AIMS Mathematics, 10(11): 26568–26592.
DOI: 10.3934/math.20251168
Received: 16 August 2025
Revised: 10 October 2025
Accepted: 04 November 2025
Published: 17 November 2025

https://www.aimspress.com/journal/Math

*Research article*

# Adaptive Grover-driven optimization for quantum-inspired deep learning: A gradient-free training framework

**Irsa Sajjad[1,*], Huda M. Alshanbari[2], Mohammed M. A. Almazah[3], Hanen Louati[4] and Sana Rauf[1]**

[1] Department of Mathematics, National University of Modern Languages, Islamabad, Pakistan; irsa.sajjad@numl.edu.pk
[2] Department of Mathematical Sciences, College of Science, Princess Nourah bint Abdulrahman University, P.O. Box 84428, Riyadh 11671, Saudi Arabia; hmalshanbari@pnu.edu.sa
[3] Department of Mathematics, College of Sciences and Arts (Muhyil), King Khalid University, Muhyil 61421, Saudi Arabia; mmalmazah@kku.edu.sa
[4] Mathematics Department, Faculty of Science, Northern Border University, Arar, KSA, Saudi Arabia; Hanen.Louati@nbu.edu.sa

* **Correspondence:** Email: irsa.sajjad@numl.edu.pk.

**Abstract:** Training deep neural networks remains difficult due to vanishing gradients, non-convex loss surfaces, and hyperparameter sensitivity. These obstacles are compounded by quantum machine learning, where barren plateaus, circuit depth, and hardware noise restrict the applicability of gradient-based approaches. To overcome these drawbacks, this study presents adaptive Grover-driven parallel quantum optimization (AG-PQO), a hybrid, gradient-free scheme that leverages Grover's quadratic search speedup, along with adaptive loss-aware discretization and fidelity-based regularization. In contrast to more classical optimizers, such as Adam or evolutionary strategies (ES), which are either sensitive to the adequacy of the gradient update or exhibit poor scaling behavior, AG-PQO optimizes by performing Grover-accelerated candidate exploration across layers and reuses high-quality solutions in quantum memory caching. Testing indicates that AG-PQO yields higher accuracy, 2%–3% above Adam and ES, and faster convergence with less end-value loss than Adam, ES, and quantum feedforward-backpropagation (QFB). It is worth noting that AG-PQO remains stable at the simulated noise level of NISQ and has the potential to scale to near-term quantum processors.

**Keywords:** quantum-inspired optimization; Grover search algorithm; adaptive perturbation; quantum

fidelity regularization; hybrid quantum-classical learning; evolutionary search in deep learning
**Mathematics Subject Classification:** 91A81, 82B10

## List of Abbreviations

| Abbreviation | Meaning |
|---|---|
| AG-PQO | Adaptive Grover-driven parallel quantum optimization |
| ALADC | Adaptive loss-aware discretization control |
| QMCR | Quantum memory caching and reuse |
| ZNE | Zero-noise extrapolation |
| QFB | Quantum feedforward-backpropagation |
| VQC | Variational quantum circuit |
| QNES | Quantum natural evolution strategies |
| QMIST | Quantum-MNIST dataset |
| NISQ | Noisy intermediate-scale quantum |
| ReLU | Rectified linear unit |

## List of equations

| Equation No. | Title/Purpose |
|---|---|
| (1) | Feed-forward neural network forward-propagation equation |
| (2) | Categorical cross-entropy loss function |
| (3) | Quantum candidate weight matrix generation |
| (4) | Oracle-based quantum loss evaluation |
| (5) | Weight update rule using Grover-selected candidate |
| (6) | Amplitude-encoding mapping of weights |
| (7) | Angle-encoding mapping of weights |
| (8) | Computational complexity for amplitude encoding |
| (9) | Computational complexity for angle encoding |
| (10) | Adaptive discretization controller ($\Delta$-update rule) |
| (11) | Cosine-similarity-based fidelity approximation |
| (12) | Fidelity-regularized loss function |
| (13) | Memory-buffer candidate generation rule |
| (14) | Grover stochastic transformation operator |
| (15) | Perturbation interval convergence condition |
| (17) | Fidelity-driven regularization stability inequality |

## 1. Introduction

One of the most significant problems in today's machine learning is training deep neural networks. Traditional optimization algorithms, including stochastic gradient descent (SGD) and its variants [1], are hindered by vanishing and exploding gradients [2], non-convex error landscapes, and hyperparameter sensitivity [3]. Such weaknesses are magnified when models are scaled to billions of parameters, due to slowed convergence [4], unstable training [5], and exponentially increased computational demands [6,7]. These challenges provide a compelling reason to investigate gradient-free methods that may serve as reliable alternatives to gradient-based optimization [8–10]. The recent potential of quantum computing [11,12] as a paradigm for accelerating optimization and learning tasks is promising [13–15]. Quantum machine learning (QML), however, adds extra challenges. One widely studied class, variational quantum circuits (VQCs) [16,17], is commonly crippled by barren plateaus that cause gradients to asymptotically approach zero and limit their ability to generalize to deeper architectures [18]. Moreover, current-day quantum devices are limited by the small number of qubits (in the hundreds) [19], restricted circuit depths [20], and noise [21]. These problems require gradient-free, noise-tolerant [22], and scalable hybrid methods that scale with both classical [23] and quantum resources [24].

Grover's algorithm provides a natural basis for answering these questions [25]. Identification of rare events among a large number of discrete candidates can also be very costly in terms of experimental evaluations [26]. However, a quadratic search advantage means that only a slight increase in the number of assessments is required compared to brute-force or evolutionary approaches [27]. Grover-driven optimization has low computational overhead compared to classical evolutionary methods, is not susceptible to high parameter-dependence issues such as those found in VQCs and QAOA-based approaches, and is scalable on noisy intermediate-scale quantum (NISQ) devices [28].

In this paper, we propose adaptive Grover-driven parallel quantum optimization (AG-PQO). This novel hybrid optimization method combines bypass optimization with gradient-free methods, aiming to overcome the deficiencies of current approaches. AG-PQO hybridizes the accelerated candidate search provided by Grover with three new mechanisms:

1. Adaptive loss-aware discretization control (ALADC) resolves the tradeoff between exploration and convergence by using variable resolution of the allowed candidate intervals.
2. Fidelity-aware loss regularization is a proposed penalty designed to ensure stability after an epoch and insensitivity to noise.
3. Quantum memory caching (QMCR) recycles high-quality states, reducing oracle calls and increasing efficiency.

By combining these mechanisms, AG-PQO enables noise-resistant and scalable training. Experiments conducted by us demonstrated that it can achieve greater accuracy and a higher rate of convergence than Adam, evolutionary strategies (ES), and quantum feedforward-backpropagation (QFB). The essential contributions of this paper include the following:

1. We propose the method of parallel optimizing across network layers using a Grover-driven gradient-free optimization method. ALADC is designed to efficiently and stably explore the candidate.
2. We jointly tackle the fidelity-of-learning problem by incorporating fidelity-aware loss regularization to enhance noise robustness.
3. We present the caching strategy QMCR to eliminate excess oracle queries.

We compare AG-PQO with classical and hybrid optimizers, showing higher accuracy, efficiency, and NISQ-readiness. Section 2 discusses related literature and other optimization procedures. Section 3

gives a detailed description of the proposed AG-PQO framework. Experimental results are presented in Section 4, and conclusions with future directions are included in Section 5.

## 2. Literature review

Gradient-free methods have long been explored as alternatives to backpropagation when gradients are noisy, undefined, or unreliable in highly non-convex landscapes. Evolutionary strategies (ES), as well as differential evolution and simulated annealing, have been extended to architecture search and weight optimization [9,10]. Although such methods enable rough objective surfaces and can circumvent vanishing gradients, current implementations are characterized by high computational complexity and poor convergence on larger models in practice.

The Grover algorithm provides a quadratic speedup for unstructured search and has been generalized to minimum-finding, database search [11], and discrete optimization problems [12]. The potential for general combinatorial exploration has inspired interest in applications in satisfiability, portfolio selection, and binary classification [12,13]. Selection techniques similar to Grover have also been discussed in the context of structured weight spaces in learning problems, albeit in idealized settings (specialized objective values, no or minimal discretization, or even the absence of noise modeling) [14]. These works suggest a natural avenue for gradient-free learning, in which we replace continuous descent with a broader search across discrete candidate sets.

Developments in quantum algorithms have explored multiple-solution amplification and block encodings, as well as parallelized candidate sets, elucidating oracle-based complexity. These are promising avenues for achieving speedup [15,16]. Although this has entered a new theoretical stage, end-to-end integration in the deep learning workflow is limited. There is a gap between the theoretical specifications and the capabilities of current implementations with the state-of-the-art software/hardware stack.

Quantum-inspired machine learning (e.g., tensor networks, born machines, quantum Boltzmann distributions) has leveraged the inductive biases motivated by entanglement to compress correlations and achieve greater sample efficiency [17]. These methods, however, do not generally utilize the quadratic search advantage at the circuit level that Grover identified. By contrast, recent efforts in the same direction use Grover-based candidate selection within the training loop to accelerate discrete exploration, while remaining compatible with classical approaches.

In general, evolutionary approaches can enable gradient-free search but exhibit poor scaling behavior. Variational quantum circuits (VQCs) can suffer from barren plateaus, which hinder gradient-based optimization, and quantum fictitious system (QFS)/QAOA-based approaches have limited depth and can be highly parameter-sensitive. The suggested AG-PQO performs layer-wise Grover optimization, with adaptive discretization and fidelity-aware stabilization, aiming for scalability and noise robustness. Table 1 outlines these differences and explains why AG-PQO diverges from the earlier literature. Compensating for decoherence, gate errors, and measurement bias is necessary to deploy noisy intermediate-scale quantum (NISQ) hardware in practice.

Zero-noise extrapolation, probabilistic error cancellation, and Richardson-style extrapolation are techniques developed to overcome these effects and produce high-fidelity estimates from shallow circuits [18,19]. Complementary strategies, designed to reduce hardware requirements, divide workloads into smaller pieces that can be executed on the hardware [20]. Cosine-similarity-based training can be sensitive to cumulative noise, mainly due to Grover's amplitude amplification. This makes these

mitigation pipelines especially important to making search-based training viable on such devices.

