

Global Optimization for Parametrized Quantum Circuits

Iosif Sakos¹, Antonios Varvitsiotis^{1,3,5}, Georgios Korpas^{2,4,5}, and Wayne Lin¹

¹Singapore University of Technology and Design, Singapore

²Quantum Technologies Group, Innovation & Ventures, HSBC, Singapore

³Centre for Quantum Technologies, Singapore

⁴Czech Technical University in Prague, Czechia

⁵Archimedes Research Unit on AI, Marousi, Greece

Main text available as [arXiv:2603.21757](https://arxiv.org/abs/2603.21757).

Abstract

In the absence of error correction, noisy intermediate-scale quantum devices are typically operated using parametrized quantum circuits (PQCs) trained to minimize a loss function. Finding the circuits' optimal parameters is a hard optimization problem, where global guarantees are known only for highly structured cases of limited practical relevance, and first-order methods can fail even locally due to barren plateaus. In this work, we study the training of practical classes of PQCs, namely polynomial-depth circuits with a constant number of trainable parameters. This captures widely used PQC families, including fixed-depth QAOA, hardware-efficient ansätze with a controlled number of parameters per layer, and Fixed Parameter Count QAOA. Our main technical result is an additive fully polynomial randomized approximation scheme (FPRAS), which, for every $\epsilon > 0$, returns an ϵ -approximate solution to the problem's global optimum with high probability, and has runtime and query complexity polynomial in $1/\epsilon$ and the number of qubits. Unlike standard VQA training, where the quantum device is queried repeatedly throughout the training, our approach has two stages: (1) an initial quantum data-acquisition phase, followed by (2) a classical global-optimization phase based on the trigonometric moment/sum-of-squares hierarchies of semidefinite programs. Under a standard flat-extension condition, which can be checked numerically, the method also supports the extraction of globally optimal circuit parameters. The FPRAS implies that the promise problem associated with the optimization of poly-depth constant-parameter PQCs is tractable. This limits the class's expressive power, namely, it cannot encode combinatorial optimization problems whose objective values are separated by an inverse-polynomial gap.

1 Introduction

The main difficulty in ground-state computation and combinatorial optimization is that the underlying Hilbert space is exponentially large. A standard way to cope with this is to use a parametrized quantum circuit (PQC) defined by a problem-specific or hardware-aware ansatz, thereby restricting the search to the subset of quantum states reachable from an easily preparable reference state. This gives rise to the parametrized quantum circuit optimization problem, which seeks the minimum value $f^* = \min_{\theta} f(\theta)$, where $f(\theta)$ is the expectation of the observable under the state prepared by the circuit with parameters θ . From an algorithmic perspective, this problem is challenging. It is strongly NP-hard in general, even under favorable assumptions [1], and even computing a local minimum with gradient-based methods is often difficult in practice because the objective landscape may exhibit barren plateaus, i.e., regions where the function is nearly flat and gradients carry essentially no useful optimization information [2, 3]. Existing global guarantees are largely limited either to very low-dimensional settings [4, 5] or to highly structured regimes [6, 7], and therefore do not fully capture practical settings of interest.

In this work, we depart from the restrictive settings considered in the existing global optimization literature [4–7], and study a class of ansätze well aligned with practical implementations. We consider PQCs with a constant number of independent parameters, but a polynomially large number of parametrized gates, and show that the **PQC optimization problem admits a fully polynomial randomized approximation scheme (FPRAS) [8] under standard assumptions**. This means that, for every $\epsilon > 0$, the algorithm

returns an ϵ -approximate solution to the problem's global optimum with high probability, and has runtime and query complexity polynomial in $1/\epsilon$ and the number of qubits. This is a strong positive result for the global optimization of PQC, and it is applicable to a broad range of practical ansätze, including fixed levels of the Quantum Approximate Optimization Algorithm [9], certain hardware-efficient ansätze with a controlled number of parameters per layer [10, 11], and Fixed Parameter Count QAOA [12]. At a high level, the FPRAS applies the complex trigonometric moment/sum-of-squares (SOS) hierarchy [13–15] to the approximated PQC optimization problem, obtaining a convergent sequence of monotone lower bounds that approach the optimum in the limit, with each level computable via semidefinite programming.

2 An FPRAS for PQC optimization

We consider a PQC $\mathbf{U}(\theta)$ formed by a sequence of fixed and parametrized unitary gates acting on the reference state $|0\rangle$. Given a Hermitian observable $\mathbf{O}$, the corresponding objective, i.e., the loss function, is

$$f(\theta) = \langle 0 | \mathbf{U}^\dagger(\theta) \mathbf{O} \mathbf{U}(\theta) | 0 \rangle. \quad (1)$$

Here, $f(\theta)$ is the expectation of the observable $\mathbf{O}$ under the state prepared by the PQC with parameters θ . To show the existence of an FPRAS for PQC optimization, the main observation is that the objective can be formulated as a Hermitian trigonometric polynomial (HTP), i.e., a function of the form

$$f(\theta) = \sum_{\alpha \in \mathbb{Z}^M} f_\alpha \exp(i\langle \alpha, \theta \rangle), \quad \theta \in [-\pi, \pi]^M, \quad \overline{f_\alpha} = f_{-\alpha}, \quad (2)$$

where M is the number of independent parameters (which can be shared across gates), α is an integer frequency vector, and f_α is the corresponding Fourier coefficient. This casts the PQC optimization problem as a trigonometric optimization problem, a class of problems that has been extensively studied in the optimization literature; see, e.g., [13–17] and references therein. Then our main result is stated as follows.

