

Advancing Differentiable Quantum Architecture Search via Topological Constraints, Gate Memory and Surrogate Neural Networks

David Kral, Thomas Blumensath

June 12, 2026

Abstract

Variational quantum algorithms are widely used in quantum computing but rely heavily on manually designed circuit architectures. Quantum Architecture Search (QAS) addresses this bottleneck by automating the circuit design process using classical machine learning. In this work, we propose novel QAS approaches that advance differentiable QAS by integrating architectural constraints, sequential gate memory, and surrogate neural networks. We evaluate a greedy (G-QAS) and a Differentiable QAS (DQAS) baseline against Architecture-Constrained DQAS (AC-DQAS), Memory DQAS (M-DQAS), and a Combined method (C-DQAS) which includes both architecture constraints and gate memory. Across three distinct benchmarks, we demonstrate that AC-DQAS yields the highest overall accuracy. Additionally, C-DQAS offers a highly effective compromise, having comparable accuracy to AC-DQAS whilst reducing both computational runtime and required circuit evaluations. Furthermore, we incorporate a surrogate neural network in the four differentiable methods, creating NN-DQAS, NN-AC-DQAS, NN-M-DQAS, and NN-C-DQAS. This surrogate is used to predict the best candidates circuit performances during the search, and only those best are evaluated using quantum circuits. We show that including this surrogate slightly decreases performance for all methods, but drastically reduces the needed number of quantum computer evaluations by at least a factor of five, making it more available for deployment on quantum hardware.

1 Introduction

The limitations of Noisy Intermediate-Scale Quantum (NISQ) devices, such as hardware noise, decoherence, and restricted circuit depths, motivate the development of Variational Quantum Algorithms (VQAs), hybrid frameworks explicitly tailored for NISQ hardware [1, 2]. In these algorithms, we optimise the parameters of continuous parametric gates to minimise a problem-specific cost function. However, a crucial vulnerability of VQAs is their reliance on a pre-determined ansatz. An unsuitable architecture restricts expressive power and prevents convergence, whilst an over-parametrised architecture can lead to barren plateaus, where gradients vanish exponentially as the system size increases [3].

Quantum Architecture Search (QAS) mitigates this by simultaneously optimising these continuous gate parameters and the topological arrangement of the circuit. Many approaches have been used, including the differentiable method [4] which we focus on in this work, evolutionary methods [5, 6, 7], deep neural networks [8, 9, 10], and various other machine learning approaches [11, 12, 13]. These methods rely heavily on evaluating a vast number of quantum circuits during the search phase. This requirement creates a significant computational bottleneck that our research specifically addresses.

2 Developed Algorithms

We developed five distinct baseline QAS algorithms, evaluated using a predefined pool of gates ($R_X, R_Y, R_Z, \text{CNOT}$) across various benchmarks. In short, these algorithms were:

- **Greedy QAS (G-QAS):** This method is based on ADAPT-VQE from Grimsley et al. [14]. It starts with an empty circuit and samples all available gates, appending the one that yields the largest improvement. As it only appends a gate if an immediate improvement is observed, it is susceptible to getting trapped in local minima for more complex problems.
- **Differentiable QAS (DQAS):** Based on the work of Zhang et al. [4], this method addresses the discrete nature of QAS by relaxing the architecture space into a continuous domain. A learnable weight matrix assigns probabilities to every candidate gate at each position. We then sample circuits probabilistically from these weights and update them based on the circuit performance. Our contributions include rigorous testing, hyperparameter modifications, and the introduction of a penalty for repeated gates.
- **Architecture-Constrained DQAS (AC-DQAS):** Building upon the DQAS algorithm, we enforce a strict topological constraint by constructing the circuit via alternating layers of single-qubit and two-qubit operations. This was inspired by Du et al. [11], who explored a similar concept but with random circuit generators. These constraints ensure the spread of entanglement and promote higher expressiveness.
- **Memory DQAS (M-DQAS):** This novel method introduces sequential dependence, using transition weight matrices that conditioned gate probabilities on the most recent operation applied to each qubit. Microscopic patterns, such as following R_X with a CNOT gate, could therefore be learned natively.
- **Combined DQAS (C-DQAS):** Our final novel baseline method merges the rigid architectural restrictions of AC-DQAS with the sequential gate memory of M-DQAS.

2.1 Surrogate Neural Network Integration

To further mitigate the need for physical quantum circuit evaluations, we integrate a classical surrogate neural network (NN) into the differentiable QAS frameworks, yielding four novel variants: NN-DQAS, NN-AC-DQAS, NN-M-DQAS, and NN-C-DQAS. Rather than calculating the exact cost function for every proposed quantum circuit, the NN is trained to predict the performance of candidate architectures. From a large generated batch, the NN identifies the best performing circuits, and only these candidates are evaluated directly on quantum hardware. These exact evaluations are subsequently appended to a dynamic dataset to continuously train the NN. Consequently, as the architecture probability matrices converge toward optimal gate configurations, this iterative retraining enables the NN to specialise within the narrowing architectural sub-space.

3 Results and Discussion

The algorithms were evaluated across three benchmarks. First, we adopted the unitary approximation dataset from Lu et al. [15]. Additionally, we generated custom datasets across two domains: Max-Cut graph instances (8, 12, and 16 nodes) formulated using standard QAOA methods established in [4, 16], and a quantum chemistry dataset of common molecular targets, similar to frameworks from other VQE studies [8, 14]. Figure 1 illustrates a subset of these results, plotting task-specific accuracy against the required physical quantum circuit evaluations. For unitary approximation, accuracy is defined by the reconstruction fidelity (f); for Max-Cut, it is represented by the ratio of the found result against the exact solution.

