






Evaluating Parameter Transfer in FALQON Across Graph Families

Alisson dos Passos Fumaco¹, Marcos Vinicius Reballo¹,
Fernando Augusto Caletti de Barros¹, Gabriel Fernandes Thomaz¹,
Eduardo I. Duzzioni^{1,2}

¹Instituto de Pesquisas Eldorado – Porto Alegre – RS – Brazil

²Departamento de Física

Universidade Federal de Santa Catarina (UFSC), 88040-900 – Florianópolis, SC – Brazil

{alisson.passos,marcos.reballo,fernando.barros}@eldorado.org.br

{gabriel.thomaz,eduardo.duzzioni.BE}@eldorado.org.br

eduardo.duzzioni@ufsc.br

Abstract. We evaluate FALQON parameter transfer for Max-Cut, transferring sequences from small donors ($n \in \{8, 10, 12\}$) to 14-node recipients. Using 3-regular and Erdős–Rényi families, we show transfer success is dictated by the recipient graph, not the donor. Transfer excels for dense recipients—achieving high approximation ratios regardless of the donor—but remains challenging in sparse cross-family cases. Crucially, performance is highly resilient to donor size, with 8-node donors matching larger instances. Thus, cheap small graphs can provide robust parameters for larger targets, significantly reducing the measurement overhead of the feedback loop.

1. Introduction

Near-term quantum devices have several issues that prevent them from being used to execute algorithms applied to practical problems. In this context, Variational Quantum Algorithms (VQAs) have emerged as a leading paradigm for exploiting the currently available quantum devices, particularly in the context of combinatorial optimization and quantum simulation. Among these, the Quantum Approximate Optimization Algorithm (QAOA) is attracting substantial attention due to the demonstrated performance across a range of optimization problems encoded as Ising Hamiltonians through QUBO formulation [Farhi et al. 2014, Lucas 2014, Zhou et al. 2020]. At the same time, ideas from control theory have increasingly been used to design new quantum optimization strategies, emphasizing feedback, stability, and dynamical structure [Brif et al. 2010, Dong and Petersen 2013, Grivopoulos and Bamieh 2003].

A promising example of this convergence of quantum optimization and control theory is the Feedback-based Algorithm for Quantum Optimization (FALQON), which reformulates variational optimization as a closed-loop control problem [Magann et al. 2022a]. In this approach, instead of relying on classical gradient descent over a fixed-depth ansatz, FALQON iteratively constructs the quantum circuit by measuring a control observable extracted from the cost Hamiltonian and updating a feedback gain that governs the next layer of the circuit [Magann et al. 2022b]. In this way, it is

possible to completely eliminate the explicit classical optimization loop and provide a mechanism for driving the system toward low-energy states.

Despite the mentioned advantages, FALQON has a fundamental challenge shared with other depth-adaptive VQAs, which is: the current attempts to apply FALQON to real world problems resulted in high-depth circuits, giving the intuition that the depth seems to increase with the problem size and the desired solution quality. Meaning that, as the number of required qubits grows, the required number of measurements makes calculating the control parameters increasingly expensive. This challenge motivates the exploration of new methods and strategies that could improve the algorithm scalability for larger problems by reusing or transferring information across related problem instances to mitigate the optimization overhead by using smaller problems to compute the control parameters.

In FALQON, the analogue of variational angles is the sequence of feedback gains, which determine how the control Hamiltonian is applied at each iteration. These gains are not arbitrary optimization variables. Instead, they arise from measured expectation values and encode dynamical information linking the current quantum state to the cost Hamiltonian [Magann et al. 2022a]. From a control-theoretic perspective, they reflect system-specific response characteristics and therefore may generalize across problem families with shared structural features.

The central motivation of this work is to investigate whether such feedback parameters can be transferred between related FALQON instances – across increasing problem sizes or within ensembles of structurally similar Hamiltonians – in order to reduce the optimization time, suppress transient instability, and accelerate convergence. If successful, parameter transfer would significantly enhance the scalability and practicality of FALQON for large combinatorial problems.