Theorem 1 (An FPRAS for PQC optimization). *Consider a PQC $\mathbf{U}(\theta)$ with n qubits whose generators have integer spectral differences, a constant number M of independent parameters, a polynomial number of parametrized gates, and spectral diameters bounded by a polynomial in n . Let $\mathbf{O}$ be a Hermitian observable whose operator norm $\|\mathbf{O}\|$ is bounded by a polynomial in n . Then, for every error $\epsilon > 0$ and failure probability $\delta > 0$, there exists a randomized algorithm with oracle access to a quantum device running in time polynomial in n , $1/\epsilon$, and $\log(1/\delta)$ that outputs an estimate $\hat{f}^*$ of f^* such that, with probability at least $1 - \delta$,*

$$|\hat{f}^* - f^*| \leq \epsilon.$$

Furthermore, the algorithm requires a number of queries to the quantum device that is a priori known and polynomial in n , $1/\epsilon$, and $\log(1/\delta)$.

The construction proceeds in three steps. First, we show that the objective $f(\theta)$ can be formulated as an HTP. Second, we approximate the HTP formulation using the quantum hardware to compute estimates of $f(\theta)$ on a uniform sampling grid of the parameter space, and the multidimensional Fast Fourier Transform to compute the Fourier coefficients from these estimates. Finally, we apply the complex trigonometric moment/SOS hierarchy [13–15] to the approximated PQC optimization problem, obtaining a convergent sequence of monotone lower bounds that approach the optimum in the limit, and show that the required order of the SOS relaxation is at most polynomial in n and $1/\epsilon$. The left panel of Figure 1 shows an example of a two-parameter HTP objective $f(\theta_1, \theta_2)$, and the right panel shows how successive levels of the SOS hierarchy give lower bounds that approach the global optimum f^* .




Figure 1: A two-parameter PQC objective and two levels of the trigonometric SOS hierarchy applied to the same objective.

Parameter extraction. The proof of [Theorem 1](#) also provides insights on how to extract an approximate optimizer of the objective function $f(\theta)$ from the solution of the SOS relaxation. This is particularly relevant in practice, where one is often interested not only in approximating the optimal value f^* , but also in finding parameters θ that approximately minimize the PQC optimization problem. Under the standard flat-extension assumption for the dual moment sequence, which can be checked numerically, the solution at a finite level is optimal and an approximate global minimizer of $f(\theta)$ can be extracted from the optimal moment sequence.

Complexity consequences. The existence of the FPRAS implies that the gap-promise problem corresponding to this family of PQCs is in the complexity class BPP^{BQP} , and therefore in BQP [18]. Essentially, this means that solving the promise problem associated with this class of PQC optimization problems is no harder than solving a problem in BQP . At the same time, this imposes strong restrictions on the expressivity of polynomial-depth constant-parameter PQCs. Unless one expects an implausible collapse such as $\text{NP} \subseteq \text{BPP}^{\text{BQP}}$, they cannot represent NP -hard problems, even approximately. In conjunction with the results of Bittel et al. [1], which show that PQC training is, in general, NP -hard, this suggests that the barrier to tractability lies in the number of trainable parameters.

3 Relevance to QTML

Our talk will be of broad interest to the QTML community because it provides a global optimization guarantee for a practical circuit regime and offers a concrete algorithmic framework for training PQCs that differs fundamentally from the standard hybrid quantum–classical training loop in variational quantum algorithms. In addition, our results highlight limitations on the expressive power of polynomial-depth constant-parameter PQCs, which is of independent interest. Taken together, these contributions represent substantial progress on the trainability and expressivity of PQCs and will be accessible and engaging to a broad range of QTML attendees.

