Li et al. *EPJ Quantum Technology* (2025) 12:106
https://doi.org/10.1140/epjqt/s40507-025-00410-w

EPJ Quantum Technology
a SpringerOpen Journal

EPJ.org

RESEARCH Open Access

Quantum architecture search with neural predictor based on ZX-calculus

Shanchuan Li[1], Daisuke Tsukayama[1], Jun-ichi Shirakashi[1*], Tetsuo Shibuya[2] and Hiroshi Imai[3]

*Correspondence:
shrakash@cc.tuat.ac.jp
[1]Department of Electrical Engineering and Computer Science, Tokyo University of Agriculture & Technology, Koganei, Tokyo, 184-8588, Japan
Full list of author information is available at the end of the article

Abstract

With the ongoing advances in noisy intermediate-scale quantum hardware, variational quantum algorithms have demonstrated significant potential in a range of quantum applications. However, obtaining high-performance, shallow-parameterized quantum circuits typically requires repeated optimization of the gate parameters over a large set of candidate circuits, resulting in prohibitively high evaluation costs. To address this challenge, this study proposes a novel predictor-based quantum architecture search (PQAS-ZX) method that leverages ZX-calculus. In this approach, a quantum circuit is first represented as a ZX diagram that supports multi-step equivalent simplifications at the diagram level. By applying these equivalence transformations, multiple circuit variants that share the same performance metric are generated, thereby significantly expanding the training dataset and enhancing the ability of the predictor to manage diverse circuit structures. ZX diagrams offer more flexible characterizations of multi-qubit entanglement and phase interactions, as well as higher-level equivalent transformations, compared with the state-of-the-art predictor-based quantum architecture search with graph measures (PQAS-GM). Numerical simulations of three variational quantum eigensolver tasks, namely the transverse-field Ising, Heisenberg, and BeH_2 molecular models, demonstrated that PQAS-ZX required only approximately 80.9%, 82.9%, and 76.1% of the queries required by PQAS-GM, respectively, to achieve the same probability of reaching the target ground-state energy. These results highlight the advantage of using ZX diagrams to identify high-quality circuits efficiently and alleviate the evaluation burden of quantum architecture searches.

Keywords: Quantum Architecture Search; Variational Quantum Algorithm; ZX-Calculus

1 Introduction

With the continuous improvements in qubit numbers and gate fidelity, quantum computing has evolved from theoretical studies towards practical applications [1, 2]. Nevertheless, current experimental platforms remain in the noisy intermediate-scale quantum (NISQ) era, constrained by limited qubit numbers, restricted connectivity, and hardware noise. These factors hinder reliable execution of deep circuits and large-scale quantum algorithms [3].

Hybrid quantum-classical frameworks, particularly variational quantum algorithms (VQAs), have emerged to leverage NISQ hardware effectively. These algorithms iteratively

Springer

employ shallow parameterized quantum circuits (PQCs) and classical optimizers, mitigating noise accumulation while maintaining high performance even at limited circuit depths [4, 5]. VQAs have shown promising results in quantum chemistry [6, 7], combinatorial optimization [8, 9], and quantum machine learning [10, 11].

However, the effectiveness of VQAs is strongly dependent on the PQC design. Simply increasing the circuit depth to enhance expressivity exacerbates hardware noise, intensifying gate errors and triggering the "barren plateau" phenomenon, in which gradients vanish rapidly in high-dimensional parameter spaces [12, 13]. Although heuristic circuit templates such as hardware-efficient circuits and UCCSD partially alleviate these issues [3, 14], effectively balancing trainability and expressivity under realistic hardware conditions remains challenging [15–17]. Thus, the selection of appropriate circuit architectures is critical.

Quantum architecture search (QAS), inspired by neural architecture searches in machine learning [18–20], has emerged as an automated optimization approach. QAS typically comprises a search module that iteratively generates or modifies candidate architectures, and an evaluation module that measures or simulates their performance. Common search techniques include reinforcement learning [21, 22], evolutionary algorithms [23, 24], greedy strategies [25], and differentiable searches [26]. However, evaluating numerous candidate circuits often demands repeated training, leading to significant computational and experimental overheads, thereby limiting the QAS scalability.

Zhang et al. [27] introduced a predictor-based QAS (PQAS) framework to reduce evaluation costs. Their method unfolds quantum circuits into linear sequences, training neural networks to estimate circuit performance quickly and filter out suboptimal candidates efficiently before expensive evaluations. He et al. [28] proposed the use of directed acyclic graphs (DAGs) combined with graph neural networks to capture multi-qubit entanglement and inter-gate dependencies more effectively beyond linear representations. Nevertheless, neither linear lists nor DAGs sufficiently represent complex entanglement structures and intricate phase interactions. In particular, DAG topologies may increase exponentially with the circuit size, complicating advanced optimization and equivalence checking. Consequently, high-level reasoning and effective simplification of quantum circuit representations remain significant challenges.

This study introduces ZX-calculus [29–32] in QAS for the first time to overcome the limitations of traditional representation methods in capturing multi-qubit entanglement and complex phase interactions. ZX-calculus is an intuitive and flexible graphical language for quantum computation, which uses graphical rewriting rules to transform and optimize quantum circuits at a high-level topology effectively. ZX-calculus has demonstrated substantial advantages in quantum circuit optimization, entangled-state analysis, and measurement-circuit design by clearly representing multi-qubit entanglement and intricate phase relationships.

A ZX diagram consists primarily of two fundamental nodes, namely Z-spiders and X-spiders, which represent operations or measurements in the Z-basis and X-basis, respectively. These spiders are interconnected via Hadamard edges, creating a flexible and intuitive graph structure. ZX diagrams preserve the circuit functionality under structural transformations (such as spider fusion and color changing), making them more expressive than traditional linear-list and DAG representations. This unified graphical framework enables a circuit and its multiple structurally equivalent variants to serve as a rich,

consistent dataset, thereby significantly enhancing the training efficiency of the predictor, especially when labeled data are limited. The explicit representation of entanglement and phase interactions in ZX diagrams naturally improves the prediction accuracy and model generalization.

This study proposes a novel predictor-based quantum architecture search method based on ZX-calculus, known as PQAS-ZX. PQAS-ZX leverages the flexible topological transformations of ZX diagrams for effective data augmentation, significantly enhancing the predictive model generalization. Experimental results on three representative variational quantum eigensolver (VQE) tasks—the transverse-field Ising model (TFIM), Heisenberg model, and BeH_2 molecular model—demonstrate that PQAS-ZX substantially reduces the number of required circuit evaluations compared with the state-of-the-art graph-measure-based method (PQAS-GM) [33], requiring only approximately 80.9%, 82.9%, and 76.1% of the evaluations, respectively, to achieve the target ground-state energy. These results highlight the distinct advantage of ZX-calculus in capturing critical circuit features and significantly reducing the computational overhead in QAS. With the continued advances in quantum hardware capabilities, the proposed PQAS-ZX framework provides essential theoretical support and practical guidance towards scalable, cost-effective, and highly efficient quantum circuit design.

