Learning theory in the quantum universe

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.

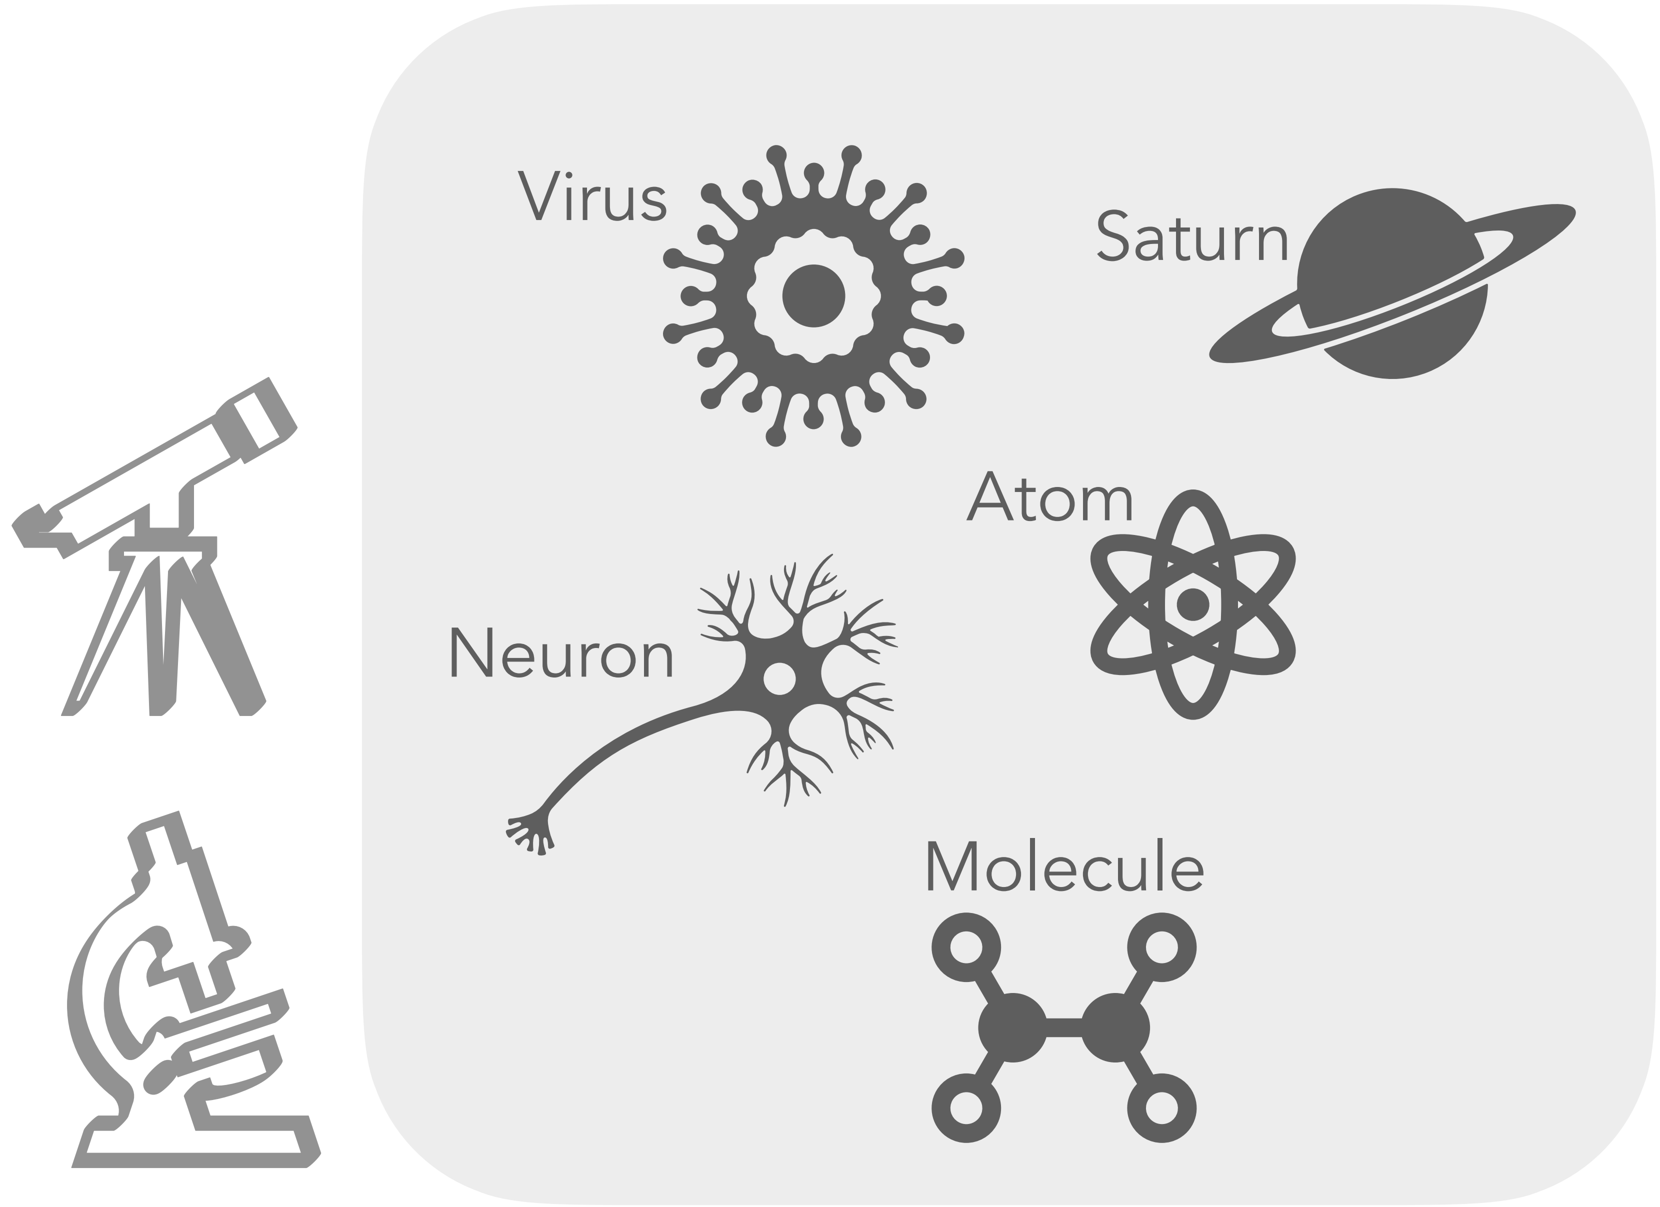

Examples of scientific disciplines

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.

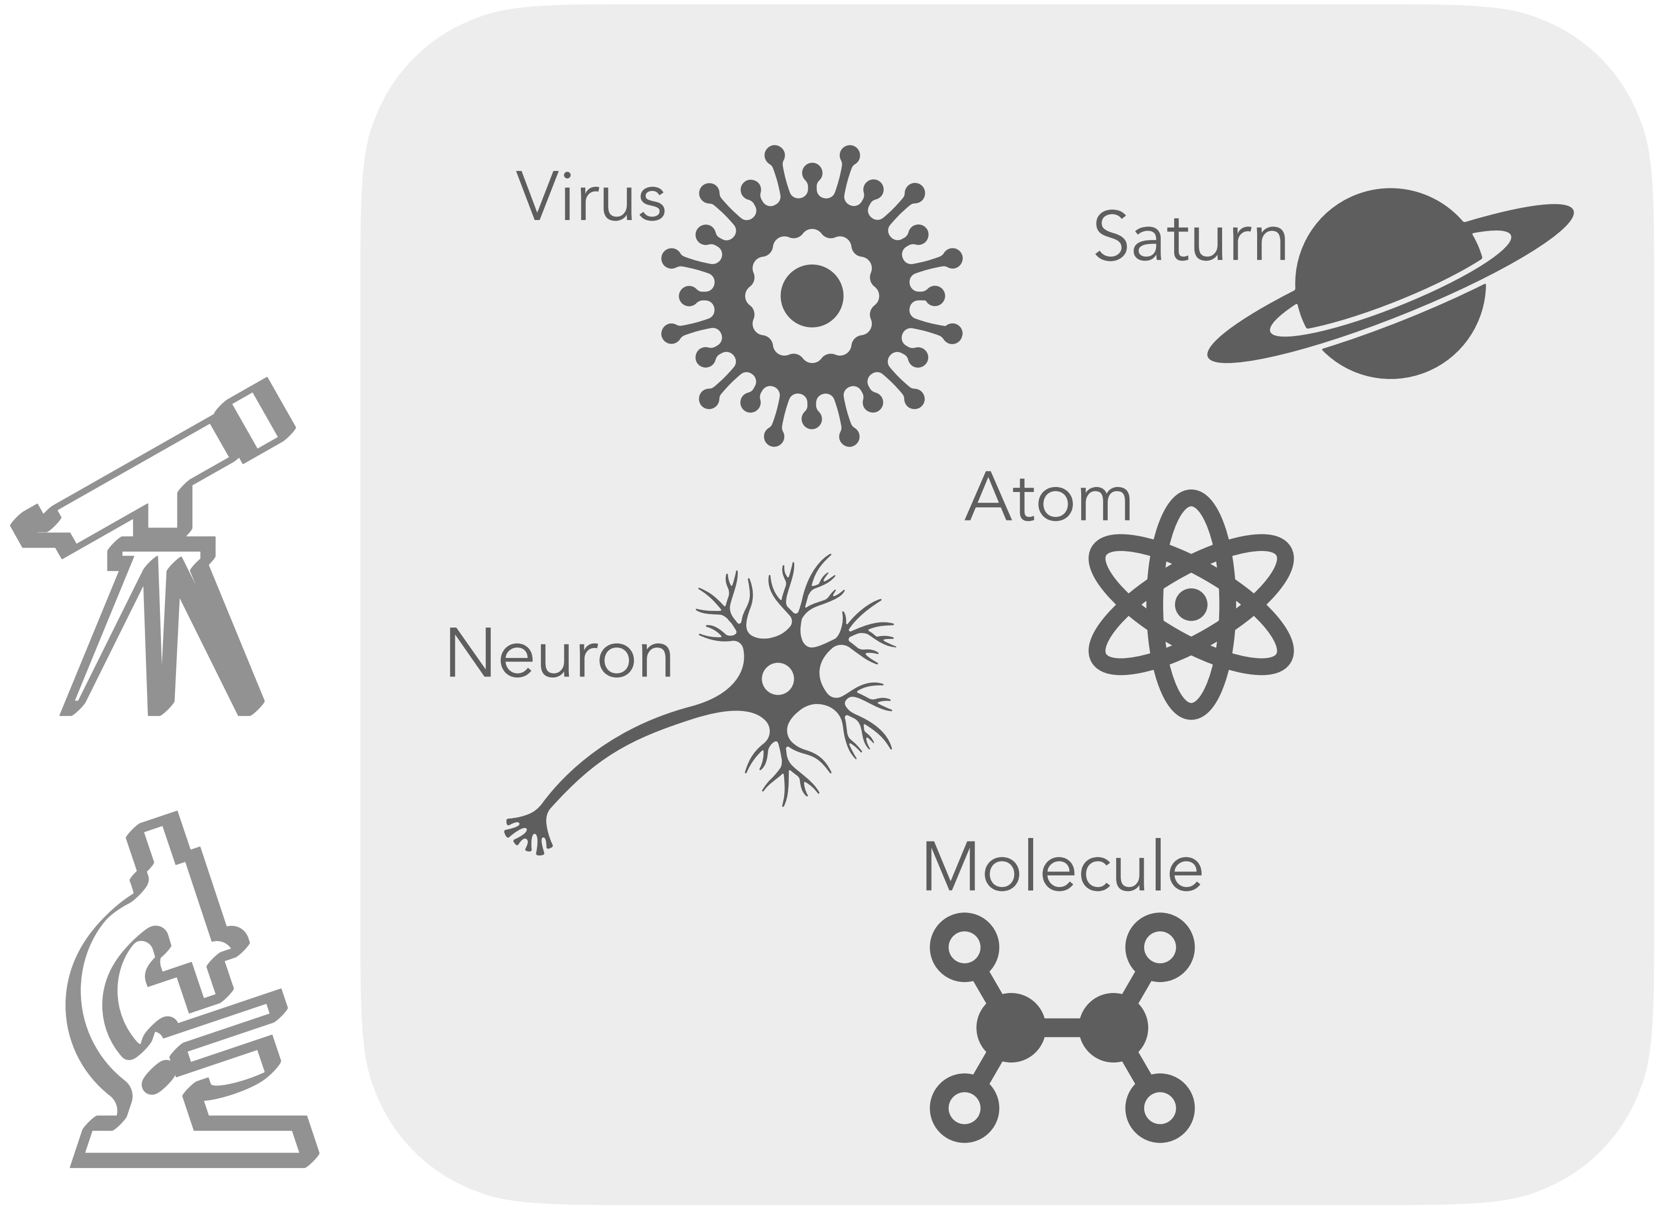

Examples of scientific disciplines

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.

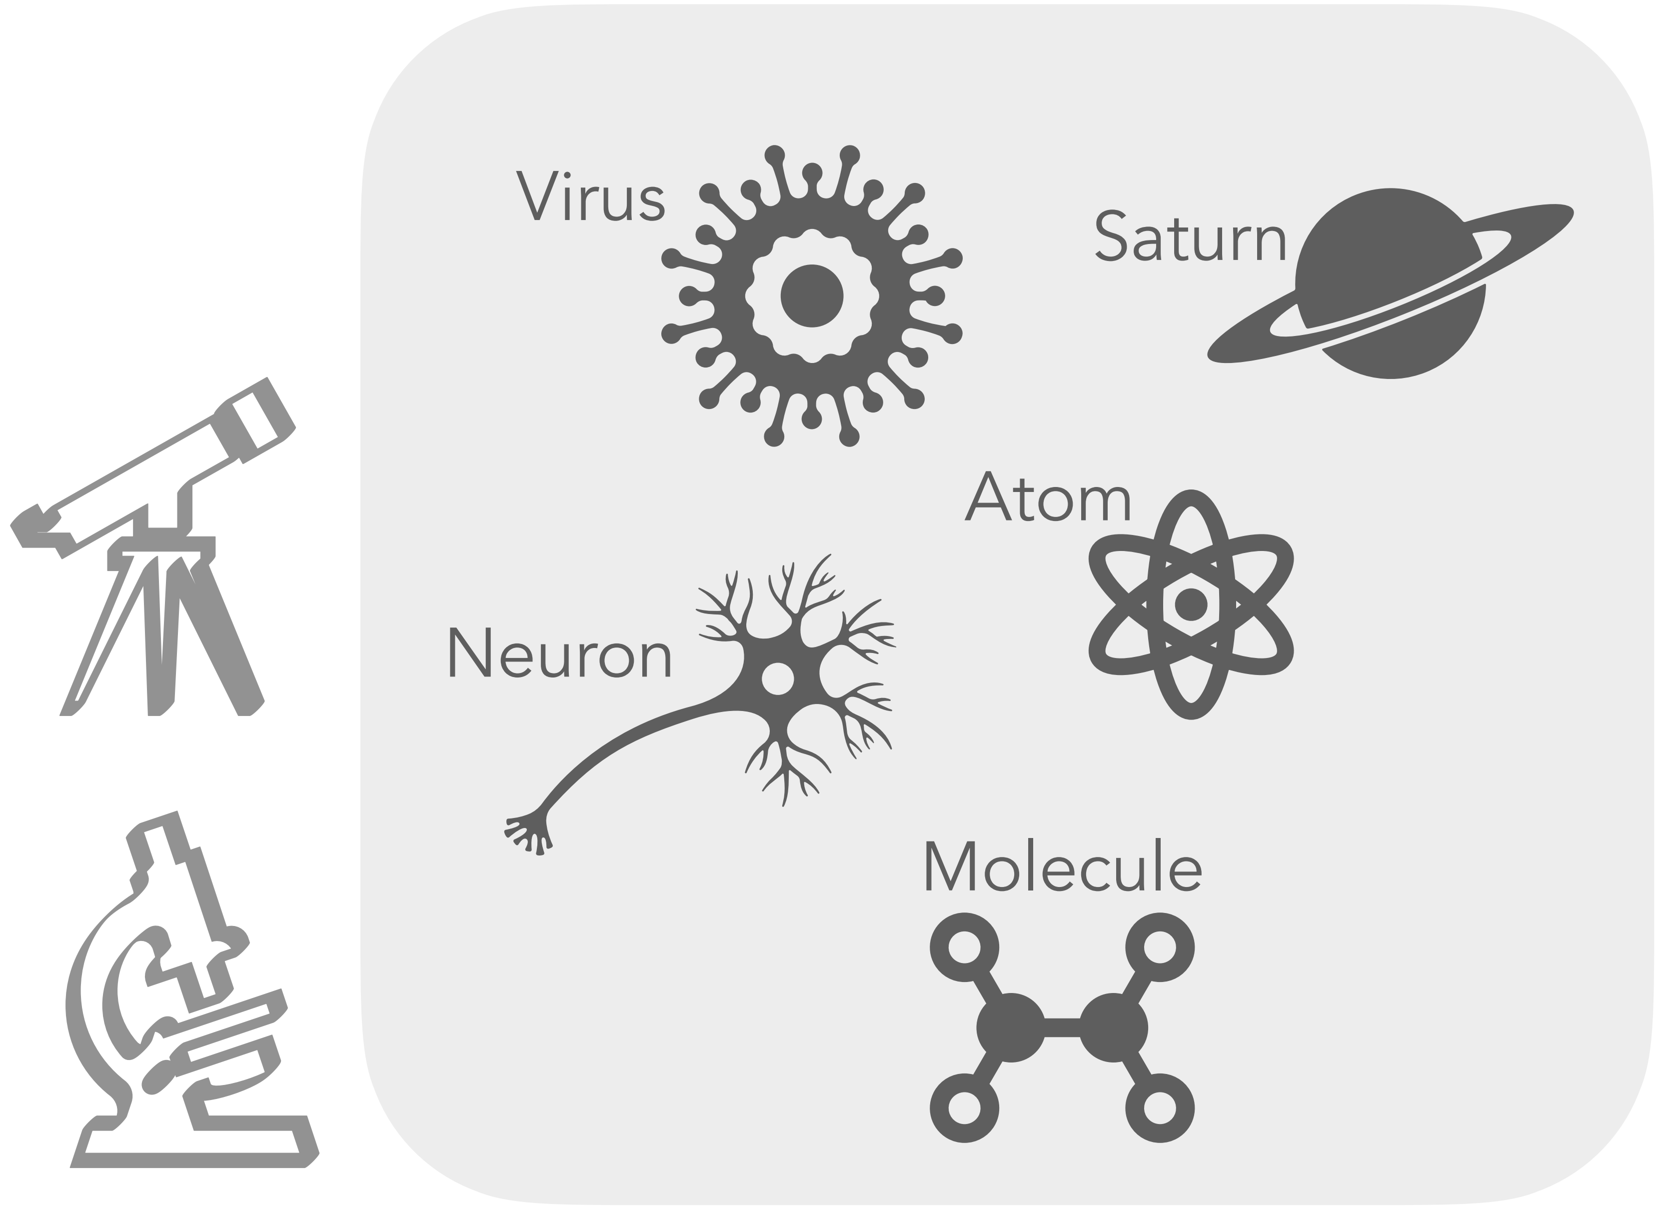

Examples of scientific disciplines

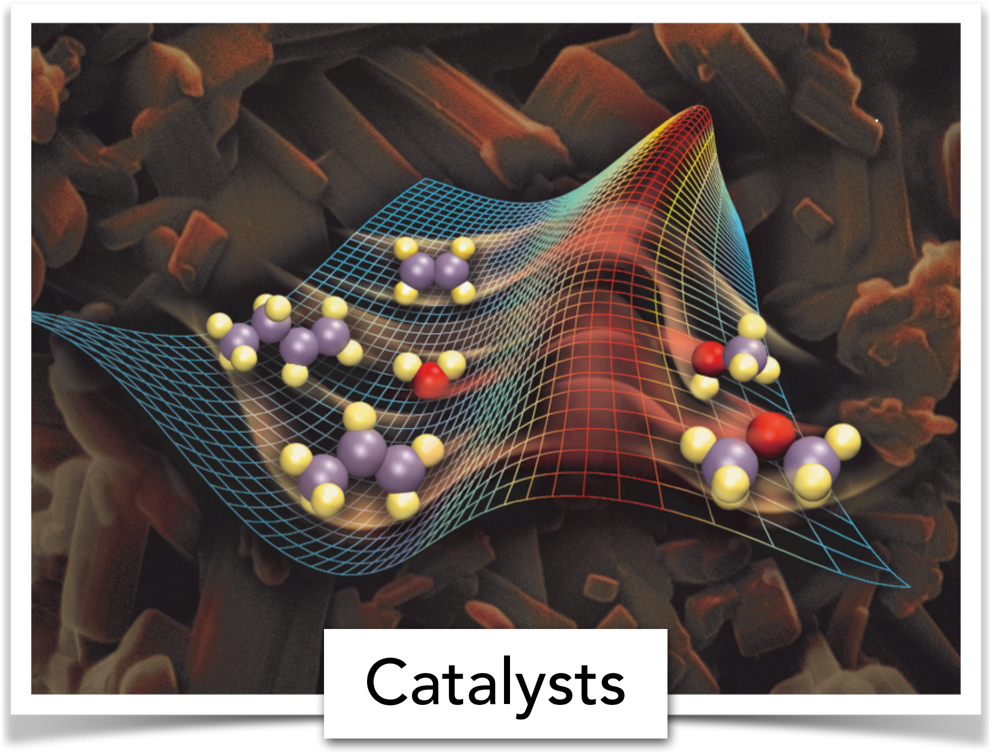

Catalysts

Pharmaceutics

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.

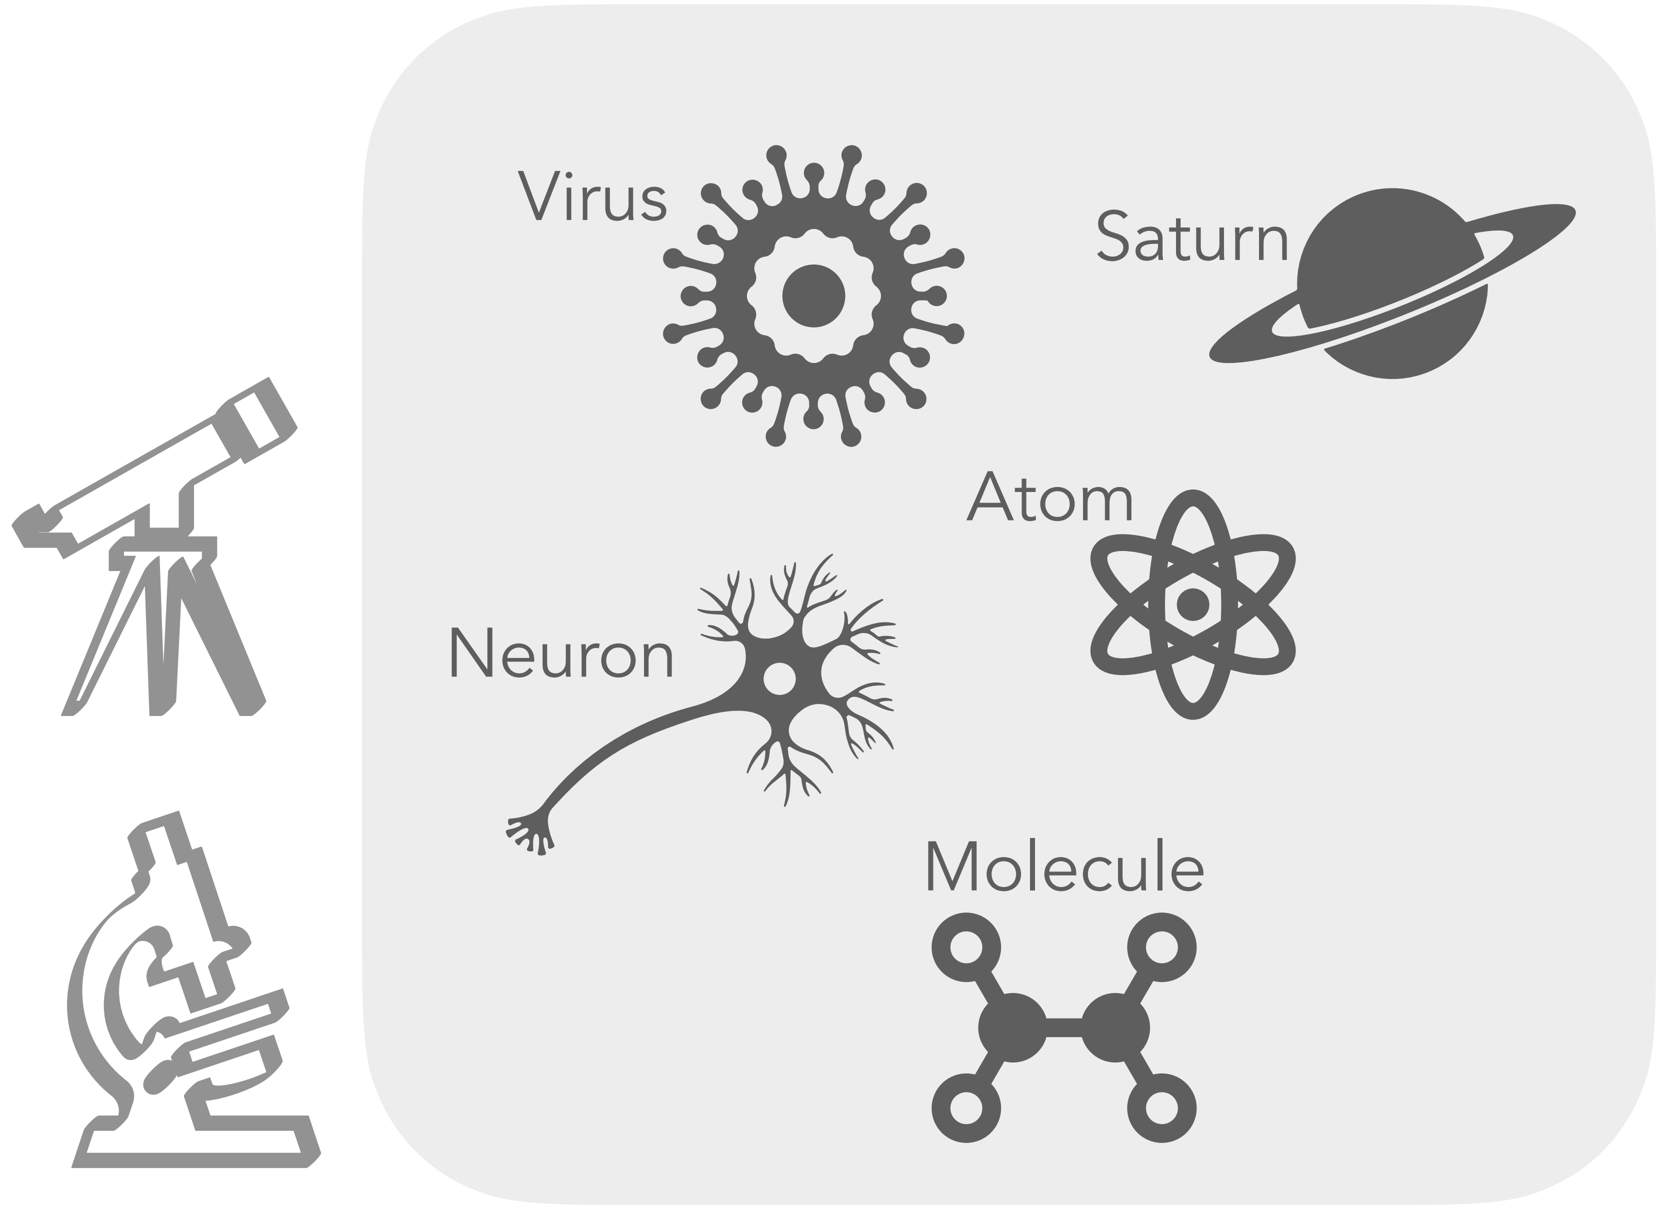

Examples of scientific disciplines

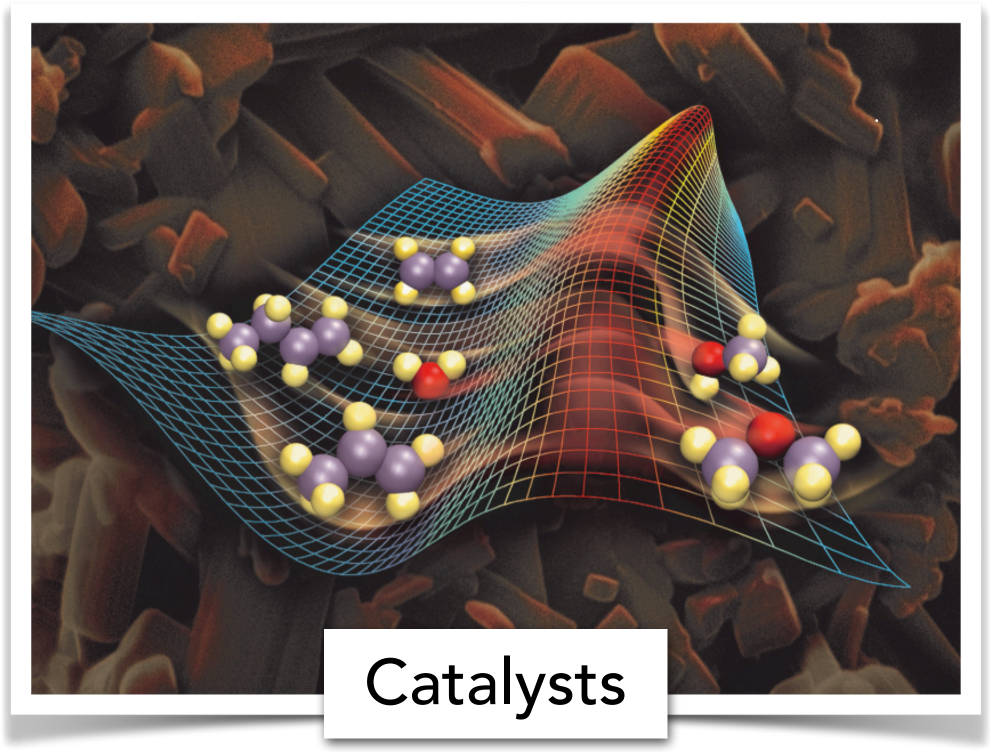

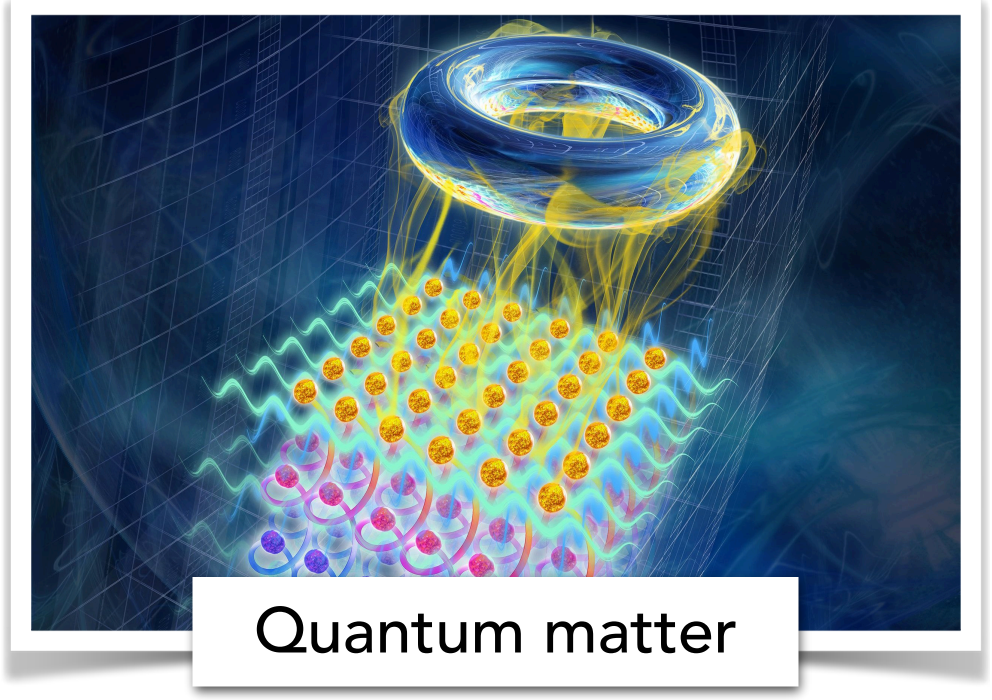

Image credits: (Top left) https://www.energy.gov/science/doe-explainscatalysts (Top right) https://theconversation.com/as-pharmaceutical-use-continues-to-rise-side-effects-are-becoming-a-costly-health-issue-105494 (Bottom left) https://news.mit.edu/2019/ultra-quantum-matter-uqm-research-given-8m-boost-0529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

Motivation

- 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.

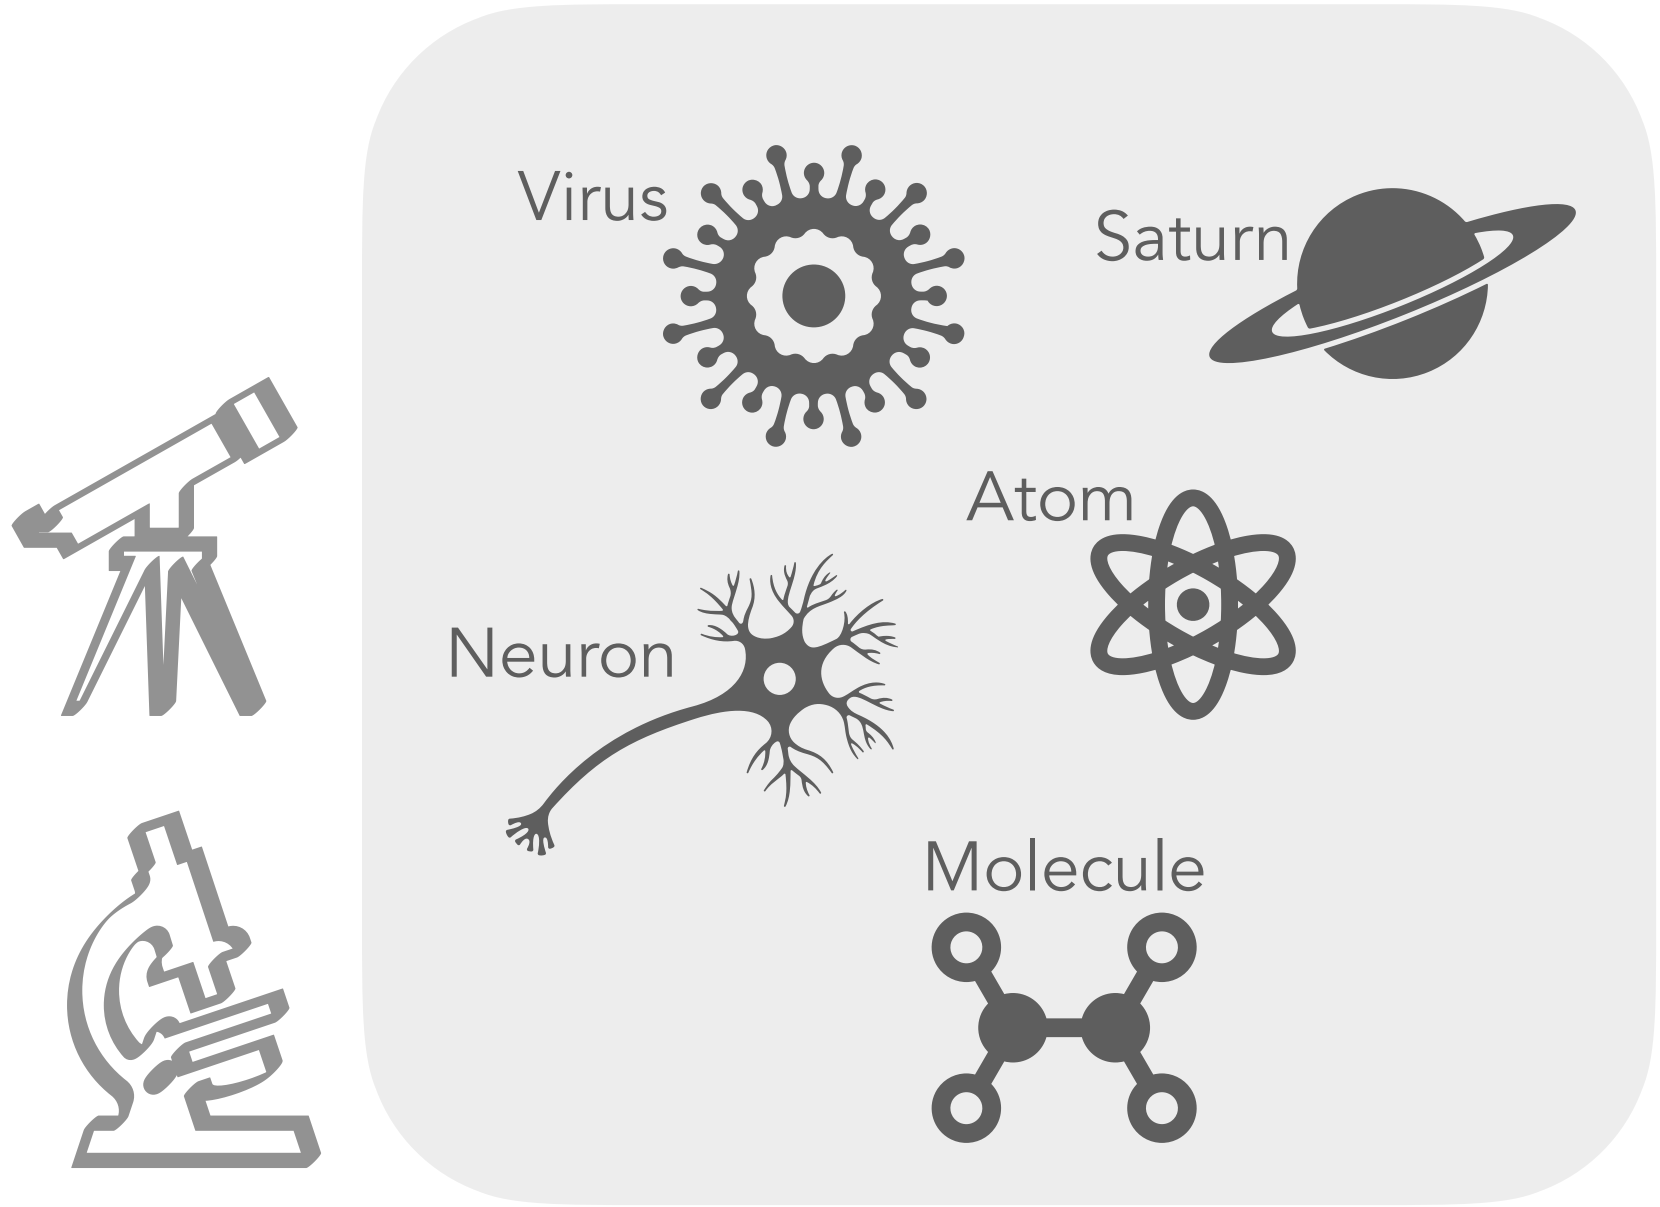

Examples of scientific disciplines

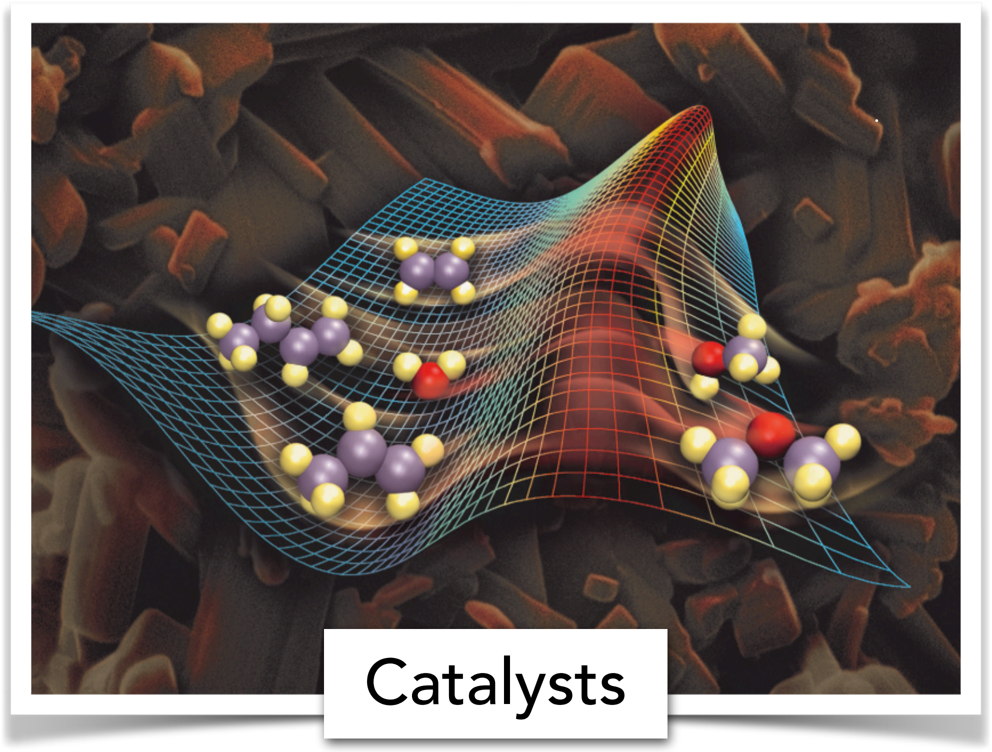

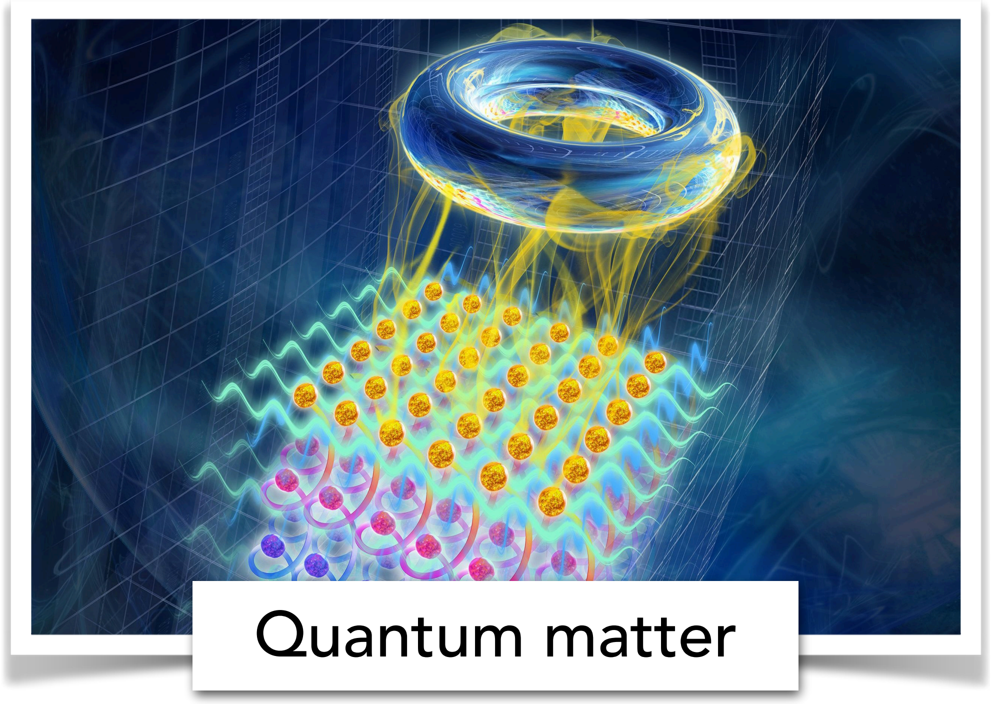

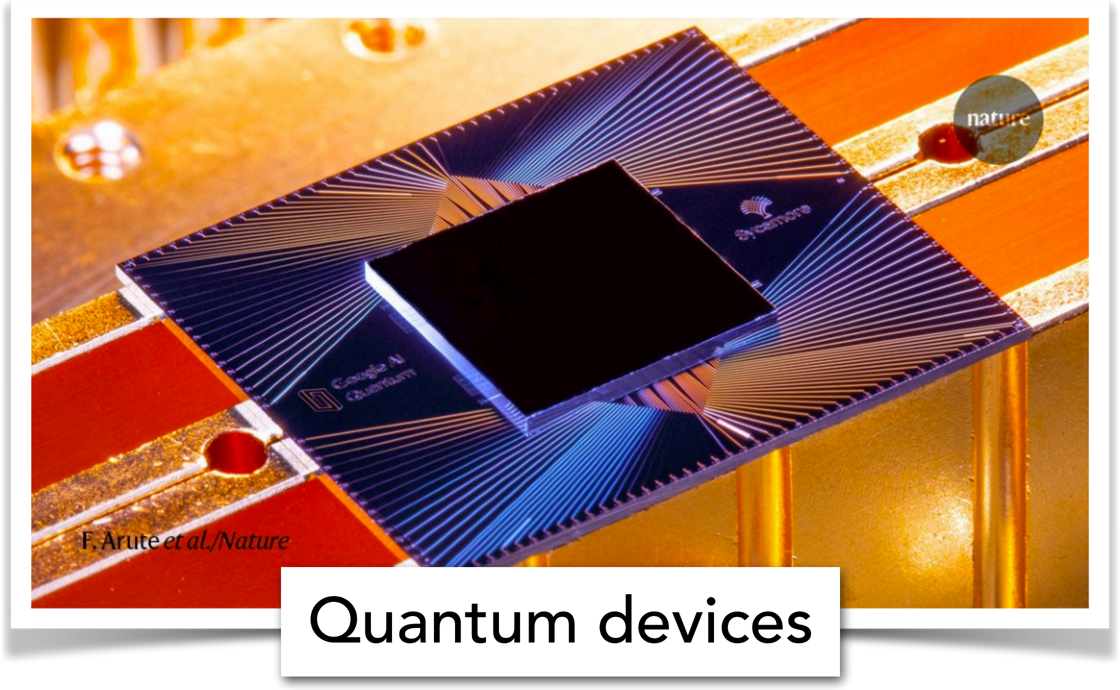

Image credits: (Top left) https://www.energy.gov/science/doe-explainscatalysts (Top right) https://theconversation.com/as-pharmaceutical-use-continues-to-rise-side-effects-are-becoming-a-costly-health-issue-105494 (Bottom left) https://news.mit.edu/2019/ultra-quantum-matter-uqm-research-given-8m-boost-0529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

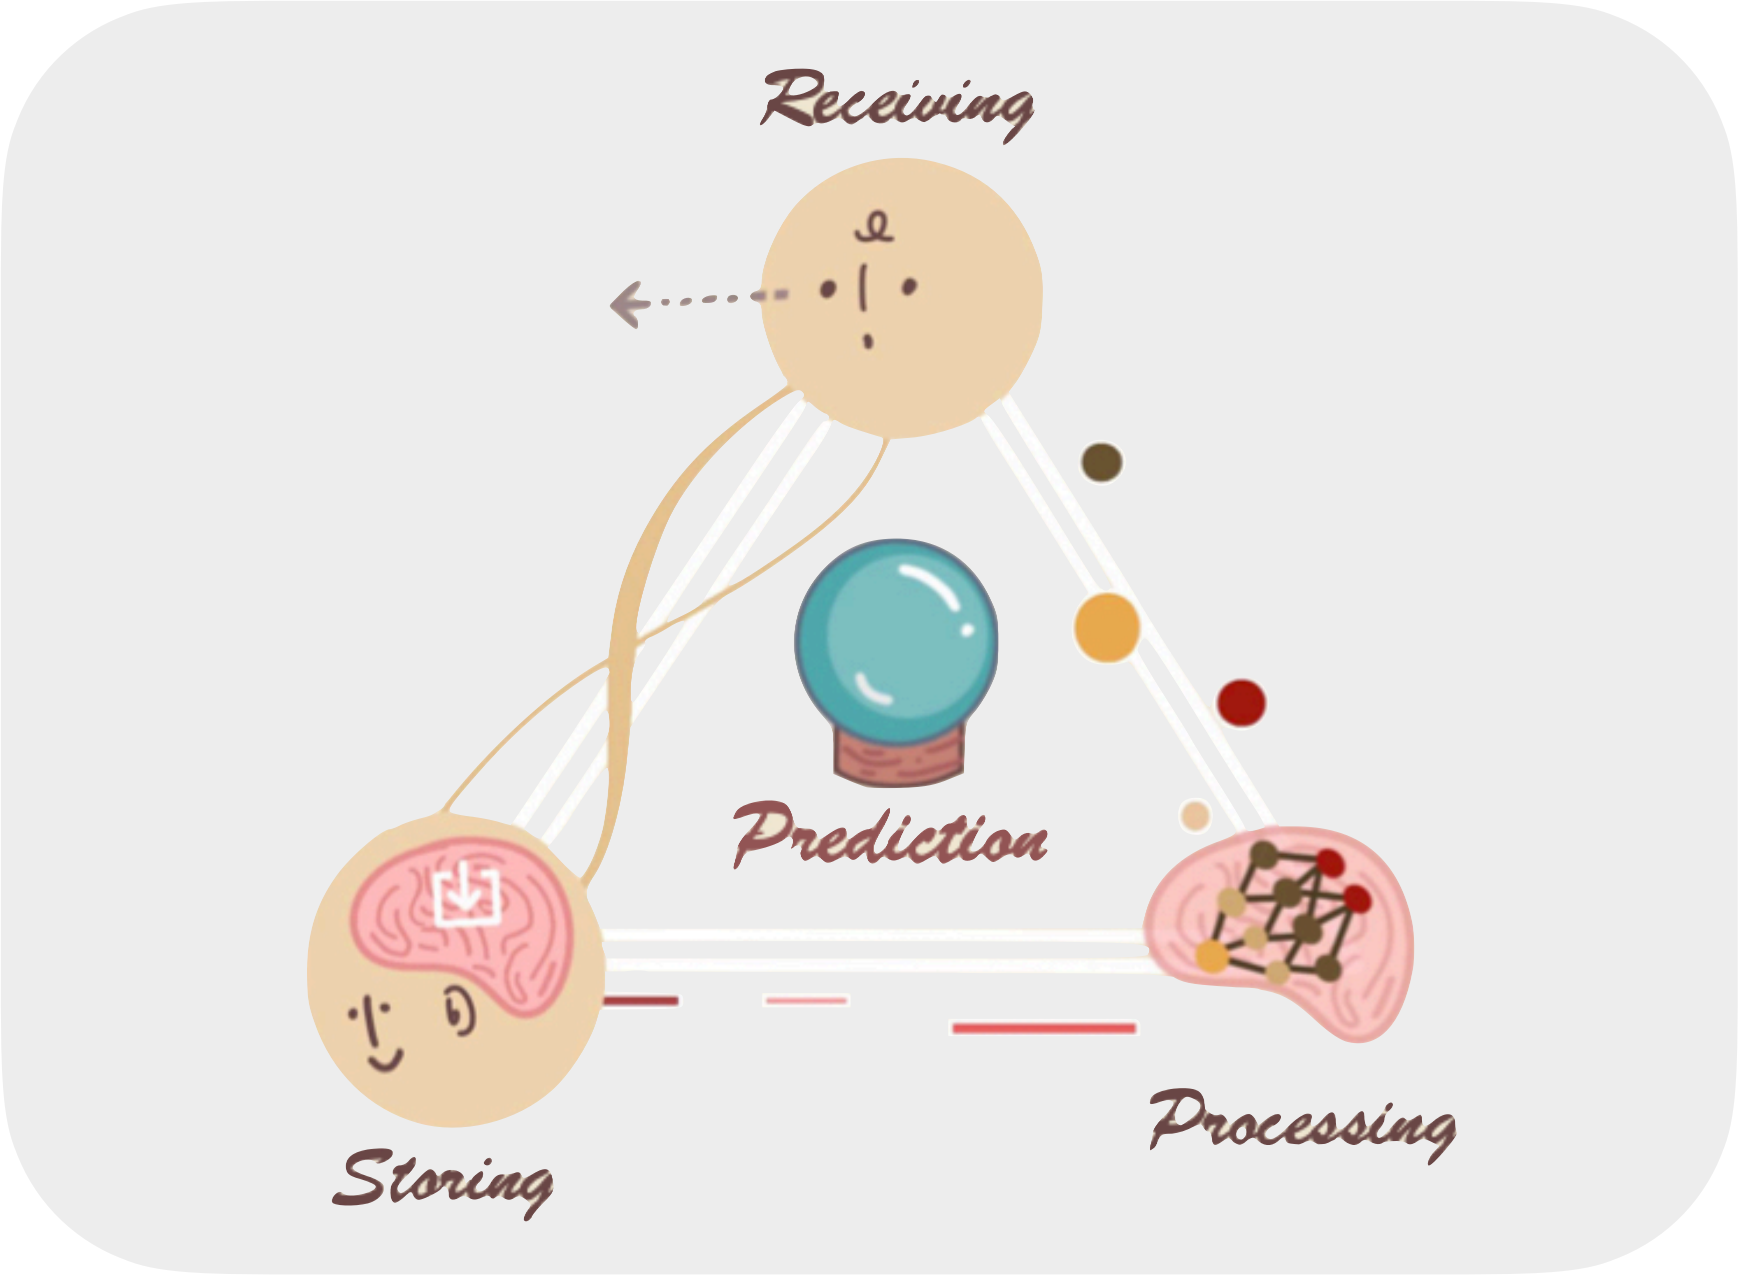

A cartoon depiction of learning

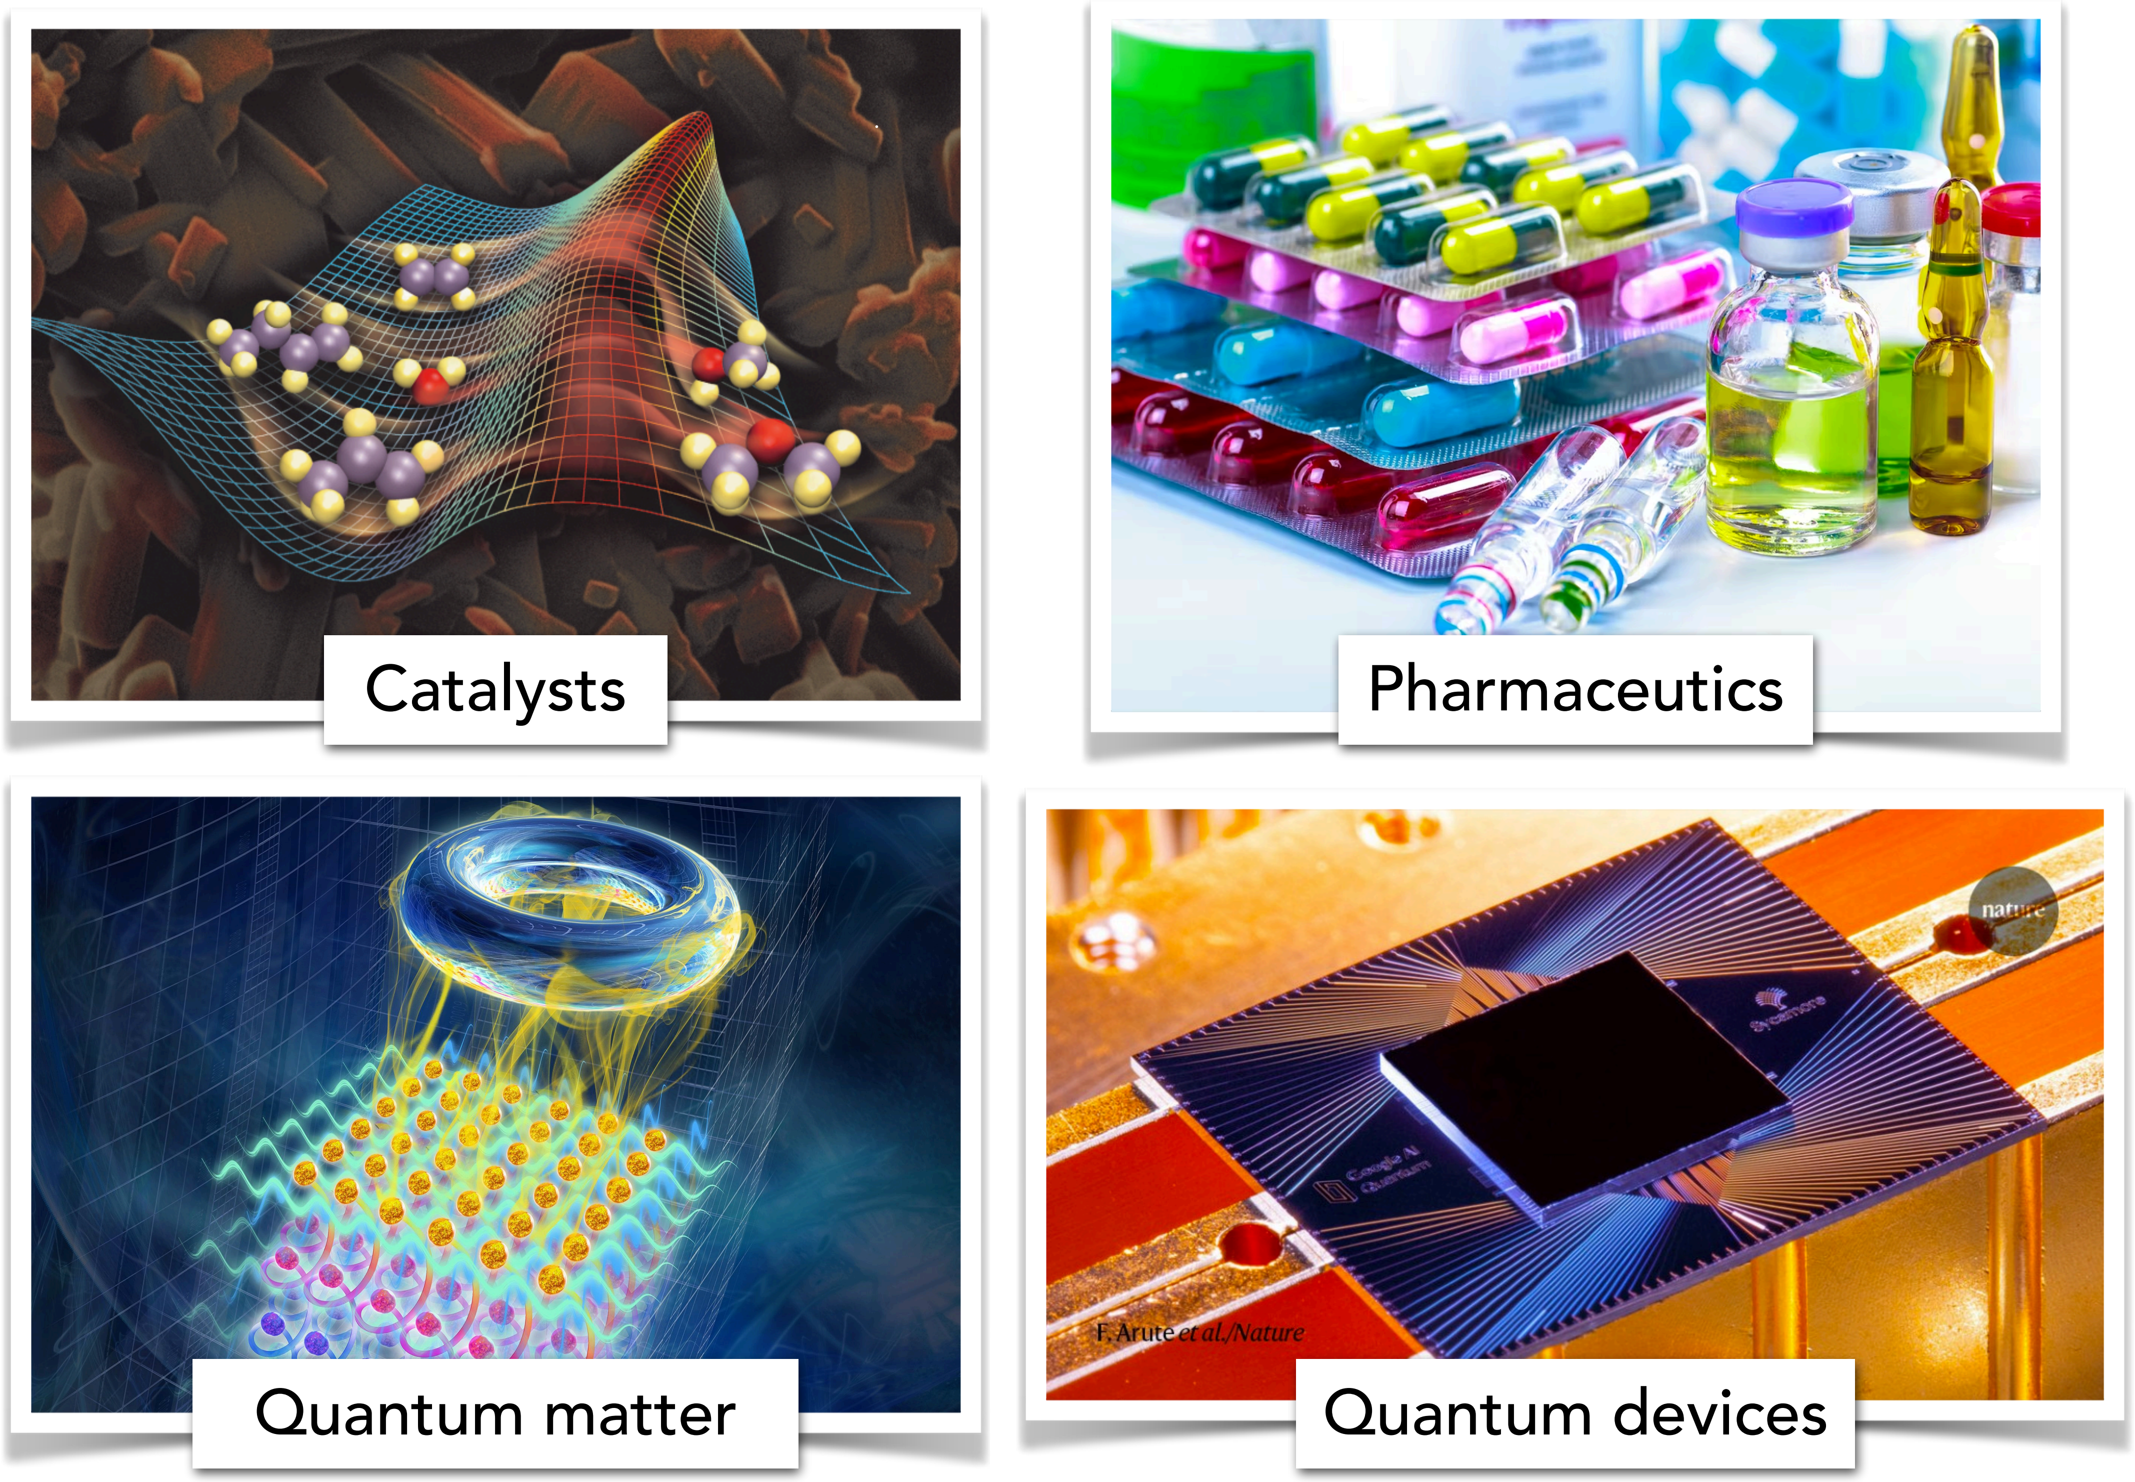

Image credits: (Top left) https://www.energy.gov/science/doe-explainscatalysts (Top right) https://theconversation.com/as-pharmaceutical-use-continues-to-rise-side-effects-are-becoming-a-costly-health-issue-105494 (Bottom left) https://news.mit.edu/2019/ultra-quantum-matter-uqm-research-given-8m-boost-0529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

Motivation

- Learning is the combination of:

  1. receiving information about the universe, 2. processing that information to form models,
  3. storing the models and, subsequently, 4. using the models to predict in new scenarios.

