally, the processed feature vectors were transformed into quantum-compatible inputs and encoded into the variational quantum circuit through parameterized quantum data encoding [22,23], allowing each feature component to be mapped onto the corresponding qubit for subsequent quantum feature extraction. This preprocessing strategy effectively reduces computational complexity, preserves the most discriminative information, and enables efficient integration of high-dimensional image data with hybrid quantum–classical learning architectures.

4.3 Variational Quantum Circuit Design

The proposed hybrid quantum–classical framework employs **Variational Quantum Circuits (VQCs)** as the quantum feature extraction component to learn expressive low-dimensional representations from the preprocessed image features [3,16,22]. To investigate the influence of quantum circuit capacity on classification performance, three baseline VQC configurations consisting of **4-, 6-, and 8-qubit circuits** were developed, corresponding to the **Hybrid VQFL (4Q)**, **Hybrid VQFL (6Q)**, and **Hybrid VQFL (8Q)** models, respectively. The proposed

UVQFL and **Enhanced UVQFL** models utilize the same quantum circuit architecture while incorporating uncertainty-aware feature refinement before quantum encoding.

Following dimensionality reduction, each feature vector is transformed into a quantum-compatible representation through **angle encoding**, where each normalized feature component is mapped to the rotation angle of its corresponding qubit [22,23]. This encoding strategy efficiently represents classical information within the limited qubit resources available on **Noisy Intermediate-Scale Quantum (NISQ)** devices [12,16].

After data encoding, each qubit undergoes a sequence of trainable single-qubit rotation operations using parameterized **RY** and **RZ** gates, enabling the circuit to learn nonlinear feature transformations during optimization [22,26]. To capture correlations among different feature dimensions, adjacent qubits are connected using **Controlled-NOT (CNOT)** gates arranged in a linear entanglement topology, allowing quantum information to propagate across the circuit and generate richer feature representations [24,26].

The parameters of the rotation gates are optimized jointly with the classical neural network using gradient-based backpropagation, resulting in an end-to-end trainable hybrid learning architecture [17,18,22]. Upon completion of the variational quantum operations, the expectation values of the **Pauli-Z** operator are measured for each qubit to obtain a low-dimensional quantum feature vector. These measured expectation values constitute the learned quantum representation and are subsequently forwarded to the classical classifier for final prediction.

By varying the number of qubits from **4 to 8**, the proposed framework systematically investigates the effect of quantum feature dimensionality on classification performance while maintaining a consistent circuit design and training strategy across all models. This variational quantum circuit design enables efficient quantum feature extraction, improved representation learning, and seamless integration with the proposed uncertainty-aware hybrid quantum–classical classification framework [22,23,24,26].

4.4 Model Training

The proposed **Hybrid VQFL**, **UVQFL**, and **Enhanced UVQFL** models were trained using an end-to-end hybrid quantum–classical optimization strategy, where the parameters of both the classical neural network and the variational quantum circuit were jointly optimized through gradient-based backpropagation [17,18,22]. The **Adam optimizer** was employed owing to its efficient adaptive learning capability and stable convergence characteristics for hybrid quantum learning [30]. The initial learning rate was set to **0.001**, providing a suitable balance between convergence speed and optimization stability.

All models were trained for **50 epochs** with a **batch size of 64**, allowing efficient utilization of computational resources while maintaining stable gradient estimation. Since the task involves multi-class image classification, the **Cross-Entropy Loss** function was adopted to minimize the discrepancy between the predicted class probabilities and the ground-truth labels. During training, the model parameters were iteratively updated to minimize the classification loss,

enabling simultaneous optimization of the classical feature extractor, uncertainty-aware modules, variational quantum circuit, and classifier.

An early stopping strategy was not employed because the validation performance remained stable throughout the training process and the predefined number of epochs was sufficient for model convergence. The proposed framework was implemented using **PyTorch** for classical deep learning components [31] and **PennyLane** for quantum circuit simulation and hybrid quantum–classical optimization [32]. All experiments were conducted in the **Google Colab** environment using **Python**, with GPU acceleration provided by an **NVIDIA Tesla T4** graphics processor whenever available.

The software environment consisted of **PyTorch**, **TorchVision**, **PennyLane**, **NumPy**, **Scikit-learn**, and related scientific computing libraries for data preprocessing, model development, training, and performance evaluation. This experimental configuration provides an efficient and reproducible platform for evaluating the proposed hybrid quantum–classical learning framework across multiple benchmark datasets.

