QAS^{2}QD: LEARNING TO DESIGN CIRCUITS FOR SCALABLE SAMPLE-BASED QUANTUM DIAGONALIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Ground-state energy estimation (GSE) is one of the most promising applications of quantum computers for understanding quantum systems across physics, chemistry, and materials science, yet remains challenging at large scales. Recent advances in sample-based quantum diagonalization (SQD) offer a promising route to this task. However, its scalability critically depends on the sampling circuit used to construct the reduced subspace. To address this bottleneck, here we first reformulate GSE tasks as a family-level learning problem across related Hamiltonians, a setting that naturally arises in many-body physics and quantum chemistry. Building on this formulation, we develop Qas^{2}QD, a learning-guided quantum architecture search framework for automatically designing effective SQD sampling circuits. At its core is a multimodal predictor that jointly learns from circuit structure, Hamiltonian information, and low-shot observations to predict high-shot SQD performance, enabling circuit-performance knowledge to transfer across related Hamiltonians. We systematically evaluate Qas^{2}QD on quantum systems with up to 40 qubits. On TFIM and XXZ benchmarks, the multimodal predictor preserves strong circuit-performance rankings using only 1% of the reference shots per circuit, enabling Qas^{2}QD to reduce the error by up to 91.4% over vanilla SQD. These results demonstrate the potential of Qas^{2}QD for tackling challenging quantum systems and advancing quantum computing toward practical utility.

1 INTRODUCTION

Understanding the ground-state behavior of quantum many-body systems is of great importance for explaining and utilizing a wide range of phenomena in quantum physics (Lanyon et al., 2010; Bauer et al., 2020; Lee et al., 2023; Wang et al., 2024). More precisely, progress in this direction can bring new capabilities for predicting quantum materials, molecular properties, and other strongly correlated phenomena (Bauer et al., 2020; McArdle et al., 2020). Despite the significance, the exponential growth of the Hilbert space makes accurate *ground-state energy estimation* (GSE) rapidly intractable for classical computation as the system size increases (Georgescu et al., 2014; Feynman et al., 2018). By leveraging quantum superposition and entanglement, quantum computers provide a natural route to this challenge by directly accessing and manipulating states in exponentially large Hilbert spaces. This capability has motivated several major paradigms for GSE, including variational quantum eigensolvers, quantum phase estimation, quantum imaginary-time evolution, and adiabatic ground-state preparation Aspuru-Guzik et al. (2005); Kandala et al. (2017); Motta et al. (2020); van Vreumingen & Schoutens (2023). However, these approaches typically either require deep quantum circuits and fault-tolerant hardware or remain practically limited to relatively small-scale problems, precluding the demonstration of their utility on near-term quantum devices.

To better exploit current quantum hardware for GSE, quantum diagonalization (McClean et al., 2017; Epperly et al., 2022) offers an alternative paradigm. The core of quantum diagonalization is to use quantum experiments to construct a tractable subspace that captures the relevant low-energy physics, followed by a classical solution of the reduced eigenvalue problem. Existing approaches to quantum diagonalization can be broadly grouped into two families, depending on how this subspace is constructed. The first class is quantum-state-based methods, e.g., quantum subspace expansion and quantum Krylov approaches (Frame et al., 2018; Parrish & McMahon, 2019; Yeter-Aydeniz et al., 2020; Cohn et al., 2021; Cortes & Gray, 2022; Yoshioka et al., 2025), which construct the

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

subspace from quantum states generated through suitable transformations or time evolution. The second class is sample-based methods that construct it from computational-basis configurations obtained through quantum measurements, as exemplified by *sample-based quantum diagonalization* (SQD) (Yu et al., 2025; Nogaki et al., 2025; Barison et al., 2025; Duriez et al., 2026). Compared with the former, SQD is particularly well suited to current quantum hardware, especially when the ground state is sparse in the computational basis (Shirakawa et al., 2026; Merz Jr et al., 2026). However, its scalability hinges critically on the employed sampling circuit. A recent 77-qubit experiment on an iron–sulfur cluster required approximately 2.46 million measurements and classical subspaces of up to 10^8 configurations (Robledo-Moreno et al., 2025). Such results indicate the importance of sampling circuit design, as an inefficient sampler can waste a limited measurement budget on irrelevant configurations and substantially increase computational costs. To this end, it naturally asks:

Can we automatically discover sampling circuits that make SQD more scalable and effective?

Here we answer this question positively by introducing Qas^{2}QD, a learning-guided quantum architecture search (QAS) framework for scalable SQD. To this end, we first reframe SQD from independently solving individual Hamiltonians into a *family-level learning problem*, a setting that naturally arises in many-body physics and quantum chemistry (Huang et al., 2022; Lewis et al., 2024; Du et al., 2026). More concretely, many GSE tasks involve families of Hamiltonians parameterized by couplings, external fields, or molecular geometries. Building on this formulation, Qas^{2}QD introduces *a multimodal neural predictor* that jointly learns from three complementary sources of information, i.e., the candidate circuit, target Hamiltonian, and low-shot measurement outcomes, to predict otherwise expensive high-shot SQD performance. The learned predictor then guides a joint search over circuit structures and parameters using only low-shot feedback, efficiently identifying sampling circuits that yield informative subspaces. In this way, Qas^{2}QD turns sampling circuit design from an instance-specific and high-cost optimization into a transferable learning problem, opening a new route toward scaling SQD to larger systems and broader Hamiltonian families.

Extensive experiments on quantum many-body systems and chemical systems up to 40 qubits demonstrate that Qas^{2}QD consistently discovers substantially better sampling circuits under limited measurement budgets. Using only 1% of the high-shot evaluation budget, the learned predictor already preserves strong circuit performance rankings and effectively guides architecture search, reducing the estimation error by up to 91.4% over vanilla SQD. More importantly, these gains generalize to unseen Hamiltonians and remain robust across different measurement budgets, system sizes, and retained-subspace dimensions, highlighting the potential of Qas^{2}QD to enable scalable circuit design for SQD and advance quantum computing toward practical utility.

Our contributions are summarized as follows. (i) We introduce the first *family-level learning formulation* for SQD circuit optimization, replacing instance-wise circuit design with transferable learning across related Hamiltonians and opening a new avenue toward scalable sampling circuit design for quantum diagonalization. (ii) We develop Qas^{2}QD, the first learning-guided QAS framework for SQD. The proposed surrogate-learning principle and multi-modal neural predictor are general and can be extended to other quantum diagonalization settings where evaluating candidate subspaces is expensive. (iii) We demonstrate substantial performance gains and scalability of Qas^{2}QD across diverse settings, with experiments scaling to 40 qubits. All source code is available at `https://anonymous.4open.science/r/QAS2QD-03DB`.

2 PRELIMINARIES

In this section, we first introduce the necessary background. After that, we briefly review the most closely related work to Qas^{2}QD. More details are deferred to Appendix A.

Ground state energy estimation (GSE). Denote an n-qubit Hamiltonian as

$$H(\boldsymbol{\alpha}) = \alpha_0 I^{\otimes n} + \sum_m \alpha_m P_m, \tag{1}$$

where $\alpha_m \in \mathbb{R}$, I is the single-qubit identity operator, and each P_m is a non-identity Pauli string in $\{I, X, Y, Z\}^{\otimes n}$. Here, $X = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, $Y = \left(\begin{smallmatrix} 0 & -i \\ i & 0 \end{smallmatrix}\right)$, and $Z = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ are the single-qubit Pauli operators. The ground-state energy of $H(\boldsymbol{\alpha})$ takes the form as

$$E_0(H(\boldsymbol{\alpha})) = \lambda_{\min}(H(\boldsymbol{\alpha})), \tag{2}$$

where $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue of a Hermitian matrix. Since $H(\boldsymbol{\alpha})$ acts on a 2^n-dimensional space, directly solving E_0 becomes computationally prohibitive as n increases. *When no confusion arises, we omit $\boldsymbol{\alpha}$ and write H and $E_0(H)$, or simply E_0, for notational simplicity.*

Sample-based quantum diagonalization (SQD). SQD has emerged as a powerful approach to large-scale GSE by combining quantum sampling with classical subspace diagonalization (Robledo-Moreno et al., 2025). To estimate E_0 in Eqn. (2), SQD collects S computational-basis measurement outcomes $\mathcal{S} = \{\boldsymbol{x}^{(s)} : \boldsymbol{x}^{(s)} \in \{0,1\}^n\}_{s=1}^{S}$ from repeated executions of a sampling circuit C. From these outcomes, at most K distinct configurations are retained in a subset $\mathcal{S}_K \subseteq \mathcal{S}$ with $|\mathcal{S}_K| \leq K$, defining a reduced subspace for subsequent classical diagonalization. Accordingly, the Hamiltonian H in Eqn. (1) is projected onto this subspace with entries $[H_K]_{\boldsymbol{x},\boldsymbol{x}'} = \langle \boldsymbol{x}|H|\boldsymbol{x}'\rangle$ for $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{S}_K$. SQD then solves the reduced eigenvalue problem, i.e.,

$$H_K \boldsymbol{v} = E\boldsymbol{v}, \qquad \text{and} \quad E_{\mathrm{SQD}}(\mathsf{C}; H, S, K) = \lambda_{\min}(H_K), \tag{3}$$

where $\boldsymbol{v}$ contains the coefficients in the selected basis. By the Rayleigh–Ritz variational principle (Saad, 2011), $E_{\mathrm{SQD}}(\mathsf{C}; H, S, K) \geq E_0(H)$, provided that the projected Hamiltonian is evaluated exactly and E_{SQD} is taken as its smallest eigenvalue.

Quantum architecture search (QAS). Variational quantum algorithms (VQAs) solve computational problems using parameterized quantum circuits $U(\boldsymbol{\theta})$, whose parameters $\boldsymbol{\theta}$ are optimized with respect to a task-specific objective Cerezo et al. (2021). For example, the variational quantum eigensolver (VQE) estimates the ground-state energy of a Hamiltonian H by solving $E_{\mathrm{VQE}} = \min_{\boldsymbol{\theta}} \langle 0| U^\dagger(\boldsymbol{\theta}) H U(\boldsymbol{\theta}) |0\rangle \geq E_0$. In conventional VQAs, the circuit architecture underlying $U(\boldsymbol{\theta})$ is predefined, and the optimization is performed only over $\boldsymbol{\theta}$.

QAS extends this paradigm by optimizing the circuit architecture $U(\boldsymbol{\theta})$ together with its parameters. Given an admissible gate set $\mathcal{G}$, such as single-qubit rotations and two-qubit entangling gates with flexible placements, a candidate circuit is represented as $\mathsf{C} = (\mathcal{A}, \boldsymbol{\theta})$, where $\mathcal{A}$ specifies the selected gates and their placements, and $\boldsymbol{\theta}$ denotes the associated parameters. Mathematically, we have

$$\mathsf{C}^\star = \arg\min_{\mathsf{C} \in \mathcal{C}} \langle 0| U_{\mathsf{C}}^\dagger H U_{\mathsf{C}} |0\rangle , \tag{4}$$

where the search space $\mathcal{C}$ is subject to constraints such as circuit depth and qubit connectivity.

Related work. The proposed Qas^{2}QD lies at the intersection of SQD and QAS. We therefore focus on these two topics here and defer a broader discussion of related work to Appendix A.2.

SQD. SQD constructs a variational subspace from sampled computational-basis configurations and diagonalize the Hamiltonian within it (Kanno et al., 2026; Robledo-Moreno et al., 2025). Subsequent developments improve sampling circuits through time evolution and circuit adaptation (Yu et al., 2025; Nakagawa et al., 2024; Kemmoku et al., 2026), increase measurement efficiency through history-dependent sampling and amplitude amplification (Stockinger et al., 2026; Kamoshita & Mitarai, 2026), and refine or enlarge the retained subspace $\mathcal{S}_K$ through classical feedback (Miura, 2026; Wang et al., 2026). The proposed Qas^{2}QD ***complements*** these advances by shifting the focus from optimizing SQD for individual Hamiltonians to learning transferable circuit-performance knowledge across a family of Hamiltonians.

QAS. QAS aims to automatically adapt quantum circuits to achieve higher performance for specified tasks, with much of the literature focusing on ansatz design for variational quantum algorithms (VQAs). For parameterized circuits, this can be formulated as joint optimization of discrete architectural choices and continuous gate parameters. Representative search strategies include heuristic search (Du et al., 2022), differentiable optimization (Zhang et al., 2022), reinforcement learning (Kuo et al., 2021), and generative methods (Nakaji et al., 2025). *For the first time*, our framework extends QAS to SQD, where circuit quality is determined by the E_{SQD} obtained from $\mathcal{S}_K$. Moreover, the multimodal predictor of Qas^{2}QD is tailored to this new search setting.

3 Implementation of Qas^{2}QD

Here we present the proposed Qas^{2}QD. We begin by reformulating GSE as a family-level learning problem in Sec. 3.1. We then introduce the overall framework of Qas^{2}QD in Sec. 3.2, followed by the implementation of its two key components in Sec. 3.3.

