Finding local minima and certifying quantum states

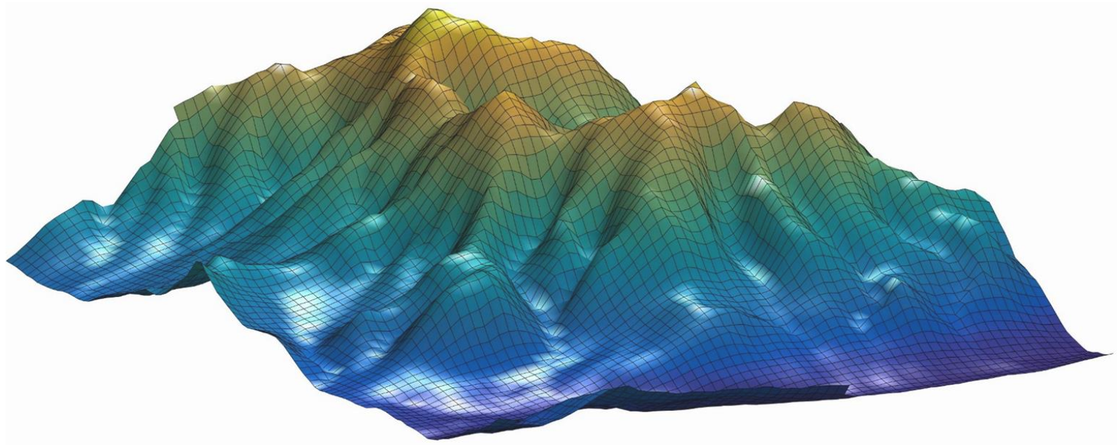

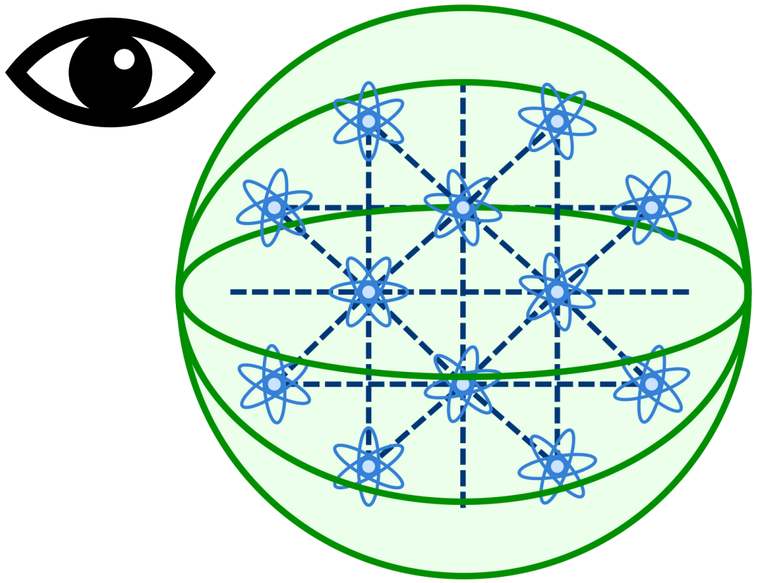

arxiv:2309.16596 with Anthony Chen, Robert Huang, and Leo Zhou
arXiv:2404.07281 with Robert Huang and Mehdi Soleimanifar

John Preskill, Caltech
NQI Joint Algorithms Workshop
20 May 2024

@preskill

Quantum computing applications

Dirac (1929): “The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble.”

Feynman (1981): “You can simulate this with a quantum system, with quantum computer elements. It’s not a Turing machine, but a machine of a different kind.”

Artificial intelligence may drive future progress in (strongly correlated) chemistry and materials science. Eventually, quantum computers can accelerate progress by providing abundant training data.

Ground states in chemistry and materials

Dirac: "… equations much too complicated to be soluble."

Yet, heuristic classical algorithms have been very successful.

We are targeting the relatively small "strongly correlated" corner of chemistry and materials science, where such methods falter.

Can quantum computers efficiently solve for ground states in cases where classical methods fail?

Quantum computers cannot find ground states for QMA-hard cases, but that's okay. Nature does not find these states either.

How useful are quantum computers in physically relevant situations that are beyond the reach of classical methods?

Ground states in chemistry and materials

We are seeking problems that are (1) quantumly easy, (2) classically hard, (3) physically relevant.

A patchwork of heuristic classical methods including: HF, DFT, CC, QMC, DMRG, TN, NN, … These lack performance guarantees, but often work. Cost need not scale exponentially with problem size.

Quantum algorithms are heuristic, too. We need an initial state that has sufficient overlap with the ground state.

Strong correlations can result in competing phases, first-order quantum phase transitions, … Adiabatic state preparation may fail.

Garnet Chan et al., Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry (2023)

Ground states in chemistry and materials

We are seeking problems that are (1) quantumly easy, (2) classically hard, (3) physically relevant.

Perhaps exponential quantum advantage should not be expected.

But a significant polynomial advantage is a reasonable expectation and could be quite impactful.

These applications require deep quantum circuits. Fault-tolerant quantum computation will be needed, at a high cost in physical qubits and gates.

Garnet Chan et al., Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry (2023)

Near-term applications

Variational methods like VQE are hampered by barren plateaus.

Simulations of quantum dynamics using near-quantum analog and digital quantum platforms can be scientifically informative. Surprises are likely.

Here, early fault-tolerant methods and error mitigation methods can be helpful.

Finding a *local* minimum

We are seeking problems that are (1) quantumly easy, (2) classically hard, (3) physically relevant.

Sometimes nature is unable to find the global minimum of the energy. That is true even classically (spin glasses, combinatorial optimization).

Instead, nature easily finds a local minimum when we cool the system.

Result: cooling to a local minimum is universal for quantum computation and there is a quantum algorithm that cools efficiently.

Hence, finding a local minimum is quantumly easy and classically hard (assuming BPP ≠ BQP).

Chen, Huang, Preskill, Zhou, arXiv:2309.16596

What is a local minimum?

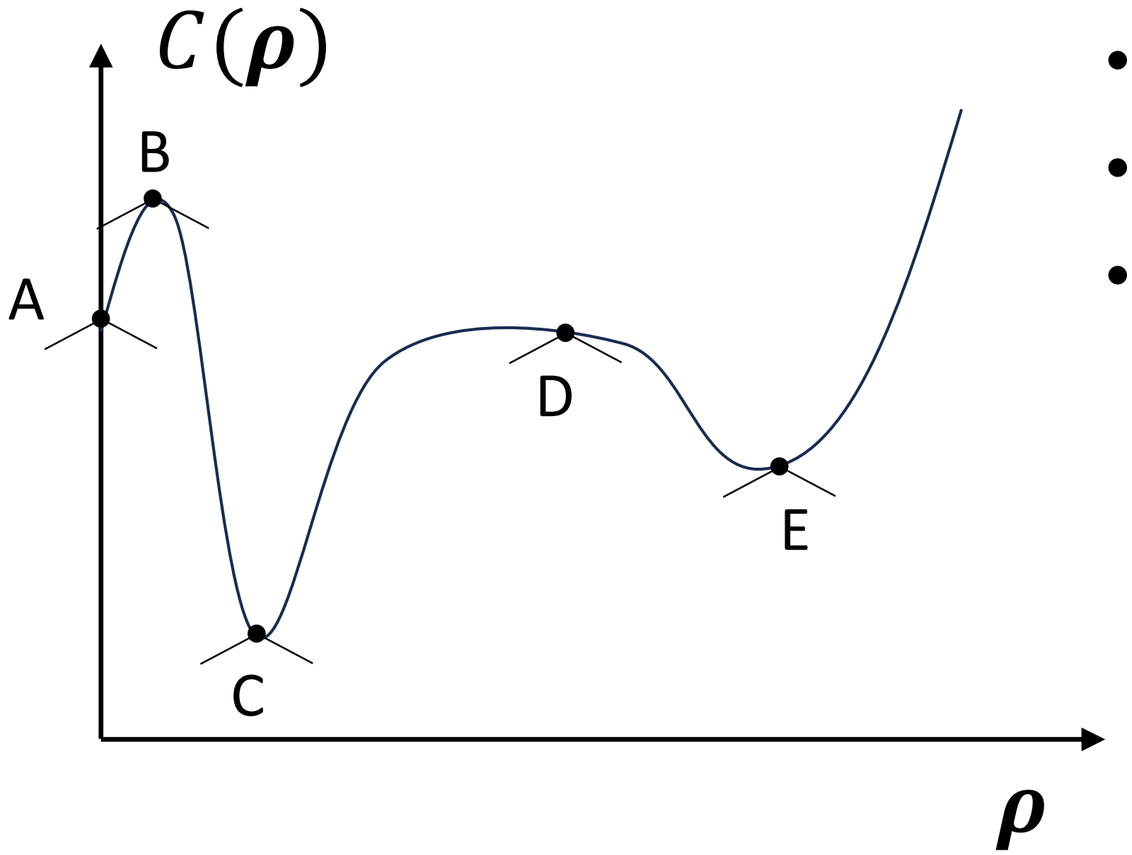

- Domain: n-qubit state $\boldsymbol{\rho}$
- Energy function: $C(\boldsymbol{\rho}) = \mathrm{tr}(\boldsymbol{H}\boldsymbol{\rho})$
- A family of perturbations: $\boldsymbol{\rho} \to \mathcal{P}_{\boldsymbol{\theta}}[\boldsymbol{\rho}]$
- $\boldsymbol{\rho}$ is an ***ϵ-approximate local minimum*** if