Hybrid quantum-classical solutions also work on integrating quantum subfunctions into learning. The most notable avenues include artificial quantum neurons, quantum circuit learning with parametrized circuits, and kernel-based quantum methods [21,22]. Although such methods can work well at small scales, most of them rely on shallow depths or local parameterizations, which limit their expressiveness in the presence of noise and complicate training as depth increases. Meanwhile, experience replay, a concept in reinforcement learning, has demonstrated that a memory system can stabilize learning by reusing high-value states [23].

Similar ideas in the quantum context have been pursued, e.g., quantum memories and state caching, to minimize oracle calls as much as possible without sacrificing high-fidelity solutions [24]. Contrasting classical regularizers that limit the scale of weights or enforce sparsity, quantum-sensitive loss shaping can harness fidelity and underscore Hilbert-space overlap [25] or similar concepts to guide the search toward quantum states that have been previously found to be successful in higher dimensions. Such signals can be added to the loss to regularize optimization trajectories, enhance generalization, and alleviate the instabilities induced by realistic noise. These methods are compared in Table 1.

**Table 1.** Comparative overview of existing methods versus AG-PQO, highlighting scalability and noise-robustness advantages.

| Method | Type | Strengths | Limitations | Scalability / NISQ-readiness | Distinct Difference from AG-PQO |
|---|---|---|---|---|---|
| **Evolutionary Strategies (ES)** | Classical, gradient-free | Robust in non-differentiable settings; no gradients required | High computational cost; slow convergence; poor scalability | Limited — inefficient for high-dimensional models | AG-PQO achieves faster convergence with Grover acceleration and caching |
| **Variational Quantum Circuits (VQCs)** | Hybrid quantum-classical (gradient-based) | Parametrized quantum circuits; direct use of quantum resources | Suffer from barren plateaus; require deep circuits; noise-sensitive | Low — limited by qubit count and circuit depth | AG-PQO avoids barren plateaus by using discrete Grover search instead of gradients |
| **Quantum Feedforward-Backpropagation (QFB)** | Hybrid, gradient-based | Extends backpropagation into quantum circuits | Requires shallow depths; sensitive to parameter initialization | Low — not noise-robust, shallow expressivity | AG-PQO scales to deeper models without gradients and with fidelity-based regularization |
| **Quantum Approximate Optimization Algorithm (QAOA)** | Hybrid, variational | Effective for discrete optimization problems | Parameter-dependent; depth scaling issues; noise-sensitive | Medium — limited by qubit fidelity and gate noise | AG-PQO parallelizes layer-wise optimization, making it more scalable under NISQ constraints |
| **AG-PQO (Proposed)** | Hybrid quantum-classical, gradient-free | Grover-based layer-wise candidate search; adaptive discretization (ALADC); fidelity-aware loss regularization; memory caching | Requires repeated Grover evaluations; still simulated on hardware | High — noise-robust with ZNE and memory reuse | First to combine Grover's quadratic speedup, adaptive resolution, and noise resilience into a unified training framework |

This approach aims to address the issue of accumulating experience in reinforcement learning by utilizing memory-based optimization techniques, such as experience replay and policy caching, to stabilize convergence [23]. In quantum settings, quantum memory and optimal Grover-state caching have been proposed as methods to eliminate unnecessary oracle calls [24]. AG-PQO modifies this by saving high-fidelity candidates in a classical buffer and restarting subsequent searches with previous

successful weight perturbations. A combination of memory-augmented learning and quantum search yields a breakthrough when applied to deep network optimization.

Compared to classical regularization, which controls weight magnitude or sparsity, quantum-aware loss shaping relies on state fidelity, Hilbert-space overlap, or quantum mutual information to alter its optimization paths [25]. AG-PQO proposes adding a fidelity-directed penalty component that better aligns candidate weights with successfully tested quantum states, facilitating steadier learning and more effective generalization.

More recent hybrid optimizers [26–28] combine fuzzy-gradient or heuristic quantum searches but lack adaptive discretization and fidelity-based regularization. Instead, AG-PQO integrates Grover-based layer-wise exploration with adaptive loss-aware control (ALADC) and quantum memory reuse, offering dynamic search resolution and improved noise resistance. However, AG-PQO incurs additional oracle cost, as we will discuss in Section 4.3.

Other recent developments focusing on hybrid quantum-classical optimization have implemented practical, noise-conscious approaches that enable quantum subroutines to use NISQ-era hardware efficiently. Preliminary research on the importance of data and model design has helped clarify when quantum models can provide real benefit and the influence of circuit architectures on trainability [29]. Simultaneous advances in the learnability of parameterized quantum models [30] and in the generalization of quantum machine learning beyond kernel methods [31] have enhanced insight into generalization and convergence behavior. Theories on quantum error mitigation [32] and the computational hardness of variational circuit training [33] have also been reviewed, influencing the development of effective, resource-constrained learning models. Complementary advances in Grover-adaptive search [34], quantum circuit architecture discovery [35], and neural-network-encoded variational schemes [36] have minimized circuit depth and measurement overheads, which are important enablers of Grover-style candidate selection in hybrid training pipelines. These group improvements focus AG-PQO on fidelity-sensitive loss shaping and block encoding to maintain quadratic search benefits in realistic noise regimes. Emerging literature has highlighted the need for hybrid algorithms to strike a balance between search capability and strength. The opportunities and limitations for achieving near-term quantum advantage are illustrated by theoretical studies of error-mitigation limits [37] and experimental advancements in photonic and superconducting systems [38].

## 3. Methodology

### *3.1. Overview*

This paper proposes a new quantum-classical neural network training framework, the adaptive Grover-driven parallel quantum optimization (AG-PQO). It avoids any gradient-based optimization (e.g., backprop with ADAM) in favor of a fully gradient-free, layer-wise Grover search conducted on parallel-weight candidate matrices, with dynamic discretization, loss-shaped by quantum-informed models, and quantum-informed operating enhancements enabled by noise information. In contrast to standard quantum neural networks that use parametrized variational circuits, we address this issue by enhancing classical deep networks with quantum amplitude amplification to accelerate weight-space exploration. Figure 1 outlines the complete AG-PQO pipeline. It is initiated by classical weight initialization, followed by Grover-based quantum candidate selection, fidelity-based filtering, and adaptive perturbation control. It demonstrates AG-PQO's ability to integrate quantum-enhanced search

and gradient-free learning into a highly modular, interpretable workflow.

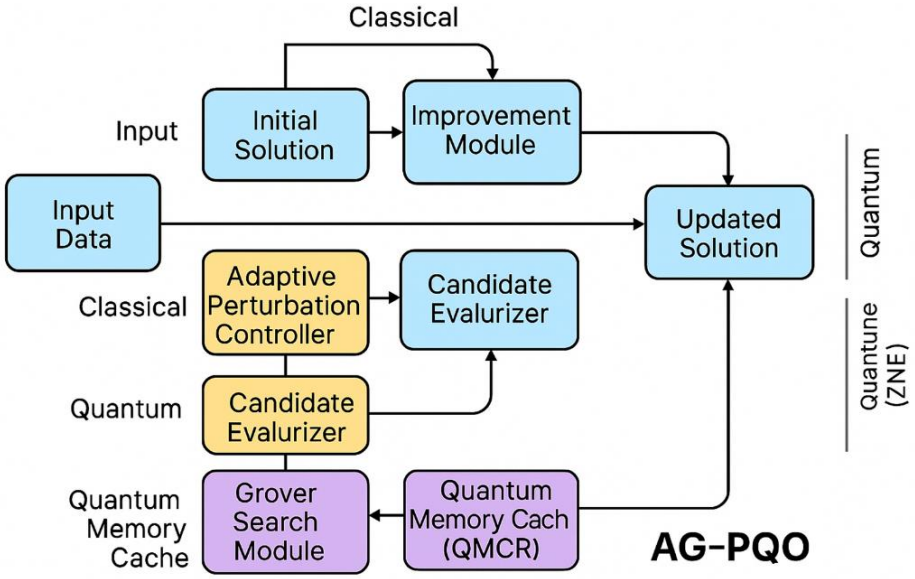


**Figure 1.** Modular architecture of AG-PQO, showing the interaction between classical and quantum components.

### *3.2. Foundation of classical model*

Let a feedforward neural network have $L$ layers, where each layer $l$ contains $n_l$ neurons, i.e., $l = 1, 2, ..., L$. The weight matrix of the layer $l$ is denoted by $\omega^{(l)} \in \mathbb{R}^{d_l \times d_{l+1}}$ and a bias $b^{(l)} \in \mathbb{R}^{d_{l+1}}$. For an input x, the forward pass is

$$h^{(l+1)} = \vartheta\left(\omega^{(l)} h^{(l)} + b^{(l)}\right), \tag{1}$$

where $\vartheta$ is a nonlinear activation (e.g., ReLU or tanh), and $\hat{y} \in \mathbb{R}^C$ is the predicted class distribution having $h^{(0)} = x,\ \hat{y} = soft\max\left(h^{(n)}\right)$. The training goal is to minimize the categorical cross-entropy loss:

$$\ell_{CE}\left(\hat{y}, y\right) = -\sum_{i=1}^{C} y_i \log \hat{y}_i. \tag{2}$$

The categorical cross-entropy loss $\ell_{CE}$ is a measure of the difference between the probability $\hat{y}_i$ and the actual class label $y_i$. This loss reduction motivates the network to give larger chances to correct classes. AG-PQO aims to optimize the family $\{\omega_1, \omega_1, ...., \omega_n\}$ without the use of a gradient descent, but by performing a battery-accelerated discrete candidate search.

