Hybrid Multi-Objective Genetic Programming for Parameterized Quantum Operator Discovery

Felix Gemeinhardt
Johannes Kepler University
Institute for Business Informatics - Software Engineering, CDL-MINT
Linz, Austria
felix.gemeinhardt@jku.at

Stefan Klikovits
Johannes Kepler University
Institute for Business Informatics - Software Engineering, CDL-MINT
Linz, Austria
stefan.klikovits@jku.at

Manuel Wimmer
Johannes Kepler University
Institute for Business Informatics - Software Engineering, CDL-MINT
Linz, Austria
manuel.wimmer@jku.at

ABSTRACT

The processing of quantum information is defined by quantum circuits. For applications on current quantum devices, these are usually parameterized, i.e., they contain operations with variable parameters. The design of such quantum circuits and aggregated higher-level quantum operators is a challenging task which requires significant knowledge in quantum information theory, provided a polynomial-sized solution can be found analytically at all. Moreover, finding an accurate solution with low computational cost represents a significant trade-off, particularly for the current generation of quantum computers. To tackle these challenges, we propose a multi-objective genetic programming approach—hybridized with a numerical parameter optimizer—to automate the synthesis of parameterized quantum operators. To demonstrate the benefits of the proposed approach, it is applied to a quantum circuit of a hybrid quantum-classical algorithm, and then compared to an analytical solution as well as a non-hybrid version. The results show that, compared to the non-hybrid version, our method produces more diverse solutions and more accurate quantum operators which even reach the quality of the analytical baseline.

CCS CONCEPTS

• **Computer systems organization → Quantum computing**; • **Mathematics of computing → Evolutionary algorithms**.

KEYWORDS

Quantum Circuit Synthesis, Genetic Programming, Hybrid Search, Search-Based Quantum Software Engineering

ACM Reference Format:
Felix Gemeinhardt, Stefan Klikovits, and Manuel Wimmer. 2023. Hybrid Multi-Objective Genetic Programming for Parameterized Quantum Operator Discovery. In *Genetic and Evolutionary Computation Conference Companion (GECCO '23 Companion), July 15–19, 2023, Lisbon, Portugal.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3583133.3590696

GECCO '23 Companion, July 15–19, 2023, Lisbon, Portugal

ACM ISBN 979-8-4007-0120-7/ 23/07.
https://doi.org/10.1145/3583133.3590696

1 INTRODUCTION

Quantum Computing. The current era of *Quantum Computing* (QC) is referred to as the *Noisy Intermediate-Scale Quantum* (NISQ) era, where the limitations of quantum hardware are mitigated by significant means of classical computation [16]. Analogously to logic gates for classical computation, in QC, quantum information is processed with operations called quantum gates. The most commonly used realistic model of QC is the so-called quantum circuit model [14]. Quantum gates can be parameterized, where the use of parameterized quantum circuits is common in the NISQ-era. This is because classical optimization of the parameters, which constitutes an NP-hard problem, allows to cope with the noise present in current quantum hardware [3]. For this reason, numerical parameter optimizers constitute a central element of NISQ-era quantum algorithms [3–5]. There is ongoing research on quantum-aware optimizers, which are particularly capable of coping with specific requirements of parameterized quantum circuits [3–5, 12].

Within the circuit model of QC, quantum gates are sequentially applied to qubits—the elementary quantum information carrier.

A simple example of a parameterized quantum circuit is depicted in Figure 1, which shows parts of the four-qubit circuit for the hybrid quantum classical GM-QAOA algorithm [2]. In general, the parameters (e.g., β, γ) are float-type rotation angles that specify the concrete action of the quantum gates. As illustrated in Figure 1, any number of quantum gates in a circuit can be composed to *quantum operators* which serve a certain higher-level functionality (e.g., *Oracle*) and result in a certain quantum state.