References

- [1] Lennart Bittel and Martin Kliesch. “Training Variational Quantum Algorithms Is NP-Hard.” In: *Physical Review Letters* 127.12 (2021), p. 120502. ISSN: 1079-7114. DOI: [10.1103/PhysRevLett.127.120502](https://doi.org/10.1103/PhysRevLett.127.120502).
- [2] Jarrod Ryan McClean et al. “Barren plateaus in quantum neural network training landscapes.” In: *Nature Communications* 9.1 (2018), p. 4812. ISSN: 2041-1723. DOI: [10.1038/s41467-018-07090-4](https://doi.org/10.1038/s41467-018-07090-4).
- [3] Marco Cerezo et al. “Does provable absence of barren plateaus imply classical simulability?” In: *Nature Communications* 16 (2025), p. 7907. DOI: [10.1038/s41467-025-63099-6](https://doi.org/10.1038/s41467-025-63099-6).
- [4] Zhihui Wang et al. “Quantum Approximate Optimization Algorithm for MaxCut: A Fermionic View.” In: *Physical Review A* 97.2 (2018), p. 022304. DOI: [10.1103/physreva.97.022304](https://doi.org/10.1103/physreva.97.022304).
- [5] Shree Hari Sureshabu et al. “Parameter Setting in Quantum Approximate Optimization of Weighted Problems.” In: *Quantum* 8 (2024), p. 1231. DOI: [10.22331/q-2024-01-18-1231](https://doi.org/10.22331/q-2024-01-18-1231).
- [6] Xuchen You, Shouvanik Chakrabarti, and Xiaodi Wu. *A Convergence Theory for Over-Parameterized Variational Quantum Eigensolvers*. arXiv:2205.12481 [quant-ph]. Version 1. 2023. DOI: [10.48550/arXiv.2205.12481](https://doi.org/10.48550/arXiv.2205.12481).
- [7] Marco Wiedmann et al. *On the Convergence of the Variational Quantum Eigensolver and Quantum Optimal Control*. arXiv:2509.05295 [quant-ph]. Version 2. 2025. DOI: [10.48550/arXiv.2509.05295](https://doi.org/10.48550/arXiv.2509.05295).
- [8] Richard M. Karp, Michael Luby, and Neal Madras. “Monte-Carlo Approximation Algorithms for Enumeration Problems.” In: *Journal of Algorithms* 10.3 (1989), pp. 429–448. DOI: [10.1016/0196-6774\(89\)90038-2](https://doi.org/10.1016/0196-6774(89)90038-2).
- [9] Edward Henry Farhi, Jeffrey Goldstone, and Samuel Gutmann. *A Quantum Approximate Optimization Algorithm*. Nov. 14, 2014. DOI: [10.48550/arXiv.1411.4028](https://doi.org/10.48550/arXiv.1411.4028). arXiv: [1411.4028](https://arxiv.org/abs/1411.4028) [quant-ph].
- [10] Juan M. Pino et al. “Demonstration of the Trapped-Ion Quantum CCD Computer Architecture.” In: *Nature* 592.7853 (2021), pp. 209–213. DOI: [10.1038/s41586-021-03318-4](https://doi.org/10.1038/s41586-021-03318-4).
- [11] J.-Z. Zhuang, Y.-K. Wu, and Luming Duan. “Hardware-Efficient Variational Quantum Algorithm in a Trapped-Ion Quantum Computer.” In: *Physical Review A* 110.6 (2024), p. 062414. DOI: [10.1103/physreva.110.062414](https://doi.org/10.1103/physreva.110.062414).
- [12] Sebastián Saavedra-Pino et al. *Quantum Approximate Optimization Algorithm with Fixed Number of Parameters*. arXiv:2512.21181 [quant-ph]. Version 1. 2025. DOI: [10.48550/arXiv.2512.21181](https://doi.org/10.48550/arXiv.2512.21181).
- [13] Bogdan Dumitrescu. *Positive Trigonometric Polynomials and Signal Processing Applications*. Springer, 2007. ISBN: 978-1-4020-5125-8.
- [14] Cedric Josz and Daniel K. Molzahn. “Lasserre Hierarchy for Large Scale Polynomial Optimization in Real and Complex Variables.” In: *SIAM Journal on Optimization* 28.2 (2018), pp. 1017–1048. DOI: [10.1137/15m1034386](https://doi.org/10.1137/15m1034386).
- [15] Francis Bach and Alessandro Rudi. “Exponential Convergence of Sum-of-Squares Hierarchies for Trigonometric Polynomials.” In: *SIAM Journal on Optimization* 33.3 (2023), pp. 2137–2159. ISSN: 1052-6234. DOI: [10.1137/22M1540818](https://doi.org/10.1137/22M1540818).
- [16] Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. “Effect of data encoding on the expressive power of variational quantum-machine-learning models.” In: *Physical Review A* 103.3 (2021), p. 032430. ISSN: 2469-9934. DOI: [10.1103/PhysRevA.103.032430](https://doi.org/10.1103/PhysRevA.103.032430).
- [17] Enrico Fontana et al. *Efficient Recovery of Variational Quantum Algorithms Landscapes Using Classical Signal Processing*. arXiv:2208.05958v1 [quant-ph]. 2022. DOI: [10.48550/arXiv.2208.05958](https://doi.org/10.48550/arXiv.2208.05958).
- [18] Charles H. Bennett et al. “Strengths and weaknesses of quantum computing.” In: *SIAM Journal on Computing* 26.5 (1997), pp. 1510–1523. DOI: [10.1137/s0097539796300933](https://doi.org/10.1137/s0097539796300933).