

QuantumSEA: In-Time Sparse Exploration for Noise Adaptive Quantum Circuits

Tianlong Chen¹, Zhenyu Zhang¹, Hanrui Wang², Jiaqi Gu^{1,5}, Zirui Li³,
 David Z. Pan¹, Frederic T. Chong⁴, Song Han², Zhangyang Wang¹
¹University of Texas at Austin ²Massachusetts Institute of Technology
³Rutgers University ⁴University of Chicago ⁵Arizona State University

Abstract—Parameterized Quantum Circuits (PQC) have obtained increasing popularity thanks to their great potential for near-term Noisy Intermediate-Scale Quantum (NISQ) computers. Achieving quantum advantages usually requires a large number of qubits and quantum circuits with enough capacity. However, limited coherence time and massive quantum noises severely constrain the size of quantum circuits that can be executed reliably on real machines. To address these two pain points, we propose QuantumSEA, an in-time sparse exploration for noise-adaptive quantum circuits, aiming to achieve two key objectives: (1) *implicit circuits capacity during training* - by dynamically exploring the circuit’s sparse connectivity and sticking a fixed *small* number of quantum gates throughout the training which satisfies the coherence time and enjoy light noises, enabling feasible executions on real quantum devices; (2) *noise robustness* - by jointly optimizing the topology and parameters of quantum circuits under real device noise models. In each update step of sparsity, we leverage the moving average of historical gradients to grow necessary gates and utilize salience-based pruning to eliminate insignificant gates. Extensive experiments are conducted with 7 Quantum Machine Learning (QML) and Variational Quantum Eigensolver (VQE) benchmarks on 6 simulated or real quantum computers, where QuantumSEA consistently surpasses noise-aware search, human-designed, and randomly generated quantum circuit baselines by a clear performance margin. For example, even in the most challenging on-chip training regime, our method establishes state-of-the-art results with only half the number of quantum gates and $\sim 2\times$ time saving of circuit executions. Codes are available at <https://github.com/VITA-Group/QuantumSEA>.

Index Terms—Quantum Machine Learning, Variational Quantum Eigensolver, Quantum Circuits

I. INTRODUCTION

As a rising computational paradigm, Quantum computing (QC) demonstrates great potential in solving traditionally intractable problems with exponential acceleration over classical computers. Numerous domains have witnessed rapid progress of QC, including cryptography [50], database search [20], chemistry [7], [28], [47], and machine learning [4], [18], [22], [35], [48]. Among these, parameterized Quantum Circuits (PQCs) [1] on the noisy intermediate-scale quantum (NISQ) devices [3] is one of the most promising near-term applications showing impressive benefits from QC.

To reach PQC’s quantum advantage, it demands a large number of qubits and quantum circuits with sufficient capacity [60]. Most of the existing works [22], [38] that perform both training and inference on classical computers via software simulations, can not satisfy these conditions, since

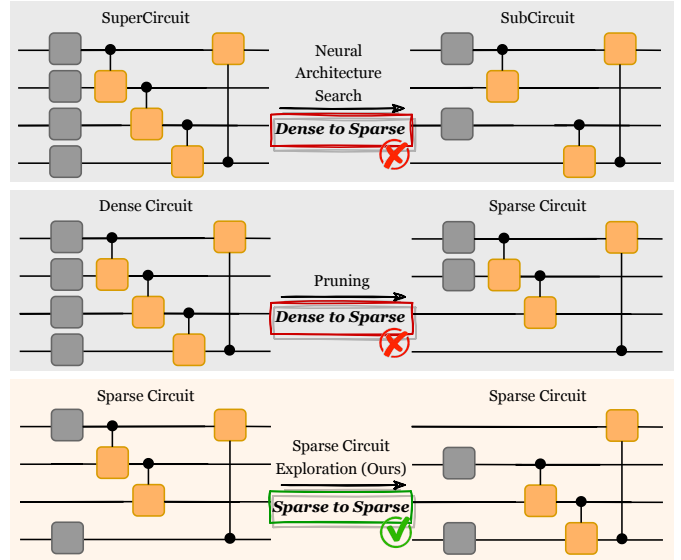


Fig. 1. Overview of method comparisons. Both neural architecture search and pruning maintain dense circuits during training and produce their sparse variants in the end. And their training suffers from massive decoherence and operation noises due to excess quantum gates. Our proposal enables the circuit design from sparse to sparse, which only deals with a small portion of quantum gates at each training iteration.

they have poor scalability and are failed to deal with exponentially boosted time and memory costs ($\mathcal{O}(2^{\#\text{qubits}})$) [60]. Unfortunately, two challenging impediments arise when we consider the on-chip training or evaluation on real quantum machines. Firstly, the unwanted interactions between qubits and the environment or other qubits cause qubit decoherence. The *coherence time* statistically indicates how long a qubit retains its information, i.e., a qubit’s “lifetime”. Therefore, all operations are required to be completed before significant decoherence happens. Otherwise, the decoherence will cause severe infidelity of quantum-information processing [62], and stultify the computing results. Therefore, the *depth* of PQC should be small enough to alleviate the decoherence impact. Meanwhile, the non-negligible *quantum errors* in the current NISQ era also make operations on qubits highly unreliable. For instance, the quantum gates commonly suffer from a high error rate of 10^{-3} to 10^{-2} [58], [60]. In the context of PQC on real QC devices, such noises will result in unstable training

and also significantly degraded performance [60] compared to the ideal noise-free simulation. Thus, to bridge quantum algorithms to realistic QC hardware, *robust circuits topology* against quantum noises are on the urgent call.

To address these obstacles, we propose QuantumSEA, an in-time **S**parse **E**xploration for noise-**A**daptive robust quantum circuits. On one hand, our proposal starts with a randomly sparsified circuit. It dynamically explores its connectivity across the training time, while sticking to a fixed small number of quantum gates. The reduced quantum operations effectively decrease noises and alleviate the negative effects of decoherence. It is also one of the key benefits of our proposal compared to the conventional circuit designs from neural architecture search or pruning, as demonstrated in Figure 1. Meantime, the dynamic topology exploration brings extra implicit circuit capacity in the space-time manifold, achieving significant improvements in the expressibility of sparse quantum circuits. On the other hand, QuantumSEA develops the robust circuit by *simultaneously* optimizing its sparse connectivities and parameters under real QC noises. Moreover, identified topologies are capable of serving an implicit regularization and promoting noise resistance, as widely demonstrated in classical machine learning [14], [21], [66]. In other words, QuantumSEA kills two “birds” (short coherent time and large quantum noises) with one “stone” (sparsity), i.e., the located noise-adaptive sparse circuit topology enables robust & efficient inference and training of PQCs on real QC machines with substantially improved performance and reduced execution time, as presented in Figure 2 (*Right*). Our innovations are unfolded along with three thrusts below:

- ★ **In-time sparse explorations** grant the quantum circuit extra implicit capacity, which significantly improves its expressiveness without introducing additional gates. Compared to the dense counterpart, our proposal is capable of reducing **over** $2\times$ execution time (Figure 2 (*Right*)), lowering the risk of severe decoherence, and improving the fidelity of predictions.
- ★ **Noise adaptive and jointly optimized topology & weights** lead to a superior noise-resilient quantum circuit, outperforming noise-aware searched, random generated, and human-designed quantum circuit variants, as demonstrated in Figure 2 (*Left and Middle*).
- ★ **Extensive experiments conducted on real quantum devices** consistently demonstrate the effectiveness of our proposal. In total, 7 representative benchmarks in QML and VQE on 6 real quantum machines are evaluated, where we establish a new start-of-the-art on-chip training performance with significant running time savings based on the recent advance [60]. Take MNIST-2 as an example. We achieve $1.33\% \sim 5.32\%$ accuracy gains and $\sim 50\%$ execution time reduction.

