Title: Learning in the quantum universe

Speakers: Hsin-Yuan Huang

Series: Colloquium

Date: November 23, 2022 - 2:00 PM

URL: https://pirsa.org/22110052

Abstract: I will present recent progress in building a rigorous theory to understand how scientists, machines, and future quantum computers could learn models of our quantum universe. The talk will begin with an experimentally feasible procedure for converting a quantum many-body system into a succinct classical description of the system, its classical shadow. Classical shadows can be applied to efficiently predict many properties of interest, including expectation values of local observables and few-body correlation functions. I will then build on the classical shadow formalism to answer two fundamental questions at the intersection of machine learning and quantum physics: Can classical machines learn to solve challenging problems in quantum physics? And can quantum machines learn exponentially faster than classical machines?

Zoom link: https://pitp.zoom.us/j/97994359596?pwd=UlBwc2hoSkNzWlZvM1o1RWErU1U2QT09

Learning in the quantum universe

Presenter: Hsin-Yuan Huang (Robert)

Collaborators: Richard Kueng, Giacomo Torlai, Victor Albert, John Preskill,
Sitan Chen, Jordan Cotler, Jerry Li, Michael Broughton, Jarrod McClean, and more

Motivation

- A central goal of science is to learn how our universe operates.

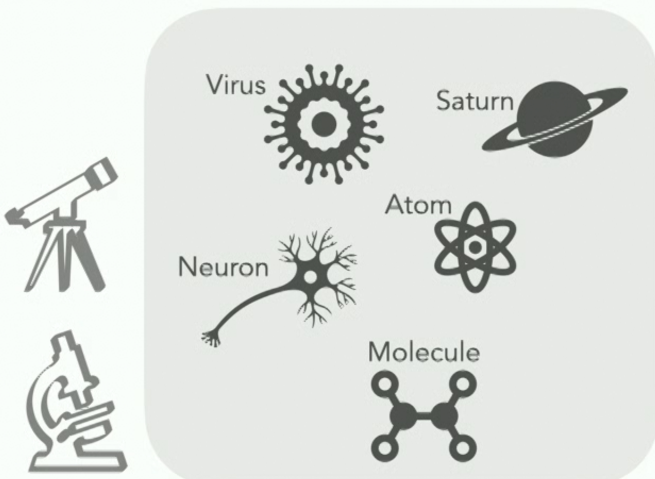

Examples of scientific disciplines

Motivation

- A central goal of science is to learn how our universe operates.
- Because our universe is **inherently quantum**, the ability to efficiently learn in the quantum world could lead to many advances.

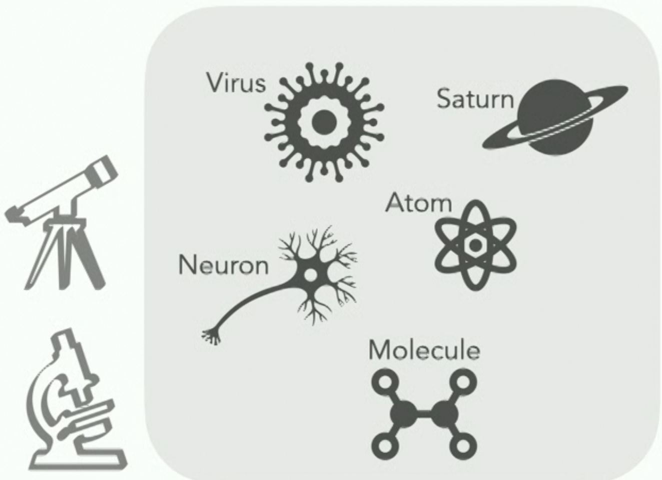

Examples of scientific disciplines

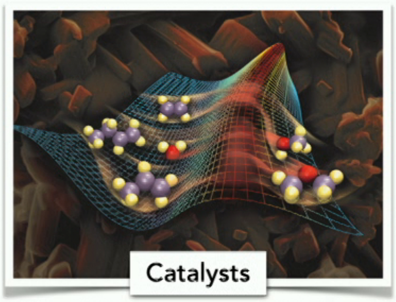

Image credits: (Top left) https://www.energy.gov/science/doe-explainscatalysts (Top right) https://theconversation.com/as-pharmaceutical-use-continues-to-rise-side-effects-are-becoming-a-costly-health-issue-105494 (Bottom left) https://news.mit.edu/2019/ultra-quantum-matter-uqm-research-given-8m-boost-0529 (Bottom right) https://www.nature.com/articles/d41586-019-03213-z

Motivation

- A central goal of science is to learn how our universe operates.
- Because our universe is **inherently quantum**, the ability to efficiently learn in the quantum world could lead to many advances.

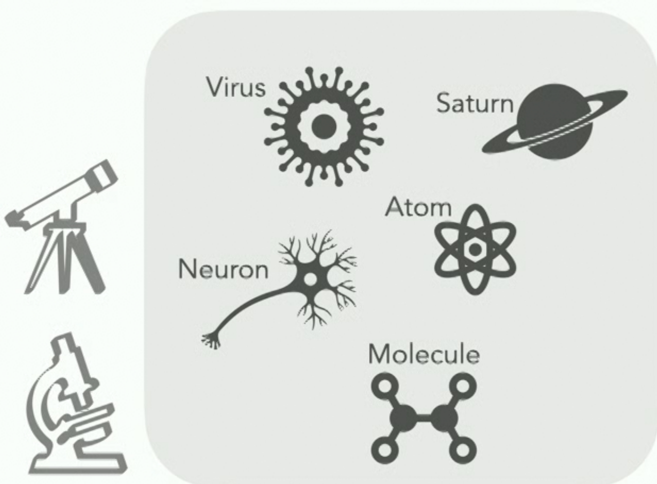

Examples of scientific disciplines

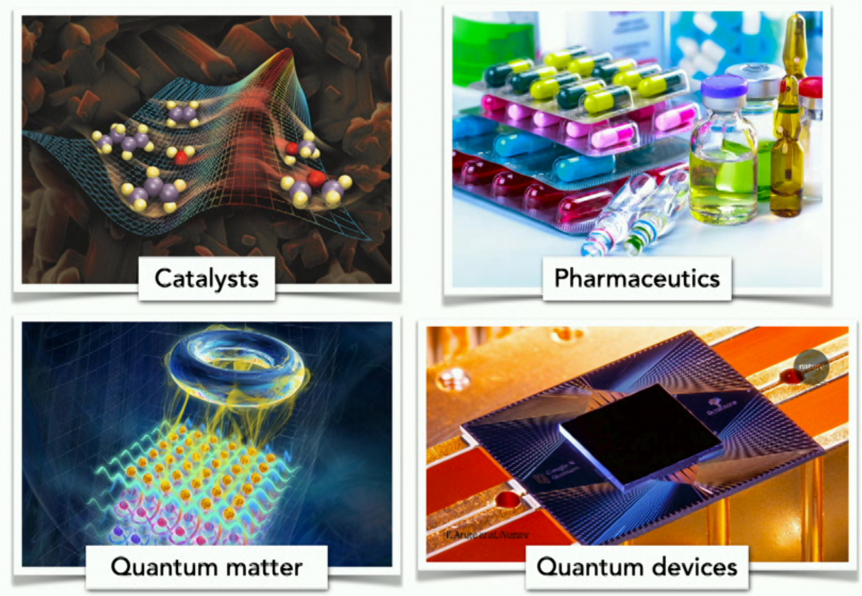

Image credits: (Top left) https://www.energy.gov/science/doe-explainscatalysts (Top right) https://theconversation.com/as-pharmaceutical-use-continues-to-rise-side-effects-are-becoming-a-costly-health-issue-105494 (Bottom left) https://news.mit.edu/2019/ultra-quantum-matter-uqm-research-given-8m-boost-0529 (Bottom right) https://www.nature.com/articles/d41586-019-03213-z

Motivation

- To accelerate and automate the development of (quantum) science, it is important to understand how to design better algorithms to **learn in the quantum universe**.

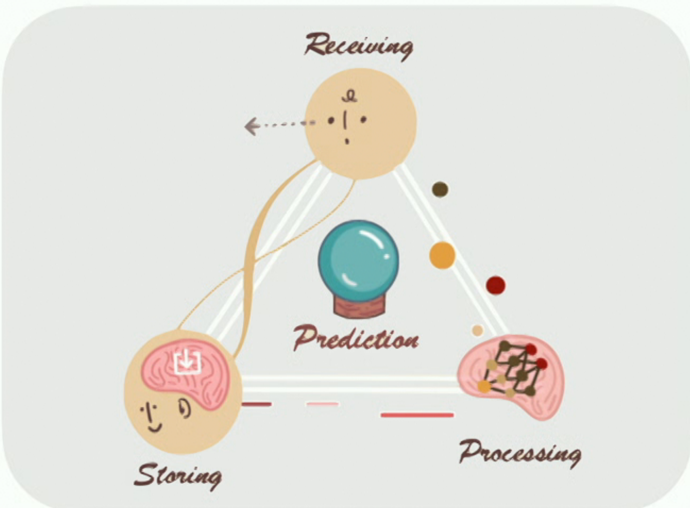

A cartoon depiction of learning

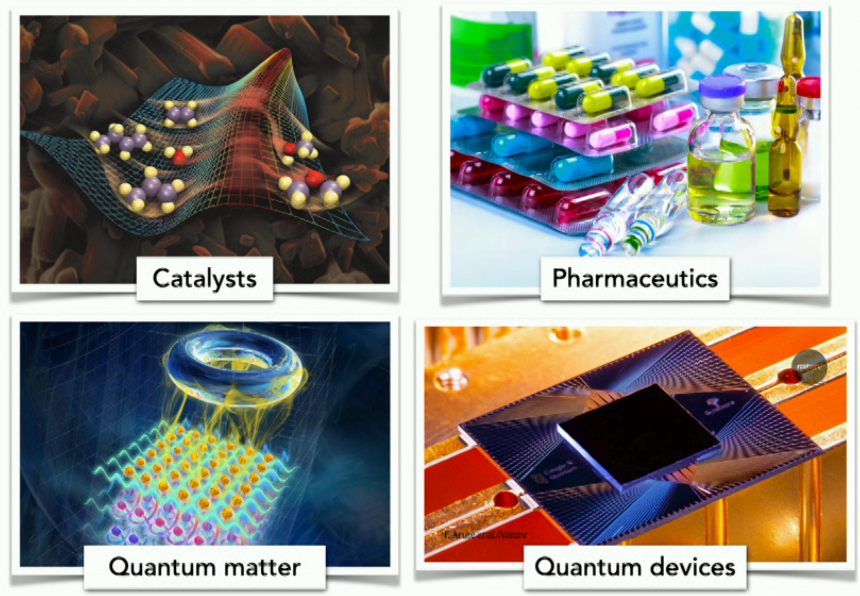

Image credits: (Top left) https://www.energy.gov/science/doe-explainscatalysts (Top right) https://theconversation.com/as-pharmaceutical-use-continues-to-rise-side-effects-are-becoming-a-costly-health-issue-105494 (Bottom left) https://news.mit.edu/2019/ultra-quantum-matter-uqm-research-given-8m-boost-0529 (Bottom right) https://www.nature.com/articles/d41586-019-03213-z

Overview

How to efficiently learn in the quantum universe?

