Certifying and Learning Local Quantum Hamiltonians

arXiv:2603.29809 | TQC 2026

Junseo Lee

SNU & Harvard

junseolee@fas.harvard.edu

QNRC Workshop
May 11, 2026

서울대학교 컴퓨터연구소
INSTITUTE OF COMPUTER TECHNOLOGY

HARVARD UNIVERSITY

Harvard Quantum Initiative
IN SCIENCE AND ENGINEERING

In collaboration with wonderful people:

Andreas Bluhm

UGA Université Grenoble Alpes

Matthias C. Caro

University of Warwick

Francisco Escudero Gutiérrez

République Française — Liberté Égalité Fraternité — Inria

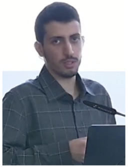

Aadil Oufkir

UM6P University Mohammed VI Polytechnic

Cambyse Rouzé

République Française — Liberté Égalité Fraternité — Inria

Myeongjin Shin

KAIST

Contents

Motivation

Why learn Hamiltonians?

Quantum computers and simulators are expected to solve problems beyond the reach of classical computation.

Why learn Hamiltonians?

Quantum computers and simulators are expected to solve problems beyond the reach of classical computation.

But this promise depends on building and controlling large-scale quantum hardware.

Why learn Hamiltonians?

Quantum computers and simulators are expected to solve problems beyond the reach of classical computation.

But this promise depends on building and controlling large-scale quantum hardware.

In particular, we need to accurately implement engineered quantum dynamics:

Why learn Hamiltonians?

Quantum computers and simulators are expected to solve problems beyond the reach of classical computation.

But this promise depends on building and controlling large-scale quantum hardware.

In particular, we need to accurately implement engineered quantum dynamics:

- to realize the intended computation,

Why learn Hamiltonians?

Quantum computers and simulators are expected to solve problems beyond the reach of classical computation.

But this promise depends on building and controlling large-scale quantum hardware.

In particular, we need to accurately implement engineered quantum dynamics:

- to realize the intended computation,
- to characterize the underlying device,

Why learn Hamiltonians?

Quantum computers and simulators are expected to solve problems beyond the reach of classical computation.

But this promise depends on building and controlling large-scale quantum hardware.

In particular, we need to accurately implement engineered quantum dynamics:

- to realize the intended computation,
- to characterize the underlying device,
- and to mitigate errors.

Why learn Hamiltonians?

Quantum computers and simulators are expected to solve problems beyond the reach of classical computation.

But this promise depends on building and controlling large-scale quantum hardware.

In particular, we need to accurately implement engineered quantum dynamics:

- to realize the intended computation,
- to characterize the underlying device,
- and to mitigate errors.

This motivates the problem of *Hamiltonian learning*:

Why learn Hamiltonians?

Quantum computers and simulators are expected to solve problems beyond the reach of classical computation.

But this promise depends on building and controlling large-scale quantum hardware.

In particular, we need to accurately implement engineered quantum dynamics:

- to realize the intended computation,
- to characterize the underlying device,
- and to mitigate errors.

This motivates the problem of *Hamiltonian learning*:

Can we learn the Hamiltonian governing an unknown quantum system?

From learning to certification

Hamiltonian learning is powerful, but sometimes it is just too much.

From learning to certification

Hamiltonian learning is powerful, but sometimes it is just too much.

In many situations, we do not need to learn the full Hamiltonian...

From learning to certification

Hamiltonian learning is powerful, but sometimes it is just too much.

In many situations, we do not need to learn the full Hamiltonian...

Suppose a provider claims that their device evolves according to a target Hamiltonian H_0.

From learning to certification

Hamiltonian learning is powerful, but sometimes it is just too much.

In many situations, we do not need to learn the full Hamiltonian...

Suppose a provider claims that their device evolves according to a target Hamiltonian H_0.

Then the relevant question is simply:

From learning to certification

Hamiltonian learning is powerful, but sometimes it is just too much.

In many situations, we do not need to learn the full Hamiltonian...

Suppose a provider claims that their device evolves according to a target Hamiltonian H_0.

Then the relevant question is simply:

Is the implemented Hamiltonian actually H_0?

Hamiltonian certification

This leads to *quantum Hamiltonian certification*!

Hamiltonian certification

This leads to *quantum Hamiltonian certification*!

Instead of reconstructing the entire Hamiltonian, we want to distinguish between:

Hamiltonian certification

This leads to *quantum Hamiltonian certification*!

Instead of reconstructing the entire Hamiltonian, we want to distinguish between:

$$H = H_0$$

Hamiltonian certification

This leads to *quantum Hamiltonian certification*!

Instead of reconstructing the entire Hamiltonian, we want to distinguish between:

$$H = H_0$$

and

$$H \text{ is } \varepsilon\text{-far from } H_0.$$

Hamiltonian certification

This leads to *quantum Hamiltonian certification*!

Instead of reconstructing the entire Hamiltonian, we want to distinguish between:

$$H = H_0$$

and

$$H \text{ is } \varepsilon\text{-far from } H_0.$$

This can be much less resource-intensive than full Hamiltonian learning.

Our setting

We study Hamiltonian certification from access to the dynamics

$$U_H(t) = e^{-itH}.$$

Our setting

We study Hamiltonian certification from access to the dynamics

$$U_H(t) = e^{-itH}.$$

The unknown Hamiltonian H acts on n quantum particles.

Our setting

We study Hamiltonian certification from access to the dynamics

$$U_H(t) = e^{-itH}.$$

The unknown Hamiltonian H acts on n quantum particles.

We focus on one of the most physically important classes:

Our setting

We study Hamiltonian certification from access to the dynamics

$$U_H(t) = e^{-itH}.$$

The unknown Hamiltonian H acts on n quantum particles.

We focus on one of the most physically important classes:

k-local Hamiltonians.

k-local Hamiltonians

A Hamiltonian is k-local if it can be written as a sum of terms, each acting non-trivially on at most k particles.

k-local Hamiltonians

A Hamiltonian is k-local if it can be written as a sum of terms, each acting non-trivially on at most k particles.

For example, on n qubits, we may write

$$H = \sum_{\substack{P \in \{I,X,Y,Z\}^{\otimes n} \\ |P| \leq k}} h_P P.$$

k-local Hamiltonians

A Hamiltonian is k-local if it can be written as a sum of terms, each acting non-trivially on at most k particles.

For example, on n qubits, we may write

$$H = \sum_{\substack{P \in \{I,X,Y,Z\}^{\otimes n} \\ |P| \leq k}} h_P P.$$

Here $|P|$ denotes the number of non-identity tensor factors in the Pauli string P.

k-local Hamiltonians

A Hamiltonian is k-local if it can be written as a sum of terms, each acting non-trivially on at most k particles.

For example, on n qubits, we may write

$$H = \sum_{\substack{P \in \{I,X,Y,Z\}^{\otimes n} \\ |P| \leq k}} h_P P.$$

Here $|P|$ denotes the number of non-identity tensor factors in the Pauli string P.

Throughout the talk, k should be thought of as a constant.

Why local Hamiltonians?

Local Hamiltonians are central objects in quantum science.

Why local Hamiltonians?

Local Hamiltonians are central objects in quantum science.

They include many physically relevant models:

Why local Hamiltonians?

Local Hamiltonians are central objects in quantum science.

They include many physically relevant models:

- geometrically local systems,

Why local Hamiltonians?

Local Hamiltonians are central objects in quantum science.

They include many physically relevant models:

- geometrically local systems,
- nearest-neighbor interaction models,

Why local Hamiltonians?

Local Hamiltonians are central objects in quantum science.

They include many physically relevant models:

- geometrically local systems,
- nearest-neighbor interaction models,
- SYK-type models,

Why local Hamiltonians?

Local Hamiltonians are central objects in quantum science.

They include many physically relevant models:

- geometrically local systems,
- nearest-neighbor interaction models,
- SYK-type models,
- classical and quantum spin systems.

Why local Hamiltonians?

Local Hamiltonians are central objects in quantum science.

They include many physically relevant models:

- geometrically local systems,
- nearest-neighbor interaction models,
- SYK-type models,
- classical and quantum spin systems.

They are also fundamental in quantum complexity theory.

Why local Hamiltonians?

Local Hamiltonians are central objects in quantum science.

They include many physically relevant models:

- geometrically local systems,
- nearest-neighbor interaction models,
- SYK-type models,
- classical and quantum spin systems.

They are also fundamental in quantum complexity theory.

For instance, the k-local Hamiltonian problem is QMA-complete, and local Hamiltonians are universal for quantum simulation.

Main question

The central question of this work is:

Main question

The central question of this work is:

What is the complexity of
certifying local Hamiltonians from dynamics?

Main question

The central question of this work is:

What is the complexity of
certifying local Hamiltonians from dynamics?

More concretely:

How much total evolution time is necessary and sufficient?

Main result: optimal certification

We give a certification algorithm for k-local Hamiltonians using total evolution time

$$O(1/\varepsilon).$$

Main result: optimal certification

We give a certification algorithm for k-local Hamiltonians using total evolution time

$$O(1/\varepsilon).$$

The algorithm distinguishes between:

$$H = H_0$$

Main result: optimal certification

We give a certification algorithm for k-local Hamiltonians using total evolution time

$$O(1/\varepsilon).$$

The algorithm distinguishes between:

$$H = H_0$$

and

$$H \text{ is } \varepsilon\text{-far from } H_0.$$

Main result: optimal certification

We give a certification algorithm for k-local Hamiltonians using total evolution time

$$O(1/\varepsilon).$$

The algorithm distinguishes between:

$$H = H_0$$

and

$$H \text{ is } \varepsilon\text{-far from } H_0.$$

This achieves *Heisenberg scaling* in the accuracy parameter ε.

