EVOLVING HYBRID QUANTUM–CLASSICAL ARCHITECTURES FOR IMAGE CLASSIFICATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Hybrid quantum classical neural networks integrate parameterized quantum circuits (PQCs) with established deep learning architectures, but their performance depends strongly on the choice of quantum circuit architecture, a choice that remains largely manual. Most existing approaches rely on hand-designed or fixed circuit ansätze, requiring circuit structure, gate composition, and qubit connectivity to be specified in advance with no guarantee that they suit the task. This limitation is especially acute in image classification, where quantum circuits must transform features extracted by classical networks while remaining compact enough for practical training, requirements that generic, task-agnostic ansätze are unlikely to satisfy simultaneously. We extend EXAQC, an evolutionary framework for automated quantum circuit discovery, to image classification. EXAQC evolves PQCs as intermediate processing modules while retaining classical feature-extraction and prediction layers. On MNIST, Fashion-MNIST, and CIFAR-10, EXAQC achieves 98.42%, 90.62%, and 85.47% accuracy, respectively, while using comparable gate counts to other quantum architecture-search methods. Against classical networks, evolved hybrid models maintain comparable accuracy with substantially fewer trainable parameters, reaching 85.68% on CIFAR-10 with over 25$\times$ fewer parameters than a 10-layer CNN. Encoding choice also matters: rotation-based encodings (RX, RY, U3) outperform amplitude encoding by 22-25 points on CIFAR-10. These results demonstrate that automated circuit discovery yields compact quantum modules that can replace larger classical components in vision architectures while retaining competitive accuracy.

1 INTRODUCTION

Parameterized quantum circuits (PQCs) are widely studied for quantum machine learning (QML), where classical data is encoded into quantum states and processed by trainable quantum operations Havlíček et al. (2019); Schuld et al. (2020); Pérez-Salinas et al. (2020); Wang et al. (2024); Ding et al. (2024); Tudisco et al. (2026). However, the effectiveness of the concerned circuit depends strongly on architectural choices such as gate composition, circuit depth, connectivity, and data encoding Sim et al. (2019); Hubregtsen et al. (2021); Benítez-Buenache & Portell-Montserrat (2025); Liu et al. (2025). Poor circuit design can also produce barren plateaus, where gradients vanish exponentially with qubit count McClean et al. (2018); Cerezo et al. (2021), stalling training and rendering even expressive circuits effectively untrainable. Architecture choice therefore determines whether a hybrid model can learn at all. Selecting an effective PQC architecture therefore remains an important challenge. The difficulty of manually selecting circuits that are both well-suited to the task and trainable has motivated quantum architecture search (QAS), which automatically optimizes circuit structure through approaches including differentiable, predictor-based, and evolutionary search Zhang & Deng (2020); Wu et al. (2023); Zhang & Deng (2021); Ding & Spector (2022); Zhang & Zhao (2023); Kar et al. (2026b).

Evolutionary circuit search has shown promise on low-dimensional benchmarks Kar et al. (2026a), but practical machine learning rarely operates at that scale. Image classification offers a more demanding test: MNIST and Fashion-MNIST contain 784 input values per image, and CIFAR-10 contains 3,072, far more than the handful of qubits available on near-term hardware. Hybrid architectures bridge this gap by using classical feature extractors to compress images into compact representations for quantum processing Wang et al. (2022); Liu et al. (2021). As a result, a harder

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

dataset does not require a larger quantum circuit, since the circuit's input can stay small even as the source images grow more complex. The question is therefore not whether evolved circuits can process thousands of raw pixels, but whether they can still find useful structure in compressed features as the underlying task becomes harder. Due to space constraints, a detailed discussion of related work on hybrid quantum classical image classification, quantum architecture search, and evolutionary circuit design is provided in Appendix A.

This work extends EXAQC (Kar et al., 2026b)[1] to image classification, introducing hybrid quantum circuits using encoders to extract image features. This works also uses U3 gates as a quantum encoding mechanism unlike previous works (Refer Appendix A). Because circuit structure is discovered rather than fixed in advance, the evolved circuits can be tailored to the learned features and kept compact, avoiding both mismatched ansätze and the trainability problems of larger circuits. We investigate both scalability and whether evolved quantum modules can replace larger classical components while maintaining competitive predictive performance with fewer trainable parameters. Specifically, we address the following research questions:

- **RQ1: Scalability and Hybrid Effectiveness.** Can EXAQC scale from low-dimensional learning problems to complex image-classification tasks by discovering useful quantum transformations of learned image representations within a hybrid quantum classical framework? *Yes: EXAQC reaches 98.42%, 90.62%, and 85.47% on the full 10-class MNIST, Fashion-MNIST, and CIFAR-10 tasks, exceeding the compared architecture-search methods on CIFAR-10 while using fewer gates than EQNAS and AQEA-QAS.*
- **RQ2: Representation Dimensionality.** How does increasing the dimensionality of the quantum circuit inputs affect predictive performance, circuit complexity, convergence, and computational requirements? *The quantum representation clearly affects performance and circuit size, since CIFAR-10 required more qubits (8-10), larger circuits (42-60 gates), and a larger search budget.*
- **RQ3: Quantum Encoding Effectiveness.** How does the choice of quantum input encoding (amplitude, RX, RY, or U3) affect the accuracy and complexity of evolved circuits, and does this effect depend on the dataset? *The effect is strongly dataset-dependent: encodings differ by under one point on MNIST, whereas on CIFAR-10 rotation-based encodings (83.75-86.68%) far outperform amplitude encoding (61.38%), with the best-performing U3 also producing the largest circuits.*
- **RQ4: Parameter Efficiency.** Do evolved quantum components maintain competitive classification performance relative to larger classical counterparts, despite having fewer trainable parameters? *Yes: EXAQC matches the 10-layer CNN on CIFAR-10 (85.68% vs. 85.50%) with 96.1% fewer parameters and exceeds the 1-layer CNN on Fashion-MNIST (91.11% vs. 89.47%) with 54.8% fewer parameters.*