To improve the reproducibility of the proposed Enhanced UVQFL architecture, Table 2 summarizes the layer-wise parameterization of the model. The network contains a total of 3,106 trainable parameters distributed across the Adaptive Feature Confidence Module (AFCM), Variational Quantum Circuit (VQC), and the enhanced classical classifier. This detailed parameter breakdown facilitates implementation, comparison with related hybrid quantum–classical models, and experimental reproducibility.

Table 2. Layer-wise Architecture and Trainable Parameter Configuration of the Enhanced UVQFL Model

Component	Layer Configuration	Trainable Parameters
Adaptive Feature Confidence Module (AFCM)	Linear (8 $\rightarrow$ 8) + Sigmoid	72
Uncertainty Estimation Module (UEM)	Confidence margin computation	0
Variational Quantum Circuit (VQC)	8 qubits, 2 variational layers	48
Enhanced Classical Classifier	Linear (8 $\rightarrow$ 64) + ReLU + Dropout	576
	Linear (64 $\rightarrow$ 32) + ReLU + Dropout	2,080
	Linear (32 $\rightarrow$ 10)	330
Total		3,106

4.5 Performance Evaluation Metrics

The performance of the proposed **Hybrid VQFL**, **UVQFL**, and **Enhanced UVQFL** models was evaluated using four widely adopted classification metrics: **Accuracy**, **Precision**, **Recall**, and **F1-score**. These metrics provide a comprehensive assessment of the classification performance by measuring the overall prediction accuracy as well as the model's ability to correctly identify

positive samples while minimizing false predictions. Let **True Positives (TP)**, **True Negatives (TN)**, **False Positives (FP)**, and **False Negatives (FN)** denote the outcomes of the classification process.

The **Accuracy** metric measures the proportion of correctly classified samples among all predictions and is defined as:

Accuracy provides an overall measure of classification performance by considering both correctly identified positive and negative samples.

The **Precision** metric evaluates the proportion of correctly predicted positive samples among all samples predicted as positive and is computed as:

A higher precision indicates that the model produces fewer false-positive predictions.

The **Recall**, also referred to as **Sensitivity**, measures the proportion of actual positive samples that are correctly identified by the classifier and is defined as:

Higher recall values indicate that the model successfully detects a greater percentage of positive instances.

The **F1-score** represents the harmonic mean of precision and recall, providing a balanced measure of classification performance, particularly when class distributions are imbalanced. It is calculated as:

These four evaluation metrics collectively provide a comprehensive assessment of the effectiveness, robustness, and generalization capability of the proposed hybrid quantum–classical learning framework across the **MNIST**, **Fashion-MNIST**, and **PathMNIST** benchmark datasets. The same evaluation protocol was consistently applied to all baseline and proposed models to ensure a fair and objective performance comparison.

4.6 Experimental Configuration

All experiments were conducted under a consistent experimental setup to ensure a fair comparison among the **Hybrid VQFL (4Q, 6Q, and 8Q)**, **UVQFL**, and **Enhanced UVQFL** models. A fixed **random seed of 42** was used throughout the experiments to initialize the model parameters, shuffle the datasets, and control the training process, thereby improving the reproducibility of the obtained results. The datasets were divided according to their standard benchmark protocols, where the predefined training and testing splits of **MNIST** and **Fashion-MNIST** were utilized, while **PathMNIST** employed the official training, validation, and testing partitions provided by the MedMNIST benchmark. All models were trained using identical optimization settings to ensure an unbiased performance comparison. The principal hyperparameters included an **Adam optimizer**, an **initial learning rate of 0.001**, a **batch size of 64**, **50 training epochs**, and the **Cross-Entropy Loss** function for multi-class classification. The number of qubits was varied according to the evaluated model, with **4**, **6**, and **8** qubits used for the Hybrid VQFL baselines, while the proposed **UVQFL** and **Enhanced UVQFL** employed the

optimal quantum circuit configuration identified during the baseline evaluation. The **Enhanced UVQFL** model contained **3,106 trainable parameters**, jointly optimized across the uncertainty-aware modules, variational quantum circuit, and enhanced classical classifier. The implementation was developed using **Python, PyTorch, TorchVision, PennyLane, NumPy**, and **Scikit-learn**, and all experiments were executed in the **Google Colab** environment with **NVIDIA Tesla T4 GPU** acceleration when available. To facilitate reproducibility, identical preprocessing procedures, quantum circuit architectures, optimization strategies, hyperparameter settings, and evaluation metrics were maintained across all experiments. Furthermore, all models were trained using the same software environment and experimental protocol, ensuring that the reported performance differences are attributable solely to the proposed architectural improvements rather than variations in the experimental configuration.