Learning with classical machines

What can classical machines learn?
Can classical ML perform
better than non-ML algorithms?

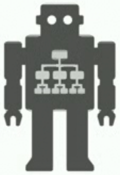

Learning with quantum machines

Can quantum machines learn faster
and/or predict more accurately
than classical machines?

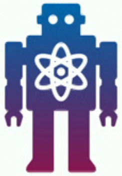

Overview

***Predicting many properties of a quantum system from very few measurements**, Nature Physics*

***Power of data in quantum machine learning**, Nature Communications*

***Provably efficient machine learning for quantum many-body problems**, Science*

Related works:

Efficient estimation of Pauli observables by derandomization, Physical Review Letter
Mixed-state entanglement from local randomized measurements, Physical Review Letter
Emergent randomness and benchmarking from many-body quantum chaos, to appear on Nature

***Information-theoretic bounds on quantum advantage in machine learning**, Physical Review Letter*

***Exponential separations between learning with and without quantum memory**, FOCS*

***Quantum advantage in learning from experiments**, Science*

Related works:

Generalization in quantum machine learning from few training data, Nature Communications
Foundations for learning from noisy quantum experiments, QIP 2022
Learning many-body Hamiltonians with Heisenberg-limited scaling, submitted

Classical shadow formalism

- How can classical machines "see" quantum many-body systems?

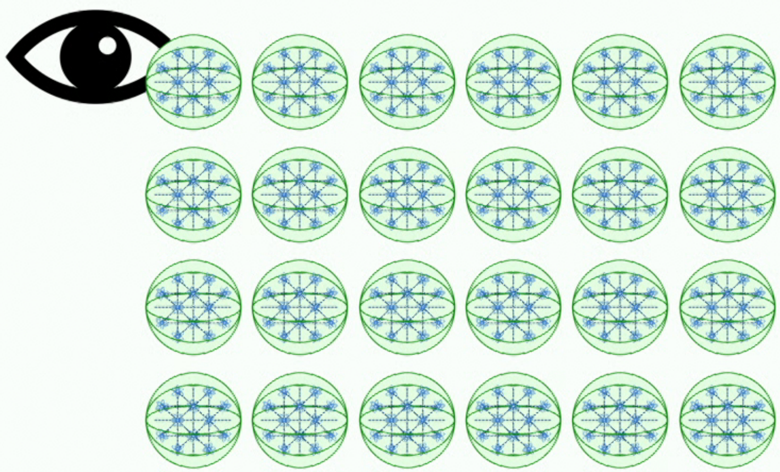

[HKP20] Hsin-Yuan Huang, Richard Kueng, John Preskill. *Predicting many properties of a quantum system from very few measurements*, Nature Physics, 2020.

Classical shadow formalism

- What do we mean by "seeing" a quantum system?
- Converting the quantum system to a classical form that accurately captures many properties of the quantum system.

[HKP20] Hsin-Yuan Huang, Richard Kueng, John Preskill. *Predicting many properties of a quantum system from very few measurements*, Nature Physics, 2020.

Classical shadow formalism

- Why do we want to construct classical representations of quantum systems?
  - ✦ We often want to know what the quantum system is.
  - ✦ Many quantum applications require an interface between the classical and the quantum realm (e.g., variational algorithms).

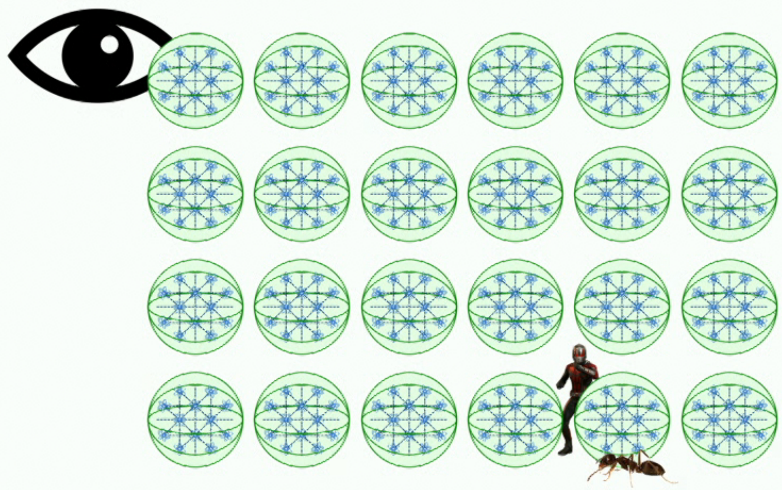

[HKP20] Hsin-Yuan Huang, Richard Kueng, John Preskill. *Predicting many properties of a quantum system from very few measurements*, Nature Physics, 2020.

Standard approach

- **Quantum state tomography:**

  Learn a <u>complete representation</u> of an n-qubit quantum state.
  $(d \times d \text{ matrix}, \; d = 2^n)$

- Sample-optimal protocol (Haah et al.; O'Donnel, Wright):

  - Sample complexity: $\Theta(2^{2n})$
  - Quantum resource: $\Theta(n2^{2n})$ qubits + exponentially long circuits
  - Classical storage: $\Omega(2^{2n})$
  - Classical post-processing: $\Omega(2^{2n})$

Classical shadow formalism

Theorem 1 [HKP20]

There exists procedure that guarantees the following.

1. Given $B, M, \epsilon > 0$, the procedure learns a classical representation of an unknown quantum state ρ from

$$T = \mathcal{O}(B \log(M)/\epsilon^2) \text{ measurements.}$$

2. Subsequently, given any $O_1, \ldots, O_M$ with $B \geq \max \|O_i\|^2_{\text{shadow}}$, the procedure can use the classical representation to predict $\hat{o}_1, \ldots, \hat{o}_M$,

where $|\hat{o}_i - \mathrm{tr}(O_i \rho)| < \epsilon$, for all i.

For example:

- $M = 10^6$, $B = 1$, then naively we need $10^6/\epsilon^2$ measurements.
- This theorem shows that we only need $6 \log(10)/\epsilon^2$ measurements.

Furthermore, we don't need to know $O_1, \ldots, O_M$ in advance.

The Procedure:
Data Acquisition Phase

Repeat the following T times:

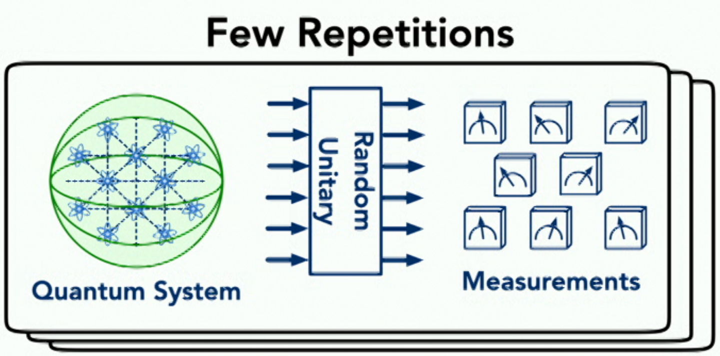

Data Acquisition Phase

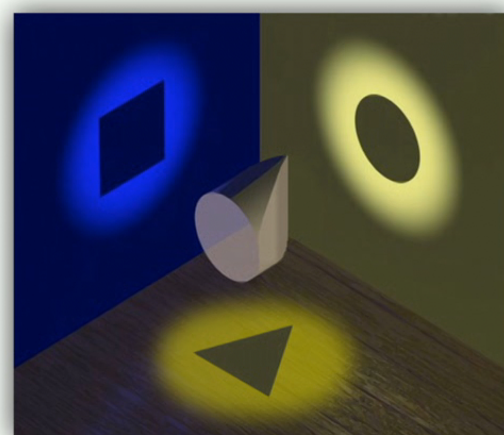

The Procedure:
Data Acquisition Phase

Repeat the following T times:

- Sample a random unitary U_i to rotate the quantum system.

E.g., measure each qubit in a random basis

Few Repetitions

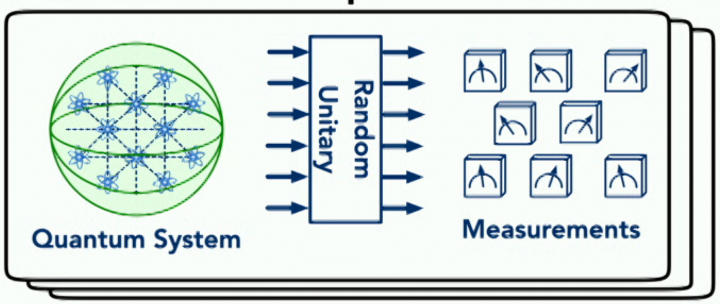

Data Acquisition Phase

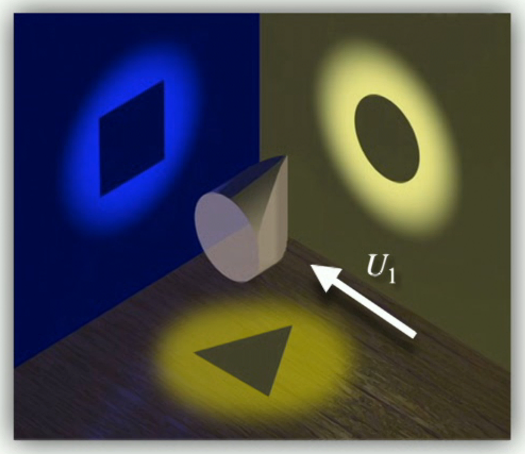

The Procedure: Data Acquisition Phase

Repeat the following T times:

- Sample a random unitary U_i to rotate the quantum system.
- Measure the system in the computational basis $|b_i\rangle \in \{0,1\}^n$.
- Store the "classical shadow": $|s_i\rangle = U_i^\dagger |b_i\rangle$.