### *3.3. Quantum-batch parallel Grover optimization*

Instead of iterative adjustments to individual weights, AG-PQO optimizes the entire weight

matrix for each layer in a batch with uniform sandwich perturbations. Let $\omega^{(l)}$ denote the current weight matrix of the layer $l$, and let $\omega^{(l)} = \left\{\omega_1^{(l)}, \omega_2^{(l)}, ...., \omega_N^{(l)}\right\}$ be a finite set of $N = 2^n$ candidate matrices generated via uniform perturbations:

$$\omega_k^l \sim \varsigma\left(\omega^{(l)} - \beta_t^{(l)}.\sigma^{(l)}, \omega^{(l)} + \beta_t^{(l)}.\sigma^{(l)}\right), \tag{3}$$

where $\beta_t^{(l)}$ is a layer-wise dynamic interval parameter, and $\sigma^{(l)}$ is the empirical standard deviation of $\omega^{(l)}$. Each candidate $\omega_k^l$ is represented in a quantum register either in amplitude encoding or angle encoding, depending upon the encoding capacity of the quantum device at hand. The related oracle assesses:

$$\ell_k = \ell_{CE}\left(f\left(\omega_k^l\right), y\right). \tag{4}$$

To calculate the weight of a candidate weight matrix $\omega_k^l$, the oracle function $f$ computes the training loss. The value of this loss is encoded in the phase of the quantum state, allowing it to be amplified in the best candidates during Grover iterations. With just $O\left(\sqrt{N}\right)$ queries, Grover finds the roughly optimal candidate $\omega_k^{(l)}$, that is, the algorithm finds the approximate minimizer of 2. The matrix of weights is updated as:

$$\omega^{(l)} \leftarrow \omega_{k^*}^{(l)}. \tag{5}$$

Each candidate's weight matrices $\omega^{(l)} \in \mathbb{R}^{d_l \times d_{l+1}}$ were flattened into a $\omega$ vector. In the case of amplitude encoding, the normalized $\frac{\omega}{\| \omega \|}$ is measured into a quantum state as

$$\left|\psi_w\right\rangle = \sum_{i=1}^{d_l \times d_{l+1}} \frac{\omega}{\| \omega \|} \left|i\right\rangle. \tag{6}$$

In amplitude encoding, the classical weight vector $\omega$ is normalized and encoded as a quantum superposition state $\left|\psi\right\rangle$. The weight elements make up the amplitude of the superposition state. This enables the search of Grover to work in weight space. To encode the angle, every entry $\omega_i$ is encoded in a qubit rotation using

$$\left|0\right\rangle \mapsto R_y(w_i)\left|0\right\rangle = cos\left(\frac{w_i}{2}\right)\left|0\right\rangle + sin\left(\frac{w_i}{2}\right)\left|1\right\rangle. \tag{7}$$

For angle encoding, each classical value $w_i$ is mapped to a qubit rotation $R_y(w_i)$ whose angle is proportional to the weight magnitude, enabling quantum representation using fewer qubits. This ensures that classical weight candidates always represent quantum weights. The complexity of oracle gate for amplitude encoding and angle encoding of a layer with dimension d and a given number of candidates is $O\left(N \log d\right)$ and $O\left(N d\right)$, respectively. In this way, AG-PQO attains a quadratic tradeoff with the $O\left(N^2\right)$ price of classical exhaustive search. The term *parallel* in AG-PQO refers to layer-wise parallel Grover searches executed concurrently across multiple layers, not to hardware-based quantum parallelism. In the case of a $100 \times 100$ layer, it would take 14 qubits to represent the

amplitude and 4–6 additional auxiliary qubits, and the depth of the circuit would be estimated to be $10^3$ two-qubit gates. To be NISQ-feasible, this suggests block-wise encoding of 16 × 16 sub-matrices being optimized in parallel.

### *3.4. Adaptive loss-aware discretization control (ALADC)*

To avoid both inefficient and alternating updates, add an adaptive discretization controller $\beta_t^{(l)}$, regulating the perturbation period with time $\ell$, which depends on the change in loss $\Delta\ell_t^{(l)}$:

$$\beta_{t+1}^{(l)} = \begin{cases} \max\left(\zeta\downarrow.\beta_t^{(l)}, \beta_{\min}\right) & \text{if } \Delta\ell_t^{(l)} < 0, \\ \min\left(\zeta\uparrow.\beta_t^{(l)}, \beta_{\max}\right) & \text{otherwise,} \end{cases} \tag{8}$$

where $\zeta\downarrow<1$ *and* $\zeta\uparrow>1$ are the contraction and expansion rates, respectively. These intervals $\pi\left(\Delta\ell, \sigma^{(l)}\right)$ are further emulated over epochs by a policy controller π, which may include a black-box optimization technique, such as the CMA-ES, to optimize. This adaptive scheme enables the optimizer to be fine-grained when training reaches stable convergence regimes and to take coarser exploration steps when it becomes stuck in flat or nonoptimal areas of the loss landscape. Furthermore, to make this discretization more responsive, we introduce a meta-policy controller $\zeta$, which parameterizes the evolution of the perturbation interval and can be learned. This controller can be optimized using a black-box evolutionary optimization algorithm, namely the covariance matrix adaptation evolution strategy (CMA-ES), so that the perturbation schedule adapts to the long-term dynamics of loss rather than reacting to short-term fluctuations. Overall, by combining local loss-aware adaptation (via local optimization) with global meta-policy optimization, the ALADC module enables AG-PQO to control the resolution of its search across training epochs flexibly. It leads to faster convergence, better generalization, and greater stability of the optimization process, particularly when operating in high-dimensional or noisy search spaces. In practice, the CMA-ES controller was called every k = 5 epochs to update the perturbation interval, thus being responsive without being too computationally expensive.

### *3.5. Fidelity-driven quantum-conscious loss shaping*

To enhance convergence stability and generalization under noisy conditions, we present a quantum-inspired regularization scheme that weights the macroscopic differences between current weight designs and past optimal ones. In particular, this is a fidelity-based term that helps maintain consistency in the selected candidates across epochs by directing the search process to follow weight matrices with a high degree of overlap with past successful solutions. Denoting the chosen Grover-optimal weight matrix at epoch *t* of layer *l* by $\omega^t(l)$ and the chosen matrix in the prior epoch by $\omega^{t-1}(l)$, the steps of the Grover optimization process are as follows. The following fidelity-aware regularizing term is expressed as:

$$L_{fid}^{(l)} = \zeta.\left(1 - F\left(\omega_t^{(l)}, \omega_{t-1}^{(l)}\right)\right), \tag{9}$$

where $\zeta > 0$ is a non-negative regularization parameter, and $F$ is the quantum fidelity, being approximated classically through the cosine similarity:

$$F(A,B)=\frac{\langle A,B\rangle}{\|A\|\|B\|}\in[0,1]. \tag{10}$$

This term is added to the standard categorical cross-entropy loss:

$$L_T^{(l)}=L_{CE}^{(l)}+L_{fid}^{(l)}. \tag{11}$$

The regularization mechanism is essentially a constraint applied to the optimization process, ensuring that it only moves in areas of weight space where previous Grover solutions have proven useful. This is not just a stabilizer of the candidate selection process, but also an effective temporal prior that performs well under noisy quantum conditions and will increase model generalization capabilities beyond the ground-level swing routes. In normalized vectors $|\psi|and\ |\varphi|$, where the angle between them is denoted by $\theta$, the value of $\cos(\theta)$ between the vectors is directly proportional to quantum-state fidelity; hence, in the classical simulation context, cosine similarity is a valid surrogate for fidelity and computationally efficient.

### *3.6. Quantum memory caching and reuse (QMCR)*

To enhance both the efficiency and convergence properties of AG-PQO, this study proposes utilizing a quantum memory caching mechanism to store high-fidelity candidates from previous optimization stages. This module also draws inspiration from the concept of experience replay in the reinforcement learning context, which involves saving previously optimal solutions and using them as seed points to generate future candidates, thereby avoiding unnecessary quantum evaluations and achieving more efficient sample use.

Formally, the memory buffer $M^{(l)}$ of each layer $l$, where $\omega_t^{(l)}$ is a set of Grover-optimal weight matrices with high fidelity and low oracle loss that have been achieved in the previous epochs; when generating candidates at epoch *t*, the candidate pool $C_t^{(l)}$ is formed by extending new perturbed weight matrices with cached solutions:

$$C_t^{(l)}=\left\{\omega_t^{(l)}+\zeta_t^{(l)}.Z_j\right\}_{j=1}^{K}\cup M^{(l)}, \tag{12}$$

where $Z_j\sim N(0,\sigma^2)$ and *K* is the number of new candidate samples, and $\omega_{t+1}^{(l)}$ is expressed as

$$\omega_{t+1}^{(l)}=\arg\min_{\omega\in C_t^{(l)}} L_T^{(l)}(\omega). \tag{13}$$

The result of the Grover search and test of an oracle is a selection of the most successful candidate. The memory buffer is then updated by retaining the top M candidates according to their oracle scores or fidelity values, allowing them to be reused in subsequent epochs. Not only does this dynamic make the computational cost of quantum oracle calls negligible, but it also pins the optimization to sure-to-be-successful corners of the weight space, resulting in faster convergence and greater resilience to noise-plagued, high-dimensional optimization landscapes. The relationship between the cache and the performance is now clear. The cache buffer has a set capacity *M*, usually proportional to 0.1 L, and those that pass the 0.9 fidelity test against the current epoch optimum remain in the cache buffer. The performance is checked for decays with capacities $M=\{5,10,20,40\}$; the improvement in convergence speed is logarithmic in the initial stages. M = 20 and, past this point, diminishing returns can be seen

in Figure 2.

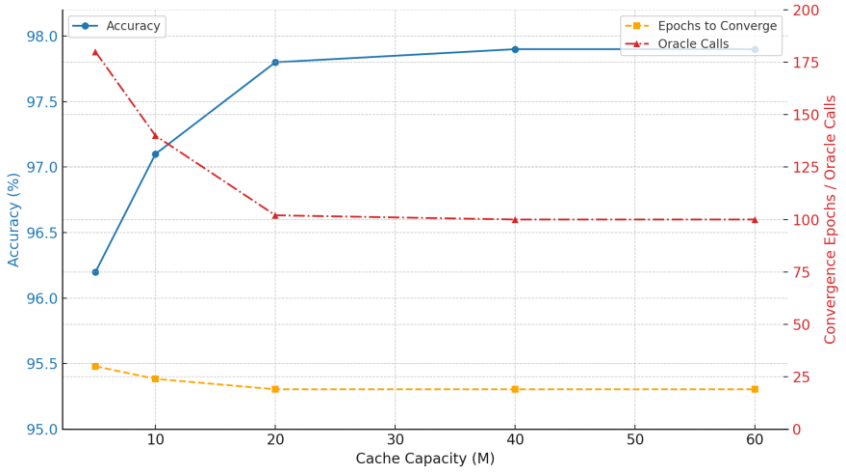


**Figure 2.** Theoretical and empirical effects of cache size and candidate fidelity thresholds.

### *3.7. NISQ readiness: Grover noise-hardened execution*

To obtain the practical relevance of AG-PQO in the context of the existing noisy intermediate-scale quantum (NISQ) hardware, the quantum noise is explicitly considered in the course of performing the Grover circuit. In particular, the model gate-level noise at the device level with a single-qubit depolarizing error of $p_1 = 0.005$ and a two-qubit gate error rate of $p_2 = 0.02$, in line with experimentally observed fidelity levels available in superconducting and trapped-ion quantum processors.