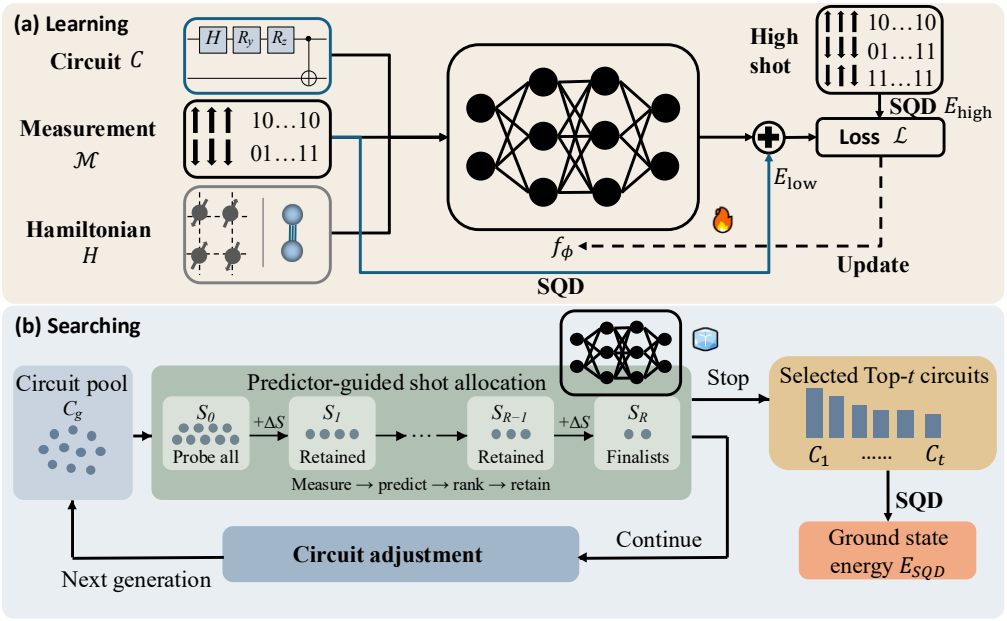

Figure 1: **Overview of the proposed Qas^{2}QD**. (a) Stage I learns a multimodal predictor from circuit structures, Hamiltonians, and low-shot measurements to predict high-shot SQD performance. (b) For an unseen Hamiltonian, Stage II uses the learned predictor to guide evolutionary QAS, progressively allocating measurements to promising circuits and selecting the best candidates for final SQD evaluation.

3.1 SQD Circuit Optimization for Family-Level GSE tasks

Ground-state estimation rarely concerns an isolated Hamiltonian. In quantum many-body physics, one typically studies how ground-state properties vary with physical parameters to characterize different physical regimes and quantum phase transitions (Huang et al., 2022; Lewis et al., 2024). Similarly, in quantum chemistry, varying molecular geometries gives rise to families of Hamiltonians whose ground-state energies define potential-energy surfaces (Cervera-Lierta et al., 2021). These problems naturally motivate a family of parameterized Hamiltonians

$$\mathcal{H} = \{H(\boldsymbol{\alpha}) : \boldsymbol{\alpha} \in \Lambda\}, \tag{5}$$

where $\boldsymbol{\alpha}$ denotes the physical parameters introduced in Eqn. (1).

Existing SQD approaches predominantly optimize or prescribe sampling circuits for individual Hamiltonian instances, without explicitly exploiting the relations among Hamiltonians within $\mathcal{H}$. Following Eqn. (3), for a given $H \in \mathcal{H}$, the ideal SQD sampling circuit can be formulated as

$$\mathsf{C}^\star(H) \in \arg\min_{\mathsf{C}\in\mathcal{C}} y(\mathsf{C}, H), \qquad y(\mathsf{C}, H) = \mathbb{E}\left[E_{\text{SQD}}(\mathsf{C}; H, S_{\text{ref}}, K)\right], \tag{6}$$

where $\mathcal{C}$ is the space of admissible sampling circuits defined by the prescribed gate set and circuit constraints, and the expectation in $y(\mathsf{C}, H)$ is taken over finite-shot measurements. Here, S_{ref} is a sufficiently large reference measurement budget for reliably assessing circuit performance. Solving Eqn. (6) independently for every target Hamiltonian requires repeatedly searching over $\mathcal{C}$ and evaluating candidate circuits with S_{ref} measurements, resulting in *substantial measurement costs*.

We instead formulate SQD circuit design over $\mathcal{H}$ as a family-level learning problem. Let $\mathcal{H}_{\text{train}} \subset \mathcal{H}$ denote a subset of Hamiltonians for which circuit-performance information has been collected, and let $\mathcal{D}$ denote the resulting training dataset. Given a previously unseen Hamiltonian $H \in \mathcal{H} \setminus \mathcal{H}_{\text{train}}$ and a limited measurement budget S_{budget} for the search, under which each candidate circuit can only be measured with far fewer than S_{ref} shots, the goal is to learn a circuit-design algorithm

$$\text{Alg}_{\boldsymbol{\phi}} : (\mathcal{D}, H, S_{\text{budget}}) \longmapsto \widehat{\mathsf{C}}(H) \in \mathcal{C}, \tag{7}$$

whose learnable parameters $\boldsymbol{\phi}$ are fitted once on $\mathcal{D}$ and reused for every unseen H, such that the resulting circuit approaches the SQD performance of the ideal circuit, i.e., $y(\widehat{\mathsf{C}}(H), H) \approx$

$y(\mathsf{C}^\star(H), H)$. This formulation shifts SQD circuit design from repeatedly optimizing each Hamiltonian from scratch to learning transferable knowledge across related GSE problems. By amortizing circuit-design knowledge across a Hamiltonian family, it provides a path toward efficient SQD for broader parameter regimes and increasingly challenging quantum systems.

3.2 The Framework of Qas^{2}QD

Qas^{2}QD realizes the family-level learning objective formulated above. A central challenge is that reliable SQD evaluation of a candidate circuit requires a large measurement budget S_{ref} in Eqn. (6), making repeated evaluations prohibitively expensive during architecture search. Simply reducing the number of shots to S_{low}, however, produces noisy SQD energies that can misrank candidate circuits and thereby misguide the search.

Qas^{2}QD addresses this trade-off by learning a surrogate for *circuit ranking*. Specifically, using only $S_{\mathrm{low}} \ll S_{\mathrm{ref}}$ measurements, the surrogate assigns each candidate $\mathsf{C} \in \mathcal{C}$ a score $\widehat{y}(\mathsf{C}, H)$ that approximates its reference-shot circuit quality $y(\mathsf{C}, H)$. Rather than requiring exact prediction, the key objective is to *preserve the reference-shot ordering of candidate circuits*, i.e., for any $\mathsf{C}_i, \mathsf{C}_j \in \mathcal{C}$,

$$y(\mathsf{C}_i, H) < y(\mathsf{C}_j, H) \quad \Longrightarrow \quad \widehat{y}(\mathsf{C}_i, H) < \widehat{y}(\mathsf{C}_j, H). \tag{8}$$

The implementation of Qas^{2}QD is illustrated in Fig. 1 and consists of two stages: multimodal predictor learning and predictor-guided circuit search. We next introduce the overall mechanism below, with further details of the two key components provided in the following subsections.

Stage I: Multimodal predictor learning. We construct the training dataset $\mathcal{D}$ by sampling candidate circuits $\mathsf{C}^{(i)} \in \mathcal{C}$ and Hamiltonians $H^{(i)} \in \mathcal{H}_{\mathrm{train}}$. For each pair, a low-shot measurement set $\mathcal{S}^{(i)}$ with S_{low} shots is collected, and the GSE are evaluated by SQD under both the low-shot and the reference-shot budgets, i.e., $E^{(i)}_{\mathrm{low}} = E_{\mathrm{SQD}}(\mathsf{C}^{(i)}; H^{(i)}, S_{\mathrm{low}}, K)$ and $E^{(i)}_{\mathrm{high}} = E_{\mathrm{SQD}}(\mathsf{C}^{(i)}; H^{(i)}, S_{\mathrm{ref}}, K)$ with $S_{\mathrm{low}} \ll S_{\mathrm{ref}}$. This yields $\mathcal{D} = \{(\mathsf{C}^{(i)}, H^{(i)}, \mathcal{S}^{(i)}, E^{(i)}_{\mathrm{low}}, E^{(i)}_{\mathrm{high}})\}_{i=1}^{N}$. Given $\mathcal{D}$, the predictor $f(\cdot;\boldsymbol{\phi})$ takes $(\mathsf{C}^{(i)}, H^{(i)}, \mathcal{S}^{(i)})$ as input and predicts the correction from the low-shot SQD energy to its reference-shot counterpart, i.e., $\widehat{E}^{(i)}_{\mathrm{high}} = E^{(i)}_{\mathrm{low}} + f(\mathsf{C}^{(i)}, H^{(i)}, \mathcal{S}^{(i)}; \boldsymbol{\phi})$, where $\boldsymbol{\phi}$ denotes the trainable parameters. The loss function to be minimized is

$$\mathcal{L}(\mathcal{D}; \boldsymbol{\phi}) = \frac{1}{N}\sum_{i=1}^{N} \left(E^{(i)}_{\mathrm{high}} - \widehat{E}^{(i)}_{\mathrm{high}}\right)^2 + \lambda_{\mathrm{rank}} \frac{1}{|\mathcal{P}|} \sum_{(i,j)\in\mathcal{P}} \max\left(0, m + \widehat{E}^{(i)}_{\mathrm{high}} - \widehat{E}^{(j)}_{\mathrm{high}}\right), \tag{9}$$

where the first term measures the prediction error relative to the reference-shot SQD energy, and the second term encourages the predictor to preserve the reference-shot ordering of candidate circuits, with $\lambda_{\mathrm{rank}} \geq 0$ balancing the two terms. Here, $\mathcal{P}$ contains pairs of circuits evaluated on the same Hamiltonian satisfying $E^{(i)}_{\mathrm{high}} < E^{(j)}_{\mathrm{high}}$, and $m > 0$ is the ranking margin.

Stage II: Predictor-guided QAS. For an unseen target Hamiltonian $H \in \mathcal{H} \setminus \mathcal{H}_{\mathrm{train}}$, Qas^{2}QD initializes a set of candidate circuits $\{\mathsf{C}_j\} \subset \mathcal{C}$ and collects a low-shot measurement set $\mathcal{S}^{(j)}$ for each candidate, whose number of shots $S_r \ll S_{\mathrm{ref}}$ is increased progressively over the search stages. The frozen predictor then outputs the surrogate score $\widehat{E}_{\mathrm{high}}(\mathsf{C}_j; H) = E_{\mathrm{SQD}}(\mathsf{C}_j; H, S_r, K) + f(\mathsf{C}_j, H, \mathcal{S}^{(j)}; \boldsymbol{\phi})$, which instantiates the surrogate in Eqn. (8) as $\widehat{y}(\mathsf{C}_j, H) = \widehat{E}_{\mathrm{high}}(\mathsf{C}_j; H)$ and is used for ranking candidate circuits and guiding QAS. As shown in Fig. 1(b), based on these scores, the search iteratively updates both circuit structures and gate parameters under a prescribed search budget S_{budget}, and only the selected circuits are evaluated with S_{ref} shots.

3.3 Two Key Components in Qas^{2}QD

Multimodal neural predictor. The predictor takes a candidate circuit description and low-shot measurement outcomes as multimodal inputs, conditioned on descriptors of the target Hamiltonian, shown in Fig. 1(a). Two Transformer encoders (Vaswani et al., 2017) separately process the circuit and measurement modalities to obtain their respective representations. A lightweight context encoder processes the Hamiltonian descriptors together with global statistics of the low-shot measurement distribution. The three representations are concatenated and passed to a multilayer perceptron

to estimate the correction term $f(\cdot;\phi)$. Detailed feature representations, network architectures, and training settings are provided in Appendix B.2.

Predictor-guided QAS for circuit design. We use the learned predictor to guide an evolutionary QAS (Du et al., 2022; Linghu et al., 2024; Martyniuk et al., 2024) for solving Eqn. (6). Starting from a population of candidate circuits, each candidate is first evaluated with a small probing budget S_0. The accumulated low-shot measurements are fed into the predictor to estimate its reference-shot SQD performance. In particular, Qas^{2}QD progressively concentrates measurements on promising circuits through successive halving (Jamieson & Talwalkar, 2016; Li et al., 2017). Specifically, at stage r, candidates are ranked by their predicted reference-shot SQD energies, and only the top-performing candidates proceed to the next stage with an increased cumulative budget $S_{r+1} > S_r$. This process continues until the final stage R or the prescribed search budget is reached.

The retained circuits then form the parent population for the next generation, where their circuit structures and rotation parameters are updated through evolutionary operations. The resulting candidates are again evaluated and filtered by the predictor, forming an iterative *predict–allocate–evolve* search process. Across all generations, the measurement cost of the search is constrained by S_{budget}. After search, the top-ranked circuits are evaluated independently, and the circuit achieving the lowest SQD energy is selected as the final design. Refer to Appendix B.3 for more details.

4 Experiments

We conduct experiments to evaluate Qas^{2}QD. Following conventions in GSE (Wang et al., 2023; 2024; Huang et al., 2026), here we focus on standard many-body quantum systems, while additional experimental results on chemical systems and the omitted details are provided in Appendix C.

4.1 Experimental setup

Two GSE tasks in many-body quantum systems. We evaluate Qas^{2}QD on two representative quantum many-body models: the transverse-field Ising model (TFIM) and the XXZ model, i.e.,

$$H_{\text{TFIM}}(J,h) = -J\sum_{i=1}^{n-1} Z_iZ_{i+1} - h\sum_i X_i,\ H_{\text{XXZ}}(J,\Delta) = J\sum_{i=1}^{n-1}\big(X_iX_{i+1}+Y_iY_{i+1}+\Delta Z_iZ_{i+1}\big).$$

where J is the coupling strength, h is the transverse-field strength, and Δ is the dimensionless XXZ anisotropy parameter. Following the notation in Eqn. (5), we have $\boldsymbol{\alpha} = (J,h)$ or $\boldsymbol{\alpha} = (J,\Delta)$.

Hyper-parameter settings of Qas^{2}QD. The training dataset $\mathcal{D}$ of Qas^{2}QD contains $N = 1.2\times10^4$ training examples. The validation and test sets are independently sampled from the same distribution, comprising 1.8×10^3 and 3×10^3 examples, respectively. Unless stated otherwise, we set the training shot budget to $S_{\text{low}} = 10^4$, the reference shot budget to $S_{\text{ref}} = 10^6$, and the retained-subspace size to $K = 2048$. Experiments requiring statistical analysis are repeated five times, unless otherwise stated.

