nature physics

Article

https://doi.org/10.1038/s41567-025-02781-4

Local minima in quantum systems

Received: 14 January 2024

Accepted: 8 January 2025

Published online: 18 February 2025

Check for updates

Chi-Fang Chen [1,2] ✉, Hsin-Yuan Huang [1,3,4] ✉, John Preskill [1,2] ✉ & Leo Zhou [1,5] ✉

Finding ground states of quantum many-body systems is known to be hard for both classical and quantum computers. Consequently, when a quantum system is cooled in a low-temperature thermal bath, the ground state cannot always be found efficiently. Instead, the system may become trapped in a local minimum of the energy. In this work, we study the problem of finding local minima in quantum systems under thermal perturbations. Although local minima are much easier to find than ground states, we show that finding a local minimum is hard on classical computers, even when the task is merely to output a single-qubit observable at any local minimum. By contrast, we prove that a quantum computer can always find a local minimum efficiently using a thermal gradient descent algorithm that mimics natural cooling processes. To establish the classical hardness of finding local minima, we construct 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 that quantum computation is more powerful than classical computation, finding local minima is classically hard but quantumly easy.

Finding ground states and other low-energy states of quantum many-body systems stands as a cornerstone challenge in physics, materials science and chemistry. To tackle this problem, researchers have developed an array of powerful computational methods. These include density functional theory[1,2], quantum Monte Carlo[3–5], variational optimization with tensor network ansatzes[6–12] or neural network ansatzes[13–15], and data-driven machine learning approaches[16–19]. These methods, which can be implemented using classical computers, work well for many physically relevant problems but fail badly in other cases.

Quantum computers offer a promising avenue for efficiently finding ground states in cases that are intractable for classical methods. However, conclusively identifying such cases presents its own set of challenges. Even when known classical methods falter, ruling out all possible classical approaches remains difficult. Moreover, finding ground states of local Hamiltonians is known to be QMA-hard[20,21], suggesting that this task is intractable even for quantum computers in some quantum many-body systems. Indeed, all known quantum algorithms for the ground-state problem[22–29] are computationally inefficient for some quantum Hamiltonians.

If finding ground states is hard for quantum computers, the task must be hard for Nature as well, if Nature can be efficiently simulated by quantum computers. In our search for physics problems in which quantum computers might demonstrate an advantage over classical counterparts, we should look to Nature for guidance. With this perspective, we observe that when a quantum system is placed in a low-temperature thermal bath[30–34], it seeks a local minimum of energy, which may be far from the global minimum. This phenomenon is particularly evident in certain physical systems, such as spin glasses[35–38], where finding a ground state is notoriously hard. In these cases, the system, when cooled, almost always finds a local minimum instead of the ground state. Consequently, the ground state of the Hamiltonian becomes physically irrelevant, as it is unlikely to be observed in experiments.

This perspective encourages us to formulate and study a physically motivated alternative to the ground-state problem, which we refer to

[1]California Institute of Technology, Pasadena, CA, USA. [2]AWS Center for Quantum Computing, Pasadena, CA, USA. [3]Google Quantum AI, Venice, CA, USA. [4]Massachusetts Institute of Technology, Cambridge, MA, USA. [5]University of California Los Angeles, Los Angeles, CA, USA. ✉e-mail: achifchen@gmail.com; hsinyuan@caltech.edu; preskill@caltech.edu; leoxzhou@ucla.edu

https://doi.org/10.1038/s41567-025-02781-4

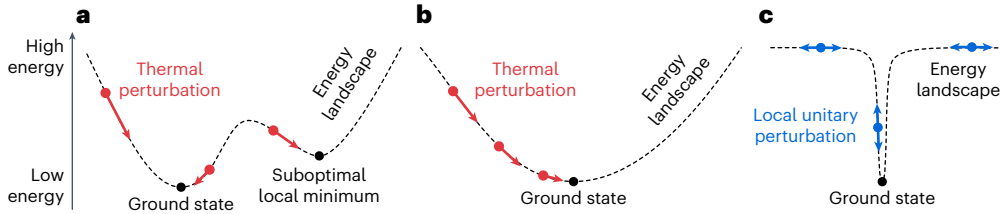

Fig. 1 | Illustration of energy landscapes. a, Energy landscape with several local minima. Some Hamiltonians, such as those for magnets in a small external magnetic field, have several local minima. One is the ground state whereas the other local minima are suboptimal. **b**, Energy landscape with one local minimum. For some local Hamiltonians, such as the family of BQP-hard Hamiltonians, the energy landscape over the entire n-qubit state space has a nice bowl shape, and the only local minimum is the global minimum. However, for QMA-hard Hamiltonians, the energy landscape necessarily contains many suboptimal local minima. **c**, Energy landscape under local unitary perturbations. For any local Hamiltonian H, there are always doubly exponentially many local minima within the n-qubit state space arising from a large barren plateau.

as the local-minima problem. We ask: How tractable is the problem of finding local minima of the energy in quantum systems using classical and quantum computers?

We address this question by proving that a machine capable of cooling any system to local minima is a universal quantum computer. In other words, by finding local minima, one can solve any problem that is efficiently solvable by a quantum computer, and furthermore, a quantum computer can always find local minima efficiently for any physical system. Hence, under the widely accepted assumption that quantum computing is more powerful than classical computing, finding local minima is both classically hard and quantumly easy. This result raises the hope that exploring local minima in quantum systems arising from physics, chemistry and materials science may offer new opportunities for solving classically intractable and physically relevant problems using quantum computers.

Results

We now present our main results concerning the tractability of finding local minima in quantum systems. A collection of notational conventions and some background on quantum thermodynamics can be found in Supplementary Information Section 1. Based on the standard definition in mathematical optimization[39–43], 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. For a brief review of mathematical optimization, see Supplementary Information Section 2.1.

Let $\mathcal{P}_{\boldsymbol{\alpha}}$ be a perturbation parameterized by a small vector $\boldsymbol{\alpha}$ 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 $\mathrm{tr}(H\rho)$ that is an approximate minimum under perturbations, that is,

$$\mathrm{tr}(H\rho) \leq \mathrm{tr}(H\mathcal{P}_{\boldsymbol{\alpha}}(\rho)) + \varepsilon \|\boldsymbol{\alpha}\|, \tag{1}$$

for all small enough $\boldsymbol{\alpha}$. The formal definition is given in Supplementary Information Section 2.2. The local minima of H form a subset of the entire quantum state space, which contains the global minima, namely the ground states of H. See Fig. 1 for energy landscapes with different numbers of local minima. We say an algorithm $\mathcal{A}$ has solved the problem of finding local minima under perturbation $\mathcal{P}$ if given any n-qubit local Hamiltonian H, written as a sum of few-qubit Hermitian operators with an efficient classical description, and any few-qubit observable O, the algorithm $\mathcal{A}$ can output a real value $\mathrm{tr}(O\rho)$ corresponding to any approximate local minimum ρ of H under perturbations $\mathcal{P}$ up to a small error. If there are several local minima, we consider finding any one instance to be sufficient.