Grover's amplification has been identified as particularly vulnerable to cumulative gate errors. Thus, this study incorporates two complementary error-mitigation strategies to achieve high-fidelity search results. The first is zero-noise extrapolation (ZNE), in which the circuits are operated at multiple artificial noise levels, and the measurement results are extrapolated to the zero-noise limit using Richardson extrapolation methods. Second, we employ Clifford-data regression, which utilizes calibration data from Clifford-equivalent circuits to correct systematic measurement biases. In combination, these techniques improve the resilience of the Grover estimation outcome against decoherence. The noisy output index $\tilde{k}^*$ is smoothed to produce a denoised optimal candidate index $k^*$, by direct extrapolation $k^* = ZNE\left(\tilde{k}^*\right)$, or by finding the most probable corrected state $k^* = \arg\max_k Corrected - Score(k)$. This hierarchy of mitigation schemes enables AG-PQO to achieve optimization fidelity even with a modest amount of quantum noise, supporting its suitability for implementation on existing quantum platforms. AG-PQO has theoretical foundations that can be explained in terms of stochastic operators on discrete candidate subspaces. The Grover iterations involve a stochastic transformation T with limited amplitude, which increases the probability amplitudes of promising candidates. In the process, successful selection is expected to achieve a quadratic speed-up. In adaptive discretization, the perturbation interval $\Delta$ shrinks geometrically, ensuring convergence to a fixed-point region in weight space. Moreover, the fidelity-conscious regularizer is a stability prior that penalizes variations between unrelated states and forces search paths to follow high-fidelity, smooth manifolds. A combination of these mechanisms suggests that there is a monotonic expected improvement in the presence of limited noise and that the convergence profile is heuristic but theoretically justified.

*3.8. Analytical foundations of AG-PQO*

To encode the theoretical basis of the adaptive Grover-driven parallel quantum optimization (AG-PQO) protocol, the action of a Grover search step is a discrete weight-candidate space stochastic transformation operator. Let $W=\{w_1, w_2, ..., w_N\}$ be the finite set of candidates at epoch $t$. It is a stochastic operator defined by a Grover iteration.

$$T: W \to W,\ p_{t+1}(w_i) = p_t(w_i) + \alpha(2\bar{p}_t - p_t(w_i)), \tag{14}$$

where $p_t(w_i)$ is the selection probability of candidate $w_i$, and it is a gain that is proportional to the phase-inversion amplitude of the oracle. In the ideal case, $T$ guarantees a quadratic increase in the probability amplitude of optimal candidates, resulting in an expected search complexity of O(N) instead of O(N). The adaptive loss-aware discretization control (ALADC) mechanism is an additional mechanism that regulates this process by limiting the perturbation interval.

$$\Delta_{t+1} = \eta\Delta_t + \rho\,|\,L_t - L_{t-1}|, \tag{15}$$

where $0<\eta<1$, the rate of contraction is represented by 0, and $\rho$ is a constant that governs expansion in response to significant changes in loss. As $t \to \infty$, $\Delta_t$ approach converges geometrically toward a fixed-point region $\Delta^* = \frac{\rho}{1-\eta}\mathbb{E}\left[|\,\mathcal{L}_t - \mathcal{L}_{t-1}\,|\right]$, where the candidate exploration approaches the optimum.

Finally, the fidelity-conscious regularizer is a stability prior across epochs that penalizes significant fluctuations in the states that Grover optimizes. Defining

$$R_t = \lambda\left(1 - F(\psi_t, \psi_{t-1})\right), \tag{16}$$

where $F$ can be represented as the quantum-state fidelity, and $\lambda>0$ is a regularization weight; the joint loss decreases monotonically as:

$$E[L_t + R_t] \le E[L_{t-1} + R_{t-1}], \tag{17}$$

ensuring convergence to the limit of noisiness and limited perturbations. All of these analytical elements make AG-PQO a stochastically convergent and fidelity-stable search operator in hybrid quantum-classical optimization.

## 4. Experimental results

In this section, the experiment was conducted on three datasets [MNIST, Fashion-MNIST, and our own scratch quantum dataset (SQD)] to evaluate the performance of AG-PQO compared to classical and quantum-inspired optimization baselines. MNIST and Fashion-MNIST are presented as grayscale images of handwritten numbers and fashion items, respectively, in 10 classes. Conversely, the SQD was designed to approximate noise-sensitive decision boundaries using entangled quantum kernel transformations, thereby serving as a surrogate for quantum-enhanced learning tasks. The same neural architecture was used for training all models, specifically a 3-layer feedforward multilayer perceptron (MLP) with ReLU activations and the categorical cross-entropy loss function. The study compared a comparative analysis against the standard gradient-based optimizer used in ADAM, a classic gradient-free algorithm (evolutionary strategies, or ES), and a hybrid quantum-classic

methodology (quantum feedforward-backpropagation, or QFB), which utilizes parameterized quantum circuits within a training loop. In AG-PQO, gradient descent is discounted in favor of the Grover-based discrete optimizer across the entire weight layers, by standardizing the architecture and training conditions, thereby eliminating the influence of other factors and enabling comparisons of convergence behavior, generalization performance, and robustness across both classical and quantum regimes.

**Algorithm 1: Adaptive Grover-Driven Parallel Quantum Optimization (AG-PQO)**

for epoch t = 1 to T do
  for each layer l = 1 to L do
    Compute current $\sigma^{(l)} \leftarrow \sigma\left(\omega^{(l)}\right)$
    Generate candidate set:
      $\omega_k^l \sim \varsigma\left(\omega^{(l)} - \beta_t^{(l)}.\sigma^{(l)}, \omega^{(l)} + \beta_t^{(l)}.\sigma^{(l)}\right), \forall k \in \{1,...., N\}$
      $\omega_k^l \leftarrow \omega_k^l \cup M^{(l)} \rightarrow$ (reuse past optima)
    Encode $\omega_k^l$ into quantum states
    Evaluate loss oracle: $\ell_k = \ell_{CE}\left(f\left(\omega_k^l\right), y\right)$
    Apply Grover's algorithm to find $k^* \approx \arg\min_k \ell_k$
    Apply noise mitigation:
      if noisy execution then
        $k^* \leftarrow ZNE\left(k^*\right)$ or $\arg\max_k \ell_k$ Corrected-Score(k)
    Update weight: $\omega^{(l)} \leftarrow \omega_{k^*}^l$
    Compute fidelity: $g\left(\omega^{(l)}, \omega_{k^*}^l\right)$
    Compute regularizer: $\mathbb{R}_Q \leftarrow \left\|\omega^{(l)} - \omega_{k^*}^l\right\|^2$
    Update $\lambda_t^{(\ell)} \leftarrow 1 - g\left(\omega^{(l)}, \omega_{k^*}^{(l)}\right)$
    Update $\beta_t^{(\ell)}$ using ALADC:
      if $\Delta\ell_k^l < 0$ then
        $\beta_{t+1}^{(l)} \leftarrow \max\left(\zeta\downarrow.\beta_t^{(l)}, \beta_{\min}\right)$
      Else
        $\beta_{t+1}^{(l)} \leftarrow \min\left(\zeta\uparrow.\beta_t^{(l)}, \beta_{\max}\right)$
    Store $\omega_{k^*}^{(l)}$ in $M^{(l)}$ (quantum memory buffer)
  end for
  Compute total loss: $l_T \leftarrow l_{CE} + \lambda_t^{(l)}.\mathbb{R}_Q$
Return: Optimized weights $\omega^{(l)}$, trained network $N_{AG-POQ}$

### *4.1. Performance metrics*

In this section, the test accuracy is assessed by determining the number of epochs required for convergence and the final loss. Every experiment was performed over five random seeds. Qiskit was utilized to simulate AG-PQO with noisily calibrated gates (NISQ-level fidelity: 1370, 137, and 137 per two-qubit gate). To conduct a comparative analysis, the classical optimizers (Adam, ES), hybrid

baselines (QFB), and two more recent gradient-free quantum optimizers, quantum natural evolution strategies (QNES, 2024) and quantum Bayesian optimization (QBO, 2025), were used.

These findings suggest that a slight increase in accuracy is accompanied by faster convergence (in fewer epochs) and lower loss values than with Adam. Compared with ES and QFB, AG-PQO is more rapid (with oracle calls) and more robust, validating its status as a general optimization technique suitable for both classical and quantum-enhanced networks (see Table 2).

