
Meta Variational Monte Carlo

Tianchen Zhao¹, James Stokes², Oliver Knitter¹, Brian Chen¹, Shravan Veerapaneni^{1,2}

¹Department of Mathematics, University of Michigan, Ann Arbor, MI 48109

²Flatiron Institute, Simons Foundation, New York, NY 10010

{ericolon, knitter, chenbri, shravan}@umich.edu
jstokes@flatironinstitute.org

Abstract

An identification is found between meta-learning and the problem of determining the ground state of a randomly generated Hamiltonian drawn from a known ensemble. A model-agnostic meta-learning approach is proposed to solve the associated learning problem and a preliminary experimental study of random Max-Cut problems indicates that the resulting Meta Variational Monte Carlo accelerates training and improves convergence.

1 Introduction

This paper concerns the development of heuristic approximation algorithms for determining the ground state (minimal eigenpair) of certain large random Hermitian matrices drawn from known ensembles of relevance to Physics. Concretely, we focus on the real vector subspace of Hermitian matrices consisting of 2-local, n -qubit Hamiltonians, which are enumerated by a number of parameters only polynomial in n . These 2-local Hamiltonians in particular are expressive enough to capture all quadratic unconstrained binary optimization (QUBO) problems.

The Variational Monte Carlo (VMC) [9] is a heuristic algorithm which, given a local Hamiltonian H , produces an estimate for the minimal eigenvalue $\lambda_{\min}(H)$ and a description of an associated eigenvector. By exploiting neural networks as trial wavefunctions, Carleo and Troyer [3] showed that VMC can achieve state-of-the-art results for the ground state energies of physically important magnetic spin models. The domain of applicability of so-called neural-network quantum states has since been expanded considerably. It was recently shown, for example, that in the case of binary optimization, VMC is equivalent to Natural Evolution Strategies (NES) [7] and state-of-the-art results can be achieved for the Max-Cut Hamiltonian [16], albeit at the expense of significantly increased computation time compared to the best known classical heuristic approximation algorithms. The slowdown is attributed to the gradient-based training loop within the VMC. In this step, the neural network parameters $\theta \in \mathbb{R}^d$ are guided in the direction of steepest descent of the following unbiased estimator of the Rayleigh quotient which upper bounds the lowest eigenvalue:

$$L(\theta) := \frac{\langle \psi_\theta | H | \psi_\theta \rangle}{\langle \psi_\theta | \psi_\theta \rangle} = \mathbb{E}_{x \sim |\psi_\theta(\cdot)|^2} \left[\frac{\langle x | H | \psi_\theta \rangle}{\psi_\theta(x)} \right] \geq \lambda_{\min}(H) , \quad (1)$$

where $|x\rangle$ denotes the computational basis of $(\mathbb{C}^2)^{\otimes n}$ and the vector $|\psi_\theta\rangle \in (\mathbb{C}^2)^{\otimes n}$ denotes a neural-network quantum state in the sense that the amplitude $\psi_\theta(x) := \langle x | \psi_\theta \rangle$ is computed by a neural network with variational parameters θ . The expectation value is taken with respect to the associated Born probabilities, which are proportional to $|\psi_\theta(x)|^2$. In practice, the objective function is optimized using a variant of stochastic minibatch gradient descent called Stochastic Reconfiguration [12].

In contrast to the canonical formulation of VMC, which accepts a single Hamiltonian as input, in this paper we propose the *meta-VMC*, which asks for an approximation of the ground energy for

an *ensemble* of 2-local Hamiltonia H_τ indexed by a random disorder parameter τ sampled from a known distribution \mathcal{T} . The simplest strategy of retraining a separate neural-network quantum state from scratch for each realization of the disorder parameter τ is clearly impractical. The goal is thus shifted to finding a neural network that is maximally adaptive to new realizations of the disorder. For experimental support, we focus in this paper on the special case of QUBO problems, which can be identified with a Hamiltonian that is diagonal in the Pauli- Z basis. These preliminary results are viewed as a stepping stone to the more physically interesting problem of random quantum spin models, in which we anticipate similar optimization considerations to apply.

The formulation of meta-VMC exhibits obvious parallels with *meta-learning* or *learning to learn* in the Machine Learning literature [13], where data from previously encountered learning tasks is employed to accelerate performance on new tasks, drawn from an underlying task distribution \mathcal{T} . In the language of meta-learning, τ indexes the learning task and \mathcal{T} denotes the distribution over all tasks. In meta-VMC, we assume the task distribution is known to the learner; in contrast, conventional meta-learning assumes \mathcal{T} is unknown but possesses sufficient regularity to render meta-learning feasible. A second, less significant distinction is that the per-task objective function for meta-VMC is an unbiased estimator for the population objective, whereas meta-learning typically focuses on Empirical Risk Minimization objectives, which suffer from nonzero bias.