Inspired by the physical process that Nature uses to cool quantum systems, we define local minima based on thermal perturbations rooted in open system thermodynamics[27,30–34]. Thermal perturbations are fundamentally different from local unitary perturbations, which are perturbations that arise from perturbing the states using few-qubit gates, as may arise from a quantum computing viewpoint. These two types of perturbations yield distinct energy landscapes, as illustrated in Fig. 1, leading to vastly different computational complexities for the task of finding a local minimum. Indeed, we will see that, in contrast to thermal perturbations, finding a local minimum under local unitary perturbations is classically easy.

Local minima under thermal perturbations

In this section, we consider local minima under thermal perturbations induced by a heat bath, as formally defined in Supplementary Information Section 2.2. Our starting point is a Markovian model of Nature's cooling process, a cornerstone in the study of open system thermodynamics (for example, ref. 32). When the coupling between an n-qubit system and a thermal bath is weak and the bath is memoryless, the complex joint system–bath Hamiltonian dynamics simplify to a Markovian Lindbladian evolution of the system alone, $\rho(t) = \mathrm{e}^{\mathcal{L}t}[\rho]$, where $\mathcal{L}$ is a continuous time generator. We focus on a modern variant[33] of the prototypical Davies generator[44], which has a quantitative dependence on the system–bath parameters and desirable analytic and algorithmic properties. Although there is flexibility in choosing the Markovian generator, many alternatives have been derived with limited physical rigour or lack a Lindbladian structure[45–48]. For our purposes, the continuous time generator $\mathcal{L}$ can be defined by merely the system Hamiltonian H, the jump operators A^a through which the bath interacts with the system and the thermodynamic properties of the bath: inverse temperature β and a characteristic timescale τ. Usually, A^a are local operators such as single-qubit or two-qubit Paulis. See Methods for an explicit description of $\mathcal{L}$ and Supplementary Information Section 1.3 for more details. Under these assumptions, we may effectively consider a thermal perturbation of n-qubit state ρ to be

$$(\text{thermal perturbation}): \qquad \rho \to \exp\left(\sum_{a=1}^{m} \alpha_a \mathcal{L}_a^{\beta,\tau,H}\right)(\rho), \tag{2}$$

where $\mathcal{L}_a^{\beta,\tau,H}$ is the thermal Lindbladian associated with each jump operator A^a acting on a few qubits, $m = \mathrm{poly}(n)$ is the number of jump operators, and $\boldsymbol{\alpha} = \sum_a \alpha_a \hat{\mathbf{e}}_a \in \mathbb{R}^m_{\geq 0}$ is a non-negative vector close to zero, where $\hat{e}_a$ is the unit vector along the a-th coordinate and α_a is any non-negative real number. Here, the vector is non-negative because thermodynamic processes are generally irreversible. The irreversibility in thermal perturbations is crucial for ensuring that there are fewer than doubly exponentially many local minima in the energy landscape (see the discussion in Supplementary Information Section 3.3).

We define a local minimum under thermal perturbations to be a state ρ with the minimum energy $\mathrm{tr}(H\rho)$ under the thermal perturbations given by equation (2). Given the definition, we next study the tractability of finding a local minimum under thermal perturbations. Our complexity-theoretic results show that finding local minima under thermal perturbations is both quantumly easy and classically hard, the

latter under the well-accepted assumption that not all quantum circuits can be efficiently simulated on classical computers (BPP ≠ BQP).

Finding local minima is easy for quantum computers. If our notion of local minima properly captures how a quantum system behaves in a cold environment, we would expect that finding local minima will be quantumly easy. Indeed, in the following theorem, we prove that a quantum computer can always efficiently find a local minimum of H under thermal perturbations starting from any initial state.

Theorem 1. (Quantumly easy-to-find local minima under thermal perturbations; informal). *The problem of finding an ε-approximate local minimum of an n-qubit local Hamiltonian H under thermal perturbations with jump operators $\{A^a\}$, inverse temperature β and timescale τ can be solved in poly$(n, 1/\varepsilon, \beta, \tau)$ quantum computational time*.

The formal statement is given in Theorem 4 of the Supplementary Information and is proven in Supplementary Information Section 4.2. To establish the theorem, we propose a quantum thermal gradient descent algorithm that cools down a quantum system by performing gradient descent based on quantum thermodynamics. This result builds on the recent algorithmic developments for quantum Markov chain Monte Carlo methods[25,26,28,29].

Our analysis establishes that the gate complexity of the quantum thermal gradient descent algorithm scales polynomially with all relevant parameters. For a Hamiltonian H with spectral norm $\|H\|_\infty \leq B$ bounded above by a positive number B, we demonstrate in Supplementary Information Section 4.2.2 that the algorithm requires at most $\mathcal{O}(B^3/\varepsilon^2)$ gradient steps to prepare one copy of an ε-approximate local minimum. Each step involves a $\mathcal{O}(1/B)$-time thermal Lindbladian simulation. Consequently, the overall algorithmic cost is primarily determined by the efficiency of simulating thermal Lindbladian dynamics. Recent advances in thermal Lindbladian simulations and quantum Gibbs sampling[27,49–51] have substantially improved these techniques. Notably, any future enhancements in this area will directly improve the gate complexity of the quantum thermal gradient descent algorithm.

Finding local minima is hard for classical computers. Although local minima are easier for quantum computers to find than ground states, how hard is it to find them? As our second main result, we address this question by considering a class of geometrically local Hamiltonians $\{H_C\}$ on a two-dimensional (2D) lattice such that the ground state of H_C encodes the outcome of a polynomial-sized quantum circuit C using a modified Kitaev's circuit-to-Hamiltonian construction[20,52,53]. To understand the computational hardness, we provide the following characterization of the energy landscape.

Theorem 2. (No suboptimal local minimum in BQP-hard Hamiltonians; informal). *For any quantum circuit C with size $|C|$, all approximate local minima of the geometrically local 2D Hamiltonian H_C under thermal perturbations with inverse temperature $\beta = \text{poly}(|C|)$ and timescale $\tau = \text{poly}(|C|)$ are close to the ground state*.

This theorem is the most technically involved contribution of this work. The formal statement is given in Theorem 5 of the Supplementary Information and is proven in Supplementary Information Section 6. Conceptually, the landscape of these 2D Hamiltonians has a nice bowl shape, as in convex optimization[40]. Therefore, performing thermal gradient descent (Theorem 1) allows us to prepare the ground state starting from an arbitrary initial state. To prove the theorem, we developed a set of techniques for establishing that a Hamiltonian H has no suboptimal local minima, that is, all local minima of H are global minima.

