IOP Publishing Mach. Learn.: Sci. Technol. 6 (2025) 035061 https://doi.org/10.1088/2632-2153/ae056d

MACHINE LEARNING
Science and Technology

OPEN ACCESS

RECEIVED
28 July 2025

REVISED
28 August 2025

ACCEPTED FOR PUBLICATION
4 September 2025

PUBLISHED
22 September 2025

PAPER

Quantum architecture search for optimizing quantum generators in quantum GAN

Quangong Ma[1], **Chaolong Hao**[1,*], **NianWen Si**[2,1] and **Dan Qu**[1,3]

[1] School of Information Systems Engineering, Information Engineering University, Zhengzhou, People's Republic of China
[2] Department of Electronic Engineering, Tsinghua University, Beijing, People's Republic of China
[3] Laboratory for Advanced Computing and Intelligence Engineering, Wuxi, People's Republic of China
[*] Author to whom any correspondence should be addressed.

E-mail: hcl_xdspeechlab@aliyun.com

Keywords: quantum architecture search, quantum generative adversarial network, quantum circuits

Abstract

As quantum computing continues to advance alongside machine learning, Quantum Generative Adversarial Networks (QGANs) have gained attention as a compelling generative modeling approach in quantum machine learning. Despite their potential, most existing QGAN implementations rely on manually designed quantum circuits, which often suffer from high complexity and limited scalability. To overcome these limitations, we propose a Quantum Architecture Search (QAS) method built upon the SuperCircuit framework. This approach enables the automatic discovery of efficient quantum circuit structures tailored for generative tasks, thereby reducing gate count. However, applying SuperCircuit-based QAS directly to QGANs presents two primary challenges: an expansive search space and a lack of guarantees regarding circuit efficiency. To mitigate these issues, we incorporate Principal Component Analysis (PCA) and a PCA-guided feature distribution mechanism during preprocessing. This strategy both compresses the search space and ensures balanced allocation of principal components across sub-generators. Extensive experiments confirm the viability of the proposed framework. When employing the discovered quantum circuits, QAS-QGAN achieves image quality comparable to a state-of-the-art (SOTA) baseline (Silver *et al* 2023 *2023 IEEE/CVF Int. Conf. on Computer Vision (ICCV)* pp 7007–16) on both the MNIST and Fashion MNIST datasets. Notably, our model reduces the number of two-qubit gates by 52.08% on MNIST and 52.92% on Fashion MNIST, highlighting substantial improvements in quantum resource efficiency. Furthermore, the model successfully generates high-resolution facial images (109×89) on the CelebA dataset, demonstrating strong scalability and practical applicability.

1. Introduction

The field of quantum machine learning (QML) has seen growing interest owing to its potential to significantly reduce training time and boost model effectiveness. Extensive research has demonstrated the practical viability of various QML architectures, including quantum neural networks [2–4], quantum autoencoders [5, 6], QGANs [1, 7–9], and quantum diffusion models [10–12].

QGANs are built upon the adversarial interplay between a quantum generator and a discriminator, forming a competitive learning paradigm. The generator transforms stochastic noise into outputs intended to mimic samples from a target distribution, whereas the discriminator evaluates both synthetic and real samples to distinguish between them. In recent years, QGANs have been effectively employed in diverse domains such as image synthesis [1, 7–9, 13], small-molecule drug discovery [14, 15], and anomaly detection [16]. To improve image fidelity and training robustness, numerous enhancements have been introduced. These include hybrid generator designs that integrate quantum circuits with classical neural networks [8, 14, 16], the adoption of advanced training strategies from classical GAN frameworks [8, 14], the development of quantum-specific loss functions [17], and the fusion of QGANs with other quantum

architectures such as Quantum Implicit Neural Representations [9], Quantum Long Short-Term Memory networks [18], and unsupervised learning techniques [1, 13]. Collectively, these innovations have markedly advanced the capabilities of QGANs, laying a robust groundwork for the continued evolution of quantum generative modeling and the broader field of QML.

It is widely recognized that the generative performance of QGANs is highly influenced by the architecture of the generator. Nonetheless, most existing works rely on manually crafted quantum circuits, which often suffer from limited generalizability and elevated circuit complexity. To overcome these constraints, we explore a promising paradigm in QML–Quantum Architecture Search (QAS). QAS facilitates the automated discovery of near-optimal ansätze (i.e. quantum circuit configurations) tailored to specific tasks, effectively balancing performance with quantum resource efficiency [19–21]. Compared to hand-designed circuits, QAS offers enhanced adaptability to diverse tasks and significantly reduces the effort associated with hyperparameter tuning. This makes it particularly advantageous for deployment on resource-constrained platforms, such as Noisy Intermediate-Scale Quantum (NISQ) devices. Prior research has validated the effectiveness of QAS in optimizing quantum neural network architectures for classification tasks [22], and its applicability has also been demonstrated in quantum chemistry domains [20, 23, 24]. Thus, incorporating QAS into the design process of QGAN generators presents a compelling strategy for structural refinement and circuit simplification, ultimately contributing to improved performance and scalability in quantum generative modeling.

Incorporating QAS into the design of QGAN generators presents a promising avenue for advancing QGAN development. The primary objective of this integration is to optimize quantum circuit structures without compromising image generation quality, while simultaneously reducing quantum resource demands–this constitutes the core motivation behind the proposed QAS-enhanced QGAN framework. In this study, we investigate the application of a SuperCircuit-based QAS algorithm [20, 21, 24] to guide the architectural design of the QGAN generator. However, directly deploying the SuperCircuit-based QAS approach in this context introduces two notable challenges: (1) The inherently high dimensionality of QGAN-generated data leads to an expansive search space, severely impacting the efficiency of the architecture search process. (2) The quantum circuits derived through QAS are not intrinsically aligned with the discriminator's decision boundaries, which can result in training instabilities such as mode collapse. To address these limitations, we implement Principal Component Analysis (PCA) and a PCA-based feature distribution strategy during data preprocessing. This not only reduces the dimensionality of the input space but also balances the distribution of principal components across different sub-generators, thereby improving search tractability. Furthermore, we incorporate the adaptive input noise method proposed in [1] to enhance image quality and further mitigate mode collapse during adversarial training.

The overall architecture of the proposed QAS-QGAN framework is depicted in figure 1. The upper section of the figure outlines the data preprocessing pipeline and the quantum generator architecture search procedure. Initially, PCA is applied to the input dataset for dimensionality reduction. The resulting principal components are ranked based on variance, and five representative dimensions are selected to preserve the most informative features of the data distribution. These selected components serve as the modeling basis for the generator. Subsequently, a parameterized SuperCircuit is constructed, and QAS is employed to evaluate a set of candidate circuit architectures. The circuit that achieves the best performance according to the evaluation criteria is then chosen to instantiate the quantum generator. The lower portion of figure 1 illustrates the image generation phase of QAS-QGAN. The optimized quantum generator–based on the selected circuit–accepts a latent vector as input and generates the corresponding principal components of an image. These components are then assessed by a classical discriminator to determine authenticity. The generator and discriminator are jointly trained via iterative adversarial optimization. Finally, inverse PCA is applied to reconstruct the full-resolution image from its principal components, yielding the final visual output.

We evaluate the generative performance of the proposed model on the MNIST [25] and Fashion MNIST [26] datasets, both widely used benchmarks for image synthesis. To ensure a comprehensive and objective assessment of image quality, we employ four complementary evaluation metrics: Fréchet Inception Distance (FID) [27], Structural Similarity Index Measure (SSIM) [28], Cosine Similarity [29], and Peak Signal-to-Noise Ratio (PSNR) [30]. These metrics collectively capture various aspects of image fidelity, perceptual similarity, and structural integrity. Beyond standard benchmarks, we further investigate the scalability of QAS-QGAN by applying it to generate facial images at a resolution of 109×89 using the CelebA dataset [31], thereby validating the model's capacity to handle higher-resolution and more complex data distributions.

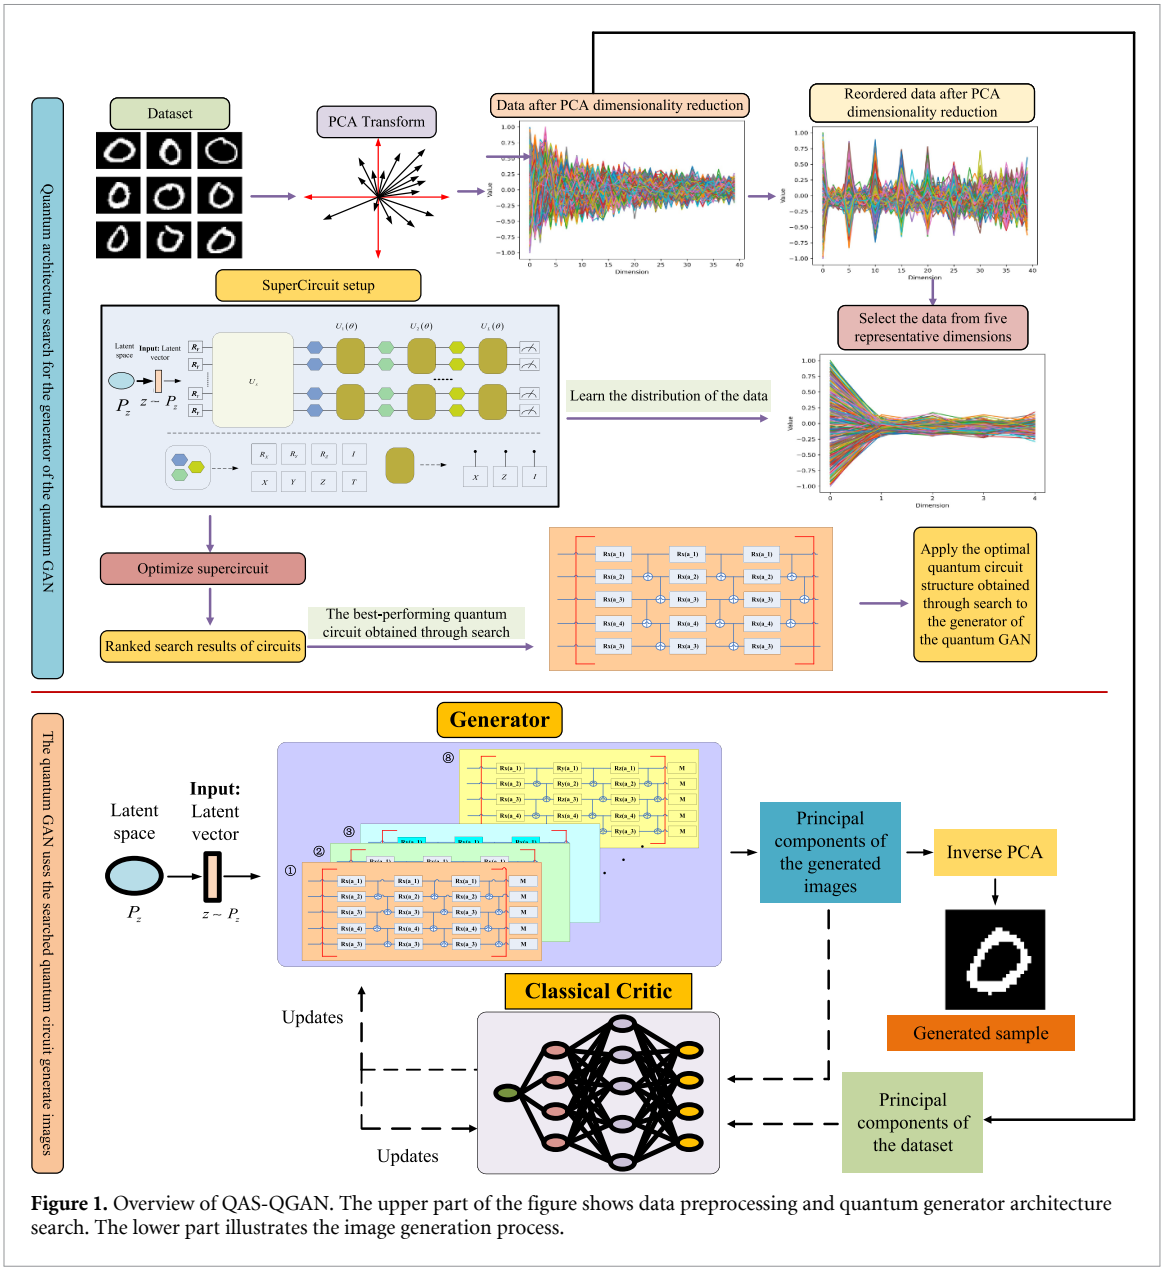

Figure 1. Overview of QAS-QGAN. The upper part of the figure shows data preprocessing and quantum generator architecture search. The lower part illustrates the image generation process.

The key contributions of this work are summarized as follows:

1. We propose a novel QAS-enhanced QGAN framework (QAS-QGAN), which leverages QAS to automatically discover optimal quantum generator structures tailored for image generation tasks. This approach significantly improves both the adaptability and performance of QGANs.
2. We integrate PCA and a PCA-based feature redistribution strategy into the architecture search and generation pipeline. This integration effectively reduces the dimensionality of the QAS search space and alleviates potential imbalances caused by uneven principal component allocation across sub-generators. As a result, the model achieves improved image quality and reduced susceptibility to mode collapse.
3. Our framework introduces a resource-efficient QGAN design that minimizes quantum computational cost while maintaining high-fidelity image synthesis. Additionally, we demonstrate the scalability of QAS-QGAN by extending the resolution of generated images to 109×89, thereby paving the way for high-resolution quantum image generation.

The structure of this paper is organized as follows: section 2 surveys related work, including developments in QGANs, QAS, Neural Architecture Search (NAS), and their integration. Section 3 introduces the proposed QAS-QGAN framework and elaborates on its design. Section 4 details the experimental settings and presents the corresponding results. In section 5, we discuss the implications of our findings, and section 6 concludes with future research directions.

2. Related work

2.1. Quantum generative adversarial network

In classical generative modeling, numerous powerful techniques have been extensively explored, including GANs [32], Autoencoders [33], and, more recently, Diffusion Models [34], which have attracted considerable interest. The success of these methods in a wide range of data generation applications offers meaningful guidance for pushing forward innovations in QML. These advancements have laid important groundwork for ongoing efforts in quantum generative modeling. Drawing inspiration from classical generative methods, QGANs were initially proposed in 2018 [35, 36], and have since been extended to a variety of application domains, such as image synthesis [1, 7–9, 13], molecular design for drug discovery [14, 15], and anomaly detection [36]. Architecturally, QGANs follow the fundamental design of classical GANs, consisting of a generator and a discriminator engaged in an adversarial min-max game, as formalized in equation (1)

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{\text{data}}(x)}\left[\log D(x)\right] + \mathbb{E}_{z\sim p_z(z)}\left[\log\left(1 - D(G(z))\right)\right]. \tag{1}$$

Here, D represents the discriminator, G the generator, p_{data} the real data distribution, and p_z the prior distribution of the input noise. Owing to the inherently adversarial dynamics of the training process, QGANs–similar to their classical counterparts–are susceptible to training instability. A primary set of challenges includes mode collapse–where the generator fails to produce diverse outputs–and the vanishing gradient problem, which limits the model's ability to update parameters effectively.

To address the limitations and challenges inherent in QGANs, a wide range of enhancement strategies have been explored. Some studies have adapted training techniques from classical GANs to improve training stability [8, 14], while others have introduced quantum-specific loss functions designed to alleviate optimization instability [17]. Additional approaches involve integrating QGANs with advanced quantum architectures, such as Quantum Implicit Neural Representations [9], QLSTM networks [18], and unsupervised learning frameworks [1, 13], thereby augmenting their representational and generative power. Moreover, pretrained autoencoders have been employed to boost image generation quality [37]. Hybrid approaches that combine quantum circuits and classical neural networks have demonstrated significant promise, paving the way for the creation of efficient quantum–classical generative adversarial networks [8, 14, 16].

2.2. Architecture search for GANs and QGANs

The development of classical machine learning significantly predates that of its quantum counterpart, and many classical methodologies provide valuable guidance for advancing quantum approaches. In classical settings, NAS focuses on the automated discovery of effective neural network architectures [38]. Analogously, in the quantum domain, QAS aims to identify near-optimal ansätze (i.e. quantum circuit structures) that are well-suited to specific computational tasks [20]. Both NAS and QAS utilize a diverse range of search strategies, including reinforcement learning-based methods [39, 40], supernet-based techniques [20, 21, 24, 41, 42], evolutionary algorithms [22, 43–48], and gradient-based optimization approaches [22, 49–53], among others.