Metric. The performance of Qas^{2}QD is evaluated by two measures. Specifically, let N_{sel} denote the number of top-ranked circuits returned by the search in Stage II, each of which is evaluated independently with S_{ref} shots. The mean high-shot SQD error of these selected circuits is

$$\varepsilon = \frac{1}{N_{\text{sel}}}\sum_{i=1}^{N_{\text{sel}}}\Big|E_{\text{SQD}}(\mathsf{C}^{(i)};H,S_{\text{ref}},K) - E_0(H)\Big|, \tag{10}$$

where $E_0(H)$ denotes the exact ground-state energy of H. This metric measures the quality of the selected circuits based on their high-shot SQD performance. We denote the per-circuit shot budgets for low-shot observations, test-time prediction or candidate ranking, each independent final SQD evaluation, and high-shot reference evaluation by S_{low}, S_{test}, S_{eval}, and S_{ref}, respectively.

Baseline. We compare Qas^{2}QD against the following typical baselines. (i) **Vanilla SQD**, which performs SQD using a predefined sampling circuit without QAS (Robledo-Moreno et al., 2025). (ii) **VQE**, a widely used variational approach for GSE (Kandala et al., 2017). Besides, we consider the *low-shot SQD-guided QAS* and *high-shot SQD-guided QAS* as reference models, which directly use SQD energies evaluated with S_{low} and S_{ref} shots, respectively, to guide circuit search in Stage II.

4.2 NUMERICAL RESULTS

Reliable performance of Qas^{2}QD from low-shot measurements. We first explore whether the proposed neural predictor of Qas^{2}QD can provide a reliable performance signal. For both TFIM and XXZ, we fix the qubit number as $n = 20$.

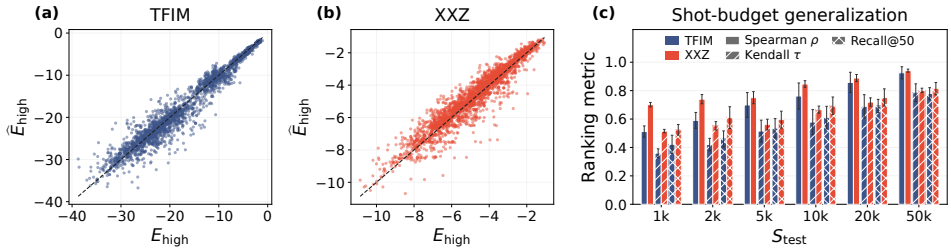

Figure 2: **Prediction of high-shot SQD energies from low-shot measurements.** Panels (a) and (b) compare the predicted and high-shot reference SQD energies for 20-qubit TFIM and XXZ, respectively, using $S_{\text{low}} = 10^4$ low-shot measurements. The dashed diagonal indicates perfect prediction. Panel (c) reports the Spearman correlation, Kendall's τ, and top-50 recall across different inference-time shot budgets S_{test}.

Fig. 2 evaluates the prediction of reference-budget SQD energies on held-out circuits. As shown in Figs. 2(a) and (b), with only $S_{\text{test}} = S_{\text{low}} = 10^4$ and $S_{\text{ref}} = 10^6$ shots, the predictions achieve mean Spearman correlations of 0.762 and 0.847 with the reference values for TFIM and XXZ, respectively. The strong rank correlations demonstrate that the learned predictor can reliably distinguish high- and low-performing sampling circuits using only low-shot observations.

We further examine how the inference-time measurement budget affects the ranking quality, i.e., the agreement between circuit rankings predicted from low-shot observations and those obtained from high-shot reference evaluations. As shown in Fig. 2(c), increasing S_{test} from 10^3 to 5×10^4 shots improves the Spearman correlation from 0.535 to 0.962 for TFIM and from 0.707 to 0.948 for XXZ, respectively. Besides, Kendall's τ and top-50 recall exhibit similar improvements. These results imply that given unseen Hamiltonians, the trained neural predictor effectively exploits additional measurements to produce increasingly reliable circuit rankings without retraining.

Qas^{2}QD designs sampling circuits to enhance performance with lower measurement cost. We next evaluate the complete Qas^{2}QD framework by using the trained predictor to guide circuit search under different evaluation measurement budgets S_{eval}. Following the previous experiments, we consider $n = 20$ for both TFIM with $J = 0.8, h = 0.4$ and XXZ with $J = 0.125, \Delta = 2.8$ and vary the total measurement budget S_{eval} from 500 to 2×10^4 shots. We compare Qas^{2}QD with two baselines. That is, *vanilla SQD* uses a randomly initialized fixed sampling circuit without architecture search, while *low-shot SQD-guided QAS* performs architecture search by directly using SQD energies estimated from S_{low} measurements as the search signal. In contrast, Qas^{2}QD uses the same S_{low} measurements to guide the architecture search.

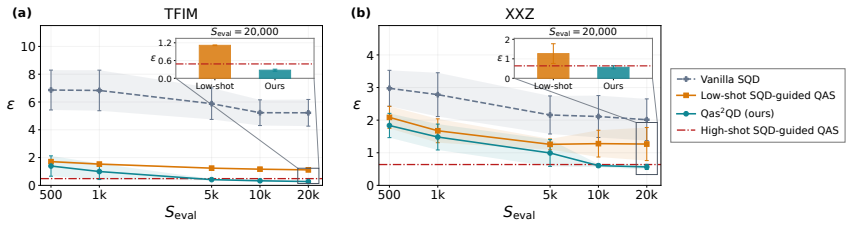

Figure 3: **SQD error ε versus the evaluation measurement budget S_{eval} for 20-qubit (a) TFIM and (b) XXZ.** Vanilla SQD, low-shot SQD-guided QAS, and Qas^{2}QD are evaluated under matched measurement budgets. Each point reports the mean over five independent circuits, with error bars indicating the standard error. The dashed red lines show the high-shot-guided search results as a separate reference.

As shown in Fig. 3, Qas^{2}QD consistently outperforms both the random fixed circuit and low-shot-SQD-guided QAS across different measurement budgets. Moreover, its advantage becomes increasingly obvious as S_{eval} increases, indicating that the learned predictor effectively exploits additional measurements to identify better sampling circuits. At $S_{\text{eval}} = 2 \times 10^4$, Qas^{2}QD achieves

SQD errors of approximately 0.27 and 0.57 for TFIM and XXZ, respectively, compared with 1.10 and 1.27 for low-shot-SQD-guided QAS.

Generalization of Qas^{2}QD across Hamiltonians and measurement budgets. Recall that a central goal of Qas^{2}QD is to learn transferable circuit-performance knowledge across a family of Hamiltonians. To this end, we evaluate the trained neural predictor under two forms of generalization. First, we test it on unseen TFIM Hamiltonians with (J, h) values drawn from the test set. Second, we vary the inference-time measurement budget S_{test} to examine how the predictor generalizes across measurement budgets, including those unseen during training.

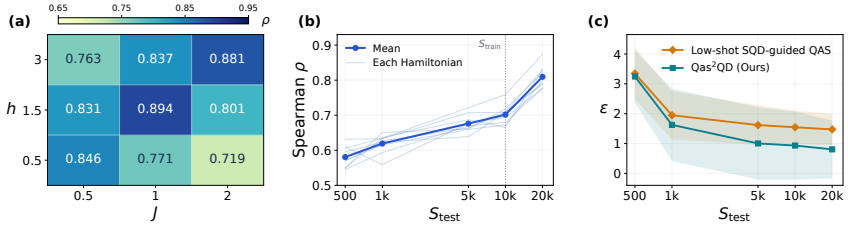

Figure 4: **Generalization of the Qas^{2}QD predictor across Hamiltonians and measurement budgets.** (a) Circuit-ranking performance across unseen 20-qubit TFIMs. (b) Ranking versus S_{test}: thin curves represent individual Hamiltonians and the thick curve their mean; (c) Circuit-selection performance across different inference-time measurement budgets S_{test}.

The achieved results validate the effectiveness of Qas^{2}QD. As shown in Fig. 4(a), the predictor maintains strong circuit-ranking performance across the unseen TFIM parameter grid, with Spearman's ρ ranging from 0.719 to 0.894 and averaging 0.816. Fig. 4(b) further shows that the same neural predictor effectively adapts to different measurement budgets: increasing S_{test} from 500 to 2×10^4 improves the Spearman correlation from 0.581 to 0.810. This improvement translates directly into better circuit selection, shown in Fig. 4(c). At $S_{\text{test}} = 10^4$, predictor-guided selection reduces the SQD error from 1.545 to 0.932 compared with direct low-shot SQD guidance, corresponding to a 39.7% reduction. At $S_{\text{test}} = 2 \times 10^4$, the error is further reduced from 1.474 to 0.804, or 45.4%. These results validate that Qas^{2}QD transfers across both Hamiltonians and measurement budgets, while effectively exploiting additional measurement information without retraining.

Ablation studies. We ablate the multimodal predictor in Qas^{2}QD to examine how its circuit, measurement, and Hamiltonian inputs contribute to prediction and the resulting circuit search. We consider a fixed 20-qubit TFIM instance with $(J, h) = (0.8, 0.4)$ and run the complete Qas^{2}QD search independently three times for each variant, with finite-shot measurements independently sampled in each run. We report both predictor-level metrics (MAE and Spearman correlation) and the final SQD error ε in Eqn. (10) achieved by the circuits identified through search.

Table 1: **Ablation of the multimodal predictor in Qas^{2}QD on TFIM with $n = 20$ and $(J, h) = (0.8, 0.4)$.** Different input modalities are compared in terms of prediction MAE, circuit-ranking correlation, and final SQD error. Results are reported as mean ± standard deviation over three independent runs.

Model	MAE ↓	Spearman ↑	Final SQD error ↓
Circuit only	4.680 ± 0.169	0.008 ± 0.014	6.911 ± 1.240
Measurement only	1.673 ± 0.055	0.692 ± 0.015	1.422 ± 0.488
Circuit + measurement	1.648 ± 0.036	0.688 ± 0.005	0.921 ± 0.466
No Hamiltonian input	1.502 ± 0.059	0.684 ± 0.019	0.681 ± 0.055
Full model	$\mathbf{1.167 \pm 0.035}$	$\mathbf{0.695 \pm 0.014}$	$\mathbf{0.336 \pm 0.095}$

As shown in Table 1, the circuit-only model yields near-zero rank correlation, whereas the measurement-only model achieves a Spearman correlation of 0.692, indicating that low-shot observations provide the primary predictive signal. Adding circuit information has little effect on Spearman correlation but reduces the mean final SQD error from 1.422 to 0.921. Incorporating non-Hamiltonian context statistics further reduces the error to 0.681, while the full model achieves 0.336. The full model obtains the best mean performance on every reported metric, reducing final SQD error by 76.4% relative to measurement-only guidance and by 50.7% relative to the no-Hamiltonian variant. These results support the value of Hamiltonian-oriented circuit selection.

Scaling behavior of Qas²QD. We next investigate whether the advantages of Qas²QD persist as the system size increases. We consider XXZ systems with $n \in \{10, 20, 30, 40\}$ qubits and fix the retained-subspace size to $K = 1024$. To enable scaling to larger systems, we employ a physically motivated circuit space (see Appendix C for the details). We compare Qas²QD with vanilla SQD and low-shot SQD-guided QAS under the same evaluation measurement budget $S_{\text{eval}} = S_{\text{ref}} = 10^6$. Results are averaged over five independent runs.

As shown in Table 2, circuit search consistently improves over vanilla SQD across all system sizes, while Qas²QD achieves the lowest mean error under the matched measurement budget. Compared with vanilla SQD, Qas²QD reduces the mean SQD error by 72.8%, 91.4%, 68.8%, and 59.2% for 10, 20, 30, and 40 qubits, respectively. Compared with low-shot SQD-guided QAS, the corresponding reductions are 14.5%, 12.4%, 6.2%, and 21.8%. Notably, Qas²QD maintains SQD errors at approximately 10^{-3} or below for 10–20 qubits and at the 10^{-2} level for systems up to 40 qubits. These results demonstrate that the benefit of predictor-guided circuit search persists with increasing system size n, supporting the scalability of Qas²QD. Additional results on the effect of the subspace dimension K are provided in Appendix C.

Table 2: **Mean absolute high-shot SQD error ε across system sizes at $K = 1024$.** Entries report means $\pm$ sample standard deviations over five paired search seeds. Bold values indicate the lower mean error among the two methods with matched reference measurement budgets S_{ref}.

n	Vanilla SQD	Low-shot SQD-guided QAS	Qas²QD (ours)	High-shot SQD-guided QAS
10	$(4.713 \pm 2.197) \times 10^{-3}$	$(1.500 \pm 0.042) \times 10^{-3}$	$(\mathbf{1.282} \pm 0.445) \times 10^{-3}$	$(4.482 \pm 4.033) \times 10^{-4}$
20	$(9.235 \pm 3.840) \times 10^{-3}$	$(9.031 \pm 2.759) \times 10^{-4}$	$(\mathbf{7.912} \pm 2.546) \times 10^{-4}$	$(3.113 \pm 0.218) \times 10^{-4}$
30	$(3.233 \pm 0.431) \times 10^{-2}$	$(1.074 \pm 0.282) \times 10^{-2}$	$(\mathbf{1.008} \pm 0.233) \times 10^{-2}$	$(6.548 \pm 0.163) \times 10^{-3}$
40	$(5.728 \pm 1.274) \times 10^{-2}$	$(2.988 \pm 0.312) \times 10^{-2}$	$(\mathbf{2.337} \pm 0.267) \times 10^{-2}$	$(1.640 \pm 0.022) \times 10^{-2}$

Qas²QD outperforms VQE. We compare our method with VQE on the same 40-qubit XXZ model, using the configuration described in the preceding paragraph on scaling behavior. The employed VQE uses finite-shot energy estimates and either a four-layer Hamiltonian variational ansatz (HVA) or a four-layer hardware-efficient ansatz (HEA) (Kandala et al., 2017), optimized using simultaneous perturbation stochastic approximation (SPSA) (Spall, 2000). For a controlled comparison, VQE and Qas²QD each use $77{,}440$ shots per run for circuit optimization and selection, with independent final evaluation accounted for separately. For VQE, we simultaneously vary the number of shots per energy evaluation and the number of optimization updates while keeping the number of total measurement fixed. As shown in Table 3, Qas²QD achieves a lower mean GSE error than both VQE variants across all tested measurement allocations. The best VQE result is obtained by HVA with 110 shots per energy evaluation and 320 SPSA updates, achieving a mean error of 0.0346. Qas²QD achieves an error of 0.0234, corresponding to a 32.5% reduction. These results demonstrate that Qas²QD remains competitive with variational ground-state estimation as the system scales to 40 qubits under the same finite-shot measurement budget.

