

BenchRL-QAS: Benchmarking Reinforcement Learning Algorithms for Quantum Architecture Search

Azhar Ikhtiarudin^{*1}, Aditi Das^{*2}, Param Thakkar³, Akash Kundu^{4†}

¹ Computational Materials Design & Quantum Engineering Group, Bandung Institute of Technology, Bandung, Indonesia

² Department of Physics, National Institute of Technology, Karnataka, India

³ Veermata Jijabai Technological Institute, Mumbai, India

⁴ QTF Centre of Excellence, Department of Physics, University of Helsinki, Helsinki, Finland
azharikhtiarudin@gmail.com, aditi.das1601@gmail.com, akash.kundu@helsinki.fi

Abstract

We present BenchRL-QAS, a unified benchmarking framework for reinforcement learning (RL) in quantum architecture search (QAS) across a spectrum of variational quantum algorithm tasks on 2- to 8-qubit systems. Our study systematically evaluates 9 different RL agents, including both value-based and policy-gradient methods, on quantum problems such as variational eigensolver, quantum state diagonalization, variational quantum classification (VQC), and state preparation, under both noiseless and noisy execution settings. To ensure fair comparison, we propose a weighted ranking metric that integrates accuracy, circuit depth, gate count, and training time. Results demonstrate that no single RL method dominates universally, the performance depends on task type, qubit count, and noise conditions providing strong evidence of no free lunch principle in RL-QAS. As a byproduct we observe that a carefully chosen RL algorithm in RL-based VQC outperforms baseline VQCs. BenchRL-QAS establishes the most extensive benchmark for RL-based QAS to date, codes and experimental made publicly available for reproducibility and future advances.

Code — <https://github.com/azhar-ikhtiarudin/bench-rlqas>

Introduction

Quantum computing offers new possibilities for tackling computational challenges that are beyond the reach of classical systems. Yet, the promise of quantum advantage remains difficult to realize, largely due to the practical limitations of current noisy intermediate-scale quantum (NISQ) devices (Preskill 2018). Constraints such as limited qubit counts, hardware connectivity issues, and high error rates collectively limit both the scale and complexity of quantum circuits that can be executed. As a consequence, many theoretically powerful quantum algorithms are not yet implementable on today’s hardware (Monz et al. 2016), underscoring the need for new algorithms specifically tailored to current quantum technologies.

To address these limitations, hybrid quantum-classical strategies have become increasingly popular in the NISQ

era. Among these, variational quantum algorithms (VQAs) stand out as a leading framework for utilizing NISQ hardware efficiently (Peruzzo et al. 2014; Cerezo et al. 2021; Bharti et al. 2022). VQAs leverage classical optimization to tune the parameters of parameterized quantum circuits (PQCs), minimizing a task-specific cost function, often related to the expectation value of a Hamiltonian. The success of this approach depends crucially on the choice of *ansatz*, which defines the PQC structure. Traditionally, these circuit layouts are selected in advance, guided either by hardware constraints (Kandala et al. 2017) or informed by heuristic, problem-specific criteria (Peruzzo et al. 2014). However, such manually designed circuits frequently face an inherent tension between expressiveness and noise resilience, limiting the scalability of VQAs (Cerezo et al. 2021; Bharti et al. 2022; Larocca et al. 2025).

Recent advances have focused on quantum architecture search (QAS) (Zhang et al. 2022), aiming to automate the discovery of optimal PQC structures from a broad set of quantum gates. QAS methodologies adaptively construct circuit architectures that are tailored both to the computational task and to hardware restrictions, searching systematically for gate sequences and placements that maximize performance. Among various QAS techniques, reinforcement learning (RL) has become a particularly promising tool for navigating the large and discrete design space of quantum circuits (Kuo, Fang, and Chen 2021; Fösel et al. 2021; Ostaszewski et al. 2021). In this context, RL agents build PQCs step-by-step, making sequential gate selections and improving their policies based on feedback from quantum performance metrics. Although there have been successful demonstrations of RL-based QAS for circuits with up to 20 qubits (Kundu and Mangini 2025), a broad and systematic understanding of which RL algorithms are most effective for different quantum optimization tasks remains lacking. Closing this gap is critical for advancing both QAS methodologies and the broader goal of practical quantum advantage with near-term hardware.

A wide variety of RL agents have recently been deployed in QAS for VQAs (see Tab. 2). While these studies highlight the potential of RL techniques for quantum circuit construction, the field is still missing a thorough and comprehensive benchmarking of these algorithms. Most prior evalua-

^{*}These authors contributed equally.

[†]Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

RL algorithms	State preparation*				VQSD				VQE (H_2O)				VQC			
	C	D	A	T	C	D	A	T	C	D	A	T	C	D	A	T
A2C (Sutton, Barto et al. 1998)			✓	✓									✓	✓		
A3C (Mnih et al. 2016)	✓	✓	✓													✓
DQN (Mnih et al. 2013)			✓		✓	✓	✓	✓	✓				✓			
DDQN (Van Hasselt 2016)	✓	✓	✓	✓	✓		✓	✓	✓	✓		✓	✓	✓		
DQN_rank (Schaul et al. 2015)			✓		✓		✓			✓		✓	✓			
DQN_per (Schaul et al. 2015)			✓						✓		✓	✓		✓		
Dueling DQN (Wang et al. 2016)			✓				✓	✓		✓			✓	✓	✓	✓
PPO (Schulman et al. 2017)			✓	✓			✓	✓	✓	✓		✓				
TPPO (Wang, He, and Tan 2020)			✓	✓	✓	✓		✓	✓	✓						

Table 1: **Weighted ranking performance (see main text) of RL algorithms is reported for four quantum tasks:** state preparation, VQSD, VQE (H_2O), and VQC. Metrics include smallest circuit (C), depth (D), accuracy (A), and episode time (T), averaged over 5–10 different neural network initializations. A (✓) marks the best result, showing the task and criterion dependence of RL methods for quantum circuit design. "*" indicates a non-parameterized action space. For completeness, we also benchmarked VQE ($B@H_2$); results showed similar trends across RL methods and are omitted here for compactness. DQN_per uses prioritized replay (Schaul et al. 2015); DQN_rank applies rank-based prioritization within DQN_per.

tions focus on only a narrow subset of RL agents, usually in isolated settings and without standardized evaluation criteria. This makes it challenging to determine which algorithms consistently achieve desirable outcomes, be it minimizing circuit depth, reducing 1- and 2-qubit gate usage, or optimizing solution accuracy on quantum processing units (QPUs). The absence of extensive benchmarking keeps open the question of the best-suited RL algorithm for different QAS objectives. Notably, recent works such as (Zhu and Hou 2023) and (Altmann et al. 2024) have compared only a few RL variants, and then solely on non-parameterized action spaces, highlighting the need for systematic benchmarking of parameterized action spaces, particularly in the NISQ context. Beyond variational quantum algorithms, the RL framework has potential applications across a wide range of scientific and engineering disciplines. RL has been used in quantum control (Bukov et al. 2018) and error correction (Nautrup et al. 2019; Olle et al. 2024). In quantum chemistry, RL-assisted ansatz synthesis accelerates ground-state energy estimation of molecules while reducing circuit depth and gate count (Ostaszewski et al. 2021). Moreover, RL is applied to higher energy models (Biswas et al. 2024; Kundu 2025).