Previous efforts to enhance the training stability of GANs have primarily relied on manual modifications to network architectures. However, such methods are labor-intensive, heavily dependent on expert intuition, and often require extensive trial-and-error. To address these limitations, NAS has been introduced to automatically discover effective GAN architectures for image generation tasks [54–58]. Inspired by this, the concept of architecture search has gradually extended into quantum machine learning. Nevertheless, this area remains in its infancy.As far as we are aware, the application of QAS to QGANs has been explored in only a limited number of studies [59, 60], most of which utilize genetic algorithms. Overall, a systematic framework that integrates QAS into QGANs for the automatic design of optimal quantum generator architectures is still largely underexplored in the current literature.

3. QAS-QGAN framework

Figure 1 illustrates the overall structure of the proposed QAS-QGAN framework, comprising two primary stages. In the first stage, a SuperCircuit-based QAS approach is employed to automatically identify a quantum generator architecture that is well-suited for the QGAN's image generation task. In the second stage, the optimized generator is integrated into the QGAN framework and trained end-to-end to generate images. The following sections provide a detailed description of the methods and processes for each stage of the framework.

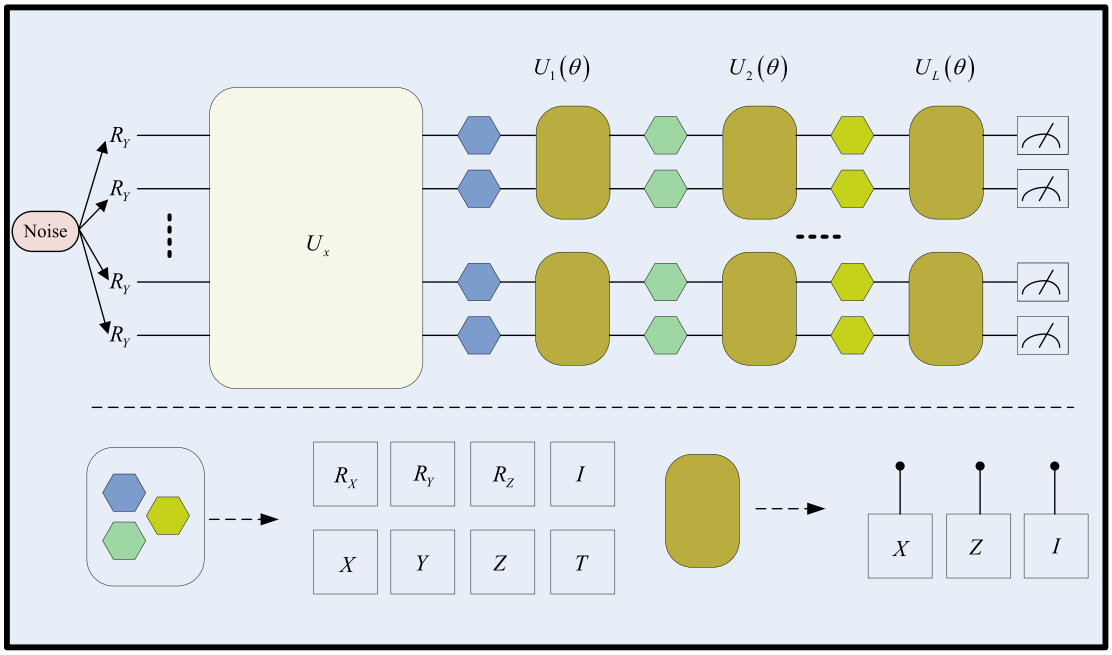

Figure 2. Paradigm of the SuperCircuit [20]. Reproduced from [21], with permission from Springer Nature.

3.1. Searching for the quantum generator

3.1.1. Construction of the supercircuit

In the quantum architecture search stage, we begin by specifying the input, target output, and loss function for the SuperCircuit, which serve as the basis for optimizing its internal parameters during training. Prior to outlining the training configuration, we first describe the construction of the SuperCircuit. A schematic diagram of its structure is shown in figure 2. The choice of a 5-qubit, 6-depth SuperCircuit in our work was motivated both by prior evidence and by practical considerations of training cost and hardware feasibility. Specifically, our choice of 5 qubits was inspired by our previous work CEQAS [21]. Regarding the circuit depth, we referred to prior studies (e.g. [61]) which demonstrated that increasing the depth of parameterized quantum circuits enhances their expressive power up to a certain point, beyond which the improvement saturates. In our setting, given that we increased the number of qubits relative to [61], we chose a depth of 6 as a reasonable balance to ensure sufficient expressivity while avoiding unnecessary complexity.

The SuperCircuit defines a searchable pool of quantum circuit candidates, where each circuit shares parameters according to a predefined parameter-sharing strategy. In the schematic, all potential single-qubit gates are represented by bright hexagonal nodes, while possible two-qubit gates are illustrated as orange rectangles. Building on prior research [20, 21], the input to the SuperCircuit consists of a noise vector sampled from a standard normal distribution. This noise is first encoded using an R_Y gate and then processed through a data encoding layer U_x, which embeds the input into the circuit for further transformation.

3.1.2. PCA feature distribution mechanism

The target output of the SuperCircuit is obtained from the original dataset after dimensionality reduction via PCA, followed by reordering using a tailored PCA Feature Distribution Mechanism. This mechanism serves to reduce the search space for the quantum architecture search algorithm by reorganizing the principal components in a variance-balanced manner. For instance, consider a dataset reduced to 40 dimensions via PCA. Directly generating such outputs would require a 40-qubit quantum circuit, resulting in an intractably large search space. To address this, the 40 components are partitioned into 8 subsets, each containing 5 dimensions, such that the explained variance is approximately balanced across groups. One such subset–comprising 5 principal components–is then used as the target output for a single SuperCircuit, which requires only 5 qubits to implement. By implementing this grouping strategy, the overall search space is significantly reduced, leading to a decrease in the computational cost of the architecture search process.

The process begins by allocating the highest principal component from the PCA-reduced feature vector to the first group, followed by the second component to the second group, and continuing in this manner to ensure that each group is initially assigned one of the most significant components. Subsequently, the remaining $n-1$ least significant principal components are assigned in reverse order: the subsequent $n-1$

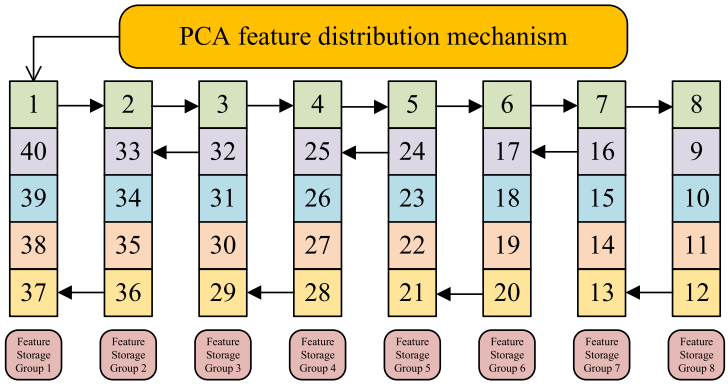

Figure 3. The Mechanism of PCA Feature Allocation: A Case Study of Dividing 40-Dimensional Data into 8 Groups.

components are allocated to the first group, followed by the second-to-last $n-1$ components to the second group, and continuing in this manner until all components are distributed, with n indicating the desired dimensionality for each group. This allocation strategy, illustrated in figure 3, ensures that each feature group contains both high-variance and low-variance components. As a result, feature importance is more uniformly distributed across substructures, preserving representational diversity while maintaining balanced information content.

3.1.3. Training details of the supercircuit

In this phase, we apply the SuperCircuit-based QAS framework to address a regression task. The dataset, denoted as D_{data}, contains 500 training samples. For each sample $\left(x^{(i)}, y^{(i)}\right)$, the input $x^{(i)}$ is sampled from a standard normal distribution, $x^{(i)} \sim \mathcal{N}(0,1)$, while the target label $y^{(i)}$ is selected from Feature Storage Group 1, as depicted in figure 3.

The dataset is divided into training and testing sets with a 4:1 ratio. The optimization objective is to reduce the discrepancy between the output of the SuperCircuit and the target labels. The specific form of the objective function used to optimize the SuperCircuit is:

$$L = \frac{1}{N}\sum_{i=1}^{N} \text{MSE}\left[S\left(x^{(i)}, \theta\right), y^{(i)}\right]. \tag{2}$$

In the above expression, N corresponds to the number of samples in either the training or testing set, while MSE refers to the Mean Squared Error function. The term $S(x^{(i)}, \theta)$ represents the output generated by the SuperCircuit for input $x^{(i)}$, with θ being the set of trainable parameters within the circuit.

The training of the SuperCircuit follows the hyperparameter configuration outlined in previous studies [20, 21], which have shown that expanding the number of parameter-sharing configurations enhances the architecture search efficiency. The number of shared parameter configurations in the SuperCircuit is set to $W = 5$, with both the input and output dimensions being 5. To meet these dimensional constraints, we employ 5 qubits. The circuit is trained for 10 000 iterations using a learning rate of 0.05, and the SuperCircuit's depth is set to $d = 6$, consistent with the quantum generator's depth in the QGAN. These settings form the primary configuration for training the SuperCircuit.

After training the SuperCircuit, we proceed to sample SubCircuits from the trained architecture. Each SubCircuit inherits its gate parameters from the optimized parameter space obtained during training. These candidate circuits are then evaluated on the test dataset to assess their performance. Rather than employing random sampling, we utilize the NSGA-II genetic algorithm, which has been demonstrated in prior studies [20, 21] to be more effective in identifying high-performing quantum circuits. The NSGA-II configuration involves 4000 circuit evaluations, a population size of 50, and 80 generations. This procedure ultimately yields a set of well-performing quantum circuits, from which the generator architecture for the QGAN is constructed.

3.2. Training phase for image generation with QGAN

This section outlines the utilization of optimal quantum circuits, obtained through QAS, in the QGAN framework for image generation. The corresponding data processing pipeline is depicted in figure 4. As introduced in section 3.1.2, the PCA Feature Distribution Mechanism ensures that each feature group contains a combination of high-variance principal components and low-variance secondary components.

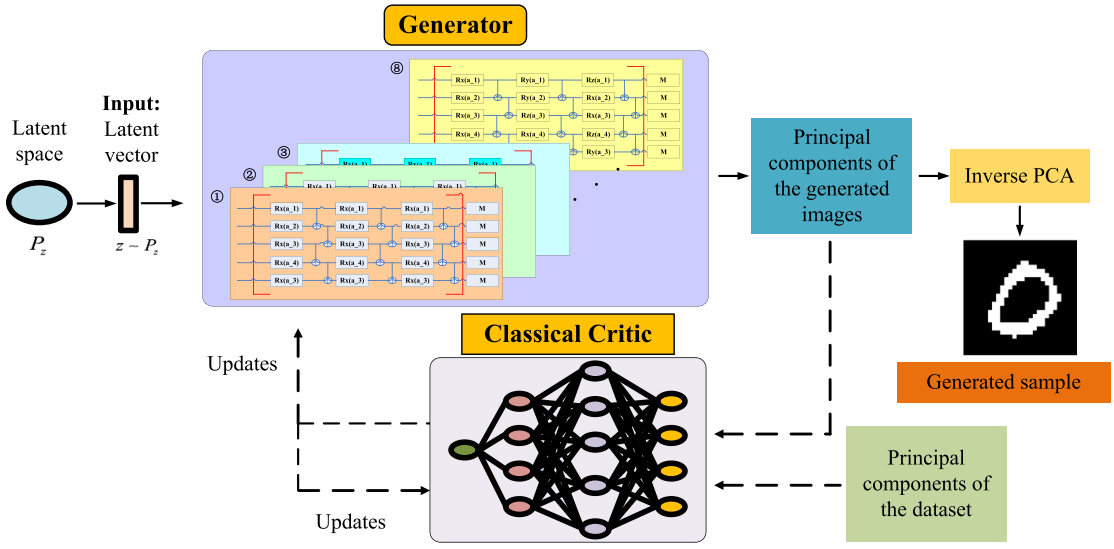

Figure 4. The QGAN employs the searched quantum circuit to generate images.

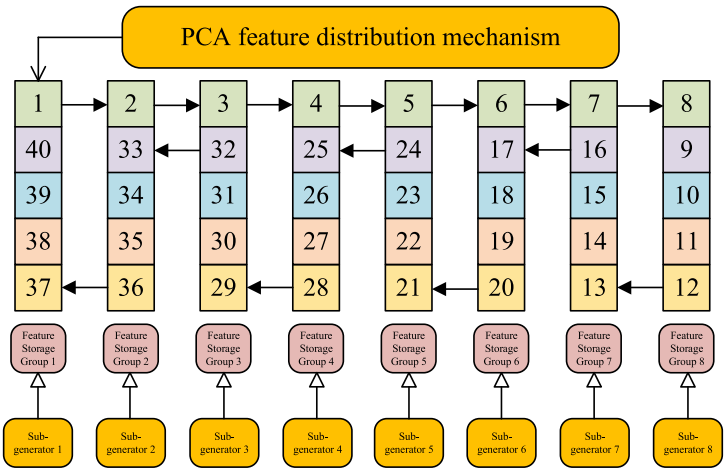

Figure 5. Task-specific feature allocation for each sub-generator.

This design maintains feature diversity while promoting a balanced distribution of feature importance across substructures. Visualization of the transformed data confirms that the distributions across feature groups exhibit high similarity. Motivated by prior work [1, 7, 8, 15], we adopt a modular generator architecture, wherein the quantum generator is decomposed into multiple sub-generators. In accordance with the dimensional structure defined by the PCA Feature Distribution Mechanism, the generator is partitioned into 8 sub-generators, with each handling the generation of a distinct feature group. The feature allocation handled by each sub-generator is illustrated in figure 5.

In this architecture, each sub-generator is tasked with independently modeling and generating features within a specific subspace of the overall feature space. As shown in figure 5, Sub-generator 1 corresponds to Feature Storage Group 1, Sub-generator 2 to Group 2, and so on, up to Sub-generator 8, which handles Group 8. By decomposing the high-dimensional feature space into multiple subspaces and enabling parallel generation through distinct sub-generators, the architecture preserves the expressive capacity of the full model while enhancing its ability to capture fine-grained, localized features. This modular feature distribution strategy ensures a balanced allocation of principal component information across sub-generators, mitigating training bias and preventing the dominance of any single component. It also guarantees that each sub-generator plays an active and meaningful role during training. Furthermore, this mechanism addresses hardware-related limitations in NISQ devices. Concentrating critical features within a single sub-generator that maps to qubits with high error rates could severely impact performance. By

promoting a more uniform feature allocation, the model reduces dependency on specific qubits, thereby improving training stability and overall generation quality [1].

To further mitigate mode collapse, we adopt the adaptive noise input strategy introduced in previous works [1, 62]. In this approach, the input noise is dynamically adjusted based on the training state, following the formulation defined in equation (3)

$$\textbf{Noise}_{\text{adaptive}} = \frac{\pi}{8} + \frac{5\pi}{8} \cdot \text{ReLU}\left(\tanh\left(\frac{D_L}{G_L} - \frac{G_{L_0}}{D_{L_0}}\right)\right). \tag{3}$$

ReLU is a widely adopted nonlinear activation function in machine learning, which sets all negative input values to zero while maintaining positive values. The Tanh function is used to scale the adaptive noise into the $[-1, 1]$ range. Here, G_L and D_L refer to the current losses of the generator and discriminator, respectively, while G_{L_0} and D_{L_0} represent their initial losses in the early stages of training. These values are then used to compute the generator–discriminator loss ratio, which facilitates the adaptive adjustment of the noise input. As highlighted in previous QGAN studies [62], various noise ranges and activation functions have been tested, and these components can be adapted to meet the specific requirements of different tasks.

4. Experiments

4.1. Preliminaries

The experiments were conducted using Python 3, with support from PyTorch [63], PennyLane [64], and Qiskit [65]. We assessed the image generation capabilities of the proposed QAS-QGAN framework on three widely used benchmark datasets: MNIST [25], Fashion MNIST [26], and CelebA [31]. All experiments reported in this section were conducted in a noiseless quantum simulation environment. Additional results obtained under noisy conditions are presented in appendix A.