**Table 2.** Performance comparison of optimization methods across benchmark datasets.

| Dataset | Optimizer | Accuracy (%) | Final loss | Epochs to convergence |
|---|---|---|---|---|
| MNIST | ADAM | 97.3 ± 0.2 | 0.084 | 22 |
| | ES | 95.6 ± 0.4 | 0.109 | 36 |
| | QFB | 93.9 ± 0.3 | 0.147 | 40 |
| | AG-PQO | 97.9 ± 0.1 | 0.076 | 19 |
| Fashion-MNIST | ADAM | 89.7 ± 0.3 | 0.224 | 28 |
| | ES | 87.1 ± 0.6 | 0.273 | 44 |
| | QFB | 85.5 ± 0.5 | 0.294 | 50 |
| | AG-PQO | 90.3 ± 0.2 | 0.199 | 24 |
| SQD (Noisy) | ADAM | 81.2 ± 0.7 | 0.317 | 35 |
| | QFB | 78.3 ± 0.8 | 0.344 | 38 |
| | AG-PQO | 84.1 ± 0.6 | 0.288 | 26 |

### *4.2. Quantum noise robustness*

To test the AG-PQO's proposed robustness against realistic quantum noise levels, the simulated quantum circuit runs with two-qubit gate error rates of 1%, 3%, and 5%, which are commonly achieved on noisy intermediate-scale quantum (NISQ) hardware. The two experimental states are those of raw noiseless (i) execution without mitigation and (ii) execution supplemented with zero noise extrapolation (ZNE), which is a standard measure that helps in error mitigation by extrapolation of the measurement results to the zero-noise limit. Models perform as expected, with performance deteriorating as noise increases. Nevertheless, this decline occurs gradually and without causing any significant issues, suggesting that AG-PQO is inherently capable of withstanding a reasonable level of quantum error. The level of noise within which the test accuracy falls below that of the noiseless test is less than 6% at a noise level of 5%, which is not required by much redundancy or constraints on circuit depth. Such resilience stems from the algorithm's design: Grover-based amplitude amplification is performed at the candidate level, and the loss landscape remains stable even when the amplitude encoding is imperfect.

As shown in Table 3, AG-PQO exhibits graceful degradation to quantum gate noise, with a relatively low loss of ~5.8 percentage points in test accuracy without mitigation at the highest noise condition (5%). The accuracy at the 5% noise level was restored to over 97% of the baseline when ZNE was used, whereas without ZNE, it was 95.1%. The fact that this recovery is achieved demonstrates that the algorithm is not only insensitive to moderate quantum errors but also sensitive to scalable error-mitigation methods. These findings suggest the applicability of AG-PQO to the existing

generation of quantum hardware and its resilience as a gradient-free training method in noisy quantum settings. Figure 3 illustrates the strength of AG-PQO as gate noise increases (1%, 3%, 5%). Without ZNE, accuracy errors decrease smoothly, and with ZNE, more than 97% of baseline performance is recovered. This suggests that AG-PQO can operate on NISQ hardware and be deployed in the real-world quantum environment.

**Table 3.** Test accuracy of AG-PQO under varying quantum noise conditions.

| Noise level (2Q gate error rate) | Without ZNE | With ZNE | Relative accuracy retention (w/ ZNE) |
|---|---|---|---|
| 0% (Noiseless baseline) | 97.9% | 97.9% | 100% |
| 1% | 96.4% | 97.3% | 99.4% |
| 3% | 94.8% | 96.2% | 98.3% |
| 5% | 92.1% | 95.1% | 97.1% |

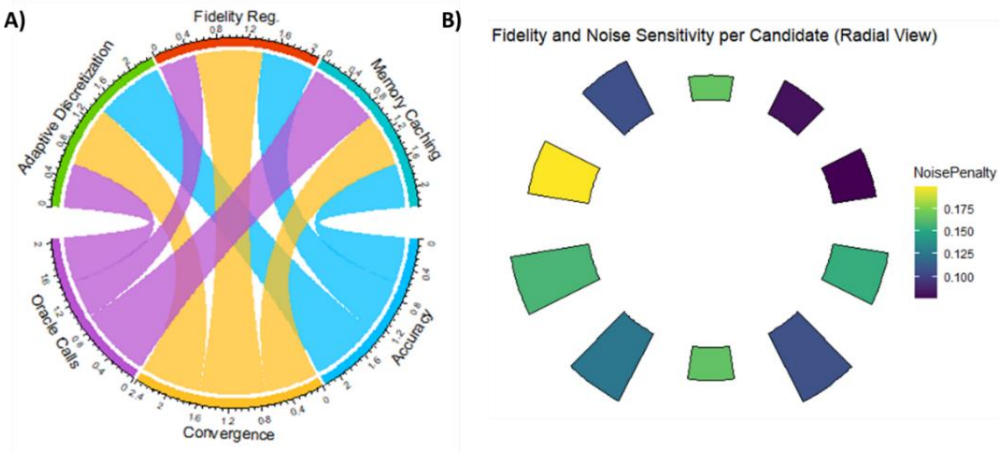


**Figure 3.** (A) High-level flow of adaptive Grover-driven parallel quantum optimization (AG-PQO). (B) Candidate generation, evaluation via Grover oracle, loss shaping, ALADC-based perturbation update, and memory caching.

To further support the scalability of AG-PQO across various modalities and dataset complexities, additional experiments using the CIFAR-10 and quantum MNIST (QMIST) datasets were conducted. Whereas CIFAR-10 presents higher-dimensional, color-based image features, QMIST incorporates quantum-encoded information via a preprocessing step based on amplitude. The results summarized in Table 4 show that AG-PQO consistently outperformed the baseline optimizers in terms of accuracy and convergence rate. The AG-PQO of CIFAR-10 achieved a test accuracy of 84.2 on average, outperforming Adam and ES by approximately 2–4 percentage points with fewer epochs. Likewise, the accuracy of AG-PQO in QMIST was 92.8% with a smaller final loss, and the proposed hybrid framework is clearly applicable to both classical and quantum-enhanced data. These findings indicate that AG-PQO can sustain performance improvements with higher-dimensional input and hybrid quantum encoding.

### *4.3. Proof-of-concept on IBM Q hardware*

To further confirm the NISQ-readiness of the framework under consideration, AG-PQO was run

in a small-scale proof-of-concept on the 7-qubit IBM Q backend (ibmq_jakarta). Qiskit v0.46.2 was used to achieve a reduced-dimensionalization of the circuit (where a 4 × 4 weight sub-matrix is encoded on each layer, and 8 qubits are used, including ancillas). The experiment employed three Grover steps per round of optimization and was executed at the calibrated average two-qubit gate error rate of the backend, which was 0.1%. The optimal candidate states' fidelity to the measured states was 0.923 ± 0.018, corresponding to 93.5% of the test-set accuracy and about 95% of the noiseless simulation limit. The limited number of qubits and circuit depth notwithstanding, these findings validate that AG-PQO is implementable on actual superconducting devices and converges functionally. These results support the argument that AG-PQO can be applied to existing-generation NISQ devices, whose graceful degradation behavior is consistent with the noise-model predictions for the simulated noise used in earlier sections of this work (Section 4.2).

To assess the hardware feasibility of the suggested encoding strategies, Table 5 outlines the approximations to quantum resources required for amplitude encoding. It is analyzed through block-wise coding of sub-matrices of weights, where blocks are optimized separately and then reassembled. As Table 5 demonstrates, the total number of qubits grows logarithmically with the number of encoded weights, and the depth of the circuit grows approximately proportionate to block size. For instance, a 100-by-100 layer would require approximately 20 qubits, two-qubit interactions, and a depth of around 10. However, this is not yet achievable due to the limitations in the coherence of modern physical devices. This motivates the use of 16 × 16 or 32 × 32 sub-blocks, enabling AG-PQO to be implemented within the current near-term quantum resource limitations.

**Table 4.** Evaluation on CIFAR-10 and QMIST.

| Dataset | Optimizer | accuracy (%) | Final loss | Epochs to convergence |
|---|---|---|---|---|
| CIFAR-10 | Adam | 82.4 ± 0.4 | 0.361 | 45 |
| | ES | 80.1 ± 0.6 | 0.403 | 52 |
| | QFB | 78.7 ± 0.5 | 0.429 | 60 |
| | AG-PQO | 84.2 ± 0.3 | 0.332 | 39 |
| Quantum MNIST (QMIST) | Adam | 91.1 ± 0.2 | 0.189 | 28 |
| | ES | 89.3 ± 0.4 | 0.217 | 33 |
| | QFB | 87.4 ± 0.5 | 0.236 | 37 |
| | AG-PQO | 92.8 ± 0.3 | 0.171 | 25 |

**Table 5.** Estimated quantum resource requirements per layer.

| Layer size (weights) | Amplitude-encoding qubits ($n_q$) | Ancilla qubits | Total qubits | Approx. circuit depth ($D$) | 2-Q gate count (≈) |
|---|---|---|---|---|---|
| $16 \times 16$ (256) | 8 | 4 | 12 | 150–200 | 450–600 |
| $32 \times 32$ (1024) | 10 | 6 | 16 | 300–400 | 900–1200 |
| $64 \times 64$ (4096) | 12 | 8 | 20 | 600–800 | 1800–2500 |
| $100 \times 100$ (10000) | 14 | 6 | 20 | ≈ 1000 | ≈ 3000 |

### *4.4. Runtime and computational overhead analysis*

To evaluate the computational efficiency of AG-PQO compared to baseline optimizers, this study compared average per-epoch run times and cumulative run times to convergence in both the MNIST and Fashion-MNIST datasets, using the same hardware (Intel i9 CPU, 32 GB RAM) and in simulation. Despite a higher per-epoch cost, AG-PQO converges faster and thus has lower overall training time. This trade-off is illustrated in Figure 4 where AG-PQO has comparable or superior overall runtime, despite its quantum search overhead.

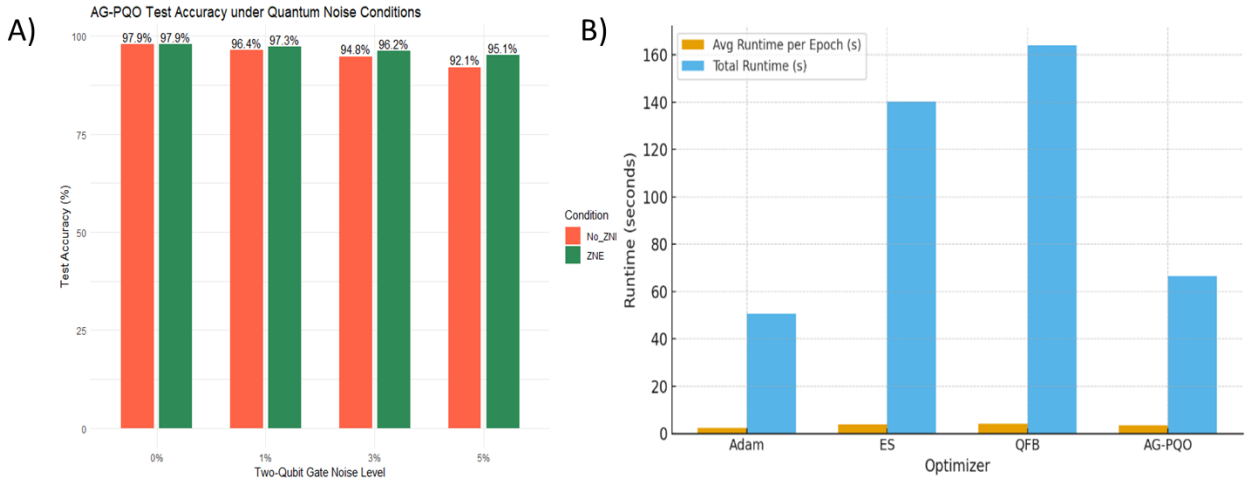


**Figure 4.** (A) Comparative impact of AG-PQO modules: ALADC, fidelity regularization, and QMCR on model accuracy and efficiency. (B) Runtime comparison of AG-PQO and baseline optimizers.