Although we have proved that H_C has no suboptimal local minima, general local Hamiltonians may well have many local minima. Finding the ground state of a local Hamiltonian is a QMA-hard problem; hence, we do not expect it to be solved efficiently by the quantum thermal gradient descent algorithm or by any other quantum algorithm. Fo a quantum circuit that verifies the witness for a problem in QMA, Kitaev's corresponding local Hamiltonian contains a term, often denoted H_{in}, that specifies some of the input qubits and leaves the input qubits corresponding to the witness unspecified, and a term, often denoted H_{out}, that checks whether the witness is accepted. Due to the unspecified witness qubits in H_{in}, the energy landscape contains a large number of local minima corresponding to all possible witnesses. Furthermore, most of these local minima correspond to rejected witnesses and are suboptimal because of the energy penalty from H_{out}. For these QMA-complete Hamiltonians, quantum thermal gradient descent will probably get stuck at a suboptimal local minimum. In our H_C, the term H_{in} specifies all input qubits and the term H_{out} is absent, which greatly simplifies the energy landscape and removes all suboptimal local minima.

Using this energy landscape characterization, we can show that finding a local minimum under thermal perturbations is classically intractable, assuming quantum computation is more powerful than classical computation. For a contradiction, suppose that a classical computer can efficiently find any local minima under thermal perturbations. Then by Theorem 2, the classical computer can efficiently find the ground state of H_C and, thus, simulate any efficient quantum circuit C, which is widely believed to be impossible. See Theorem 6 of the Supplementary Information for a formal statement and the proof.

Theorem 3. (Classically hard-to-find local minima under thermal perturbations; informal). *Assume the widely believed conjecture that* BPP ≠ BQP. *The problem of finding an approximate local minimum of an n-qubit local Hamiltonian H under thermal perturbations is universal for quantum computation and is, thus, classically hard*.

We emphasize that our classical hardness results are tailored to a carefully constructed family of Hamiltonians, which enables us to leverage the conjecture that BPP ≠ BQP to provide a rigorous statement of quantum advantage. For Hamiltonians arising in physics, our quantum thermal gradient descent algorithm always succeeds in finding local minima. However, the classical hardness of finding these local minima may vary on case by case. By using more advanced circuit-to-Hamiltonian constructions, refs. 54–57 have demonstrated that low-energy states of various Hamiltonian families, including one-dimensional systems, 2D Heisenberg models and 2D XY models, can encode low-energy states of all local Hamiltonians. It seems plausible that our proof techniques could be combined with these more sophisticated circuit-to-Hamiltonian constructions to establish quantum advantages in finding local minima for a broader range of Hamiltonian families. This presents intriguing opportunities for future research on the potential quantum advantage in finding local minima across diverse Hamiltonian families.

Local minima under local unitary perturbations

An alternative approach to defining local minima involves considering local unitary perturbations instead of thermal perturbations. These local unitary perturbations are mathematically simple and naturally align with a quantum computing perspective. However, our analysis reveals that finding local minima under such perturbations is classically tractable. Consequently, this approach falls short in delineating a problem that decisively separates quantum and classical computation.

Local unitary perturbations are short-time unitary evolutions under a sum of few-body Hermitian operators (for example, a quantum circuit consisting of near-identity two-qubit gates). A local unitary perturbation of an n-qubit pure state $|\psi\rangle$ is given by

$$\text{(local unitary perturbation):} \qquad |\psi\rangle \to \exp\left(-\mathrm{i}\sum_{a=1}^{m}\alpha_a h^a\right)|\psi\rangle, \quad (3)$$

where h^a is a Hermitian operator acting on a few qubits, $m = \text{poly}(n)$ is the number of such Hermitian operators and $\boldsymbol{\alpha} = \sum_a \alpha_a \hat{\mathbf{e}}_a \in \mathbb{R}^m$ is a

vector close to zero. This definition is the state version of the Riemannian geometry of quantum computation defined in ref. 58.

We prove that for any local Hamiltonian H, a random n-qubit pure state is almost always a local minimum of H under local unitary perturbations; see Lemma 1 in Methods. Hence, there are doubly exponentially many suboptimal local minima in the energy landscape. This is closely related to the barren plateau phenomenon for a sufficiently expressive quantum circuit ansatz[59]. Because the number of local minima is enormous, the properties of these local minima are concentrated on fixed values, making this problem classically easy.

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

See Proposition 3.2 of the Supplementary Information for the detailed statement and Supplementary Information Section 3.2 for its proof.

A common strategy considered in adaptive variational quantum eigensolvers[60–62] is to variationally minimize the energy by applying few-qubit gates to the current state. This algorithm naturally finds a local minimum under local unitary perturbations. Although Proposition 1 shows that finding one such local minimum is classically easy for any given H, there also exist local minima that are classically hard (for example, the ground state of H). By making a clever selection of the circuit ansatze and employing a careful initialization of the parameters, a variational quantum eigensolver still has the potential to find a classically hard local minimum, even though the landscape contains many classically easy local minima. This is like strategies employed in addressing the barren plateau issue[59].

Our results reveal several distinct advantages of thermal perturbations over local unitary perturbations: (1) Thermal perturbations yield a more favourable landscape than local unitary perturbations that could flow to the ground state from arbitrary initial states; see Theorem 2. (2) Unlike variational quantum eigensolvers, where overly powerful circuit ansatze with too many parameterized gates can lead to barren plateaus[59], introducing extra jump operators in thermal perturbations always improves the energy landscape. (3) Thermal perturbations more accurately reflect physical reality: systems in thermal contact naturally evolve through open system dynamics rather than unitary evolution.

Discussion

There are strong theoretical arguments that scalable fault-tolerant quantum computers will be more powerful than classical computers, but identifying problems of practical interest with provable superpolynomial quantum advantage remains an open challenge. Although quantum computers might substantially accelerate the characterization of ground states for local Hamiltonians in physics, chemistry and materials science, it remains unclear which specific problems will exhibit meaningful speed-ups[63]. Some problems admit efficient classical solutions, whereas others remain intractable even for quantum computers due to the inherent difficulty in finding ground states.

We examine a physically motivated alternative: finding the local rather than the global minima of a Hamiltonian. This problem naturally emerges from fundamental physical processes. When quantum systems are placed in a cold thermal bath, they efficiently evolve towards local energy minima. We demonstrate that our proposed quantum thermal gradient descent algorithm efficiently solves this problem. Moreover, we prove that finding any local minimum is classically hard in general (assuming BPP $\neq$ BQP), establishing this as a concrete example where quantum computers can achieve superpolynomial advantage.