In this work, we address this critical gap by presenting a unified benchmark encompassing a broad range of RL algorithms-including both value-based and policy-gradient methods, across several quantum optimization tasks. In particular, we investigate tasks including variational quantum state diagonalization (VQSD), variational quantum eigensolver (VQE), variational quantum classification (VQC), and state preparation, covering system sizes from 2- to 8-qubit. Our benchmark systematically evaluates nine distinct RL agents on each problem. To ensure statistical robustness, each experiment is repeated 5–10 times with independent neural network initializations, resulting in a dataset of 325 separate quantum optimization problem instances. By evaluating all RL approaches under consistent conditions, we provide a comprehensive account of the strengths and limitations of each algorithm for QAS. The highlights of our results are presented in Tab. 1 (noiseless setting) and Tab. 4

(noisy settings), where we employ a *weighted ranking performance estimator*, fully described later in the paper, to identify the best RL agent for each problem type.

Contributions

- We present **BenchRL-QAS**, a platform for systematic benchmarking of RL methods in quantum architecture search (QAS) across diverse variational quantum algorithms (VQAs).
- We provide the largest benchmark to date of RL algorithms for QAS, evaluating parameterized and non-parameterized action spaces under both noiseless and noisy settings. Results show no RL algorithm is universally optimal, demonstrating the *no free lunch* (Wolpert and Macready 1997) principle in QAS.
- We offer a comparative analysis of RL effectiveness, identifying strengths of specific algorithms and giving recommendations to guide future RL-QAS research.

Related Work

Quantum Architecture Search (QAS) One of the early efforts to automate quantum circuit design involves using search heuristics based on evolutionary genetics to generate more efficient and novel quantum circuits (Williams and Gray 1998). Despite their effectiveness in exploring large problem spaces, evolutionary algorithms often yield problem-specific circuits (Rattew et al. 2019; Chivilikhin et al. 2020; Huang et al. 2022). Another approach is differentiable quantum architecture search (Zhang et al. 2022; Chen et al. 2025), which is inspired by the DARTS method in neural architecture search (Liu, Simonyan, and Yang 2018). Grimsley et al. (Grimsley et al. 2019) introduced an adaptive approach for constructing quantum circuits, known as ADAPT-VQE, which has since been further improved (Ramôa et al. 2024). Other techniques have also been explored, including generative models (Nakaji et al. 2024; Furrutter, Muñoz-Gil, and Briegel 2024), Bayesian optimization (Duong et al. 2022; He et al. 2025), and Monte

Related works	RL algorithms
(Ostaszewski et al. 2021; Patel et al. 2024b; Kundu et al. 2024; Kundu and Mangini 2025; Kundu and Sarra 2024)	DDQN
(Ye and Chen 2021)	PPR-DQN
(Olle, Yevtushenko, and Marquardt 2025; Foderà et al. 2024)	PPO
(Kuo, Fang, and Chen 2021; Fösel et al. 2021)	A2C, PPO
(Lockwood 2021)	TD3, SAC, PPO
(Patel et al. 2024a)	REINFORCE
(Altmann et al. 2024)	A2C, PPO, SAC, TD3
(Zhu and Hou 2023)	A2C, PPO, TRPO
(Dutta et al. 2025)	A2C, PPO, DDQN, TD3
BenchRL-QAS (this work)	A2C, A3C, DQN, DQN_per, DQN_rank, Dueling DQN, DDQN, PPO, TPPO

Table 2: **Summary of RL algorithms used in QAS.** In our work, DQN_per refers to DQN with prioritized experience replay (Schaul et al. 2015) and DQN_rank is the rank based prioritization in DQN_per.

Carlo tree search (Wang et al. 2023). Recent work has proposed QAS methods driven by circuit topology (Su et al. 2025) and by landscape fluctuation analysis (Zhu et al. 2025), as well as training-free approaches that rank circuit feature maps via fast proxy metrics (Gujju et al. 2025).

Reinforcement Learning for QAS Numerous studies have successfully applied reinforcement learning (RL) techniques to optimize quantum circuit architectures. These applications include variational quantum eigensolver (VQE) optimization (Ostaszewski et al. 2021), the construction of multi-qubit maximally entangled states (Ye and Chen 2021), and the optimization of quantum machine learning models (Lockwood 2021). Building on these efforts, Patel et al. (Patel et al. 2024b) introduced curriculum RL combined with advanced pruning techniques, which were later utilized by Dutta et al. (Dutta et al. 2025), while Tang et al. (Tang et al. 2024) integrated RL with Monte Carlo tree search and Ruiz et al. (Ruiz et al. 2025) employed tensor decomposition methods. RL-based variational quantum algorithms (VQAs) have also been applied to solve combinatorial optimization problems (Patel et al. 2024a), quantum state diagonalization tasks (Kundu et al. 2024), and Maximum Cut problems (Foderà et al. 2024; Fösel et al. 2021). More recently, RL-based quantum architecture search (RL-QAS) has been further enhanced through the incorporation of tensor network techniques (Kundu and Mangini 2025) and composite gate constructions (gadgets) (Olle, Yevtushenko, and Marquardt 2025; Kundu and Sarra 2024). A summary of the reinforcement learning algorithms used in previous works is provided in Table 2.

BenchRL-QAS

The BenchRL-QAS framework provides a systematic and reproducible platform for benchmarking reinforcement learning (BenchRL) algorithms in quantum architecture search (QAS) across a diverse set of variational quantum algorithms. It addresses the need for standardized evaluation in quantum circuit design, supporting key challenges such as variational quantum state diagonalization (VQSD), variational quantum eigensolver (VQE), variational quantum classification (VQC), and GHZ state preparation. Each task is formulated as a circuit optimization problem, enabling RL agents to discover efficient circuit architectures based on metrics like accuracy, depth, and gate count. The framework is defined in Algorithm 1. In the following we define the RL-state, action space, reward function, illegal actions to accelerate the performance of the agent and a weighted ranking metric to evaluate RL agents.