2 Theory

A simple strategy that has proven successful in meta-learning of deep neural networks is multi-task transfer learning [4, 2], which aims to learn an initialization for subsequent tasks by jointly optimizing the learning objective of multiple tasks simultaneously, using a minibatch training strategy that interleaves batches across the tasks. Multi-task learning is, however, prone to catastrophic interference [8], making it unsuitable for generalization to the VMC. The problem is exemplified by some of the simplest examples of disordered spin systems: suppose H_τ is a random Hamiltonian whose expected value under the disorder parameter vanishes $\mathbb{E}_{\tau \sim \mathcal{T}}[H_\tau] = 0$. As a concrete example, consider the Sherrington-Kirkpatrick Hamiltonian, in which τ represents a collection of i.i.d. centered Gaussian random variables $J_{ij} \sim N(0, 1)$ representing the exchange energies. If we denote by L_τ the objective function corresponding to disorder parameter τ , then the multi-task learning objective function, expressed in the population limit, is given by

$$L_{\text{MTL}}(\theta) := \mathbb{E}_{\tau \sim \mathcal{T}} [L_\tau(\theta)] = \mathbb{E}_{\tau \sim \mathcal{T}} \left\{ \mathbb{E}_{x \sim |\psi_\theta(\cdot)|^2} \left[\frac{\langle x | H_\tau | \psi_\theta \rangle}{\psi_\theta(x)} \right] \right\} = \frac{\langle \psi_\theta | \mathbb{E}[H_\tau] | \psi_\theta \rangle}{\langle \psi_\theta | \psi_\theta \rangle} = 0 . \quad (2)$$

The fact that the multi-task learning objective loses dependence on θ in the population limit implies that the associated minibatch algorithm makes no progress asymptotically.

In order to define an objective function which is asymptotically non-vacuous and which promotes adaptation to new realizations of disorder, we propose to optimize the following meta-learning objective function, again presented in population form for simplicity [1],

$$L_{\text{ML}}(\theta) := \mathbb{E}_{\tau \sim \mathcal{T}} [L_\tau(U_\tau^t(\theta))] = \mathbb{E}_{\tau \sim \mathcal{T}} [L_\tau(\underbrace{U_\tau \circ \dots \circ U_\tau}_{t \text{ times}}(\theta))] , \quad (3)$$

where $U_\tau^t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ denotes the t -fold application of a task adaptation operator U_τ , which in the simplest case of gradient descent with step size β , is given by $U_\tau(\theta) = \theta - \beta \nabla L_\tau(\theta)$. Optimization of the meta-learning objective L_{ML} ensures that when a new realization of the disorder parameter is drawn, the initialization performs well after performing one or more steps of gradient descent. Loosely speaking, meta-learning can be justified when one has a budget for running a few steps of gradient descent.

The fact that the meta-learning objective function manages to avoid the catastrophic interference phenomenon can be illustrated by the following toy model¹. Rather than considering the Rayleigh quotient (1), consider the following ensemble of quadratic functions specified by a random positive-definite matrix $A \in \mathbb{R}^{d \times d}$ and a random vector $b \in \mathbb{R}^d$,

$$L_\tau(\theta) = \frac{1}{2} \langle \theta, A \theta \rangle - \langle b, \theta \rangle , \quad (4)$$

¹This quadratic model has also been analyzed in the context of convergence theory in [5].

where the random variable $\tau = (A, b)$ now corresponds to the task label. In the simplest setting of single-step ($t = 1$) meta-learning with vanilla update operator $U_\tau(\theta) = \theta - \beta \nabla L_\tau(\theta)$, the optimal solution of the multi-task and meta-learning objectives can be found in closed form,

$$\arg \min_{\theta \in \mathbb{R}^d} L_{\text{ML}}(\theta) = \mathbb{E}[A(I - \beta A)^2]^{-1} \mathbb{E}[(I - \beta A)^2 b] . \quad (5)$$

In the limit $\beta \rightarrow 0$ corresponding to multi-task learning, the optimal solution is found to only depend on the mean value of the random variable τ , whereas the meta-learner ($\beta > 0$) exploits information in the higher-order moments of τ .

