

# FAQNAS: FLOPs-aware Hybrid Quantum Neural Architecture Search using Genetic Algorithm

Muhammad Kashif<sup>\*†</sup>, Shaf Khalid<sup>\*†</sup>, Alberto Marchisio<sup>\*†</sup>, Nouhaila Innan<sup>\*†</sup>, Muhammad Shafique<sup>\*†</sup>

<sup>\*</sup> eBrain Lab, Division of Engineering, New York University Abu Dhabi, PO Box 129188, Abu Dhabi, UAE

<sup>†</sup> Center for Quantum and Topological Systems, NYUAD Research Institute, New York University Abu Dhabi, UAE

Emails: {muhammadrashid, sk10741, alberto.marchisio, nouhaila.innan, muhammad.shafique}@nyu.edu

**Abstract**—Hybrid Quantum Neural Networks (HQNNs), which combine parameterized quantum circuits with classical neural layers, are emerging as promising models in the noisy intermediate-scale quantum (NISQ) era. While quantum circuits are not naturally measured in floating point operations (FLOPs), most HQNNs (in NISQ era) are still trained on classical simulators where FLOPs directly dictate runtime and scalability. Hence, FLOPs represent a practical and viable metric to measure the computational complexity of HQNNs. In this work, we introduce FAQNAS, a FLOPs-aware neural architecture search (NAS) framework that formulates HQNN design as a multi-objective optimization problem balancing accuracy and FLOPs. Unlike traditional approaches, FAQNAS explicitly incorporates FLOPs into the optimization objective, enabling the discovery of architectures that achieve strong performance while minimizing computational cost. Experiments on five benchmark datasets (MNIST, Digits, Wine, Breast Cancer, and Iris) show that quantum FLOPs dominate accuracy improvements, while classical FLOPs remain largely fixed. Pareto-optimal solutions reveal that competitive accuracy can often be achieved with significantly reduced computational cost compared to FLOPs-agnostic baselines. Our results establish FLOPs-awareness as a practical criterion for HQNN design in the NISQ era and as a scalable principle for future HQNN systems.

**Index Terms**—Quantum Machine Learning, Hybrid Quantum Neural Networks, FLOPs, Quantum Neural Architecture Search

## I. INTRODUCTION

The rapid advancement of quantum computing has spurred significant interest in Quantum Machine Learning (QML) as a paradigm that may offer computational advantages over classical methods in the near term [1]–[4]. Hybrid Quantum Neural Networks (HQNNs), which integrate parameterized quantum circuits (PQCs) with classical neural layers (Fig. 1), have emerged as promising candidates for leveraging noisy intermediate-scale quantum (NISQ) devices [5]. These architectures combine the representational power of quantum states with the scalability of classical models, making them attractive for practical learning tasks [6]. However, the design of effective HQNN architectures remains a non-trivial challenge. Unlike classical deep learning, where architecture search has been extensively studied [7], [8], architecture design in HQNNs is still relatively underexplored. Existing efforts in quantum architecture search have primarily targeted performance optimization, either by customizing circuit designs for specific devices or by developing strategies to mitigate hardware noise and scalability constraints, with the aim of achieving higher accuracy [9]–[16]. As most of these approaches are performance-oriented, our methodology fills this gap by giving equal importance to predictive accuracy and computational resource efficiency, thereby enabling a more balanced design.

As HQNNs continue to evolve, a central question emerges: how can we design architectures that deliver competitive performance while keeping computational demands manageable? In practice, model selection for HQNNs involves conflicting objectives. On one hand, increasing the number of trainable parameters or circuit depth can improve expressivity and classification performance [17]. On the other hand, deeper circuits and larger parameter spaces are costly as they exacerbate barren plateaus in the optimization landscape and introduce significant computational overhead [18]–[22]. This overhead is particularly critical in the NISQ era, where most HQNNs are trained and benchmarked on classical simulators. In such settings, floating-point operations (FLOPs) directly determine runtime




Fig. 1: An illustration of a Hybrid Quantum Neural Network architecture. Typically Parametrized Quantum Circuits are sandwiched between classical neuron layers with classical optimization loop.

feasibility, memory footprint, and scalability [23]. Even when real quantum hardware is used for circuit execution, the hybrid training loop, including gradient estimation, optimization, and loss evaluation, remains largely classical and FLOPs-intensive. Thus, FLOPs provide a practical and unifying measure of computational cost in HQNNs.

Beyond their current role in simulation, FLOPs-aware principles also have forward-looking relevance. As hybrid workloads increasingly integrate quantum processors with large-scale classical infrastructures, the classical component will remain an indispensable bottleneck, from data encoding to iterative optimization. FLOPs-aware design therefore serves as a hardware-agnostic efficiency metric, guiding architectures that are not only viable for present-day simulators but also scalable for future hybrid quantum-classical deployments. Motivated by these observations, this work investigates the FLOPs-aware search of HQNN architectures under two objectives: minimizing FLOPs and maximizing accuracy. By framing HQNN design as a multi-objective optimization problem, we use genetic algorithm [24], and aim to uncover trade-offs that illuminate the balance between efficiency and performance in hybrid models. To demonstrate generality, we evaluate our approach on five benchmark datasets of varying complexity: Iris and Wine (structured, small-scale datasets), Breast Cancer and Digits (medium-scale datasets), and MNIST (a canonical image classification task)<sup>1</sup>. These datasets allow us to examine how optimal HQNN architectures adapt across problem domains and scales, providing insights into resource-efficient design for both simulators and emerging hardware.