II. RELATED WORKS

a) Quantum machine learning (QML) and variational quantum eigensolver (VQE): QML [4], [34], [36], [37] studies

machine learning algorithms on quantum computers, which shows an impressive ability in generating and estimating highly complex kernels [23]. Recently, a group of QML models, QNN [1], emerges which parameterizes quantum circuits with trainable weights. Although diverse theoretical formulations of QNNs have been established, most of their implementations are limited to small-scale simulations of quantum systems via classical computing [19]. VQE [41], [47], [56] is another representative example of NISQ algorithms, whose goal is to estimate the Hamiltonian’s ground-state energy, serving as the first step towards investigating the energetic properties of molecules [12], [25], [63] and materials [9], [57]. VQE leverages the superposition principle [51] of quantum mechanics to encode exponentially increased information with only a linearly grown number of qubits.

b) Quantum circuit connectivity design with QML:

Both [24] and [58] adopt neural architecture search approaches to automate the connectivity design of quantum circuits. The former [24] introduces a space pruning algorithm to trim down the search space during architecture exploration by progressively removing unpromising candidate gates; the later [58] further uses magnitude-based pruning methods to eliminate insignificant quantum gates of searched quantum circuits. More possibilities for the circuit design via advanced search algorithms have been extensively discussed [5], [16], [29], [46], [67]. It is worth mentioning that our QuantumSEA is fundamentally different from neural architecture based approaches. Since we always optimize a highly reduced number of quantum gates, QuantumSEA (1) avoids the massive operation noises from additional quantum gates (2) and makes quantum circuits more feasible for on-chip training given the limitation of qubit’s coherence time.

To tackle the challenging optimization in the vast parameter space of large quantum circuits, [53] proposes a parameter-efficient circuit training framework by leveraging the adaptive pruning techniques from classical machine learning. In each iteration, it activates and optimizes a selected subset of circuit parameters, which dynamically alters the circuit connectivity. But they adopt a random strategy to update the circuit structure, resulting in a sub-optimal optimization process. [45] also evolves the circuit structure together with the optimization of circuit parameters by simultaneously selecting the gate and finding the direction & angle of rotation that yield minimum object values, in which they maintain all the gates during training, under the risk of more gate errors. Moreover, these methods only demonstrate the effectiveness of their proposal in small-scale cases like parameterized single-qubit gates and fixed two-qubit gates.

c) Quantum noise and its mitigation: Quantum noises are widely existing in real quantum computer devices, resulting from unwanted interactions between qubits, imperfect control signals, or interference from the environment [26]. And quantum noise mitigation has been extensively developed recently. Extrapolation methods [31], [55] extrapolate the ideal noise-free measurement from multiple measurements of a quantum circuit with different noise rates. Quasiprobabil-




Fig. 2. (Left and Middle) Real QC testing accuracy (%) over the number of parameters in quantum circuits. Evaluations are conducted on IBMQ-Santiago with RXYZ and MNIST-2. The accuracy of conventionally designed circuits is quickly saturated and then degraded. Our QuantumSEA outperforms other approaches in quantum noise mitigation, allowing larger circuit capacity and superior accuracy. (Right) Real QC testing accuracy (%) over the training time of quantum circuits, *i.e.*, RXYZ) on IBMQ-Santiago.

ity [27], [55], [59] inserts X, Y, and Z quantum gates to the circuit in probability and then calculates the summation to cancel out the noise effects. Quantum subspace expansion [40], learning-based mitigation [11], [54], [61], post-measurement normalization and quantization [59] are invented to address the noise issue. In our work, we proposed QuantumSEA to neutralize noise impacts from a fundamentally different perspective: high-quality sparse typologies of quantum circuits.

III. METHODOLOGY

A. Preliminary

a) Qubits: A quantum bit $|\psi\rangle$ is defined as a linear combination of two basis states $|0\rangle$ and $|1\rangle$: $|\psi\rangle = \alpha \times |0\rangle + \beta \times |1\rangle$, where $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$. Leveraging the superposition principles [51], an n -qubits quantum system is capable of representing a linear combination of 2^n basis states, while the classical register requires n -bits to store only one of the 2^n states.

b) Quantum circuits: Figure 3 (bottom) presents an example QNN quantum circuit for QML. It first encodes the raw input through rotation gates like \mathcal{R}_X , \mathcal{R}_Y , and \mathcal{R}_Z . To further manipulate the information in quantum systems, we execute several sequences of (trainable) quantum gates on qubits for desired computations. Specifically, the quantum gates convert one quantum state to another via unitary transformations, *i.e.*, $|\psi\rangle \rightarrow \mathcal{U}|\psi\rangle$, where \mathcal{U} is a unitary matrix and can also be parameterized by θ for adaptive optimizations. Lastly, a qubit readout operation, *i.e.*, measurement, is applied to obtain the prediction from quantum circuits. It collapses a qubit $|\psi\rangle$ from its superposition to either $|0\rangle$ or $|1\rangle$ in probability based on their amplitudes α or β , respectively. Given a task, designing the best, realistic, noise-robust quantum circuits is highly non-trivial [15], [58], since it needs to express the required transformation (*accurate*), complete the whole execution within an efficient number of steps (*coherent time*), and achieve satisfactory fidelity on the real quantum hardware (*noise adaptive*).

c) Variational quantum circuits: If the quantum gates are parameterized by trainable variables, *i.e.*, angles in the

quantum rotation gate, this kind of quantum circuit is named the variational circuit. Such parameterized quantum circuits $\Phi(x, \theta)$ usually serve as one of the efficient and robust implementations for diverse quantum applications, which is also our focus in this paper. $\Phi(x, \theta)$ can be adopted to prepare a variational quantum state (*i.e.*, ansatz) as $|\psi(x, \theta)\rangle = \Phi(x, \theta)|0, \dots, 0\rangle$, where x is the input samples of certain quantum application and θ is the trainable variable. It has demonstrated a great potential in QML [2], [4], [32], [49], [64], numerical analysis [34], [36], quantum simulation [28], [41], [47], and optimizations [18], [44].

The trainability of variational quantum circuits is one of the key bottlenecks. As presented in [39], a variational circuit might be untrainable with gradient-based approaches due to the excessive flat loss surface (*i.e.*, “barren plateau”). Meanwhile, the training quality of these circuits directly determines whether the variational quantum algorithm is successful or not. A classic training pipeline of variational quantum circuits contains two parts: (i) selecting a handcrafted or searched circuit structure given a computational quantum application; (ii) performing the hybrid quantum-classical optimization to update an optimal set of circuit parameters.

The parameterized quantum circuit (PQC) is a popular application of variational quantum circuits, which is broadly used for VQE and QML tasks. Specifically, for VQE tasks, PQC is adopted for preparing states and is measured based on the given molecule. The state preparation and measurement will be executed multiple times on different qubits and bases. An expectation value is calculated as the final result for the estimation of the molecule’s ground state energy, as shown in Figure 4. As for QML tasks like image recognition, it first encodes the raw pixel with rotation gates (*e.g.*, \mathcal{R}_X , \mathcal{R}_Y , and \mathcal{R}_Z); and then applies parameterized quantum gates to transform the information towards minimum object values; finally, measures the qubits on Z-basis and calculates the according probability for each class with a `softmax` function. An illustration of the circuit structure is provided in Figure 3.