RL-State The state of the agent is defined by a tensor-based encoding of the current quantum circuit (Patel et al. 2024b). This encoding scheme captures both the structural arrangement of quantum gates, depth, their parameter values, and current achieved accuracy providing a compact yet expressive representation of the ansatz and its performance. The ansatz is expressed as a tensor of dimension $[D_{\max} \times ((N + 3) \times N)]$, where N number of qubits and D_{\max} is the considered maximum depth of the ansatz.

The Reward Function The reward function R in the BenchRL-QAS framework (for VQE, VQSD and VQC problems), is given by (Ostaszewski et al. 2021):

$$R = \begin{cases} 5 & \text{if } C_t \leq \zeta, \\ -5 & \text{if } t \geq D_{\max} \text{ and } C_t \geq \zeta, \\ \max\left(\frac{C_{t-1} - C_t}{C_{t-1} - E_{\min}}, -1\right) & \text{otherwise.} \end{cases} \quad (1)$$

Here $C_t(\vec{\theta})$ is the cost function at step t , and ζ is a predefined convergence threshold i.e. the ansatz accuracy and is treated as a hyperparameter. The ζ is problem specific and discussed in the experimental settings. For the state preparation task the reward is described by (Kundu, Sarkar, and Sadhu 2024):

$$R = \begin{cases} \mathcal{R}, & \text{if } F(s_t) \geq 0.98 \\ F(s_t), & \text{otherwise} \end{cases} \quad (2)$$

\mathcal{R} is a hyperparameter reward with $\mathcal{R} \gg F(s_t)$ and $F(s_t)$ is the fidelity of state at step t .

Action Space The action space in BenchRL-QAS is tailored to the quantum task, supporting both parameterized and non-parameterized circuit construction. For parameterized tasks such as VQE, VQSD, and VQC, it is hybrid: each action specifies a gate type (RX, RY, RZ, CX), its target qubit(s), and, for parameterized gates, a continuous parameter value. This enables the RL agent to simultaneously optimize both the discrete circuit structure and continuous gate parameters. In contrast, for non-parameterized tasks

Algorithm 1: BenchRL-QAS

Require: Quantum tasks \mathcal{T} , RL algorithms \mathcal{A} , encoding scheme, illegal action handling, curriculum (if used)

- 1: **for** task t in \mathcal{T} **do**
- 2: Initialize environment \mathcal{E}_t for t
- 3: **for** RL algorithm a in \mathcal{A} **do**
- 4: Init agent \mathcal{R}_a and empty quantum circuit C
- 5: **while** not converged **do**
- 6: $s \leftarrow \text{EncodeState}(C)$
- 7: $u \leftarrow \mathcal{R}_a.\text{SelectAction}(s)$
- 8: **if** $\text{IsIllegalAction}(u)$ **then**
- 9: Penalize or mask, update agent; **continue**
- 10: **end if**
- 11: Update $C \leftarrow \text{ApplyAction}(C, u)$
- 12: $r \leftarrow \mathcal{E}_t.\text{Evaluate}(C)$
- 13: $s' \leftarrow \text{EncodeState}(C)$
- 14: $\mathcal{R}_a.\text{Update}(s, u, r, s')$
- 15: **end while**
- 16: Log metrics for (t, a)
- 17: **end for**
- 18: [Curriculum] Increase difficulty, if enabled
- 19: **end for**
- 20: Aggregate and compare results

like GHZ state preparation, the action space is limited to a discrete set of gates (CX, X, Y, Z, H, T), where actions involve selecting and placing gates on specific qubits, without continuous parameters. This design allows the framework to benchmark RL agents across a wide range of quantum circuit design problems.

Illegal Actions In BenchRL-QAS, we focus on enforcing only the two most critical illegal action constraints during quantum architecture search. Specifically, an action a is considered illegal in state s if it satisfies either of the following analytical conditions:

$$\begin{aligned} \text{(Redundancy)} \quad & G_{m,q} = G_{m-1,q} \\ \text{(CX repetition)} \quad & G_{m,(q_1,q_2)} = \text{CX} \wedge G_{m-1,(q_1,q_2)} = \text{CX} \end{aligned}$$

where $G_{m,q}$ denotes the gate applied to qubit q at moment m , and $G_{m,(q_1,q_2)}$ denotes a 2-qubit gate between control q_1 and target q_2 . Other potential constraints, such as hardware connectivity, circuit depth, or parameter validity, are not enforced in our implementation. To ensure the agent avoids these illegal actions during training, we assign $Q(a, s) = -\infty$ whenever either of the above conditions is met, effectively masking such actions from the policy optimization.

Weighted Ranking Approach for Agent Evaluation To objectively compare RL algorithms across multiple criteria, we use a *weighted ranking* scheme (Ayan, Abacıoğlu, and Basilio 2023). The key metrics, average circuit error (E), number of gates (G), circuit depth (D), and time per episode (T) are each normalized to $[0, 1]$ (lower is better): $X_{\text{norm}} = \frac{X - X_{\min}}{X_{\max} - X_{\min}}$, where X is the metric value for a given algorithm, and X_{\min} and X_{\max} are the minimum and maximum values of that metric across all algorithms. The composite score for

each algorithm is then computed as

$$S = w_E E_{\text{norm}} + w_G G_{\text{norm}} + w_D D_{\text{norm}} + w_T T_{\text{norm}}. \quad (3)$$

Lower S indicates better overall performance. Algorithms are ranked by ascending S , ensuring that accuracy is the most influential criterion in the final ranking, while still incorporating resource and efficiency considerations. We explicitly choose the weights $[w_E, w_G, w_D, w_T] = [0.5, 0.2, 0.2, 0.1]$ for noiseless results in Tab. 1 and $[0.6, 0.1, 0.3, 0.0]$ for noisy in Tab. 4, placing the greatest emphasis on problem accuracy (circuit error), while still accounting for circuit size, depth, and computational efficiency.

