



Local Minima in Quantum Systems

Chi-Fang Chen

California Institute of Technology
Pasadena, USA
AWS Center for Quantum Computing
Pasadena, USA
chifang@caltech.edu

Hsin-Yuan Huang

California Institute of Technology
Pasadena, USA
Google Quantum AI
Venice, USA
Massachusetts Institute of Technology
Cambridge, USA
hsinyuan@caltech.edu

John Preskill

California Institute of Technology
Pasadena, USA
AWS Center for Quantum Computing
Pasadena, USA
preskill@caltech.edu

Leo Zhou

California Institute of Technology
Pasadena, USA
leozhou@caltech.edu

ABSTRACT

Finding ground states of quantum many-body systems is known to be hard for both classical and quantum computers. As a result, when Nature cools a quantum system in a low-temperature thermal bath, the ground state cannot always be found efficiently. Instead, Nature finds a local minimum of the energy. In this work, we study the problem of finding local minima in quantum systems under thermal perturbations. While local minima are much easier to find than ground states, we show that finding a local minimum is computationally hard for classical computers, even when the task is to output a single-qubit observable at any local minimum. In contrast, we prove that a quantum computer can always find a local minimum efficiently using a *thermal gradient descent algorithm* that mimics the cooling process in Nature. To establish the classical hardness of finding local minima, we consider a family of two-dimensional Hamiltonians such that any problem solvable by polynomial-time quantum algorithms can be reduced to finding local minima of these Hamiltonians. Therefore, cooling systems to local minima is universal for quantum computation, and, assuming quantum computation is more powerful than classical computation, finding local minima is classically hard and quantumly easy.

CCS CONCEPTS

• **Theory of computation** → Quantum computation theory; Quantum complexity theory; • **Mathematics of computing** → Mathematical optimization; • **Applied computing** → Physics.

KEYWORDS

quantum computing, many-body systems, quantum complexity theory, quantum thermodynamics, quantum algorithm, optimization

ACM Reference Format:

Chi-Fang Chen, Hsin-Yuan Huang, John Preskill, and Leo Zhou. 2024. Local Minima in Quantum Systems. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24)*, June 24–28, 2024, Vancouver, BC, Canada. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3618260.3649675>

1 INTRODUCTION

Finding ground states and other low-energy states of quantum many-body systems is a central problem in physics, materials science, and chemistry. To address this problem, many powerful computational methods, such as density functional theory (DFT) [32, 40], quantum Monte Carlo (QMC) [7, 12, 58], variational optimization with tensor network ansatzes [30, 34, 55, 59, 64, 67, 68] or neural network ansatzes [11, 22, 31], and data-driven machine learning approaches [27, 33, 43, 56], have been developed. These methods work well for many physically relevant problem instances but fail badly in other cases. One hopes that scalable fault-tolerant quantum computers will be able to solve a broader array of problem instances, but finding ground states of local Hamiltonians is known to be QMA-hard [37, 39], and therefore is expected to be intractable even for quantum computers.

Under the widely accepted conjecture that Nature can be efficiently simulated on a quantum computer, the hardness of finding ground states on quantum computers implies that Nature cannot find ground states in general. When a quantum system with Hamiltonian H is placed in a low-temperature thermal bath, the system seeks a local minimum of the energy, which may not be the ground state of H . For some physical systems, such as spin glasses [8, 24, 38, 49], finding a ground state is indeed known to be computationally hard; such systems, when cooled, almost always find a local minimum instead. In these cases, the ground state of the Hamiltonian is physically irrelevant in that it is rarely observed in experiments.



This work is licensed under a Creative Commons Attribution 4.0 International License.

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0383-6/24/06

<https://doi.org/10.1145/3618260.3649675>

Motivated by this perspective, in this work we study the problem of finding local minima in quantum many-body systems. For concreteness, we consider an n -qubit system governed by a local Hamiltonian H . The central question we are interested in is:

QUESTION. *How tractable is the problem of finding local minima of the energy in quantum systems using classical and quantum computers?*

To begin to answer this question, we need a mathematical definition of local minima in quantum systems. Based on the standard definition in mathematical optimization [2, 4, 9, 35, 54], we consider a local minimum in a quantum system governed by Hamiltonian H to be a quantum state such that the expectation value of H does not decrease under any small perturbation applied to the state. More formally, let $\mathcal{P}_\alpha$ be a perturbation parameterized by a small vector α that maps quantum states to quantum states. An ϵ -approximate local minimum of an n -qubit Hamiltonian H under perturbation $\mathcal{P}$ is a state ρ with an energy $\text{tr}(H\rho)$ that is an approximate minimum under perturbations, i.e.,

$$\text{tr}(H\rho) \leq \text{tr}(H\mathcal{P}_\alpha(\rho)) + \epsilon \|\alpha\| \quad (1.1)$$

for all small enough α . We say an algorithm $\mathcal{A}$ has solved the problem of *finding* local minima under perturbation $\mathcal{P}$ if given any n -qubit Hamiltonian H , written as a sum of few-qubit Hermitian operators, and any few-qubit observable O , the algorithm $\mathcal{A}$ can output a real value $\text{tr}(O\rho)$ corresponding to any approximate local minimum ρ of H under perturbations $\mathcal{P}$ up to a small error.¹ The local minima of H form a subset of the entire quantum state space, which contains the global minima, the ground states of H . We will consider two definitions of perturbations for defining local minima. The first one is, in a sense, mathematically natural but turns out to be inadequate for reasons we will explain. The second one is well-motivated physically and turns out to have interesting properties which we will explore.

The first definition of perturbations we study in this work is local unitary perturbations, which can be viewed as short-time unitary evolution governed by a sum of few-body Hermitian operators, as might arise in an adaptive variational quantum eigensolver (VQE) [13, 29, 53]. A drawback of this definition is that finding a local minimum becomes so easy that even a classical computer can solve it efficiently. We prove that a random n -qubit pure state is almost always a local minimum of H under local unitary perturbations. Hence, there are $\exp(\exp(\Omega(n)))$ many local minima that are not global minima in the energy landscape. Because the number of local minima is enormous, finding a local minimum under this definition is *classically easy*. While local unitary perturbations are natural from a mathematical perspective, they are not physically motivated since the evolution of a quantum system interacting with a low-temperature thermal bath is governed by *quantum thermodynamics* and is inherently nonunitary.

Our second definition is inspired by how quantum systems actually seek out local minima in Nature. Under suitable physical assumptions², perturbations induced by a thermal bath are represented by a master equation defined by a linear combination of

thermal Lindbladians $\mathcal{L}_a$, each associated with a local system-bath interaction A^a [10, 20, 44]. In its modern formulation [15, 48], the thermal Lindbladian $\mathcal{L}_a$ depends on the system Hamiltonian H and two macroscopic bath quantities: the inverse temperature β and a characteristic time scale τ . We prove two fundamental results concerning the problem of finding local minima under thermal perturbations. We prove that a quantum computer can efficiently find a local minimum under thermal perturbations using a proposed *quantum thermal gradient descent algorithm* that mimics Nature's cooling process. And in stark contrast to the definition of a local minimum based on local unitary perturbations, we prove that finding local minima under thermal perturbations is universal for quantum computation and, hence, is *classically hard* under the standard assumption $\text{BPP} \neq \text{BQP}$.