To optimize PQCs, the standard way in a simulator is backpropagation, which computes the derivative $\frac{\partial \mathcal{L}}{\partial \theta_i}$ of each

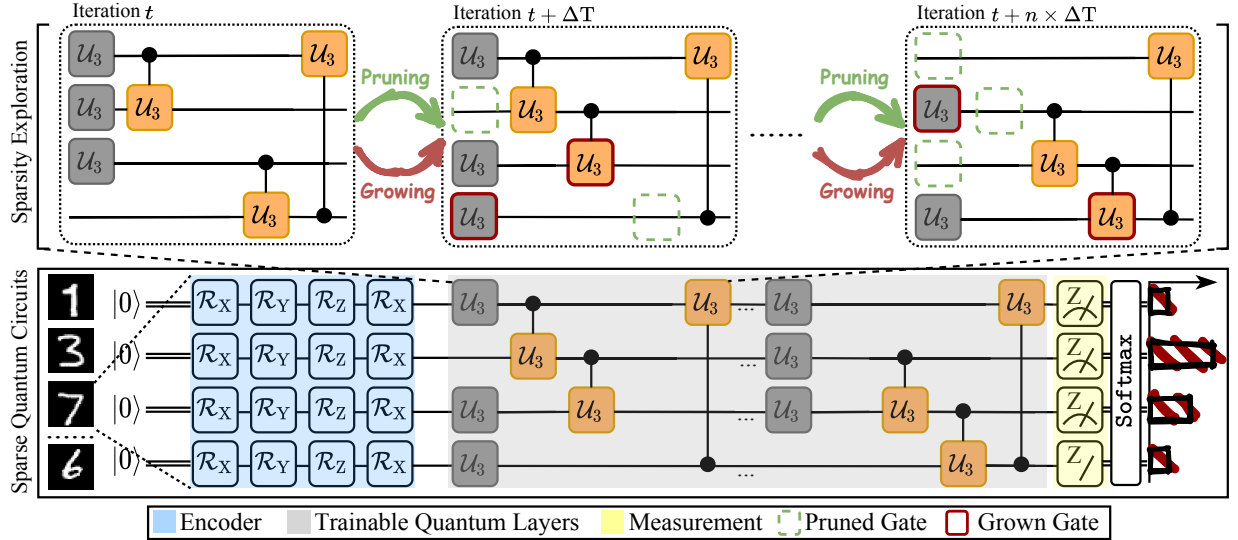


Fig. 3. The overview of QuantumSEA. The *upper* figure demonstrates the in-time sparse exploration of quantum circuits. Specifically, it first trains sparse circuit for ΔT iterations, then leverages prune-and-grow strategies to explore the crucial sparse topology, repeating until convergence. Both weight and topology updates are aware of real QC noises. The *bottom* figure presents an example sparse circuit for QML tasks, which consists of **data encoder**, **trainable quantum layers**, and **the measurement layer**.




Fig. 4. Example quantum circuits for VQE tasks.

parameter θ_i respect to the objective function \mathcal{L} , and then update the according parameter as $\theta_i = \theta_i - \eta \times \frac{\partial \mathcal{L}}{\partial \theta_i}$ where η is the learning rate. However, this approach is not feasible if we consider the real-noise model or on-chip training due to their non-differentiable nature. To enable the training of quantum circuits in these challenging scenarios, the parameter shift rule is utilized to calculate quantum gradients [60], which is followed by most of our experiments.

B. QuantumSEA: In-Time Sparse Exploration for Noise-Resistant Quantum Circuits

a) *QuantumSEA overview*: QuantumSEA produces noise-adaptive quantum circuits by conducting in-time sparse explorations. As demonstrated in Figure 3, it contains three different phases: ① QuantumSEA starts with a randomly sparsified circuit. Each parameter in quantum operations has a probability of being removed or maintained based on certain probability distributions. ② The sparse circuit is then optimized for several iterations ΔT , i.e., updating parameters as normal PQC's training. ③ Sparse topologies are changed straight after the weight optimization. Specifically, it first eliminates a portion of quantum gate weights according to the pre-defined pruning criterion and inserts new gates based on the grow indicators. Note that the gate changes from both pruning and growing happen in the logical quantum circuits which will be further compiled before the deployment.

Thus, the time cost of gate changes is irrelevant to qubits' coherent time which is an ideal constraint of execution time for physical quantum circuits. Lastly, the upgraded sparse quantum circuits are continually trained until the next topology update. A detailed procedure for this in-time exploration is also described in Algorithm 1.

Algorithm 1 In-Time Sparse Exploration for Robust Circuits

Initialize: The quantum circuit $\mathcal{C}(\theta)$, Dataset \mathcal{D} , Sparsity distribution $\mathbb{S} = \{s_1, \dots, s_B\}$, Update schedule $\{\Delta T, T_{\text{end}}, \gamma, f_{\text{decay}}\}$, Learning rate η

- 1: Randomly sparsify circuits with the sparsity \mathbb{S} *{Initial circuits with reduced quantum gates.}*
- 2: (Optional) Initialize gradient magnitude accumulator \mathcal{M}
- 3: **for** each training iteration t **do**
- 4: Sample a mini-batch from \mathcal{D}
- 5: **if** $(t \bmod \Delta T == 0)$ and $t < T_{\text{end}}$ **then**
- 6: **for** each quantum block b **do**
- 7: $\rho = f_{\text{decay}}(t, \gamma, T_{\text{end}}) \cdot (1 - s_b) \cdot N_b$
- 8: Prune unnecessary gates and grow new important gates with portion ρ based on certain criteria; Update circuit accordingly *{In-time sparse exploration of topologies}*
- 9: **end for**
- 10: **else**
- 11: Compute gradients through the parameter shift $\nabla_{\theta} \mathcal{L}(\theta) = \frac{1}{2} \times (\mathcal{C}(\theta_+) - \mathcal{C}(\theta_-)) \times \frac{\partial \mathcal{L}(\theta)}{\partial \mathcal{C}(\theta)}$ [60] *{Computing parameter updates under real noise models or on real quantum machines}*
- 12: $\theta = \theta - \eta \times \nabla_{\theta} \mathcal{L}(\theta)$ *{Updating weights}*
- 13: (Optional) Update $\mathcal{M} \leftarrow \tau \mathcal{M} + (1 - \tau) |\nabla_{\theta} \mathcal{L}(\theta)|$
- 14: **end if**
- 15: **end for**
- 16: **return** A sparse and noise-adaptive quantum circuit

In our designs, quantum gates are globally ranked according to their significance. Therefore, the sparse topology optimization of different layers is harmonized. Meantime, we do not allow the growth of extra layers, which prevents the increase of circuit depth. Interesting future work is to integrate a circuit depth-based metric during the sparse topology exploration.

The crucial factors of QuantumSEA’s performance lie in four aspects: *i) Sparsity distribution*, which determines the initial shape of quantum circuits; *ii) Update schedule*, that balances the trade-off between exploration and exploitation of sparse circuits’ topology; *iii) Pruning* and *iv) Growth criteria*, which are leveraged to guide the optimization of circuit topologies.

▷ *Sparsity distribution*. The conventional sparse distribution in sparse training of classical networks, i.e., *Erdős-Rényi* [43], determines the layer-wise sparse ratio to be proportional to the number of parameters within each layer. Directly following this convention, it will degrade to a uniform distribution if each block of quantum circuits has the same number of parameters N_b . However, a naïve uniform random sparsification may remove a full control quantum gate, which **places obstacles to the necessary communications between different qubits** and damages the trainability of resulted circuits. Therefore, we first randomly maintain one parameter for each control gate and then adopt the uniform distribution to sparsify the rest of the parameters. Another potential strategy is to introduce non-parameterized control gates like RXYZ and IBMQ Basis quantum circuits, ensuring essential qubit communication. It is also evidenced by our empirical findings from the sparse topology exploration, the performance of quantum circuits deteriorates when most control gates are pruned, due to the poor entangling capability [42].

▷ *Update schedule*. The update schedule usually contains:

- (1) An update interval ΔT that is the number of training iterations between two adjacent updates of sparse topologies.
- (2) An ending epoch T_{end} , after which the sparse connectivity of quantum circuits will be fixed but the updating of their parameters is retained.
- (3) The initial fraction γ of quantum gates that allows being pruned or grown. In our case, we set $\gamma = 50\%$ via a hyperparameter grid search. It is worth mentioning that QuantumSEA does not introduce thresholds to determine the amount of pruned/grown quantum gates. Instead, it removes or inserts operations with top γ insignificance or significance respectively.
- (4) A decay schedule $f_{\text{decay}}(t, \gamma, T_{\text{end}})$ which computes the portion of changeable quantum gate parameters. Following the standard [8], [17], [33], we select a cosine annealing schedule $\frac{\gamma}{2} \times (1 + \cos \frac{t \times \pi}{T_{\text{end}}})$.

▷ *Pruning criterion*. Several common heuristics for approximating the model parameter’s significance can be utilized, including the following three criteria:

- (1) Weight magnitude $|\theta|$ [8];
- (2) Gradient magnitude $|\nabla_{\theta} \mathcal{L}(\theta)|$ [8], [17];
- (3) Saliency $|\theta \times \nabla_{\theta} \mathcal{L}(\theta)|$ [8].