E.g., measure each qubit in a random basis

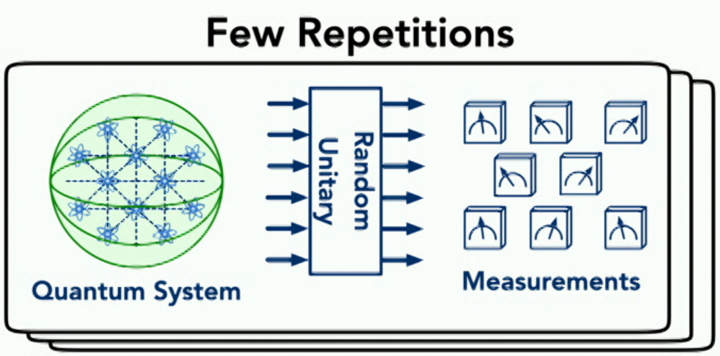

Data Acquisition Phase

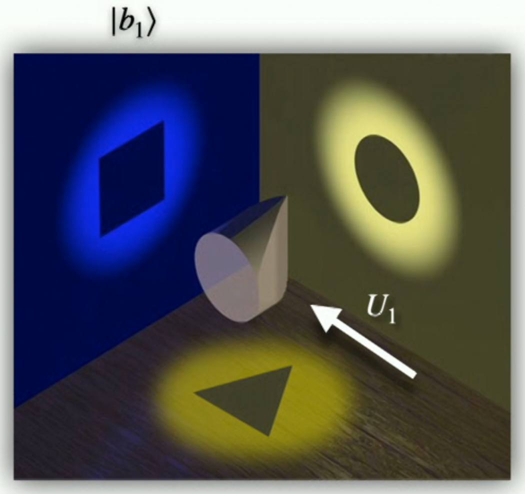

The Procedure: Prediction Phase

Given $S(\rho) = \{|s_1\rangle, \ldots, |s_T\rangle\}$ (the classical shadow),

how to predict properties of the quantum state ρ?

★ $\mathbb{E}[|s_i\rangle\langle s_i|] = \mathcal{M}(\rho)$. ($\mathcal{M}$: some CPTP map)

➡ $\rho = \mathbb{E}[\mathcal{M}^{-1}(|s_i\rangle\langle s_i|)] \implies \rho \approx \mathcal{M}^{-1}(|s_i\rangle\langle s_i|)$.

Algorithm for predicting $\mathrm{tr}(O\rho)$: (median-of-means)
Compute $X_i = \mathrm{tr}(O\mathcal{M}^{-1}(|s_i\rangle\langle s_i|)), \forall i = 1, \ldots, T$.
Predict $\hat{o} = \mathrm{median}\left(\frac{1}{T/K}\sum_{i=1}^{T/K} X_i, \ldots, \frac{1}{T/K}\sum_{i=T-T/K+1}^{N} X_i\right)$.

The Procedure: Data Acquisition Phase

Repeat the following T times:

- Sample a random unitary U_i to rotate the quantum system.
- Measure the system in the computational basis $|b_i\rangle \in \{0,1\}^n$.
- Store the "classical shadow": $|s_i\rangle = U_i^\dagger |b_i\rangle$.

E.g., measure each qubit in a random basis

Few Repetitions

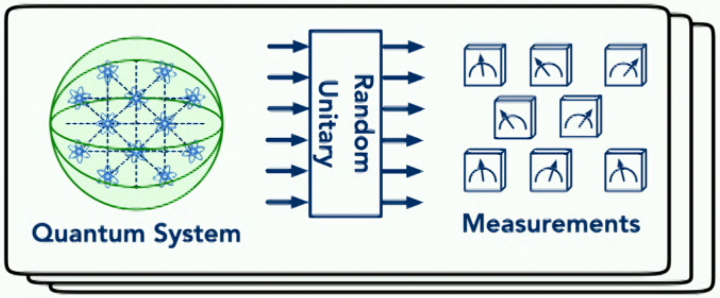

Data Acquisition Phase

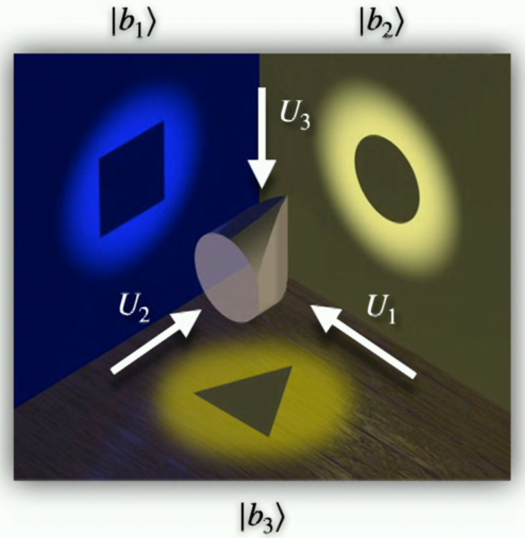

The Procedure: Prediction Phase

Given $S(\rho) = \{|s_1\rangle, \ldots, |s_T\rangle\}$ (the classical shadow),

how to predict properties of the quantum state ρ?

★ $\mathbb{E}[|s_i\rangle\langle s_i|] = \mathcal{M}(\rho)$. ($\mathcal{M}$: some CPTP map)

➡ $\rho = \mathbb{E}[\mathcal{M}^{-1}(|s_i\rangle\langle s_i|)] \implies \rho \approx \mathcal{M}^{-1}(|s_i\rangle\langle s_i|)$.

Algorithm for predicting $\mathrm{tr}(O\rho)$: (median-of-means)

Compute $X_i = \mathrm{tr}(O\mathcal{M}^{-1}(|s_i\rangle\langle s_i|)), \forall i = 1, \ldots, T$.

Predict $\hat{o} = \mathrm{median}\left(\frac{1}{T/K}\sum_{i=1}^{T/K} X_i, \ldots, \frac{1}{T/K}\sum_{i=T-T/K+1}^{N} X_i\right)$.

Classical shadow formalism

Theorem 1 [HKP20]

There exists procedure that guarantees the following.

1. Given $B, M, \epsilon > 0$, the procedure learns a classical representation of an unknown quantum state ρ from

$$T = \mathcal{O}(B \log(M)/\epsilon^2) \text{ measurements.}$$

2. Subsequently, given any $O_1, \ldots, O_M$ with $B \geq \max \|O_i\|^2_{\text{shadow}}$, the procedure can use the classical representation to predict $\hat{o}_1, \ldots, \hat{o}_M$, where $|\hat{o}_i - \text{tr}(O_i \rho)| < \epsilon$, for all i.

For example:

- $M = 10^6$, $B = 1$, then naively we need $10^6/\epsilon^2$ measurements.
- This theorem shows that we only need $6 \log(10)/\epsilon^2$ measurements.

Furthermore, we don't need to know $O_1, \ldots, O_M$ in advance.

Classical shadow formalism

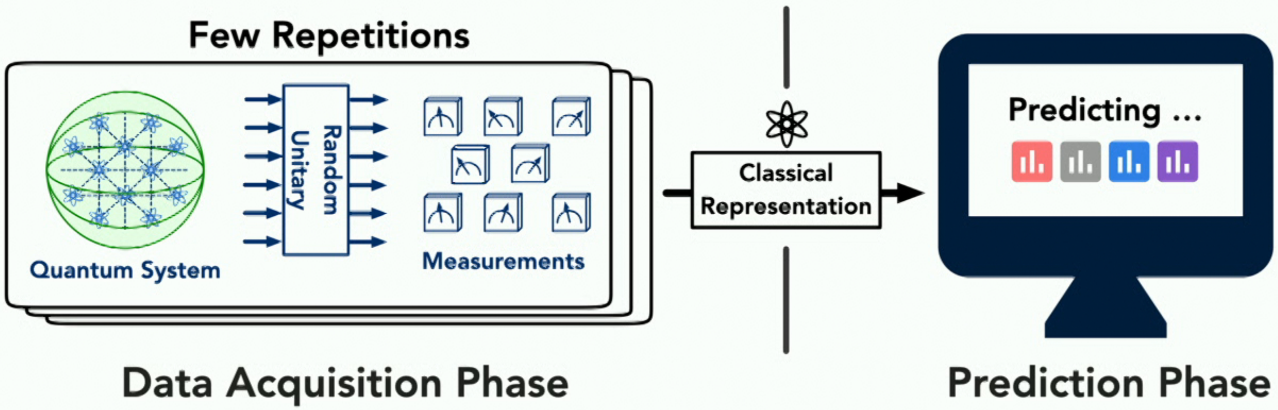

Possible Properties

- Quantum Fidelity
- Entanglement Witness
- Entanglement Entropy
- 2-point Correlations
- Hamiltonian
- Local Observables

[HKP20] Hsin-Yuan Huang, Richard Kueng, John Preskill. *Predicting many properties of a quantum system from very few measurements*, Nature Physics, 2020.

Overview

***Predicting many properties of a quantum system from very few measurements**, Nature Physics*

***Power of data in quantum machine learning**, Nature Communications*

***Provably efficient machine learning for quantum many-body problems**, Science*

Related works:

Efficient estimation of Pauli observables by derandomization, Physical Review Letter
Mixed-state entanglement from local randomized measurements, Physical Review Letter
Emergent randomness and benchmarking from many-body quantum chaos, to appear on Nature

***Information-theoretic bounds on quantum advantage in machine learning**, Physical Review Letter*

***Exponential separations between learning with and without quantum memory**, FOCS*

***Quantum advantage in learning from experiments**, Science*

Related works:

Generalization in quantum machine learning from few training data, Nature Communications
Foundations for learning from noisy quantum experiments, QIP 2022
Learning many-body Hamiltonians with Heisenberg-limited scaling, submitted

Overview

***Predicting many properties of a quantum system from very few measurements**, Nature Physics*

***Power of data in quantum machine learning**, Nature Communications*

***Provably efficient machine learning for quantum many-body problems**, Science*

Related works:

- *Efficient estimation of Pauli observables by derandomization, Physical Review Letter*
- *Mixed-state entanglement from local randomized measurements, Physical Review Letter*
- *Emergent randomness and benchmarking from many-body quantum chaos, to appear on Nature*

***Information-theoretic bounds on quantum advantage in machine learning**, Physical Review Letter*

***Exponential separations between learning with and without quantum memory**, FOCS*

***Quantum advantage in learning from experiments**, Science*

Related works:

- *Generalization in quantum machine learning from few training data, Nature Communications*
- *Foundations for learning from noisy quantum experiments, QIP 2022*
- *Learning many-body Hamiltonians with Heisenberg-limited scaling, submitted*

Classical ML for quantum problems

- Can classical machines learn to solve challenging problems in quantum physics?
- And can they yield better solutions than non-ML algorithms?

