Discovering Fault-Tolerant Quantum Circuits and Quantum Error Correction Codes via Reinforcement Learning

Remmy Zen[a], **Jan Olle**[a], **Florian Marquardt**[b]

[a] *Max Planck Institute for the Science of Light, Staudtstraße 2, 91058 Erlangen, Germany* remmy.zen@mpl.mpg.de

[a] *Department of Physics, Friedrich-Alexander Universität Erlangen-Nürnberg, Staudtstraße 5, 91058 Erlangen, Germany*

* Presenting author

1. Introduction

Reinforcement Learning (RL) [1] is a framework designed to discover optimal action sequences in decision-making problems, where the optimal strategy is unknown and often nontrivial. In this work, we investigate the use of RL in an array of tasks related to Quantum Error Correction (QEC). These tasks include the automatic discovery of QEC codes from scratch in different noise channels simultaneously [2] and the automatic discovery of fault-tolerant circuits that prepare target logical states of qubits from a known QEC code [3]. We discuss how our RL-based strategy is a versatile tool for QEC across diverse quantum hardware platforms of interest.

2. Automatic Discovery of Noise-Aware QEC Codes

Numerical techniques have already been employed to construct QEC codes. Often, this has involved greedy algorithms, which may lead to sub-optimal solutions but can be relatively fast. For instance, in [4] a greedy algorithm was implemented to extend classical linear codes, code concatenation was explored in [5], and greedy search for finding stabilizer codes was used in [6]. Often, such numerical methods for QEC code construction were restricted to finding a subclass of codes of a particular structure, e.g. using reduction to classical code search [7].

However, knowledge of a code does not automatically translate to knowing how to encode the logical states of that code in an efficient way. The reason is that standard approaches are *unconstrained*, meaning that an all-to-all connectivity between qubits is assumed as well as the noise model and a set of gates that are not necessarily native to the hardware platform of interest. This then leads to larger-than-necessary circuits when implementing them on specific devices.

The main objective of this work [2] is to automatize the discovery of QEC codes and their encoding circuits using RL. We will consider a scenario where the encoding circuit is assumed to be error-free (non fault-tolerant encoding). This is applicable to quantum communication or quantum memories, where the majority of errors happen during transmission over a noisy channel or during the time the memory is retaining the information. Eventually we will show that it is possible to train a meta-agent that is capable of adapting its strategy according to the noise model, without any retraining. This leverages the concept of transfer learning, where improvements gained in training for one scenario (here, one value of a noise parameter) carry over and accelerate the training progress for other scenarios.

The task of the RL agent is to discover a suitable encoding sequence of gates for the particular error model under consideration. The agent receives a representation of the quantum circuit and the noise model as input (as its observation) and suggests the next gate (action) to apply. In this way, an encoding circuit is built up step by step, taking into account the available gate set and connectivity for the particular hardware platform. The agent is rewarded based on the Knill-LaFlamme (KL) conditions [8]. This process terminates when the KL conditions are satisfied for the target error channel and the learned circuit can then be used to encode any state of choice.

In Fig. 1, we show how our reinforcement learning agent can discover codes and their encoding circuits with unbiased noise for various numbers of physical qubits, logical qubits, and distances with relatively small circuit size. We discover both known and novel quantum error correction codes. The RL agent can also adapt to different hardware constraints, such as different qubit connectivity, gate set, and noise model. In the case of biased noise, a noise-aware RL agent trained for different noise models outperforms noise-unaware RL agents trained for a given noise model without any training.

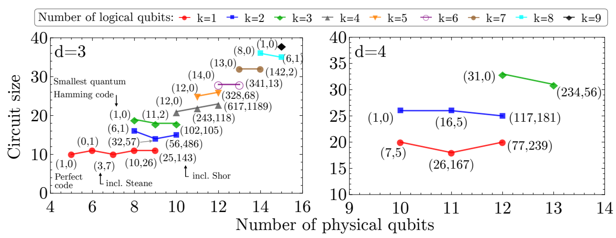

Fig. 1: Discovering codes and encoding circuits for various numbers of physical qubits, logical qubits, and distances.

3. Automatic Discovery of Fault-Tolerant Logical State Preparation

In the previous approach, we assume that the gates used to encode the logical state are perfect and the error only occurred in the communication channel. However, for quantum computation, these gates are faulty and multi-qubit gates proliferate the errors, compromising the scalability of QEC. Therefore, we want to minimize the number of possible faulty operations that can lead to harmful errors: this is achieved by designing *fault-tolerant* (FT) circuits [9]. Only by using FT schemes we can ensure systematic improvement in correction as the size of the code scales. Therefore, FT is of paramount importance in making scalable quantum computers [10, 11]. Several classes of FT protocols have been proposed.