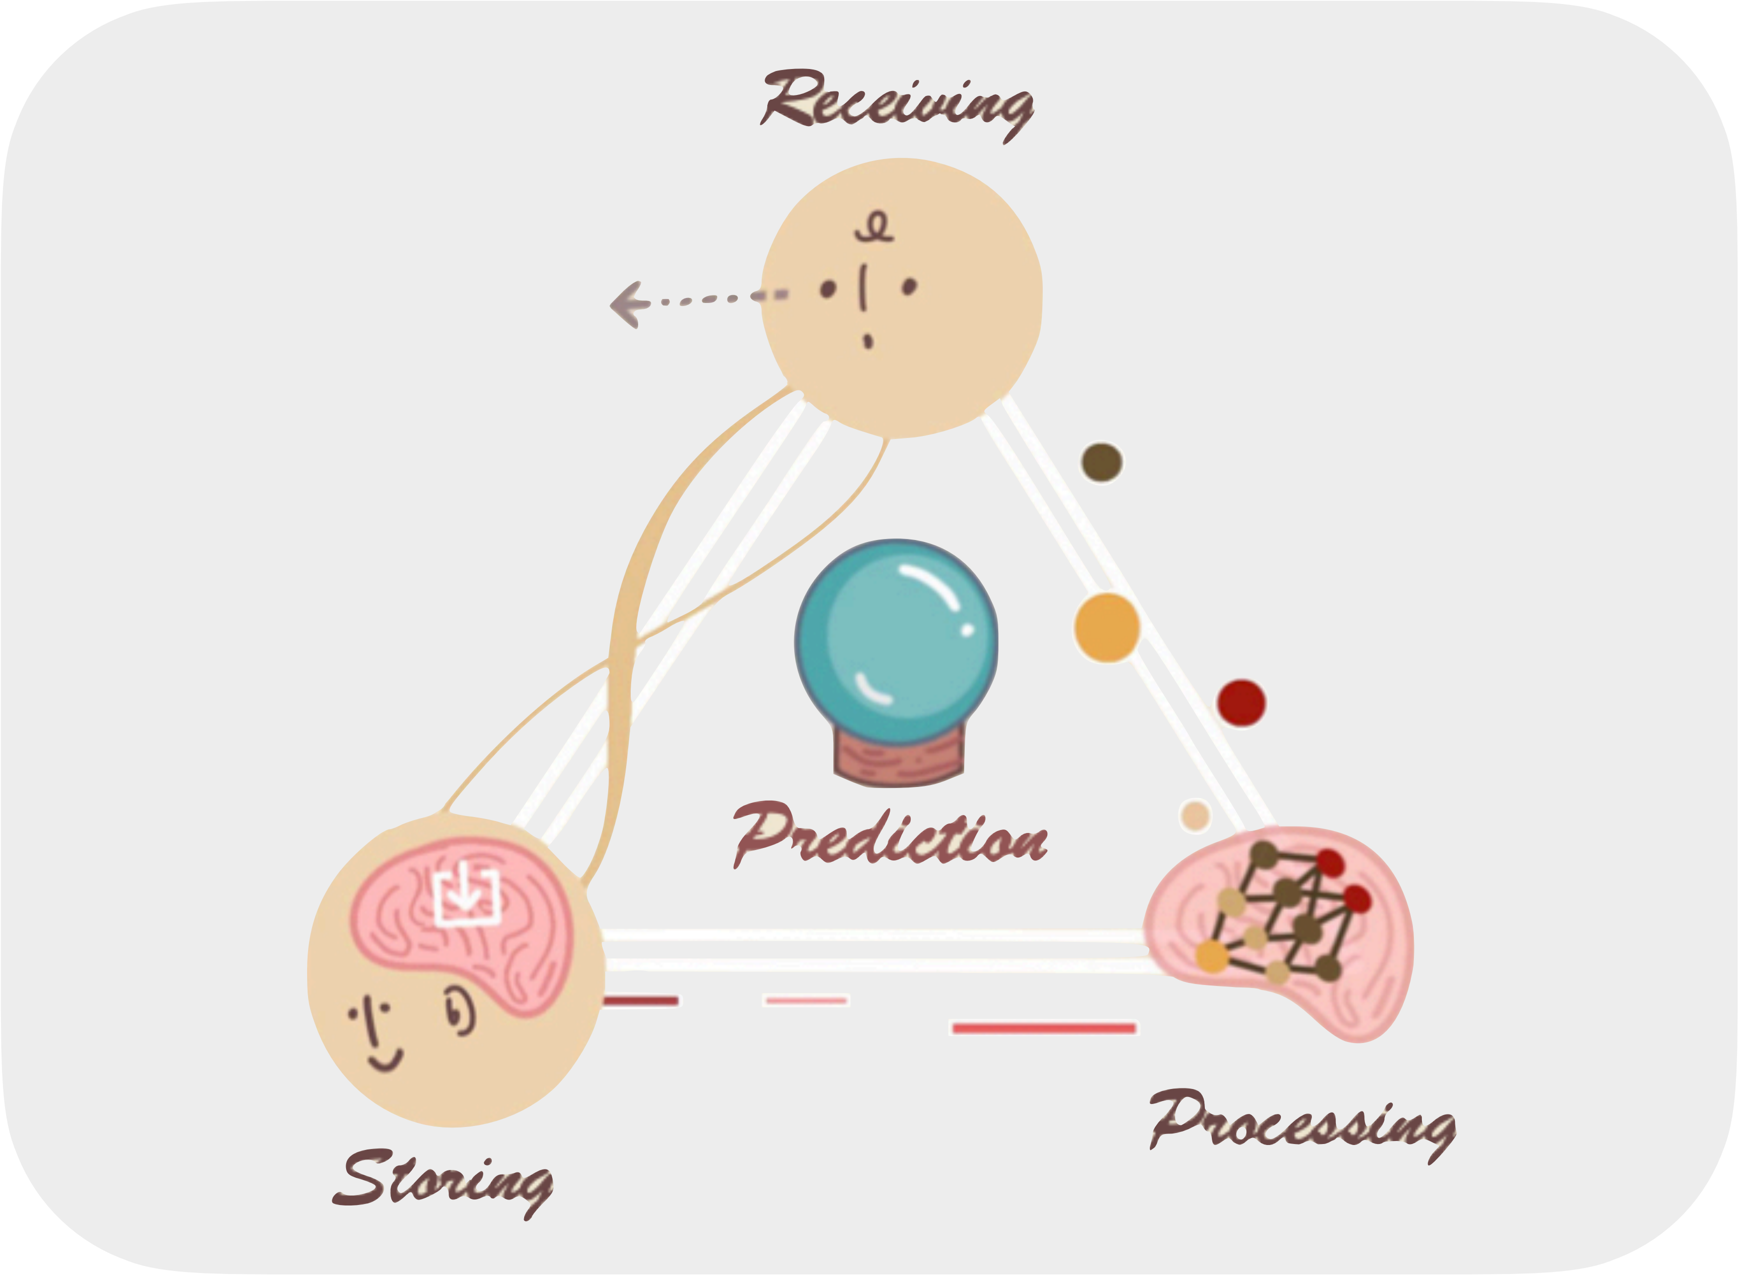

A cartoon depiction of learning

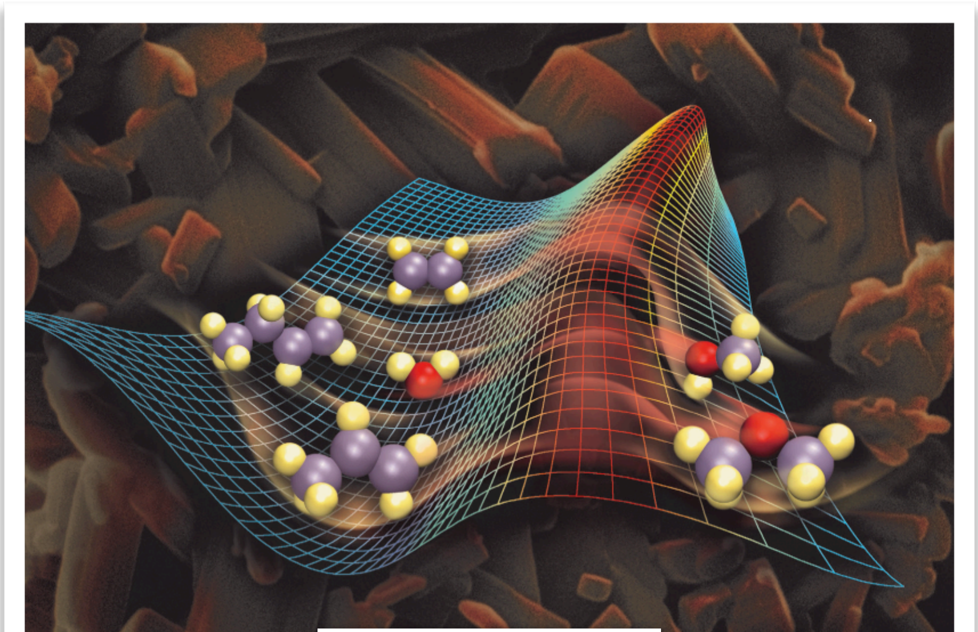
Catalysts


Pharmaceutics

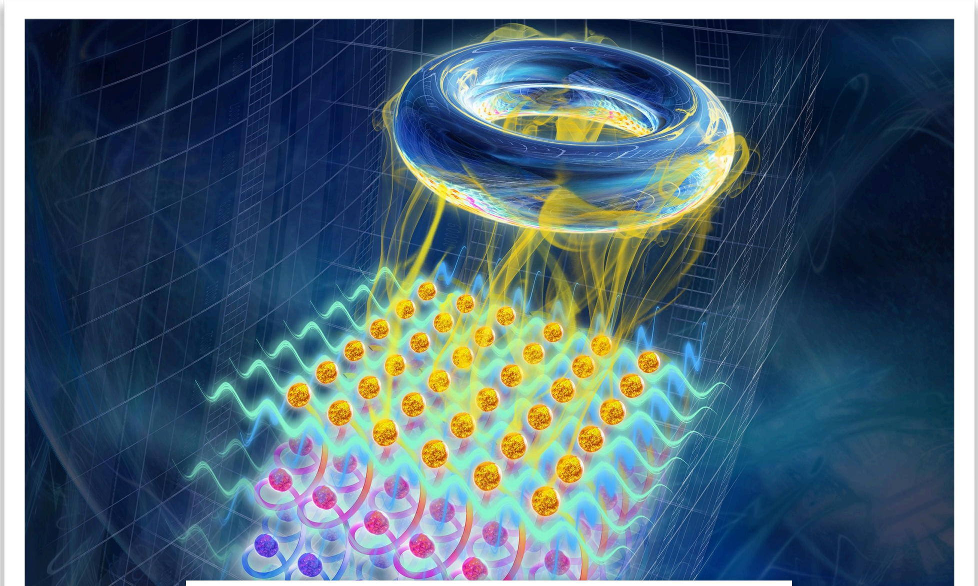
Quantum matter

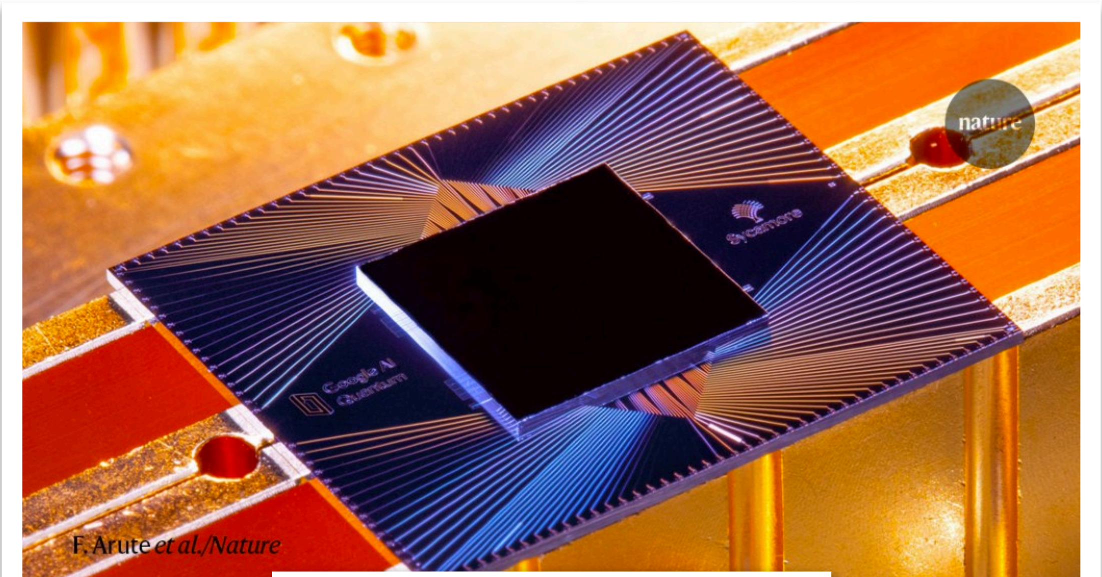

Quantum devices

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?

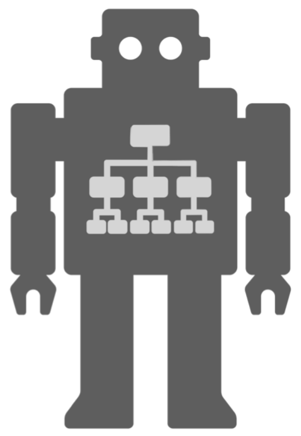

Learning with quantum machines

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

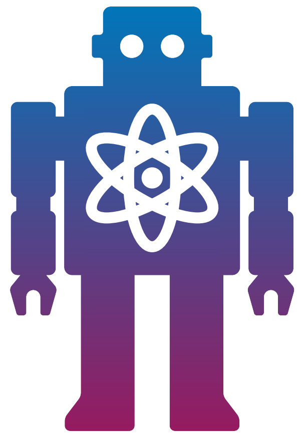

Overview

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?

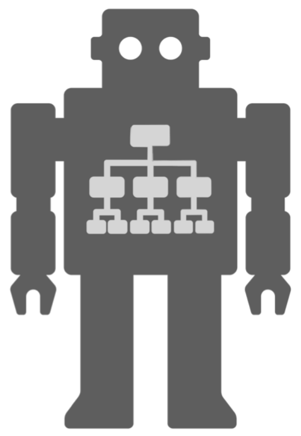

Learning with quantum machines

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

The big question

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

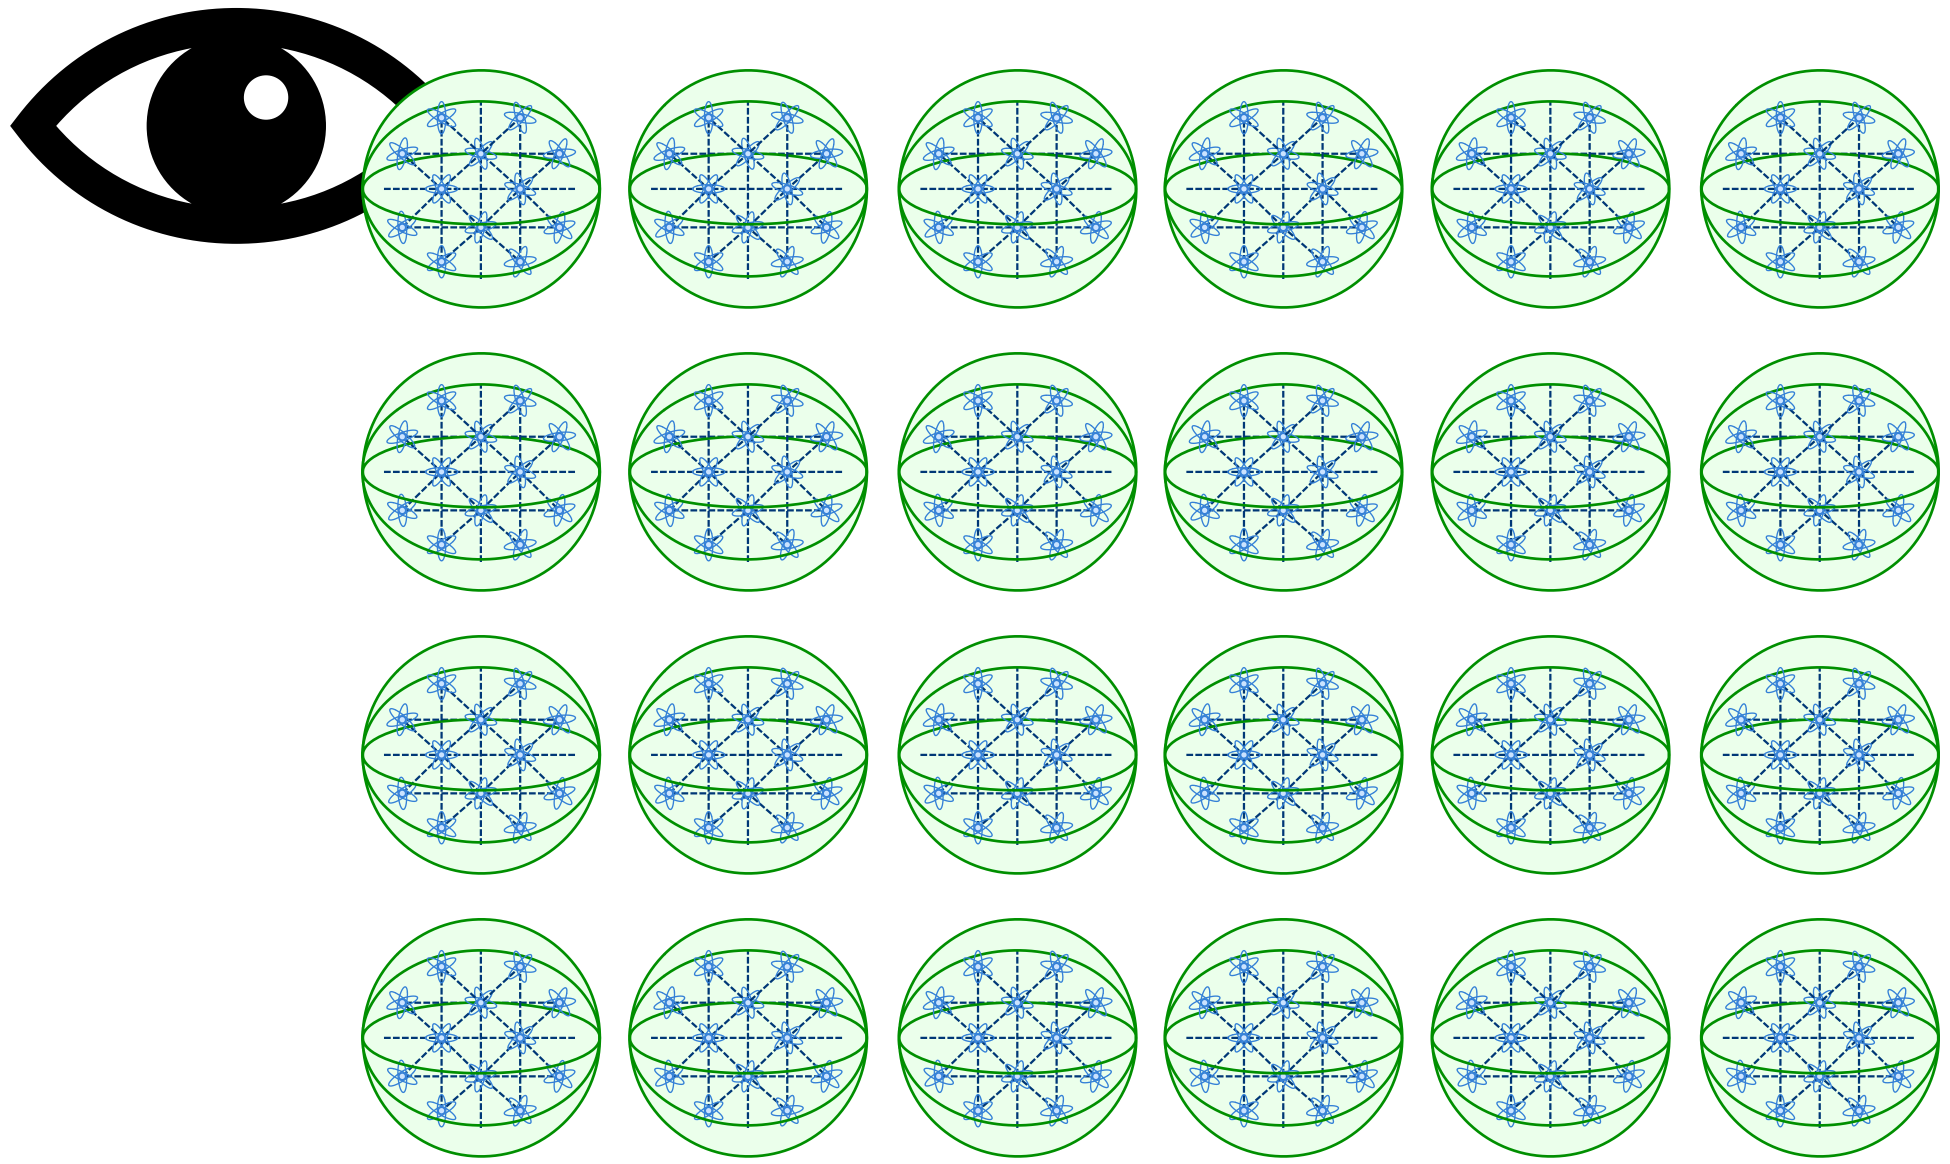

[HKP20] Hsin-Yuan Huang, Richard Kueng, John Preskill. *Predicting many properties of a quantum system from very few measurements*, Nature Physics, 2020.

The big question

- 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.

Unknown Quantum System

Efficient Classical Representation

Predict Properties

[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 the density matrix representation of the n-qubit state ρ.
  ($2^n \times 2^n$ PSD matrix with trace 1)

- 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: $\Omega(2^{2n})$

Recent approach

- **Deep learning:**

  Perform simple quantum measurements.

  Train a neural network to represent the quantum state.

- POVM Neural Network Tomography (Carrasquilla et al.):

  ? Sample complexity: Unknown (could be exponential)

  Quantum resource: Simple quantum circuit + measurements

  ? Classical storage: Unknown (depends on sample complexity)

  ? Classical post-processing: Unknown (could be very long)

What one wants

Find a provably efficient procedure that

1. Learns a classical representation of an unknown n-qubit state ρ from **very few measurements** (not exponential in n).
2. Uses the classical representation to **predict many properties** of the quantum state ρ.

Step 1:

Data Acquisition

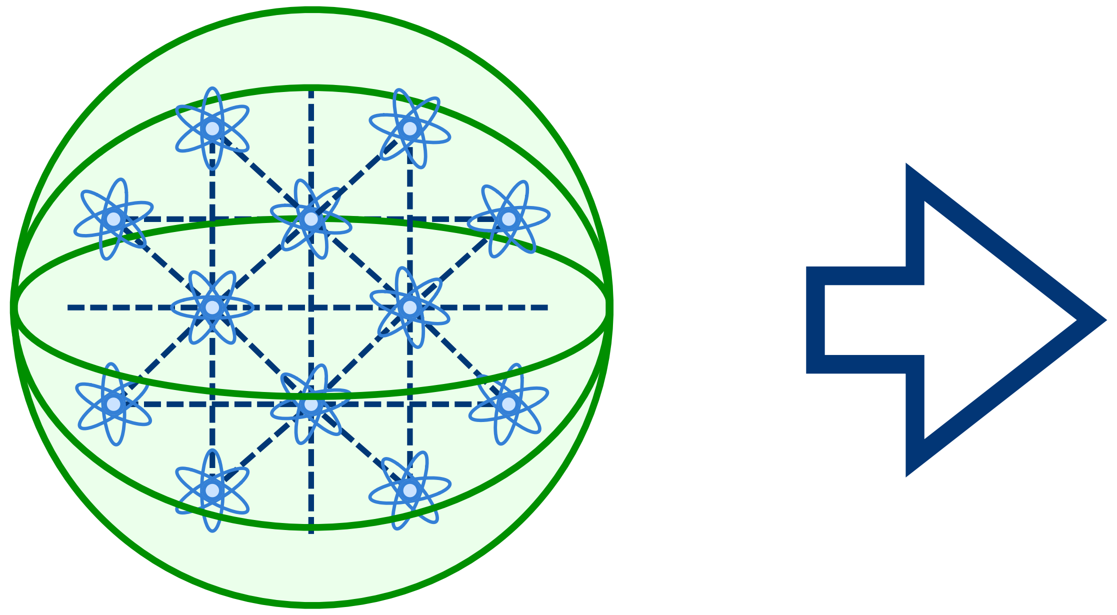

Unknown Quantum System

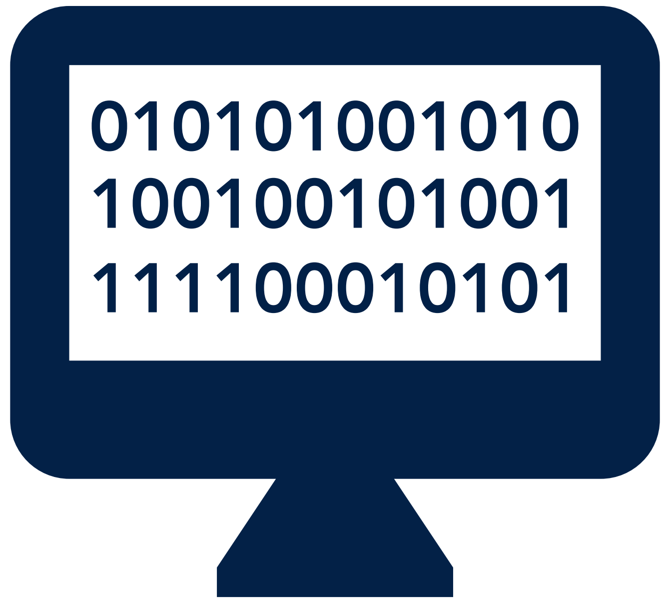

Efficient Classical Representation

What one wants

Find a provably efficient procedure that

1. Learns a classical representation of an unknown n-qubit state ρ from **very few measurements** (not exponential in n).
2. Uses the classical representation to **predict many properties** of the quantum state ρ.

Step 2:

Prediction

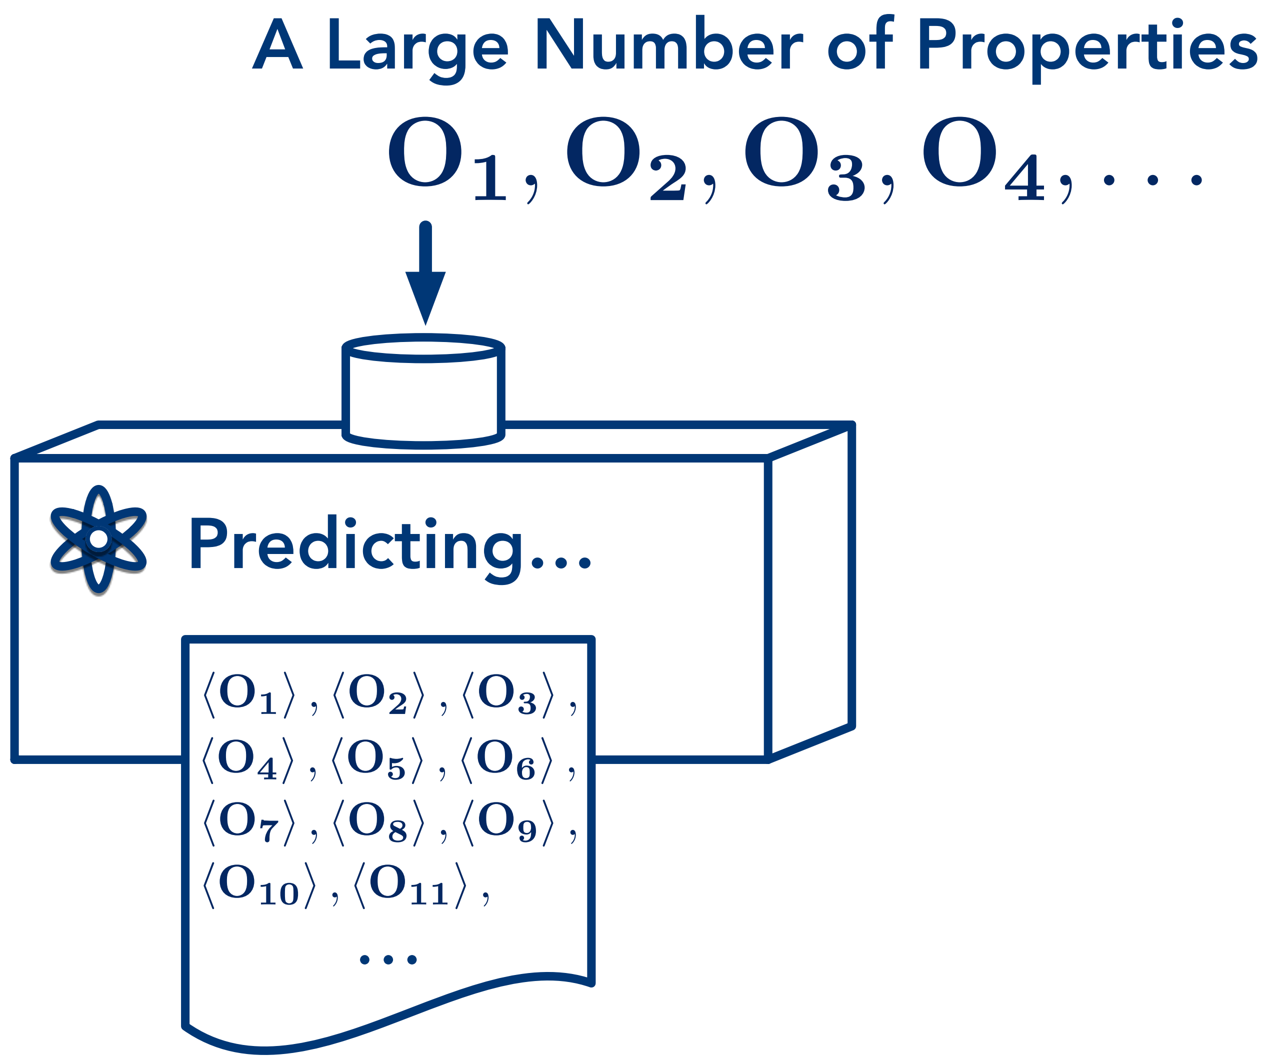

Classical shadow formalism

Theorem 1 [HKP20]

Let M = # of properties, B = norm bound, ϵ = error. ∃ a procedure that

1. Learns a classical representation of an unknown n-qubit state ρ from
$$T = \mathcal{O}(B\log(M)/\epsilon^2) \text{ measurements.}$$
2. Given any $O_1, \ldots, O_M$ with $B \geq \mathrm{Tr}(O_i^2)$, the procedure can use the classical representation to predict $\hat{o}_1, \ldots, \hat{o}_M$, where with high prob.,
$$|\hat{o}_i - \mathrm{tr}(O_i\rho)| < \epsilon, \text{ for all } i = 1,\ldots,M.$$

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

Theorem 1 [HKP20]

Let M = # of properties, B = norm bound, ϵ = error. ∃ a procedure that

1. Learns a classical representation of an unknown n-qubit state ρ from

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

2. Given any $O_1, \ldots, O_M$ with $B \geq \mathrm{Tr}(O_i^2)$, the procedure can use the classical representation to predict $\hat{o}_1, \ldots, \hat{o}_M$, where with high prob.,

$$|\hat{o}_i - \mathrm{tr}(O_i \rho)| < \epsilon, \text{ for all } i = 1, \ldots, M.$$

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:

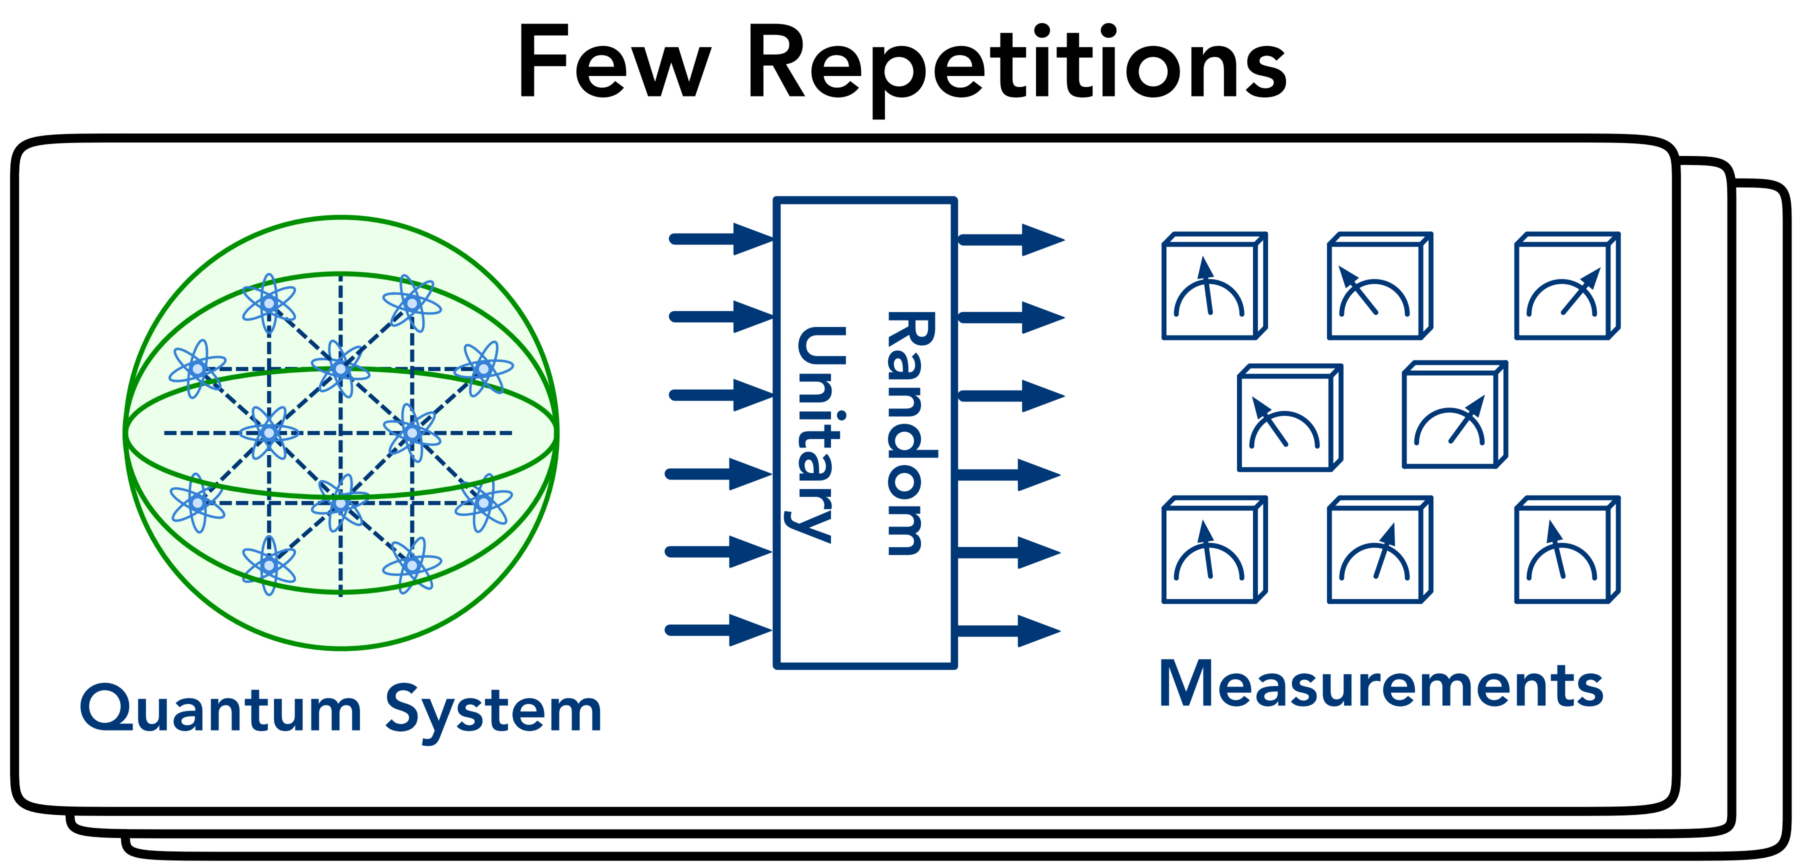

Data Acquisition Phase

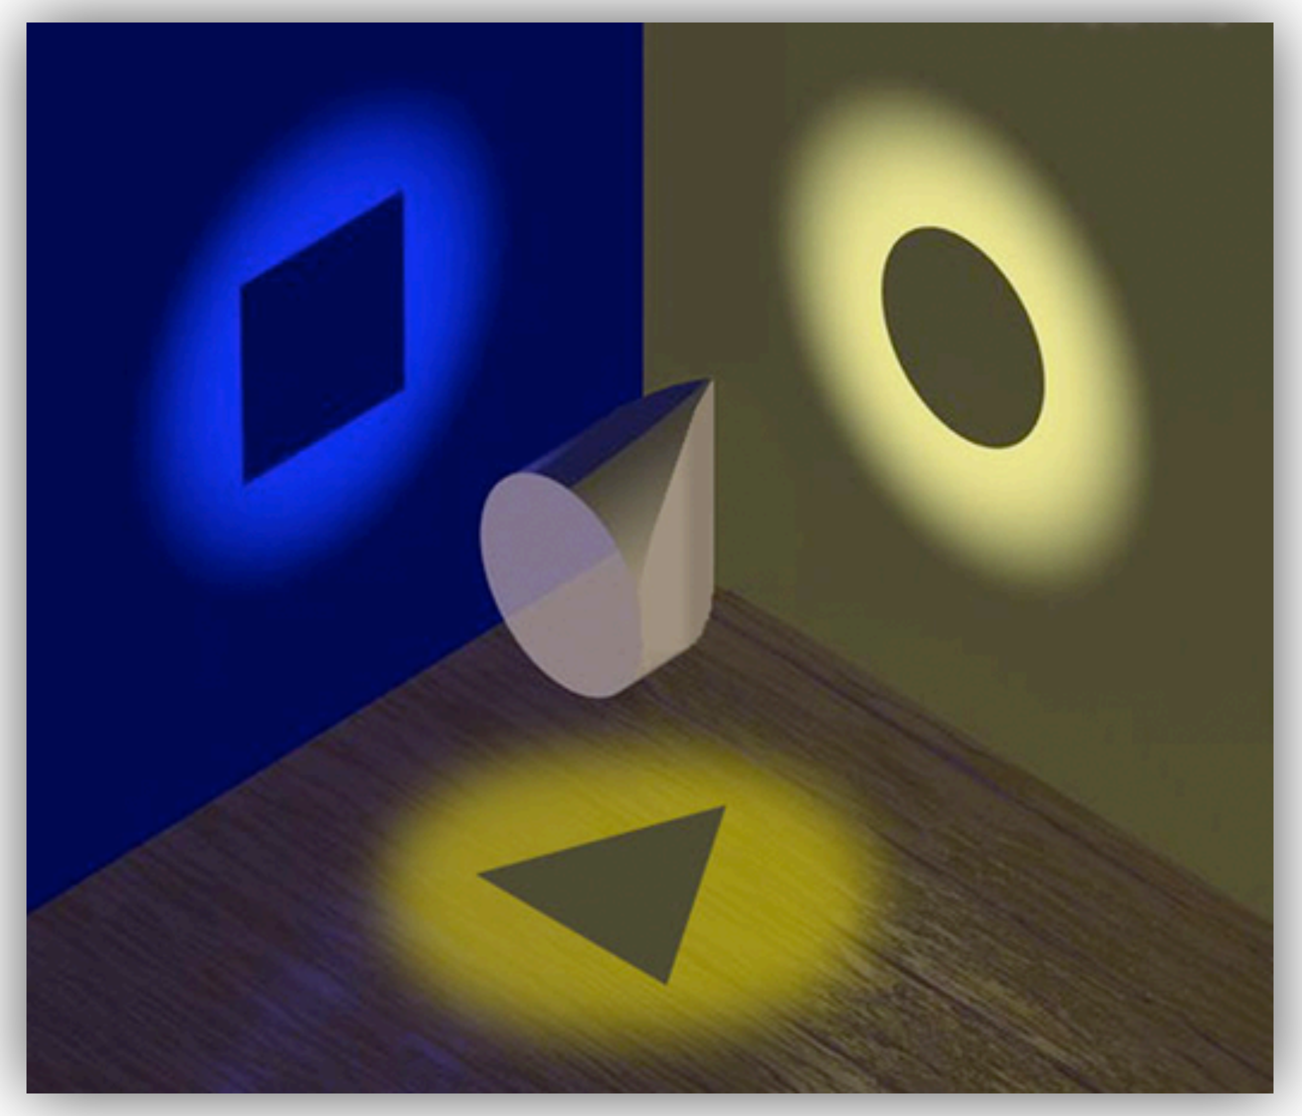

The Procedure:
Data Acquisition Phase

Repeat the following T times:

- Sample a random Clifford circuit U_i to evolve the system.

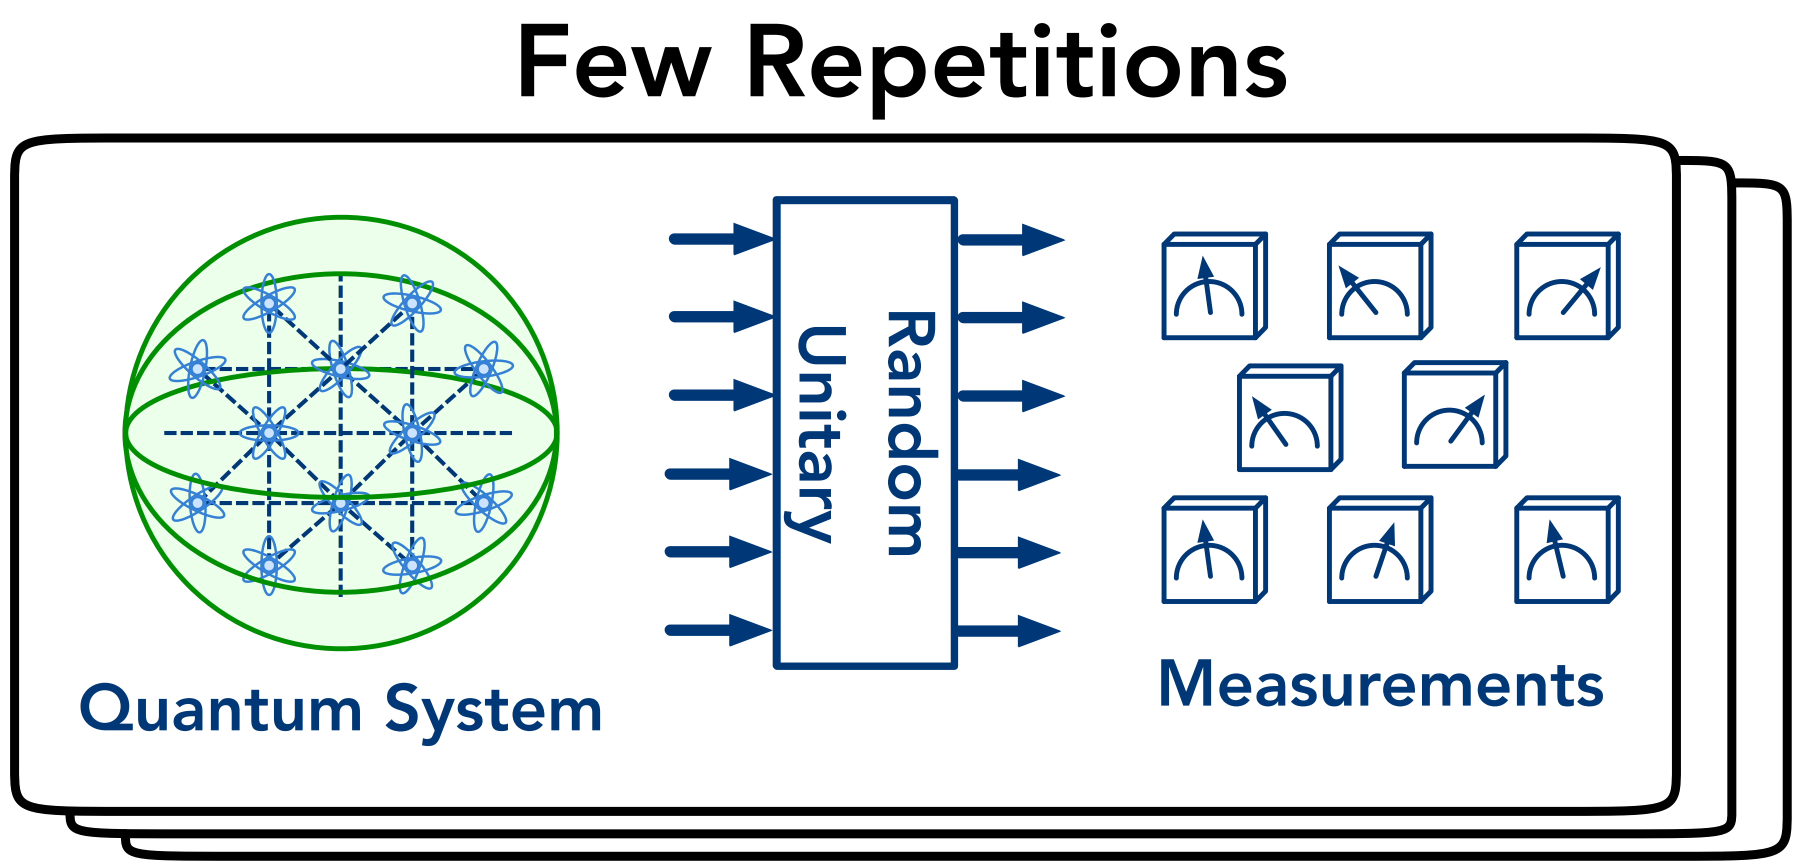

Data Acquisition Phase

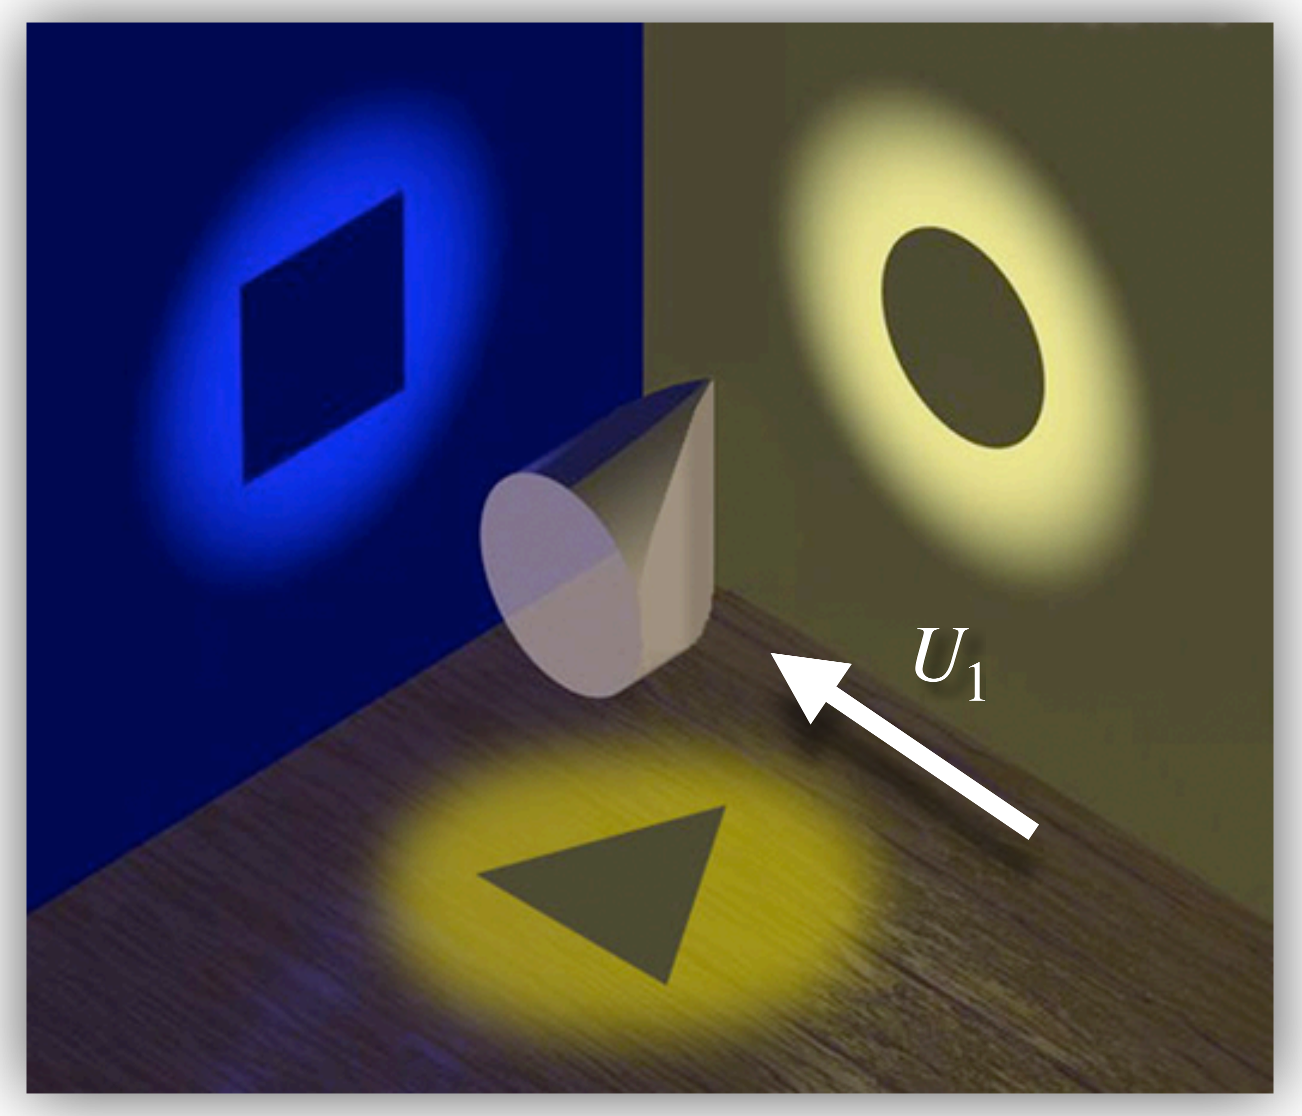

The Procedure:
Data Acquisition Phase

Repeat the following T times:

- Sample a random Clifford circuit U_i to evolve the system.
- Measure the system in the computational basis $|b_i\rangle \in \{0,1\}^n$.

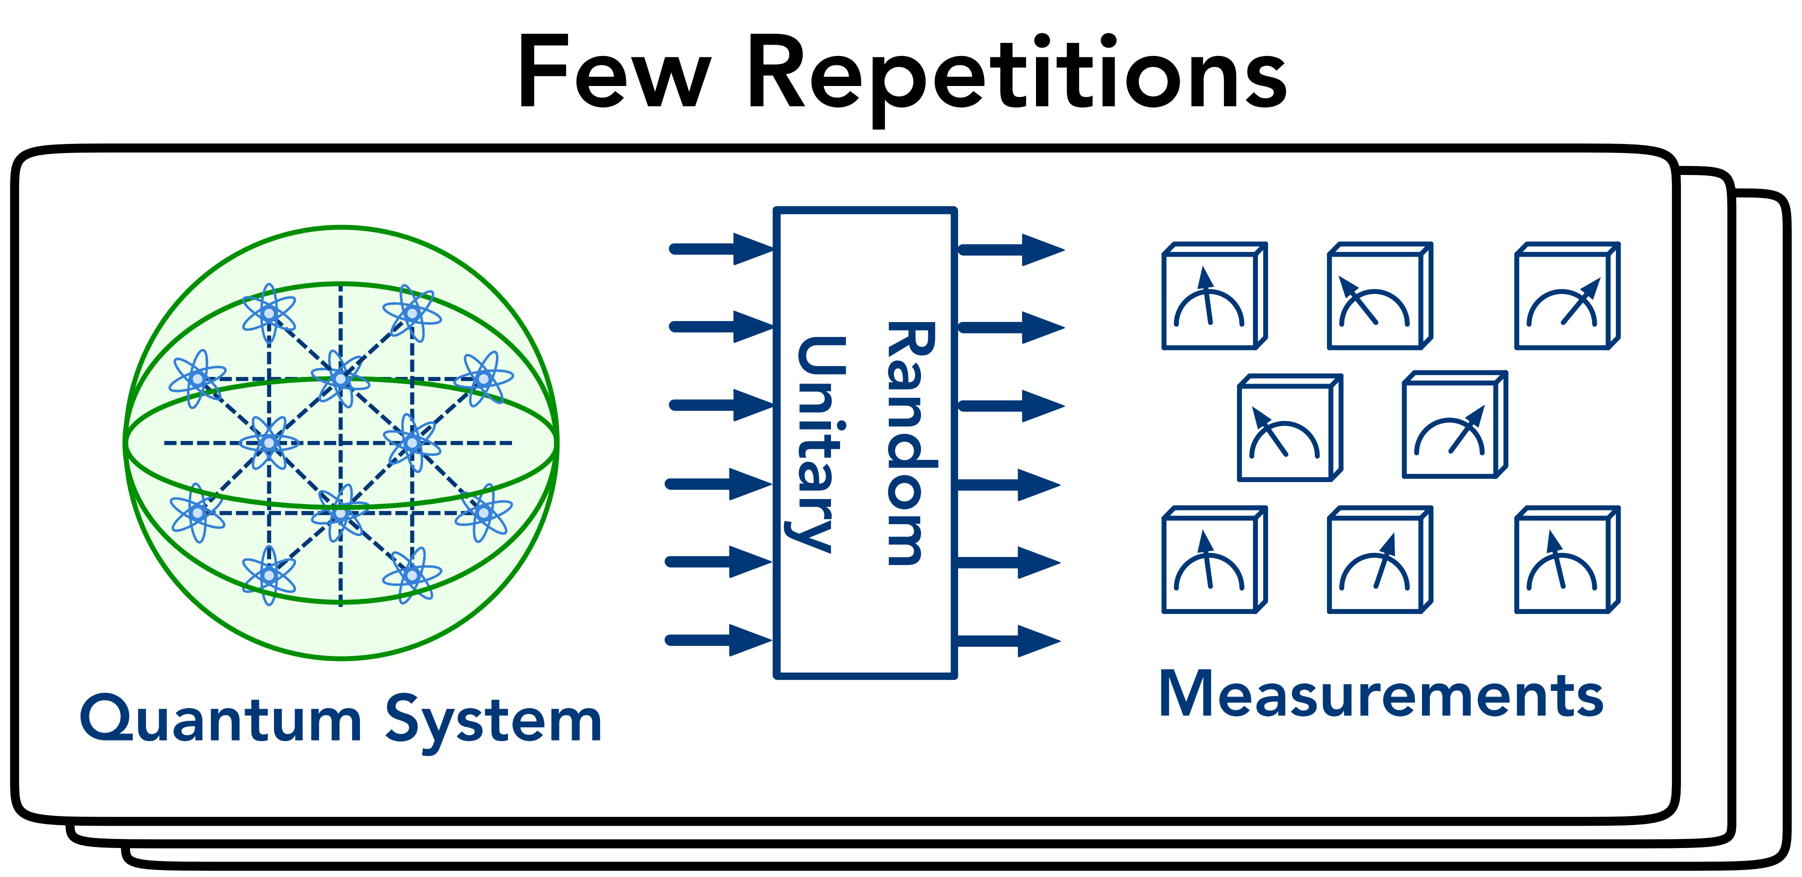

Data Acquisition Phase

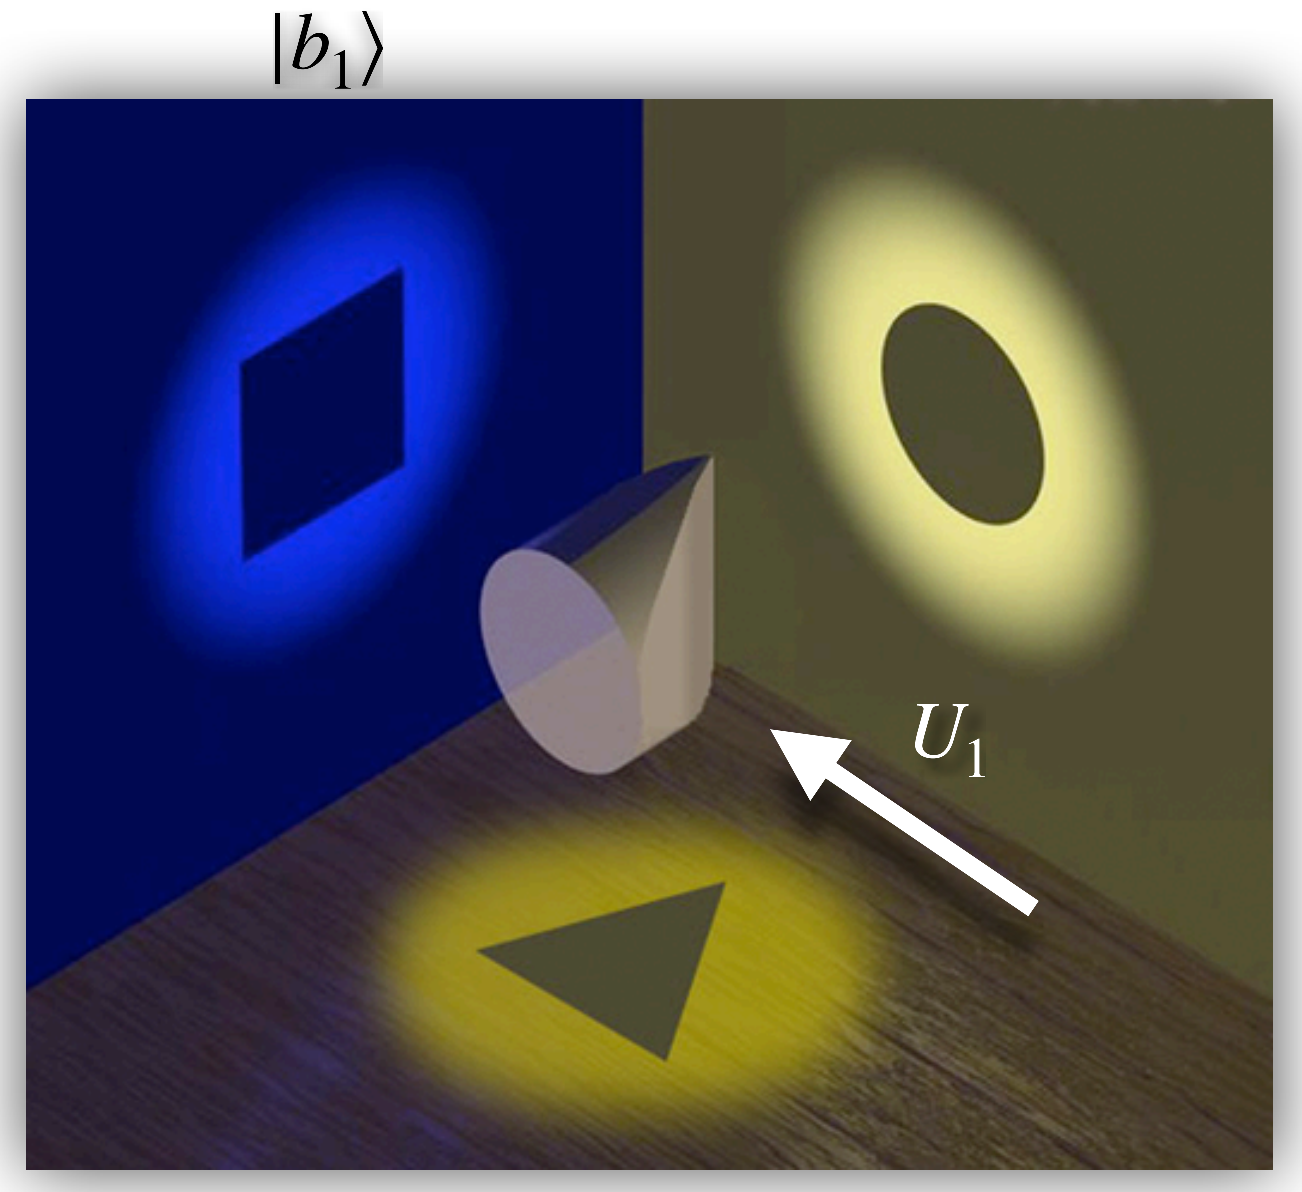

The Procedure: Data Acquisition Phase

Repeat the following T times:

- Sample a random Clifford circuit U_i to evolve the system.
- Measure the system in the computational basis $|b_i\rangle \in \{0,1\}^n$.
- Store the *classical shadow*: $\hat{\sigma}_i = (2^n + 1)U_i^\dagger |b_i\rangle\langle b_i| U_i - I$.

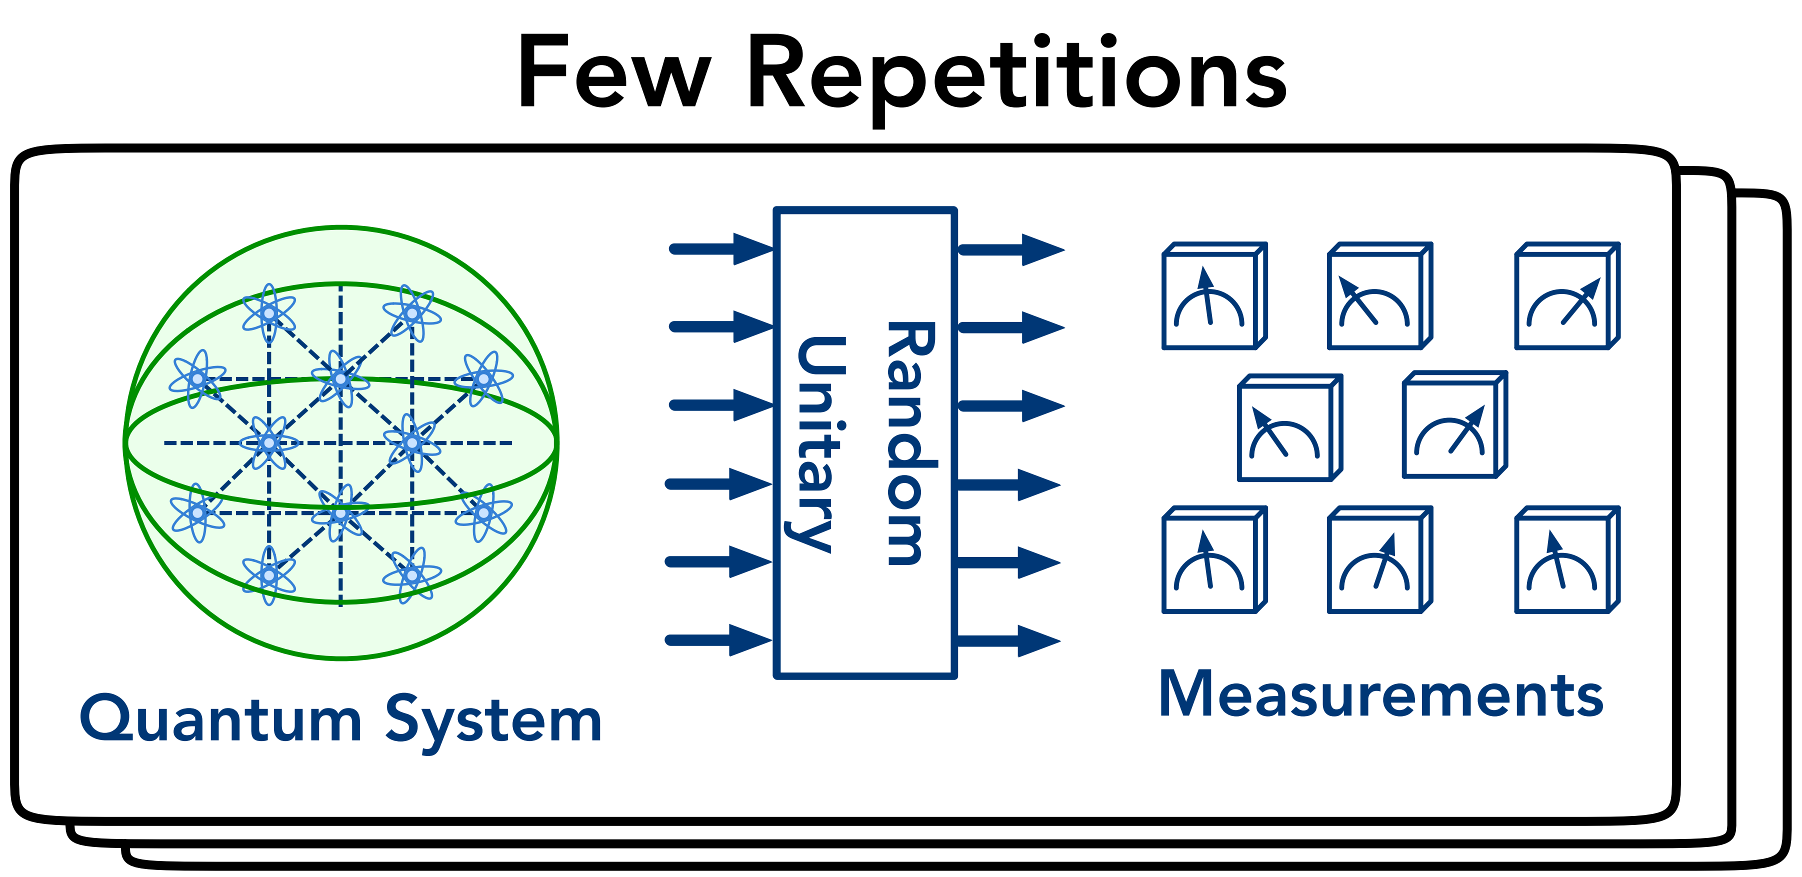

Data Acquisition Phase

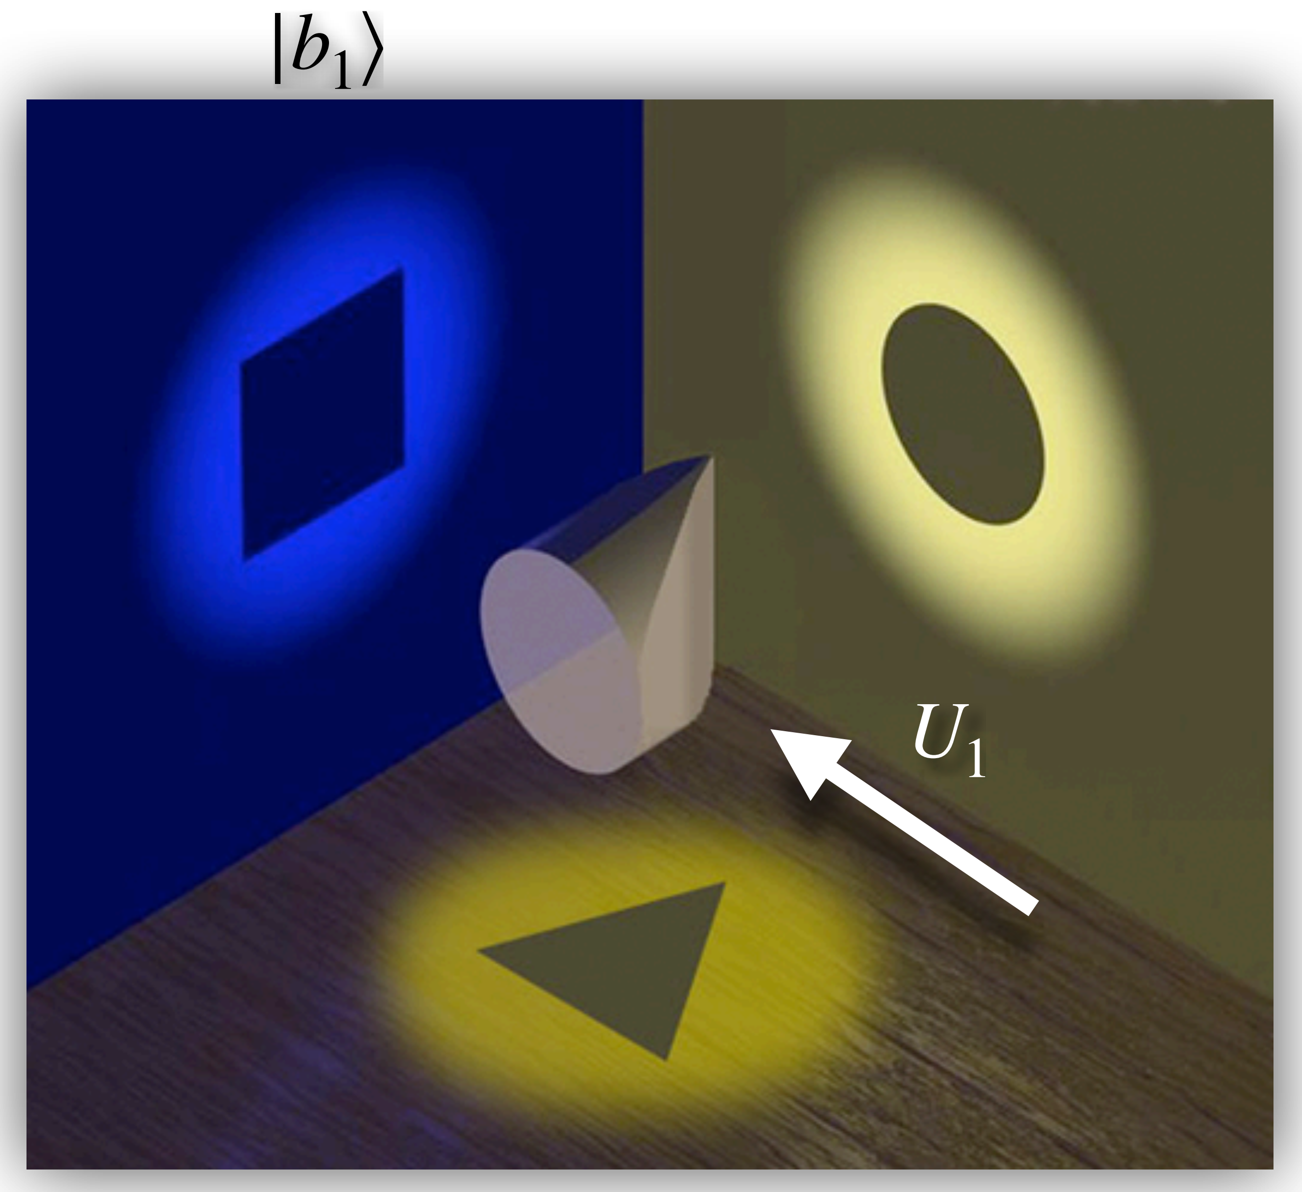

The Procedure:
Data Acquisition Phase

Repeat the following T times:

- Sample a random Clifford circuit U_i to evolve the system.
- Measure the system in the computational basis $|b_i\rangle \in \{0,1\}^n$.
- Store the *classical shadow*: $\hat{\sigma}_i = (2^n + 1)U_i^\dagger |b_i\rangle\langle b_i| U_i - I$.

