Investigating Quantum Circuit Designs Using Neuro-Evolution

Devroop Kar
Rochester Institute of technology
Rochester NY, United States
dk7405@rit.edu

Daniel Krutz
Rochester Institute of technology
Rochester NY, United States
dxkvse@rit.edu

Travis Desell
Rochester Institute of technology
Rochester NY, United States
tjdvse@rit.edu

Abstract

Designing effective quantum circuits remains a central challenge in quantum computing, as circuit structure strongly influences expressivity, trainability, and hardware feasibility. Current approaches, whether using manually designed circuit templates, fixed heuristics, or automated rules, face limitations in scalability, flexibility, and adaptability, often producing circuits that are poorly matched to the specific problem or quantum hardware. This work presents the Evolutionary eXploration of Augmenting Quantum Circuits (EXAQC), an evolutionary approach to the automated design and training of parameterized quantum circuits (PQCs) which extends on strategies from neuroevolution and genetic programming. EXAQC jointly searches over gate types, qubit connectivity, parameterization, and circuit depth while respecting hardware and noise constraints. It supports both Qiskit and Pennylane libraries, allowing the user to configure allowed architectural components. Preliminary results demonstrate that circuits evolved on classification tasks are able to achieve over 90% accuracy on most of the benchmark datasets with a limited computational budget, and are able to emulate target circuit quantum states with high fidelity scores.

CCS Concepts

• **Computing methodologies → Search methodologies**; **Genetic programming**; • **Hardware → Quantum computation**; • **Computer systems organization → Quantum computing**.

Keywords

Quantum Computing, Neuroevolution, Quantum Machine Learning

ACM Reference Format:
Devroop Kar, Daniel Krutz, and Travis Desell. 2026. Investigating Quantum Circuit Designs Using Neuro-Evolution. In *Genetic and Evolutionary Computation Conference (GECCO Companion '26), July 13–17, 2026, San Jose, Costa Rica.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3795101.3805320

1 Introduction

Quantum machine learning (QML) has emerged as a promising paradigm for leveraging quantum computational principles—such as superposition, entanglement, and interference—to enhance learning and decision-making tasks [3, 6, 9, 18]. Variational quantum circuits (VQCs), which combine parameterized quantum circuits with classical optimization, form the backbone of many contemporary QML approaches [5, 8, 10] . Despite their appeal, the practical success of VQCs depends to a great degree on the architecture [1, 2, 4, 7, 12, 16], including gate types, connectivity, parameterization, and measurement strategy. Designing such architectures remains a major challenge, particularly as most approaches rely on manually crafted templates or shallow heuristic designs that may not generalize across tasks or datasets.

This work proposes an evolutionary framework for task-driven quantum circuit discovery that directly optimizes both circuit structure and parameters for supervised learning on classical datasets as well as emulating specific circuit behavior. The framework supports both PennyLane and Qiskit, where quantum circuits are represented as mutable genomes composed of parameterized and non-parameterized quantum gates, enabling the evolution of circuit depth, gate ordering, qubit connectivity, and entanglement patterns. Circuit parameters are optimized using gradient-based learning, while structural modifications are explored through evolutionary operators, resulting in a memetic evolutionary–variational training process.

We evaluate the proposed approach on multiple UCI benchmark datasets, including the Iris, Wine, Seeds, and Breast Cancer classification tasks. Classical features are embedded into quantum states via angle-based encodings, while predictions are obtained from designated readout qubits using marginal probability distributions. Through evolution, the framework discovers nontrivial circuit topologies that increasingly entangle input and output registers, leading to improved classification performance. Notably, expressive circuit structures emerge organically from the evolutionary process rather than being imposed through predefined entangling layers.

2 Methodology

Evolutionary Exploration of Augmenting Quantum Circuits (EXAQC) extends the EXAMM [15] neuroevolution and EXA-GP [13, 14] genetic programming frameworks to the quantum circuit setting, where candidate structures are constrained by qubit availability and gate placement. EXAQC jointly evolves circuit topology and gate parameters using mutation and crossover, while leveraging Lamarckian weight inheritance [11] to reduce retraining cost. In the current implementation, candidate circuits are evaluated in a distributed asynchronous framework, enabling efficient parallel exploration of the search space.