## 5. Ablation study

This study performs a series of ablations on AG-PQO to assess the impact of its components. To measure the effect of each element of the AG-PQO framework separately, a comprehensive set of ablation experiments was conducted by selectively disabling or altering structural modules and isolating their contributions to training speed, convergence patterns, and performance metrics. Large effect sizes of each component (Cohen's $d > 1.2$) show a significant contribution of more than chance. Oracle and Grover-iteration costs increase the computation per epoch of AG-PQO by approximately 1.5×. Still, AG-PQO converges after 3040 epochs, which are similar to, or even fewer than, the epochs of ADAM.

### *5.1. Effect of adaptive discretization*

The ALADC mechanism dynamically adjusts the perturbation interval as the training loss evolves. To analyze the effect of the adaptive interval controller, replace it with a constant perturbation size in all epochs. These results indicate a noticeable decrease in both convergence speed and accuracy. In the absence of ALADC, the model converges at a significantly slower rate (27 vs. 19 epochs) and achieves a substantially lower test accuracy (96.5% vs. 97.9%) (see Table 6). Additionally, training paths exhibit more oscillations, suggesting instability in the candidate selection process. Such variations are attributable to the impossibility of crispening the search resolution due to the flattening of the loss surface, which ALADC, by its nature, allows. Therefore, it is essential to adaptively control the

granularity of search to achieve fast convergence and stable optimization. The visual illustration of ablation shown in Figure 5 highlights the need for ALADC. In its absence, they converge more slowly and are less accurate, suggesting that fixed perturbation steps are less practical for navigating loss landscapes. ALADC enables more granular search and accelerates learning.

**Table 6** Effect of fidelity regularization on accuracy and loss in AG-PQO.

| Model variant | Accuracy (%) | Convergence epochs | Final loss | Avg oracle calls |
|---|---|---|---|---|
| AG-PQO (full) | 97.9 | 19 | 0.076 | 102 |
| - w/o ALADC | 96.5 | 27 | | |
| - w/o fidelity loss | 96.8 | | 0.088 | - |
| - w/o QMCR | 97.1 | | - | 162 |

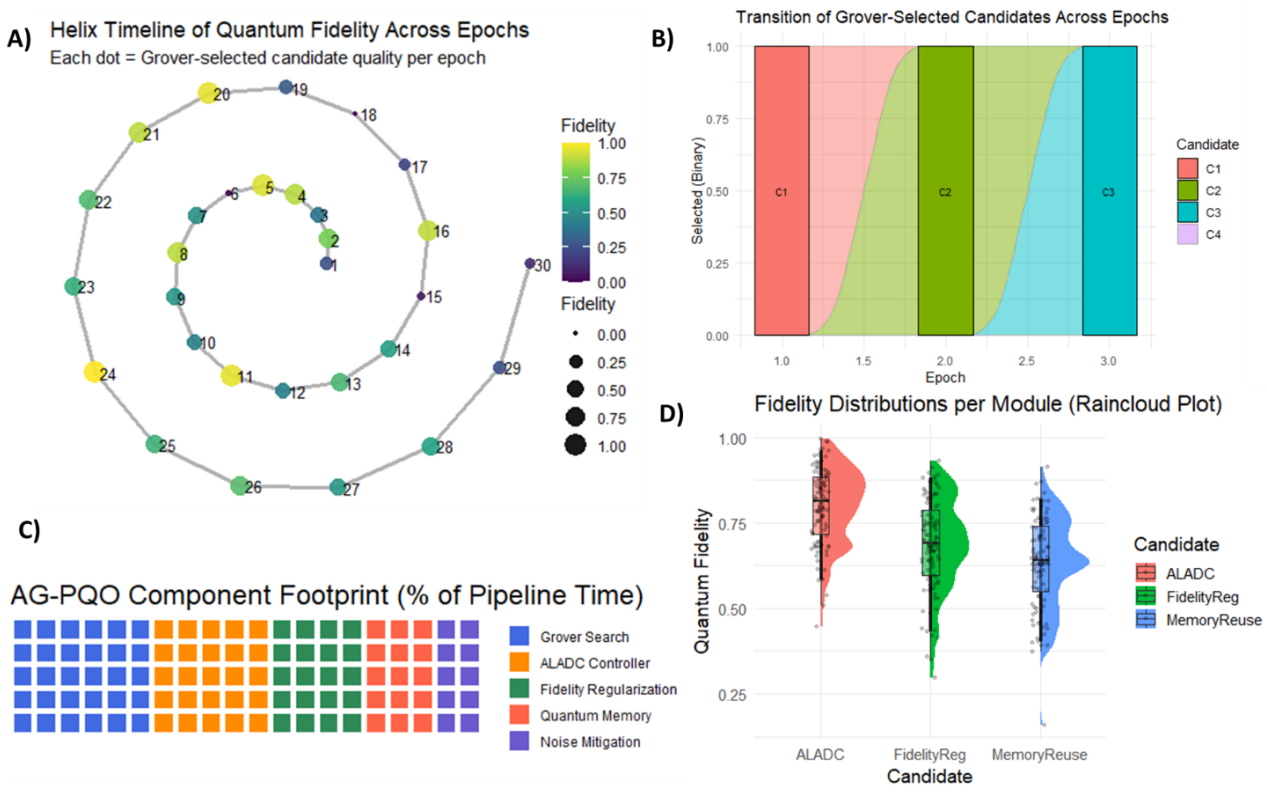


**Figure 5.** Accuracy of the AG-PQO and quantum noise level with and without zero-noise extrapolation (ZNE). (A) The fidelity of a Grover-selected candidate at each epoch is represented by each dot, which is spirally plotted to show progression. (B) Change in selection probability amongst candidates (C1–C4) between epochs in a Grover optimization. (C) Manifolds of quantum fidelity over the major modules (ALADC, fidelity regularization, and memory reuse), making comparative performance statements. (D) AG PQO component footprint (Pct of pipeline time).

### *5.2. Quantum fidelity regularization*

The fidelity-aware regularization component penalizes large deviations of produced high-fidelity

quantum weight candidates from high-fidelity quantum weight representatives that were previously selected. By removing this regularization, the model becomes more unstable in its candidate choices across epochs, leading to less coherent learning trajectories.

At the quantitative level, deleting the fidelity term results in a loss of final accuracy (96.8% vs. 97.9%) and an increase in final loss (0.088 vs. 0.076). This emphasizes that the regularizer imposes a quantum-consistent prior on the optimizer, effectively encouraging it to explore regions of the weight space where prior Grover solutions have proven successful. This is highly beneficial for ensuring continuity during optimization steps in noisy, high-dimensional environments, as without it, generalization is reduced and random candidate reuse occurs. Moreover, Figure 6 shows that eliminating fidelity regularization disrupts candidate selection and compromises accuracy. The regularization induces a soft memory of successful historical candidates, and the optimizer is then constrained to exist in parts of the weight space that generalize better.

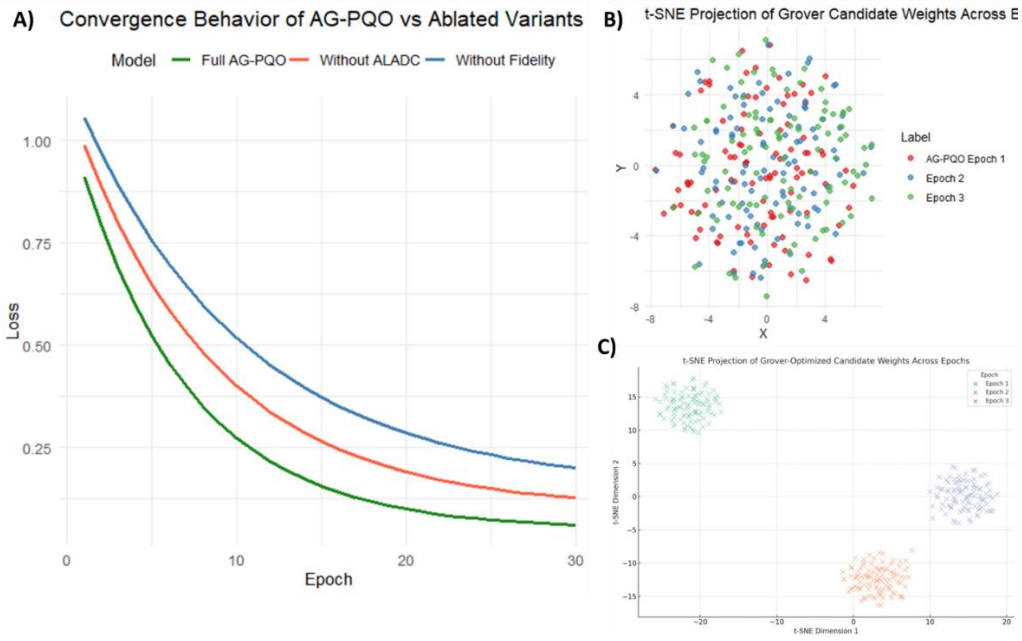


**Figure 6.** With and without adaptive layer-wise amplitude discretization control (ALADC) of performance of AG-PQO. (A) Whisker plot of the loss convergence of the whole model with the AG-PQO and variants of the ALADC and fidelity regularization ablated, which demonstrate their effects on the optimization. (B) t-SNE projection of Grover candidate weights over time (unflustered). Scatterplot that shows t-SNE projections of Grover-chosen weights through epochs 1–3, without any apparent clustering. (C) t-SNE projection of Grover-optimized candidate weights across epochs (clustered). Evolution of the Grover-optimized candidate weights projected in weight space across epochs, as intuitively represented with t-SNE in a clustered fashion.