Only $\mathcal{O}(n^2)$ bits

Few Repetitions

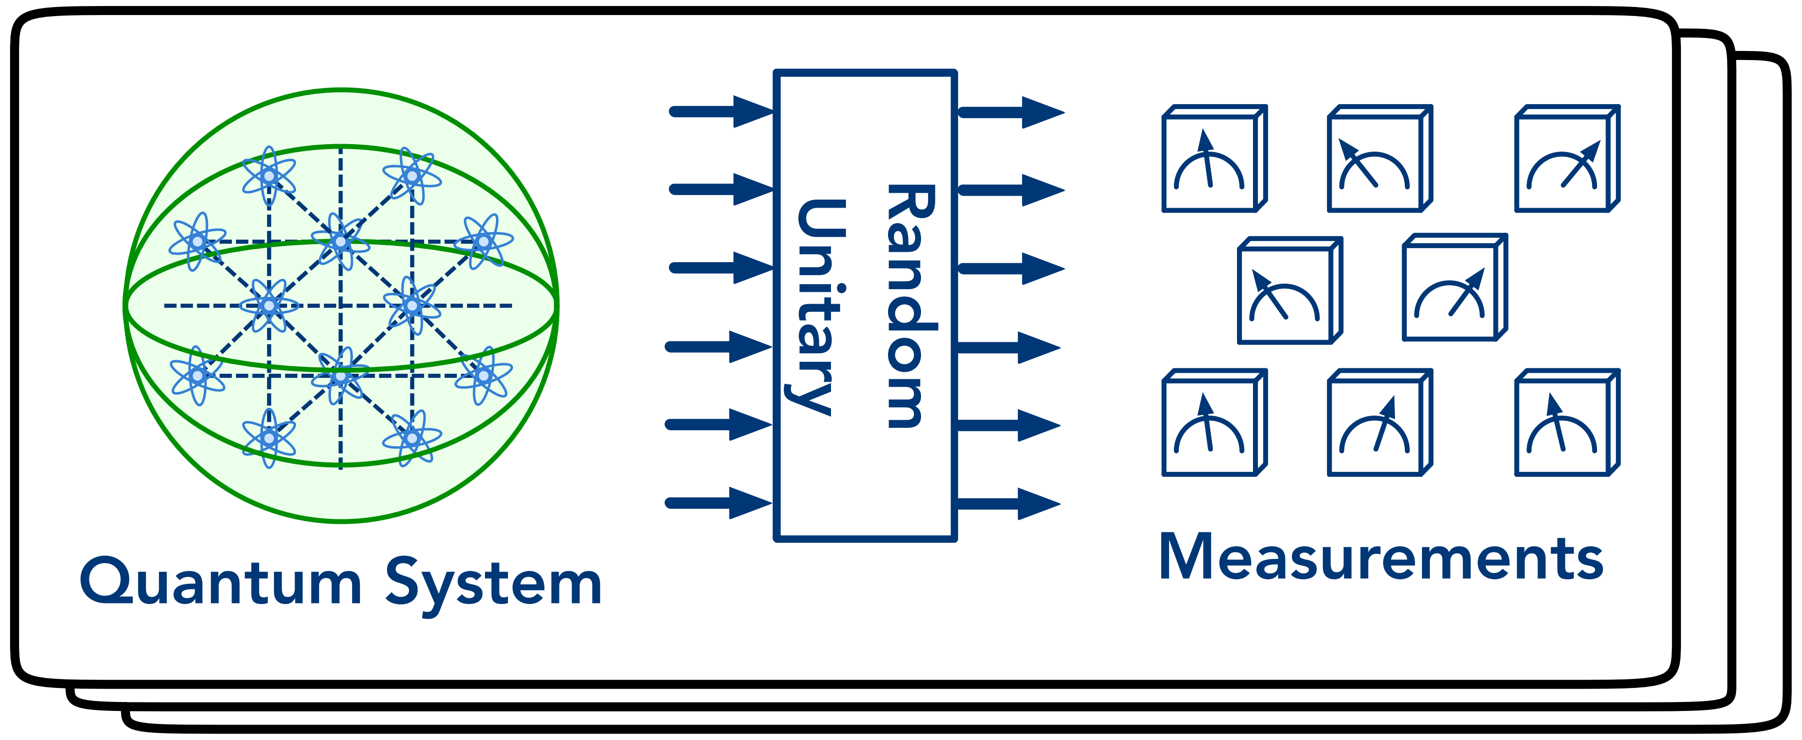

Data Acquisition Phase

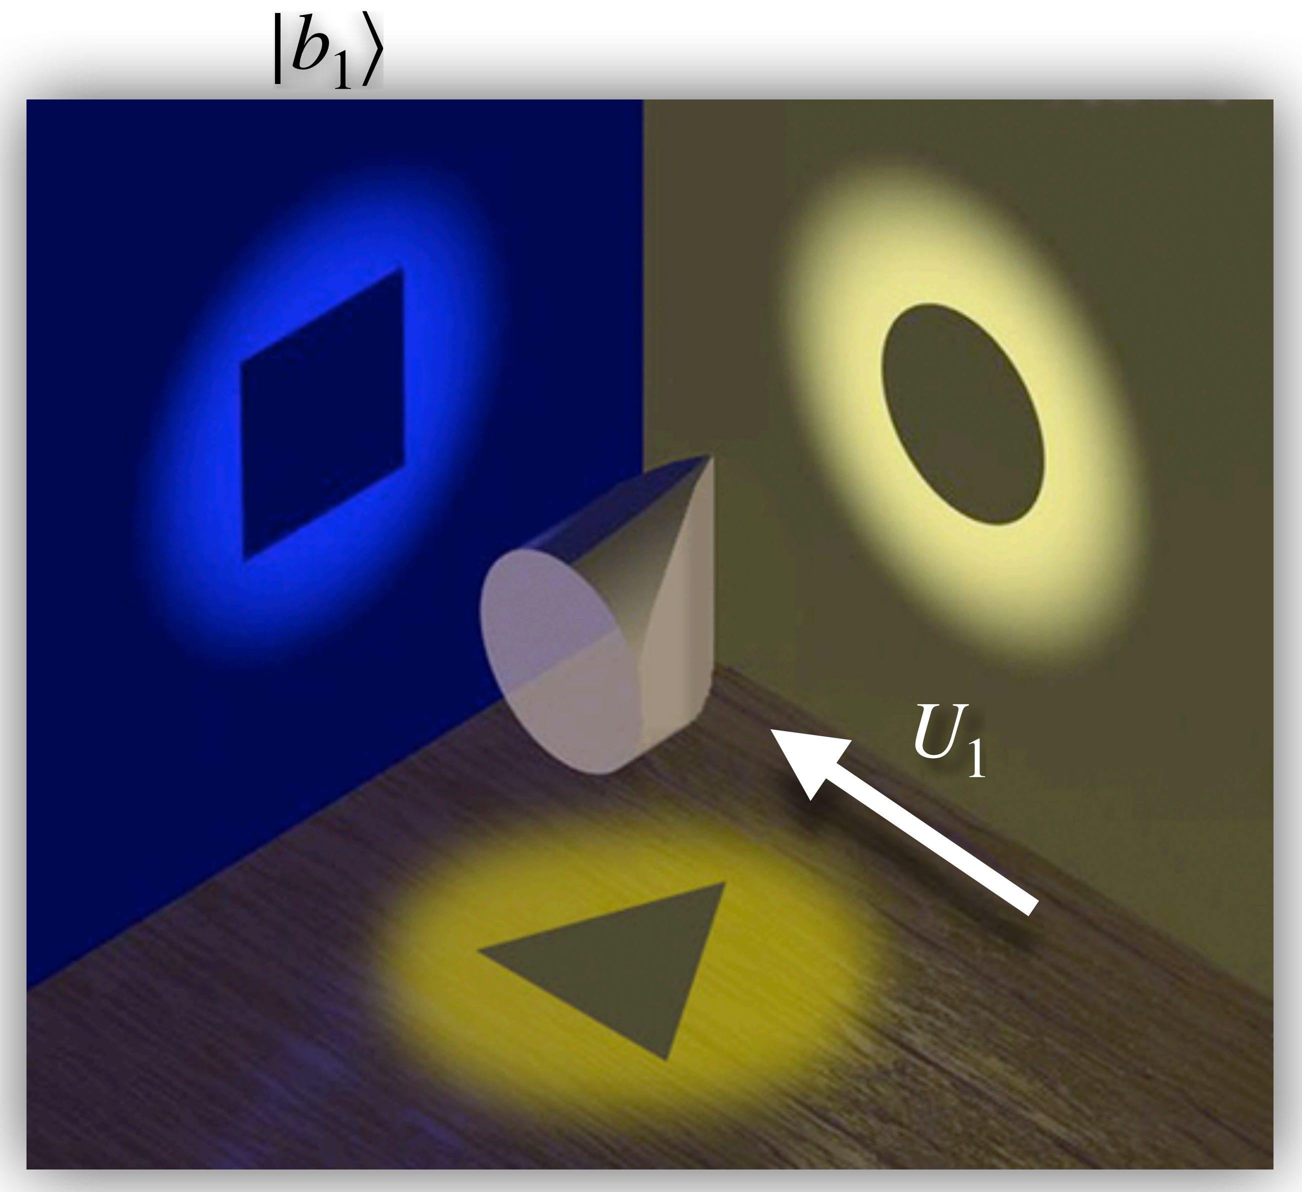

The Procedure:
Data Acquisition Phase

Repeat the following T times:

- Sample a random Clifford circuit U_i to evolve the system.
- Measure the system in the computational basis $|b_i\rangle \in \{0,1\}^n$.
- Store the *classical shadow*: $\hat{\sigma}_i = (2^n + 1)U_i^\dagger|b_i\rangle\langle b_i|U_i - I$.

Only $\mathcal{O}(n^2)$ bits

Few Repetitions

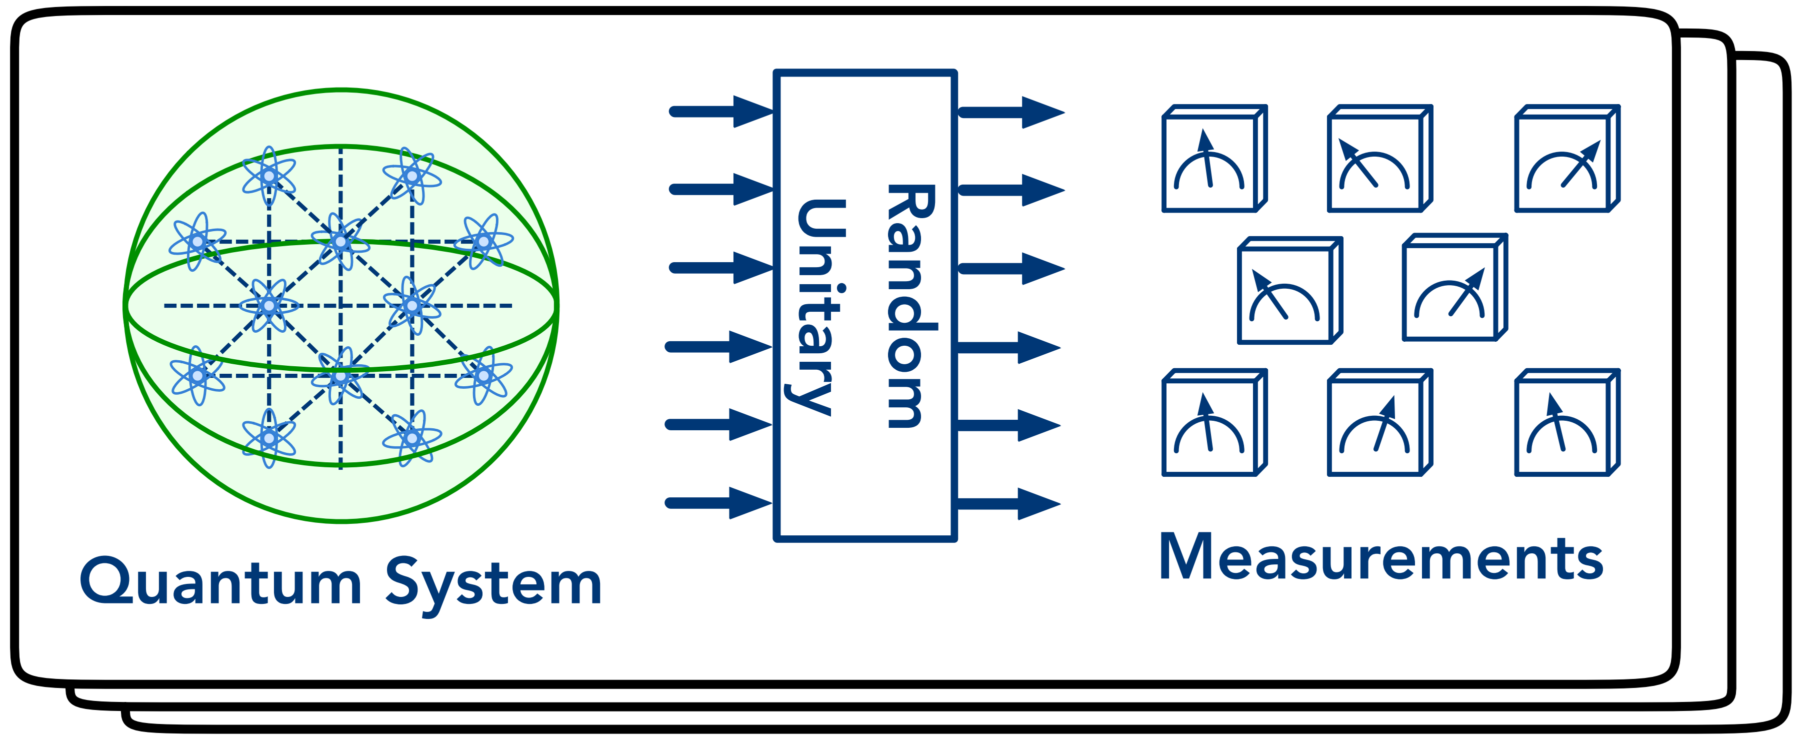

Data Acquisition Phase

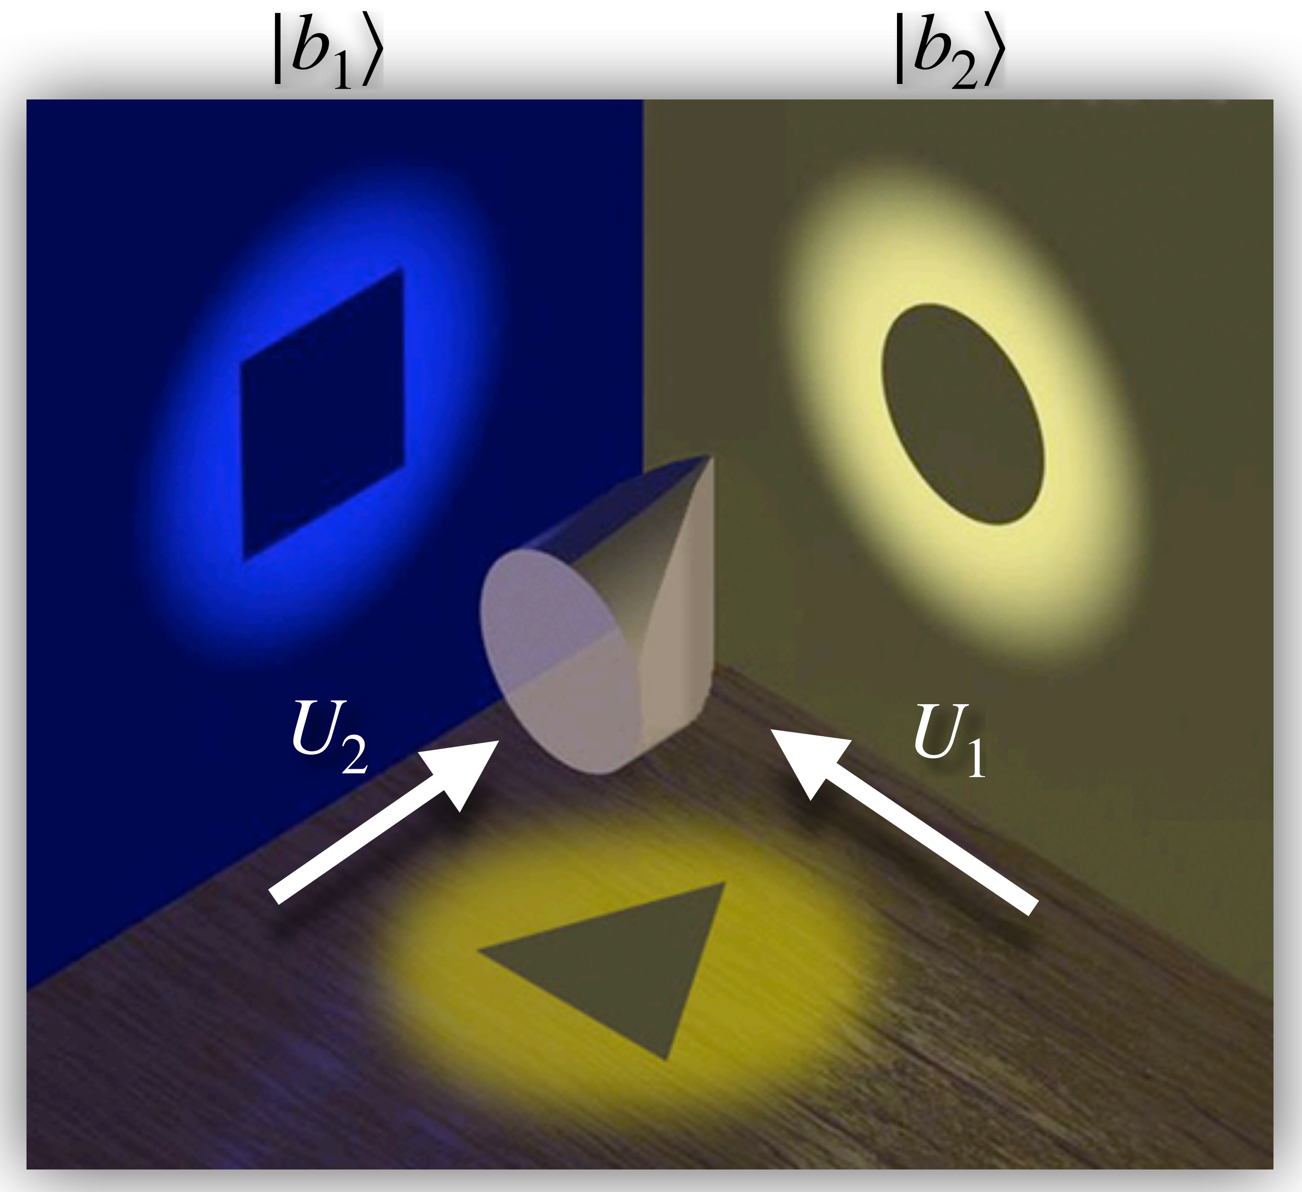

The Procedure:
Data Acquisition Phase

Repeat the following T times:

- Sample a random Clifford circuit U_i to evolve the system.
- Measure the system in the computational basis $|b_i\rangle \in \{0,1\}^n$.
- Store the *classical shadow*: $\hat{\sigma}_i = (2^n + 1)U_i^\dagger |b_i\rangle\langle b_i| U_i - I$.

Only $\mathcal{O}(n^2)$ bits

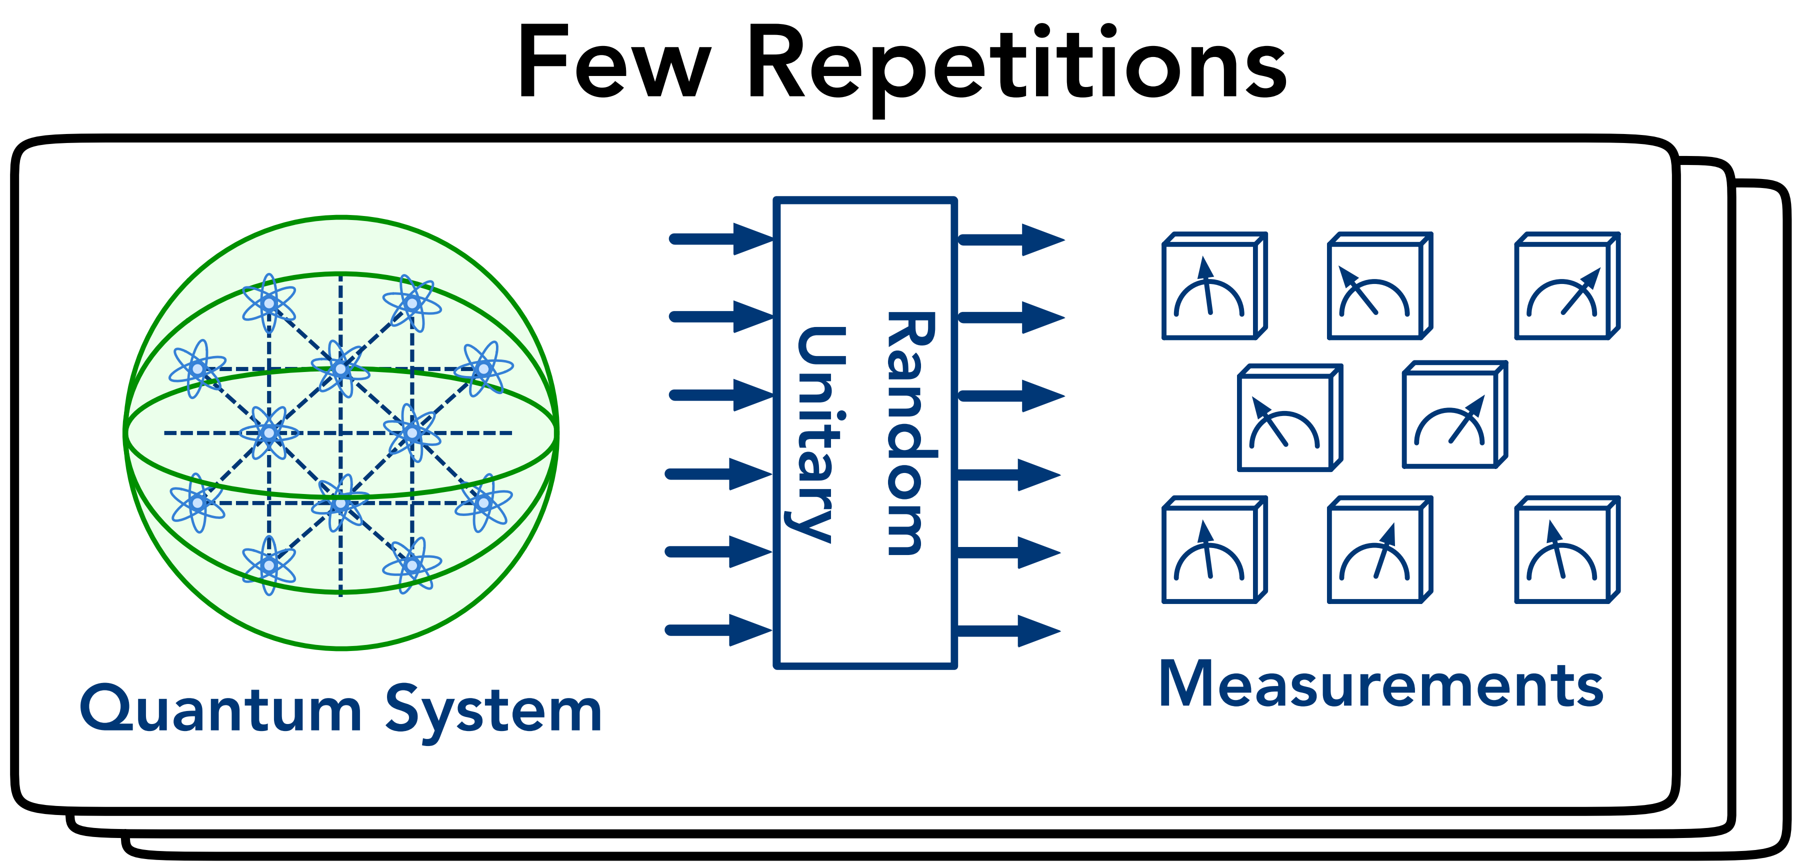

Data Acquisition Phase

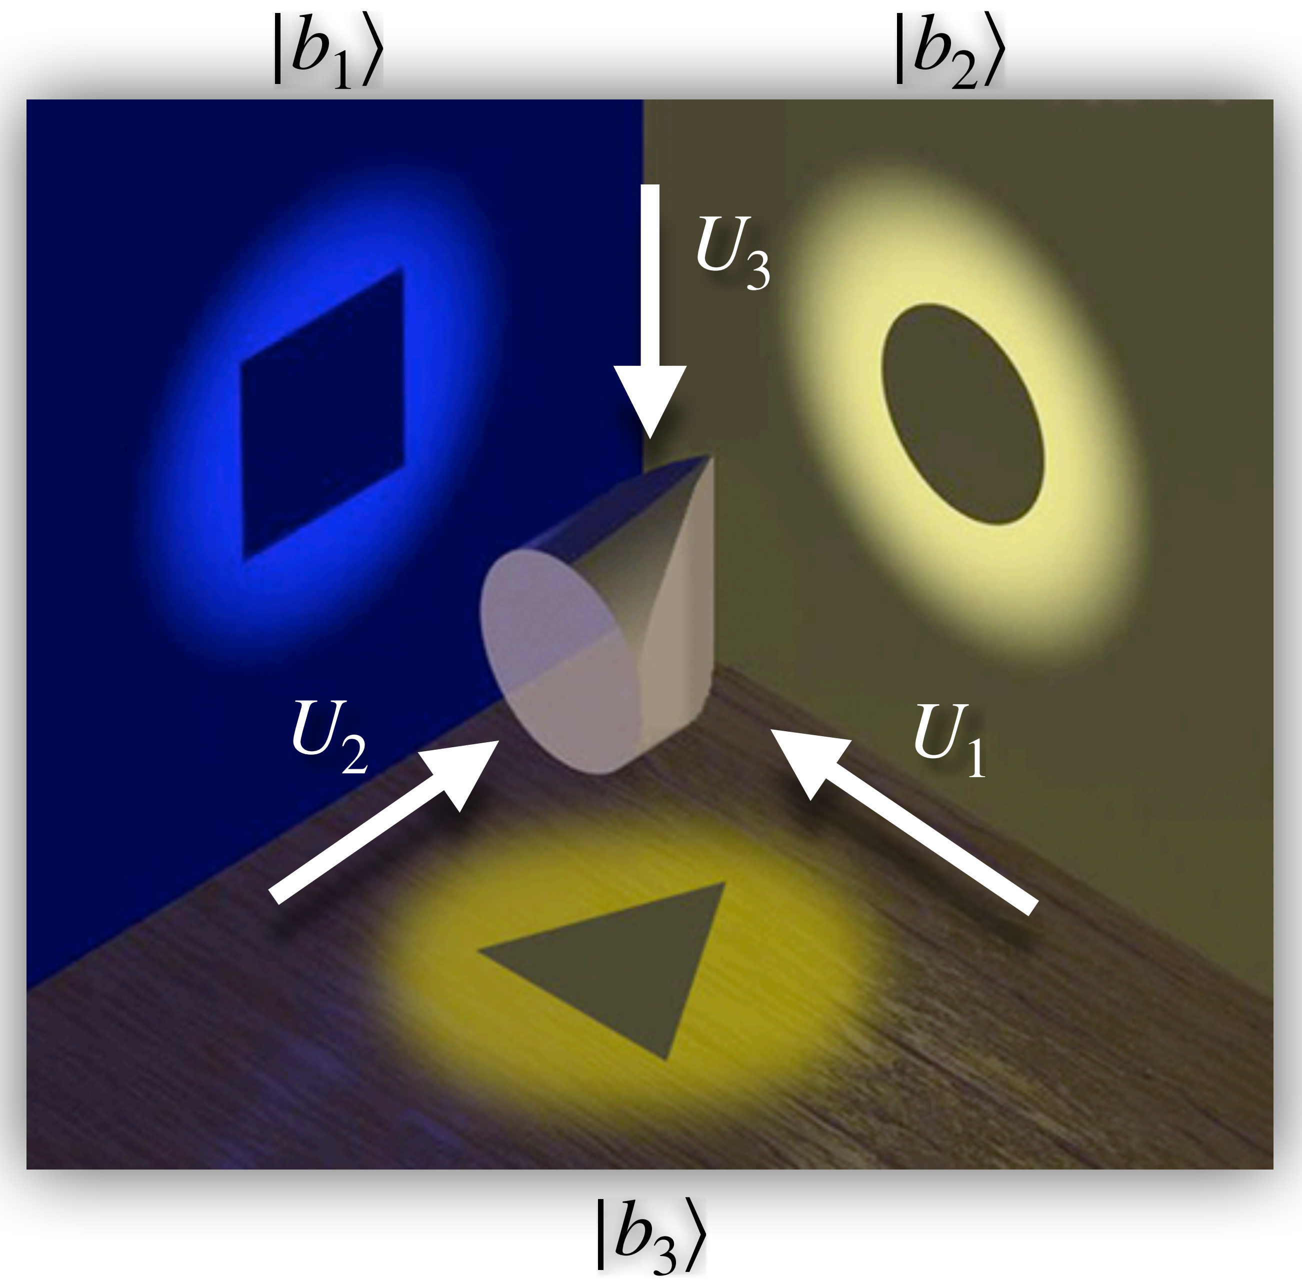

The Procedure: Prediction Phase

Classical shadow
$\hat{\sigma}_i = (2^n + 1)U_i^\dagger |b_i\rangle\langle b_i| U_i - I$

Given $S_T(\rho) = \{\hat{\sigma}_1, \ldots, \hat{\sigma}_T\}$,

how to predict properties of the quantum state ρ?

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

Compute $X_i = \mathrm{tr}(O\hat{\sigma}_i), \forall i = 1, \ldots, T.$

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

Proof Sketch: Moments

Classical shadow

$$\hat{\sigma}_i = (2^n + 1)U_i^\dagger |b_i\rangle\langle b_i| U_i - I$$

- 1st moment $\mathbb{E}[\hat{\sigma}_i]$ corresponds to the 4th moment of random Clifford circuit U_i.

$$\mathbb{E}[\hat{\sigma}_i] = \mathbb{E}_{U_i} \sum_{b_i \in \{0,1\}^n} \langle b_i | U_i \rho U_i^\dagger | b_i \rangle \left[(2^n + 1) U_i^\dagger |b_i\rangle\langle b_i| U_i - I \right]$$

- 2nd moment $\mathbb{E}[\hat{\sigma}_i \otimes \hat{\sigma}_i]$ corresponds to the 6th moment of random Clifford circuit U_i.

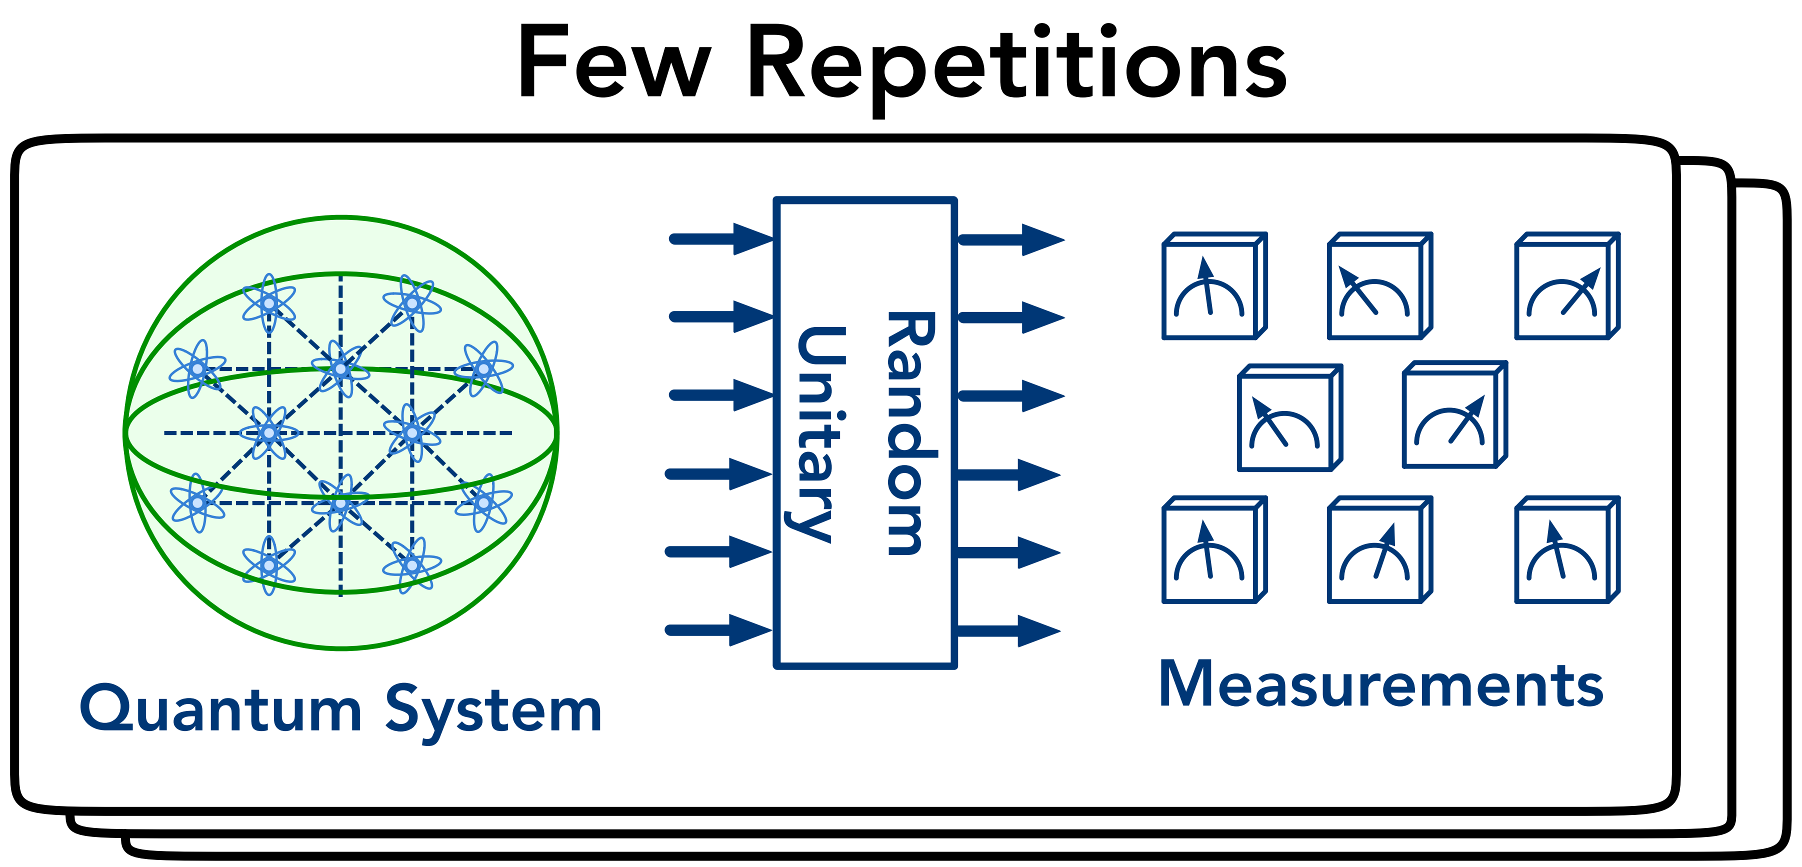

Data Acquisition Phase

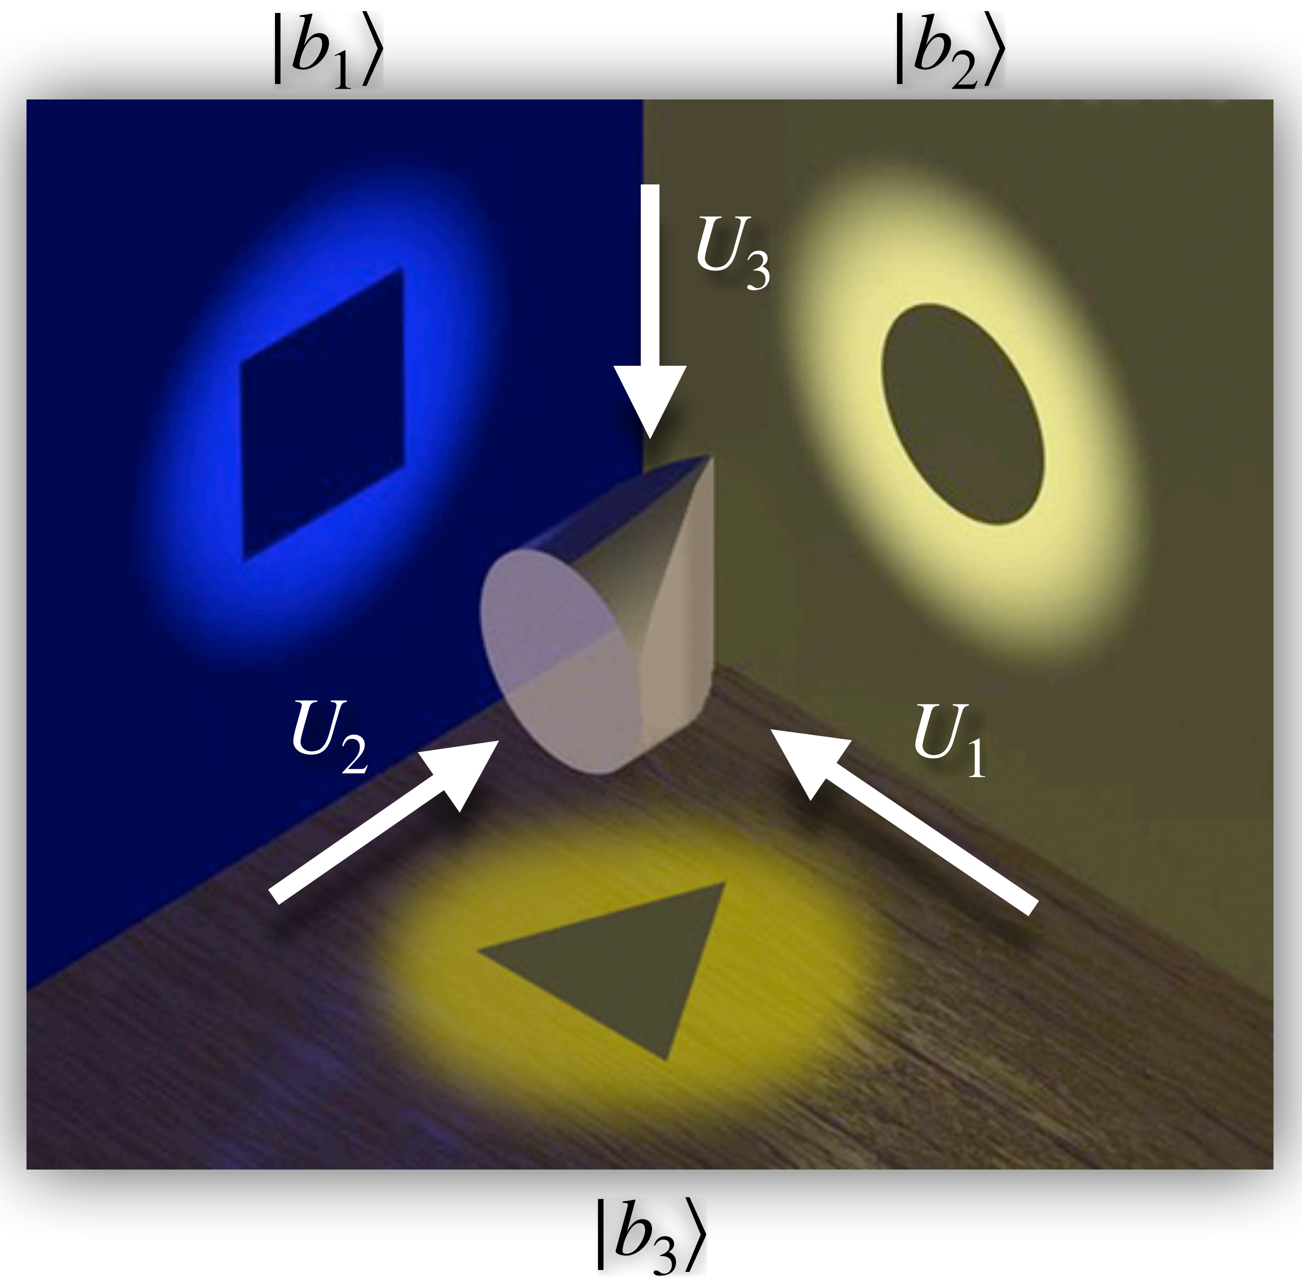

Proof Sketch: Weingarten Calculus

Classical shadow
$\hat{\sigma}_i = (2^n + 1)U_i^\dagger |b_i\rangle\langle b_i| U_i - I$

- Weingarten calculus fully characterizes moments of random $\exp(n)$-size circuit.

  The moments of $\mathrm{SU}(2^n)$ correspond to the symmetric group.

- Random Clifford circuit matches random $\exp(n)$-size circuit up to the 6th moments.

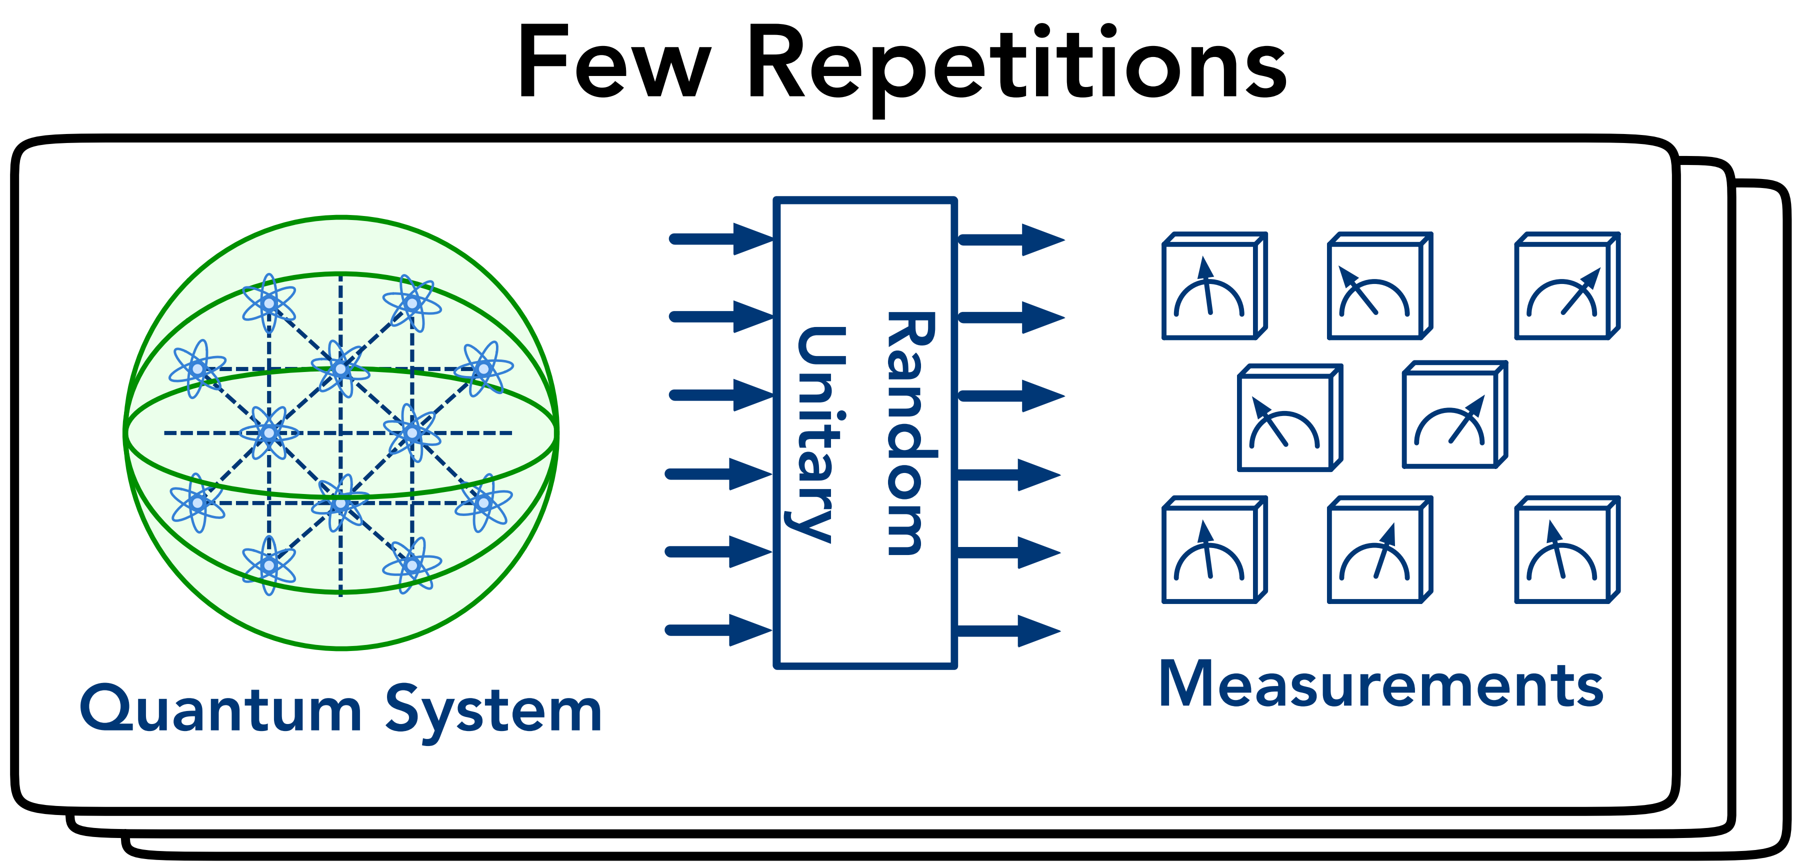

Data Acquisition Phase

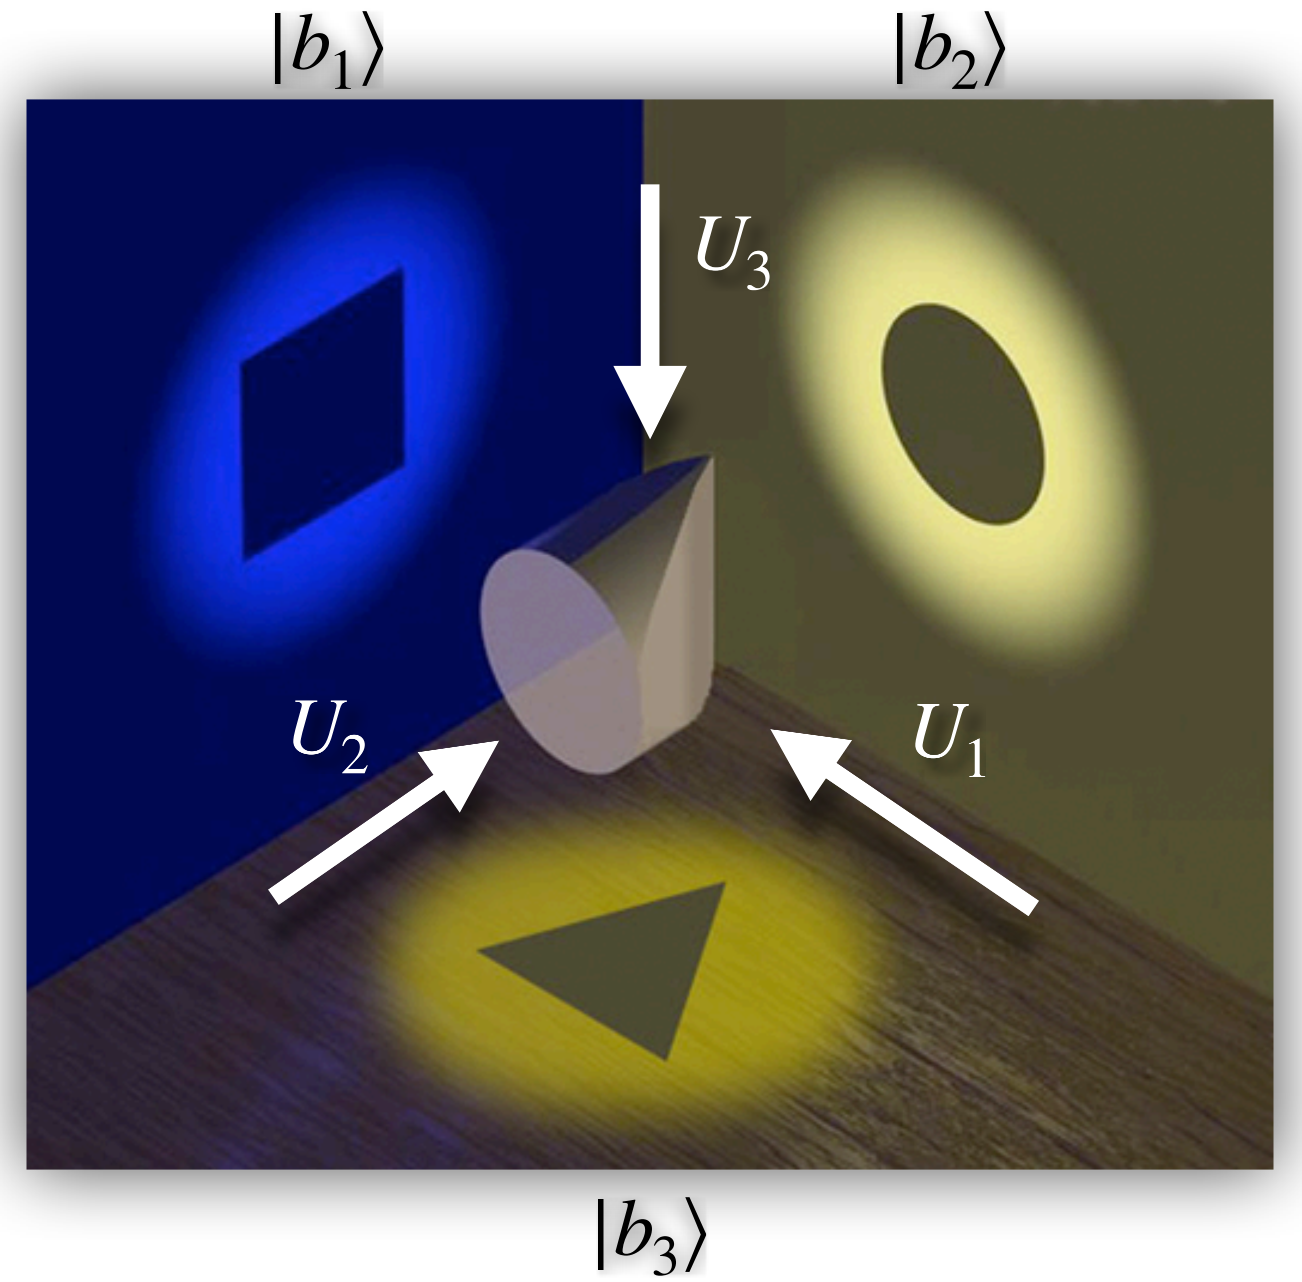

Proof Sketch: Concentration

Classical shadow
$\hat{\sigma}_i = (2^n + 1)U_i^\dagger |b_i\rangle\langle b_i| U_i - I$

- 1st moment $\mathbb{E}[\hat{\sigma}_i] = \rho$. So $\hat{\sigma}_i \approx \rho$ up to random fluctuations.
- 2nd moment $\mathbb{E}[\hat{\sigma}_i \otimes \hat{\sigma}_i]$ is well-controlled despite the $(2^n + 1)$ factor.
- Median-of-means estimator only cares about the first two moments.

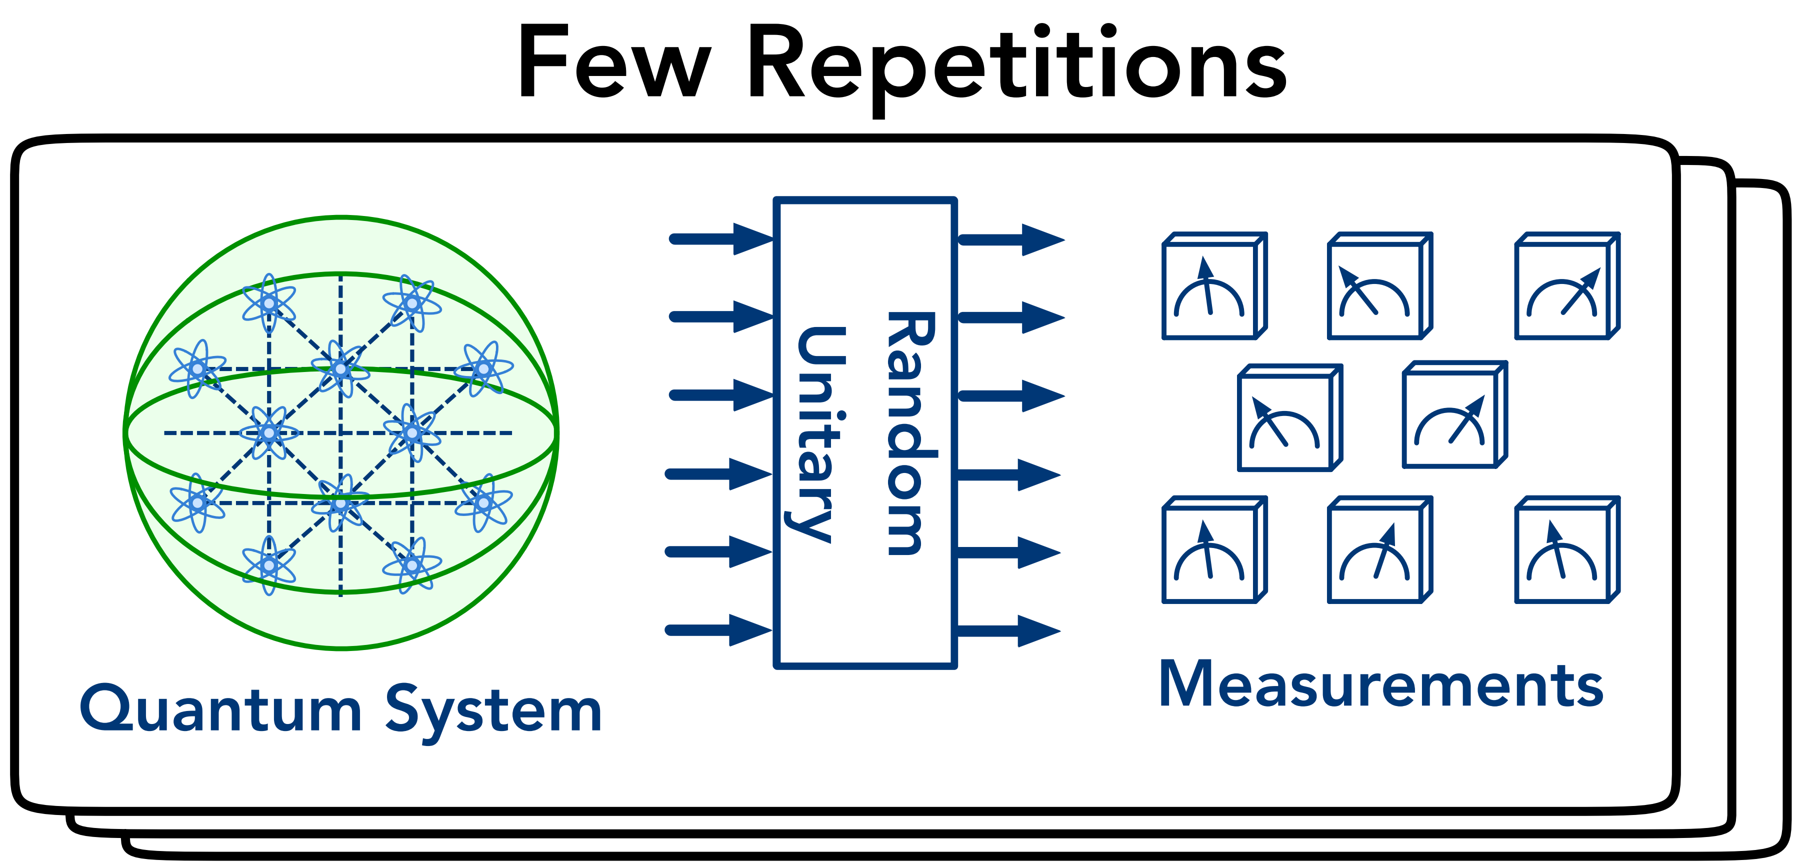

Data Acquisition Phase

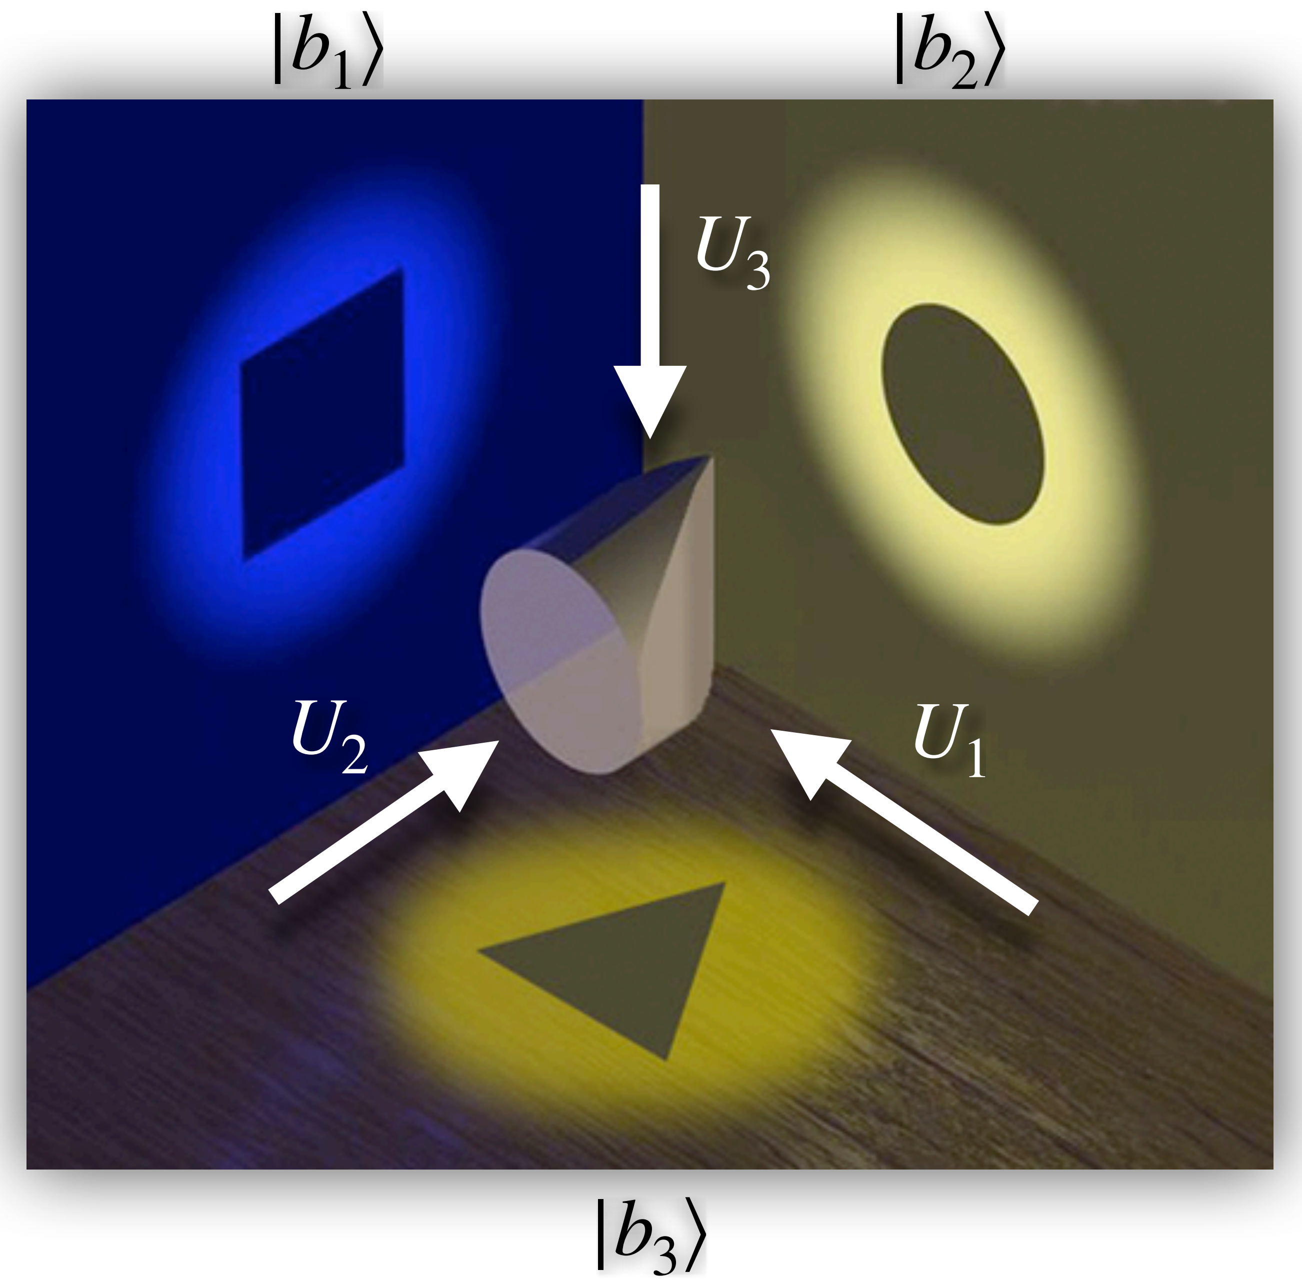

Classical shadow formalism

Theorem (Huang et al.; 2020)

We can predict any $O_1, \ldots, O_M$ with $B \geq \mathrm{Tr}(O_i^2)$ to ϵ error from $T = \mathcal{O}(B \log(M)/\epsilon^2)$ measurements.

[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

Theorem (Huang et al.; 2020)

We can predict any $O_1, \ldots, O_M$ with $B \geq \mathrm{Tr}(O_i^2)$ to ϵ error from $T = \mathcal{O}(B \log(M)/\epsilon^2)$ measurements.

Q: Could this be further improved?

Theorem (Huang et al.; 2020)

To predict any $O_1, \ldots, O_M$ with $B \geq \mathrm{Tr}(O_i^2)$ to ϵ error, we need $T = \Omega(B \log(M)/\epsilon^2)$ measurements.

- Proved by relating to quantum communication tasks.
- This lower bound applies *only* to learning using classical machines.

[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

Theorem (Huang et al.; 2020)

We can predict any $O_1, \ldots, O_M$ with $B \geq \|O_i\|^2_{\text{shadow}}$ to ϵ error from $T = \mathcal{O}(B \log(M)/\epsilon^2)$ measurements.

$$\|O\|_{\text{shadow}} = \max_{\sigma \in S_{2n}} \left(\mathbb{E}_{U \sim \mathcal{U}} \sum_{b \in \{0,1\}^n} \langle b | U \sigma U^\dagger | b \rangle \langle b | U \mathcal{M}^{-1}(O) U^\dagger | b \rangle^2 \right)^{1/2}.$$

Random Clifford Unitary
$U \in \mathrm{Cl}(2^n)$

$$\mathcal{M}_n(\rho) = (\rho + I)/(2^n + 1)$$
$$\mathcal{M}_n^{-1}(X) = (2^n + 1)X - I$$

Random Clifford Circuits

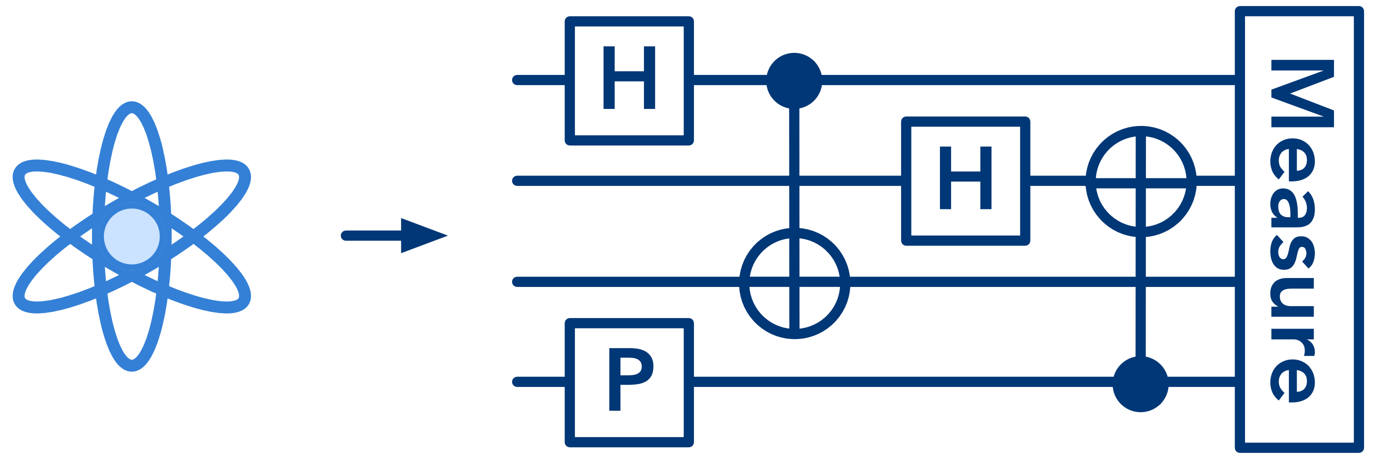

Classical shadow formalism

Theorem (Huang et al.; 2020)

We can predict any $O_1, \ldots, O_M$ with $B \geq \|O_i\|^2_{\text{shadow}}$ to ϵ error from $T = \mathcal{O}(B \log(M)/\epsilon^2)$ measurements.

$$\|O\|_{\text{shadow}} = \max_{\sigma \in S_{2n}} \left(\mathbb{E}_{U \sim \mathcal{U}} \sum_{b \in \{0,1\}^n} \langle b \,|\, U\sigma U^\dagger \,|\, b\rangle \langle b \,|\, U \mathcal{M}^{-1}(O)\, U^\dagger \,|\, b\rangle^2 \right)^{1/2}.$$

Random Clifford Unitary

$$U \in \mathrm{Cl}(2^n)$$

$$\mathcal{M}_n(\rho) = (\rho + I)/(2^n + 1)$$
$$\mathcal{M}_n^{-1}(X) = (2^n + 1)X - I$$

The corresponding norm:

$$\|O\|_{\text{shadow}} \leq \sqrt{\mathrm{tr}(O^2)}$$

Classical shadow formalism

Theorem (Huang et al.; 2020)

We can predict any $O_1, \ldots, O_M$ with $B \geq \|O_i\|^2_{\text{shadow}}$ to ϵ error from $T = \mathcal{O}(B \log(M)/\epsilon^2)$ measurements.

$$\|O\|_{\text{shadow}} = \max_{\sigma \in S_{2n}} \left(\mathbb{E}_{U \sim \mathcal{U}} \sum_{b \in \{0,1\}^n} \langle b \,|\, U \sigma U^\dagger \,|\, b \rangle \langle b \,|\, U \mathcal{M}^{-1}(O)\, U^\dagger \,|\, b \rangle^2 \right)^{1/2}.$$

Random Clifford Unitary

$U \in \mathrm{Cl}(2^n)$

$$\mathcal{M}_n(\rho) = (\rho + I)/(2^n + 1)$$
$$\mathcal{M}_n^{-1}(X) = (2^n + 1)X - I$$

The corresponding norm:

$$\|O\|_{\text{shadow}} \leq \sqrt{\mathrm{tr}(O^2)}$$

Application:

Quantum fidelity $|\psi\rangle\langle\psi|$

Classical shadow formalism

Theorem (Huang et al.; 2020)

We can predict any $O_1, \ldots, O_M$ with $B \geq \|O_i\|^2_{\text{shadow}}$ to ϵ error from $T = \mathcal{O}(B \log(M)/\epsilon^2)$ measurements.

$$\|O\|_{\text{shadow}} = \max_{\sigma \in S_{2n}} \left(\mathbb{E}_{U \sim \mathcal{U}} \sum_{b \in \{0,1\}^n} \langle b | U \sigma U^\dagger | b \rangle \langle b | U \mathcal{M}^{-1}(O) U^\dagger | b \rangle^2 \right)^{1/2}.$$

Random Local Clifford

$U \in \mathrm{Cl}(2)^{\otimes n}$

$$\mathcal{M} = \otimes_{i=1}^n \mathcal{M}_1$$
$$\mathcal{M}^{-1} = \otimes_{i=1}^n \mathcal{M}_1^{-1}$$

Random
Pauli Measurement

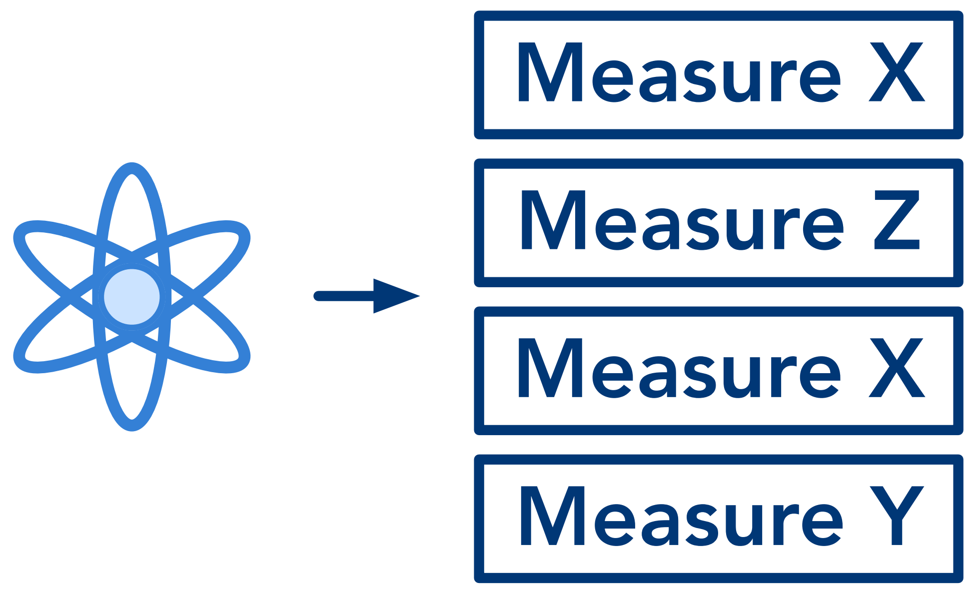

Classical shadow formalism

Theorem (Huang et al.; 2020)

We can predict any $O_1, \ldots, O_M$ with $B \geq \|O_i\|^2_{\text{shadow}}$ to ϵ error from $T = \mathcal{O}(B \log(M)/\epsilon^2)$ measurements.

$$\|O\|_{\text{shadow}} = \max_{\sigma \in S_{2n}} \left(\mathbb{E}_{U \sim \mathcal{U}} \sum_{b \in \{0,1\}^n} \langle b \,|\, U \sigma U^\dagger \,|\, b \rangle \langle b \,|\, U \mathcal{M}^{-1}(O)\, U^\dagger \,|\, b \rangle^2 \right)^{1/2}.$$

Random Local Clifford

$U \in \mathrm{Cl}(2)^{\otimes n}$

$\mathcal{M} = \otimes_{i=1}^n \mathcal{M}_1$

$\mathcal{M}^{-1} = \otimes_{i=1}^n \mathcal{M}_1^{-1}$

The corresponding norm:

$\|O\|_{\text{shadow}} \leq 2^k \|O\|_\infty$

Observable O acts on k qubits

Classical shadow formalism

Theorem (Huang et al.; 2020)

We can predict any $O_1, \ldots, O_M$ with $B \geq \|O_i\|^2_{\text{shadow}}$ to ϵ error from $T = \mathcal{O}(B \log(M)/\epsilon^2)$ measurements.

$$\|O\|_{\text{shadow}} = \max_{\sigma \in S_{2n}} \left(\mathbb{E}_{U \sim \mathcal{U}} \sum_{b \in \{0,1\}^n} \langle b \,|\, U \sigma U^\dagger \,|\, b \rangle \langle b \,|\, U \mathcal{M}^{-1}(O)\, U^\dagger \,|\, b \rangle^2 \right)^{1/2}.$$

Random Local Clifford

$U \in \mathrm{Cl}(2)^{\otimes n}$

$\mathcal{M} = \otimes_{i=1}^{n} \mathcal{M}_1$

$\mathcal{M}^{-1} = \otimes_{i=1}^{n} \mathcal{M}_1^{-1}$

The corresponding norm:

$\|O\|_{\text{shadow}} \leq 2^k \|O\|_\infty$

Observable O acts on k qubits

Application:
2-point correlation,
local Hamiltonian.

