

Topic: Quantum Circuit Optimization via Hierarchical Reinforcement Learning

Motivation / Related Work

Optimizing quantum circuits is essential for various applications, such as chemical simulations, due to the inherent complexity of quantum systems and the current limitations of quantum hardware. Effective quantum simulation demands careful attention to key aspects of circuit optimization, including the reduction of gate depth and gate count to minimize errors and decoherence, as well as the efficient mapping of quantum algorithms onto available hardware.

Reinforcement learning (RL) has emerged as a powerful tool for quantum circuit optimization, leveraging its capability to navigate and exploit complex solution spaces to enhance circuit efficiency. By treating the optimization problem as a sequential decision-making process, RL algorithms iteratively refine circuit designs through trial and error, guided by a reward signal that reflects performance metrics such as gate depth, gate count, and fidelity [1-3].

This methodology enables RL to uncover novel optimization strategies that traditional approaches may overlook, ultimately leading to more efficient quantum computations and improved use of quantum hardware resources.

Overall goal

The goal of this work is to train an agent via RL for the optimization of quantum circuits for chemical simulations with respect to a defined metric as reward, e.g., gate depth and gate count. The existing environment is a quantum circuit simulator that is initialized with randomly generated circuits, which the agent will then alter and optimize.

The agent will operate within a discrete, hierarchically structured action space. First, it must select a location within the quantum circuit, and then determine the appropriate modification to apply at that location. To achieve this, the RL framework detailed in [4] will be employed, which trains agents to navigate hierarchical action spaces using an adapted version of Gumbel AlphaZero's [5] action selection process.

The primary task of this work involves integrating the approach from [4] with the existing quantum circuit optimization environment. This includes identifying and implementing a suitable neural network architecture (e.g., a graph neural network or a transformer) that is capable of handling variable-sized, discrete action spaces. Finally, the performance of the RL agent will be evaluated through a comparative study against other quantum circuit optimization techniques. The study should also contain an outlook on the integration of continuous actions (e.g., the rotation angle of the gates) into the framework.

Timetable (6 months, in 24 person weeks (PW))

- 3 PW Familiarization with relevant work in the subject areas and literature.
- 3 PW Conceptualization of a suitable neural network architecture for the agent.
- 8 PW Implementation of the approach in Python. The RL environment and the code for the action selection from [4] is provided to the student.
- 4 PW Experiments and refinement of the method.
- 6 PW Writing of the final transcript.

Advisors: Maniraman Periyasamy (M. Sc.) (IIS), Dr.-Ing. Quirin Göttl (IIS), Dr.-Ing. Christopher Mutschler (IIS),
Supervising Professor Prof. Dr. Björn Eskofier (FAU)

Student: Haris Asif

Start—End: 01.09.2024—28.02.2025

References

- [1] Kundu, A., Sarkar, A., & Sadhu, A., 2024. KANQAS: Kolmogorov-Arnold Network for Quantum Architecture Search. arXiv:2406.17630v2.
- [2] Ostaszewski, M., Trenkwalder, L.M., Masarczyk, W., Scerri, E., & Dunjko, V., 2021. Reinforcement learning for optimization of variational quantum circuit architectures. Advances in Neural Information Processing Systems 34 (NeurIPS).
- [3] Preti, F., Schilling, M., Jerbi, S., Trenkwalder, L., Nautrup, H., Motzoi, F., & Briegel, H., 2024. Hybrid discrete-continuous compilation of trapped-ion quantum circuits with deep reinforcement learning. Quantum, 8, 1343.
- [4] Göttl, Q., Asif, H., Mattick, A., Marzilger, R., & Plinge, A., 2024. Automated Design in Hybrid Action Spaces by Reinforcement Learning and Differential Evolution. 47th German Conference on Artificial Intelligence.
- [5] Danihelka, I., Guez, A., Schrittwieser, J., & Silver, D., 2022. Policy improvement by planning with Gumbel. International Conference on Learning Representations (ICLR).