### A. Our Contributions

Our contributions are summarized as follows:

- **Motivation for FLOPs-aware HQNNs:** We argue for the relevance of FLOPs-aware design in HQNNs, showing how FLOPs serve as a practical metric for simulator-based training in the NISQ era and as a guiding principle for scalable quantum-classical integration.
- **FLOPs-aware NAS framework:** We develop a multi-objective neural architecture search (NAS) framework that explicitly incorporates FLOPs into the optimization process, balancing predictive accuracy with computational efficiency for HQNNs.

<sup>1</sup>Although relatively small in scale, these datasets are widely adopted in the QML community as standard benchmarks. The use of larger datasets is generally infeasible due to the prohibitive computational overhead associated with simulating quantum circuits, where the cost grows exponentially with the number of qubits and training iterations.

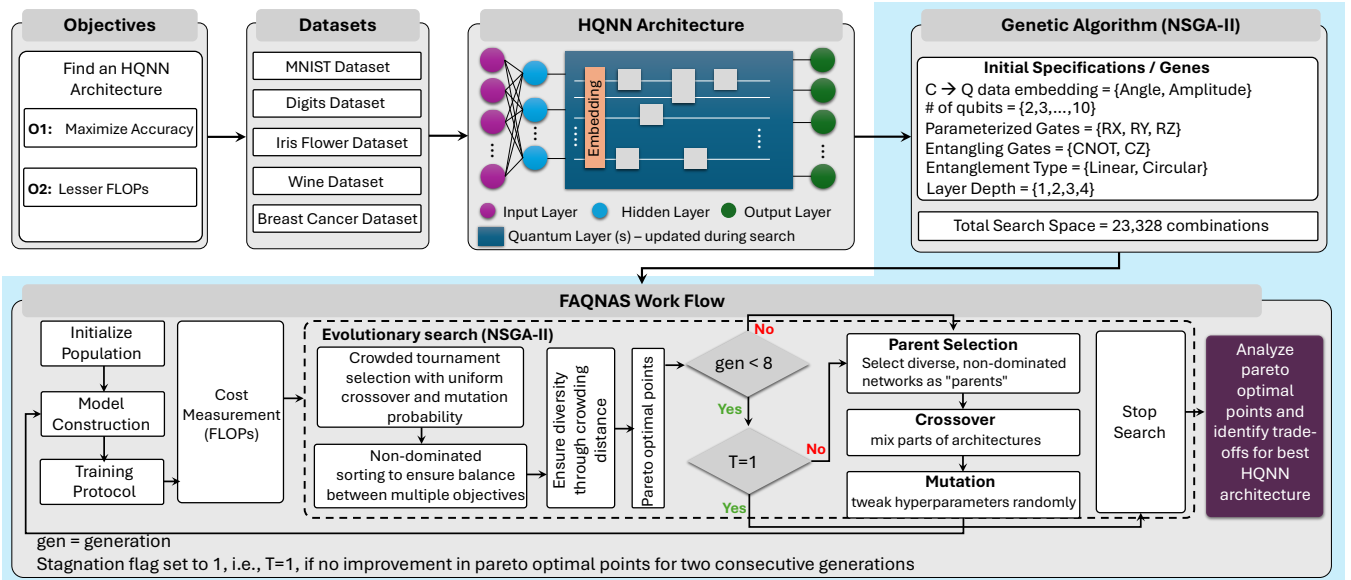


Fig. 2: Our Methodology for FLOPs-aware Hybrid Quantum Neural Network (HQNN) architecture search. Candidate HQNNs are evaluated on accuracy and FLOPs across multiple datasets. The search evolves quantum layer configurations (encoding, qubits, gates, entanglement, depth) from a space of 23,328 combinations using non-dominated sorting, crossover, and mutation. Pareto-optimal architectures are identified to balance performance and efficiency.

- **Dataset-driven evaluation:** We apply our FLOPs-aware HQNN search to five benchmarks (MNIST, Digits, Wine, Breast Cancer, and Iris), demonstrating how FLOPs-aware optimization adapts to dataset complexity and identifies Pareto-optimal architectures.
- **Insights into FLOPs-accuracy trade-offs:** Our analysis highlights that classical FLOPs remain mostly fixed by design, while quantum FLOPs drive accuracy improvements, and shows how FLOPs-awareness helps avoid over-provisioned architectures.
- **Blueprint for resource-efficient HQNNs:** We provide a principled approach for designing hybrid networks that balance accuracy and efficiency, offering practical benefits for simulation-based exploration today and scalable deployment strategies for future quantum-classical computing.

By addressing both performance and resource considerations, our study sheds light on the practical design space of HQNNs. We anticipate that this work will contribute to the growing discourse on making QML models not only accurate but also hardware-efficient and trainable in realistic settings.