Classical shadow formalism

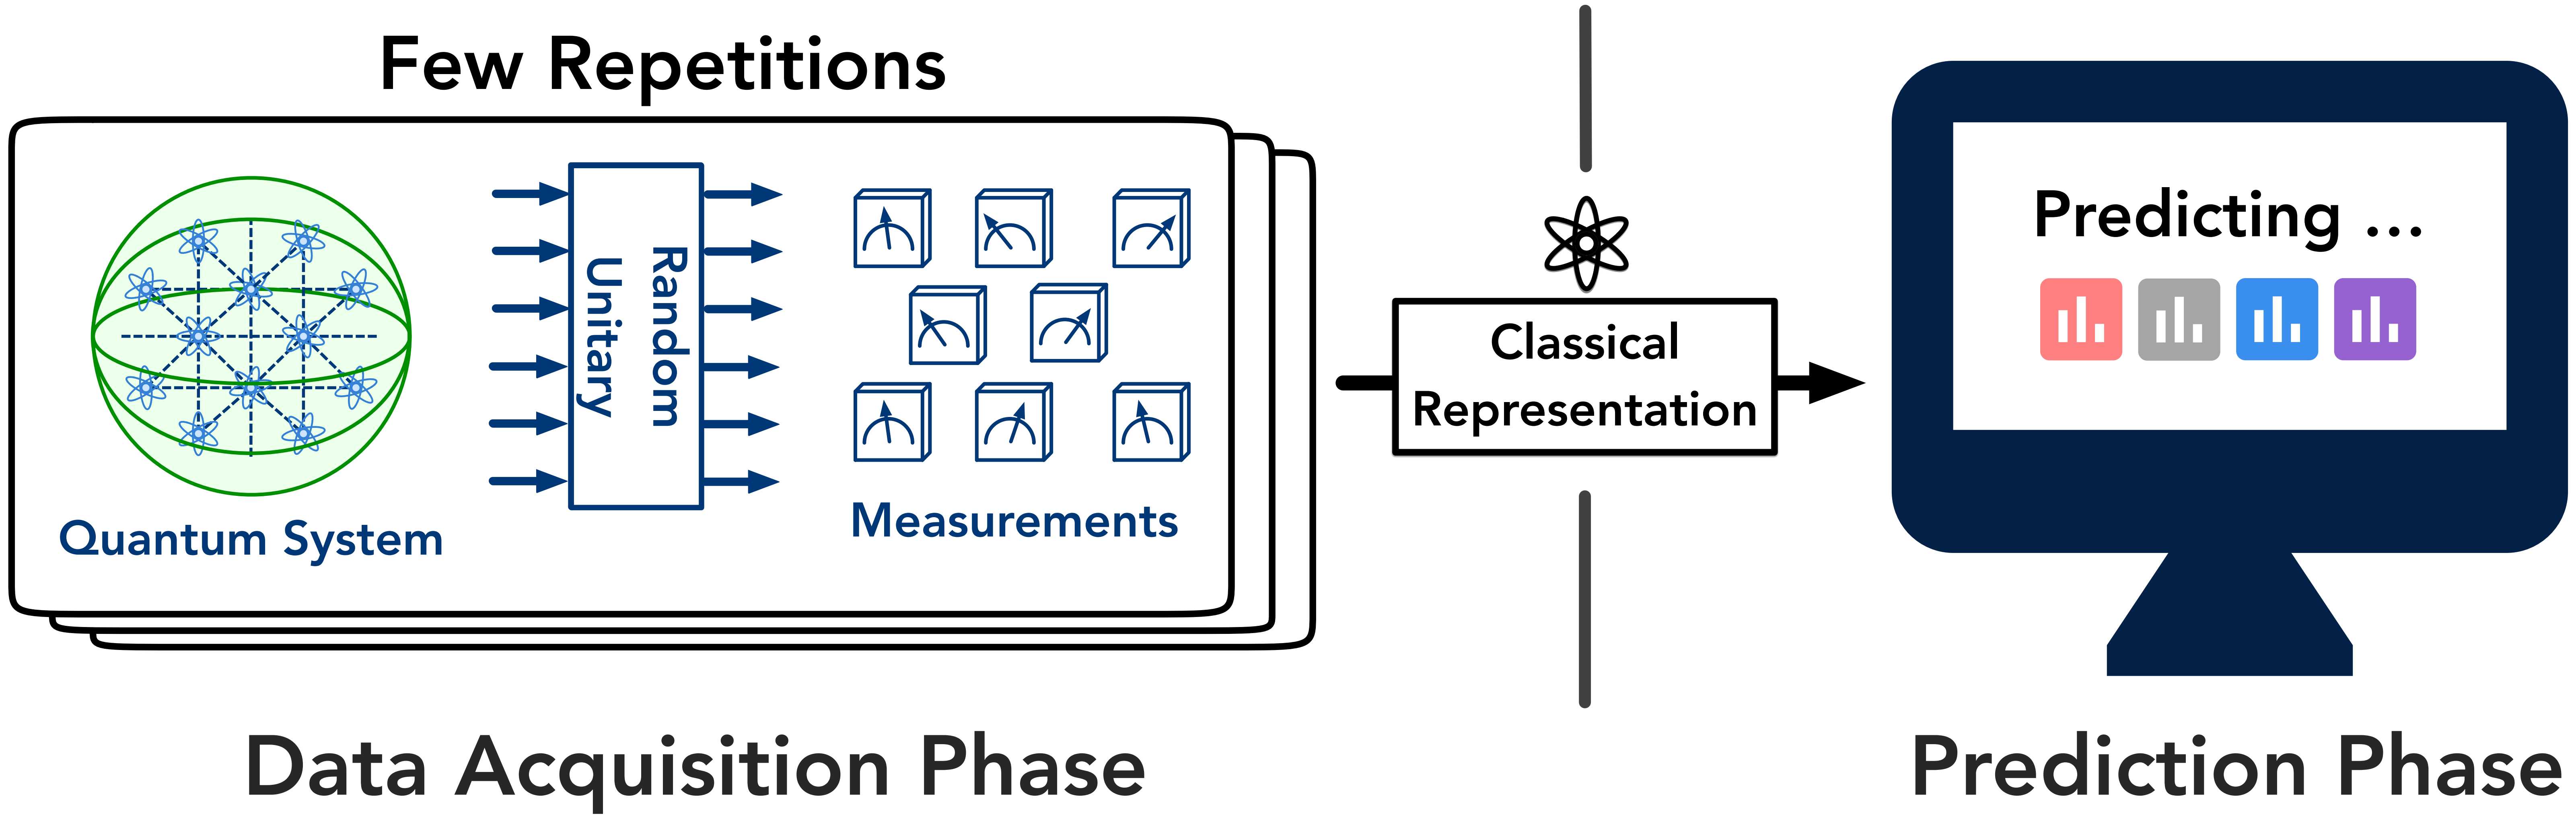

Possible Properties

- 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.

Applications of classical shadows

Benchmarking quantum systems

The ability to estimate quantum fidelity and verify entanglement enables efficient benchmarking.

[a] Huang et al. *Nature Physics* (2020).
[b] Elben et al. *Physical Review Letters* (2020).

Applications of classical shadows

Benchmarking quantum systems

The ability to estimate quantum fidelity and verify entanglement enables efficient benchmarking.

[a] Huang et al. *Nature Physics* (2020).
[b] Elben et al. *Physical Review Letters* (2020).

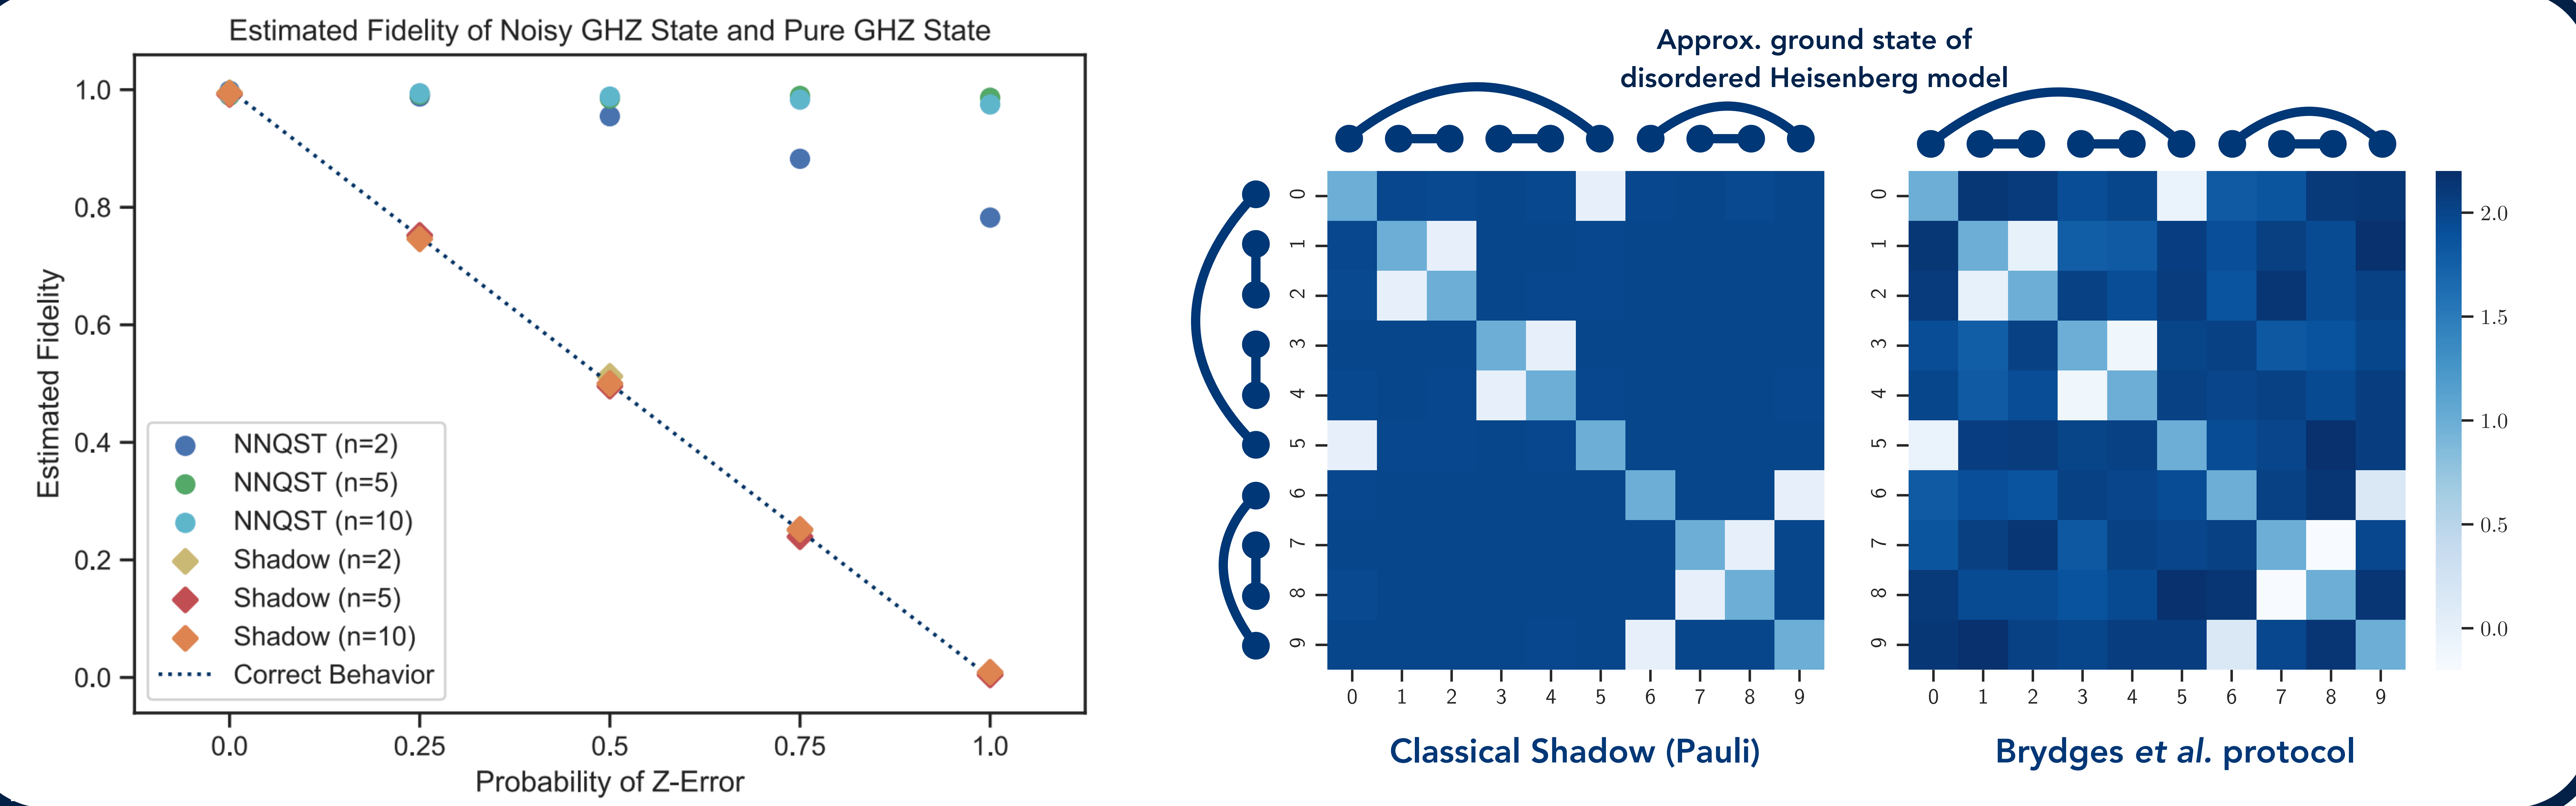

Applications of classical shadows

Benchmarking quantum systems

The ability to estimate quantum fidelity and verify entanglement enables efficient benchmarking.

[a] Huang et al. *Nature Physics* (2020).
[b] Elben et al. *Physical Review Letters* (2020).

Quantum chemistry simulation

Quantum simulation algorithms often need to estimate many properties (e.g., local observables, Hamiltonian).

[c] Zhao et al. *Physical Review Letters* (2021).
[d] Huggins et al. *Nature* (2022).

Applications of classical shadows

Benchmarking quantum systems

The ability to estimate quantum fidelity and verify entanglement enables efficient benchmarking.

[a] Huang et al. *Nature Physics* (2020).
[b] Elben et al. *Physical Review Letters* (2020).

Quantum chemistry simulation

Quantum simulation algorithms often need to estimate many properties (e.g., local observables, Hamiltonian).

[c] Zhao et al. *Physical Review Letters* (2021).
[d] Huggins et al. *Nature* (2022).

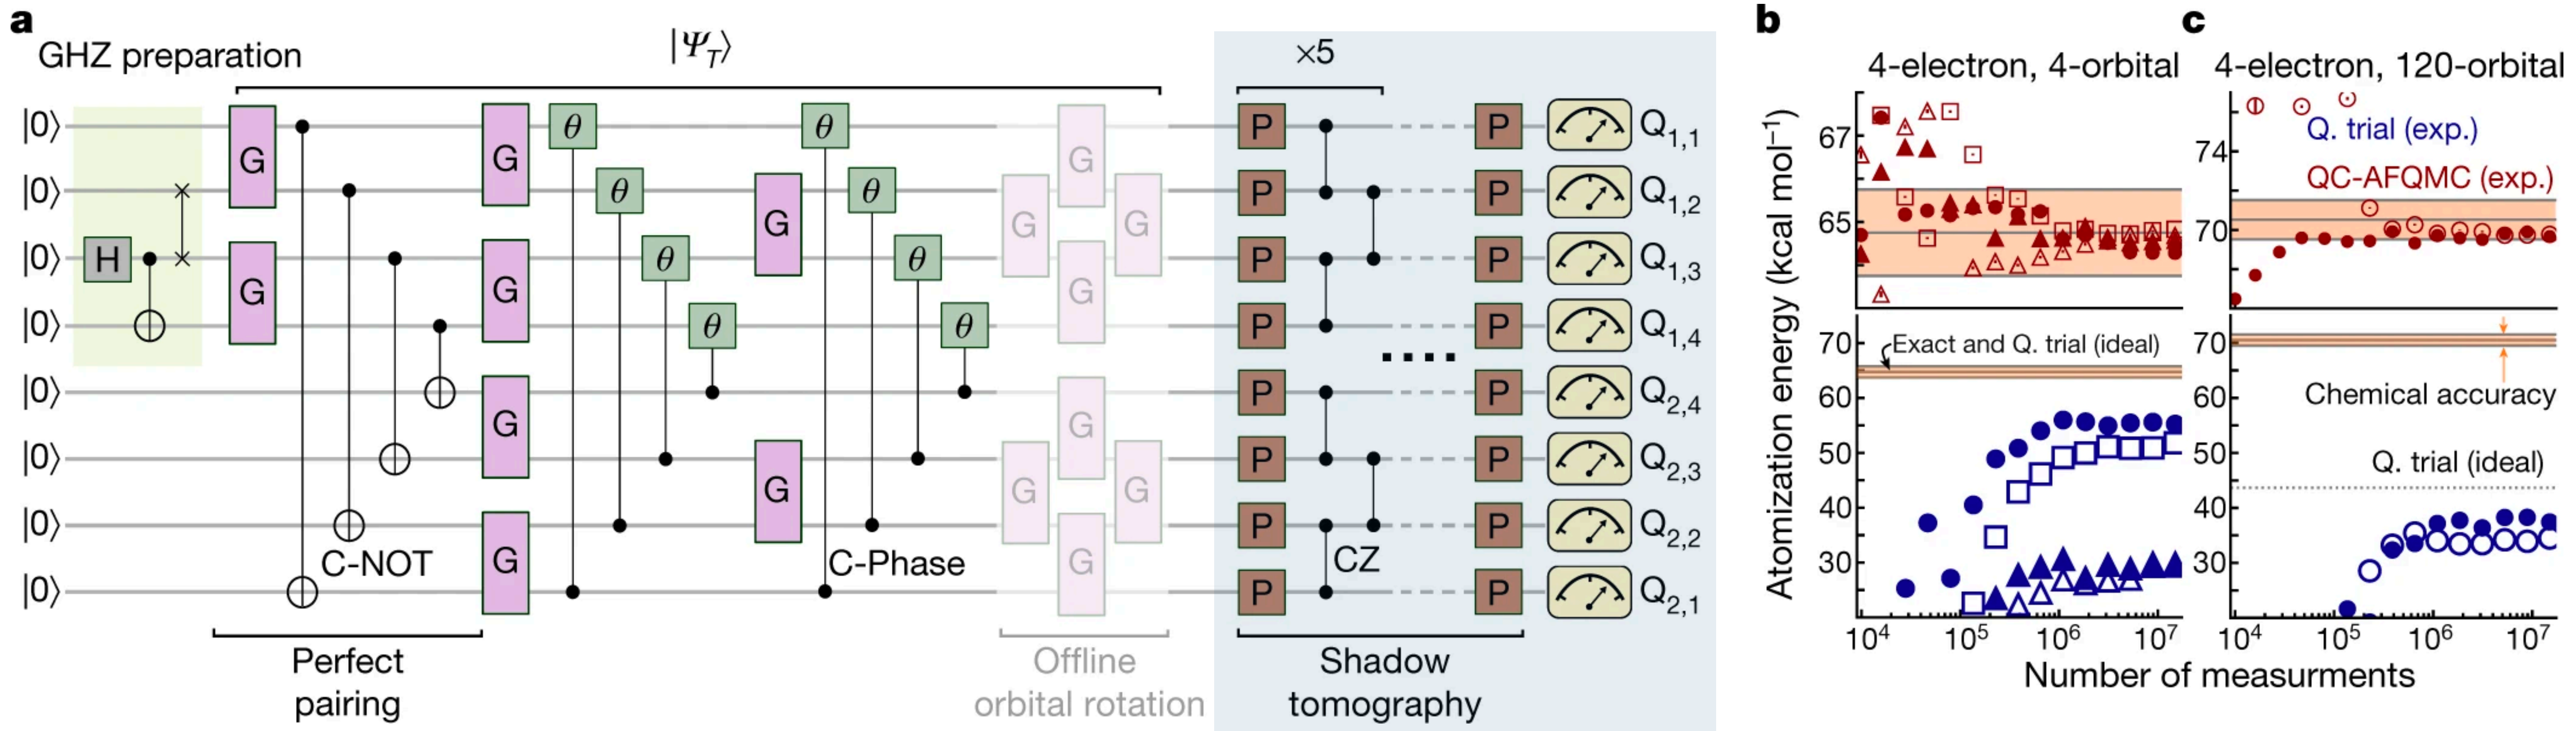

Applications of classical shadows

Benchmarking quantum systems

The ability to estimate quantum fidelity and verify entanglement enables efficient benchmarking.

[a] Huang et al. *Nature Physics* (2020).
[b] Elben et al. *Physical Review Letters* (2020).

Quantum chemistry simulation

Quantum simulation algorithms often need to estimate many properties (e.g., local observables, Hamiltonian).

[c] Zhao et al. *Physical Review Letters* (2021).
[d] Huggins et al. *Nature* (2022).

Universal quantum-to-classical converter

Enable a variety of classical computational techniques (e.g., ML) for addressing quantum problems.

[e] Huang et al. *Nature Communications* (2021).
[f] Huang et al. *Science* (2022).

Applications of classical shadows

Benchmarking quantum systems

The ability to estimate quantum fidelity and verify entanglement enables efficient benchmarking.

[a] Huang et al. *Nature Physics* (2020).
[b] Elben et al. *Physical Review Letters* (2020).

Quantum chemistry simulation

Quantum simulation algorithms often need to estimate many properties (e.g., local observables, Hamiltonian).

[c] Zhao et al. *Physical Review Letters* (2021).
[d] Huggins et al. *Nature* (2022).

Universal quantum-to-classical converter

Enable a variety of classical computational techniques (e.g., ML) for addressing quantum problems.

[e] Huang et al. *Nature Communications* (2021).
[f] Huang et al. *Science* (2022).

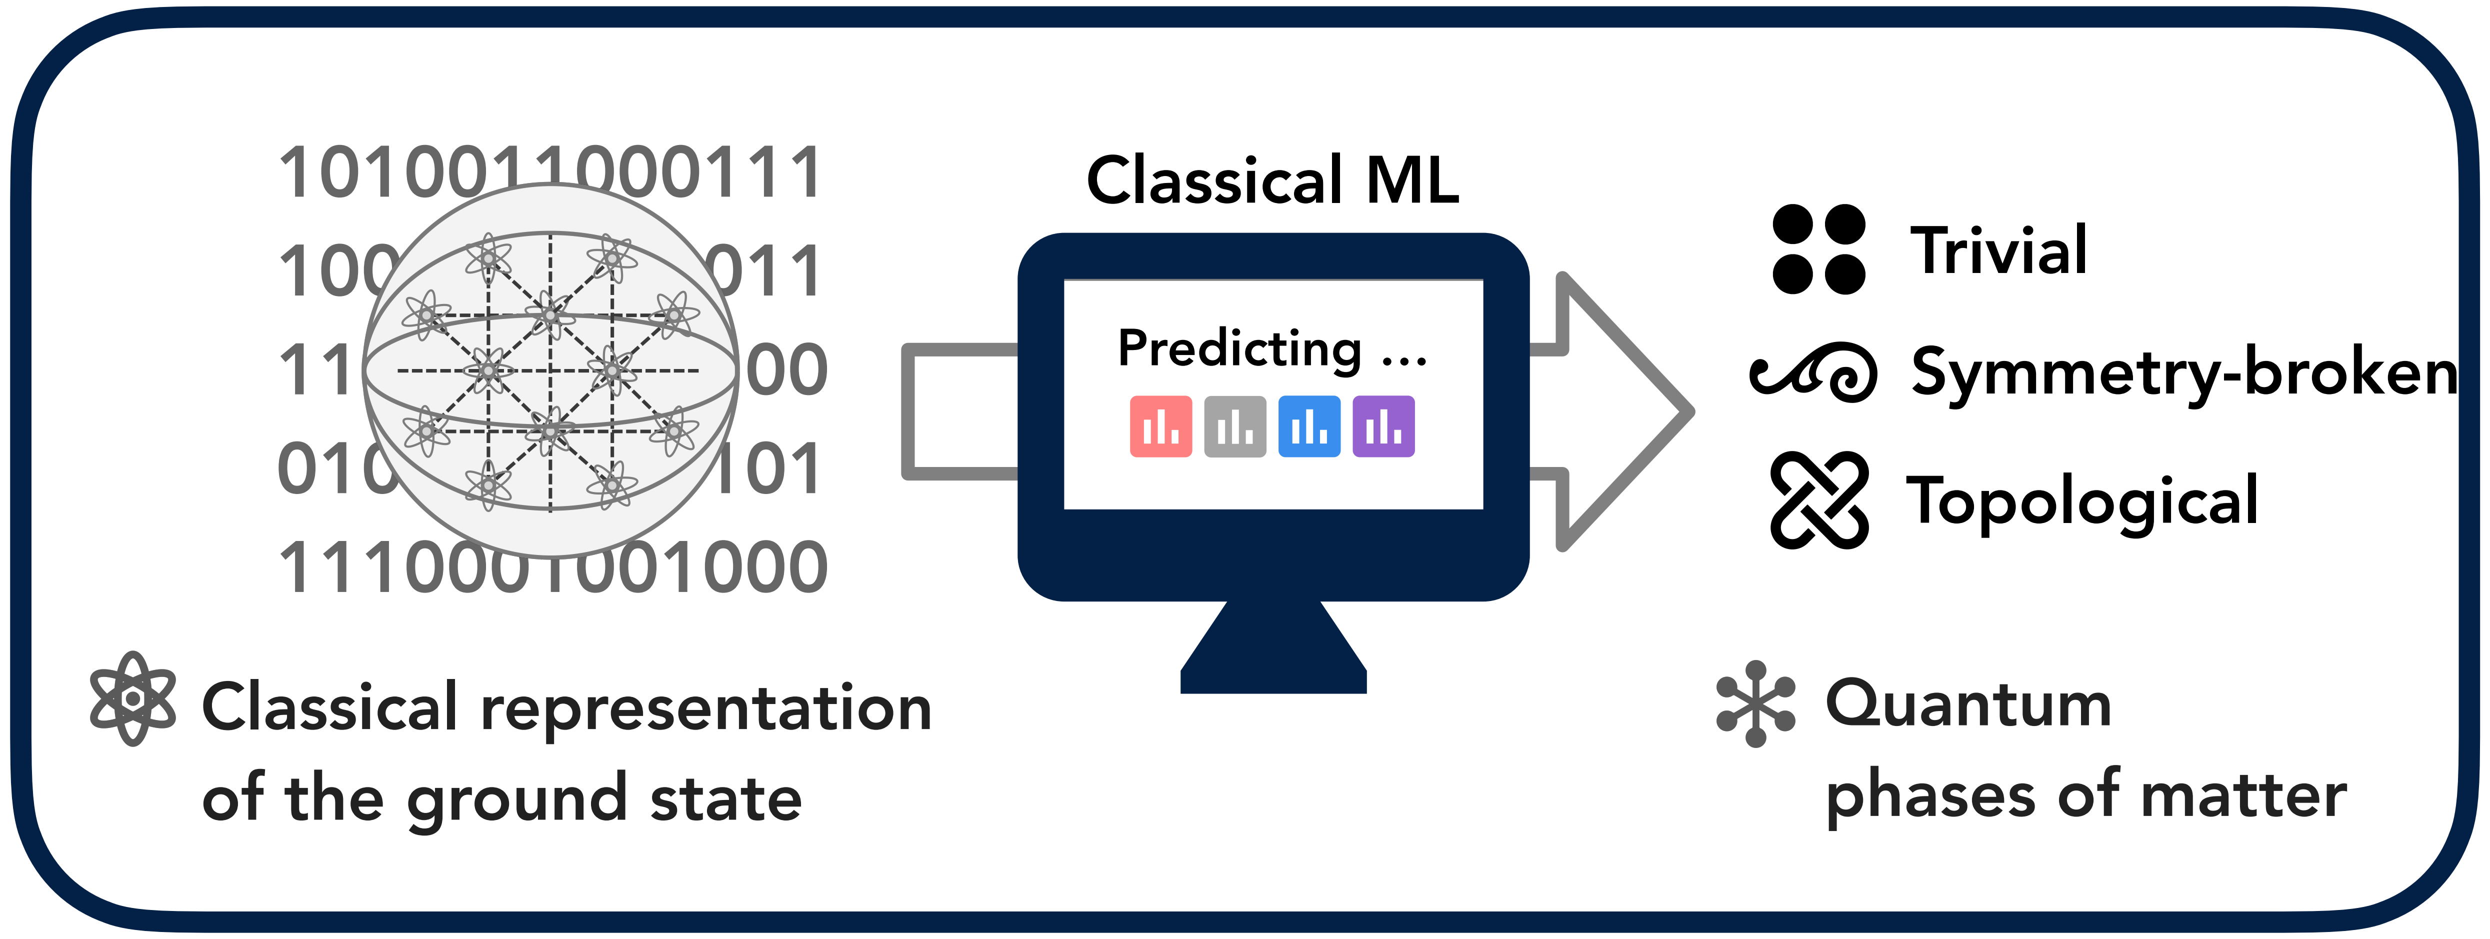

Applications of classical shadows

Benchmarking quantum systems

The ability to estimate quantum fidelity and verify entanglement enables efficient benchmarking.

[a] Huang et al. *Nature Physics* (2020).
[b] Elben et al. *Physical Review Letters* (2020).

Quantum chemistry simulation

Quantum simulation algorithms often need to estimate many properties (e.g., local observables, Hamiltonian).

[c] Zhao et al. *Physical Review Letters* (2021).
[d] Huggins et al. *Nature* (2022).

Universal quantum-to-classical converter

Enable a variety of classical computational techniques (e.g., ML) for addressing quantum problems.

[e] Huang et al. *Nature Communications* (2021).
[f] Huang et al. *Science* (2022).

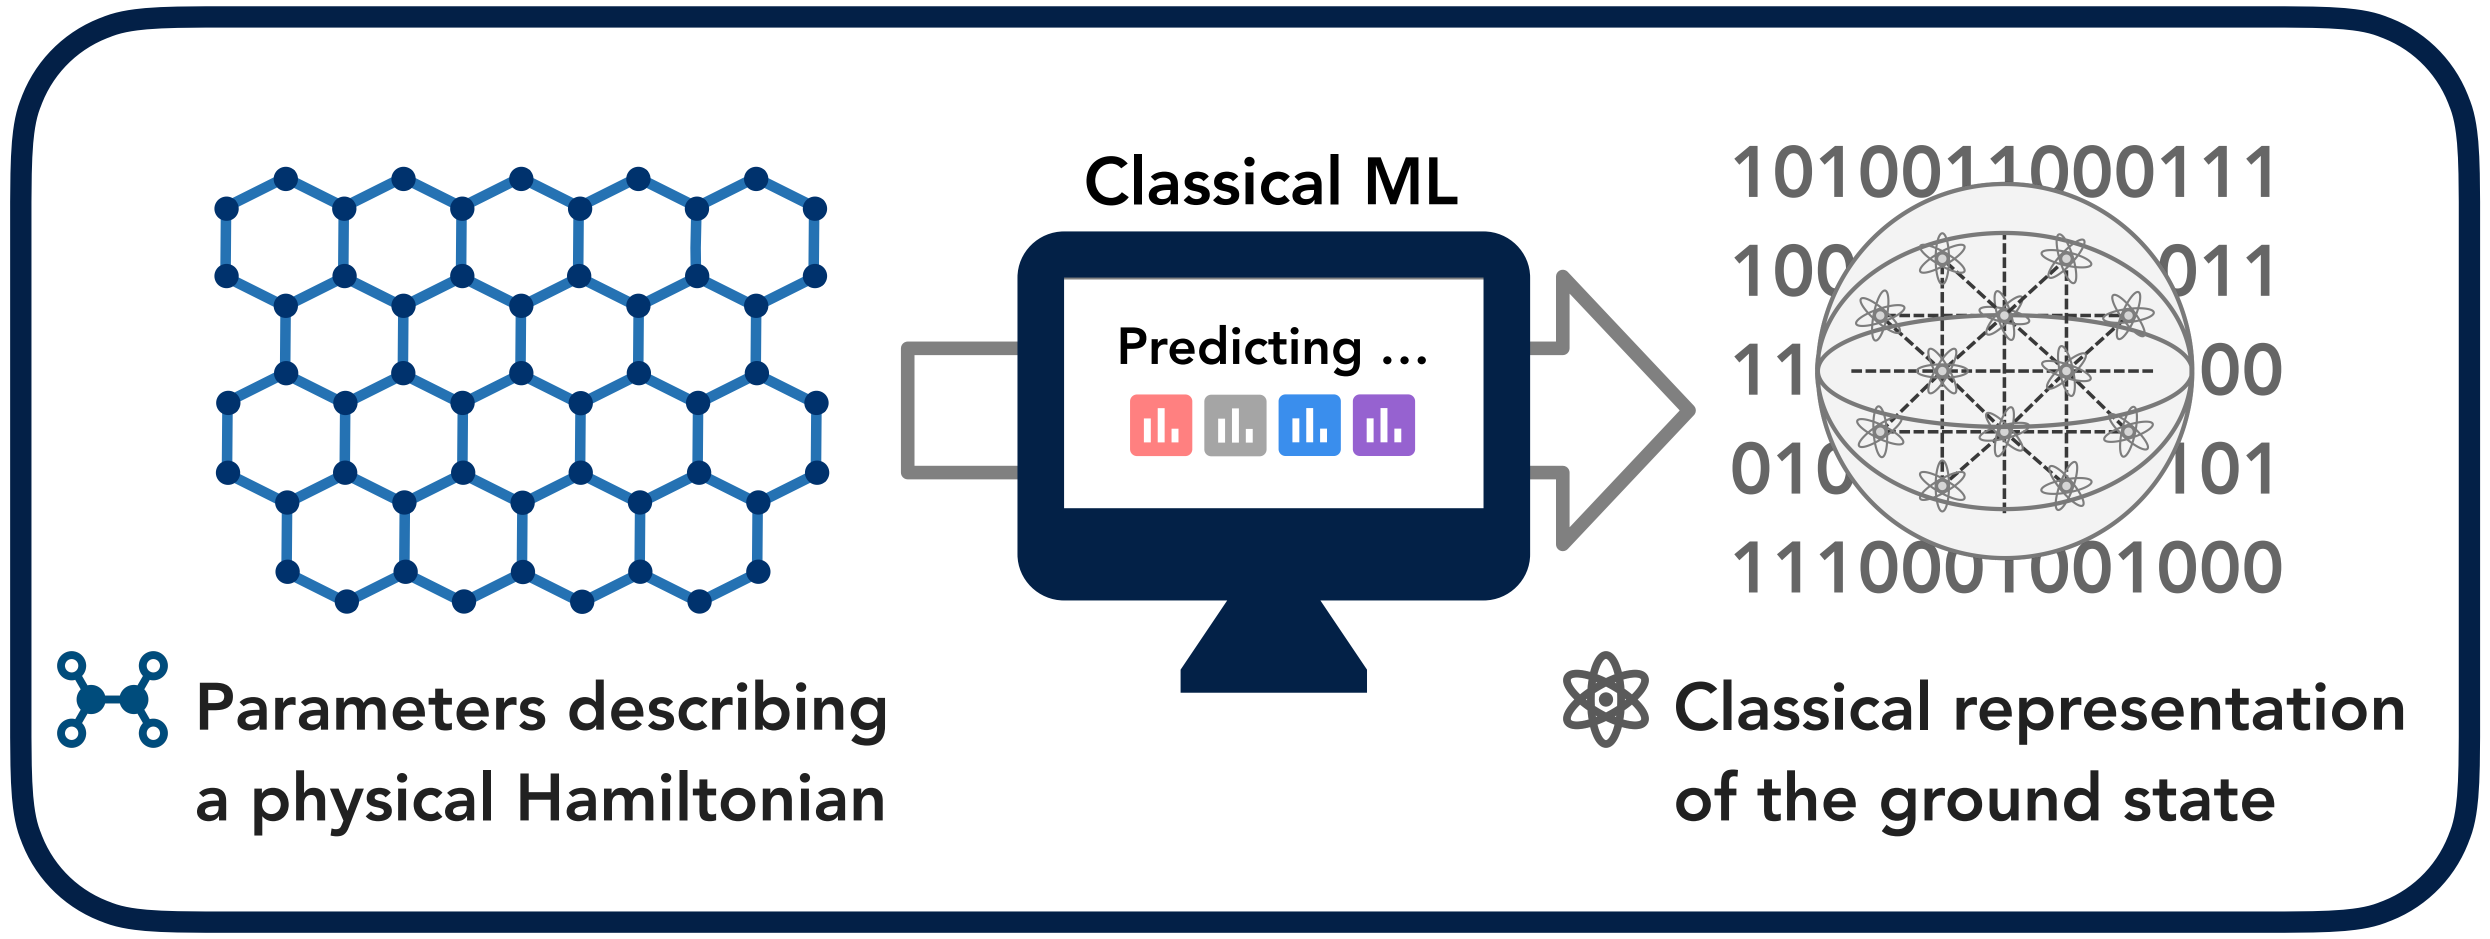

Classical ML for quantum problems

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

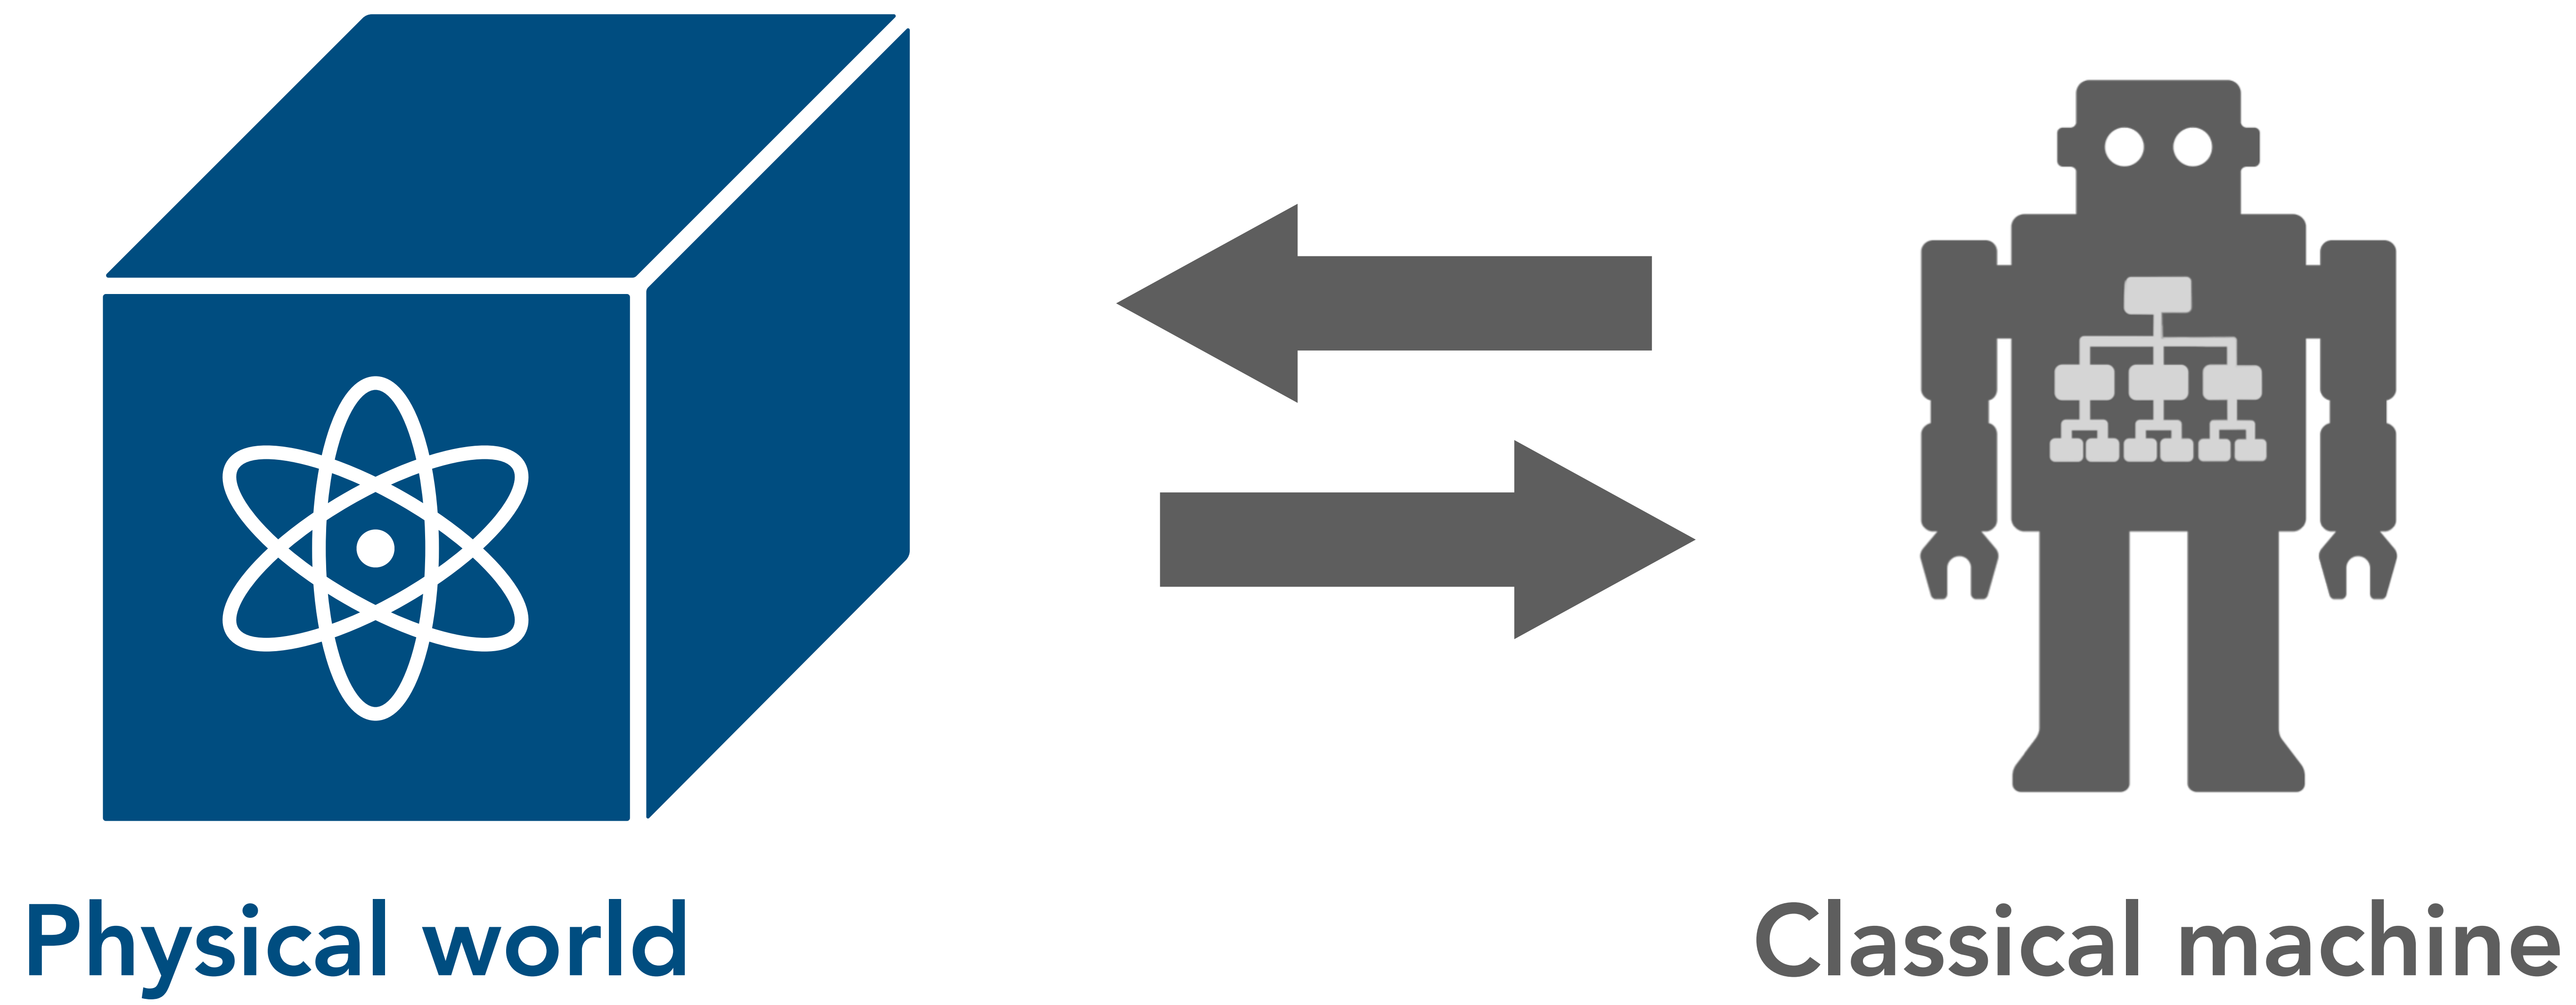

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

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.

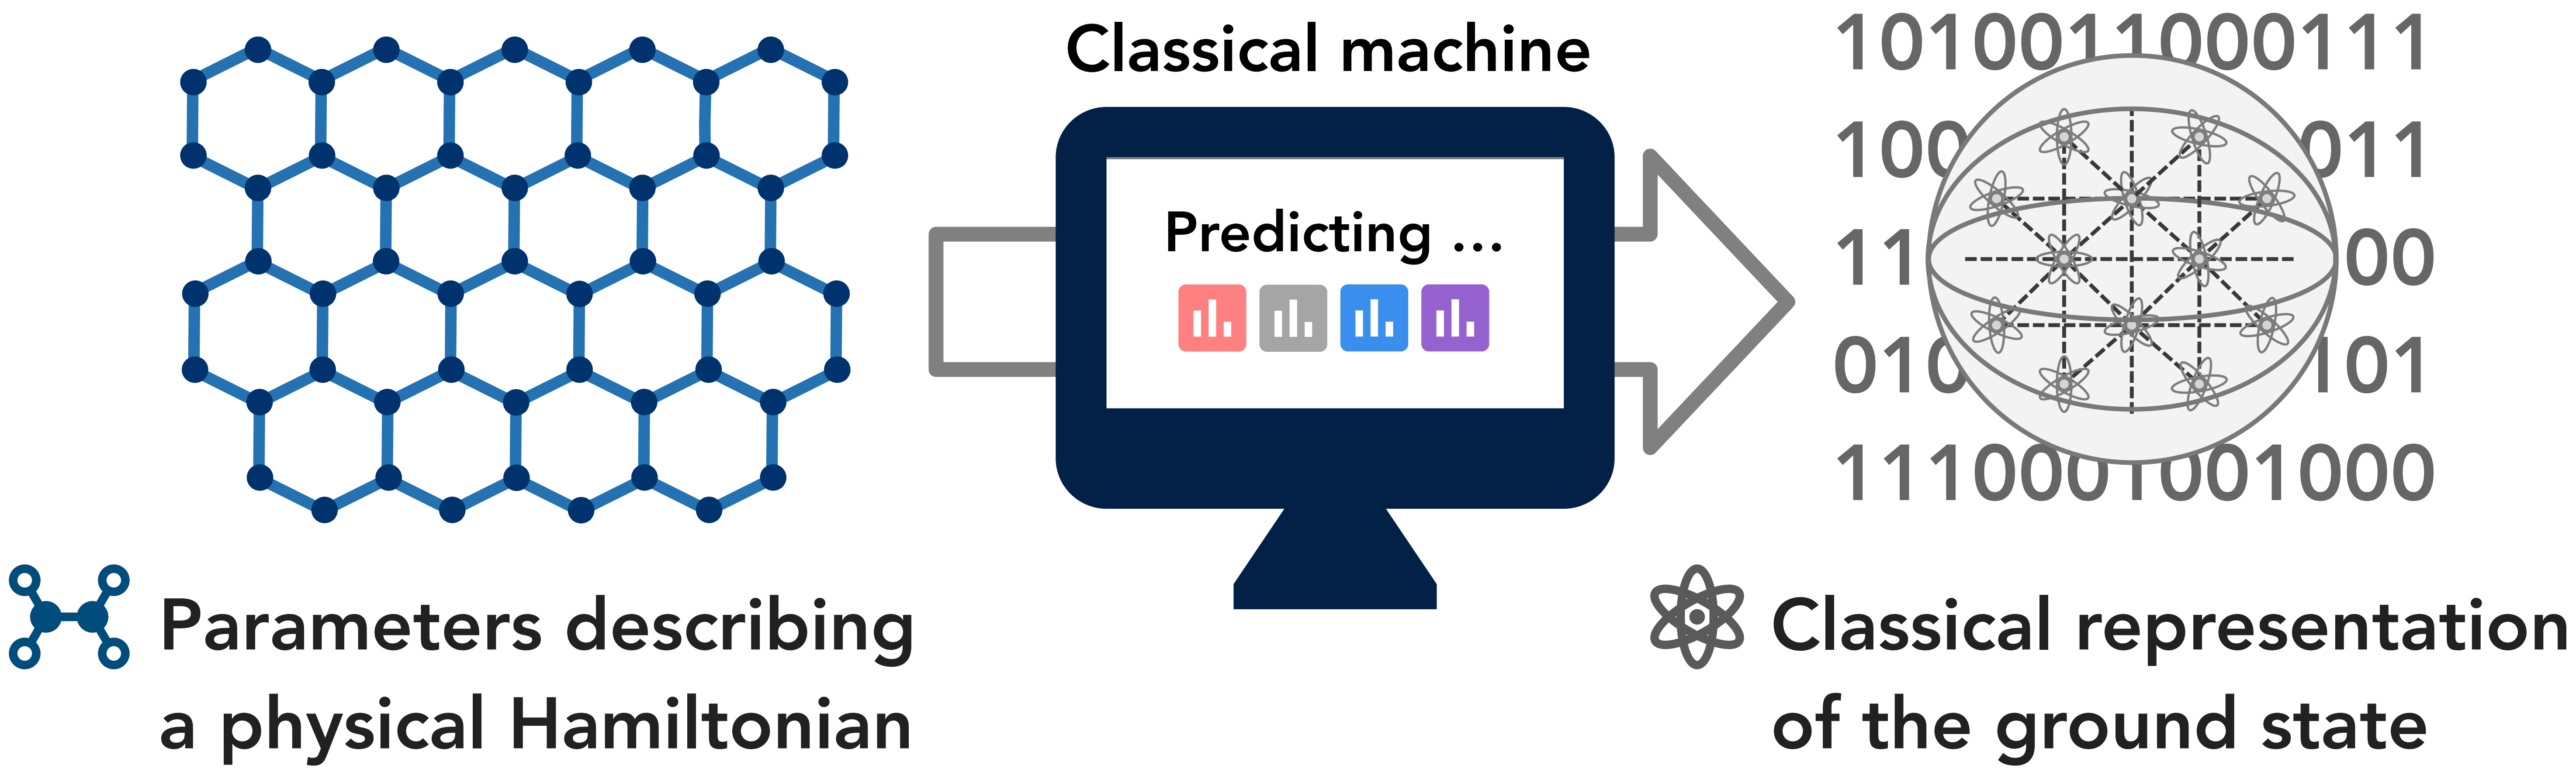

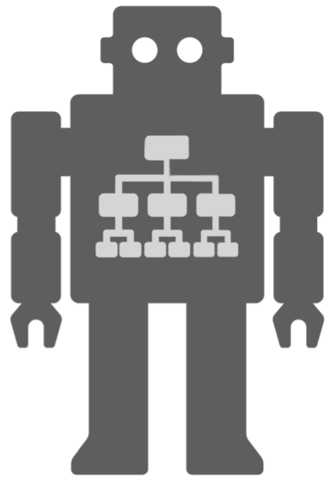

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

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.

1D

2D

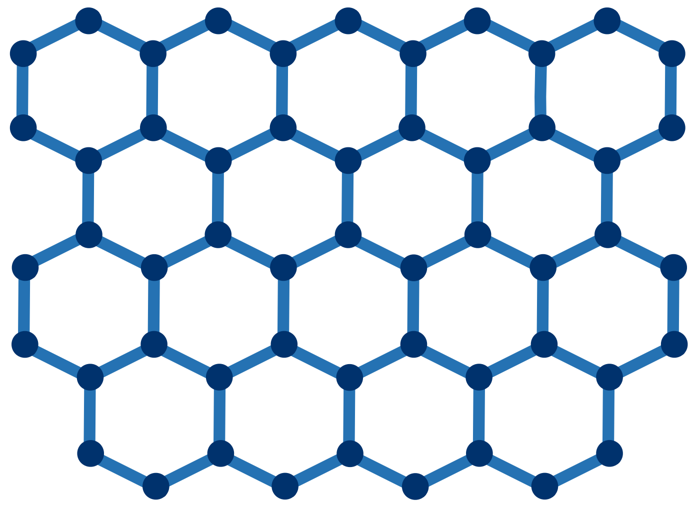

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

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.