By contrast, our approach of finding local minima emerges directly from the fundamental physical processes that Nature employs to produce low-energy states in physics, chemistry and materials science. This connection to physics provides inherent advantages. The local-minima problem inherits the robust features of thermodynamics: it depends only on macroscopic bath parameters β and τ and remains insensitive to the microscopic details. Furthermore, the framework exhibits a form of error tolerance: the addition of jump operators, even unwanted ones, can only improve the energy landscape by eliminating suboptimal local minima.

Previous research has explored various approaches for solving classically hard problems by finding suitable quantum states, such as designing a gapped adiabatic path for Hamiltonians to find ground states[53], engineering Lindbladians to have rapid dissipative evolution towards steady states[64] and performing quantum phase estimation on an initial state with a high ground-state overlap[23]. Although these approaches draw inspiration from physical principles to develop algorithms for both analogue and digital quantum devices, they do not directly correspond to naturally occurring physical processes. By contrast, the problem of finding local minima emerges directly from the fundamental physical processes that Nature employs to produce low-energy states in physics, chemistry and materials science. Moreover, the local-minima problem inherits the robustness of thermodynamics: one needs to specify only the macroscopic bath quantities β and τ without worrying about the microscopic details, and the choice of jump operators can be flexible, as adding more jumps (even unwanted ones) removes only suboptimal local minima.

Classical algorithms for minimizing the energies of quantum systems often rely on variational approaches with structured ansatzes, including tensor networks[6–12,65–74] and neural quantum states[13–15,75–83]. Although these algorithms successfully find local minima within their prescribed ansatz families, such minima may not represent a local minimum under thermal perturbations. In such cases, we can initialize the state found by the classical algorithm on a quantum computer and find a lower energy state by running the quantum thermal gradient descent algorithm. This observation leads to the following result:

Corollary 1. (Quantum enhancement of variational optimization; informal). *Assume that not all polynomial-sized quantum circuits can be efficiently simulated classically. Then there are 2D geometrically local Hamiltonians such that given any classical ansatz that allows efficient estimation of single-qubit observables and an output state $\boldsymbol{\rho}^{\#}$ of any efficient classical algorithm that optimizes the classical ansatz, running quantum thermal gradient descent starting at $\boldsymbol{\rho}^{\#}$* strictly lowers the energy.

A formal statement and detailed proof can be found in Supplementary Corollary 4.1 and Supplementary Information Section 4.2. In many cases, we can determine if the classically optimized state $\rho^{\#}$ represents a local minimum under thermal perturbation by performing an efficient classical computation of the energy gradient. A negative gradient provides a rigorous certificate that quantum optimization initialized at $\rho^{\#}$ will yield strictly better solutions than the classical approach.

Our framework extends beyond mimicking quantum thermodynamics. By engineering the thermal Lindbladians to have enhanced analytic properties or reduced computational costs[27], we can develop synthetic quantum processes that may potentially outperform Nature. These synthetic Lindbladians form part of a broader family of quantum Markov chain Monte Carlo algorithms[25–29,84]. Beyond Lindbladians, other families of perturbations, such as unitary perturbations accompanied by mid-circuit measurements or qubit resets, may generate favourable energy landscapes devoid of suboptimal local minima, leading to more powerful quantum optimization algorithms for finding low-energy states or for other applications.

This work raises several intriguing questions for further investigation. 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 (see, for example, case 1 in Supplementary Information Section 5.5).

Can we characterize when ground states can be found efficiently despite having many suboptimal local minima? Our conclusion that finding local minima under thermal perturbations is classically hard relied on the complexity-theoretic conjecture that BPP ≠ BQP. Can we establish unconditionally that finding local minima is hard for classical algorithms, perhaps within a black-box oracle model? Having demonstrated quantum advantages for finding local minima of quantum systems, might there also be quantum advantages in finding better local minima in classical optimization problems through some variant of quantum thermal gradient descent? See related recent works[85–87] studying the potential for quantum dynamics to yield advantages in classical optimization.

Although finding ground states is generally intractable, many physical systems efficiently cool to their ground states. This suggests that Hamiltonians relevant to physics, chemistry and materials science may possess energy landscapes without suboptimal local minima. We have proven in Theorem 2 that a specific family of BQP-hard Hamiltonians exhibits this property. A key challenge is to characterize broader classes of Hamiltonians with similarly favourable landscapes. Although our negative-gradient condition (Supplementary Lemma 4.3) provides a sufficient condition for excluding suboptimal local minima, verifying this condition for physical systems—whether involving spins, fermions or bosons—remains computationally challenging. Progress in this direction would substantially advance our understanding of the energy landscapes of quantum systems and illuminate promising paths towards quantum advantages in physically relevant problems.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41567-025-02781-4.