Table 3: **VQE versus Qas²QD for 40-qubit XXZ**. S denotes shots per energy estimate and T denotes SPSA updates.

S	T	VQE–HVA	VQE–HEA
22	1,600	0.047 ± 0.019	0.481 ± 0.033
55	640	0.042 ± 0.010	0.489 ± 0.019
110	320	0.035 ± 0.011	0.480 ± 0.008
220	160	0.065 ± 0.024	0.467 ± 0.003
440	80	0.054 ± 0.009	0.463 ± 0.003
880	40	0.122 ± 0.068	0.460 ± 0.001
Qas²QD ($K = 1024$)		$\mathbf{0.02337 \pm 0.00267}$	

5 Conclusion

In this work, we reframed SQD sampling-circuit design from instance-wise optimization as a family-level learning problem across related Hamiltonians. Building on this formulation, we introduced Qas²QD, a learning-guided QAS framework that transfers circuit-performance knowledge across Hamiltonians to automatically design effective sampling circuits. Experiments on diverse quantum systems with up to 40 qubits demonstrate substantial improvements over the evaluated baselines, together with effective generalization to unseen Hamiltonians and measurement budgets. More broadly, our results establish architecture search as a promising new route toward scalable and adaptive quantum diagonalization.

AI USE STATEMENT

Generative AI tools were used to assist with the preparation and revision of the manuscript, including language polishing, refinement of academic expression, and improvements to clarity and readability. The authors are responsible for ensuring that AI-assisted revisions accurately reflect the technical content and intended meaning of the work. The authors take full responsibility for the accuracy, originality, and integrity of the final manuscript, including all methods, results, and conclusions.

ETHICS STATEMENT

This work develops learning-guided quantum circuit design for ground state energy estimation in quantum many-body systems. Both the method and numerical experiments do not involve human participants or personal data. Therefore, we do not identify specific ethical concerns or anticipate direct adverse impacts on individuals or society.

REPRODUCIBILITY STATEMENT

We are committed to supporting the independent verification and reproduction of the results presented in this work. During peer review, the source code and supporting experimental materials are provided at `https://anonymous.4open.science/r/QAS2QD-03DB`. Upon acceptance, we will publicly release the source code, benchmark specifications, experimental configurations, optimized circuits, and the result data and scripts needed to reproduce the reported tables and figures.

REFERENCES

Alán Aspuru-Guzik, Anthony D Dutoi, Peter J Love, and Martin Head-Gordon. Simulated quantum computation of molecular energies. *Science*, 309(5741):1704–1707, 2005.

Stefano Barison, Javier Robledo Moreno, and Mario Motta. Quantum-centric computation of molecular excited states with extended sample-based quantum diagonalization. *Quantum Science and Technology*, 10(2):025034, 2025.

Bela Bauer, Sergey Bravyi, Mario Motta, and Garnet Kin-Lic Chan. Quantum algorithms for quantum chemistry and quantum materials science. *Chemical reviews*, 120(22):12685–12717, 2020.

Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, 2021.

Alba Cervera-Lierta, Jakob S. Kottmann, and Alán Aspuru-Guzik. Meta-variational quantum eigensolver: Learning energy profiles of parameterized Hamiltonians for quantum simulation. *PRX Quantum*, 2(2):020329, 2021.

Jeffrey Cohn, Mario Motta, and Robert M Parrish. Quantum filter diagonalization with compressed double-factorized hamiltonians. *PRX Quantum*, 2(4):040352, 2021.

Cristian L Cortes and Stephen K Gray. Quantum krylov subspace algorithms for ground-and excited-state energy estimation. *Physical Review A*, 105(2):022417, 2022.

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

Yuxuan Du, Yan Zhu, Yuan-Hang Zhang, Min-Hsiu Hsieh, Patrick Rebentrost, Weibo Gao, Yi-Zhuang You, Jens Eisert, Giulio Chiribella, Dacheng Tao, et al. Artificial intelligence for representing and characterizing quantum systems. *Nature Reviews Physics*, pp. 1–18, 2026.

Alan Duriez, Pamela C Carvalho, Marco Antonio Barroca, Federico Zipoli, Ben Jaderberg, Rodrigo Neumann Barros Ferreira, Kunal Sharma, Antonio Mezzacapo, Benjamin Wunsch, and Mathias Steiner. Computing band gaps of periodic materials via sample-based quantum diagonalization. *npj Computational Materials*, 2026.

Ethan N Epperly, Lin Lin, and Yuji Nakatsukasa. A theory of quantum subspace diagonalization. *SIAM Journal on Matrix Analysis and Applications*, 43(3):1263–1290, 2022.

Richard P Feynman et al. Simulating physics with computers. *Feynman and computation*, 22: 133–153, 2018.

Dillon Frame, Rongzheng He, Ilse Ipsen, Daniel Lee, Dean Lee, and Ermal Rrapaj. Eigenvector continuation with subspace learning. *Physical review letters*, 121(3):032501, 2018.

Iulia M Georgescu, Sahel Ashhab, and Franco Nori. Quantum simulation. *Reviews of Modern Physics*, 86(1):153–185, 2014.

Zhimin He, Chuangtao Chen, Lvzhou Li, Shenggen Zheng, and Haozhen Situ. Quantum architecture search with meta-learning. *Advanced Quantum Technologies*, 5(8):2100134, 2022.

Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V Albert, and John Preskill. Provably efficient machine learning for quantum many-body problems. *Science*, 377(6613):eabk3333, 2022. doi: 10.1126/science.abk3333.

Yiming Huang, Yajie Hao, Yuxuan Du, Jing Zhou, Xiao Yuan, and Xiaoting Wang. Palqo: Physics-informed model for accelerating large-scale quantum optimization. *Advances in Neural Information Processing Systems*, 38:58808–58830, 2026.

Kevin Jamieson and Ameet Talwalkar. Non-stochastic best arm identification and hyperparameter optimization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pp. 240–248. PMLR, 2016. URL `https://proceedings.mlr.press/v51/jamieson16.html`.

Masahiko Kamoshita and Kosuke Mitarai. Qsci-cmp: Quantum-selected configuration interaction with chemically motivated preselection. *arXiv preprint arXiv:2608.05766*, 2026.

Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *nature*, 549(7671):242–246, 2017.

Keita Kanno, Masaya Kohda, Ryosuke Imai, Sho Koh, Kosuke Mitarai, Wataru Mizukami, and Yuya O Nakagawa. Quantum-selected configuration interaction: Classical diagonalization of hamiltonians in subspaces selected by quantum computers. *Physical Review Research*, 8(2): 023268, 2026.

Ryota Kemmoku, Qi Gao, Shu Kanno, Kimberlee Keithley, Ikko Hamamura, Naoki Yamamoto, and Kouhei Nakaji. Generative circuit design for quantum-selected configuration interaction. *arXiv preprint arXiv:2604.09756*, 2026.

En-Jui Kuo, Yao-Lung L. Fang, and Samuel Yen-Chi Chen. Quantum architecture search via deep reinforcement learning. *arXiv preprint arXiv:2104.07715*, 2021. doi: 10.48550/arXiv.2104.07715.

Benjamin P Lanyon, James D Whitfield, Geoff G Gillett, Michael E Goggin, Marcelo P Almeida, Ivan Kassal, Jacob D Biamonte, Masoud Mohseni, Ben J Powell, Marco Barbieri, et al. Towards quantum chemistry on a quantum computer. *Nature chemistry*, 2(2):106–111, 2010.

Seunghoon Lee, Joonho Lee, Huanchen Zhai, Yu Tong, Alexander M Dalzell, Ashutosh Kumar, Phillip Helms, Johnnie Gray, Zhi-Hao Cui, Wenyuan Liu, et al. Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nature communications*, 14 (1):1952, 2023.

Laura Lewis, Hsin-Yuan Huang, Viet T. Tran, Sebastian Lehner, Richard Kueng, and John Preskill. Improved machine learning algorithm for predicting ground state properties. *Nat. Commun.*, 15 (1):895, 2024. doi: 10.1038/s41467-024-45014-7.

Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*, 18(185):1–52, 2017.

Kehuan Linghu, Yang Qian, Ruixia Wang, Meng-Jun Hu, Zhiyuan Li, Xuegang Li, Huikai Xu, Jingning Zhang, Teng Ma, Peng Zhao, et al. Quantum circuit architecture search on a superconducting processor. *Entropy*, 26(12):1025, 2024.

Xudong Lu, Kaisen Pan, Ge Yan, Jiaming Shan, Wenjie Wu, and Junchi Yan. Qas-bench: rethinking quantum architecture search and a benchmark. In *International Conference on Machine Learning*, pp. 22880–22898. PMLR, 2023.

Darya Martyniuk, Johannes Jung, and Adrian Paschke. Quantum architecture search: a survey. In *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, volume 1, pp. 1695–1706. IEEE, 2024.

Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C Benjamin, and Xiao Yuan. Quantum computational chemistry. *Reviews of Modern Physics*, 92(1):015003, 2020.

Jarrod R McClean, Mollie E Kimchi-Schwartz, Jonathan Carter, and Wibe A De Jong. Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states. *Physical Review A*, 95(4):042308, 2017.

Fan-Xu Meng, Ze-Tong Li, Xu-Tao Yu, and Zai-Chen Zhang. Quantum circuit architecture optimization for variational quantum eigensolver via monto carlo tree search. *IEEE Transactions on Quantum Engineering*, 2:1–10, 2021.

Kenneth M Merz Jr, Akhil Shajan, Danil Kaliakin, Fangchun Liang, Yuichi Otsuka, Tomonori Shirakawa, Lukas Broers, Han Xu, Miwako Tsuji, Mitsuhisa Sato, et al. Crossing the 12,000-atom barrier with heterogeneous quantum-classical supercomputing: quantum chemistry of protein-ligand complexes. *arXiv preprint arXiv:2605.01138*, 2026.

Mathias Mikkelsen and Yuya O Nakagawa. Quantum-selected configuration interaction with time-evolved state. *Physical Review Research*, 7(4):043043, 2025.

Rinka Miura. Active sampling sample-based quantum diagonalization from finite-shot measurements. In *2026 International Conference on Quantum Communications, Networking, and Computing (QCNC)*, pp. 1065–1070. IEEE, 2026.

Mario Motta, Chong Sun, Adrian TK Tan, Matthew J O'Rourke, Erika Ye, Austin J Minnich, Fernando GSL Brandao, and Garnet Kin-Lic Chan. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nature Physics*, 16(2):205–210, 2020.

Yuya O Nakagawa, Masahiko Kamoshita, Wataru Mizukami, Shotaro Sudo, and Yu-ya Ohnishi. Adapt-qsci: Adaptive construction of an input state for quantum-selected configuration interaction. *Journal of Chemical Theory and Computation*, 20(24):10817–10825, 2024.

Kouhei Nakaji, Lasse Bjorn Kristensen, Ryota Kemmoku, et al. The generative quantum eigensolver and its application for ground state search. *arXiv preprint arXiv:2401.09253*, 2025. doi: 10.48550/arXiv.2401.09253.

Michael A Nielsen, Isaac L Chuang, et al. *Quantum computation and quantum information*, volume 1. Cambridge university press Cambridge, 2000.

Kosuke Nogaki, Steffen Backes, Tomonori Shirakawa, Seiji Yunoki, and Ryotaro Arita. Symmetry-adapted sample-based quantum diagonalization: Application to lattice model. *arXiv preprint arXiv:2505.00914*, 2025.

Mateusz Ostaszewski, Lea M Trenkwalder, Wojciech Masarczyk, Eleanor Scerri, and Vedran Dunjko. Reinforcement learning for optimization of variational quantum circuit architectures. *Advances in Neural Information Processing Systems*, 34:18182–18194, 2021.

Robert M Parrish and Peter L McMahon. Quantum filter diagonalization: Quantum eigendecomposition without full quantum phase estimation. *arXiv preprint arXiv:1909.08925*, 2019.

Samuele Piccinelli, Alberto Baiardi, Stefano Barison, Max Rossmannek, Almudena Carrera Vazquez, Francesco Tacchino, Stefano Mensa, Edoardo Altamura, Ali Alavi, Mario Motta, et al. Quantum chemistry with provable convergence via randomized sample-based krylov quantum diagonalization. *arXiv preprint arXiv:2508.02578*, 2025.

Javier Robledo-Moreno, Mario Motta, Holger Haas, Ali Javadi-Abhari, Petar Jurcevic, William Kirby, Simon Martiel, Kunal Sharma, Sandeep Sharma, Tomonori Shirakawa, et al. Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer. *Science Advances*, 11(25):eadu9991, 2025.