Experimental Settings

Our experimental settings span a range of quantum circuit design tasks, each presenting a distinct challenge for RL-based quantum architecture search. VQSD requires agents to diagonalize arbitrary quantum states without Hamiltonian dependence; VQE targets ground state energy approximation for specific Hamiltonians; VQC trains parameterized circuits for supervised learning on synthetic data; and GHZ state preparation constructs non-parameterized circuits to generate maximally entangled states essential for quantum information. By casting these as circuit optimization problems, BenchRL-QAS enables systematic evaluation of RL agents on accuracy, circuit depth, and gate count, providing a unified and versatile benchmarking platform. Throughout the paper BenchRL-QAS utilizes vanilla curriculum (Ostaszewski et al. 2021). Moreover we utilize 9 different RL-agents and benchmark the performance of these agents. In the following we utilize neural network consists of L layers, each containing 1000 neurons, a batch size of 1000, and a replay memory of 20000 transitions. The learning rate is set to 3×10^{-4} using the ADAM (Kingma 2014), with no dropout, and the target network is updated every 500 steps. Exploration is controlled by an epsilon-greedy policy with epsilon decaying from 1.0 to a minimum of 0.05 at a rate of 0.99995 per step, and a discount factor gamma of 0.88 is used. Circuit parameters are optimized using COBYLA with a maximum 500 iterations. we detail the algorithms and their agent-environment specifications. Furthermore, specifically for the A3C agent we utilize 3 workers for all optimization problems. All the results are obtained utilizing AMD Rome 7H12 CPU based on AMD Zen 2 architecture and an Nvidia Ampere A100 GPU. We utilized 2 CPUs with maximum 4GB and maximum 40 GB GPU memory.

RL-VQSD

Variational quantum state diagonalization (VQSD) is a type of variational quantum algorithm that aims to find a unitary transformation which diagonalizes a quantum state in the computational basis (LaRose et al. 2019). More specifically, given a quantum state ρ , the method optimizes parameters $\vec{\theta}$ of a parameterized unitary $U(\vec{\theta})$ such that:

$$\rho' = U(\vec{\theta}_{\text{opt}}) \rho U(\vec{\theta}_{\text{opt}})^\dagger = \rho_{\text{diag}}, \quad (4)$$

where ρ_{diag} is the diagonal form of ρ in its eigenbasis. Compared to quantum principal component analysis (Lloyd, Mohseni, and Rebentrost 2014), which promises exponential speedup but requires many qubits and deep circuits, VQSD is more suitable for current noisy quantum devices due to its shallow circuits and reduced hardware demands. VQSD has applications in fidelity estimation of quantum states (Cerezo et al. 2020), and quantum device certification (Kundu and Miszczak 2022). A central challenge in implementing VQSD lies in designing a hardware-efficient and scalable ansatz. To address this challenge, reinforcement learning has recently been introduced as a tool to automate and optimize the construction of ansatz circuits for variational quantum algorithms (RL-VQSD) (Kundu et al. 2024).

Under the BenchRL-QAS framework we consider several 2-qubit Haar random mixed quantum states and diagonalize them utilizing the VQSD algorithm. The agent-environment setup consists of maximum depth $D_{\text{max}} = 40$, with the agent aiming to minimize a cost function below a threshold $\zeta = 5 \times 10^{-2}$ in reward (see Eq. 1) with $L = 4$ layers of neural network.

RL-VQE

The variational quantum eigensolver (VQE) is designed to find the eigenstates and eigenvalues of a given Hamiltonian, with particular emphasis on estimating the ground state energy. VQE operates by minimizing a cost function defined as:

$$C(\vec{\theta}) \equiv E(\vec{\theta}) = \langle \psi(\vec{\theta}) | \hat{H}_q | \psi(\vec{\theta}) \rangle, \quad (5)$$

where \hat{H}_q is the qubit Hamiltonian and $|\psi(\vec{\theta})\rangle$ is a parameterized wavefunction. The qubit Hamiltonian \hat{H}_q is derived from the molecular system, which includes information such as molecular geometry and atomic charges. This leads to a fermionic Hamiltonian expressed in the second quantization formalism. The fermionic Hamiltonian is then mapped to a qubit representation using techniques such as the Jordan-Wigner (Fradkin 1989), Bravyi-Kitaev (Bravyi and Kitaev 2002), and the qubit tapering (Bravyi et al. 2017). VQE is a leading method in quantum chemistry and materials science for estimating molecular ground state energies, enabling insights into system properties, reaction prediction, and materials design (Peruzzo et al. 2014; Tilly et al. 2022; Fedorov et al. 2021). Unlike quantum phase estimation (Kitaev 1995), which demands deep, complex circuits impractical for today’s hardware, VQE is well-suited to current quantum devices. However, traditional ansatzes like UCCSD and ADAPT-VQE (Grimsley et al. 2019) often result in circuits too deep for near-term quantum processors, while hardware-efficient ansatzes face scalability and trainability challenges. To overcome these limitations, reinforcement learning-based approaches (Ostaszewski et al. 2021; Patel et al. 2024b) have been developed to automate and scale the construction of efficient VQE ansatzes.

In BenchRL-QAS the environments utilize 4- H_2 , 6- BeH_2 , and 8- H_2O with D_{max} , 40, 70, and 250 steps each mapped using the Jordan-Wigner transformation. The reward function is similar to that defined in Eq. 1, but the cost function at each step is obtained by Eq. 5. The main goal of

the agent is to achieve chemical accuracy $\zeta = 1.6 \times 10^{-3}$. The L is of 3-, 4-, and 5- layers for 4-, 6-, and 8-qubit problem. The circuit optimization is performed using COBYLA with 100 (4-qubit), 200 (6-qubit), or 500 (8-qubit) iterations.

RL-VQC

Variational quantum classifier (VQC) is a hybrid quantum-classical algorithm for supervised learning problems such as classification (Gil Fuster 2019). Like other variational quantum algorithms, VQC consists of a parameterized quantum circuit used in conjunction with a classical optimizer to learn from labeled data. For a given input $x = (x_1, x_2) \in \mathbb{R}^2$, the quantum circuit initially encodes the data using single qubit rotations $R_y(x_i\pi)$. The encoded state is subsequently processed through a trainable ansatz comprising further rotation gates and optional entangling gates such as CX or CZ. The circuit ends with measurements on all qubits, and the outcomes are used to calculate a prediction-dependent cost function. The behavior of the circuit is based on the trainable parameters $\vec{\theta} = (\theta_1, \dots, \theta_l)$ that are learned by a classical algorithm (usually gradient descent) to minimize a quadratic cost function:

$$C = \frac{1}{2n} \sum_x |y(x; \vec{\theta}) - a(x)|^2, \quad (6)$$

where $y(x; \vec{\theta})$ is the circuit’s output for input x and $a(x)$ is the true label. The model continues to improve iteratively to minimize this cost. After the VQC outputs measurement results interpreted as class labels, these predicted labels are compared against the true labels to compute accuracy. VQC have demonstrated better accuracy and noise robustness over classical neural networks in accelerator physics (Yin et al. 2025) but the performance of VQC is highly dependent on diligent ansatz design, circuit depth, and task-dependent encoding schemes.