To establish the classical hardness of finding local minima under thermal perturbations, we consider geometrically local Hamiltonians on a 2D lattice, such that the ground state encodes the outcome of any efficient quantum computation using a modified version of Kitaev's circuit-to-Hamiltonian construction [3, 39, 52]. The most technically involved result of this work is a theorem stating that for these 2D Hamiltonians, all local minima under low-temperature thermal perturbations are global minima, i.e., ground states. That is, the energy landscape for these Hamiltonians has a nice bowl shape over the entire n -qubit state space such that quantum gradient descent efficiently finds the ground state starting with *arbitrary* initial states. Meanwhile, if a classical computer can always efficiently find any local minima under thermal perturbations, then the classical computer can efficiently simulate quantum computation, which is widely believed to be impossible. To prove the theorem, we develop a set of techniques for establishing that a Hamiltonian H has *no suboptimal local minima*, i.e., all local minima of H are global minima.

We now present, in more details, our main results concerning the tractability of finding local minima in quantum systems. The results are organized into the complexity of finding local minima under local unitary perturbations and under thermal perturbations.

2 LOCAL MINIMA UNDER LOCAL UNITARY PERTURBATIONS

We begin by presenting our results studying local minima under local unitary perturbations in more details. Local unitary perturbations are short-time unitary evolutions under a sum of few-body Hermitian operators. A quantum circuit consisting of near-identity two-qubit gates induces a local unitary perturbation. Consider an n -qubit pure state $|\psi\rangle$. A local unitary perturbation of $|\psi\rangle$ is given by

$$(\text{local unitary perturbation}): \quad |\psi\rangle \rightarrow \exp\left(-i \sum_{a=1}^m \alpha_a \mathbf{h}^a\right) |\psi\rangle, \quad (2.1)$$

where $\mathbf{h}^a$ is a Hermitian operator acting on a few qubits, $m = \text{poly}(n)$ is the number of such Hermitian operators, and $\alpha = \sum_a \alpha_a \mathbf{e}_a \in \mathbb{R}^m$ is a vector close to zero. This definition is inspired by adaptive variational quantum eigensolvers [13, 29, 53], and is the state version of the Riemannian geometry of quantum computation defined in [50]. When one variationally minimizes the energy

¹Since there could be multiple local minima and we consider finding one instance to be sufficient, this problem is closer to a *relational* problem than to a *decision* problem.

²Typical assumptions are that the system-bath coupling is weak and the thermal bath is memoryless.

by applying unitary gates, one finds a local minimum under local unitary perturbations.

To understand how easy the problem of finding local minima under local unitary perturbations is, we need to characterize the energy landscape. The following lemma provides a universal characterization of the structure of the energy landscape under the geometry defined by local unitary perturbations for any local Hamiltonian H . See the full version of this paper [14] for the formal statement and the proof.

Lemma 2.1 (Barren plateau; informal). *Given any n -qubit local Hamiltonian H , a random pure n -qubit state $|\psi\rangle$ is an approximate local minimum of H under local unitary perturbations.*

Furthermore, the proof of the above lemma illustrates the following physical picture: the energy landscape in the pure state space defined based on local unitary perturbations consists of a large barren plateau [46] with doubly-exponentially many approximate local minima having exponentially small energy gradient. Additionally, almost all of the local minima have local properties that are exponentially close to that of the maximally mixed state. As a result, while finding ground states is classically hard, finding local minima under local unitary perturbations is classically trivial.

Theorem 1 (Classically easy to find local minima under local unitary perturbations; informal). *The problem of finding approximate local minima of n -qubit local Hamiltonian H under local unitary perturbations is classically easy.*

The presence of barren plateaus in the energy landscape under local unitary perturbations causes the problem of finding local minima to be classically easy. However, a definition of local minima based on local unitary perturbation is not physically well motivated since Nature cools a physical system via open-system dynamics by coupling to a thermal bath rather than by unitary dynamics.

3 LOCAL MINIMA UNDER THERMAL PERTURBATIONS

Next, we consider local minima under thermal perturbations, which are short-time open-system evolution undergone by a system coupled to a heat bath. Under suitable physical assumptions, the thermal perturbations are represented by a master equation defined by a linear combination of *thermal Lindbladians* $\mathcal{L}_a$, each associated with a system-bath interaction A^a acting on a few qubits [10, 20, 44].

3.1 Thermal Lindbladians

In its modern formulation [15, 48], the thermal Lindbladian $\mathcal{L}_a$ depends on the system Hamiltonian H and two macroscopic bath quantities: the inverse temperature β and a characteristic time scale τ . Mathematically, one can think of the thermal Lindbladian as the quantum analog of Glauber dynamics, generating a continuous-time quantum Markov chain.

For each system-bath interaction specified by a jump operator A^a , the thermal Lindbladian is a sum over the coherent part and

the dissipative part:

$$\begin{aligned} \mathcal{L}_a^{\beta, \tau, H}(\rho) := & -i[H_{LS,a}^{\beta, \tau, H}, \rho] \\ & + \int_{-\infty}^{\infty} \gamma_\beta(\omega) \left[\hat{A}^a(\omega) \rho \hat{A}^a(\omega)^\dagger - \frac{1}{2} \{ \hat{A}^a(\omega)^\dagger \hat{A}^a(\omega), \rho \} \right] d\omega. \end{aligned} \quad (3.1)$$

We will focus on explaining the dissipative part in the second line.

Transition weight. At a fixed inverse temperature β , the *transition weight* $\gamma_\beta(\omega)$ tells us how strong the rate of a transition/jump should be, depending on the energy difference ω . In particular, the transition weight satisfies the following *Kubo-Martin-Schwinger (KMS) condition* and convenient normalization:

$$\gamma_\beta(\omega)/\gamma_\beta(-\omega) = e^{-\beta\omega} \quad \text{and} \quad 0 \leq \gamma_\beta(\omega) \leq 1 \quad \text{for any } \omega \in \mathbb{R}, \quad (3.2)$$

which is reminiscent of how detailed balance is imposed in classical Markov chains.