1D

2D

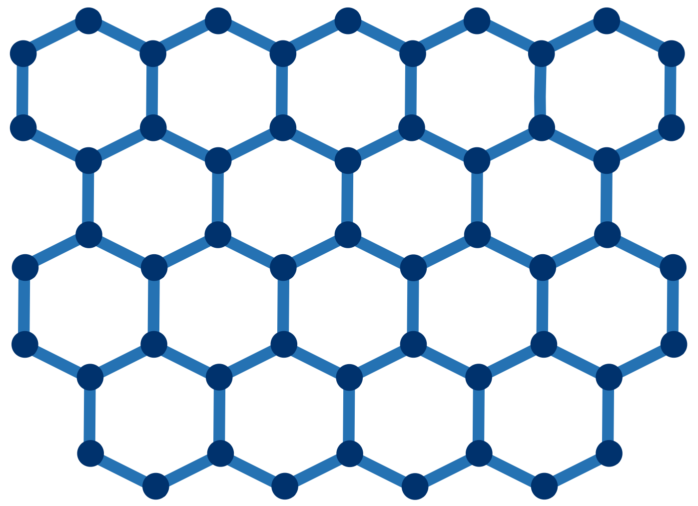

Proposition 1

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

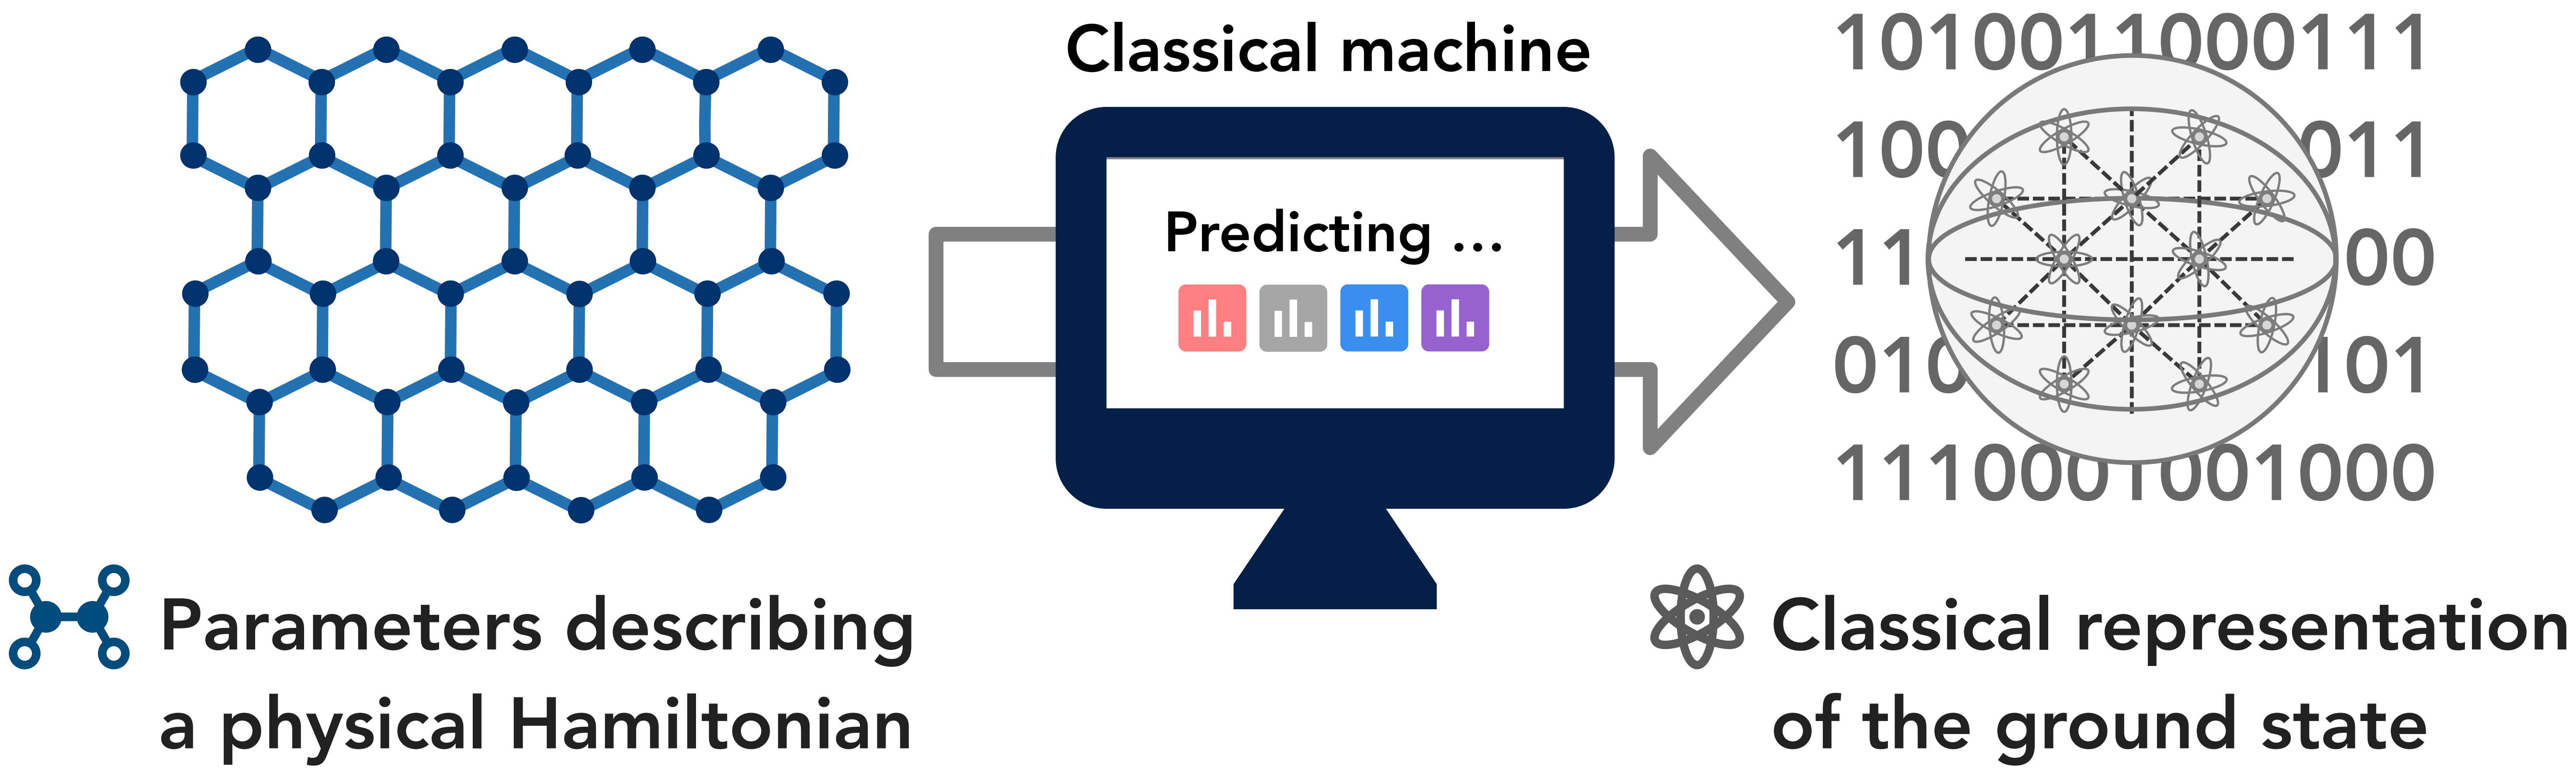

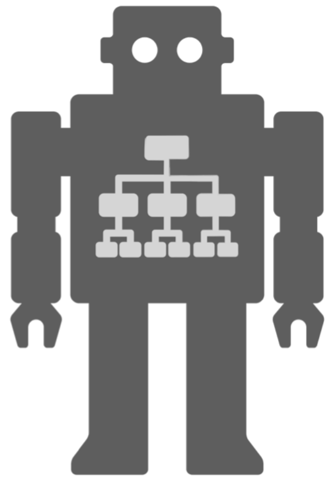

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Task

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

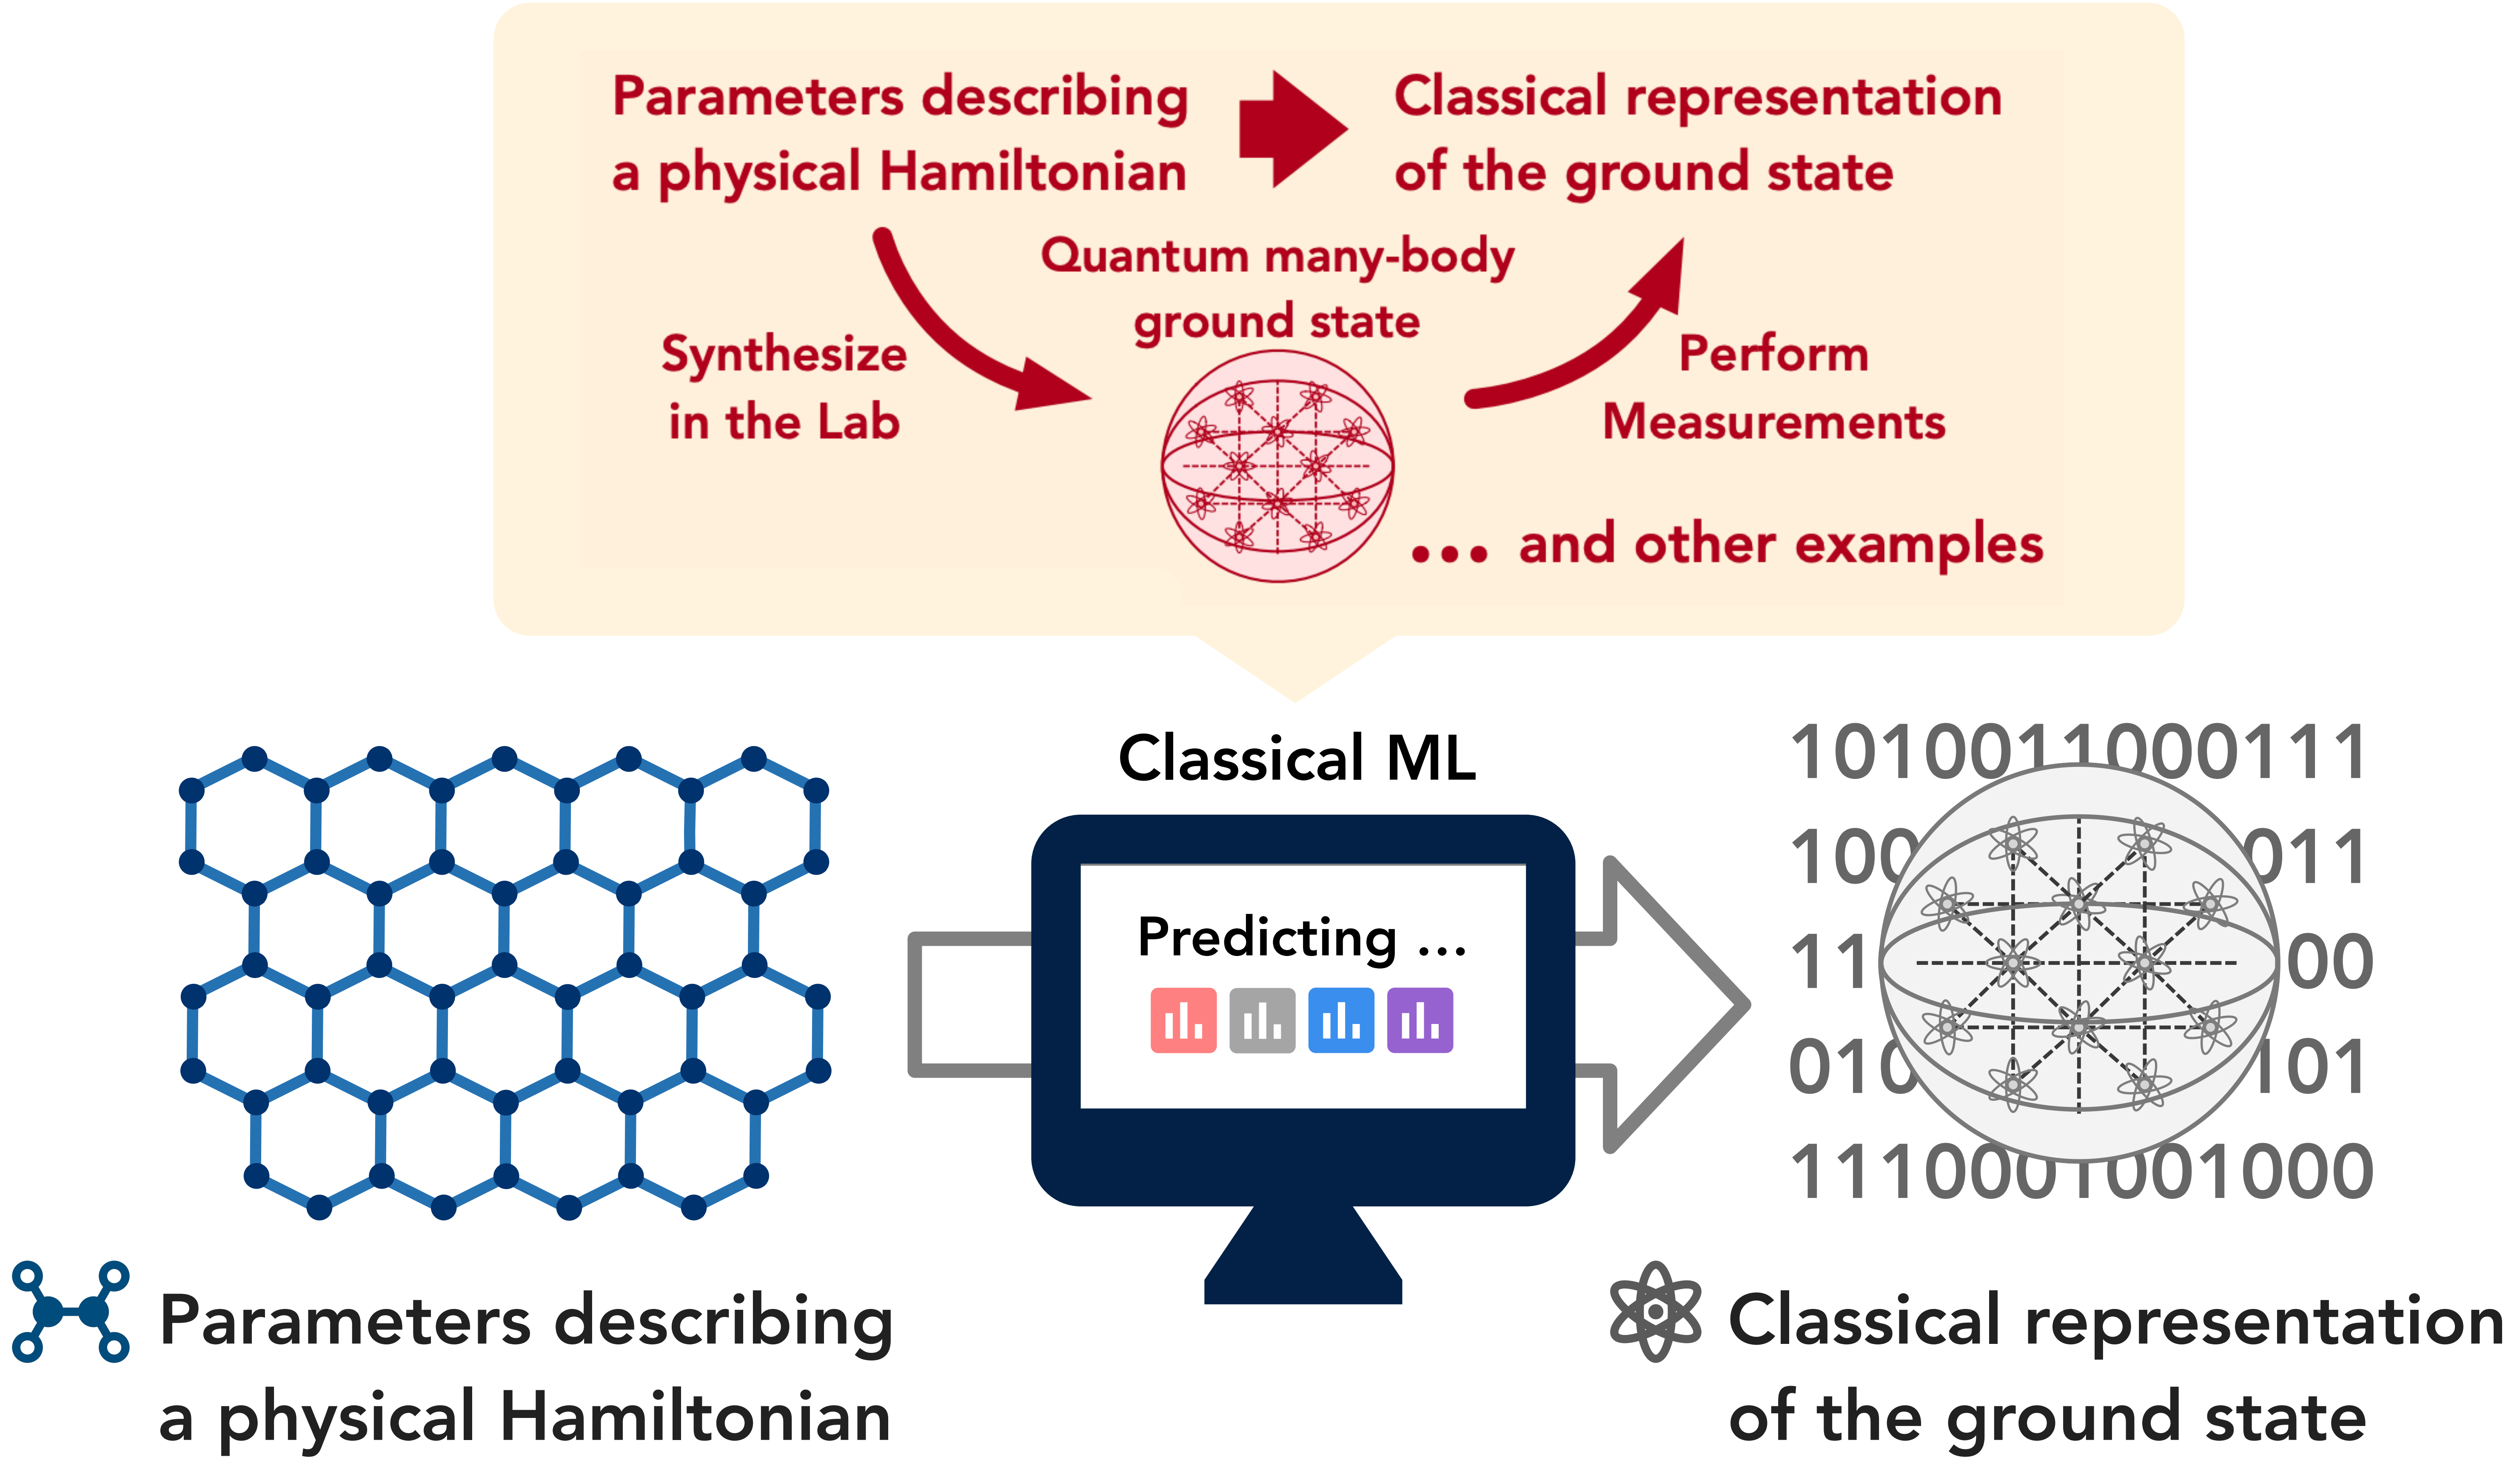

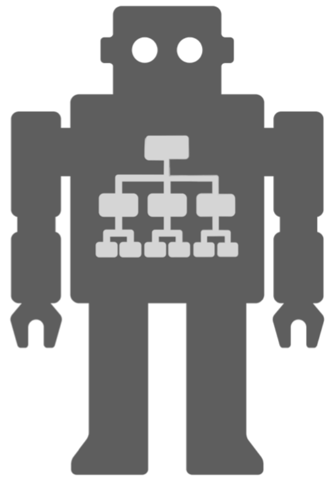

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Task

Classical shadow

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

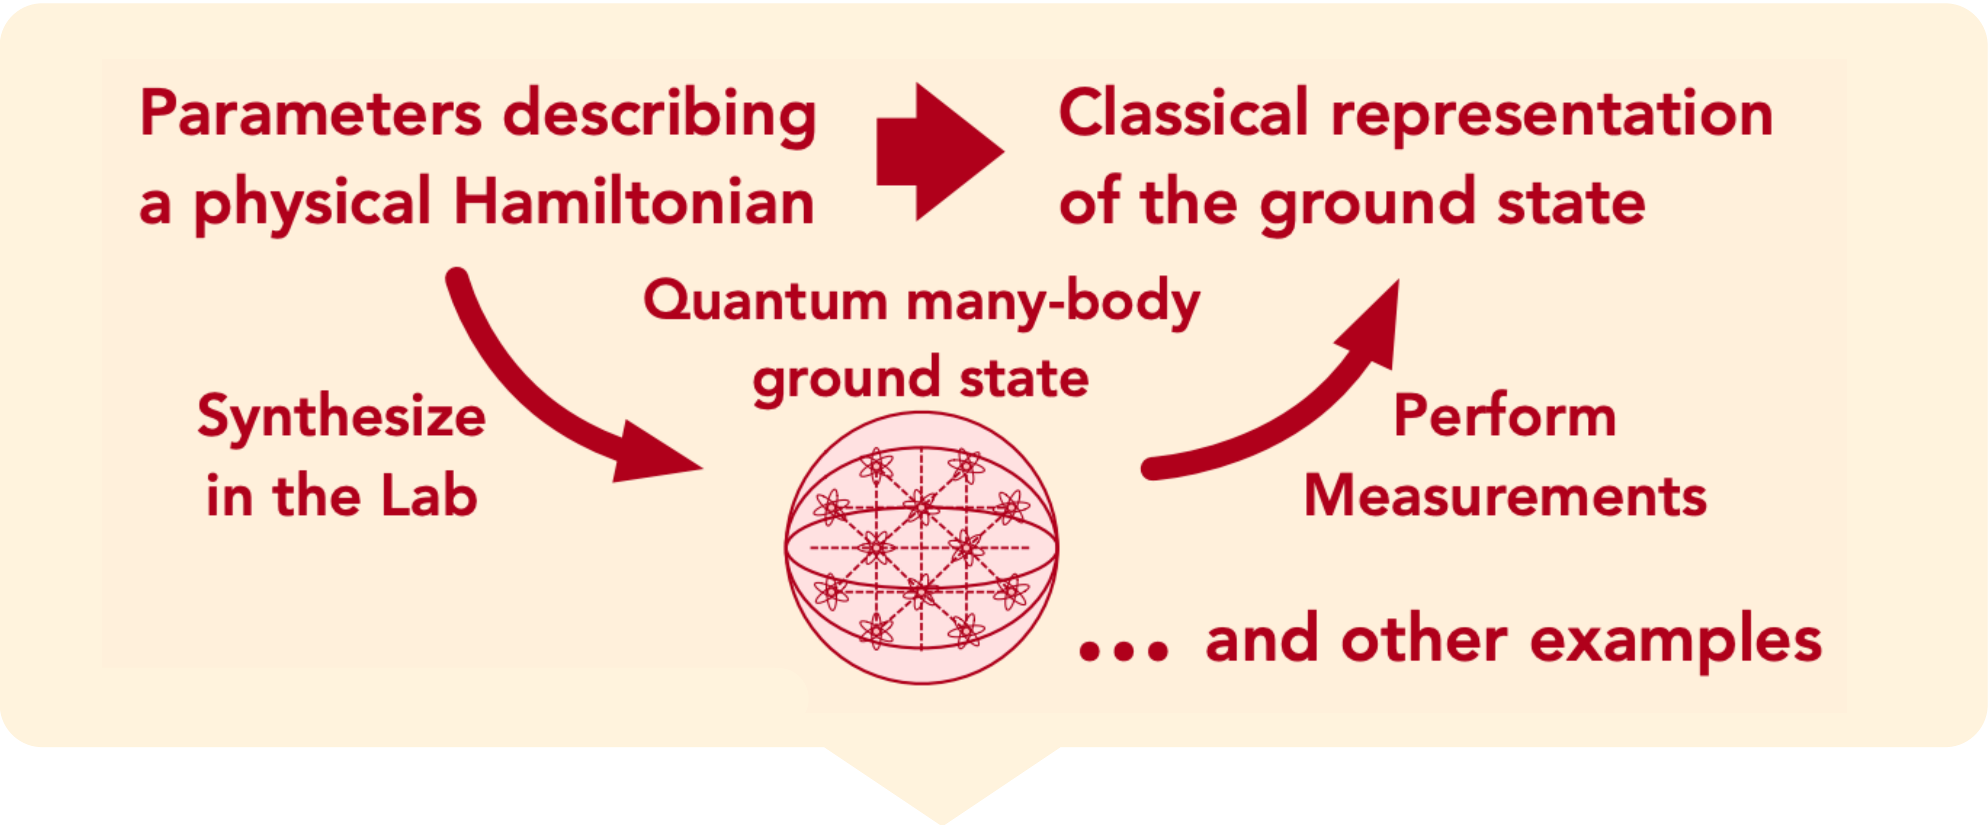

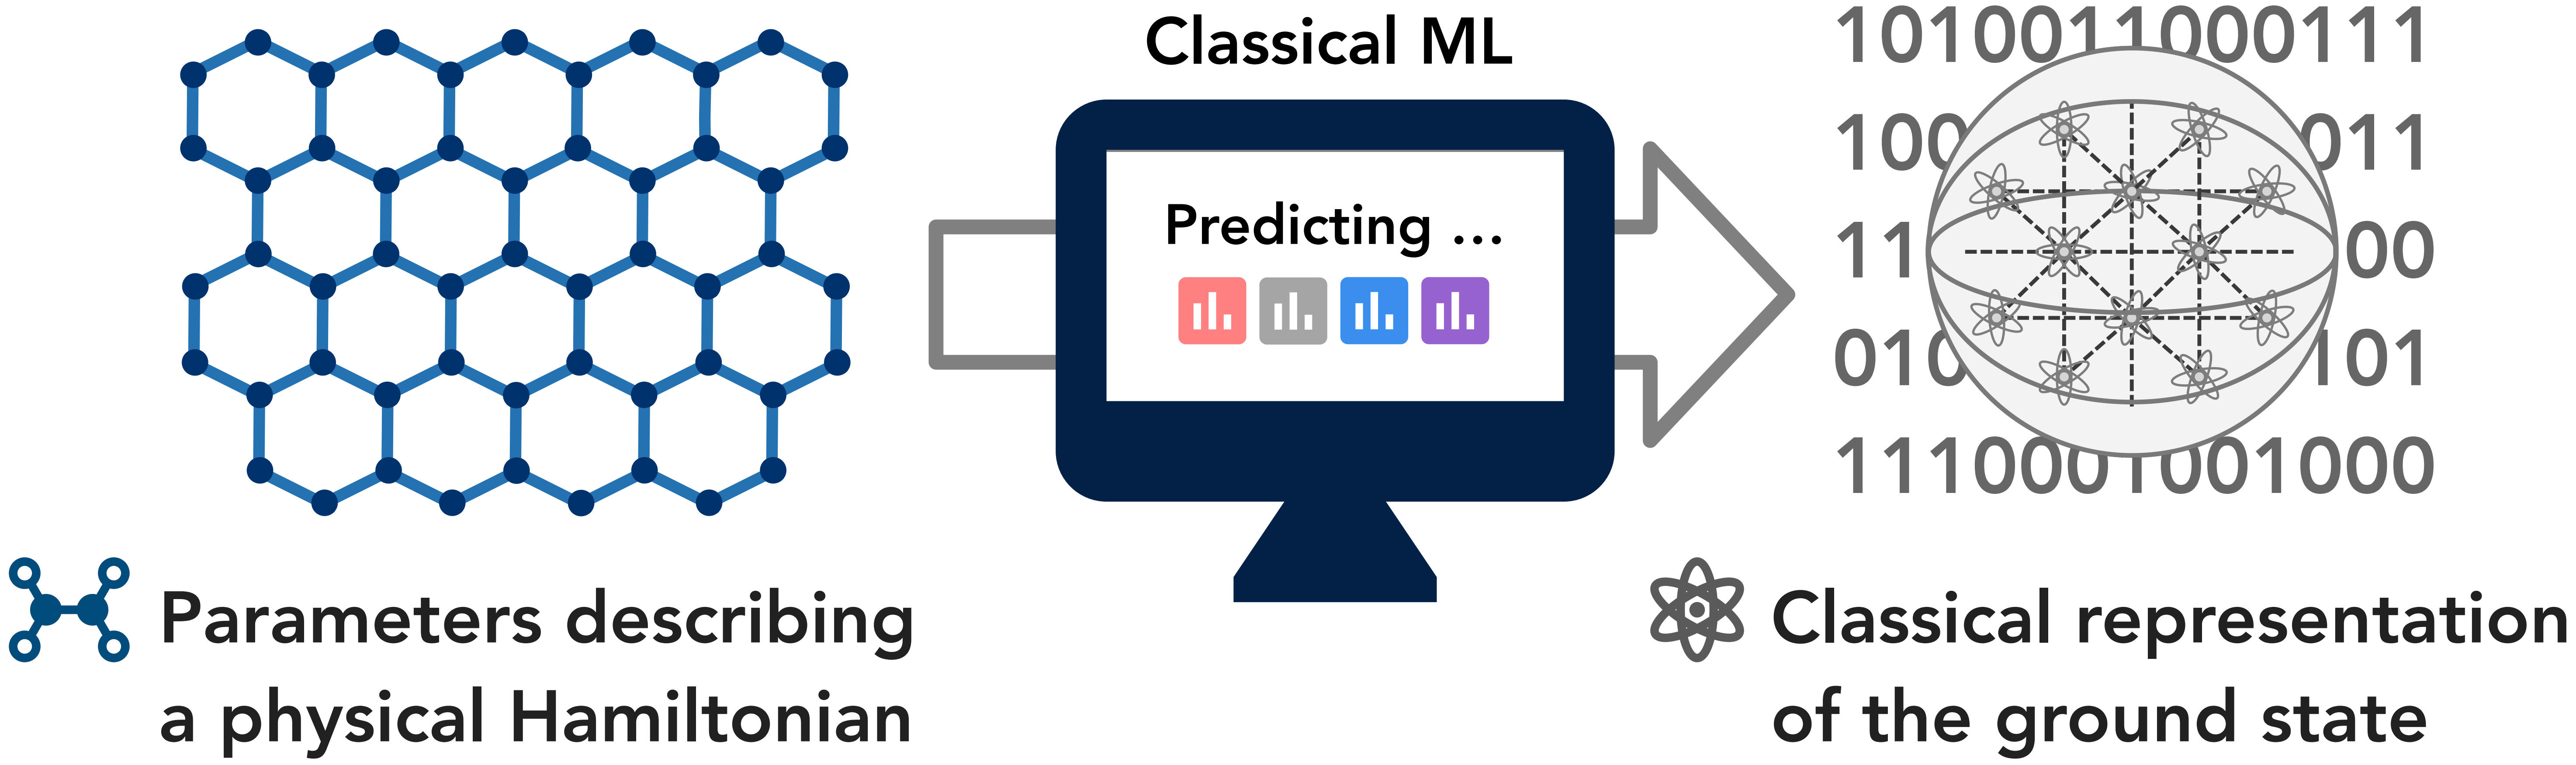

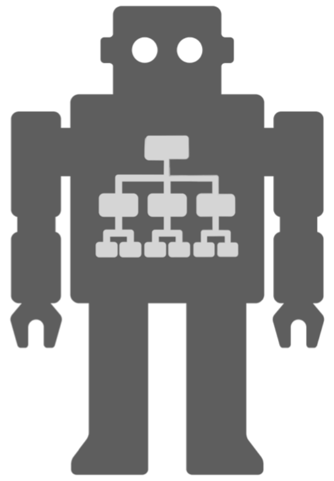

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Proposition 1

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

Classical algorithm

2D

spectral gap 1

1-body observable

average prediction error 1/4

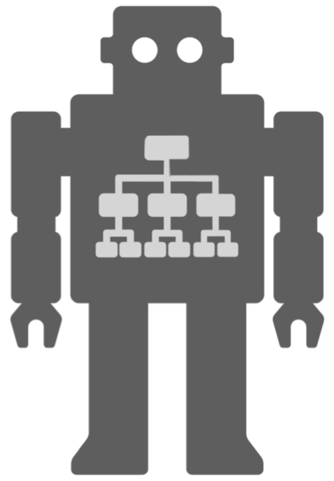

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Theorem 1: Improved version from [LHT+23]

A classical ML algorithm can achieve an average prediction error $\leq \epsilon$ using $\log(n)$ training data and $n\log(n)$ 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 $\epsilon = \mathcal{O}(1)$

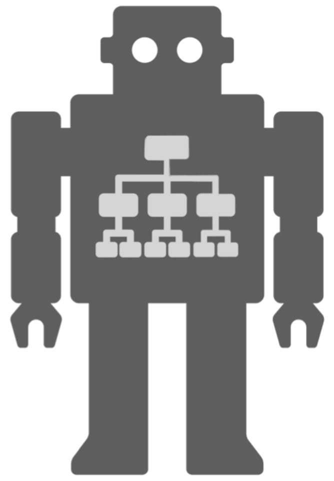

[LHT+23] Lewis, Huang, Tran, Lehner, Kueng, Preskill. Improved machine learning algorithm for predicting ground state properties, *Submitted*, 2023.
[HKT+22] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Theorem 1: Improved version from [LHT+23]

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

A classical ML model

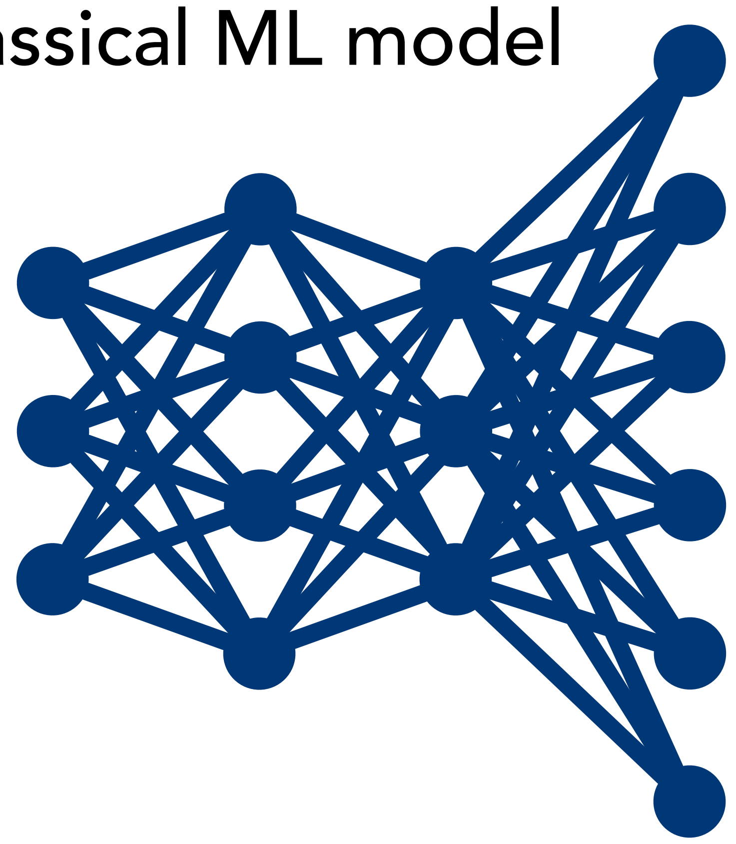

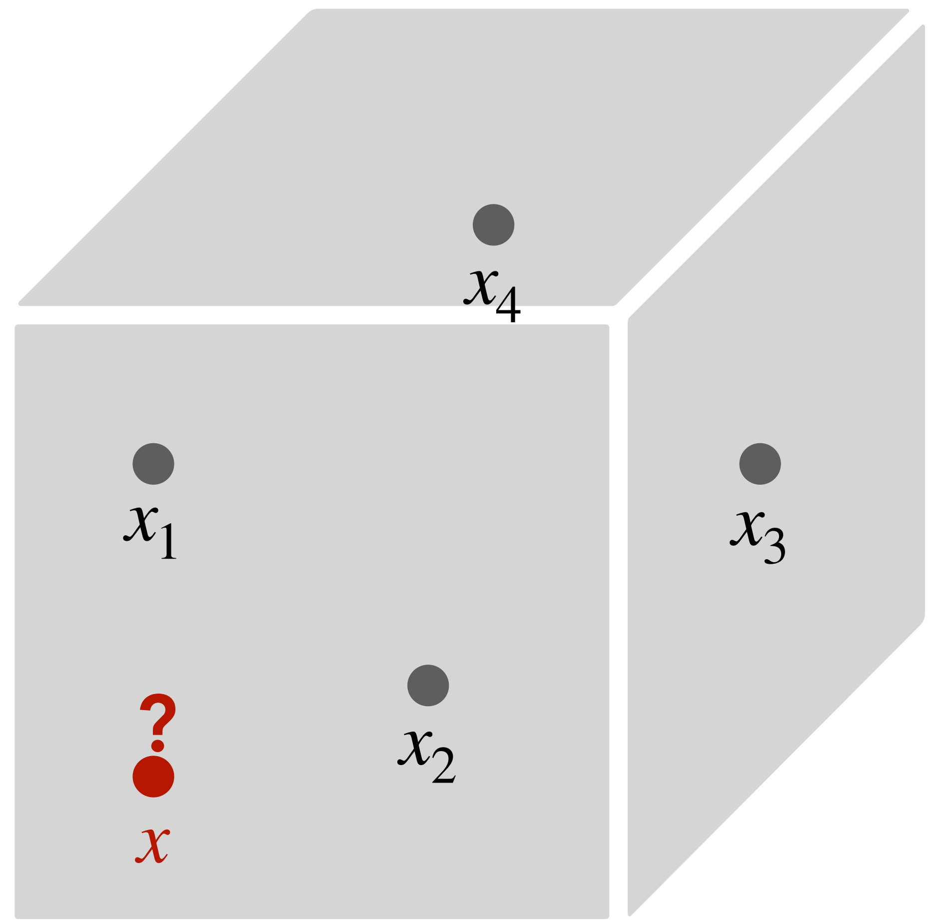

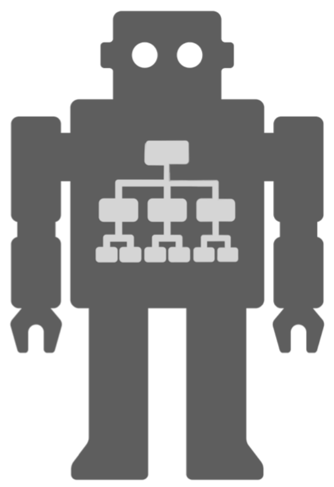

[HKT+22] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Theorem 1: Improved version from [LHT+23]

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

A classical ML model

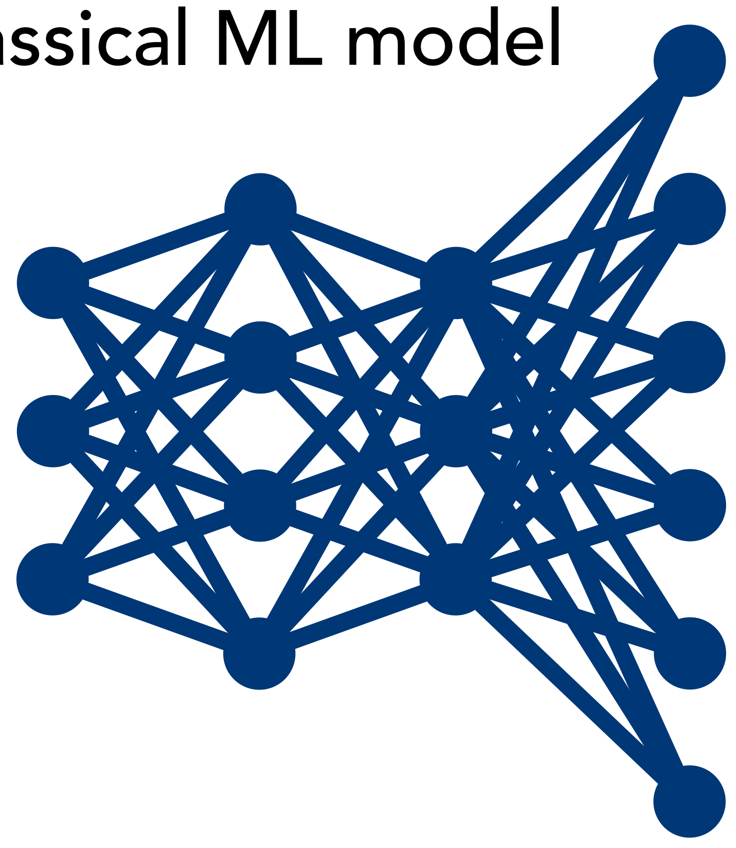

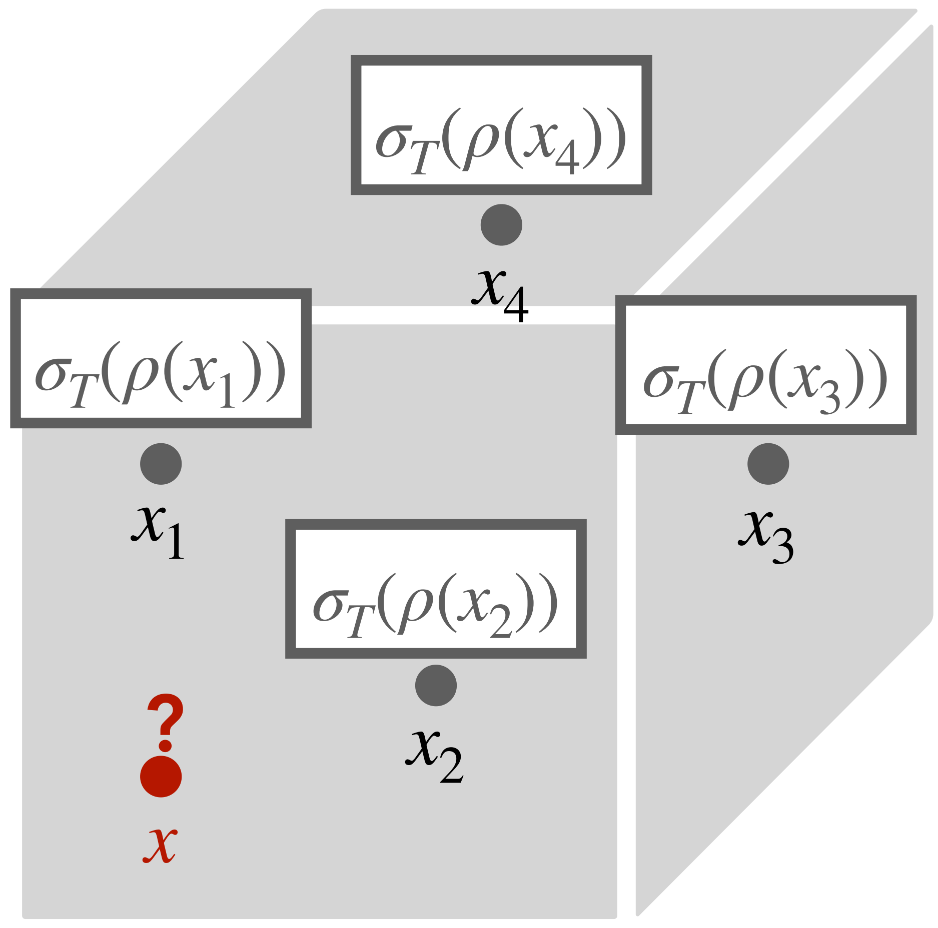

[HKT+22] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Theorem 1: Improved version from [LHT+23]

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

A classical ML model

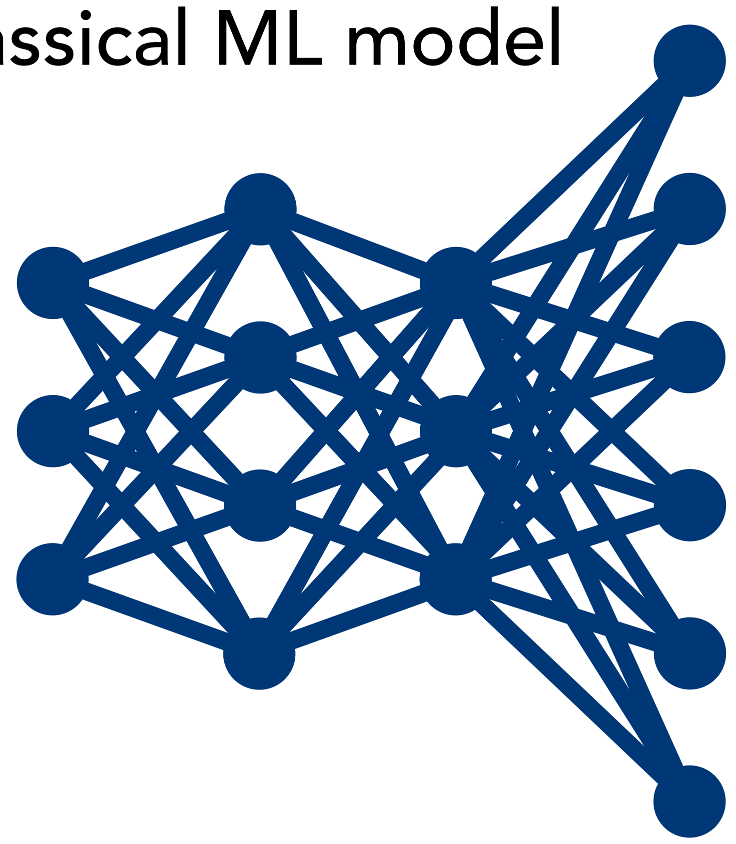

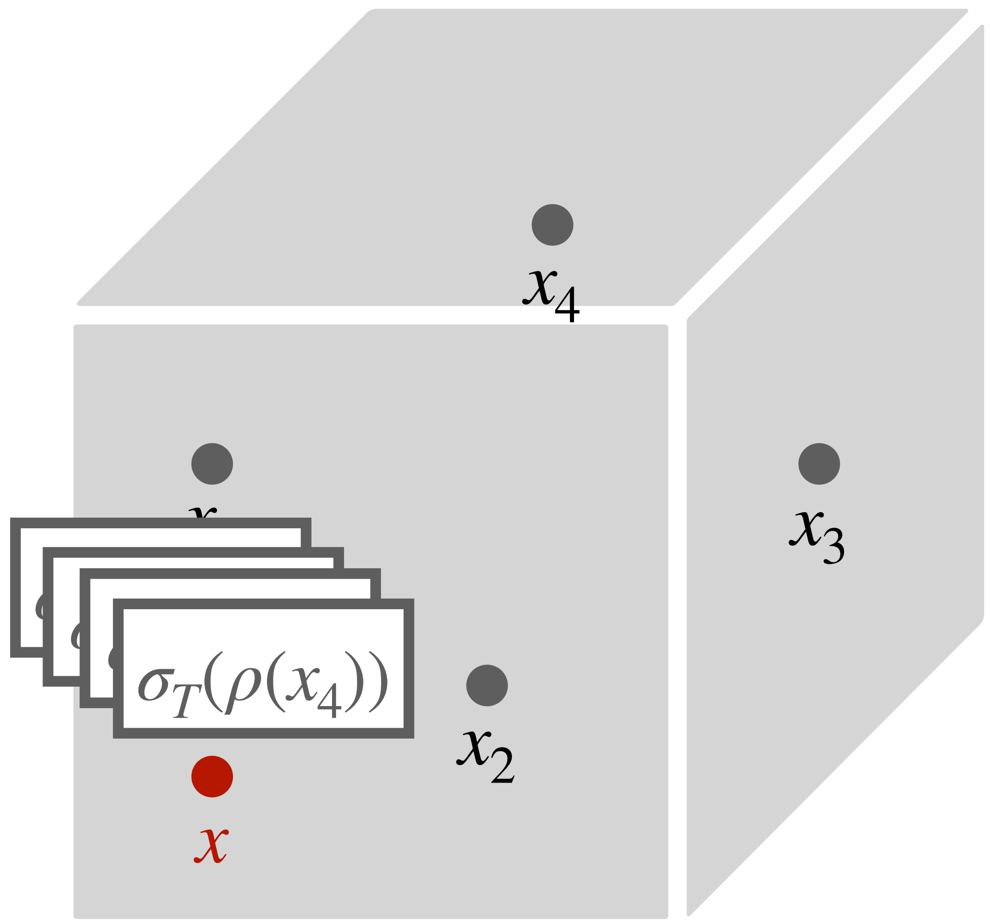

[HKT+22] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

Theorem 1: Improved version from [LHT+23]

A classical ML algorithm can achieve an average prediction error $\leq \epsilon$ using $\log(n)$ training data and $n\log(n)$ 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 $\epsilon = \mathcal{O}(1)$

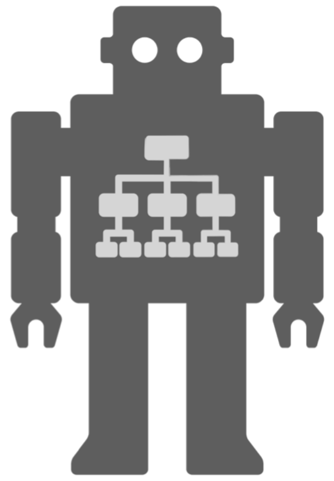

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

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

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 $\epsilon = \mathcal{O}(1)$

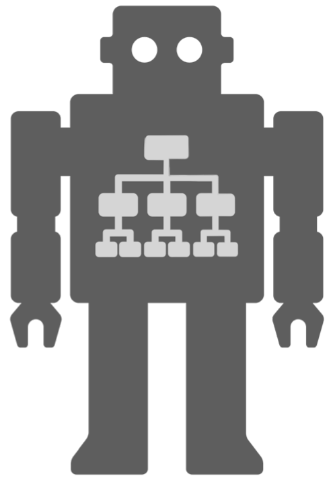

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

Predicting ground states: Theorem

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

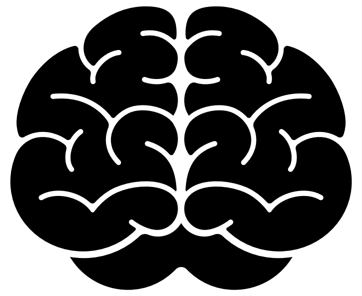

The answer:
Generalizing from data can be easier than computing everything

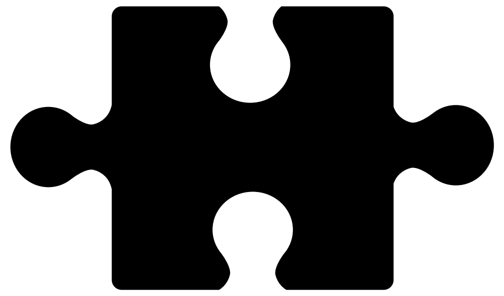

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

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?

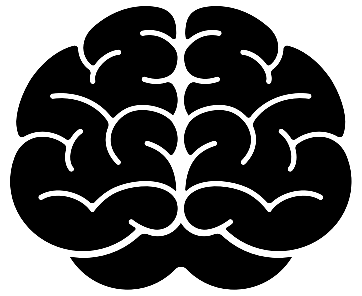

The answer:
Generalizing from data can be easier than computing everything

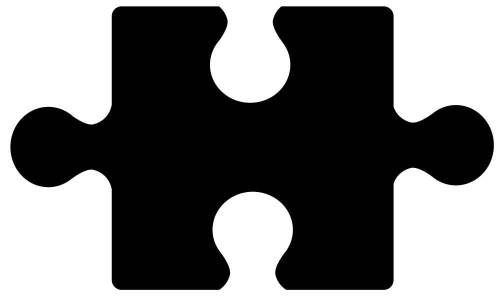

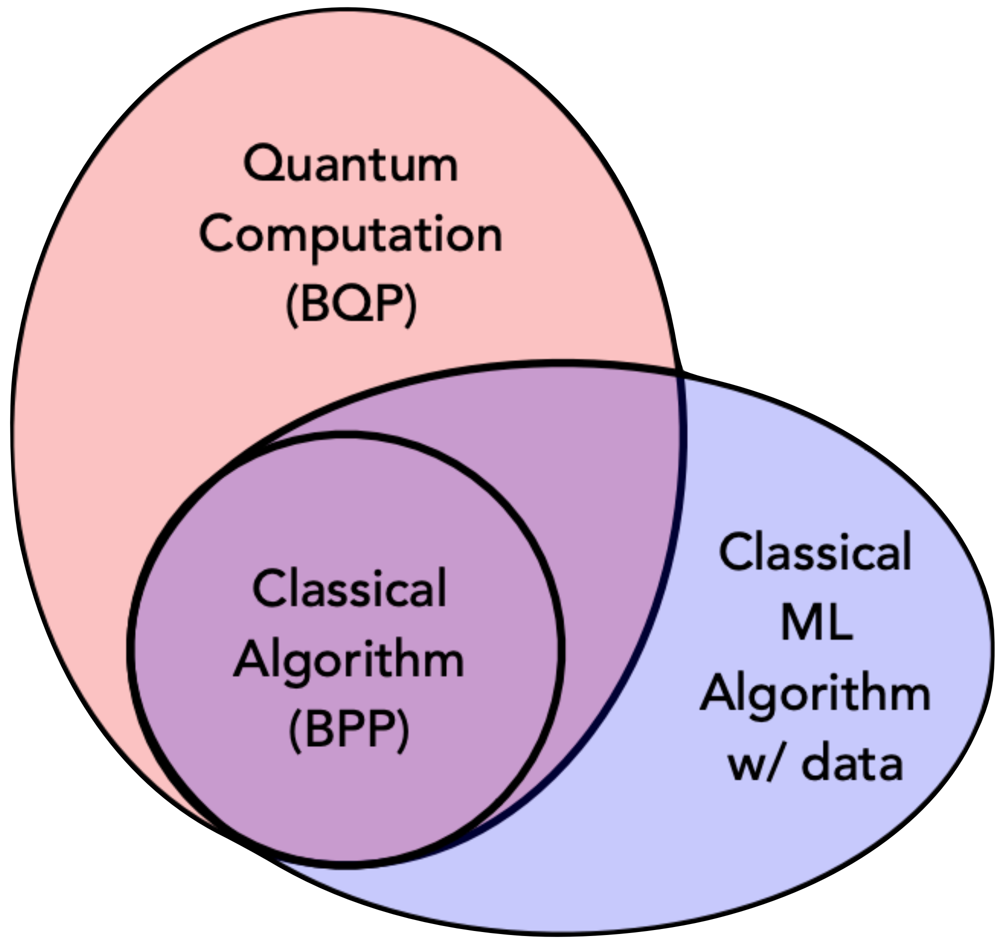

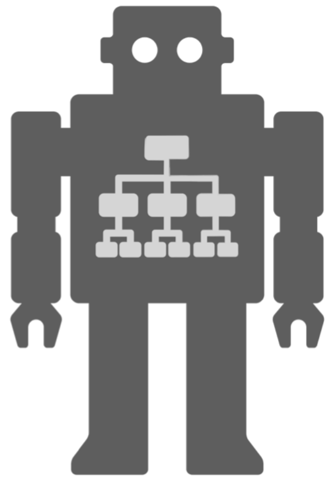

[HBM+21] Huang, Broughton, et al. Power of data in quantum machine learning. Nature Communications, 2021.
[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

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.

(a) 2D anti-ferromagnetic random Heisenberg model

$$H = \sum_{\langle ij \rangle} J_{ij}(X_i X_j + Y_i Y_j + Z_i Z_j)$$

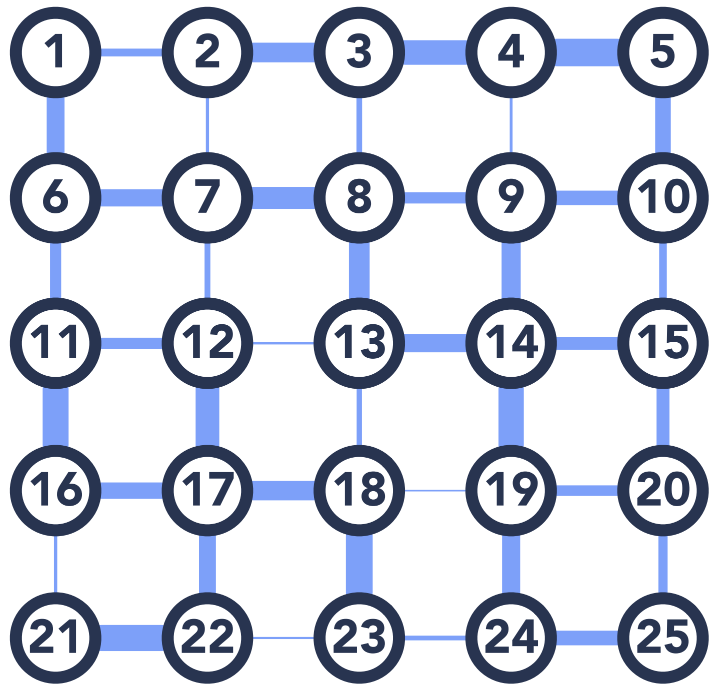

*The random J considered in **(b)**

(b)

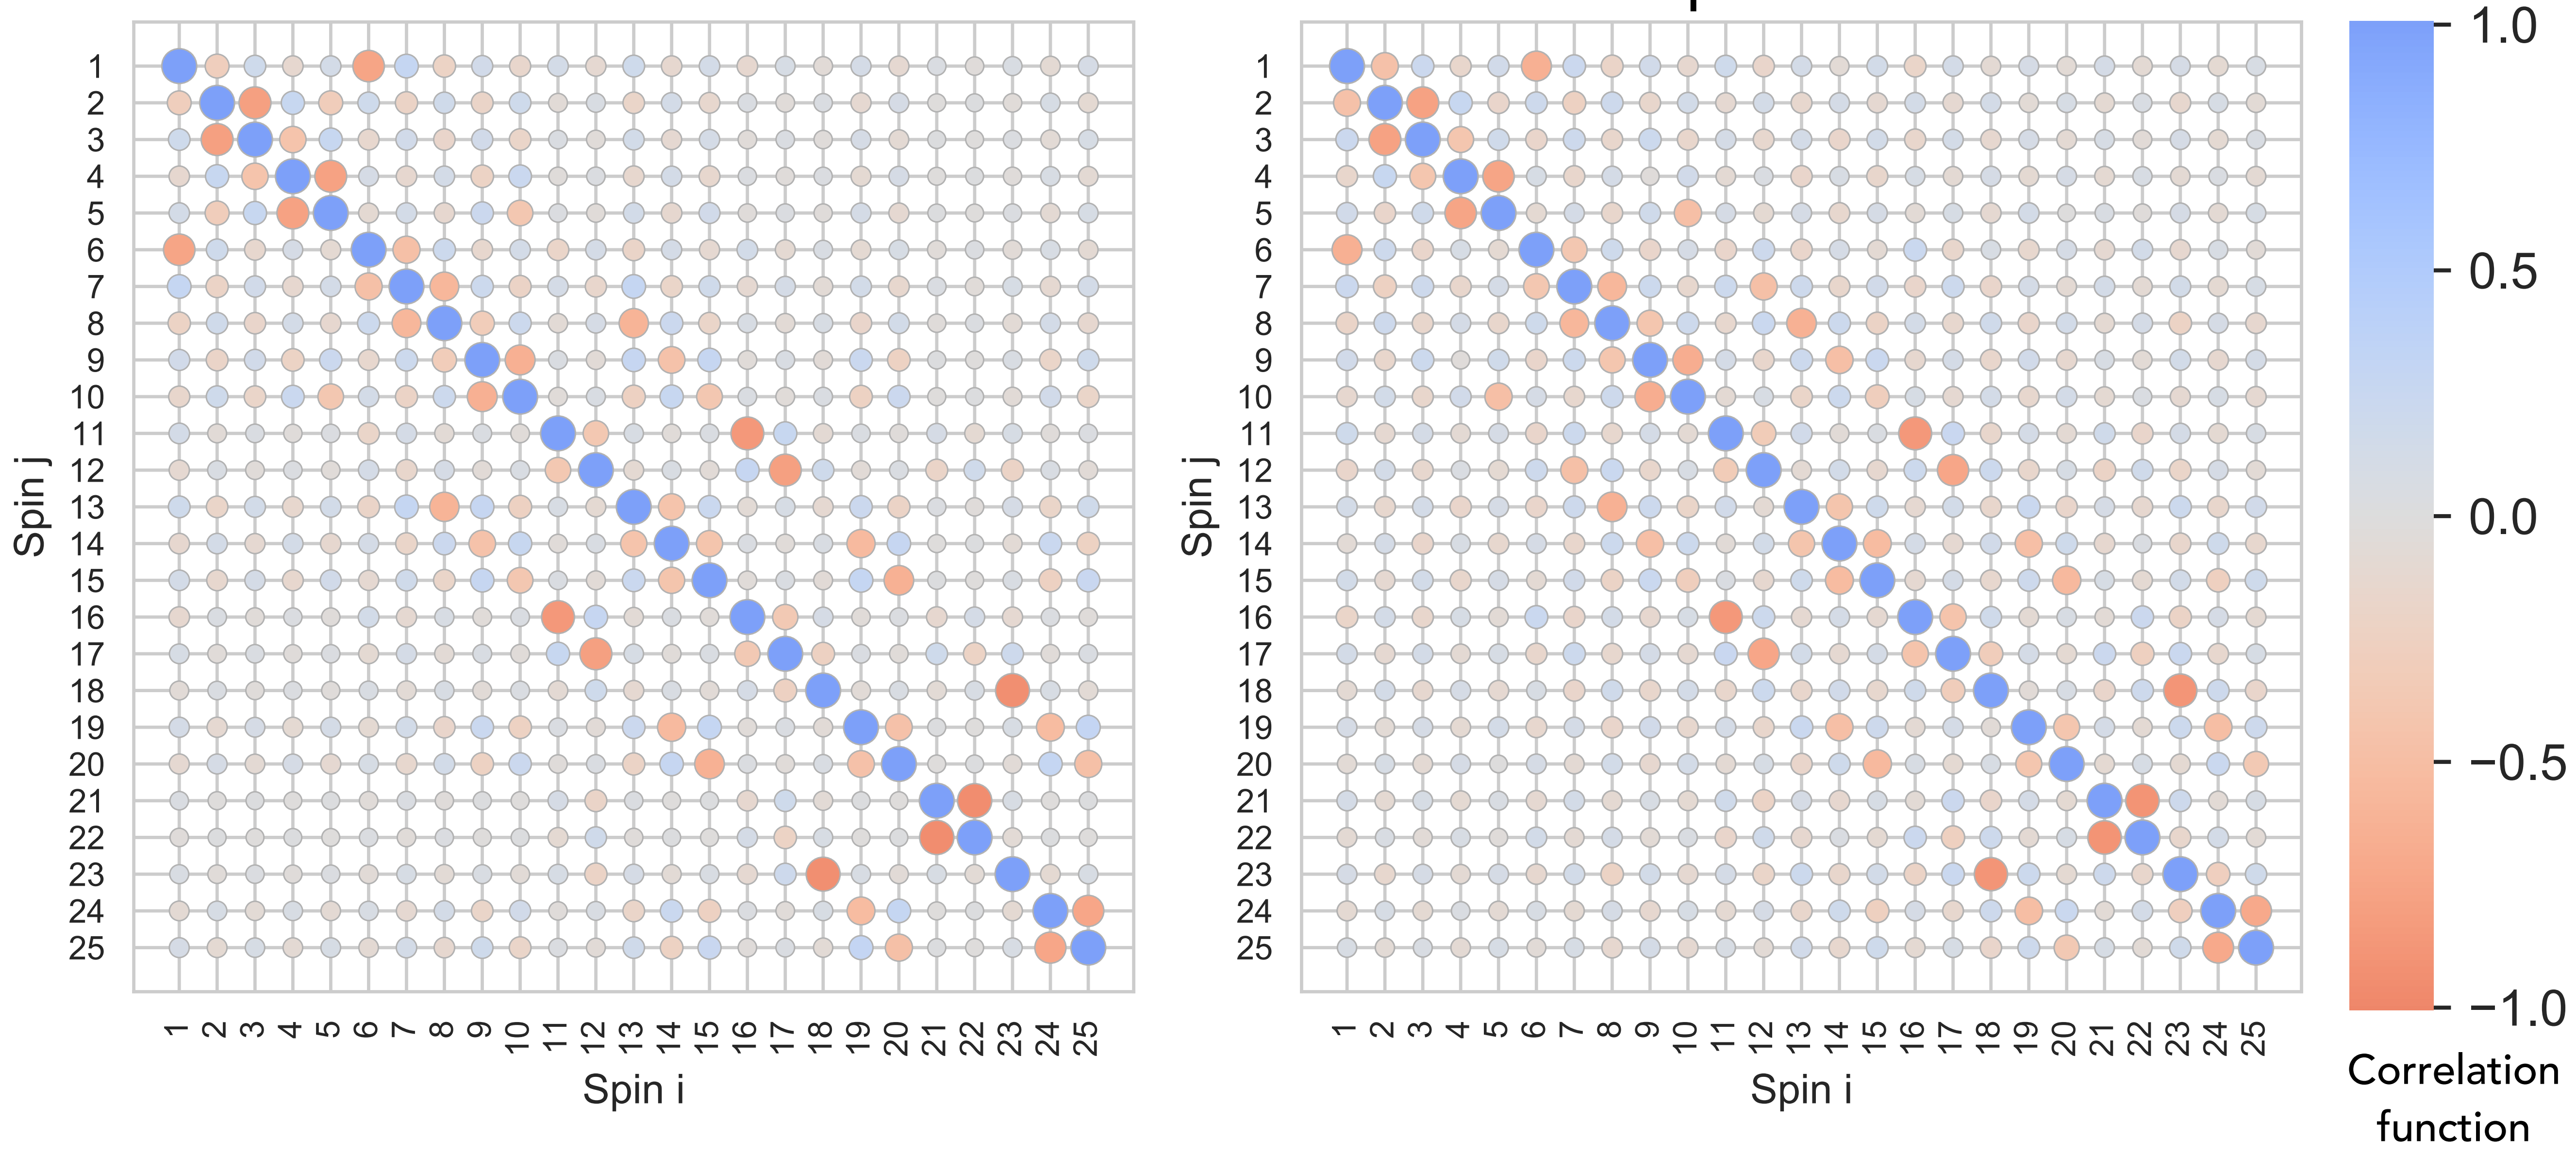

[HKT+21] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

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.

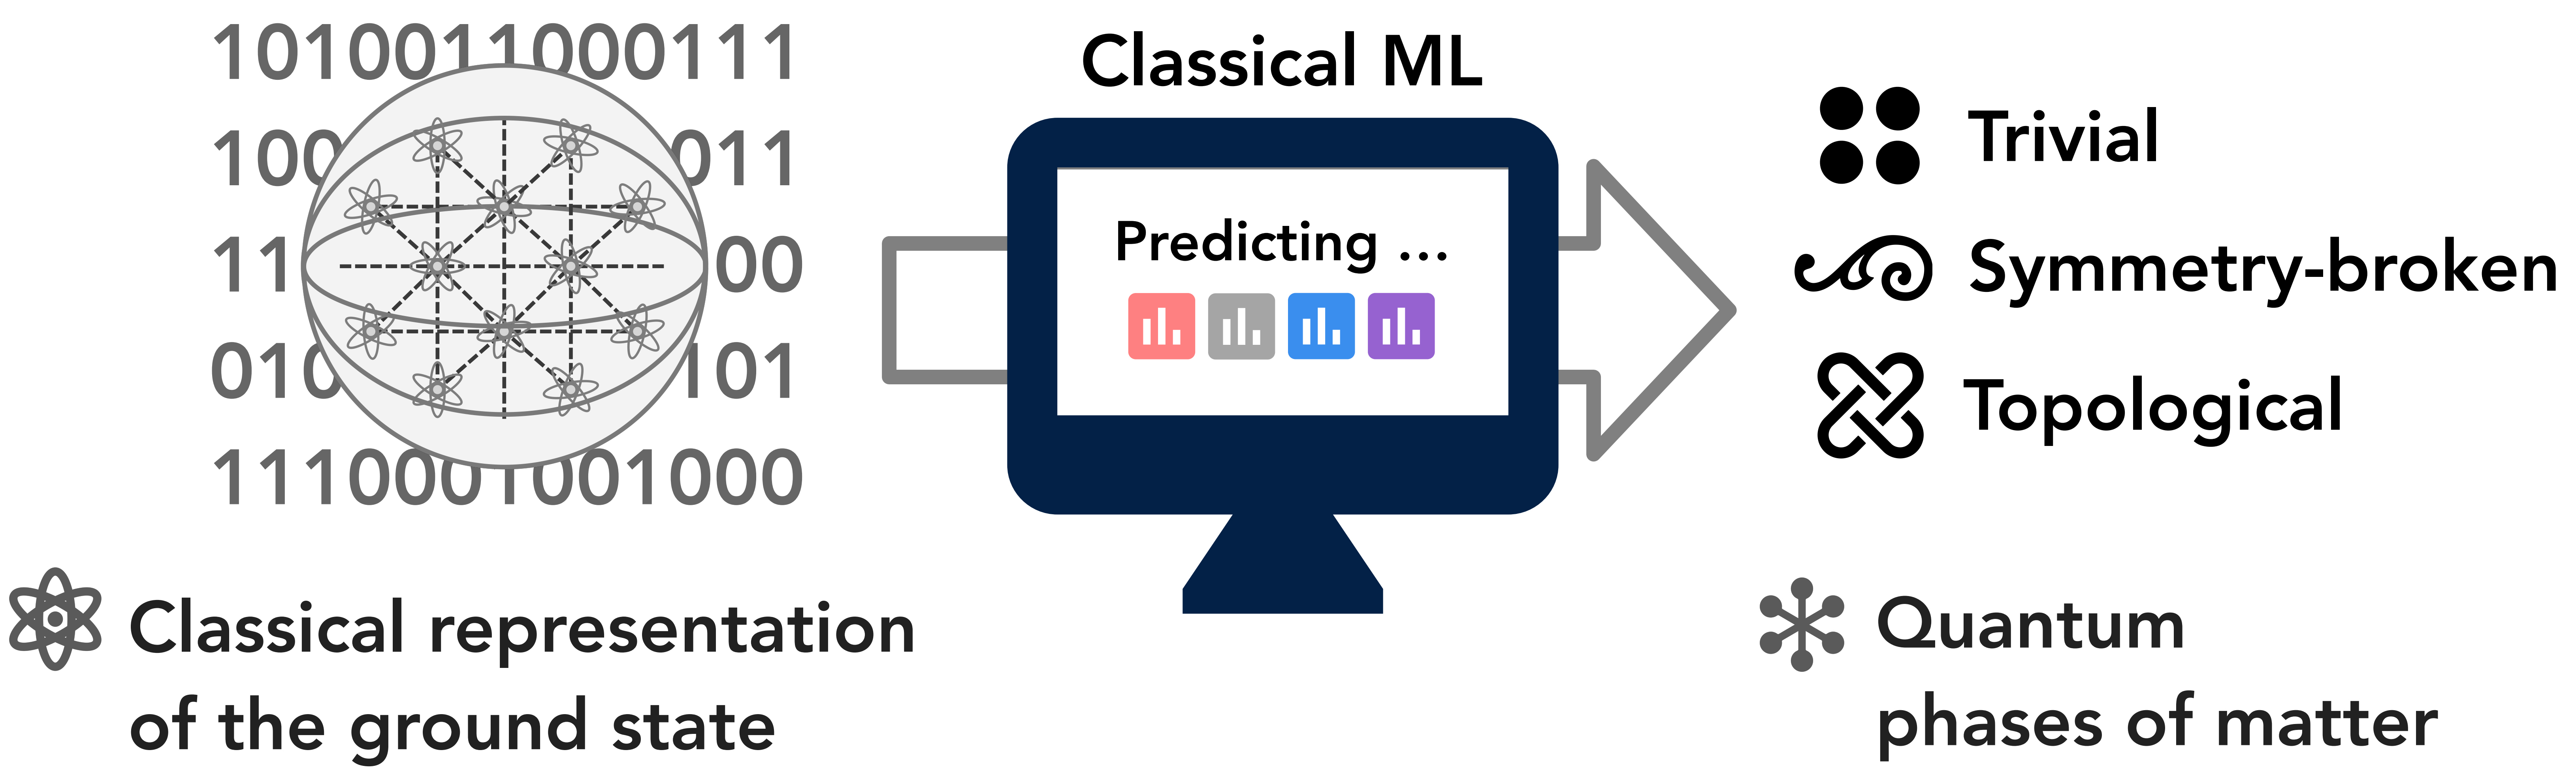

[HKT+22] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

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.

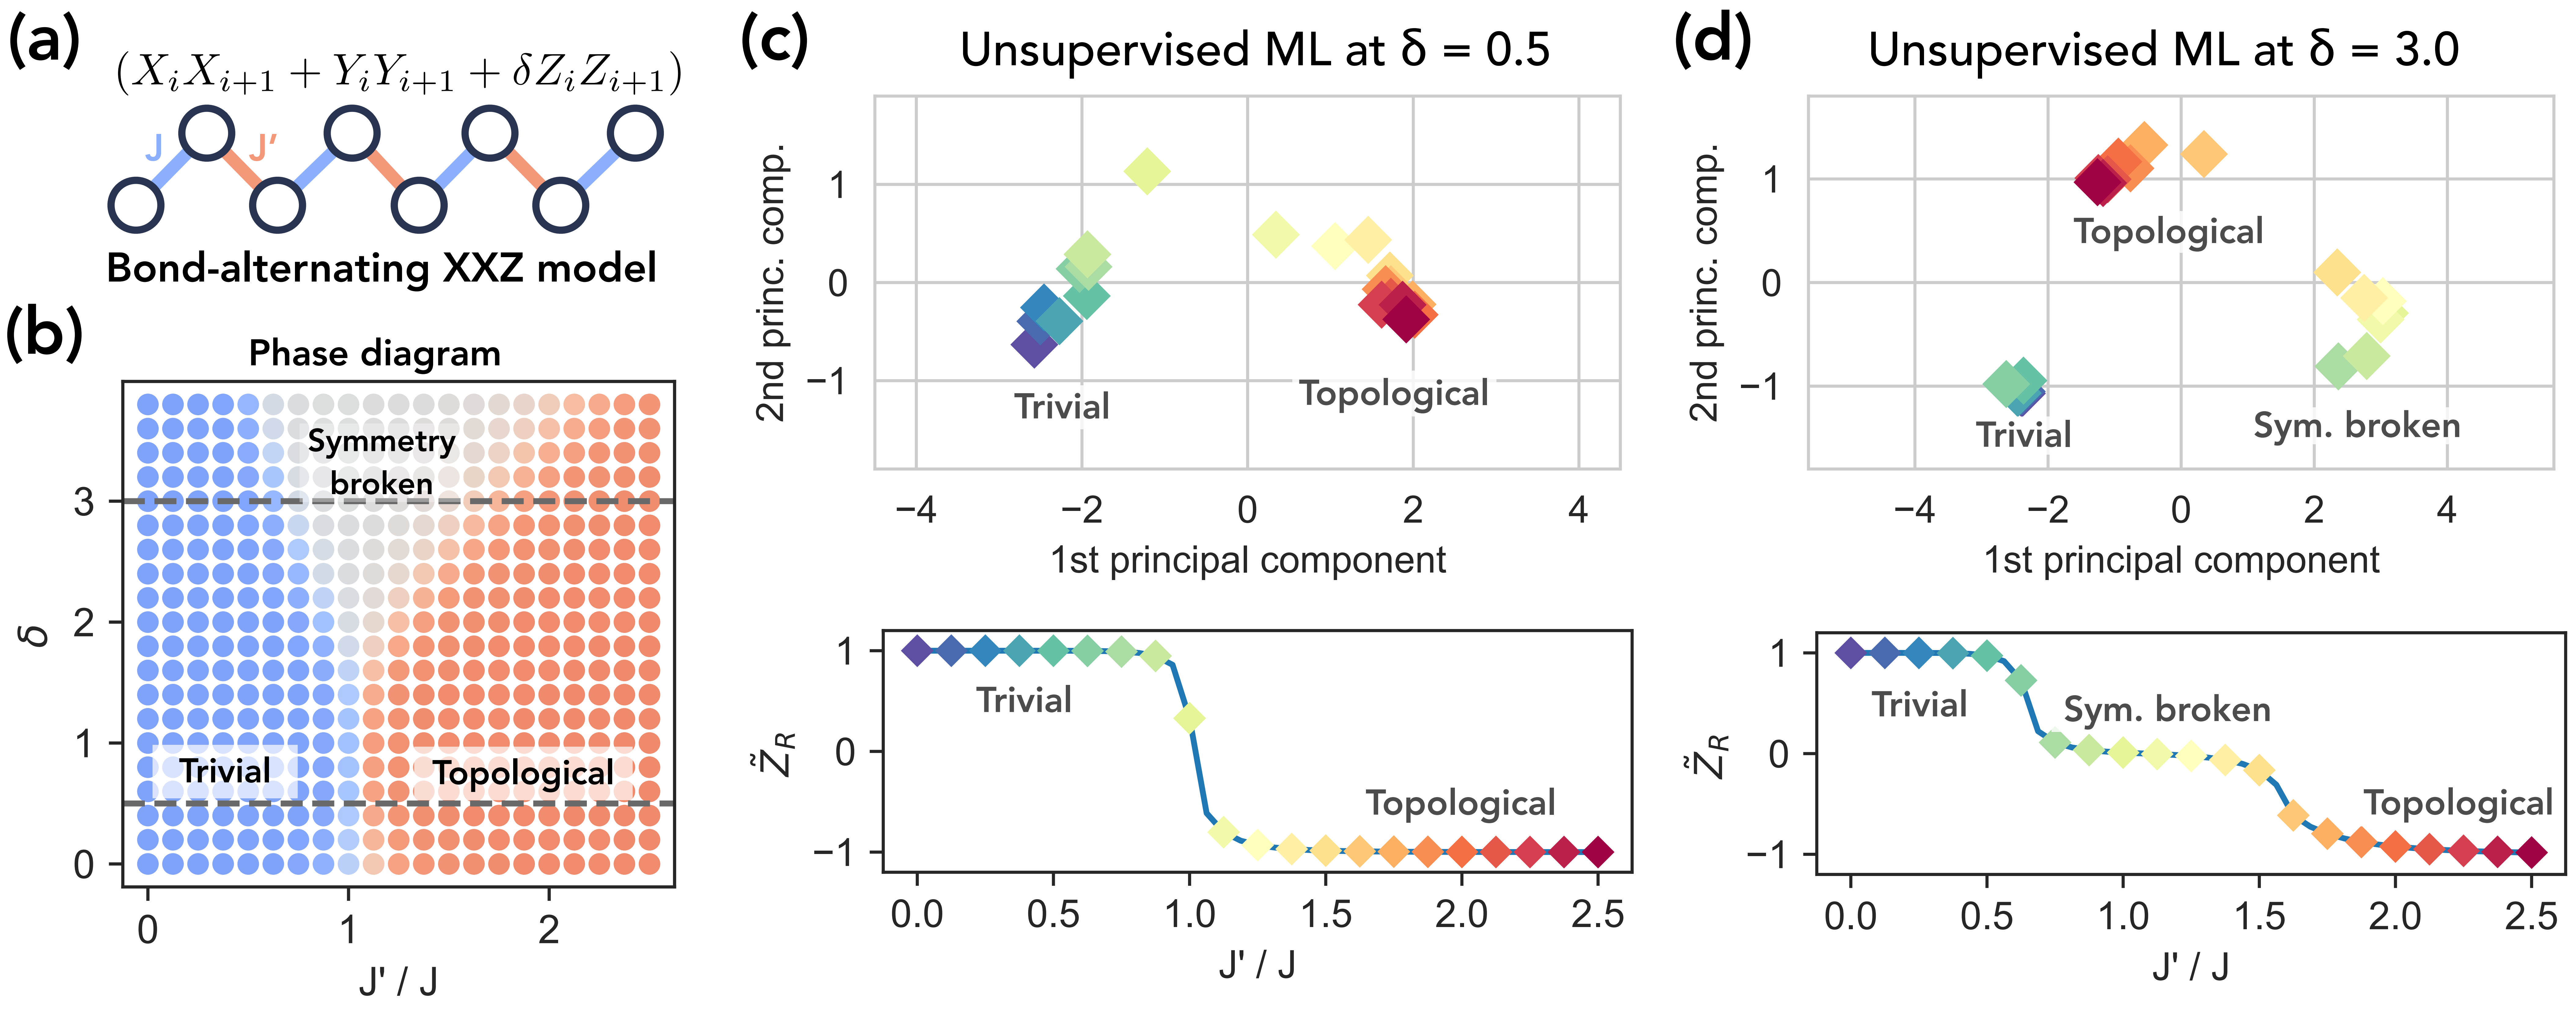

[HKT+22] Huang, Kueng, Torlai, Albert, Preskill. Provably efficient machine learning for quantum many-body problems, *Science*, 2022.