To tackle this, we incorporate RL to automate the design of quantum circuit for VQC namely RL-VQC. We consider 3-qubit synthetic binary classification data with a predefined training error threshold $\zeta = 0.2$ i.e. achieving at least 80%

Method	Train accuracy	Test accuracy
Hardware efficient ansatz		
2 layers	90%	85%
3 layers	90%	85%
4 layers	92%	92%
(Du et al. 2022)		
W = 1, T = 10	N.A.	50% – 60%
W = 1, T = 400	N.A.	> 90%
W = 5, T = 400	N.A.	> 90%
RL-VQC (This work)	99.996%	99.991%

Table 3: **Training and testing accuracy comparison:** RL-VQC with DQN and rank-based prioritized replay achieves the highest accuracy, outperforming hardware-efficient ansatz and net-based methods (Du et al. 2022) in the noiseless scenario without retraining.




Figure 1: **Benchmark of circuit depth, CX, and single-qubit gate counts for RL-based quantum architecture search (QAS) across VQSD (2-qubit), VQC (3-qubit), and VQE (4-H₂, 6-BeH₂ and 8-H₂O) in noiseless scenario.** Our RL approach delivers circuits for 6-BeH₂ with errors over three orders of magnitude lower and gate counts less than half those of the recent TF-QAS method, establishing one of the new standards for QAS.

accuracy in the training loss $\mathcal{L}_{\text{training}}$ within a maximum allowable depth $D_{\text{max}} = 25$ and a neural net of $L = 3$ layers. In the reward in Eq.1, the cost function is evaluated according to Eq. 6. The quantum circuit parameters are optimized using COBYLA with 1000 iterations per step. Under BenchRL-QAS we benchmark RL-VQC for 9 different agents. The preliminary results in Tab. 3 show that the DQN_rank has the optimal performance, with 99.996% training and 99.991% test accuracy, outperforming both the hardware-efficient ansatz (HEA), as well as net-based approaches (Du et al. 2022). The HEA used here is a layered variational circuit of iterated single-qubit rotations RY and nearest-neighbor CNOT gates applied after a data-dependent RY feature encoding. A broader benchmark is provided in results section.

RL State Preparation

Beyond the VQAs, we also utilize the BenchRL-QAS quantum state preparation, which in this case preparing

the Greenberger–Horne–Zeilinger (GHZ) state, defined as: $|GHZ\rangle = \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle)$, with non-parameterized gate-set (CX, X, Y, Z, H, T). The aim of BenchRL-QAS is to construct a quantum circuit that reproduces the $|GHZ\rangle$ state, the closer the RL outcome is to the $|GHZ\rangle$ state the higher the reward to the agent, following the reward in Eq. 2.

Results

Noiseless

The benchmarking study in Fig. 1 systematically evaluates reinforcement learning (RL) algorithms for quantum architecture search (QAS) across three representative variational quantum algorithms: variational quantum state diagonalization (VQSD, 2-qubit), variational quantum classifier (VQC, 3-qubit), and variational quantum eigensolver (VQE, 4-H₂, 6-BeH₂, 8-H₂O). The VQE figure highlights that DDQN, DQN and DQN_rank consistently minimize circuit depth, CX count, and single-qubit gate count as problem size increases, while PPO and TPPO emerge as strong contenders

RL algorithm	4-H ₂			6-BeH ₂			8-H ₂ O		
	Error	Gates	Depth	Error	Gates	Depth	Error	Gates	Depth
A2C	7.29×10^{-6}	15	8	8.25×10^{-5}	26	15	4.13×10^{-4}	105	44
A3C	3.33×10^{-5}	19	12	6.99×10^{-7}	32	14	6.87×10^{-4}	85	34
DDQN	3.72×10^{-6}	7	4	5.04×10^{-6}	65	35	1.82×10^{-4}	173	158
DQN	1.11×10^{-6}	17	11	1.44×10^{-5}	62	30	2.71×10^{-5}	137	68
DQN_PER	4.57×10^{-6}	8	7	2.73×10^{-6}	34	17	5.67×10^{-5}	64	32
DQN_rank	4.99×10^{-6}	24	14	3.09×10^{-6}	69	31	3.13×10^{-6}	168	105
Dueling_DQN	8.43×10^{-6}	2	2	8.92×10^{-6}	28	17	1.50×10^{-5}	216	112
PPO	1.46×10^{-6}	25	21	1.21×10^{-4}	11	6	1.22×10^{-5}	97	41
TPPO	3.86×10^{-6}	15	11	1.01×10^{-5}	26	12	4.77×10^{-5}	34	21

Table 4: **Benchmarking RL agents for QAS under realistic noise applies 0.1% single-qubit and 0.01% two-qubit depolarizing noise after each gate.** Results report the best-performing seeds across runs, highlighting agent robustness and efficiency in hardware-like conditions.

for intermediate and larger problem sizes, often matching or outperforming A3C and DDQN in both depth and gate efficiency. For VQSD and VQC, both value-based (DQN, Dueling DQN) and advanced policy-gradient methods (A3C, TPPO) yield compact, expressive circuits, with robust exploration and stable policy improvement proving advantageous in smaller and moderately complex settings. DDQN and similar variants generally lag in larger VQE problems due to conservative updates. Overall, no single RL paradigm is universally optimal; rather, algorithmic features such as asynchronous updates, trust-region constraints, and action prioritization are critical for hardware-efficient quantum circuits, and the RL algorithm should be selected based on the structure and requirements of the quantum problem.

Notably, *BenchRL-QAS on the 6-BeH₂ outperforms TF-QAS (He et al. 2024), which achieves an error 0.0018 with 57 gates. In contrast, we obtain an error on the order of 10^{-6} over three orders of magnitude improvement, while reducing the total gate to 16 – 22 depending on the RL-agent. This improvement in both accuracy and hardware efficiency establishes a new state-of-the-art for VQE ansatz synthesis.*

Noisy

The performance results in Table 4 demonstrate significant variability among RL algorithms when tasked with quantum circuit design under realistic noise settings. We consider 0.1% 1-qubit depolarizing noise and 0.01% of 2-qubit depolarizing noise to benchmark the performance of RL agents. To make the setting more realistic the noise is applied after the application of each gate while constructing the ansatz. Some algorithms, such as *DQN and PPO, achieve notably low error rates for certain qubit sizes, while others, such as DDQN and Dueling DQN, excel in minimizing gate counts or circuit depth. This diversity in strengths highlights that no single algorithm consistently outperforms all others across every metric (error, gates, depth) and problem size.*

If we consider the weighted ranking metric in Eq. 3 with $w_E = 0.6$, $w_G = 0.1$, $w_D = 0.3$, and $w_T = 0$ the best RL agents are DDQN (4-H₂), PPO (6-BeH₂), and TPPO (8-H₂O). This weighting is well-suited for realistic noisy quantum devices, as it prioritizes accuracy and depth crucial for noise resilience, while still considering gate efficiency.