2 METHODOLOGY

Raw images are far too large to encode directly on near-term quantum hardware, and fixed circuit ansätze may be poorly matched to the features a classical network extracts. To address both issues, we extend EXAQC to higher-dimensional image classification using a hybrid quantum classical architecture consisting of a classical encoder[2], an evolved parameterized quantum circuit (PQC), and a classical decoder. The encoder extracts a compact image representation, EXAQC evolves the quantum transformation applied to this representation, and the decoder maps quantum measurements to class predictions. This design separates the dimensionality of the original image from that processed by the quantum component, allowing a small, evolved circuit to operate on complex image data.

2.1 HYBRID QUANTUM CLASSICAL ARCHITECTURE

Given an image $\mathbf{x}$, the hybrid model is $\hat{\mathbf{y}} = D_\phi\left(Q_{\theta,\mathcal{A}}\left(E_\psi(\mathbf{x})\right)\right)$, where E_ψ is the classical encoder, $Q_{\theta,\mathcal{A}}$ is a quantum circuit with trainable parameters θ and evolved architecture $\mathcal{A}$, and D_ϕ is the classical decoder. For $\mathbf{x} \in \mathbb{R}^{C\times H\times W}$, the encoder maps the raw $d_{\text{raw}} = C \times H \times W$ input

[1] Source code publicly available at: [hidden]

[2] Please note that encoder in the context of this work means feature extractor

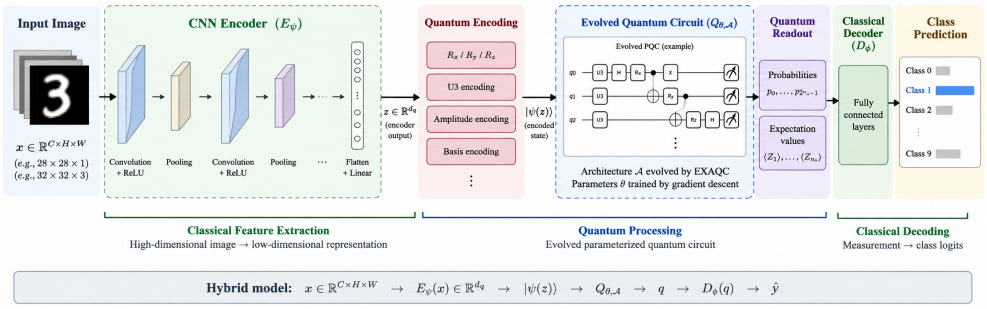


Figure 1: Overview of the proposed EXAQC hybrid quantum classical architecture. A CNN encoder compresses the input image into a lower-dimensional representation that is encoded and processed by an evolved PQC. Quantum measurements are subsequently mapped to class predictions by a classical decoder.

dimensions to $E_\psi : \mathbb{R}^{C\times H\times W} \rightarrow \mathbb{R}^{d_q}$, where d_q denotes the representation supplied to the quantum encoding stage and C, H, W represent the image channel, height and width. This separation allows raw input complexity and quantum representation dimensionality to be studied independently.

2.2 Classical Feature and Quantum Encoding

The encoder E_ψ can be identity block, a linear block or multi-layer perceptron (MLP), or consist of configurable convolutional blocks for spatial feature extraction followed by a fully connected projection to the required quantum input dimension. Its convolutional channels, pooling, activation, normalization, and dropout are configurable. EXAQC supports rotation, U3, and amplitude input encodings Rath & Date (2024). For single-angle R_x, or R_y encoding, each of n_q input qubits receives one feature from d_q input features and therefore $n_q = d_q$. For U3 encoding, three features parameterize each input qubit using

$$U_3(\theta,\phi,\lambda) = \begin{bmatrix} \cos(\theta/2) & -e^{i\lambda}\sin(\theta/2) \\ e^{i\phi}\sin(\theta/2) & e^{i(\phi+\lambda)}\cos(\theta/2) \end{bmatrix}, \tag{1}$$

giving $n_q = d_q/3$. U3 encoding therefore increases the number of classical features presented to a fixed number of input qubits. Amplitude encoding provides an alternative mapping between classical representations and quantum states where $n_q = \lceil log_2(d_q) \rceil$.

2.3 Evolutionary Quantum Circuit Search

EXAQC represents each PQC as a variable-length circuit genome containing its input and output qubits and an ordered collection of quantum gates. Each gate specifies its operation, associated qubits, circuit position, parameters, innovation number, and enabled state. An architecture is therefore $\mathcal{A} = \{g_1, g_2, \ldots, g_m\}$, where both m and the arrangement of gates are determined through evolution. Candidate circuits can consequently differ in gate composition, depth, parameter count, and qubit connectivity rather than being constrained to a fixed ansätz.

Fig. 2 demonstrates how evolution modifies the architecture $\mathcal{A}$ while gradient-based learning optimizes the continuous model parameters. Structural mutation can add, enable, disable, or reorder gates and modify their associated qubits, while binary, exponential, and n-ary crossover combine information from parent circuits. Compatible learned parameters are transferred to offspring through Lamarckian parameter inheritance (Lyu et al. (2021); Kar et al. (2026b)). Candidate circuits are evaluated in parallel using EXAQC's asynchronous master-worker framework. Under steady-state evolution, evaluated offspring immediately compete for population insertion rather than waiting for generational replacement. EXAQC also supports island-based steady-state evolution, in which multiple subpopulations explore different regions of the architecture space while allowing inter-island information exchange.