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?

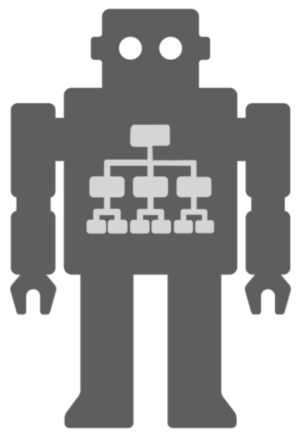

Learning with quantum machines

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?

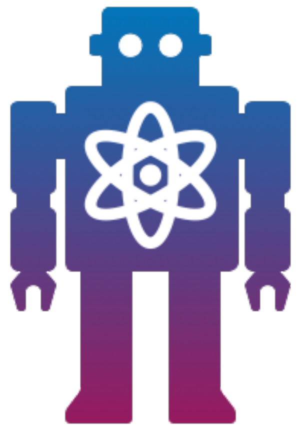

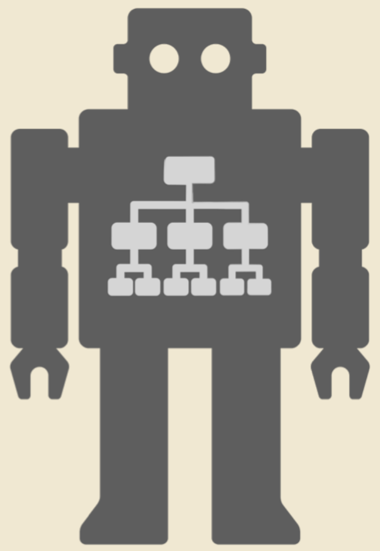

Classical agent

Receive, process, and store classical information

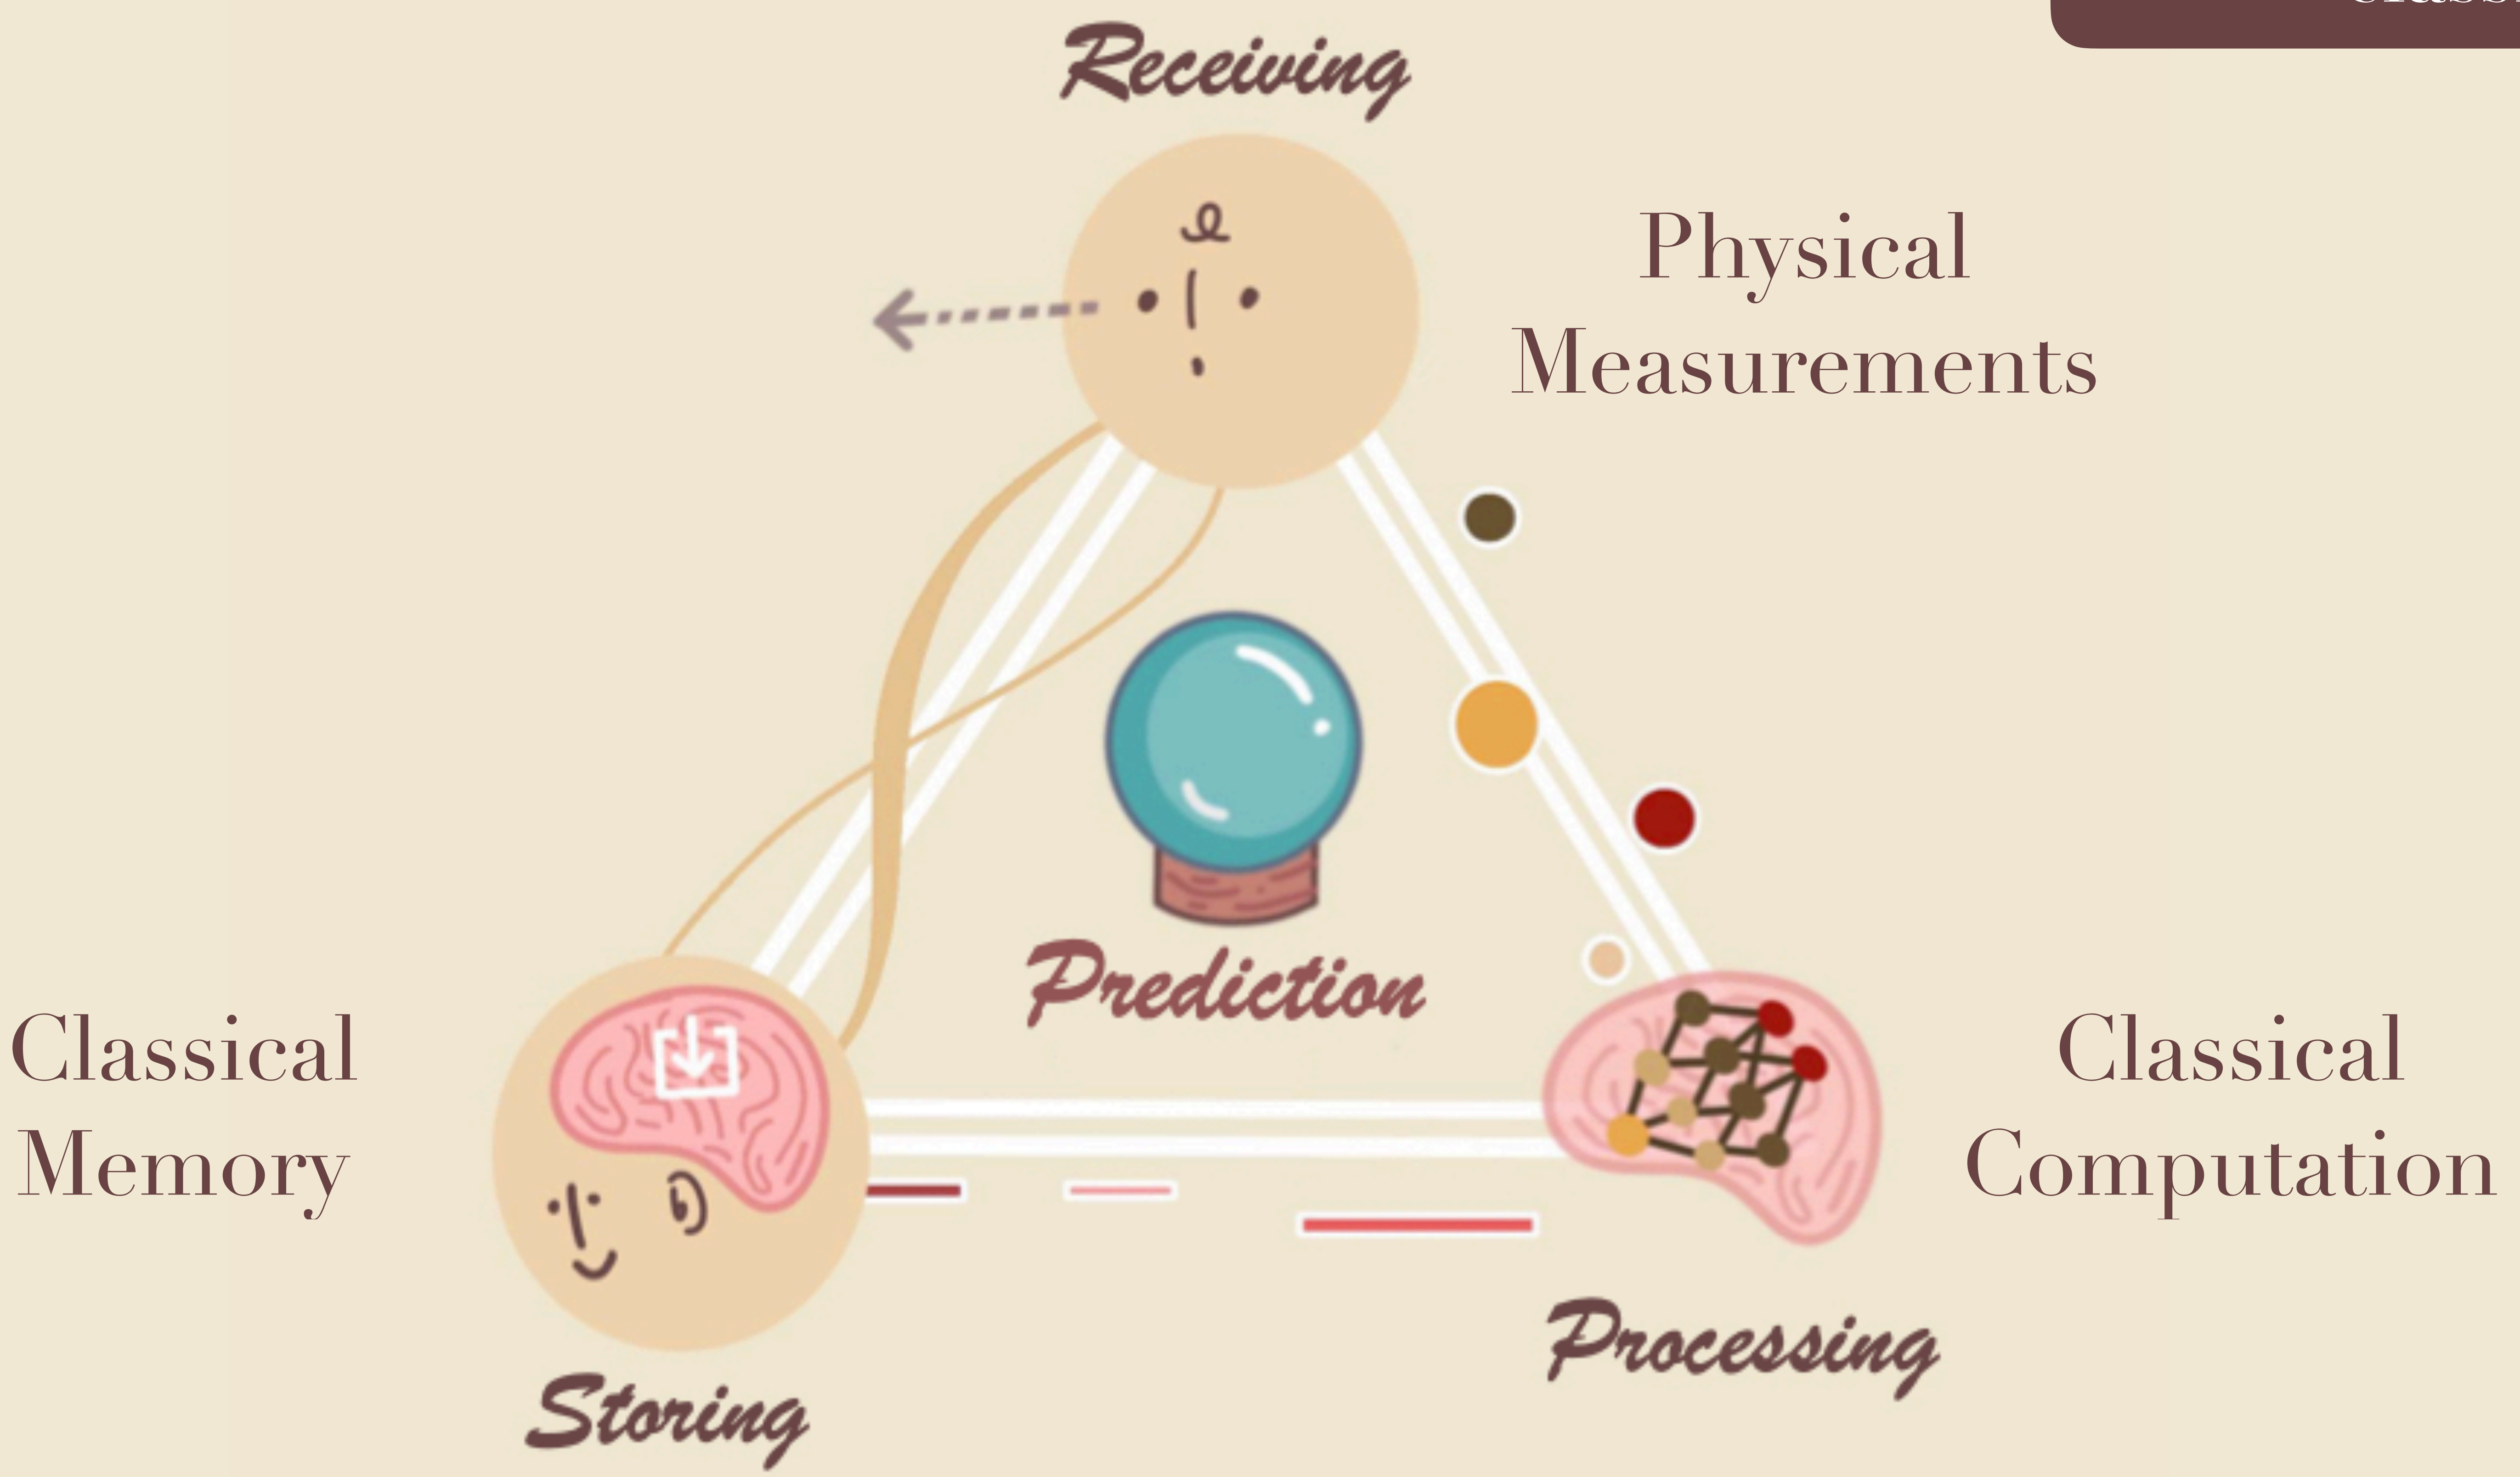

[HKP21] Huang, Kueng, Preskill. Information-theoretic bounds on quantum advantage in machine learning, *Physical Review Letters*, 2021.
[CCHL21] Chen, Cotler, Huang, Li. Exponential separations in learning with and without quantum memory, *FOCS*, 2021.
[HBC+] Huang, et all. Quantum advantage in learning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

Quantum agent

Receive, process, and store quantum information

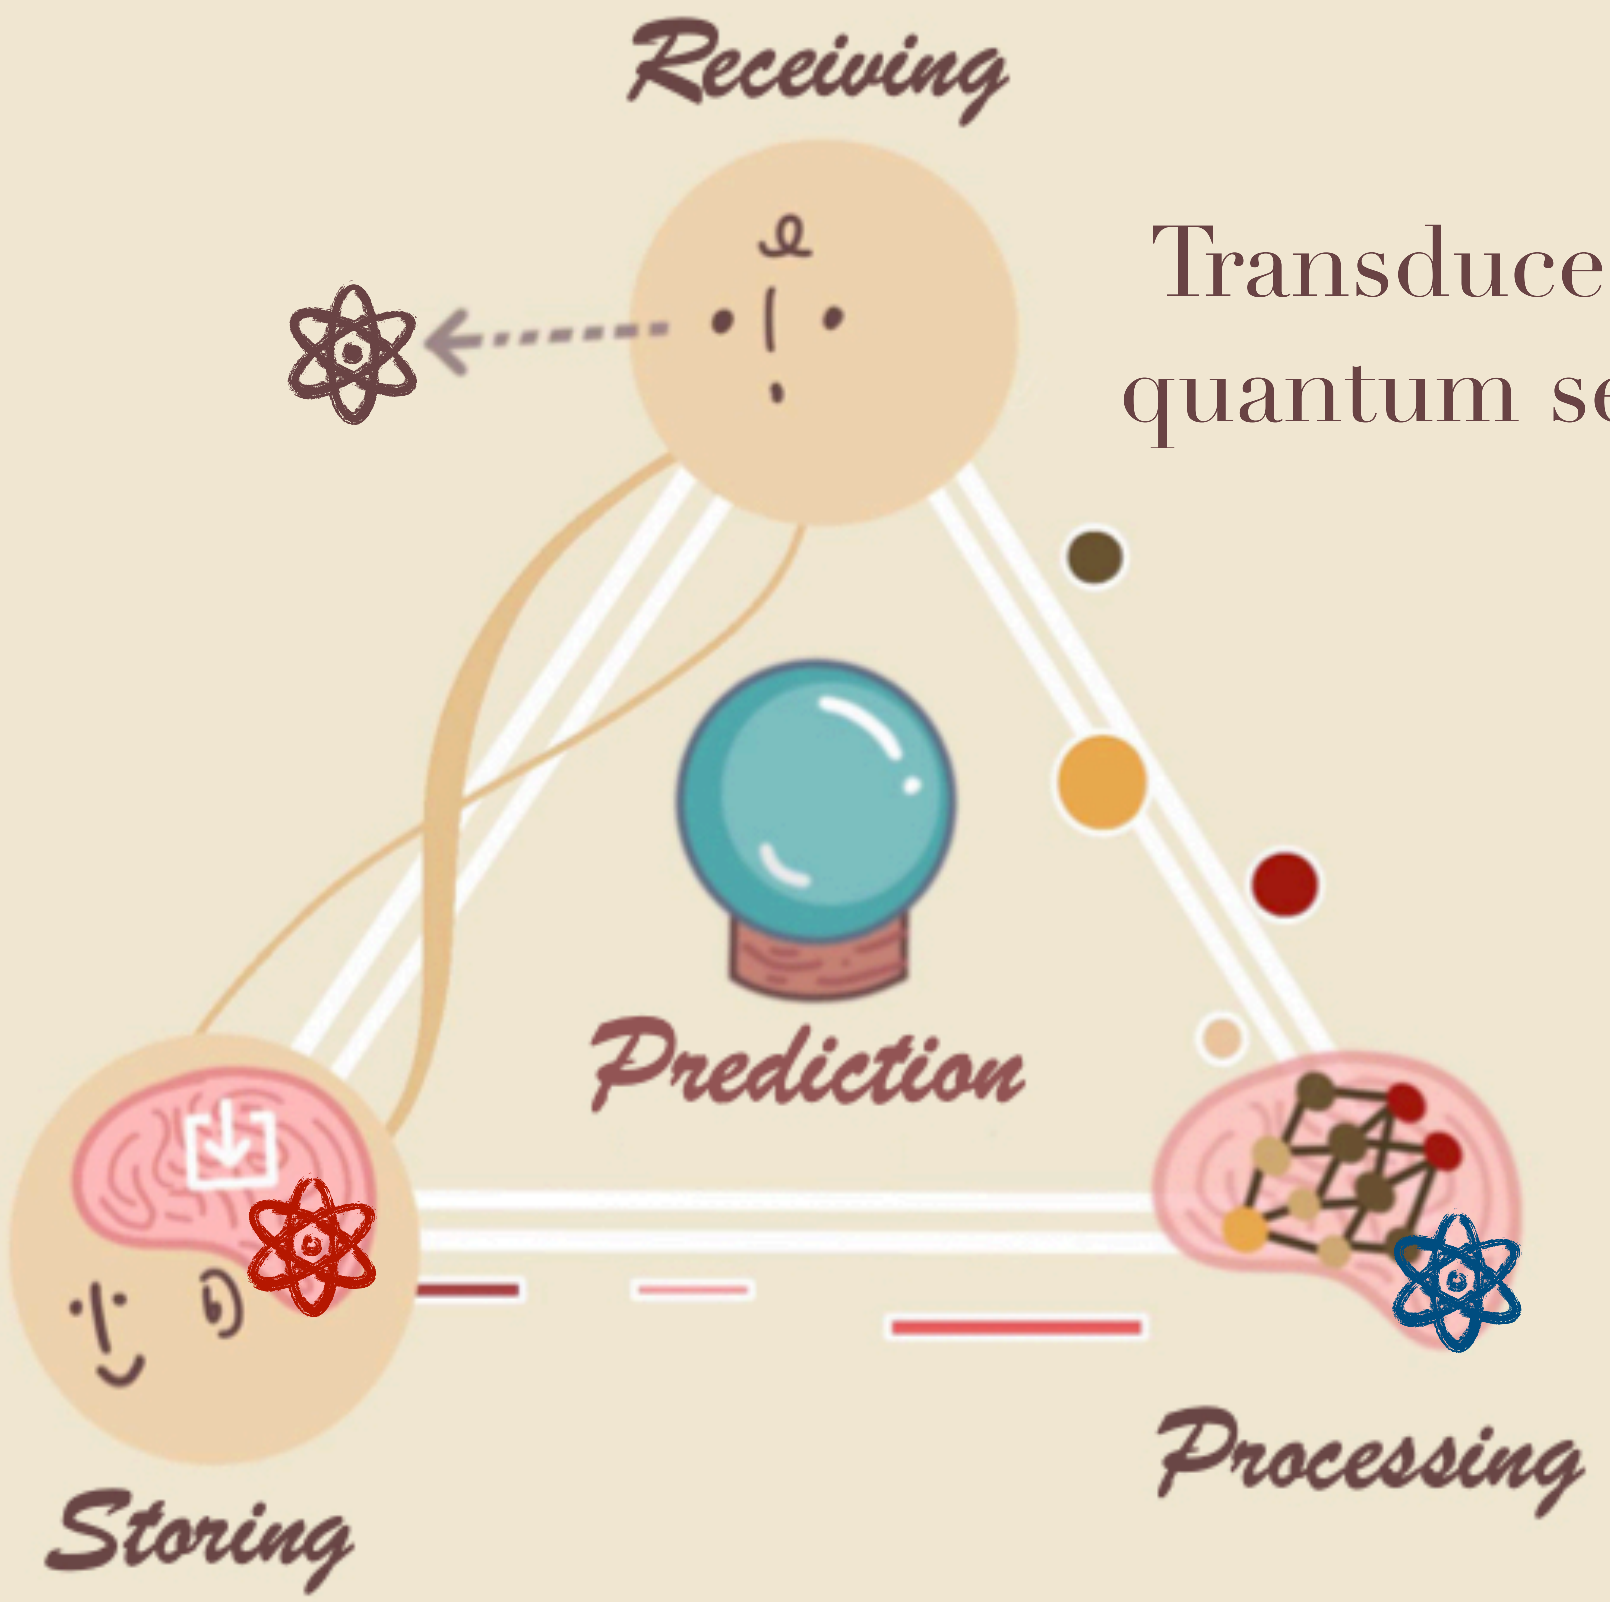

Transduce from quantum sensors

Quantum Memory

Quantum Computation

[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, et all. 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?

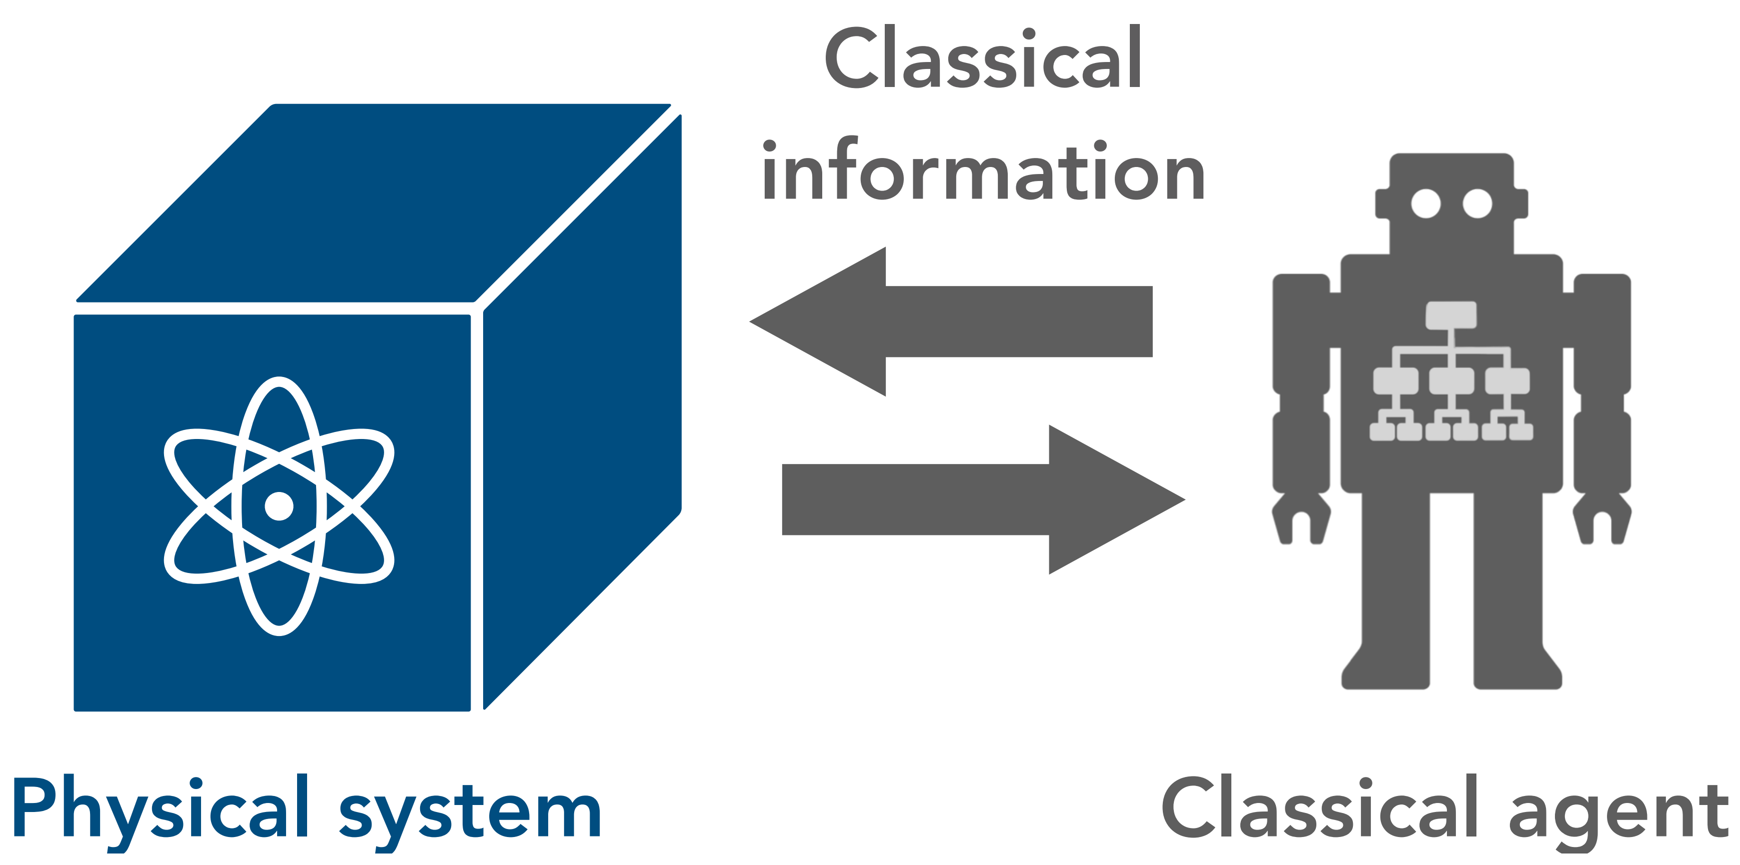

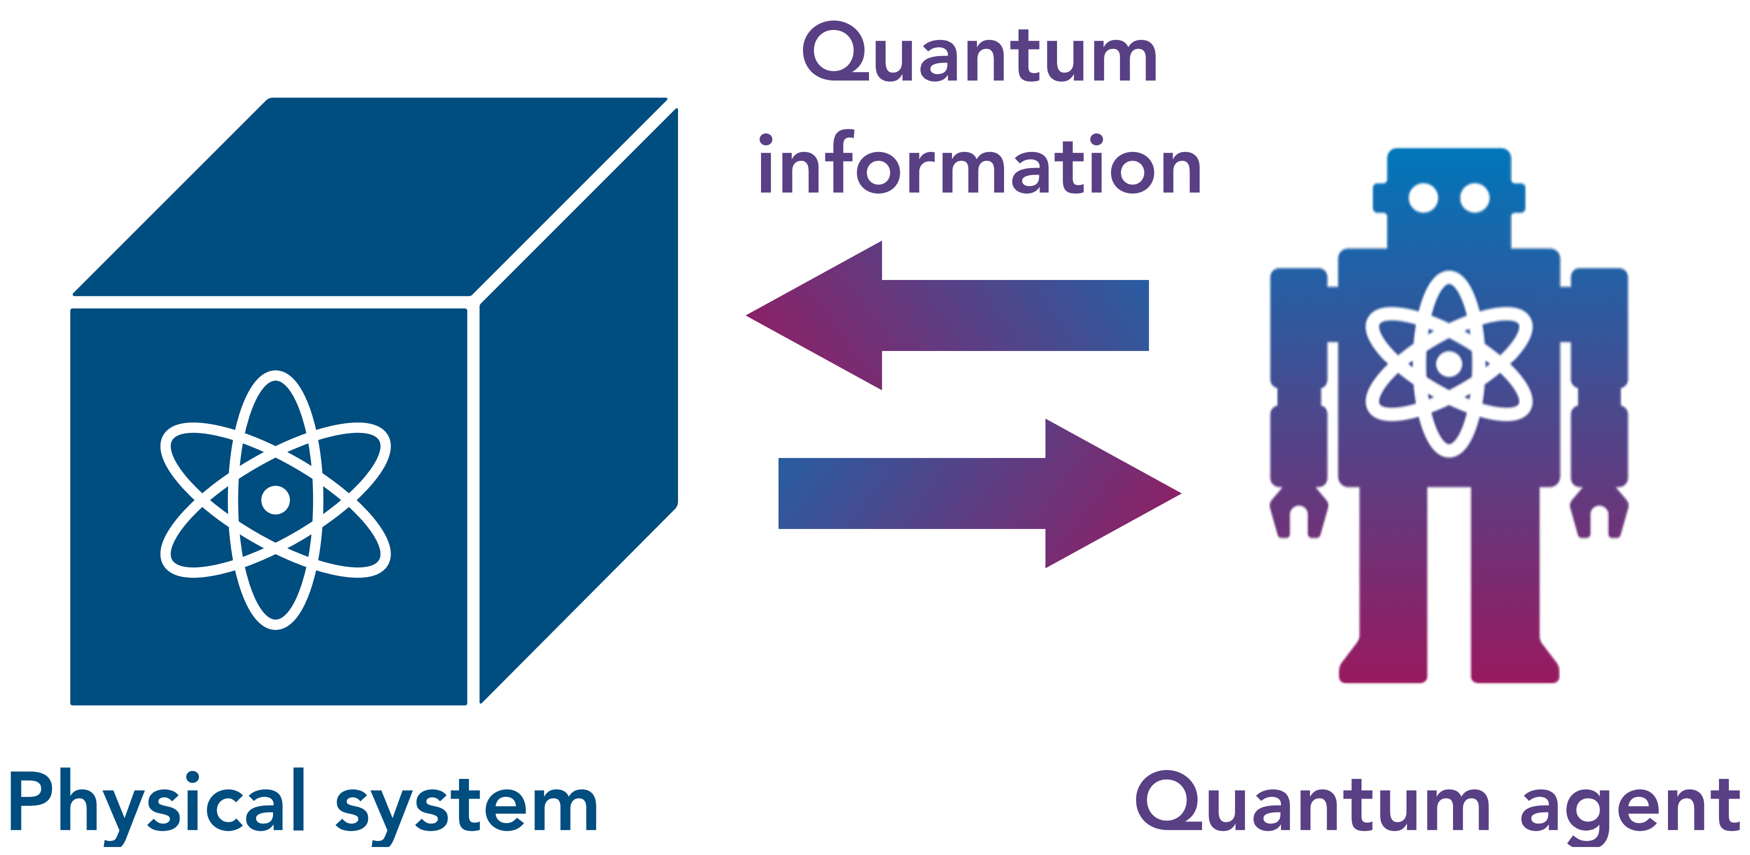

[HKP21] Huang, Kueng, Preskill. Information-theoretic bounds on quantum advantage in machine learning, *Physical Review Letters*, 2021.
[CCHL21] Chen, Cotler, Huang, Li. Exponential separations in learning with and without quantum memory, *FOCS*, 2021.
[HBC+] Huang, et all. Quantum advantage in learning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

Learning a state

- Assume the only unknown that we care about is an n-qubit state ρ.
- Classical agents can perform any measurement on ρ in each experiment.
- Quantum agents can obtain and store ρ coherently from each experiment.