In this work, we present a systematic study of parameter transferability in quantum optimization algorithms based on FALQON. Our main contributions are as follows:

1. **Problem-size transferability:** We analyze the reuse of FALQON control parameter learned on instances of size $n \in \{8, 10, 12\}$ when initializing or guiding optimization on larger instances with $n = 14$, quantifying performance degradation and convergence behavior.
2. **Hamiltonian family dependence:** We investigate transfer across canonical random graph ensembles, including random regular graphs $\text{RRG}(k)$ and Erdős-Rényi graphs $G(n, p)$, which are widely used benchmarks in QAOA and graph optimization studies [Bollobás 2001, Herrman et al. 2021].

The information presented in this work is organized in a sequence of 5 sections, starting with **Section 2**, in which a summary of the current state-of-the-art of Variational and Feedback-based algorithms is presented, the landscape of Parameter Transfer experiments on Quantum Optimization, and the known challenges faced when scaling and generalizing Quantum Algorithms for practical problems. Then, in **Section 3**, the base information required to proceed with the discussions brought by the experiments conducted in this work is presented, describing the FALQON formalism and specific characteristics of Max-Cut problems mapped to Hamiltonians, along with the specific definitions of random graph ensembles considered in the experiments. After that, in **Section 4**, the

framework used for the investigation and experiments is described, along with the specific experiments chosen to evaluate the characteristics of Parameter Transfer used with FALQON in different scenarios. In **Section 5**, the achieved results are presented, along with all the specific information related to each execution and all the findings brought by each of them. After presenting the experiments, the results are explored, analyzing the Problem-size Transferability and the Hamiltonian Family dependence on different scenarios. By the end, in **Section 6**, the main findings of the work are presented, providing more information about transfer quality across recipient regimes, findings about the transfer effectiveness resilience to the size gap between donor and recipient graphs, and showing that Parameter Transfer is a promising strategy to enhance FALQON-based optimization.

2. Related Work

Hybrid quantum algorithms have emerged as leading candidates for exploiting near-term quantum devices in combinatorial optimization. Among these approaches, *variational quantum algorithms* (VQAs) combine shallow parameterized circuits with classical optimization routines. Despite their conceptual appeal, VQAs face well-known challenges that hinder scalability and practical robustness, such as the barren plateau phenomenon and hardware-induced noise [Cerezo et al. 2021, McClean et al. 2018]. Motivated by these limitations, alternative strategies leveraging ideas from quantum control and measurement-based feedback have been proposed.

Accordingly, this section reviews related work in three main areas: the formulation and properties of variational and feedback-based algorithms (Section 2.1); parameter transfer techniques exploiting structural regularities across problem instances (Section 2.2); and strategies addressing scalability, generalization, and optimization stability (Section 2.3).

2.1. Variational & Feedback-Based Algorithms

Variational quantum algorithms constitute a central paradigm for near-term quantum optimization, with the *Quantum Approximate Optimization Algorithm* (QAOA) as a canonical example [Farhi et al. 2014, Zhou et al. 2020]. QAOA alternates between cost and mixer Hamiltonians in a parameterized circuit, relying on a classical outer loop to navigate non-convex energy landscapes. However, practical performance is often limited by challenging parameter optimization, the presence of barren plateaus, and sensitivity to noise, particularly at increasing circuit depths.

To mitigate these limitations, algorithms inspired by Lyapunov control introduce deterministic, sequential parameter updates. A prominent example is the *Feedback-Based Algorithm for Quantum Optimization* (FALQON), which assigns circuit parameters iteratively using expectation values of commutators between the problem and driver Hamiltonians [Magann et al. 2022a]. This construction guarantees a monotonic decrease of the cost Hamiltonian expectation value and eliminates the need for a classical optimizer, shifting the focus to locally informed, dynamics-driven parameter assignment while retaining a circuit structure closely related to QAOA.