2.1 Genome Representation

Each circuit is represented by a genome consisting of a set of input qubits, a set of output qubits, and an ordered list of gates. Every gate stores its operation type, qubit assignments, trainable parameters, depth $d \in [0, 1]$, and an innovation number following NEAT-style historical markings [19]. This variable-length representation allows circuits to grow or shrink during evolution while preserving meaningful structural alignment during crossover.

ACM ISBN 979-8-4007-2488-6/2026/07
https://doi.org/10.1145/3795101.3805320

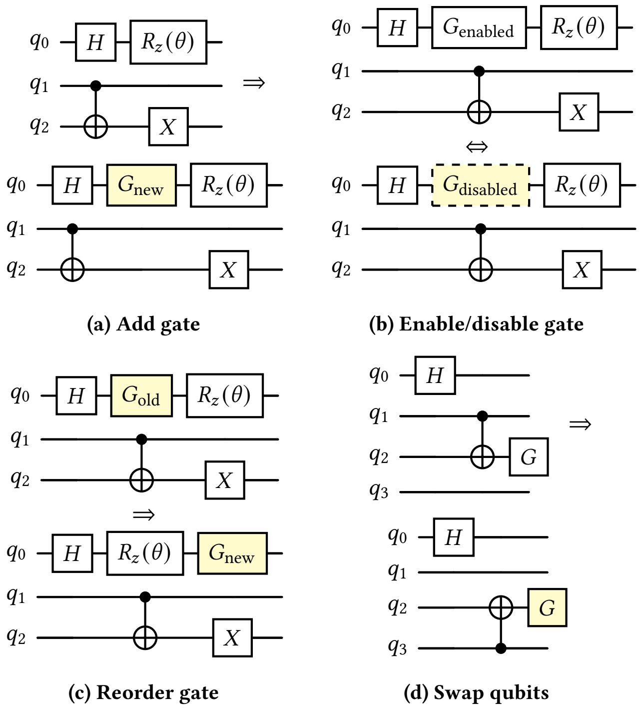

Figure 1: Structural mutation operators used in EXAQC. Highlighted gates indicate newly inserted or modified operations.

2.2 Mutation Operators

Mutation is the primary mechanism for structural exploration in EXAQC. Four mutation operators are used to modify circuit topology while preserving input–output connectivity. *Add gate* inserts a new gate at a sampled depth $d \sim U(0, 1)$, with qubits selected from the set of currently valid upstream and downstream qubits so that the new operation remains relevant to the computational graph. If the gate is parameterized, its parameters are initialized from $U(-\pi, \pi)$. *Enable/disable gate* toggles an existing gate, allowing the search to prune or reactivate structure without permanently deleting it. *Reorder gate* relocates a selected gate to a new depth, thereby changing the computational ordering of operations. *Swap qubits* reassigns one or more qubits used by a gate, allowing the search to explore alternative connectivity and entanglement patterns.

Because some mutations can disconnect inputs from outputs, every newly created child is validated by traversing the circuit in depth order to verify that at least one input influences at least one output. Invalid offspring are discarded. Together, these operators allow EXAQC to expand expressivity, alter gate placement, and refine qubit-level interactions while controlling unproductive growth.

2.3 Crossover

EXAQC uses three crossover operators: *binary crossover*, *n-ary crossover*, and *exponential crossover*. All of them exploit innovation numbers to identify homologous gates across parents.

Binary crossover. Binary crossover recombines two parents by always copying shared gates and probabilistically inheriting non-shared gates, with a stronger bias toward the fitter parent. If a gate appears only in the better parent, it is retained with probability p_b; if it appears only in the other parent, it is retained with probability p_o, where $p_b > p_o$. For shared parameterized gates, offspring parameters are inherited through randomized line interpolation:

$$p_{\text{new}} = p_{\text{other}} + r\,(p_{\text{best}} - p_{\text{other}}), \tag{1}$$

where r is randomly sampled. This preserves useful structure while allowing mild extrapolation in parameter space.

n-ary crossover. N-ary crossover generalizes binary crossover to multiple parents. One parent is designated as the primary parent and the remaining parents serve as auxiliary sources of structure. Gates shared between the primary parent and at least one auxiliary parent are favored, while gates appearing only in auxiliary parents are inherited with lower probability. For shared parameterized gates, offspring parameters are computed by combining the primary-parent value with the average value from the auxiliary parents:

$$p_{\text{new}} = p_{\text{avg}} + r\,(p_{\text{best}} - p_{\text{avg}}), \tag{2}$$

where p_{avg} is the average of the corresponding parameters across the auxiliary parents. This operator allows EXAQC to aggregate useful substructures from several lineages while still biasing inheritance toward the most fit parent.