Runtime Analysis of Agents

Table 5 presents the average per-episode runtime (in seconds) for RL algorithms across the RL-VQE, RL-VQC, and RL-VQSD tasks, with a final column highlighting 3-qubit state preparation, thereby illustrating each method’s computational efficiency. Policy-gradient methods, particularly PPO, consistently achieve the fastest times in most VQE scenarios (1.56s for 6-BeH₂, 35.06s for 8-H₂O) and the lowest state preparation runtime (0.0419s), while TPPO is also efficient (e.g., 1.84s for 6-BeH₂ and 0.95s for 2-qubit VQSD); DDQN excels in 4-H₂(0.24s) and is competitive in VQC (0.81s) and state preparation (0.0714s). In contrast, value-based methods tend to incur greater costs as system size grows, exemplified by DQN_rank’s 159.15s per episode on 8-H₂O and generally slower performance from DQN_PER and A3C, with the slowest recorded runtime from A2C (265.87s on 8-H₂O). Notably, state preparation is uniformly swift across all methods (all < 0.26s), suggesting computational bottlenecks arise from agent-environment interaction and neural network training rather than quantum simulation overhead. These results indicate that policy-gradient approaches like PPO and TPPO scale more efficiently and are thus better suited to large-scale quantum circuit design, while value-based methods require further optimization for practical use in large systems.

Summary

Our benchmarking across noiseless and noisy regimes reveals a central result: *no single RL algorithm is universally optimal for quantum circuit design.* This aligns with the No Free Lunch Theorem (Wolpert and Macready 1997), which states that no optimizer consistently outperforms others across all tasks. Empirically, we find that algorithmic performance, measured by circuit depth, gate count, and error which varies with problem structure and noise. Value-based methods (e.g., DQN, DQN_rank) excel for VQE in noiseless settings, while policy-gradient methods (e.g., A3C, TPPO) are more effective for VQSD and VQC. Under noise, the best-performing algorithm depends on the specific metric and qubit size. *These results highlight that RL algorithm selection must be tailored to each quantum problem and noise regime. Systematic benchmarking, as performed here, is essential for identifying the most suitable approach.* In

RL Algorithm	RL-VQE			RL-VQC 3-qubit	RL-VQSD 2-qubit	RL-State preparation* 3-qubit
	4-H ₂	6-BeH ₂	8-H ₂ O			
PPO	0.44	1.56	35.06	2.96	1.11	0.04194
Dueling_DQN	0.81	2.57	72.51	4.60	1.53	0.15503
DDQN	0.24	2.64	66.58	0.81	1.70	0.07143
A3C	0.78	2.68	56.11	8.27	2.43	0.16253
DQN	0.76	2.58	66.44	4.62	1.58	0.14595
DQN_PER	1.03	3.11	85.55	4.76	1.85	0.23197
DQN_rank	1.04	3.27	159.15	4.64	1.90	0.25066
TPPO	0.49	1.84	132.52	1.51	0.95	0.04844
A2C	0.75	2.19	265.87	2.88	1.66	0.04729

Table 5: Average per-episode runtime (in seconds) for each RL algorithm across quantum architecture search tasks (RL-VQE, RL-VQC, RL-VQSD and RL state preparation). Lowest times per task are highlighted in green color. * denotes the action space for this task is non-parameterized.

summary, our study provides practical evidence for the NFL principle in RL-based quantum circuit design: algorithmic strengths are context-dependent, and no single method dominates across all scenarios.

Discussion and Conclusion

We present **BenchRL-QAS**, a unified benchmarking framework for reinforcement learning (RL) algorithms in quantum architecture search (QAS), systematically evaluating both policy-gradient and value-based agents across a diverse set of quantum variational tasks and system sizes (2- to 8-qubit). BenchRL-QAS introduces a weighted performance ranking for fair comparison and makes all code and data publicly available for reproducibility. Our results demonstrate that no single RL agent is universally optimal which is a direct reflection of the “No Free Lunch” principle (Wolpert and Macready 1997): value-based methods (e.g., DQN, DQN rank, Dueling DQN) excel in deeper VQE tasks, while policy-gradient methods (e.g., A3C, TPPO) are more effective for structured or smaller problems like VQSD and VQC; under noise, optimal performance is metric- and problem-dependent. This highlights the need for tailored algorithm selection and comprehensive benchmarking in RL-driven quantum circuit design.

Limitations and future work. While BenchRL-QAS represents the most comprehensive RL-QAS benchmarking effort to date, several limitations remain. Our current study evaluates a utilizes vanilla curriculum RL-QAS, which is not the best performing if we consider advanced methods such as TensorRL-QAS (Kundu and Mangini 2025) or Gadget RL-QAS (Kundu and Sarra 2024). Moreover, while the tasks in this study span multiple VQA classes such as quantum approximate optimization algorithm (Farhi, Goldstone, and Gutmann 2014) were not included. Additionally, the noisy scenario was limited to the VQE algorithm. Incorporating a broader set of quantum optimization problems under realistic QPU constraints (e.g., qubit connectivity and/or gate fidelity) would enhance the generality of the framework.

Acknowledgments

We thank QOSF for providing the platform for collaboration, we also thank Maria Demidik for fruitful discussions.

Finally we thank Center of Science (CSC) IT centre in Finland for providing us with the CPU and GPU resources.

References

- Altmann, P.; Stein, J.; Kölle, M.; Bärligea, A.; Zorn, M.; Gabor, T.; Phan, T.; Feld, S.; and Linnhoff-Popien, C. 2024. Challenges for reinforcement learning in quantum circuit design. In *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, volume 1, 1600–1610. IEEE.
- Ayan, B.; Abacıoğlu, S.; and Basilio, M. P. 2023. A comprehensive review of the novel weighting methods for multi-criteria decision-making. *Information*, 14(5): 285.
- Bharti, K.; Cervera-Lierta, A.; Kyaw, T. H.; Haug, T.; Alperin-Lea, S.; Anand, A.; Degroote, M.; Heimonen, H.; Kottmann, J. S.; Menke, T.; et al. 2022. Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, 94(1): 015004.
- Biswas, A.; Acquarone, M.; Wang, H.; Miretti, F.; Misul, D. A.; and Emadi, A. 2024. Safe reinforcement learning for energy management of electrified vehicle with novel physics-informed exploration strategy. *IEEE Transactions on Transportation Electrification*, 10(4): 9814–9828.
- Bravyi, S.; Gambetta, J. M.; Mezzacapo, A.; and Temme, K. 2017. Tapering off qubits to simulate fermionic Hamiltonians. *arXiv preprint arXiv:1701.08213*.
- Bravyi, S. B.; and Kitaev, A. Y. 2002. Fermionic quantum computation. *Annals of Physics*, 298(1): 210–226.
- Bukov, M.; Day, A. G.; Sels, D.; Weinberg, P.; Polkovnikov, A.; and Mehta, P. 2018. Reinforcement learning in different phases of quantum control. *Physical Review X*, 8(3): 031086.
- Cerezo, M.; Arrasmith, A.; Babbush, R.; Benjamin, S. C.; Endo, S.; Fujii, K.; McClean, J. R.; Mitarai, K.; Yuan, X.; Cincio, L.; et al. 2021. Variational quantum algorithms. *Nature Reviews Physics*, 3(9): 625–644.
- Cerezo, M.; Poremba, A.; Cincio, L.; and Coles, P. J. 2020. Variational quantum fidelity estimation. *Quantum*, 4: 248.
- Chen, S. Y.-C.; Liu, C.-Y.; Chen, K.-C.; Huang, W.-J.; Chang, Y.-J.; and Huang, W.-H. 2025. Differentiable quantum architecture search in quantum-enhanced