References

1. Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Phys. Rev.* **136**, B864 (1964).
2. Kohn, W. Nobel lecture: electronic structure of matter—wave functions and density functionals. *Rev. Mod. Phys.* **71**, 1253 (1999).
3. Ceperley, D. & Alder, B. Quantum Monte Carlo. *Science* **231**, 555–560 (1986).
4. Sandvik, A. W. Stochastic series expansion method with operator-loop update. *Phys. Rev. B* **59**, R14157–R14160 (1999).
5. Becca, F. & Sorella, S. *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge Univ. Press, 2017).
6. White, S. R. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.* **69**, 2863–2866 (1992).
7. White, S. R. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B* **48**, 10345–10356 (1993).
8. Perez-Garcia, D., Verstraete, F., Wolf, M. M. & Cirac, J. I. Matrix product state representations. *Quantum Inf. Comput.* **7**, 401–430 (2007).
9. Verstraete, F., Murg, V. & Cirac, J. I. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Adv. Phys.* **57**, 143–224 (2008).
10. Schollwock, U. The density-matrix renormalization group in the age of matrix product states. *Ann. Phys.* **326**, 96–192 (2011).
11. Haghshenas, R., O'Rourke, M. J. & Chan, G. K.-L. Conversion of projected entangled pair states into a canonical form. *Phys. Rev. B* **100**, 054404 (2019).
12. Hyatt, K. & Stoudenmire, E. M. DMRG approach to optimizing two-dimensional tensor networks. Preprint at https://arxiv.org/abs/1908.08833 (2019).
13. Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
14. Hibat-Allah, M., Ganahl, M., Hayward, L. E., Melko, R. G. & Carrasquilla, J. Recurrent neural network wave functions. *Phys. Rev. Res.* **2**, 023358 (2020).
15. Deng, D.-L., Li, X. & Sarma, S. D. Quantum entanglement in neural network states. *Phys. Rev. X* **7**, 021021 (2017).
16. Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O. & Dahl, G. E. Neural message passing for quantum chemistry. In *Proc. 34th International Conference on Machine Learning* (eds Precup, D. & Teh, Y. W.) 1263–1272 (JMLR, 2017).
17. Qiao, Z., Welborn, M., Anandkumar, A., Manby, F. R. & Miller, T. F. OrbNet: Deep learning for quantum chemistry using symmetry-adapted atomic-orbital features. *J. Chem. Phys.* **153**, 124111 (2020).
18. Huang, H.-Y., Kueng, R., Torlai, G., Albert, V. V. & Preskill, J. Provably efficient machine learning for quantum many-body problems. *Science* https://doi.org/10.1126/science.abk3333 (2022).
19. Lewis, L. et al. Improved machine learning algorithm for predicting ground state properties. *Nat. Commun.* **15**, 895 (2024).
20. Kitaev, A. Y., Shen, A. & Vyalyi, M. N. *Classical and Quantum Computation* (American Mathematical Society, 2002).
21. Kempe, J., Kitaev, A. & Regev, O. The complexity of the local Hamiltonian problem. *SIAM J. Comput.* **35**, 1070–1097 (2006).
22. Lin, L. & Tong, Y. Heisenberg-limited ground-state energy estimation for early fault-tolerant quantum computers. *PRX Quantum* **3**, 010318 (2022).
23. Gharibian, S. & Le Gall, F. 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 (STOC 2022)* 19–32 (Association for Computing Machinery, 2022).
24. Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. Quantum computation by adiabatic evolution. Preprint at https://arxiv.org/abs/quant-ph/0001106 (2000).
25. Temme, K., Osborne, T. J., Vollbrecht, K. G., Poulin, D. & Verstraete, F. Quantum Metropolis sampling. *Nature* **471**, 87–90 (2011).
26. Rall, P., Wang, C. & Wocjan, P. Thermal state preparation via rounding promises. *Quantum* **7**, 1132 (2023).
27. Chen, C.-F., Kastoryano, M. J., Brandão, F. G. S. L. & Gilyén, A. Quantum thermal state preparation. Preprint at https://arxiv.org/abs/2303.18224 (2023).
28. Ding, Z., Chen, C.-F. & Lin, L. Single-ancilla ground state preparation via Lindbladians. *Phys. Rev. Res.* **6**, 033147(2024).
29. Chen, C.-F., Kastoryano, M. J. & Gilyén, A. An efficient and exact noncommutative quantum Gibbs sampler. Preprint at https://arxiv.org/abs/2311.09207 (2023).
30. Lindblad, G. On the generators of quantum dynamical semigroups. *Commun. Math. Phys.* **48**, 119–130 (1976).
31. Davies, E. B. Generators of dynamical semigroups. *J. Funct. Anal.* **34**, 421–432 (1979).
32. Breuer, H.-P. & Petruccione, F. *The Theory of Open Quantum Systems* (Oxford Univ. Press, 2002).
33. Mozgunov, E. & Lidar, D. Completely positive master equation for arbitrary driving and small level spacing. *Quantum* **4**, 227 (2020).
34. Nathan, F. & Rudner, M. S. Universal Lindblad equation for open quantum systems. *Phys. Rev. B* **102**, 115109 (2020).
35. Edwards, S. F. & Anderson, P. W. Theory of spin glasses. *J. Phys. F* **5**, 965 (1975).
36. Kirkpatrick, S. & Sherrington, D. Infinite-ranged models of spin-glasses. *Phys. Rev. B* **17**, 4384 (1978).
37. Binder, K. & Young, A. P. Spin glasses: experimental facts, theoretical concepts, and open questions. *Rev. Mod. Phys.* **58**, 801 (1986).