2.2. Parameter Transfer in Quantum Optimization

A growing body of work has shown that optimal parameters of variational quantum algorithms often exhibit significant structure, including smoothness and concentration

around characteristic values for classes of related problem instances. In the context of QAOA, this has led to *parameter transfer*, where parameters optimized for a given instance are reused as initializations or fixed schedules for structurally similar problems. Early studies demonstrated successful transfer between graph instances sharing local structural properties, such as degree distribution, with only minor degradation in performance [Galda et al. 2023].

Parameter reuse across related problem instances has been extensively studied. For instance, near-optimal parameters for weighted Max-Cut have been shown to transfer across graph sizes and weight patterns with minimal performance degradation using suitable ansatzes and rescaling heuristics [Shaydulin et al. 2023], and related data-driven approaches exploit previously optimized instances to avoid re-optimization [Self et al. 2021, Jing et al. 2023]. Subsequent work leveraged graph representation learning to systematically identify suitable donor instances, further improving robustness [Falla et al. 2024]. These approaches leverage the observation that variational parameters implicitly encode structural properties of the cost Hamiltonian [Wang et al. 2018, Herrman et al. 2021]. Theoretical analyses have complemented these findings by identifying symmetries and transferable domains in the QAOA parameter landscape, clarifying when and why parameter reuse is effective [Lyngfelt and García-Álvarez 2025].

In contrast, existing approaches to parameter reuse in feedback-based optimization rely on learning-mediated strategies. Pérez *et al.* trained classical machine learning models to predict FALQON parameter schedules, reporting accurate performance for instances close to the training distribution but reduced generalization to larger graphs [Pérez et al. 2026]. This motivates the exploration of feedback-inspired optimization schemes that enable direct and scalable parameter transfer without reliance on machine learning models.

2.3. Scalability and Generalization in Quantum Algorithms

Scalability and generalization pose fundamental challenges for near-term quantum optimization algorithms. Although increasing circuit depth can improve solution quality, deeper variational circuits exacerbate optimization difficulty and sensitivity to noise, often leading to unstable training dynamics.

One successful strategy to mitigate these effects is *warm-starting*, where quantum optimization algorithms are initialized using approximate solutions obtained from classical heuristics or relaxations. In the context of QAOA, warm-starting significantly enhances performance at low depths by biasing the initial state toward promising regions of the solution space [Egger et al. 2021], with the classical approximation quality strongly correlating with quantum performance [Okada et al. 2023].

Beyond initialization, several works have examined the stability of variational optimization under realistic noise and finite sampling, showing that noise can significantly distort optimization landscapes and hinder convergence. These observations reinforce the motivation for optimization schemes that reduce reliance on iterative classical optimization and instead exploit structural priors.

Taken together, these works position FALQON as a compelling alternative to fully variational optimization for near-term quantum devices. By replacing global classical optimization with feedback-driven, dynamics-informed parameter updates, FALQON miti-

gates key limitations of variational quantum algorithms, including optimization instability and noise sensitivity [Magann et al. 2022a, Magann et al. 2022b], while opening new avenues for robust parameter transfer.

3. Preliminaries

3.1. FALQON formalism

FALQON [Magann et al. 2022a] is governed by a time-dependent Schrödinger equation, $i\frac{d}{dt}|\psi(t)\rangle = (H_C + \beta(t)H_M)|\psi(t)\rangle$, where the cost Hamiltonian H_C encodes the problem, the mixer Hamiltonian H_M drives transitions, and $\beta(t)$ is a real-valued control signal.

FALQON employs a Lyapunov control strategy to monotonically decrease the expected cost $V(t) = \langle\psi(t)|H_C|\psi(t)\rangle$. Differentiating $V(t)$ yields $\frac{d}{dt}V(t) = A(t)\beta(t)$, where $A(t) = \langle\psi(t)|i[H_M, H_C]|\psi(t)\rangle$. To guarantee $\frac{d}{dt}V(t) \leq 0$, FALQON sets the feedback law to $\beta(t) = -A(t)$ [Brif et al. 2010, Dong and Petersen 2013, Magann et al. 2022a].