Operator Fourier transform. Given a jump operator A^a , we consider the *operator Fourier Transform* [15] for the Heisenberg-evolved jump operator $A^a(t) = e^{iHt} A^a e^{-iHt}$ characterized by a time scale $\tau \in \mathbb{R}$ of the heat bath

$$\hat{A}^a(\omega) := \frac{1}{\sqrt{2\pi\tau}} \int_{-\tau/2}^{\tau/2} e^{iHt} A^a e^{-iHt} e^{-i\omega t} dt. \quad (3.3)$$

The operator $\hat{A}^a(\omega)$ corresponds to matrix elements in A^a that induce jumps between energy eigenstates with an energy difference approximately $\omega \pm O(1/\tau)$. The bigger τ is, the more precise ω corresponds to the true energy difference.

Altogether, the effect of a thermal Lindbladian is to weight the transitions due to a jump A^a by the energy difference ω ; the heating transition $\omega > 0$ is relatively suppressed compared to the cooling transitions $\omega < 0$. Remarkably, this is not a carefully designed optimization algorithm to help with the analysis, but rather something that arises naturally from first principles in thermodynamics; it's Nature's optimization algorithm.

3.2 Formulating Local Minima

Now, we consider a thermal perturbation of n -qubit state ρ to be

$$(\text{thermal perturbation}): \quad \rho \rightarrow \exp \left(\sum_{a=1}^m \alpha_a \mathcal{L}_a^{\beta, \tau, H} \right) (\rho), \quad (3.4)$$

where $\mathcal{L}_a^{\beta, \tau, H}$ is the thermal Lindbladian, $m = \text{poly}(n)$ is the number of jump operators, and $\alpha = \sum_a \alpha_a \hat{e}_a \in \mathbb{R}_{\geq 0}^m$ is a *nonnegative* vector close to zero. Here, the vector is nonnegative because thermodynamic processes are generally irreversible.

A local minimum under thermal perturbations is a state ρ with the minimum energy $\text{tr}(H\rho)$ under thermal perturbations given in Eq. (3.4). More precisely, we will consider ϵ -approximate local minima as in Eq. (1.1). A central concept that enables us to understand the energy landscape and establishes the computational complexity of finding local minima under thermal perturbations is the *energy gradient operator*,

$$(\text{energy gradient operator}): \quad \sum_{a=1}^m \mathcal{L}_a^{\dagger \beta, \tau, H}(H) \hat{e}_a, \quad (3.5)$$

where the adjoint $\mathcal{L}^\dagger$ is the Heisenberg-picture Lindbladian, i.e., $\text{tr}(\mathcal{L}^\dagger[\mathcal{O}]\rho) = \text{tr}(\mathcal{O}\mathcal{L}[\rho])$. The energy gradient operator is a vector of individual gradient operators³ associated with each jump operator A^a . Indeed, the energy gradient operator naturally emerges by taking an infinitesimal perturbation, i.e., the gradient of the energy $\text{tr}(\mathbf{H}\rho)$,

$$\begin{aligned} & \text{tr}\left(\mathbf{H} \exp\left(\sum_{a=1}^m \alpha_a \mathcal{L}_a^{\beta, \tau, \mathbf{H}}\right)(\rho)\right) \\ &= \text{tr}(\mathbf{H}\rho) + \alpha \cdot \sum_{a=1}^m \text{tr}\left(\mathcal{L}_a^{\dagger \beta, \tau, \mathbf{H}}(\mathbf{H})\rho\right) \hat{e}_a + O(\|\alpha\|^2). \end{aligned} \quad (3.6)$$

The expectation of the energy gradient operator on a state ρ gives the gradient of the energy $\text{tr}(\mathbf{H}\rho)$ under thermal perturbations.

With the definition of local minima under thermal perturbations in hand, we next study how tractable is the problem of finding a local minimum under thermal perturbations.

3.3 Finding Local Minima is Easy for Quantum Computers

In practice, quantum systems find local minima easily when coupled to a cold thermal bath. Therefore, if our definition of a local minimum properly captures how a quantum system behaves in a cold environment, we expect finding local minima to be quantumly easy. Indeed, in the following theorem, we prove that a quantum computer can always efficiently find a local minimum of $\mathbf{H}$ under thermal perturbations by simulating thermal Lindbladians.

Theorem 2 (Quantumly easy to find a local minimum under thermal perturbations). *Let n be the problem size. There is a $\text{poly}(n)$ -time quantum algorithm that guarantees the following: Suppose we are given error $\epsilon = 1/\text{poly}(n)$, inverse temperature $0 \leq \beta \leq \text{poly}(n)$, time scale $\tau = \text{poly}(n)$, an n -qubit local Hamiltonian $\mathbf{H}$ with $\|\mathbf{H}\|_\infty = \text{poly}(n)$, m local jump operators $\{A^a\}_{a=1}^m$ with $m = \text{poly}(n)$, and a local observable $\mathcal{O}$ with $\|\mathcal{O}\|_\infty \leq 1$.*

Then, the quantum algorithm outputs a real value $v \in [-1, 1]$, such that v is ϵ -close to $\text{tr}(\mathcal{O}\rho)$ for an ϵ -approximate local minimum ρ of $\mathbf{H}$ under thermal perturbations with an inverse temperature β , a time scale τ , and system-bath interactions generated by $\{A^a\}_a$.

PROOF IDEA. We prove this theorem by developing a *quantum thermal gradient descent algorithm* based on the energy gradient operator. Gradient descent is necessary when the inverse temperature β and time scale τ are not infinite. When $\beta = \tau = \infty$, the energy gradient $\mathcal{L}_a^{\dagger \infty, \infty, \mathbf{H}}(\mathbf{H}) \leq 0$. In this case, the algorithm can just perform a random walk along random directions because no perturbations increase energy. But when β and τ are finite, the energy gradient can be positive and the algorithm needs to carefully walk in directions with negative energy gradients.

The quantum thermal gradient descent algorithm mimics how Nature cools the quantum system when the system is interacting locally and weakly with a low-temperature heat bath. The algorithm starts with an arbitrary initial state ρ_0 . For each step $t = 0, 1, 2, \dots$, the algorithm considers the current state ρ_t and proposes the next state ρ_{t+1} . The tangent space at ρ_t is high dimensional with many

possible directions/dynamics depending on the system-bath interaction. The algorithm chooses a direction that lowers the energy as fast as possible by computing the gradient of the energy and proposes ρ_{t+1} by performing gradient descent. As long as the current state ρ_t is not an ϵ -approximate local minimum of $\mathbf{H}$ under thermal perturbations, the energy will decrease by a sufficiently large amount

$$\text{tr}(\mathbf{H}\rho_{t+1}) < \text{tr}(\mathbf{H}\rho_t) - \frac{1}{\text{poly}(n)}. \quad (3.7)$$

Because the energy is bounded from below, there are, at most, a polynomial number of steps $t \leq \text{poly}(n)$ until the algorithm arrives at an ϵ -approximate local minimum of $\mathbf{H}$ under thermal perturbations.