Usually, criteria (2-3) lead to better performance, which is consistent with our experimental results in Section IV-C. A possible reason is that pruning with criteria (2-3) tends to have high gradients in the proceeding iteration of circuit training, leading to a quick reduction of objective values.

▷ *Growth criterion*. We adopt the superior indicator from literature [8], [17], [33], i.e., gradient magnitude, to guide the growth of newly added quantum gates. Such criterion normally requires the full model’s gradient at each stage for updating topology. However, it is not feasible in quantum circuit training, because increased quantum gates will introduce extra gate noises and may violate the constraint of quantum coherent time. Thus, we are unable to get the gradient information of gates that are not contained in the current sparse circuit, via parameter shift. To tackle this dilemma, we propose a historical gradient accumulator \mathcal{M} , which records the gradient of a quantum gate if it has appeared once during the in-time sparse exploration. The accumulator \mathcal{M} is updated by $\tau \mathcal{M} + (1 - \tau) |\nabla_{\theta} \mathcal{L}(\theta)|$, where τ controls the contribution of historical gradients and we use $\tau = 0.9$. This design enables a *fixed* and *efficient* number (a small portion) of quantum gates across the overall procedure. Because ① the size of the accumulator \mathcal{M} is only linearly enlarged along with the increase of quantum qubits, and this negligible memory overhead can be easily afforded by classical computers; ② The historical gradient information offers an approximated importance of quantum operations, discarding the need of computing the significance for all quantum gates at each training iteration.

Meanwhile, to enhance the exploration ability of QuantumSEA, we further inject randomness into the growth indicator as $\kappa \times r + (1 - \kappa) \times |\mathcal{M}|$ (i.e., “Random + Historical Gradient”). r is a random score drawn from the uniform distribution $\mathcal{U}(0, 1)$, $|\mathcal{M}|$ presents the normalized gradient score lied in $[0, 1]$, and the coefficient κ balances the exploration and exploitation of identifying robust sparse quantum circuits. In our case, κ is linearly decayed from 1 to 0 along with the training. Note that the number of newly activated quantum gates or parameters is equal to the number of pruned ones, and the new gates are initialized to zero since it produces better results as suggested by [8], [17], [33].

Note that both pruning and growth are conducted between training iterations when the quantum states have already collapsed. Thus such operations aren’t constrained by coherence time. Also, the number of updated gates via pruning and growth is the same. Hence the scale of sparse circuits won’t gradually increase and so as the risk of qubit decoherence.

b) Implicit capacity: We measure the additional implicit circuit capacity \mathcal{I} gained from in-time sparse exploration, which refers to the overall dimensionality of quantum circuits explored in the space-time manifold. To be specific, $\mathcal{I} = \{\# \text{ the number of gates ever appeared}\} / \{\# \text{ the number of gates in the sparse circuit}\} - 1$, and \mathcal{I} lies in the range $[0, \frac{1}{1-s} - 1]$ where $s \in (0, 1]$ is the circuit’s sparsity. Our extensive experiments demonstrate that a larger \mathcal{I} empowers circuits with improved expressiveness and robustness towards

quantum noises. Additionally, the key to implicit capacity is dynamic exploration and this metric can naturally be generalized beyond QuantumSEA.

c) Circuits' sparsity: A quantum circuit with higher sparsity is less affected by quantum noise as it has a smaller number of gates and thus less inserting noise on real QC devices, which makes it more trainable in terms of noise tolerance. However, fewer gates limit the expressiveness of circuits in the meantime, which cause the circuits less trainable. Fortunately, in our QuantumSEA, we can maintain the circuits with a small number of gates but grant the circuits considerable implicit capacity. Thus it reaches a better trade-off between expressiveness and noise tolerance.

IV. EXPERIMENT

A. Implementation Details

a) Benchmarks: We experiment on 5 QML and 2 VQE common benchmarks. QML tasks are the image recognition including MNIST [30] 2-class (3, 6), 4-class (0, 1, 2, 3); Fashion-MNIST [65] 2-class (dress, shirt), 4-class (t-shirt/top, trouser, pullover, dress), and Vowel [13] 4-class (hid, hId, had, hOd). All these tasks require 4 logical qubits. Due to limited real quantum computing resources, we report the test accuracy on 300 randomly selected test examples, following the conventions in [58]. A center crop is applied to trim input images from size 28×28 to 24×24 , and the averaging pooling is utilized to further downsample them to 4×4 for 2-class and 4-class recognition. As for Vowel-4 dataset, it contains 990 samples, of which 60% data are used for training, 10% for validation, and 30% for testing, following [58]. We compress the original vowel features via principal component analysis (PCA) and only keep the top-10 significant dimensions.

To convert the classical images and vowel features into their quantum representation, we flatten and encode them with rotation gates, where 4 logical qubits are used for 2- and 4-class image recognition, respectively. Specifically, for 4×4 inputs, we use 4 qubits and 4 layers with 4 \mathcal{R}_X , 4 \mathcal{R}_Y , 4 \mathcal{R}_Z , and 4 \mathcal{R}_X quantum gates on each layer. In this way, the 16 classical values are encoded by the 16 gates as their rotation phases. Similarly, we encode the 10 vowel features with 4 qubits and 3 layers, using 4 \mathcal{R}_X , 4 \mathcal{R}_Y , and 2 \mathcal{R}_Z for each layer. Note that different input encoding schemes, *e.g.*, qubit and amplitude embedding, can affect the final performance. For a fair comparison, we follow the same encoding method in [58]. During the measurement stage, we compute the expectation values on the Pauli-Z basis, obtain a value in $[-1, 1]$ for each qubit, and further process them with `Softmax` to get probabilities. Particularly for the 2-class image recognition, we sum the qubits {0 and 1}, {2 and 3} to obtain two values before `Softmax`.

For VQE, the objective is to find the low-energy eigenvalue of a target molecule, which calculates the expectation of the molecule's Hamiltonian by a repetitive measurement. The investigated molecules in this paper are H_2 and H_2O , which are encoded with 2 and 6 logical qubits. Their Hamiltonians

are transformed from the fermionic form to the qubit form by Bravyi-Kitaev mapping [6].

b) Baseline circuit architectures: We inherit the representative circuit designs from [58]. *i) U3+CU3* - Each block contains a U3 and a CU3 layer. The former has one U3 gate for each qubit; the latter is constructed with ring connections of CU3 (0, 1), CU3 (1, 2), CU3 (2, 3), and CU3 (3, 0). *ii) ZZ+RY* - Each block has one ZZ layer with ring connections and one RY layer. *iii) RXYZ* - Each block contains one RX, RY, RZ and CZ layer, in which the CZ layer is formed with ring connections. Besides, there is one \sqrt{H} layer before the blocks. *iv) ZX+XX* - Each block contains one ZX and XX layer. *v) IBMQ Basis* - Each block has RZ, X, RZ, SX, RZ, and CNOT 6 layers in order. For circuits $i \sim iv$, we by default use 8 blocks, except 20 blocks for circuit *v*). And we follow the same human design architecture as in [58].

c) Quantum noise and associated training configurations: We examine four different quantum noise levels from the ideal to realistic cases.

- ❶ *Noise-free Simulation.* It does not involve any noise information in training.
- ❷ *Differentiable Noise Simulation.* Training circuits with approximated noise information [59], which includes quantum error gate insertion via Pauli Twirling [59] and readout noise injection.
- ❸ *Real Noise Models.* It adopts the practical noise models from real quantum computers during training.
- ❹ *On-Chip Training.* Training quantum circuits on real quantum devices.