38. Mydosh, J. A. *Spin Glasses: An Experimental Introduction* (CRC Press, 1993).
39. Pardalos, P. M. & Vavasis, S. A. Quadratic programming with one negative eigenvalue is np-hard. *J. Glob. Optimiz.* **1**, 15–22 (1991).
40. Boyd, S., Boyd, S. P. & Vandenberghe, L. *Convex Optimization* (Cambridge Univ. Press, 2004).
41. Agarwal, N., Allen-Zhu, Z., Bullins, B., Hazan, E. & Ma, T. Finding approximate local minima faster than gradient descent. In *Proc. 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)* 1195–1199 (Association for Computing Machinery, 2017).
42. Jin, C., Netrapalli, P. & Jordan, M.I. Accelerated gradient descent escapes saddle points faster than gradient descent. In *Proc. Machine Learning Research* 1042–1085 (MLR Press, 2018).
43. Ahmadi, A. A. & Zhang, J. On the complexity of finding a local minimizer of a quadratic function over a polytope. *Math. Program.* **195**, 783–792 (2022).
44. Davies, E. B. *Quantum Theory of Open Systems* (Academic Press, 1976).
45. Redfield, A. G. The theory of relaxation processes. *Adv. Magn. Opt. Reson.* **1**, 1–32 (1965).
46. Terhal, B. M. & DiVincenzo, D. P. Problem of equilibration and the computation of correlation functions on a quantum computer. *Phys. Rev. A* **61**, 022301 (2000).
47. Chen, C.-F. & Brandao, F. G. S. L. Fast thermalization from the eigenstate thermalization hypothesis. Preprint at https://arxiv.org/abs/2112.07646 (2023).
48. Shtanko, O. & Movassagh, R. Preparing thermal states on noiseless and noisy programmable quantum processors. Preprint at https://arxiv.org/abs/2112.14688 (2023).
49. Gilyén, A., Chen, C.-F., Doriguello, J. F. & Kastoryano, M. J. Quantum generalizations of Glauber and Metropolis dynamics. Preprint at https://arxiv.org/abs/2405.20322 (2024).
50. Ding, Z., Li, B. & Lin, L. Efficient quantum Gibbs samplers with Kubo–Martin–Schwinger detailed balance condition. Preprint at https://arxiv.org/abs/2404.05998 (2024).
51. Jiang, J. & Irani, S. Quantum Metropolis sampling via weak measurement. Preprint at https://arxiv.org/abs/2406.16023 (2024).
52. Oliveira, R. & Terhal, B. M. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Inf. Comput.* **8**, 900–924 (2008).
53. Aharonov, D. et al. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM J. Comput.* **37**, 166–194 (2007).
54. Piddock, S. & Montanaro, A. The complexity of antiferromagnetic interactions and 2D lattices. *Quantum Inf. Comput.* **17**, 636–672 (2017).
55. Cubitt, T. & Montanaro, A. Complexity classification of local Hamiltonian problems. *SIAM J. Comput.* **45**, 268–316 (2016).
56. Cubitt, T. S., Montanaro, A. & Piddock, S. Universal quantum Hamiltonians. *Proc. Natl Acad. Sci. USA* **115**, 9497–9502 (2018).
57. Zhou, L. & Aharonov, D. Strongly universal Hamiltonian simulators. Preprint at https://arxiv.org/abs/2102.02991 (2021).
58. Nielsen, M. A., Dowling, M. R., Gu, M. & Doherty, A. C. Quantum computation as geometry. *Science* **311**, 1133–1135 (2006).
59. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018).
60. O'Malley, P. J. J. et al. Scalable quantum simulation of molecular energies. *Phys. Rev. X* **6**, 031007 (2016).
61. Grimsley, H. R., Economou, S. E., Barnes, E. & Mayhall, N. J. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nat. Commun.* **10**, 3007 (2019).
62. Cerezo, M. et al. Variational quantum algorithms. *Nat. Rev. Phys.* **3**, 625–644 (2021).
63. Lee, S. et al. Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nat. Commun.* **14**, 1952 (2023).
64. Verstraete, F., Wolf, M. M. & Cirac, J. I. Quantum computation and quantum-state engineering driven by dissipation. *Nat. Phys.* **5**, 633–636 (2009).
65. Jordan, J., Orús, R., Vidal, G., Verstraete, F. & Cirac, J. I. Classical simulation of infinite-size quantum lattice systems in two spatial dimensions. *Phys. Rev. Lett.* **101**, 250602 (2008).
66. Vanderstraeten, L., Haegeman, J., Corboz, P. & Verstraete, F. Gradient methods for variational optimization of projected entangled-pair states. *Phys. Rev. B* **94**, 155123 (2016).
67. Corboz, P. Variational optimization with infinite projected entangled-pair states. *Phys. Rev. B* **94**, 035133 (2016).
68. Landau, Z., Vazirani, U. & Vidick, T. A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians. *Nat. Phys.* **11**, 566–569 (2015).
69. Arad, I., Landau, Z., Vazirani, U. & Vidick, T. Rigorous RG algorithms and area laws for low energy eigenstates in 1D. *Commun. Math. Phys.* **356**, 65–105 (2017).
70. Abrahamsen, N. A polynomial-time algorithm for ground states of spin trees. Preprint at https://arxiv.org/abs/1907.04862 (2019).
71. Stoudenmire, E. M. & White, S. R. Studying two-dimensional systems with the density matrix renormalization group. *Annu. Rev. Condens. Matter Phys.* **3**, 111–128 (2012).
72. Wu, H.-Q., Gong, S.-S. & Sheng, D. N. Randomness-induced spin-liquid-like phase in the spin-$\frac{1}{2}$ $J_1 - J_2$ triangular Heisenberg model. *Phys. Rev. B* **99**, 085141 (2019).
73. Szasz, A., Motruk, J., Zaletel, M. P. & Moore, J. E. Chiral spin liquid phase of the triangular lattice Hubbard model: a density matrix renormalization group study. *Phys. Rev. X* **10**, 021042 (2020).
74. Zaletel, M. P. & Pollmann, F. Isometric tensor network states in two dimensions. *Phys. Rev. Lett.* **124**, 037201 (2020).
75. Deng, D.-L., Li, X. & Sarma, S. D. Machine learning topological states. *Phys. Rev. B* **96**, 195145 (2017).
76. Nomura, Y., Darmawan, A. S., Yamaji, Y. & Imada, M. Restricted Boltzmann machine learning for solving strongly correlated quantum systems. *Phys. Rev. B* **96**, 205152 (2017).
77. Choo, K., Mezzacapo, A. & Carleo, G. Fermionic neural-network states for ab-initio electronic structure. *Nat. Commun.* **11**, 2368 (2020).
78. Ferrari, F., Becca, F. & Carrasquilla, J. Neural Gutzwiller-projected variational wave functions. *Phys. Rev. B* **100**, 125131 (2019).
79. Glasser, I., Pancotti, N., August, M., Rodriguez, I. D. & Cirac, J. I. Neural-network quantum states, string-bond states, and chiral topological states. *Phys. Rev. X* **8**, 011006 (2018).
80. Choo, K., Carleo, G., Regnault, N. & Neupert, T. Symmetries and many-body excitations with neural-network quantum states. *Phys. Rev. Lett.* **121**, 167204 (2018).
81. Vieijra, T. et al. Restricted Boltzmann machines for quantum states with non-abelian or anyonic symmetries. *Phys. Rev. Lett.* **124**, 097201 (2020).
82. Morawetz, S., De Vlugt, I. J. S., Carrasquilla, J. & Melko, R. G. U(1)-symmetric recurrent neural networks for quantum state reconstruction. *Phys. Rev. A* **104**, 012401 (2021).
83. Luo, D. et al. Gauge-invariant and anyonic-symmetric autoregressive neural network for quantum lattice models. *Phys. Rev. Res.* **5**, 013216 (2023).
84. Basso, J., Chen, C. F., & Dalzell, A. M. Optimizing random local Hamiltonians by dissipation. Preprint at https://arxiv.org/abs/2411.02578 (2024).
85. Leng, J., Hickman, E., Li, J. & Wu, X. Quantum Hamiltonian descent. Preprint at https://arxiv.org/abs/2303.01471 (2023).

86. Leng, J., Zheng, Y. & Wu, X. A quantum-classical performance separation in nonconvex optimization. Preprint at https://arxiv.org/abs/2311.00811 (2023).
87. Chen, Z., Lu, Y., Wang, H., Liu, Y. & Li, T. Quantum Langevin dynamics for optimization. Preprint at https://arxiv.org/abs/2311.15587 (2023).

Methods

Here, we describe in more detail the notion of local minima, the precise form of the thermal Lindbladian and our strategy for proving the main theorems.

Local minima under thermal perturbations

A central concept that enables us to understand the energy landscape and establish the computational complexity of finding local minima under thermal perturbations in Theorems 1 and 3 is the energy gradient operator,

$$\text{(energy gradient operator):} \qquad \sum_{a=1}^{m} \mathcal{L}_a^{\dagger\beta,\tau,H}(H)\hat{\mathbf{e}}_a, \tag{4}$$

where the adjoint $\mathcal{L}^\dagger$ is the Heisenberg-picture Lindbladian, that is, $\mathrm{tr}(\mathcal{L}^\dagger[O]\rho) = \mathrm{tr}(O\mathcal{L}[\rho])$. The energy gradient operator is a vector of individual gradient operators associated with each jump operator A^a. This is like the spin operator $\boldsymbol{\sigma} = \sigma^x\hat{\mathbf{x}} + \sigma^y\hat{\mathbf{y}} + \sigma^z\hat{\mathbf{z}}$, which is a vector of Hermitian observables. Indeed, the energy gradient operator naturally emerges by taking an infinitesimal perturbation, that is, the gradient of the energy $\mathrm{tr}(H\rho)$,

$$\mathrm{tr}\left(H\exp\left(\sum_{a=1}^{m}\alpha_a\mathcal{L}_a^{\beta,\tau,H}\right)(\rho)\right) = \mathrm{tr}(H\rho) + \boldsymbol{\alpha}\cdot\sum_{a=1}^{m}\mathrm{tr}\left(\mathcal{L}_a^{\dagger\beta,\tau,H}(H)\rho\right)\hat{\mathbf{e}}_a + \mathcal{O}\left(\|\boldsymbol{\alpha}\|^2\right). \tag{5}$$

In Supplementary Information Section 4.1, we describe the formal definition and some properties of the energy gradient. We provide a concrete example by showing the sets of local minima for ferromagnetic Ising chains under different longitudinal field strengths in Supplementary Information Section 5.5.