- neural network parameter generation. *arXiv preprint arXiv:2505.09653*.
- Chivilikhin, D.; Samarin, A.; Ulyantsev, V.; Iorsh, I.; Oganov, A. R.; and Kyriienko, O. 2020. MoG-VQE: Multiobjective genetic variational quantum eigensolver. *arXiv preprint arXiv:2007.04424*.
- Du, Y.; Huang, T.; You, S.; Hsieh, M.-H.; and Tao, D. 2022. Quantum circuit architecture search for variational quantum algorithms. *npj Quantum Information*, 8(1): 62.
- Duong, T.; Truong, S. T.; Tam, M.; Bach, B.; Ryu, J.-Y.; and Rhee, J.-K. K. 2022. Quantum neural architecture search with quantum circuits metric and bayesian optimization. *arXiv preprint arXiv:2206.14115*.
- Dutta, S.; Innan, N.; Yahia, S. B.; and Shafique, M. 2025. QAS-QTNs: Curriculum Reinforcement Learning-Driven Quantum Architecture Search for Quantum Tensor Networks. *arXiv preprint arXiv:2507.12013*.
- Farhi, E.; Goldstone, J.; and Gutmann, S. 2014. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*.
- Fedorov, D. A.; Peng, B.; Govind, N.; and Alexeev, Y. 2021. VQE Method: A Short Survey and Recent Developments.
- Foderà, S.; Turati, G.; Nembrini, R.; Dacrema, M. F.; and Cremonesi, P. 2024. Reinforcement Learning for Variational Quantum Circuits Design. *arXiv preprint arXiv:2409.05475*.
- Fösel, T.; Niu, M. Y.; Marquardt, F.; and Li, L. 2021. Quantum circuit optimization with deep reinforcement learning. *arXiv preprint arXiv:2103.07585*.
- Fradkin, E. 1989. Jordan-Wigner transformation for quantum-spin systems in two dimensions and fractional statistics. *Physical review letters*, 63(3): 322.
- Fürrutter, F.; Muñoz-Gil, G.; and Briegel, H. J. 2024. Quantum circuit synthesis with diffusion models. *Nature Machine Intelligence*, 6(5): 515–524.
- Gil Fuster, E. M. 2019. Variational quantum classifier.
- Grimsley, H. R.; Economou, S. E.; Barnes, E.; and Mayhall, N. J. 2019. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature communications*, 10(1): 3007.
- Gujju, Y.; Harang, R.; Li, C.; Shibuya, T.; and Zhao, Q. 2025. QuProFS: An Evolutionary Training-free Approach to Efficient Quantum Feature Map Search. *arXiv preprint arXiv:2508.07104*.
- He, Z.; Chen, H.; Zhou, Y.; Situ, H.; Li, Y.; and Li, L. 2025. Self-supervised representation learning for Bayesian quantum architecture search. *Phys. Rev. A*, 111: 032403.
- He, Z.; Deng, M.; Zheng, S.; Li, L.; and Situ, H. 2024. Training-free quantum architecture search. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 12430–12438.
- Huang, Y.; Li, Q.; Hou, X.; Wu, R.; Yung, M.-H.; Bayat, A.; and Wang, X. 2022. Robust resource-efficient quantum variational ansatz through an evolutionary algorithm. *Physical Review A*, 105(5): 052414.
- Kandala, A.; Mezzacapo, A.; Temme, K.; Takita, M.; Brink, M.; Chow, J. M.; and Gambetta, J. M. 2017. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *nature*, 549(7671): 242–246.
- Kingma, D. P. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kitaev, A. Y. 1995. Quantum measurements and the Abelian stabilizer problem. *arXiv preprint quant-ph/9511026*.
- Kundu, A. 2025. Improving thermal state preparation of Sachdev-Ye-Kitaev model with reinforcement learning on quantum hardware. *Machine Learning: Science and Technology*.
- Kundu, A.; Bedelek, P.; Ostaszewski, M.; Danaci, O.; Patel, Y. J.; Dunjko, V.; and Miszczak, J. A. 2024. Enhancing variational quantum state diagonalization using reinforcement learning techniques. *New Journal of Physics*, 26(1): 013034.
- Kundu, A.; and Mangini, S. 2025. TensorRL-QAS: Reinforcement learning with tensor networks for scalable quantum architecture search. *arXiv preprint arXiv:2505.09371*.
- Kundu, A.; and Miszczak, J. A. 2022. Variational certification of quantum devices. *Quantum Science and Technology*, 7(4): 045017.
- Kundu, A.; Sarkar, A.; and Sadhu, A. 2024. Kanqas: Kolmogorov-arnold network for quantum architecture search. *EPJ Quantum Technology*, 11(1): 76.
- Kundu, A.; and Sarra, L. 2024. Reinforcement learning with learned gadgets to tackle hard quantum problems on real hardware. *arXiv preprint arXiv:2411.00230*.
- Kuo, E.-J.; Fang, Y.-L. L.; and Chen, S. Y.-C. 2021. Quantum architecture search via deep reinforcement learning. *arXiv preprint arXiv:2104.07715*.
- Larocca, M.; Thanasilp, S.; Wang, S.; Sharma, K.; Biamente, J.; Coles, P. J.; Cincio, L.; McClean, J. R.; Holmes, Z.; and Cerezo, M. 2025. Barren plateaus in variational quantum computing. *Nature Reviews Physics*, 7(4): 174–189.
- LaRose, R.; Tikku, A.; O’Neel-Judy, É.; Cincio, L.; and Coles, P. J. 2019. Variational quantum state diagonalization. *npj Quantum Information*, 5(1): 57.
- Liu, H.; Simonyan, K.; and Yang, Y. 2018. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*.
- Lloyd, S.; Mohseni, M.; and Rebentrost, P. 2014. Quantum principal component analysis. *Nature physics*, 10(9): 631–633.
- Lockwood, O. 2021. Optimizing quantum variational circuits with deep reinforcement learning. *arXiv preprint arXiv:2109.03188*.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 1928–1937. PmLR.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.