For ❶ *Noise-free* and ❷ *Differentiable Noise Simulations*, we use an Adam optimizer with 1×10^{-4} weight decay and train the quantum circuits via back-propagation on full training datasets. The learning rate starts from 0 and linearly increases to 5×10^{-3} in the first 30 epochs for QML tasks and 150 steps for VQE tasks, then decreases by cosine annealing. In total, we train 200 epochs with batch size 256 for QML and 1,000 steps with batch size 1 for VQE tasks. In QuantumSEA, we update the circuit topology every 10 epoch for QML and 50 steps for VQE. For ❸ *Real Noise Models* and ❹ *On-Chip Training*, we adopt parameter shift [10], [60] to optimize quantum circuits. The learning rate starts from 0.3 and decays to 0.03, following a cosine annealing scheduler. We train 50 epochs and update the structure of circuits every 5 epoch for QML, while training 250 steps and updating every 25 step for VQE. In all settings, we perform pruning and growth only during the front 70% training period, leaving enough training budget for the final sparse circuit to converge. Following [60], MNIST and Fashion 2-class adopt the front 500 images for training; MNIST and Fashion 4-class choose the front 100 images as the training set. It is worth mentioning that most of our experiments ($> 90\%$) are conducted on practical, and challenging noise-adaptive scenarios and all evaluations are implemented on *real* quantum devices.

d) Quantum devices and compiler configurations: TorchQuantum library [58] is used for circuit construction




Fig. 5. Real QC testing accuracy (%) over the number of parameters in quantum circuits. Evaluation are conducted on *IBMQ-Santiago* with configurations of $\{(IBMQ \text{ Basis, MNIST-2}), (ZX+XX, \text{Fashion-MNIST-2}), (ZZ+RY, \text{MNIST-4})\}$ and real noise models [58].

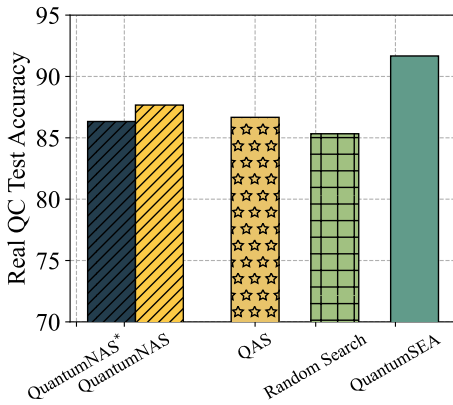


Fig. 6. QuantumSEA v.s. Search Methods like QuantumNAS [58], QAS [16], and Random Search.

and training. *IBMQ* quantum computers with Qiskit¹ APIs are adopted in experiments.

B. In-Time Sparse Exploration Enables Robust Quantum Circuits

Q1: What are the key benefits of QuantumSEA? Improved (implicit) capacity and topology of circuits.

To figure out the key benefits of QuantumSEA, we evaluate it on several QML tasks with ZZ+RY and *IBMQ* Basis quantum circuits under practical noise models of *IBMQ-Santiago*. As shown in Figure 5, we compare our proposal to circuits with various capacities, obtained from noise-aware searching (QuantumNAS [58]), random generation, and human design. Based on the results, several consistent observations can be drawn:

- ❶ *QuantumSEA* delays the accuracy peak. Along with the rise of circuit parameters, measured real QC testing accuracy first increases to a peak value and then degrades. It is within the expectation of the performance improvements in the beginning, thanks to the enlarged circuit capacity.

Notably, excessive parameters of further added quantum gates amplify the noises, leading to deteriorated circuits. However, QuantumSEA maintains fewer gates over the whole training, and leverages dynamic topology updates to explore high-quality sparse quantum circuits. In this way, it expands the in-time circuit capacity implicitly while avoiding additional quantum noises, which allows higher accuracies and postponed performance saturation.

- ❷ *Comparison with existing approaches.* As indicated by dark curves in Figure 5, QuantumSEA consistently surpasses the other three baselines at all parameter levels for diverse scenarios. For example, up to $\{30.66\%, 30.34\%, 17.33\%\}$ accuracy boosts given the same number of parameters and $\{3.00\% \sim 7.61\%, 5.33\% \sim 6.00\%, 1.33\% \sim 6.00\%\}$ improvements for the best achievable accuracy on the settings of $\{(IBMQ \text{ Basis, MNIST-2}), (ZX+XX, \text{Fashion-MNIST-2}), (ZZ+RY, \text{MNIST-4})\}$. It suggests the superiority of located circuit topology against noises, compared to the searched, randomly generated, and human-designed architectures. Therefore, we conclude that the improved implicit capacity and topology of circuits make major contributions.

Q2: Does QuantumSEA perform better than neural architecture search-based methods?

We study our found sparse circuits versus searched circuits, to reveal the advantages of QuantumSEA. The comparison results of (RXYZ, MNIST-2) with 36 quantum gates, are collected in Figure 6 where * denotes *noise-free* simulation, and the others are under practical noise models from *IBMQ-Santiago*. We observe that:

- ❶ The noise-aware optimization brings improved performance, echoed with the finding in [58];
- ❷ Compared to all other three types of searched quantum circuits, QuantumSEA establishes a superior quantum circuit ($\geq 4.00\%$ higher accuracy) by jointly learning topology and weights, which demonstrates the effectiveness of our proposed framework.

Note that QuantumNAS and another neural architecture

¹<https://quantum-computing.ibm.com/>

TABLE I

REAL QC TESTING ACCURACY (%) AND ENERGY ACROSS DIVERSE QUANTUM TASKS (QML AND VQE) AND CIRCUIT (8 BLOCKS) TYPES. *Real Noise Simulation* USES REAL NOISE MODELS [58] AND THE PARAMETER-SHIFT UPDATE RULE DURING TRAINING. *On-Chip Training* OPTIMIZES QUANTUM CIRCUITS WITH THE PARAMETER-SHIFT UPDATE RULE ON REAL QUANTUM DEVICES (THE *bottom right* TABLE). N. A. STANDS FOR “NOT APPLICABLE”. A LARGER ACCURACY OR A SMALLER ENERGY VALUE IMPLIES BETTER PERFORMANCE.

Circuit	Method (Sparsity)	MNIST-2	MNIST-4	F-MNIST-2	F-MNIST-4	Vowel-4	H ₂	H ₂ O
RXYZ on IBMQ-Quito	Human Design (0%)	88.67	47.67	84.33	64.00	28.45	-1.7263	-36.2711
	Human Design (50%)	82.00	55.00	81.67	44.33	31.03	-1.7275	-33.7164
	Random (50%)	89.00	56.67	84.00	65.00	44.83	-1.7285	-39.4637
	QuantumNAS (50%) [58]	88.33	56.33	85.33	66.67	40.52	-	-
	QuantumSEA (50%)	90.67	57.67	86.00	72.33	50.86	-1.7444	-39.8873
	Human Design (0%)	88.67	47.67	84.33	64.00	28.45	-1.7263	-36.2711
	Human Design (80%)	82.67	45.33	75.33	50.00	43.10	-1.7710	-36.4497
	Random (80%)	86.67	21.00	88.00	58.00	40.52	-1.6298	-24.4218
	QuantumNAS (80%) [58]	86.00	46.67	83.00	51.67	31.03	-1.3362	-31.3603
	QuantumSEA (80%)	89.33	54.00	89.33	62.67	49.14	-1.7760	-40.8411
Task	Method (Sparsity)	ZZ+RY	RXYZ	ZX+XX	IBMQ Basis	Circuit	MNIST-2	MNIST-4
MNIST-4 on IBMQ-Santiago	Human Design (0%)	27.67	60.00	34.33	46.00	RXYZ	85.00	58.00
	Human Design (50%)	46.67	52.67	22.33	30.33		81.00	53.33
	Random (50%)	44.67	50.33	42.00	33.00		78.67	54.67
	QuantumNAS (50%) [58]	51.67	58.67	44.67	40.00		N. A.	N. A.
	QuantumSEA (50%)	50.33	62.33	45.67	60.67		86.33	55.33

TABLE II
REAL QC TESTING ACCURACY (%) ACROSS TWO 6-CLASS CLASSIFICATION TASKS ON IBM-OSLO WITH ZZ+RY CIRCUITS.

Circuit	Method (Sparsity)	MNIST-6	F-MNIST-6
ZZ+RY on IBM-OSLO	Human Design (0%)	44.67	50.67
	Human Design (50%)	39.00	49.00
	Random (50%)	43.67	51.00
	QuantumNAS (50%) [58]	44.33	52.67
	QuantumSEA (50%)	45.33	54.33

search-based method QAS [16] require a densely optimized super-circuit, suffering from the scalability issue for on-chip training. In contrast, our method always deals with a sparse circuit. Moreover, we update both weights and topologies under real quantum noise, while QuantumNAS and QAS only perform noise-adaptive topology updates.