To prove the convergence of quantum thermal gradient descent, we show that every small gradient step decreases the energy. To establish this claim, we derive analytic properties of thermal Lindbladians based on a smoothness bound on the second derivatives in [15]. To implement a gradient step based on thermal perturbations, we build on a recently developed efficient quantum algorithm that simulates thermal Lindbladian evolution using a quantum circuit augmented by mid-circuit measurements [15]. $\square$

3.4 Finding Local Minima is Hard for Classical Computers

Given that finding local minima under local unitary perturbations is classically trivial, it is natural to wonder whether finding local minima under thermal perturbations is also classically easy. What does the corresponding energy landscape look like? And what computational problems can be solved using quantum thermal gradient descent (or, roughly speaking, by just putting the system in a cold bath)? As our second main result, we address these questions for a class of geometrically local Hamiltonians $\{\mathbf{H}_C\}$ on two-dimensional lattices, where the ground state $|\eta_0\rangle$ encodes the output of quantum circuit C .

Theorem 3 (All local minima are global in BQP-hard Hamiltonians). *Let $P_G(\mathbf{H}_C) = |\eta_0\rangle\langle\eta_0|$ be the ground state projector of the 2D Hamiltonian $\mathbf{H}_C$ acting on $n + T = \text{poly}(n)$ qubits. There is a choice of $m = \text{poly}(n)$ two-qubit jump operators $\{A^a\}_a$ satisfying the following.*

Suppose we are given $0 < \delta < 1$. Then, for any small error $\epsilon = 1/\text{poly}(n, 1/\delta)$, any ϵ -approximate local minimum ρ of $\mathbf{H}_C$ under thermal perturbations generated by $\{A^a\}_a$ with a large inverse temperature $\beta = \text{poly}(n, 1/\delta)$, a large time scale $\tau = \text{poly}(n, 1/\delta)$ is an exact global minimum with high probability, i.e., we have $\text{tr}(P_G(\mathbf{H}_C)\rho) \geq 1 - \delta$.

This theorem is the most technically involved contribution of this work. Conceptually, the landscape of these 2D Hamiltonians has a nice *bowl shape*, like in convex optimization [9]. Therefore, performing thermal gradient descent (Theorem 2) allows us to prepare the ground state starting from an *arbitrary* initial state. For a choice of inverse temperature that grows polynomially with $|C|$, thermal fluctuations in the cooling process do not kill the power of quantum computation.

Showing that no suboptimal local minimum exists seems daunting due to the doubly exponentially large space of possible quantum

³This is similar to the spin operator $\vec{\sigma} = \sigma^x \hat{e}_x + \sigma^y \hat{e}_y + \sigma^z \hat{e}_z$, which is a vector of Hermitian observables.

states as well as the complex expression for the thermal Lindbladian $\mathcal{L}_a^{\beta, \tau, H}$. While previous studies on circuit-to-Hamiltonian mappings focused mainly on the lowest energy states, here we need to worry about potential local minima from states in superposition with *arbitrarily high energies*. To make progress, we propose a sufficient condition based on the energy gradient operator in Eq. (3.5) that captures the nice landscape of H_C and rules out the presence of *any* suboptimal local minimum: Let $P_G(H)$ be the projector onto the ground state space of H . The condition requires the existence of a unit vector $\hat{\alpha} \in \mathbb{R}_{\geq 0}^m$ and $r > 0$ with

$$-\sum_{a=1}^m \hat{\alpha}_a \mathcal{L}_a^{\dagger \beta, \tau, H}(H) \geq r(I - P_G(H)). \quad (3.8)$$

This operator inequality, which we call the “*negative gradient condition*,” implies that any state with a small ground state overlap must experience a substantially negative energy gradient, i.e., it must not be a local minimum.

To prove that H_C satisfies the sufficient condition, we propose a series of mathematical techniques for characterizing energy gradients in few-qubit systems, in commuting Hamiltonians, in subspaces of the Hamiltonian, and in perturbed Hamiltonians. Controlling perturbations of the energy gradient is surprisingly challenging. The perturbative errors are not suppressed by the spectral gap of the Hamiltonian as seen in more standard settings but instead by the Bohr-frequency gap, which can be much smaller. These new techniques build on the operator Fourier transform, and the secular approximation in [15]. We emphasize that while we proved that H_C has no suboptimal local minima, this is not true for any local Hamiltonian. In the case of Kitaev’s QMA-hard Hamiltonian, the energy landscape contains a large number of suboptimal local minima corresponding to all possible rejected witnesses.

In particular, the family of Hamiltonians $\{H_C\}$ on two-dimensional lattices that we consider yields the following proposition:

Proposition 3.1 (BQP-hardness for estimating properties of the ground state of H_C). *If there is a classical algorithm that can estimate any single-qubit observable on the unique ground state of the geometrically local Hamiltonian H_C in time polynomial in the number of qubits in H_C to error $1/4$ for any H_C in the class, then BPP = BQP.*

Building on the energy landscape characterization in Theorem 3 and the proposition above, we establish the following:

Theorem 4 (Classically hard to find a local minimum under thermal perturbations). *Let n be the problem size. Suppose there is a poly(n)-time classical algorithm guaranteeing the following:*

Given error $\epsilon = 1/\text{poly}(n)$, inverse temperature $0 \leq \beta \leq \text{poly}(n)$, time scale $0 \leq \tau \leq \text{poly}(n)$, an n -qubit local Hamiltonian H with $\|H\|_\infty = \text{poly}(n)$, m local jump operators $\{A^a\}_{a=1}^m$ with $m = \text{poly}(n)$, and a single-qubit observable O with $\|O\|_\infty \leq 1$. Then, the classical algorithm outputs a real value $v \in [-1, 1]$, such that v is ϵ -close to $\text{tr}(Op)$ for an ϵ -approximate local minimum p of the Hamiltonian H under thermal perturbations with an inverse temperature β , a time scale τ , and system-bath interactions generated by $\{A^a\}_a$.

Then, BPP = BQP.