Why is this optimal?

The dependence

$$O(1/\varepsilon)$$

is optimal due to known lower bounds.

Why is this optimal?

The dependence

$$O(1/\varepsilon)$$

is optimal due to known lower bounds.

Thus our protocol achieves the best possible scaling in total evolution time.

Why is this optimal?

The dependence

$$O(1/\varepsilon)$$

is optimal due to known lower bounds.

Thus our protocol achieves the best possible scaling in total evolution time.

To the best of our knowledge, this gives the first provably optimal result for a Hamiltonian property testing task!

Why is this optimal?

The dependence

$$O(1/\varepsilon)$$

is optimal due to known lower bounds.

Thus our protocol achieves the best possible scaling in total evolution time.

To the best of our knowledge, this gives the first provably optimal result for a Hamiltonian property testing task!

We also show that the protocol can be made *tolerant*.

Beyond dynamics: Gibbs-state access

We also study a different access model.

Beyond dynamics: Gibbs-state access

We also study a different access model.

Instead of accessing the dynamics e^{-itH}, suppose we are given copies of the Gibbs state

$$\rho(\beta) = \frac{e^{-\beta H}}{\mathrm{Tr}(e^{-\beta H})},$$

where β is the inverse temperature.

Beyond dynamics: Gibbs-state access

We also study a different access model.

Instead of accessing the dynamics e^{-itH}, suppose we are given copies of the Gibbs state

$$\rho(\beta) = \frac{e^{-\beta H}}{\mathrm{Tr}(e^{-\beta H})},$$

where β is the inverse temperature.

This is the natural thermal state associated with H.

Beyond dynamics: Gibbs-state access

We also study a different access model.

Instead of accessing the dynamics e^{-itH}, suppose we are given copies of the Gibbs state

$$\rho(\beta) = \frac{e^{-\beta H}}{\mathrm{Tr}(e^{-\beta H})},$$

where β is the inverse temperature.

This is the natural thermal state associated with H.

The question becomes:

Can we learn or certify this thermal state efficiently?

Why Gibbs states are challenging

A natural approach is to first learn the Hamiltonian H.

Why Gibbs states are challenging

A natural approach is to first learn the Hamiltonian H.

Once H is known, one can in principle reconstruct the Gibbs state $\rho(\beta)$.

Why Gibbs states are challenging

A natural approach is to first learn the Hamiltonian H.

Once H is known, one can in principle reconstruct the Gibbs state $\rho(\beta)$.

However... learning H from Gibbs-state access is known to require exponential dependence on β.

Why Gibbs states are challenging

A natural approach is to first learn the Hamiltonian H.

Once H is known, one can in principle reconstruct the Gibbs state $\rho(\beta)$.

However... learning H from Gibbs-state access is known to require exponential dependence on β.

Therefore, this approach is not efficient at low temperatures.

Why Gibbs states are challenging

A natural approach is to first learn the Hamiltonian H.

Once H is known, one can in principle reconstruct the Gibbs state $\rho(\beta)$.

However... learning H from Gibbs-state access is known to require exponential dependence on β.

Therefore, this approach is not efficient at low temperatures.

This suggests that Gibbs-state learning and certification should be treated as independent tasks.

Gibbs-state learning and certification

For learning thermal states of local Hamiltonians, we give an algorithm that is sample-efficient in **all** relevant parameters!

Gibbs-state learning and certification

For learning thermal states of local Hamiltonians, we give an algorithm that is sample-efficient in **all** relevant parameters!

For certifying such states, we give an algorithm that is both:

Gibbs-state learning and certification

For learning thermal states of local Hamiltonians, we give an algorithm that is sample-efficient in **all** relevant parameters!

For certifying such states, we give an algorithm that is both:

- sample-efficient,

Gibbs-state learning and certification

For learning thermal states of local Hamiltonians, we give an algorithm that is sample-efficient in **all** relevant parameters!

For certifying such states, we give an algorithm that is both:

- sample-efficient,
- and time-efficient.

Gibbs-state learning and certification

For learning thermal states of local Hamiltonians, we give an algorithm that is sample-efficient in **all** relevant parameters!

For certifying such states, we give an algorithm that is both:

- sample-efficient,
- and time-efficient.

This resolves the question of whether efficient certification of Gibbs states is possible in this setting.

Conceptual message

The main conceptual message is:

Conceptual message

The main conceptual message is:

Certification can be strictly easier than learning!

Conceptual message

The main conceptual message is:

Certification can be strictly easier than learning!

For dynamics, this leads to an optimal Heisenberg-limited certification protocol.

Conceptual message

The main conceptual message is:

Certification can be strictly easier than learning!

For dynamics, this leads to an optimal Heisenberg-limited certification protocol.

For thermal states, this lets us avoid the exponential-in-β barrier for Hamiltonian learning.

Setting

We consider n-qubit Hamiltonians H.

Setting

We consider n-qubit Hamiltonians H.

Throughout, H is assumed to be k-local, for a constant $k \in \mathbb{N}$.

Setting

We consider n-qubit Hamiltonians H.

Throughout, H is assumed to be k-local, for a constant $k \in \mathbb{N}$.

Thus, its Pauli expansion has the form

$$H = \sum_{\substack{P \in \{I,X,Y,Z\}^{\otimes n} \\ |P| \leq k}} h_P P.$$

Setting

We consider n-qubit Hamiltonians H.

Throughout, H is assumed to be k-local, for a constant $k \in \mathbb{N}$.

Thus, its Pauli expansion has the form

$$H = \sum_{\substack{P \in \{I,X,Y,Z\}^{\otimes n} \\ |P| \leq k}} h_P P.$$

Here $|P|$ denotes the number of sites where the Pauli string P differs from identity.

Setting

We consider n-qubit Hamiltonians H.

Throughout, H is assumed to be k-local, for a constant $k \in \mathbb{N}$.

Thus, its Pauli expansion has the form

$$H = \sum_{\substack{P \in \{I,X,Y,Z\}^{\otimes n} \\ |P| \leq k}} h_P P.$$

Here $|P|$ denotes the number of sites where the Pauli string P differs from identity.

We assume without loss of generality that

$$h_{I^{\otimes n}} = \frac{\mathrm{Tr}[H]}{2^n} = 0.$$

Why can we ignore the identity?

Two Hamiltonians that differ only by a multiple of the identity induce the same physical objects relevant to us.

Why can we ignore the identity?

Two Hamiltonians that differ only by a multiple of the identity induce the same physical objects relevant to us.

For dynamics,

$$e^{-it(H+cI)} = e^{-ict}e^{-itH},$$

so the difference is only a global phase.

Why can we ignore the identity?

Two Hamiltonians that differ only by a multiple of the identity induce the same physical objects relevant to us.

For dynamics,

$$e^{-it(H+cI)} = e^{-ict}e^{-itH},$$

so the difference is only a global phase.

For Gibbs states,

$$\frac{e^{-\beta(H+cI)}}{\mathrm{Tr}[e^{-\beta(H+cI)}]} = \frac{e^{-\beta H}}{\mathrm{Tr}[e^{-\beta H}]}.$$

Why can we ignore the identity?

Two Hamiltonians that differ only by a multiple of the identity induce the same physical objects relevant to us.

For dynamics,

$$e^{-it(H+cI)} = e^{-ict} e^{-itH},$$

so the difference is only a global phase.

For Gibbs states,

$$\frac{e^{-\beta(H+cI)}}{\mathrm{Tr}[e^{-\beta(H+cI)}]} = \frac{e^{-\beta H}}{\mathrm{Tr}[e^{-\beta H}]}.$$

Thus, assuming traceless Hamiltonians is without loss of generality.

Certification via access to the dynamics

Access via dynamics

If a quantum system is governed by H, then its time evolution is

$$U_H(t) = e^{-itH}.$$

Access via dynamics

If a quantum system is governed by H, then its time evolution is

$$U_H(t) = e^{-itH}.$$

If the initial state is ρ, then after time t it becomes

$$U_H(t)\rho U_H(t)^\dagger.$$

Access via dynamics

If a quantum system is governed by H, then its time evolution is

$$U_H(t) = e^{-itH}.$$

If the initial state is ρ, then after time t it becomes

$$U_H(t)\rho U_H(t)^\dagger.$$

Thus, a natural access model is query access to

$$U_H(t).$$

Access via dynamics

If a quantum system is governed by H, then its time evolution is

$$U_H(t) = e^{-itH}.$$

If the initial state is ρ, then after time t it becomes

$$U_H(t)\rho U_H(t)^\dagger.$$

Thus, a natural access model is query access to

$$U_H(t).$$

In the lab, one query to $U_H(t)$ corresponds to letting the system evolve for time t.

Experiments

An experiment can interleave unknown evolution and known controls:

$$\rho \longrightarrow U_H(t_1) \longrightarrow V_1 \longrightarrow U_H(t_2) \longrightarrow V_2 \longrightarrow \cdots \longrightarrow \text{measure.}$$

Experiments

An experiment can interleave unknown evolution and known controls:

$$\rho \longrightarrow U_H(t_1) \longrightarrow V_1 \longrightarrow U_H(t_2) \longrightarrow V_2 \longrightarrow \cdots \longrightarrow \text{measure.}$$

Here each V_i is a known unitary independent of H.

Experiments

An experiment can interleave unknown evolution and known controls:

$$\rho \longrightarrow U_H(t_1) \longrightarrow V_1 \longrightarrow U_H(t_2) \longrightarrow V_2 \longrightarrow \cdots \longrightarrow \text{measure.}$$

Here each V_i is a known unitary independent of H.

The main figure of merit is the *total evolution time* $\sum_i t_i$.