Q3: Does QuantumSEA achieve superior performance across quantum tasks and circuit types?

Comprehensive experiments about different quantum tasks (QML and VQE) and circuit types have been carried out to support the great generalization ability of QuantumSEA. According to the results in Table I and Table II, we see that:

- 1 *Win other baselines with improved in-time capacity \mathcal{I} and circuit topologies.* On 5 QML tasks, 2 VQE tasks, and 4 circuit types, QuantumSEA reaches first-rate performance compared to other circuits with the same parameters (e.g., 50% sparsity), implying the better quality of our explored circuit connectivities. At most cases like RXYZ on F-MNIST-2/4 and H₂ etc., our sparse circuits outperform their dense counterparts (0% sparsity) significantly. The main reasons stem from two aspects: *i*) excessive quantum gates in dense human-designed circuits induce overmuch noise and hurt the QNN prediction; *ii*) in these

cases, QuantumSEA normally obtains sufficient implicit capacity \mathcal{I} (≥ 0.95 and ≥ 3.95 for 50% and 80% sparsity circuits), signifying a thorough exploration of space-time manifold and enhanced expressiveness.

- 2 *Superior performance in on-chip training.* Based on a recent state-of-the-art on-chip training research [60], we examine our approaches on MNIST-2/4 with RXYZ circuits. **Marked** results (Table I) suggest that ours surpass human-designed circuits [60] (both 50% and 100%) by a margin of 1.33% \sim 5.32% performance improvements on MNIST-2, and win 1.99% accuracy gains on MNIST-4 compared to human-designed circuits with the same number of quantum gates.
- 3 *For a more complicated 6-class classification task.* QuantumSEA also achieves superior performance compared with other circuits with the same parameters. Specifically, as shown in Table II, QuantumSEA obtains an extra performance improvement by up to 6.33% and 5.33% on MNIST-6 and F-MNIST-6, respectively.

Note that for a fair comparison, we adopt the same circuit design style from [58], which explores the circuit topology within a basic block and duplicates it to construct the full quantum circuits. As shown in Table I, our methods achieve consistent performance gains in most cases.

Q4: Does QuantumSEA achieve superior performance across diverse real quantum devices?

We further evidence the superiority of QuantumSEA across six different QC machines. For each quantum device, we train circuits with their associated practical noise model and evaluate optimized circuits on real QC. Table III tells that:

- 1 *Given the same number of quantum gates (or sparsity), our proposals show clear advantages in accuracy gains and energy decrements, indicated by **bold** numbers.*

TABLE III
REAL QC TESTING ACCURACY (%) AND ENERGY ACROSS DIVERSE QUANTUM COMPUTING DEVICES WITH THE CIRCUIT RXYZ. REAL NOISE MODELS AND THE PARAMETER-SHIFT UPDATE ARE ADOPTED DURING TRAINING.

Quantum Computing		Sparsity	IBMQ-Santiago	IBMQ-Bogota	IBMQ-Lima	IBMQ-Belem	IBMQ-Quito	IBMQ-Manila
MNIST-4	Human Design	0%	60.00	65.00	54.00	23.67	47.67	68.33
	Human Design	50%	52.67	53.00	46.33	23.67	55.00	67.67
	Random	50%	50.33	59.67	56.67	42.00	56.67	58.67
	QuantumNAS [58]	50%	58.67	57.67	55.67	51.33	56.33	61.00
	QuantumSEA	50%	62.33	62.67	58.67	52.67	57.67	69.67
H ₂	Human Design	0%	-1.6823	-1.6520	-1.7600	-1.4168	-1.7263	-1.7202
	Human Design	50%	-1.6673	-1.6792	-1.7561	-1.4535	-1.7275	-1.7097
	Random	50%	-1.6742	-1.6585	-1.7357	-1.4356	-1.7285	-1.7482
	QuantumSEA	50%	-1.6822	-1.6989	-1.7717	-1.4919	-1.7444	-1.7587




Fig. 7. Ablation study of update intervals.

Although the dense (0% sparsity) human-designed circuits moderately beat QuantumSEA in some QCs like IBMQ-Bogota, our sparse ones still enjoy around 50% time savings for both training and inference, as also depicted in Figure 2 (Right) and 8. Determined by the physical environments of quantum machines, the operation error of executing quantum gates has varied intensities [58]. It is why certain real QCs might be more amenable to handling circuits with more parameters.

C. Ablation Study and Visualization

a) *Ablation on pruning and growth criteria:* We perform the ablation study on pruning and growth indicators, which play crucial roles in QuantumSEA. From Table IV, experiments with different sparsity levels show consistent preferences of the criteria. Specifically, using “Salience” to estimate weight significance for pruning and adopting the fused score of “Historical (His.) Gradient + Random” for growing new connections, demonstrates the best performance, i.e., 1.00% ~ 12.67% better accuracy compared to other options. Thus, we inherit this optimal setup for all other experiments. Additionally, to evaluate the impact of the historical gradient, we replace it with the vanilla gradient in the optimal setup. We can observe that the historical gradient significantly improves the final accuracy by up to 4.34%, which is expectable because using the vanilla gradient, the gates that are not contained in the current circuits can only grow by a random chance.

b) *Ablation on the update interval of QuantumSEA:* The length of the update interval ΔT controls the critical exploration and exploitation trade-off [8], [33] in QuantumSEA. For instance, a larger ΔT emphasizes exploitation since more epochs are taken to approximate weight significance; on the

TABLE IV
ABLATION OF PRUNING AND GROW CRITERIA IN QUANTUMSEA. EVALUATIONS ARE CONDUCTED ON IBMQ-SANTIAGO WITH (ZZ+RY, MNIST-4) AND PRACTICAL NOISE MODELS. “OURS” DENOTES “HIS. GRADIENT + RANDOM”.

Sparsity	Pruning	Grow	Accuracy (%)
50%	Random	Ours	43.67
50%	Weight	Ours	46.67
50%	His. Gradient	Ours	42.33
50%	Salience	Ours w.o. His.	43.33
50%	Salience	Ours	47.67
50%	Salience	His. Gradient	42.33
50%	Salience	Random	35.00
80%	Random	Ours	47.33
80%	Weight	Ours	49.67
80%	His. Gradient	Ours	45.00
80%	Salience	Ours w.o. His.	48.33
80%	Salience	Ours	51.67
80%	Salience	His. Gradient	49.33
80%	Salience	Random	45.00

TABLE V
COMPATIBLE WITH EXISTING TECHNIQUES OF NOISE MITIGATION [59]. WE USE THE SAME EXPERIMENT CONFIGURATIONS IN [59].

Method	F-MNIST-4 (Full Training Data)
Ours (50% Sparsity)	72.33
+ Normalization	83.00
+ Quantization	74.00
Ours (80% Sparsity)	54.67
+ Normalization	69.00
+ Quantization	66.33

contrary, a smaller ΔT allows more topology updates and enhances the exploration ability (usually comes with a larger implicit capacity \mathcal{I}). We conduct the ablation with fine-grained ΔT choices in Figure 7. It suggests $\Delta T = 2$ is the “sweet-point” setup of our in-time sparse circuit explorations.

c) *Compatibility with existing noise mitigation:* QuantumSEA functions in orthogonal to existing noise mitigation mechanisms, such as post-measurement normalization and

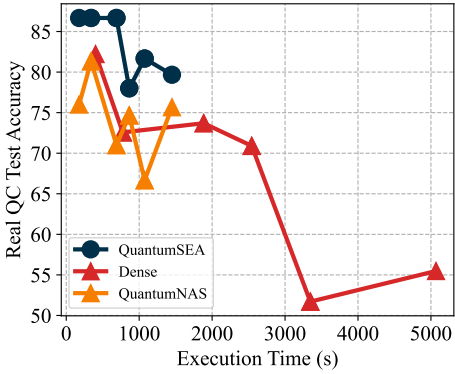


Fig. 8. Real QC testing accuracy (%) of Fashion-MNIST-2 over the training time of quantum circuit ZX+XX on IBMQ-Santiago. Our proposed QuantumSEA achieves a better Pareto Frontier approximation of the trade-off between accuracy and execution time.