The remainder of this paper is organized as follows. Section 2 reviews related work on QAS and ZX-calculus. Section 3 introduces the PQAS-ZX framework and its integration into QAS. Section 4 presents the experimental validations on representative VQE tasks. Finally, Sect. 5 concludes the study and outlines future research directions.

2 Related work

The aim of QAS is to automate the design of high-performance quantum circuits for VQAs, improving algorithmic efficiency under noisy hardware conditions [27, 34]. A typical QAS framework comprises a search module, which iteratively generates or modifies circuit architectures from a large candidate space, and an evaluation module that assesses their performance through quantum simulations or hardware measurements. Previous studies have demonstrated QAS to be superior to manually designed circuits in terms of flexibility and performance across various VQA tasks. However, evaluating numerous candidate circuits typically requires repeated optimization of the gate parameters, resulting in substantial computational and experimental overheads. Thus, reducing the evaluation costs while maintaining the search efficacy is a central focus in QAS research.

Various acceleration strategies have been proposed to address the high costs of large-scale evaluations. Representative methods include the following: (1) One-shot supercircuits train a "supercircuit" encompassing all candidate gate configurations. Each subcircuit inherits optimized parameters from the supercircuit, thereby reducing repetitive training [23]. Although this approach reduces the training overhead, optimizing the supercircuit itself is highly complex and lacks theoretical guarantees regarding individual subcircuit performance. (2) Predictor-based QAS trains neural networks on small sets of circuit performance data to predict the circuit quality rapidly. Thus, low-performing candidates can be filtered out, allowing intensive evaluation only for the most promising circuits [27]. Although this method efficiently reduces the number of evaluations, it demands high accuracy and generalization capabilities from predictors, as well as effective circuit representations that capture the operational and structural features. (3) Training-free fast

evaluation methods include universal or heuristic metrics (e.g., circuit path lengths and expressibility measures) that quickly estimate the circuit quality without full parameter optimization [35]. Despite reducing the overhead further, these methods often lack task-specific precision and require careful fine-tuning or custom metric design, limiting their generalizability across different quantum tasks.

The circuit representation critically influences the performance prediction accuracy in predictor-based QAS frameworks. Previous studies have explored various representations, including linear sequences, DAGs, and general computational graphs [27, 28, 33]. Although linear-list representations are realistic and simple, they often fail to represent multi-qubit entanglement structures and complex phase relationships adequately. DAG-based methods effectively encode inter-gate dependencies, but their complexity increases exponentially with the circuit scale and connectivity, complicating high-level circuit simplifications and equivalence checks. Some methods, such as PQAS-GM [33], employ computational graphs enriched with node-level features like node degree, distance, or clustering coefficients to capture the internal correlations more effectively and reduce the reliance on large labeled datasets. Although this partially alleviates the data requirements, fine-grained gate-level details and specific operational characteristics are still insufficiently captured.

In summary, despite significant progress in QAS methodologies, accelerating large-scale evaluations and devising expressive circuit representations for effective predictor training remain major challenges. Traditional methods such as linear lists and DAGs increasingly exhibit limitations in capturing complex entanglement structures and phase interactions with greater circuit complexity. Therefore, recent research has explored alternative circuit representations that facilitate high-level simplification and equivalence transformations. ZX-calculus, which is a graphical language based on intuitive and flexible graph representations, has emerged as a promising approach. For instance, Ewen et al. [36] recently combined ZX-calculus with genetic programming to optimize quantum circuits through topological mutations such as local complementation and pivoting. However, their work primarily focused on evolving ZX diagrams via topological operations, without explicitly leveraging ZX transformations for systematic data augmentation to enhance predictor generalization. To address these limitations, this study proposes PQAS-ZX, which systematically leverages the high-level graphical rewriting and flexible equivalence transformations of ZX diagrams to generate structurally diverse yet logically equivalent circuits for effective data augmentation. This approach explicitly captures essential circuit features, improving the prediction accuracy and generalization performance while significantly reducing the computational overhead associated with performance evaluations. PQAS-ZX offers significant advantages in quantum circuit optimization by exploiting the rich structural representations provided by ZX diagrams, establishing an essential theoretical foundation and practical roadmap for scalable and efficient QAS.

3 Method

In the previous section, we discussed challenges that are faced by current QAS approaches and introduced ZX-calculus as an effective alternative to conventional gate-level and DAG-based representations. In this section, we present the proposed PQAS-ZX method in detail, highlighting its core principles, key workflow, and main innovations. Specifically, we discuss how multi-step equivalence transformations of quantum circuits in their

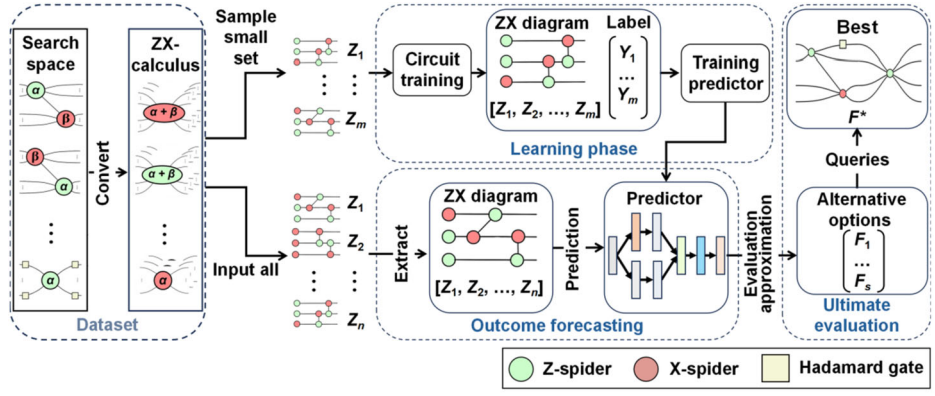

Figure 1 Schematic overview of the PQAS-ZX workflow. (1) Candidate quantum circuits are converted into ZX diagrams, followed by iterative simplification to yield multiple equivalent circuit variants. (2) A small subset of circuits undergoes full VQA training to obtain accurate performance labels. (3) The remaining candidates are rapidly evaluated and screened using the trained predictor. (4) Parameter optimization is performed on the screened subset to determine the true performance and select the optimal quantum circuit architecture

ZX representation can generate multiple functionally equivalent variants, thereby significantly enriching the training dataset to enhance the accuracy and generalization of predictor models.