Exponential crossover. Exponential crossover recombines two parents using a sampled crossover depth $d_c \sim U(0, 1)$. All gates from the first parent with depth $d < d_c$ are copied into the child, and all gates from the second parent with depth $d \geq d_c$ are appended. Unlike innovation-based recombination, this operator performs depth-wise structural splicing and is intended to preserve useful early- and late-stage circuit segments from different parents.

2.4 Objective Functions

Although EXAQC supports several loss functions, two objectives are central in this work. When a target circuit is available, we use a fidelity-based state loss. Given a predicted state $|\psi\rangle$ and target state $|\phi\rangle$, the fidelity is

$$F(\phi, \psi) = |\langle \phi \mid \psi \rangle|^2, \tag{3}$$

and the corresponding loss is

$$\mathcal{L}_{\text{fid}} = 1 - F(\phi, \psi). \tag{4}$$

This objective directly measures quantum state similarity and is well suited for teacher–student circuit matching, where the goal is to reproduce the behavior of a reference circuit under simulation.

For supervised learning tasks, circuit outputs are mapped to class probabilities over designated readout states and optimized using cross-entropy,

$$\mathcal{L}_{\text{CE}} = -\sum_{k=1}^{K} y_k \log p_k, \tag{5}$$

where p_k is the predicted probability of class k and y_k is the corresponding one-hot target. Cross-entropy provides strong discriminative pressure and naturally aligns the quantum model with standard classical classification objectives.

Overall, EXAQC combines variable-length circuit genomes, connectivity aware mutation, innovation-guided crossover, and complementary fidelity- and task-based objectives into a unified framework for evolutionary quantum circuit optimization.

3 Experimental Setup

We evaluate EXAQC on both classical supervised datasets and synthetic teacher quantum circuits, enabling assessment of task-driven learning and purely quantum structural imitation.

3.1 Benchmarks

We use standard UCI benchmarks spanning increasing dimensionality: Iris (4 features), Seeds (7), Wine (13), and Breast Cancer (30). These datasets test scalability of encoding, circuit depth, and robustness to class imbalance.

To isolate architectural learning, we evolve circuits to imitate fixed teacher circuits using state-level objectives. Benchmarks include: (i) single-gate transformations, (ii) Bell-state generators, (iii) input-controlled entangling circuits, and (iv) multi-layer circuits with parameterized rotations. These tasks emphasize entanglement learning and input–output quantum state mapping.

Supervised datasets evaluate classification performance, while teacher circuits evaluate the ability to recover quantum transformations independent of labels.

3.2 Experimental Settings

All experiments use a steady-state population of 50 genomes implemented in Pennylane. Initial populations are generated via mutation only. During evolution, mutation is applied with 70% probability and crossover (binary, n-ary with $n = 4$, and exponential) with 30% total probability. Mutation operators are weighted toward *add gate* (70%), with smaller contributions from reorder and qubit-swap (10% each) and enable/disable (5% each). Two mutations are applied per step to accelerate search. Each run evaluates 500 genomes using a distributed setup (12 processes). Parameterized gates are trained for 200 epochs using Adam with learning rate 10^{-3} and weight decay 10^{-4}.

4 Experimental Results

4.1 Classification Performance

Across all datasets (Table 1), EXAQC rapidly discovers high performing circuits within a small evaluation budget. Strong accuracy is achieved even with fewer than 500 evaluated genomes, indicating efficient exploration of the search space. Larger circuits generally yield better performance, suggesting that additional evolutionary steps could further improve results.

We also observe that weaker runs often converge early, indicating premature stagnation. In several high-performing solutions, some qubits are not connected to measured outputs, implying that improved connectivity or longer evolution could yield further gains.

4.2 Teacher Circuit Results

Results in Table 2 show that EXAQC achieves high fidelity across all teacher circuits. Simple transformations are learned almost immediately (within a few genomes), while more complex entangled circuits require deeper exploration. Fidelity decreases as circuit complexity increases, but remains above 90%, demonstrating strong capability in capturing quantum state transformations.