Experiments

An experiment can interleave unknown evolution and known controls:

$$\rho \longrightarrow U_H(t_1) \longrightarrow V_1 \longrightarrow U_H(t_2) \longrightarrow V_2 \longrightarrow \cdots \longrightarrow \text{measure.}$$

Here each V_i is a known unitary independent of H.

The main figure of merit is the *total evolution time* $\sum_i t_i$.

This is the resource that captures how long we need to let the unknown system evolve.

Experiments

An experiment can interleave unknown evolution and known controls:

$$\rho \longrightarrow U_H(t_1) \longrightarrow V_1 \longrightarrow U_H(t_2) \longrightarrow V_2 \longrightarrow \cdots \longrightarrow \text{measure}.$$

Here each V_i is a known unitary independent of H.

The main figure of merit is the *total evolution time* $\sum_i t_i$.

This is the resource that captures how long we need to let the unknown system evolve.

Spoiler: very short evolution times are often hard to implement.
So whether Hamiltonian learning can work with only long-time dynamics is an important question.

Experiments

An experiment can interleave unknown evolution and known controls:

$$\rho \longrightarrow U_H(t_1) \longrightarrow V_1 \longrightarrow U_H(t_2) \longrightarrow V_2 \longrightarrow \cdots \longrightarrow \text{measure.}$$

Here each V_i is a known unitary independent of H.

The main figure of merit is the *total evolution time* $\sum_i t_i$.

This is the resource that captures how long we need to let the unknown system evolve.

Spoiler: very short evolution times are often hard to implement.
So whether Hamiltonian learning can work with only long-time dynamics is an important question.

For that, please come to Myeongjin's talk tomorrow!
(Joint work with Myeongjin Shin, Changhun Oh, and myself.)

Distance measure

We assume that the Hamiltonians have bounded operator norm.

Distance measure

We assume that the Hamiltonians have bounded operator norm.

However, the total evolution time of our protocol will be independent of this operator norm bound.

Distance measure

We assume that the Hamiltonians have bounded operator norm.

However, the total evolution time of our protocol will be independent of this operator norm bound.

For the testing distance, we use the normalized Frobenius norm:

$$\|H - H'\|_F = \sqrt{\frac{\mathrm{Tr}[(H - H')^2]}{2^n}}.$$

Distance measure

We assume that the Hamiltonians have bounded operator norm.

However, the total evolution time of our protocol will be independent of this operator norm bound.

For the testing distance, we use the normalized Frobenius norm:

$$\|H - H'\|_F = \sqrt{\frac{\mathrm{Tr}[(H - H')^2]}{2^n}}.$$

This is an average-case notion of distance between Hamiltonians.

Why Frobenius distance?

The normalized Frobenius norm is natural for Hamiltonian testing.

Why Frobenius distance?

The normalized Frobenius norm is natural for Hamiltonian testing.

If

$$\|H - H'\|_F$$

is small, then the two Hamiltonians are close on average.

Why Frobenius distance?

The normalized Frobenius norm is natural for Hamiltonian testing.

If

$$\left\|H - H'\right\|_F$$

is small, then the two Hamiltonians are close on average.

More concretely, the expected values of observables generated by evolving Haar-random states under H and H' are close.

Why Frobenius distance?

The normalized Frobenius norm is natural for Hamiltonian testing.

If

$$\left\|H - H'\right\|_F$$

is small, then the two Hamiltonians are close on average.

More concretely, the expected values of observables generated by evolving Haar-random states under H and H' are close.

So our setting is:

bounded in a worst-case norm, tested in an average-case norm.

Result: Hamiltonian certification

Let H_0 be a known target Hamiltonian.

Result: Hamiltonian certification

Let H_0 be a known target Hamiltonian.

Let H be an unknown n-qubit k-local Hamiltonian.

Result: Hamiltonian certification

Let H_0 be a known target Hamiltonian.

Let H be an unknown n-qubit k-local Hamiltonian.

There is an algorithm with access to a description of H_0 and to the time evolution of H that uses total evolution time

$$O(1/\varepsilon).$$

Result: Hamiltonian certification

Let H_0 be a known target Hamiltonian.

Let H be an unknown n-qubit k-local Hamiltonian.

There is an algorithm with access to a description of H_0 and to the time evolution of H that uses total evolution time

$$O(1/\varepsilon).$$

With high probability, it distinguishes between

$$\|H - H_0\|_F \leq \frac{\varepsilon}{8 \cdot 3^k}$$

and

$$\|H - H_0\|_F \geq \varepsilon.$$

Main result 1 (Formal)

Theorem (Tolerant certification of k-local Hamiltonians)

Main result 1 (Formal)

Theorem (Tolerant certification of k-local Hamiltonians)

Let H and H_0 be k-local n-qubit Hamiltonians, where H_0 is known and H can be accessed via its time evolution operator. Let $\varepsilon, \delta > 0$. Let $C_{\mathrm{op}} \geq 1$ be such that $\|H_0\|_{\mathrm{op}}, \|H\|_{\mathrm{op}} \leq C_{\mathrm{op}}$.

Main result 1 (Formal)

Theorem (Tolerant certification of k-local Hamiltonians)

Let H and H_0 be k-local n-qubit Hamiltonians, where H_0 is known and H can be accessed via its time evolution operator. Let $\varepsilon, \delta > 0$. Let $C_{\mathrm{op}} \geq 1$ be such that $\|H_0\|_{\mathrm{op}}, \|H\|_{\mathrm{op}} \leq C_{\mathrm{op}}$.

- *There is an algorithm that uses $O(9^{2k}\log(1/\delta)/\varepsilon)$ total evolution time to test, with success probability $\geq 1-\delta$, whether $\|H - H_0\|_F \leq \varepsilon/(8 \cdot 3^k)$ or $\|H - H_0\|_F \geq \varepsilon$.*

Main result 1 (Formal)

Theorem (Tolerant certification of k-local Hamiltonians)

Let H and H_0 be k-local n-qubit Hamiltonians, where H_0 is known and H can be accessed via its time evolution operator. Let $\varepsilon, \delta > 0$. Let $C_{\text{op}} \geq 1$ be such that $\|H_0\|_{\text{op}}, \|H\|_{\text{op}} \leq C_{\text{op}}$.

- *There is an algorithm that uses $O(9^{2k} \log(1/\delta)/\varepsilon)$ total evolution time to test, with success probability $\geq 1 - \delta$, whether $\|H - H_0\|_F \leq \varepsilon/(8 \cdot 3^k)$ or $\|H - H_0\|_F \geq \varepsilon$.*
- *Furthermore, the algorithm uses no ancilla qubits, it makes $O(9^{2k} \log(1/\delta))$ experiments, and it makes $O(3^{5k}(C_{\text{op}}/\varepsilon)^{3/2} \log(1/\delta))$ time evolution queries.*

Main result 1 (Formal)

Theorem (Tolerant certification of k-local Hamiltonians)

Let H and H_0 be k-local n-qubit Hamiltonians, where H_0 is known and H can be accessed via its time evolution operator. Let $\varepsilon, \delta > 0$. Let $C_{\mathrm{op}} \geq 1$ be such that $\|H_0\|_{\mathrm{op}}, \|H\|_{\mathrm{op}} \leq C_{\mathrm{op}}$.

- *There is an algorithm that uses $O(9^{2k}\log(1/\delta)/\varepsilon)$ total evolution time to test, with success probability $\geq 1-\delta$, whether $\|H - H_0\|_F \leq \varepsilon/(8\cdot 3^k)$ or $\|H - H_0\|_F \geq \varepsilon$.*
- *Furthermore, the algorithm uses no ancilla qubits, it makes $O(9^{2k}\log(1/\delta))$ experiments, and it makes $O(3^{5k}(C_{\mathrm{op}}/\varepsilon)^{3/2}\log(1/\delta))$ time evolution queries.*
- *The algorithm is robust to at least $\Omega(9^{-k})$ SPAM errors, and the classical post-processing time is $O(9^{2k}\log(1/\delta))$.*

Optimality

The total evolution time

$$O(1/\varepsilon)$$

is optimal up to constant factors.

Optimality

The total evolution time

$$O(1/\varepsilon)$$

is optimal up to constant factors.

Indeed, $\Omega(1/\varepsilon)$ evolution time is already necessary to distinguish

$$H = \varepsilon X \qquad \text{from} \qquad H = -\varepsilon X.$$

Optimality

The total evolution time

$$O(1/\varepsilon)$$

is optimal up to constant factors.

Indeed, $\Omega(1/\varepsilon)$ evolution time is already necessary to distinguish

$$H = \varepsilon X \qquad \text{from} \qquad H = -\varepsilon X.$$

Thus, our result achieves the optimal Heisenberg scaling.

Optimality

The total evolution time

$$O(1/\varepsilon)$$

is optimal up to constant factors.

Indeed, $\Omega(1/\varepsilon)$ evolution time is already necessary to distinguish

$$H = \varepsilon X \qquad \text{from} \qquad H = -\varepsilon X.$$

Thus, our result achieves the optimal Heisenberg scaling.

To our knowledge, this gives the first optimal algorithm for testing a property of quantum Hamiltonians.

Comparison with prior certification

Previous Hamiltonian property testing algorithms were not optimal in total evolution time.

Comparison with prior certification

Previous Hamiltonian property testing algorithms were not optimal in total evolution time.

For example, the closest-to-optimal bound for a Hamiltonian property testing task was $O(1/\varepsilon^2)$, which is quadratically worse than the lower bound $\Omega(1/\varepsilon)$.

Comparison with prior certification

Previous Hamiltonian property testing algorithms were not optimal in total evolution time.

For example, the closest-to-optimal bound for a Hamiltonian property testing task was $O(1/\varepsilon^2)$, which is quadratically worse than the lower bound $\Omega(1/\varepsilon)$.