For gate-based quantum hardware, this continuous evolution is discretized into L steps of duration Δt . The circuit is approximated by alternating unitaries $U_C = \exp(-iH_C\Delta t)$ and $U_M(\beta_k) = \exp(-i\beta_k H_M\Delta t)$, yielding the overall unitary

$$U(\boldsymbol{\beta}) = \prod_{k=1}^L U_M(\beta_k)U_C. \quad (1)$$

At each step k , the commutator expectation $A_k = \langle\psi_k|i[H_M, H_C]|\psi_k\rangle$ is measured to determine the next control parameter via the update rule $\beta_{k+1} = -A_k$. This ensures the discrete-time cost decreases monotonically across layers, subject to finite sampling noise [Magann et al. 2022a].

3.2. Max-Cut and random graph ensembles

The Max-Cut problem on an undirected, unweighted graph $G = (V, E)$, seeks a vertex bipartition that maximizes the number of crossing edges. By assigning spin variables $z_i \in \{+1, -1\}$ to each vertex, the contribution of an edge (i, j) to the cut is $\frac{1}{2}(1 - z_i z_j)$. Promoting these classical spins to Pauli Z operators yields the cost Hamiltonian [Lucas 2014]:

$$H_C = \frac{1}{2} \sum_{(i,j) \in E} (1 - Z_i Z_j). \quad (2)$$

The ground state of H_C encodes the maximum cut of G , and this operator serves as the cost Hamiltonian in our experiments.

Following standard practice in the literature, we consider two random graph ensembles as target instances [Farhi et al. 2014, Wang et al. 2018, Herrman et al. 2021, Magann et al. 2022a, Bollobás 2001]:

- **3-Regular Graphs ($\mathcal{G}_{n,3}$):** Graphs sampled uniformly where every vertex has degree exactly three, providing highly symmetric test cases.
- **Erdős–Rényi Graphs ($G(n, p)$):** Graphs where each of the $\binom{n}{2}$ possible edges exists independently with probability p , allowing us to probe performance across a broad range of densities.

4. Methodology

We investigate parameter transfer in FALQON using a donor–recipient protocol, where a parameter sequence $\beta^{(D)} = (\beta_1^{(D)}, \dots, \beta_L^{(D)})$ optimized on a smaller *donor* (D) graph is directly applied to a moderately larger *recipient* (R) graph. We employ a strict one-to-one layer mapping, setting $\beta_t^{(R)} = \beta_t^{(D)}$ for $t = 1, \dots, L$. This position-wise transfer preserves the temporal structure of the learned control sequence without any layer re-indexing, interpolation, or post-transfer adjustments.

To assess how well parameters generalize across sizes and structural connectivities, donor graphs are generated with sizes $n \in \{8, 10, 12\}$, while recipient graphs are fixed at $n' = 14$. We evaluate transfers within and across two random graph families: 3-regular graphs and Erdős–Rényi graphs $G(n, p)$. For the Erdős–Rényi family, donor instances use edge probabilities $p \in \{0.2, 0.3, 0.4, 0.5\}$, whereas recipient instances span $p \in \{0.2, 0.3, \dots, 1.0\}$. This comprehensive evaluation allows us to isolate the effects of donor size and structural similarity on transfer quality.

For each configuration, the transferred parameters are evaluated on 10 independently sampled recipient graphs, and the results are averaged. We assess the viability of the transfer using the approximation ratio, which measures the solution’s cut value relative to the exact Max-Cut optimum. To quantify how well the control behavior is preserved, the mean approximation ratio achieved on the recipient instances is directly compared against the reference performance obtained on the original donor graph.