## II. OUR METHODOLOGY

### A. Objective

The primary objective of the proposed framework is to systematically explore and identify HQNN architectures that achieve a favorable balance between accuracy and computational efficiency. Specifically, the framework aims to maximize model performance while simultaneously minimizing resource consumption in terms of FLOPs. By incorporating Pareto-based multi-objective optimization, the methodology ensures that the selected architectures are not only accurate but also scalable and practical for deployment in resource-constrained environments.

### B. Datasets

To evaluate the effectiveness and generalizability of the proposed framework, we employed five widely used benchmark datasets that vary in complexity and dimensionality:

#### 1) MNIST Handwritten Digits Dataset

The MNIST dataset [25] is a large-scale benchmark consisting of 70,000 grayscale images of handwritten digits (0–9), each with a resolution of  $28 \times 28$  pixels. It includes 10 output classes, corresponding to the ten digit categories. The models are trained on

60,000 images and tested on 10,000 images. It serves as a standard dataset for image classification tasks and provides a relatively high-dimensional feature space, which is suitable for testing the scalability of HQNN architectures.

#### 2) Digits Dataset

The Digits dataset [26] is a medium-scale benchmark comprising 1,797 grayscale images of handwritten digits (0–9). Each image has a resolution of  $8 \times 8$  pixels, yielding a significantly lower-dimensional feature space compared to MNIST’s  $28 \times 28$  resolution. The dataset includes 10 output classes, corresponding to the digit categories. Unlike MNIST, which contains 70,000 samples and serves as a high-dimensional benchmark for large-scale image classification, the Digits dataset is relatively small in size and dimensionality. This makes it particularly useful for rapid prototyping and low-resource experimentation with HQNN architectures, while MNIST is better suited for evaluating scalability under more complex input spaces.

#### 3) Iris Flower Dataset

The Iris dataset [26] is a classical benchmark in ML comprising 150 samples from three species of Iris flowers (setosa, versicolor, and virginica). Each sample is represented by four numerical features: sepal length, sepal width, petal length, and petal width. Its small size (3 classes) and low dimensionality make it particularly useful for testing the behavior of HQNNs in constrained feature spaces.

#### 4) Wine Dataset

The Wine dataset [26] contains 178 samples of three wine cultivars. Each sample is described by 13 continuous physicochemical features, such as alcohol content, flavanoids, and color intensity. This dataset includes 3 classes, corresponding to the wine cultivars, and represents a medium-sized classification challenge and enables assessment of the framework on structured tabular data.

#### 5) Breast Cancer Dataset

The Breast Cancer Wisconsin Diagnostic dataset [26] consists of 569 samples describing cell nuclei features computed from digitized images of breast tissue. Each sample is characterized by 30 real-valued features, which capture properties such as radius, texture, perimeter, area, and smoothness of the cell nuclei. The dataset is divided into two output classes: malignant and benign tumors. It is a well-established benchmark for binary classification tasks in biomedical ML. The




Fig. 3: An Illustration of Entanglement Topology used in this paper. (a) Linear Topology (b) Circular Topology.

relatively small number of samples, combined with a moderate feature dimension, makes it a useful dataset for testing the generalization ability of HQNNs under constrained data availability.

Collectively, these datasets cover a diverse range of modalities, images, low-dimensional tabular data, and medium-dimensional physicochemical data, thereby providing a comprehensive testbed for evaluating the trade-offs between accuracy and resource efficiency in HQNN architectures.

### C. HQNN Architecture

The HQNN architecture that we consider has three components (as shown in Fig. 1), details of which are presented below:

- 1) **Classical Pre-processing:** The classical front-end consists of a fully connected (linear) layer that prepares raw input data for quantum processing. Given an input sample represented as a feature vector of dimensionality  $d$ , this layer projects it onto an  $n_{\text{qubits}}$ -dimensional feature space through multiplication with a weight matrix of size  $d \times n_{\text{qubits}}$ . The resulting vector of length  $n_{\text{qubits}}$  serves as the embedding that is fed into the quantum layer. This linear transformation serves a dual purpose: dimensionality reduction for high-dimensional inputs and feature expansion for low-dimensional inputs, thereby aligning the data representation with the quantum circuit size.
- 2) **Quantum Layers:** The quantum layer forms the core of the hybrid architecture and is implemented using PennyLane as a QNode wrapped as a PyTorch layer<sup>2</sup>, enabling seamless integration with classical optimization routines. The input vector from the front-end is embedded into the quantum state using one of two encoding strategies: angle embedding or amplitude embedding (Eq. (1) and Eq. (2), respectively), as selected by the Genetic Algorithm (GA). The PQC then applies up to a configurable number of layers of parameterized single-qubit rotations. For each qubit and each layer, the GA selects the gate type from the set  $\{R_X(\theta), R_Y(\theta), R_Z(\theta)\}$ . The rotation angles are trainable parameters organized in a tensor of shape  $(\text{max\_layers}, n_{\text{qubits}})$ .

$$|x\rangle = \bigotimes_{i=1}^N \cos(x_i) |0\rangle + \sin(x_i) |1\rangle \quad (1)$$

$$\begin{pmatrix} x_1 \\ \vdots \\ x_{2^n} \end{pmatrix} \leftrightarrow |\psi_x\rangle = \sum_{j=1}^{2^n} x_j |j\rangle \quad (2)$$

Entanglement is introduced to capture correlations across qubits. Depending on the GA configuration, one or more rounds of entangling operations are applied in either a linear or circular topology (Fig. 3), using two-qubit gates chosen from  $\{CNOT, CZ\}$ .