In this section, we present a thorough analysis of the experimental results derived from the QAS-QGAN framework. As described in section 3.1.1, the SuperCircuit was constructed using five qubits. Single-qubit operations are performed using R_X, R_Y, and R_Z gates, while two-qubit interactions are realized through CNOT gates. The hyperparameter settings used for training the SuperCircuit are provided in section 3.1.3. With this configuration, the QAS mechanism is able to effectively search for quantum circuit structures that are well-suited for the generator component of the QGAN model.

To evaluate the effectiveness of the quantum generator discovered via our QAS method, we compare QAS-QGAN against several recent models, including classical and hybrid quantum–classical GANs. The baseline models considered are PQWGAN [8], WGAN-GP [66], Quantum AnoGAN [16], MosaiQGAN [1], and QINR-QGAN [9]. PQWGAN features a hybrid architecture, with a quantum circuit serving as the generator and a classical fully connected neural network as the discriminator. Conversely, WGAN-GP is a classical GAN model, where both the generator and discriminator are based on classical neural networks. We implement two variants of WGAN-GP for comparison: (1) WGAN-GP-Little, featuring a generator with two hidden layers of 512 and 256 neurons, respectively; and (2) WGAN-GP-Big, whose generator contains four hidden layers with 512, 256, 256, and 128 neurons. Quantum AnoGAN is another hybrid model with a generator composed of both quantum and classical layers, while its discriminator remains entirely classical. MosaiQGAN incorporates PCA into the image generation pipeline, similar to our method. QINR-QGAN extends the generator with Quantum Implicit Neural Representation, yet retains a classical discriminator. In all models, including our proposed QAS-QGAN, the discriminator is implemented as a fully connected neural network. A summary of the architectural configurations (Learning Rates) for the models under comparison is presented in table 1.

For a comprehensive evaluation of the generated images, we use four well-established metrics: FID [27], SSIM [28], Cosine Similarity [29], and PSNR [30]. Together, these metrics ensure a thorough and objective assessment. FID evaluates the distance between the generated and real image distributions, with lower values reflecting higher fidelity. SSIM measures structural similarity, where higher values suggest better preservation of spatial patterns. Cosine Similarity gauges the alignment of feature vectors, with values closer to 1 indicating stronger similarity. PSNR assesses image quality by comparing signal power to noise, with higher values indicating superior visual quality.

4.2. Experimental results on the MNIST dataset

We conducted experiments on the MNIST dataset, starting with the application of QAS to identify quantum circuit architectures optimized for generating specific digits. For instance, to design a generator tailored for digit '0,' we first extract all samples labeled as '0' from the dataset. These samples are then processed using PCA, followed by the PCA Feature Distribution Mechanism, to produce feature representations aligned with

Table 1. Hyperparameter Configurations (Learning Rates) for Baseline Models Across Datasets (G: Generator, D: Discriminator or Critic).

	MNIST		Fashion MNIST		CelebA	
Models	G	D	G	D	G	D
PQWGAN [8]	0.01	0.0002	0.01	0.0002	—	—
WGAN-GP-Little [66]	0.000 07	0.000 12	0.000 07	0.000 12	—	—
WGAN-GP-Big [66]	0.000 07	0.000 12	0.000 05	0.000 12	—	—
Quantum AnoGAN [16]	0.000 01	0.0001	0.000 01	0.0001	—	—
MosaiQGAN [1]	0.3	0.05	0.3	0.05	—	—
QINR-QGAN [9]	0.000 01	0.000 12	0.000 01	0.000 12	0.000 01	0.000 12
QAS-QGAN (Ours)	0.3	0.05	0.3	0.05	0.3	0.05

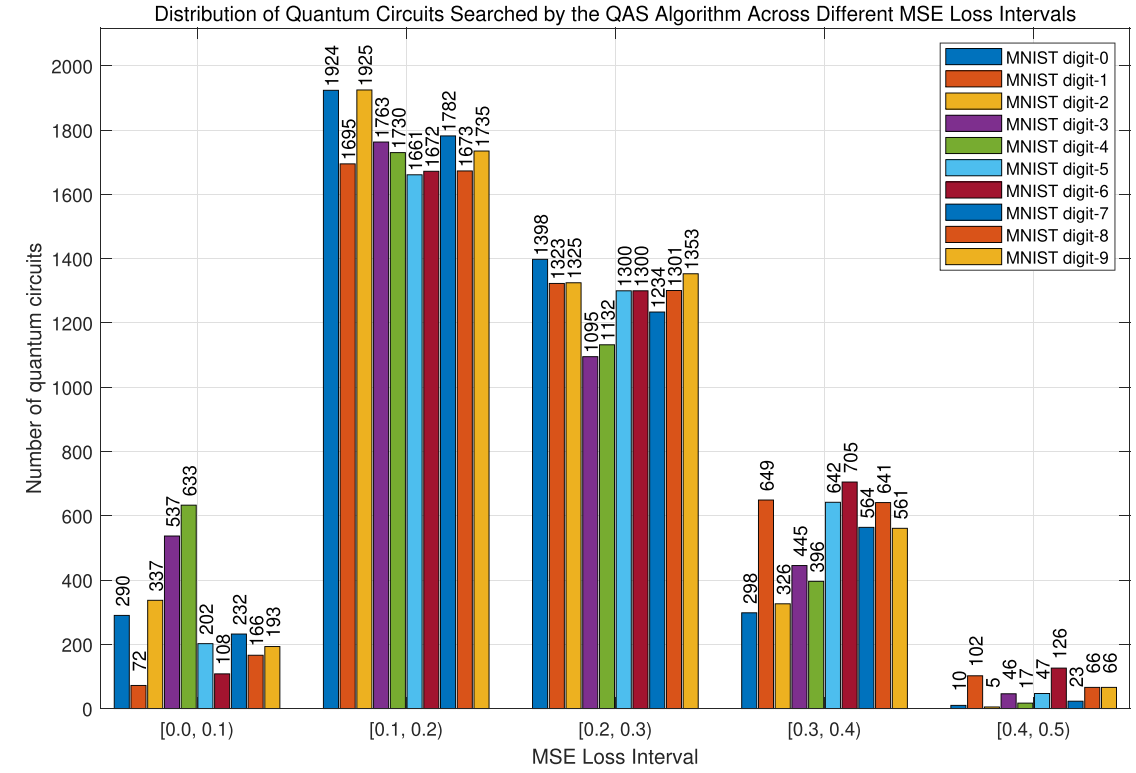

Figure 6. The number of quantum circuits searched by QAS within different MSE loss intervals on the MNIST dataset.

the output dimensionality requirements of the QAS stage. A total of 500 samples were selected in this experiment for preprocessing and the subsequent circuit search process.

Figure 6 presents the distribution of quantum circuits discovered by the QAS algorithm on the MNIST dataset, categorized by MSE intervals. The results demonstrate that, aided by our tailored data preprocessing pipeline, QAS efficiently identifies circuits that meet the desired performance threshold. Notably, a substantial proportion of the circuits fall within the [0.0, 0.1] MSE range, indicating high-quality solutions. As the MSE threshold increases, the number of corresponding circuits declines progressively, further validating the effectiveness and precision of the proposed approach.

Figure 7 shows the best-performing quantum circuit architecture identified by the QAS algorithm, trained on data derived from the MNIST dataset. While the figure illustrates the circuit used by a single sub-generator, this same architecture is uniformly applied across all eight sub-generators in the QGAN framework. It should be noted that the gate parameters depicted in the figure are randomly generated for visualization purposes and do not correspond to the actual trained values obtained from the SuperCircuit. This architectural consistency aligns with our design rationale described earlier in the PCA Feature Distribution section. Additionally, the best-performing quantum circuits obtained for each MNIST digit

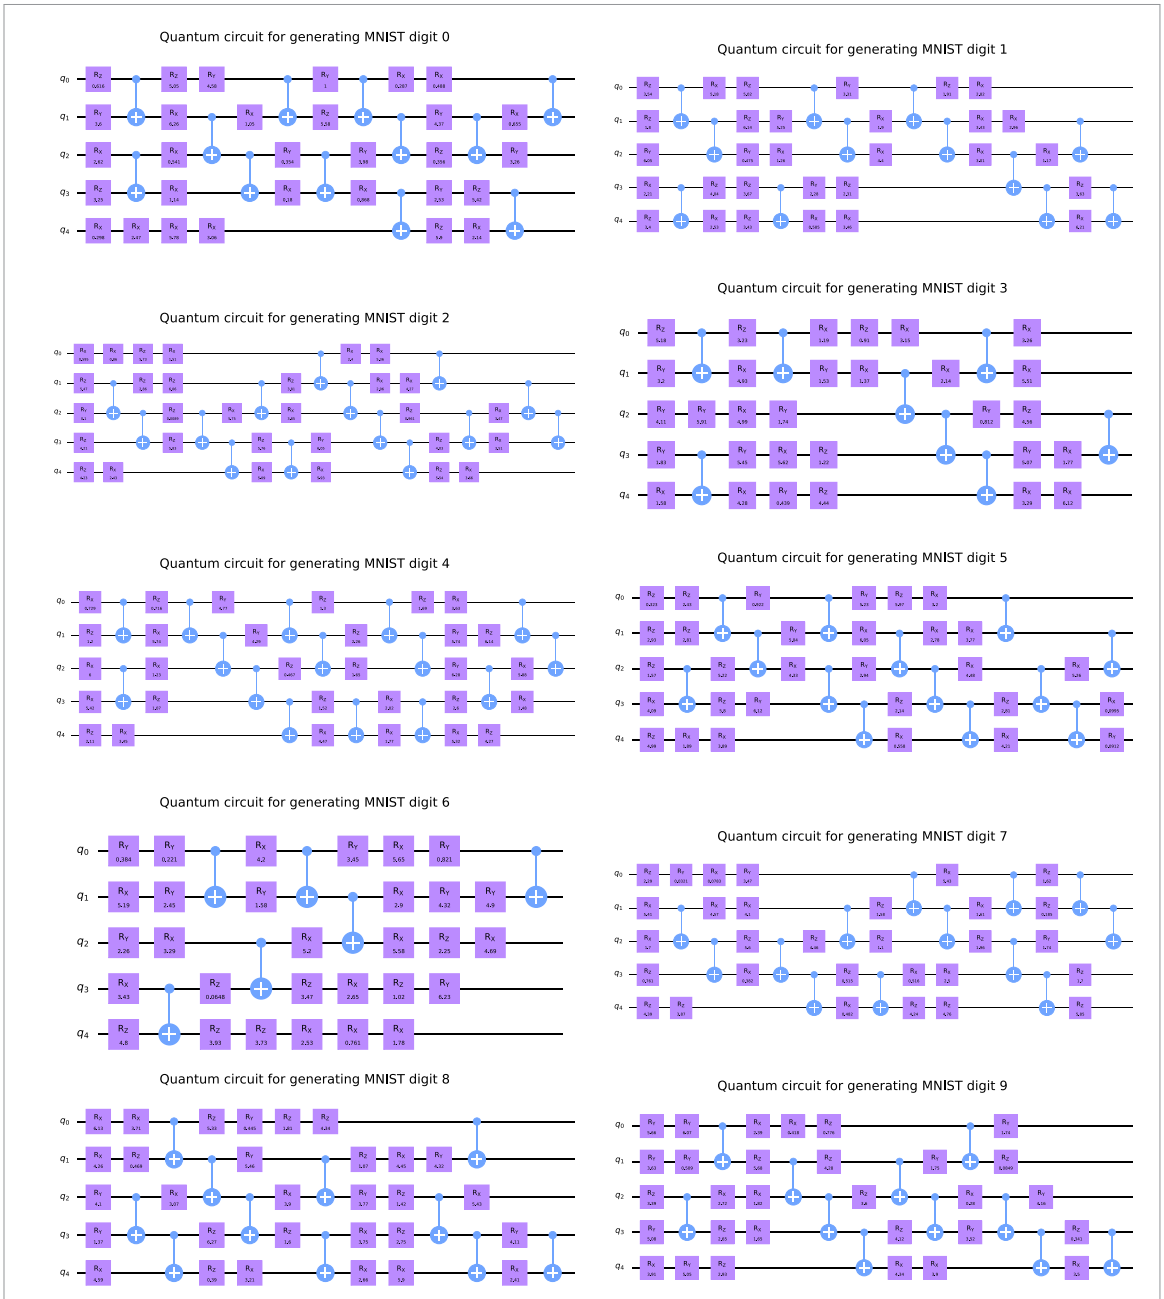

Figure 7. The best-performing quantum circuit obtained by QAS using the training data constructed from the MNIST dataset.

class under noisy simulation conditions are provided in appendix B. The selected optimal circuit is subsequently used as the generator core for image synthesis within the QGAN.

Figure 8 presents the experimental results of all compared models on the MNIST dataset. As shown in figure 8, each model successfully completes the image generation task following training convergence. For illustration, we display a subset of digit classes–specifically 0, 1, 3, 6, 7, and 9. The first row ('Train data') contains real handwritten digits used as ground-truth references, while the subsequent rows display images generated by each model under identical input conditions. The results reveal noticeable differences in visual quality, structural fidelity, and digit recognizability across models. WGAN-GP-Little generates noisy and blurry outputs, with substantial distortion particularly evident in digits such as 6 and 9. In comparison, WGAN-GP-Big shows improved stability and produces sharper digit contours. Hybrid models including PQWGAN, Quantum AnoGAN, and QINR-QGAN demonstrate stronger preservation of image structure and contrast, yielding more recognizable digits. However, PQWGAN still suffers from noticeable fuzziness–some digit strokes appear incomplete or overly smoothed. Both MosaiQGAN and QAS-QGAN exhibit superior structural preservation, producing outputs with high contrast and clearly defined edges.

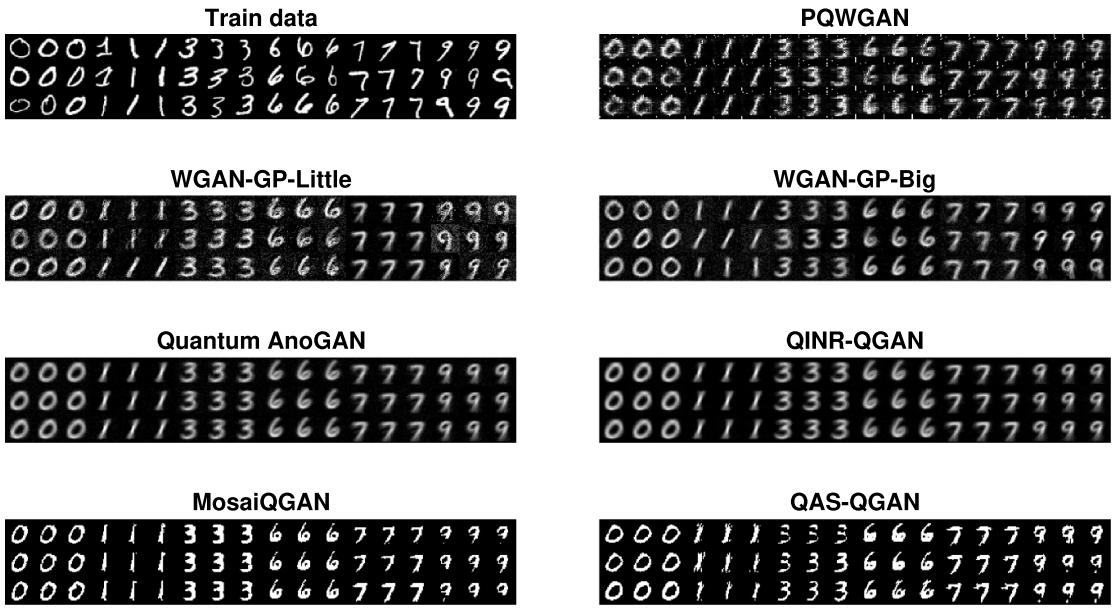

Figure 8. Original MNIST images and those generated by PQWGAN, WGAN-GP, Quantum AnoGAN, QINR-QGAN, MosaiQGAN, and QAS-QGAN.

While QAS-QGAN occasionally shows minor deformation (e.g. in digit 9), it consistently achieves high visual quality and strong digit fidelity overall.

In addition to visual analysis, we carried out a quantitative evaluation of image quality across all models to provide a comprehensive and objective performance assessment. Figure 9 compares the results of four standard metrics–FID, SSIM, Cosine Similarity, and PSNR–across all ten digit classes in the MNIST dataset. Overall, QAS-QGAN outperforms the baseline models in FID, SSIM, and PSNR, indicating superior fidelity, structural similarity, and signal clarity. While its performance in Cosine Similarity is marginally lower than that of a few competing models, the difference is relatively small and does not significantly affect the overall evaluation. Notably, QAS-QGAN achieves comparable or better performance than MosaiQGAN, despite the latter employing a non-optimized quantum generator structure. This result demonstrates that circuit architectures discovered through QAS can be effectively optimized without compromising image quality. Moreover, QAS-QGAN achieves this with fewer quantum gates, highlighting its advantage in quantum resource efficiency. A detailed analysis of this aspect is provided in the subsequent discussion section.