All quantum algorithmic evaluations are conducted via exact state-vector simulation. On real quantum devices, the sequential nature of FALQON requires a continuous classical feedback loop to estimate the commutator expectation values. This loop faces significant challenges, such as finite sampling noise, readout errors, and communication latency between the quantum processor and classical hardware. Therefore, using exact simulation ensures that the feedback dynamics are evaluated deterministically, isolating the fundamental transferability of the parameter sequences from these hardware-induced overheads and decoherence. Furthermore, the maximum recipient size of $n' = 14$ serves two crucial purposes. First, it allows us to exactly compute the global Max-Cut optimum for every generated instance using classical brute-force evaluation, establishing an absolute performance baseline devoid of heuristic approximations.

5. Results and Discussion

We organize the analysis around three empirical patterns observed for 14-node recipient graphs. The discussion focuses on the relationship between the transferred approximation-ratio trajectories and the donor-side training trajectory to characterize whether FALQON parameter sequences preserve, degrade, or improve after transfer.

Transfer is more effective for dense Erdős–Rényi recipient graphs. As shown in Figure 1, transferred parameter sequences achieve high approximation ratios on dense 14-node Erdős–Rényi recipients, frequently matching or exceeding the donor-side training trajectory. This stability is most pronounced for recipient densities $p \in \{0.8, 0.9, 1.0\}$. In these cases, the final approximation ratio reaches approximately 0.95 for $p = 0.8$, 0.96 for $p = 0.9$, and 0.98 for the complete graph ($p = 1.0$). Crucially, this high performance

is achieved with remarkably low variance across the recipient instances, demonstrating that transfer to structurally homogeneous, dense graphs is not only effective but highly reliable.

This behavior can be interpreted by translating the quantum dynamics of the FALQON commutator into the structural properties of the target graphs. In FALQON, the commutator expectation value $\langle i[H_M, H_C] \rangle$ acts as a feedback signal—analogueous to a gradient in classical optimization—that dictates the next parameter update. For highly dense graphs, the adjacency matrix is heavily populated, meaning local structural irregularities (such as specific missing edges) have a vanishingly small impact on the overall cost function. The optimization landscape becomes uniform, governed by the average degree of the nodes rather than specific local connectivities. Because of this structural homogeneity, the feedback signals generated by a small dense donor proportionally mirror the macroscopic behavior of a larger dense recipient. Consequently, the control parameters learned on the smaller instance remain optimal for the larger one, bypassing the need for instance-specific recalculation.