$$C(\boldsymbol{\rho}) \leq C(\mathcal{P}_{\boldsymbol{\theta}}[\boldsymbol{\rho}]) + \epsilon\|\boldsymbol{\theta}\|$$

for all small enough $\boldsymbol{\theta}$

A, C, D, E are ϵ-approx LM.
B is not.

Finding a local minimum: the problem

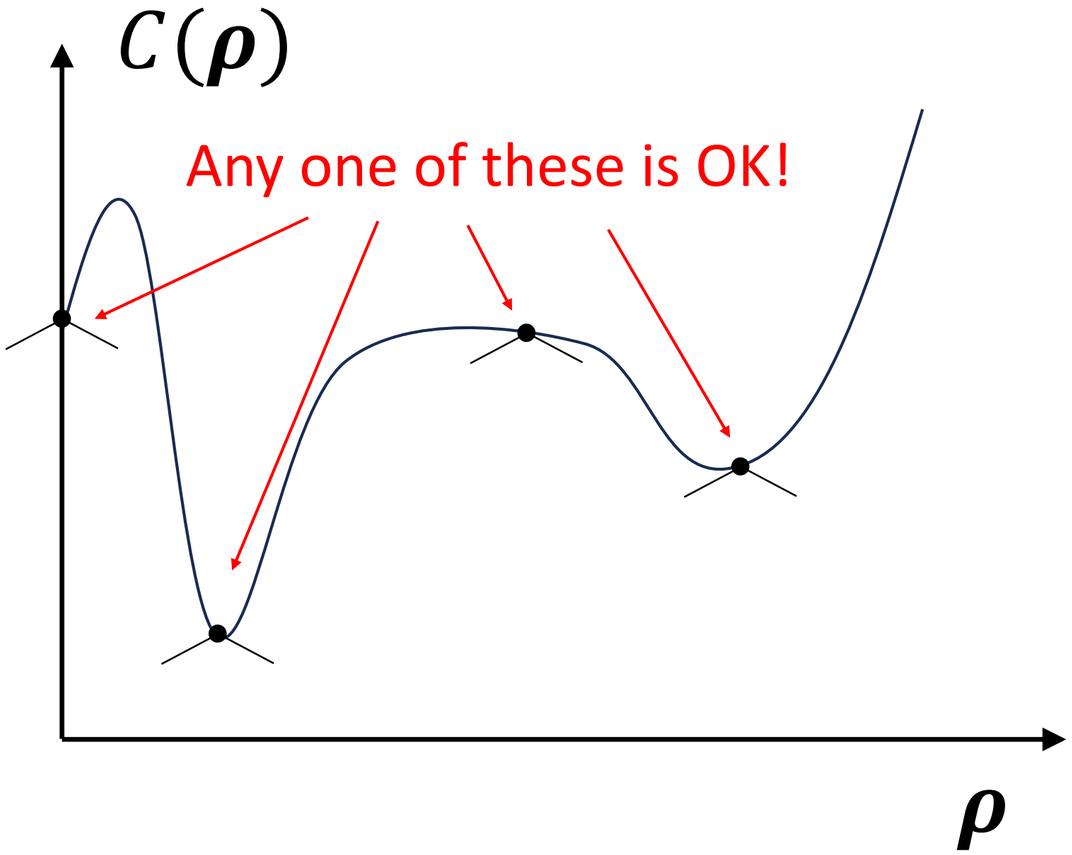

- **Input**:
  1. $\boldsymbol{H}$, where $\|\boldsymbol{H}\| = \mathrm{poly}(n)$
  2. a family of perturbation $\{\mathcal{P}_{\boldsymbol{\theta}}\}_{\boldsymbol{\theta}}$
  3. some $\epsilon > 1/\mathrm{poly}(n)$
  4. a (local) observable $\boldsymbol{O}$
- **Problem**: Output estimated $\mathrm{Tr}(\boldsymbol{O}\boldsymbol{\rho}^*)$ within ϵ error for *any* ϵ-approx local minimum $\boldsymbol{\rho}^*$ under the perturbations.

Note: purely classical input + output.

Complexity of finding a local minimum

Local unitary perturbation

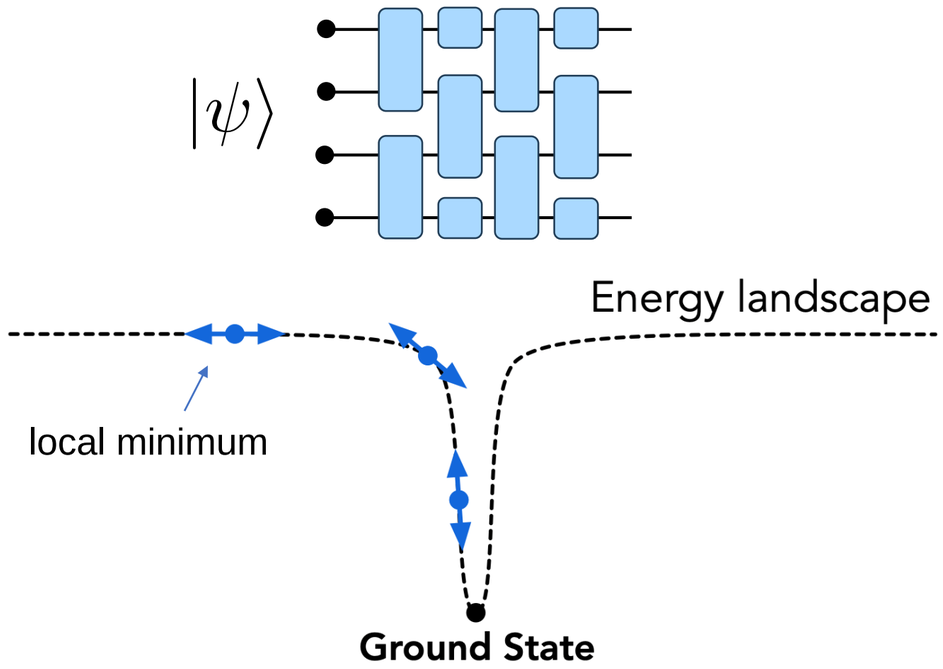

- There are $\exp(\exp(n))$ many local minima for all quantum systems!
- Finding a local minimum is always classically **EASY**.

Thermal perturbation

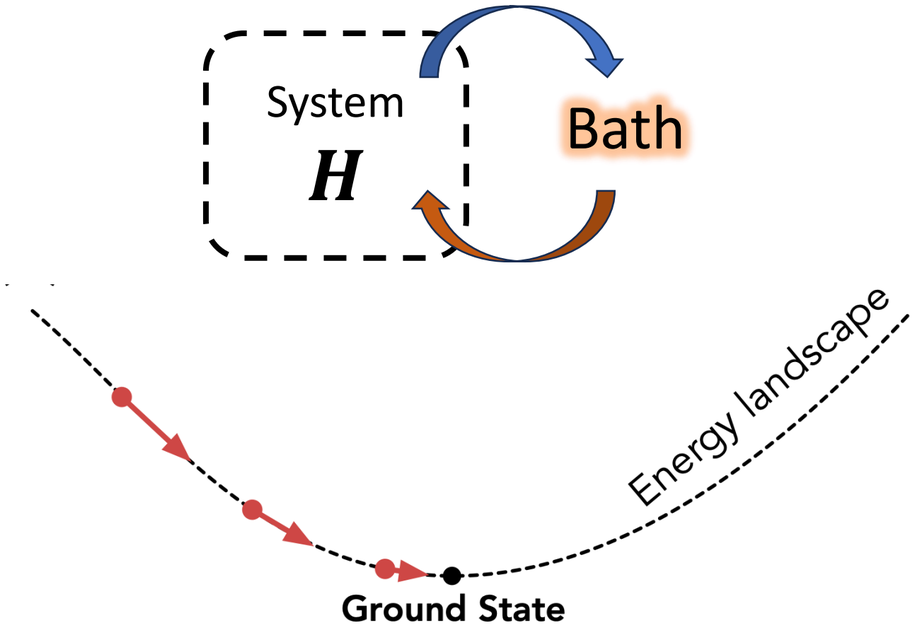

For some 2D systems, finding a local minimum is classically **HARD** and quantumly **EASY**!

