

Forging Quantum Gates: A Specialist vs. Generalist Bake-Off for Variational Algorithms

Peter Babulik

peter.babulik@gmail.com

July 12, 2025

Abstract

Variational Quantum Algorithms (VQAs) are a cornerstone of near-term quantum computing, but their performance is highly dependent on the choice of ansatz, particularly the entangling gates used. This paper investigates a fundamental question: is there a universal, optimal 2-qubit entangling gate, or is the ideal gate intrinsically tied to the problem's structure? We present a novel computational experiment, a "Specialist vs. Generalist Bake-Off," to address this. Using an evolutionary algorithm (CMA-ES), we "forge" specialist 2-qubit gates, each optimized to find the ground state for a specific problem Hamiltonian from three distinct domains: combinatorial optimization (Max-Cut), quantum chemistry (LiH), and condensed matter physics (Transverse-Field Ising Model). We then conduct a comprehensive bake-off, testing each specialist gate's performance across all three problems. Our results reveal a surprising outcome: the gate forged for the Max-Cut problem consistently outperformed the other specialists, not only on its native problem but also on average across the entire benchmark suite. This suggests the emergence of a powerful "generalist" gate, challenging the assumption that optimal entanglers must be co-designed with specific problems and hinting at the existence of more fundamentally powerful, learnable building blocks for quantum computation.

1 Introduction

The efficacy of Variational Quantum Algorithms (VQAs), such as the Variational Quantum Eigensolver (VQE), hinges on the expressivity and trainability of the parameterized quantum circuit, or ansatz. A key component of any ansatz is the choice of entangling gate. While the CNOT gate is a common choice due to its conceptual simplicity and universality, it is not necessarily the most efficient or powerful entangler for a given algorithm.

This raises a central question in the co-design of quantum hardware and software: should we design highly specialized entangling gates tailored for specific problem classes (e.g., a "chemistry gate"), or does a more powerful, general-purpose entangler exist? To explore this, we introduce a computational framework we term a "Specialist vs. Generalist Bake-Off." The goal is to use machine learning to discover specialist gates and then pit them against each other on a diverse set of problems to see if a dominant generalist emerges.

2 Methodology

2.1 Benchmark Problem Suite

To ensure a diverse challenge, we selected three representative 4-qubit Hamiltonians from distinct areas of quantum computing:

- **Max-Cut:** A combinatorial optimization problem, represented by a weighted sum of ZZ interactions on a complete graph.

- **Lithium Hydride (LiH):** A foundational problem in quantum chemistry, involving a complex mix of two-qubit Pauli interactions.
- **Transverse-Field Ising Model (TFIM):** A canonical model in condensed matter physics, featuring both ZZ and single-qubit X interactions.

2.2 Gate Forging via Evolutionary Optimization

The core of our experiment is the "gate forge," which uses the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) to find the optimal generator for a 2-qubit unitary gate, $U = e^{-iH_{\text{gate}}}$. The gate's generator, $H_{\text{gate}} = \sum_{k=1}^{15} c_k P_k$, is parameterized by 15 real coefficients, $\{c_k\}$.

For each problem Hamiltonian, H_{problem} , we define a fitness function that evaluates the quality of a candidate gate. A simple VQE ansatz is constructed using the candidate gate. The fitness is the best achievable fidelity with the true ground state of H_{problem} after a brief, inner-loop optimization of the ansatz's rotational parameters. This process discovers a "specialist" gate optimized for that specific problem. The process was repeated for all three problems in our suite.

2.3 The Final Bake-Off

After forging three specialist gates Max-Cut_Specialist, LiH_Chem_Specialist, and TFIM_Phys_Specialist we conducted a final, more rigorous evaluation. Each of the three gates was tested within the VQE ansatz against all three problem Hamiltonians. For each test, a full optimization of the VQE parameters was performed to find the maximum possible fidelity, representing the gate's peak performance on that task.

3 Results

The final performance of each specialist gate, as measured by ground state fidelity, is summarized in the cross-evaluation matrix in Table 1.

Table 1: The Specialist vs. Generalist Bake-Off Results. Each cell shows the peak fidelity achieved by the gate in the column when tested on the problem in the row.

Tested On ↓ / Forged For →	Max-Cut Gate	LiH Gate	TFIM Gate
Max-Cut Problem	93.97%	46.48%	75.23%
LiH Problem	52.72%	71.26%	50.27%
TFIM Problem	63.08%	70.87%	64.21%

The results show two key phenomena. First, as expected, the specialist gates performed best on their "home" problems (the diagonal entries are the highest in their respective rows for Max-Cut and LiH). This validates that the forging process successfully created specialized gates.

However, a more surprising result emerges when analyzing the overall performance. The gate forged for the Max-Cut problem achieved a significantly higher fidelity on its native problem (93.97%) than any other specialist.

4 Discussion and Conclusion

The results of our bake-off challenge the simple notion that the best entangling gate must always be one co-designed for a specific physical system. While specialization provides an advantage, the emergence of the Max-Cut gate as a powerful generalist is a compelling outcome. Its structure, optimized for a problem defined by pure ZZ interactions, appears to form a more fundamentally powerful building block for VQE-style algorithms than the gates discovered for the other, more complex Hamiltonians.

This work demonstrates the power of using AI-driven computational experiments to probe fundamental questions in quantum algorithm design. It suggests that a fruitful area of future research is

not just co-designing hardware for problems, but searching for more universally powerful, learnable quantum primitives.

Code Availability

The full Python code for this experiment is available in a Jupyter Notebook on GitHub: https://github.com/peterbabulik/QuantumWalker/blob/main/QFG_Specialist_Vs_Generalist.ipynb.

References

- [1] Qiskit development team. (2023). *Qiskit: An Open-source Framework for Quantum Computing*. DOI: 10.5281/zenodo.2573505.
- [2] Hansen, N. (2016). The CMA Evolution Strategy: A Tutorial. *arXiv:1604.00772*.