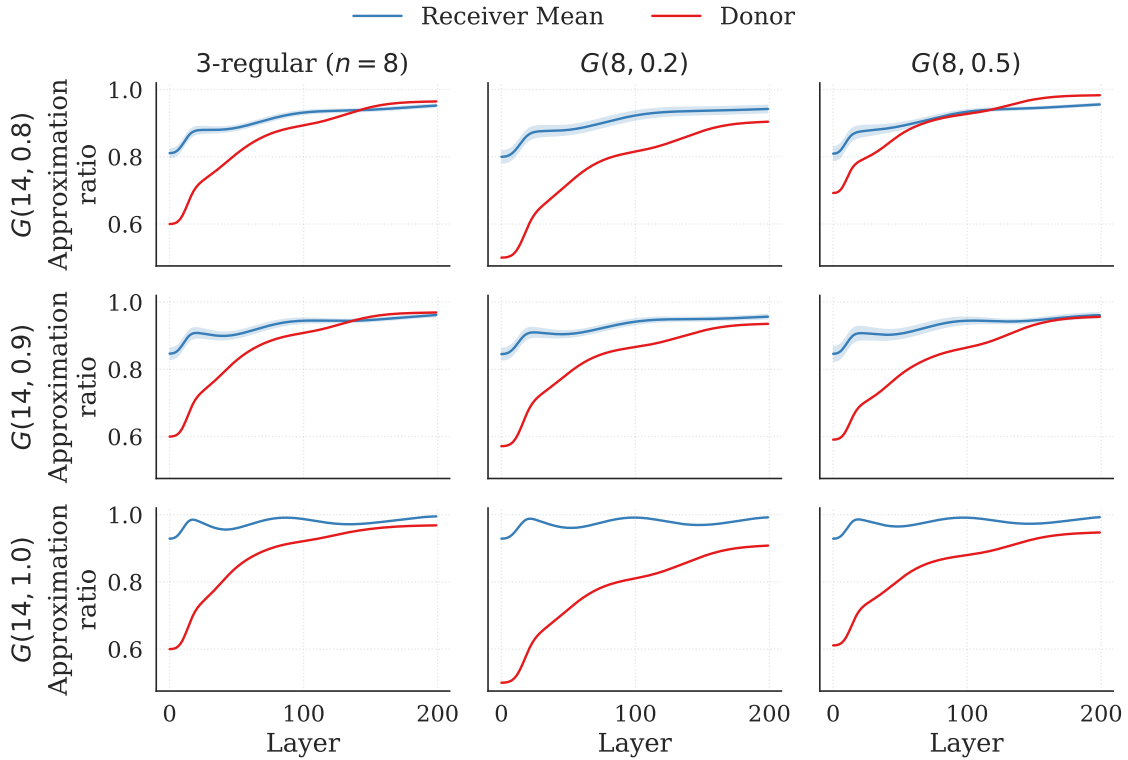


Figure 1. Representative transfer trajectories for dense 14-node Erdős–Rényi recipients ($p \in \{0.8, 0.9, 1.0\}$) from 8-node donor graphs. The transferred mean approximation-ratio curves (blue) exhibit narrow variance and monotonic growth, consistently matching or exceeding the donor training trajectories (red).

Sparse cross-family transfer is substantially harder. A contrasting pattern appears in cross-family transfers from 3-regular donor graphs to sparse 14-node Erdős–Rényi recipients (Figure 2). For highly sparse recipients ($p = 0.2$), transferred curves consistently

remain below the donor-side reference. As recipient density increases to $p = 0.3$ and $p = 0.4$, this performance gap narrows. This sensitivity indicates that highly sparse target graphs present greater structural variability and a stronger mismatch with 3-regular donors. Consequently, sparse cross-family settings act as a stress test, revealing that transfer quality depends heavily on the structural compatibility between source and target Hamiltonians.




Figure 2. Representative transfer trajectories for sparse 14-node Erdős–Rényi recipients ($p \in \{0.2, 0.3, 0.4\}$) from 3-regular donor graphs of sizes 8, 10, and 12. The transferred curves (blue) remain below the donor references (red) for $p = 0.2$, with the gap narrowing as recipient density increases.

Transfer performance is resilient to donor size. The transfer behavior illustrated in Figure 3 demonstrates significant resilience across donor sizes. For both 3-regular and Erdős–Rényi ($p = 0.5$) training families, the final approximation ratios for 14-node recipients remain stable across $n \in \{8, 10, 12\}$. While mean performance slightly fluctuates across these donor sizes, the substantial overlap in their standard deviations indicates no statistically significant advantage to using larger donor instances. This confirms that the smallest, most computationally efficient 8-node donors capture control features robust enough for effective reuse, yielding performance statistically indistinguishable from larger donors.

Taken together, these results indicate that parameter transfer in FALQON is primarily governed by the structural properties of the recipient graph rather than the donor size. Transfer quality and reliability improve significantly for denser recipient graphs, while sparse cross-family settings pose the greatest challenge. Crucially, the ro-