Finally, measurement is performed in the computational basis: the expectation value of the Pauli-Z operator is obtained for each qubit, resulting in an  $n_{\text{qubits}}$ -dimensional feature vector. This vector represents the quantum-transformed embedding and is passed to the classical post-processing stage.

- 3) **Classical Post-Processing:** The classical post-processing module is responsible for mapping the quantum-transformed feature vector to the final prediction space. A fully connected layer projects the  $n_{\text{qubits}}$ -dimensional output into a vector of logits corresponding to the number of classes in the target dataset. A log-softmax activation is applied to convert the logits into log-probabilities, which are then used for training and evaluation through a negative log-likelihood (NLL) loss.

### D. Genetic Algorithm

Genetic algorithms (GAs) are powerful evolutionary optimization techniques inspired by natural selection and survival of the fittest. Among them, the Non-dominated Sorting Genetic Algorithm II (NSGA-II) has emerged as one of the most widely adopted approaches for solving multi-objective optimization problems [27]. NSGA-II efficiently balances convergence toward the Pareto-optimal front with diversity preservation using mechanisms such as fast non-dominated sorting, elitism, and crowding distance. Its computational efficiency and robustness make it particularly suitable for complex, real-world applications where trade-offs among conflicting objectives must be carefully managed. By maintaining a diverse set of high-quality solutions, NSGA-II provides researchers and practitioners with deeper insights into the optimization landscape.

TABLE I: Search space specifications for genetic algorithm encoding

Parameter	Values
Number of Qubits ( $n_{\text{qubits}}$ )	{2, 3, 4, 5, 6, 7, 8, 9, 10}
Embedding Method	{Angle, Amplitude}
Parameterized Rotation Gates	$\{R_X(\theta), R_Y(\theta), R_Z(\theta)\}$
Entangling Gate	{CNOT, CZ}
Entanglement Topology	{Linear, Circular}
Number of Layers	{1, 2, 3, 4}

#### 1) Initial Specifications

In the proposed FAQNAS framework, each individual in the population is encoded as a chromosome where the genes correspond to quantum circuit hyperparameters. The search space is defined over several quantum-specific parameters that form the genetic material. The number of qubits is varied from 2 to 10, enabling exploration of scalability across circuit widths. For classical to quantum data embedding, two schemes are considered: angle embedding and amplitude embedding, which encode classical data into quantum states with different resource requirements.

Entangling operations are introduced by both gate type, either CNOT or CZ, and entanglement topology, chosen from linear or circular connectivity. The number of entangling layers is allowed to vary from 1 to 4, thereby controlling circuit depth and entanglement capacity. Additionally, parameterized rotation gates are randomly drawn from the set  $\{R_X(\theta), R_Y(\theta), R_Z(\theta)\}$ . Together, these parameters form the genes of the evolutionary search, where different combinations represent candidate architectures subject to optimization. The combination of these parameters yields a total search space of 23,328 unique quantum circuit configurations. This structured representation ensures that the genetic algorithm can systematically evolve candidate HQNN architectures by recombining and mutating different parameter sets, while NSGA-II guides the search toward Pareto-optimal solutions with respect to accuracy and FLOPs. A summary of the specifications that govern the GA search space is presented in Table I.

#### 2) FAQNAS Workflow

The evolutionary process for identifying optimal HQNN architectures follows the Non-dominated Sorting Genetic Algorithm II (NSGA-II), the complete pseudocode of FAQNAS is presented below followed by a step-by-step explanation:

<sup>2</sup>[https://pennylane.ai/qml/demos/tutorial\\_qnn\\_module\\_torch](https://pennylane.ai/qml/demos/tutorial_qnn_module_torch)