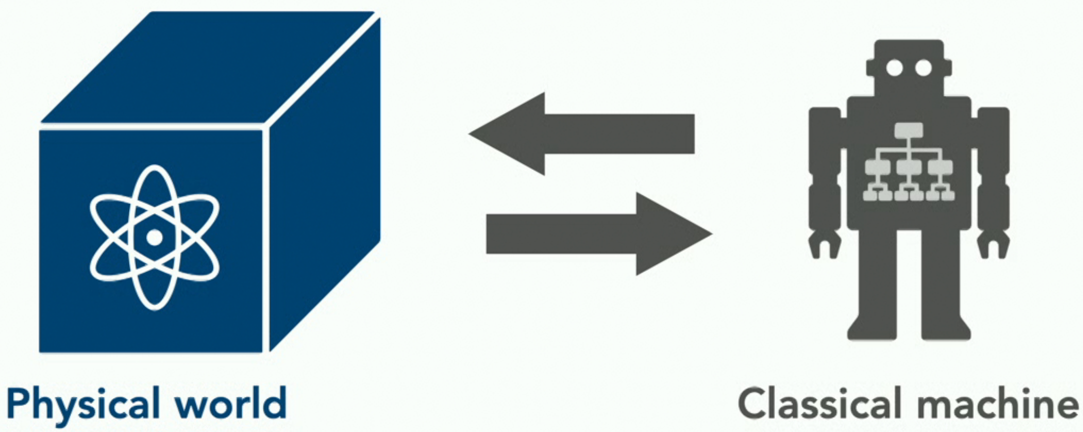

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Task

- Given $x \in [-1,1]^m$ that describes an n-qubit Hamiltonian $H(x)$, the machine predicts a classical representation (e.g., classical shadow) of the ground state $\rho(x)$ of $H(x)$.

- Vector x specifies laser intensities, few-body interactions, magnetic fields, etc.

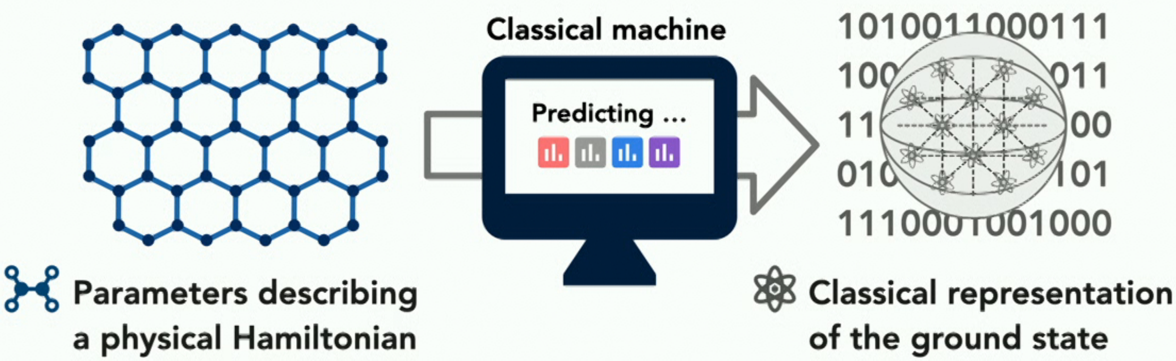

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Computational hardness

- This problem is *extremely* hard!
- Consider a smooth class of n-qubit **2D** Hamiltonians $H(x)$ with **spectral gap 1**, and the machine only predicts **1-body** observable O in ground state $\rho(x)$.
- Furthermore, we only care about average prediction error.

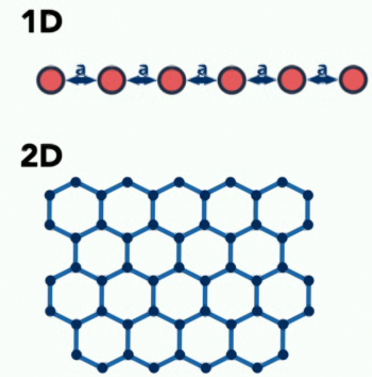

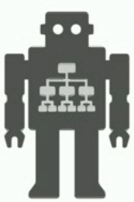

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Computational hardness

- This problem is *extremely* hard!
- Consider a smooth class of n-qubit **2D** Hamiltonians $H(x)$ with **spectral gap 1**, and the machine only predicts **1-body** observable O in ground state $\rho(x)$.
- Furthermore, we only care about average prediction error.

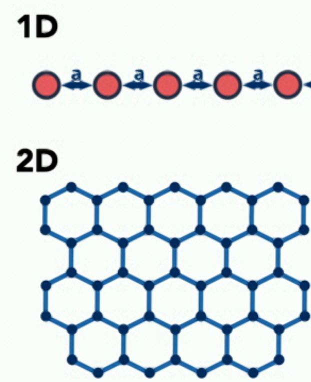

Proposition 1

Assuming RP $\neq$ NP, then no randomized classical algorithm can achieve an average prediction error $\leq 1/4$ within $\text{poly}(n)$ time.

RP $\neq$ NP: NP-complete problems cannot be solved in randomized polynomial time.

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Task

- Can classical ML algorithms do something useful for this challenging problem?

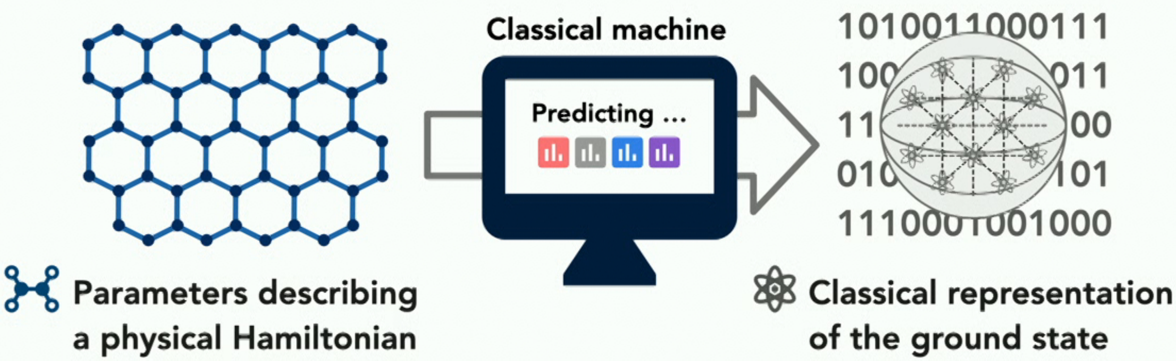

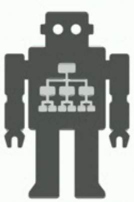

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Task

Training data: $\{x_\ell \to \sigma_T(\rho(x_\ell))\}_{\ell=1}^{N}$

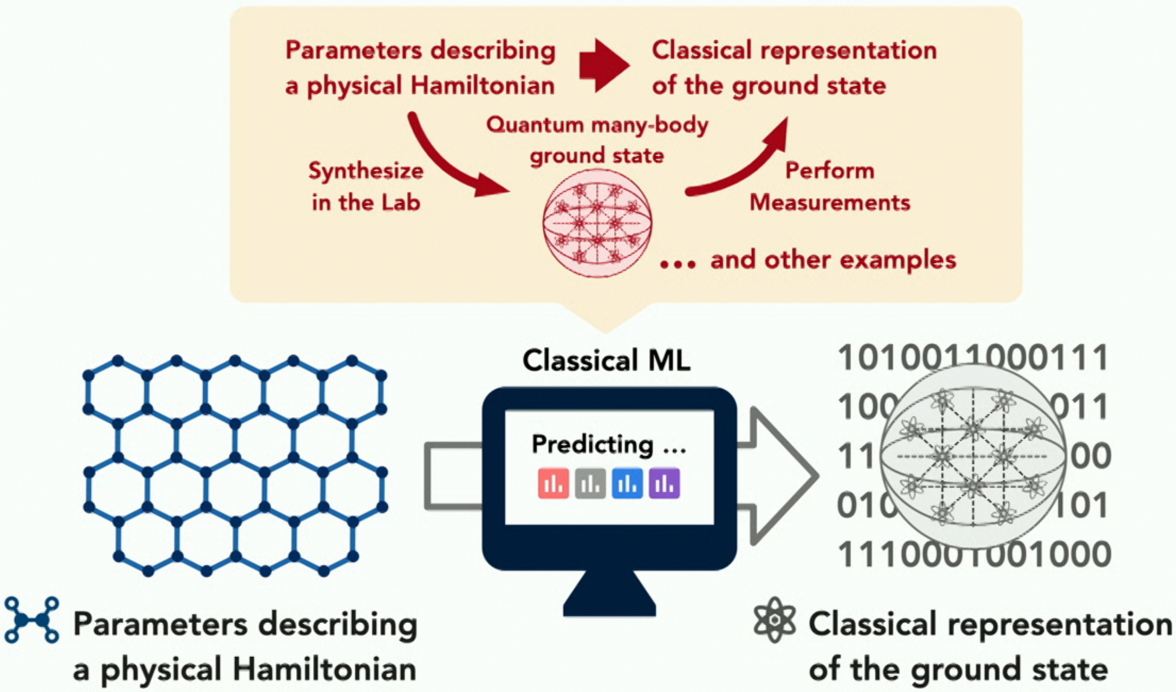

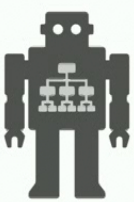

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Task

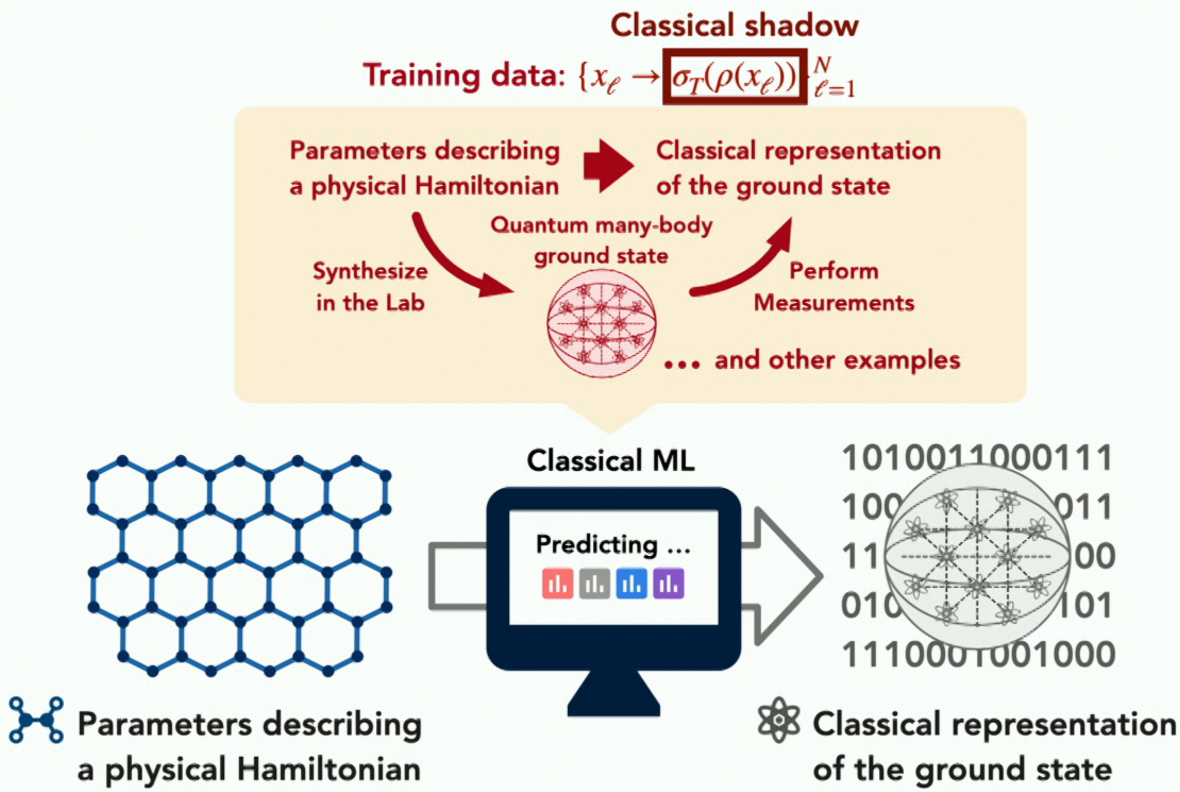

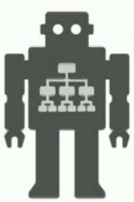

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Proposition 1

Assuming RP $\neq$ NP, then no randomized classical algorithm can achieve an average prediction error $\leq 1/4$ within $\text{poly}(n)$ time.