Chen, Huang, Preskill, Zhou, arXiv:2309.16596

Thermal perturbations: inspired by nature

$$\boldsymbol{\rho} \to \mathcal{P}_{\boldsymbol{\theta}}[\boldsymbol{\rho}] = e^{\sum_a \theta_a \mathcal{L}_a}[\boldsymbol{\rho}] \qquad \theta_a \in \mathbb{R}_{\geq 0}$$

where $\mathcal{L}_a$ is a thermal Lindbladian for the system weakly coupled to a bath

Parameters

β	inverse temperature
τ	coarse-grain timescale
$\{\boldsymbol{A}^a\}_a$	local jump operators

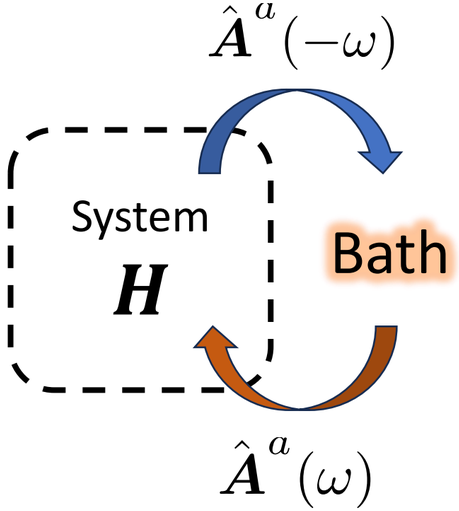

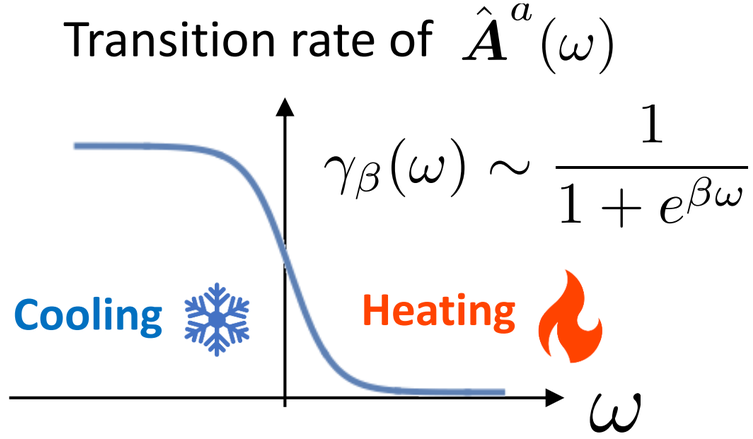

Based on rigorous version of Davies equation

Mozgunov, Lidar 2020

More on thermal Lindbladian

- Rigorous version of “Davies generator”

$$\mathcal{L}_a[\boldsymbol{\rho}] = \int_{-\infty}^{\infty} \gamma_\beta(\omega)\left[\hat{\boldsymbol{A}}^a(\omega)\boldsymbol{\rho}\hat{\boldsymbol{A}}^a(\omega)^\dagger - \frac{1}{2}\{\hat{\boldsymbol{A}}^a(\omega)^\dagger\hat{\boldsymbol{A}}^a(\omega), \boldsymbol{\rho}\}\right]d\omega$$

“Filtered” local jump operator

$$\hat{\boldsymbol{A}}^a(\omega) = \frac{1}{\sqrt{2\pi\tau}}\int_{-\frac{\tau}{2}}^{\frac{\tau}{2}} e^{i\omega t}e^{i\boldsymbol{H}t}\boldsymbol{A}^a e^{-i\boldsymbol{H}t}dt \approx \sum_{\substack{E_1-E_2\\ \approx\omega\pm 1/\tau}} \boldsymbol{P}_{E_1}\boldsymbol{A}^a\boldsymbol{P}_{E_2}$$

$$\boldsymbol{H} = \sum_E E\boldsymbol{P}_E$$

Transition weight

$$\gamma_\beta(\omega) \sim \frac{1}{1+e^{\beta\omega}}$$

Mozgunov, Lidar 2020

Chen, Kastoryano, Brandão, Gilyén 2023

Finding a local mimimum under *thermal perturbations is quantumly easy*

Quantum thermal gradient descent algorithm:

- Consider any n-qubit Hamiltonian $\boldsymbol{H}$ where $\|\boldsymbol{H}\| \leq B$
- To find an ϵ-local minimum under thermal perturbations by $\{\mathcal{L}_a\}_{a=1}^m$:

Initialize at any state $\boldsymbol{\rho}^{(0)}$, and for each step $t = 1, 2, 3, \ldots$

1. Estimate $g_a = \partial_a \langle \boldsymbol{H} \rangle = \mathrm{tr}\left(\boldsymbol{\rho} \mathcal{L}_a^\dagger[\boldsymbol{H}]\right)$ to 0.01ϵ precision

2. If all $g_a > -0.99\epsilon$, STOP. Otherwise evolve $\boldsymbol{\rho}^{(t)} = \exp(\sum_a \theta_a \mathcal{L}_a)\,[\boldsymbol{\rho}^{(t-1)}]$ where $\theta_a = -\min(0, g_a)/9B^2$

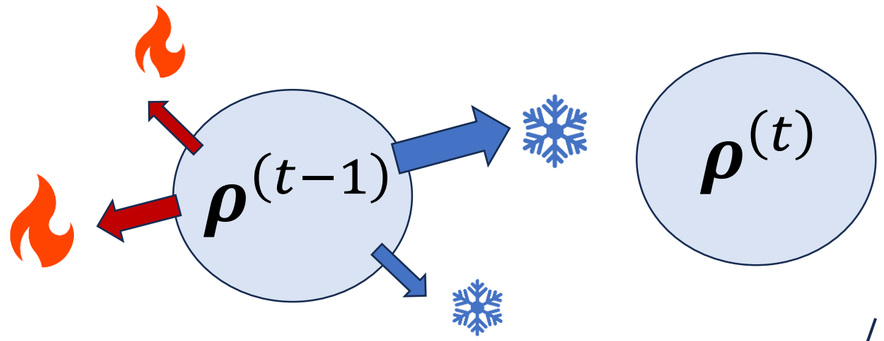

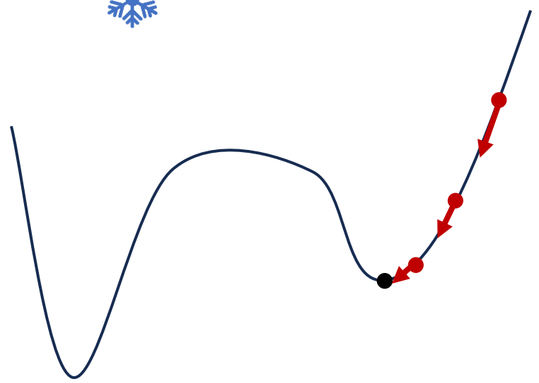

This **provably converges** within $O(B^3/\epsilon^2)$ steps!

Finding a local minimum is **quantumly easy**!