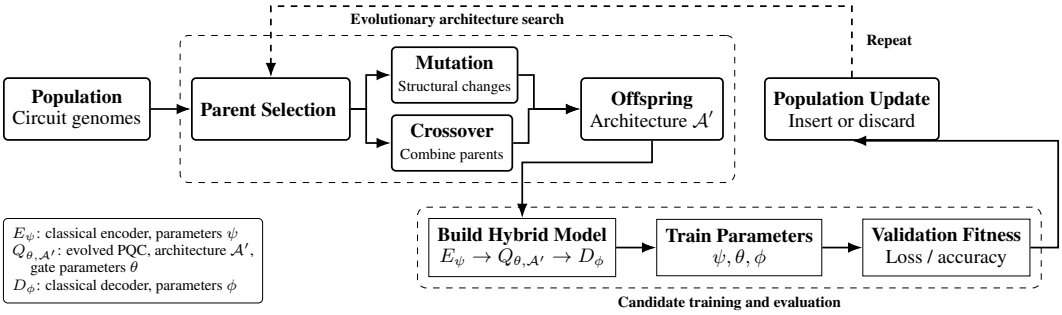


Figure 2: EXAQC evolutionary search workflow for hybrid image classification. Parent circuit genomes are selected and modified through mutation or crossover to create offspring architectures. Each offspring is instantiated within the hybrid model and its classical and quantum parameters are optimized using gradient-based training. Validation performance determines whether the offspring is inserted into the steady-state population.

2.4 Training and Quantum Readout

For each candidate architecture, the encoder parameters ψ, quantum parameters θ, and decoder parameters ϕ, denoted jointly as $\Theta = \{\psi, \theta, \phi\}$, are optimized by minimizing the cross-entropy loss $\mathcal{L}_{CE} = -\frac{1}{N}\sum_{i=1}^{N} \log p_\Theta(y_i \mid \mathbf{x}_i)$, where N is the number of training examples, $\mathbf{x}_i$ is the i-th input image, y_i is its class label, and $p_\Theta(y_i \mid \mathbf{x}_i)$ is the probability the hybrid model assigns to the correct class, obtained by applying a softmax to the decoder's output logits. Validation performance provides the fitness used by evolutionary selection. After the evolved circuit is applied, designated output qubits are measured and supplied to the decoder. For probability readout over n_o output qubits, $\mathbf{q} = [p_0, p_1, \ldots, p_{2^{n_o}-1}]$, where p_k is the probability of measuring computational basis state k, whereas expectation-value readout produces $\mathbf{q} = [\langle Z_1\rangle, \ldots, \langle Z_{n_o}\rangle]$, where $\langle Z_j\rangle$ is the expectation value of the Pauli-Z observable on output qubit j. The decoder D_ϕ then maps $\mathbf{q}$ to the final class logits.

3 Experimental Design

As evolutionary search is stochastic, we report test performance across circuit evaluations and variability across 10 independent runs, indicating how quickly and reliably good circuits are found. Circuit complexity is measured by gate count, and trainable quantum parameters, which determine practicality on near-term hardware and match the resources reported by prior methods. All evolutionary experiments use the same EXAQC implementation and quantum simulation backend. Configurations specify population strategy and size, evaluation budget, mutation and crossover settings, optimizer, learning rate, training steps, batch size, encoding and readout modes, and random seeds. Because EXAQC uses asynchronous steady-state evolution, search progress is reported by the number of evaluated circuits rather than conventional generations. Aaggregate statistics are reported across runs. Representative evolved circuits are selected using test performance.

Baselines and Model Configurations EXAQC hybrid models are compared against fixed-architecture classical baselines. Classical models use the same image-processing pipeline, replacing the evolved PQC with classical fully connected transformations. Where possible, encoder and decoder configurations are held constant across corresponding models.

Scalability and Representation Studies To address **RQ1**, EXAQC is evaluated on MNIST Deng (2012), Fashion-MNIST Xiao et al. (2017), and CIFAR-10 Krizhevsky et al. (2009) , comparing accuracy, circuit resources, and convergence against prior quantum approaches. To address **RQ2**, we vary the classical encoder that produces the quantum circuit's input (identity, MLP, or CNN) while holding the evolutionary configuration fixed, isolating how the supplied representation affects accuracy and circuit complexity. Accuracy, convergence, circuit depth, gate and parameter counts, and computational cost are measured as d_q increases. Representation dimensionality and qubit

count are reported separately when the selected encoding maps multiple features to each qubit. This controlled experiment distinguishes scalability with respect to the representation processed by EXAQC from the raw dimensionality of the source image.

Encoding and Parameter Efficiency To address **RQ3**, we compare amplitude, RX, RY, and U3 input encodings on each dataset, holding the encoder and evolutionary configuration fixed, and report test accuracy and evolved gate and parameter counts to assess whether encoding effects are dataset-dependent. To address **RQ4**, EXAQC-based hybrid models are compared with classical MLP, CNN, and ResNet baselines. Given a classical model with P_C trainable parameters and a hybrid model with P_H, parameter reduction is $R_P = \frac{P_C - P_H}{P_C} \times 100\%$. Because fewer parameters are only meaningful at comparable accuracy, parameter reduction is always assessed jointly with predictive performance.

4 RESULTS

We evaluate EXAQC along three dimensions: predictive performance and circuit resources relative to prior quantum approaches (RQ1 and RQ2), sensitivity to quantum encoding (RQ3), and the accuracy parameter trade-off relative to classical architectures (RQ4).