Yousef Saad. *Numerical methods for large eigenvalue problems: revised edition*. SIAM, 2011.

Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3):032331, 2019.

Tomonori Shirakawa, Javier Robledo-Moreno, Toshinari Itoko, Vinay Tripathi, Kento Ueda, Yukio Kawashima, Lukas Broers, William Kirby, Himadri Pathak, Hanhee Paik, et al. Closed-loop calculations of electronic structure on a quantum processor and a classical supercomputer at full scale. *Future Generation Computer Systems*, pp. 108731, 2026.

James C Spall. Adaptive stochastic approximation by the simultaneous perturbation method. *IEEE transactions on automatic control*, 45(10):1839–1853, 2000.

Nina Stockinger, Ludwig Nützel, and Michael J Hartmann. Sample-based quantum diagonalization with amplitude amplification. *arXiv preprint arXiv:2605.02565*, 2026.

Kenji Sugisaki, Shu Kanno, Toshinari Itoko, Rei Sakuma, and Naoki Yamamoto. Hamiltonian simulation-based quantum-selected configuration interaction for large-scale electronic structure calculations with a quantum computer. *Physical Chemistry Chemical Physics*, 27(38):20869–20884, 2025.

Giacomo Torlai and Matthew Fishman. Pastaq: A package for simulation, tomography and analysis of quantum computers, 2020. *URL https://github. com/GTorlai/PastaQ. jl*.

Dyon van Vreumingen and Kareljan Schoutens. Adiabatic ground-state preparation of fermionic many-body systems from a two-body perspective. *Physical Review A*, 108(6):062603, 2023.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Jie Wang, Jacopo Surace, Irénée Frérot, Benoît Legat, Marc-Olivier Renou, Victor Magron, and Antonio Acín. Certifying ground-state properties of many-body systems. *Physical Review X*, 14 (3):031006, 2024.

Qiaohong Wang, Kevin J Sung, Ruhee D'Cunha, Matthew R Hermes, Tanvi Gujarati, Yukio Kawashima, Yu-ya Ohnishi, Gavin O Jones, Mario Motta, and Laura Gagliardi. Localized sample-based quantum diagonalization for strongly correlated chemistry. *Proceedings of the National Academy of Sciences*, 123(28):e2603914123, 2026.

Xinbiao Wang, Junyu Liu, Tongliang Liu, Yong Luo, Yuxuan Du, and Dacheng Tao. Symmetric pruning in quantum neural networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=K96AogLDT2K.

Wenjie Wu, Ge Yan, Xudong Lu, Kaisen Pan, and Junchi Yan. Quantumdarts: differentiable quantum architecture search for variational quantum algorithms. In *International Conference on Machine Learning*, pp. 37745–37764. PMLR, 2023.

Esther Ye and Samuel Yen-Chi Chen. Quantum architecture search via continual reinforcement learning. *arXiv preprint arXiv:2112.05779*, 2021.

Kübra Yeter-Aydeniz, Raphael C Pooser, and George Siopsis. Practical quantum computation of chemical and nuclear energy levels using quantum imaginary time evolution and lanczos algorithms. *npj Quantum Information*, 6(1):63, 2020.

Nobuyuki Yoshioka, Mirko Amico, William Kirby, Petar Jurcevic, Arkopal Dutt, Bryce Fuller, Shelly Garion, Holger Haas, Ikko Hamamura, Alexander Ivrii, et al. Krylov diagonalization of large many-body hamiltonians on a quantum processor. *Nature Communications*, 16(1):5014, 2025.

Jeffery Yu, Javier Robledo Moreno, Joseph T Iosue, Luke Bertels, Daniel Claudino, Bryce Fuller, Peter Groszkowski, Travis S Humble, Petar Jurcevic, William Kirby, et al. Quantum-centric algorithm for sample-based krylov diagonalization. *arXiv preprint arXiv:2501.09702*, 2025.

Shi-Xin Zhang, Chang-Yu Hsieh, Shengyu Zhang, and Hong Yao. Neural predictor based quantum architecture search. *Machine Learning: Science and Technology*, 2(4):045027, 2021.

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

A PRELIMINARY

In this section, we first introduce the background needed to understand our framework in Appendix A.1, then provide the literature review related to $\mathsf{Qas}^2\mathsf{QD}$ in Appendix A.2.

A.1 BACKGROUND

Quantum computing basics. The fundamental element in quantum computation is a qubit, which has two computational-basis states $|0\rangle$ and $|1\rangle$ using Dirac notation. An n-qubit pure state is a normalized vector in $(\mathbb{C}^2)^{\otimes n}$ and can be written as

$$|\psi\rangle = \sum_{\boldsymbol{x}\in\{0,1\}^n} a_{\boldsymbol{x}} |\boldsymbol{x}\rangle$$

with $\sum_{\boldsymbol{x}} |a_{\boldsymbol{x}}|^2 = 1$. Here $|\boldsymbol{x}\rangle = |x_1\rangle \cdots \otimes |x_n\rangle$. To process or prepare a quantum state, we construct a quantum circuit $\mathsf{C} = (\mathcal{A}, \boldsymbol{\theta})$, where $\mathcal{A}$ and $\boldsymbol{\theta}$ denote the circuit layout and the corresponding parameters, respectively. Let U_{C} denote the unitary implemented by the circuit C. The final n-qubit state can be written as $|\psi_{\mathsf{C}}\rangle = U_{\mathsf{C}} |0\rangle$.

The construction of the quantum circuit U_{C} is completed by a set of *quantum gates* (Nielsen et al., 2000). Namely, an n-qubit gate $U \in \mathcal{U}(2^n)$ obeys $UU^\dagger = U^\dagger U = I_{2^n}$, where $\mathcal{U}(2^n)$ refers to the unitary group in dimension 2^n. Typical single-qubit quantum gates include the Pauli gates, which can be written as Pauli matrices:

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \tag{11}$$

The more general quantum gates are their corresponding rotation gates $R_x(\theta) = e^{-i\frac{\theta}{2}X}, R_y(\theta) = e^{-i\frac{\theta}{2}Y}$, and $R_z(\theta) = e^{-i\frac{\theta}{2}Z}$ with a tunable parameter θ, which can be written in the matrix form as

$$R_x(\theta) = \begin{bmatrix} \cos\frac{\theta}{2} & -i\sin\frac{\theta}{2} \\ -i\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{bmatrix}, R_y(\theta) = \begin{bmatrix} \cos\frac{\theta}{2} & -\sin\frac{\theta}{2} \\ \sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{bmatrix}, R_z(\theta) = \begin{bmatrix} e^{-i\frac{\theta}{2}} & 0 \\ 0 & e^{i\frac{\theta}{2}} \end{bmatrix}. \tag{12}$$

They are equivalent to rotating a tunable angle θ around x, y, and z axes of the Bloch sphere, and recovering the Pauli gates X, Y, and Z when $\theta = \pi$ up to a global phase. Moreover, a multi-qubit gate can be either an individual gate (e.g., CNOT gate with $\text{CNOT} = \begin{pmatrix} 1&0&0&0 \\ 0&1&0&0 \\ 0&0&0&1 \\ 0&0&1&0 \end{pmatrix}$) or a tensor product of multiple single-qubit gates. By the universality of quantum gate sets (Nielsen et al., 2000), an arbitrary unitary operation U can be approximated to any desired accuracy by a circuit composed of gates from a universal set, such as $\{\mathsf{H}, R_z, \text{CNOT}\}$ with $\mathsf{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$.

Measurement along the computational bases returns $\boldsymbol{x}$ with probability $p_{\mathsf{C}}(\boldsymbol{x}) = |\langle \boldsymbol{x}|\psi_{\mathsf{C}}\rangle|^2$. After conducting multiple measurements, we count the frequencies of each configuration $\boldsymbol{x}$ as the estimation of the corresponding probabilities.

Ground state energy estimation (GSE). A quantum system can be described by a Hamiltonian H, whose Pauli decomposition is given in Eqn. (1).

The mathematical formulation of a Hamiltonian H in Eqn. (1) indicates that it is a Hermitian operator. Consequently, it can always be expressed through its eigen-decomposition as

$$H(\boldsymbol{\alpha}) = \sum_{k=1}^{2^n} \lambda_k |\psi_k\rangle\langle\psi_k| \in \mathbb{C}^{2^n \times 2^n}. \tag{13}$$

Suppose that the eigenvalues are arranged in ascending order, i.e., $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{2^n}$. Then the eigenvalue λ_k represents the k-th energy level and the eigenvector $|\psi_k\rangle$ is the corresponding energy state. The **ground state** and the **ground state energy** of $H(\boldsymbol{\alpha})$ refer to $|\psi_1\rangle$ and λ_1, respectively. Throughout the study, we interchangeably use $|\psi(\boldsymbol{\alpha})\rangle$ and $|\psi\rangle$ (or $E_0(\boldsymbol{\alpha})$ and E_0) to specify the ground state (or ground state energy).

Sample-based quantum diagonalization (SQD). SQD is a hybrid approach for GSE, following the procedure of sampling on quantum devices and diagonalization on classical devices. This procedure is also central to quantum-selected configuration interaction (QSCI) (Kanno et al., 2026; Robledo-Moreno et al., 2025). Specifically, the procedure of SQD in this work consists of three steps:

Step 1 *State preparation and sampling.* Let C denote a sampling circuit that prepares the state $|\psi_{\mathsf{C}}\rangle = U_{\mathsf{C}}|0\rangle$. Measuring $|\psi_{\mathsf{C}}\rangle$ in the computational basis for S shots yields a collection of bit strings $\mathcal{S} = \{\boldsymbol{x}^{(s)}\}_{s=1}^{S}$, where $\boldsymbol{x}^{(s)} \in \{0,1\}^n$. Let $\nu_{\boldsymbol{x}}$ denote the number of times a bit string $\boldsymbol{x}$ appears in $\mathcal{S}$, such that $S = \sum_{\boldsymbol{x}} \nu_{\boldsymbol{x}}$. The empirical sampling probability of each configuration (i.e., bitstring) $\boldsymbol{x}$ is then given by

$$\widehat{p}_{\mathsf{C}}^{(S)}(\boldsymbol{x}) = \frac{\nu_{\boldsymbol{x}}}{S}.$$

Step 2 *Configuration selection.* Given the S measurement outcomes, we retain at most K distinct configurations with the largest empirical probabilities $\widehat{p}_{\mathsf{C}}^{(S)}(\boldsymbol{x})$. Let $\mathcal{S}_K$ denote the resulting set of configurations. The retained configurations define the SQD subspace

$$\mathcal{V}(\mathsf{C}; \mathcal{S}, K) = \mathrm{span}\{|\boldsymbol{x}\rangle : \boldsymbol{x} \in \mathcal{S}_K\}. \tag{14}$$

If fewer than K configurations are observed, all are retained.

Step 3 *Subspace construction and diagonalization.* Classically compute the matrix elements of the target Hamiltonian H within the selected subspace. Let $M = |\mathcal{S}_K| \leq K$ denote the number of retained configurations, and index them as $\mathcal{S}_K = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_M\}$. The resulting projected Hamiltonian $H_K \in \mathbb{C}^{M \times M}$ has matrix elements

$$[H_K]_{i,j} = \langle \boldsymbol{x}_i | H | \boldsymbol{x}_j \rangle, \qquad i, j \in \{1, \ldots, M\}.$$

After that, SQD solves the reduced eigenvalue problem with

$$E_{\mathrm{SQD}}(\mathsf{C}; H, S, K) = \lambda_{\min}(H_K), \tag{15}$$

which coincides with Eqn. (3) in the main text; the dependence on the sampled set $\mathcal{S}$ enters through $\mathcal{S}_K$.

Variational Quantum Algorithms Variational Quantum Algorithms (VQAs) are hybrid quantum-classical algorithms designed to utilize noisy intermediate-scale quantum (NISQ) hardware (Cerezo et al., 2021).

The primary objective of VQAs is to optimize a loss function, typically defined as

$$E(\boldsymbol{\theta}) = \langle \psi(\boldsymbol{\theta}) | H | \psi(\boldsymbol{\theta}) \rangle,$$

where $|\psi(\boldsymbol{\theta})\rangle$ is a tunable quantum state generated by applying a parameterized quantum circuit to a reference state and H represents a Hamiltonian encoding the target problem. To minimize $E(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$, a common approach is to apply a gradient-based method that iteratively adjusts the parameters according to the partial derivative $\partial_{\boldsymbol{\theta}} E$, which can be estimated using the Parameter Shift Rule (Schuld et al., 2019), i.e.,

$$\frac{\partial E}{\partial \boldsymbol{\theta}_j} = \frac{1}{2}\left(E(\boldsymbol{\theta}_j + \frac{\pi}{2}) - E(\boldsymbol{\theta}_j - \frac{\pi}{2})\right).$$

This method provides a closed-form expression for gradient estimation without requiring additional qubits. However, it necessitates evaluating $E(\boldsymbol{\theta})$ twice with shifted parameter values at the same position to estimate the gradient of a single parameter, which becomes computationally prohibitive for large-scale tasks.

A.2 Extended related work

We expand the brief discussion in Section 2. We first position our method with sample-based quantum diagonalization, followed by introduction about other approaches for GSE.

A.2.1 Variants of SQD

The accuracy of SQD introduced in Appendix A.1 is governed by how well the retained subspace $\mathcal{V}$ captures the target eigenstates. This, in turn, depends on the sampling circuit C, measurement budget S, and device noise, which determine the configurations available for selection, as well as the classical post-processing strategy used to construct the subspace. A larger subspace can improve representational capacity but increases classical diagonalization costs.