Fig. 9. Real QC testing accuracy (%) of Fashion-MNIST-2 and MNIST-2 over the implicit capacity of quantum circuit IBMQ Basis on the IBMQ-Quito quantum machine.

quantization from the recent advance [59]. The compatibility is shown by obtaining extra performance improvements from their combinations, as recorded in Table V.

d) Discussion on quantum coherence time: Since the runtime on real quantum devices scales nearly linearly to # qubits and # quantum gates, QuantumSEA exhibits great advantages in the trade-off between real QC accuracy and execution time, as shown in Figure 8. It is mainly because our methods enable superior performance with a small number of quantum gates with the assistance of in-time implicit capacity.

e) Ablation on implicit capacity: Our QuantumSEA provides the quantum circuit training with extra implicit capacity \mathcal{I} . As shown in Figure 9, the granted implicit capacity \mathcal{I} is positively related to the final achievable accuracy. A larger \mathcal{I} promotes circuits with enhanced expressiveness and robustness against quantum errors, which is a potential reason for QuantumSEA’s superior performance. On the contrary, as other competitive methods won’t dynamically adjust the sparse circuits, their implicit capacity equals zero.

f) Visualization of dense and sparse quantum circuits: Figure 10 presents identified sparse quantum circuit topologies from QuantumSEA (Middle) & QuantumNAS (Right) and their dense counterpart (Left). Several observations can be drawn from these results:

- 1 Compared with the dense circuit in Figure 10 (Left), the

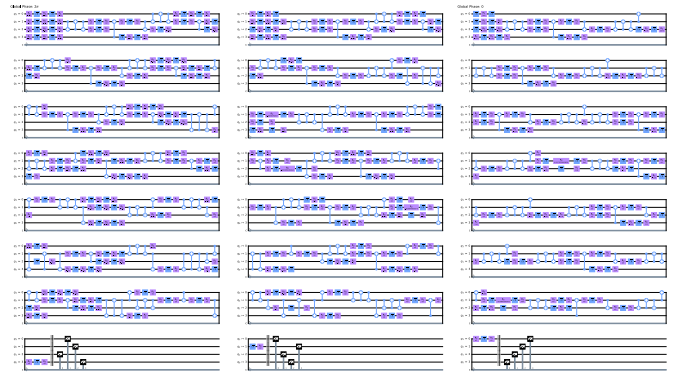


Fig. 10. Visualization of a dense circuit (Left) and sparse circuits from our QuantumSEA (Middle) & QuantumNAS (Right) with 8 blocks RXYZ on IBMQ-Quito. Please kindly zoom-in for better readability.

sparse circuit from QuantumSEA (Figure 10 (Middle)) preserves a similar number of remaining quantum gates for each qubit. In the meantime, most of the control gates across different qubits are maintained. Such “undamaged” communication ability between qubits could be one possible reason for the superiority of QuantumSEA.

- 2 Quantitative results of {# depth, # X/Y/Z gate, # CZ gate, expressibility [52]} for QuantumNAS and our QuantumSEA are {32, 96, 24, 0.0064} and {31, 48, 24, 0.0047}, respectively. It shows the improved expressibility (i.e., smaller Kullback Leibler divergence) and reduced gate error of our generated sparse circuits.

V. CONCLUSION AND DISCUSSION

In this paper, we introduce an in-time sparse exploration framework, QuantumSEA, which jointly optimizes the topology and weight of quantum circuits under practical quantum noises. Its advantages lie in: (1) Sparse circuits with reduced quantum gates over the full training and inference processes, substantially alleviating the effects of decoherence; (2) The in-time dynamic connectivity exploration produces extra implicit capacity, greatly enhancing the circuit’s expressiveness; (3) Identified noise-adaptive sparse topology regularizes circuit training and boosts the noise resistance. Comprehensive experiments validate our proposals regarding superior performance and significant training time savings. Our approaches currently only consider trimming quantum gates in the circuit. We plan to investigate the possibility of sparsifying both quantum qubits and gates in future works. Although we see our work as scientific in nature, it might amplify the existing negative societal impacts of using quantum computing because we have no control over anyone who can get access to our QuantumSEA and its generated quantum circuits.

When and what tasks do PQCs can achieve quantum advantages? Currently, quantum approaches have demonstrated advantages in applications like large unitary matrix multiplication. To outperform classical neural networks, we think it requires PQCs with more than 50 qubits and effective error correction mechanisms (or fault-tolerant machines).

ACKNOWLEDGMENT

This work is funded in part by EPIQC, an NSF Expedition in Computing, under award CCF-1730449; in part by STAQ under award NSF Phy-1818914; in part by the US Department of Energy Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing Program; and in part by the NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks (NSF Award 2016136) and in part based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. FTC is Chief Scientist for Quantum Software at Infleqtion and an advisor to Quantum Circuits, Inc. We thank MIT-IBM Watson AI Lab, Qualcomm Innovation Fellowship for supporting this research. We acknowledge the use of IBM Quantum services for this work.