Explicit form of the thermal Lindbladian. Let us we review the explicit form of the thermal Lindbladian derived in ref. 33. Although other candidates exist in the open system literature (and which we believe are qualitatively similar), the following mathematical form gives a rigorous starting point for our analysis and algorithmic guarantees. For each jump operator A^a, the thermal Lindbladian $\mathcal{L}_a^{\beta,\tau,H}$ is given by:

$$\mathcal{L}_a^{\beta,\tau,H}(\rho) := -\mathrm{i}\left[H_{LS,a}^{\beta,\tau,H},\rho\right] + \mathcal{D}_a^{\beta,\tau,H}[\rho], \tag{6}$$

where $\mathcal{D}_a^{\beta,\tau,H}$ is the purely dissipative part:

$$\mathcal{D}_a^{\beta,\tau,H}[\rho] := \int_{-\infty}^{\infty}\gamma_\beta(\omega)\left[\hat{A}^a(\omega)\rho\hat{A}^a(\omega)^\dagger - \frac{1}{2}\{\hat{A}^a(\omega)^\dagger\hat{A}^a(\omega),\rho\}\right]\mathrm{d}\omega. \tag{7}$$

Key components of this formulation include:

- **Operator Fourier transform**: The operator $\hat{A}^a(\omega)$ is defined for the Heisenberg-evolved jump operator $A^a(t) = \mathrm{e}^{\mathrm{i}Ht}A^a\mathrm{e}^{-\mathrm{i}Ht}$, characterized by a timescale $\tau \in \mathbb{R}$ of the heat bath:

$$\hat{A}^a(\omega) := \frac{1}{\sqrt{2\pi\tau}}\int_{-\tau/2}^{\tau/2}A^a(t)\,\mathrm{e}^{-\mathrm{i}\omega t}\,\mathrm{d}t. \tag{8}$$

Roughly, the Fourier transform selects transitions induced by A^a that change the energy by ω up to uncertainty $1/\tau$.

- **Transition weight**: The function $\gamma_\beta(\omega)$ determines the strength of transitions with energy difference ω. It satisfies the Kubo–Martin–Schwinger condition and normalization:

$$\gamma_\beta(\omega)/\gamma_\beta(-\omega) = \mathrm{e}^{-\beta\omega} \text{ and } 0 \le \gamma_\beta(\omega) \le 1 \text{ for any } \beta \ge 0 \text{ and any } \omega \in \mathbb{R}. \tag{9}$$

Roughly, this function assigns a weight depending on whether the energy is increased ($\omega > 0$) or decreased ($\omega < 0$). In particular, we use the Glauber dynamics form with cutoff frequency Λ_0:

$$\gamma_\beta(\omega) = \frac{1}{2+\ln(1+\beta\Lambda_0)}\frac{\mathrm{e}^{-\omega^2/2\Lambda_0^2}}{1+\mathrm{e}^{\beta\omega}}. \tag{10}$$

- **Lamb-shift Hamiltonian**: The coherent term $H_{LS,a}^{\beta,\tau,H}$ depends on the bath correlation function $c_\beta(t)$:

$$H_{LS,a}^{\beta,\tau,H} := \frac{\mathrm{i}}{2\sqrt{2\pi\tau}}\int_{-\tau/2}^{\tau/2}\int_{-\tau/2}^{\tau/2}\mathrm{sgn}(t_1-t_2)c_\beta(t_2-t_1)A^a(t_2)A^a(t_1)\,\mathrm{d}t_2\,\mathrm{d}t_1, \tag{11}$$

where $c_\beta(t)$ is the Fourier transform of $\gamma_\beta(\omega)$:

$$c_\beta(t) := \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\gamma_\beta(\omega)\,\mathrm{e}^{+\mathrm{i}\omega t}\,\mathrm{d}t. \tag{12}$$

The Lamb-shift term commutes with the Hamiltonian in the large $\tau \to \infty$ limit, and for our purposes, its particular form is irrelevant and will be swept into the error bounds.

These precisely defined components encapsulate the fundamental physical principles underlying quantum thermodynamics. For a more comprehensive review of these concepts, we direct the reader to Supplementary Information Section 1.3.

Quantum thermal gradient descent. To establish that finding local minima under thermal perturbations is quantumly easy (Theorem 1), we propose an algorithm called quantum thermal gradient descent. For a low temperature and long timescale $\beta, \tau \to \infty$, the energy gradient $\mathcal{L}_a^{\dagger\infty,\infty,H}(H) \preccurlyeq 0$ is non-positive. In this case, the algorithm can just perform a random walk along random directions because no perturbations increase the energy and the use of gradient descent is not necessary. However, for arbitrary finite β and τ, the energy gradient can be positive. To find a local minimum that is an ε-minimum under all thermal perturbations for sufficiently small ε, the algorithm needs to carefully walk in directions with negative energy gradients.

To prove the convergence of quantum thermal gradient descent to a local minimum, we show that every small gradient step starting from a state that is not an approximate local minimum will decrease the energy. Because the energy of a quantum system is bounded from below, the algorithm must terminate at some point, which means it will find an approximate local minimum. To establish this claim, we derive analytic properties of thermal Lindbladians based on a smoothness bound on the second derivatives in ref. 27. To implement a gradient step based on thermal perturbations on a quantum computer, we build on a recently developed efficient quantum algorithm that simulates thermal Lindbladian evolution using a quantum circuit augmented by mid-circuit measurements[27].

All local minima are global. We now highlight the proof idea for Theorem 2 showing that all approximate local minima are close to ground states in a family of Hamiltonians with classically hard ground states. The family of geometrically local n-qubit Hamiltonians $\{H_C\}$ in a 2D lattice studied in this work is a modification of Kitaev's circuit-to-Hamiltonian construction[20,52] where the ground state encodes the computation of a quantum circuit $U_C = U_T\ldots U_1$. In particular, we design the ground state of H_C to be

$$\sum_{t=0}^{T}\sqrt{\xi_t}\,(U_t\cdots U_1\,|0^n\rangle)\otimes|1^t0^{T-t}\rangle, \qquad \text{where} \quad \xi_t := \frac{1}{2^T}\binom{T}{t}. \tag{13}$$

The binomial coefficient ξ_t is our modification of Kitaev's construction and is chosen to ensure that the desired properties hold for the spectrum and the energy gradients. The binomial distribution ensures that the Bohr-frequency gap is sufficiently large, which is central to the robustness of energy gradients under errors due to the finite temperature and small perturbations. We believe that the standard

circuit-to-Hamiltonian construction also has a large Bohr-frequency gap, but the proof seems more difficult. Estimating local properties of the ground state of H_C is equivalent to simulating the quantum circuit C, which is BQP-hard. Hence, under the standard assumption that not all quantum circuits can be simulated by classical computations, estimating the local properties of the ground state of H_C is classically hard.