Hyperparameters EXAQC used a population of 30 and a search budget of 800 candidate genomes for MNIST/Fashion-MNIST and 1000 genomes for CIFAR-10. Candidate parameters were trained for 100 epochs using Adam with a learning rate of 0.001 and cross-entropy loss. Mutation and Crossover operations were performed adopting threshold values from the classification runs in Kar et al. (2026a). Experiments were repeated across 10 independent runs. We used 5-6 qubits for MNIST and Fashion-MNIST and 8-10 qubits for CIFAR-10. Quantum input encoding involved amplitude, angle encoding using *RX* and *RY* gates and, U3 encoding where each gate parameter is fed by an encoder output. Dimensionality representation was configured using encoder architectures involving identity, MLP and CNN based feature extractors. The results reported in Tables 1 utilize 1-layered CNN encoders for MNIST and Fashion-MNIST; and a 3-layered CNN encoder for CIFAR-10 to generate best possible results uisng fine-grained features. Table 2 utilizes linear encoders (MLPs) for MNIST and Fashion-MNIST to drive circuit exploration based on raw features rather than overly depending on the encoder ; CIFAR-10 utilizes convolutional layers to extract finer features since simple encoders were not able to so efficiently. 50K+10K images were used for training and validation, for MNIST and Fashion MNIST with 10K images for testing. 40K+10K images were used for training and validation with CIFAR-10, and 10K images for testing.

4.1 COMPARISON WITH PRIOR IMAGE-CLASSIFICATION APPROACHES

Table 1 compares EXAQC with prior quantum image-classification and architecture-search approaches. EXAQC achieves $98.42 \pm 1.03\%$, $90.62 \pm 0.8\%$, and $85.47 \pm 1.37\%$ accuracy on the full 10-class MNIST, Fashion-MNIST, and CIFAR-10 tasks, respectively. On CIFAR-10, this exceeds the reported results of the architecture-search methods included in the comparison, as some approaches use simplified versions of the datasets that are either downscaled in size and/or reduced to fewer classes. Table 1 further compares quantum circuit resources. EXAQC uses 20 gates and 14 parameters on MNIST and 30 gates and 21 parameters on Fashion-MNIST and CIFAR-10. Its evolved circuits use fewer gates than EQNAS and AQEA-QAS architectures, who follow evolutionary search strategies, across all three datasets, although parameter counts are lower only on MNIST.

4.2 EFFECT OF QUANTUM ENCODING STRATEGY

Table 2 shows that encoding choice has a dataset-dependent effect on both accuracy and circuit complexity. Differences are small on MNIST, where accuracy ranges from $95.65 \pm 0.10\%$ to $96.42 \pm 0.20\%$, but become more pronounced on Fashion-MNIST and CIFAR-10. On CIFAR-10, amplitude encoding achieves $61.38 \pm 0.80\%$, compared with $83.75 \pm 1.40\%$, $84.37 \pm 1.30\%$, and $86.68 \pm 1.50\%$ for RX, RY, and U3, respectively. CIFAR-10 also produces larger evolved circuits, containing 42-60 gates and 22-40 parameters, which is expected due to the additional complexity of the task. Because U3 encodes three features per qubit whereas RX and RY encode one, the encodings cannot be

Table 1: Comparison with prior quantum and quantum-architecture-search approaches on image classification: test accuracy (%) and quantum circuit resources (G/P = gates / quantum gate parameters). Cls.- specifies how many classes were actually used from the dataset. Accuracies are not necessarily directly comparable, as prior studies differ in number of classes, dataset subsets, preprocessing, and quantum input representations. "–" indicates a result or resource count that was not reported.

		MNIST		Fashion-MNIST		CIFAR-10	
Method	**Cls.**	**Acc.**	**G/P**	**Acc.**	**G/P**	**Acc.**	**G/P**
Hybrid DQAS[a] (Afane et al., 2026)	2	97.97 ± 0.40	–	93.86 ± 1.70	–	71.82 ± 1.60	–
DQAS[b] (Zhang et al., 2022)	2	96.28 ± 3.90	–	93.40 ± 0.09	–	71.12 ± 5.70	–
QuantumDARTS[b] (Wu et al., 2023)	2	95.66 ± 3.30	–	92.20 ± 0.06	–	70.14 ± 1.80	–
EQAS (Zhang & Zhao, 2023)	2	94.00	22/**9**	75.00	**24/20**	–	–
EQNAS[e] (Li et al., 2023a)	2	87.35	43/20	84.15	41/18	59.65	43/20
AQEA-QAS (Li et al., 2025)	2	91.92	39/16	87.15	39/16	61.65	39/**16**
Quantum K-NN[c,d] (Maldonado-Romo et al., 2024)	2	94.00 ± 2.11	–	–	–	–	–
HQNN-Parallel[d] (Senokosov et al., 2024)	10	$\mathbf{99.21}$	–	–	–	82.64	–
QNAS[f] (Maleki et al., 2026)	10	97.16	68/–	87.38	42/–	–	–
EXAQC (ours)	10	$\mathbf{98.42 \pm 1.03}$	**20**/14	$\mathbf{90.62 \pm 0.8}$	30/21	$\mathbf{85.47 \pm 1.37}$	**30**/21

[a]Images resized for 4-, 9-, and 16-qubit configurations. [b]Results from the comparative evaluation in Hybrid DQAS. [c]MNIST digits 0 and 1 only. [d]Not an architecture-search method. [e]Cross-dataset results from the comparative evaluation in AQEA-QAS. [f]8 qubits/2 layers (MNIST), 5 qubits/2 layers (Fashion-MNIST); gate count follows the original work's transpilation convention.

compared at both a fixed qubit count and a fixed input dimension. We therefore hold the number of qubits fixed, so U3 receives three times as many input features; matching the input dimension instead would require tripling the qubit count for RX and RY, making the process more computationally expensive.