These trade-offs have motivated several directions for improving SQD. Here, we focus on three that are most relevant to our framework, i.e., *designing effective sampling circuits*, *improving measurement efficiency*, and *adaptively constructing the retained subspace*.

Sampling-state construction and circuit adaptation. One line of work changes the states from which configurations are sampled. Hamiltonian simulation-based quantum-selected configuration interaction (QSCI), and time-evolved QSCI sample states obtained by evolving an initial reference under the target Hamiltonian, thereby accessing configurations generated by its couplings (Sugisaki et al., 2025; Mikkelsen & Nakagawa, 2025). Sample-based Krylov quantum diagonalization (SKQD) collects configurations from a sequence of time-evolved states, and SqDRIFT uses randomized Hamiltonian simulation to reduce the circuit cost of generating these samples (Yu et al., 2025; Piccinelli et al., 2025). The convergence guarantees for SKQD and SqDRIFT rely on assumptions about the reference-state overlap and concentration of the ground state. Other methods adapt the sampling circuit C through feedback. For example, ADAPT-QSCI (Nakagawa et al., 2024) grows the circuit using operator-selection information computed from the classically reconstructed state. Closed-loop electronic-structure calculations optimize circuit parameters (Shirakawa et al., 2026), while generative circuit design trains a Transformer policy using QSCI energies to propose operator sequences (Kemmoku et al., 2026)

Measurement-efficient and history-dependent sampling. Another line of work aims to reduce redundant measurements of previously observed configurations. To be concrete, SQD with amplitude amplification (SQD-AA) leverages the measurement history to suppress previously sampled outcomes, thereby increasing the probability of discovering new configurations (Stockinger et al., 2026). Similarly, QSCI with chemically motivated preselection (QSCI-CMP) incorporates configurations identified from prior chemical knowledge directly into the classical subspace and excludes them from subsequent amplification, avoiding quantum resources spent on rediscovering known configurations (Kamoshita & Mitarai, 2026). These approaches improve the efficiency of configuration discovery, albeit at the cost of additional coherent quantum operations.

Adaptive subspace selection and enlargement. Classical feedback can improve which configurations enter the projected problem, even after quantum sampling has finished. Active Sampling SQD (AS-SQD) initializes a subspace from measured bitstrings, generates additional configurations connected to the current support by H, and selects them using an importance score (Miura, 2026). Its iterative expansion adapts the contents of the subspace while increasing its dimension. Carryover and refinement procedures retain important configurations across iterations or combine the supports of several approximate wavefunctions (Wang et al., 2026; Kemmoku et al., 2026).

Our work is complementary to these SQD extensions but focuses on a different dimension by enabling generalizable sampling-circuit design across a family of Hamiltonians. Going beyond instance-wise sampling circuit optimization, $\mathsf{Qas}^2\mathsf{QD}$ learns transferable circuit-performance knowledge across related Hamiltonians, enabling efficient circuit design for unseen instances.

A.2.2 OTHER APPROACHES FOR GSE

Besides SQD, there are other approaches to address GSE tasks. Quantum phase estimation (QPE) estimates energy eigenvalues through controlled Hamiltonian evolution (Aspuru-Guzik et al., 2005). Its probability of returning the ground-state energy depends on the initial state's ground-state overlap, while higher precision requires longer coherent evolution. The variational quantum eigensolver (VQE) instead minimizes measured energy expectations through classical optimization of a parameterized state (Kandala et al., 2017). Its accuracy depends on the expressivity of the ansatz, the optimization procedure, and the measurement budget.

A.3 QUANTUM ARCHITECTURE SEARCH

Quantum architecture search (QAS) (Zhang et al., 2021; Ye & Chen, 2021; Ostaszewski et al., 2021; Kuo et al., 2021; Meng et al., 2021; Du et al., 2022; Linghu et al., 2024; He et al., 2022; Zhang et al., 2022; Wu et al., 2023; Lu et al., 2023; Martyniuk et al., 2024) automates the design of quantum circuits by selecting gate operations, connectivity, and circuit depth, often together with optimizing continuous gate parameters. This approach is proposed to improve the performance of variational quantum algorithms Cerezo et al. (2021).

Population-based and parameter-sharing search. Population-based methods iteratively select promising circuits and generate new candidates through mutation or recombination. To reduce the cost of optimizing each candidate, parameter-sharing methods reuse trained parameters across architectures. For example, Du et al. (2022) organize candidates within a supernet and share parameters between compatible circuit layers. Candidates are then ranked using the inherited parameters, and the selected architecture is fine-tuned.

Differentiable search. Differentiable QAS represents discrete gate choices through continuously parameterized probability distributions. It optimizes these distributions alongside gate parameters using objective values from sampled circuits (Zhang et al., 2022). This formulation allows architecture selection and parameter optimization to proceed within a common gradient-based framework.

Reinforcement-learning and generative search. Reinforcement-learning methods formulate circuit construction as a sequence of actions, training a policy using task-dependent rewards (Kuo et al., 2021). Generative approaches learn distributions over gate sequences from which new candidates can be sampled. The generative quantum eigensolver, for example, trains a Transformer to propose circuits using energy feedback and supports pretraining on existing circuit data (Nakaji et al., 2025). Both approaches learn how to propose promising circuits from previous evaluations.

Qas^{2}QD differs from prior QAS studies in two key aspects. First, it introduces architecture search to SQD for the first time. Second, it develops a new multimodal neural predictor tailored to this setting. These advances open a new direction for integrating QAS with SQD.

B IMPLEMENTATION DETAILS OF QAS^{2}QD

In this section, we introduce the implementation details of Qas^{2}QD, including dataset construction, model architecture and training strategy, and inference pipeline.

B.1 DATASET CONSTRUCTION

Recall that the training dataset in Qas^{2}QD is constructed from a collection of candidate circuits and Hamiltonians. Specifically, each circuit–Hamiltonian pair (C_i, H_j) with $H_j = H(\boldsymbol{\alpha}^{(j)}) \in \mathcal{H}_{\text{train}}$ defines a training instance, where the same circuit is paired with multiple Hamiltonians; the main text indexes these pairs by a single index i. For each pair, we construct a low-shot input from S_{low} measurements and obtain its supervision target from a larger reference budget $S_{\text{ref}} \gg S_{\text{low}}$. In this subsection, we describe the construction of the candidate circuit space and the corresponding measurement–label pairs.

Candidate circuit space. Qas^{2}QD operates over a configurable circuit space $\mathcal{C}$ defined by the allowed quantum gate set $\mathcal{G}$. The space can be further constrained by the native gates and connectivity of the target quantum device, as well as user-specified requirements on circuit depth or structure. A

candidate circuit $\mathsf{C} \in \mathcal{C}$ consists of a sequence of L layers,

$$U_{\mathsf{C}} = U_L U_{L-1} \cdots U_1, \qquad U_\ell = \prod_{g \in \mathcal{G}_\ell} U_g(\theta_g), \tag{16}$$

where $\mathcal{G}_\ell \subseteq \mathcal{G}$ specifies the gates selected at layer ℓ, and θ_g denotes the associated parameter when g is parameterized. Gate placements are restricted to valid qubits or qubit pairs under the prescribed connectivity. Candidate architectures and their parameters are then sampled from $\mathcal{C}$ to construct the offline training set.

Measurement and label collection. For each candidate circuit C_i, we perform a reference measurement with S_{ref} shots and obtain the count vector $\boldsymbol{\nu}_i^{\text{ref}}$. Since the prepared state depends only on C_i, the same measurement record can be paired with different Hamiltonians without additional quantum measurements. For a target Hamiltonian H_j, the corresponding high-shot SQD energy

$$E_{\text{high}}^{(i,j)} = E_{\text{SQD}}(\mathsf{C}_i; H_j, S_{\text{ref}}, K) \tag{17}$$

serves as the supervision target.

To construct the low-shot predictor input, we draw S_{low} samples from the reference measurement distribution, where $S_{\text{low}} \ll S_{\text{ref}}$. Defining $\widehat{p}_i^{\text{ref}}(\boldsymbol{x}) = \nu_{i,\boldsymbol{x}}^{\text{ref}}/S_{\text{ref}}$, the resulting count vector is sampled as

$$\boldsymbol{\nu}_i^{(S_{\text{low}})} \sim \text{Multinomial}\left(S_{\text{low}}, \widehat{\mathbf{p}}_i^{\text{ref}}\right). \tag{18}$$

The corresponding low-shot SQD energy is

$$E_{\text{low}}^{(i,j)} = E_{\text{SQD}}(\mathsf{C}_i; H_j, S_{\text{low}}, K). \tag{19}$$

B.2 Neural circuit performance predictor

The neural circuit performance predictor consists of three branches: circuit branch for processing circuit information, measurement branch for processing the measurement results, and context branch for processing Hamiltonian descriptions together with statistical information about the measurement outcome distribution. The concrete architecture is shown in Fig. 5.

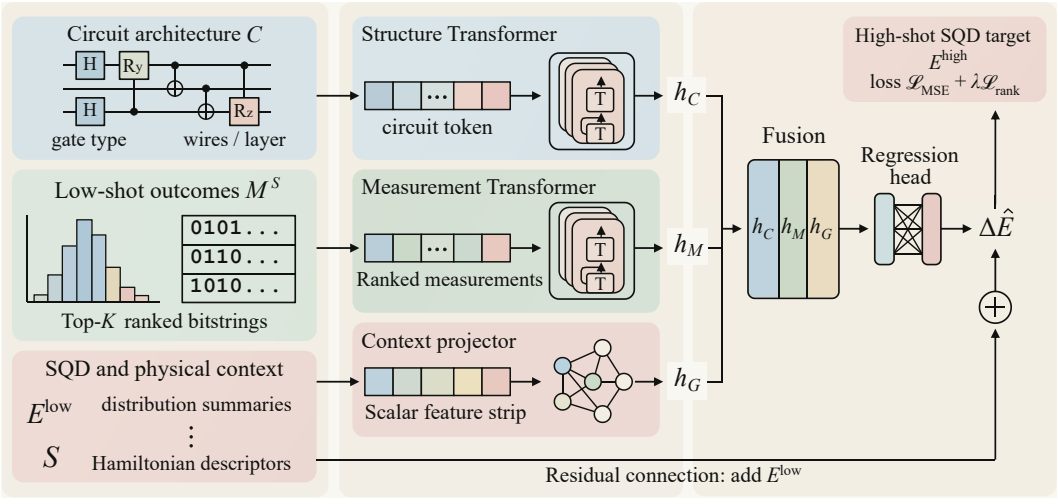

Figure 5: **Architecture of the multimodal residual predictor.**

Circuit branch. We encode the circuit layout in Eqn. (16) as an ordered sequence of discrete tokens. The sequence begins with $\boldsymbol{\tau}_{\text{CLS}} = (\text{CLS}, 0, 0, 0)$. For each layer ℓ, every single-qubit slot contributes $\boldsymbol{\tau}_{\ell,q} = (V_{\ell,q}, q, 0, \ell)$, including slots for which $V_{\ell,q} = I$. Each active directed pair $(q, q') \in \mathcal{E}_\ell$ then contributes $\boldsymbol{\tau}_{\ell,(q,q')} = (\text{CNOT}, q, q', \ell)$. Thus, the four entries of a non-padding token specify its gate identity, first qubit, second qubit, and layer index. Within each layer, the n

single-qubit tokens are ordered by qubit index and followed by the CNOT tokens in the stored pair order. Flattening the layers produces

$$\boldsymbol{\tau}_i^C = \left[\boldsymbol{\tau}_{\mathrm{CLS}}, \left\{\boldsymbol{\tau}_{\ell,1}, \ldots, \boldsymbol{\tau}_{\ell,n}, \left\{\boldsymbol{\tau}_{\ell,(q,q')}\right\}_{(q,q')\in\mathcal{E}_\ell}\right\}_{\ell=1}^{L_i}\right], \tag{20}$$

where $L_i \le L_{\max}$ is the number of layers in circuit C_i. For a token $t = (g_t, q_{1,t}, q_{2,t}, \ell_t)$, its vector representation is

$$\mathbf{e}_t^C = \mathbf{e}_{\mathrm{gate}}(g_t) + \mathbf{e}_{q_1}(q_{1,t}) + \mathbf{e}_{q_2}(q_{2,t}) + \mathbf{e}_{\mathrm{layer}}(\ell_t), \tag{21}$$

where g_t, $q_{1,t}$, $q_{2,t}$, and ℓ_t denote the gate identity, first qubit, second qubit, and layer index, respectively. The four functions $\mathbf{e}_{\mathrm{gate}}$, $\mathbf{e}_{q_1}$, $\mathbf{e}_{q_2}$, and $\mathbf{e}_{\mathrm{layer}}$ are trainable lookup embeddings that map these discrete identifiers to the common d-dimensional latent space. The second-qubit index is set to zero for a single-qubit operation. Applying Eqn. (21) to Eqn. (20) and appending $\boldsymbol{\tau}_{\mathrm{PAD}} = (\mathrm{PAD}, 0, 0, 0)$ gives the fixed-length embedded sequence $\mathbf{Z}_i^C = [\mathbf{e}_{\mathrm{CLS}}^C, \mathbf{e}_1^C, \ldots, \mathbf{e}_{T_i-1}^C, \mathbf{e}_{\mathrm{PAD}}^C, \ldots, \mathbf{e}_{\mathrm{PAD}}^C] \in \mathbb{R}^{T_{\max}\times d}$, where T_i is the unpadded length and $T_{\max} = 1 + 2nL_{\max}$. A binary attention mask $\mathbf{a}_i^C \in \{0,1\}^{T_{\max}}$ identifies the non-padding positions. A pre-norm Transformer encoder with GELU feed-forward blocks then computes $\mathbf{H}_i^C = \mathrm{Transformer}_C(\mathbf{Z}_i^C; \mathbf{a}_i^C)$, and the normalized output at the CLS position defines $\mathbf{h}_C = \mathrm{LayerNorm}(\mathbf{H}_{i,0}^C)$. Continuous gate angles are retained in the source record but are not separately embedded by this branch; their influence remains observable through the circuit's measurement distribution.

