PARAMETERIZED HAMILTONIANS FOR COMBINATORIAL OPTIMIZATION PROBLEMS - OPTIMIZATION AND/OR CHARACTERIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Combinatorial optimization problems (COPs) represent a promising application domain for quantum computing, yet current quantum optimization approaches treat each problem instance independently, requiring expensive re-optimization for every configuration. In this paper we propose a different paradigm inspired by quantum many-body physics, where parameterized Hamiltonians naturally encode system variations under changing global conditions. Our parametrized COPs formulation, where a global parameter changes the problem configuration, allows to model parameterized problems and opens access to problem classes that were previously difficult and inefficient to formulate. Second, we provide a concrete algorithmic framework, using implicit differentiation to solve these parameterized COPs classes efficiently. Drawing from techniques used in quantum susceptibility calculations, our method propagates optimal circuit parameters across different Hamiltonian configurations without expensive re-optimization. We demonstrate this approach by finding globally optimal configurations in Max-Cut problems, where the Hamiltonian parameter controls edge weight distributions. Our implementation systematically generates parameterized problem families from Max-Cut, Knapsack, and Portfolio Optimization domains and translates them into quantum formulations suitable for variational algorithms. Experiments on simulated quantum hardware demonstrate substantial computational speedups compared to independent optimization approaches.

1 INTRODUCTION

Quantum computing holds the promise of accelerating selected computational tasks, from factoring integers to simulating quantum systems, yet current devices are noisy and resource-limited (Preskill, 2018). To work within these constraints, researchers have developed *Variational Quantum Algorithms* (VQAs), which hybridize quantum and classical computation: a parameterized quantum circuit (PQC) is evaluated on quantum hardware and a classical optimizer updates its parameters (Cerezo et al., 2021) on a classical computer. VQAs have proven versatile across domains ranging from quantum chemistry to machine learning.

A prominent VQA is the *Variational Quantum Eigensolver* (VQE) (Peruzzo et al., 2014), which aims to find the ground state of a target Hamiltonian and has been widely applied in chemistry and optimization problems. Another widely studied VQA is the *Quantum Approximate Optimization Algorithm* (QAOA) (Farhi et al., 2014), which alternates problem and mixer evolutions and shows promise on NP-hard combinatorial tasks.

In physics applications, Hamiltonians often include *global parameters* (such as external fields or interacting forces), allowing a *single* parameterized Hamiltonian to represent a family of related systems, each described by a specific parametrization. Recent meta-learning variants of VQE explicitly exploit this formulation: *Meta-VQE* encodes Hamiltonian parameters directly into the ansatz to learn energy profiles across a continuous domain (Cervera-Lierta et al., 2021), while *NN-VQE* uses a small neural network to generate PQC parameters as a function of the Hamiltonian parameters (Miao et al., 2024).

By contrast, in standard treatments of *combinatorial optimization problems* (COPs) such as the Portfolio Optimization or Knapsack problems, each instance is mapped to a *fixed*, instance-specific Ising/QUBO Hamiltonian with no free global parameter (Lucas, 2014). This instance-wise modeling overlooks practical situations where families of COP instances are driven by global factors (e.g., overall demand, market volatility, external penalties) that could be abstracted as continuous parameters. Incorporating such global parameters into the Hamiltonian would allow a single model to span a family of related instances and opens the door to *reusing* quantum optimization results across them.

In this paper we address this gap by introducing a formal framework for **parameterized combinatorial optimization**, defining COP Hamiltonians $H(\lambda)$ with one or more continuous global parameters λ that smoothly reshape the cost landscape.