As shown in Figure 1, AC-DQAS consistently achieved the highest accuracy across all benchmarks. C-DQAS emerged as a highly practical compromise, maintaining accuracy superior to unconstrained DQAS while significantly reducing the number of required hardware evaluations. Conversely, G-QAS performed poorly and scaled inefficiently on complex tasks. M-DQAS also




Figure 1: Performance comparison of accuracy versus required quantum circuit evaluations for Unitary Approximation (top) and Max-Cut (bottom). Ideal performance lies in the top-left quadrant. Models using an NN are denoted by stars; baseline models by dots.

struggled regarding accuracy; lacking macroscopic constraints, it failed to capture global structural context and exhibited highly inefficient time scaling.

Crucially, integrating surrogate neural networks reduced the required physical quantum circuit evaluations by at least a factor of five. By shifting the evaluative burden onto classical processors, the NN-assisted variants also mitigated the severe time complexity seen in standard methods, achieving highly efficient time scaling. While this classical pre-filtering caused a small decrease in absolute accuracy, NN-AC-DQAS and NN-C-DQAS still consistently outperformed standard DQAS baselines. For practical near-term hardware deployment, this massive reduction in execution cost and classical runtime overwhelmingly can justify the slight performance penalty.

4 Conclusion

Our results show that applying architectural constraints yields the highest accuracy improvements for Differentiable QAS. While introducing sequential gate dependencies slightly reduces this accuracy, it significantly decreases the required quantum circuit evaluations. Consequently, combining both ideas provides a framework that optimally balances high accuracy with reduced computational runtime. Furthermore, integrating a surrogate neural network into these frameworks reduces the necessary physical quantum evaluations by a factor of five or more with only a minimal loss in accuracy, offering a highly efficient and scalable solution for future NISQ applications.

References

- [1] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018. arXiv:1801.00862 [quant-ph].
- [2] Edward Farhi and Aram W. Harrow. Quantum Supremacy through the Quantum Approximate Optimization Algorithm, October 2019. arXiv:1602.07674 [quant-ph].
- [3] Martín Larocca, Supanut Thanasilp, Samson Wang, Kunal Sharma, Jacob Biamonte, Patrick J. Coles, Lukasz Cincio, Jarrod R. McClean, Zoë Holmes, and M. Cerezo. Barren plateaus in variational quantum computing. *Nature Reviews Physics*, 7(4):174–189, April 2025.
- [4] Shi-Xin Zhang, Chang-Yu Hsieh, Shengyu Zhang, and Hong Yao. Differentiable Quantum Architecture Search. *Quantum Science and Technology*, 7(4):045023, October 2022. arXiv:2010.08561 [quant-ph].
- [5] Hanrui Wang, Yongshan Ding, Jiaqi Gu, Zirui Li, Yujun Lin, David Z. Pan, Frederic T. Chong, and Song Han. QuantumNAS: Noise-Adaptive Search for Robust Quantum Circuits. In *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pages 692–708, April 2022. arXiv:2107.10845 [quant-ph].
- [6] Yuhan Huang, Qingyu Li, Xiaokai Hou, Rebing Wu, Man-Hong Yung, Abolfazl Bayat, and Xiaoting Wang. Robust resource-efficient quantum variational ansatz through evolutionary algorithm. *Physical Review A*, 105(5):052414, May 2022. arXiv:2202.13714 [quant-ph].
- [7] Darya Martyniuk, Johannes Jung, and Adrian Paschke. Quantum Architecture Search: A Survey. In *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 1695–1706, September 2024. arXiv:2406.06210 [quant-ph].
- [8] Mateusz Ostaszewski, Lea M Trenkwalder, Wojciech Masarczyk, Eleanor Scerri, and Vedran Dunjko. Reinforcement Learning for Optimization of Variational Quantum Circuit Architectures. *arXiv*, 2021.
- [9] En-Jui Kuo, Yao-Lung L. Fang, and Samuel Yen-Chi Chen. Quantum Architecture Search via Deep Reinforcement Learning, April 2021. arXiv:2104.07715 [quant-ph].
- [10] Abhishek Sadhu, Aritra Sarkar, and Akash Kundu. A quantum information theoretic analysis of reinforcement learning-assisted quantum architecture search. *Quantum Machine Intelligence*, 6(2):49, August 2024.
- [11] 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, May 2022. arXiv:2010.10217 [quant-ph].
- [12] Junjian Su, Jiacheng Fan, Shengyao Wu, Guanghui Li, Sujuan Qin, and Fei Gao. Topology-driven quantum architecture search framework. *Science China Information Sciences*, 68(8):180507, July 2025.
- [13] Zhimin He, Maijie Deng, Shenggen Zheng, Lvzhou Li, and Haozhen Situ. Training-Free Quantum Architecture Search. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(11):12430–12438, March 2024.

- [14] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature Communications*, 10(1):3007, July 2019. arXiv:1812.11173 [quant-ph].
- [15] Xudong Lu, Kaisen Pan, Ge Yan, Jiaming Shan, Wenjie Wu, and Junchi Yan. QAS-Bench: Rethinking Quantum Architecture Search and A Benchmark. In *Proceedings of the 40th International Conference on Machine Learning*, pages 22880–22898. PMLR, July 2023.
- [16] Trong Duong, Sang T. Truong, Minh Tam, Bao Bach, Ju-Young Ryu, and June-Koo Kevin Rhee. Quantum Neural Architecture Search with Quantum Circuits Metric and Bayesian Optimization, June 2022. arXiv:2206.14115 [quant-ph].