5 Results

5.1 Experimental Evaluation Overview

The experimental results presented in this section focus on the MNIST benchmark dataset, providing a detailed evaluation of the proposed UVQFL framework through baseline comparisons, convergence analysis, confusion matrix evaluation, and overall classification performance. The benchmark datasets described in Section 4 were used during framework development, while MNIST was selected for the comprehensive experimental analysis reported in this study.

All experiments followed the experimental protocol described in Section 4. The input images were normalized and reduced in dimensionality using Principal Component Analysis (PCA) before quantum feature encoding. Baseline Hybrid Variational Quantum Feature Learning (Hybrid VQFL) models employing 4-, 6-, and 8-qubit variational quantum circuits were first investigated to study the effect of quantum feature dimensionality on classification performance. Based on the observations from these baseline experiments, the UVQFL model was developed by incorporating the Adaptive Feature Confidence Module (AFCM) and the Uncertainty Estimation Module (UEM). An enhanced version with a deeper classical classifier was subsequently evaluated to examine its contribution to predictive performance.

The experimental evaluation was carried out using four standard performance metrics: accuracy, precision, recall, and F1-score. Training and validation curves were analysed to examine the convergence behaviour of each model, while confusion matrix analysis was used to assess class-wise prediction performance for the final model.

The following sections first present the performance of the baseline Hybrid VQFL models, followed by the evaluation of the proposed UVQFL models. A comparative analysis is then provided to highlight the improvements achieved through uncertainty-aware quantum feature learning.


5.2 Baseline Hybrid Variational Quantum Feature Learning Performance

The baseline experiments investigated the influence of quantum feature dimensionality on classification performance. Three Hybrid Variational Quantum Feature Learning (Hybrid VQFL) models employing 4-, 6-, and 8-qubit variational quantum circuits were evaluated using the common experimental protocol described in Section 4, with only the number of qubits varied across the models.


As summarized in Table 3, increasing the number of qubits consistently improved classification performance. The 8-qubit Hybrid VQFL model achieved the highest baseline accuracy (81.29%) together with corresponding improvements in precision, recall, and F1-score. Figures 2 and 3 show stable convergence of all baseline models, with the 8-qubit configuration exhibiting the best optimization behaviour. Based on these observations, the 8-qubit model was selected as the baseline for the proposed UVQFL framework.

Table 3. Performance Comparison of Baseline Hybrid VQFL Models


Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
Hybrid VQFL (4 Qubits)	65.80	64.88	65.80	64.95
Hybrid VQFL (6 Qubits)	67.44	66.66	67.44	66.63
Hybrid VQFL (8 Qubits)	81.29	81.23	81.29	81.21



(a) 4-qubit Hybrid Variational Quantum Feature Learning (Hybrid VQFL)




(b) 6-qubit Hybrid Variational Quantum Feature Learning (Hybrid VQFL) model




(c) 8-qubit Hybrid Variational Quantum Feature Learning (Hybrid VQFL)

Figure 2. Training and validation loss curves of the baseline Hybrid VQFL models (4-, 6-, and 8-qubit configurations)



(a) 4-qubit Hybrid Variational Quantum Feature Learning (Hybrid VQFL)



(b) 6-qubit Hybrid Variational Quantum Feature Learning (Hybrid VQFL) **model**



(C) 8-qubit Hybrid Variational Quantum Feature Learning (Hybrid VQFL)

Figure 3. Training and validation accuracy curves of the baseline Hybrid VQFL models (4-, 6-, and 8-qubit configurations)

5.3 Performance of the Proposed UVQFL Models

To investigate whether uncertainty-aware feature learning improves the performance of hybrid quantum–classical classification, two variants of the proposed model were evaluated. The first model integrates the **Adaptive Feature Confidence Module (AFCM)** and the **Uncertainty Estimation Module (UEM)** into the baseline Hybrid VQFL architecture. The second model, referred to as the **Enhanced UVQFL**, retains the same quantum feature extraction mechanism while employing a deeper classical classifier to improve feature discrimination and classification capability. Both models were trained and evaluated using the same experimental protocol as the baseline models.

5.3.1 Proposed UVQFL Model


The proposed UVQFL model extends the baseline Hybrid VQFL architecture by incorporating the Adaptive Feature Confidence Module (AFCM) and the Uncertainty Estimation Module (UEM) to refine feature representations before quantum classification, emphasizing reliable features while attenuating uncertain information.

As summarized in Table 4, the proposed UVQFL model achieved competitive classification performance while introducing uncertainty-aware feature refinement without increasing the complexity of the variational quantum circuit. The training and validation curves shown in


Figure 4 demonstrate stable convergence throughout optimization, indicating that the proposed uncertainty-aware modules improve feature reliability while maintaining robust learning behaviour.