Measurement branch. Let K_{meas} be the maximum number of measurement outcomes provided to the predictor, and let S be the number of shots currently available for the circuit. The observed bitstrings are sorted by decreasing count. For the outcome of zero-based rank r, $\boldsymbol{x}_r = (x_{r,1}, \ldots, x_{r,n}) \in \{0,1\}^n$ denotes its bitstring and ν_r denotes its observed count, with $\sum_r \nu_r = S$. Its input row is

$$\mathbf{m}_r = \left[x_{r,1}, \ldots, x_{r,n}, \frac{\nu_r}{S}, \frac{\log(1+\nu_r)}{\log(1+S)}, \frac{r}{\max(1, K_{\mathrm{meas}}-1)}, \frac{\|\boldsymbol{x}_r\|_1}{n}\right]. \tag{22}$$

The four quantities appended to the n binary entries are the empirical frequency, normalized log-count, relative frequency rank, and normalized Hamming weight. We use $K_{\mathrm{meas}} = K$, aligning the maximum measurement-sequence length with the SQD subspace cap. If fewer than K_{meas} unique outcomes are observed, the unused rows are zero-padded and masked.

Each row is encoded as $\mathbf{z}_r^M = \mathbf{W}_M\mathbf{m}_r + \mathbf{b}_M + \mathbf{e}_{\mathrm{rank}}(r+1)$, where $\mathbf{W}_M$ and $\mathbf{b}_M$ are trainable projection parameters and $\mathbf{e}_{\mathrm{rank}}$ is a learned rank embedding. A trainable CLS vector is prepended to the encoded rows, and the resulting sequence is processed by a second pre-norm Transformer encoder. The normalized output at the CLS position defines the measurement representation $\mathbf{h}_M$. The attention mask ensures that unobserved outcomes and padded rows do not contribute to this representation.

Context branch and multimodal fusion. Let $\widehat{p}(\boldsymbol{x}) = \nu_{\boldsymbol{x}}/S$ denote the empirical probability of observed bitstring $\boldsymbol{x}$, where $\nu_{\boldsymbol{x}}$ is its count, and let E_{low} be the SQD energy computed from the same S-shot record. The context branch receives the eight-dimensional vector

$$\mathbf{g} = \left[E_{\mathrm{low}}, \mathrm{Ent}(\widehat{\mathbf{p}}), \frac{N_{\mathrm{eff}}}{S}, w_{\lfloor n/2\rfloor}, \log(1+S), \alpha_1, \alpha_2, \frac{\alpha_2}{\alpha_1}\right], \tag{23}$$

where $\mathrm{Ent}(\widehat{\mathbf{p}}) = -\sum_{\boldsymbol{x}} \widehat{p}(\boldsymbol{x}) \log \widehat{p}(\boldsymbol{x})$ is the Shannon entropy, $N_{\mathrm{eff}} = 1/\sum_{\boldsymbol{x}} \widehat{p}(\boldsymbol{x})^2$, and $w_{\lfloor n/2\rfloor} = S^{-1}\sum_{\|\boldsymbol{x}\|_1=\lfloor n/2\rfloor} \nu_{\boldsymbol{x}}$ is the fraction of counts in the half-filled Hamming-weight sector. Here, (α_1, α_2) are the Hamiltonian couplings entering the context vector: for TFIM, $(\alpha_1, \alpha_2) = (J, h)$; for XXZ, $(\alpha_1, \alpha_2) = (J_{xy}, J_z) = (J, J\Delta)$, which is in one-to-one correspondence with $\boldsymbol{\alpha} = (J, \Delta)$ in the main text. The coupling ratio is set to zero when $\alpha_1 = 0$. A trainable affine projection followed by GELU activation and layer normalization maps $\mathbf{g}$ to the context representation $\mathbf{h}_G$.

The three modality vectors are concatenated and passed through a multilayer perceptron g,

$$f(\mathsf{C}, H, \mathcal{S}; \boldsymbol{\phi}) = g(\mathbf{h}_C \,\|\, \mathbf{h}_M \,\|\, \mathbf{h}_G), \tag{24}$$

whose scalar output is the correction term added to E_{low} in Sec. 3.2, and $\boldsymbol{\phi}$ collects the trainable parameters of the three branches and of g. For the experiments reported in the main text, all three

Algorithm 1 Predictor-guided quantum architecture search

Require: Hamiltonian H, frozen predictor, candidate circuit family,
stage schedule $\{(S_r, N_r)\}_{r=0}^{R}$, number of updates G,
ranking dimension K_{rank}, selection sizes $N_{\mathrm{parent}}, N_{\mathrm{sel}}$
Ensure: Selected circuits and independently evaluated SQD energies

```
Sample an initial population of N_0 circuits; initialize an empty archive
Set S_{-1} = 0
for g = 0, ..., G do
    Set all population members active and reset their histograms ν_C
    for r = 0, ..., R do
        Collect S_r − S_{r−1} additional shots for each active circuit and update ν_C
        Compute low-shot SQD energies at K_rank from the accumulated measurements
        Predict high-shot SQD energies from the circuits, H, and low-shot observations
        Rank active circuits by increasing predicted energy
        if r < R then
            Retain the best N_{r+1} circuits for the next measurement stage
        end if
    end for
    Select the best N_parent active circuits as parents
    Update the archive with these circuits and their scores
    if g < G then
        Generate offspring from the parents
        Form the next population from elites, offspring, and newly sampled circuits
    end if
end for
Select the N_sel distinct circuits with the lowest archived scores
Evaluate the selected circuits using independent measurements
return selected circuits and their evaluated SQD energies
```

representations have dimension 128. The circuit and measurement Transformers use four attention heads and respectively four and two encoder layers; their feed-forward dimension is 256 and dropout is 0.1. The fusion head maps the concatenated 384-dimensional vector through hidden dimensions 256 and 128 before producing the scalar residual. Because the qubit and layer embeddings depend on n and $L_{\max}$, a separate predictor checkpoint is trained for each system size. Each checkpoint stores n, $L_{\max}$, K_{meas}, and the network hyperparameters; inference uses the same fixed gate vocabulary as its training data pipeline.

Training configuration. We optimize the objective in Eqn. (9) using AdamW with batch size 64, learning rate 10^{-3}, and weight decay 10^{-4}. The ranking weight is $\lambda_{\mathrm{rank}} = 0.1$ and the margin is $m = 0.05$. Ranking pairs are formed within the same Hamiltonian. Tied labels are omitted, and the ranking loss is zero if no eligible pair exists. Training batches are shuffled and gradient norms are clipped to 1. Training lasts at most 200 epochs, with early stopping after 20 epochs without improvement in validation MSE. The lowest-MSE validation checkpoint is retained. Epoch losses are averaged over samples, and test data are not used for model selection. These settings apply to the primary problem-agnostic experiments; the scaling and molecular studies specify their own training settings.

B.3 Predictor-guided QAS for circuit design

The trained predictor serves as an effective circuit performance indicator to guide the QAS to find circuits that align well with the target Hamiltonian. The detailed procedure is implemented in Algorithm 1. Raw low-shot QAS follows the same procedure but ranks candidates by their measured SQD energies. All methods share the initial population within each search seed. In Algorithm 1, N_r denotes the number of active circuits at measurement stage r, and K_{rank} denotes the retained-subspace dimension used for the low-shot SQD energies during ranking, which equals K unless otherwise stated.

Hyperparameters of QAS. Table 4 lists the settings used for the main 20-qubit adaptive-search experiments. The reported 20 evolutionary updates are preceded by the evaluation of the generation-zero population. With a reduction factor of two, the five fidelity levels evaluate respectively 64,

Table 4: Implementation settings for the main adaptive QAS experiments.

Setting	Value
Population size N_0	64
Evolutionary updates	20
Cumulative shot levels	$500, 1000, 5000, 10000, 20000$
Reduction factor η	2
Generation finalists (parents) N_{parent}	5
Per-generation shot-budget cap	1.28×10^6
Elite / tournament size	4 / 3
Crossover probability	0.7
Architecture mutation probability	0.35
Per-angle mutation probability	0.2
Final archived circuits N_{sel}	5
Final shots per selected circuit S_{ref}	10^6
SQD subspace cap K	2048

32, 16, 8, and 5 active candidates before returning five generation finalists. Unless an experiment explicitly studies a different schedule, these settings are shared by the TFIM and XXZ searches.

C More experimental results

C.1 More detailed settings of the experiments in the main text

Dataset construction. The standard construction uses 2,000 source circuits per model. Circuit identities are split $60\%/15\%/25\%$ before expansion over ten Hamiltonians or low-shot resampling. This gives 1,200, 300, and 500 source circuits, or 12,000, 3,000, and 5,000 circuit–Hamiltonian records. All variants of a source circuit remain in the same partition.

Evaluation and statistical scope. When evaluating whether the learned predictor is reliable, we vary the measurement budget after the circuits have been fixed. In particular, we use five circuits per method and compute the mean and sample standard deviation of their SQD errors at each shot level. When comparing the overall performance of our method, we instead evaluate circuits promoted at each low-shot stage using their high-shot reference errors, averages within each generation, and then summarize 21 generation means.

Hamiltonian splits and generalization settings. For TFIM, the predictor is trained on ten Hamiltonians,

$$\mathcal{H}_{\mathrm{train}} = \{H(J,h) : (J,h) \in \{(0.4,0.8),\ (0.6,2.0),\ (0.8,1.0),\ (1.0,2.5),\ (1.2,0.6),\ (1.4,1.8),\ (1.6,2.8),\ (1.8,1.2),\ (2.0,2.2),\ (2.1,0.9)\}\}. \tag{25}$$

These Hamiltonians span $J \in [0.4, 2.1]$ and $h \in [0.6, 2.8]$. The training, validation, and test sets described in Sec. 4.1 are split by *circuit identity* while sharing the same $\mathcal{H}_{\mathrm{train}}$. Therefore, Fig. 2 evaluates generalization to unseen circuits rather than unseen Hamiltonians.

Generalization to unseen Hamiltonians is evaluated separately in Figs. 4(a) and (b) using nine TFIM instances,

$$\mathcal{H}_{\mathrm{test}} = \{H(J,h) : J \in \{0.5, 1, 2\},\ h \in \{0.5, 1.5, 3\}\}\,, \tag{26}$$

all of which are excluded from predictor training, such that $\mathcal{H}_{\mathrm{train}} \cap \mathcal{H}_{\mathrm{test}} = \varnothing$. The test Hamiltonians cover both interpolation within and moderate extrapolation beyond the parameter values represented by $\mathcal{H}_{\mathrm{train}}$: most test points lie within the training domain, while $h = 0.5$ and $h = 3$ slightly extend beyond the minimum and maximum training values, respectively.

The 20-qubit predictor-guided search experiments in Fig. 3, Fig. 4(c), and Table 1 similarly use target Hamiltonians excluded from predictor training. In particular, the TFIM target $(J,h) = (0.8, 0.4)$ does not belong to $\mathcal{H}_{\mathrm{train}}$. For XXZ, the target $(J,\Delta) = (0.125, 2.8)$ is likewise excluded from the Hamiltonians used for predictor training.

For the system-size scaling, VQE comparison, noise, and molecular experiments (Table 2, Table 3, and Appendix C), we instead train a separate predictor for each target Hamiltonian. These exper-

iments therefore evaluate generalization across candidate circuits and the resulting search performance, rather than generalization across Hamiltonians.

C.2 Physically motivated circuits for the scaling experiments

The scaling study uses the physically motivated circuit pool, distinct from the problem-agnostic pool used in the primary XXZ experiments. We specify this family and its experimental protocol below, then evaluate the selected circuits at different retained-subspace dimensions.

C.2.1 Experimental setup

Candidate circuit architecture. We consider the XXZ model introduced in the main text with open boundaries, $J = 0.0625$, and $\Delta = 5.6$ for all system sizes $n \in \{10, 20, 30, 40\}$. Starting from $|0\rangle^{\otimes n}$, we apply X gates to either the odd-indexed or even-indexed qubits, preparing $|1010\cdots\rangle$ or $|0101\cdots\rangle$, respectively. Each circuit contains $L \in \{2, 4, 6, 8\}$ blocks. Each block consists of two layers of nearest-neighbor exchange gates, acting on pairs $(1, 2), (3, 4), \ldots$ and $(2, 3), (4, 5), \ldots$, respectively. The two layers may appear in either order, and individual gate positions may be omitted. Each enabled position applies

$$U_{\mathrm{ex}}(\theta) = \exp[-i\theta(XX + YY)/2]. \tag{27}$$

Every block except the last is followed by single-qubit $R_z(\theta)$ rotations on all qubits. The final block is followed directly by computational-basis measurement. Both gate types preserve the number of ones in the computational basis, so every measured configuration contains exactly $n/2$ ones.

Dataset generation. Each system size has 24 independently seeded circuit families, split into 16 training, four validation, and four test families. A family begins with the 64-circuit scaling initialization in Appendix B.3. In each of eight offline mutation rounds, the 16 lowest-energy circuits accumulated so far form a parent set; eight uniformly sampled parents produce eight children. This offline selection uses high-shot labels. Duplicate circuits are skipped, giving 2,047 training circuits for 10 qubits and 2,048 for each other size, with 512 validation and 512 test circuits per size. Each circuit receives an independent 10^6-shot label at $K = 1024$ and nested low-shot observations at $\{10, 20, 100, 200, 400\}$ shots. Entire families remain within one partition, and evaluation search seeds are separate.