4.3 Parameter Reduction

Table 3 compares the accuracy and trainable parameter counts of EXAQC-based hybrid models with classical architectures. On MNIST, EXAQC achieves 96.13% accuracy with 101,458 parameters, compared with 97.48% using 468,874 parameters for the MLP. On Fashion-MNIST, the CNN-based EXAQC model achieves 91.11% with 33,844 parameters, compared with 89.47% and 74,922 parameters for the one-layer CNN. On CIFAR-10, EXAQC achieves 85.68% with 100,173 parameters, compared with 85.50% and 2,569,962 parameters for the 10-layer CNN. These results show substantial reductions in trainable parameters while retaining comparable classification performance.

Table 2: Comparison of quantum encoding strategies for EXAQC across the image-classification benchmarks. Results are reported over multiple independent runs as mean ± standard deviation.

Dataset	Quantum Encoding	Test Acc. (%)	# Gates	# Params.
MNIST	Amplitude	95.65 ± 0.1	8	5
	RX	96.22 ± 0.1	11	6
	RY	96.27 ± 0.2	17	9
	U3	$\mathbf{96.42 \pm 0.2}$	14	14
Fashion-MNIST	Amplitude	82.19 ± 0.2	5	6
	RX	$\mathbf{88.01 \pm 0.8}$	12	6
	RY	81.90 ± 1.2	13	8
	U3	87.26 ± 0.5	4	4
CIFAR-10	Amplitude	61.38 ± 0.8	42	22
	RX	83.75 ± 1.4	43	26
	RY	84.37 ± 1.3	45	27
	U3	$\mathbf{86.68 \pm 1.5}$	60	40

Table 3: Comparison of classical and EXAQC-based hybrid models in terms of classification performance and parameter efficiency. Parameter reduction is calculated relative to the corresponding classical baseline for each dataset. Best Results are reported over 10 independent runs using RY angle encoding.

Dataset	Model	Test Acc. (%)	# Trainable Params.	Param. Reduction ($\mathbf{R_P}\%$)
MNIST	MLP ($784 \times 512 \times 128 \times 10$)	97.48	468874	78.36 (↓)
	EXAQC + MLP	96.83	101458	–
Fashion-MNIST	MLP ($784 \times 512 \times 128 \times 10$)	88.03	468874	76.67 (↓)
	EXAQC + MLP	86.95	109421	–
	1-layer CNN	89.47	74922	54.83 (↓)
	2-layer CNN	90.87	79418	57.38 (↓)
	EXAQC + CNN	91.11	33844	–
CIFAR-10	3-layer CNN	81.11	258986	61.32 (↓)
	10-layer CNN	85.50	2569962	96.10 (↓)
	ResNet18	87.08	11173962	99.10 (↓)
	EXAQC + CNN	85.68	100173	–

4.4 Circuit Visualization

Figure 3 shows a representative evolved hybrid architecture for CIFAR-10, the most complex of the three. The encoder applies three convolutional blocks ($3 \to 32 \to 64 \to 128$ feature maps, each followed by 2×2 max pooling), and a linear layer then compresses the flattened features to only 8 dimensions. This reduces the 3,072-dimensional input by a factor of 384 before quantum processing. The 8 features are RY-encoded onto 8 qubits, one feature per qubit. The evolved circuit is substantially larger and more densely entangled than the MNIST and Fashion-MNIST circuits. It combines single-qubit rotations, many controlled two-qubit gates, and multi-qubit Ising interactions spanning distant qubits in an irregular, non-layered arrangement that differs from the repeated templates of hand-designed ansätze. Probability readout over all 8 qubits produces 256 values, which a linear decoder maps to the 10 class logits. The classical encoder thus performs heavy feature extraction and compression, while the evolved circuit supplies a richer quantum transformation of the

compact representation than is needed for the simpler datasets. Representations of circuit genomes as evolved for MNIST and Fashion MNIST dataset have been added in Appendix B.6.

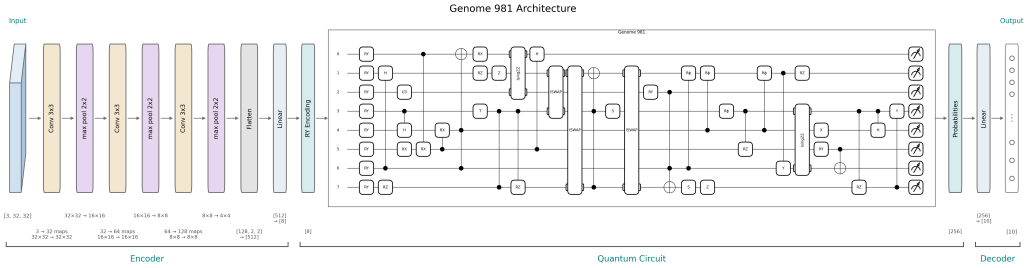


Figure 3: Evolved EXAQC hybrid architecture for CIFAR-10. A three-block CNN encoder compresses each image to 8 features, which are RY-encoded onto 8 qubits.

5 Discussion

The results demonstrate that EXAQC can be extended to hybrid quantum classical image classification while maintaining compact evolved circuits and competitive predictive performance. Across the three benchmarks, the experiments reveal differences in performance, circuit complexity, encoding behavior, and parameter requirements as task complexity increases. The convergence and operator analyses further characterize how EXAQC searches and adapts circuit architectures during evolution. We discuss these findings below in the context of the four research questions.