For certifying sparse Hamiltonians, a previous approach gives evolution time roughly $\widetilde{O}(n^{3k/2}/\varepsilon)$ when specialized to k-local Hamiltonians.

Comparison with prior certification

Previous Hamiltonian property testing algorithms were not optimal in total evolution time.

For example, the closest-to-optimal bound for a Hamiltonian property testing task was $O(1/\varepsilon^2)$, which is quadratically worse than the lower bound $\Omega(1/\varepsilon)$.

For certifying sparse Hamiltonians, a previous approach gives evolution time roughly $\widetilde{O}(n^{3k/2}/\varepsilon)$ when specialized to k-local Hamiltonians.

Our algorithm removes this polynomial dependence on n from the total evolution time!

Technical overview

For simplicity, let us first consider the intolerant version.

Technical overview

For simplicity, let us first consider the intolerant version.

We want to distinguish between

$$H = H_0$$

and

$$\|H - H_0\|_F \geq \varepsilon.$$

Technical overview

For simplicity, let us first consider the intolerant version.

We want to distinguish between

$$H = H_0$$

and

$$\|H - H_0\|_F \geq \varepsilon.$$

Define

$$\Delta H = H - H_0.$$

Technical overview

For simplicity, let us first consider the intolerant version.

We want to distinguish between

$$H = H_0$$

and

$$\|H - H_0\|_F \geq \varepsilon.$$

Define

$$\Delta H = H - H_0.$$

The goal is to detect whether ΔH is zero or has large Frobenius norm.

Spectral decomposition

Write the spectral decomposition of ΔH as

$$\Delta H = \sum_{s \in \{0,1\}^n} \lambda_s \left|\psi_s\right\rangle \left\langle\psi_s\right| .$$

Spectral decomposition

Write the spectral decomposition of ΔH as

$$\Delta H = \sum_{s \in \{0,1\}^n} \lambda_s \left| \psi_s \right\rangle \left\langle \psi_s \right| .$$

The key object will be the eigenvalue differences

$$\lambda_r - \lambda_s .$$

Spectral decomposition

Write the spectral decomposition of ΔH as

$$\Delta H = \sum_{s \in \{0,1\}^n} \lambda_s \left|\psi_s\right\rangle \left\langle\psi_s\right| .$$

The key object will be the eigenvalue differences

$$\lambda_r - \lambda_s .$$

Our test will detect whether many such eigenvalue differences are large.

Spectral decomposition

Write the spectral decomposition of ΔH as

$$\Delta H = \sum_{s \in \{0,1\}^n} \lambda_s \left|\psi_s\right\rangle \left\langle\psi_s\right| .$$

The key object will be the eigenvalue differences

$$\lambda_r - \lambda_s .$$

Our test will detect whether many such eigenvalue differences are large.

This is where locality of ΔH becomes crucial.

Simulating evolution under ΔH

We do not directly have access to

$$e^{-it\Delta H}.$$

Simulating evolution under ΔH

We do not directly have access to

$$e^{-it\Delta H}.$$

But we have access to

$$e^{-itH} \qquad \text{and} \qquad e^{-itH_0}.$$

Simulating evolution under ΔH

We do not directly have access to

$$e^{-it\Delta H}.$$

But we have access to

$$e^{-itH} \quad \text{and} \quad e^{-itH_0}.$$

Using Trotterization, we can approximate

$$e^{-it(H-H_0)} = e^{-it\Delta H}$$

to arbitrary accuracy.

Simulating evolution under ΔH

We do not directly have access to

$$e^{-it\Delta H}.$$

But we have access to

$$e^{-itH} \quad \text{and} \quad e^{-itH_0}.$$

Using Trotterization, we can approximate

$$e^{-it(H-H_0)} = e^{-it\Delta H}$$

to arbitrary accuracy.

This lets us perform tests on the difference Hamiltonian ΔH.

Bell sampling

Consider Bell sampling applied to

$$e^{-it\Delta H}.$$

Bell sampling

Consider Bell sampling applied to

$$e^{-it\Delta H}.$$

Equivalently, this samples Pauli strings according to the squared Pauli coefficients of

$$e^{-it\Delta H}.$$

Bell sampling

Consider Bell sampling applied to

$$e^{-it\Delta H}.$$

Equivalently, this samples Pauli strings according to the squared Pauli coefficients of

$$e^{-it\Delta H}.$$

The probability of measuring the identity Pauli is

$$I(t) = \frac{1}{4^n} \sum_{r,s \in \{0,1\}^n} \cos\left((\lambda_r - \lambda_s)t\right).$$

Bell sampling

Consider Bell sampling applied to

$$e^{-it\Delta H}.$$

Equivalently, this samples Pauli strings according to the squared Pauli coefficients of

$$e^{-it\Delta H}.$$

The probability of measuring the identity Pauli is

$$I(t) = \frac{1}{4^n} \sum_{r,s \in \{0,1\}^n} \cos\left((\lambda_r - \lambda_s)t\right).$$

So $I(t)$ depends only on eigenvalue differences of ΔH.

The key spectral quantity

Define the proportion of ε-separated eigenvalue pairs:

$$\Lambda(\Delta H, \varepsilon) = \frac{1}{4^n} \sum_{\substack{r,s: \\ |\lambda_r - \lambda_s| \geq \varepsilon}} 1.$$

The key spectral quantity

Define the proportion of ε-separated eigenvalue pairs:

$$\Lambda(\Delta H, \varepsilon) = \frac{1}{4^n} \sum_{\substack{r,s: \\ |\lambda_r - \lambda_s| \geq \varepsilon}} 1.$$

If many eigenvalue pairs are well separated, then Bell sampling is likely to output a non-identity Pauli.

The key spectral quantity

Define the proportion of ε-separated eigenvalue pairs:

$$\Lambda(\Delta H, \varepsilon) = \frac{1}{4^n} \sum_{\substack{r,s: \\ |\lambda_r - \lambda_s| \geq \varepsilon}} 1.$$

If many eigenvalue pairs are well separated, then Bell sampling is likely to output a non-identity Pauli.

For a uniformly random

$$t \in [0, 2/\varepsilon],$$

we show that

$$I(t) \leq 1 - \frac{\Lambda}{4}$$

with probability at least $1/3$.

From eigenvalue gaps to certification

The previous bound says:

$$\text{many large gaps} \quad \Longrightarrow \quad I(t) < 1.$$

From eigenvalue gaps to certification

The previous bound says:

$$\text{many large gaps} \quad \Longrightarrow \quad I(t) < 1.$$

But $I(t)$ is the probability of measuring the identity.

From eigenvalue gaps to certification

The previous bound says:

$$\text{many large gaps} \quad \Longrightarrow \quad I(t) < 1.$$

But $I(t)$ is the probability of measuring the identity.

Thus,

$$I(t) < 1$$

means that Bell sampling outputs non-identity Paulis with noticeable probability.

From eigenvalue gaps to certification

The previous bound says:

$$\text{many large gaps} \quad \Longrightarrow \quad I(t) < 1.$$

But $I(t)$ is the probability of measuring the identity.

Thus,

$$I(t) < 1$$

means that Bell sampling outputs non-identity Paulis with noticeable probability.

So the certification problem reduces to proving that

$$\|\Delta H\|_F \geq \varepsilon$$

implies many large eigenvalue gaps.

Moment calculation

Let

$$F(r, s) = \lambda_r - \lambda_s.$$

Moment calculation

Let

$$F(r, s) = \lambda_r - \lambda_s.$$

The second moment of F is

$$\frac{1}{4^n} \sum_{r,s} |F(r, s)|^2 = 2 \|\Delta H\|_2^2 = 2 \|\Delta H\|_F^2 .$$

Moment calculation

Let

$$F(r,s) = \lambda_r - \lambda_s.$$

The second moment of F is

$$\frac{1}{4^n}\sum_{r,s}|F(r,s)|^2 = 2\left\|\Delta H\right\|_2^2 = 2\left\|\Delta H\right\|_F^2.$$

The fourth moment is

$$\frac{1}{4^n}\sum_{r,s}|F(r,s)|^4 = 2\left\|\Delta H\right\|_4^4 + 6\left\|\Delta H\right\|_2^4.$$

Using hypercontractivity

Because ΔH is k-local, the quantum Bonami lemma gives

$$\|\Delta H\|_4^4 \leq 9^k \|\Delta H\|_2^4 .$$

Using hypercontractivity

Because ΔH is k-local, the quantum Bonami lemma gives

$$\|\Delta H\|_4^4 \leq 9^k \|\Delta H\|_2^4 .$$

Therefore, the fourth moment of

$$F(r, s) = \lambda_r - \lambda_s$$

is controlled by its second moment.

Using hypercontractivity

Because ΔH is k-local, the quantum Bonami lemma gives

$$\|\Delta H\|_4^4 \leq 9^k \|\Delta H\|_2^4 .$$

Therefore, the fourth moment of

$$F(r, s) = \lambda_r - \lambda_s$$

is controlled by its second moment.

This prevents the eigenvalue-gap distribution from being too concentrated near zero.

Paley–Zygmund step

Applying the Paley–Zygmund inequality gives

$$\Lambda(\Delta H, \|\Delta H\|_F) = \Pr_{r,s}\left[|\lambda_r - \lambda_s| \geq \|\Delta H\|_F\right] \geq \frac{1}{3} \cdot 9^{-k}.$$

Paley–Zygmund step

Applying the Paley–Zygmund inequality gives

$$\Lambda(\Delta H, \|\Delta H\|_F) = \Pr_{r,s}\left[|\lambda_r - \lambda_s| \geq \|\Delta H\|_F\right] \geq \frac{1}{3} \cdot 9^{-k}.$$

Hence,

$$\|\Delta H\|_F \geq \varepsilon \implies \Lambda(\Delta H, \varepsilon) = \Omega(9^{-k}).$$

Paley–Zygmund step

Applying the Paley–Zygmund inequality gives

$$\Lambda(\Delta H, \|\Delta H\|_F) = \Pr_{r,s}\left[|\lambda_r - \lambda_s| \geq \|\Delta H\|_F\right] \geq \frac{1}{3} \cdot 9^{-k}.$$

Hence,

$$\|\Delta H\|_F \geq \varepsilon \implies \Lambda(\Delta H, \varepsilon) = \Omega(9^{-k}).$$

This is the key structural fact behind the algorithm.