4.3. Experimental results on the fashion MNIST dataset

We further evaluated the QAS-QGAN framework on the Fashion MNIST dataset, beginning with quantum circuit discovery via QAS. For example, to design a generator specialized in producing 'T-shirt' images, we first extracted corresponding samples from the dataset. These were then processed using PCA and the PCA Feature Distribution Mechanism to construct feature vectors aligned with the output dimensionality required by QAS. Based on this preprocessed data, QAS was used to search for an optimal circuit architecture tailored to the 'T-shirt' category. In this experiment, 500 samples from the target class were selected for training and evaluation.

Figure 10 illustrates the distribution of quantum circuits identified by the QAS algorithm on the Fashion MNIST dataset, grouped by MSE loss intervals. Although certain categories exhibit a relatively small number of high-performing circuits–defined as those with MSE loss in the range $[0, 0.12)$–the overall quantity remains sufficient to support subsequent model training and practical deployment.

Figure 11 shows the optimal quantum circuit structure identified by QAS for the Fashion MNIST dataset, following training data construction. (Note: The parameters shown in the figure are randomly generated for visualization purposes and do not reflect the actual values learned by the SuperCircuit.) Under noisy simulation condition, the optimal circuits corresponding to each Fashion MNIST category were summarized and are provided in appendix B. The selected circuit was subsequently integrated into the generator of QGAN to carry out the image generation task.

Figure 12 displays the results of image generation from different models on the Fashion MNIST dataset. In terms of visual quality, both WGAN-GP-Little and PQWGAN produce images with noticeable noise and

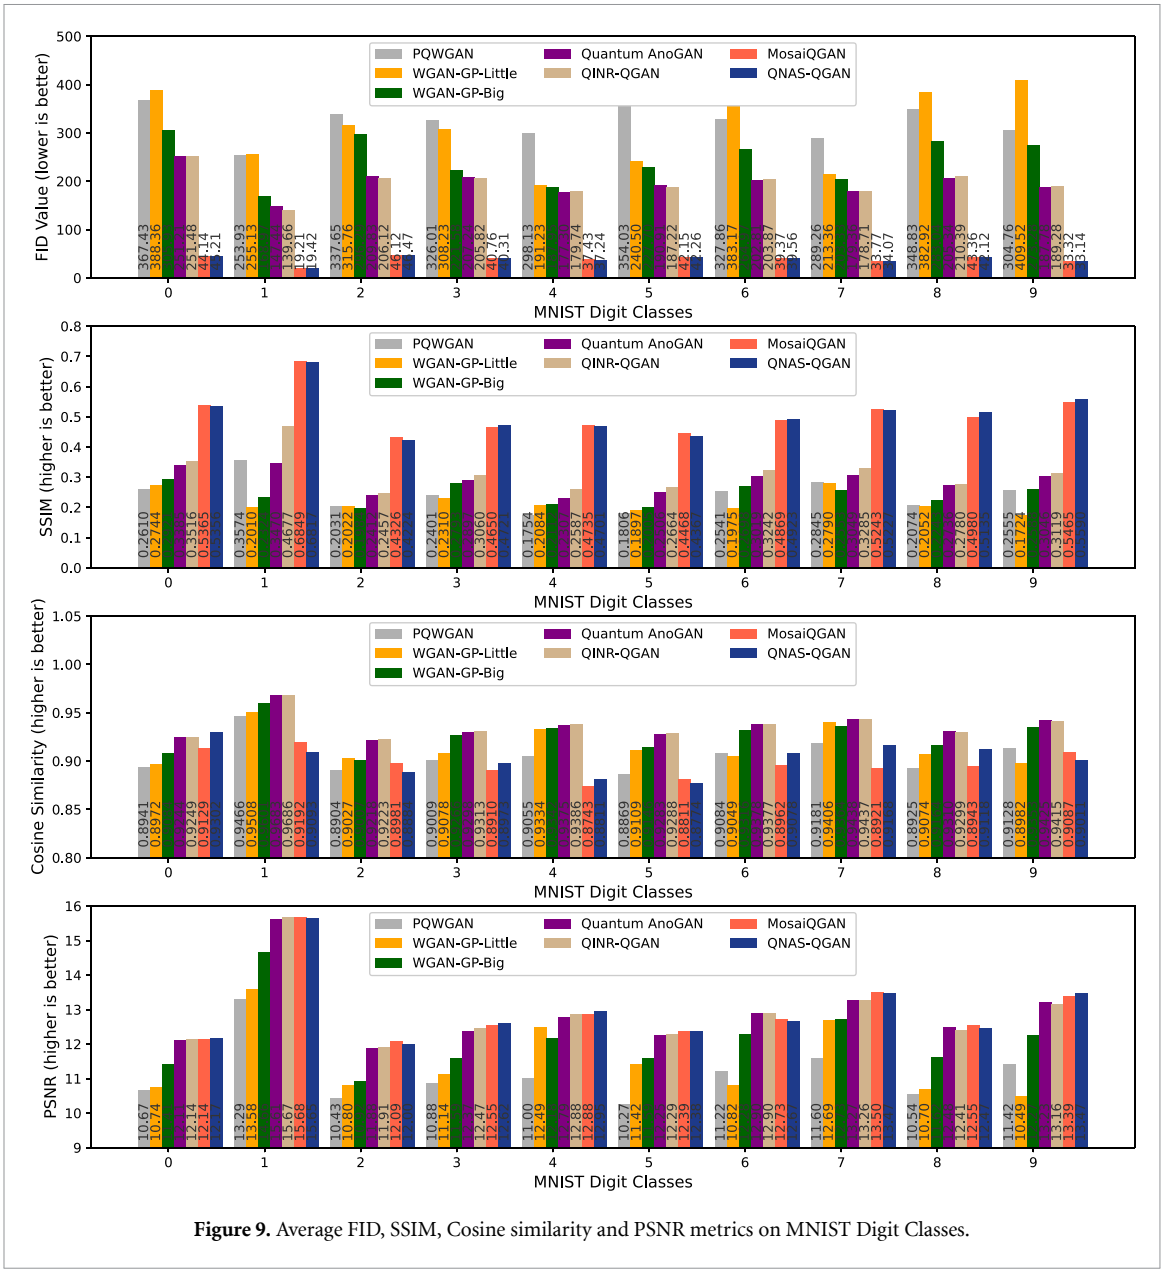

Figure 9. Average FID, SSIM, Cosine similarity and PSNR metrics on MNIST Digit Classes.

blurriness. Although WGAN-GP-Big alleviates these issues to some extent, it still exhibits instability in rendering fine-grained details. In contrast, Quantum AnoGAN and QINR-QGAN demonstrate more consistent performance, producing images with well-defined contours. Their outputs are particularly effective for structured clothing items such as shirts and trousers, preserving symmetry and edge continuity. MosaiQGAN and QAS-QGAN further distinguish themselves by generating images with high contrast and sharp contours. While QAS-QGAN may slightly underperform MosaiQGAN in reproducing intricate textures, it achieves higher overall recognizability and maintains strong visual clarity across categories.

As shown in the experimental results in figure 13, QAS-QGAN outperforms all baseline models across the full range of evaluation metrics, demonstrating consistently strong performance. In contrast to its relatively weaker results in Cosine Similarity on the MNIST dataset, QAS-QGAN exhibits a marked improvement on the Fashion MNIST dataset.

This enhancement is particularly evident in clothing-related categories such as ‘T-shirt,’ ‘Dress,’ and ‘Coat,’ where QAS-QGAN achieves significantly higher Cosine Similarity scores. These findings suggest greater consistency in the feature representation space and more accurate modeling of structural characteristics in complex visual patterns.

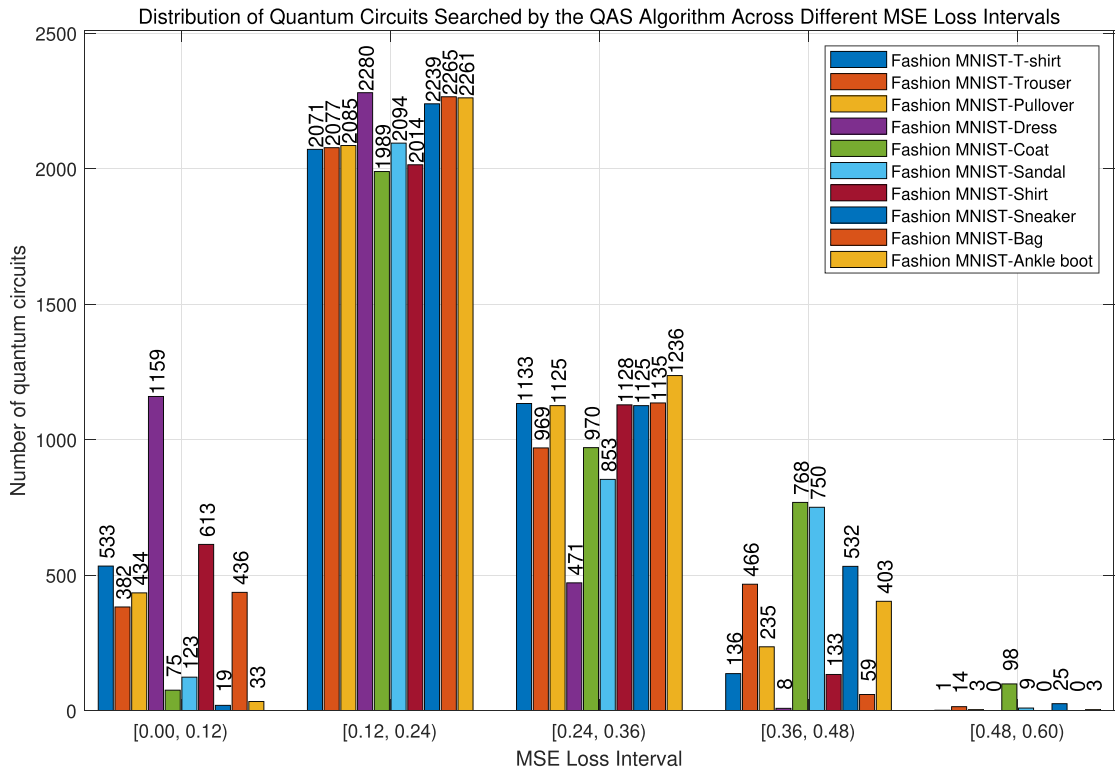

Figure 10. The number of quantum circuits searched by QAS within different MSE loss intervals on the Fashion MNIST dataset.

4.4. Experimental results on the celebA dataset

We further evaluated the performance of QAS-QGAN on the CelebA dataset by generating higher-resolution facial images at 109×89 pixels. Compared to previous studies that typically generated facial images at 78×64 resolution, this experiment marks a substantial increase in image scale.

During data preprocessing, facial images with the same attribute label were grouped into training categories. For classes with fewer than 500 samples, the dataset was augmented by duplicating and randomly shuffling existing images until each category contained exactly 500 samples–ensuring consistency in training size across categories.

Figure 14 shows the distribution of quantum circuits discovered by QAS across different MSE loss intervals on the CelebA dataset. For the selected facial attribute categories, QAS successfully identifies a substantial number of high-quality circuits with MSE losses in the range $[0, 0.13)$, demonstrating its robust search capability even in the presence of complex, high-dimensional data distributions.

Figure 15 illustrates representative optimal quantum circuits obtained by QAS for several facial categories. Each circuit consists of multiple rotation gates–R_X, R_Y, and R_Z–as well as CNOT gates. The resulting architectures are compact and structurally stable, aligning effectively with the specific generative tasks. Detailed circuit configurations under noisy conditions are provided in appendix B.

Figure 16 compares the facial image generation performance of QAS-QGAN and QINR-QGAN on the CelebA dataset. Real training samples are shown on the left, with images generated by each model on the right. QAS-QGAN produces higher-resolution outputs (109×89 pixels), surpassing the 78×64 pixel images generated by QINR-QGAN, and more effectively reconstructs facial structures. Although minor blurriness and localized noise can still be observed around facial features, the overall contour and semantic integrity of the faces are well preserved.

Moreover, QAS-QGAN achieves a superior trade-off between detail preservation and noise suppression, offering notable improvements in visual consistency, structural fidelity, and reduced distortion. While the increased resolution introduces minor noise artifacts, these do not significantly impair the perceptual quality of the generated images.

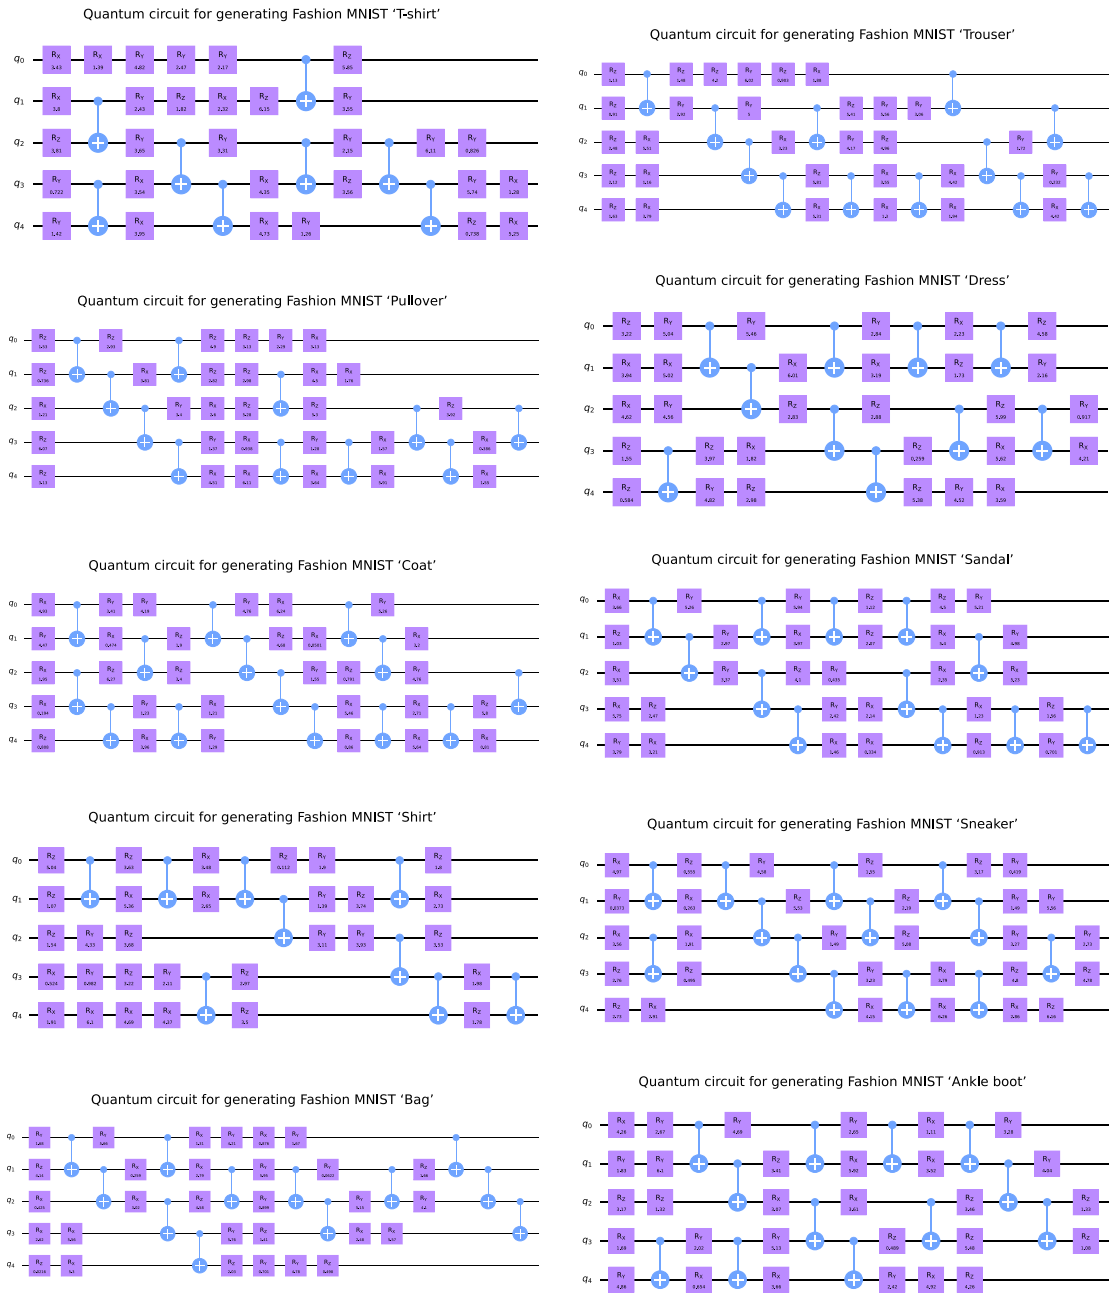

Figure 11. The best-performing quantum circuit obtained by QAS using the training data constructed from the Fashion MNIST dataset.