Finding a local minimum under thermal perturbations is *classically hard*

For any circuit $\boldsymbol{U}_C = \boldsymbol{U}_T \cdots \boldsymbol{U}_1$

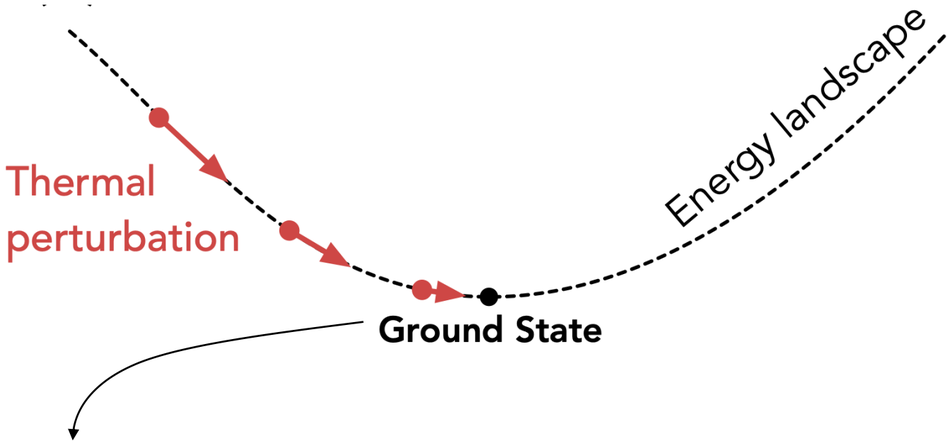

$$|\mathrm{GS}\rangle = \sum_{t=0}^{T} \sqrt{\xi_t} (\boldsymbol{U}_t \cdots \boldsymbol{U}_2 \boldsymbol{U}_1 |0^n\rangle) \otimes |t\rangle$$

$$\langle \mathrm{GS} | \boldsymbol{Z}_j | \mathrm{GS} \rangle \approx \langle 0^n | \boldsymbol{U}_C^\dagger \, \boldsymbol{Z}_j \, \boldsymbol{U}_C | 0^n \rangle$$

- **<u>Theorem</u>**: Certain 2D Hamiltonians whose ground states encode universal quantum computation **have no suboptimal local minima** when $\epsilon^{-1}, \beta, \tau \geq \mathrm{poly}(n)$.

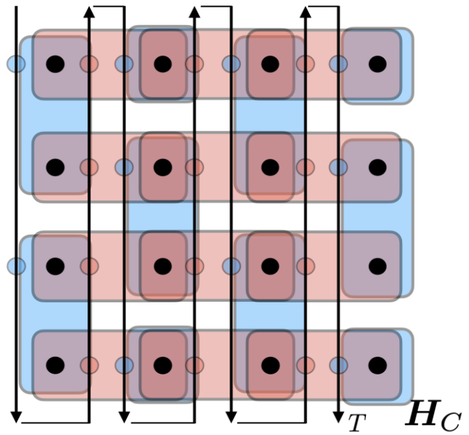

$$\boldsymbol{H} = \sum_t U_t \otimes |011\rangle\langle 001| + \cdots$$

If quantum circuits cannot be efficiently simulated classically, then finding a local minimum of $\boldsymbol{H}$ is *classically hard*.

Quantum advantage in cooling to local minima

For a 2D quantum system, finding a local minimum under thermal perturbations is classically hard and quantumly easy!

For typical physical systems, classical algorithms are nonetheless used routinely (e.g. DFT, DMRG, etc.), often with great success.

How to identify more systems where *quantum beats classical*?

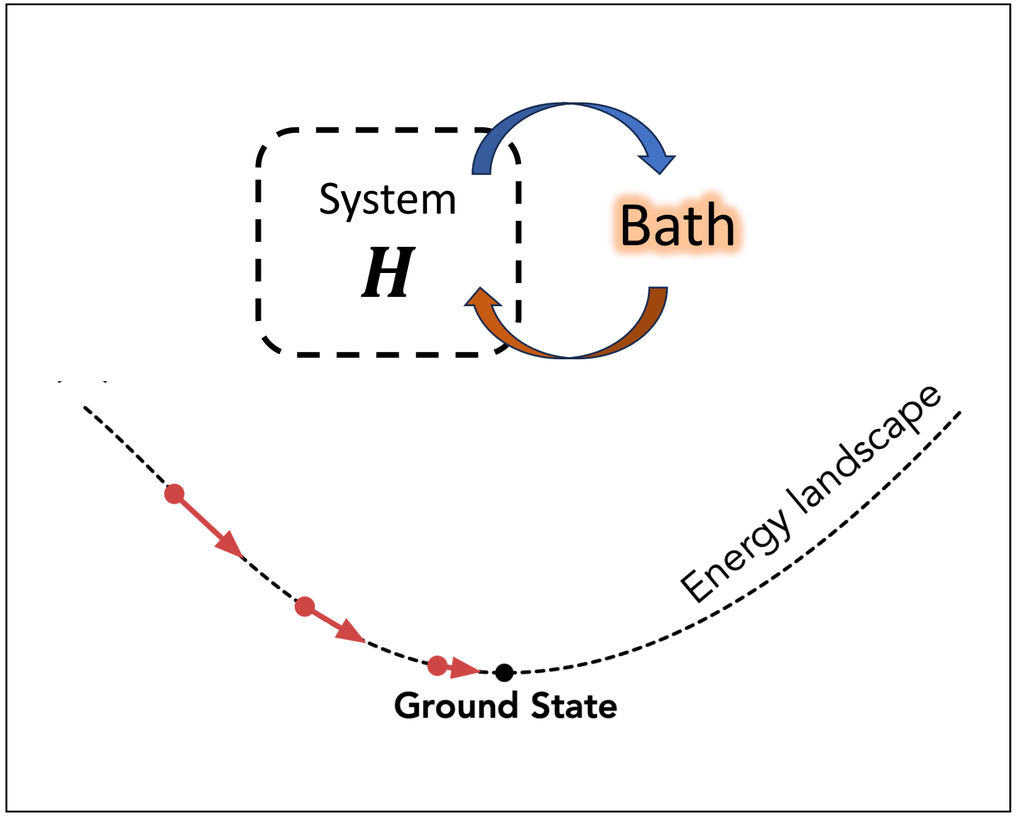

Quantum advantage in cooling to local minima

A possible method to detect quantum advantage

$g_a < -\epsilon \rightarrow$
quantum advantage!