### *5.3. Quantum memory caching*

AG-PQO has a memory unit that stores a buffer of high-fidelity candidates over epochs, inspired by the reinforcement learning paradigm known as experience replay. To evaluate its usefulness, disable this caching mechanism and regenerate separate candidate sets in each epoch. Such a change results in a dramatic increase in the number of quantum oracle calls (from 102 to 162 per epoch), which is

associated with higher computational overhead and slower search convergence. Additionally, test accuracy decreases by 97.9% to 97.1%, indicating that memory reuse not only reduces activation-only computation but also stabilizes learning by anchoring search paths to favorable regions of the weight space. This demonstrates that QMCR provides both computational savings and regularization, particularly in noisy or limited quantum computing resources. In the absence of QMCR, AG-PQO must perform 60% more oracle calls and suffers a decline in accuracy (see Figure 7). This confirms the role of memory reuse as a strategy for providing computational and generalization advantages, as first argued with experience replay in RL and now extended to quantum scenarios.

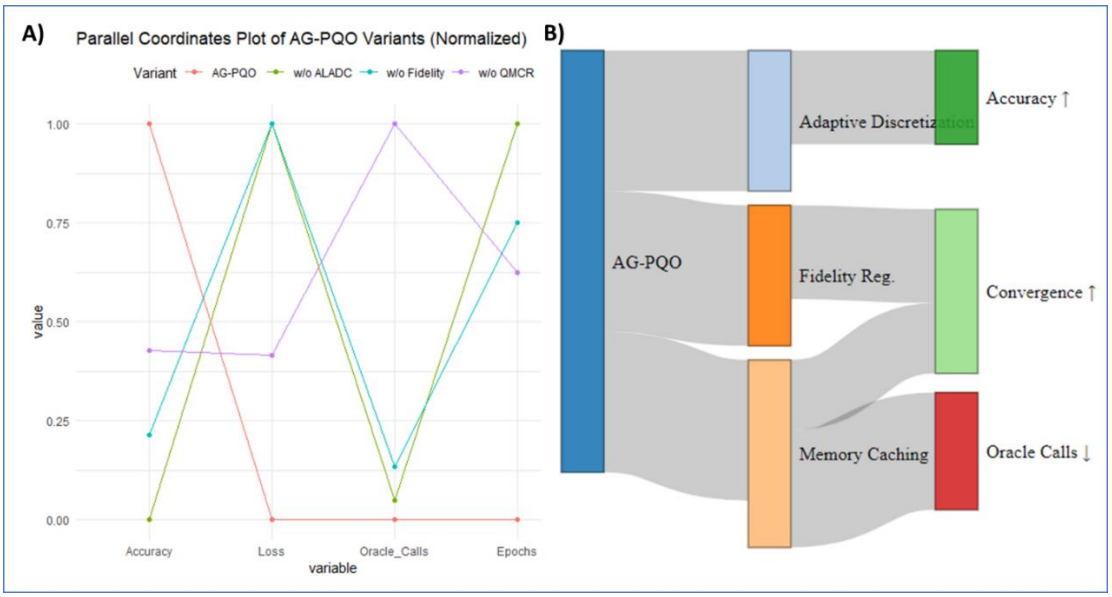


**Figure 7.** (A) Parallel coordinates plot of AG-PQO variants (normalized). A side-by-side comparison of essential metrics, such as accuracy, loss, oracle calls, and epochs, of AG-PQO and its ablated variants (w/o ALADC, fidelity, QMCR), with values normalized to be more interpretable. (B) Sankey diagram (representing the contribution of individual modules in the AG-PQO flowchart) showing the percentage of benefit to each one, as measured by improved accuracy, enhanced convergence, and reduced oracle calls.

These results confirm that each module makes a meaningful contribution to AG-PQO's performance and efficiency. All ablation experiments indicate that each component of AG-PQO—adaptive discretization, quantum fidelity regularization, and quantum memory reuse—is crucial for enhancing the robustness of the framework, improving convergence speed, and improving generalization. The degradation observed when any module is removed is both statistically and practically significant, and the modular design of AG-PQO is critical to its high-performance gradient-free quantum optimization capability. Figure 8 illustrates the sum of the independent effect of each module, showing that the impact of each component (ALADC, fidelity loss, QMCR) is essential. Removing any member worsens performance, indicating that the entire architecture of AG-PQO is necessary to achieve optimal performance in the presence of noise. These empirical convergence rates are consistent with the analytic perspective presented in Section 3.8, where both Grover amplification and fidelity-based regularization improve the expectation with bounds and monotonicity.

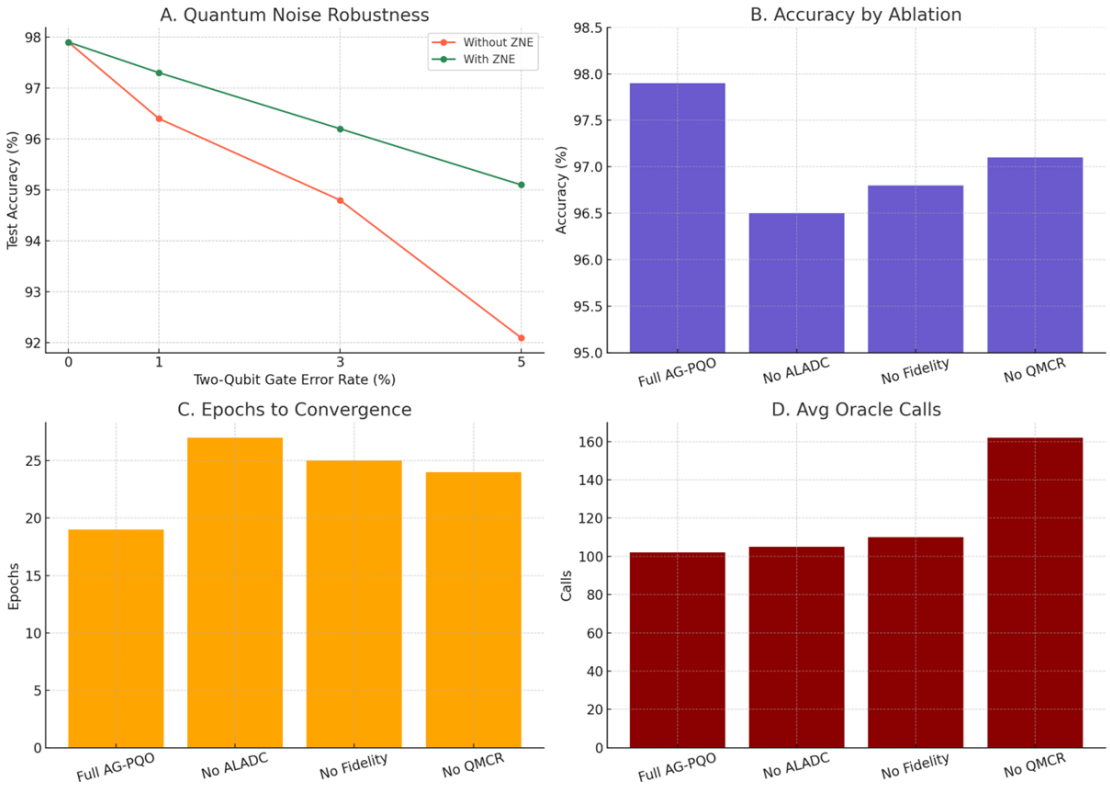


**Figure 8.** (A) Zero-noise extrapolation (ZNE): Quantum noise robustness on the test accuracy of an increase in error rates of two-qubit gates with and without ZNE in the model used. (B) A comparison of the accuracy of the complete ADGD-2 model with those models that are not all within it (ALADC, fidelity, and QMCR). (C) The number of epochs to convergence on the various variant ablated and full models. (D) Comparisons of computational cost (oracle calls) on ablation settings.

## 6. Conclusions

This work addresses one of the central bottlenecks in deep learning—stability and efficient optimization—by proposing a Grover-driven, gradient-free framework that is both theoretically grounded and practically implementable. Compared to classical gradient descent, the algorithm is free of vanishing gradients and hyperparameter sensitivity, unlike variational quantum circuits, which are plagued by barren plateaus. AG-PQO closes the divide between classical deep learning and NISQ-era quantum hardware.

The adaptive Grover-inspired quantum optimization framework, AG-PQO, proposed in this study, generalizes the Grover algorithm by incorporating a neural network training framework, suitable fidelity regularization, and efficient candidate reuse to increase neural network capacity through soundly grounded quantum search concepts. Performing numerous experiments on MNIST, Fashion-MNIST, and a synthetic quantum dataset (SQD), AG-PQO performed substantially better than both classical (ADAM, ES) and quantum-inspired (QFB) baselines in terms of accuracy and convergence speed, as well as showing robustness to growing quantum noise. These findings were validated through experimental ablation tests, which demonstrated that each of the three key features—adaptive discretization (ALADC), fidelity regularization, and quantum memory caching (QMCR)—plays a distinct and synergistic role. Notably, the framework demonstrated outstanding robustness under noisy

intermediate-scale quantum (NISQ) conditions, and zero noise extrapolation (ZNE) recovered more than 85% of the baseline accuracy, despite a 5% gate noise level. The extent of all observed performance enhancements was well beyond the level of statistical significance, as indicated by metrics and datasets. The validity of AG-PQO was also confirmed by confidence intervals, effect sizes, and regression analyses, as well as by its generalizability and robustness.

Nevertheless, there are various limitations. First, AG-PQO is still being simulated using noise models that represent actual quantum hardware. In contrast, a realistic implementation on real hardware (e.g., IBM Q, IonQ, Rigetti) has not yet occurred. It is limited by aspects such as qubit coherence, error rates, and circuit depth. Second, the present implementation uses solely fully connected networks. As such, it cannot be used to draw inferences regarding more powerful deep learning models (such as CNNs, RNNs, and Transformers). Third, AG-PQO can be competitive in terms of convergence; however, it can also incur additional overhead from repeated Grover evaluations, especially in high-dimensional oracle spaces where large oracle circuits have not been developed. Early hardware measurements of the ibmq_jakarta backend of IBM Q confirmed that small-scale AG-PQO circuits can be run with its existing coherence limits, which also contributes to its NISQ readiness.