3.1 Overview of PQAS-ZX

As shown in **Fig. 1,** our method comprises the following four main stages:

1. **ZX Diagram Generation and Multi-step Simplification**

Each candidate quantum circuit is first converted into a ZX diagram. Graphical rewriting rules such as spider fusion, the identity rule, phase-free simplification, and scalar reduction are iteratively applied to yield multiple functionally equivalent yet structurally distinct circuit representations.

2. **Predictor Training**

A small subset of candidate circuits is selected from the pool for comprehensive VQA training to determine accurate performance labels. The ZX-derived variants of each original circuit inherit the same ground-truth performance, thereby significantly enriching the training dataset without requiring additional expensive evaluations. These circuit-performance pairs are subsequently used to train a neural predictor model that can efficiently estimate the performance of arbitrary ZX diagram representations.

3. **Candidate Screening**

The trained predictor is employed for the remaining large set of candidate circuits (and their corresponding ZX diagram variants) to estimate the circuit performance rapidly. Only candidates that are predicted to have high potential undergo comprehensive parameter optimization, substantially reducing the computational overhead associated with circuit evaluations.

4. **Final Selection**

Finally, complete VQA training is conducted on the selected high-potential subset to determine their true performance accurately, as a result of which the best-performing quantum circuit architecture is identified.

Compared with the classical PQAS pipeline, PQAS-ZX systematically exploits ZX-calculus to generate multiple topologically distinct yet logically equivalent circuit repre-

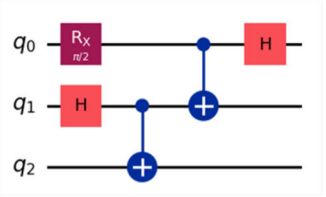

(a) Gate-level circuit

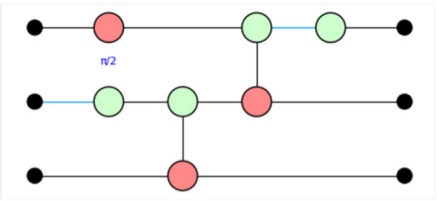

(b) ZX diagram

	Type	Phase	Qubit	Boundary
0	0	0.00	0	0
1	0	0.00	1	0
2	0	0.00	2	0
3	1	0.50	0	1
4	2	0.00	1	1
5	1	0.00	2	2
6	2	0.00	1	2
7	1	0.00	1	2
8	2	0.00	0	2
9	2	0.00	0	1
10	0	0.00	0	0
11	0	0.00	1	0
12	0	0.00	2	0

(c) Node features matrix

	Node_u	Node_v	Phase-coupling	Weight
0	0	1	0.50	1
1	0	2	0.00	2
2	0	1	0.00	1
3	1	2	0.50	1
4	2	2	0.00	1
5	1	2	0.00	1
6	1	0	0.00	1
7	2	1	0.00	1
8	1	2	0.00	1
9	1	0	0.00	1
10	2	2	1.00	2
11	2	0	0.00	1

(d) Edge features matrix

Figure 2 Construction of node and edge feature matrices from a ZX diagram. (a) Gate-level representation of a candidate quantum circuit. (b) Corresponding ZX diagram composed of Z-spider and X-spider nodes interconnected by Hadamard edges. (c) Node feature matrix constructed from the attributes of spider nodes, capturing the node type, phase values, and local edge attributes. (d) Edge feature matrix representing the connectivity patterns, topological relationships, and phase interactions among spider nodes. These matrices constitute structured numerical inputs for the neural predictor

sentations from a single original circuit. PQAS-ZX significantly improves predictor accuracy and generalization, particularly in scenarios with limited labeled data, by effectively enriching the training data and explicitly capturing the entanglement and phase relationships. Moreover, the proportion of circuits that are retained during the candidate screening stage can be flexibly adjusted to balance the evaluation costs and prediction accuracy.

3.2 Feature extraction from ZX diagrams

Each candidate quantum circuit is first converted into structured numerical representations that are compatible with the prediction model to integrate ZX diagram representations into the neural predictor effectively. As illustrated in Fig. 2, the initial step involves transforming quantum circuits into their corresponding ZX diagrams, which consist of interconnected Z-spiders and X-spiders. ZX diagrams explicitly encode essential multi-qubit entanglement structures and intricate phase interactions within a compact graphical form, offering intuitive visualization and structurally rich information that are crucial for accurate circuit performance prediction.

Subsequently, structured numerical information is extracted from these ZX diagrams by systematically constructing two numerical feature matrices—the node and edge feature matrices—as depicted in Fig. 2(c) and (d), respectively. Specifically, the node feature matrix encodes key attributes of each spider node, including the node type (Z- or X-spider), associated phase values, and local connectivity details such as the count and types of directly connected edges (e.g., Hadamard or standard edges). The edge feature matrix explicitly represents inter-node relationships, capturing the detailed connection patterns, phase interactions, and structural dependencies that are intrinsic to ZX diagrams. Both

the local operational characteristics and global topological features of quantum circuits are comprehensively encapsulated by jointly leveraging these node and edge matrices.

The resulting numerical feature matrices serve as direct inputs to the neural predictor, the architecture and training methodology of which are detailed in Sect. 3.4.

3.3 Equivalent circuit generation and data augmentation

We propose a systematic data augmentation strategy based on ZX-calculus simplifications to enrich the training dataset and enhance the generalization capability of our neural predictor. The initial quantum circuit is first converted into its equivalent ZX diagram. Starting with this diagram, we iteratively apply four primary ZX-calculus simplification rules: spider fusion, identity removal, phase-free simplification, and scalar reduction. These transformations significantly alter the topology and structure of the ZX diagram while preserving its logical functionality. We explain each simplification rule below, with illustrative examples depicted in Fig. 3.

Spider fusion (Fig. 3(a)): Adjacent spiders of the same color (either Z- or X-type) are merged into a single spider by combining their phases, thereby reducing the node count and simplifying the diagram structure.

Identity removal (Fig. 3(b)): A spider node with exactly two edges (one input and one output) and no phase represents an identity operation, and can thus be removed directly, simplifying the diagram without altering the logical functionality.

Phase-free simplification (Fig. 3(c)): This rule systematically redistributes internal node phases towards the boundaries, minimizing or eliminating phases in central nodes, thereby yielding a simplified canonical diagram form that facilitates further optimization.

Scalar reduction (Fig. 3(d)): This rule eliminates global scalar factors (e.g., redundant loops), simplifying the overall diagram. Although these scalar components do not affect the measurement outcomes, their removal streamlines the circuit representation.

Multiple logically equivalent yet structurally diverse circuit variants are produced from a single initial circuit by iteratively applying these ZX-calculus rules in various combinations and orders. As illustrated in Fig. 3, each simplification step generates distinct ZX diagram representations exhibiting significant topological and structural variations. This "one-circuit, multi-representation" strategy significantly enhances the dataset diversity for the subsequent neural predictor training.