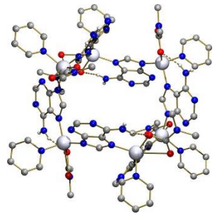

Energy landscape of
classical ansatz

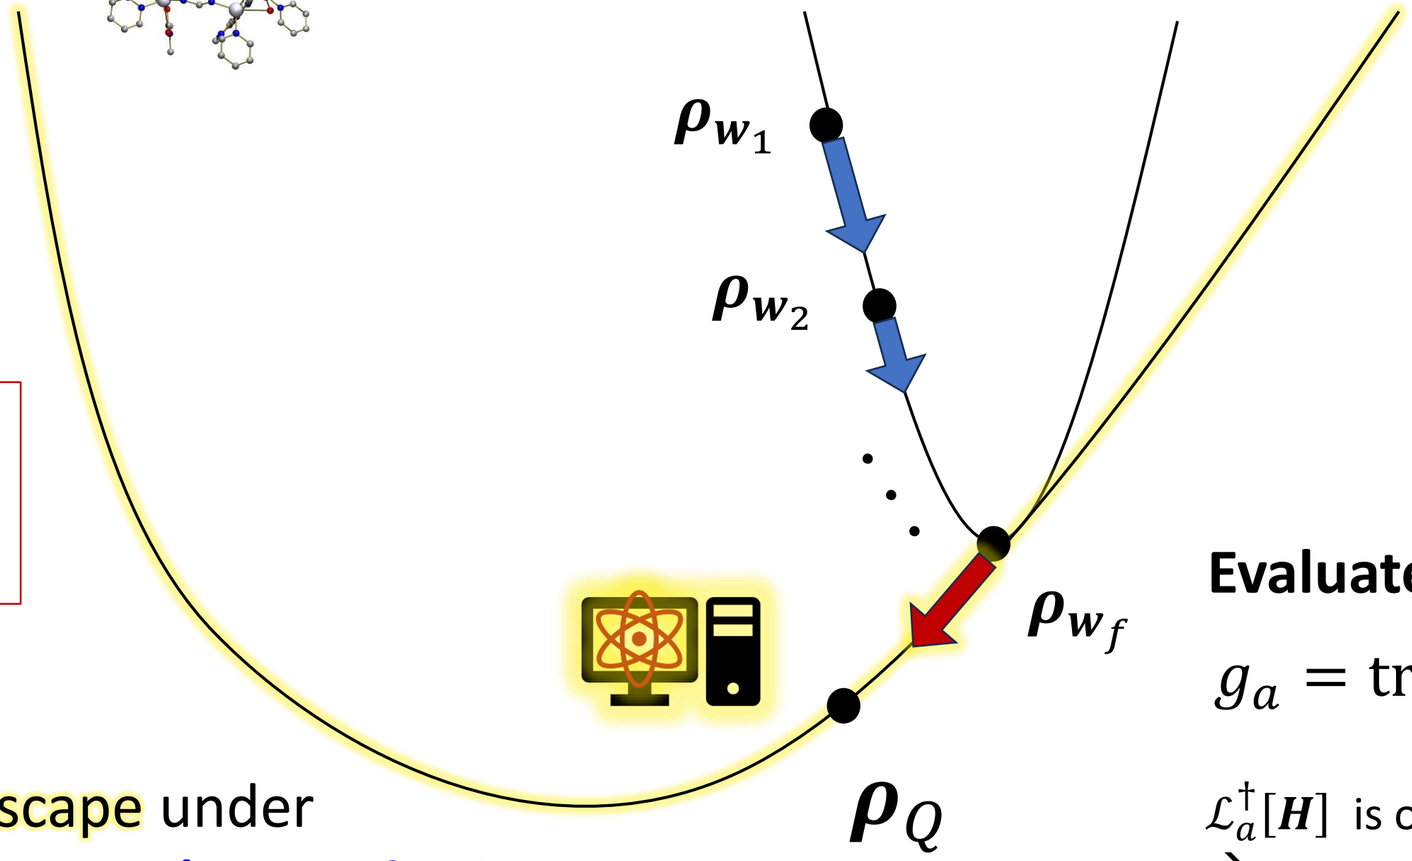

Energy landscape under
thermal (quantum) perturbations

Evaluate gradient

$$g_a = \mathrm{tr}\left(\mathcal{L}_a^\dagger[\boldsymbol{H}]\boldsymbol{\rho}_{w_f}\right)$$

$\mathcal{L}_a^\dagger[\boldsymbol{H}]$ is often quasi-local
→ can evaluate classically

Certifying a quantum state

We have a classical description of an *n*-qubit target pure state $|\psi\rangle$.

We have access to multiple copies of an *n*-qubit state ρ in the lab.

We want to check whether the fidelity $F = \langle\psi|\rho|\psi\rangle$ is close to 1.

For previously proposed certification methods, one of these is true:

-- deep quantum circuits required.
-- exponentially many measurements needed.
-- only works for a class of states with low entanglement.
-- no rigorous guarantee of accuracy.

Certifying a quantum state

Classical shadows from random Pauli measurement. Easy measurements, but number of copies needed to predict fidelity accurately is exponential in n.

Classical shadows from random Clifford measurement. Circuits with depth comparable to n needed. Not currently practical for $n \approx 100$ qubits.

Cross-entropy benchmark (XEB). Easy measurements, but can access only the terms in ρ that are diagonal in the Z basis. States far from the target can obtain a high score.

Result: Efficient certification using single-qubit measurements with performance guarantees.

Huang, Preskill, Soleimanifar, arXiv:2404.07281

What we do in the lab: Single-qubit measurements

Select one of the *n* qubits uniformly at random.

Measure the selected qubit in the *X*, *Y*, or *Z* basis, chosen uniformly at random.

Measure the remaining *n*-1 qubits in the *Z* basis.

Repeat many times.

That's all!

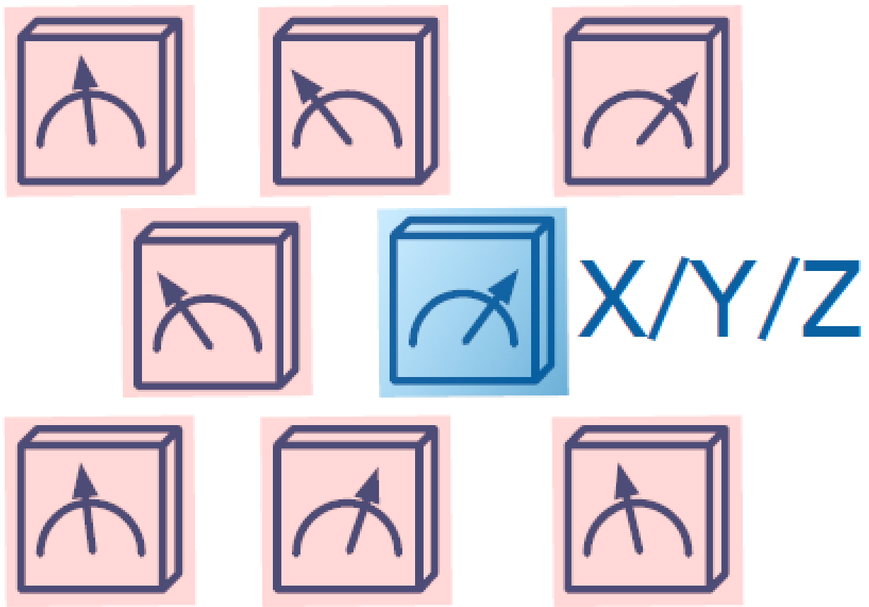

Huang, Preskill, Soleimanifar, arXiv:2404.07281

What we do with the data: The *shadow overlap*

Two *n*-qubits strings b_0, b_1 are compatible with the (*n*-1) *Z*-basis measurement outcomes.

If the *n*-1 qubits were measured in the ideal target state $|\psi\rangle$, the (unnormalized) post-measurement state would be

$$|\psi_{b_0, b_1}\rangle = |0\rangle\langle b_0|\psi\rangle + |1\rangle\langle b_1|\psi\rangle.$$

Using the randomized X/Y/Z measurement outcome and the classical shadow formalism, estimate the fidelity ω of that ideal one-qubit state with the postselected one-qubit lab state.

Estimated fidelity of the *n*-qubit lab state with the target state is: $F = \langle\psi \,|\, \rho \,|\, \psi\rangle \approx \mathbb{E}[\omega]$.

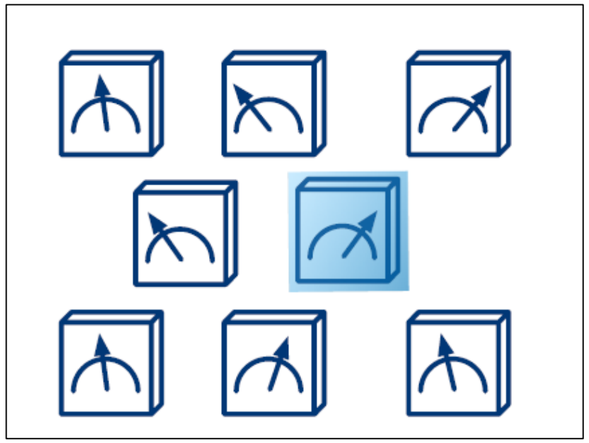

That's all!

How the shadow overlap tracks the fidelity

$$F = \langle \psi | \rho | \psi \rangle \geq 1 - \epsilon \text{ implies } \mathbb{E}[\omega] \geq 1 - \epsilon.$$

$$\mathbb{E}[\omega] \geq 1 - \epsilon \text{ implies } F = \langle \psi | \rho | \psi \rangle \geq 1 - \tau\epsilon.$$

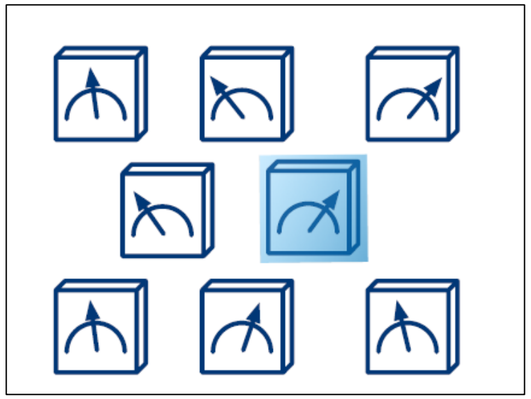

Here τ is the relaxation time of a (lazy) random walk on *n*-bit strings that converges to the stationary distribution $\pi(b) = |\langle b|\psi\rangle|^2$.

(This walk is not used in the protocol itself, only in the analysis of the protocol.)

Thus we can certify that ρ is ϵ-close to $|\psi\rangle$ by measuring $O(\tau^2/\epsilon^2)$ samples (actually $O(\tau/\epsilon)$ for optimal choice of measurement basis).

Furthermore, $\tau = O(n^2)$ for Haar-random target states.