## FLOPs-Aware NSGA-II for HQNN Architecture Search

```

1: Inputs: pop size  $P$ , generations  $G$ ,  $p_c$ ,  $p_m$ , stagnation  $T$ , search
   space  $\mathcal{S}$ 
2: Objectives:  $f(\mathbf{c}) = (F^Q, 1 - \text{Acc}_{\text{val}}, \text{Params})$ , with  $F^Q =
   \max\{F^{\text{Total}} - F^{\text{Classical}}, 0\}$ 
3:  $\mathcal{P} \leftarrow \text{INITIALIZEPOPULATION}(P, \mathcal{S})$ 
4: for all  $\mathbf{c} \in \mathcal{P}$  do
5:    $\mathbf{c}.f \leftarrow \text{EVALUATE}(\mathbf{c})$ 
6: end for
7:  $\mathcal{F} \leftarrow \text{FASTNONDOMINATEDSORT}(\mathcal{P})$ ;
    $\text{COMPUTECROWDINGDISTANCE}(F), \forall F \in \mathcal{F}$ 
8:  $\mathcal{F}_{\text{prev}} \leftarrow \text{rank-0 objective tuples}$ ;  $\text{stagn} \leftarrow 0$ 
9: for  $g = 1$  to  $G$  do
10:   $\mathcal{O} \leftarrow \emptyset$ 
11:  for  $i = 1$  to  $P$  do
12:     $p_1 \leftarrow \text{CROWDEDTOURNAMENTSELECT}(\mathcal{P})$ ;  $p_2 \leftarrow
    \text{CROWDEDTOURNAMENTSELECT}(\mathcal{P})$ 
13:     $\mathbf{c} \leftarrow \text{UNIFORMCROSSOVER}(p_1.\mathbf{c}, p_2.\mathbf{c})$  w.p.  $p_c$  else  $p_1.\mathbf{c}$ 
14:     $\mathbf{c} \leftarrow \text{POINTWISEMUTATE}(\mathbf{c}, p_m, \mathcal{S})$ 
15:     $\mathbf{c}.f \leftarrow \text{EVALUATE}(\mathbf{c})$ ;  $\mathcal{O} \leftarrow \mathcal{O} \cup \{\mathbf{c}\}$ 
16:  end for
17:   $\mathcal{U} \leftarrow \mathcal{P} \cup \mathcal{O}$ ;  $\mathcal{F} \leftarrow \text{FASTNONDOMINATEDSORT}(\mathcal{U})$ 
18:   $\mathcal{P}_{\text{new}} \leftarrow \emptyset$ 
19:  for all  $F \in \mathcal{F}$  do
20:     $\text{COMPUTECROWDINGDISTANCE}(F)$ 
21:    if  $|\mathcal{P}_{\text{new}}| + |F| \leq P$  then
22:       $\mathcal{P}_{\text{new}} \leftarrow \mathcal{P}_{\text{new}} \cup F$ 
23:    else
24:       $\text{SORTBYCROWDINGDISTANCEDESC}(F)$ ; take first
        $P - |\mathcal{P}_{\text{new}}|$ ; break
25:    end if
26:  end for
27:   $\mathcal{P} \leftarrow \mathcal{P}_{\text{new}}$ 
28:   $\mathcal{F}^* \leftarrow \text{rank-0 tuples of } \mathcal{P}$ ;  $\text{stagn} \leftarrow \text{stagn} + 1$  if  $\mathcal{F}^* = \mathcal{F}_{\text{prev}}^*$ 
else 0;  $\mathcal{F}_{\text{prev}}^* \leftarrow \mathcal{F}^*$ 
29:  if  $\text{stagn} \geq T$  then
30:    break
31:  end if
32: end for
33: Return: rank-0 Pareto set; Top- $K$ ; optional single compromise
   (min  $\ell_2$  after max-norm)

```

### a) Initialize Population

To start the search, an initial population is created, where each genome  $\mathbf{c}$  independently samples each gene from Table I to produce a diverse starting pool. As the search progresses, all the candidate architectures are defined by a unique set of hyperparameters (embedding type, number of qubits, gate choices, entanglement structure, and layer depth) sampled from the search space.

### b) Model Construction

For a given  $\mathbf{c}$ , we build a three-stage HQNN similar to that shown in Fig. 1 and explained in Section II-C: (i) a classical linear (pre-processing) layer that flattens the input (e.g.,  $28 \times 28$  for MNIST) and projects to  $\mathbb{R}^{n_{\text{qubits}}}$ , (ii) a quantum module with the chosen feature map (Angle Embedding or Amplitude Embedding; for the latter we zero-pad and normalize to length  $2^{n_{\text{qubits}}}$ ),  $L$  layers of per-wire single-qubit rotations with the specified gate types, and  $n_{\text{ent}}$  entangling layers under the selected topology; measurements are  $Z$ -expectations per qubit, and (iii) a linear classical (postprocessing) layer mapping  $\mathbb{R}^{n_{\text{qubits}}} \rightarrow \mathbb{R}^n$  (where  $n$  is number of classes in the dataset) followed by a negative log-likelihood (NLL) loss.

### c) Training Protocol for Candidate Architectures

Each architecture in the population is trained on the selected datasets to evaluate its learning ability and generalization performance. We use Adam with a fixed learning rate (e.g.,  $10^{-3}$ ), batch size 64, and train each candidate 5 epochs. We report the best validation

accuracy ( $\text{Acc}_{\text{val}}$ ) over epochs, to reduce variance from early-epoch noise. Random seeds are fixed for reproducibility.

### d) Cost measurement (FLOPs)

To make FLOPs comparable across candidates, we measure *one-sample* forward+backward cost with `PyTorch profiler`<sup>3</sup>. For each genome, (i) We run the full hybrid model to obtain total number FLOPs ( $F^{\text{Total}}$ ) and (ii) run a classical-only baseline (pre and post-processing classical layers) to obtain the FLOPs consumption of classical layers  $F^{\text{Classical}}$ . The quantum contribution is  $F^Q = \max\{F^{\text{Total}} - F^{\text{Classical}}, 0\}$ . This isolates the incremental cost of the quantum layer while holding classical work constant except for the width of hidden layer coupling to  $n_{\text{qubits}}$ .

### e) Multi-objective formulation

The trained networks are then assessed with respect to the optimization objectives: maximizing classification accuracy and minimizing FLOPs. We minimize the vector  $(F^Q, 1 - \text{Acc}_{\text{val}}, \text{Params})$ . The first term embodies the FLOPs-aware design principle and the second preserves predictive performance.