The ZX-calculus-based multi-equivalent representation strategy provides several notable advantages for data augmentation compared with conventional PQAS pipelines. First, it supports multi-step simplifications and unified performance labeling. Once an original circuit (denoted as Z_0) is comprehensively trained and evaluated (e.g., by obtaining the ground-state energy via VQA), the same ground-truth performance label is consistently assigned to all logically equivalent ZX variants $\{Z_0^0, Z_0^1, \ldots, Z_0^k\}$. This approach significantly increases the dataset diversity without additional costly evaluations. Second, it provides rich structural variations. Equivalent circuits exhibit substantial diversity in node types, phase configurations, and subgraph structures, greatly enhancing the generalization capabilities of the predictor. Third, label scarcity is mitigated. Generating labeled data via quantum evaluations and parameter optimizations is resource intensive. Our multi-representation, single-performance strategy effectively addresses the label scarcity issues that are commonly encountered in large-scale QAS tasks, maximizing information utilization from each evaluation, and thus, improving the predictor robustness.

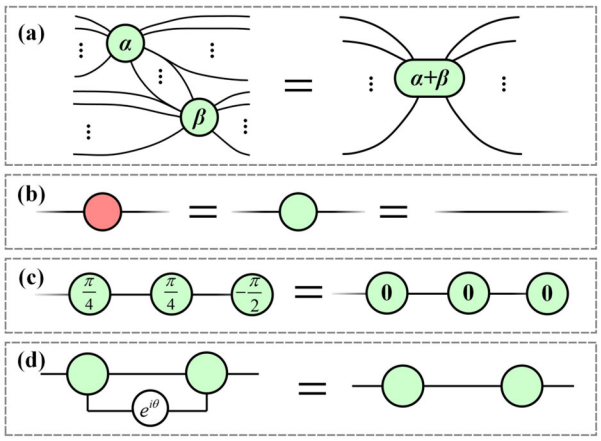

Figure 3 Illustrations of the primary ZX-calculus simplification rules used in PQAS-ZX for data augmentation. (a) Spider fusion merges adjacent spiders of the same color by combining their phases. (b) Identity removal eliminates identity operations represented by phase-free spiders with exactly two edges. (c) Phase-free simplification redistributes the internal node phases towards the diagram boundaries, simplifying the diagram structure. (d) Scalar reduction removes global scalar components. These rules hold equally for spiders of the opposite color despite being explicitly illustrated using spiders of one color. Each step preserves logical functionality, producing structurally diverse yet functionally equivalent ZX diagram variants

In summary, the proposed ZX-calculus-based data augmentation strategy significantly enriches the training dataset, enhancing the predictor accuracy and generalization. The number and sequence of applied simplification rules can be flexibly adjusted according to the circuit complexity and computational resources, providing a practical balance between the simplification depth and overheads. The enriched and structurally diverse dataset obtained in this manner serves directly as input for the neural predictor described in Sect. 3.4.

3.4 Neural predictor construction

A neural predictor is constructed based on the structured numerical representations described in Sects. 3.2 and 3.3 to estimate quantum circuit performance accurately and efficiently. As illustrated in Fig. 4, our predictor employs a dual-branch neural architecture that independently processes the node and edge feature matrices extracted from the ZX diagrams.

Specifically, the node and edge feature matrices serve as inputs to two parallel processing branches. Each branch comprises a long short-term memory (LSTM) layer followed by convolutional (Conv) layers, effectively capturing both the temporal dependencies and local structural patterns within each type of feature. The resulting feature embeddings from both branches at each of the five timesteps (t_0 to t_4) are temporally aggregated and concatenated into a unified representation. This integrated representation is further refined through a fully connected (FC) layer, ultimately producing a scalar output to predict performance metrics of the quantum circuit, such as the ground-state energy.

During training, structurally diverse yet logically equivalent circuit variants that are derived from the same original quantum circuit share identical ground-truth performance labels. The predictor is trained by minimizing the mean squared error between the predicted values and actual circuit performance metrics. Leveraging diverse ZX diagram rep-

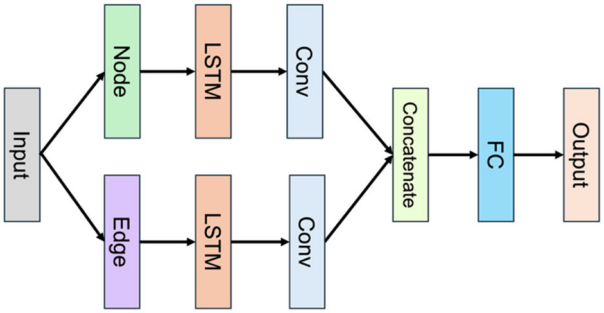

Figure 4 Schematic of the neural predictor architecture used in PQAS-ZX. The node and edge feature matrices extracted from the ZX diagrams are separately processed by two parallel branches, each consisting of an LSTM layer followed by Conv layers. Subsequently, the outputs of these branches are concatenated and refined through an FC layer, resulting in a final scalar prediction of the quantum circuit performance

resentations enriches the training dataset, enhancing the generalization capabilities of the predictor across unseen circuit architectures.

4 Results and discussion

To evaluate the effectiveness of the proposed PQAS-ZX method, numerical simulations were performed on three representative VQE tasks using the Qiskit platform: (1) six-qubit TFIM, (2) five-qubit Heisenberg model, and (3) ground-state energy estimation of a six-qubit BeH_2 molecule. We adopted task definitions, parameter settings, and search-space configurations identical to those used in Ref. [33] to ensure a direct comparison with the state-of-the-art approach PQAS-GM [33]. The core Hamiltonians and success criteria are summarized in the following, and further details can be found in Ref. [33].

4.1 Tasks

TFIM:

The target Hamiltonian is

$$H_{TFIM} = \sum_{i=1}^{n} Z_i Z_{i+1} + \sum_{i=1}^{n} X_i, \quad (n = 6), \tag{1}$$

where the energy $E \leq -7.72$ is assumed to have reached the ground state.

Heisenberg:

The target Hamiltonian is

$$H_{Heis} = \sum_{i=1}^{n} (X_i X_{i+1} + Y_i Y_{i+1} + Z_i Z_{i+1}) + \sum_{i=1}^{n} Z_i, \quad (n = 5) \tag{2}$$

under periodic boundary conditions. The estimated energy $E \leq -8.47$ is assumed to be the ground state.

BeH_2 Molecular Ground State:

We used the Hamiltonian coefficients and Pauli operators provided in Ref. [3]. This result is considered to achieve chemical accuracy if it satisfies the condition

$$|E - E_0| \leq 1.6 \times 10^{-3} Hartree\ (Ha).$$

We performed 10 independent simulations for each task to reduce variability arising from random parameter initialization. We recorded the final converged energies and number of optimization iterations to evaluate the convergence efficiency and accuracy.