Classical algorithm

2D

spectral gap 1

1-body observable

average prediction error 1/4

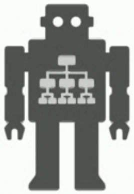

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Theorem 1

A classical ML algorithm can achieve an average prediction error $\leq \epsilon$ using $\mathrm{poly}(n)$ training data and computational time.

Classical algorithm

2D
spectral gap 1
1-body observable
average prediction error 1/4

Classical ML (trained with data)

any constant dimension
any constant spectral gap
any local observable
any average prediction error $\epsilon = \mathcal{O}(1)$

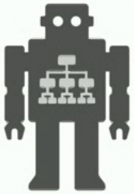

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Theorem 1

A classical ML algorithm can achieve an average prediction error $\leq \epsilon$ using $\mathrm{poly}(n)$ training data and computational time.

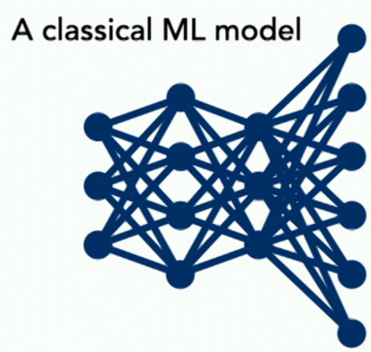

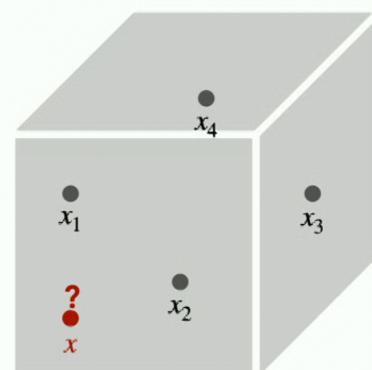

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Classical algorithm

2D
spectral gap 1
1-body observable
average prediction error 1/4

Classical ML (trained with data)

any constant dimension
any constant spectral gap
any local observable
any average prediction error $\epsilon = \mathcal{O}(1)$

We proved that a poly-time classical ML algorithm (w/ data) can predict **much better** than **any** poly-time classical algorithm.

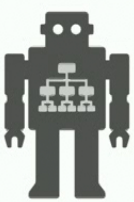

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

The question:
Why ML can be more useful than non-ML algorithms?

The answer:
Generalizing from data can be easier than computing everything

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Data contain computational power
(e.g., nature operates quantumly)

The question:
Why ML can be more useful than non-ML algorithms?

The answer:
Generalizing from data can be easier than computing everything

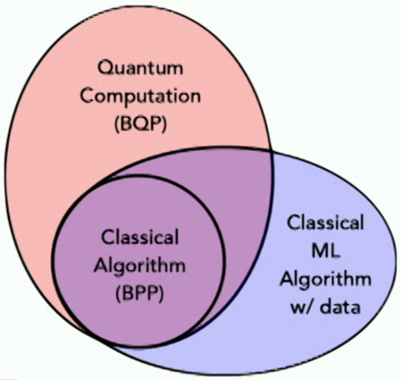

[HBM+21] Huang, Broughton, et al. Power of data in quantum machine learning. Nature Communications, 2021.
[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

2D random Heisenberg model

We consider training data size $N = 100$, $T = 500$ randomized measurements for constructing classical shadows.
The best ML model is chosen from Gaussian kernel method, infinite-width neural networks, and l_2-Dirichlet kernel.

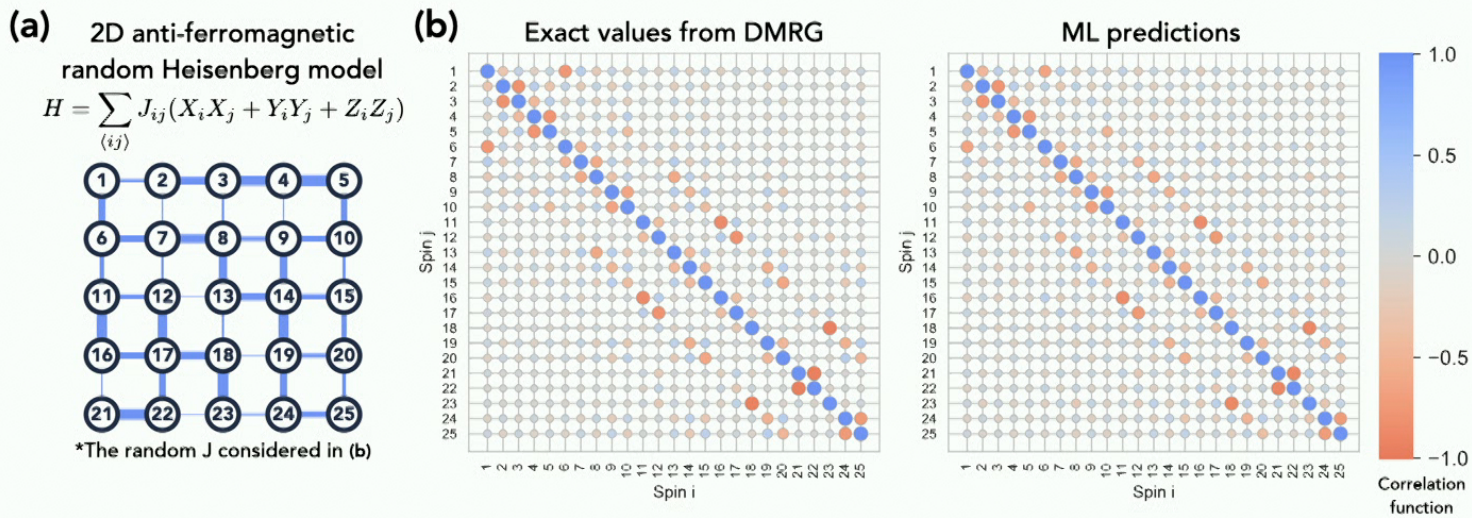

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Classifying quantum phases: Theorem

Theorem 2

If there exists a **nonlinear** function of **few-body** reduced density matrices for classifying the phases, then the classical ML algorithm can **efficiently** learn to classify these phases.

- The assumption is believed to hold for gapped quantum systems.

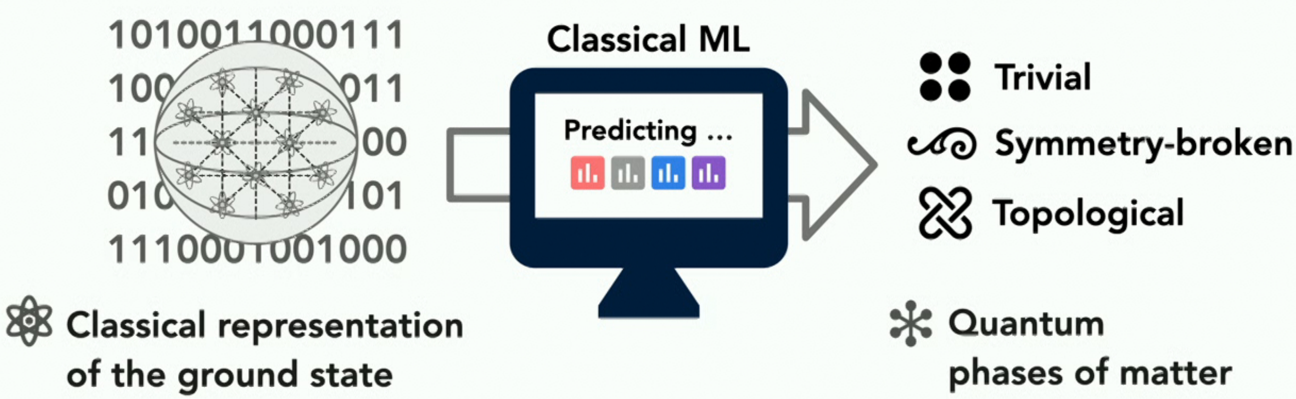

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

1D Symmetry protected topological phases

We consider $T = 500$ randomized measurements to construct classical shadows for each state.
The classical **unsupervised** ML model is a kernel PCA using the shadow kernel.