- Markus

2 Background

2.1 Combinatorial Optimization Problems - the Knapsack Problem

- Combinatorial optimization problems involve finding an optimal solution from a finite but often exponentially large set of possibilities - therefore combinatorial. They typically appear in scheduling, resource allocation, logistics, and machine learning.
- Many combinatorial optimization problems are **NP-hard**, meaning that no known classical algorithm can solve them efficiently in the general case.
- A classical example is the Knapsack problem, which models how to select items with given values and weights under a capacity constraint so as to maximize the total value.
- Mathematically, given n items with values v_i and weights w_i, and a maximum weight capacity W, the problem is: $\max \sum_{i=1}^{n} v_i x_i \quad \text{subject to} \quad \sum_{i=1}^{n} w_i x_i \leq W, \quad x_i \in \{0, 1\}$. Here, $x_i = 1$ if item i is chosen, and $x_i = 0$ otherwise.
- The binary variables x_i make the Knapsack problem a **discrete optimization problem** that grows exponentially with the number of items.
- In the context of quantum computing, the Knapsack problem (like other combinatorial problems) can be mapped to a Hamiltonian, where solutions correspond to basis states, and the energy of each state encodes the quality of the solution.
- This mapping allows the use of quantum algorithms (e.g., adiabatic quantum computing, variational approaches) to search for low-energy states, which correspond to near-optimal or optimal solutions of the original problem.

2.2 Hamiltonians and their parameterized Form

- A Hamiltonian in quantum mechanics is a Hermitian operator or function $\hat{H}$ that represents the total energy of a system
- Hamiltonians can be built up from different terms and represented as a sum, e.g. the sum of the kinetic energy $\hat{T}$ and the potential energy $\hat{V}$ -¿ e.g. $\hat{H} = \hat{T} + \hat{V}$.
- At the same time the Hamiltonian determines the time evolution of the underlying system via the Schrödinger equation: $i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = \hat{H} |\psi(t)\rangle$
- its eigenvalues correspond to measurable energy levels; its eigenstates are stationary states of the system
- The Hamiltonian spectrum describes all possible energy outcomes of the system.
- The operators can thereby be dependent on external parameters that can present different influences, like Coupling strength, external magnetic fields, or control variables used in tunable gate parameters.
- While the Hamiltonains represent energy functions and time evolutions, they often inhere time as parameter
- they can be used to translate/represent cost functions.

- Here, the expectation value of the Hamiltonian is minimized, being similar to a cost function
- One special Hamiltonian is the Ising Hamiltonian derived from the Ising model
- todo: explanation Ising Hamiltonian (two to four sentences)

2.3 Variational Quantum Algorithms - specific what we use

- Variational Quantum Algorithms (VQAs) are a class of hybrid quantum–classical algorithms designed to leverage near-term quantum hardware (NISQ devices).
- The key idea is to use a parameterized quantum circuit (Ansatz) that prepares a trial quantum state $|\psi(\boldsymbol{\theta})\rangle$, where $\boldsymbol{\theta}$ denotes a set of tunable parameters.
- A cost function is defined as the expectation value of a Hamiltonian: $C(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|\hat{H}|\psi(\boldsymbol{\theta})\rangle$. Minimizing this cost function corresponds to finding low-energy states of the Hamiltonian.
- The optimization is performed in a hybrid loop the quantum device evaluates $C(\boldsymbol{\theta})$ for given parameters, while a classical optimizer updates $\boldsymbol{\theta}$ to minimize the cost.
- VQAs are flexible: the Hamiltonian $\hat{H}$ can represent physical systems (ground-state energy in quantum chemistry) or abstract cost functions from optimization problems.
- Well-known examples include:
- Variational Quantum Eigensolver (VQE) for estimating ground-state energies and Quantum Approximate Optimization Algorithm (QAOA).
- The advantage of VQAs lies in their ability to approximate solutions using relatively shallow circuits, making them suitable for noisy intermediate-scale quantum (NISQ) devices.
- Challenges include designing efficient ansätze, mitigating noise, and avoiding barren plateaus (regions of flat optimization landscapes).

3 Related Work

- Meta-VQE? https://journals.aps.org/prxquantum/pdf/10.1103/PRXQuantum.2.020329

VQAs for families of Hamiltonians. A key inspiration for our work comes from efficiently solving *parameterized* Hamiltonians in chemistry and physics. *Meta-VQE* augments the VQE ansatz with an encoding layer for Hamiltonian parameters and, after training on a small grid, generalizes ground-state predictions across a continuous domain (Cervera-Lierta et al., 2021). More recently, *NN-VQE* employs a small neural network that maps Hamiltonian parameters directly to PQC angles, enabling near-instant parameter prediction for unseen instances once trained (Miao et al., 2024). Both approaches amortize VQA optimization cost across related problems described by different parametrizations of the same Hamiltonian.