Certification algorithm

Repeat the following procedure $O(1)$ times:

Certification algorithm

Repeat the following procedure $O(1)$ times:

1. Sample a uniformly random time $t \in [0, 2/\varepsilon]$.

Certification algorithm

Repeat the following procedure $O(1)$ times:

1. Sample a uniformly random time $t \in [0, 2/\varepsilon]$.
2. Perform Bell sampling of $e^{-it\Delta H}$ a total of $O(9^{2k})$ times.

Certification algorithm

Repeat the following procedure $O(1)$ times:

1. Sample a uniformly random time $t \in [0, 2/\varepsilon]$.
2. Perform Bell sampling of $e^{-it\Delta H}$ a total of $O(9^{2k})$ times.
3. Output FAR if many non-identity Paulis are observed.

Certification algorithm

Repeat the following procedure $O(1)$ times:

1. Sample a uniformly random time $t \in [0, 2/\varepsilon]$.
2. Perform Bell sampling of $e^{-it\Delta H}$ a total of $O(9^{2k})$ times.
3. Output FAR if many non-identity Paulis are observed.

If no repetition outputs FAR, output EQUAL.

Total evolution time

Each sampled time satisfies

$$t \leq \frac{2}{\varepsilon}.$$

Total evolution time

Each sampled time satisfies

$$t \leq \frac{2}{\varepsilon}.$$

Each repetition uses

$$O(9^{2k})$$

Bell-sampling experiments.

Total evolution time

Each sampled time satisfies

$$t \le \frac{2}{\varepsilon}.$$

Each repetition uses

$$O(9^{2k})$$

Bell-sampling experiments.

Thus the total evolution time is bounded by

$$O(9^{2k}) \cdot \frac{2}{\varepsilon} = O_k(1/\varepsilon).$$

Total evolution time

Each sampled time satisfies

$$t \leq \frac{2}{\varepsilon}.$$

Each repetition uses

$$O(9^{2k})$$

Bell-sampling experiments.

Thus the total evolution time is bounded by

$$O(9^{2k}) \cdot \frac{2}{\varepsilon} = O_k(1/\varepsilon).$$

For constant k, this is exactly Heisenberg scaling.

Correctness: EQUAL case

Suppose

$$\Delta H = 0.$$

Correctness: EQUAL case

Suppose

$$\Delta H = 0.$$

Then

$$e^{-it\Delta H} = I$$

for every t.

Correctness: EQUAL case

Suppose

$$\Delta H = 0.$$

Then

$$e^{-it\Delta H} = I$$

for every t.

Therefore, Bell sampling outputs the identity Pauli with probability one.

Correctness: EQUAL case

Suppose

$$\Delta H = 0.$$

Then

$$e^{-it\Delta H} = I$$

for every t.

Therefore, Bell sampling outputs the identity Pauli with probability one.

Up to small Trotter error, the algorithm observes only a few non-identity outcomes.

Correctness: EQUAL case

Suppose

$$\Delta H = 0.$$

Then

$$e^{-it\Delta H} = I$$

for every t.

Therefore, Bell sampling outputs the identity Pauli with probability one.

Up to small Trotter error, the algorithm observes only a few non-identity outcomes.

Hence it outputs EQUAL.

Correctness: FAR case

Suppose

$$\|\Delta H\|_F \geq \varepsilon.$$

Correctness: FAR case

Suppose

$$\|\Delta H\|_F \geq \varepsilon.$$

Then

$$\Lambda(\Delta H, \varepsilon) = \Omega(9^{-k}).$$

Correctness: FAR case

Suppose

$$\|\Delta H\|_F \geq \varepsilon.$$

Then

$$\Lambda(\Delta H, \varepsilon) = \Omega(9^{-k}).$$

For random

$$t \in [0, 2/\varepsilon],$$

we get

$$I(t) \leq 1 - \Omega(9^{-k})$$

with constant probability.

Correctness: FAR case

Suppose

$$\|\Delta H\|_F \geq \varepsilon.$$

Then

$$\Lambda(\Delta H, \varepsilon) = \Omega(9^{-k}).$$

For random

$$t \in [0, 2/\varepsilon],$$

we get

$$I(t) \leq 1 - \Omega(9^{-k})$$

with constant probability.

Thus Bell sampling sees non-identity Paulis with noticeable probability, and the algorithm outputs FAR with high probability.

Additional features

Our certification algorithm is efficient beyond total evolution time.

Additional features

Our certification algorithm is efficient beyond total evolution time.

In particular:

- the number of experiments is independent of the operator norm bound;

Additional features

Our certification algorithm is efficient beyond total evolution time.

In particular:

- the number of experiments is independent of the operator norm bound;
- the total evolution time is independent of the operator norm bound;

Additional features

Our certification algorithm is efficient beyond total evolution time.

In particular:

- the number of experiments is independent of the operator norm bound;
- the total evolution time is independent of the operator norm bound;
- the algorithm is robust to constant SPAM error;

Additional features

Our certification algorithm is efficient beyond total evolution time.

In particular:

- the number of experiments is independent of the operator norm bound;
- the total evolution time is independent of the operator norm bound;
- the algorithm is robust to constant SPAM error;
- Bell sampling can be replaced by identity-probability estimation;

Additional features

Our certification algorithm is efficient beyond total evolution time.

In particular:

- the number of experiments is independent of the operator norm bound;
- the total evolution time is independent of the operator norm bound;
- the algorithm is robust to constant SPAM error;
- Bell sampling can be replaced by identity-probability estimation;
- hence, the algorithm can be implemented without quantum memory.

Learning thermal states

Learning thermal states

Now we move to Gibbs-state access.

Learning thermal states

Now we move to Gibbs-state access.

Let

$$\rho_H(\beta) = \frac{e^{-\beta H}}{\mathrm{Tr}[e^{-\beta H}]}$$

be the Gibbs state of an unknown local Hamiltonian H.

Learning thermal states

Now we move to Gibbs-state access.

Let

$$\rho_H(\beta) = \frac{e^{-\beta H}}{\mathrm{Tr}[e^{-\beta H}]}$$

be the Gibbs state of an unknown local Hamiltonian H.

A common approach is:

$$\text{learn } H \quad \Longrightarrow \quad \text{learn } \rho_H(\beta).$$

Learning thermal states

Now we move to Gibbs-state access.

Let

$$\rho_H(\beta) = \frac{e^{-\beta H}}{\mathrm{Tr}[e^{-\beta H}]}$$

be the Gibbs state of an unknown local Hamiltonian H.

A common approach is:

$$\text{learn } H \quad \Longrightarrow \quad \text{learn } \rho_H(\beta).$$

But Hamiltonian learning from Gibbs states can require sample complexity exponential in β.

Our approach

We avoid learning the Hamiltonian.

Our approach

We avoid learning the Hamiltonian.

Instead, we learn the Gibbs state directly!

Our approach

We avoid learning the Hamiltonian.

Instead, we learn the Gibbs state directly!

This allows us to obtain sample complexity polynomial in all relevant parameters.

Our approach

We avoid learning the Hamiltonian.

Instead, we learn the Gibbs state directly!

This allows us to obtain sample complexity polynomial in all relevant parameters.

The price is that the learning algorithm is not time-efficient, due to a covering-net search.

Result: Gibbs-state learning

Let $\rho_H(\beta)$ be the Gibbs state of an unknown n-qubit local Hamiltonian H.

Result: Gibbs-state learning

Let $\rho_H(\beta)$ be the Gibbs state of an unknown n-qubit local Hamiltonian H.

Assume

$$|h_P| \leq 1 \qquad \text{for every } P.$$

Result: Gibbs-state learning

Let $\rho_H(\beta)$ be the Gibbs state of an unknown n-qubit local Hamiltonian H.

Assume

$$|h_P| \leq 1 \qquad \text{for every } P.$$

Then there is an algorithm that ε-learns $\rho_H(\beta)$ in trace norm using

$$\widetilde{O}\left(\frac{\mathrm{poly}(n)\beta^2}{\varepsilon^4}\right)$$

single copies of the state.

Result: Gibbs-state learning

Let $\rho_H(\beta)$ be the Gibbs state of an unknown n-qubit local Hamiltonian H.

Assume

$$|h_P| \leq 1 \qquad \text{for every } P.$$

Then there is an algorithm that ε-learns $\rho_H(\beta)$ in trace norm using

$$\widetilde{O}\left(\frac{\mathrm{poly}(n)\beta^2}{\varepsilon^4}\right)$$

single copies of the state.

The guarantee holds with high probability.

Main result 2 (Formal)

Theorem (Learning Gibbs states)

Main result 2 (Formal)

Theorem (Learning Gibbs states)

Let ρ be the Gibbs state at inverse temperature β of an n-qubit and k-local Hamiltonian H with $|h_P| \leq 1$ for every P. Let $\delta, \varepsilon \in (0, 1)$.

Main result 2 (Formal)

Theorem (Learning Gibbs states)

Let ρ be the Gibbs state at inverse temperature β of an n-qubit and k-local Hamiltonian H with $|h_P| \leq 1$ for every P. Let $\delta, \varepsilon \in (0,1)$.