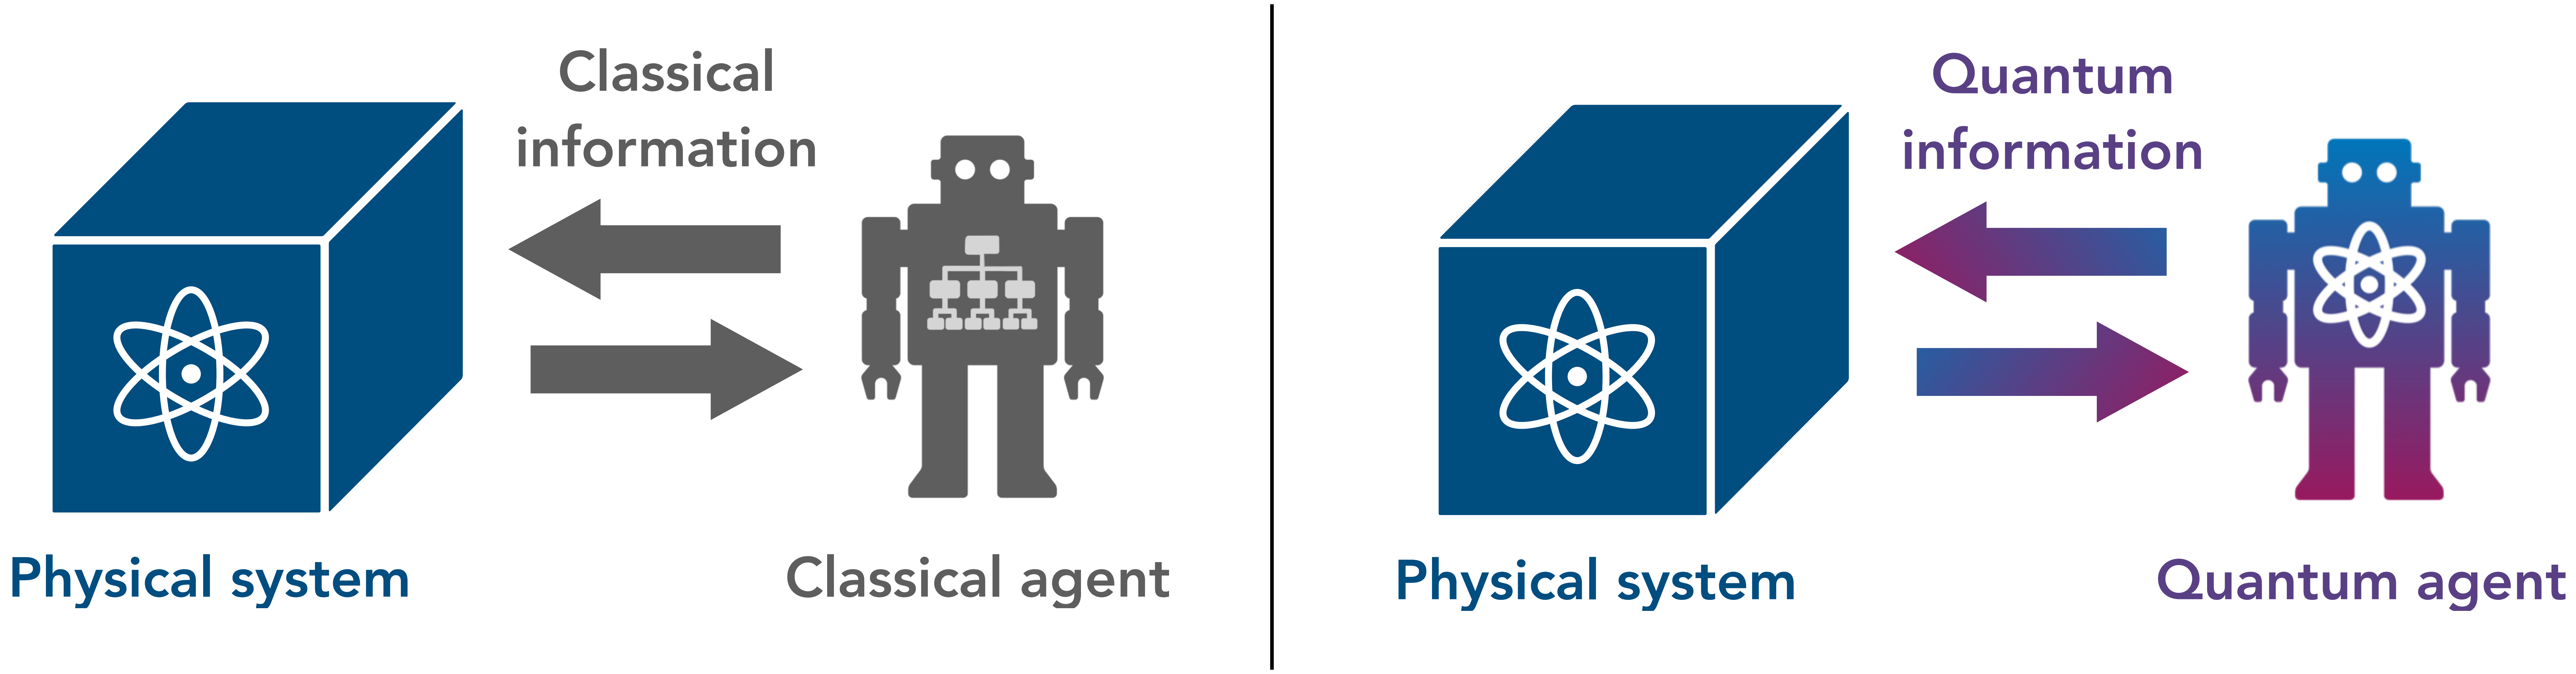

What can classical agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

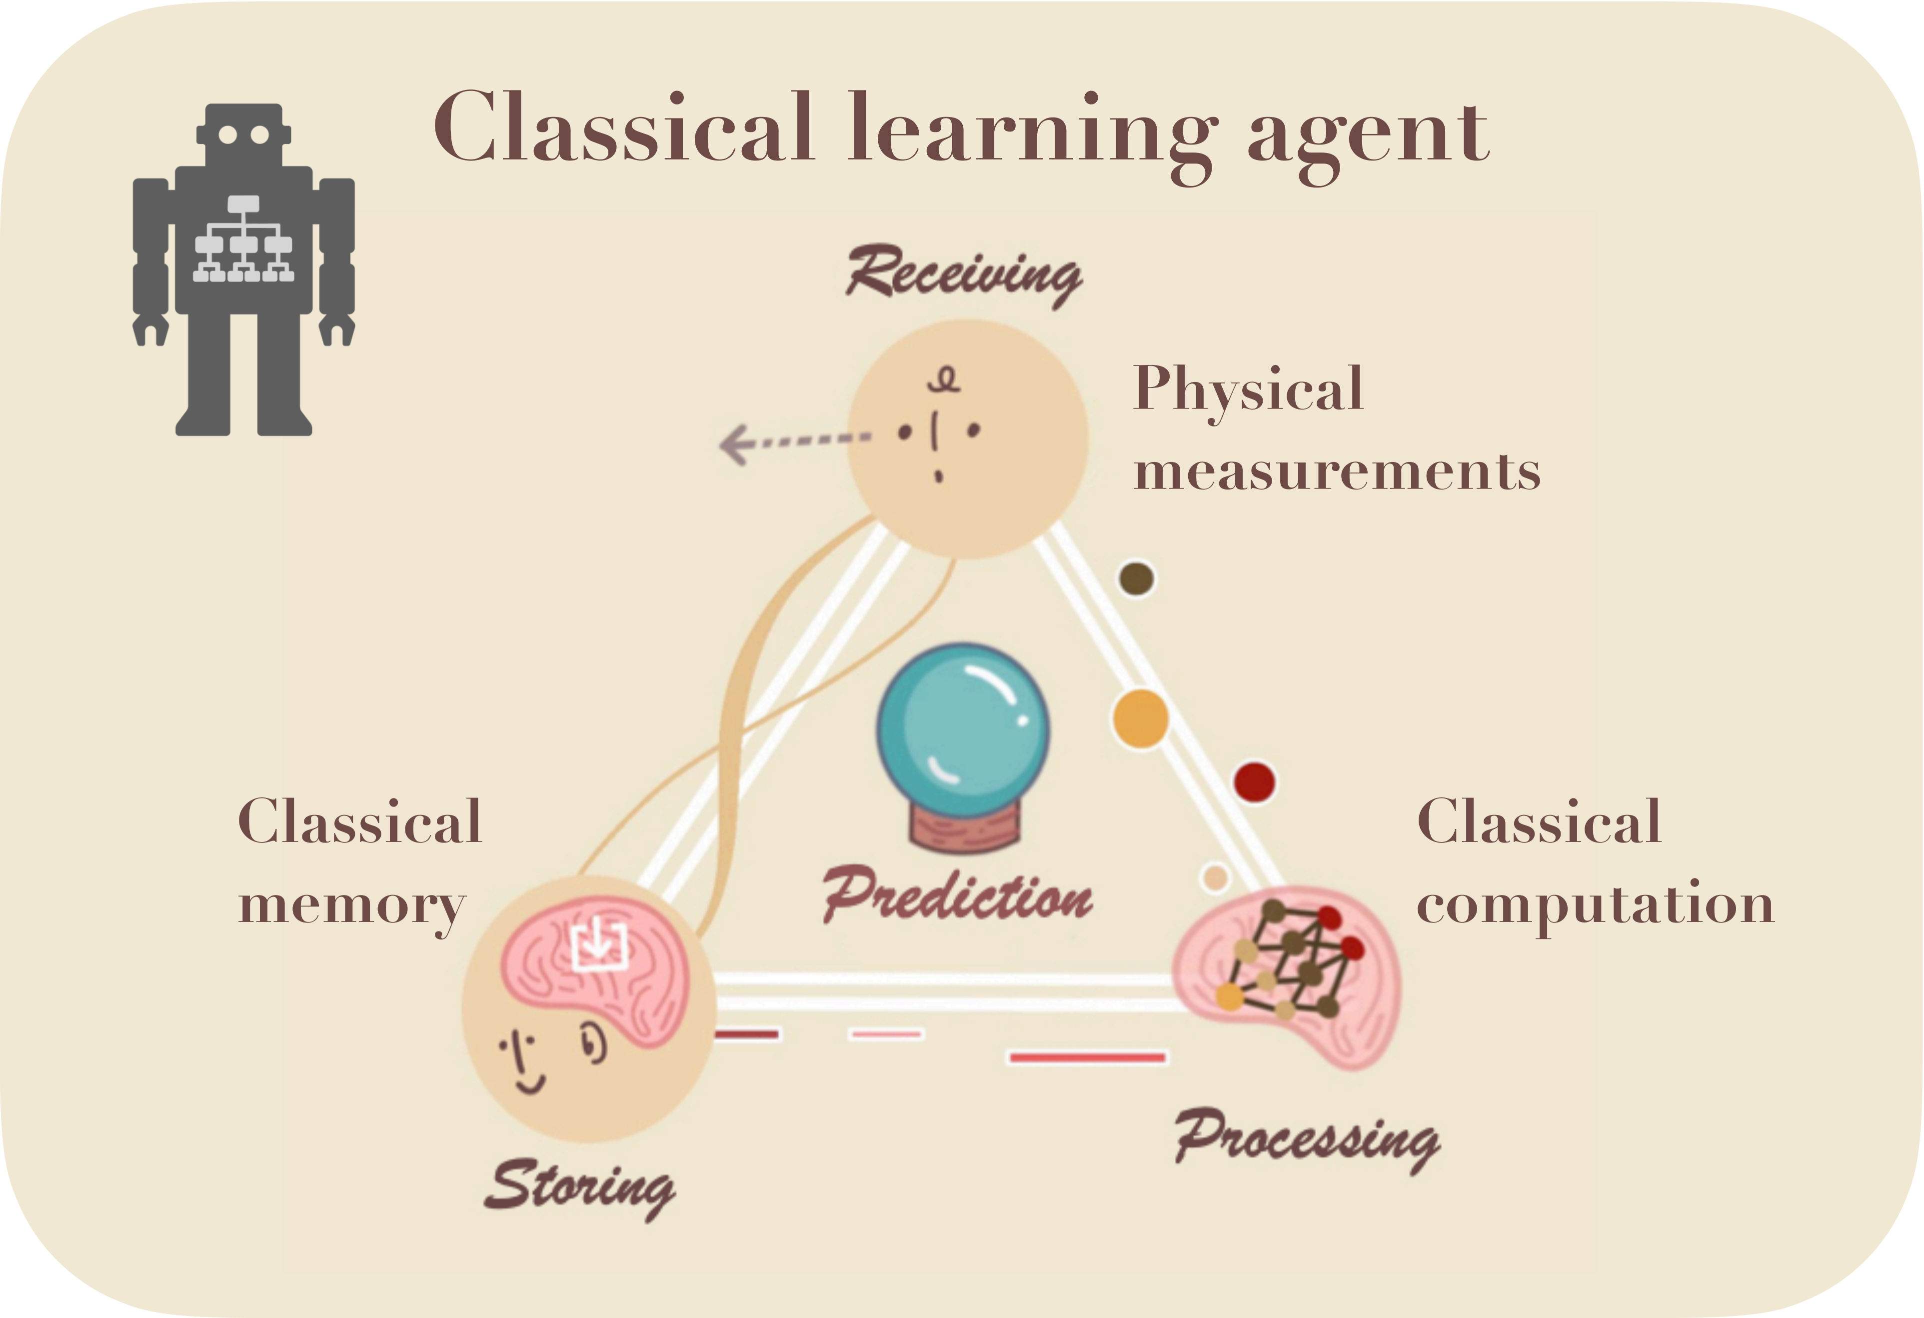

What can classical agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

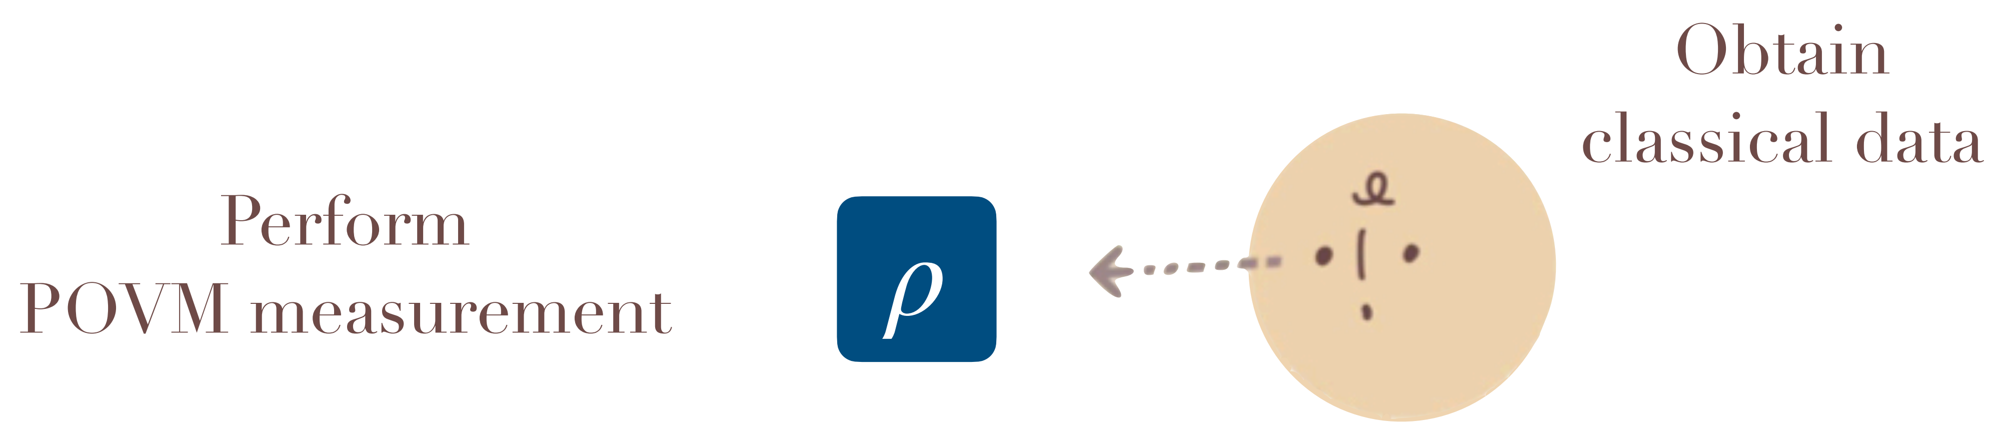

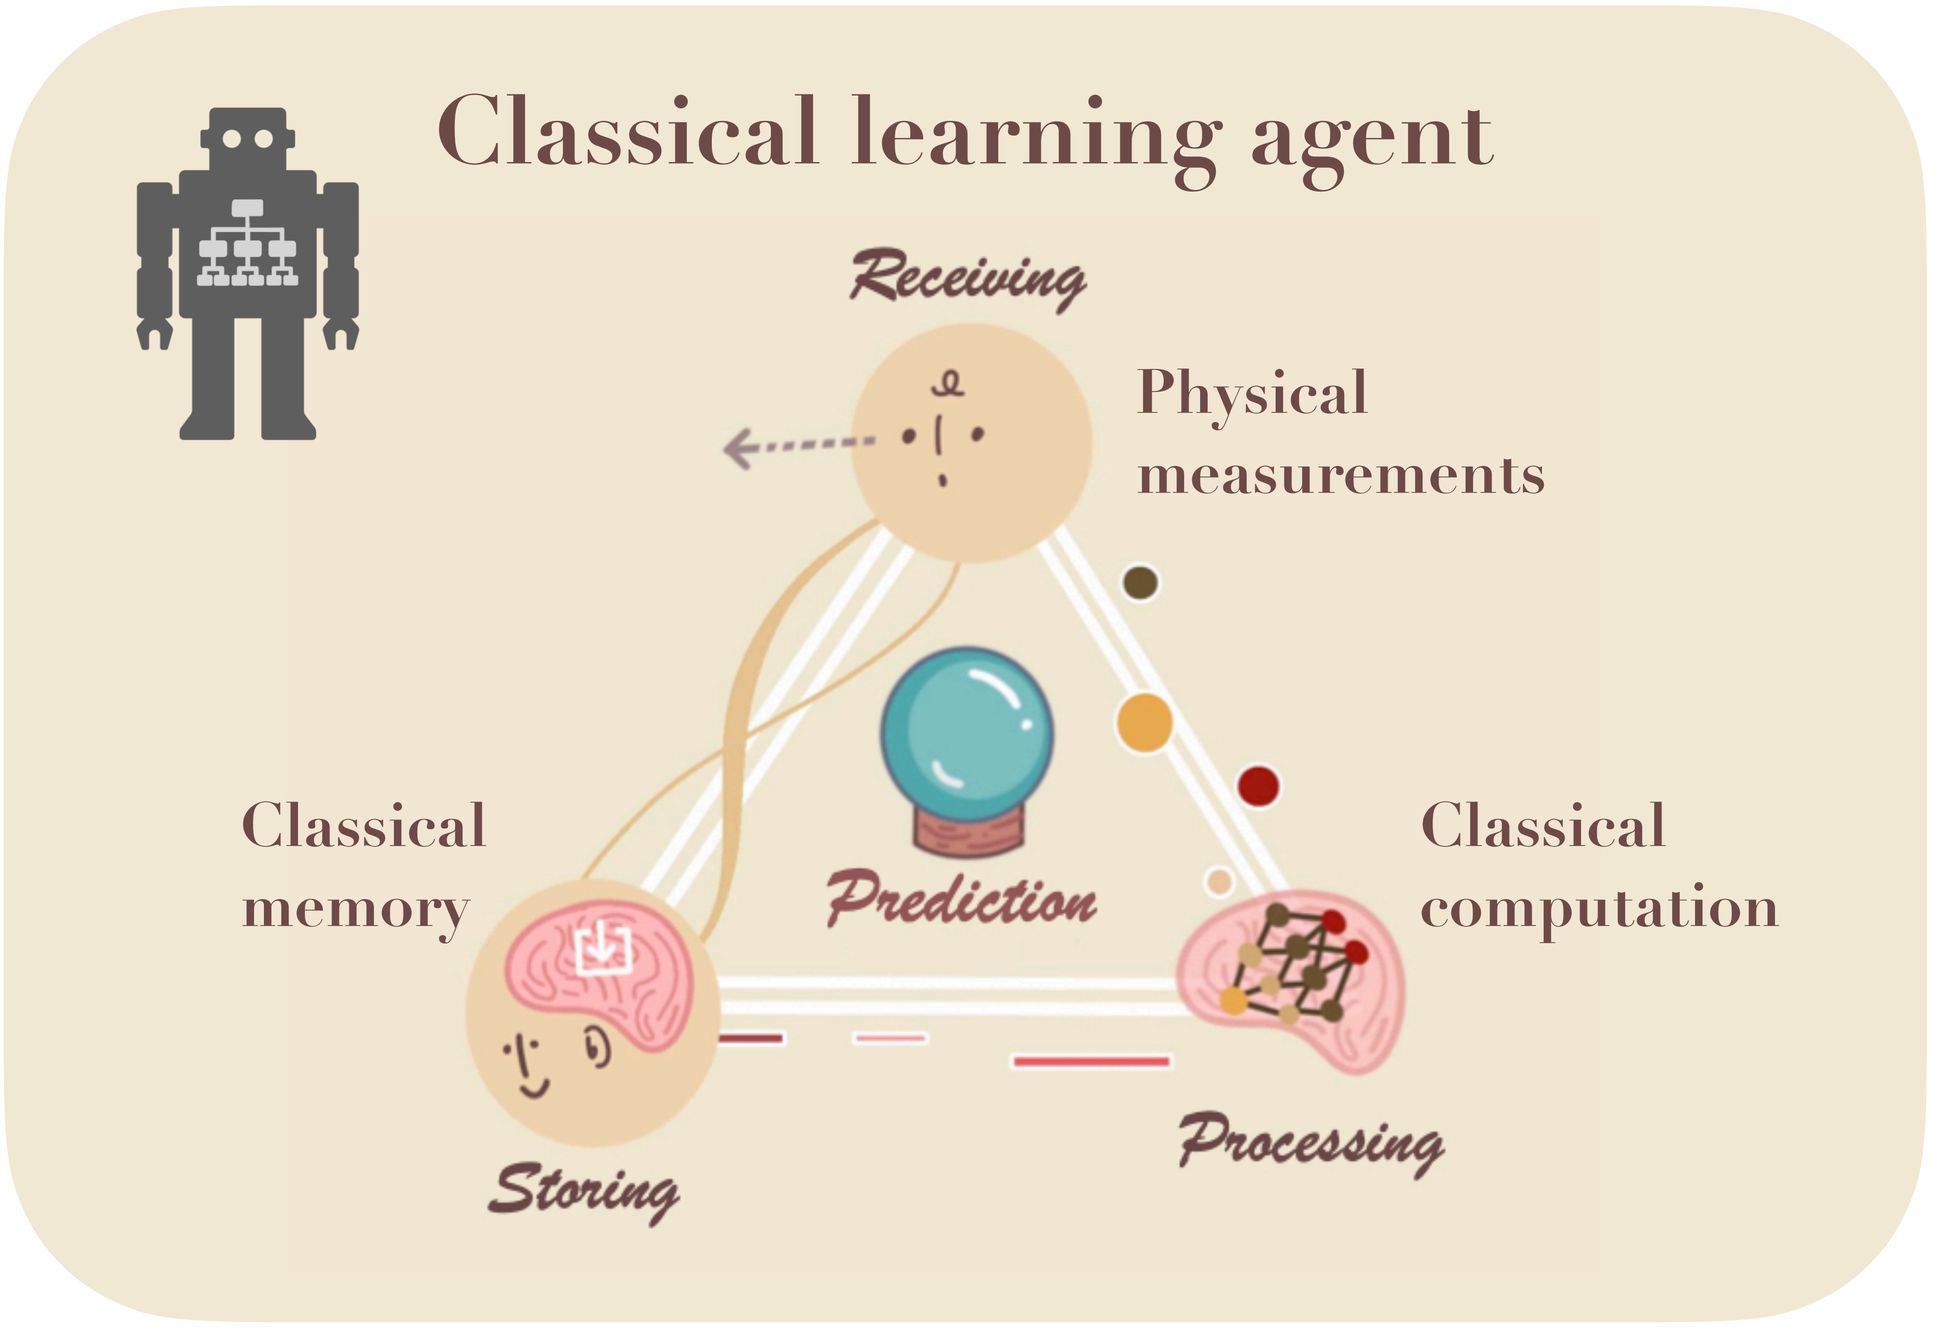

What can classical agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

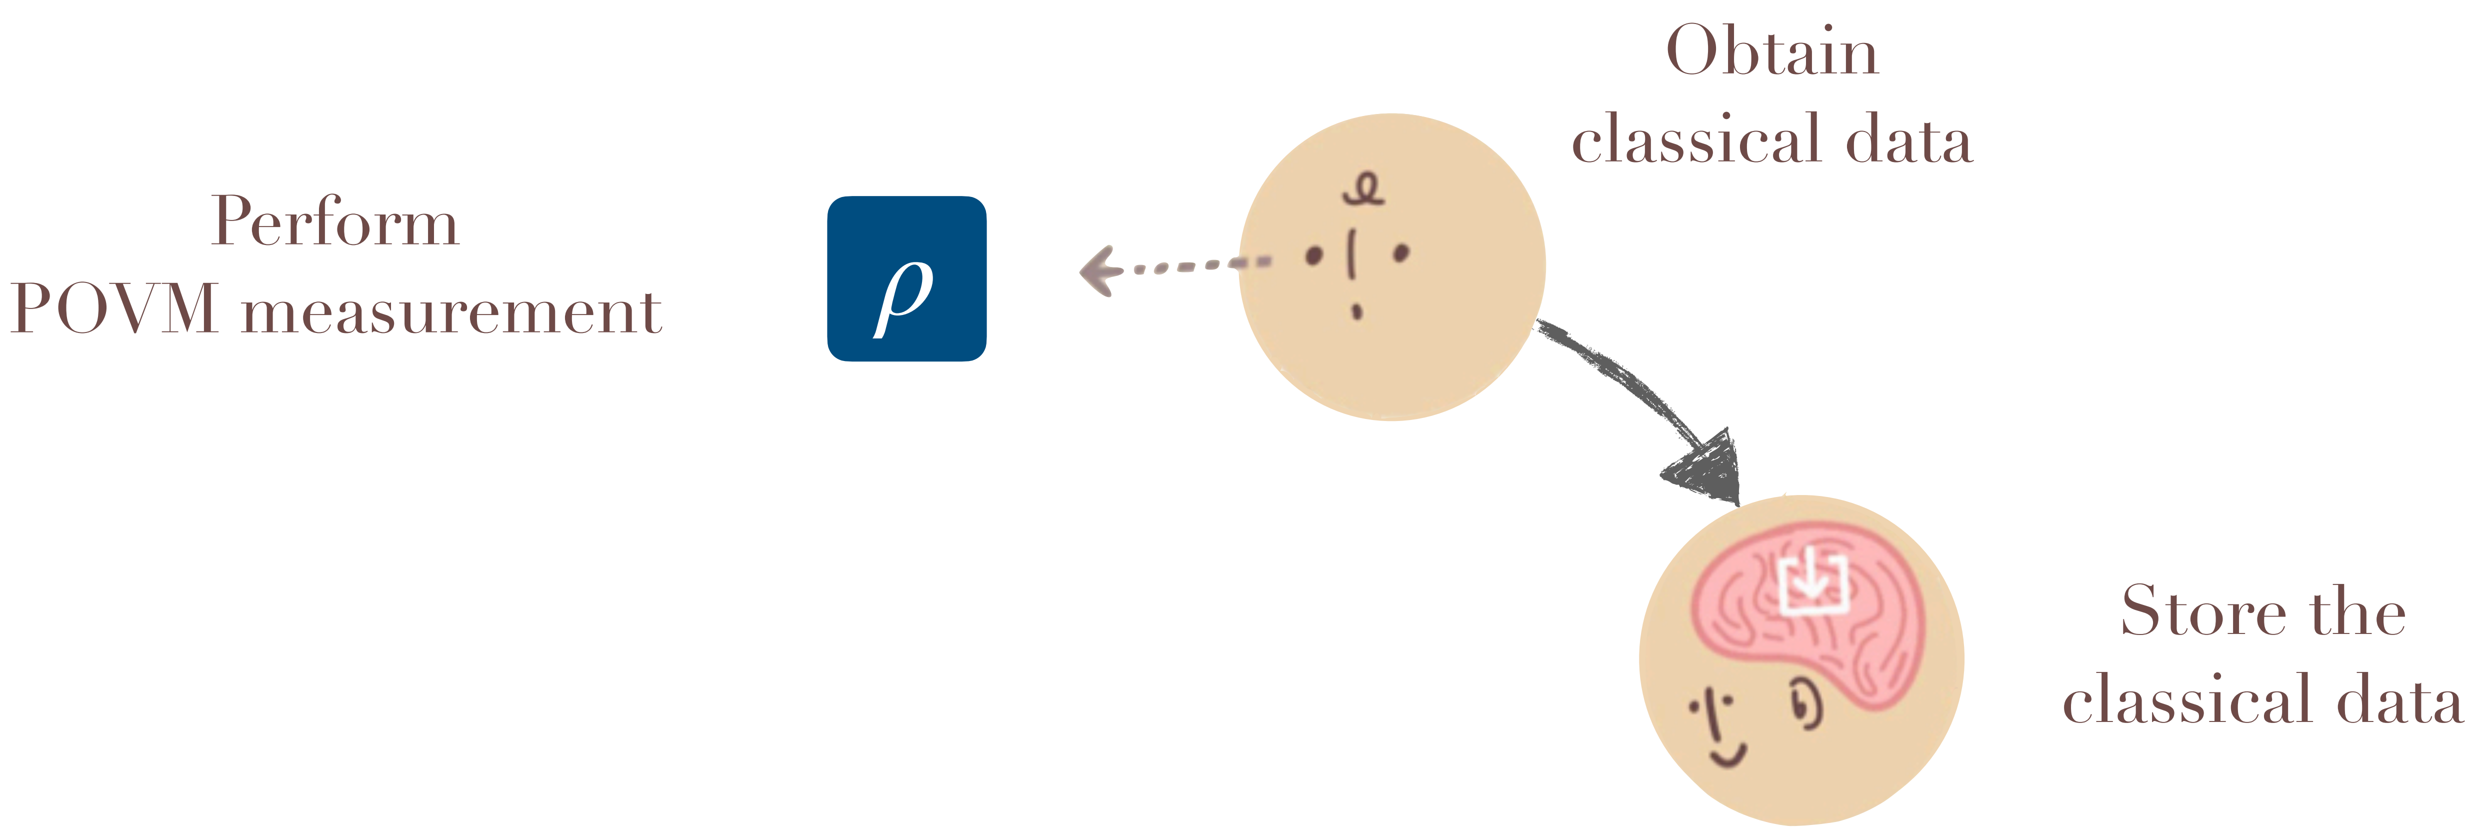

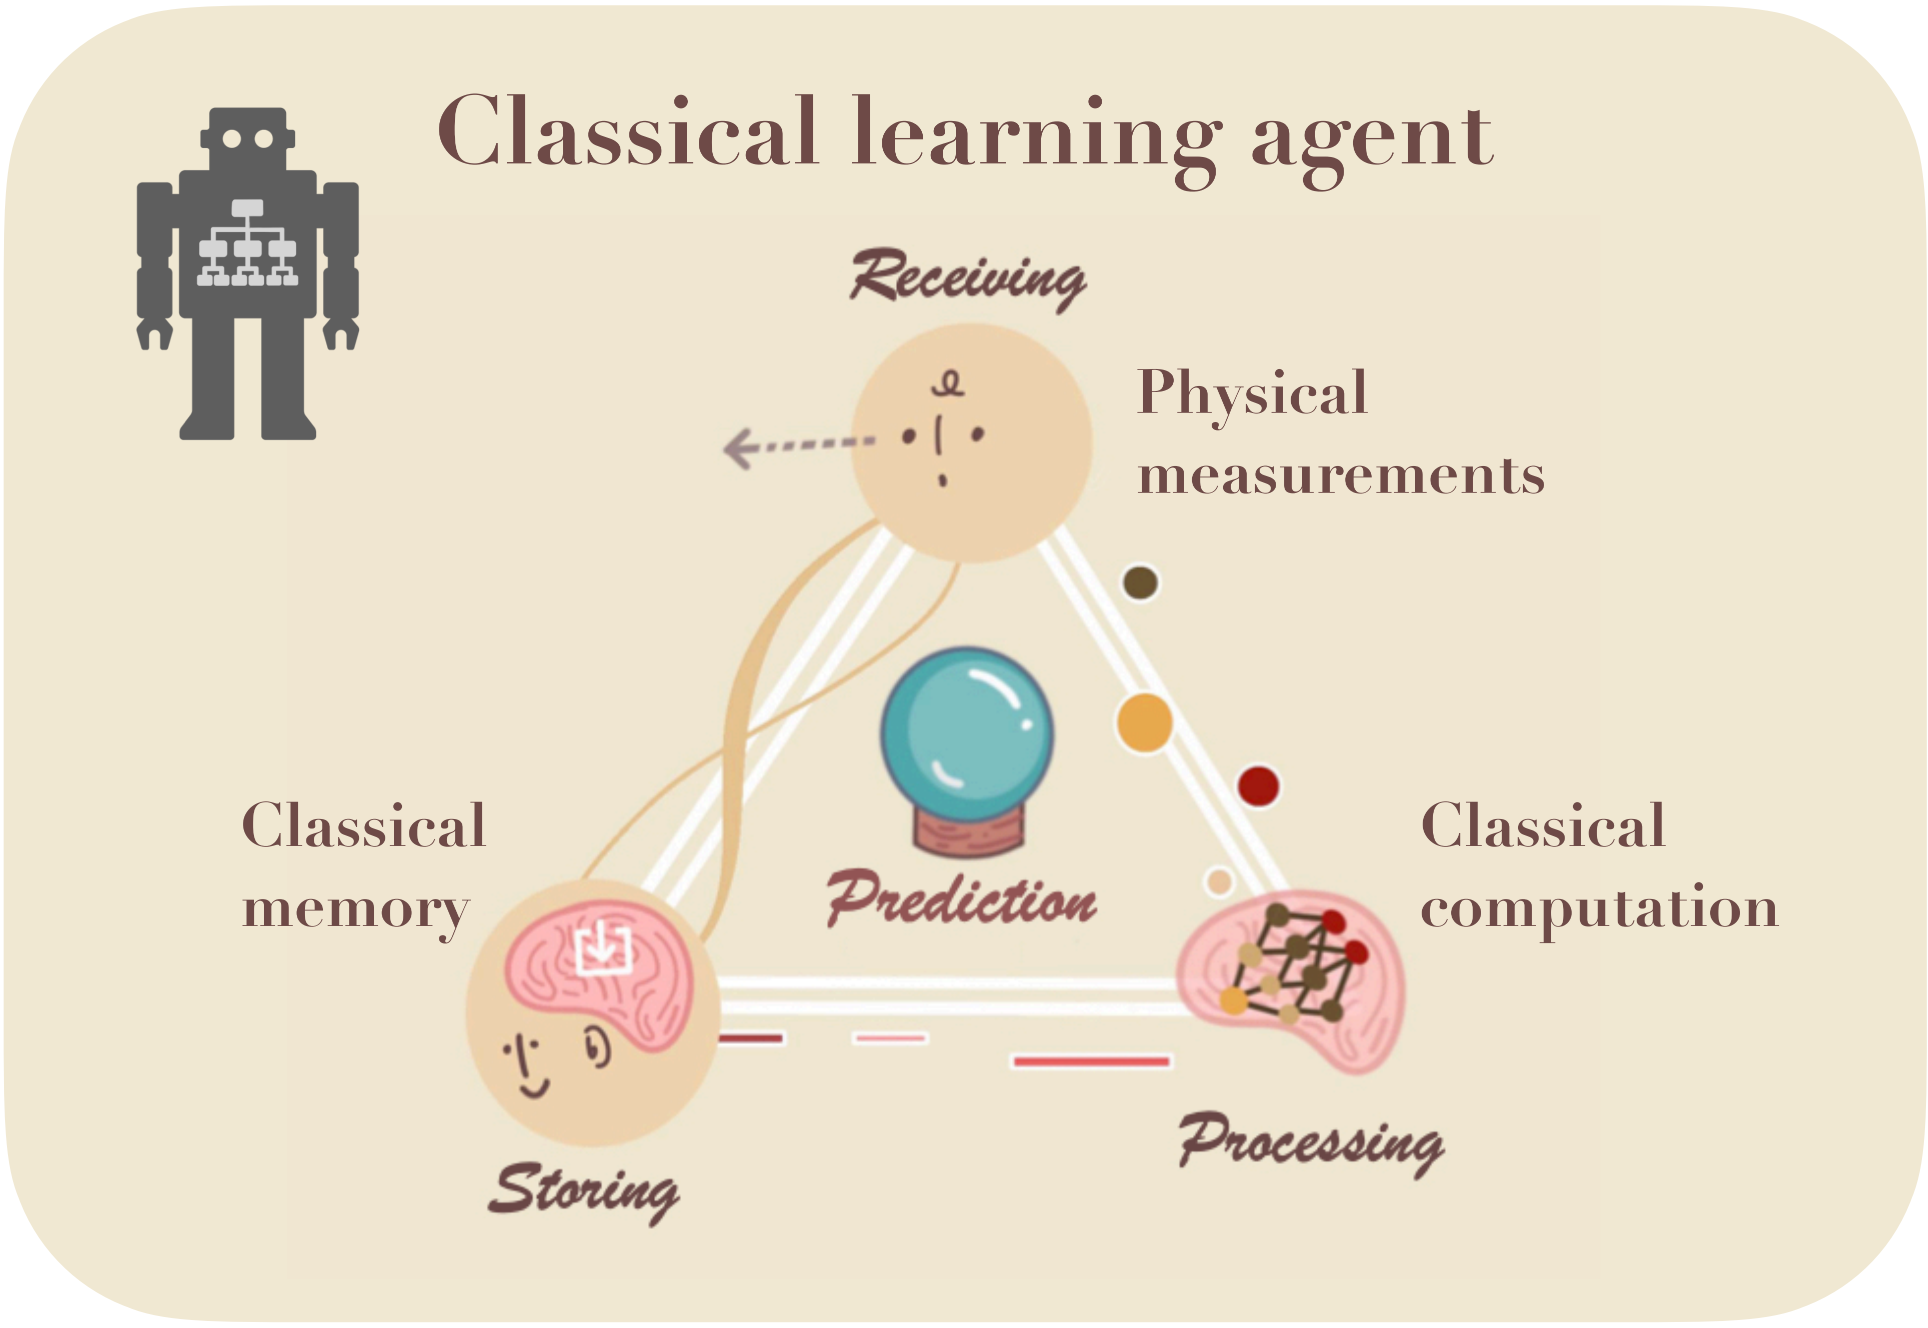

What can classical agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

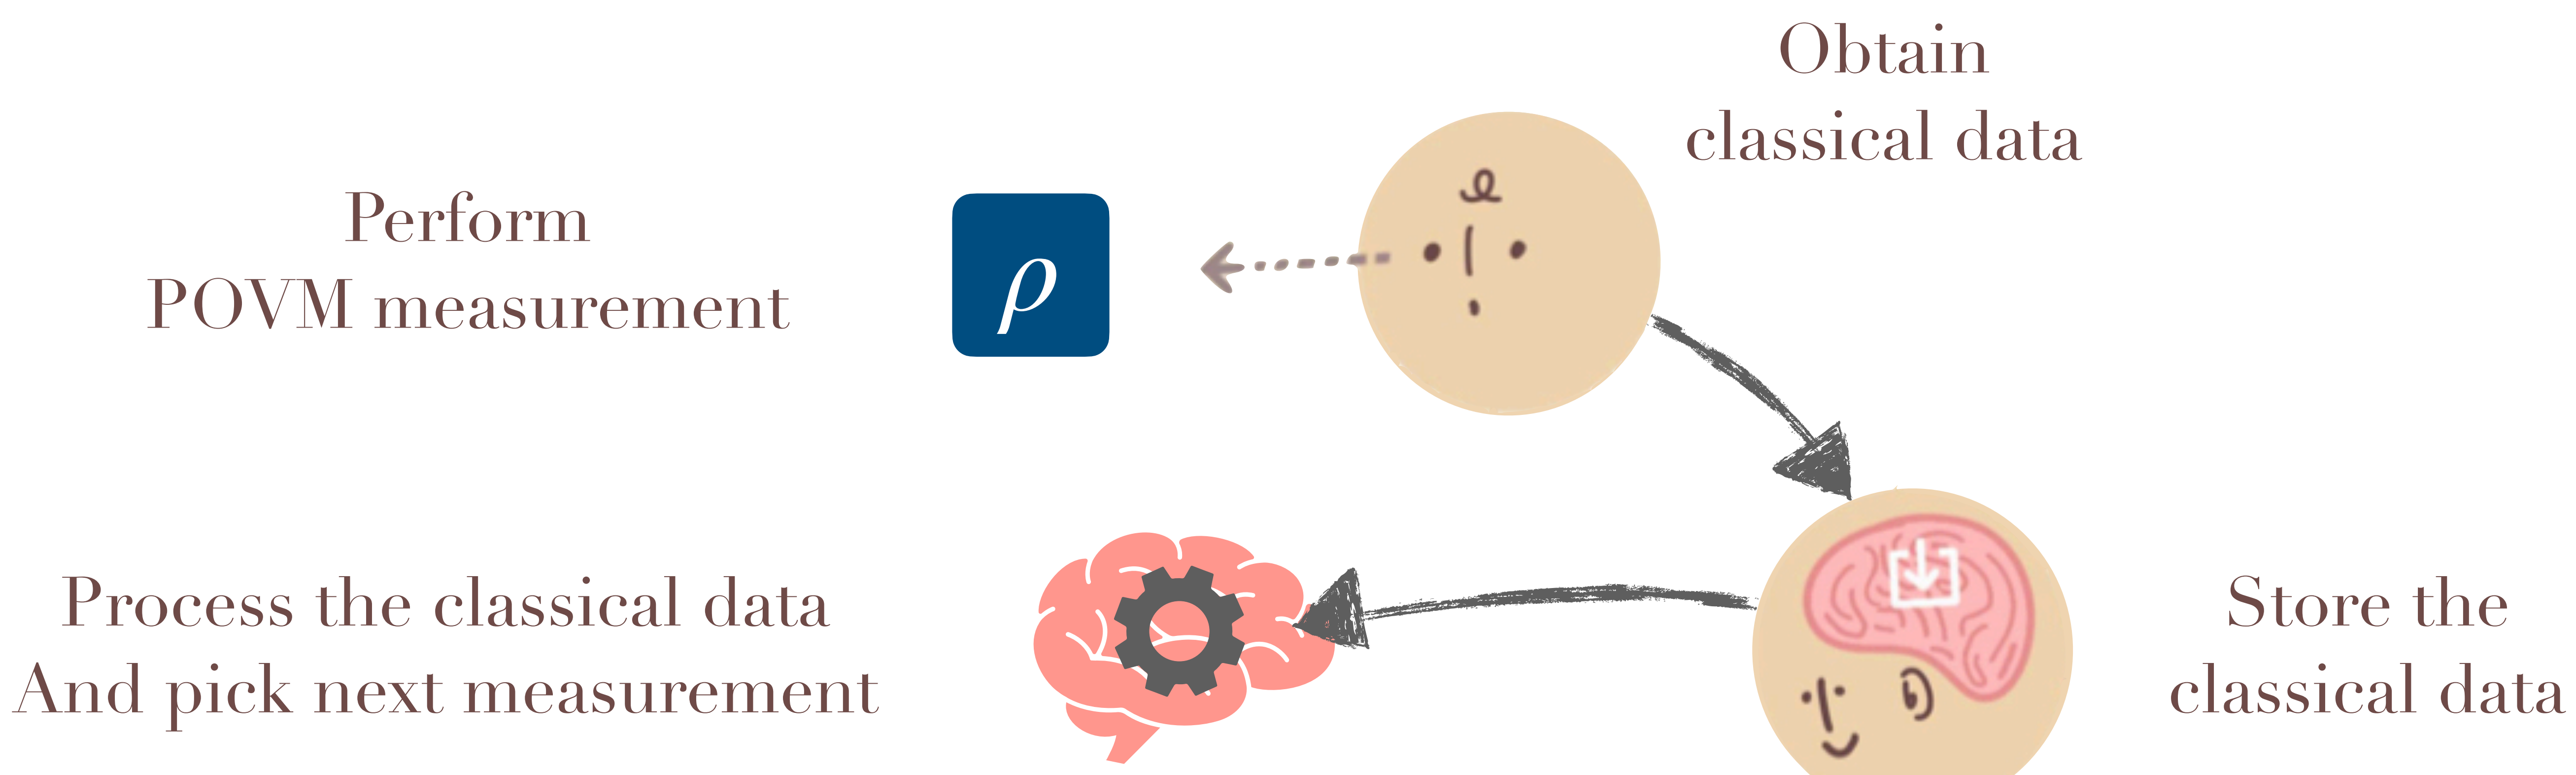

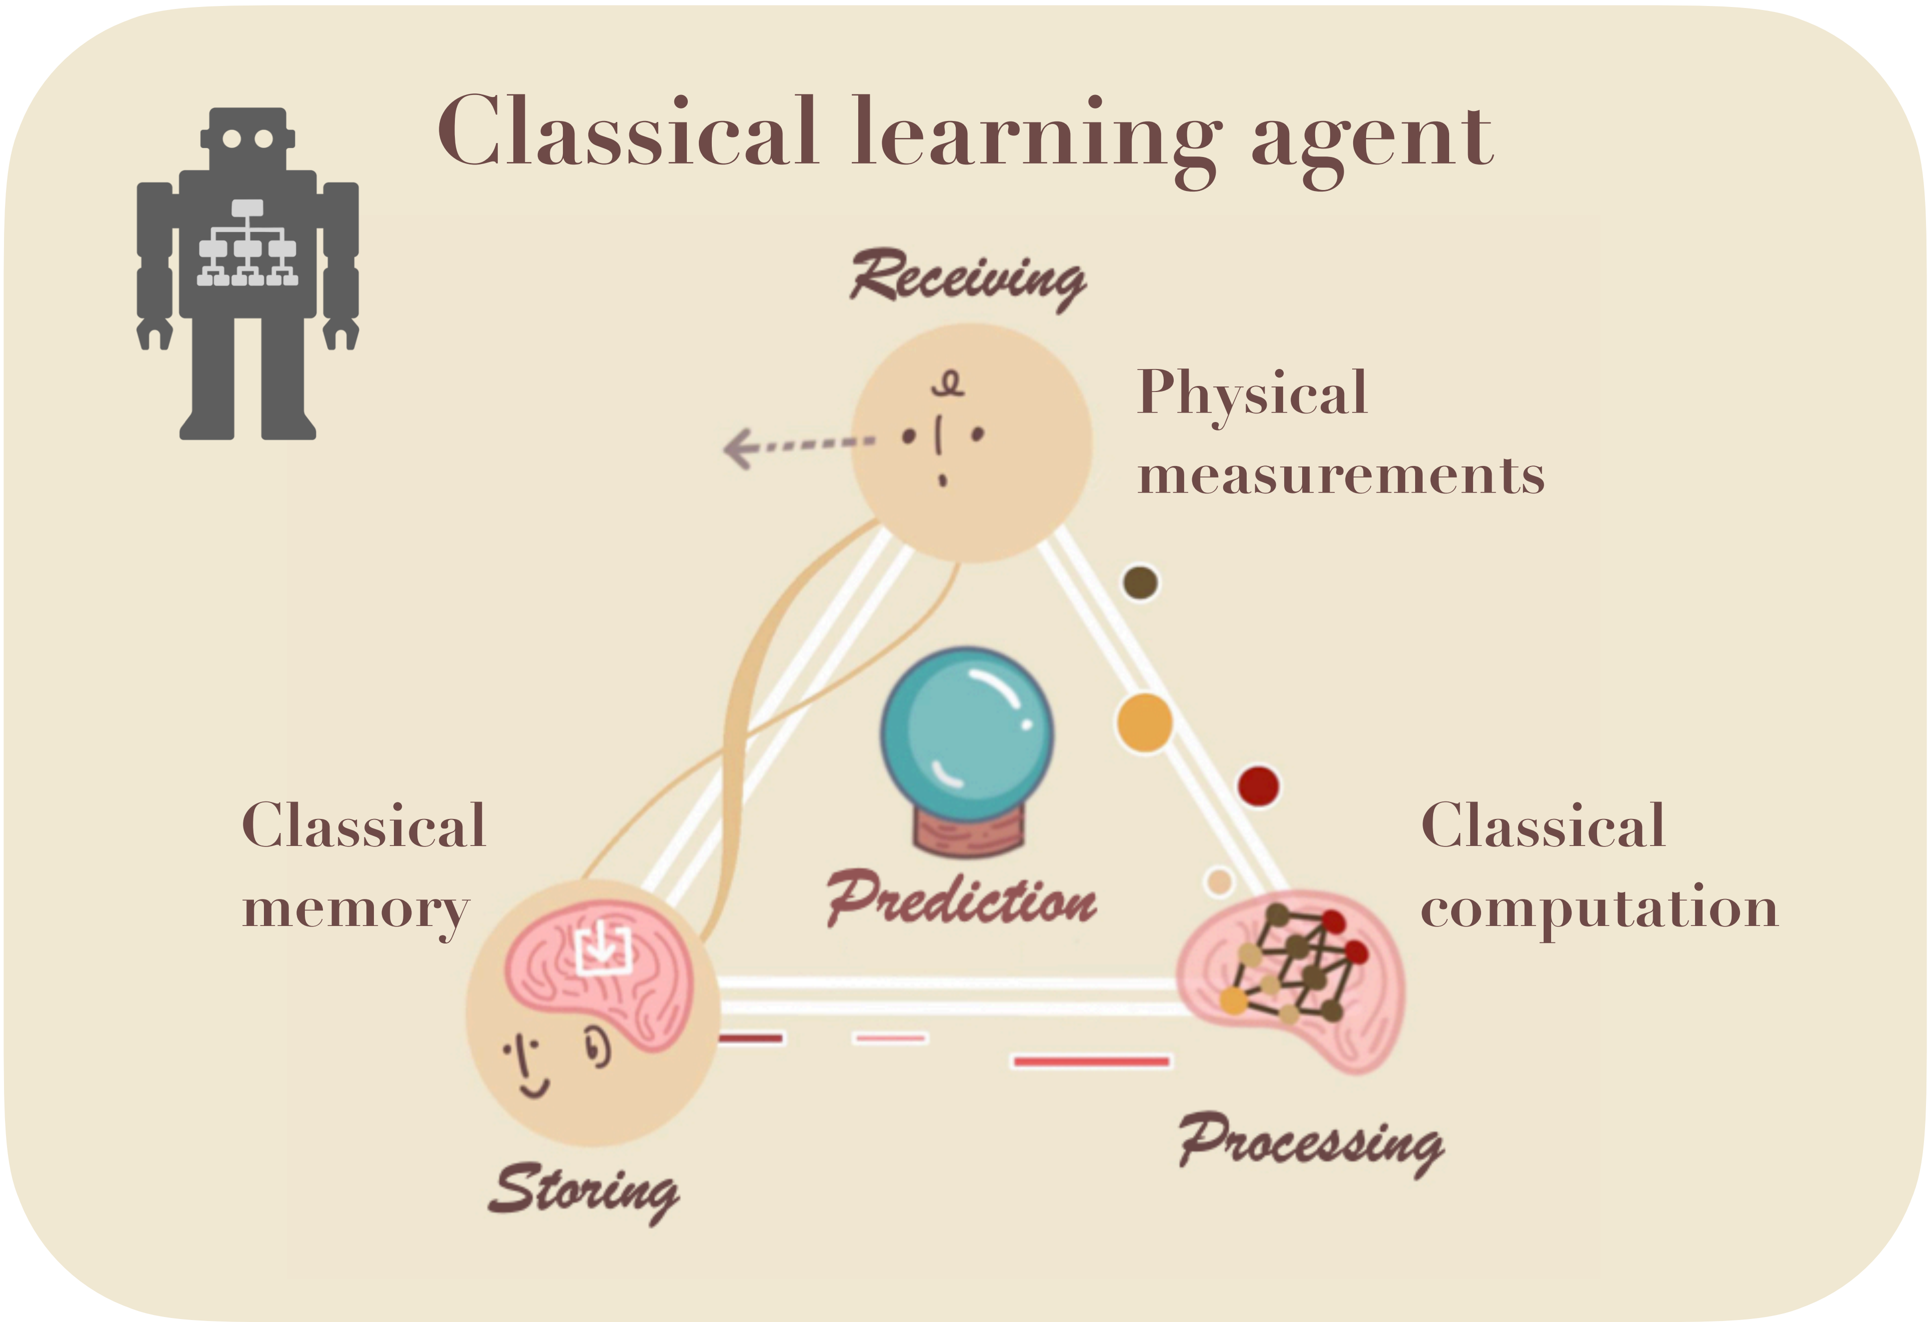

What can classical agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

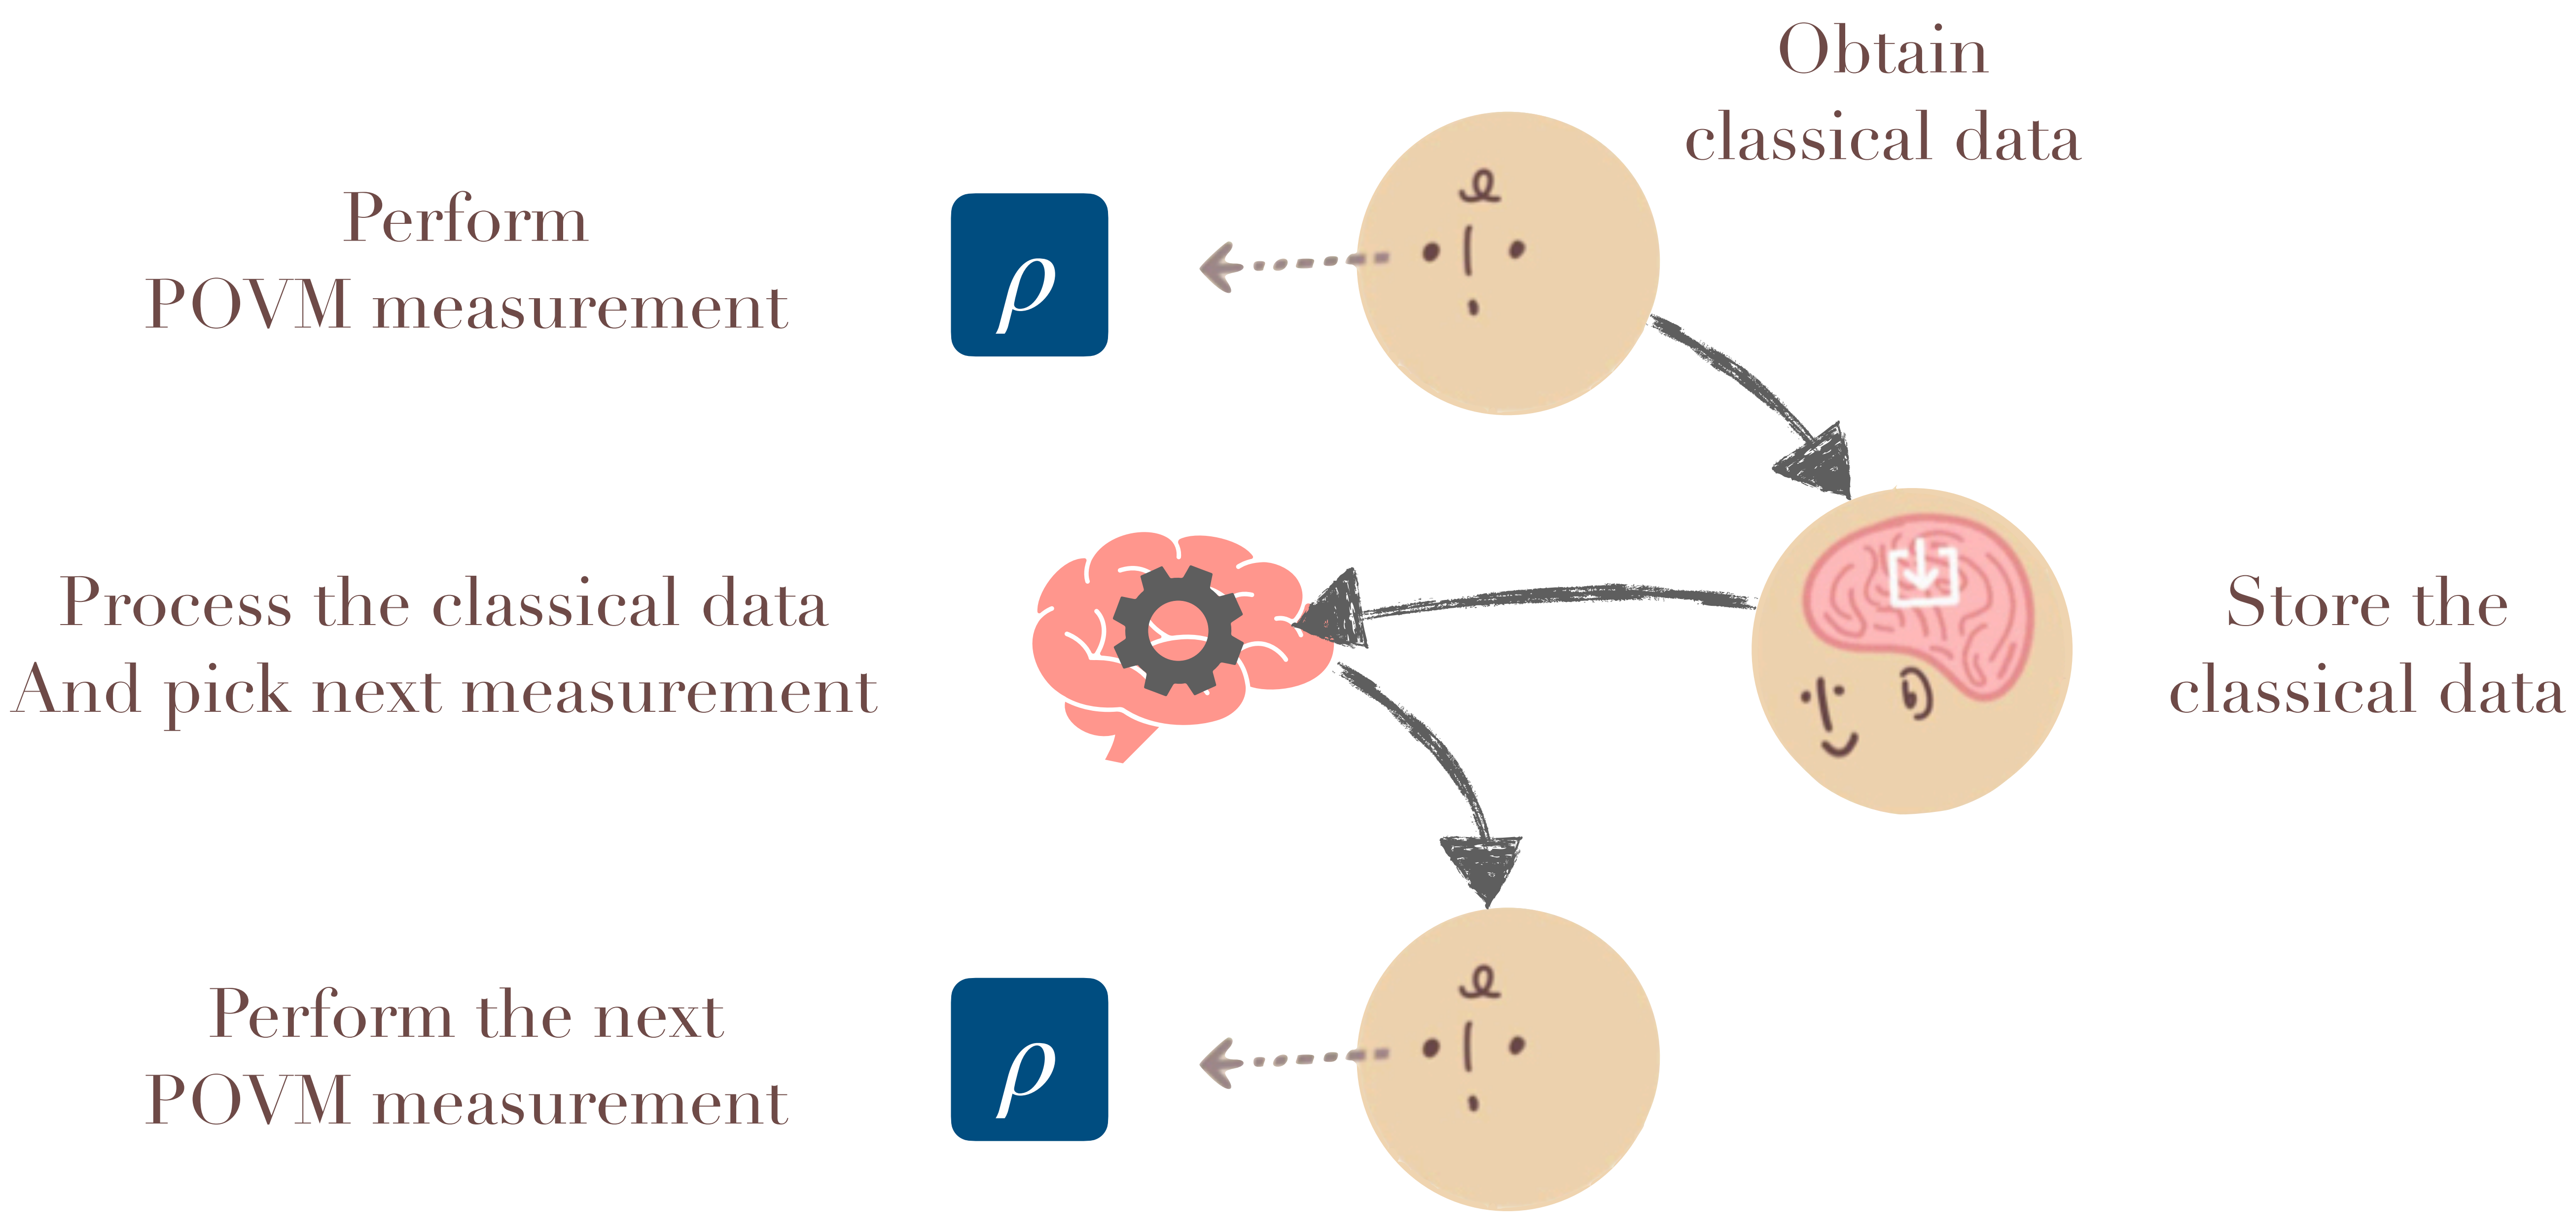

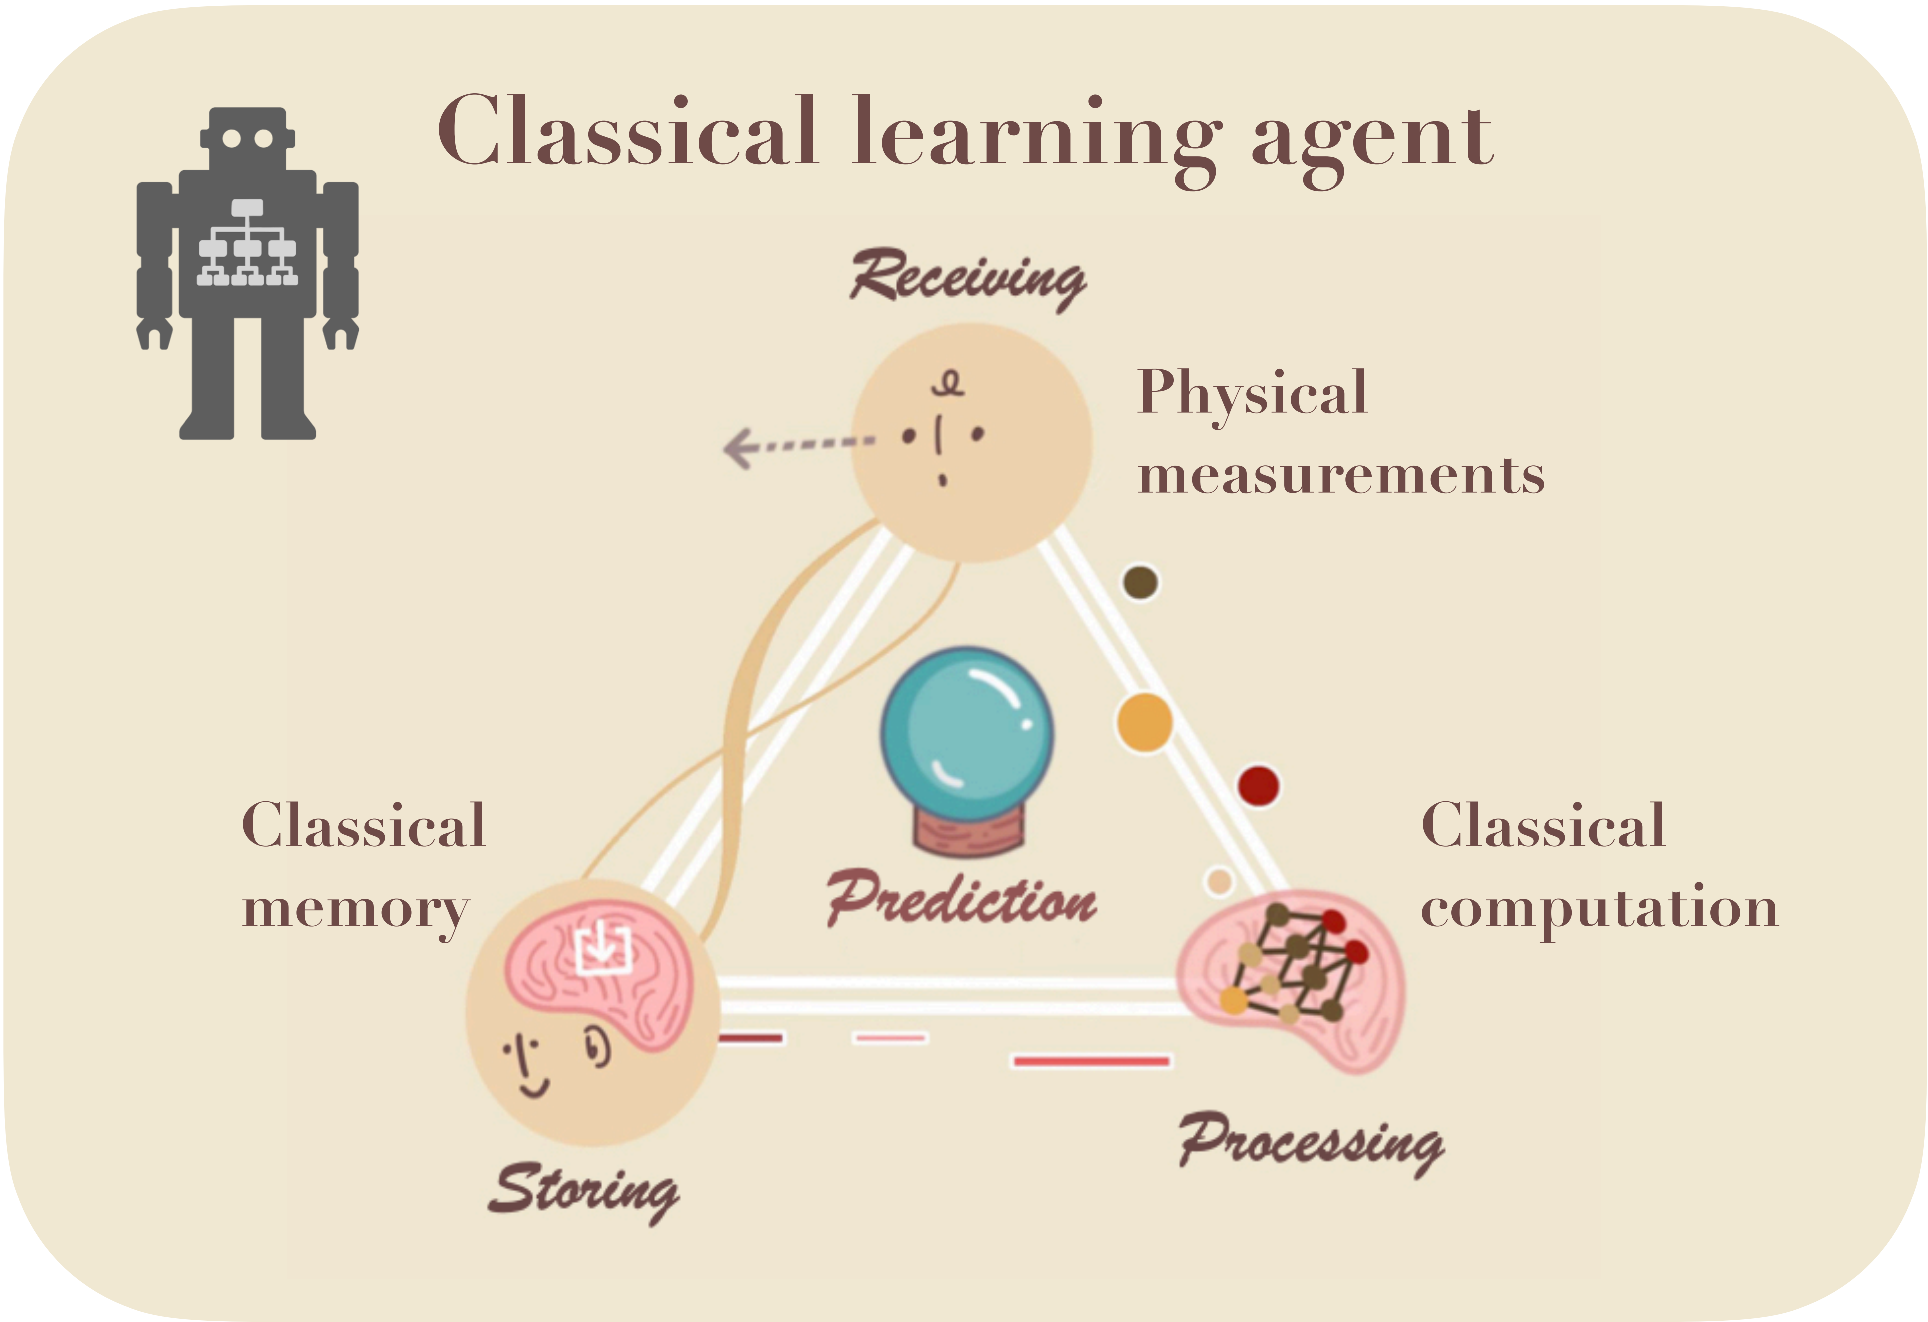

What can classical agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

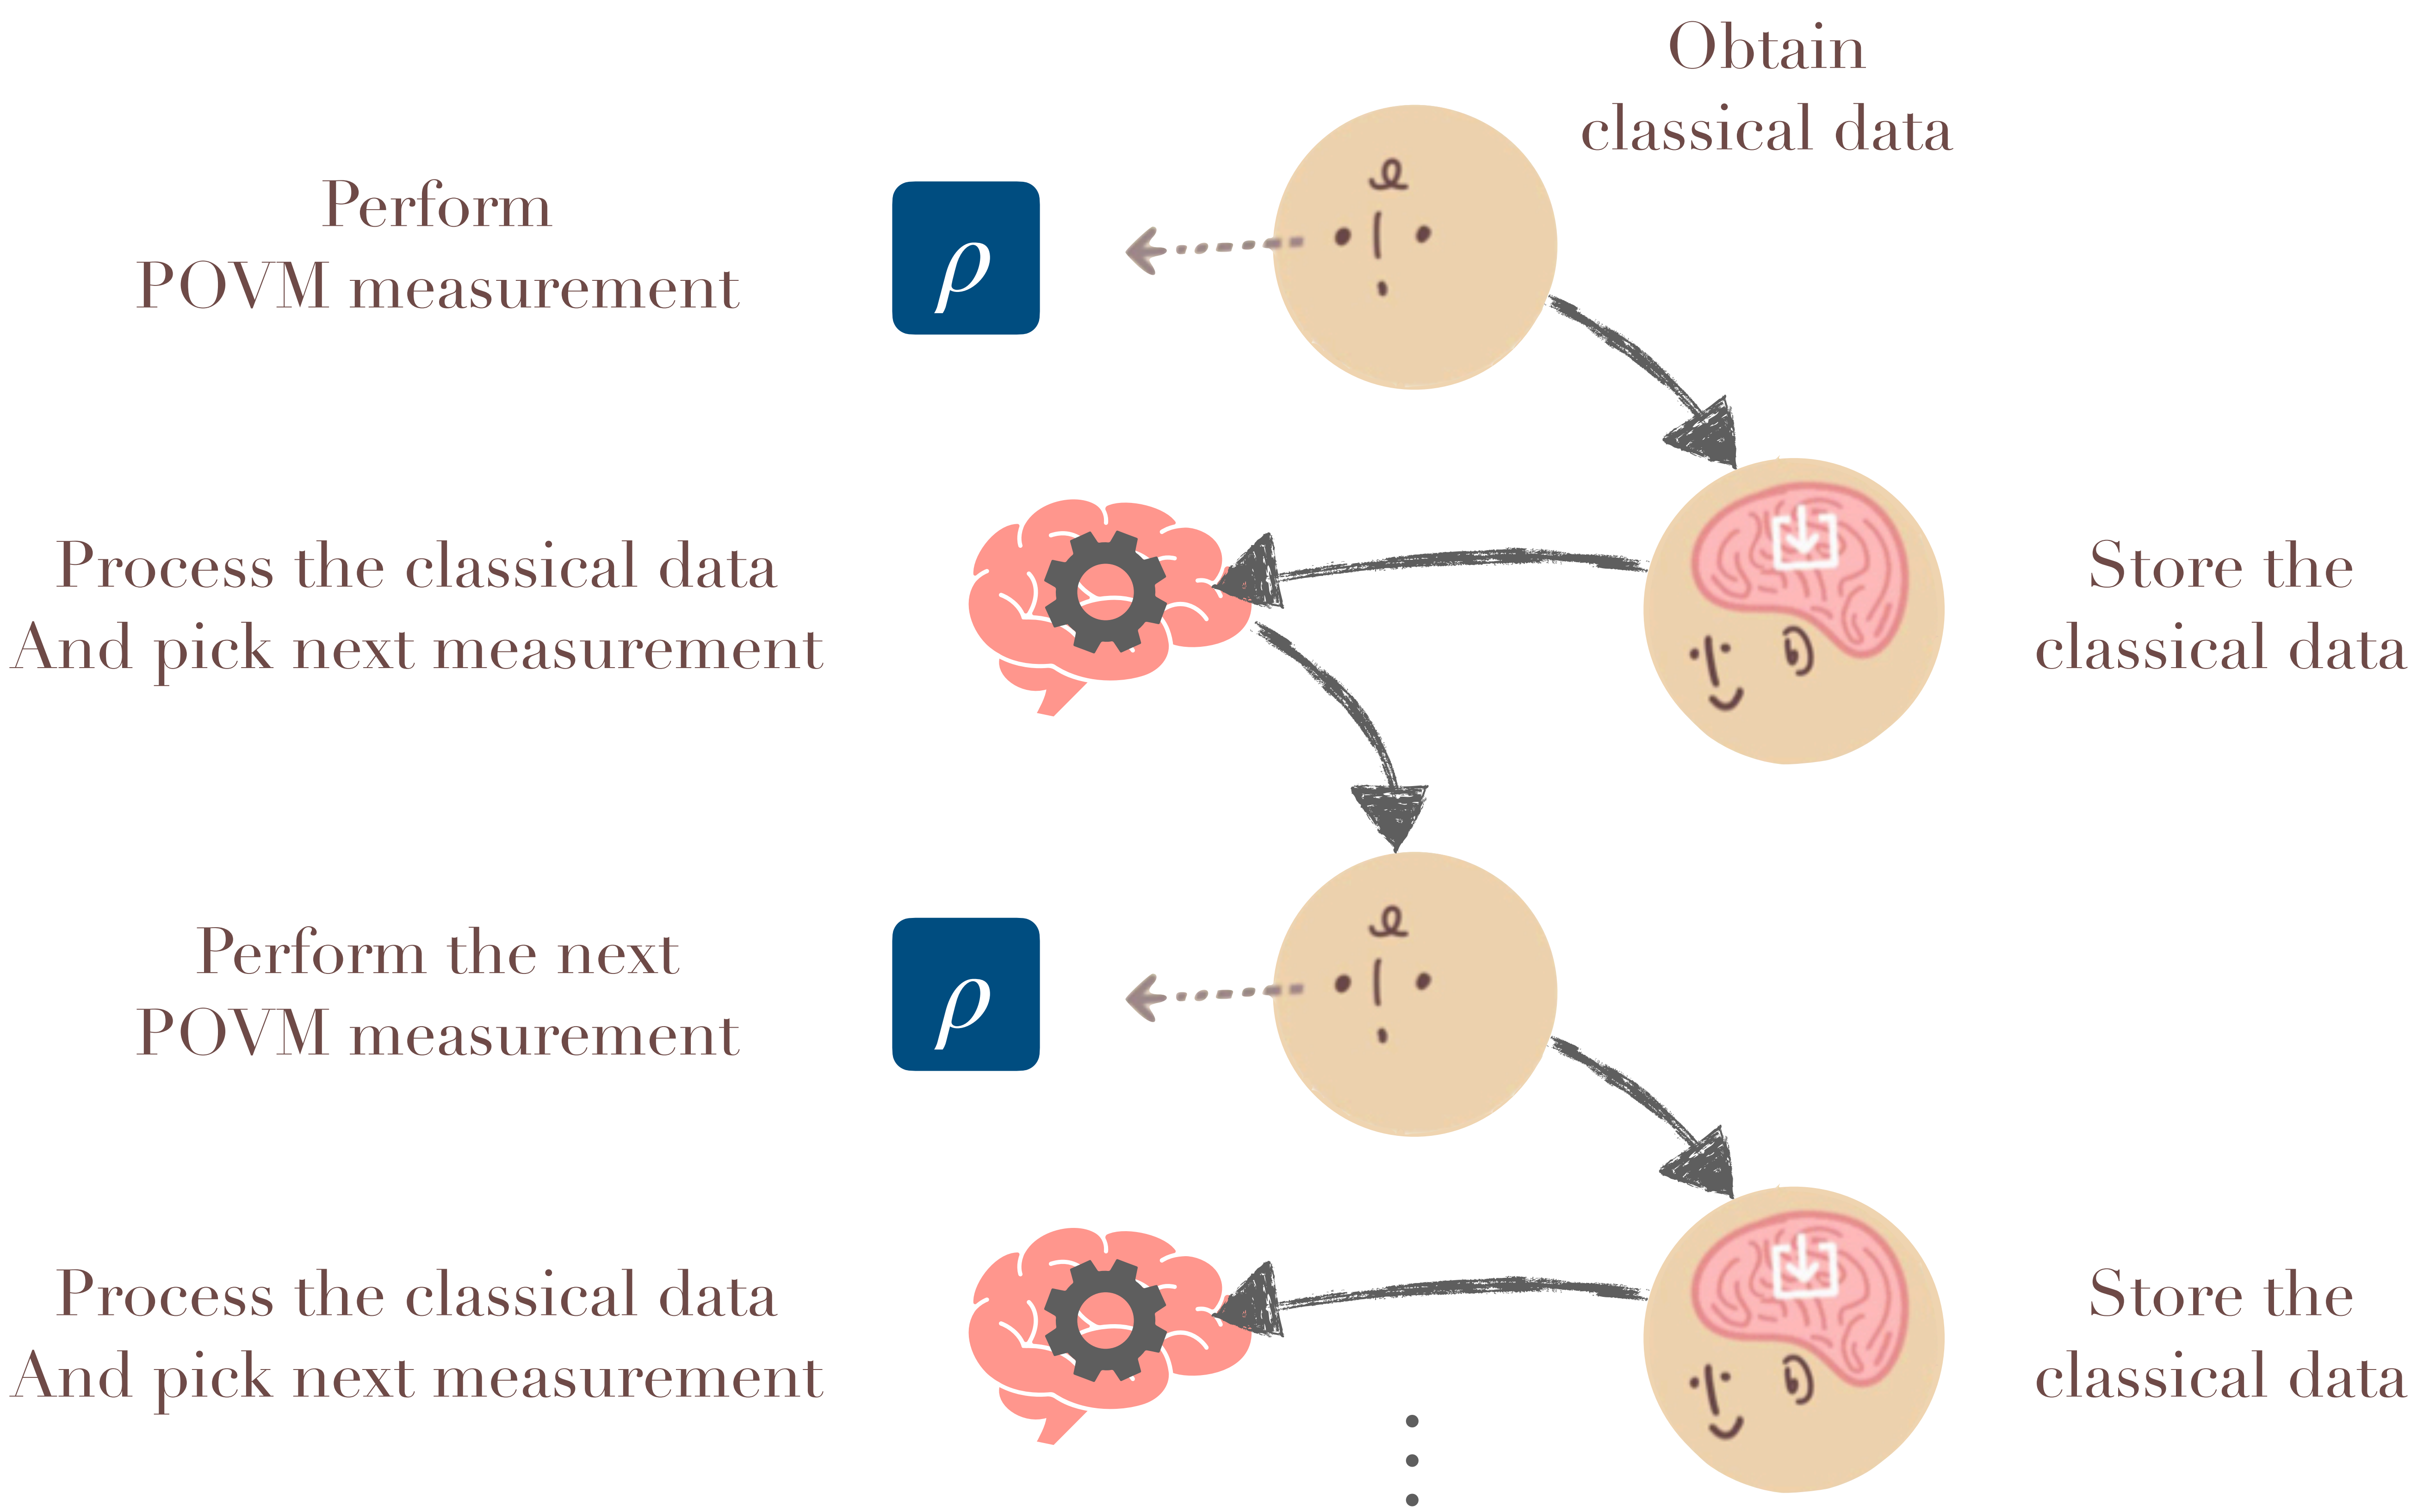

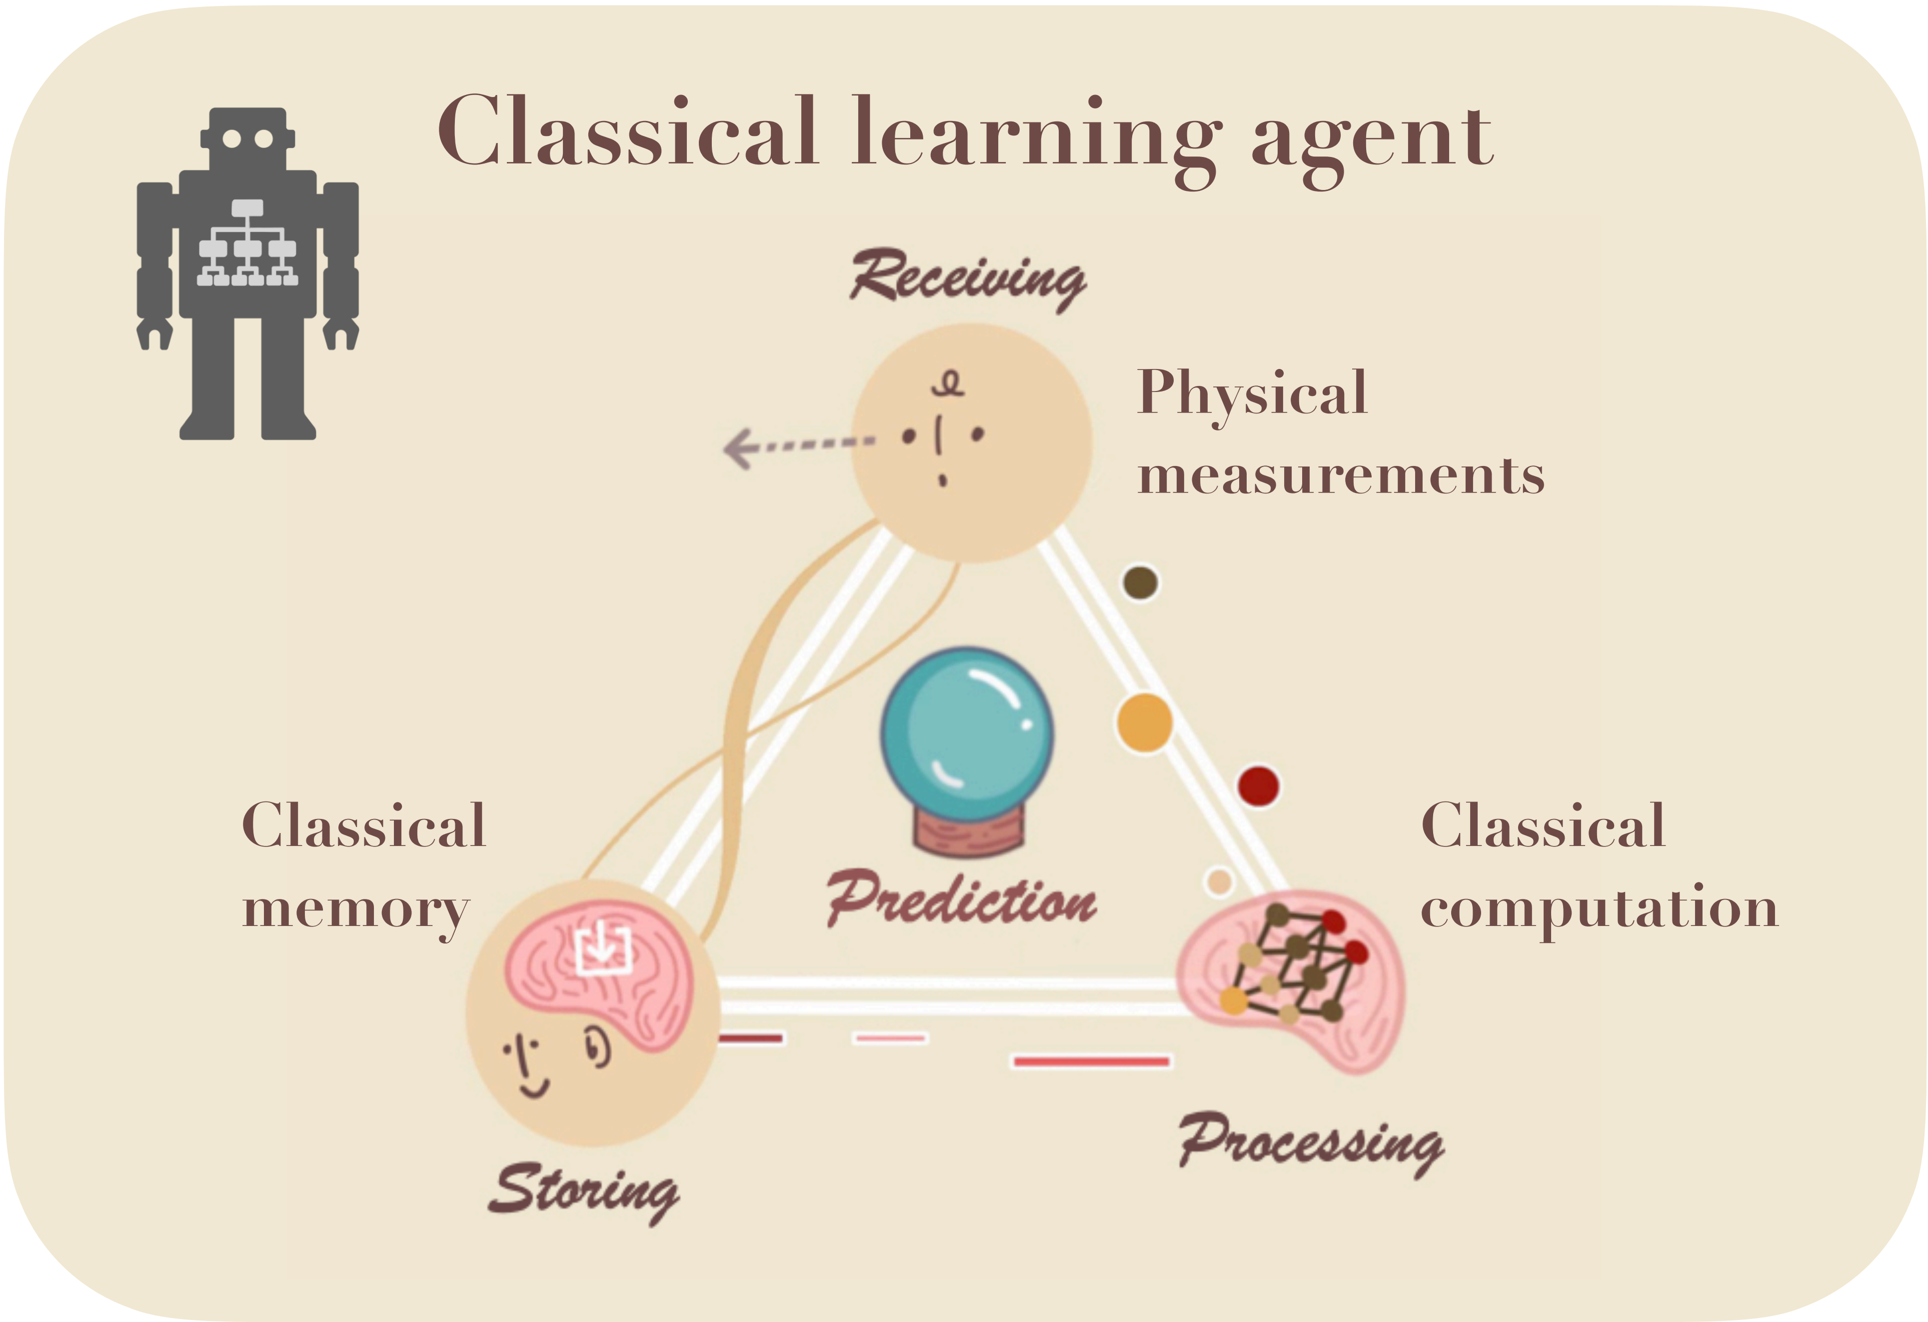

What can classical agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

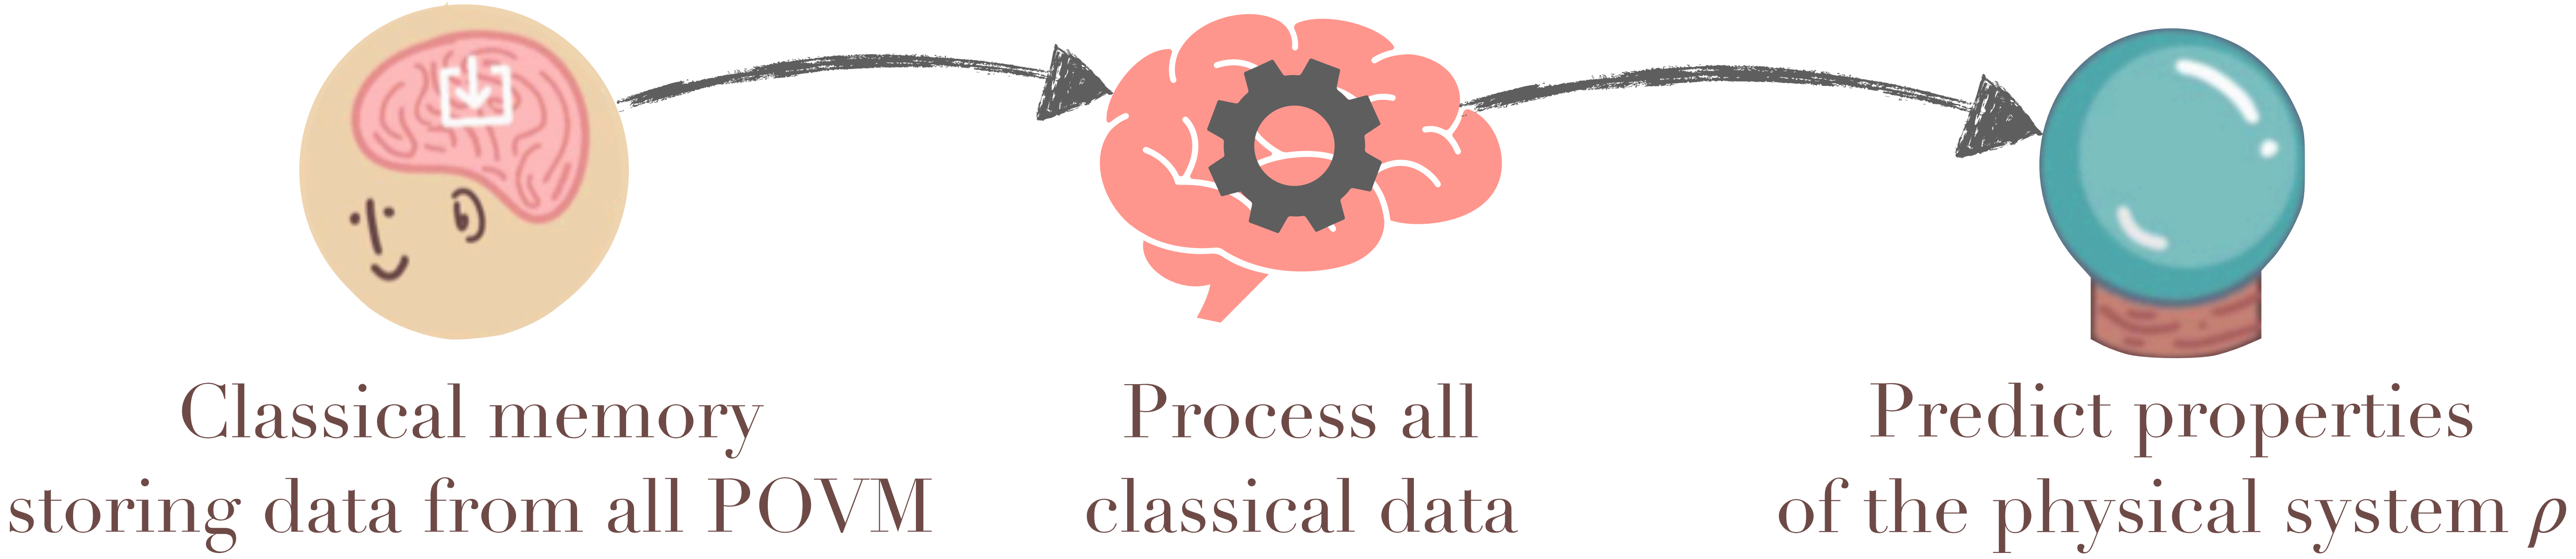

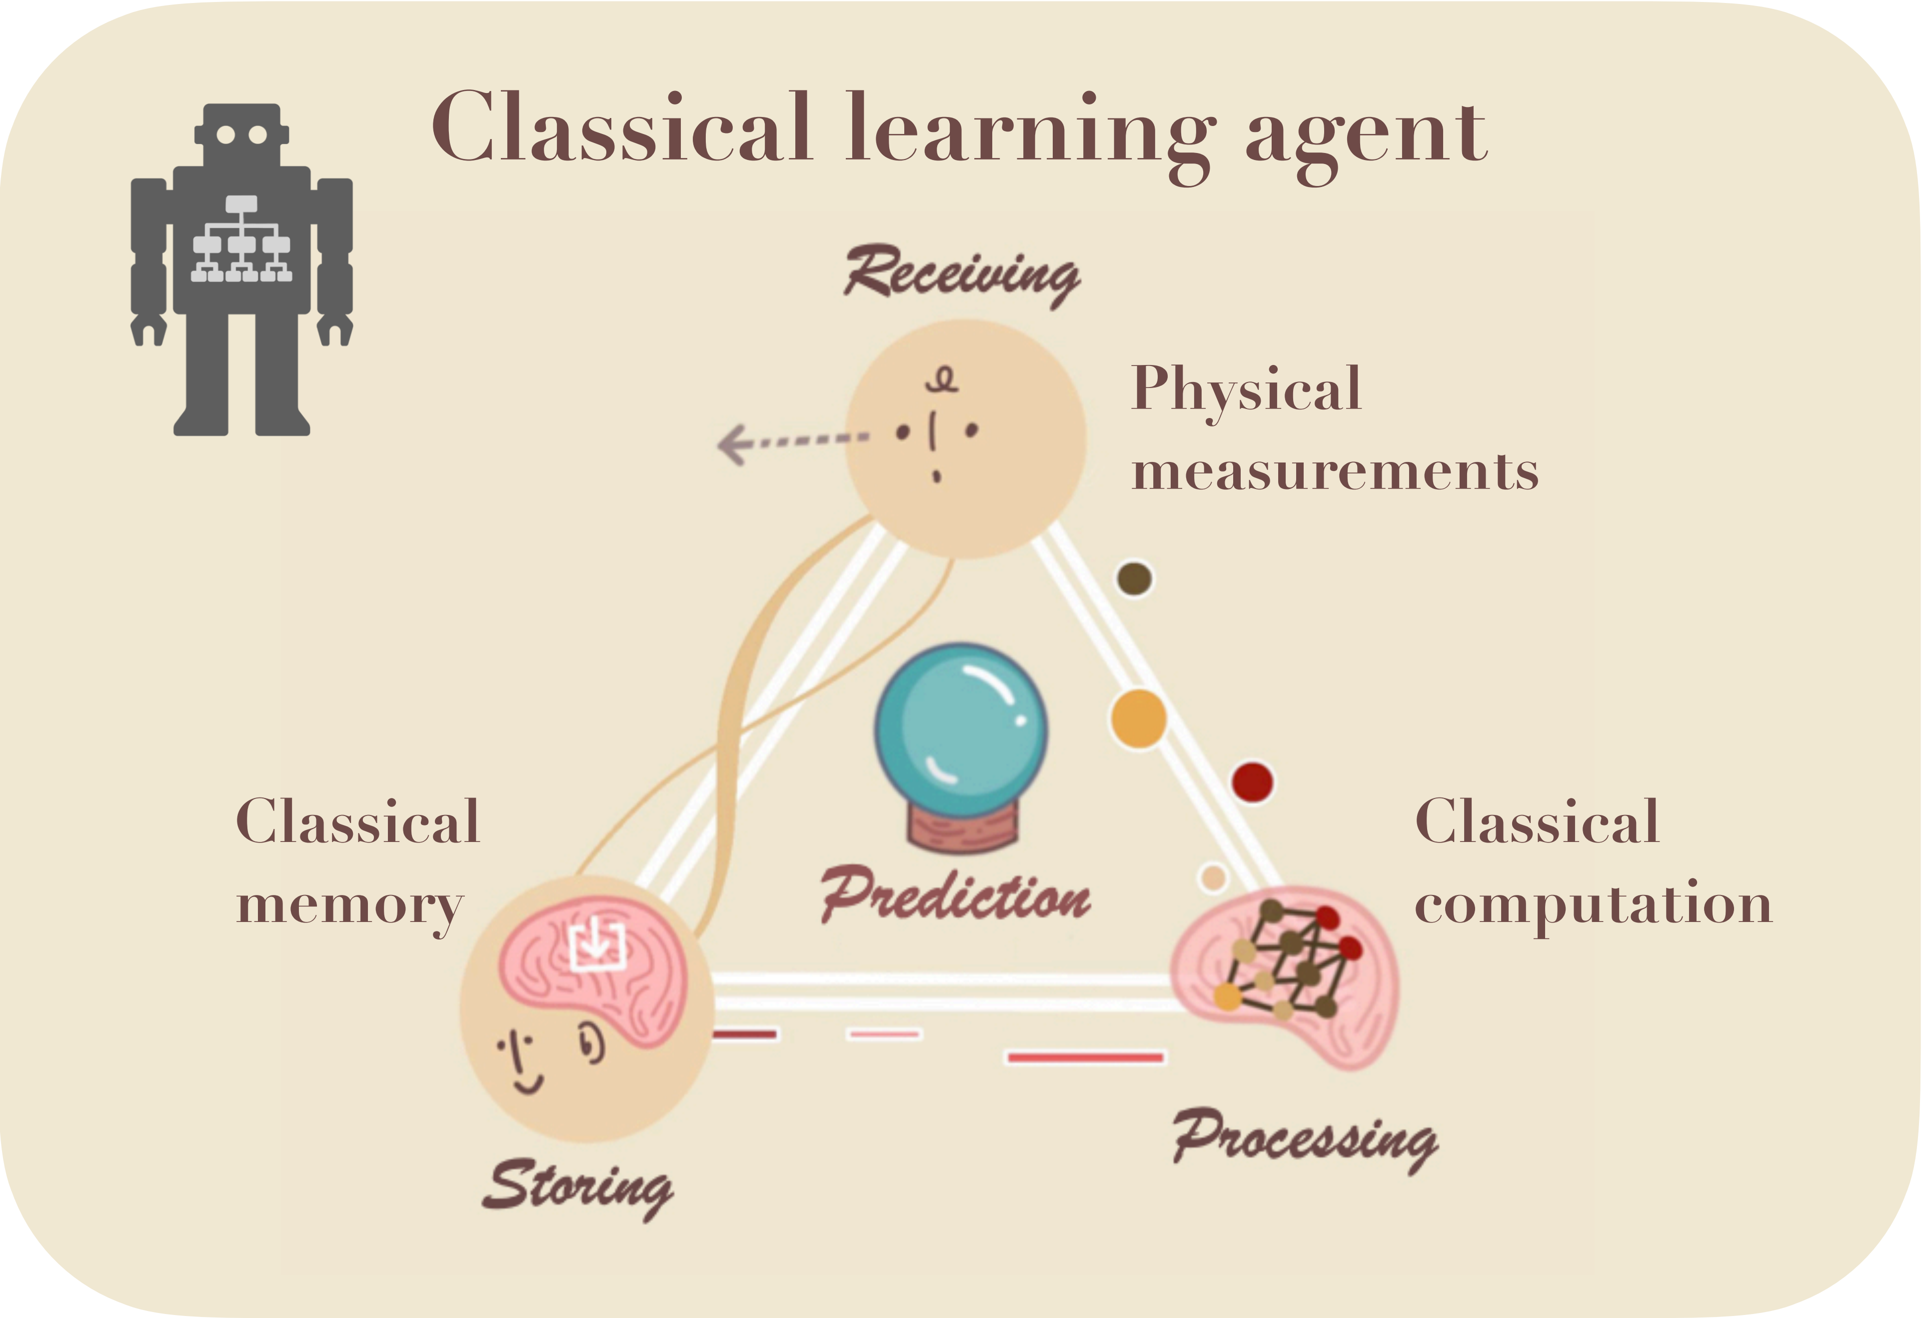