Table 4. Performance of the Proposed UVQFL Model

Metric	Value (%)
Accuracy	81.06
Precision	81.09
Recall	81.06
F1-score	81.04



(a) Training and Validation Loss Curves of the Proposed UVQFL Framework



(b) Training and Validation Accuracy Curves of the Proposed UVQFL Framework

Figure 4. Training and validation performance of the proposed UVQFL model: (a) loss curves and (b) accuracy curves.


5.3.2 Enhanced UVQFL Model

To further improve classification performance, the proposed UVQFL architecture was extended by incorporating a deeper multi-layer classical classifier while preserving the same uncertainty-aware quantum feature extraction process. The enhanced classifier was designed to exploit the refined quantum features more effectively during the final classification stage.


As summarized in Table 5, the Enhanced UVQFL model achieved the highest performance among all evaluated models. The learning curves shown in Figure 5 indicate stable optimization throughout training, while the confusion matrix in Figure 6 demonstrates strong class-wise prediction performance with relatively few misclassifications. These results confirm that combining uncertainty-aware quantum feature learning with a more expressive classical classifier enhances overall hybrid quantum–classical classification performance.

Table 5. Performance of the Enhanced UVQFL Model

Metric	Value (%)
Accuracy	84.62
Precision	84.60
Recall	84.62
F1-score	84.59



(a) Training and Validation Loss Curves of the Enhanced UVQFL Framework



(b) Training and Validation Accuracy Curves of the Enhanced UVQFL Framework

Figure 5. Training and validation performance of the Enhanced UVQFL model: (a) loss curves and (b) accuracy curves.




Figure 6. Confusion matrix of the Enhanced UVQFL model illustrating class-wise classification performance.

5.4 Comparative Performance Analysis

A comparative evaluation was conducted to examine the progression in classification performance from the baseline Hybrid Variational Quantum Feature Learning (Hybrid VQFL) models to the proposed uncertainty-aware architectures. Table 8 summarizes the performance of all evaluated models using four standard evaluation metrics: accuracy, precision, recall, and F1-score.

The baseline experiments demonstrated that increasing the number of qubits improved the quality of the learned quantum feature representation. The **4-qubit Hybrid VQFL** model achieved an accuracy of **65.80%**, which increased to **67.44%** with the **6-qubit** configuration. The **8-qubit Hybrid VQFL** model produced a substantial improvement, reaching an accuracy of **81.29%**, indicating that higher-dimensional quantum feature representations contributed to better classification performance.

The proposed **UVQFL** model incorporated adaptive feature confidence estimation and uncertainty-aware feature refinement while maintaining a comparable level of predictive performance (**81.06%** accuracy). Although the overall accuracy was close to that of the strongest baseline model, the proposed architecture introduced a mechanism for evaluating the reliability of quantum features before classification, providing a more robust learning strategy.

The **Enhanced UVQFL** model achieved the best overall performance, with an **accuracy of 84.62%**, **precision of 84.60%**, **recall of 84.62%**, and **F1-score of 84.59%**. The consistent improvement across all evaluation metrics indicates that combining uncertainty-aware quantum feature learning with a deeper classical classifier produced more discriminative feature representations and improved classification performance.

Overall, the results demonstrate that integrating uncertainty-aware feature refinement with an enhanced classical classifier consistently improves hybrid quantum–classical classification performance over the baseline Hybrid VQFL models.

A comprehensive comparative analysis of the baseline models and the proposed architectures across accuracy, precision, recall, and F1-score is presented in **Table 6**.

Table 6. Comparative Performance of All Evaluated Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
Hybrid VQFL (4 Qubits)	65.80	64.88	65.80	64.95
Hybrid VQFL (6 Qubits)	67.44	66.66	67.44	66.63
Hybrid VQFL (8 Qubits)	81.29	81.23	81.29	81.21
Proposed UVQFL	81.06	81.09	81.06	81.04
Enhanced UVQFL	84.62	84.60	84.62	84.59

5.5 Summary of Experimental Results

The experimental evaluation demonstrated progressive improvements from the baseline Hybrid VQFL models to the proposed Enhanced UVQFL architecture. The enhanced model achieved the best overall classification performance, indicating that uncertainty-aware feature refinement

combined with an improved classical classifier enhances hybrid quantum–classical learning. The broader implications and limitations of these findings are discussed in the following section.