Shadow overlap vs. XEB as proxies for fidelity

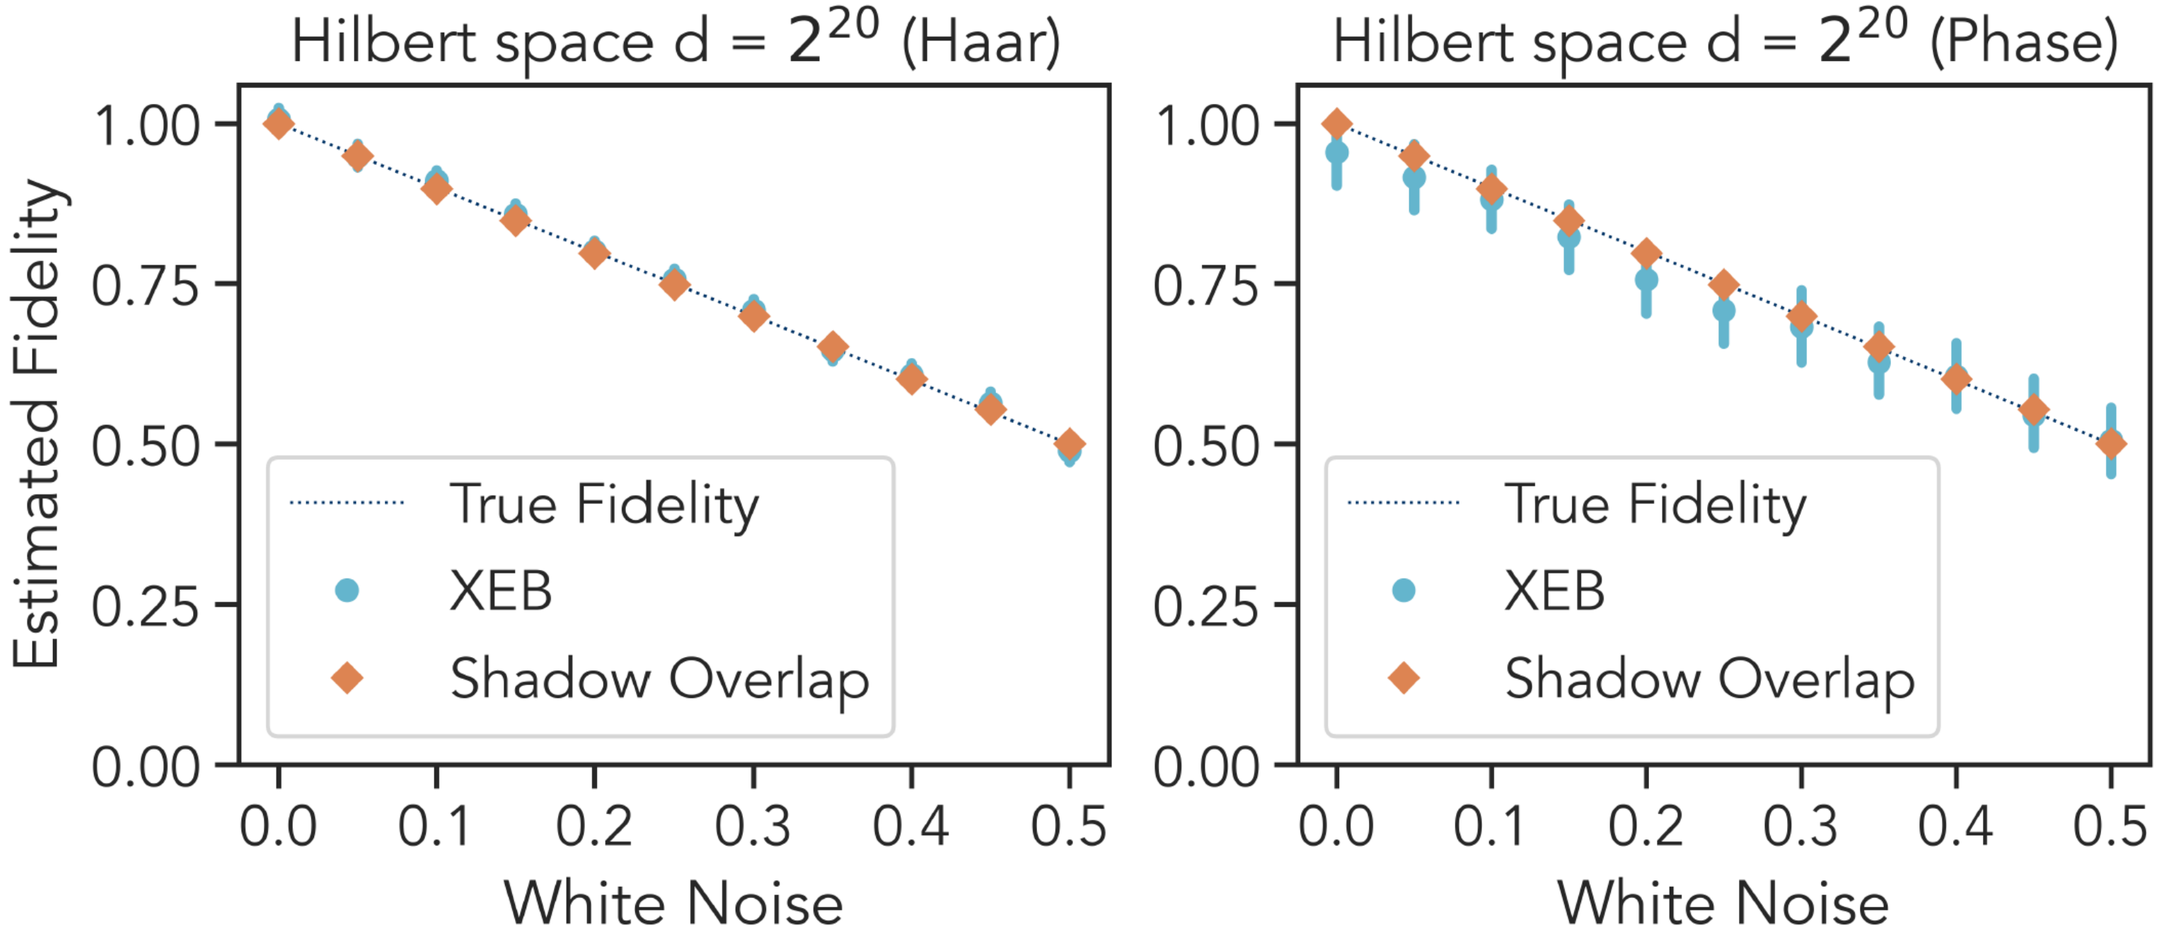

Target phase state: $|\psi\rangle = U_{phase} \otimes |\psi_i\rangle$.

White noise: global depolarizing noise.