Recently, *flag FT error correction* [12, 13, 14] was introduced as a way to achieve FT protocols with a minimal number of ancilla qubits. For instance, in the specific case of preparing a state FT, a flag FT protocol uses a *verification circuit* after the encoding circuit that utilizes a few extra ancilla qubits, known as *flag qubits*, to flag harmful errors (errors that the code cannot correct) while keeping the logical state intact. There are already examples of flag verification circuits in state preparation on several QEC codes [15, 16] and they have also been shown to be effective in reducing logical error rates in experimental realizations [17, 18]. Despite their success, flag-based protocols are typically handcrafted. Furthermore, flag-based protocols have so far been implemented in devices with all-to-all qubit connectivity. In other words, the automatic compilation of FT circuits has not been widely explored yet.

In this work [3], we present a novel approach to automatically discover hardware-adapted FT quantum circuits for QEC using RL. Hardware-adapted means that we can constrain the qubit connectivity and the available gate set based on the quantum platforms of interest, such that the discovered circuits can be directly realized in experiments. We apply our method to the task of logical state preparation, that is the automatic discovery of quantum circuits that fault-tolerantly prepare the target logical state of a given QEC code under given available gate set and qubit connectivity.

Similar to previous section, an RL agent is trained to output circuits (suggesting a sequence of gates). The input to the RL agent is the target logical state, the number of flag qubits, the available qubit connectivity and gate set. At each step, the RL agent observes the state of a quantum circuit and applies a gate to the quantum circuit as an action. All possible errors happening to the gate are tracked every time. The qubit connectivity and gate set determine the set of possible actions that the agent can take. The reward consists of three components: how far the state of the current quantum circuit to the target logical state, how many harmful errors are flagged by the flag qubits, and whether the data qubits and flag qubits are disentangled. The latter ensure that when we measure the flag qubits the state in the data qubits are not destroyed.

Table 1: The comparison of FT logical state preparation circuits on all-to-all qubit connectivity. We show the minimum number of two-qubit gates and the number of flag qubits in parentheses.

Code	State	RL	Existing
$[[5,1,3]]$ perfect	$\lvert 1\rangle_L$	**12 (2)**	-
	$\lvert -\rangle_L$	**12 (2)**	20 (6) [14]
$[[7,1,3]]$ Steane	$\lvert 0\rangle_L$	11 (1)	11 (1) [15]
	$\lvert +\rangle_L$	11 (1)	-
$[[9,1,3]]$ Shor	$\lvert 0\rangle_L$	6 (0)	-
	$\lvert +\rangle_L$	11 (1)	-
$[[9,1,3]]$ Surface-17	$\lvert 0\rangle_L$	11 (1)	-
	$\lvert +\rangle_L$	11 (1)	-
$[[15,1,3]]$ Reed-Muller	$\lvert 0\rangle_L$	**25 (1)**	**25 (1)** [20]
	$\lvert +\rangle_L$	**31 (1)**	32 (1) [20]

We apply our RL method to several known codes in all-to-all qubit connectivity, as shown in Table 1. We see that RL finds new circuits that were not known before, and always finds circuits with fewer or the same number of two-qubit gates as existing circuits. With restricted connectivity, we find for the first time the FT preparation of the $\lvert 0\rangle_L$ for the $[[7,1,3]]$ Steane code on a 3×3 grid based on the Google Sycamore device, the $\lvert 1\rangle_L$ for the $[[5,1,3]]$ perfect code on a 2D grid based on the IBMQ Tokyo device, and $\lvert 0\rangle_L$ of the $[[9,1,3]]$ Surface-17 code with the connectivity and qubit placement from [19].

4. Conclusion

This work presents a novel approach to automatically discover compact, hardware-adapted, and noise-aware quantum circuits for encoding quantum error correction codes and fault-tolerant logical state preparation. We employ reinforcement learning (RL) as an enabling tool, leveraging a fast parallelized stabilizer quantum circuit simulator and a nontrivial reward function specifically adapted to the problem. We first show that the RL agent can discover known and new quantum error correction codes with hardware constraints. We then show that in the task of fault-tolerant logical state preparation, RL discovers not only circuits with fewer gates and ancillary qubits than published results, but also novel circuits without and with hardware constraints. More generally, our work sets the framework towards the use of RL or other machine learning techniques for quantum circuit discovery with hardware constraints to make real progress towards the realization of large-scale quantum computers.

Acknowledgments

This research is part of the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus.

References

[1] Richard S. Sutton and Andrew G. Barto. *Reinforcement learning: an introduction*. Adaptive computation and machine learning series. The MIT Press, Cambridge, Massachusetts, second edition, 2018.

[2] Jan Olle, Remmy Zen, Matteo Puviani, and Florian Marquardt. Simultaneous discovery of quantum error correction codes and encoders with a noise-aware reinforcement learning agent. *npj Quantum Information*, 10(1):1–17, 2024.