4.2 Search space definition and experimental setup

We defined the QAS search space according to Ref. [27], which includes three single-qubit rotation gates $\{Rx(\theta), Ry(\theta), Rz(\theta)\}$ and three two-qubit rotation gates $\{\mathrm{XX}(\theta), \mathrm{YY}(\theta), \mathrm{ZZ}(\theta)\}$. For instance, $Rx(\theta) = \mathrm{e}^{-i\theta X/2}$ and $XX(\theta) = \mathrm{e}^{-i\theta XiXj/2}$. We applied an initial layer of Hadamard gates to all qubits in the six-qubit TFIM task. This method is consistent with the approach used in Ref. [27]. As single-qubit gates generally exhibit high robustness to noisy hardware, we prioritized their placement to minimize noise accumulation while ensuring sufficient multi-qubit entanglement.

We adopted the hierarchical sampling strategy introduced in Ref. [27], generating circuits through two pipelines, namely *gatewise* and *layerwise*, to avoid excessively random gate selections and placements that could result in highly irregular layouts or severe barren plateau effects. In the former, each gate slot is assigned either a single-qubit or two-qubit gate, with two-qubit gates restricted to adjacent qubit pairs. In the latter, the circuit depth is organized into alternating "half layers," each comprising n / 2 operations for an n-qubit system, which are applied to alternating pairs of qubits. This approach yields structured and diverse circuit layouts.

Rotation angles θ for each gate were uniformly sampled from the interval $[-2\pi, 2\pi]$ for parameter initialization and training. We employed the *L-BFGS-B* optimizer for VQE training and recorded the final energies upon convergence or reaching a predefined maximum number of iterations. Unless otherwise specified, all other experimental settings followed those described in Ref. [27]. This setup balanced hardware constraints, noise considerations, and exploration diversity, enabling effective performance evaluation for each candidate circuit.

4.3 Queries required to achieve ground state

In this experiment, we sampled 40,000 circuits from a predefined search space and used a trained predictor to estimate their performance. Subsequently, we ranked all circuits in terms of their predicted performance (from best to worst) and sequentially queried the true energy of each circuit until one that achieved the ground state was identified. Figure 5 shows the number of queries required to achieve the ground-state energy (curves), along with the total number of queries (bars), including those from both the training and search phases, across varying training set sizes.

As the training-set size increased, the predictor became more accurate in identifying high-performance circuits, thereby reducing the number of queries required. However, beyond a certain point, using additional training data resulted in diminishing returns, indicating that further expansion yielded minimal improvements. Notably, PQAS-ZX consistently required fewer total queries than PQAS-GM [33] across most training set sizes, probably because the ZX-based representation captured excellent structural features that were beneficial to the predictor. The total number of queries exhibited a "decrease-then-increase" pattern as the training set size increased, implying that extremely small or large training sets were suboptimal in terms of efficiency.

PQAS-ZX achieved the optimal query cost with training set sizes of 10, 20, and 60 for the TFIM, Heisenberg, and BeH_2 tasks, respectively, whereas PQAS-GM [33] typically

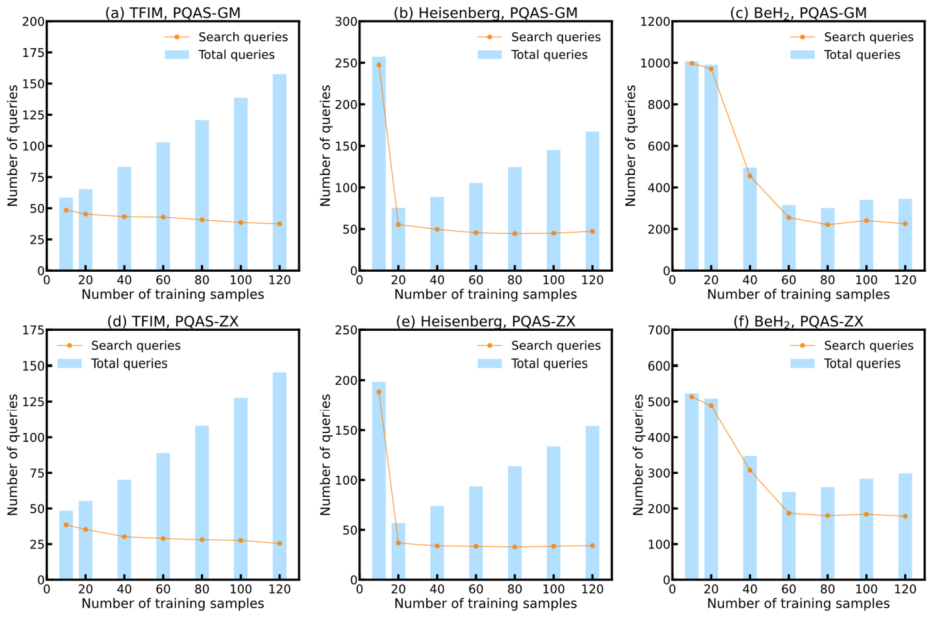

Figure 5 Comparison of the number of circuit evaluations (queries) required to reach the target ground-state energy (orange curves, labeled "Search queries") and total cumulative queries including initial training samples (blue bars, labeled "Total queries"), for PQAS-GM [33] (top row: a–c) and the proposed PQAS-ZX (bottom row: d–f). Results are presented across three representative VQE tasks: (a, d) the TFIM, (b, e) Heisenberg model, and (c, f) BeH_2 molecular model. Each data point represents the mean value calculated over 10 independent runs. The number of queries required to achieve the target ground-state energy generally decreased as the training set size increased. PQAS-ZX consistently achieved lower total query costs compared with PQAS-GM [33], indicating greater efficiency

required 10, 20, and 80 training samples. Unless otherwise noted, these optimal settings were used in subsequent experiments.

Figure 6 compares the minimum number of queries required by PQAS-ZX, PQAS-GM [33], and random search (RS). RS randomly selected circuits as a naive baseline. Compared with PQAS-GM [33], PQAS-ZX required only approximately 80.9%, 82.9%, and 76.1% of the number of queries for the TFIM, Heisenberg, and BeH_2 tasks, respectively, highlighting the advantage of ZX-based representations for efficiently identifying top-performing circuits. PQAS-ZX could improve the sampling efficiency by over 85.6 times on the TFIM task compared with RS.

4.4 Performance distribution and comparative analysis

We sampled 40,000 candidate circuits from the search spaces of the TFIM, Heisenberg, and BeH_2 VQE tasks to assess the effectiveness of our predictor in identifying high-performance quantum circuits further. The trained predictor was used to estimate the energies of these candidates rapidly, following which the circuits were ranked from best to worst. Subsequently, the top 300 circuits identified by the predictor underwent full parameter optimization and precise energy evaluations. We employed RS, which uniformly sampled 300 circuits at random from the same set of 40,000 candidates, as a baseline for the comparison.