Then, from $O(3^k n^{2k} k \log(n/\delta)(\max\{\beta, 1\})^2/\varepsilon^4)$ single copies of ρ, our algorithm obtains $\rho' \in \mathcal{S}_{\varepsilon', n, k, \beta}$ such that $\|\rho' - \rho\|_{\mathrm{tr}} \leq \varepsilon$ with probability $\geq 1 - \delta$.

Main result 2 (Formal)

Theorem (Learning Gibbs states)

Let ρ be the Gibbs state at inverse temperature β of an n-qubit and k-local Hamiltonian H with $|h_P| \leq 1$ for every P. Let $\delta, \varepsilon \in (0,1)$.

Then, from $O(3^k n^{2k} k \log(n/\delta)(\max\{\beta, 1\})^2/\varepsilon^4)$ single copies of ρ, our algorithm obtains $\rho' \in \mathcal{S}_{\varepsilon', n, k, \beta}$ such that $\|\rho' - \rho\|_{\mathrm{tr}} \leq \varepsilon$ with probability $\geq 1 - \delta$.

The classical post-processing time of the protocol is $(n^k \max\{\beta, 1\}/\varepsilon)^{O(n^k)}$.

Why is this useful?

When

$$\beta = \mathrm{poly}(n),$$

this gives Gibbs-state tomography with polynomial sample complexity.

Why is this useful?

When

$$\beta = \mathrm{poly}(n),$$

this gives Gibbs-state tomography with polynomial sample complexity.

General state tomography requires

$$\Theta(4^n)$$

copies.

Why is this useful?

When

$$\beta = \mathrm{poly}(n),$$

this gives Gibbs-state tomography with polynomial sample complexity.

General state tomography requires

$$\Theta(4^n)$$

copies.

Thus, for local Gibbs states, we obtain an exponential sample-complexity improvement over general tomography.

Why is this useful?

When

$$\beta = \mathrm{poly}(n),$$

this gives Gibbs-state tomography with polynomial sample complexity.

General state tomography requires

$$\Theta(4^n)$$

copies.

Thus, for local Gibbs states, we obtain an exponential sample-complexity improvement over general tomography.

The only structural assumption is locality of the Hamiltonian.

Gibbs-state continuity

The proof begins with the inequality

$$\|\rho_H(\beta) - \rho_{H'}(\beta)\|_{\mathrm{tr}} \leq \sqrt{2\beta \operatorname{Tr}[(\rho_H(\beta) - \rho_{H'}(\beta))(H' - H)]}.$$

Gibbs-state continuity

The proof begins with the inequality

$$\|\rho_H(\beta) - \rho_{H'}(\beta)\|_{\mathrm{tr}} \leq \sqrt{2\beta \operatorname{Tr}[(\rho_H(\beta) - \rho_{H'}(\beta))(H' - H)]}.$$

For k-local Hamiltonians, this implies

$$\|\rho_H(\beta) - \rho_{H'}(\beta)\|_{\mathrm{tr}} = O\left(\beta n^k \max_{|P| \leq k} |h_P - h'_P|\right).$$

Gibbs-state continuity

The proof begins with the inequality

$$\|\rho_H(\beta) - \rho_{H'}(\beta)\|_{\mathrm{tr}} \leq \sqrt{2\beta \operatorname{Tr}[(\rho_H(\beta) - \rho_{H'}(\beta))(H' - H)]}.$$

For k-local Hamiltonians, this implies

$$\|\rho_H(\beta) - \rho_{H'}(\beta)\|_{\mathrm{tr}} = O\left(\beta n^k \max_{|P| \leq k} |h_P - h'_P|\right).$$

So coefficient-wise approximation of the Hamiltonian gives trace-norm approximation of the Gibbs state.

Covering net

Choose a grid spacing

$$\eta \asymp \frac{\varepsilon}{\beta n^k}.$$

Covering net

Choose a grid spacing

$$\eta \asymp \frac{\varepsilon}{\beta n^k}.$$

Define a finite net of k-local Hamiltonians:

$$\mathcal{H}_\eta = \{H : h_P \in \eta\mathbb{Z} \cap [-1, 1] \text{ for all } |P| \leq k\}.$$

Covering net

Choose a grid spacing

$$\eta \asymp \frac{\varepsilon}{\beta n^k}.$$

Define a finite net of k-local Hamiltonians:

$$\mathcal{H}_\eta = \{H : h_P \in \eta\mathbb{Z} \cap [-1, 1] \text{ for all } |P| \leq k\}.$$

Then

$$\mathcal{S}_\eta = \{\rho_H(\beta) : H \in \mathcal{H}_\eta\}$$

is an ε-covering net for local Gibbs states.

Using classical shadows

For $H, H' \in \mathcal{H}_\eta$, the observable

$$H - H'$$

is a sum of k-local Pauli strings.

Using classical shadows

For $H, H' \in \mathcal{H}_\eta$, the observable

$$H - H'$$

is a sum of k-local Pauli strings.

Therefore, classical shadows can estimate all quantities

$$\mathrm{Tr}[\rho(H - H')]$$

sample-efficiently.

Using classical shadows

For $H, H' \in \mathcal{H}_\eta$, the observable

$$H - H'$$

is a sum of k-local Pauli strings.

Therefore, classical shadows can estimate all quantities

$$\mathrm{Tr}[\rho(H - H')]$$

sample-efficiently.

Let $\Delta_{H,H'}$ denote these estimates.

Using classical shadows

For $H, H' \in \mathcal{H}_\eta$, the observable

$$H - H'$$

is a sum of k-local Pauli strings.

Therefore, classical shadows can estimate all quantities

$$\mathrm{Tr}[\rho(H - H')]$$

sample-efficiently.

Let $\Delta_{H,H'}$ denote these estimates.

These estimates allow us to compare the unknown state with every candidate Gibbs state in the net.

Selecting a hypothesis

We output a net point that best matches the observed estimates.

Selecting a hypothesis

We output a net point that best matches the observed estimates.

Formally, choose

$$\rho' = \operatorname*{argmin}_{\rho' \in \mathcal{S}_\eta} \max_{H, H' \in \mathcal{H}_\eta} \left| \Delta_{H,H'} - \mathrm{Tr}[\rho'(H - H')] \right| .$$

Selecting a hypothesis

We output a net point that best matches the observed estimates.

Formally, choose

$$\rho' = \operatorname*{argmin}_{\rho' \in \mathcal{S}_\eta} \max_{H, H' \in \mathcal{H}_\eta} \left| \Delta_{H,H'} - \mathrm{Tr}[\rho'(H - H')] \right| .$$

The analysis shows that, even with estimation error and net discretization error,

$$\left\| \rho - \rho' \right\|_{\mathrm{tr}} \leq \varepsilon$$

with high probability.

Why learning is not time-efficient

The sample complexity is polynomial.

Why learning is not time-efficient

The sample complexity is polynomial.

However, the algorithm searches over the covering net

$$\mathcal{S}_\eta.$$

Why learning is not time-efficient

The sample complexity is polynomial.

However, the algorithm searches over the covering net

$$\mathcal{S}_\eta.$$

This net is exponentially large.

Why learning is not time-efficient

The sample complexity is polynomial.

However, the algorithm searches over the covering net

$$\mathcal{S}_\eta.$$

This net is exponentially large.

Therefore, the Gibbs-state learning algorithm is sample-efficient, but not time-efficient.

Why learning is not time-efficient

The sample complexity is polynomial.

However, the algorithm searches over the covering net

$$\mathcal{S}_\eta.$$

This net is exponentially large.

Therefore, the Gibbs-state learning algorithm is sample-efficient, but not time-efficient.

This motivates the certification problem, where a target state is known.

Certifying thermal states

Certifying thermal states

Now suppose we are given a target Hamiltonian H_0.

Certifying thermal states

Now suppose we are given a target Hamiltonian H_0.

We want to certify whether an unknown Gibbs state equals

$$\rho_{H_0}(\beta).$$

Certifying thermal states

Now suppose we are given a target Hamiltonian H_0.

We want to certify whether an unknown Gibbs state equals

$$\rho_{H_0}(\beta).$$

The task is to distinguish

$$\rho_H(\beta) = \rho_{H_0}(\beta)$$

from

$$\|\rho_H(\beta) - \rho_{H_0}(\beta)\|_{\mathrm{tr}} \geq \varepsilon.$$

Certifying thermal states

Now suppose we are given a target Hamiltonian H_0.

We want to certify whether an unknown Gibbs state equals

$$\rho_{H_0}(\beta).$$

The task is to distinguish

$$\rho_H(\beta) = \rho_{H_0}(\beta)$$

from

$$\|\rho_H(\beta) - \rho_{H_0}(\beta)\|_{\mathrm{tr}} \geq \varepsilon.$$

Here the goal is both sample efficiency and time efficiency.

Result: Gibbs-state certification

Let $\rho_H(\beta)$ and $\rho_{H_0}(\beta)$ be Gibbs states of local Hamiltonians H and H_0.

Result: Gibbs-state certification

Let $\rho_H(\beta)$ and $\rho_{H_0}(\beta)$ be Gibbs states of local Hamiltonians H and H_0.

Assume

$$|h_P|,\ |(h_0)_P| \leq 1 \qquad \text{for every } P.$$

Result: Gibbs-state certification

Let $\rho_H(\beta)$ and $\rho_{H_0}(\beta)$ be Gibbs states of local Hamiltonians H and H_0.

Assume

$$|h_P|,\ |(h_0)_P| \leq 1 \qquad \text{for every } P.$$

There is an algorithm that decides whether

$$\rho_H(\beta) = \rho_{H_0}(\beta)$$

or

$$\|\rho_H(\beta) - \rho_{H_0}(\beta)\|_{\mathrm{tr}} \geq \varepsilon$$

using

$$\widetilde{O}\left(\frac{\mathrm{poly}(n)\beta^2}{\varepsilon^4}\right)$$

single copies.