REFERENCES

- [1] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf, "Training deep quantum neural networks," *Nature communications*, vol. 11, no. 1, pp. 1–6, 2020.
- [2] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," *Quantum Science and Technology*, vol. 4, no. 4, p. 043001, 2019.
- [3] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke *et al.*, "Noisy intermediate-scale quantum (nisq) algorithms," *arXiv preprint arXiv:2101.08448*, 2021.
- [4] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [5] M. Bilkis, M. Cerezo, G. Verdon, P. J. Coles, and L. Cincio, "A semi-agnostic ansatz with variable structure for quantum machine learning," *arXiv preprint arXiv:2103.06712*, 2021.
- [6] S. B. Bravyi and A. Y. Kitaev, "Fermionic quantum computation," *Annals of Physics*, vol. 298, no. 1, pp. 210–226, 2002.
- [7] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. Sawaya *et al.*, "Quantum chemistry in the age of quantum computing," *Chemical reviews*, vol. 119, no. 19, pp. 10856–10915, 2019.
- [8] T. Chen, Y. Cheng, Z. Gan, L. Yuan, L. Zhang, and Z. Wang, "Chasing sparsity in vision transformers: An end-to-end exploration," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [9] M. A. Continentino, *Key Methods and Concepts in Condensed Matter Physics*. IOP Publishing, 2021.
- [10] G. E. Crooks, "Gradients of parameterized quantum gates using the parameter-shift rule and gate decomposition," *arXiv preprint arXiv:1905.13311*, 2019.
- [11] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, "Error mitigation with clifford quantum-circuit data," *Quantum*, vol. 5, p. 592, 2021.
- [12] P. Deglmann, A. Schäfer, and C. Lennartz, "Application of quantum calculations in the chemical industry—an overview," *International Journal of Quantum Chemistry*, vol. 115, no. 3, pp. 107–136, 2015.
- [13] D. H. Deterding, "Speaker normalisation for automatic speech recognition." Ph.D. dissertation, University of Cambridge, 1990.
- [14] S. Ding, T. Chen, and Z. Wang, "Audio lottery: Speech recognition made ultra-lightweight, noise-robust, and transferable," in *International Conference on Learning Representations, 2022*. [Online]. Available: <https://openreview.net/forum?id=9Nk6AJkVYB>
- [15] Y. Ding and F. T. Chong, "Quantum computer systems: Research for noisy intermediate-scale quantum computers," *Synthesis Lectures on Computer Architecture*, vol. 15, no. 2, pp. 1–227, 2020.
- [16] Y. Du, T. Huang, S. You, M.-H. Hsieh, and D. Tao, "Quantum circuit architecture search for variational quantum algorithms," *npj Quantum Information*, vol. 8, no. 1, pp. 1–8, 2022.
- [17] U. Evci, T. Gale, J. Menick, P. S. Castro, and E. Elsen, "Rigging the lottery: Making all tickets winners," in *International Conference on Machine Learning*. PMLR, 2020, pp. 2943–2952.
- [18] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.
- [19] E. Farhi and H. Neven, "Classification with quantum neural networks on near term processors," *arXiv preprint arXiv:1802.06002*, 2018.
- [20] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 1996, pp. 212–219.
- [21] S. Gui, H. Wang, H. Yang, C. Yu, Z. Wang, and J. Liu, "Model compression with adversarial robustness: A unified optimization framework," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [22] A. W. Harrow, A. Hassidim, and S. Lloyd, "Quantum algorithm for linear systems of equations," *Physical review letters*, vol. 103, no. 15, p. 150502, 2009.
- [23] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.
- [24] Z. He, J. Su, C. Chen, M. Pan, and H. Situ, "Search space pruning for quantum architecture search," *The European Physical Journal Plus*, vol. 137, no. 4, p. 491, 2022.
- [25] A. Heifetz, *Quantum mechanics in drug discovery*. Springer, 2020.
- [26] C.-Y. Hsieh, C.-H. Wu, C.-H. Huang, H.-S. Goan, and J. C. M. Li, "Realistic fault models and fault simulation for quantum dot quantum circuits," in *2020 57th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2020, pp. 1–6.
- [27] M. Huo and Y. Li, "Self-consistent tomography of temporally correlated errors," *Communications in Theoretical Physics*, vol. 73, no. 7, p. 075101, 2021.
- [28] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, 2017.
- [29] E.-J. Kuo, Y.-L. L. Fang, and S. Y.-C. Chen, "Quantum architecture search via deep reinforcement learning," *arXiv preprint arXiv:2104.07715*, 2021.
- [30] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [31] Y. Li and S. C. Benjamin, "Efficient variational quantum simulator incorporating active error minimization," *Physical Review X*, vol. 7, no. 2, p. 021050, 2017.
- [32] Z. Liang, H. Wang, J. Cheng, Y. Ding, H. Ren, Z. Gao, Z. Hu, D. S. Boning, X. Qian, S. Han *et al.*, "Variational quantum pulse learning," in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2022, pp. 556–565.
- [33] S. Liu, L. Yin, D. C. Mochan, and M. Pechenizkiy, "Do we actually need dense over-parameterization? in-time over-parameterization in sparse training," *arXiv preprint arXiv:2102.02887*, 2021.
- [34] S. Lloyd, S. Garnerone, and P. Zanardi, "Quantum algorithms for topological and geometric analysis of data," *Nature communications*, vol. 7, no. 1, pp. 1–7, 2016.
- [35] S. Lloyd, M. Mohseni, and P. Rebentrost, "Quantum algorithms for supervised and unsupervised machine learning," *arXiv preprint arXiv:1307.0411*, 2013.
- [36] —, "Quantum principal component analysis," *Nature Physics*, vol. 10, no. 9, pp. 631–633, 2014.
- [37] S. Lloyd, M. Schuld, A. Ijaz, J. Izaac, and N. Killoran, "Quantum embeddings for machine learning," *arXiv preprint arXiv:2001.03622*, 2020.
- [38] E. Magesan, J. M. Gambetta, and J. Emerson, "Characterizing quantum gates via randomized benchmarking," *Physical Review A*, vol. 85, no. 4, p. 042311, 2012.
- [39] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes," *Nature communications*, vol. 9, no. 1, pp. 1–6, 2018.
- [40] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. De Jong, "Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states," *Physical Review A*, vol. 95, no. 4, p. 042308, 2017.
- [41] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," *New Journal of Physics*, vol. 18, no. 2, p. 023023, 2016.
- [42] D. A. Meyer and N. R. Wallach, "Global entanglement in multiparticle systems," *Journal of Mathematical Physics*, vol. 43, no. 9, pp. 4273–4278, 2002.

- [43] D. C. Mocanu, E. Mocanu, P. Stone, P. H. Nguyen, M. Gibescu, and A. Liotta, "Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science," *Nature communications*, vol. 9, no. 1, pp. 1–12, 2018.
- [44] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn *et al.*, "Quantum optimization using variational algorithms on near-term quantum devices," *Quantum Science and Technology*, vol. 3, no. 3, p. 030503, 2018.
- [45] M. Ostaszewski, E. Grant, and M. Benedetti, "Structure optimization for parameterized quantum circuits," *Quantum*, vol. 5, p. 391, 2021.
- [46] M. Ostaszewski, L. M. Trenkwalder, W. Masarczyk, E. Scerri, and V. Dunjko, "Reinforcement learning for optimization of variational quantum circuit architectures," *Advances in Neural Information Processing Systems*, vol. 34, pp. 18 182–18 194, 2021.
- [47] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," *Nature communications*, vol. 5, no. 1, pp. 1–7, 2014.
- [48] P. Rebentrost, M. Mohseni, and S. Lloyd, "Quantum support vector machine for big data classification," *Physical review letters*, vol. 113, no. 13, p. 130503, 2014.
- [49] M. Schuld and F. Petruccione, *Supervised learning with quantum computers*. Springer, 2018, vol. 17.
- [50] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.
- [51] M. P. Silverman, *Quantum superposition: counterintuitive consequences of coherence, entanglement, and interference*. Springer Science & Business Media, 2008.
- [52] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, "Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms," *Advanced Quantum Technologies*, vol. 2, no. 12, p. 1900070, 2019.
- [53] S. Sim, J. Romero, J. F. Gonthier, and A. A. Kunitsa, "Adaptive pruning-based optimization of parameterized quantum circuits," *Quantum Science and Technology*, vol. 6, no. 2, p. 025019, 2021.
- [54] A. Strikis, D. Qin, Y. Chen, S. C. Benjamin, and Y. Li, "Learning-based quantum error mitigation," *PRX Quantum*, vol. 2, no. 4, p. 040330, 2021.
- [55] K. Temme, S. Bravyi, and J. M. Gambetta, "Error mitigation for short-depth quantum circuits," *Physical review letters*, vol. 119, no. 18, p. 180509, 2017.
- [56] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth *et al.*, "The variational quantum eigensolver: a review of methods and best practices," *arXiv preprint arXiv:2111.05176*, 2021.
- [57] A. Van der Ven, Z. Deng, S. Banerjee, and S. P. Ong, "Rechargeable alkali-ion battery materials: theory and computation," *Chemical Reviews*, vol. 120, no. 14, pp. 6977–7019, 2020.
- [58] H. Wang, Y. Ding, J. Gu, Z. Li, Y. Lin, D. Z. Pan, F. T. Chong, and S. Han, "Quantumnas: Noise-adaptive search for robust quantum circuits," *HPCA*, 2022.
- [59] H. Wang, J. Gu, Y. Ding, Z. Li, F. T. Chong, D. Z. Pan, and S. Han, "Quantumnat: Quantum noise-aware training with noise injection, quantization and normalization," *DAC*, 2022.
- [60] H. Wang, Z. Li, J. Gu, Y. Ding, D. Z. Pan, and S. Han, "Qoc: Quantum on-chip training with parameter shift and gradient pruning," *DAC*, 2022.
- [61] H. Wang, P. Liu, J. Cheng, Z. Liang, J. Gu, Z. Li, Y. Ding, W. Jiang, Y. Shi, X. Qian *et al.*, "Quest: Graph transformer for quantum circuit reliability estimation," *arXiv preprint arXiv:2210.16724*, 2022.
- [62] P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, "Single ion qubit with estimated coherence time exceeding one hour," *Nature communications*, vol. 12, no. 1, pp. 1–8, 2021.
- [63] B. J. Williams-Noonan, E. Yuriev, and D. K. Chalmers, "Free energy methods in drug design: Prospects of "alchemical perturbation" in medicinal chemistry: miniperspective," *Journal of medicinal chemistry*, vol. 61, no. 3, pp. 638–649, 2018.
- [64] P. Wittek, *Quantum machine learning: what quantum computing means to data mining*. Academic Press, 2014.
- [65] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
- [66] S. Ye, K. Xu, S. Liu, H. Cheng, J.-H. Lambrechts, H. Zhang, A. Zhou, K. Ma, Y. Wang, and X. Lin, "Adversarial robustness vs. model compression, or both?" in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 111–120.
- [67] S.-X. Zhang, C.-Y. Hsieh, S. Zhang, and H. Yao, "Differentiable quantum architecture search," *arXiv preprint arXiv:2010.08561*, 2020.