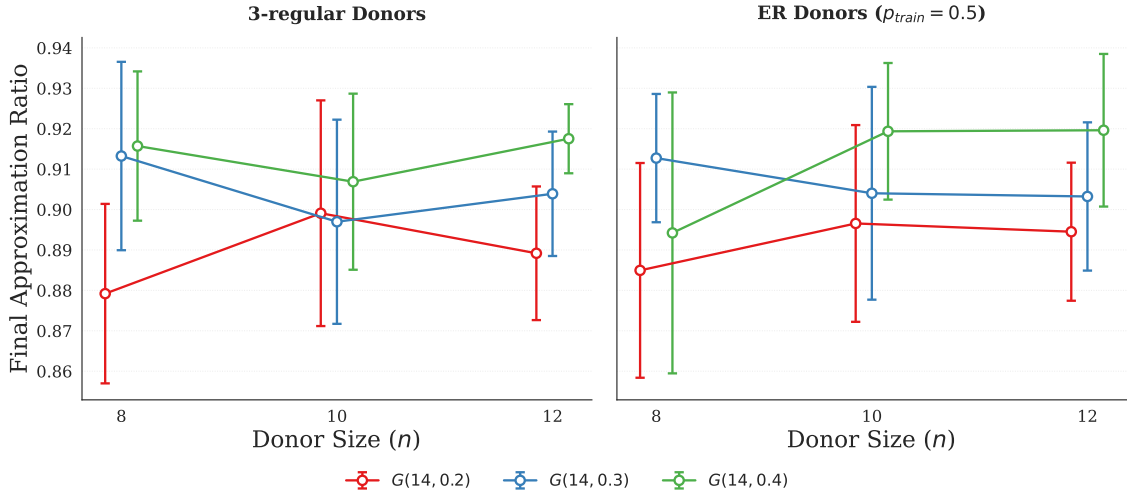


Figure 3. Comparison of transfer performance across donor sizes for (a) 3-regular and (b) Erdős–Rényi donors (trained at $p = 0.5$). The substantial overlap in error bars indicates that transfer quality is resilient to donor size, achieving highly competitive and statistically comparable performance even with 8-node donors.

bust performance of 8-node donors—validated by overlapping variance bounds across scales—suggests that FALQON learns scalable control sequences independent of specific training instance sizes. This supports the viability of leveraging small, easily optimized donors to accelerate optimization on larger target Hamiltonians.

6. Conclusion

This work presented an empirical study of parameter transfer in FALQON for Max-Cut on random graphs. We investigated whether parameter sequences learned on smaller donor graphs can be reused on moderately larger 14-node recipient graphs across different graph families and densities. The empirical results reveal three main behaviors. First, transferring to an easier problem instance yields strong results: as the Erdős–Rényi recipient approaches a complete graph, parameters from virtually any donor achieve high approximation ratios. This is likely because cutting a nearly complete graph is an inherently easier Max-Cut instance. Second, transferring to a harder problem results in degraded performance: sparse cross-family transfer from 3-regular donors to highly sparse Erdős–Rényi recipients remains substantially challenging.

Third, the success of the transfer is primarily dictated by the recipient graph, with the donor having minimal impact. Transfer performance is remarkably resilient to the size gap; parameter sequences learned on smaller 8-node donor graphs transfer competitively to 14-node recipients when compared to those learned on 10- and 12-node donors. This resilience indicates that the specific scale of the donor instance is a secondary factor. Ultimately, if the target graph is an easy instance, the transfer will succeed regardless of the donor.

Overall, these findings indicate that parameter transfer is a highly promising strategy for reducing the cost of FALQON-based optimization, as computationally cheap, small donor instances can yield effective parameter sequences for larger targets. As fu-

ture work, it would be valuable to investigate the absolute limits of this scale invariance and whether some form of parameter normalization or rescaling can extend transferability across much more extreme donor–recipient size gaps. Additional directions include extending the analysis to significantly larger graph sizes compared against standard heuristic baselines, evaluating transfer under noisy hardware constraints, and explicitly investigating how the spectral properties and the energy landscape of the problem Hamiltonians dictate transferability bounds, with a particular focus on understanding and overcoming the limitations observed in sparse cross-family regimes.

Acknowledgements

This work was executed under the TIC26 – Brazil Quantum Camp project, funded within the scope of the Prioritized Informatics Programs and Projects (PPI), Process No. 01245.008254/2025-22, under the responsibility of the Ministry of Science, Technology and Innovation (MCTI), with operational coordination by the Association for the Promotion of Brazilian Software Excellence (SOFTEX), and executed by CESAR and the Instituto de Pesquisas ELDORADO.