- Monz, T.; Nigg, D.; Martinez, E. A.; Brandl, M. F.; Schindler, P.; Rines, R.; Wang, S. X.; Chuang, I. L.; and Blatt, R. 2016. Realization of a scalable Shor algorithm. *Science*, 351(6277): 1068–1070.
- Nakaji, K.; Kristensen, L. B.; Campos-Gonzalez-Angulo, J. A.; Vakili, M. G.; Huang, H.; Bagherimehrab, M.; Gorgulla, C.; Wong, F.; McCaskey, A.; Kim, J.-S.; et al. 2024. The generative quantum eigensolver (GQE) and its application for ground state search. *arXiv preprint arXiv:2401.09253*.
- Nautrup, H. P.; Delfosse, N.; Dunjko, V.; Briegel, H. J.; and Friis, N. 2019. Optimizing quantum error correction codes with reinforcement learning. *Quantum*, 3: 215.
- Olle, J.; Yevtushenko, O. M.; and Marquardt, F. 2025. Scaling the Automated Discovery of Quantum Circuits via Reinforcement Learning with Gadgets. *arXiv preprint arXiv:2503.11638*.
- Olle, J.; Zen, R.; Puviani, M.; and Marquardt, F. 2024. Simultaneous discovery of quantum error correction codes and encoders with a noise-aware reinforcement learning agent. *npj Quantum Information*, 10(1): 126.
- Ostaszewski, M.; Trenkwalder, L. M.; Masarczyk, W.; Scerri, E.; and Dunjko, V. 2021. Reinforcement learning for optimization of variational quantum circuit architectures. *Advances in Neural Information Processing Systems*, 34: 18182–18194.
- Patel, Y. J.; Jerbi, S.; Bäck, T.; and Dunjko, V. 2024a. Reinforcement learning assisted recursive QAOA. *EPJ Quantum Technology*, 11(1): 6.
- Patel, Y. J.; Kundu, A.; Ostaszewski, M.; Bonet-Monroig, X.; Dunjko, V.; and Danaci, O. 2024b. Curriculum reinforcement learning for quantum architecture search under hardware errors. *arXiv preprint arXiv:2402.03500*.
- Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.-H.; Zhou, X.-Q.; Love, P. J.; Aspuru-Guzik, A.; and O’Brien, J. L. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, 5(1): 4213.
- Preskill, J. 2018. Quantum computing in the NISQ era and beyond. *Quantum*, 2: 79.
- Ramôa, M.; Anastasiou, P. G.; Santos, L. P.; Mayhall, N. J.; Barnes, E.; and Economou, S. E. 2024. Reducing the Resources Required by ADAPT-VQE Using Coupled Exchange Operators and Improved Subroutines. *arXiv preprint arXiv:2407.08696*.
- Rattew, A. G.; Hu, S.; Pistoia, M.; Chen, R.; and Wood, S. 2019. A domain-agnostic, noise-resistant, hardware-efficient evolutionary variational quantum eigensolver. *arXiv preprint arXiv:1910.09694*.
- Ruiz, F. J.; Laakkonen, T.; Bausch, J.; Balog, M.; Barekatin, M.; Heras, F. J.; Novikov, A.; Fitzpatrick, N.; Romera-Paredes, B.; van de Wetering, J.; et al. 2025. Quantum circuit optimization with alphasensor. *Nature Machine Intelligence*, 1–12.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Su, J.; Fan, J.; Wu, S.; Li, G.; Qin, S.; and Gao, F. 2025. Topology-driven quantum architecture search framework. *Science China Information Sciences*, 68(8): 180507.
- Sutton, R. S.; Barto, A. G.; et al. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.
- Tang, W.; Duan, Y.; Kharkov, Y.; Fakoore, R.; Kessler, E.; and Shi, Y. 2024. AlphaRouter: Quantum Circuit Routing with Reinforcement Learning and Tree Search. *arXiv preprint arXiv:2410.05115*.
- Tilly, J.; Chen, H.; Cao, S.; Picozzi, D.; Setia, K.; Li, Y.; Grant, E.; Wossnig, L.; Rungger, I.; Booth, G. H.; et al. 2022. The variational quantum eigensolver: a review of methods and best practices. *Physics Reports*, 986: 1–128.
- Van Hasselt, D. S., Arthur Guez. 2016. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.
- Wang, P.; Usman, M.; Parampalli, U.; Hollenberg, L. C.; and Myers, C. R. 2023. Automated quantum circuit design with nested monte carlo tree search. *IEEE Transactions on Quantum Engineering*, 4: 1–20.
- Wang, Y.; He, H.; and Tan, X. 2020. Truly proximal policy optimization. In *Uncertainty in artificial intelligence*, 113–122. PMLR.
- Wang, Z.; Schaul, T.; Hessel, M.; Hasselt, H.; Lanctot, M.; and Freitas, N. 2016. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, 1995–2003. PMLR.
- Williams, C. P.; and Gray, A. G. 1998. Automated design of quantum circuits. In *NASA International Conference on Quantum Computing and Quantum Communications*, 113–125. Springer.
- Wolpert, D. H.; and Macready, W. G. 1997. No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, 1(1): 67–82.
- Ye, E.; and Chen, S. Y.-C. 2021. Quantum architecture search via continual reinforcement learning. *arXiv preprint arXiv:2112.05779*.
- Yin, H.-X.; Hu, Z.-Y.; Zeng, H.-H.; Guan, J.-B.; and Wang, J.-k. 2025. Application of quantum machine learning using variational quantum classifier in accelerator physics. *arXiv preprint arXiv:2506.06662*.
- Zhang, S.-X.; Hsieh, C.-Y.; Zhang, S.; and Yao, H. 2022. Differentiable quantum architecture search. *Quantum Science and Technology*, 7(4): 045023.
- Zhu, C.; Wu, X.; Zhang, H.-K.; Wu, S.; Li, G.; and Wang, X. 2025. Scalable Quantum Architecture Search via Landscape Analysis. *arXiv preprint arXiv:2505.05380*.
- Zhu, X.; and Hou, X. 2023. Quantum architecture search via truly proximal policy optimization. *Scientific Reports*, 13(1): 5157.