Quantum Circuit Synthesis. Designing suitable quantum operators and quantum circuits for a desired computational task requires significant amounts of expert knowledge in quantum physics, quantum information theory, and linear algebra. This motivates the use of intelligent search approaches to automatically synthesise quantum circuits and quantum operators, where, additionally, the complex, vast and unknown search spaces make *genetic programming* (GP) favourable over traditional search techniques [10]. Early work on automated non-parameterized quantum circuit synthesis has already been conducted in the late 1990s. These early attempts can be understood as an answer to a lack of invention after promising quantum algorithms have been developed. A review on such approaches can be found in [10]. In contrast to our work, those approaches do not allow treating parameterized quantum circuits.

Particularly the given hardware limitations of the NISQ-era call for multi-objective search methods to account for existing trade-offs. The arguably most important metrics in this regard are the *accuracy* and *computational cost* of the solutions. Accuracy, which represents

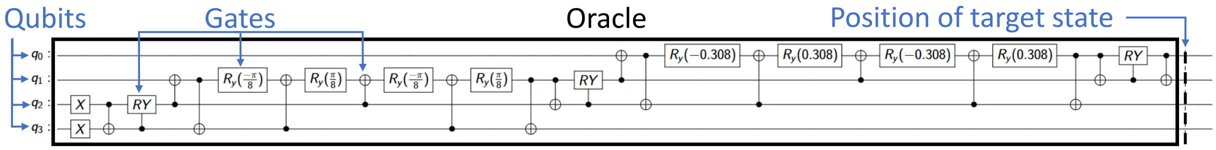

Figure 1: 4-qubit circuit for state preparation operator of GM-QAOA; dashed line: position of prepared quantum state

the functionality of a quantum operation—i.e., the extent to which the produced output resembles an expected target output—, is typically the most important objective in quantum algorithm design. However, excessive computational cost may deteriorate any theoretically given accuracy in practical settings. As NISQ-era gates are fragile, long sequences of quantum gates increase the cumulative error and, thereby, make possible advantages in terms of theoretical accuracy obsolete. Therefore, a computationally inexpensive quantum circuit of imperfect accuracy might be favoured over a perfect but infeasible circuit. Additionally, for many tasks, the existence of a polynomial-sized quantum circuit is not guaranteed, and in certain cases non-existence can even be proved mathematically [2]. Existing multi-objective approaches allow in principle for a treatment of parameterized quantum circuits using a global optimization of the parameters within the GP itself [15]. However, to the best of our knowledge, the potentially beneficial combination of evolutionary algorithms with quantum-aware numerical parameter optimizers has not been studied yet.

Contribution. The main contribution of our work is the combination of multi-objective GP with a quantum-aware parameter optimizer to a hybrid approach. Thereby, we provide quantum algorithm developers with a tool to automatically synthesise and explore suitable quantum operators on representative instances of quantum circuits. We evaluate the proposed approach and compare the obtained results with an analytical solution [2] as well as a non-hybridized version where the parameters are optimized by the evolutionary algorithm.

We find that the hybridization with a quantum-aware optimizer yields (i) Pareto-fronts of higher diversity compared to the non-hybrid version, (ii) quantum operators of higher accuracy compared to the non-hybrid version, and (iii) solutions which are similar to the analytical solution.

2 APPROACH

Encoding and Search Space. Within the proposed genetic programming approach, we represent an individual—i.e., a quantum operator—as a list of genes, which resembles a stackless linear representation [18]. Each gene holds information regarding (i) an individual quantum gate, (ii) the qubits on which an individual gate is applied, and (iii) the parameters of the individual gate if applicable. A quantum operator can comprise a variable number of individual gates. We define two sets of gates. First, the set of non-parameterized gates $\mathcal{NP} = \{H, X, Y, Z, CX, CY, CZ, SWAP\}$, and second, the set of parameterized gates with variable parameters $\mathcal{P}_{var} = \{RX_\theta, RY_\theta, RZ_\theta, CU_{\theta,\lambda,\phi,\gamma}, RXX_\theta, RYY_\theta, RZZ_\theta\}$. Based on these two sets, we define the gateset as $\mathcal{NP} \cup \mathcal{P}_{var}$.