Main result 3 (Formal)

Theorem (Certifying Gibbs states)

Main result 3 (Formal)

Theorem (Certifying Gibbs states)

Let ρ and ρ_0 be the Gibbs states at inverse temperature β of n-qubit and k-local Hamiltonians H and H_0 with $|h_P|, |(h_0)_P| \leq 1$ for every P, respectively. Assume that H_0 is known. Let $\delta, \varepsilon \in (0, 1)$.

Main result 3 (Formal)

Theorem (Certifying Gibbs states)

Let ρ and ρ_0 be the Gibbs states at inverse temperature β of n-qubit and k-local Hamiltonians H and H_0 with $|h_P|, |(h_0)_P| \leq 1$ for every P, respectively. Assume that H_0 is known. Let $\delta, \varepsilon \in (0, 1)$.

Then, our algorithm decides, with success probability $\geq 1 - \delta$, whether $\|\rho - \rho_0\|_{\mathrm{tr}} \leq \varepsilon^2/(400\beta n^k)$ or $\|\rho - \rho_0\|_{\mathrm{tr}} \geq 2\varepsilon$ with $O\left(\beta^2 n^{2k} 3^k k \log(n/\delta)/\varepsilon^4\right)$ single copies of ρ and ρ_0.

Main result 3 (Formal)

Theorem (Certifying Gibbs states)

Let ρ and ρ_0 be the Gibbs states at inverse temperature β of n-qubit and k-local Hamiltonians H and H_0 with $|h_P|, |(h_0)_P| \leq 1$ for every P, respectively. Assume that H_0 is known. Let $\delta, \varepsilon \in (0,1)$.

Then, our algorithm decides, with success probability $\geq 1-\delta$, whether $\|\rho - \rho_0\|_{\mathrm{tr}} \leq \varepsilon^2/(400\beta n^k)$ or $\|\rho - \rho_0\|_{\mathrm{tr}} \geq 2\varepsilon$ with $O\left(\beta^2 n^{2k} 3^k k \log(n/\delta)/\varepsilon^4\right)$ single copies of ρ and ρ_0.

Moreover, the protocol only requires Pauli measurements, and a classical post-processing time of order $O\left(\beta^2 n^{3k} 3^k k \log(n/\delta)/\varepsilon^4\right)$. The same conclusion holds if ρ and ρ_0 are both unknown and we are given copy access to both.

Why certification is easier

Unlike learning, certification does not require searching over all possible Gibbs states.

Why certification is easier

Unlike learning, certification does not require searching over all possible Gibbs states.

The target state is fixed:

$$\rho_{H_0}(\beta).$$

Why certification is easier

Unlike learning, certification does not require searching over all possible Gibbs states.

The target state is fixed:

$$\rho_{H_0}(\beta).$$

So the algorithm only needs to test whether the unknown state is close to this target.

Why certification is easier

Unlike learning, certification does not require searching over all possible Gibbs states.

The target state is fixed:

$$\rho_{H_0}(\beta).$$

So the algorithm only needs to test whether the unknown state is close to this target.

This avoids the time-inefficient covering-net search.

Why certification is easier

Unlike learning, certification does not require searching over all possible Gibbs states.

The target state is fixed:

$$\rho_{H_0}(\beta).$$

So the algorithm only needs to test whether the unknown state is close to this target.

This avoids the time-inefficient covering-net search.

As a result, the certification algorithm is both sample-efficient and time-efficient.

Speedup over general certification

When

$$\beta = \mathrm{poly}(n),$$

our Gibbs-state certification algorithm uses polynomially many copies.

Speedup over general certification

When

$$\beta = \mathrm{poly}(n),$$

our Gibbs-state certification algorithm uses polynomially many copies.

General quantum state certification requires

$$\Theta(2^n)$$

copies.

Speedup over general certification

When

$$\beta = \mathrm{poly}(n),$$

our Gibbs-state certification algorithm uses polynomially many copies.

General quantum state certification requires

$$\Theta(2^n)$$

copies.

Thus, local Gibbs-state certification admits an exponential improvement over general state certification.

Speedup over general certification

When

$$\beta = \mathrm{poly}(n),$$

our Gibbs-state certification algorithm uses polynomially many copies.

General quantum state certification requires

$$\Theta(2^n)$$

copies.

Thus, local Gibbs-state certification admits an exponential improvement over general state certification.

This resolves the question of whether efficient certification is possible for local Gibbs states.

Overall picture

We obtain three main results.

Overall picture

We obtain three main results.

1. *Dynamics access:*

$$O(1/\varepsilon)$$

total evolution time for local Hamiltonian certification.

Overall picture

We obtain three main results.

1. *Dynamics access:*

$$O(1/\varepsilon)$$

total evolution time for local Hamiltonian certification.

2. *Gibbs-state access:*

$$\widetilde{O}\left(\frac{\mathrm{poly}(n)\beta^2}{\varepsilon^4}\right)$$

samples for learning local Gibbs states.

Overall picture

We obtain three main results.

1. *Dynamics access:*
$$O(1/\varepsilon)$$
total evolution time for local Hamiltonian certification.
2. *Gibbs-state access:*
$$\widetilde{O}\left(\frac{\mathrm{poly}(n)\beta^2}{\varepsilon^4}\right)$$
samples for learning local Gibbs states.
3. *Gibbs-state certification:* same sample complexity, but also time-efficient.

Discussion and open problems

What we have shown

In this work, we studied certification and learning problems for local Hamiltonians.

What we have shown

In this work, we studied certification and learning problems for local Hamiltonians.

First, we considered certification from time-evolution access.

What we have shown

In this work, we studied certification and learning problems for local Hamiltonians.

First, we considered certification from time-evolution access.

We gave an algorithm for local Hamiltonian certification with optimal total evolution time

$$O(1/\varepsilon).$$

What we have shown

In this work, we studied certification and learning problems for local Hamiltonians.

First, we considered certification from time-evolution access.

We gave an algorithm for local Hamiltonian certification with optimal total evolution time

$$O(1/\varepsilon).$$

To our knowledge, this gives the first optimal bound for any Hamiltonian property testing task in the time-evolution access model.

From Hamiltonians to Gibbs states

We then shifted focus from Hamiltonians themselves to their associated Gibbs states.

From Hamiltonians to Gibbs states

We then shifted focus from Hamiltonians themselves to their associated Gibbs states.

For Gibbs-state learning, we obtained a sample-efficient algorithm in all relevant parameters.

From Hamiltonians to Gibbs states

We then shifted focus from Hamiltonians themselves to their associated Gibbs states.

For Gibbs-state learning, we obtained a sample-efficient algorithm in all relevant parameters.

For Gibbs-state certification, we obtained an algorithm that is both sample-efficient and time-efficient.

From Hamiltonians to Gibbs states

We then shifted focus from Hamiltonians themselves to their associated Gibbs states.

For Gibbs-state learning, we obtained a sample-efficient algorithm in all relevant parameters.

For Gibbs-state certification, we obtained an algorithm that is both sample-efficient and time-efficient.

This overcomes the known exponential-in-β barrier that appears when one first tries to learn the Hamiltonian from Gibbs-state access.

Conceptual message

The results suggest that Gibbs-state learning and testing should be viewed as tasks of independent interest.

Conceptual message

The results suggest that Gibbs-state learning and testing should be viewed as tasks of independent interest.

They are not merely consequences of Hamiltonian learning.

Conceptual message

The results suggest that Gibbs-state learning and testing should be viewed as tasks of independent interest.

They are not merely consequences of Hamiltonian learning.

Sometimes, learning the state is easier than learning the Hamiltonian.

Conceptual message

The results suggest that Gibbs-state learning and testing should be viewed as tasks of independent interest.

They are not merely consequences of Hamiltonian learning.

Sometimes, learning the state is easier than learning the Hamiltonian.

Similarly, Hamiltonian certification can be much easier than full Hamiltonian learning.

Open problem I: beyond locality

Our optimal certification algorithm uses locality in an essential way.

Open problem I: beyond locality

Our optimal certification algorithm uses locality in an essential way.

For $O(1)$-local Hamiltonians, we achieve optimal total evolution time

$$O(1/\varepsilon)$$

even without inverse or controlled time evolution.

Open problem I: beyond locality

Our optimal certification algorithm uses locality in an essential way.

For $O(1)$-local Hamiltonians, we achieve optimal total evolution time

$$O(1/\varepsilon)$$

even without inverse or controlled time evolution.

A fundamental open question is:

Open problem I: beyond locality

Our optimal certification algorithm uses locality in an essential way.

For $O(1)$-local Hamiltonians, we achieve optimal total evolution time

$$O(1/\varepsilon)$$

even without inverse or controlled time evolution.

A fundamental open question is:

Can one achieve $O(1/\varepsilon)$ certification without any locality assumption?

Is locality fundamental?

For general non-local Hamiltonians, previous optimal certification protocols typically rely on stronger access models.

Is locality fundamental?

For general non-local Hamiltonians, previous optimal certification protocols typically rely on stronger access models.

For example, they may require inverse evolution or controlled evolution.

Is locality fundamental?

For general non-local Hamiltonians, previous optimal certification protocols typically rely on stronger access models.

For example, they may require inverse evolution or controlled evolution.

Our protocol avoids these assumptions, but only in the local setting.

Is locality fundamental?

For general non-local Hamiltonians, previous optimal certification protocols typically rely on stronger access models.

For example, they may require inverse evolution or controlled evolution.

Our protocol avoids these assumptions, but only in the local setting.