References

- Bollobás, B. (2001). *Random Graphs*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge.
- Brif, C., Chakrabarti, R., and Rabitz, H. (2010). Control of quantum phenomena: past, present and future. *New Journal of Physics*, 12(7):075008.
- Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., and Coles, P. J. (2021). Variational quantum algorithms. *Nature Reviews Physics*, 3:625–644.
- Dong, D. and Petersen, I. R. (2013). A survey of quantum lyapunov control methods. *Annual Reviews in Control*, 37(1):1–17.
- Egger, D. J., Mareček, J., and Woerner, S. (2021). Warm-starting quantum optimization. *Quantum*, 5:479.
- Falla, J., Langfitt, Q., Alexeev, Y., and Safro, I. (2024). Graph representation learning for parameter transferability in the quantum approximate optimization algorithm. *Quantum Machine Intelligence*, 6(46).
- Farhi, E., Goldstone, J., and Gutmann, S. (2014). A quantum approximate optimization algorithm. *arXiv:1411.4028*.
- Galda, A., Gupta, E., Falla, J., Liu, X., Lykov, D., Alexeev, Y., and Safro, I. (2023). Similarity-based parameter transferability in the quantum approximate optimization algorithm. *Frontiers in Quantum Science and Technology*, 2:1200975.
- Grivopoulos, S. and Bamieh, B. (2003). Lyapunov-based control of quantum systems. *Proceedings of the 42nd IEEE Conference on Decision and Control*, pages 434–438.
- Herrman, R., Treffert, L., Ostrowski, J., Lotshaw, P. C., Humble, T. S., and Siopsis, G. (2021). Impact of graph structures for qaoa on maxcut. *arXiv:2102.05997*.
- Jing, H., Wang, Y., and Li, Y. (2023). Data-driven quantum approximate optimization algorithm for power systems. *Communications Engineering*, 2(1):12.

- Lucas, A. (2014). Ising formulations of many np problems. *Frontiers in Physics*, 2:5.
- Lyngfelt, I. and García-Álvarez, L. (2025). Symmetry-informed transferability of optimal parameters in the quantum approximate optimization algorithm. *Physical Review A*, 111:022418.
- Magann, A. B., Rudinger, K. M., Grace, M. D., and Sarovar, M. (2022a). Feedback-based quantum optimization. *Physical Review Letters*, 129:250502.
- Magann, A. B., Rudinger, K. M., Grace, M. D., and Sarovar, M. (2022b). Lyapunov-control-inspired strategies for quantum combinatorial optimization. *Physical Review A*, 106:062414.
- McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., and Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1).
- Okada, K. N., Nishi, H., Kosugi, T., and Matsushita, Y.-i. (2023). Systematic study on the dependence of the warm-start quantum approximate optimization algorithm on approximate solutions. *Scientific Reports*, 13:50406.
- Pérez, V. P., Grace, M. D., Arenz, C., and Magann, A. B. (2026). Learning parameter curves in feedback-based quantum optimization algorithms. *arXiv preprint arXiv:2601.08085*.
- Self, C. N., Khosla, K. E., Smith, A. W. R., Sauvage, F., Haynes, P. D., Knolle, J., Mintert, F., and Kim, M. S. (2021). Variational quantum algorithm with information sharing. *npj Quantum Information*, 7(1):116.
- Shaydulin, R., Lotshaw, P. C., Larson, J., Ostrowski, J., and Humble, T. S. (2023). Parameter transfer for quantum approximate optimization of weighted maxcut. *ACM Transactions on Quantum Computing*, 4(3):1–15.
- Wang, Z., Hadfield, S., Jiang, Z., and Rieffel, E. (2018). Quantum approximate optimization algorithm for maxcut: A fermionic view. *Physical Review A*, 97(2):022304.
- Zhou, L., Wang, S.-T., Choi, S., Pichler, H., and Lukin, M. D. (2020). Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Physical Review X*, 10:021067.