### f) Evolutionary search (NSGA-II)

- We initialize  $P$  genomes by uniform sampling of search space  $\mathcal{S}$  and evaluate their objectives.
- Each generation uses crowded tournament selection (preferring lower rank, then larger crowding distance), uniform crossover with probability  $p_c = 0.8$ , and per-gene pointwise mutation with probability  $p_m = 0.2$  drawn from the gene's domain.
- Offspring are evaluated and merged with parents; we apply fast non-dominated sorting where the candidate solutions are ranked based on Pareto dominance to balance the trade-offs between multiple objectives. This ensures that solutions closer to the Pareto front are favored.
- Then the truncation is applied based on crowding-distance to retain  $P$  elites. The crowding distance measure preserves the diversity among non-dominated solutions and avoid premature convergence to a narrow set of architectures.
- We terminate after  $G = 8$  generations or early if the rank-0 front remains unchanged for  $T = 2$  consecutive generations.

## III. RESULTS AND DISCUSSION

We evaluated the proposed FLOPs-aware HQNN architecture search on five benchmark datasets: MNIST, Wine, Breast Cancer, Iris, and Digits. In all cases, the input and output layers were fixed, which are dependent on input size and number of classes in the dataset, respectively. The neurons in the first and only hidden layer scales with the number of qubits in the PQC. This ensures that the majority of variability explored during search was due to the quantum circuits, making FLOPs-awareness a central criterion in guiding efficient architecture discovery. All candidate architectures were trained for 5 epochs using the Adam optimizer with a learning rate of 0.003, and a batch size of 64 across all datasets. Figures 4-8 present the results, where accuracy is plotted against classical, quantum, and total FLOPs. Candidate architectures are shown in purple, while Pareto-optimal points are highlighted in gold.

### A. Results on MNIST Dataset

For the MNIST dataset, a clear dependence on computational cost is observed, as shown in Fig. 4. Architectures with low FLOPs, particularly on the classical side, achieve accuracies below 50%, but performance improves steadily with increasing computational budgets. Accuracy rises above 80% once classical FLOPs exceed  $10^4$ , with further gains obtained primarily through quantum FLOPs. The

<sup>3</sup><https://docs.pytorch.org/docs/stable/profiler.html>




Fig. 4: Accuracy versus computational cost for candidate HQNN architectures on the MNIST dataset. Each panel shows accuracy as a function of (left) classical FLOPs, (middle) quantum FLOPs, and (right) total FLOPs. Purple dots represent all candidate architectures while gold stars denote Pareto-optimal solutions.




Fig. 5: Accuracy versus computational cost for candidate HQNN architectures on the Digits dataset. Each panel shows accuracy as a function of (left) classical FLOPs, (middle) quantum FLOPs, and (right) total FLOPs. Purple dots represent all candidate architectures while gold stars denote Pareto-optimal solutions.

quantum FLOPs plot reveals a sharper trend, where accuracy increases significantly around  $10^4$  FLOPs and plateaus near 90% beyond  $10^5$  FLOPs. The total FLOPs plot closely follows the quantum FLOPs trend, underscoring that accuracy gains in MNIST are driven primarily by quantum circuit complexity. The Pareto-optimal points highlight that high accuracy can be obtained without resorting to the most resource-intensive solutions, confirming the benefit of FLOPs-aware optimization in balancing accuracy and efficiency.

#### B. Results on Digits Dataset

The Digits dataset presents a distinct behavior compared to the MNIST, as shown in Fig. 5. Accuracy improves rapidly with small increases in quantum FLOPs, surpassing 90% at budgets of  $10^3 - 10^4$ . Beyond this point, performance quickly saturates near 100%, and larger circuits provide no additional benefit. The classical FLOPs distribution shows Pareto-optimal points at both low and high FLOPs, but with no strong dependency between accuracy and cost. The total FLOPs plot is sharply defined, with a clear Pareto front demonstrating that many architectures achieve maximal accuracy at substantially lower FLOPs. These results indicate that for moderately complex datasets like Digits, FLOPs-aware search is highly effective at identifying compact and efficient HQNN architectures without compromising performance.

#### C. Results on Iris Dataset

For the Iris dataset, accuracy outcomes vary widely, ranging from 40% to above 90% across candidate architectures, as presented in Fig. 6. Interestingly, the classical FLOPs plot shows that several Pareto-optimal architectures achieving high accuracy are found at very low computational budgets, suggesting that classical overhead is not a limiting factor for this dataset. The quantum FLOPs results reveal that strong performance can be achieved at both low and high FLOPs, although many mid-range FLOPs architectures perform inconsistently. The total FLOPs distribution reinforces the importance of careful search: FLOPs-aware optimization identifies efficient solutions that avoid wasted resources while still delivering high accuracy. This variability illustrates that, for simple datasets such as Iris, careful

selection of quantum circuits plays a disproportionately important role in determining performance.