Search Process. The proposed hybrid genetic programming framework takes as input a user-defined overall quantum circuit as, e.g., the one shown in Figure 2. The search aims to find a computationally inexpensive implementation for the *Oracle* with high accuracy, meaning that the observed output quantum state matches a user-defined target quantum state.

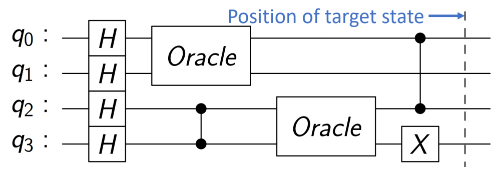

Figure 2: Scope of operator search; *Oracles*: searched quantum operator; dashed line: position of target quantum state

Our genetic programming framework utilizes the NSGA-III algorithm [7]. In each fitness evaluation, a numerical optimizer is applied to find those parameters of the individual that maximize the similarity between observed output and expected target quantum state. Hereby, we build on existing knowledge on numerical optimizers for parameterized quantum circuits [4, 5]. Note, that the quantum states are always accessible as our approach is based on quantum simulation. Further user configurations allow to customize the search process[1].

Mutator and Crossover Operations. The following mutation operators inspired from previous work [1, 15] have been implemented: (M1) change of qubits to which the gate is applied, (M2–M5) insertion, deletion, replacement, or movement of an individual gate in the sequence, (M6) swap of 2 individual gates, and (M7–M9) replacement, swap, or permutation of a whole subsequence of individual gates. The implemented crossover operators are: (C1) one-point crossover where the size of the parent individuals is preserved, (C2) one-point crossover where the size of the children may be different to those of the parent individuals, (C3) two-point crossover where the size of the parent individuals is preserved, and (C4) two-point crossover where the size of the children may be different to those of the parent individuals. We further apply constraints to the individuals, which ensure a size between two and the user-defined maximum size. Finally, duplicates are removed in each generation.

Fitness Evaluation. Especially in the current NISQ-era of quantum computing, we face a trade-off between the accuracy of the obtained solution and its size and complexity. Therefore, we implemented the following fitness objectives: (1) the *overlap*—a commonly known measure in quantum physics—evaluates the accuracy of a quantum operation by measuring the similarity between the obtained output quantum state and the user-defined target quantum state, based on the scalar product of the two complex valued vectors [14]; (2) the *number of gates* denotes the size of the individual; (3) the *depth*,

[1] github.com/jku-win-se/Genetic-Programming-for-Quantum-Operator-Discovery

which denotes the maximum number of gates applied to a single qubit in the circuit; (4) the *number of non-local gates* denotes the number of multi-qubit gates; (5) a low *number of parameters* is supposed to speed-up the genetic programming approach and is taken as a rough measure for the complexity of the numerical optimization.

With the exception of the *overlap*, all fitness objectives should be minimised. The latter three fitness values (depth, number of non-local gates, number of parameters) have been identified in [17] as the key measures for quantum circuit cost, and the overall number of gates, the depth and the number of non-local gates have been stated in [6] as key metrics for characterization of quantum circuits.

3 EVALUATION

We provide a quantitative comparison of the proposed framework, coined `HYBRID`, with a non-hybridized version where parameters are optimized within the evolutionary algorithm itself, referred to as `NON-HYBRID`. For a quantitative comparison with `NON-HYBRID`, we realized the concept of optimizing the parameters within the GP by implementing an additional parameter mutation operator [15]. We also compare the search-based results to an analytical solution (`ANA`). The trade-off between accuracy and computational cost allows in principle for a manifold of possible solutions, where the actual feasibility can only be determined by using information from the specific quantum device for execution. As this information is not known a-priori, we view the diversity of the obtained solutions as a crucial metric. In the rest of this section, we answer the following research questions:

- *RQ1:* How does `HYBRID` perform regarding the diversity of the resulting Pareto-fronts when compared to `NON-HYBRID`?
- *RQ2:* How does `HYBRID` perform regarding the resulting accuracy (i.e., overlap) when compared to `NON-HYBRID`?
- *RQ3:* How do the search-based solutions perform when compared to `ANA`?