In the case of meta-VMC, we consider gradient-based optimization. Specifically, we focus on model-agnostic meta-learning (MAML) [6] which is a gradient-based algorithm that has been proposed for optimizing the meta-learning objective. Straightforward application of the chain rule gives rise to the following gradient estimator for the meta-learning objective,

$$\nabla L_{\text{ML}}(\theta) = \mathbb{E}_{\tau \sim \mathcal{T}} [(U_\tau^t)'(\theta) \nabla L_\tau(U_\tau^t(\theta))] , \quad (6)$$

where $(U_\tau^t)'(\theta)$ denotes the Jacobian matrix of the function $U_\tau^t : \mathbb{R}^d \rightarrow \mathbb{R}^d$. The pseudocode for MAML is outlined in Algorithm 1. In order to facilitate readability, we have presented the algorithm with batching only in the task index, leaving the remaining expectation values (with respect to Born probabilities) in population form. In a practical algorithm, the intermediate variables θ_τ and ∇_τ are estimated stochastically using independent batches of data generated by the same task τ . Since the computation of the Jacobian involves an expensive back-propagation, first-order MAML (foMAML) has been proposed (e.g., [6, 10]), which is a simplification of MAML in which the Jacobian matrix is approximated by the identity matrix.

Algorithm 1: MAML [6] adapted to meta-VMC (batched over tasks).

Input: Matrix ensemble \mathcal{T} , adaptation operator U_τ , adaptation steps t
Initialize θ

while *not done* **do**

	Sample batch of disorder parameters $B \stackrel{\text{iid}}{\sim} \mathcal{T}$
	for <i>each</i> disorder parameter $\tau \in B$ do
	$\theta_\tau = U_\tau^t(\theta)$
	$\nabla_\tau = (U_\tau^t)'(\theta) \nabla L_\tau(\theta_\tau)$
	$\nabla = \frac{1}{ B } \sum_{\tau \in B} \nabla_\tau$
	$\theta \leftarrow \text{OPTIMIZER}(\theta, \nabla)$

3 Relationship with previous work

In this section, we differentiate our proposal from the uses of meta-learning that have been proposed in the quantum computing literature. In [15], for example, meta-learning has been proposed to mitigate various sources of noise, specifically shot noise and parameter noise. In the context of VMC, shot noise is analogous to variance associated with finite minibatches, whereas parameter noise has no clear analogue. Ref. [14] is the most similar to ours in that they consider meta-learning from known distributions. They differ by the choice to focus on variational quantum algorithms such as VQE and QAOA and by the fact that they do not use model-agnostic meta-learning. Instead, the meta-learning outer-loop involves training a separate recurrent neural network, similar to [1].

4 Experiments

The experiments focus on the Max-Cut problem which is defined in terms of a simple, undirected graph $G = (V, E)$ with binary-indicator adjacency matrix J . In order to attack the Max-Cut problem with the VMC, we encode the solution in the ground state of the following classical antiferromagnetic Ising Hamiltonian with exchange interaction energy matrix given by J ,

$$H_\tau = \sum_{1 \leq i < j \leq n} J_{ij} Z_i Z_j , \quad (7)$$

where Z_i denotes the Pauli- Z operator acting locally to the i th qubit. Since the Max-Cut Hamiltonian is diagonal in the Pauli- Z basis, it acts as a multiplication operator and the ground state can be chosen as a computational basis vector $|x\rangle$ corresponding to a maximal cut. The variational wavefunction was chosen to be a real-valued Boltzmann machine with n hidden units. The task distribution \mathcal{T} defining the ensemble of Max-Cut Hamiltonians was defined by the following procedure. The adjacency matrix A for a Bernoulli random graph with edge probability 0.5 was first chosen and fixed throughout the experiments. Sampling an adjacency matrix from the task distribution \mathcal{T} is performed by rounding the matrix $A + (X + X^T)/2$ to a 0/1 matrix, where X denotes an $n \times n$ matrix with entrywise Gaussian noise $X_{ij} \sim N(0, \sigma^2)$. The hyperparameter controlling the task diversity for this ensemble is thus the variance σ^2 .

During training, each iteration of the meta-learning loop involved independently sampling a batch of 16 tasks from \mathcal{T} . During testing, $N_{\text{test}} = 32$ adjacency matrices were sampled from \mathcal{T} and fixed for evaluation purposes. The inner loop used $t = 15$ iterations of vanilla SGD (step size $\beta = 0.01$ batch size 128), while the outer loop training used 100 iterations of vanilla SGD (step size $\alpha = 0.01$, batch size 16). The meta-learning experiments were conducted using MAML [6] and first-order MAML (foMAML) [6, 10]. For baselines, we compared against training a neural-network quantum state from scratch as well as a pre-trained initialization with fine-tuning. The learning curves, illustrated for different values of σ in Fig. 1, clearly show that model-agnostic meta-learning dramatically accelerates training compared to the baselines on the N_{test} testing Max-Cut instances, consistent with the goal of MAML in promoting adaptivity. The networks trained using MAML also found larger cut values compared to the baselines which appear to converge prematurely to suboptimal states.