RQ1: ***Can EXAQC scale from low-dimensional learning problems to complex image-classification tasks by discovering useful quantum transformations of learned image representations within a hybrid quantum–classical framework?*** EXAQC remains effective across MNIST, Fashion-MNIST, and CIFAR-10, achieving 98.42%, 90.62%, and 85.47% test accuracy on the full 10-class problems. On CIFAR-10, this exceeds the reported accuracies of the architecture-search methods included in Table 1, while the evolved circuits use fewer gates than EQNAS and AQEA-QAS on all three datasets. The convergence trajectories in Figure 5 (Refer Appendix B) show a consistent pattern across datasets, with rapid initial loss reduction followed by gradual stabilization. Together, these results indicate that evolved PQCs can effectively transform classically extracted image representations as task complexity increases.

RQ2: ***How does the classical representation supplied to the quantum circuit affect predictive performance and circuit complexity?*** The choice of classical encoder substantially affects the evolved hybrid models. On Fashion-MNIST, EXAQC with a CNN encoder achieves 91.11% accuracy with 33,844 parameters, compared with 86.95% and 109,421 parameters using an MLP encoder (Table 3), so the richer representation improves accuracy while reducing model size. On CIFAR-10, simple encoders did not provide sufficiently informative features, and a CNN encoder was required to reach competitive accuracy. These results suggest that EXAQC performs best when the quantum circuit operates on learned, spatially informed features rather than raw or linearly projected inputs.

RQ3: ***How does the choice of quantum input encoding affect the accuracy and complexity of evolved circuits, and does this effect depend on the dataset?*** Encoding effects are strongly dataset-dependent (Table 2). On MNIST, all encodings fall within one percentage point (95.65-96.42%), while on Fashion-MNIST, RX (88.01%) and U3 (87.26%) outperform amplitude (82.19%) and RY (81.90%) encoding, with U3 achieving this using only 4 gates. On CIFAR-10, the rotation-based encodings (83.75-86.68%) substantially outperform amplitude encoding (61.38%), although the best-performing U3 encoding also yields the largest circuits (60 gates, 40 parameters). Encoding choice therefore introduces a dataset-dependent trade-off between accuracy and circuit resources, and its importance grows with task complexity.

RQ4: ***Do evolved quantum components maintain competitive classification performance relative to larger classical counterparts, despite having fewer trainable parameters?*** The hybrid models substantially reduce trainable parameters while retaining comparable accuracy (Table 3). On CIFAR-10, EXAQC achieves 85.68% with 100,173 parameters, matching the 10-layer CNN

(85.50%) with 96.1% fewer parameters and remaining within 1.4 points of ResNet18 with 99.1% fewer. On Fashion-MNIST, EXAQC exceeds the one-layer CNN (91.11% vs. 89.47%) with 54.8% fewer parameters. On MNIST, the accuracy cost is small (96.83% vs. 97.48%) for a 78.4% reduction. The evolved PQCs are also compact at the circuit level, using 20-30 gates. These results demonstrate a favorable accuracy parameter trade-off for the evolved hybrid models.

6 Limitations and Future Work

While the results demonstrate the potential of EXAQC for compact hybrid image-classification models, they also identify several areas requiring further investigation. First, comparisons with prior work are not strictly controlled because of differences in preprocessing, classification settings, quantum representations, and resource-counting conventions. Prior works also rely on simplified formulations of the classification datasets. Standardized evaluations would enable more direct comparisons. Second, the effectiveness of the evolved PQC depends on the quality of the features supplied by the classical encoder. High-dimensional image classification therefore requires feature extractors that produce compact quantum-facing representations while preserving task-relevant information. Future work should investigate joint optimization of the encoder, quantum encoding, and evolved circuit architecture. Third, the current evaluation primarily relies on simulation. Experiments under realistic noise and on quantum hardware are needed to assess the robustness of evolved circuits, with noise-aware evolutionary objectives providing a natural extension. Finally, evolutionary search becomes increasingly expensive as the circuit and representation spaces grow. Future work will explore more efficient distributed search, multi-objective optimization of performance and circuit resources, and evolution of additional hybrid components such as quantum convolutional blocks.

7 Conclusion

This work extends EXAQC to evolutionary discovery of quantum circuits for hybrid image classification. Across MNIST, Fashion MNIST, and CIFAR-10, EXAQC evolves compact quantum components that achieve competitive classification performance while substantially reducing the trainable parameter requirements of the surrounding hybrid models. EXAQC achieves 98.42%, 90.62%, and 85.47% accuracy on the three full 10-class benchmarks, respectively, while the evolved circuits contain only 20-30 gates in the resource comparison. Unlike most prior quantum architecture search methods, which rely on fixed input encodings, EXAQC trains the classical encoder jointly with each evolved circuit, learning the representation supplied to the PQC. This is important as task complexity grows: on CIFAR-10, learned representations with rotation-based encodings reach 83.75-86.68% accuracy, compared with only 61.38% using amplitude encoding. To our knowledge, this is the first work to combine learned U3 encodings with evolutionary quantum architecture search, with U3 achieving the highest CIFAR-10 accuracy (86.68%) among the evaluated encodings. In particular, applying quantum circuits to high-dimensional image data requires effective classical feature extractors that produce compact representations while retaining task-relevant information. On CIFAR-10, the hybrid model achieves 85.68% accuracy with 100,173 trainable parameters, compared with 85.50% using 2,569,962 parameters for the 10-layer CNN. Overall, the results support evolutionary search as an approach for discovering compact quantum components, while highlighting the importance of jointly considering feature extraction, quantum encoding, and circuit architecture when designing parameter-efficient hybrid models.