Dataset	Test Acc. %	# Gates	Genome #
Iris	86.7%	21	735
	70.0%	7	592
	83.3%	12	754
	70.0%	12	683
	90.0%	21	800
Seeds	92.9%	12	790
	92.9%	8	888
	95.2%	13	787
	90.5%	12	898
	92.9%	8	733
Wine	75.0%	11	625
	75.0%	15	804
	77.8%	12	876
	88.9%	16	851
	86.1%	15	734
Breast Cancer	91.2%	13	849
	89.2%	10	765
	90.8%	11	735
	88.6%	11	860
	91.0%	12	831

Table 1: Experimental results for each classification benchmark.

Circuit	Fidelity	Angular Distance
Baseline	100.00%	0.0
Bell-State Generator	98.32%	0.014
Input Controlled	94.11%	0.058
Multi-Layered	91.73%	0.101

Table 2: Fidelity and Angular Distance measure comparison against teacher circuits.

4.3 Discussion

From Table 1, EXAQC achieves strong classification performance with relatively few evaluated genomes, indicating efficient search; in a synchronous setting, these results would correspond to fewer than 10 generations. Performance is closely tied to mutation dynamics. Early evolution is dominated by *add gate* operations to establish input–output connectivity, while later stages benefit from refinement via reordering, qubit reassignment, and selective enabling/disabling. This suggests that adapting mutation operator probabilities over time could further improve results.

Compared to prior work such as Sünkel et al. [20], which restricts search to a small gate set and synthetic tasks, EXAQC operates over a broader gate space (Pennylane/Qiskit) and generalizes across both supervised and quantum objectives. Future work could further improve efficiency through adaptive gate selection. Results on teacher benchmarks (Table 2) show that EXAQC reliably captures quantum

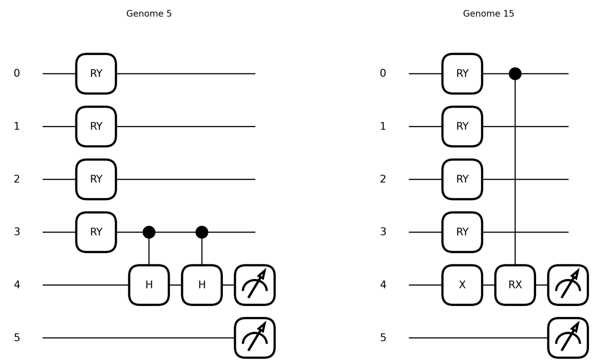

Figure 2: An example baseline identity gate as learned by two different genomes.

transformations with high fidelity. Notably, evolved circuits often achieve equivalent functionality through structurally different constructions (e.g., repeated Hadamards or redundant rotations in Figure 2), highlighting the flexibility and robustness of evolutionary search.

5 Conclusion

In this paper, we propose the Evolutionary eXploration of Augmenting Quantum Circuits (EXAQC) algorithm to design and train quantum circuits. Our methodology is generic, and can evolve circuits for both the Qiskit and Pennylane quantum frameworks. It is parallelized for improved performance and leverages advanced crossover and Lamarckian weight inheritance methodolgies adapted from neuroevolution and genetic programming strategies. Our results demonstrate that the evolved circuits can be learn to solve classical classification problems with improving accuracy and even generate quantum circuits to emulate pre-conceived circuits and their input output quantum with high fidelity. This framework is also not limited to any fixed set of gates and supports nearly all gates as supported by prevalent quantum computing libraries.

Results show the promise of the EXAQC methodology and lay the groundwork for significant future work. EXAQC allows for any quantum objective function to be plugged into the search process, which can allow circuits can be trained to potentially model any optimization problem. Expanding the library to include examples for reinforcement learning, time series forecasting and more complicated classification tasks (such as computer vision) will highlight its ability as a generic quantum circuit optimization framework. Further, EXAQC currently only utilizes a single steady state population and utilizes a single objective for optimization. Future work will involve investigating utilizing multiple islands and different speciation strategies to enhance optimization, as well as adapting these population strategies to incorporate multi-objective optimization to better utilize the variety of loss metrics EXAQC utilizes.

Acknowledgments

This work has been supported by the U.S. National Science Foundation under Grant Number #2225354. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. We also thank RIT's Research Computing team for their assistance and support [17].