- Markus

Alternative objective functions for COPs. Standard VQE minimizes the expectation value given by $\langle\psi(\theta)|H|\psi(\theta)\rangle$, which can be ill-suited for discrete objectives where optima are computational basis states. Barkoutsos et al. (2020) introduced the *CVaR* objective that averages the best α-fraction of measured energies, increasing the probability of sampling optimal solutions on problems such as MaxCut and Portfolio Optimization. Building on this, Kolotouros & Wallden (2022) proposed an *Ascending-CVaR* schedule that increases α during training, initially enforcing a sharp focus on elite samples, then broadening to refine solutions, yielding substantial success-rate gains on challenging COP landscapes.

- Markus

4 FORMULATING COP CONFIGURATIONS AS A PARAMETERIZED HAMILTONIAN

Here we explain how one can formulate the COP as a PH. This should be done with respect to one single Problem Category. In the Appendix we can add several other Problem categories. However, the presented single Problem Category could be split up into three subchapter: one linear variable, two linear variables, higher-order variables.

My proposal (Fede):

In this section, we describe how a single Parametrized Hamiltonian (PH) can model different configurations of a COP, eliminating the need for multiple instance-specific Hamiltonians. For different COPs, we show how a corresponding PH can be constructed, where the PH-parameter represents a global factor with a meaningful impact on the configuration of the problem. We denote the PH-parameter by $\lambda \in \mathbb{R}^n$, as multiple global factors can influence a single COP.

In particular, we exemplarily construct a PH for the Knapsack problem. We describe three scenarios: (1) a single global parameter ($\lambda \in \mathbb{R}$) has a linear effect on the Knapsack configuration, (2) a single global parameter has a non-linear effect on the Knapsack configuration, and (3) multiple global parameters ($\lambda \in \mathbb{R}^3$) have a non-linear effect on the Knapsack configuration.

- Fede

4.1 ONE LINEAR VARIABLE

Parameterized Weighted MaxCut. Given an undirected weighted graph $G = (V, E)$ with edge weights w_{ij} on $(i, j) \in E$, a standard Ising form (up to an additive constant) is

$$H_{\text{MaxCut}}(G) \;=\; \frac{1}{2} \sum_{(i,j)\in E} w_{ij} \left(1 - Z_i Z_j\right), \tag{1}$$

where Z_i are Pauli-Z operators on qubit i; the ground state encodes the optimal cut. We now introduce a *global parameter* $\lambda \in \mathbb{R}$ that modifies all edges in a controlled way, e.g. to model network load:

$$w_{ij}(\lambda) \;=\; w_{ij}^{\text{init}} + b_{ij}\,\lambda, \tag{2}$$

with base weights w_{ij}^{init} and fixed sensitivities b_{ij}. Plugging equation 2 into equation 1 yields the *parameterized* family

$$H_{\text{MaxCut}}(\lambda) \;=\; \frac{1}{2} \sum_{(i,j)\in E} \left(w_{ij}^{\text{init}} + b_{ij}\lambda\right) \left(1 - Z_i Z_j\right). \tag{3}$$

- Markus

Parametrized Knapsack with Integer Weights. Given a Knapsack with a maximum weight capacity W and N items to be considered, labeled by indices α, each with an integer weight w_α and a value c_α, find the most valuable set of items which can be carried in the Knapsack (i.e. whose combined weight is smaller than W).

As shown in Lucas (2014), the problem can be modeled by an Ising Hamiltonian: Let y_n for $1 \leq n \leq W$ denote a binary variable which is 1 if the final weight of the Knapsack is n, and 0 otherwise. The Knapsack Hamiltonian is given by $H_{KS} = H_A + H_B$, with

$$H_A = A\left(1 - \sum_{n=1}^{W} y_n\right)^2 + A\left(\sum_{n=1}^{W} n y_n - \sum_{\alpha} w_\alpha x_\alpha\right)^2, \tag{4}$$

where the first term enforces the total weight of the selected items to take only one value, and the second term ensures that this value is indeed the sum of the weights of the selected items. Finally, the value of the selected items is maximized

$$H_B = -B \sum_{\alpha} c_\alpha x_\alpha \tag{5}$$

(a negative term is used as the Hamiltonian itself is minimized). To avoid solutions where H_A is weakly violated at the expense of H_B becoming more negative, we require

$$0 < B \max(c_\alpha) < A$$

(adding one item to the Knapsack, which makes it too heavy, is not allowed).