5. Discussion

This paper presents QAS-QGAN, a novel framework that integrates QAS into the generator design of QGANs. Experimental results on datasets such as MNIST, Fashion MNIST, and CelebA demonstrate that QAS-QGAN effectively identifies a wide variety of high-performing quantum circuit architectures. Even in experiments with certain data labels, where fewer high-performing circuits were found, the number of suitable circuits remained sufficient for the tasks. When applied to image generation tasks, the model consistently exhibits clear advantages in image quality, training stability, and quantum resource efficiency.

In this section, we present a deeper analysis of our findings, focusing on the following key aspects as revealed through the experimental results:

The Value of Automated Structure Search for QGANs: Traditional QGAN models often use manually crafted quantum circuits for their generators, which restricts scalability and flexibility across different tasks. In this work, we present a QAS method based on the SuperCircuit framework, which facilitates the automatic

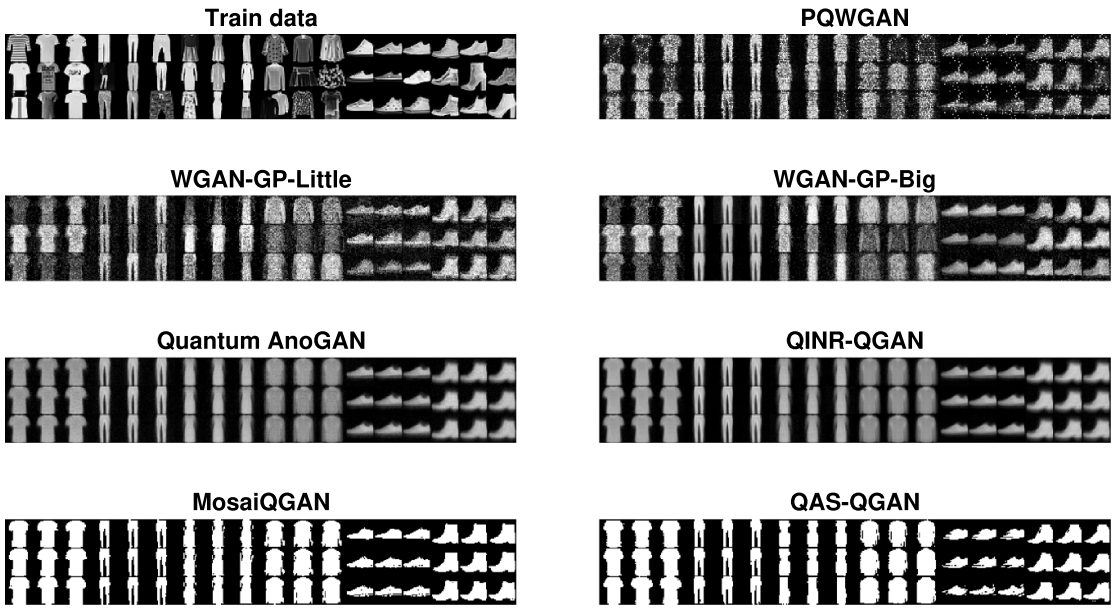

Figure 12. Original Fashion MNIST images and images generated by PQWGAN, WGAN-GP, Quantum AnoGAN, QINR-QGAN, MosaiQGAN, and QAS-QGAN.

creation of optimal quantum circuits customized for generating images in QGANs. While QAS shows some limitations in improving the generalization capability of QGANs, it significantly reduces the complexity of circuit design. During the training of the SuperCircuit in QAS, we ensure that each layer and each line includes single-qubit rotation gates. The circuits identified by QAS contain a similar number of single-qubit gates as those in MosaiQGAN. However, in contrast to MosaiQGAN, which exclusively uses the R_Y gate, QAS investigates a more diverse set of single-qubit gates, such as R_X, R_Y, and R_Z. We present in table 2 the number of two-qubit quantum gates used in the quantum generators of both MosaiQGAN and QAS-QGAN during image generation across various datasets. Figure 17 also presents the general quantum generator circuit architecture used in MosaiQGAN.

As shown in table 2, the number of two-qubit gates used in the quantum circuits discovered by QAS-QGAN is reduced by 52.08% on the MNIST dataset and 52.92% on the Fashion MNIST dataset, compared to MosaiQGAN, significantly lowering resource consumption. The experimental findings demonstrate that QAS-QGAN results in more compact circuit architectures with fewer quantum gates, all while maintaining or potentially improving the quality of generated images. Taking MNIST and Fashion MNIST as examples, although the image quality produced by QAS-QGAN is comparable to that of MosaiQGAN–and in some classes even slightly better–QAS-QGAN uses fewer two-qubit gates, confirming its advantage in quantum resource efficiency. Overall, QAS-QGAN effectively reduces the complexity of the generator's quantum circuits without sacrificing generative performance. This reduction in gate count lessens dependence on NISQ devices, demonstrating both strong practicality and hardware adaptability. This is further supported by experiments conducted under noisy environments (see appendix A).

The integration of QAS with QGAN marks a significant advancement in QML, particularly in generative tasks. By automating the search for quantum circuits optimized for generation tasks, QAS not only reduces the manual effort required in traditional quantum circuit design but also enhances the adaptability of QGANs to diverse tasks and data distributions. This integration allows QGAN to automatically adjust its generator architecture, enabling the generation of high-quality images while reducing the computational complexity inherent in manually designing these circuits. Additionally, the use of QAS reduces the complexity of the generator circuit, making QGAN more efficient and suitable for implementation on quantum hardware with limited resources.

Balancing Image Generation Quality and Quantum Circuit Architecture: Across different datasets, experimental results show that QAS-QGAN surpasses baseline models in most image quality evaluation metrics. Although it falls slightly behind in cosine similarity on the MNIST dataset, it shows a reversal on the Fashion MNIST dataset–particularly excelling in clothing categories such as T-shirt, Coat, and Dress–demonstrating higher representational consistency. This observation suggests that the circuit

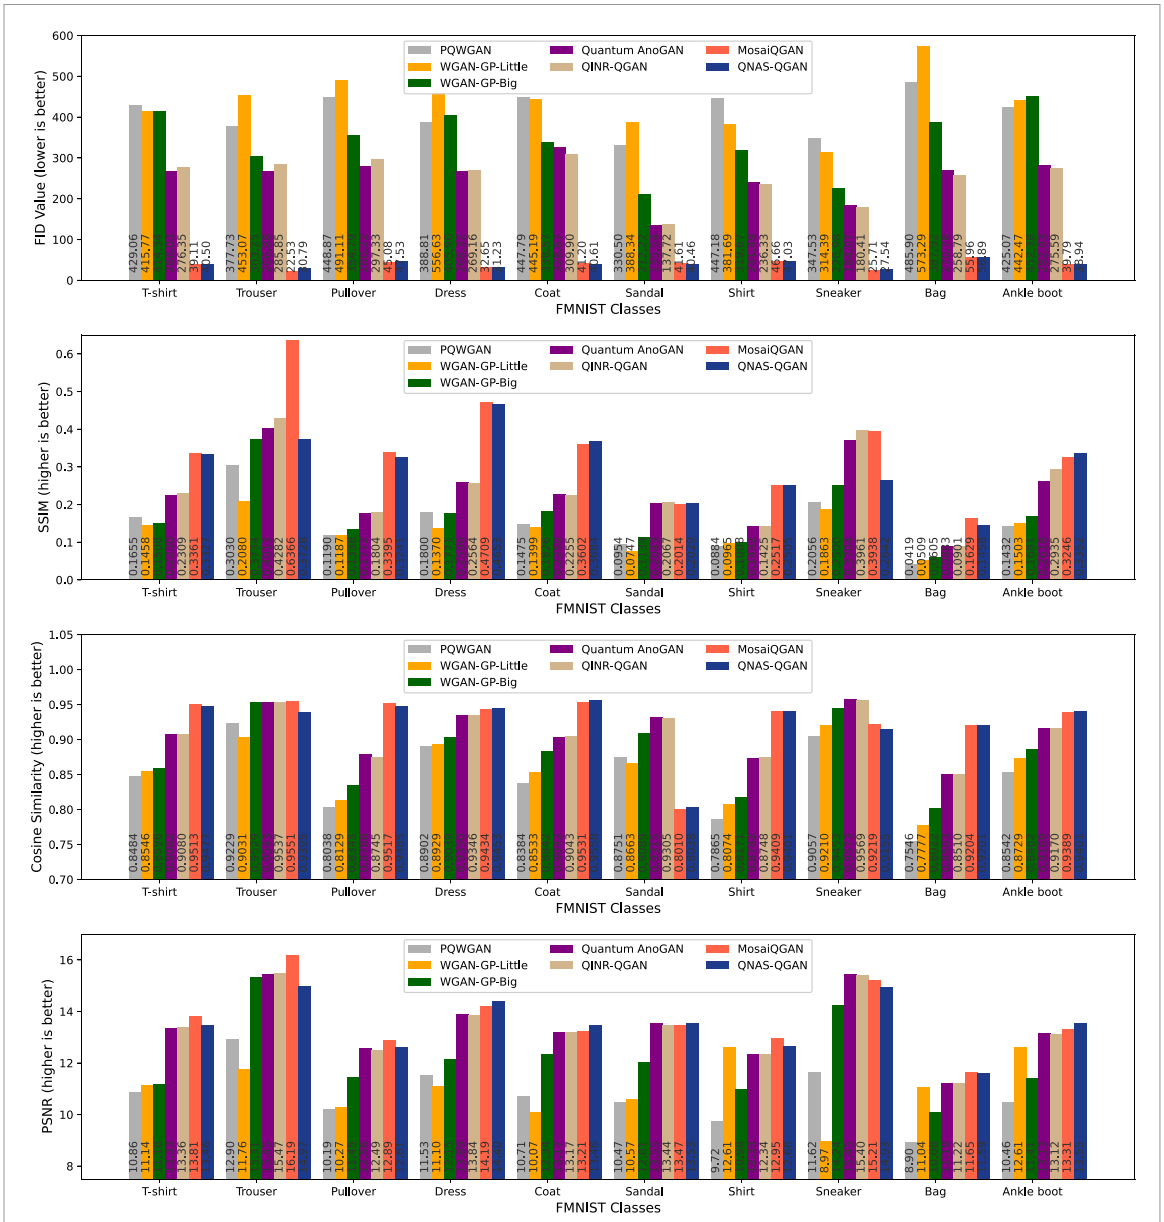

Figure 13. Average FID, SSIM, Cosine similarity and PSNR metrics on Fashion MNIST Classes.

structures generated by QAS possess a certain adaptability to data distribution, enabling the generator to automatically adjust and capture the core features of different image categories.

Improvement in Training Stability: In addition to enhancing quantum circuit design, another key benefit of integrating QAS into QGAN is the improvement in training stability. Traditional QGANs often face challenges such as mode collapse, where the generator fails to produce diverse outputs. This issue is especially prominent when quantum circuits are highly complex and manually designed. By using QAS to automatically optimize the generator architecture, combined with PCA and the PCA feature distribution mechanism, QAS-QGAN seems to alleviate these issues by finding circuit structures that better align with task-specific data distributions.

Feasibility and Challenges of High-Resolution Image Generation: We further explore the generation of higher-resolution facial images using the CelebA dataset, targeting a resolution of 109×89. Compared to the previously typical resolution of 78×64 in previous QGAN study, this represents a significant step forward. Experimental results show that QAS-QGAN performs comparably to QINR-QGAN in terms of detail preservation, with facial contours and structures well retained, despite the presence of some blurriness and noise. These findings indicate that the proposed method retains its modeling capacity at higher dimensions. Although the results indicate good performance, they also expose limitations in generating fine-grained

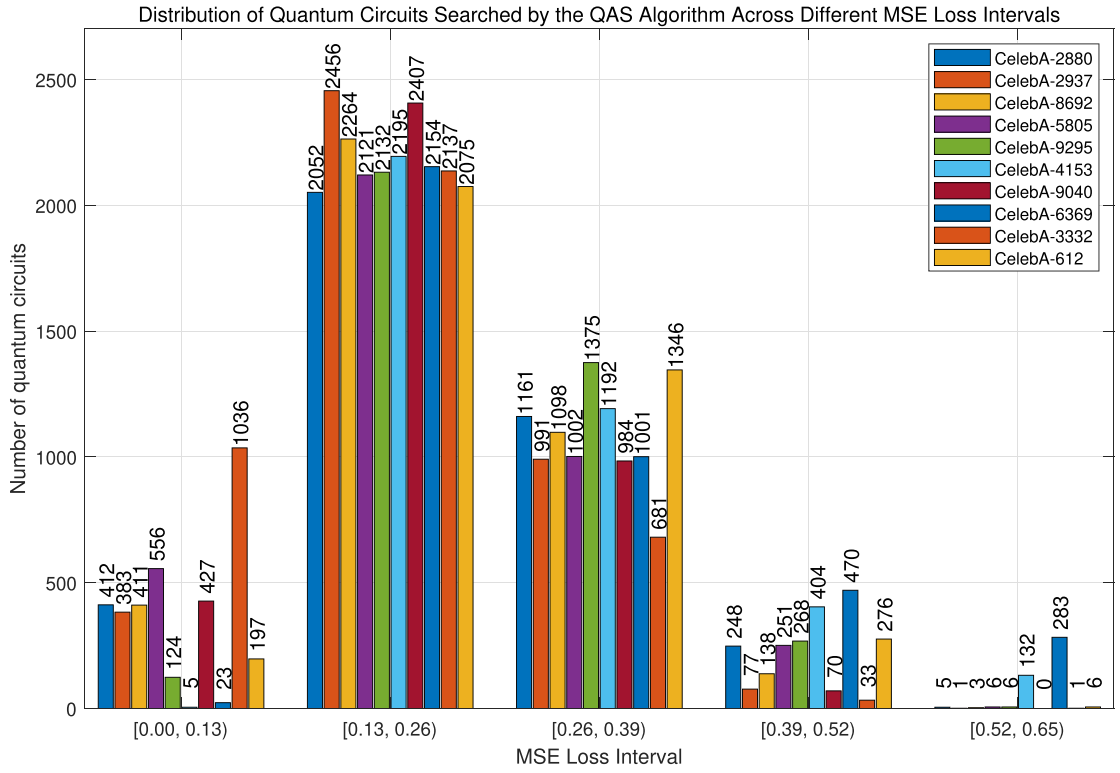

Figure 14. The number of quantum circuits searched by QAS within different MSE loss intervals on the CelebA dataset.

details and, more importantly, a lack of diversity in the generated samples. The model displays low sensitivity to variations in the latent space, which hinders its ability to generate diverse outputs. In future work, we aim to introduce quantitative diversity metrics–such as LPIPS, perceptual path length, and intra-class FID–to provide a more thorough evaluation of the model's generalization and sample diversity.

Other Dimensionality Reduction Methods: If we consider nonlinear dimensionality reduction techniques such as t-SNE or autoencoders, several trade-offs arise. t-SNE is effective for visualization; however, its stochastic nature and high computational cost limit its scalability in iterative QAS settings. In our practical deployment of QAS-QGAN, we employed PCA not only for dimensionality reduction but also for dimensionality restoration via inverse PCA, which can accurately preserve data information as long as eigenvalues are not discarded. By contrast, t-SNE lacks a direct and precise inverse transformation, making it less suitable for this purpose. Autoencoders can provide more expressive low-dimensional embeddings and may further improve generator performance or circuit compactness. Nonetheless, they introduce additional training overhead, since both encoder and decoder must be optimized to achieve reduction and restoration, whereas PCA accomplishes this without extra training.