Shadow overlap vs. XEB as proxies for fidelity

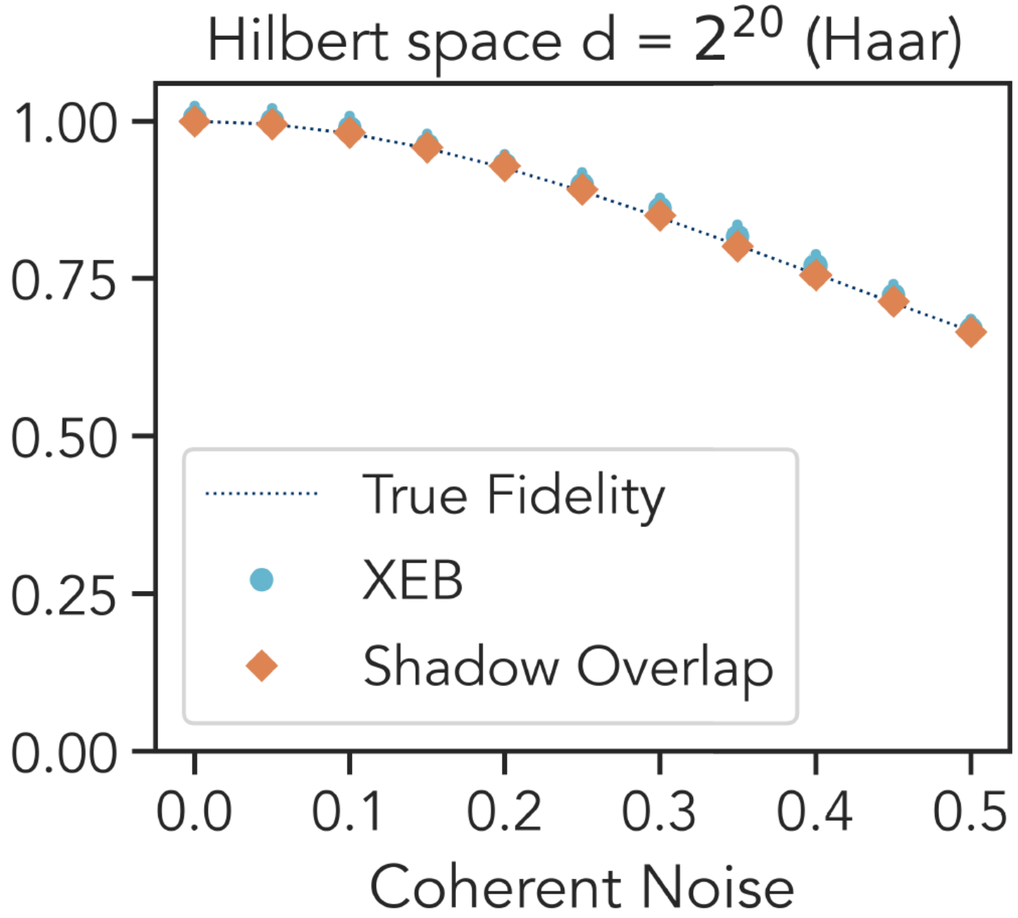

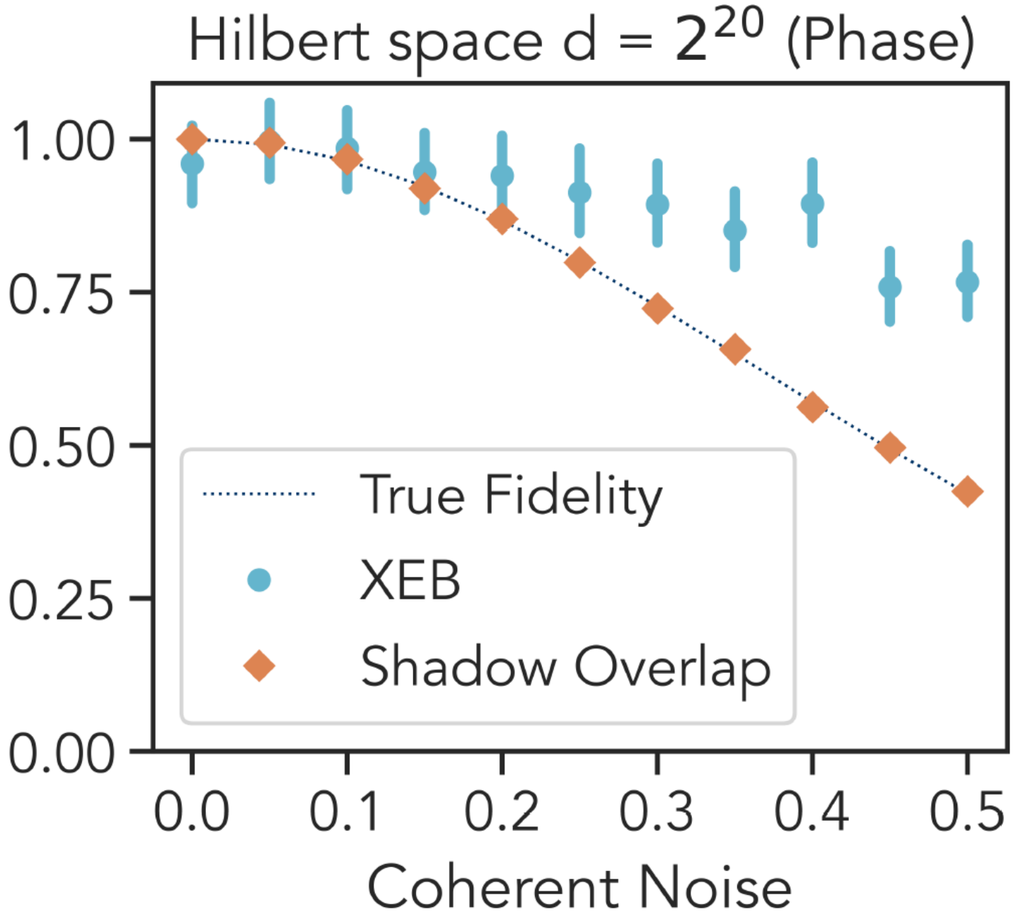

Target phase state: $|\psi\rangle = U_{phase} \otimes |\psi_i\rangle$.

Coherent noise: Gaussian-distributed shifts of amplitudes and phases.

Shadow overlap is more forgiving than fidelity

$|+++\ldots++\rangle$ and $|---\ldots--\rangle$

Fidelity is 0, shadow overlap is 0.

$|+++\ldots++\rangle$ and $|+++\ldots+-\rangle$

Fidelity is 0, shadow overlap is $(n\text{-}1)/n$.

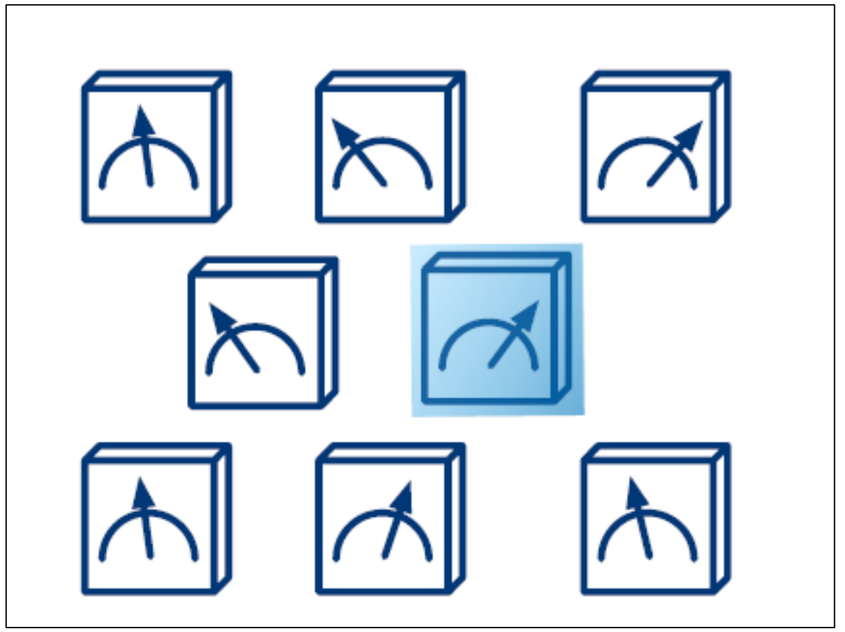

Fidelity remains stuck at zero, while the shadow overlap steadily improves as we flip more bits. (The shadow overlap behaves like the Hamming distance between bit strings.) Therefore the shadow overlap can be advantageous when we are training quantum circuits.