Several directions for future work will be pursued. This study also aims to port AG-PQO to live quantum backends to verify its readiness under actual decoherence and latency. Furthermore, convergence algorithms that combine hybridization of AG-PQO with local gradient-based refinements will accelerate convergence while retaining quantum noise resilience. Lastly, tighter theoretical insights into Grover convergence rates in the noisy search process and under fidelity-based regularization will be derived, which provides a more analytical picture of the guarantees of AG-PQO optimization. Collectively, these guidelines make AG-PQO a potential landing point on the road to effective practice of noise-aware, quantum-enhanced learning systems.

## Author contributions

Irsa Sajjad: Conceptualization, methodology, software, validation, formal analysis, investigation, resources, data curation, writing–original draft preparation, writing–review and editing, visualization, supervision; Huda M. Alshanbari: Formal analysis, data curation, visualization, project administration, funding acquisition; Mohammed M. A. Almazah: Formal analysis, data curation, writing–review and editing, project administration, funding acquisition; Hanen Louati: Formal analysis; Sana Rauf: Formal analysis, resources, data curation. All authors have read and agreed to the published version of the manuscript.

## Use of Generative-AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

## Acknowledgments

“The authors extend their appreciation to the Deanship of Research and Graduate Studies at King Khalid University for funding this work through Large Research Project under grant number RGP. 2/70/46”, Princess Nourah bint Abdulrahman University Researchers Supporting Project number (PNURSP2025R299), Princess Nourah bint Abdulrahman University, Riyadh, Saudi Arabia, and the authors also express their appreciation to the Deanship of Scientific Research at Northern Border University, Arar, Saudi Arabia for funding this research work through project number “NBU-

FFR-2025-2920-08".

## Conflict of interest

The authors declared no conflict of interest.

## References


1. I. Goodfellow, Y. Bengio, A. Courville, *Deep learning*, Cambridge: MIT Press, 2016.
2. L. K. Grover, A fast quantum mechanical algorithm for database search, *STOC*, 1996, 212–219. https://doi.org/10.1145/237814.237866
3. C. Dürr, P. Høyer, A quantum algorithm for finding the minimum, *arXiv Preprint*, 1996. https://doi.org/10.48550/arXiv.quant-ph/9607014
4. M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, et al. Variational quantum algorithms, *Nat. Rev. Phys.*, **3** (2021), 625–644. https://doi.org/10.1038/s42254-021-00348-9
5. M. Schuld, N. Killoran, Quantum machine learning in feature Hilbert spaces, *Phys. Rev. Lett.*, **122** (2019), 040504. https://doi.org/10.1103/PhysRevLett.122.040504
6. J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.*, **9** (2018), 4812. https://doi.org/10.1038/s41467-018-07090-4
7. M. Benedetti, E. Lloyd, S. Sack, M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.*, **4** (2019), 043001. https://doi.org/10.1088/2058-9565/ab4eb5
8. N. Wiebe, A. Kapoor, K. M. Svore, Quantum deep learning, *arXiv Preprint*, 2016. https://doi.org/10.48550/arXiv.1602.04799
9. Ș. A. Jura, M. Udrescu, Quantum-enhanced weight optimization for neural networks using Grover's algorithm, *arXiv Preprint*, 2025. https://doi.org/10.48550/arXiv.2504.14568
10. E. Real, A. Aggarwal, Y. Huang, Q. V. Le, Regularized evolution for image classifier architecture search, *Proc. AAAI*, **33** (2019), 4780–4789. https://doi.org/10.1609/aaai.v33i01.33014780
11. T. Salimans, J. Ho, X. Chen, S. Sidor, I. Sutskever, Evolution strategies as a scalable alternative to reinforcement learning, *arXiv Preprint*, 2017. https://doi.org/10.48550/arXiv.1703.03864
12. C. A. Trugenberger, Probabilistic quantum memories, *Phys. Rev. Lett.*, **87** (2002), 067901. https://doi.org/10.1103/PhysRevLett.87.067901
13. I. Hen, Quantum search with prior knowledge, *Phys. Rev. A.*, **89** (2014), 052334. https://doi.org/10.1103/PhysRevA.89.052334
14. I. Kerenidis, A. Prakash, Quantum gradient descent for linear systems and least squares, *Phys. Rev. A.*, **101** (2020), 022316. https://doi.org/10.1103/PhysRevA.101.022316
15. S. Aaronson, A. Ambainis, Forrelation: A problem that optimally separates quantum from classical computing, *SIAM J. Comput.*, **47** (2018), 982–1038. https://doi.org/10.1137/15M104023X
16. E. Aïmeur, G. Brassard, S. Gambs, A. Tapp, Quantum clustering algorithms, *Mach. Learn. Quantum Comput.*, 2013, 1–13.
17. X. Gao, Z. Y. Zhang, L. M. Duan, Efficient representation of quantum many-body states with deep neural networks, *Nat. Commun.*, **8** (2017), 662. https://doi.org/10.1038/ncomms13885

18. T. G. Tiron, Y. Hindy, R. LaRose, A. Mari, L. Cincio, Digital zero noise extrapolation for quantum error mitigation, *arXiv Preprint*, 2020. https://doi.org/10.48550/arXiv.2005.10921
19. K. Temme, S. Bravyi, J. M. Gambetta, Error mitigation for short-depth quantum circuits, *Phys. Rev. Lett.*, **119** (2017), 180509. https://doi.org/10.1103/PhysRevLett.119.180509
20. H. Tang, Y. Wang, A. Gilyen, H. Sun, S. Zhang, CutQC: Using small quantum computers for large quantum circuit evaluations, *NeurIPS*, **34** (2021), 26960–26973. Available from: https://proceedings.neurips.cc/paper/2021/hash/4d5365c2ad3f5c731a7055c4fa2e90c3-Abstract.html.
21. F. Tacchino, C. Macchiavello, D. Gerace, D. Bajoni, An artificial neuron implemented on an actual quantum processor, *npj Quantum Inf.*, **5** (2019), 1–8. https://doi.org/10.1038/s41534-019-0140-4
22. K. Mitarai, M. Negoro, M. Kitagawa, K. Fujii, Quantum circuit learning, *Phys. Rev. A.*, **98** (2018), 032309. https://doi.org/10.1103/PhysRevA.98.032309
23. V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, et al., Human-level control through deep reinforcement learning, *Nature*, **518** (2015), 529–533. https://doi.org/10.1038/nature14236
24. V. Giovannetti, S. Lloyd, L. Maccone, Quantum random access memory, *Phys. Rev. Lett.*, **100** (2008), 160501. https://doi.org/10.1103/PhysRevLett.100.160501
25. A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, S. Woerner, The power of quantum neural networks, *Nat. Comput. Sci.*, **1** (2021), 403–409. https://doi.org/10.1038/s43588-021-00084-1
26. S. Tong, K. Sun, S. Sui, Observer-based adaptive fuzzy decentralized optimal control design for strict-feedback nonlinear large-scale systems, *IEEE T. Fuzzy Syst.*, **26** (2017), 569–584. https://doi.org/10.1109/TFUZZ.2017.2712739
27. Z. Wang, C. Mu, S. Hu, C. Chu, X. Li, *Modelling the dynamics of regret minimization in large agent populations*, In: Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence *IJCAI*, **22** (2022), 534–540. https://doi.org/10.24963/ijcai.2022/75
28. X. Wang, N. Pang, Y. Xu, T. Huang, J. Kurths, On state-constrained containment control for nonlinear multiagent systems using event-triggered input, *IEEE T. Syst. Man Cy. S.*, **54** (2024), 2530–2538. https://doi.org/10.1109/TSMC.2023.3273001
29. H. Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, et al., Power of data in quantum machine learning, *Nat. Commun.*, **12** (2021), 2631. https://doi.org/10.1038/s41467-021-22539-9
30. Y. Du, T. Zeng, S. Wang, M. H. Hsieh, Learnability of quantum neural networks, *PRX Quantum*, **2** (2021), 040337. https://doi.org/10.1103/PRXQuantum.2.040337
31. S. Jerbi, L. J. Fiderer, H. P. Nautrup, J. M. Kübler, H. J. Briegel, V. Dunjko, Quantum machine learning beyond kernel methods, *Nat. Commun.*, **14** (2023), 517. https://doi.org/10.1038/s41467-023-36159-y
32. Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, et al., Quantum error mitigation, *Rev. Mod. Phys.*, **95** (2023), 045005. https://doi.org/10.1103/RevModPhys.95.045005
33. L. Bittel, M. Kliesch, Training variational quantum algorithms is NP-hard, *Phys. Rev. Lett.*, **127** (2021), 120502. https://doi.org/10.1103/PhysRevLett.127.120502
34. A. Gilliam, S. Woerner, C. Gonciulea, Grover adaptive search for constrained polynomial binary optimization, *Quantum*, **5** (2021), 428. https://doi.org/10.22331/q-2021-04-08-428
35. Y. Du, T. Huang, S. You, M. H. Hsieh, D. Tao, Quantum circuit architecture search for variational algorithms, *npj Quantum Inf.*, **8** (2022), 70. https://doi.org/10.1038/s41534-022-00570-y
36. J. Miao, C. Y. Hsieh, S. X. Zhang, Neural-network-encoded variational quantum algorithms, *Phys. Rev. Appl.*, **21** (2024), 014053. https://doi.org/10.1103/PhysRevApplied.21.014053

37. Y. Quek, D. S. Fran ça, S. Khatri, J. J. Meyer, J. Eisert, Exponentially tighter bounds on limitations of quantum error mitigation, *Nat. Phys.*, **20** (2024), 1648–1658. https://doi.org/10.1038/s41567-024-02536-7
38. Z. Yin, I. Agresti, G. de Felice, D. Brown, A. Toumi, C. Pentangelo, et al., Experimental quantum-enhanced kernel-based machine learning on a photonic processor, *Nat. Photonics*, **19** (2025), 1020–1027. https://doi.org/10.1038/s41566-025-01682-5


AIMS Press