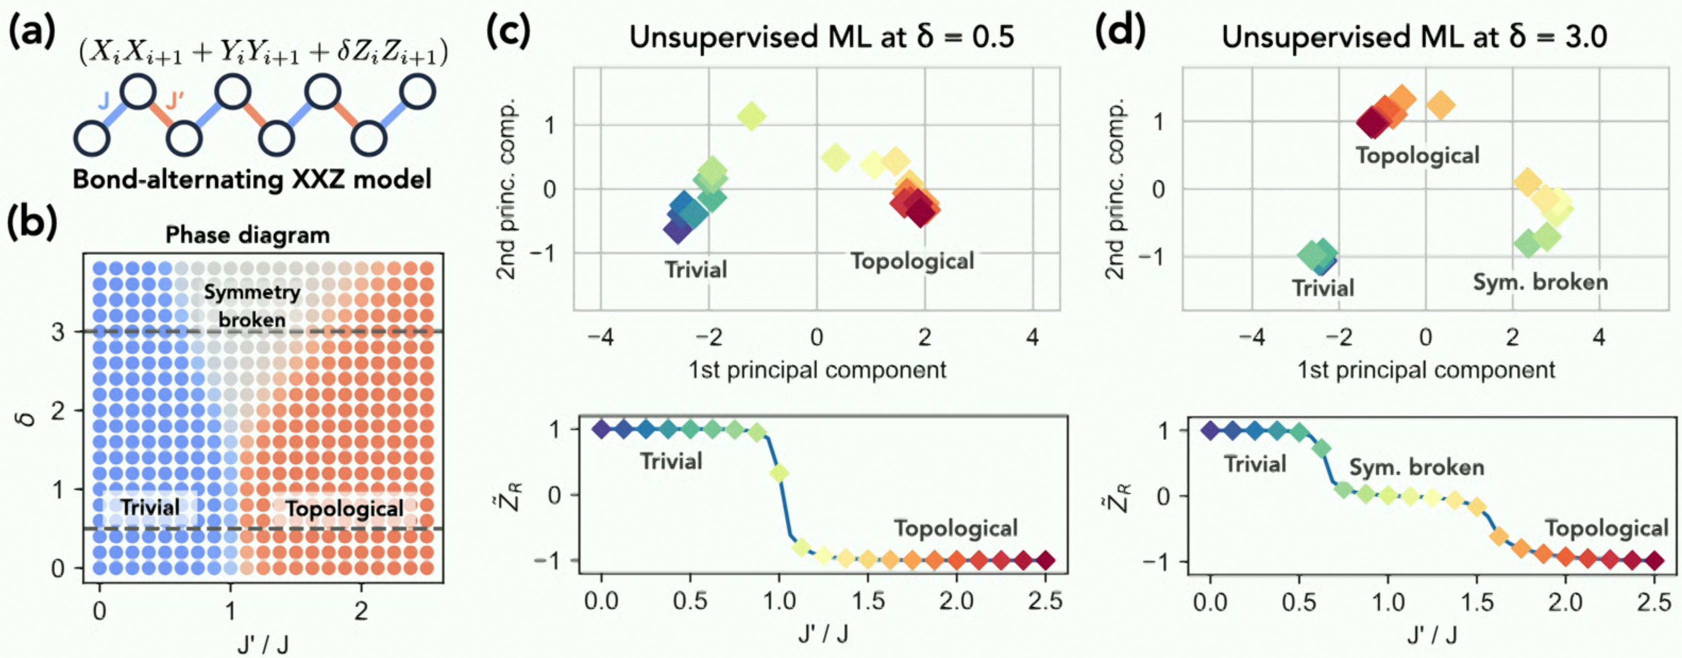

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Overview

How to efficiently learn in the quantum universe?

Learning with classical machines

What can classical machines learn?
Can classical ML perform
better than non-ML algorithms?

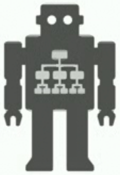

Learning with quantum machines

Can quantum machines learn faster
and/or predict more accurately
than classical machines?

Overview

How to efficiently learn in the quantum universe?

Learning with classical machines

What can classical machines learn?
Can classical ML perform
better than non-ML algorithms?

Learning with quantum machines

Can quantum machines learn faster
and/or predict more accurately
than classical machines?

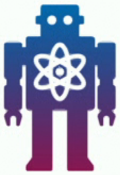

Classical agent

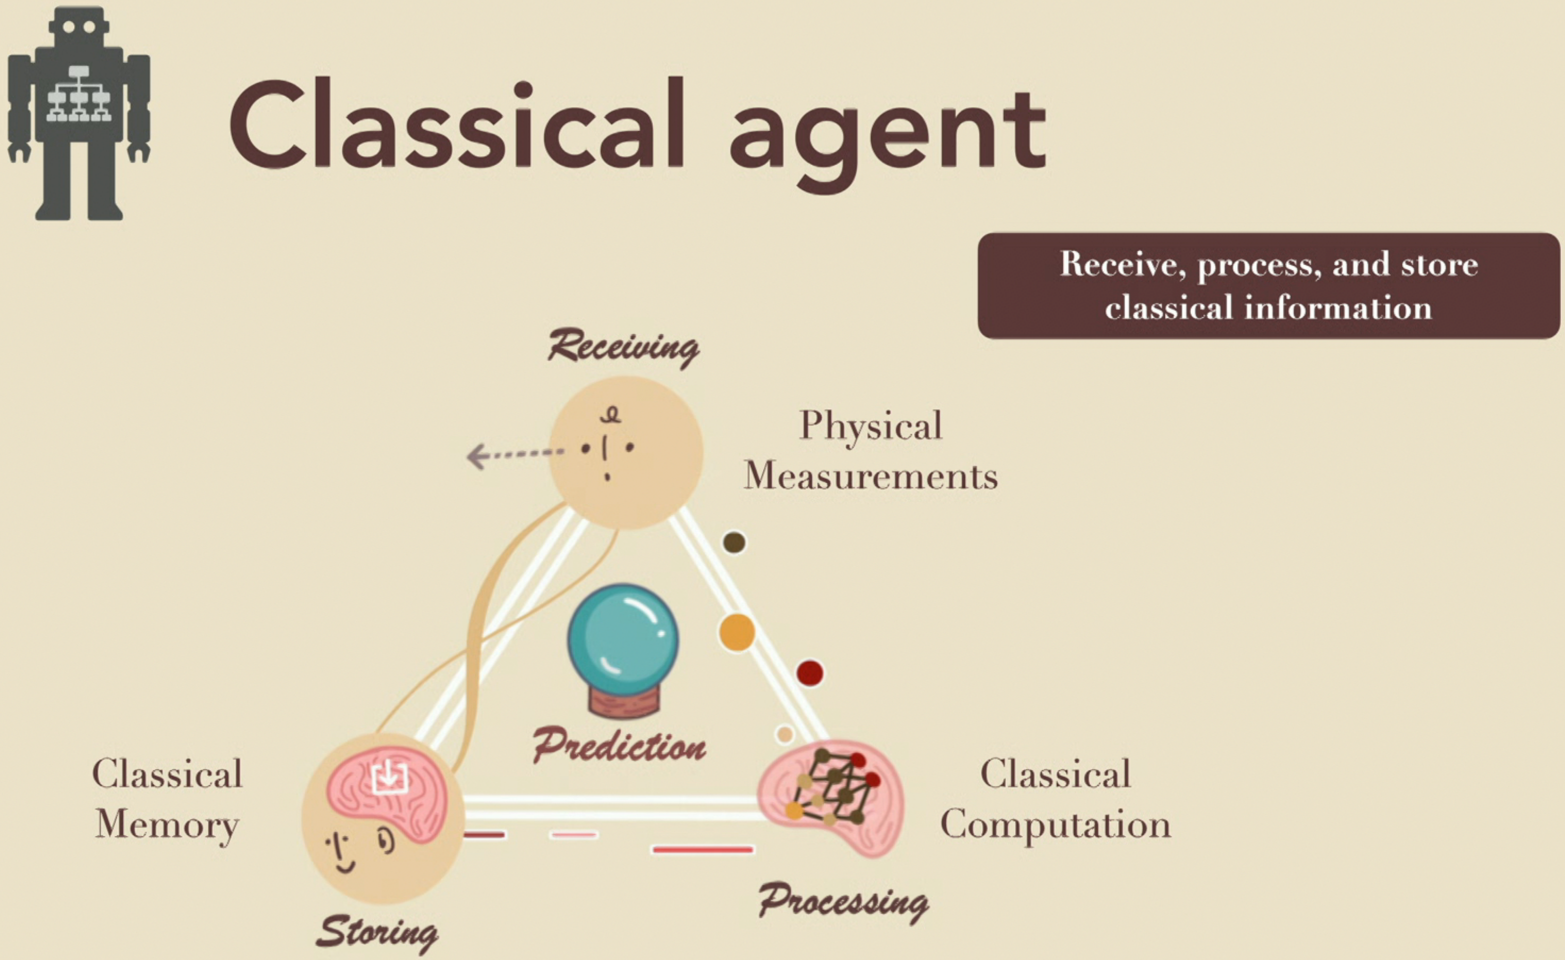

[HKP21] Huang, Kueng, Preskill. Information-theoretic bounds on quantum advantage in machine learning, *Physical Review Letters*, 2021.
[CCHL21] Chen, Cotler, Huang, Li. Exponential separations in learning with and without quantum memory, *FOCS*, 2021.
[HBC+] Huang, Broughton, Cotler, Chen, Li, Mohseni, Neven, Babbush, Kueng, Preskill, McClean. Quantum advantage in learning from experiments, *Science*, 2022.

Classical vs Quantum

- What are the advantage of a quantum agent over a classical agent?
- Could quantum technology fundamentally alter our ability to learn about the physical world?

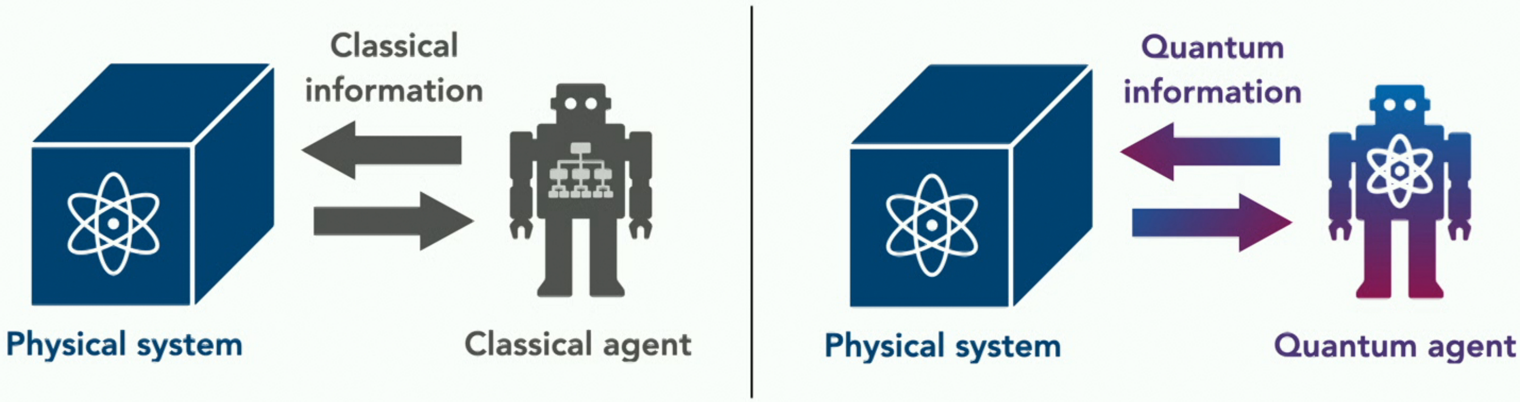

Learning a state

- Assume the only unknown in the entire universe is an n-qubit state ρ.
- Classical agent can perform any measurement on ρ in each experiment.
- Quantum agent can obtain and store ρ coherently from each experiment.