Resolving this would clarify whether locality is merely a technical condition in our analysis, or whether it is fundamentally needed for evolution-time-optimal certification.

Open problem II: time-efficient Gibbs-state learning

Our Gibbs-state learning algorithm is sample-efficient in all relevant parameters.

Open problem II: time-efficient Gibbs-state learning

Our Gibbs-state learning algorithm is sample-efficient in all relevant parameters.

However, it is not time-efficient, because the proof relies on a covering-net search.

Open problem II: time-efficient Gibbs-state learning

Our Gibbs-state learning algorithm is sample-efficient in all relevant parameters.

However, it is not time-efficient, because the proof relies on a covering-net search.

This leads to the natural question:

Open problem II: time-efficient Gibbs-state learning

Our Gibbs-state learning algorithm is sample-efficient in all relevant parameters.

However, it is not time-efficient, because the proof relies on a covering-net search.

This leads to the natural question:

Can Gibbs states of k-local Hamiltonians be learned both sample-efficiently and time-efficiently?

Why this question is natural

There is a useful analogy with Hamiltonian learning from Gibbs states.

Why this question is natural

There is a useful analogy with Hamiltonian learning from Gibbs states.

Early algorithms were sample-efficient but not time-efficient.

Why this question is natural

There is a useful analogy with Hamiltonian learning from Gibbs states.

Early algorithms were sample-efficient but not time-efficient.

Later works made Hamiltonian learning time-efficient under appropriate assumptions.

Why this question is natural

There is a useful analogy with Hamiltonian learning from Gibbs states.

Early algorithms were sample-efficient but not time-efficient.

Later works made Hamiltonian learning time-efficient under appropriate assumptions.

It is natural to ask whether the same evolution is possible for Gibbs-state learning itself.

Open problem III: optimal Gibbs-state certification

Our Gibbs-state certification algorithm is efficient in both sample and time complexity.

Open problem III: optimal Gibbs-state certification

Our Gibbs-state certification algorithm is efficient in both sample and time complexity.

However, we do not know whether the sample complexity is optimal.

Open problem III: optimal Gibbs-state certification

Our Gibbs-state certification algorithm is efficient in both sample and time complexity.

However, we do not know whether the sample complexity is optimal.

Even in the classical setting, the precise complexity of testing local Gibbs distributions appears to be open.

Open problem III: optimal Gibbs-state certification

Our Gibbs-state certification algorithm is efficient in both sample and time complexity.

However, we do not know whether the sample complexity is optimal.

Even in the classical setting, the precise complexity of testing local Gibbs distributions appears to be open.

Thus, we ask:

Open problem III: optimal Gibbs-state certification

Our Gibbs-state certification algorithm is efficient in both sample and time complexity.

However, we do not know whether the sample complexity is optimal.

Even in the classical setting, the precise complexity of testing local Gibbs distributions appears to be open.

Thus, we ask:

What is the optimal sample complexity of certifying local Gibbs states?

Open problem IV: certification from local probes

Our dynamics-access protocol assumes global access to the time evolution.

Open problem IV: certification from local probes

Our dynamics-access protocol assumes global access to the time evolution.

But in many experimental settings, one may only observe a small part of the system.

Open problem IV: certification from local probes

Our dynamics-access protocol assumes global access to the time evolution.

But in many experimental settings, one may only observe a small part of the system.

Recent progress in Hamiltonian learning suggests that even single-site or constant-size probes can sometimes yield meaningful information.

Open problem IV: certification from local probes

Our dynamics-access protocol assumes global access to the time evolution.

But in many experimental settings, one may only observe a small part of the system.

Recent progress in Hamiltonian learning suggests that even single-site or constant-size probes can sometimes yield meaningful information.

This motivates a certification analogue.

Local probes: main questions

Can one efficiently certify a k-local Hamiltonian using only dynamics observed through a strictly local probe?

Local probes: main questions

Can one efficiently certify a k-local Hamiltonian using only dynamics observed through a strictly local probe?

More quantitatively:

Local probes: main questions

Can one efficiently certify a k-local Hamiltonian using only dynamics observed through a strictly local probe?

More quantitatively:

What is the tradeoff between probe size and certification complexity?

Local probes: main questions

Can one efficiently certify a k-local Hamiltonian using only dynamics observed through a strictly local probe?

More quantitatively:

What is the tradeoff between probe size and certification complexity?

For example, a constant-size probe may or may not suffice to achieve optimal

$$\Theta(1/\varepsilon)$$

evolution time.

Why local probes matter

Understanding certification from local probes would bring the theory closer to realistic experiments.

Why local probes matter

Understanding certification from local probes would bring the theory closer to realistic experiments.

In practice, one often has limited control and limited measurement access.

Why local probes matter

Understanding certification from local probes would bring the theory closer to realistic experiments.

In practice, one often has limited control and limited measurement access.

The key question is whether local observations can still detect the spectral features distinguishing

$$H \quad \text{from} \quad H_0.$$

Why local probes matter

Understanding certification from local probes would bring the theory closer to realistic experiments.

In practice, one often has limited control and limited measurement access.

The key question is whether local observations can still detect the spectral features distinguishing

$$H \quad \text{from} \quad H_0.$$

This remains widely open.

Open problem V: optimal sparsity testing

Our certification algorithm relies on spectral gap arguments for local Hamiltonians.

Open problem V: optimal sparsity testing

Our certification algorithm relies on spectral gap arguments for local Hamiltonians.

Similar ideas also appear in recent work on learning and testing sparse Hamiltonians.

Open problem V: optimal sparsity testing

Our certification algorithm relies on spectral gap arguments for local Hamiltonians.

Similar ideas also appear in recent work on learning and testing sparse Hamiltonians.

This suggests another question:

Open problem V: optimal sparsity testing

Our certification algorithm relies on spectral gap arguments for local Hamiltonians.

Similar ideas also appear in recent work on learning and testing sparse Hamiltonians.

This suggests another question:

Can these techniques lead to optimal algorithms for Hamiltonian sparsity testing?

Open problem V: optimal sparsity testing

Our certification algorithm relies on spectral gap arguments for local Hamiltonians.

Similar ideas also appear in recent work on learning and testing sparse Hamiltonians.

This suggests another question:

Can these techniques lead to optimal algorithms for Hamiltonian sparsity testing?

One particularly interesting setting is sparsity testing under an additional promise of locality.

Open problem VI: beyond Hamiltonians

Our analysis focuses on closed-system dynamics generated by time-independent Hamiltonians.

Open problem VI: beyond Hamiltonians

Our analysis focuses on closed-system dynamics generated by time-independent Hamiltonians.

Many realistic quantum devices are better modeled by more general dynamical maps.

Open problem VI: beyond Hamiltonians

Our analysis focuses on closed-system dynamics generated by time-independent Hamiltonians.

Many realistic quantum devices are better modeled by more general dynamical maps.

Examples include:

- noisy quantum channels;

Open problem VI: beyond Hamiltonians

Our analysis focuses on closed-system dynamics generated by time-independent Hamiltonians.

Many realistic quantum devices are better modeled by more general dynamical maps.

Examples include:

- noisy quantum channels;
- Pauli channels induced by noise processes;

Open problem VI: beyond Hamiltonians

Our analysis focuses on closed-system dynamics generated by time-independent Hamiltonians.

Many realistic quantum devices are better modeled by more general dynamical maps.

Examples include:

- noisy quantum channels;
- Pauli channels induced by noise processes;
- Lindbladian evolutions for open quantum systems.

Open problem VI: beyond Hamiltonians

Our analysis focuses on closed-system dynamics generated by time-independent Hamiltonians.

Many realistic quantum devices are better modeled by more general dynamical maps.

Examples include:

- noisy quantum channels;
- Pauli channels induced by noise processes;
- Lindbladian evolutions for open quantum systems.

A natural direction is to extend Hamiltonian certification techniques to these settings.

Broader direction

The broader goal is to develop a theory of certification for realistic quantum dynamics.

Broader direction

The broader goal is to develop a theory of certification for realistic quantum dynamics.

Such a theory should handle:

- locality and sparsity;

Broader direction

The broader goal is to develop a theory of certification for realistic quantum dynamics.

Such a theory should handle:

- locality and sparsity;
- finite-temperature states;

Broader direction

The broader goal is to develop a theory of certification for realistic quantum dynamics.

Such a theory should handle:

- locality and sparsity;
- finite-temperature states;
- limited control and local probes;

Broader direction

The broader goal is to develop a theory of certification for realistic quantum dynamics.

Such a theory should handle:

- locality and sparsity;
- finite-temperature states;
- limited control and local probes;
- noise and open-system dynamics.

Broader direction

The broader goal is to develop a theory of certification for realistic quantum dynamics.

Such a theory should handle:

- locality and sparsity;
- finite-temperature states;
- limited control and local probes;
- noise and open-system dynamics.

This would help bridge the gap between theoretically optimal protocols and experimentally available access models.

Takeaway

The main takeaway is:

Takeaway

The main takeaway is:

Certification can be substantially easier than learning.

Takeaway

The main takeaway is:

Certification can be substantially easier than learning.

For local Hamiltonians, this leads to optimal Heisenberg-limited certification.

Takeaway

The main takeaway is:

Certification can be substantially easier than learning.

For local Hamiltonians, this leads to optimal Heisenberg-limited certification.

For Gibbs states, it leads to efficient algorithms beyond what Hamiltonian learning alone would imply.

Takeaway

The main takeaway is:

Certification can be substantially easier than learning.

For local Hamiltonians, this leads to optimal Heisenberg-limited certification.

For Gibbs states, it leads to efficient algorithms beyond what Hamiltonian learning alone would imply.

The remaining challenge is to understand how far this separation extends.

Thank you!

junseolee@fas.harvard.edu