6. Discussion

6.1 Effect of Quantum Feature Dimensionality

The baseline Hybrid Variational Quantum Feature Learning (Hybrid VQFL) experiments demonstrate that the dimensionality of the quantum feature representation has a substantial influence on classification performance. Increasing the number of qubits from four to eight consistently improved the predictive capability of the hybrid quantum–classical model. While the transition from the 4-qubit to the 6-qubit configuration resulted in only a modest improvement, the 8-qubit model produced a considerably stronger performance, indicating that higher-dimensional quantum feature representations capture more discriminative information.

This behaviour suggests that larger variational quantum circuits provide a richer feature space for representing complex image patterns. As the number of qubits increases, the circuit is capable of encoding additional relationships among the extracted features, enabling the classical classifier to separate classes more effectively. Nevertheless, increasing the number of qubits alone does not explicitly address the reliability of the learned quantum representations. The baseline models treat all extracted features equally, regardless of their confidence or uncertainty, which may limit robustness when the feature quality varies across samples.

6.2 Contribution of Uncertainty-Aware Feature Learning

The proposed UVQFL model extends the baseline architecture by introducing uncertainty-aware feature refinement through the Adaptive Feature Confidence Module (AFCM) and the Uncertainty Estimation Module (UEM). Rather than relying solely on the output of the variational quantum circuit, the proposed framework evaluates the reliability of the extracted quantum features before they are forwarded to the classifier.

The experimental results indicate that uncertainty-aware feature learning maintains competitive predictive performance while introducing an additional mechanism for assessing feature reliability. This observation suggests that incorporating confidence estimation into the hybrid learning pipeline enables the model to focus more effectively on informative quantum representations while reducing the influence of uncertain features. Such behaviour is particularly valuable for hybrid quantum–classical models, where the quality of encoded quantum features directly affects downstream classification.

Unlike conventional hybrid quantum learning approaches that optimise only classification accuracy, the proposed framework introduces an additional level of feature assessment without increasing the complexity of the quantum circuit itself. Consequently, improvements are achieved through more effective feature refinement rather than simply enlarging the quantum model.

The proposed UEM employs a deterministic confidence margin derived from the model prediction rather than probabilistic uncertainty estimation techniques such as Monte Carlo Dropout, Bayesian neural networks, or ensemble learning. Consequently, the proposed module is intended as a lightweight confidence-guided feature refinement mechanism that introduces negligible computational overhead while improving feature reliability before quantum encoding. Future work will investigate probabilistic uncertainty estimation within hybrid quantum learning.

6.3 Influence of the Enhanced Classical Classifier

The Enhanced UVQFL model achieved the best overall performance among all evaluated configurations. The primary architectural difference between the proposed UVQFL model and the enhanced version lies in the increased capacity of the classical classifier, while the uncertainty-aware quantum feature extraction process remains unchanged.

These results indicate that the refined quantum features generated by the uncertainty-aware modules contain additional discriminative information that can be exploited more effectively by a deeper classifier. The enhanced classifier provides greater representational capacity, enabling more complex decision boundaries during the final classification stage. The corresponding learning curves presented in the Results section also demonstrate stable optimisation behaviour, suggesting that the increase in model capacity does not introduce noticeable training instability.

The consistent improvement observed across accuracy, precision, recall, and F1-score indicates that the enhanced architecture improves overall predictive performance rather than benefiting only a particular evaluation metric. This balanced improvement suggests that the model generalises well to unseen samples while maintaining reliable class discrimination.

6.4 Overall Interpretation of Experimental Findings

The experimental investigation demonstrates that improvements in hybrid quantum–classical image classification arise from the combined influence of quantum feature representation, uncertainty-aware feature refinement, and classical decision learning. Rather than relying on a single architectural modification, the proposed framework progressively enhances the learning process by integrating complementary components at different stages of feature extraction and classification.

The baseline experiments established that increasing the number of qubits improved the expressive capability of the variational quantum circuit, leading to richer feature representations. Building upon this observation, the proposed UVQFL model introduced uncertainty-aware feature refinement through the Adaptive Feature Confidence Module (AFCM) and the Uncertainty Estimation Module (UEM), enabling the model to assess the reliability of quantum features before classification. Although the predictive performance remained comparable to the strongest baseline configuration, the uncertainty-aware design provides an additional mechanism for improving the quality of the learned feature representations.

Further improvement was achieved by replacing the original classifier with a deeper multi-layer architecture while preserving the same uncertainty-aware quantum feature extraction process. The consistent improvement observed across all evaluation metrics indicates that refined quantum features can be exploited more effectively when coupled with a classifier having greater representational capacity. These findings suggest that uncertainty-aware learning and classical feature modelling complement each other, leading to a more reliable and effective hybrid quantum–classical classification framework.