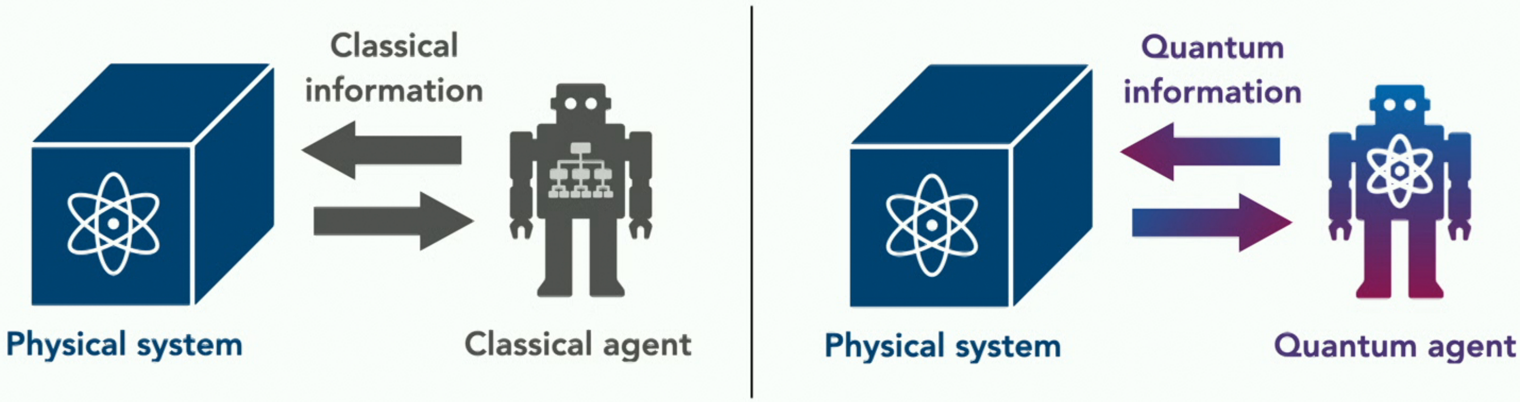

Quantum advantage in predicting Pauli observables

- The classical/quantum agent learns about the unknown n-qubit state ρ.
- Subsequently, the agent predicts $\mathrm{Tr}(P\rho)$ for any observable $P \in \{I, X, Y, Z\}^{\otimes n}$.

Theorem

Classical agent needs $\Omega(2^n)$ experiments to predict an adversarially chosen P, but quantum agent only needs $\mathcal{O}(n)$ experiments to predict all 4^n observables.

Quantum advantage in predicting Pauli observables

- The classical/quantum agent learns about the unknown n-qubit state ρ.
- Subsequently, the agent predicts $\mathrm{Tr}(P\rho)$ for any observable $P \in \{I, X, Y, Z\}^{\otimes n}$.

Theorem

Classical agent needs $\Omega(2^n)$ experiments to predict an adversarially chosen P, but quantum agent only needs $\mathcal{O}(n)$ experiments to predict all 4^n observables.

Exponential quantum advantage is present even when the state ρ is a classical distribution over product states (no entanglement!).

Quantum advantage in predicting Pauli observables

- The classical/quantum agent learns about the unknown n-qubit state ρ.
- Subsequently, the agent predicts $\mathrm{Tr}(P\rho)$ for any observable $P \in \{I, X, Y, Z\}^{\otimes n}$.

Theorem

Classical agent needs $\Omega(2^n)$ experiments to predict an adversarially chosen P, but quantum agent only needs $\mathcal{O}(n)$ experiments to predict all 4^n observables.

Uncertainty principle significantly hinders the learning ability of classical agents, but surprisingly not the ability of a quantum agent.

Exponential quantum advantage

Predicting many incompatible observables

To predict non-commuting observables $O_1, O_2, \ldots$,
classical agent suffers from uncertainty principle,
quantum agent does not.

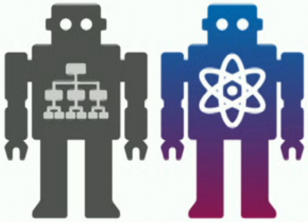

Exponential quantum advantage

Predicting many incompatible observables

To predict non-commuting observables $O_1, O_2, \ldots,$ classical agent suffers from uncertainty principle, quantum agent does not.

Performing quantum PCA

To estimate property of principal component, classical agent needs exponential time, quantum agent needs polynomial.

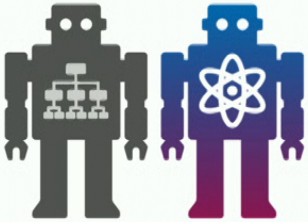

Exponential quantum advantage

Predicting many incompatible observables

To predict non-commuting observables $O_1, O_2, \ldots$, classical agent suffers from uncertainty principle, quantum agent does not.

Performing quantum PCA

To estimate property of principal component, classical agent needs exponential time, quantum agent needs polynomial.

Uncovering symmetry in dynamics

Classifying dynamics with or without time-reversal symmetry

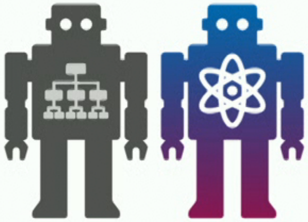

Exponential quantum advantage

Predicting many incompatible observables

To predict non-commuting observables $O_1, O_2, \ldots,$
classical agent suffers from uncertainty principle,
quantum agent does not.

Performing quantum PCA

To estimate property of principal component,
classical agent needs exponential time,
quantum agent needs polynomial.

Uncovering symmetry in dynamics

Classifying dynamics with or without
time-reversal symmetry

Learning physical dynamics

To learn a polynomial-time quantum process,
a classical agent requires exponential experiments,
a quantum agent requires polynomial.

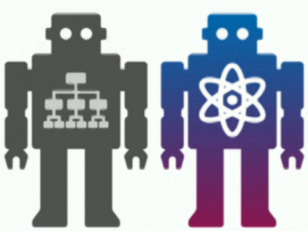

Exponential quantum advantage

Predicting many incompatible observables

Performing quantum PCA

To predict n
classical age
quantum ag

Un

Cla
tim

nponent,
ne,

rocess,
xperiments,

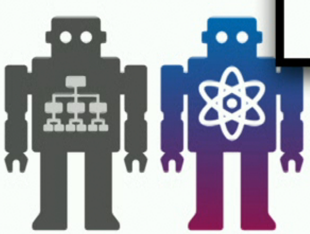

Proven using learning-theoretic techniques

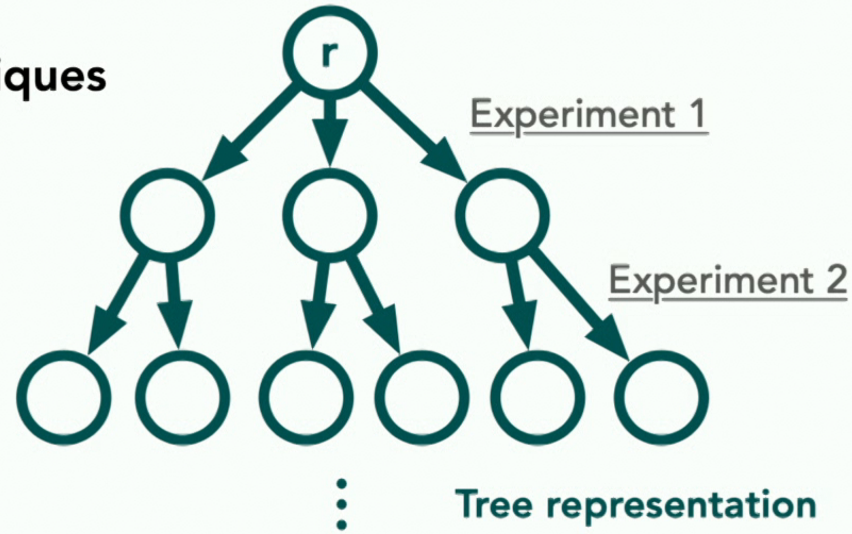

Demonstration on Sycamore: Quantum advantage in learning states

Utilizing a total of 40 qubits

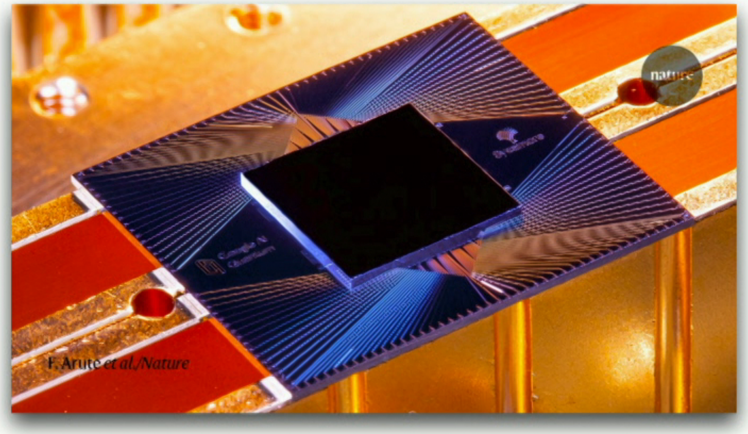

Sycamore Processor

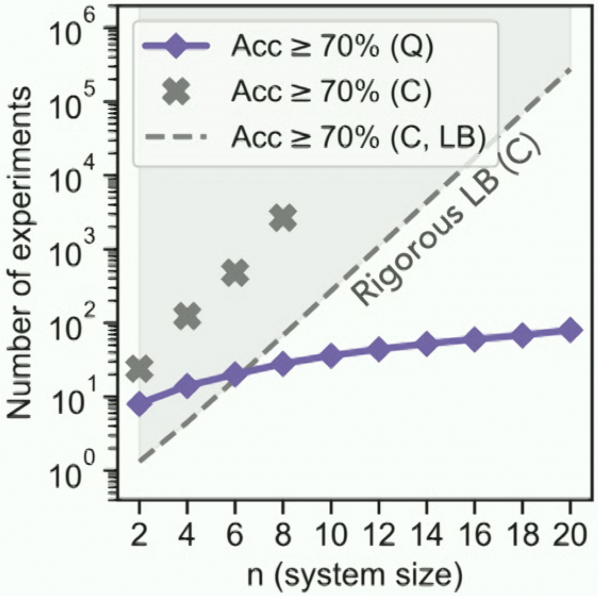

[HFP22] Huang, Flammia, Preskill. Foundations for learning from noisy quantum experiments, *QIP*, 2022.

[HBC+] Huang, Broughton, Cotler, Chen, Li, Mohseni, Neven, Babbush, Kueng, Preskill, McClean. Quantum advantage in learning from experiments, *Science*, 2022.