Reproducibility Statement

We have taken several steps to support reproducibility of the reported experiments. The methodology describes the EXAQC evolutionary search procedure, hybrid quantum classical architecture, quantum encoding and readout strategies, and evaluation metrics, while the experimental setup reports the datasets, model configurations, hyperparameters, and number of independent runs used in our evaluation across Sections 2- 5. All experiments use publicly available MNIST, Fashion-MNIST, and CIFAR-10 datasets with appropriate citations. Results are reported across multiple independent runs with measures of variation where applicable in Section 4 and the scripts are available in the repository.

AI Usage

Generative AI has been used to design some of the code utilized in these experiments, with a focus on generating unit tests and performing comprehensive code review to find and correct bugs and inconsistencies. It has also been used to polish the writing, identify and correct grammatical errors and typos in the text.

References

Mohamed Afane, Quanjiang Long, Haoting Shen, Ying Mao, Junaid Farooq, Ying Wang, and Juntao Chen. Differentiable architecture search for adversarially robust quantum computer vision. *Quantum Machine Intelligence*, 8(1):2, 2026.

Alexander Benítez-Buenache and Queralt Portell-Montserrat. Bayesian parameterized quantum circuit optimization (bpqco): A task and hardware-dependent approach. *Quantum Machine Intelligence*, 7(2):85, 2025.

Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature communications*, 12(1): 1791, 2021.

Li Deng. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE signal processing magazine*, 29(6):141–142, 2012.

Li Ding and Lee Spector. Evolutionary quantum architecture search for parametrized quantum circuits. In *Proceedings of the genetic and evolutionary computation conference companion*, pp. 2190–2195, 2022.

Li Ding and Lee Spector. Multi-objective evolutionary architecture search for parameterized quantum circuits. *Entropy*, 25(1):93, 2023.

Xiaodong Ding, Zhihui Song, Jinchen Xu, Yifan Hou, Tian Yang, and Zheng Shan. Scalable parameterized quantum circuits classifier. *Scientific reports*, 14(1):15886, 2024.

Yuxuan Du, Min-Hsiu Hsieh, Tongyang Liu, and Dacheng Tao. Quantum circuit architecture search for variational quantum algorithms. *npj Quantum Information*, 8(1), 2022a. doi: 10.1038/s41534-022-00570-y. URL https://www.nature.com/articles/s41534-022-00570-y.

Yuxuan Du, Tao Huang, Shan You, Min-Hsiu Hsieh, and Dacheng Tao. Quantum circuit architecture search for variational quantum algorithms. *npj Quantum Information*, 8(1):62, 2022b.

Alessandro Giovagnoli, Volker Tresp, Yunpu Ma, and Matthias Schubert. Qneat: Natural evolution of variational quantum circuit architecture. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*, pp. 647–650, 2023.

Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019. doi: 10.1038/s41586-019-0980-2. URL https://arxiv.org/abs/1804.11326.

Hsin-Yuan Huang, Kishor Bharti, and Xingjian Wu. Adaptive diversity-based quantum circuit architecture search. *Physical Review Research*, 6:033033, 2024. doi: 10.1103/PhysRevResearch.6.033033. URL https://link.aps.org/doi/10.1103/PhysRevResearch.6.033033.

Thomas Hubregtsen, Josef Pichlmeier, Patrick Stecher, and Koen Bertels. Evaluation of parameterized quantum circuits: on the relation between classification accuracy, expressibility, and entangling capability. *Quantum Machine Intelligence*, 3(1):9, 2021.

Devroop Kar, Daniel Krutz, and Travis Desell. Evolutionary design of parameterized quantum circuits for classification and reinforcement learning tasks. In *International Conference on Parallel Problem Solving from Nature*, pp. 284–302. Springer, 2026a.

Devroop Kar, Daniel Krutz, and Travis Desell. Investigating quantum circuit designs using neuroevolution. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 541–544, 2026b.

Michael Kölle, Tom Bintener, Maximilian Zorn, Gerhard Stenzel, Leo Sünkel, Thomas Gabor, and Claudia Linnhoff-Popien. Evaluating mutation techniques in genetic-algorithm-based quantum circuit synthesis. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 907–915, 2025.

Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

Yangyang Li, Ruijiao Liu, Xiaobin Hao, Ronghua Shang, Peixiang Zhao, and Licheng Jiao. Eqnas: Evolutionary quantum neural architecture search for image classification. *Neural Networks*, 168: 471–483, 2023a.

Yaochong Li, Jing Zhang, Rigui Zhou, Yi Qu, and Ruiqing Xu. Aqea-qas: An adaptive quantum evolutionary algorithm for quantum architecture search. *Entropy*, 27(7):733, 2025.

Yifan Li, Ming Zhang, and Dong-Ling Deng. Sa-dqas: Self-attention based differentiable quantum architecture search. *arXiv preprint arXiv:2306.04586*, 2023b. URL https://arxiv.org/abs/2306.04586.

Junhua Liu, Kwan Hui Lim, Kristin L Wood, Wei Huang, Chu Guo, and He-Liang Huang. Hybrid quantum-classical convolutional neural networks. *Science China Physics, Mechanics & Astronomy*, 64(9):290311, 2021.

Yu Liu, Kazuya Kaneko, Kentaro Baba, Jumpei Koyama, Koichi Kimura, and Naoyuki Takeda. Analysis of parameterized quantum circuits: on the connection between expressibility and types of quantum gates. *IEEE Transactions on Quantum Engineering*, 6:1–12, 2025.

Matt Lourens, Ilya Sinayskiy, Daniel K Park, Carsten Blank, and Francesco Petruccione. Hierarchical quantum circuit representations for neural architecture search. *npj Quantum Information*, 9 (1):79, 2023.