Figure 7 presents the energy distributions of the 300 selected circuits for each VQE task (TFIM, Heisenberg, and BeH_2). The orange histograms represent the energy distributions

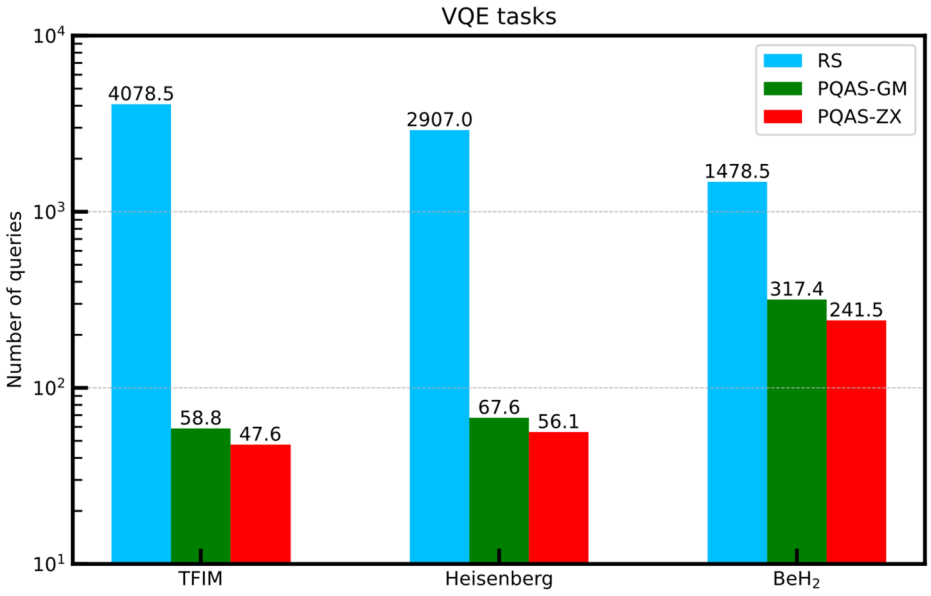

Figure 6 Comparison of the minimum number of circuit evaluations (queries) required by RS (blue), PQAS-GM [33] (green), and the proposed PQAS-ZX (red) to achieve the target ground-state energy across three representative VQE tasks: the TFIM, Heisenberg model, and BeH_2 molecular model. Each result was averaged over 10 independent experimental runs. Lower query counts correspond to higher search efficiency, indicating that PQAS-ZX consistently outperformed both PQAS-GM and RS

of circuits selected via predictor-based screening, whereas the blue histograms depict circuits that were selected randomly (RS). Each panel explicitly highlights a defined "*low-energy region*" (TFIM: $E \leq -7.50$, Heisenberg: $E \leq -8.25$, BeH_2: $E \leq -1.92$), and the exact ground-state energy for each task is indicated by a vertical red dashed line. Circuits selected by the predictor exhibited dense clustering within these low-energy regions, demonstrating the high success rate of the predictor in identifying top-performing candidates. In contrast, circuits selected by RS exhibited broader energy distributions with significantly fewer high-quality solutions within the specified low-energy range.

These findings confirm that our predictor effectively filters out poorly performing circuits from a large candidate pool, efficiently directing computational resources towards circuits with greater potential. Thus, the predictor-based screening mechanism achieves an improved balance between the training cost and search efficiency, providing a scalable pathway towards larger and more complex quantum circuit design problems.

4.5 Comparison under different query budgets

Finally, we constrained the maximum query budget to N_s / 10, where N_s denotes the number of queries required to reach the ground state in an unconstrained scenario. Each experiment was independently repeated 10 times, and the success rate (the probability of reaching the ground state within this constrained budget) was evaluated. Table 1 summarizes and compares the success rates achieved by PQAS-ZX and PQAS-GM [33] in the TFIM, Heisenberg, and BeH_2 VQE tasks. Although the success rate of both methods improved as the query budget increased, PQAS-ZX consistently achieved higher success rates under identical query constraints.

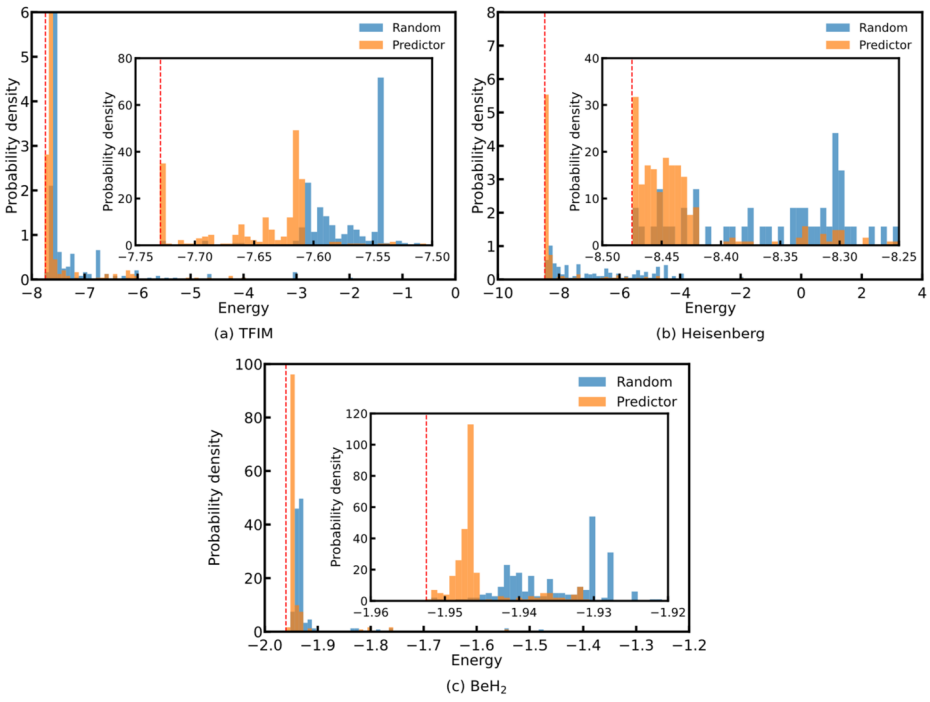

Figure 7 Energy distributions of the 300 quantum circuits selected by the predictor (orange histograms) and RS (blue histograms) across three VQE tasks: (a) TFIM, (b) Heisenberg, and (c) BeH_2. The vertical red dashed line indicates the true ground-state energy for each task. The insets provide zoomed-in views of the low-energy regions (TFIM: $E \leq -7.50$, Heisenberg: $E \leq -8.25$, and BeH_2: $E \leq -1.92$). The histogram bin sizes were set to 100 for the main plots and 50 for the insets, resulting in differing y-axis scales