Although the proposed Enhanced UVQFL model achieved an accuracy of 84.62%, the primary objective of this study was not to surpass state-of-the-art classical deep learning models on the MNIST benchmark. Modern convolutional neural networks routinely achieve accuracies exceeding 99% by exploiting substantially larger feature representations and deeper architectures. In contrast, the proposed framework investigates the feasibility of uncertainty-aware hybrid quantum–classical learning under current Noisy Intermediate-Scale Quantum (NISQ) constraints, where limited qubit resources require aggressive dimensionality reduction prior to quantum encoding.

Overall, the results demonstrate that uncertainty-aware feature refinement and enhanced classical learning complement variational quantum feature extraction within a consistent hybrid quantum–classical framework. All comparative evaluations followed the common experimental protocol described in Section 4, ensuring that the observed performance improvements are attributable to the proposed architectural enhancements. While comparisons with fundamentally different classical deep learning architectures and alternative uncertainty-aware quantum learning methods were beyond the scope of this study, they represent important directions for future investigation.

6.5 Limitations and Future Research

Although the proposed UVQFL framework achieved encouraging classification performance, all experiments in this study were conducted using the PennyLane quantum simulator under ideal (noise-free) quantum execution conditions. While simulation provides a controlled environment for evaluating the proposed uncertainty-aware learning framework, current Noisy Intermediate-Scale Quantum (NISQ) hardware is affected by gate errors, readout errors, decoherence, limited qubit connectivity, and device-specific calibration noise. These hardware imperfections are expected to degrade feature fidelity and classification performance to varying degrees depending on the noise characteristics, circuit depth, qubit connectivity, and hardware calibration. Quantifying this degradation requires dedicated noisy quantum simulations or execution on real NISQ hardware and is therefore left for future investigation. Therefore, the reported results should be interpreted as an evaluation under ideal quantum simulation rather than as a direct estimate of real quantum hardware performance. Future work will investigate the proposed UVQFL framework using realistic noisy quantum simulations incorporating depolarizing, thermal relaxation, and readout error models, followed by experimental validation on available NISQ quantum processors. Such investigations will also enable assessment of whether the proposed Adaptive Feature Confidence Module (AFCM) and Uncertainty Estimation Module (UEM) improve robustness under practical quantum noise conditions.

Second, this study focused on variational quantum circuits containing up to eight qubits. Although larger quantum circuits may provide richer feature representations, increasing the number of qubits also increases the dimensionality of the quantum feature space, the number of trainable quantum operations, and the computational cost associated with quantum circuit evaluation and gradient-based optimisation. Future research will investigate scalable uncertainty-aware quantum architectures employing larger variational quantum circuits beyond the 8-qubit configuration to evaluate scalability, optimisation stability, computational efficiency, and predictive performance under increasing quantum model complexity.

Although larger quantum circuits are expected to increase training time, execution time comparisons were beyond the scope of the present study because the primary objective was to evaluate classification performance under identical experimental conditions.

Future work will also investigate alternative confidence estimation mechanisms, adaptive variational circuit designs, hardware-efficient quantum architectures, and deployment on real NISQ devices to further improve the robustness and practical applicability of uncertainty-aware hybrid quantum–classical learning.

Future work will extend the experimental evaluation by including additional benchmark datasets such as Fashion-MNIST and PathMNIST together with comparisons against representative classical deep learning architectures and emerging uncertainty-aware quantum learning methods. These investigations will provide a broader assessment of the proposed framework under diverse image classification scenarios.

Another limitation of the present study is that the individual contributions of the Adaptive Feature Confidence Module (AFCM) and the Uncertainty Estimation Module (UEM) were not evaluated independently through dedicated ablation experiments. The proposed UVQFL framework was assessed as an integrated architecture to evaluate its overall effectiveness. Future work will perform systematic ablation studies to quantify the individual and combined contributions of AFCM, UEM, and the enhanced classical classifier.

7. Conclusion

This study presented an **Uncertainty-Aware Variational Quantum Feature Learning (UVQFL)** framework for hybrid quantum–classical image classification. The proposed approach integrates the **Adaptive Feature Confidence Module (AFCM)** and the **Uncertainty Estimation Module (UEM)** to improve the reliability of quantum feature representations before the final classification stage. Unlike conventional hybrid quantum learning models that primarily focus on increasing the expressive capability of variational quantum circuits, the proposed framework introduces uncertainty-aware feature refinement as an additional mechanism for enhancing the quality of learned representations.