#### D. Results on Wine Dataset

On the Wine dataset, the trade-off between FLOPs and accuracy is not very straightforward (Fig. 7). Candidate architectures span a broad accuracy range from 30% to nearly 100%, and high-accuracy models emerge even at modest computational budgets. The classical FLOPs plot shows that efficient solutions exist with relatively low costs, while the quantum FLOPs plot reveals a smoother trend where accuracy improves gradually but not monotonically with increased FLOPs. The total FLOPs distribution confirms that excessive computational budgets are not strictly necessary, as several Pareto-optimal solutions achieve  $>90\%$  accuracy while incurring far fewer FLOPs than the most complex architectures. These results indicate that for smaller datasets like Wine, FLOPs-aware search is critical for uncovering efficient solutions that avoid over-provisioned quantum circuits.

#### E. Results on Breast Cancer Dataset

The Breast Cancer dataset exhibits a different trend, shown in Fig. 8. Many candidate architectures already achieve strong performance, with accuracies above 70% even at low FLOPs. Both the quantum and total FLOPs plots reveal a dense cluster of solutions achieving 90-95% accuracy across a wide FLOPs range. Classical FLOPs remain relatively uninformative, as even minimal classical cost yields strong performance. The Pareto-optimal front shows that high accuracy can be sustained at low or moderate FLOPs, reflecting the dataset's relatively low complexity. In this case, FLOPs-aware search demonstrates its strength in identifying minimal architectures that retain high accuracy, confirming that more complex solutions are unnecessary.

#### F. Comparative Analysis

A comparative analysis across datasets highlights several consistent trends. Accuracy generally increases with FLOPs but saturates after a dataset-dependent threshold. Large-scale datasets such as MNIST require higher FLOPs budgets to reach strong performance, while moderate datasets like Digits achieve near-maximum accuracy with




Fig. 6: Accuracy versus computational cost for candidate HQNN architectures on the Iris dataset. Each panel shows accuracy as a function of (left) classical FLOPs, (middle) quantum FLOPs, and (right) total FLOPs. Purple dots represent all candidate architectures while gold stars denote Pareto-optimal solutions.




Fig. 7: Accuracy versus computational cost for candidate HQNN architectures on the Wine dataset. Each panel shows accuracy as a function of (left) classical FLOPs, (middle) quantum FLOPs, and (right) total FLOPs. Purple dots represent all candidate architectures while gold stars denote Pareto-optimal solutions.




Fig. 8: Accuracy versus computational cost for candidate HQNN architectures on the Breast Cancer dataset. Each panel shows accuracy as a function of (left) classical FLOPs, (middle) quantum FLOPs, and (right) total FLOPs. Purple dots represent all candidate architectures; gold stars denote Pareto-optimal solutions.

relatively few FLOPs. Smaller and simpler datasets, including Wine, Breast Cancer, and Iris, frequently yields efficient solutions that achieve high accuracy at surprisingly low FLOPs, though with more variability in performance among candidate architectures. In all cases, improvements in accuracy are primarily driven by quantum FLOPs, since classical FLOPs remain largely fixed due to the constrained design of classical layers. Overall, these results demonstrate that FLOPs-aware HQNN search consistently identifies architectures that balance accuracy and computational efficiency. The Pareto fronts across datasets confirm that strong performance can be obtained without over-provisioning FLOPs, with dataset complexity dictating the threshold at which additional FLOPs cease to provide accuracy gains. The findings underscore the importance of incorporating FLOPs-awareness into HQNN search, as it not only enables scalable optimization for large datasets but also prevents wasted computational resources for simpler ones.

#### IV. CONCLUSION

We presented a FLOPs-aware NAS framework for HQNNs, addressing the challenge of designing models that are both accurate and computationally efficient. By explicitly considering FLOPs alongside

accuracy, our approach uncovers Pareto-optimal architectures that balance performance with resource cost. Results across five datasets highlight that accuracy improvements are primarily driven by quantum FLOPs, while classical FLOPs remain fixed and less influential in a hybrid setting. Importantly, FLOPs-aware analysis identifies dataset-specific thresholds beyond which additional FLOPs yield no further accuracy gains, preventing over-provisioned designs. This reveals a key intuition about HQNNs: their efficiency depends not on maximizing circuit size, but on identifying the smallest quantum circuit that reaches the accuracy plateau for a given problem. These findings underscore that FLOPs-aware design is essential for NISQ-era HQNNs, providing immediate benefits for simulator-based exploration and forward-looking value for scalable quantum-classical computing.

#### ACKNOWLEDGMENT

This work was supported in part by the NYUAD Center for Quantum and Topological Systems (CQTS), funded by Tamkeen under the NYUAD Research Institute grant CG008.