Figure 1: VMC training curves for Max-Cut instances ($|V| = 50$) using different parameter initialization strategies. Two baselines are considered: **Random** refers to training from scratch and **Pretrain** initializes with the model pre-trained with the base graph. Experiments were conducted using three task distributions characterized by $\sigma \in \{0.3, 0.5, 0.7\}$. For each task distribution, $N_{\text{test}} = 32$ graph instances are sampled for testing. We train the models on all testing graphs for 300 VMC iterations and evaluate the performance with the averaged approximation ratio, which is defined to be the ratio between the cut number and optimal value of the SDP relaxation. Models initialized from MAML and foMAML can discover near-optimal solutions within very few iterations, whereas pre-trained models get stuck in the local maximum. Moreover, MAML and foMAML provide better model parameter initializations that outperform the baselines in the long run. The advantage diminishes, however, as σ gets large.

5 Discussion and future directions

The preliminary experimental results for Max-Cut ensembles indicate that MAML effectively solves meta-VMC by accelerating training and improving convergence. While the Max-Cut problem is exactly solvable for the graph sizes considered here (e.g., by the Branch and Bound method [11]), our work paves the way to investigate physically-interesting matrix ensembles that cannot be diagonalized by a local change of basis. The ideas presented in this paper naturally extend also to variational quantum algorithms (VQAs) such as the variational quantum eigensolver. The key difference in the case of VQAs is that the denominator in the Rayleigh quotient (1) is automatically normalized $\langle \psi_\theta | \psi_\theta \rangle = 1$ and stochastic estimation of the quantum expectation value $\langle \psi_\theta | H_\tau | \psi_\theta \rangle$ involves performing measurements in multiple bases if the Hamiltonian contains non-commuting terms. The exploration of meta-VQA and associated learning algorithms is left to future work.

Broader Impact

The authors have not identified any ethical impacts or future societal consequences of this work.

Acknowledgments and Disclosure of Funding

The authors would like to thank Giuseppe Carleo for many helpful discussions. Authors gratefully acknowledge support from NSF under grant DMS-2038030.

References

- [1] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in neural information processing systems*, pages 3981–3989, 2016.
- [2] Jonathan Baxter. A model of inductive bias learning. *Journal of artificial intelligence research*, 12:149–198, 2000.
- [3] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.
- [4] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- [5] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. On the convergence theory of gradient-based model-agnostic meta-learning algorithms. In *International Conference on Artificial Intelligence and Statistics*, pages 1082–1092, 2020.
- [6] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [7] Joseph Gomes, Keri A McKiernan, Peter Eastman, and Vijay S Pande. Classical quantum optimization with neural network quantum states. *arXiv preprint arXiv:1910.10675*, 2019.
- [8] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.
- [9] W. L. McMillan. Ground state of liquid he^4 . *Phys. Rev.*, 138:A442–A451, 1965.
- [10] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [11] Franz Rendl, Giovanni Rinaldi, and Angelika Wiegele. Solving max-cut to optimality by intersecting semidefinite and polyhedral relaxations. *Mathematical Programming*, 121(2):307, 2010.
- [12] Sandro Sorella. Green function monte carlo with stochastic reconfiguration. *Physical Review Letters*, 80(20):4558–4561, 1998.
- [13] Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012.
- [14] Guillaume Verdon, Michael Broughton, Jarrod R McClean, Kevin J Sung, Ryan Babbush, Zhang Jiang, Hartmut Neven, and Masoud Mohseni. Learning to learn with quantum neural networks via classical neural networks. *arXiv preprint arXiv:1907.05415*, 2019.
- [15] Max Wilson, Sam Stromswold, Filip Wudarski, Stuart Hadfield, Norm M Tubman, and Eleanor Rieffel. Optimizing quantum heuristics with meta-learning. *arXiv preprint arXiv:1908.03185*, 2019.
- [16] Tianchen Zhao, Giuseppe Carleo, James Stokes, and Shravan Veerapaneni. Natural evolution strategies and variational Monte Carlo. *arXiv preprint arXiv:2005.04447*, 2020.