These results highlight the significant advantage of adopting a ZX-based representation under limited query budgets. Specifically, PQAS-ZX required only approximately 52.9%, 70.0%, and 79.3% of the queries required by PQAS-GM [33] to achieve a 100% success rate for the TFIM, Heisenberg, and BeH_2 tasks, respectively. This improvement primarily arises from the multi-step equivalence transformations enabled by ZX diagrams, which allow for systematic data augmentation by generating multiple distinct yet logically equivalent circuit variants from a single evaluation. Consequently, the predictor benefits from a substantially richer and more diverse training dataset, even under limited query budgets, leading to improved accuracy in screening high-performance circuits. In contrast, PQAS-GM relies on a fixed set of graph-measure features and provides fewer opportunities for such systematic data augmentation, thereby limiting the predictor accuracy and efficiency under similar training conditions. Consequently, PQAS-ZX reliably achieved the desired success rate using significantly fewer queries within the constrained query budget of N_s / 10, highlighting its potential for rapidly identifying high-quality quantum circuits in large-scale search tasks.

5 Conclusions

We have introduced PQAS-ZX, which is a predictor-based QAS method that uses ZX-calculus to capture the rich structural information of quantum circuits and efficiently screen large candidate pools. Our experimental results demonstrated that PQAS-ZX successfully identified high-performance quantum circuits with significantly fewer evaluation

Table 1 Success rates (%) of PQAS-ZX and PQAS-GM [33] under constrained query budgets (N_s / 10) for three VQE tasks: TFIM, Heisenberg, and BeH_2. The number of training circuits, the number of queries per training circuit ("#Queries"), total query budget, and the corresponding success rates (averaged over 10 independent runs) are presented for each task and method. A higher success rate indicates superior efficiency in identifying optimal quantum circuits within limited query budgets

Task	Method	#Training circuits	#Queries	Total budget	Success rate
TFIM	PQAS-GM [33]	10	80	90	90%
		10	160	170	100%
	PQAS-ZX	10	40	50	80%
		10	80	90	100%
Heisenberg	PQAS-GM [33]	20	60	80	90%
		20	80	100	100%
	PQAS-ZX	20	40	60	80%
		20	50	70	100%
BeH_2	PQAS-GM [33]	80	400	480	90%
		80	500	580	100%
	PQAS-ZX	60	300	360	80%
		60	400	460	100%

queries, thereby accelerating convergence to the target ground-state energies. Specifically, PQAS-ZX required only approximately 80.9%, 82.9%, and 76.1% of the queries to achieve equivalent success rates for three representative VQE tasks, namely TFIM, Heisenberg, and BeH_2, respectively, compared with the state-of-the-art graph-measure-based PQAS-GM method [33]. These findings highlight the efficacy of ZX-calculus representations in accurately modeling quantum circuit structures and effectively identifying high-quality circuit candidates. The proposed PQAS-ZX method provides valuable insights for leveraging ZX-calculus in high-level quantum circuit optimization and reasoning, offering robust theoretical and practical support for scalable and efficient QAS.

In future work, we plan to enhance ZX-based representations further by incorporating advanced visualization techniques and optimized data structures to capture subtle circuit correlations more effectively. In addition, we aim to address critical challenges associated with large-scale quantum systems, such as exponential state-space growth and hardware coherence-time limitations. Potential solutions include integrating advanced quantum compilation strategies, developing novel hybrid quantum-classical algorithms, and strategically controlling the circuit depth. Furthermore, we intend to validate the performance of PQAS-ZX experimentally on real quantum hardware and systematically investigate its robustness under realistic noise models. We anticipate that these extensions will further unlock opportunities for efficient, scalable, and practical quantum circuit designs in near-term quantum computing platforms.

Abbreviations

NISQ, noisy intermediate-scale quantum; VQA, variational quantum algorithm; PQC, parameterized quantum circuit; VQE, variational quantum eigensolver; TFIM, transverse-field Ising model; QAS, quantum architecture search; DAG, directed acyclic graph; RS, random search; FC, fully connected; LSTM, long short-term memory.

Acknowledgements

We sincerely thank Mr. Pengyue Ma from Beijing Normal University, Hong Kong Baptist University United International College (UIC), for his helpful contributions in optimizing the prediction model, performing repeated experimental trials, and preparing the manuscript figures.

Author contributions

Authors' contributions Shanchuan Li: Investigation (lead), methodology (lead), experimentation (lead), data curation (lead), writing – original draft (lead). Daisuke Tsukayama: Supervision (equal); writing, review, and editing (equal). Jun-ichi Shirakashi: Funding acquisition (lead), supervision (equal), writing, review, and editing (equal). Tetsuo Shibuya: Supervision (equal), writing, review, and editing (equal). Hiroshi Imai: Supervision (equal), writing, review, and editing (equal).

Funding information
Not applicable.

Data availability
The data supporting the findings of this study are available from the corresponding author upon request.

Declarations

Competing interests
The authors declare no competing interests.

Author details
[1]Department of Electrical Engineering and Computer Science, Tokyo University of Agriculture & Technology, Koganei, Tokyo, 184-8588, Japan. [2]Division of Medical Data Informatics, Human Genome Center, The Institute of Medical Science, The University of Tokyo, Minato, Tokyo, 108-8639, Japan. [3]The Graduate School of Information Science and Technology, The University of Tokyo, Bunkyo, Tokyo, 113-8656, Japan.

Received: 11 March 2025 Accepted: 19 August 2025 Published online: 02 September 2025