PROOF. Assuming the existence of a polynomial-time classical algorithm that satisfies the properties stated in the theorem. Apply

this classical algorithm to the 2D Hamiltonian H_C considered in Theorem 3 with a sufficiently small approximation error ϵ , such that any ϵ -approximate local minimum p of H_C under thermal perturbations with polynomially-large β , τ and system-bath interactions generated by $\{A^a\}_a$ is an exact global minimum with high probability, i.e.,

$$\langle \eta_0 | p | \eta_0 \rangle = \text{tr}(P_G(H_C)p) \geq 1 - \frac{1}{16^2}, \quad (3.9)$$

where $|\eta_0\rangle$ is the unique ground state of H_C . We further consider ϵ to be small enough such that

$$\epsilon < \frac{1}{8}. \quad (3.10)$$

Let p be an ϵ -approximate local minimum of the Hamiltonian H under thermal perturbations. Consider the observable $O_j = Z_j$ from the proof of Proposition 3.1. Using the Fuchs–van de Graaf inequalities, we have

$$\|p - |\eta_0\rangle\langle\eta_0|\|_1 \leq \frac{1}{8}. \quad (3.11)$$

Because the classical algorithm can estimate $\text{tr}(O_j p)$ to error ϵ , from Eq. (3.10) and (3.11), the classical algorithm can estimate $\langle \eta_0 | O_j | \eta_0 \rangle$ to error $1/4$ in time polynomial in the number of qubits in H_C . From Prop. 3.1, this implies that BPP = BQP. $\square$

There have been other proposals for solving BQP-hard problems by finding suitable quantum states, such as designing a gapped adiabatic path for Hamiltonians to find ground states [3], engineering Lindbladians to have rapid dissipative evolution towards steady states [65] and performing quantum phase estimation on an initial state with high ground-state overlap [26]. These approaches draw inspiration from physics to motivate algorithms for solving problems on analog and digital quantum devices but do not emulate naturally occurring physical processes. In contrast, the problem of finding a local minimum is motivated by ubiquitous physical processes in Nature that produce the low-energy states studied in physics, chemistry, and materials science. Furthermore, the local minima problem enjoys the robustness of thermodynamics: one merely needs to specify macroscopic bath quantities β and τ without worrying about microscopic details, and the choice of jump operators can be flexible since adding more jumps (even unwanted ones) only *improves* the gradient and *removes* suboptimal local minima.

4 DISCUSSION

We have good reasons for believing that scalable fault-tolerant quantum computers will be more powerful than classical computers, but for what problems of practical interest should we expect a super-polynomial quantum advantage? Quantum computers might substantially speed up the task of characterizing properties of ground states for some local Hamiltonians that arise in physics, chemistry, and materials science, but it is not clear how to identify particular problems for which such speedups occur [42]. In some cases, classical methods provide good solutions, while in other cases, the problem is hard even for quantum computers.

Here we have focused on an easier problem, namely finding local minima rather than global minima of a Hamiltonian. This problem is very well motivated physically because the task of finding a local

minimum under thermal perturbations is routinely carried out by actual physical systems when in contact with a low-temperature thermal bath. We showed that this problem is solved efficiently by a proposed quantum optimization algorithm, the *quantum thermal gradient descent algorithm*. Furthermore, we showed that finding a local minimum is classically hard in general (assuming that $\text{BPP} \neq \text{BQP}$). Hence, the local minimum problem is a quantumly tractable alternative to the ground state problem for which superpolynomial quantum advantage can be achieved for some problem instances.

Our main results pertain to perturbations that arise in quantum thermodynamics [10, 15, 20, 44, 48]. We noted that the energy landscape under such thermal perturbations is much nicer than the energy landscape encountered by quantum optimization algorithms relying on local unitary perturbations such as VQE [13, 29, 53]; see Theorems 1 and 3. From an algorithmic design perspective, we are free to choose any perturbation. Indeed, we may modify the thermal Lindbladians to have nicer analytic properties or algorithmic costs [15]. While these synthetic Lindbladians may not simulate Nature, they constitute a broader class of quantum Markov chain Monte Carlo (or, quantum Gibbs sampling) algorithms [15, 23, 57, 62] that may improve upon Nature. Apart from Lindbladians, other families of perturbations, such as unitary perturbations accompanied by mid-circuit measurements and/or qubit resets, may also yield nice bowl-shaped energy landscapes without suboptimal local minima. Progress on this question could lead to more efficient quantum optimization algorithms for finding low-energy states or for other applications.

There are a plethora of classical algorithms for minimizing energies of quantum systems based on classical variational ansatzes for quantum states, such as tensor networks [1, 6, 19, 30, 34, 36, 41, 55, 59–61, 63, 64, 67–70] and neural network quantum states [11, 16, 17, 21, 22, 25, 28, 31, 45, 47, 51, 66]. These classical algorithms find a local minimum within a family of states defined by the classical variational ansatz. However, a local minimum of the energy among the set of states subject to the classical ansatz might not be a local minimum under thermal perturbations. If not, we can load the state found by the classical algorithm into a quantum computer and find a lower energy state by running the quantum thermal gradient descent algorithm. A corollary of our main results states the following.

Corollary 4.1 (Quantum advantage over variationally optimized classical ansatz). *Under the conjecture that $\text{BPP} \neq \text{BQP}$, there exists a class of n -qubit geometrically-local Hamiltonian H on a two-dimensional lattice with $\|H\|_\infty = O(n)$ that satisfies the following. Given any classical ansatz of n -qubit state ρ that can estimate the expectation value of single-qubit observables to $1/\text{poly}(n)$ error in $\text{poly}(n)$ -time on classical computers, any $\text{poly}(n)$ -time classical algorithm for minimizing the energy $\text{tr}(H\rho)$ using the classical ansatz, and samples of the state ρ represented by the optimized classical ansatz. A quantum machine can find a state $\rho^\#$ with strictly lower energy than ρ in $\text{poly}(n)$ time by running a quantum thermal gradient descent based on low-temperature cooling.*

PROOF IDEA. The central claim is that the state ρ found by an efficient classical algorithm cannot be an ϵ -approximate local minimum under low-temperature thermal perturbations. We establish this claim by contradiction. Suppose that the classical ansatz for ρ

found by the efficient classical algorithm is an ϵ -approximate local minimum. Then the classical algorithm can use the classical ansatz to predict the expectation values of single-qubit observables of an ϵ -approximate local minimum ρ of H to ϵ error. From Theorem 4, this implies that $\text{BPP} = \text{BQP}$, which is a contradiction.

Because ρ is not an ϵ -approximate local minimum under low-temperature thermal perturbations, a quantum machine can use samples of ρ to initialize at the state ρ and perform one gradient descent step based on low-temperature cooling. From a lemma on the necessary condition for local minima, there exists $a \in \{1, \dots, m\}$ such that $\text{tr}(H \mathcal{L}_a^{\beta, \tau, H}[\rho]) < -\epsilon$. From a lemma on cooling by gradient descent, one can show that a single gradient descent step yields a state $\rho^{(\text{next})}$ with a strictly lower energy than the state ρ . Hence, one establishes the desired claim. $\square$

The point is that we have proved the existence of local Hamiltonians for which finding a local minimum is quantumly easy and classically hard. For any such Hamiltonian, any quantum state $\rho^\#$ found by the efficient classical algorithm will not be a local minimum; therefore, quantum thermal gradient descent will be able to descend to a state with strictly lower energy, even with just one gradient step. Furthermore, in many cases, we can evaluate the energy gradient at the classically optimized state $\rho^\#$ by executing an efficient classical computation. A negative energy gradient confirms that a quantum algorithm starting from $\rho^\#$ could outperform the classical algorithm.