Experimental evaluation on the MNIST benchmark demonstrated that the proposed framework progressively improves hybrid quantum–classical classification through uncertainty-aware

feature refinement and an enhanced classical classifier while operating within current NISQ constraints.

Overall, the proposed UVQFL framework demonstrates that integrating uncertainty-aware feature refinement with variational quantum feature learning provides a promising direction for developing more reliable hybrid quantum–classical learning systems. Future research will focus on validation using larger quantum circuits, realistic noisy quantum hardware, broader benchmark datasets, and comprehensive ablation studies.

References

- [1] Ayhan, B., & Berens, P. (2022). Test-time data augmentation for improved aleatoric uncertainty estimation in deep learning. **Scientific Reports**, 12(1), 14297. <https://doi.org/10.1038/s41598-022-14297-7>
- [2] Carion, F., Papadimitriou, G., Gewaltig, M., & Eliasmith, C. (2020). Brain-inspired vision systems: A comparative survey. **Frontiers in Neuroscience**, 14, 591405. <https://doi.org/10.3389/fnins.2020.591405>
- [3] De Luca, G. (2021). A survey of NISQ era hybrid quantum-classical machine learning research. **Journal of Artificial Intelligence and Technology**, 2(1), 9–15. <https://doi.org/10.37965/jait.2021.12002>
- [4] Deepika, D., Vanishree, K., & Munishamaiah Krishna. (2025). A hybrid quantum–classical convolutional neural network for enhanced image classification. **International Journal of Computer Applications**, 187(66), 28–34. <https://doi.org/10.5120/ijca2025926113>
- [5] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., et al. (2021). An Image is Worth 16×16 Words: Transformers for Image Recognition at Scale. Proceedings of the **International Conference on Learning Representations (ICLR)**, <https://doi.org/10.48550/arXiv.2010.11929>
- [6] Graziani, M., Andrearczyk, V., & Müller, H. (2024). A unified review of uncertainty quantification in deep learning models for medical image analysis. **Artificial Intelligence in Medicine**, 150, 102830. <https://doi.org/10.1016/j.artmed.2024.102830>
- [7] Huang, H., Earnest-Brown, N., Onodera, T., & Xie, Y. (2024). Exponential separations between learning with and without quantum samples. **Science**, 384(6701), 1327–1332. <https://doi.org/10.1126/science.adl6141>
- [8] Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? **Advances in Neural Information Processing Systems**, 30, 5574–5584.
- [9] Chen, Y. A novel image classification framework based on variational quantum algorithms. *Quantum Inf Process* 23, 362 (2024). <https://doi.org/10.1007/s11128-024-04566-9>
- [10] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. **Nature**, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- [11] Malinin, A., & Rahimi, B. (2020). Uncertainty estimation in one-stage object detection. **IEEE Transactions on Pattern Analysis and Machine Intelligence**, 42(10), 2439–2453. <https://doi.org/10.1109/TPAMI.2020.2996931>
- [12] Preskill, J. (2023). Seeking a quantum advantage for machine learning. **Nature Machine**