3.1 Methodology

We apply `HYBRID` and `NON-HYBRID` to the state-preparation operator of the GM-QAOA algorithm [2] (cf. *Oracle* in Figure 1). Our evaluation is based on the results of 30 runs for each approach using different random *seeds*. We provide statistical information on the respective quantities of interest for the last generation of the search process. The statistical significance of the distributions is assessed using the *Wilcoxon signed-rank* test [20] with a significance level of $\alpha = 0.05$. Furthermore, the effect size is measured according to [19] and categorized according to [11], in order to provide information on the strength of the statistical difference.

We answer *RQ1* by evaluating the diversity of the whole Pareto-fronts of the last generations using the *Diversity Comparison Indicator* (DCI) [13]. Thereafter, we reduce the Pareto-fronts to only those solutions with an overlap of more than 85%. It can be expected that the cumulative error of sequential gates does not exceed the lack of accuracy for lower overlaps. Regarding *RQ2*, we select the solution with the highest overlap from the final Pareto-front to assess the respective accuracy. For a comparison to the analytical benchmark (*RQ3*), not only the overlap of the selected solutions, but all fitness values are evaluated.

The population size and number of generations have been chosen, such that the maximum execution times of the two approaches are roughly similar. Based on naive hyper-parameter tuning in form of non-exhaustive trial & error the number of generations has been set to 17 (`HYBRID`) and 700 (`NON-HYBRID`), and the population size to 33 (`HYBRID`) and 1400 (`NON-HYBRID`), respectively. The maximum size of the quantum operator individuals has been limited to 30 for both approaches. The probabilities for mutation and crossover are set to 1.0 (as provided in the Deap documentation[2]). We use the *Nelder-Mead* [8] method as a common quantum-aware parameter optimizer [3, 12].

3.2 Results and Discussion

RQ1: How does `HYBRID` *perform when compared to* `NON-HYBRID` *regarding the diversity of the resulting Pareto-fronts?* The results on the statistical evaluation of DCI-values are shown in Figure 3. The figure presents both, the diversity of the full final Pareto-fronts, as well as the diversity of the reduced Pareto-fronts, i.e., those individuals that have an overlap >85%. We see that `HYBRID` performs *largely* better, i.e., yields a largely higher diversity, than `NON-HYBRID` in both cases.

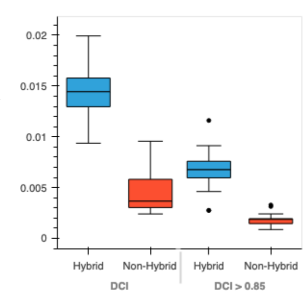

Figure 3: Box plots for DCI values of final Pareto-fronts (*DCI*), and final reduced (overlap>85%) Pareto-fronts (*DCI>0.85*)

RQ2: How does `HYBRID` *perform when compared to* `NON-HYBRID` *regarding the resulting accuracy (i.e., overlap)?* Figure 4 shows the overlap convergence of the final solution within each generation. We see, that `HYBRID` already starts with high overlap in the very first generation and converges rather fast to solutions of high accuracy. `NON-HYBRID` converges slower and fails to a large extent at achieving overlaps above 90%.

[2]https://github.com/DEAP/deap/blob/master/examples/ga/nsga3.py

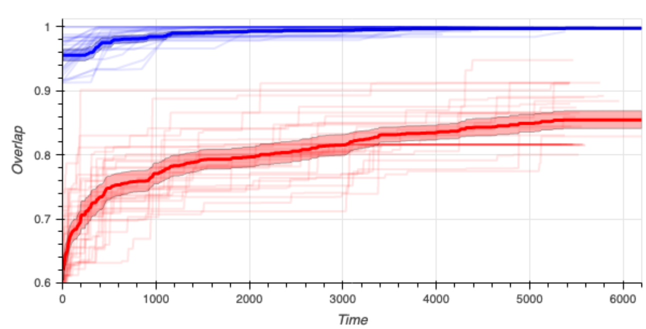

Figure 4: Overlap of final quantum operators; bright lines: mean values & standard deviations for `HYBRID` (blue) and `NON-HYBRID` (red); faint lines: results of individual runs