Many other interesting and challenging questions remain open. Theorem 3 shows that there are no suboptimal local minima in BQP-hard n -qubit Hamiltonians for inverse temperature $\beta = \text{poly}(n)$. Do there exist BQP-hard Hamiltonians with no suboptimal local minimum even for constant temperature, i.e., $\beta = O(1)$? If so, quantum advantage can be achieved by simply coupling a quantum system to a heat bath at a sufficiently low but constant temperature. Our conclusion that finding local minima under thermal perturbations is classically hard relied on the complexity-theoretic conjecture that $\text{BPP} \neq \text{BQP}$. Can we prove unconditionally that finding local minima is hard for classical algorithms, perhaps within a black-box oracle model? Sometimes, when a system performs a random walk over a large plateau of suboptimal local minima for a sufficiently long time, the system escapes the plateau and reaches the true ground state. Could we characterize when ground states can be found efficiently despite having many suboptimal local minima? We have shown that there is a quantum advantage in finding local minima of quantum systems. Might there also be a quantum advantage in finding better local minima in classical optimization problems under some variant of quantum thermal gradient descents?

While ground state problems are hard to solve in general, many experimentally observed quantum systems efficiently relax to their ground states when cooled. This physical phenomenon suggests that perhaps many Hamiltonians of interest in physics, chemistry, and materials science have no suboptimal local minima. We have shown in Theorem 3 that a particular family of BQP-hard Hamiltonians has no suboptimal local minima under thermal perturbation. An important future goal is to characterize broader classes of Hamiltonians that have a similarly good energy landscape. Our proposed negative gradient condition suffices to rule out suboptimal local minima, but checking this condition for a general Hamilton involves

highly complex calculations. It would be helpful to develop more general-purpose and efficient methods to verify this property for specified physical Hamiltonians over spins, fermions, or bosons. We hope the ideas and techniques presented here will yield a deeper understanding of the energy landscapes of quantum systems and point toward promising opportunities for achieving quantum advantage for physically relevant problems.

ACKNOWLEDGMENTS

The authors thank Anurag Anshu, Ryan Babbush, Fernando Brandão, Garnet Chan, Sitan Chen, Soonwon Choi, Jordan Cotler, Jarrod R. McClean, and Mehdi Soleimanifar for their valuable input. HH thanks Patrick Coles, Gavin Crooks, and Faris Sباهي for the inspiring discussions and for sharing their recent works on classical thermodynamics for AI applications [5, 18]. CFC is supported by the AWS Center for Quantum Computing internship. HH is supported by a Google PhD fellowship and a MediaTek Research Young Scholarship. HH acknowledges the visiting associate position at Massachusetts Institute of Technology. LZ acknowledges funding from the Walter Burke Institute for Theoretical Physics at Caltech. JP acknowledges support from the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research (DE-NA0003525, DE-SC0020290), the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator, and the National Science Foundation (PHY-1733907). The Institute for Quantum Information and Matter is an NSF Physics Frontiers Center.