Different quantum noise models: In our view, the quantum noise encountered in real devices is inevitably more complex than the simplified noise models typically used in simulations, such as bit-flip, depolarizing, or amplitude damping channels. To better validate the feasibility of our QAS algorithm on NISQ devices, it may be more appropriate to consider hybrid or composite noise models in simulation environments, which could serve as a more realistic approximation of physical noise. In the present work, our QAS-QGAN framework was primarily designed to integrate QAS with QGAN for the automatic discovery of generator circuits, with a focus on demonstrating the effectiveness of QAS in reducing two-qubit gate counts. Although we have not yet benchmarked performance under explicit noise processes, we expect that the reduction in circuit depth and two-qubit operations will inherently improve the resilience of QAS-discovered circuits compared to manually designed circuit.

Additional Limitations: We did not employ explicit diversity-oriented quantitative metrics such as entropy or coverage. We acknowledge that this represents a limitation of our study. Furthermore, in our broader investigation of quantum generative adversarial networks, we have observed that most existing QGAN works suffer from limited generalization capability due to the relatively small number of parameters in their

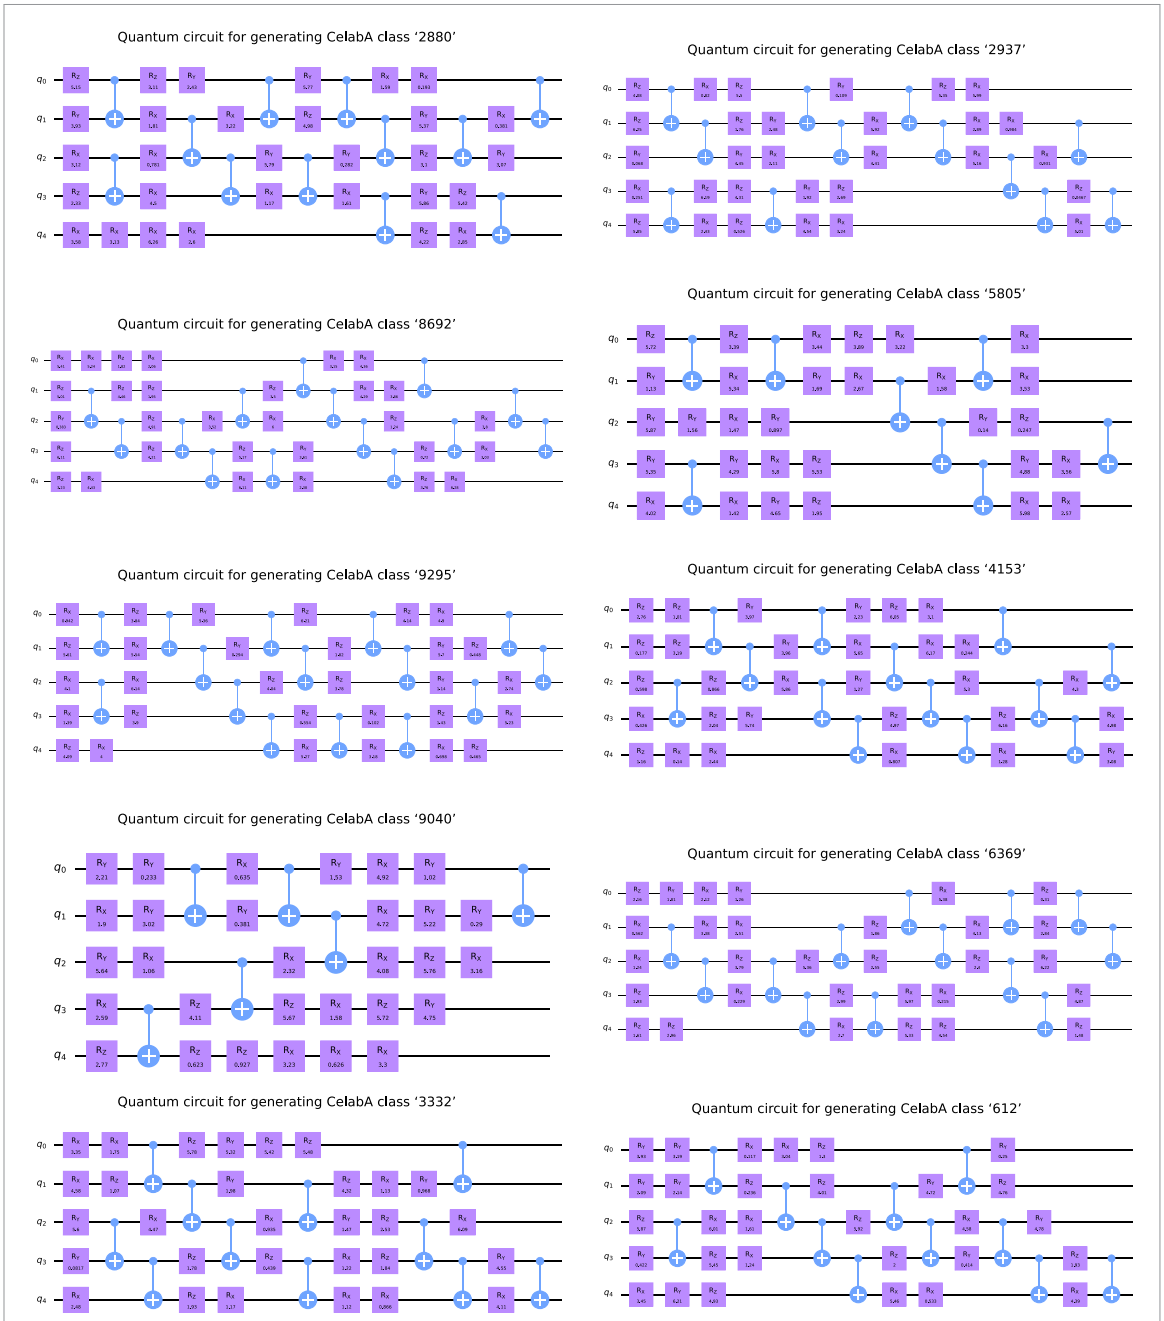

Figure 15. The best-performing quantum circuit obtained by QAS using the training data constructed from the CelebA dataset.

parameterized quantum circuits. In our current implementation, the QAS stage indeed introduces additional computational overhead compared to traditional QGAN training, since multiple candidate circuits need to be evaluated and compared during the search. For example, in the case of a 5-qubit, 6-depth SuperCircuit, under our hardware environment (40-core CPU and a single NVIDIA GeForce GTX 1080 Ti GPU), training for 3000 epochs while searching over 4000 quantum circuits required approximately 5 days. Afterward, we deployed the best-performing circuit obtained from the search into the QGAN and trained it for image generation, which took about another 2 days. Although the QAS stage incurs this one-time additional cost, its advantage lies in automatically discovering more compact generator circuits, significantly reducing the number of two-qubit gates, and improving the stability of subsequent adversarial training. In the present work, our focus was mainly on demonstrating the feasibility of integrating QAS into QGANs, and we did not conduct systematic studies on dataset diversity or training stability. Due to current hardware limitations, this was beyond the scope of the present work.

The integration of QAS into QGAN offers several advantages, as detailed above. Firstly, QAS automates the search for optimal or near-optimal generator circuit architectures, eliminating the need for manual design. This significantly reduces the complexity of quantum circuits and improves efficiency. Secondly, by

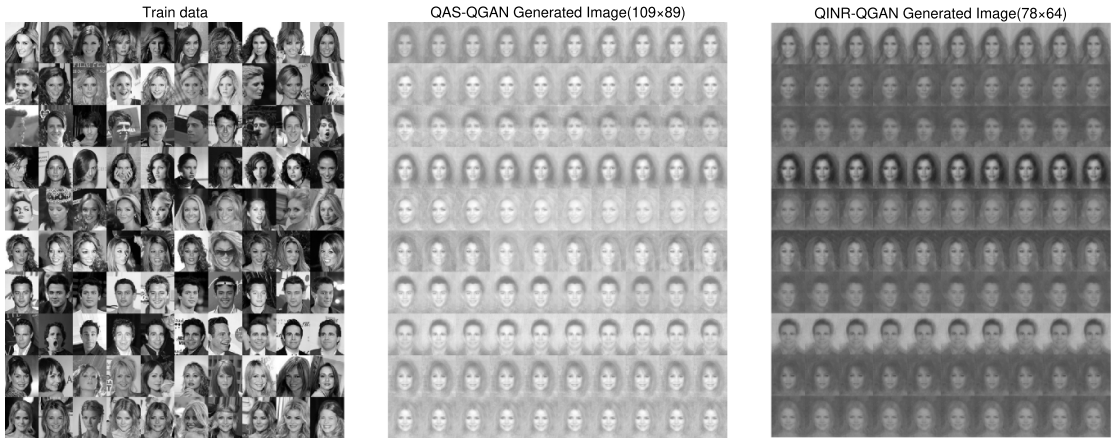

Figure 16. Facial images (109 × 89) generated by QAS-QGAN and facial images (78 × 64) generated by QINR-QGAN.

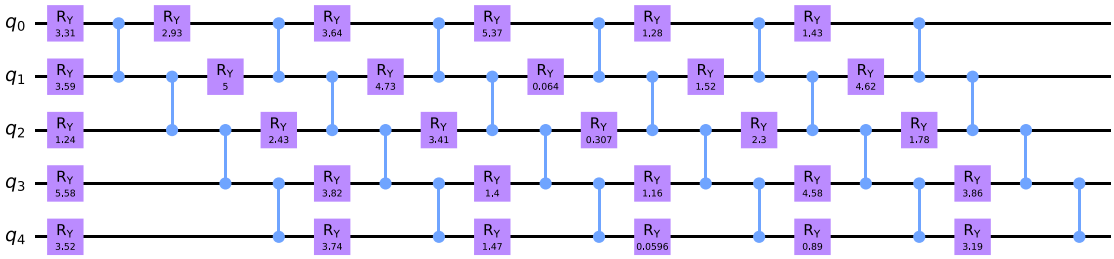

Figure 17. The general quantum generator circuit architecture used in MosaiQGAN.

Table 2. Comparison of the number of two-qubit gates used in the quantum circuits of MosaiQGAN and QAS-QGAN.

MNIST Dataset										
Model	Digit 0	Digit 1	Digit 2	Digit 3	Digit 4	Digit 5	Digit 6	Digit 7	Digit 8	Digit 9
MosaiQGAN	24	24	24	24	24	24	24	24	24	24
QAS-QGAN	12	12	14	8	15	13	6	13	11	11
Fashion MNIST Dataset										
Model	T-shirt	Trouser	Pullover	Dress	Coat	Sandal	Shirt	Sneaker	Bag	Ankle boot
MosaiQGAN	24	24	24	24	24	24	24	24	24	24
QAS-QGAN	8	12	11	10	14	12	9	13	12	12

enabling the search for quantum circuit architectures, QAS enhances the adaptability of QGANs to varying data distributions and task-specific requirements, leading to improved image generation quality. Furthermore, the simplified circuit designs reduce reliance on two-qubit gate operations, decreasing the demand for quantum hardware resources. As a result, QGANs become more feasible for implementation on current NISQ devices.

In summary, the incorporation of QAS enhances the flexibility, scalability, and practicality of QGANs, offering a promising path toward the development of efficient quantum generative models. QAS-QGAN represents a significant advancement in quantum machine learning by optimizing the design of quantum generators in QGANs. It automatically discovers optimized quantum circuits, reduces circuit design complexity, and provides a more scalable solution for implementation on current quantum devices. Experimental results show that QAS-QGAN not only competes with traditional and hybrid models but, in some cases, outperforms them while maintaining quantum resource efficiency.

6. Outlook and future work

Although QAS-QGAN has achieved promising results across multiple datasets, there remains ample room for further optimization and extension of the framework. Future research can be pursued along the following directions:

Incorporating Multi-Objective Quantum Architecture Search Mechanisms: The current QAS approach in this work primarily relies on reconstruction error (MSE) as the sole optimization objective. In future research, a multi-objective optimization framework could be introduced–jointly considering reconstruction error and quantum resource consumption metrics such as the number of single-qubit gates, two-qubit gates, and overall circuit depth. This would enable a more effective trade-off between image generation quality and resource efficiency. For example, integrating multi-objective search strategies from reinforcement learning or evolutionary algorithms–such as NSGA-III–could facilitate more efficient quantum circuit discovery.

Co-Optimization with NAS: At present, QAS-QGAN focuses solely on optimizing the quantum architecture of the generator, while the discriminator remains a fixed classical neural network. In future studies, the joint optimization of NAS and QAS could be explored to enable end-to-end co-design of both the quantum generator and classical discriminator, potentially leading to a more robust and effective hybrid quantum–classical generative model.

Data availability statement

The code of the paper can be found at: https://github.com/ma893460966/QAS-QGAN. All data that support the findings of this study are included within the article (and any supplementary files).

Funding

This work is supported by the National Natural Science Foundation of China under Grant No. 62171470, Henan Province Central Plains Science and Technology Innovation Leading Talent Project (No. 234200510019), Natural Science Foundation of Henan (No. 232300421240), Natural Science Foundation of Henan (252300420990), and Laboratory for Advanced Computing and Intelligence Engineering Fund.

Author contributions

QGM and CLH designed the experiments. NWS developed the idea and framework. QGM was responsible for deriving the theoretical framework and coding. DQ, NWS, and CLH contributed to the review, editing, and evaluation of the results. The manuscript was written with input from all authors, and all authors reviewed and approved the final version.

Appendix A. Experimental results under noisy environments

To better reflect the behavior of quantum circuits on real NISQ hardware, we incorporated a depolarizing noise model into the quantum gate operations during the QAS-QGAN architecture search phase. The results obtained under these noisy conditions are presented below. Specifically, during the circuit search phase of QAS-QGAN, we applied depolarizing noise with a probability of $p_1 = 0.05$ to single-qubit gate operations, and $p_2 = 0.2$ to two-qubit gate operations. In the image generation phase, we further added depolarizing noise with $p = 0.05$ to every layer and wire of the quantum circuits discovered through the search process, in order to closely mimic the non-idealities commonly observed in real quantum hardware. The quantum circuits that achieved the best performance for generating each image category under noisy conditions are collectively presented in appendix B.

Figure 18 shows the images generated by QAS-QGAN under noisy conditions. As observed, QAS-QGAN is still capable of successfully generating class-representative images from the target datasets despite the presence of noise.

Furthermore, figure 19 compares evaluation metrics for images generated by QAS-QGAN under both noisy and noise-free environments across different categories. The results further demonstrate QAS-QGAN's robustness to quantum noise.

Interestingly, for certain evaluation metrics, the values obtained under noisy conditions are even better than those in the noise-free setting. We attribute this to the observation that low levels of quantum noise are not necessarily detrimental to quantum machine learning algorithms; rather, they may in fact be beneficial for certain quantum learning tasks [67, 68].

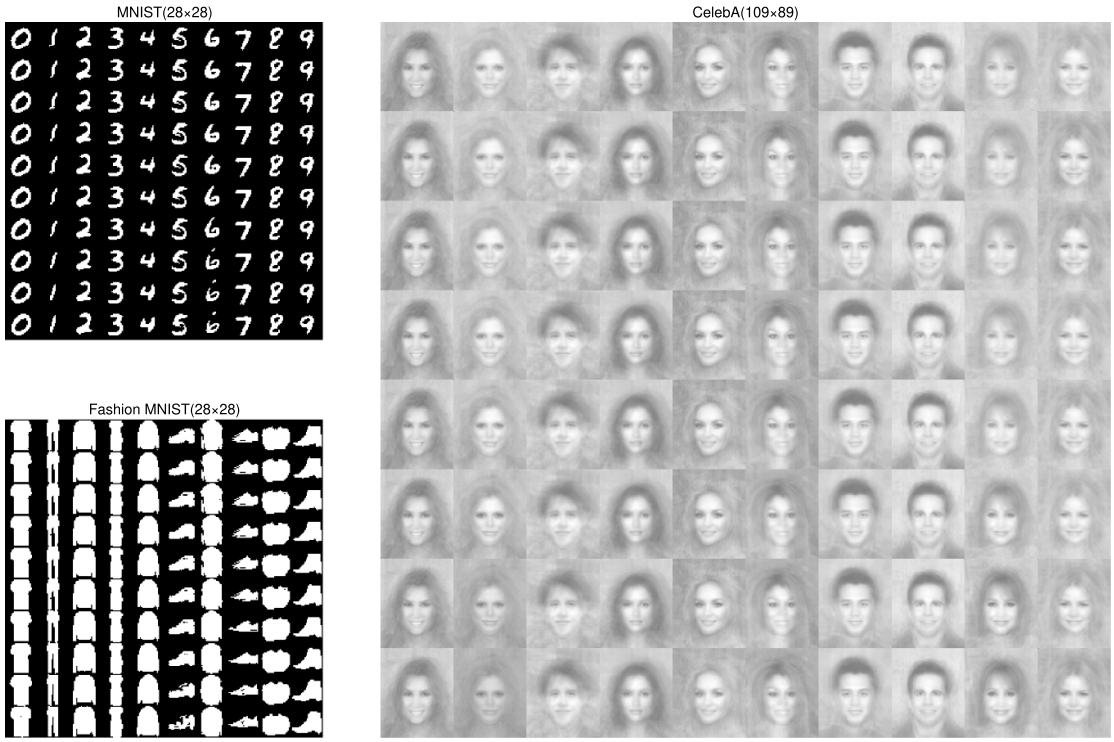

Figure 18. Images Generated by QAS-QGAN Under Noisy Environments.