References

[1] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. 2021. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature communications* 12, 1 (2021), 1791.
[2] Chih-Chieh Chen, Masaya Watabe, Kodai Shiba, Masaru Sogabe, Katsuyoshi Sakamoto, and Tomah Sogabe. 2021. On the expressibility and overfitting of quantum circuit learning. *ACM Transactions on Quantum Computing* 2, 2 (2021), 1–24.
[3] Giulio Chiribella, G Mauro D'Ariano, and Paolo Perinotti. 2008. Quantum circuit architecture. *Physical review letters* 101, 6 (2008), 060401.
[4] Yuxuan Du, Tao Huang, Shan You, Min-Hsiu Hsieh, and Dacheng Tao. 2022. Quantum circuit architecture search for variational quantum algorithms. *npj Quantum Information* 8, 1 (2022), 62.
[5] Harper R Grimsley, Sophia E Economou, Edwin Barnes, and Nicholas J Mayhall. 2019. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature communications* 10, 1 (2019), 3007.
[6] Lov K Grover. 1996. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*. 212–219.
[7] Thomas Hubregtsen, Josef Pichlmeier, Patrick Stecher, and Koen Bertels. 2021. Evaluation of parameterized quantum circuits: on the relation between classification accuracy, expressibility, and entangling capability. *Quantum Machine Intelligence* 3, 1 (2021), 9.
[8] Sami Khairy, Ruslan Shaydulin, Lukasz Cincio, Yuri Alexeev, and Prasanna Balaprakash. 2020. Learning to optimize variational quantum circuits to solve combinatorial problems. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 34. 2367–2375.
[9] Alexei Yu Kitaev, Alexander Shen, and Mikhail N Vyalyi. 2002. *Classical and quantum computation*. Number 47. American Mathematical Soc.
[10] Martin Lukac, Marek Perkowski, Hilton Goi, Mikhail Pivtoraiko, Chung Hyo Yu, Kyusik Chung, Hyunkoo Jeech, Byung-Guk Kim, and Yong-Duk Kim. 2003. Evolutionary approach to quantum and reversible circuits synthesis. *Artificial Intelligence Review* 20, 3 (2003), 361–417.
[11] Zimeng Lyu, AbdElRahman ElSaid, Joshua Karns, Mohamed Mkaouer, and Travis Desell. 2021. An Experimental Study of Weight Initialization and Lamarckian Inheritance on Neuroevolution. *The 24th International Conference on the Applications of Evolutionary Computation (EvoStar: EvoApps)* (2021).
[12] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. 2018. Barren plateaus in quantum neural network training landscapes. *Nature communications* 9, 1 (2018), 4812.
[13] Jared Murphy and Travis Desell. 2024. Minimizing the EXA-GP Graph-Based Genetic Programming Algorithm for Interpretable Time Series Forecasting. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 1686–1690.
[14] Jared Murphy, Devroop Kar, Joshua Karns, and Travis Desell. 2024. EXA-GP: Unifying Graph-Based Genetic Programming and Neuroevolution for Explainable Time Series Forecasting. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 523–526.
[15] Alexander Ororbia, AbdElRahman ElSaid, and Travis Desell. 2019. Investigating Recurrent Neural Network Memory Structures Using Neuro-evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference* (Prague, Czech Republic) *(GECCO '19)*. ACM, New York, NY, USA, 446–455. doi:10.1145/3321707.3321795
[16] Arthur Pesah, Marco Cerezo, Samson Wang, Tyler Volkoff, Andrew T Sornborger, and Patrick J Coles. 2021. Absence of barren plateaus in quantum convolutional neural networks. *Physical Review X* 11, 4 (2021), 041011.
[17] Rochester Institute of Technology. 2019. Research Computing Services. doi:10.34788/0S3G-QD15 Accessed 2026-01-23.
[18] Peter W Shor. 1999. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review* 41, 2 (1999), 303–332.
[19] Kenneth O Stanley and Risto Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary computation* 10, 2 (2002), 99–127.
[20] Leo Sünkel, Philipp Altmann, Michael Kölle, Gerhard Stenzel, Thomas Gabor, and Claudia Linnhoff-Popien. 2025. Quantum circuit construction and optimization through hybrid evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 934–942.