Martin Lukac, Marek Perkowski, Hilton Goi, Mikhail Pivtoraiko, Chung Hyo Yu, Kyusik Chung, Hyunkoo Jeech, Byung-Guk Kim, and Yong-Duk Kim. Evolutionary approach to quantum and reversible circuits synthesis. *Artificial Intelligence Review*, 20(3):361–417, 2003.

Zimeng Lyu, AbdElRahman ElSaid, Joshua Karns, Mohamed Mkaouer, and Travis Desell. An experimental study of weight initialization and lamarckian inheritance on neuroevolution. *The 24th International Conference on the Applications of Evolutionary Computation (EvoStar: EvoApps)*, 2021.

QuanGong Ma, ChaoLong Hao, XuKui Yang, LongLong Qian, Hao Zhang, NianWen Si, MinChen Xu, and Dan Qu. Continuous evolution for efficient quantum architecture search. *EPJ Quantum Technology*, 11(1):54, 2024.

Alberto Maldonado-Romo, J Yaljá Montiel-Pérez, Victor Onofre, Javier Maldonado-Romo, and Juan Humberto Sossa-Azuela. Quantum k-nearest neighbors: Utilizing qram and swap-test techniques for enhanced performance. *Mathematics*, 12(12):1872, 2024.

Kooshan Maleki, Alberto Marchisio, and Muhammad Shafique. Qnas: A neural architecture search framework for accurate and efficient quantum neural networks. *arXiv preprint arXiv:2604.07013*, 2026.

Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature communications*, 9(1):4812, 2018.

Fernando T Miranda, Pedro Paulo Balbi, and Pedro Costa. Synthesis of quantum circuits with an island genetic algorithm. *arXiv preprint arXiv:2106.03115*, 2021.

Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. Data re-uploading for a universal quantum classifier. *Quantum*, 4:226, 2020. doi: 10.22331/q-2020-02-06-226. URL https://arxiv.org/abs/1907.02085.

Arthur Pesah, Marco Cerezo, Samson Wang, Tyler Volkoff, Andrew T Sornborger, and Patrick J Coles. Absence of barren plateaus in quantum convolutional neural networks. *Physical Review X*, 11(4):041011, 2021.

Minati Rath and Hema Date. Quantum data encoding: A comparative analysis of classical-to-quantum mapping techniques and their impact on machine learning accuracy. *EPJ Quantum Technology*, 11(1):72, 2024.

Cristian Ruican, Mihai Udrescu, Lucian Prodan, and Mircea Vladutiu. A genetic algorithm framework applied to quantum circuit synthesis. In *Nature inspired cooperative strategies for optimization (NICSO 2007)*, pp. 419–429. Springer, 2008.

Maria Schuld and Nathan Killoran. Effect of data encoding on the expressive power of variational quantum machine learning models. *Physical Review A*, 99(3):032331, 2019. doi: 10.1103/PhysRevA.99.032331. URL https://arxiv.org/abs/1803.07128.

Maria Schuld, Alexey Bocharov, Krysta M. Svore, and Nathan Wiebe. Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308, 2020. doi: 10.1103/PhysRevA.101.032308. URL https://arxiv.org/abs/1804.00633.

Arsenii Senokosov, Alexandr Sedykh, Asel Sagingalieva, Basil Kyriacou, and Alexey Melnikov. Quantum machine learning for image classification. *Machine Learning: Science and Technology*, 5(1):015040, 2024.

Sukin Sim, Peter D Johnson, and Alán Aspuru-Guzik. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12):1900070, 2019.

Leo Sünkel, Darya Martyniuk, Denny Mattern, Johannes Jung, and Adrian Paschke. Ga4qco: genetic algorithm for quantum circuit optimization. *arXiv preprint arXiv:2302.01303*, 2023.

Leo Sünkel, Philipp Altmann, Michael Kölle, Gerhard Stenzel, Thomas Gabor, and Claudia Linnhoff-Popien. Quantum circuit construction and optimization through hybrid evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 934–942, 2025.

Antonio Tudisco, Andrea Marchesin, Maurizio Zamboni, Mariagrazia Graziano, and Giovanna Turvani. Evaluating angle and amplitude encoding strategies for variational quantum machine learning: Their impact on model's accuracy. *Advanced Quantum Technologies*, 9(4):e00611, 2026.

Yabo Wang, Bo Qi, Chris Ferrie, and Daoyi Dong. Trainability enhancement of parameterized quantum circuits via reduced-domain parameter initialization. *Physical Review Applied*, 22(5): 054005, 2024.

Yunqian Wang, Yufeng Wang, Chao Chen, Runcai Jiang, and Wei Huang. Development of variational quantum deep neural networks for image recognition. *Neurocomputing*, 501:566–582, 2022.

Zhenyu Wu, Zhirui Zhang, Zhihui Wang, and Runyao Duan. Quantumdarts: Differentiable quantum architecture search for variational quantum algorithms. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 2023. URL https://proceedings.mlr.press/v202/wu23v.html.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Anqi Zhang and Shengmei Zhao. Evolutionary-based searching method for quantum circuit architecture. *Quantum Information Processing*, 22(7):283, 2023.

Ming Zhang and Dong-Ling Deng. Differentiable quantum architecture search. *arXiv preprint arXiv:2010.08561*, 2020. URL https://arxiv.org/abs/2010.08561.

Ming Zhang and Dong-Ling Deng. Neural predictor based quantum architecture search. *arXiv preprint arXiv:2103.06524*, 2021. URL https://arxiv.org/abs/2103.06524.

Shi-Xin Zhang, Chang-Yu Hsieh, Shengyu Zhang, and Hong Yao. Differentiable quantum architecture search. *Quantum Science & Technology*, 7(4):045023, 2022.