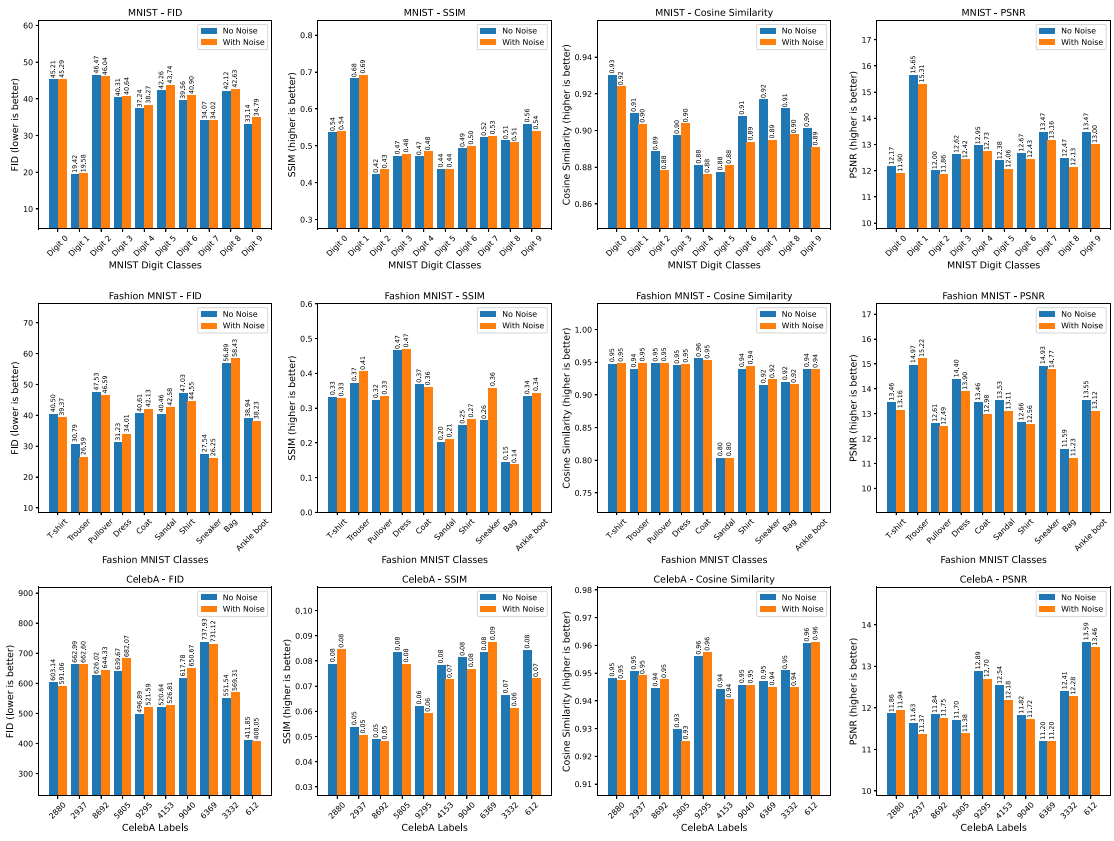

Figure 19. Image Quality Metrics Comparison: Noise-Free vs. Noisy QAS-QGAN.

Appendix B. Performance-optimal quantum circuits discovered by QAS under noisy environments

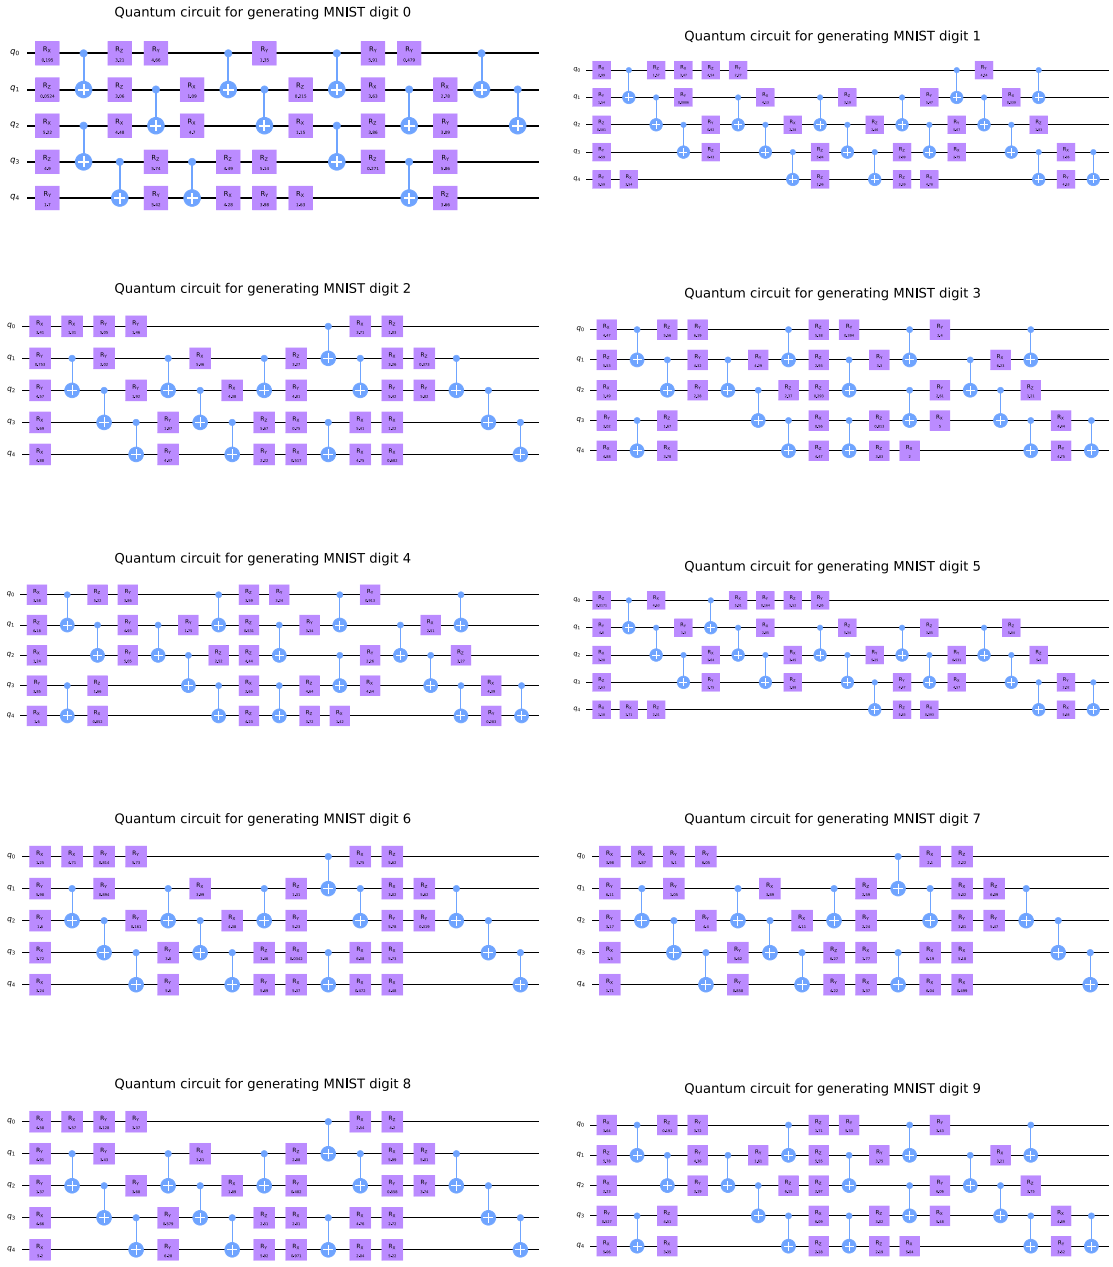

Figure 20. Optimal Quantum Circuits Discovered by QAS Under Noisy Conditions on the MNIST Dataset.

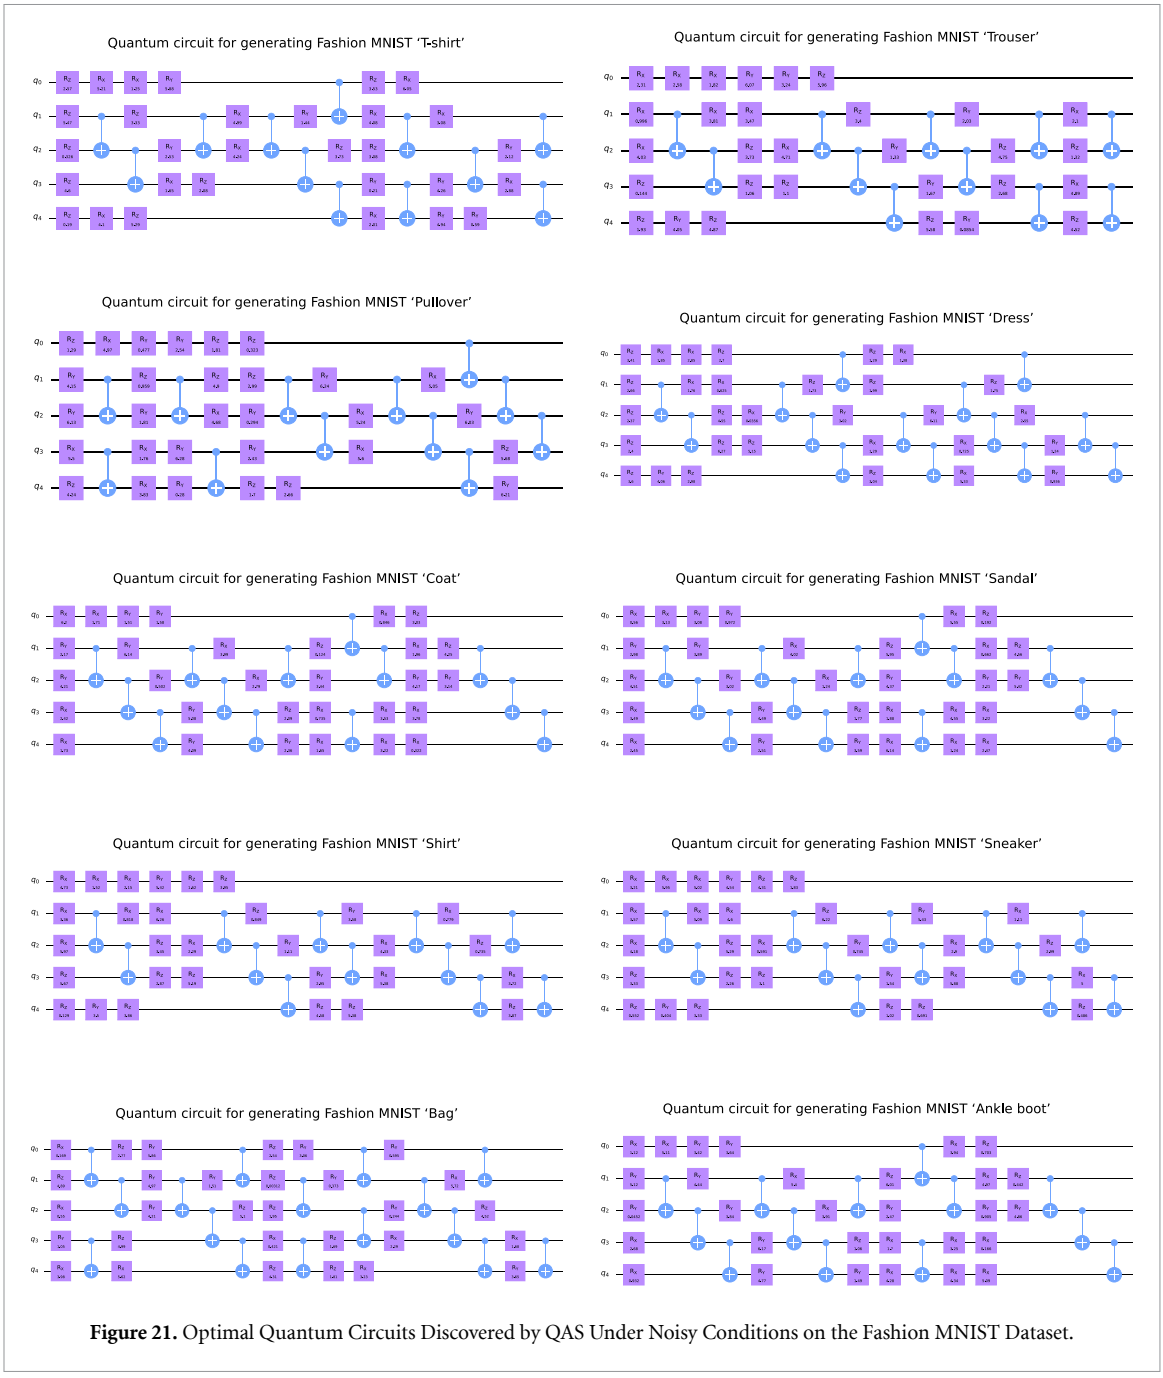

Figure 21. Optimal Quantum Circuits Discovered by QAS Under Noisy Conditions on the Fashion MNIST Dataset.

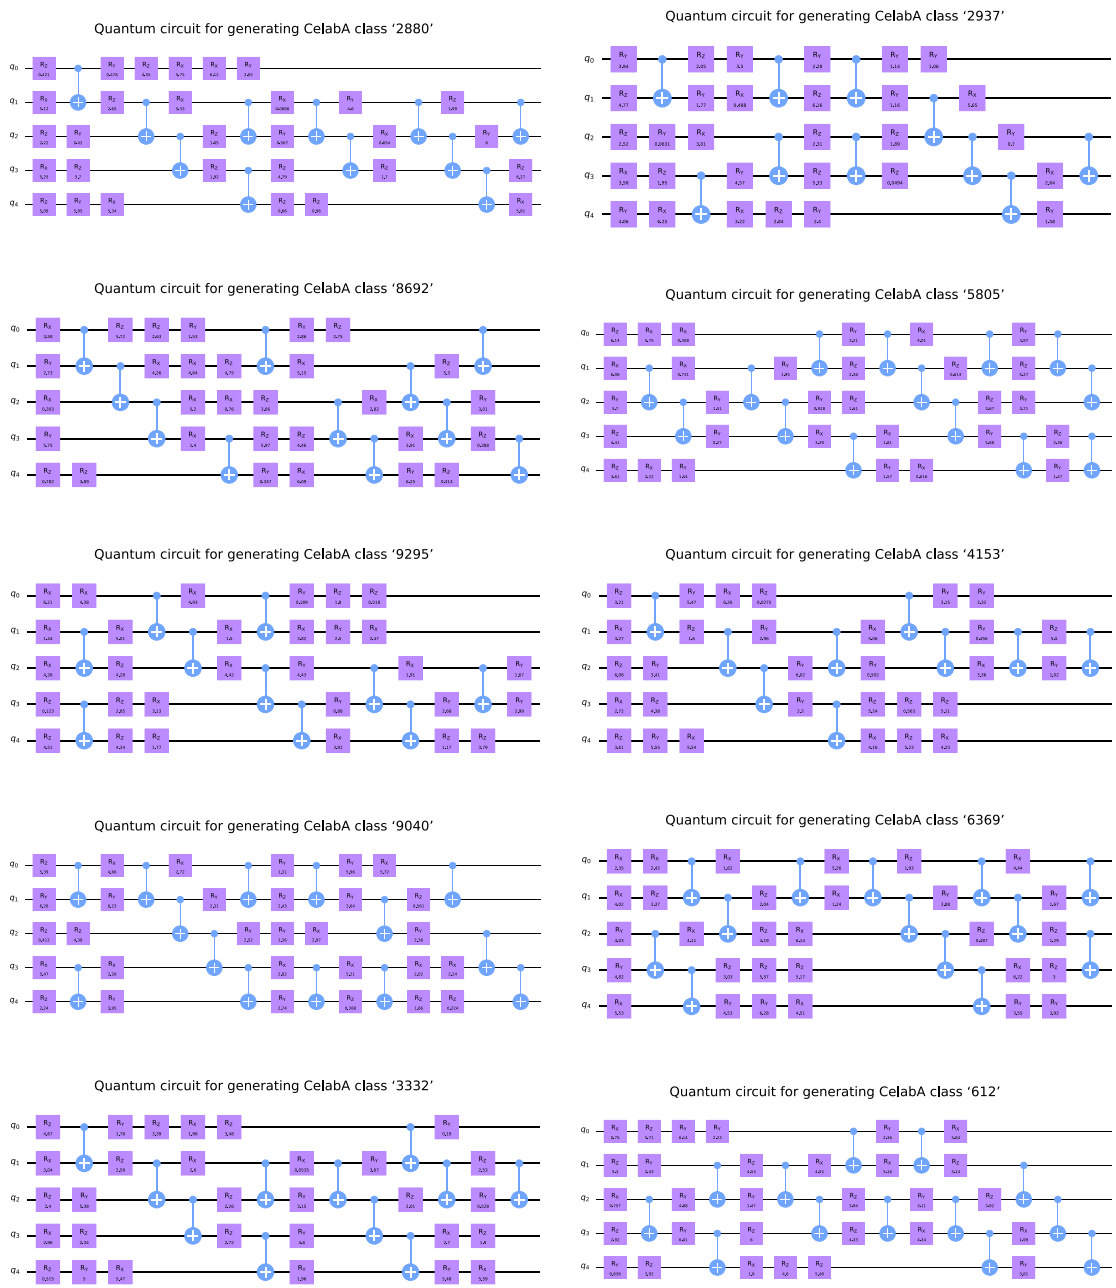

Figure 22. Optimal Quantum Circuits Discovered by QAS Under Noisy Conditions on the CelebA Dataset.