REFERENCES

- [1] Nilin Abrahamsen. 2019. A polynomial-time algorithm for ground states of spin trees. *arXiv preprint arXiv:1907.04862* (2019). arXiv:1907.04862 [quant-ph]
- [2] Naman Agarwal, Zeyuan Allen-Zhu, Brian Bullins, Elad Hazan, and Tengyu Ma. 2017. Finding approximate local minima faster than gradient descent. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 1195–1199. <https://doi.org/10.1145/3055399.3055464>
- [3] Dorit Aharonov, Wim van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. 2007. Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation. *SIAM J. Comput.* 37, 1 (April 2007), 166–194. <https://doi.org/10.1137/S0097539705447323>
- [4] Amir Ali Ahmadi and Jeffrey Zhang. 2022. On the complexity of finding a local minimizer of a quadratic function over a polytope. *Mathematical Programming* 195, 1-2 (2022), 783–792. <https://doi.org/10.1007/s10107-021-01714-2>
- [5] Maxwell Aifer, Kaelan Donatella, Max Hunter Gordon, Thomas Ahle, Daniel Simpson, Gavin E Crooks, and Patrick J Coles. 2023. Thermodynamic Linear Algebra. *arXiv preprint arXiv:2308.05660* (2023). arXiv:2308.05660 [cond-mat.stat-mech]
- [6] Itai Arad, Zeph Landau, Umesh Vazirani, and Thomas Vidick. 2017. Rigorous RG algorithms and area laws for low energy eigenstates in 1D. *Commun. Math. Phys.* 356, 1 (2017), 65–105. <https://doi.org/10.1007/s00220-017-2973-z>
- [7] Federico Becca and Sandro Sorella. 2017. *Quantum Monte Carlo Approaches for Correlated Systems*. Cambridge University Press. <https://doi.org/10.1017/9781316417041>
- [8] Kurt Binder and A Peter Young. 1986. Spin glasses: Experimental facts, theoretical concepts, and open questions. *Rev. Mod. Phys.* 58, 4 (1986), 801. <https://doi.org/10.1103/RevModPhys.58.801>
- [9] Stephen Boyd, Stephen P Boyd, and Lieven Vandenbergh. 2004. *Convex optimization*. Cambridge University Press.
- [10] Heinz-Peter Breuer and Francesco Petruccione. 2002. *The theory of open quantum systems*. Oxford University Press, USA.
- [11] Giuseppe Carleo and Matthias Troyer. 2017. Solving the quantum many-body problem with artificial neural networks. *Science* 355, 6325 (2017), 602–606. <https://doi.org/10.1126/science.aag2302>
- [12] David Ceperley and Berni Alder. 1986. Quantum Monte Carlo. *Science* 231, 4738 (1986), 555–560. <https://doi.org/10.1126/science.231.4738.555>
- [13] 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. <https://doi.org/10.1038/s42254-021-00348-9>
- [14] Chi-Fang Chen, Hsin-Yuan Huang, John Preskill, and Leo Zhou. 2023. Local minima in quantum systems. *arXiv preprint arXiv:2309.16596* (2023). arXiv:2309.16596 [quant-ph]
- [15] Chi-Fang Chen, Michael J Kastoryano, Fernando GSL Brandão, and András Gilyén. 2023. Quantum Thermal State Preparation. *arXiv preprint arXiv:2303.18224* (2023). arXiv:2303.18224 [quant-ph]
- [16] Kenny Choo, Giuseppe Carleo, Nicolas Regnault, and Titus Neupert. 2018. Symmetries and Many-Body Excitations with Neural-Network Quantum States. *Phys. Rev. Lett.* 121 (2018), 167204. Issue 16. <https://doi.org/10.1103/PhysRevLett.121.167204>
- [17] Kenny Choo, Antonio Mezzacapo, and Giuseppe Carleo. 2020. Fermionic neural-network states for ab-initio electronic structure. *Nature Communications* 11, 1 (May 2020), 2368. <https://doi.org/10.1038/s41467-020-15724-9>
- [18] Patrick J. Coles, Collin Szczepanski, Denis Melanson, Kaelan Donatella, Antonio J. Martinez, and Faris Sباهي. 2023. Thermodynamic AI and the fluctuation frontier. (2023). arXiv:2302.06584 [cs.ET]
- [19] Philippe Corboz. 2016. Variational optimization with infinite projected entangled-pair states. *Phys. Rev. B* 94 (2016), 035133. Issue 3. <https://doi.org/10.1103/PhysRevB.94.035133>
- [20] E Brian Davies. 1979. Generators of dynamical semigroups. *Journal of Functional Analysis* 34, 3 (1979), 421–432.
- [21] Dong-Ling Deng, Xiaopeng Li, and S. Das Sarma. 2017. Machine learning topological states. *Phys. Rev. B* 96 (2017), 195145. Issue 19. <https://doi.org/10.1103/PhysRevB.96.195145>
- [22] Dong-Ling Deng, Xiaopeng Li, and S. Das Sarma. 2017. Quantum Entanglement in Neural Network States. *Phys. Rev. X* 7 (2017), 021021. Issue 2. <https://doi.org/10.1103/PhysRevX.7.021021>
- [23] Zhiyan Ding, Chi-Fang Chen, and Lin Lin. 2023. Single-ancilla ground state preparation via Lindbladians. *arXiv preprint arXiv:2308.15676* (2023). arXiv:2308.15676 [quant-ph]
- [24] Samuel Frederick Edwards and Phil W Anderson. 1975. Theory of spin glasses. *Journal of Physics F: Metal Physics* 5, 5 (1975), 965.
- [25] Francesco Ferrari, Federico Becca, and Juan Carrasquilla. 2019. Neural Gutzwiller-projected variational wave functions. *Phys. Rev. B* 100 (2019), 125131. Issue 12. <https://doi.org/10.1103/PhysRevB.100.125131>
- [26] Sevag Gharibian and François Le Gall. 2022. Dequantizing the quantum singular value transformation: hardness and applications to quantum chemistry and the quantum PCP conjecture. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 19–32. <https://doi.org/10.1145/3519935.3519991>
- [27] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. 2017. Neural Message Passing for Quantum Chemistry. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 1263–1272. <https://proceedings.mlr.press/v70/gilmer17a.html>
- [28] Ivan Glasser, Nicola Pancotti, Moritz August, Ivan D. Rodriguez, and J. Ignacio Cirac. 2018. Neural-Network Quantum States, String-Bond States, and Chiral Topological States. *Phys. Rev. X* 8 (2018), 011006. Issue 1. <https://doi.org/10.1103/PhysRevX.8.011006>
- [29] Harper R Grimsley, Sophia E Economou, Edwin Barnes, and Nicholas J Mayhall. 2019. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature Communications* 10, 1 (2019), 3007. <https://doi.org/10.1038/s41467-019-10988-2>
- [30] Reza Haghsheenas, Matthew J. O’Rourke, and Garnet Kin-Lic Chan. 2019. Conversion of projected entangled pair states into a canonical form. *Phys. Rev. B* 100 (2019), 054404. Issue 5. <https://doi.org/10.1103/PhysRevB.100.054404>
- [31] Mohamed Hibat-Allah, Martin Ganahl, Lauren E Hayward, Roger G Melko, and Juan Carrasquilla. 2020. Recurrent neural network wave functions. *Phys. Rev. Res.* 2, 2 (2020), 023358. <https://doi.org/10.1103/PhysRevResearch.2.023358>
- [32] P. Hohenberg and W. Kohn. 1964. Inhomogeneous Electron Gas. *Phys. Rev.* 136 (1964), B864–B871. Issue 3B. <https://doi.org/10.1103/PhysRev.136.B864>
- [33] Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V. Albert, and John Preskill. 2022. Provably efficient machine learning for quantum many-body problems. *Science* 377, 6613 (2022). <https://doi.org/10.1126/science.abk3333>
- [34] Katharine Hyatt and E. M. Stoudenmire. 2019. DMRG Approach to Optimizing Two-Dimensional Tensor Networks. *arXiv preprint arXiv:1908.08833* (2019). arXiv:1908.08833 [cond-mat.str-el]
- [35] Chi Jin, Praneeth Netrapalli, and Michael I. Jordan. 2018. Accelerated Gradient Descent Escapes Saddle Points Faster than Gradient Descent. In *Proceedings of the 31st Conference On Learning Theory (Proceedings of Machine Learning Research, Vol. 75)*, PMLR, 1042–1085. <https://proceedings.mlr.press/v75/jin18a.html>
- [36] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac. 2008. Classical Simulation of Infinite-Size Quantum Lattice Systems in Two Spatial Dimensions. *Phys. Rev. Lett.* 101 (2008), 250602. Issue 25. <https://doi.org/10.1103/PhysRevLett.101.250602>
- [37] Julia Kempe, Alexei Kitaev, and Oded Regev. 2006. The Complexity of the Local Hamiltonian Problem. *SIAM J. Comput.* 35, 5 (2006), 1070–1097.