Given the Hamiltonian H_C, showing that all of its approximate local minima under thermal perturbations are also approximate global minima seems daunting due to the complex expression for the thermal Lindbladian $\mathcal{L}_a^{\beta,\tau,H}$ and the doubly exponentially large space of possible quantum states. Previous studies on circuit-to-Hamiltonian mappings mainly focused on the lowest energy states. Here, we need to worry about potential local minima in all excited states in any superposition. To make progress, we propose a sufficient condition in Supplementary Information Section 4.1.3 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. Assume there exists a unit vector $\vec{\alpha} \in \mathbb{R}^m_{\geq 0}$ and $r > 0$ with

$$\text{(negative-gradient condition):} \qquad -\sum_{a=1}^{m} \hat{\alpha}_a \mathcal{L}_a^{\dagger\beta,\tau,H}(H) \succcurlyeq r(I - P_G(H)). \tag{14}$$

This negative-gradient condition implies that any state with a small ground-state overlap must experience a substantially negative energy gradient, that is, it cannot be a local minimum.

To prove that H_C satisfies the negative-gradient condition, we propose a series of lemmas and mathematical techniques for characterizing energy gradients in few-qubit systems, in commuting Hamiltonians and in subspaces of the Hamiltonian, which are stated in Supplementary Information Section 5 and proven in Supplementary Information Section 7. These new techniques build on the operator Fourier transform and the secular approximation given in ref. 27 for systematically handling energy uncertainty in thermal Lindbladians, which we review and adapt for our purpose in Supplementary Information Section 1.4. Using these techniques, we analyse the energy gradient of the entire system perturbatively by considering a sequence of Hamiltonians

$$H_1 \to H_2 \to H_3 = H_C,$$

with refining ground spaces $P_1 \supset P_2 \supset P_3$ where $\|H_1\| \gg \|H_2 - H_1\| \gg \|H_3 - H_2\|$.

Through these perturbations, we sequentially rule out local minima in excited states of the Hamiltonians H_1 and H_2 and, finally, $H_3 = H_C$. For example, we show that the first Hamiltonian H_1 satisfies the negative-gradient condition and that the gradient is stable under perturbation going from $H_1 \to H_2 \to H_3$. Controlling perturbations of the energy gradient is surprisingly challenging, and it is not a priori clear why this stability property should hold due to several (possibly competing) energy scales, including β^{-1} and τ^{-1}, which are the spectral gap and the Bohr-frequency gap. Recall that the spectral gap is the minimum non-zero difference between energy eigenvalues. The Bohr-frequency gap is the minimum non-zero difference between the difference of energy eigenvalues. The perturbative errors are not suppressed by the spectral gap of the Hamiltonian, as seen in standard settings, but instead by the Bohr-frequency gap, which can be much smaller (Theorem 6 of the Supplementary Information). These techniques allow us to establish the robustness of energy gradients when perturbing a degenerate Hamiltonian with a sufficiently large Bohr-frequency gap.

Local minima under local unitary perturbations

Here we present the central idea for proving Proposition 1, which establishes the classical easiness for finding local minima 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 local unitary perturbations. The formal statement is given in Lemma 3.1 of the Supplementary Information, and the proof is given in Supplementary Information Section 3.1.

Lemma 1. (Barren plateau; informal). *Given any n-qubit local Hamiltonian H, then a random pure n-qubit state* $|\psi\rangle$ *is an approximate local minimum of H under local unitary perturbations.*

The proof of the above lemma illustrates the following physical picture. The energy landscape in the pure state space defined in terms of local unitary perturbations consists of a large barren plateau[59] with doubly exponentially many approximate local minima having an exponentially small energy gradient. Additionally, almost all the local minima have local properties that are exponentially close to that of the maximally mixed state. As a result, although finding ground states is classically hard, finding local minima under local unitary perturbations is classically trivial.

Data availability

We did not analyse or generate any datasets because our work proceeds within a theoretical and mathematical approach.

Code availability

We do not have any computer code because our work proceeds within a theoretical and mathematical approach. All algorithms are analysed mathematically in the Supplementary Information.

References

88. Coles, P. J. et al. Thermodynamic AI and the fluctuation frontier. In *2023 IEEE International Conference on Rebooting Computing (ICRC)* 1–10 (ICRC, 2023).
89. Aifer, M. et al. Thermodynamic linear algebra. *npj Unconv. Comput.* **1**, 13 (2024).
90. Chen, C.-F., Huang, H.-Y., Preskill, J. & Zhou, L. Local minima in quantum systems. In *Proce. 56th Annual ACM Symposium on Theory of Computing (STOC 2024)* 1323–1330 (Association for Computing Machinery, 2024).

Acknowledgements

We thank A. Anshu, R. Babbush, S. Bravyi, F. Brandao, G. Chan, S. Chen, S. Choi, J. Cotler, D. Gosset, J. R. McClean and M. Soleimanifar for their valuable input. H.H. thanks P. Coles, G. Crooks and F. Sbahi for the inspiring discussions and for sharing their recent works on classical thermodynamics for AI applications[88,89]. Part of the manuscript was previously published as a STOC abstract[90]. C.C. is supported by the AWS Center for Quantum Computing internship. H.H. is supported by a Google PhD fellowship and a MediaTek Research Young Scholarship. H.H. acknowledges the visiting associate position at the Massachusetts Institute of Technology. L.Z. acknowledges funding from the Walter Burke Institute for Theoretical Physics at Caltech. J.P. acknowledges support from the United States Department of Energy Office of Science, Office of Advanced Scientific Computing Research (Grant Nos. DE-NA0003525 and DE-SC0020290), the United States Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator, and the National Science Foundation (Grant No. PHY-1733907). The Institute for Quantum Information and Matter is an NSF Physics Frontiers Center.

Author contributions

H.H., J.P. and L.Z. conceived the study of local minima of quantum systems. H.H. and L.Z. developed the computational complexity framework and the analysis of local minima under local unitary perturbations. C.C. introduced the quantum thermodynamics aspect. C.C. and H.H. established the theoretical framework for quantum thermal gradient descent. C.C. and L.Z. proved the BQP-hardness

of finding local minima under thermal perturbation. All authors contributed to the writing of the manuscript.

Competing interests
The authors declare no competing interests.

Additional information
Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41567-025-02781-4.

Correspondence and requests for materials should be addressed to Chi-Fang Chen, Hsin-Yuan Huang, John Preskill or Leo Zhou.

Peer review information *Nature Physics* thanks Ángela Capel, Tongyang Li and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.