ORCID iDs

Quangong Ma 0000-0002-9093-0142
Chaolong Hao 0009-0002-1335-3457
NianWen Si 0000-0003-4619-4325
Dan Qu 0000-0001-9917-7794

References

[1] Silver D, Ranjan A, Patel T, Gandhi H, Cutler W and Tiwari D 2023 MosaiQ: quantum generative adversarial networks for image generation on NISQ computers *2023 IEEE/CVF Int. Conf. on Computer Vision (ICCV)* pp 7007–16
[2] Schuld M and Killoran N 2019 Quantum machine learning in feature hilbert spaces *Phys. Rev. Lett.* **122** 040504
[3] Beer K, Bondarenko D, Farrelly T, Osborne T J, Salzmann R, Scheiermann D and Wolf R 2020 Training deep quantum neural networks *Nat. Commun.* **11** 808
[4] Thompson J 2025 Quantum neural networks can be normal *Nat. Phys.* **21** 1042–3
[5] Romero J, Olson J P and Aspuru-Guzik A 2017 Quantum autoencoders for efficient compression of quantum data *Quantum Sci. Technol.* **2** 045001

[6] Locher D F, Cardarelli L and Müller M 2023 Quantum error correction with quantum autoencoders *Quantum* **7** 942
[7] Huang H-L *et al* 2021 Experimental quantum generative adversarial networks for image generation *Phys. Rev. Appl.* **16** 024051
[8] Tsang S L, West M T, Erfani S M and Usman M 2023 Hybrid quantum–classical generative adversarial network for high-resolution image generation *IEEE Trans. Quantum Eng.* **4** 1–19
[9] QuanGong M, Hao C, NianWen S, Chen G, Zhang J and Dan Q 2025 Quantum adversarial generation of high-resolution images *EPJ Quantum Technol.* **12** 3
[10] Chen C, Zhao Q, Zhou M, Zhimin H, Sun Z and Situ H 2024 Quantum generative diffusion model: a fully quantum-mechanical model for generating quantum state ensemble (arXiv:2401.07039)
[11] Zhang B, Peng X, Chen X and Zhuang Q 2024 Generative quantum machine learning via denoising diffusion probabilistic models *Phys. Rev. Lett.* **132** 100602
[12] Zhang J, Che X, Fan Y, Peng S, Chen G, Quangong M and Juncheng H 2025 Denoising diffusion models with optimized quantum implicit neural networks for image generation *Future Gener. Comput. Syst.* **173** 107875
[13] Chu C, Skipper G, Swany M and Chen F 2023 IQGAN: robust quantum generative adversarial network for image synthesis on NISQ devices *ICASSP 2023–2023 IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)* pp 1–5
[14] Anoshin M, Sagingalieva A, Mansell C, Zhiganov D, Shete V, Pflitsch M and Melnikov A 2024 Hybrid quantum cycle generative adversarial network for small molecule generation *IEEE Trans. on Quantum Eng.* **5** 1–14
[15] Junde Li, Topaloglu R O and Ghosh S 2021 Quantum generative models for small molecule drug discovery *IEEE Trans. on Quantum Engineering* **2** 1–8
[16] Herr D, Obert B and Rosenkranz M 2021 Anomaly detection with variational quantum generative adversarial networks *Quantum Sci. Technol.* **6** 045004
[17] Chakrabarti S, Huang Y, Tongyang Li, Feizi S and Xiaodi W 2019 Quantum wasserstein GANs *Proc. 33rd Int. Conf. on Neural Information Processing Systems (NeurIPS 2019)* (Curran Associates Inc.)
[18] Chu C, Hastak A and Chen F 2025 LSTM-QGAN: scalable NISQ generative adversarial network *ICASSP 2025 - 2025 IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)* pp 1–5
[19] Lian-Hui Y, Xiao-Yu Li, Chen G, Zhu Q-S, Hui Li and Yang G-W 2025 Knowledge-driven quantum architecture search through filtering and focusing *Inf. Sci.* **718** 122420
[20] Yuxuan D, Huang T, You S, Hsieh M-H and Tao D 2022 Quantum circuit architecture search for variational quantum algorithms *npj Quantum Inf.* **8** 62
[21] QuanGong M, Hao C, Yang X, Qian L, Zhang H, NianWen S, MinChen X and Dan Q 2024 Continuous evolution for efficient quantum architecture search *EPJ Quantum Technol.* **11** 54
[22] Yangyang Li, Liu R, Hao X, Shang R, Zhao P and Jiao L 2023 EQNAS: evolutionary quantum neural architecture search for image classification *Neural Netw.* **168** 471–83
[23] Wenjie W, Yan G, Xudong L, Pan K, Yan J and Sabato S 2023 QuantumDARTS: differentiable quantum architecture search for variational quantum algorithms *Proc. 40th Int. Conf. on Machine Learning (Proc. Machine Learning Research vol 202) (Honolulu, Hawaii, USA, 23–29 July 2023)* ed A Krause, E Brunskill, K Cho, B Engelhardt and J Scarlett (PMLR) pp 37745–3776
[24] Wang H, Ding Y, Jiaqi G, Lin Y, Pan D Z, Chong F T and Han S 2022 QuantumNAS: noise-adaptive search for robust quantum circuits *2022 IEEE Int. Symp. on High-Performance Computer Architecture (HPCA)* pp 692–708
[25] Lecun Y, Bottou L, Bengio Y and Haffner P 1998 Gradient-based learning applied to document recognition *Proc. of the IEEE* vol 86 pp 2278–324
[26] Xiao H, Rasul K, and Vollgraf R 2017 Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms (arXiv:1708.07747)
[27] Heusel M, Ramsauer H, Unterthiner T, Nessler B and Hochreiter S 2017 Gans trained by a two time-scale update rule converge to a local nash equilibrium *Advances in Neural Information Processing Systems* vol 30, ed I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan and R Garnett (Curran Associates, Inc.) pp 3730–8
[28] Wang Z, Bovik A C, Sheikh H R and Simoncelli E P 2004 Image quality assessment: from error visibility to structural similarity *IEEE Trans. Image Process.* **13** 600–12
[29] Zhu X, Shisheng S, Mingzhu F, Liu J, Zhu L, Yang W, Jing G and Guo Y 2018 A cosine similarity algorithm method for fast and accurate monitoring of dynamic droplet generation processes *Sci. Rep.* **8** 9967
[30] Sheikh H R, Sabir M F and Bovik A C 2006 A statistical evaluation of recent full reference image quality assessment algorithms *IEEE Trans. Image Process.* **15** 3440–51
[31] Liu Z, Luo P, Wang X and Tang X 2015 Deep learning face attributes in the wild *2015 IEEE Int. Conf. on Computer Vision (ICCV)* pp 3730–8
[32] Goodfellow I J, Pouget-Abadie J, Mirza M, Bing X, Warde-Farley D, Ozair S, Courville A and Bengio Y 2014 Generative adversarial nets *Proc. 28th Int. Conf. on Neural Information Processing Systems NIPS*' vol 2 (MIT Press) pp 2672–80
[33] Kingma D P and Welling M 2022 Auto-encoding variational bayes (arXiv:1312.6114)
[34] Jonathan H, Jain A and Abbeel P 2020 Denoising diffusion probabilistic models *Advances in Neural Information Processing Systems* (Red Hook, NY, USA) vol 33, ed H Larochelle, M Ranzato, R Hadsell, M F Balcan and H Lin (Curran Associates, Inc.) pp 6840–51
[35] Dallaire-Demers P-L and Killoran N 2018 Quantum generative adversarial networks *Phys. Rev. A* **98** 012324
[36] Lloyd S and Weedbrook C 2018 Quantum generative adversarial learning *Phys. Rev. Lett.* **121** 040502
[37] Chang S Y, Thanasilp S, Saux B L, Vallecorsa S, and Grossi M2024 Latent style-based quantum GAN for high-quality image generation (arXiv:2406.02668)
[38] Elsken T, Metzen J H and Hutter F 2019 Neural architecture search: a survey *J. Mach. Learn. Res.* **20** 1997–2017
[39] Ostaszewski M, Trenkwalder L M, Masarczyk W, Scerri E and Dunjko V 2021 Reinforcement learning for optimization of variational quantum circuit architectures *Proc. 35th Int. Conf. on Neural Information Processing Systems, NIPS'21 (Curran Associates Inc.)*
[40] Han Z, Hong D, Gao L, Roy S K, Zhang B and Chanussot J 2022 Reinforcement learning for neural architecture search in hyperspectral unmixing *IEEE Geosci. Remote Sens. Lett.* **19** 1–5
[41] Pham H, Guan M, Zoph B, Quoc L and Dean J 2018 Efficient neural architecture search via parameters sharing *Proc. 35th Int. Conf. on Machine Learning (Proc. of Machine Learning Research* vol 80) *(Stockholm, Sweden)* ed J Dy and A Krause (PMLR) pp 4095–104
[42] Yang Z, Wang Y, Chen X, Shi B, Chao X, Xu C, Tian Q and Chang X 2020 CARS: continuous evolution for efficient neural architecture search *2020 IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)* pp 1826–35
[43] Wang C, Chang X, Yao X and Tao D 2019 Evolutionary generative adversarial networks *IEEE Trans. Evol. Comput.* **23** 921–34

[44] Rattew A G, Shaohan H, Pistoia M, Chen R, and Wood S 2020 A domain-agnostic, noise-resistant, hardware-efficient evolutionary variational quantum eigensolver (arXiv:1910.09694)
[45] Yuhui X, Xie L, Zhang X, Chen X, Guo-Jun Q, Tian Q and Xiong H 2020 PC-DARTS: partial channel connections for memory-efficient architecture search *Int. Conf. on Learning Representations*
[46] Real E, Aggarwal A, Huang Y and Quoc V L 2019 Regularized evolution for image classifier architecture search *Proc. 33rd AAAI Conf. on Artificial Intelligence and 31st Innovative Applications of Artificial Intelligence Conf. 9th AAAI Symp. on Educational Advances in Artificial Intelligence AAAI'19/IAAI'19/EAAI'19* (AAAI Press)
[47] Ding Li and Spector L 2022 Evolutionary quantum architecture search for parameterized quantum circuits *Proc. Genetic and Evolutionary Computation Conf. Companion, GECCO'22* (Association for Computing Machinery) pp 2190–5
[48] Chivilikhin D, Samarin A, Ulyantsev V, Iorsh I, Oganov A R and Kyriienko O 2020 MoG-VQE: multiobjective genetic variational quantum eigensolver (arXiv:2007.04424)
[49] He Z, Wei J, Chen C, Huang Z, Situ H and Li L 2024 Gradient-based optimization for quantum architecture search *Neural Netw.* **179** 106508
[50] Sun Y, Yunpu M and Tresp V 2023 Differentiable quantum architecture search for quantum reinforcement learning *2023 IEEE Int. Conf. on Quantum Computing and Engineering (QCE)* vol 02 pp 15–19
[51] Liu H, Simonyan K and Yang Y 2019 DARTS: differentiable architecture search *Int. Conf. on Learning Representations*
[52] Zhang S-X, Hsieh C-Y, Zhang S and Yao H 2022 Differentiable quantum architecture search *Quantum Sci. Technol.* **7** 045023
[53] Zhang X, Yonggang Li, Zhang X, Wang Y and Sun J 2023 Differentiable architecture search with random features *2023 IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)* pp 16060–9
[54] Tian Y, Shen Li, Shen Li, Guinan S, Zhifeng Li and Liu W 2022 AlphaGAN: fully differentiable architecture search for generative adversarial networks *IEEE Trans. Pattern Anal. Mach. Intell.* **44** 6752–66
[55] Gao C, Chen Y, Liu S, Tan Z and Yan S 2020 AdversarialNAS: adversarial neural architecture search for GANs *2020 IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)* pp 5679–88
[56] Ying G, Xin H, Gao B, Han B and Chu X 2022 EAGAN: efficient two-stage evolutionary architecture search for GANs *Computer Vision—ECCV 2022: 17th European Conf. (Tel Aviv, Israel, 23–27 October 2022) (Proc., Part XVI) (Springer)* pp 37–53
[57] Doveh S and Giryes R 2021 DEGAS: differentiable efficient generator search *Neural Comput. Appl.* **33** 17173–84
[58] Ganepola V V V and Wirasingha T 2021 Automating generative adversarial networks using neural architecture search: a review *2021 Int. Conf. on Emerging Smart Computing and Informatics (ESCI)* pp 577–82
[59] Massacci B M 2024 Evolutionary optimization for hybrid quantum image generation gan *Master's Thesis* Maastricht University, Faculty of Science and Engineering, Department of Advanced Computing Sciences
[60] Xie J, Liu C and Dong Y 2024 An evolutionary quantum generative adversarial network *Quantum Mach. Intell.* **6** 84
[61] Sim S, Johnson P D and Aspuru-Guzik A 2019 Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms *Adv. Quantum Technol.* **2** 1900070
[62] Tian Y, Tian C, Fan Z, Minghao F and Hongyang M 2025 Quantum generative adversarial network with automated noise suppression mechanism based on WGAN-GP *EPJ Quantum Technol.* **12** 80
[63] Paszke A *et al* 2019 PyTorch: an imperative style, high-performance deep learning library *Advances in Neural Information Processing Systems* vol 32, ed H Wallach, H Larochelle, A Beygelzimer, F d' Alché-Buc, E Fox and R Garnett (Curran Associates, Inc.)
[64] Bergholm V et al 2022 PennyLane: automatic differentiation of hybrid quantum-classical computations (arXiv:1811.04968)
[65] Javadi-Abhari A et al 2024 Quantum computing with qiskit (arXiv:2405.08810)
[66] Gulrajani I, Ahmed F, Arjovsky M, Dumoulin V and Courville A 2017 Improved training of wasserstein GANs *Proc. 31st Int. Conf. on Neural Information Processing Systems NIPS'17 (Curran Associates Inc.)* pp 5769–79
[67] Nguyen N H, Behrman E C and Steck J E 2020 Quantum learning with noise and decoherence: a robust quantum neural network *Quantum Mach. Intell.* **2** 1
[68] Yuxuan D, Hsieh M-H, Liu T, Tao D and Liu N 2021 Quantum noise protects quantum classifiers against adversaries *Phys. Rev. Res.* **3** 023153