[3] Remmy Zen, Jan Olle, Luis Colmenarez, Matteo Puviani, Markus Müller, and Florian Marquardt. Quantum circuit discovery for fault-tolerant logical state preparation with reinforcement learning. *arXiv:2402.17761*, 2024.

[4] Markus Grassl and Sunghyu Han. Computing extensions of linear codes using a greedy algorithm. In *2012 IEEE International Symposium on Information Theory Proceedings*, pages 1568–1572, 2012.

[5] Markus Grassl, Peter W. Shor, Graeme Smith, John Smolin, and Bei Zeng. New constructions of codes for asymmetric channels via concatenation. *IEEE Transactions on Information Theory*, 61(4):1879–1886, 2015.

[6] Muyuan Li, Mauricio Gutiérrez, Stanley E. David, Alonzo Hernandez, and Kenneth R. Brown. Fault tolerance with bare ancillary qubits for a [[7,1,3]] code. *Phys. Rev. A*, 2017.

[7] Isaac Chuang, Andrew Cross, Graeme Smith, John Smolin, and Bei Zeng. Codeword stabilized quantum codes: Algorithm and structure. *Journal of Mathematical Physics*, 2009.

[8] Emanuel Knill and Raymond Laflamme. Theory of quantum error-correcting codes. *Phys. Rev. A*, 1997.

[9] Earl T. Campbell, Barbara M. Terhal, and Christophe Vuillot. Roads towards fault-tolerant universal quantum computation. *Nature*, 549(7671):172–179, September 2017.

[10] Dorit Aharonov and Michael Ben-Or. Fault-Tolerant Quantum Computation with Constant Error Rate. *SIAM J. Comput.*, 38(4):1207–1282, January 2008.

[11] A.Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, January 2003.

[12] Rui Chao and Ben W. Reichardt. Flag Fault-Tolerant Error Correction for any Stabilizer Code. *PRX Quantum*, 1(1):010302, September 2020.

[13] Christopher Chamberland and Michael E. Beverland. Flag fault-tolerant error correction with arbitrary distance codes. *Quantum*, 2:53, February 2018.

[14] Rui Chao and Ben W. Reichardt. Quantum error correction with only two extra qubits. *Phys. Rev. Lett.*, 121(5):050502, August 2018.

[15] Hayato Goto. Minimizing resource overheads for fault-tolerant preparation of encoded states of the Steane code. *Sci Rep*, 6(1):19578, January 2016.

[16] Adam Paetznick and Ben W. Reichardt. Fault-tolerant ancilla preparation and noise threshold lower bounds for the 23-qubit Golay code. *Quant. Inf. Comput.*, 12(11-12):1034–1080, 2012.

[17] Lukas Postler, Sascha Heußen, Ivan Pogorelov, Manuel Rispler, Thomas Feldker, Michael Meth, Christian D. Marciniak, Roman Stricker, Martin Ringbauer, Rainer Blatt, Philipp Schindler, Markus Müller, and Thomas Monz. Demonstration of fault-tolerant universal quantum gate operations. *Nature*, 605(7911):675–680, May 2022.

[18] C. Ryan-Anderson, N. C. Brown, M. S. Allman, B. Arkin, G. Asa-Attuah, C. Baldwin, J. Berg, J. G. Bohnet, S. Braxton, N. Burdick, J. P. Campora, A. Chernoguzov, J. Esposito, B. Evans, D. Francois, J. P. Gaebler, T. M. Gatterman, J. Gerber, K. Gilmore, D. Gresh, A. Hall, A. Hankin, J. Hostetter, D. Lucchetti, K. Mayer, J. Myers, B. Neyenhuis, J. Santiago, J. Sedlacek, T. Skripka, A. Slattery, R. P. Stutz, J. Tait, R. Tobey, G. Vittorini, J. Walker, and D. Hayes. Implementing Fault-tolerant Entangling Gates on the Five-qubit Code and the Color Code, August 2022. arXiv:2208.01863 [quant-ph].

[19] Sebastian Krinner, Nathan Lacroix, Ants Remm, Agustin Di Paolo, Elie Genois, Catherine Leroux, Christoph Hellings, Stefania Lazar, Francois Swiadek, Johannes Herrmann, Graham J. Norris, Christian Kraglund Andersen, Markus Müller, Alexandre Blais, Christopher Eichler, and Andreas Wallraff. Realizing repeated quantum error correction in a distance-three surface code. *Nature*, 605(7911):669–674, May 2022.

[20] Friederike Butt, Sascha Heußen, Manuel Rispler, and Markus Müller. Fault-Tolerant Code-Switching Protocols for Near-Term Quantum Processors. *PRX Quantum*, 5(2):020345, May 2024.