As in the Weighted MaxCut case, it is possible to parametrize H_{KS} to model a family of problem configurations. For instance, the Hamiltonian-parameter $\lambda \in \mathbb{R}$ could affect the item values, modeling the demand for a certain category of products in the market. Different items have an initial value c_α^{init} and a particular sensitivity to market demand s_α. The impact of demand fluctuations on item values can be modeled by

$$c_\alpha = c_\alpha^{init} + \lambda s_\alpha. \tag{6}$$

Following Eq. 4 and 5, the parametrized Hamiltonian H_{KS} takes form

$$H_{KS}(\lambda) = H_A + H_B(\lambda), \tag{7}$$

where

$$H_B(\lambda) = -B \sum_\alpha \left(c_\alpha^{init} + \lambda s_\alpha\right) x_\alpha. \tag{8}$$

4.2 Several Linear Variables

The formulation of parametrized Hamiltonians easily allows for more complex models. For instance, extending the Knapsack problem described above, we can add a second global parameter, modeling the taxation on a given item type. In this case, $\lambda \in \mathbb{R}^2$, with λ_1 being the demand factor as in the previous section and λ_2 being the taxation percentage on a given item type. The real value of the items becomes

$$c_\alpha = c_\alpha^{init} + \lambda_1 s_\alpha + \lambda_2 t_\alpha, \tag{9}$$

where t_α is a binary variable which is 1 if the item is to be taxed at the rate given by λ_2.

The associated cost Hamiltonian becomes

$$H_B(\lambda) = -B \sum_\alpha \left(c_\alpha^{init} + \lambda_1 s_\alpha + \lambda_2 t_\alpha\right) x_\alpha. \tag{10}$$

4.3 Several high-order variables

Of course, the effect of the Hamiltonian parameters must not necessarily be linear. For instance, consider the following sigmoidal model for the value of the items:

$$c_\alpha = V_\alpha^{\max} \frac{1}{1 + \exp\left(-k_\alpha \left(c_\alpha^{\text{init}} + \lambda_1 s_\alpha - \lambda_2 t_\alpha\right)\right)}, \tag{11}$$

where λ_1, λ_2, s_α and t_α play the same role as before, while $V_\alpha^{\max}$ represents the maximum possible value item α can haxve, and $k_\alpha > 0$ is a slope parameter that controls the speed with which the sigmoidal curve saturates for this item. With this equation, the value of the items follows a commonly observed pattern: with small demand, it goes close to zero, rising almost linearly with the demand and saturating at a maximum value.

5 Meta-QAOA Warm-Start

Goal Learn an amortized warm-start map $g_\phi(\lambda)$ that outputs QAOA angles for an entire family of weighted Max-Cut instances with λ-dependent edges, enabling strong zero-/few-shot performance without instance-specific pre-optimization.

SETUP

- **Ansatz & Baseline:** p-layer QAOA (ZZ cost evolutions + RX mixer) versus a factorized classical mean-field model (independent spins).
- **Task Family:** Erdős–Rényi graphs with *in-distribution* (ID) and *out-of-distribution* (OOD) splits in both graphs and λ.
- **Meta-Map:** Compact Poly+RBF regressor; polynomials capture global trends in λ, RBFs add local flexibility.
- **Objective:** CVaR(α) of sampled energies.

METHOD

- **Outer optimization:** SPSA on meta-weights ϕ of g_ϕ (no ES).
- **Inner (optional) few-shot:** SPSA on QAOA angles $\theta = [\beta_{1..p}, \gamma_{1..p}]$ at test time for K steps.
- **Shot model:** Sampling-based evaluation (fixed shots), CVaR aggregation.

EVALUATION PROTOCOL

- **Zero-shot quality:** Energy from $g_\phi(\lambda)$ without per-instance tuning.
- **Few-shot improvement:** Energy after K inner SPSA steps.
- **Time-to-Target (TTT):** Shots required to reach a target energy ratio $E/E^\star$.
- **Splits:** ID-graph/ID-λ, ID-graph/OOD-λ, OOD-graph/OOD-λ, OOD-graph/ID-λ.