References

1. Preskill J. Quantum computing in the NISQ era and beyond. Quantum. 2018;2:79. https://doi.org/10.22331/q-2018-08-06-79.
2. Bharti K, Cervera-Lierta A, Kyaw TH, Haug T, Alperin-Lea S, Anand A, et al. Noisy intermediate-scale quantum algorithms. Rev Mod Phys. 2022;94:015004. https://doi.org/10.1103/RevModPhys.94.015004.
3. Kandala A, Mezzacapo A, Temme K, Takita M, Brink M, Chow JM, et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. Nature. 2017;549:242–6. https://doi.org/10.1038/nature23879.
4. Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, et al. Variational quantum algorithms. Nat Rev Phys. 2021;3:625–44. https://doi.org/10.1038/s42254-021-00348-9.
5. McClean JR, Romero J, Babbush R, Aspuru-Guzik A. The theory of variational hybrid quantum-classical algorithms. New J Phys. 2016;18:023023. https://doi.org/10.1088/1367-2630/18/2/023023.
6. Peruzzo A, McClean J, Shadbolt P, Yung M-H, Zhou X-Q, Love PJ, et al. A variational eigenvalue solver on a photonic quantum processor. Nat Commun. 2014;5:4213. https://doi.org/10.1038/ncomms5213.
7. Higgott O, Wang D, Brierley S. Variational quantum computation of excited states. Quantum. 2019;3:156. https://doi.org/10.22331/q-2019-07-01-156.
8. Farhi E, Goldstone J, Gutmann S. A quantum approximate optimization algorithm. 2014. arXiv preprint. arXiv:1411.4028. https://doi.org/10.48550/arXiv.1411.4028.
9. Zhou L, Wang S-T, Choi S, Pichler H, Lukin MD. Quantum approximate optimization algorithm: performance, mechanism, and implementation on near term devices. Phys Rev X. 2020;10:021067. https://doi.org/10.1103/PhysRevX.10.021067.
10. Mitarai K, Negoro M, Kitagawa M, Fujii K. Quantum circuit learning. Phys Rev A. 2018;98:032309. https://doi.org/10.1103/PhysRevA.98.032309.
11. Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N, Lloyd S. Quantum machine learning. Nature. 2017;549:195–202. https://doi.org/10.1038/nature23474.
12. McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H. Barren plateaus in quantum neural network training landscapes. Nat Commun. 2018;9:4812. https://doi.org/10.1038/s41467-018-07090-4.
13. Cerezo M, Sone A, Volkoff T, Cincio L, Coles PJ. Cost function dependent barren plateaus in shallow parametrized quantum circuits. Nat Commun. 2021;12:1791. https://doi.org/10.1038/s41467-021-21728-w.
14. Romero J, Babbush R, McClean JR, Hempel C, Love PJ, Aspuru-Guzik A. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. Quantum Sci Technol. 2018;4:014008. https://doi.org/10.1088/2058-9565/aad3e4.
15. Sim S, Johnson PD, Aspuru-Guzik A. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. Adv Quantum Technol. 2019;2:1900070. https://doi.org/10.1002/qute.201900070.
16. De Leon NP, Itoh KM, Kim D, Mehta KK, Northup TE, Paik H, et al. Materials challenges and opportunities for quantum computing hardware. Science. 2021;372:eabb2823. https://doi.org/10.1126/science.abb2823.
17. Schuld M, Bergholm V, Gogolin C, Izaac J, Killoran N. Evaluating analytic gradients on quantum hardware. Phys Rev A. 2019;99:032331. https://doi.org/10.1103/PhysRevA.99.032331.
18. Yao Q, Wang M, Chen Y, Dai W, Li Y-F, Tu W-W, et al. Taking human out of learning applications: a survey on automated machine learning. 2018. arXiv preprint. arXiv:1810.13306. https://doi.org/10.48550/arXiv.1810.13306.
19. Elsken T, Metzen JH, Hutter F. Neural architecture search: a survey. J Mach Learn Res. 2019;20:1–21.
20. Ren P, Xiao Y, Chang X, Huang P-Y, Li Z, Chen X, et al. A comprehensive survey of neural architecture search: challenges and solutions. ACM Comput Surv. 2022;54:1–34. https://doi.org/10.1145/3447582.
21. Kuo E-J, Fang Y-LL, Chen SY-C. Quantum architecture search via deep reinforcement learning. 2021. arXiv preprint. arXiv:2104.07715. https://doi.org/10.48550/arXiv.2104.07715.
22. Fösel T, Tighineanu P, Weiss T, Marquardt F. Reinforcement learning with neural networks for quantum feedback. Phys Rev X. 2018;8:031084. https://doi.org/10.1103/PhysRevX.8.031084.
23. Wang H, Ding Y, Gu J, Lin Y, Pan DZ, Chong FT, et al. Quantumnas: noise-adaptive search for robust quantum circuits. In: IEEE international symposium on high-performance computer architecture (HPCA). New York: IEEE Press; 2022. p. 692–708. https://doi.org/10.1109/HPCA53966.2022.00057.

24. Grimsley HR, Economou SE, Barnes E, Mayhall NJ. An adaptive variational algorithm for exact molecular simulations on a quantum computer. Nat Commun. 2019;10:3007. https://doi.org/10.1038/s41467-019-10988-2.
25. Ostaszewski M, Grant E, Benedetti M. Structure optimization for parameterized quantum circuits. Quantum. 2021;5:391. https://doi.org/10.22331/q-2021-01-28-391.
26. Zhang S-X, Hsieh C-Y, Zhang S, Yao H. Differentiable quantum architecture search. Quantum Sci Technol. 2022;7:045023. https://doi.org/10.1088/2058-9565/ac87cd.
27. Zhang S-X, Hsieh C-Y, Zhang S, Yao H. Neural predictor based quantum architecture search. Mach Learn: Sci Technol. 2021;2:045027.
28. He Z, Chen H, Zhou Y, et al. Self-supervised representation learning for Bayesian quantum architecture search. Phys Rev A. 2025;111(3):032403. https://doi.org/10.1103/PhysRevA.111.032403.
29. Coecke B, Duncan R. Interacting quantum observables: categorical algebra and diagrammatics. New J Phys. 2011;13:043016. https://doi.org/10.1088/1367-2630/13/4/043016.
30. Duncan R, Perdrix S. Graph states and the necessity of Euler decomposition. In: Mathematical theory and computational practice. Proceedings 5: 5th conference on computability in Europe, CiE 2009, Heidelberg, Germany, Jul 19–24, 2009. Springer; 2009. p. 167–77.
31. Backens M. The ZX-calculus is complete for stabilizer quantum mechanics. New J Phys. 2014;16:093021. https://doi.org/10.1088/1367-2630/16/9/093021.
32. Wetering J. ZX-calculus for the working quantum computer scientist. 2020. arXiv preprint. arXiv:2012.13966. https://doi.org/10.48550/arXiv.2012.13966.
33. He Z, Li Z, Deng M, Zheng S, Situ H, Li L. Quantum architecture search with neural predictor based on graph measures. Adv Quantum Technol. 2024;7:2400223. https://doi.org/10.1002/qute.202400223.
34. Du Y, Huang T, You S, Hsieh M-H, Tao D. Quantum circuit architecture search for variational quantum algorithms. npj Quantum Inf. 2022;8:62. https://doi.org/10.1038/s41534-022-00570-y.
35. He Z, Deng M, Zheng S, Li L, Situ H. Training-free quantum architecture search. In: AAAI. vol. 38. 2024. p. 12430–8. https://doi.org/10.1609/aaai.v38i11.29135.
36. Ewen T, Turkalj I, Holzer P, Wolf M-O. Application of ZX-calculus to quantum architecture search. Quantum Mach Intell. 2025;7:34. https://doi.org/10.1007/s42484-025-00264-6.

Publisher's note