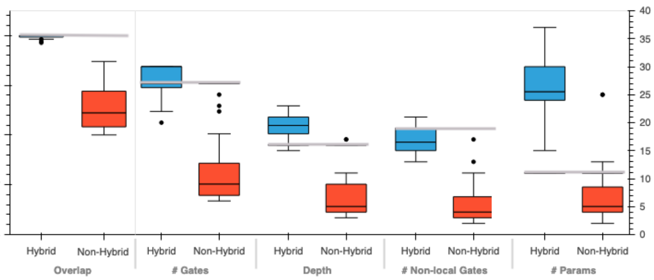

Figure 5: Box plots for fitness values of full final Pareto-fronts obtained by `HYBRID` (blue) and `NON-HYBRID` (red); left y-axis: overlaps, right y-axis: remaining objectives

RQ3: How do the search-based solutions perform when compared to `ANA`*?* The results for the final solutions are shown in Figure 5. The analytical solution `ANA` (cf. Figure 1) is indicated by bold, gray bars with fitness values of $overlap = 100\%$, $number\ of\ gates = 29$, $depth = 17$, $number\ of\ non\text{-}local\ gates = 19$, $number\ of\ parameters = 11$. As shown in Figure 5, `HYBRID` succeeds in reaching the analytical overlap limit of 100%. Whereas `NON-HYBRID` fails to produce quantum operators of high accuracy, the associated fitness objectives for the computational cost are significantly better than for `ANA` as well as `HYBRID`, respectively. We attribute this observation to the fact that `NON-HYBRID` does not find optimal solutions in the region of the search space, where individuals show high overlaps. Note, that the number of parameters merely represents the cost for the parameter optimizer within the evolutionary algorithm.

Threats to validity and limitations. Evidently, the obtained results are specific to the presented use case and may vary significantly for different problems. To mitigate this threat we have chosen the use case to be a representative example of parameterized quantum circuits for hybrid quantum-classical algorithms of the NISQ-era. Additionally, we decided not to optimize the search parameters to show "off-the-shelve performance". Finally, `HYBRID` is limited by the capabilities of quantum simulators, because in real quantum devices the quantum state information is not accessible. In combination with the exploding search space when dealing with individual gates as genes, this restricts the application of our approach to small quantum circuits.

4 CONCLUSION

In this paper, we present a multi-objective genetic programming approach for the automated synthesis of parameterized quantum operators, which is hybridized with a quantum-aware parameter optimizer and takes trade-offs regarding accuracy and computational cost of quantum operators into account. Our hybrid approach is compared to a non-hybridized version that does not employ a numerical optimizer, as well as an analytical benchmark solution. The evaluation shows that first, our approach yields Pareto-fronts of higher diversity and solutions of higher accuracy compared to the non-hybridized version, and second, the resulting quantum operators are comparable to the analytical solution. In future, we aim to apply our framework to other use cases and evaluate the effects of different search configurations.

ACKNOWLEDGMENTS

Financial support by the Austrian Federal Ministry for Digital and Economic Affairs and the National Foundation for Research, Technology and Development and by the Austrian Science Fund (P 30525-N31) is gratefully acknowledged.

DATA AVAILABILITY

All source code for the hybrid multi-objective genetic programming approach and the evaluation data is publicly available on Zenodo [9] and in the repository https://github.com/jku-win-se/Genetic-Programming-for-Quantum-Operator-Discovery.

REFERENCES

[1] Usama Ahsan et al. 2020. AutoQP: Genetic Programming for Quantum Programming. In *IBCAST*. IEEE, 378–382.