KEY RESULTS

- **Strong zero-shot:** Meta-QAOA yields high-quality initial angles that already approach target thresholds.
- **Few-shot boosts:** Additional SPSA steps reliably improve on zero-shot, often reaching targets faster than the baseline.
- **Outperforms classical:** Across ID and OOD splits, Meta-QAOA surpasses the factorized baseline in median quality and TTT.

ABLATIONS (ENTANGLEMENT)

- **Remove ZZ (γ=0) or replace by only single-qubit RZ:** The warm-start advantage largely collapses.
- **Conclusion:** Two-qubit/entangling structure is crucial for the observed gains.

INTERPRETATION

- The performance improvement appears to stem from the problem-aligned, entanglement-capable quantum ansatz combined with an amortized meta-map, not merely from parameter tuning.
- This is *evidence of a practical quantum-native benefit*, distinct from a formal proof of quantum advantage.

LIMITATIONS & OUTLOOK

- Current experiments use small N, simulated noise, and a non-entangling classical baseline.
- Future work: larger scales, hardware validation, stronger classical baselines, and alternative meta-parameterizations.

Takeaway *Meta-QAOA delivers robust zero-/few-shot warm-starts and consistently beats a strong factorized classical baseline. The gain aligns with the entanglement-enabled structure of QAOA and the flexibility of the Poly+RBF meta-map.*

6 Global Optimization

Here looking for the global best solution. , finding the optimal Parameter of the PH and its optimal solution.

Frage: Kann es sein, dass bei linearer Parametrisierung in $\boldsymbol{\lambda}$—also mit einer oder mehreren linearen Variablen—das Optimum in $\boldsymbol{\lambda}$ tatsächlich immer an den Rändern des Definitionsbereichs liegt (bzw. bei mehreren Variablen an den Box-Rändern)? Das würde auch sinn machen weil es sich um eine Affine Transformation handelt. Ich habe dies zur Intuition zusätzlich simuliert; auch dort treten die Optima durchweg an den Rändern auf. Damit würde für den Linearen Fall das Optimierungsproblem bedeutungslos werden - Markus

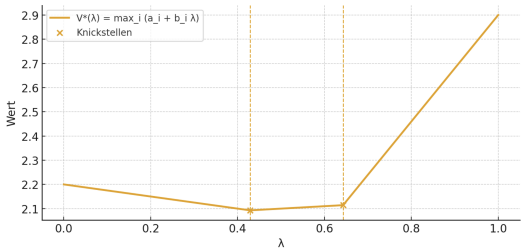

7 Discussion

8 Conclusion

References

Panagiotis Kl. Barkoutsos, Giacomo Nannicini, Peter J. Karalekas, Andrew Robert, Igor O. Sokolov, Nikolaj Moll, Antonio Mezzacapo, Piero Farquhar, Maria H. Roth, et al. Improving variational quantum optimization using CVaR. *Quantum*, 4:256, 2020. doi: 10.22331/q-2020-04-20-256.

M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles. Variational quantum algorithms. *Nature Reviews Physics*, 3 (9):625–644, 2021. doi: 10.1038/s42254-021-00348-9.

Alberto Cervera-Lierta, Jakob S. Kottmann, and Alán Aspuru-Guzik. Meta-variational quantum eigensolver: Learning energy profiles of parameterized hamiltonians. *PRX Quantum*, 2(2): 020329, 2021. doi: 10.1103/PRXQuantum.2.020329.

Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm, 2014.

Ioannis Kolotouros and Petros Wallden. Evolving objective function for improved variational quantum optimization. *Physical Review Research*, 4(2):023225, 2022. doi: 10.1103/PhysRevResearch.4.023225.

Andrew Lucas. Ising formulations of many NP problems. *Frontiers in Physics*, 2:5, 2014. doi: 10.3389/fphy.2014.00005.

Jiahao Miao, Chang-Yu Hsieh, and Sheng-Xin Zhang. Neural-network-encoded variational quantum algorithms. *Physical Review Applied*, 21(1):014053, 2024. doi: 10.1103/PhysRevApplied.21.014053.

Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O'Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5:4213, 2014. doi: 10.1038/ncomms5213.

John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018. doi: 10.22331/q-2018-08-06-79.

A APPENDIX

You may include other additional sections here.