Numerical evaluation. We simulate all QAS experiments using the matrix product state (MPS) backend implemented with PastaQ (Torlai & Fishman). Offline predictor datasets are generated using exact state vectors in the half-filling sector for 20 qubits and MPS simulations for the remaining system sizes. Ground-state reference energies are obtained by exact diagonalization for 10 and 20 qubits and by density matrix renormalization group (DMRG) calculations for 30 and 40 qubits. The corresponding energies are approximately -3.273, -6.882, -10.493, and -14.104, respectively. Tightening the DMRG settings changes the reference energies by less than 10^{-10}. Circuit simulations use an MPS truncation cutoff of 10^{-14} and a maximum bond dimension of 256. Repeating the simulations for all finalist circuits from the development and evaluation runs with a cutoff of 10^{-15} and a maximum bond dimension of 512 yields infidelities below 3×10^{-13} relative to the original simulations. Neither reference-state amplitudes nor unobserved configurations are provided to predictor training or QAS.

The measurements used for final evaluation are independent of those used for online circuit selection. Each search run returns a finalist set $\mathcal{F}$ containing the five archived circuits with the lowest stored ranking scores. Each circuit's score is fixed when it first enters the archive. We evaluate every finalist in three independent repetitions, each using $S_{\mathrm{ref}} = 10^6$ shots. Within a repetition, the same samples are reused for all values of K, retaining the most frequently observed configurations. The actual dimension is $\min(K, N_{\mathrm{unique}})$, where N_{unique} is the number of distinct observed configurations. We compute ε for each repetition and average the three repetitions within each search seed. We then report the mean and sample standard deviation across five paired seeds. This final evaluation requires 1.5×10^7 additional shots per run, shared across the subspace sweep. Independent diagnostics evaluate the five current finalists after each of the 11 ranking rounds using 10^6 shots per circuit. These diagnostics require another 5.5×10^7 shots per run and do not influence search decisions. Final evaluation and diagnostic measurements are accounted for separately from the online search budget.

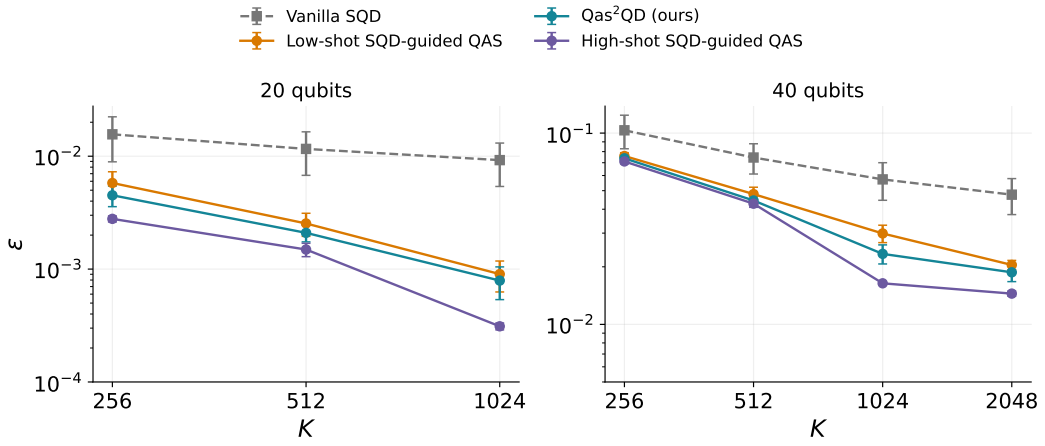

Figure 6: Mean absolute high-shot SQD error ε versus retained-subspace dimension K for 20- and 40-qubit XXZ systems. QAS finalists are selected at $K_{\mathrm{rank}} = 1024$, whereas fixed random circuits are selected without ranking or optimization. All circuits and evaluation samples are held fixed across K. Error bars show sample standard deviations over five paired seeds. Final evaluation budgets are identical across methods. Fixed random-circuit SQD has no search cost, raw and predictor-guided QAS use identical online budgets, and the high-shot reference uses a larger online budget.

C.2.2 Effect of the retained-subspace dimension

We evaluate the circuits from the scaling study at $K \in \{256, 512, 1024\}$ for 20 qubits and $K \in \{256, 512, 1024, 2048\}$ for 40 qubits. The circuits and evaluation samples remain fixed throughout the sweep, so changes in error arise from retaining different numbers of sampled configurations rather than from additional search or measurement resources.

Fig. 6 shows that both QAS methods consistently outperform vanilla SQD across all tested values of K, with Qas²QD further improving over low-shot SQD-guided QAS. For 20 qubits, the mean error of Qas²QD decreases from 4.508×10^{-3} at $K = 256$ to 7.912×10^{-4} at $K = 1024$, corresponding to improvements of 22.2%, 17.4%, and 12.4% over low-shot SQD-guided QAS at the three values of K. For 40 qubits, the error decreases from 7.382×10^{-2} at $K = 256$ to 1.872×10^{-2} at $K = 2048$, with corresponding improvements of 2.7%, 7.4%, 21.8%, and 8.5%. These results demonstrate that the advantage of Qas²QD persists across different retained-subspace dimensions without increasing the measurement budget.

C.3 Finite-shot VQE protocol

Circuit construction. The VQE comparison uses the same open-boundary 40-qubit XXZ Hamiltonian as the scaling study. Both ansatzes start from an alternating computational-basis state, with its orientation determined by seed parity. Each of the four HVA blocks applies exchange gates from Eqn. (27) on odd then even bonds, followed by $\exp(-i\theta ZZ)$ gates on odd then even bonds. A single angle is shared within each of these four groups. Each of the four HEA blocks applies independent R_y and R_z rotations to every qubit, followed by odd- and even-bond CZ layers.

Shot allocation. Both ansatzes are optimized using SPSA with finite-shot energy estimates. The measurement cost of VQE depends on the number of optimization iterations and the number of shots used per iteration. Each SPSA iteration requires two energy estimates, so T iterations with S shots per estimate consume $2ST$ shots. We vary S and adjust T while keeping $2ST = 70{,}400$ to guarantee a fair comparison. Increasing the shots per estimate therefore leaves fewer iterations available for optimization. We reserve another 7,040 shots for checkpoint selection, evaluating the circuit with 640 shots at initialization and after every $T/10$ iterations and retaining the checkpoint with the lowest measured energy.

C.4 EXPERIMENTS ON A MOLECULAR BENCHMARK

We evaluate the sampling efficiency of the proposed framework on N_2 to assess its applicability beyond spin systems. We use the STO-3G basis and an active space of ten electrons in eight spatial orbitals, corresponding to 16 qubits. The bond length is varied over $d \in \{1.1, 1.8, 2.5\}$ Å, covering equilibrium and stretched geometries. Reference ground-state energies are obtained by complete active-space configuration interaction within the same active space.

C.4.1 EXPERIMENT SETUP

Candidate circuit architecture. Each circuit starts from the Hartree–Fock configuration $|\mathrm{HF}\rangle = |11111000\rangle_\uparrow \otimes |11111000\rangle_\downarrow$, prepared by X gates on the first five qubits of each spin block. Each circuit then applies $L \in \{16, 32, 64, 96\}$ fermionic single- and double-excitation rotations, with excitation channels sampled uniformly with replacement from the allowed set. The circuit prepares

$$|\psi_{\mathsf{C}}\rangle = U_{k_L}(\theta_L) \cdots U_{k_1}(\theta_1)|\mathrm{HF}\rangle, \quad U_k(\theta) = \exp\Big[\theta(T_k - T_k^\dagger)\Big], \tag{28}$$

where k_ℓ specifies the excitation at position ℓ. The excitation operator is either $T_k = \hat{a}_p^\dagger \hat{a}_i$ or $T_k = \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_j \hat{a}_i$, where i, j label occupied spin orbitals in the reference configuration and p, q label virtual spin orbitals. The operators $\hat{a}_r^\dagger$ and $\hat{a}_r$ respectively create and annihilate an electron in spin orbital r. Only channels that preserve both spin populations are allowed.

Predictor training. For each geometry, we construct 256 training, 64 validation, and 64 test families. Each contain a random base circuit and three descendants obtained by applying 1, 3, or 8 mutations to that base, using the operators in Appendix B.3. The resulting splits contain 1,024, 256, and 256 circuits. Each has an independent 10^6-shot SQD label at $K = 256$ and nested observations at $S \in \{10, 20, 100, 200, 400\}$. Across the three datasets, offline labels use 4.608×10^9 shots and low-shot inputs use 1.843×10^6 shots.

C.4.2 RESULTS

We evaluate the five circuits selected in each run at $K = 256$, varying the shots S_{eval} from 10^2 to 10^7 per circuit. After circuit selection, we collect three independent sets of 10^7 shots per circuit and use their prefixes for the smaller shot budgets. We compute the energy error ε at each S_{eval} and average it over the three repetitions within each seed, then report the mean and sample standard deviation across multiple seeds.

Fig. 7 shows lower mean errors for both QAS methods than for fixed random-circuit SQD across varying shot budgets for evaluation. At very small evaluation budgets, the advantage over raw low-shot search is limited or absent. As more measurements become available, the quality of the selected sampling distributions becomes more apparent, yielding a larger separation between the methods, especially for stretched N_2. Note that the predictor-guided errors eventually plateau near 10^{-2} Ha due to the limitation of small fixed subspace dimension. These results demonstrate improved sampling efficiency while identifying the remaining accuracy limitations of the present circuit family and subspace setting.

C.5 SQD WITH NOISE

Experiment setup. We consider the effect of quantum gate and readout noise on the performance of our method. Specifically, we use the same 20-qubit XXZ model and candidate circuit pool as in Appendix C.2. After each gate, we apply a depolarizing channel to the qubits on which the gate acts. Mathematically, the channel acting on k qubits is defined by

$$\mathcal{E}_p^{(k)}(\rho) = (1-p)\rho + p\frac{I_{2^k}}{2^k}, \tag{29}$$

where ρ is a k-qubit density matrix. We set $p = p_2 \in \{0, 10^{-3}, 3 \times 10^{-3}, 10^{-2}\}$ for two-qubit gates and $p = p_1 = p_2/10$ for single-qubit gates. To model readout errors, each measured bit independently flips with probability p_2.

Results. We retrain one predictor jointly across all four noise levels, preserving the circuit splits used in the scaling experiment, and freeze it before QAS. Table 5 shows that our method achieves lower

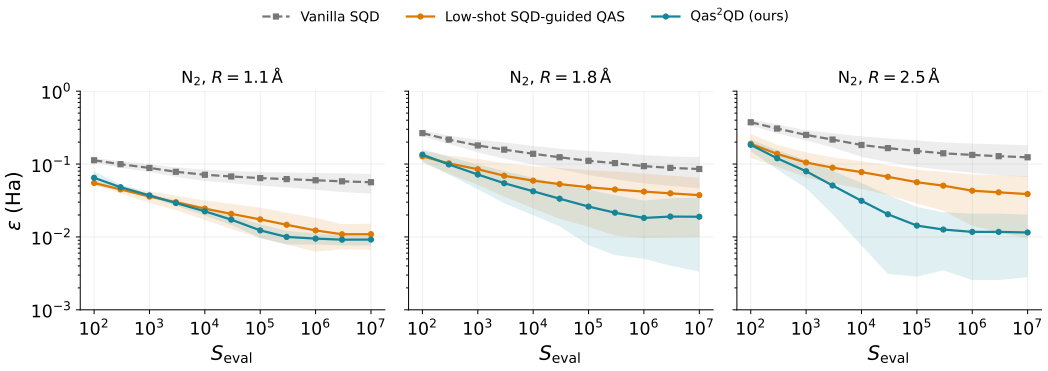

Figure 7: Sampling efficiency for N_2 at three bond lengths with $K = 256$. The circuits selected by each method are held fixed while the evaluation shots per circuit are varied. Both QAS methods use 10,560 online search shots per run, whereas fixed random-circuit SQD performs no search. Lines and shaded bands indicate means and one standard deviation over ten paired seeds.

Table 5: Mean SQD error $\varepsilon \pm$ sample standard deviation over five paired seeds on the noisy 20-qubit XXZ model. All entries are in units of 10^{-3}. Final-evaluation budgets are matched; our method additionally uses 77,440 search shots per run.

p_2	Vanilla SQD	Ours
0	8.168 ± 2.981	$\mathbf{1.326} \pm 0.390$
10^{-3}	3.845 ± 1.278	$\mathbf{1.388} \pm 0.319$
3×10^{-3}	2.921 ± 0.842	$\mathbf{1.579} \pm 0.115$
10^{-2}	3.429 ± 0.917	$\mathbf{2.082} \pm 0.238$

mean errors than vanilla SQD at every tested noise level. At nonzero noise rates, the reductions range from 39.3% to 63.9%. Although its error increases with noise, our method retains a mean error of approximately 2.1×10^{-3} at $p_2 = 10^{-2}$. These results concern the predictor retrained with noisy data, rather than direct transfer of the clean-trained predictor.

D LIMITATIONS

Our experiments use prescribed subspace dimensions K and circuit families $\mathcal{C}$ constrained by the available classical simulation resources. Larger subspaces may improve accuracy but increase the cost of classical diagonalization. A natural extension is to jointly adapt the sampling circuits, shot allocation, and subspace dimension, balancing quantum measurement and classical computational costs to achieve a target accuracy.

Another direction is to extend learned circuit evaluation to quantum Krylov diagonalization. Its quantum-generated basis states may capture low-energy structure that would require many computational-basis configurations in SQD, potentially broadening the applicability of learned circuit design. This extension would require input representations and supervision targets suited to the measurement of Hamiltonian and overlap matrices, together with an evaluation strategy that accounts for the sensitivity of the resulting generalized eigenproblem to finite-shot errors.