[2] Andreas Bärtschi and Stephan Eidenbenz. 2020. Grover mixers for QAOA: Shifting complexity from mixer design to state preparation. In *QCE*. IEEE, 72–82.

[3] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermanni Heimonen, Jakob S Kottmann, Tim Menke, et al. 2021. Noisy intermediate-scale quantum (NISQ) algorithms. *arXiv preprint* (2021).

[4] Xavier Bonet-Monroig, Hao Wang, Diederick Vermetten, Bruno Senjean, Charles Moussa, Thomas Bäck, Vedran Dunjko, and Thomas E O'Brien. 2021. Performance comparison of optimization methods on variational quantum algorithms. *arXiv preprint arXiv:2111.13454* (2021).

[5] Marco Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, et al. 2021. Variational quantum algorithms. *Nature Reviews Physics* 3, 9 (2021), 625–644.

[6] José A Cruz-Lemus, Luis A Marcelo, and Mario Piattini. 2021. Towards a Set of Metrics for Quantum Circuits Understandability. In *QUATIC*. Springer, 239–249.

[7] Kalyanmoy Deb and Himanshu Jain. 2013. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: solving problems with box constraints. *IEEE Transactions on Evolutionary Computation* 18, 4 (2013), 577–601.

[8] Fuchang Gao and Lixing Han. 2012. Implementing the Nelder-Mead simplex algorithm with adaptive parameters. *Computational Optimization and Applications* 51, 1 (2012), 259–277.

[9] Felix Gemeinhardt, Stefan Klikovits, and Manuel Wimmer. 2023. *jku-win-se/Genetic-Programming-for-Quantum- Operator-Discovery: ZenodoRelease.* https://doi.org/10.5281/zenodo.7838844

[10] Adrian Gepp and Phil Stocks. 2009. A review of procedures to evolve quantum algorithms. *Genetic programming and evolvable machines* 10, 2 (2009), 181–228.

[11] Barbara Kitchenham, Lech Madeyski, David Budgen, Jacky Keung, Pearl Brereton, Stuart Charters, Shirley Gibbs, and Amnart Pohthong. 2017. Robust Statistical Methods for Empirical Software Engineering. *Empirical Softw. Eng.* 22, 2 (2017), 579–630.

[12] Wim Lavrijsen, Ana Tudor, Juliane Müller, Costin Iancu, and Wibe De Jong. 2020. Classical optimizers for noisy intermediate-scale quantum devices. In *QCE*. IEEE, 267–277.

[13] Miqing Li, Shengxiang Yang, and Xiaohui Liu. 2014. Diversity comparison of Pareto front approximations in many-objective optimization. *IEEE Transactions on Cybernetics* 44, 12 (2014), 2568–2584.

[14] Michael A Nielsen and Isaac L Chuang. 2001. Quantum computation and quantum information. *Phys. Today* 54, 2 (2001), 60.

[15] Václav Potoček, Alan P Reynolds, Alessandro Fedrizzi, and David W Corne. 2018. Multi-objective evolutionary algorithms for quantum circuit discovery. *arXiv preprint arXiv:1812.04458* (2018).

[16] John Preskill. 2018. Quantum computing in the NISQ era and beyond. *Quantum* 2 (2018), 79.

[17] Sukin Sim, Peter D Johnson, and Alán Aspuru-Guzik. 2019. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies* 2, 12 (2019), 1900070.

[18] Lee Spector, Howard Barnum, Herbert J Bernstein, and Nikhil Swamy. 1999. Quantum computing applications of genetic programming. *Advances in genetic programming* 3 (1999), 135–160.

[19] András Vargha and Harold D Delaney. 2000. A critique and improvement of the CL common language effect size statistics of McGraw and Wong. *Journal of Educational and Behavioral Statistics* 25, 2 (2000), 101–132.

[20] Frank Wilcoxon. 1992. Individual Comparisons by Ranking Methods. In *Breakthroughs in Statistics: Methodology and Distribution*, Samuel Kotz and Norman L. Johnson (Eds.). Springer, 196–202.