Training quantum circuits using the shadow overlap

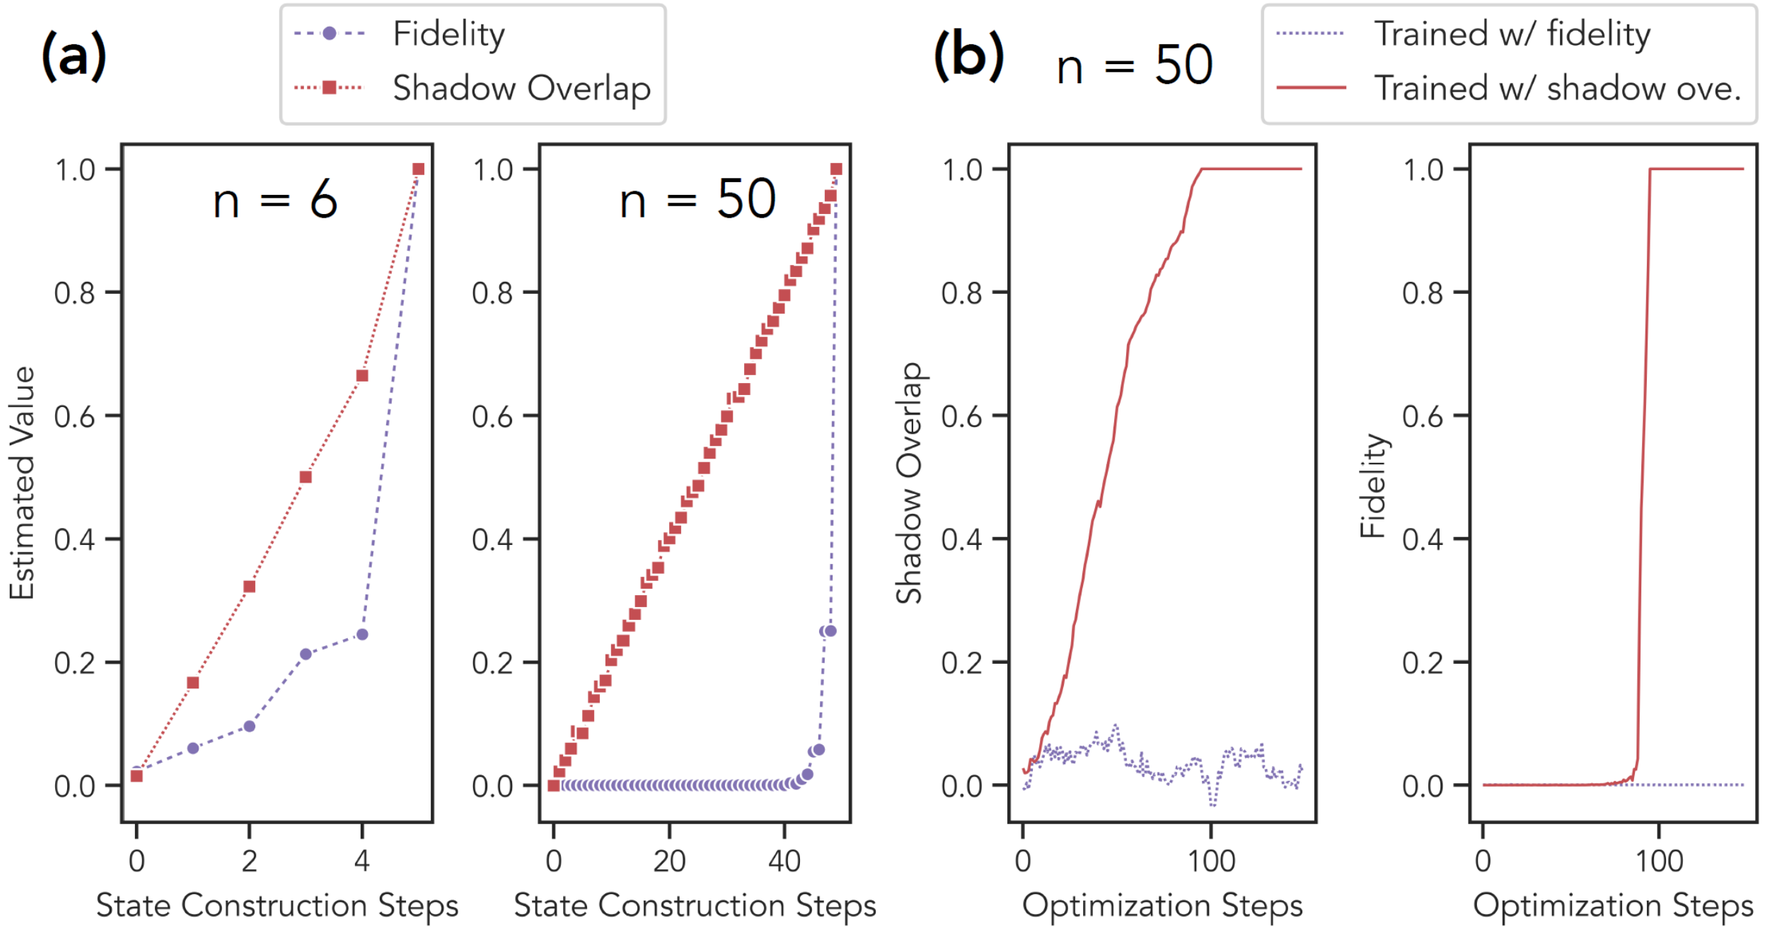

Goal: training a quantum circuit of H, CZ, T gates to prepare an MPS target state, using shadow overlap as a loss function.

(a) The shadow overlap increases steadily as circuit is executed.

(b) Training with fidelity fails, while training with shadow overlap succeeds.

Fidelity: Circuit wanders aimlessly on the barren plateau.

Shadow overlap: A high-fidelity circuit is easily found.

Certifying a quantum state

Using the *shadow overlap*, we can certify fidelity of a lab state with a target pure state using single-qubit measurements.

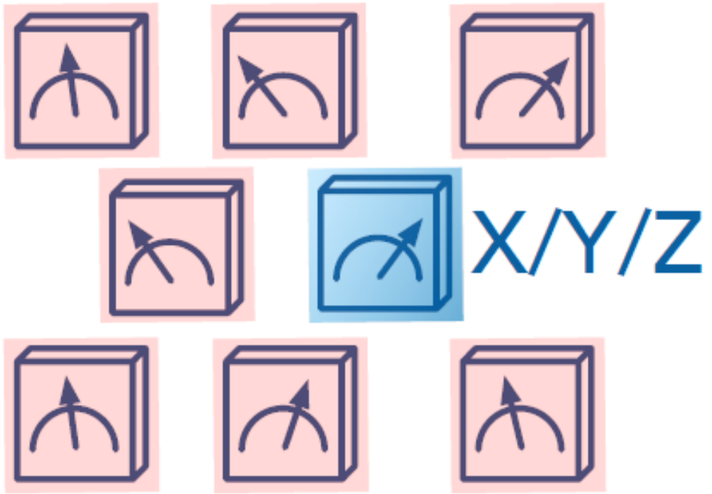

A more robust benchmarking method than, for example, the cross-entropy benchmark (XEB).

Evade barren plateaus by training quantum circuits using the shadow overlap rather than fidelity.

Can all pure quantum states be certified with single-qubit measurements of poly(n) copies?

Finding local minima and certifying quantum states

Anthony

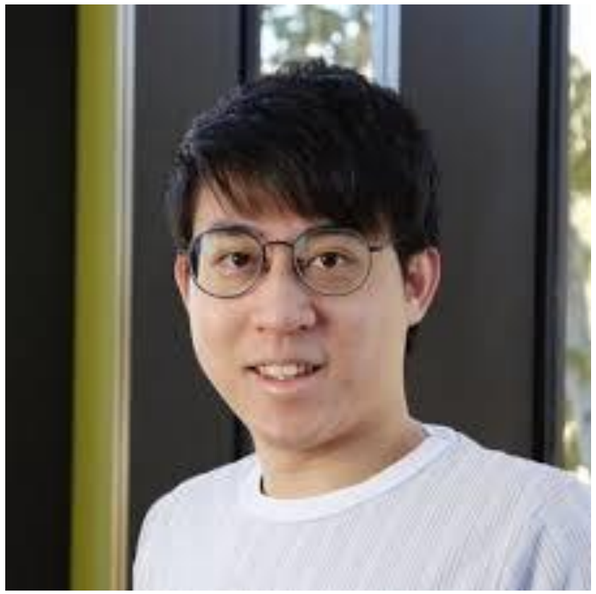

Robert

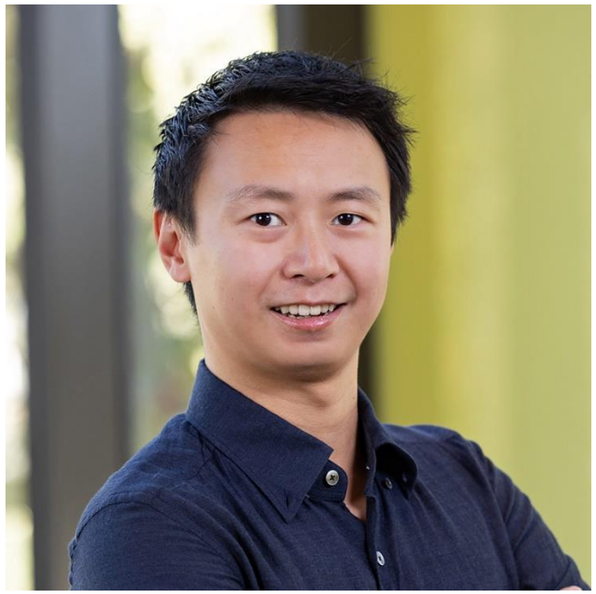

Leo

Medhi

arxiv:2309.16596 with Anthony Chen, Robert Huang, and Leo Zhou

arXiv:2404.07281 with Robert Huang and Mehdi Soleimanifar