- Intelligence**, 5, 813. <https://doi.org/10.1038/s42256-023-00710-9>
- [13] San Sebastian Sein, P., Cañizo, M., & Orús, R. (2025). Image classification with rotation-invariant variational quantum circuits. **Physical Review Research**, 7, 013082. <https://doi.org/10.1103/PhysRevResearch.7.013082>
- [14] Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. **International Conference on Learning Representations (ICLR)**. <https://doi.org/10.48550/arXiv.1409.1556>
- [15] Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. Proceedings of the 36th International Conference on Machine Learning (ICML) (pp. 6105–6114), PMLR.
- [16] Wang, Y., & Liu, J. (2024). A comprehensive review of quantum machine learning: From NISQ to fault tolerance. **Reports on Progress in Physics**, 87, 116402. <https://doi.org/10.1088/1361-6633/ad7f69>
- [17] Senokosov, A., Sedykh, A., Sagingalieva, A., Kyriacou, B., & Melnikov, A. (2024). Quantum machine learning for image classification. **Machine Learning: Science and Technology**, 5(1), 015040, DOI: <https://doi.org/10.1088/2632-2153/ad2aef>
- [18] Liu, A., Wen, C., & Wang, J. (2025). Lean Classical–Quantum Hybrid Neural Network Model for Image Classification. **Advanced Quantum Technologies**, 8(10), 2400703. DOI: <https://doi.org/10.1002/qute.202400703>
- [19] Khan, M. N., et al. (2023). Quantum Machine Learning: A Review and Case Studies. **Entropy**, 25(2), 287, DOI: <https://doi.org/10.3390/e25020287>
- [20] Kim, J., et al. (2025). Quantum Machine Learning: Towards Hybrid Quantum–Classical Vision Models. **Mathematics**, 13(16), 2645, DOI: <https://doi.org/10.3390/math13162645>
- [21] Yu, R., Zhang, X., & Ren, S. (2025). A Review of Quantum Machine Learning Algorithms Based on Variational Quantum Circuit. **Journal of Computer Research and Development**, 62(4), 821–851, DOI: <https://doi.org/10.7544/issn1000-1239.202330979>
- [22] Schuld, M., Bocharov, A., Svore, K. M., & Wiebe, N. (2020). Circuit-centric quantum classifiers. **Physical Review A**, 101(3), 032308. DOI: <https://doi.org/10.1103/PhysRevA.101.032308>
- [23] Pérez-Salinas, A., Cervera-Lierta, A., Gil-Fuster, E., & Latorre, J. I. (2020). Data re-uploading for a universal quantum classifier. **Quantum**, 4, 226. DOI: <https://doi.org/10.22331/q-2020-02-06-226>
- [24] Cong, I., Choi, S., & Lukin, M. D. (2019) Quantum Convolutional Neural Networks. **Nature Physics**, 15(12), 1273–1278, DOI: <https://doi.org/10.1038/s41567-019-0648-8>
- [25] McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. **Nature Communications**, 9, 4812. DOI: <https://doi.org/10.1038/s41467-018-07090-4>
- [26] Sim, S., Johnson, P. D., & Aspuru-Guzik, A. (2019). Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. **Advanced Quantum Technologies**, 2(12), 1900070. DOI: <https://doi.org/10.1002/qute.201900070>
- [27] LeCun, Y., Cortes, C., & Burges, C. J. C. (1998). The MNIST Database of Handwritten Digits. Available at: <http://yann.lecun.com/exdb/mnist/>
- [28] Xiao, H., Rasul, K., & Vollgraf, R. (2017). Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms. arXiv:1708.07747
- [29] Yang, J., Shi, R., Wei, D., et al. (2023). MedMNIST v2: A large-scale lightweight

- benchmark for 2D and 3D biomedical image classification. *Scientific Data*, 10, 41.
- [30] Kingma, D. P., & Ba, J. (2015). Adam: A Method for Stochastic Optimization. ICLR.
- [31] Paszke, A., et al. (2019). PyTorch: An Imperative Style, High-Performance Deep Learning Library. NeurIPS.
- [32] Bergholm, V., et al. (2022). PennyLane: Automatic differentiation of hybrid quantum-classical computations.

Chen, Y. A novel image classification framework based on variational quantum algorithms. *Quantum Inf Process* 23, 362 (2024). <https://doi.org/10.1007/s11128-024-04566-9>

Author Contributions

S.B.S. conceived the research idea, designed the methodology, developed the proposed Uncertainty-Aware Variational Quantum Feature Learning (UVQFL) framework, implemented the algorithms, conducted the experiments, analyzed the results, and prepared the original manuscript draft. **S.V.M.** contributed to the study design, methodology refinement, and critical review of the manuscript. **S.J.** assisted with the literature review, data analysis, and validation of the experimental results. **V.M.** contributed to model evaluation, interpretation of the results, and manuscript revision. **R.R.K.** assisted with software implementation, experimental verification, and performance analysis. **R.S.** contributed to the validation of the methodology, supervision of the research, and critical revision of the manuscript. **S.T.** supervised the overall study, provided conceptual guidance, reviewed the manuscript critically for important intellectual content, and approved the final version for publication. All authors reviewed, revised, and approved the final manuscript.

Data Availability Statement

The datasets used in this study are publicly available. The **MNIST** dataset is available at the MNIST database (<http://yann.lecun.com/exdb/mnist/>). The **Fashion-MNIST** dataset is available at <https://github.com/zalandoresearch/fashion-mnist>. The **PathMNIST** dataset, part of the **MedMNIST v2** benchmark, is publicly available at <https://medmnist.com/>. The processed data and source code used to support the findings of this study are available from the corresponding author upon reasonable request.

Code Availability Statement

The source code developed for the proposed Uncertainty-Aware Variational Quantum Feature Learning (UVQFL) framework is publicly available at GitHub: <https://github.com/shaikarchi/UQVQFL>. The repository contains the implementation of the Hybrid VQFL, UVQFL, and Enhanced UVQFL models, together with the scripts required for data preprocessing, model training, and performance evaluation.

Conflict of Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding Statement

The authors received no financial support for the research, authorship, and/or publication of this article.