- [38] Scott Kirkpatrick and David Sherrington. 1978. Infinite-ranged models of spin-glasses. *Phys. Rev. B* 17, 11 (1978), 4384. <https://doi.org/10.1103/PhysRevB.17.4384>
- [39] A. Yu Kitaev, A. Shen, and M. N. Vyalyi. 2002. *Classical and Quantum Computation*. American Mathematical Society.
- [40] W. Kohn. 1999. Nobel Lecture: Electronic structure of matter—wave functions and density functionals. *Rev. Mod. Phys.* 71 (1999), 1253–1266. Issue 5. <https://doi.org/10.1103/RevModPhys.71.1253>
- [41] Zeph Landau, Umesh Vazirani, and Thomas Vidick. 2015. A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians. *Nature Physics* 11, 7 (2015), 566–569. <https://doi.org/10.1038/nphys3345>
- [42] Seunghoon Lee, Joonho Lee, Huanchen Zhai, Yu Tong, Alexander M Dalzell, Ashutosh Kumar, Phillip Helms, Johnnie Gray, Zhi-Hao Cui, Wenyuan Liu, et al. 2023. Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nature Communications* 14, 1 (2023), 1952. <https://doi.org/10.1038/s41467-023-37587-6>
- [43] Laura Lewis, Hsin-Yuan Huang, Viet T. Tran, Sebastian Lehner, Richard Kueng, and John Preskill. 2024. Improved machine learning algorithm for predicting ground state properties. *Nature Communications* 15, 1 (2024). <https://doi.org/10.1038/s41467-024-45014-7>
- [44] Goran Lindblad. 1976. On the generators of quantum dynamical semigroups. *Commun. Math. Phys.* 48 (1976), 119–130.
- [45] Di Luo, Zhuo Chen, Kaiwen Hu, Zhizhen Zhao, Vera Mikyoung Hur, and Bryan K. Clark. 2023. Gauge-invariant and anyonic-symmetric autoregressive neural network for quantum lattice models. *Phys. Rev. Res.* 5 (Mar 2023), 013216. Issue 1. <https://doi.org/10.1103/PhysRevResearch.5.013216>
- [46] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. 2018. Barren plateaus in quantum neural network training landscapes. *Nature Communications* 9, 1 (2018), 1–6. <https://doi.org/10.1038/s41467-018-07090-4>
- [47] Stewart Morawetz, Isaac J. S. De Vlucht, Juan Carrasquilla, and Roger G. Melko. 2021. U(1)-symmetric recurrent neural networks for quantum state reconstruction. *Phys. Rev. A* 104 (2021), 012401. Issue 1. <https://doi.org/10.1103/PhysRevA.104.012401>
- [48] Evgeny Mozgunov and Daniel Lidar. 2020. Completely positive master equation for arbitrary driving and small level spacing. *Quantum* 4 (Feb. 2020), 227. <https://doi.org/10.22331/q-2020-02-06-227>
- [49] John A Mydosh. 1993. *Spin glasses: an experimental introduction*. CRC Press.
- [50] Michael A Nielsen, Mark R Dowling, Mile Gu, and Andrew C Doherty. 2006. Quantum computation as geometry. *Science* 311, 5764 (2006), 1133–1135. <https://doi.org/10.1126/science.1121541>
- [51] Yusuke Nomura, Andrew S. Darmawan, Youhei Yamaji, and Masatoshi Imada. 2017. Restricted Boltzmann machine learning for solving strongly correlated quantum systems. *Phys. Rev. B* 96 (2017), 205152. Issue 20. <https://doi.org/10.1103/PhysRevB.96.205152>
- [52] Roberto Oliveira and Barbara M. Terhal. 2008. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Inf. Comput.* 8 (2008), 900–924.
- [53] Peter JJ O'Malley, Ryan Babbush, Ian D Kivlichan, Jonathan Romero, Jarrod R McClean, Rami Barends, Julian Kelly, Pedram Roushan, Andrew Tranter, Nan Ding, et al. 2016. Scalable Quantum Simulation of Molecular Energies. *Phys. Rev. X* 6 (Jul 2016), 031007. Issue 3. <https://doi.org/10.1103/PhysRevX.6.031007>
- [54] Panos M Pardalos and Stephen A Vavasis. 1991. Quadratic programming with one negative eigenvalue is NP-hard. *Journal of Global optimization* 1, 1 (1991), 15–22.
- [55] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac. 2007. Matrix Product State Representations. *Quantum Info. Comput.* 7, 5 (2007), 401–430.
- [56] Zhuoran Qiao, Matthew Welborn, Animashree Anandkumar, Frederick R. Manby, and Thomas F. Miller. 2020. OrbNet: Deep learning for quantum chemistry using symmetry-adapted atomic-orbital features. *The Journal of Chemical Physics* 153, 12 (2020). <https://doi.org/10.1063/5.0021955>
- [57] Patrick Rall, Chunhao Wang, and Pawel Wocjan. 2023. Thermal State Preparation via Rounding Promises. *Quantum* 7 (2023), 1132. <https://doi.org/10.22331/q-2023-10-10-1132>
- [58] Anders W. Sandvik. 1999. Stochastic series expansion method with operator-loop update. *Phys. Rev. B* 59 (1999), R14157–R14160. Issue 22. <https://doi.org/10.1103/PhysRevB.59.R14157>
- [59] Ulrich Schollwöck. 2011. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics* 326, 1 (2011), 96–192. <https://doi.org/10.1016/j.aop.2010.09.012>
- [60] E.M. Stoudenmire and Steven R. White. 2012. Studying Two-Dimensional Systems with the Density Matrix Renormalization Group. *Annu. Rev. Condens. Phys.* 3, 1 (2012), 111–128. <https://doi.org/10.1146/annurev-conmatphys-020911-125018>
- [61] Aaron Szasz, Johannes Motruk, Michael P. Zaletel, and Joel E. Moore. 2020. Chiral Spin Liquid Phase of the Triangular Lattice Hubbard Model: A Density Matrix Renormalization Group Study. *Phys. Rev. X* 10 (2020), 021042. Issue 2. <https://doi.org/10.1103/PhysRevX.10.021042>
- [62] Kristan Temme, Tobias J. Osborne, Karl G. Vollbrecht, David Poulin, and Frank Verstraete. 2011. Quantum Metropolis sampling. *Nature* 471, 7336 (2011), 87–90. <https://doi.org/10.1038/nature09770>
- [63] Laurens Vanderstraeten, Jutho Haegeman, Philippe Corboz, and Frank Verstraete. 2016. Gradient methods for variational optimization of projected entangled-pair states. *Phys. Rev. B* 94 (2016), 155123. Issue 15. <https://doi.org/10.1103/PhysRevB.94.155123>
- [64] F. Verstraete, V. Murg, and J.I. Cirac. 2008. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Adv. Phys.* 57, 2 (2008), 143–224. <https://doi.org/10.1080/14789940801912366>
- [65] Frank Verstraete, Michael M Wolf, and J Ignacio Cirac. 2009. Quantum computation and quantum-state engineering driven by dissipation. *Nature Physics* 5, 9 (2009), 633–636. <https://doi.org/10.1038/nphys1342>
- [66] Tom Vieijra, Corneel Casert, Jannes Nys, Wesley De Neve, Jutho Haegeman, Jan Ryckebusch, and Frank Verstraete. 2020. Restricted Boltzmann Machines for Quantum States with Non-Abelian or Anyonic Symmetries. *Phys. Rev. Lett.* 124 (2020), 097201. Issue 9. <https://doi.org/10.1103/PhysRevLett.124.097201>
- [67] Steven R. White. 1992. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* 69 (1992), 2863–2866. Issue 19. <https://doi.org/10.1103/PhysRevLett.69.2863>
- [68] Steven R. White. 1993. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B* 48 (1993), 10345–10356. Issue 14. <https://doi.org/10.1103/PhysRevB.48.10345>
- [69] Han-Qing Wu, Shou-Shu Gong, and D. N. Sheng. 2019. Randomness-induced spin-liquid-like phase in the spin- $\frac{1}{2}$ $J_1 - J_2$ triangular Heisenberg model. *Phys. Rev. B* 99 (2019), 085141. Issue 8. <https://doi.org/10.1103/PhysRevB.99.085141>
- [70] Michael P. Zaletel and Frank Pollmann. 2020. Isometric Tensor Network States in Two Dimensions. *Phys. Rev. Lett.* 124 (2020), 037201. Issue 3. <https://doi.org/10.1103/PhysRevLett.124.037201>

Received 10-NOV-2023; accepted 2024-02-11