## REFERENCES

- [1] M. Schuld, I. Sinayskiy, and F. Petruccione, "An introduction to quantum machine learning," *Contemporary Physics*, vol. 56, no. 2, p. 172–185, Oct. 2014.
- [2] M. Kashif and M. Shafique, "Position paper: Quantum neural networks - a paradigm shift in ai or a theoretical promise?" in *2025 International Joint Conference on Neural Networks (IJCNN)*, 2025, pp. 1–10.
- [3] K. Zaman, A. Marchisio, M. A. Hanif, and M. Shafique, "A survey on quantum machine learning: Current trends, challenges, opportunities, and the road ahead," *arXiv preprint arXiv:2310.10315*, 2023.
- [4] N. Innan, M. Kashif, A. Marchisio, M. Bennai, and M. Shafique, "Next-generation quantum neural networks: Enhancing efficiency, security, and privacy," in *2025 IEEE 31st International Symposium on On-Line Testing and Robust System Design (IOLTS)*. IEEE, 2025, pp. 1–4.
- [5] M. Kashif and S. Al-Kuwari, "Design space exploration of hybrid quantum-classical neural networks," *Electronics*, vol. 10, no. 23, p. 2980, 2021.
- [6] M. Kashif and S. Al-Kuwari, "Demonstrating quantum advantage in hybrid quantum neural networks for model capacity," in *2022 IEEE International Conference on Rebooting Computing (ICRC)*, 2022, pp. 36–44.
- [7] S. P. Avval *et al.*, "Systematic review on neural architecture search," *Artificial Intelligence Review*, vol. 58, no. 3, p. 73, 2025.
- [8] I. Salehin *et al.*, "Automl: A systematic review on automated machine learning with neural architecture search," *Journal of Information and Intelligence*, vol. 2, no. 1, pp. 52–81, 2024.
- [9] Y. Du, T. Huang, S. You, M.-H. Hsieh, and D. Tao, "Quantum circuit architecture search for variational quantum algorithms," *npj Quantum Information*, vol. 8, no. 1, p. 62, 2022.
- [10] Y. Li, R. Liu, X. Hao, R. Shang, P. Zhao, and L. Jiao, "Eqnas: Evolutionary quantum neural architecture search for image classification," *Neural Networks*, vol. 168, pp. 471–483, 2023.
- [11] Y. Li, G. Liu, P. Zhao, R. Shang, and L. Jiao, "Balanced quantum neural architecture search," *Neurocomputing*, vol. 594, p. 127860, 2024.
- [12] M. Kashif and M. Shafique, "Nrqnn: The role of observable selection in noise-resilient quantum neural networks," in *World Congress in Computer Science, Computer Engineering & Applied Computing*. Springer, 2024, pp. 116–131.
- [13] M. Kashif and M. Shafique, "Hqnet: Harnessing quantum noise for effective training of quantum neural networks in nisq era," *arXiv preprint arXiv:2402.08475*, 2024.
- [14] S. Dutta, N. Innan, S. B. Yahia, and M. Shafique, "QAS-QTNs: Curriculum reinforcement learning-driven quantum architecture search for quantum tensor networks," in *2025 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1. IEEE, 2025, pp. 1739–1747.
- [15] T. Ahmed, M. Kashif, A. Marchisio, and M. Shafique, "A comparative analysis and noise robustness evaluation in quantum neural networks," *Scientific Reports*, vol. 15, no. 1, p. 33654, 2025.
- [16] A. Marchisio, E. Sychiuco, M. Kashif, and M. Shafique, "Cutting is all you need: Execution of large-scale quantum neural networks on limited-qubit devices," *arXiv preprint arXiv:2412.04844*, 2024.
- [17] M. Kashif and S. Al-Kuwari, "The unified effect of data encoding, ansatz expressibility and entanglement on the trainability of hqnns," *International Journal of Parallel, Emergent and Distributed Systems*, vol. 38, no. 5, pp. 362–400, 2023. [Online]. Available: <https://doi.org/10.1080/17445760.2023.2231163>
- [18] J. R. McClean *et al.*, "Barren plateaus in quantum neural network training landscapes," *Nature Communications*, vol. 9, no. 1, nov 2018. [Online]. Available: <https://doi.org/10.1038/s41467-018-07090-4>
- [19] M. Kashif and S. Al-Kuwari, "The impact of cost function globality and locality in hybrid quantum neural networks on nisq devices," *Machine Learning: Science and Technology*, vol. 4, no. 1, p. 015004, 2023.
- [20] M. Kashif and M. Shafique, "Deep quantum neural networks with enhanced trainability and gradient propagation," *Scientific Reports*, no. 15, 2025.
- [21] M. Kashif and S. Al-Kuwari, "Resqnets: a residual approach for mitigating barren plateaus in quantum neural networks," *EPJ Quantum Technology*, vol. 11, no. 1, Jan. 2024. [Online]. Available: <http://dx.doi.org/10.1140/epjqt/s40507-023-00216-8>
- [22] M. Kashif, M. Rashid, S. Al-Kuwari, and M. Shafique, "Alleviating barren plateaus in parameterized quantum machine learning circuits: Investigating advanced parameter initialization strategies," in *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2024, pp. 1–6.
- [23] M. Kashif, A. Marchisio, and M. Shafique, "Computational advantage in hybrid quantum neural networks: Myth or reality?" in *2025 62nd ACM/IEEE Design Automation Conference (DAC)*, 2025, pp. 1–7.
- [24] S. Forrest, "Genetic algorithms," *ACM computing surveys (CSUR)*, vol. 28, no. 1, pp. 77–80, 1996.
- [25] Y. LeCun, C. Cortes, and C. Burges, "Mnist handwritten digit database," *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, vol. 2, 2010.
- [26] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [27] H. Ma, Y. Zhang, S. Sun, T. Liu, and Y. Shan, "A comprehensive survey on nsga-ii for multi-objective optimization and applications," *Artificial Intelligence Review*, vol. 56, no. 12, pp. 15 217–15 270, 2023.