Demonstration on Sycamore: Quantum advantage in learning states

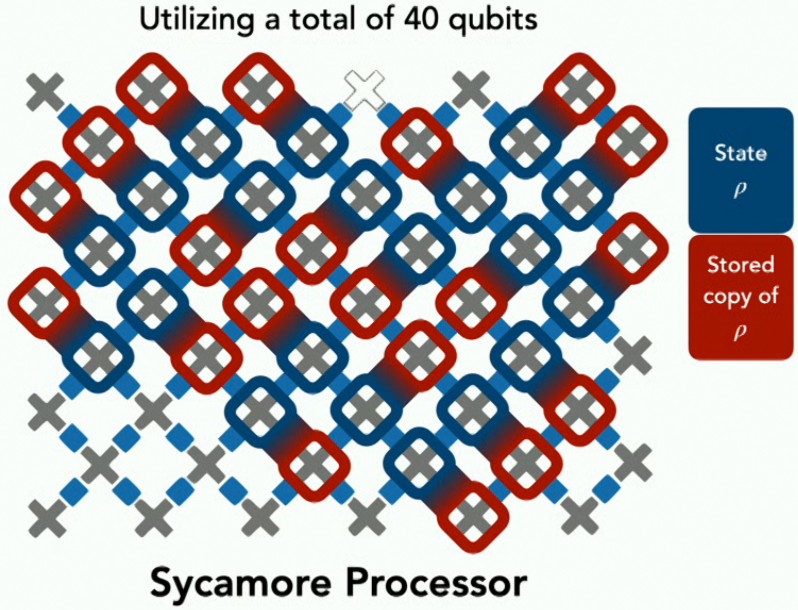

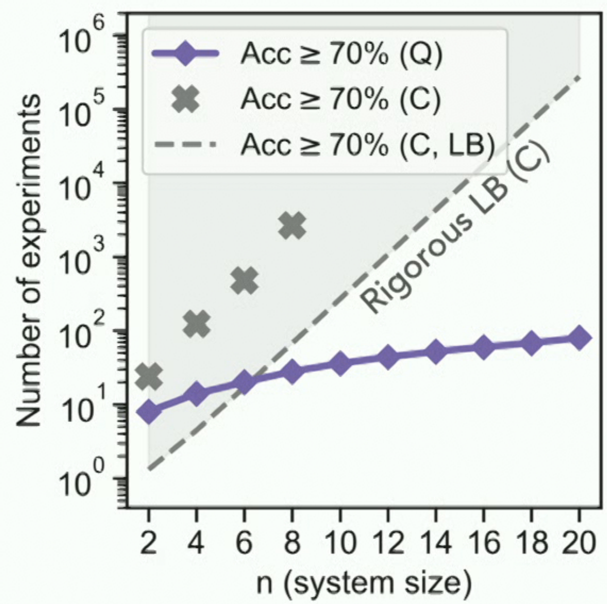

[HFP22] Huang, Flammia, Preskill. Foundations for learning from noisy quantum experiments, *QIP*, 2022.

[HBC+] Huang, Broughton, Cotler, Chen, Li, Mohseni, Neven, Babbush, Kueng, Preskill, McClean. Quantum advantage in learning from experiments, *Science*, 2022.

Demonstration on Sycamore:
Quantum advantage in learning dynamics

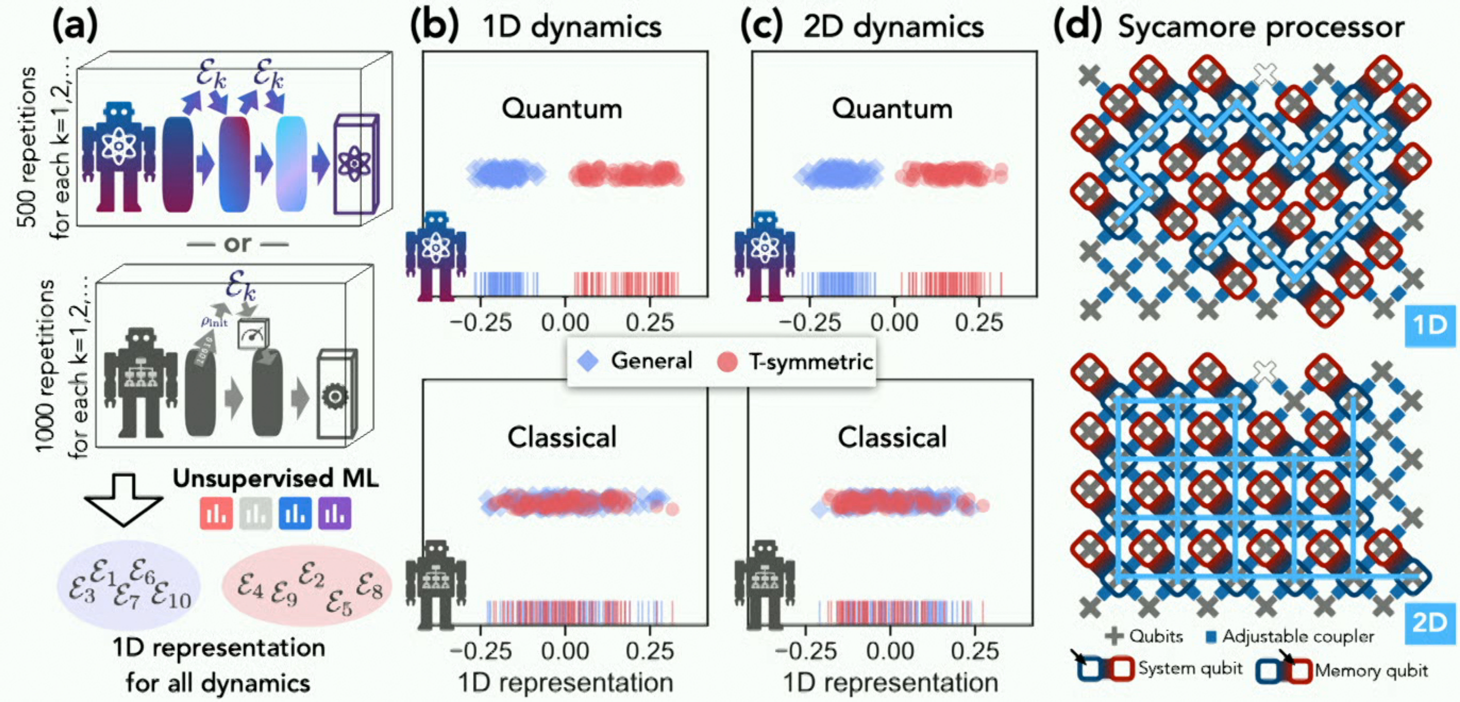

[HFP22] Huang, Flammia, Preskill. Foundations for learning from noisy quantum experiments, *QIP*, 2022.

[HBC+] Huang, Broughton, Cotler, Chen, Li, Mohseni, Neven, Babbush, Kueng, Preskill, McClean. Quantum advantage in learning from experiments, *Science*, 2022.

Overview

How to efficiently learn in the quantum universe?

Learning with classical machines

What can classical machines learn?
Can classical ML perform
better than non-ML algorithms?

Learning with quantum machines

Can quantum machines learn faster
and/or predict more accurately
than classical machines?

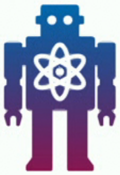

Overview

How to efficiently learn in the quantum universe?

Learning with classical machines

What can classical machines learn?
Can classical ML perform
better than non-ML algorithms?

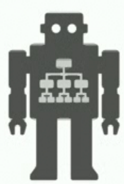

Learning with quantum machines

Can quantum machines learn faster
and/or predict more accurately
than classical machines?

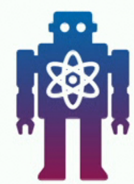

Conclusion

- Significant progress in understanding how to learn in the quantum universe. But most on lower-level tasks (predicting properties, classifying phases, etc.).
- How to create rigorous ML algorithms for higher-level tasks: designing quantum circuits / protocols / algorithms, discovering new physics?

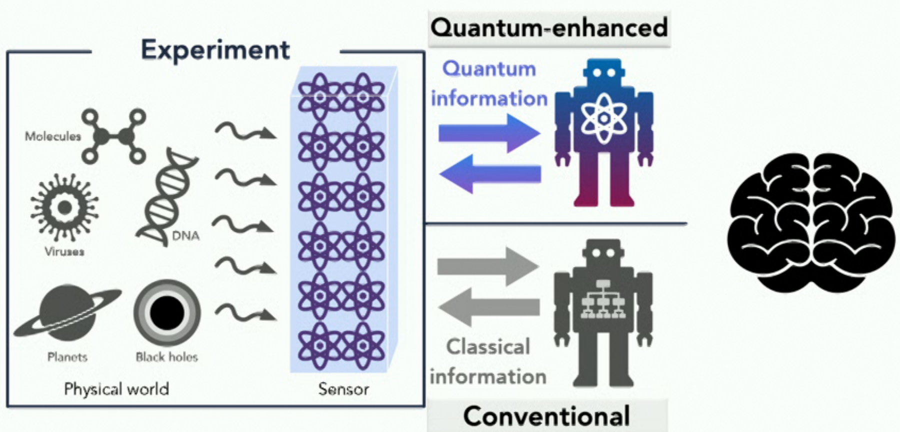

Long-term ambitions

- Develop our understanding of learning to accelerate/automate scientific development (and, perhaps, give rise to an algorithmic theory of science).
- Build a quantum machine capable of learning and discovering new facets of our universe beyond humans and classical machines.

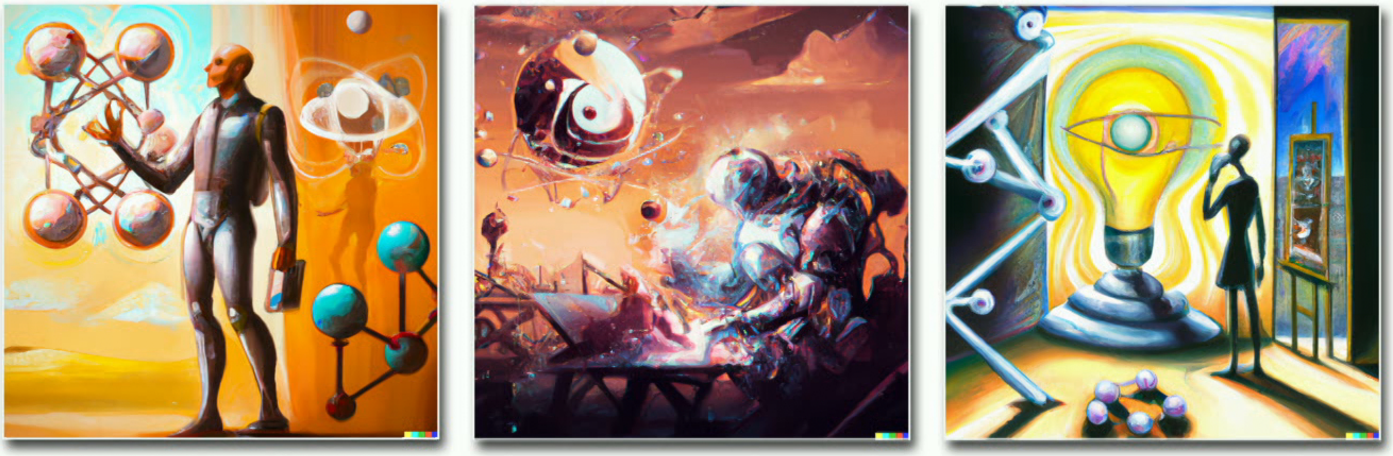

AI imagination of itself learning and discovering new facets of our quantum universe (Credit: DALL·E)