What can quantum agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

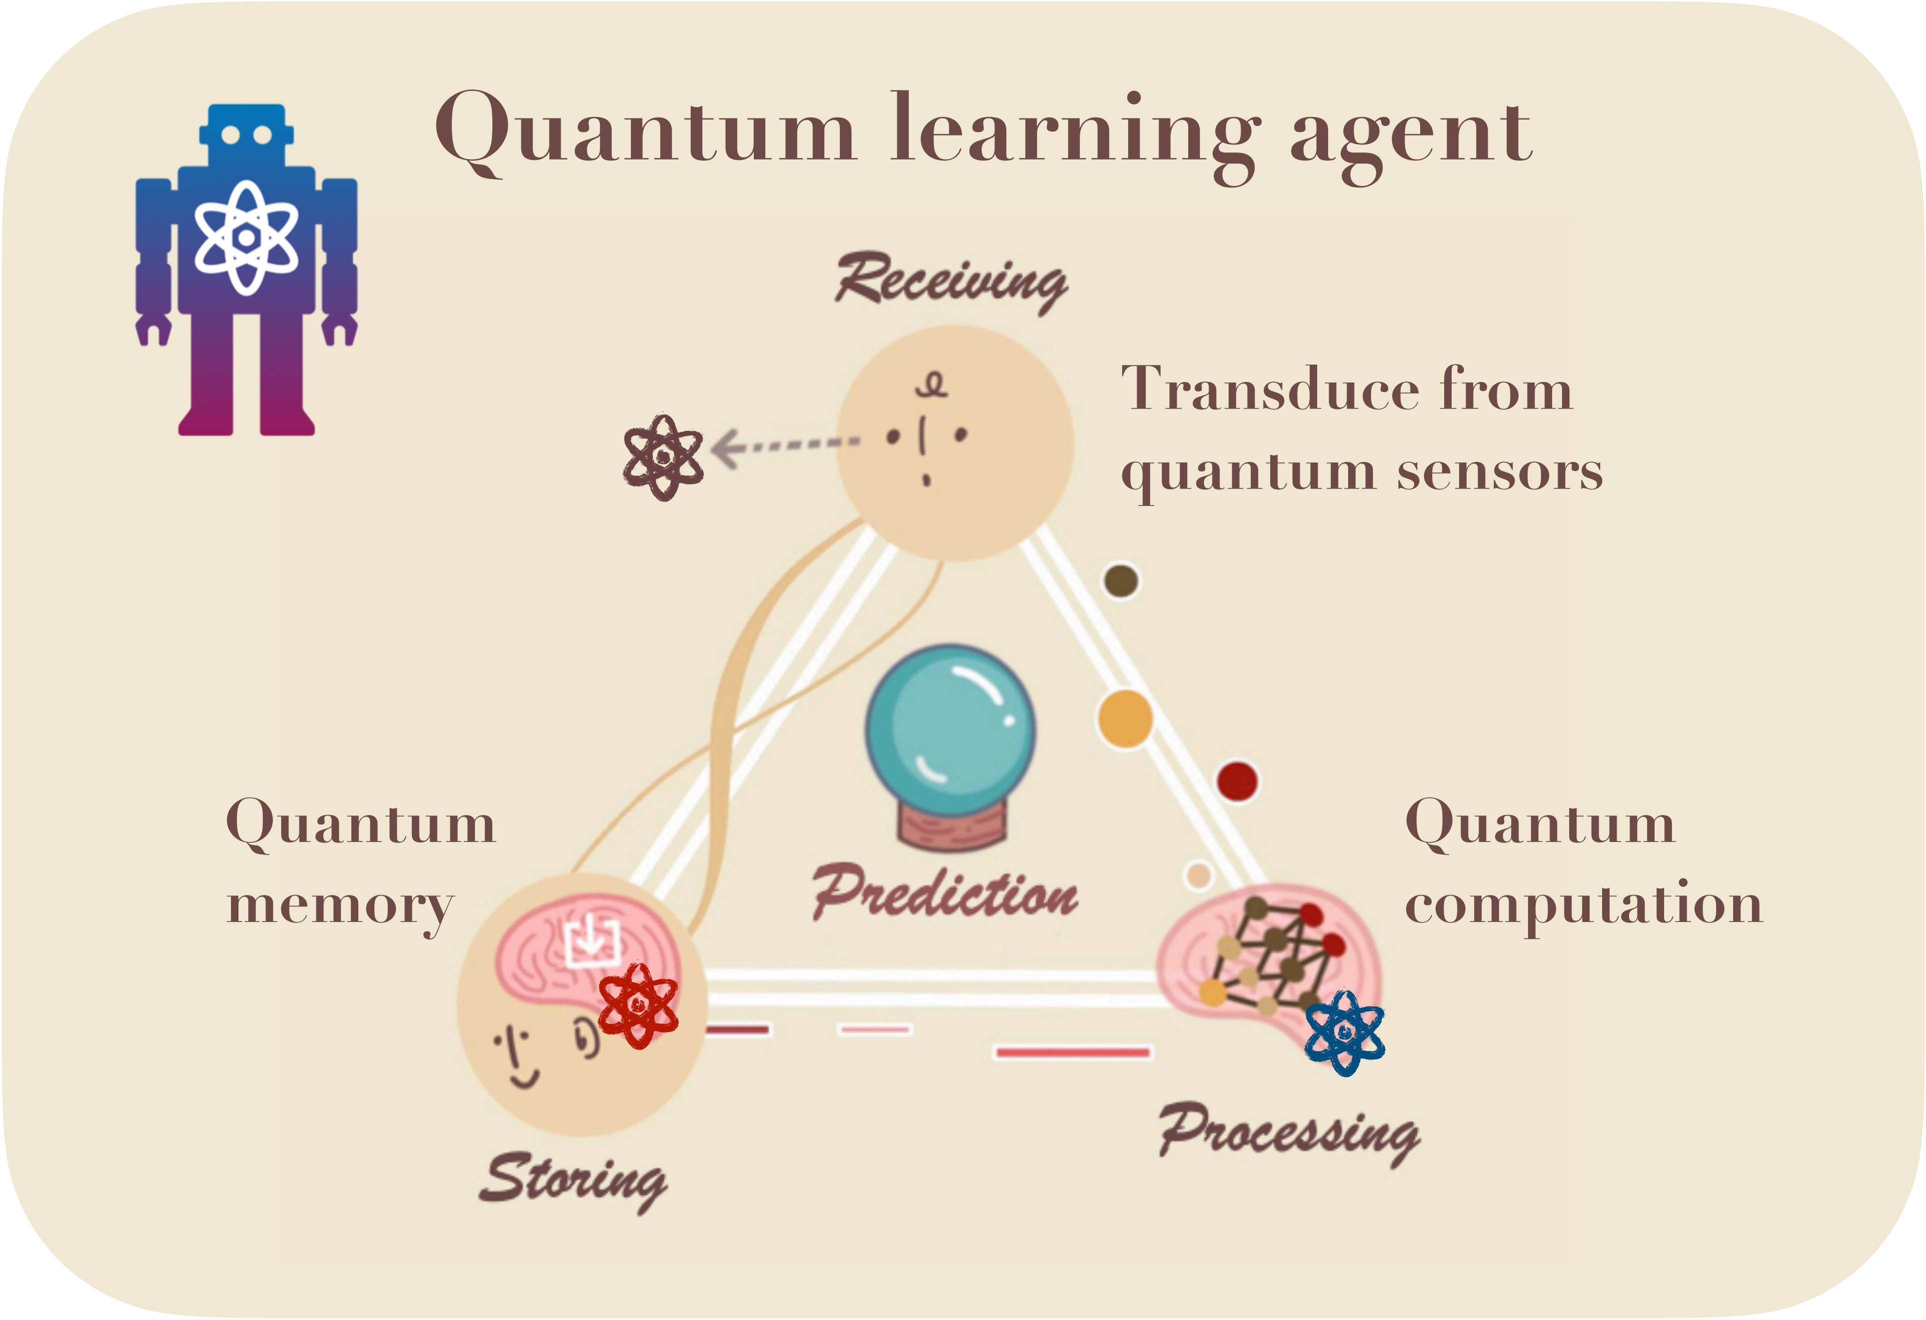

What can quantum agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

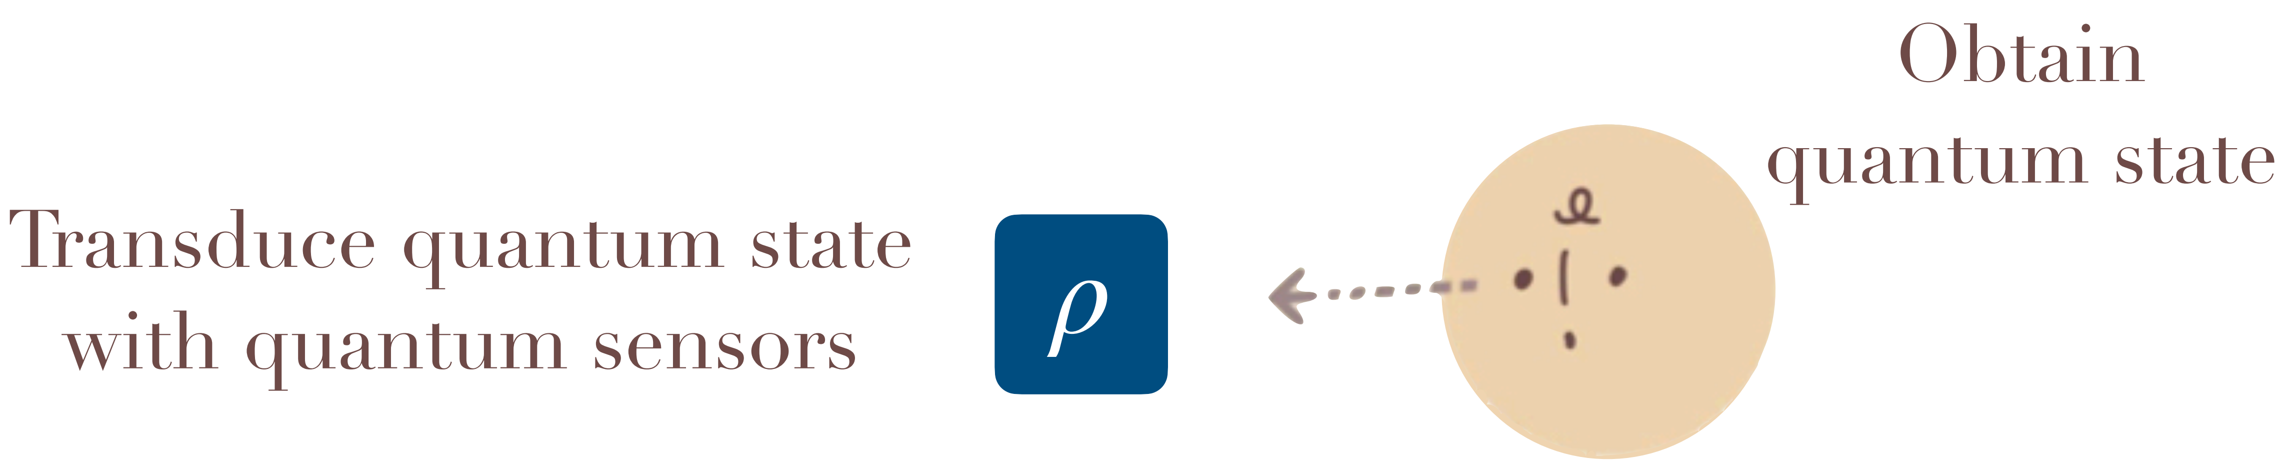

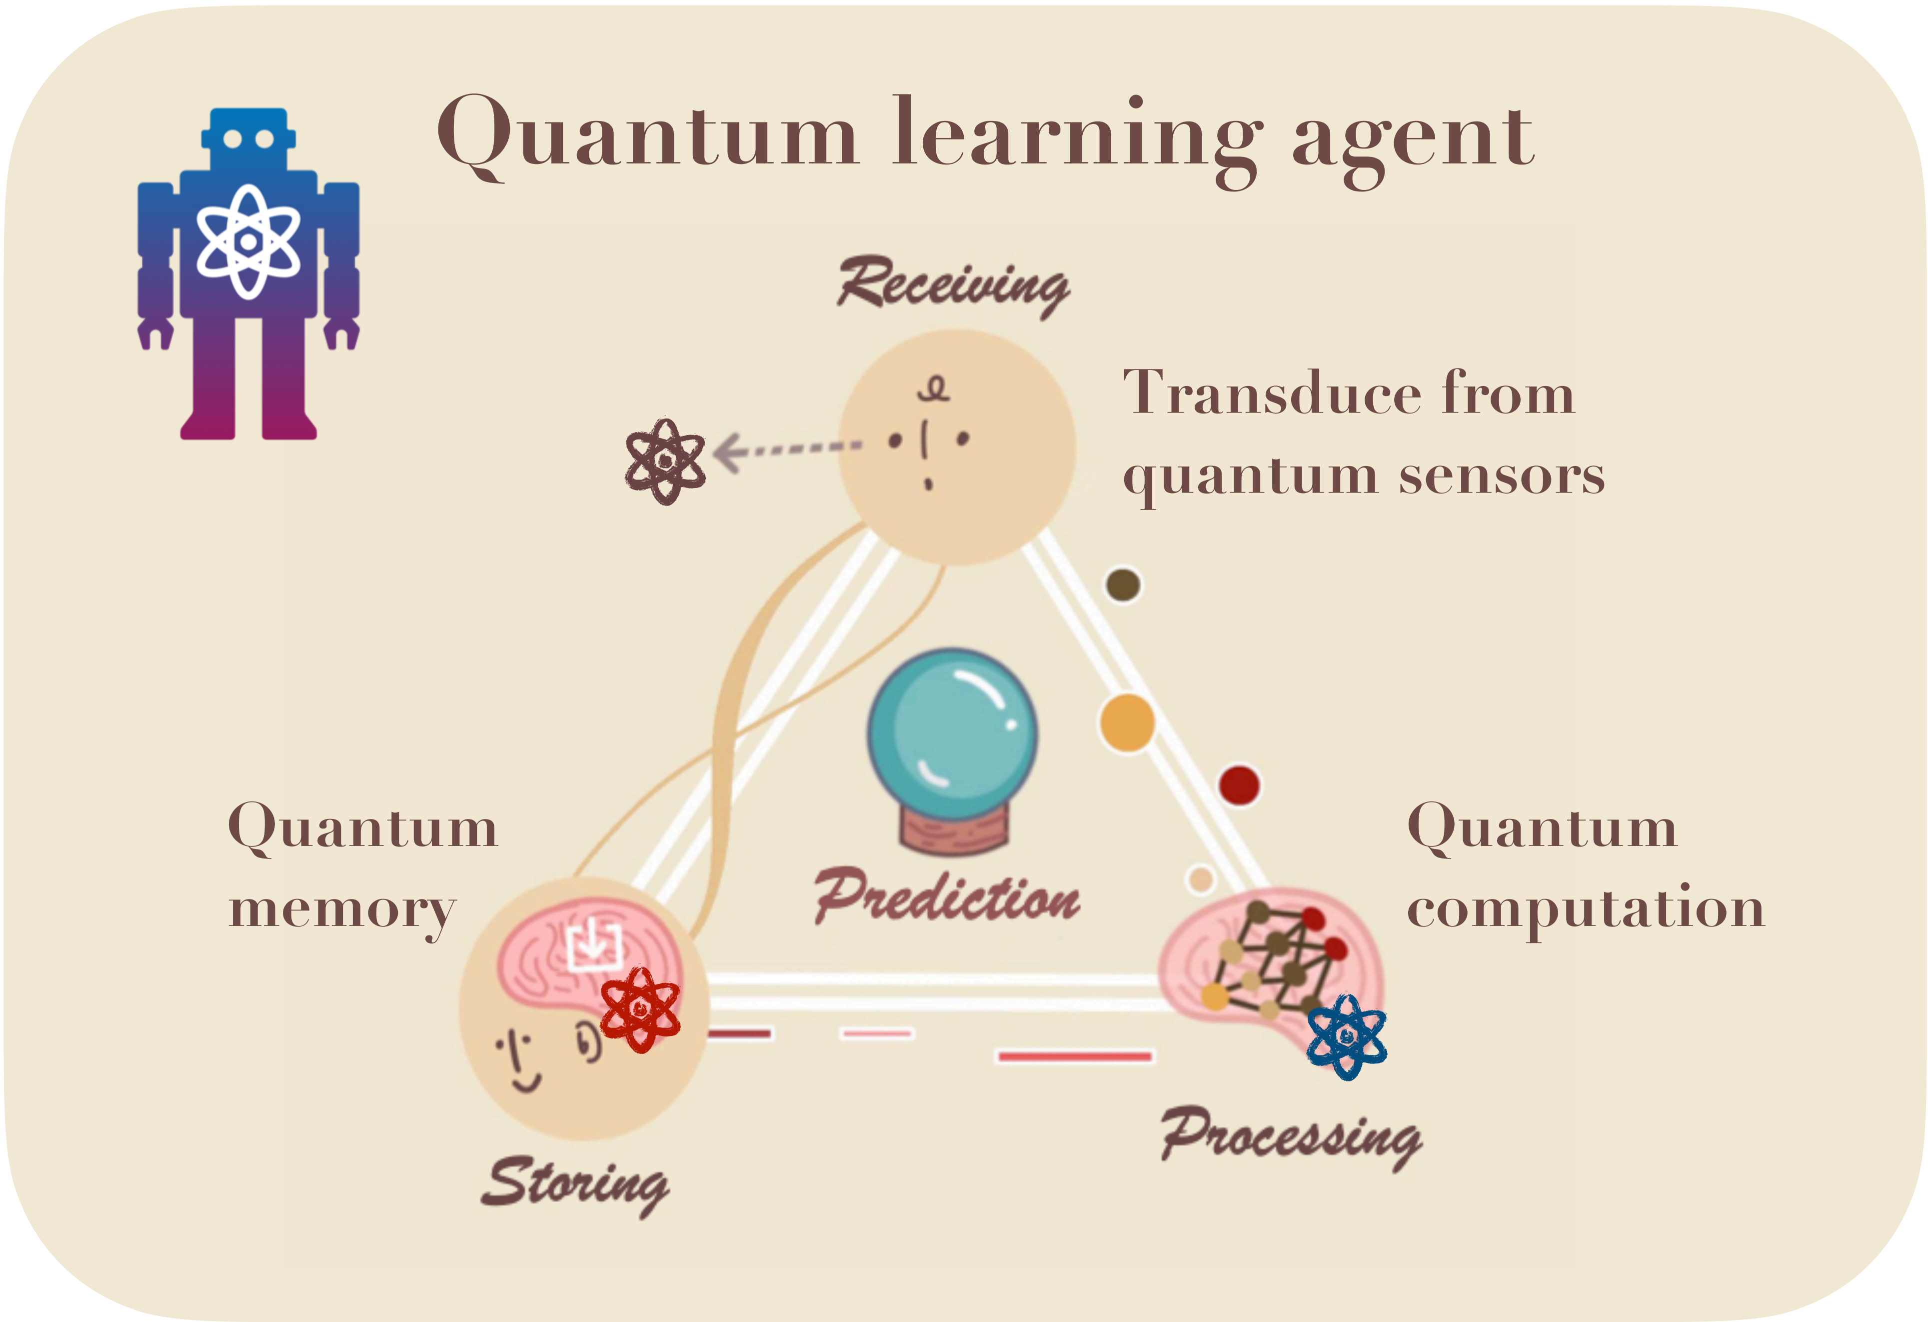

What can quantum agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

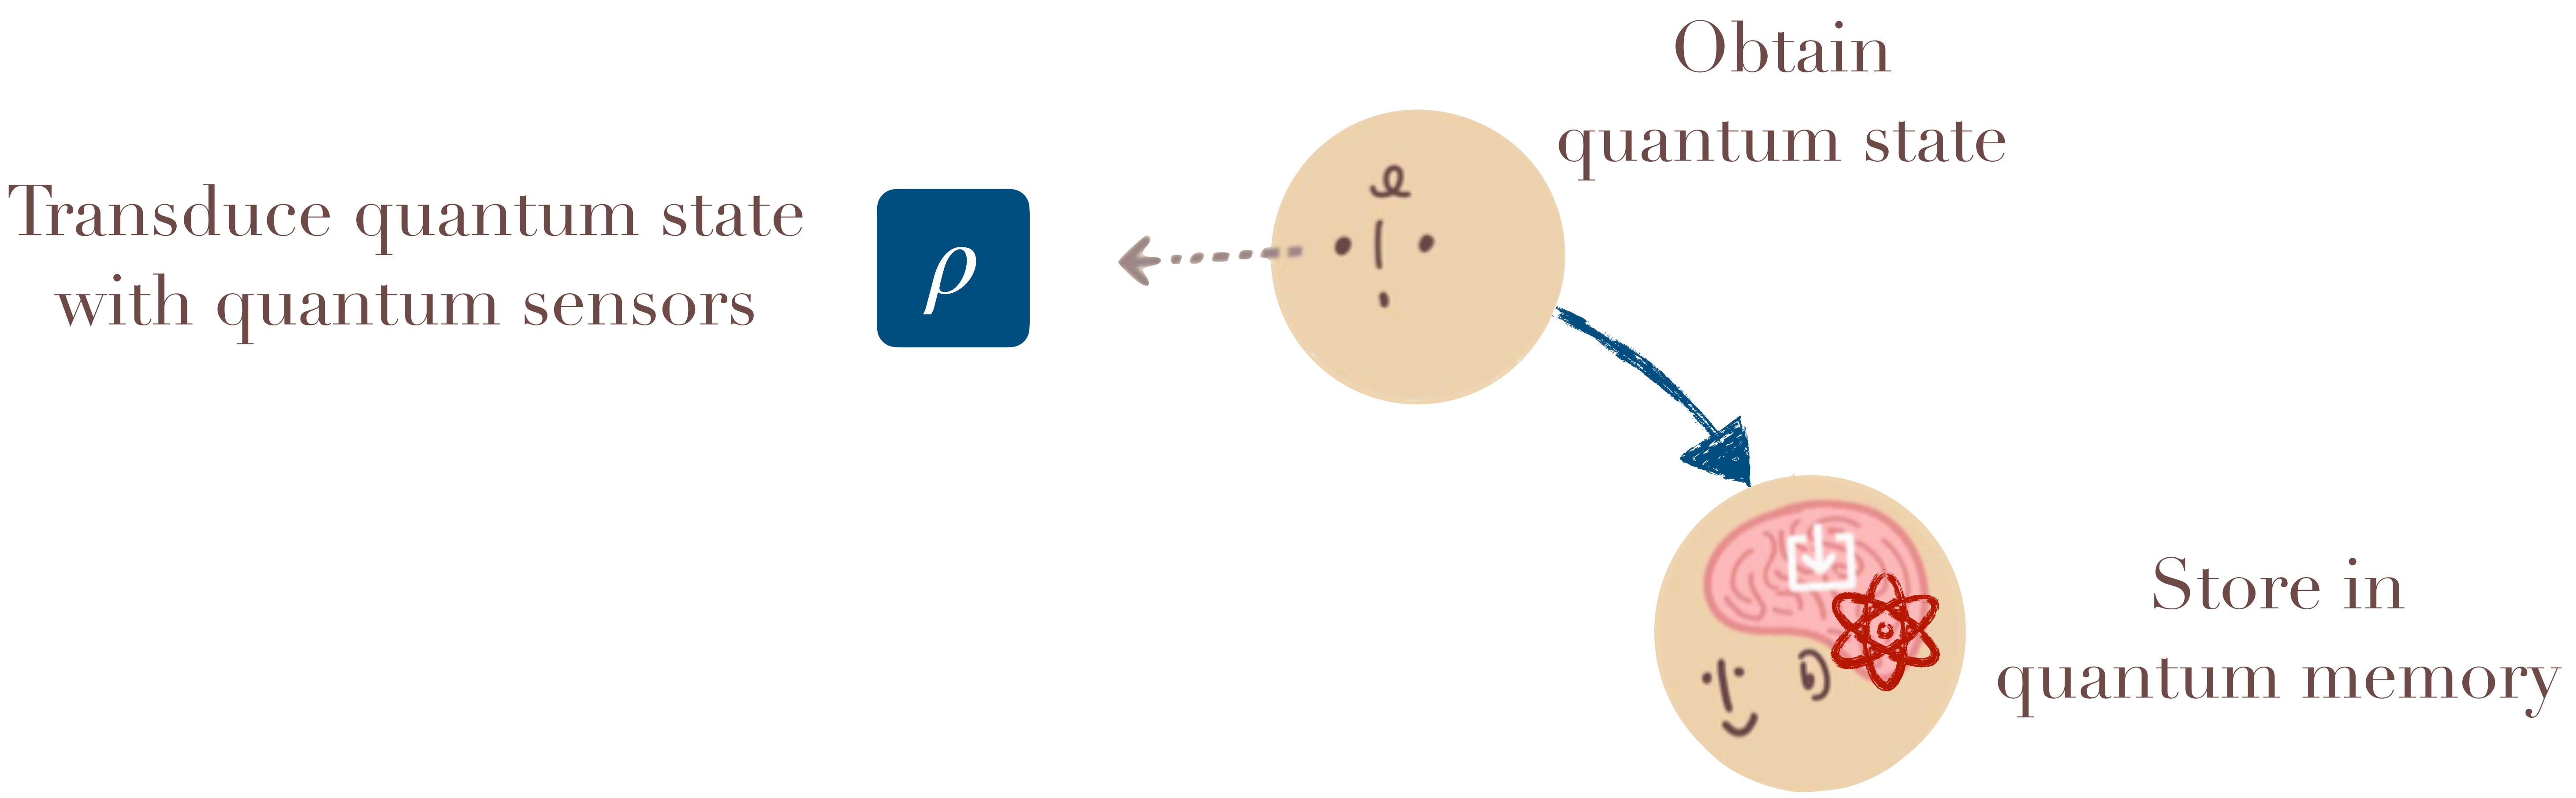

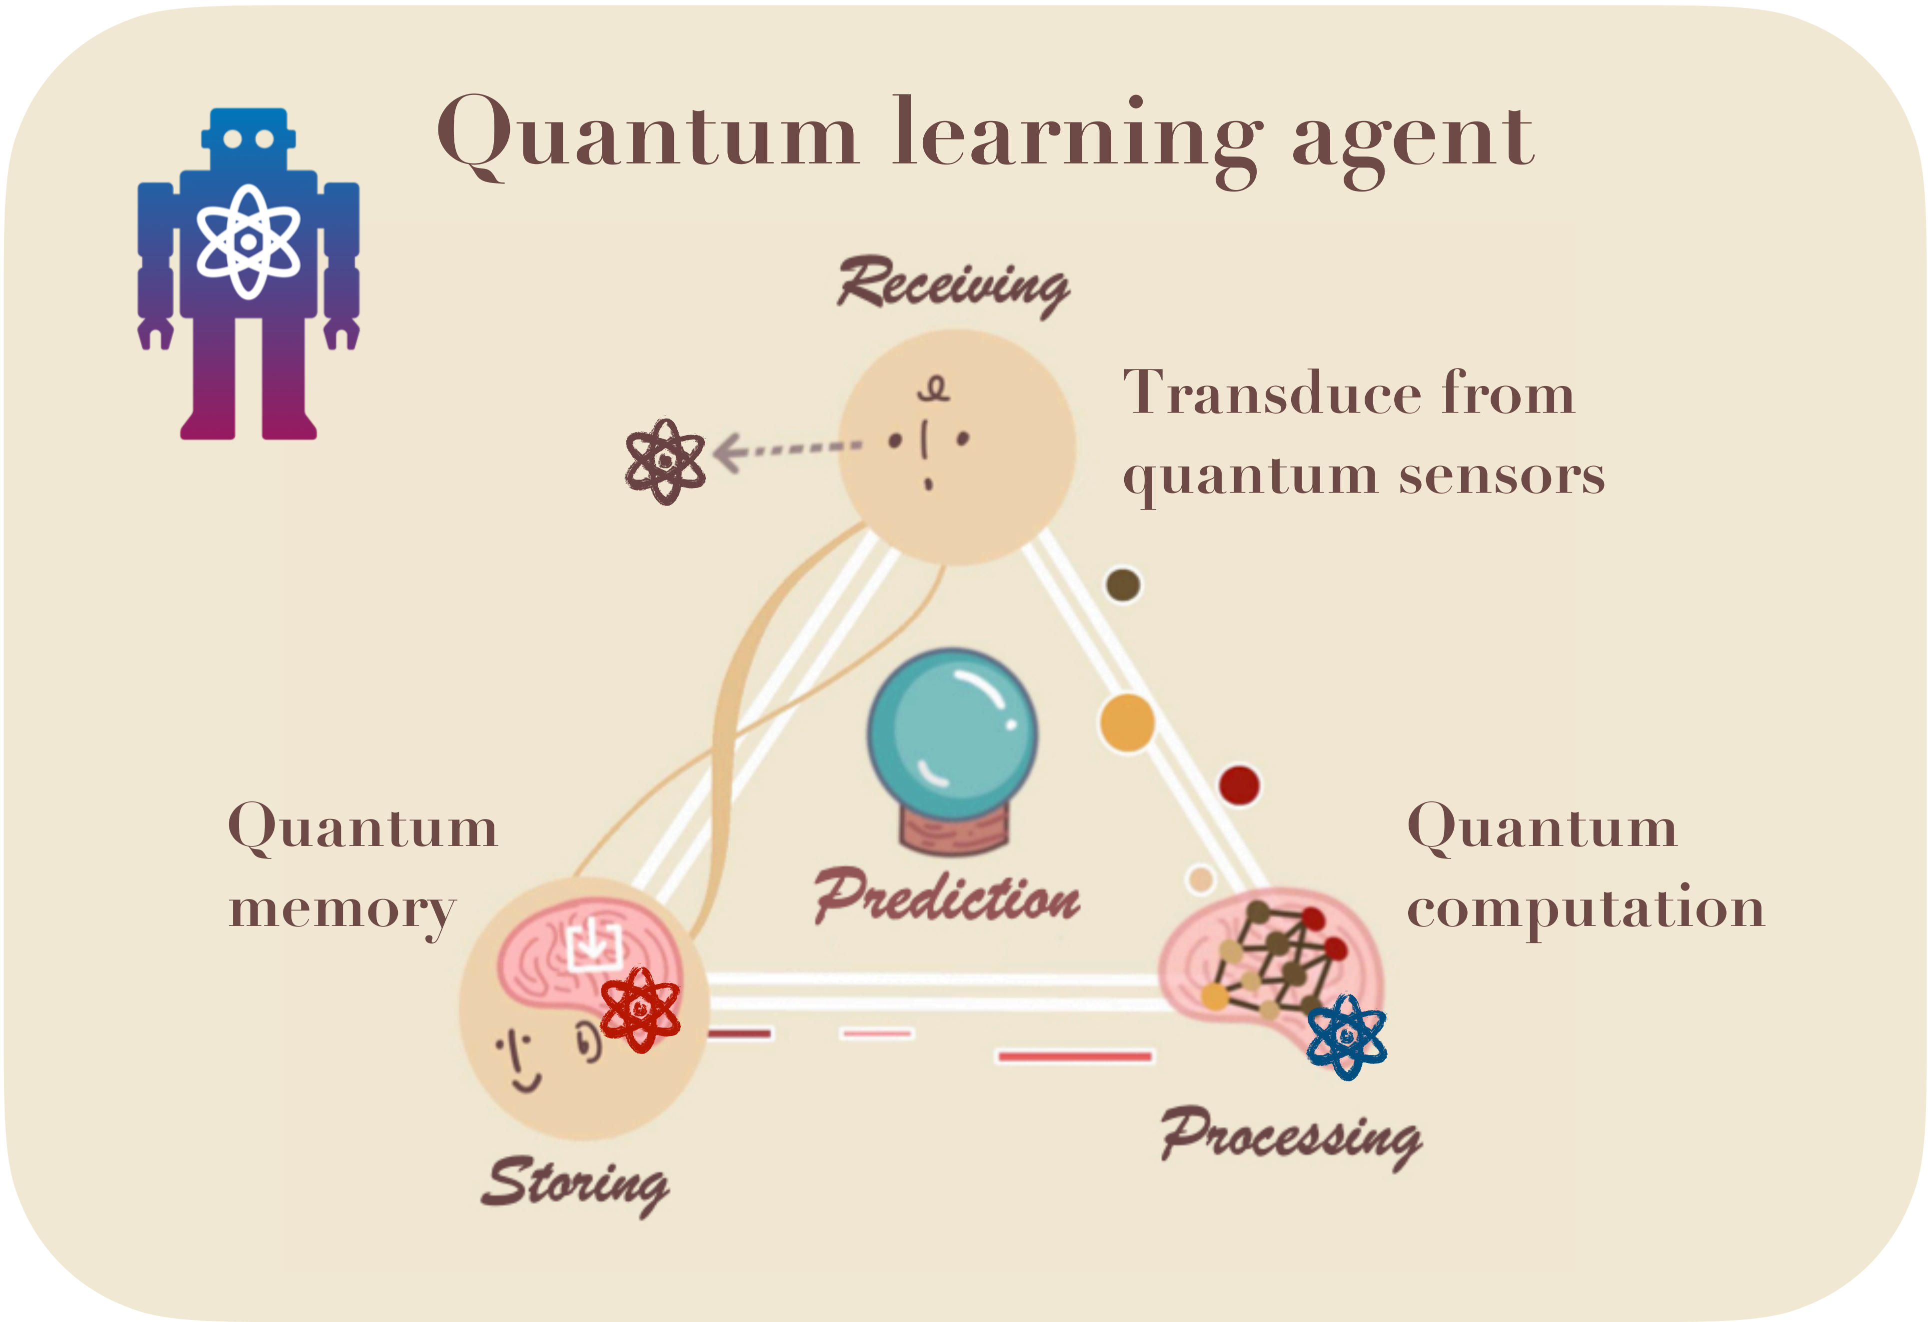

What can quantum agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

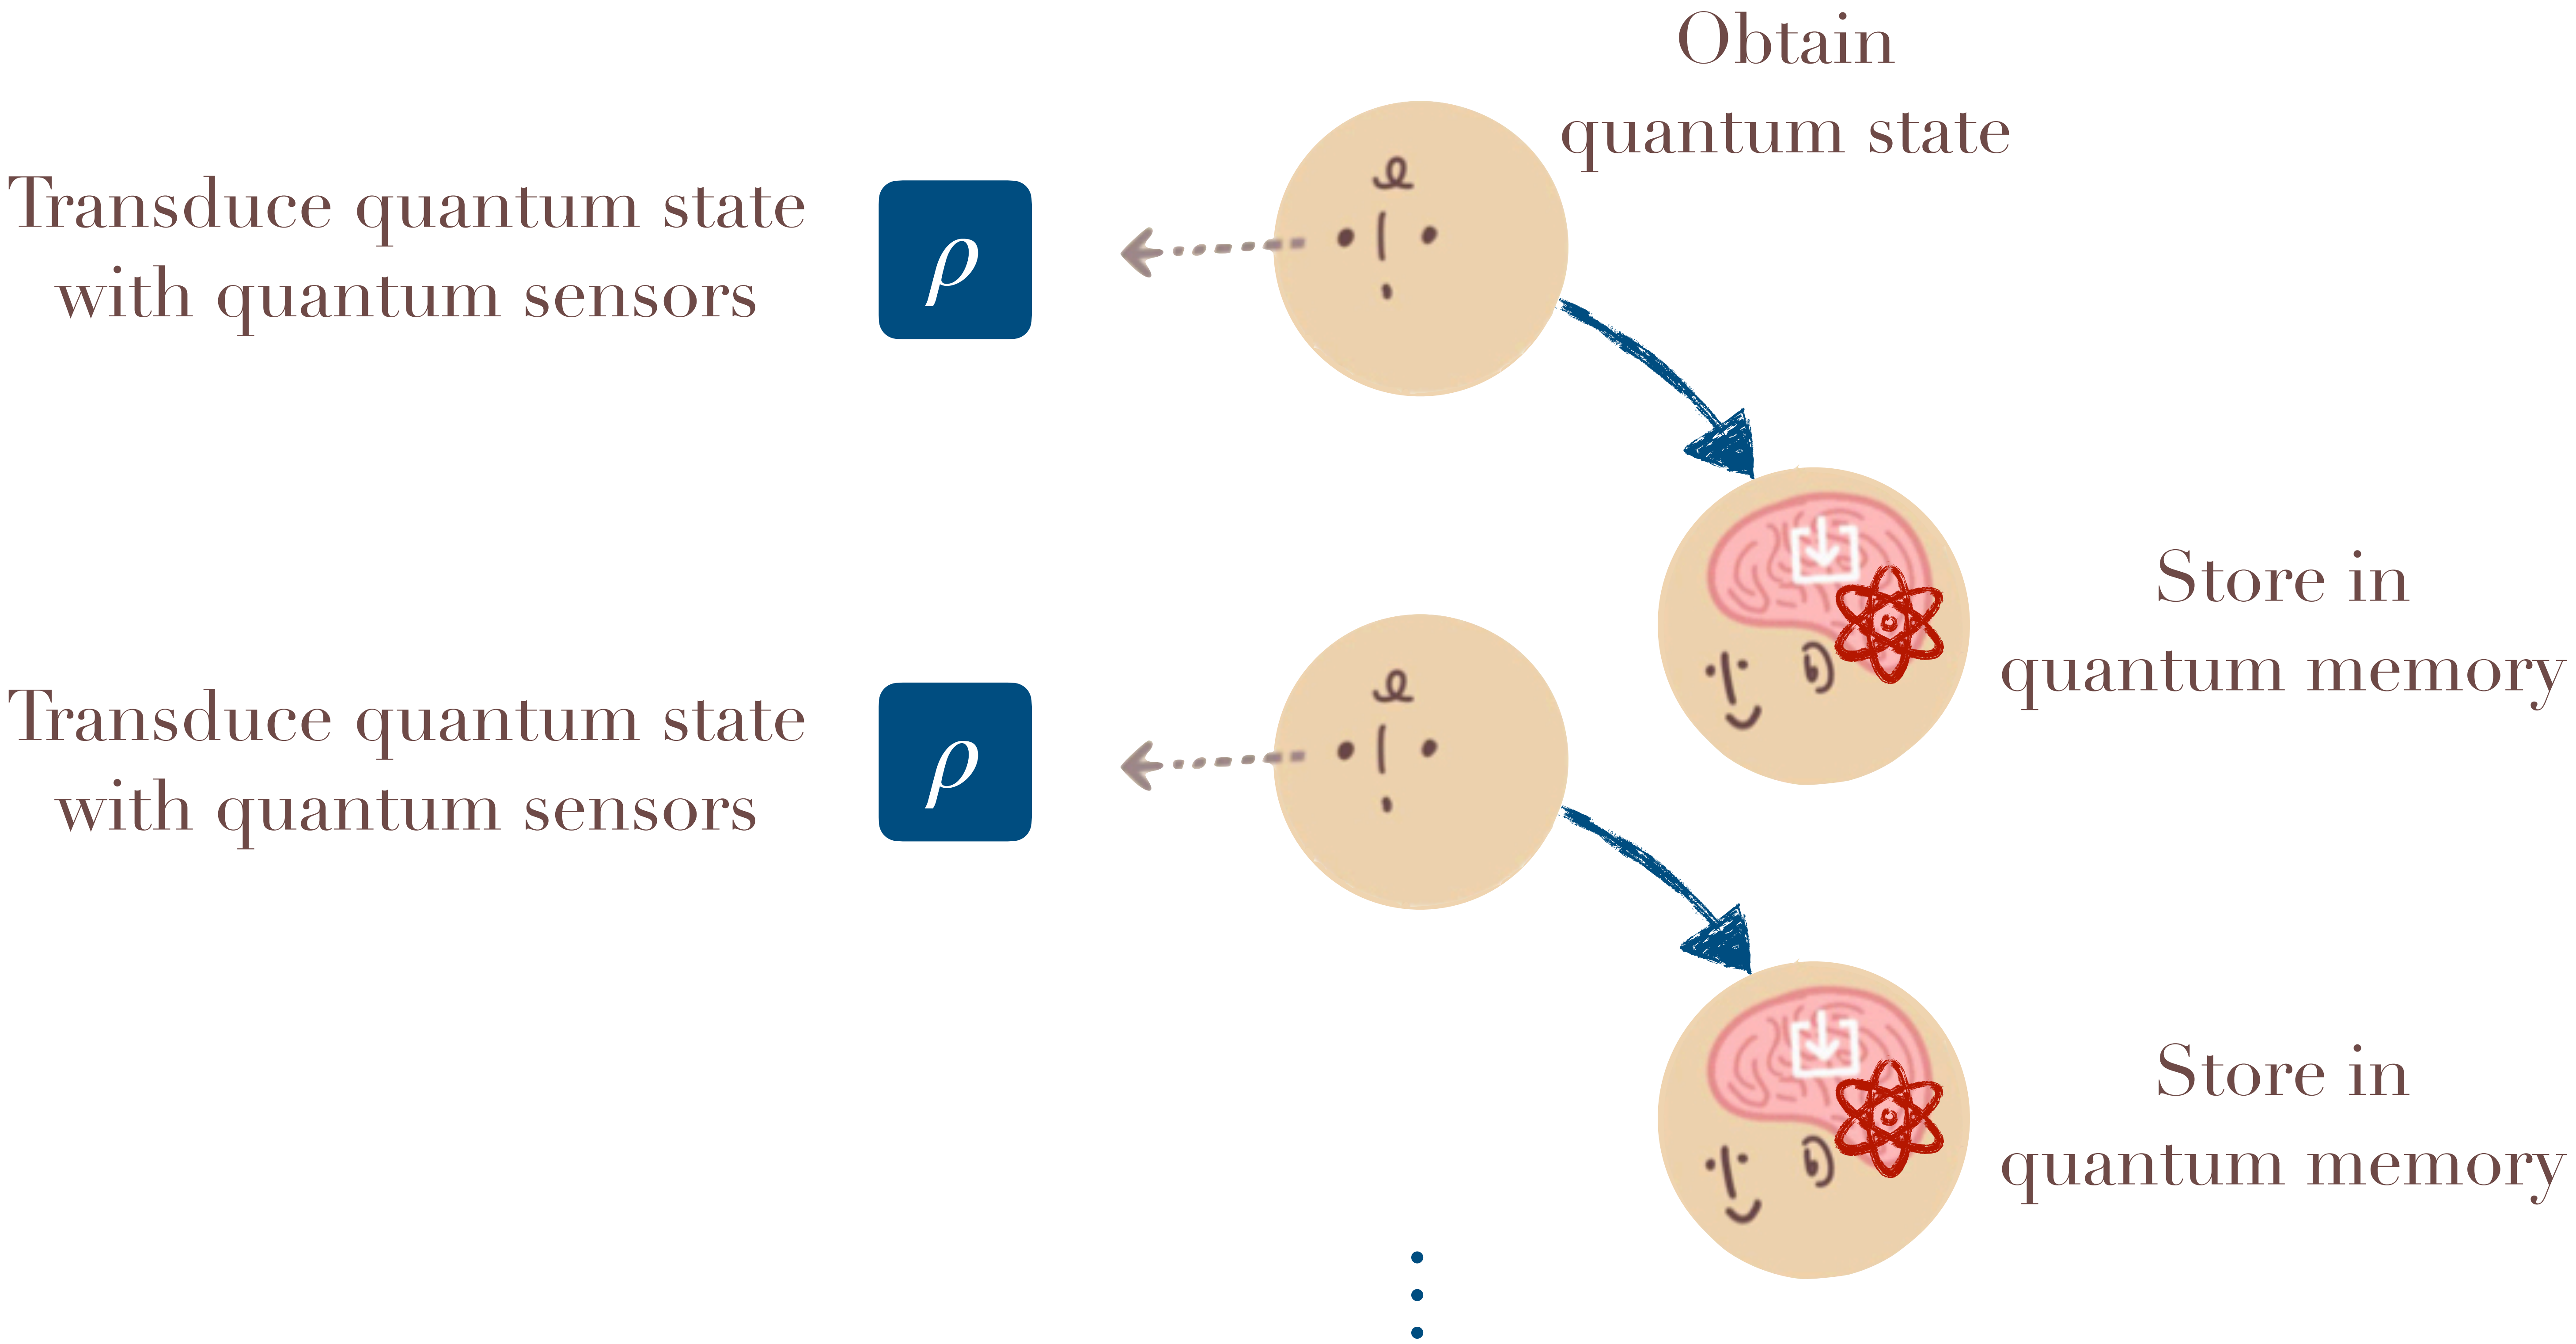

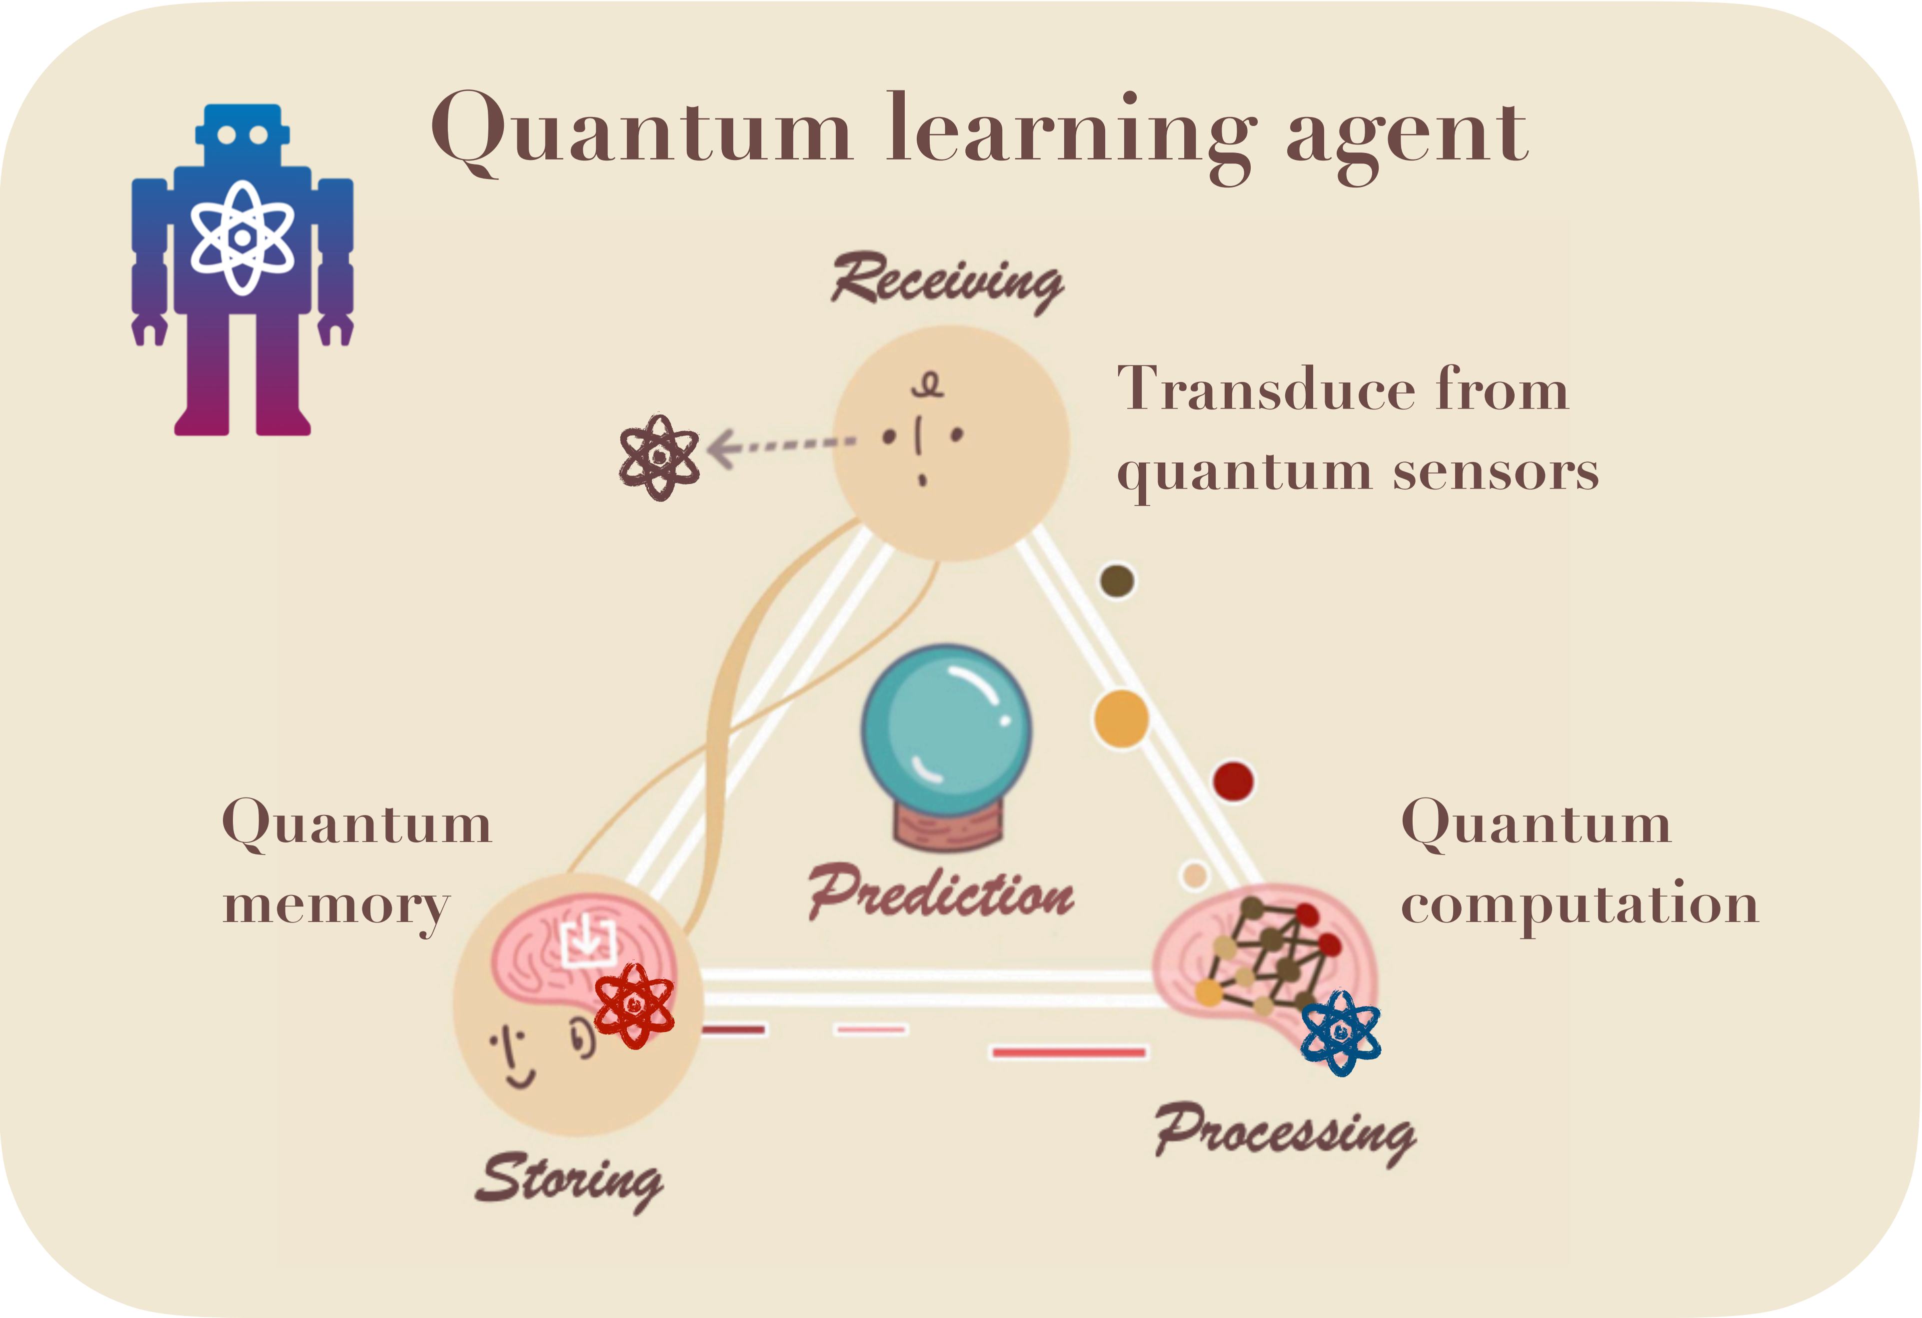

What can quantum agents do?

- We begin with the simpler task of learning about an unknown physical system ρ (density matrix).
- We consider a physical source that could generate a single copy of ρ at a time.

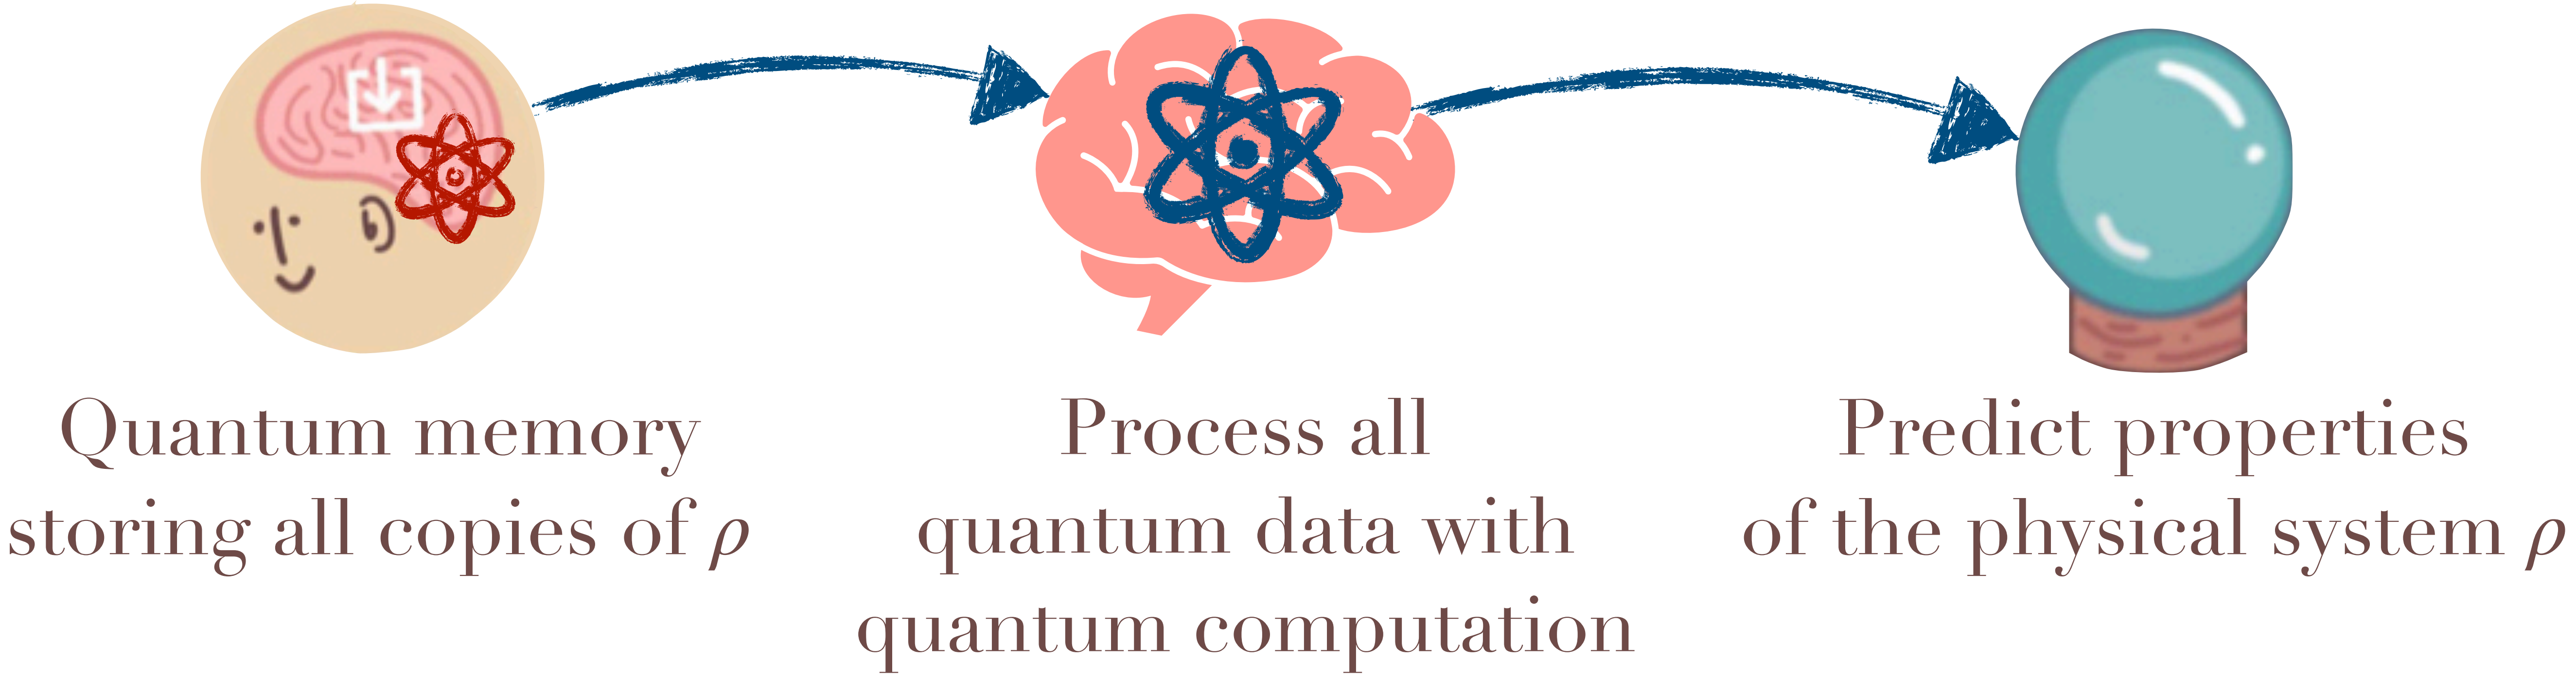

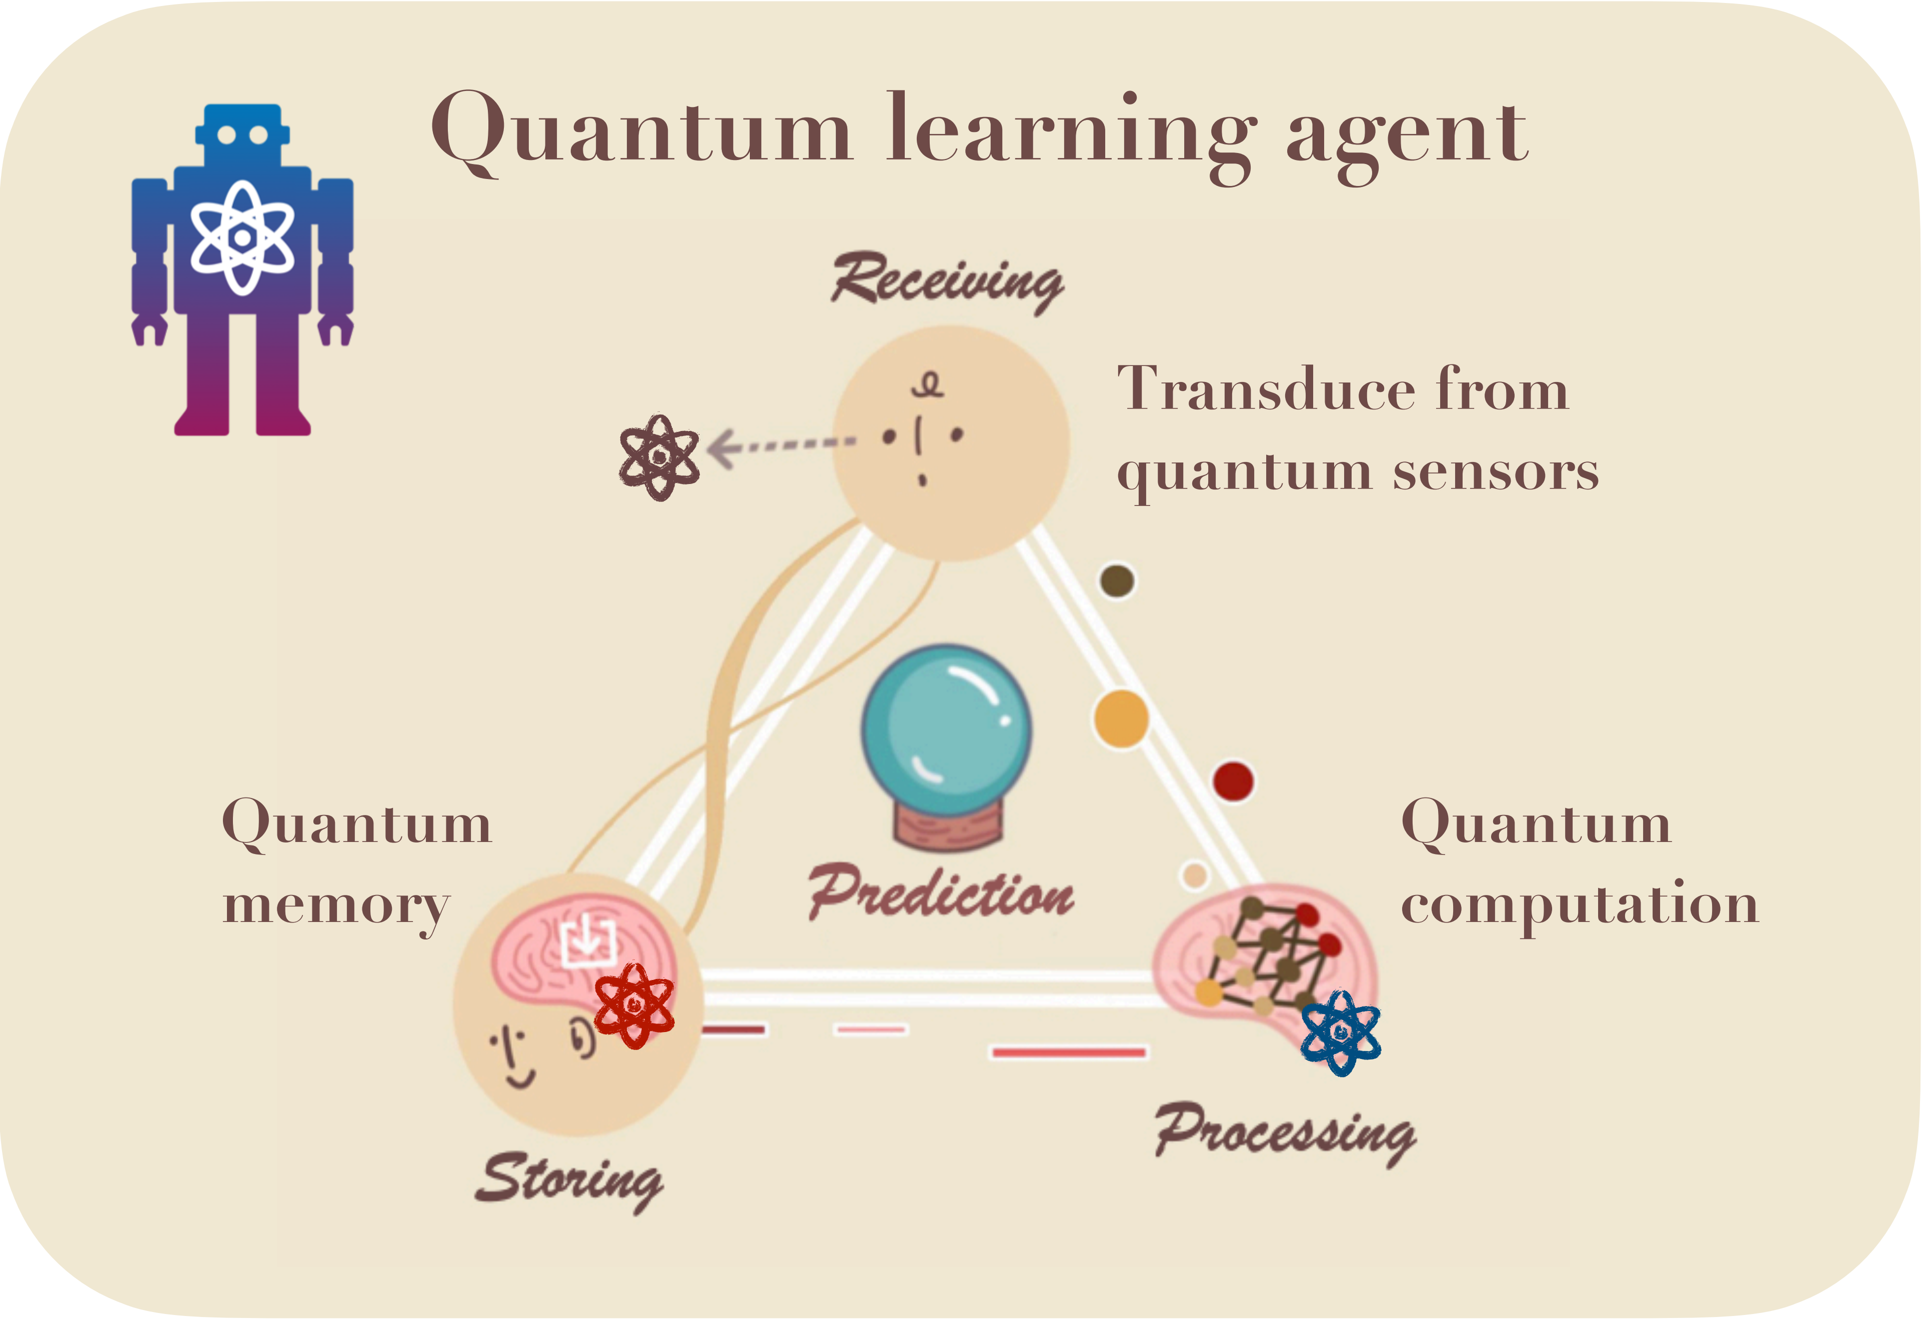

Quantum advantage in predicting Pauli observables

- 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.

Proof Sketch: Tree representation

- Consider the lower bound $\Omega(2^n)$ for classical agents.
- We consider a graphical representation of the memory state of the classical agent.

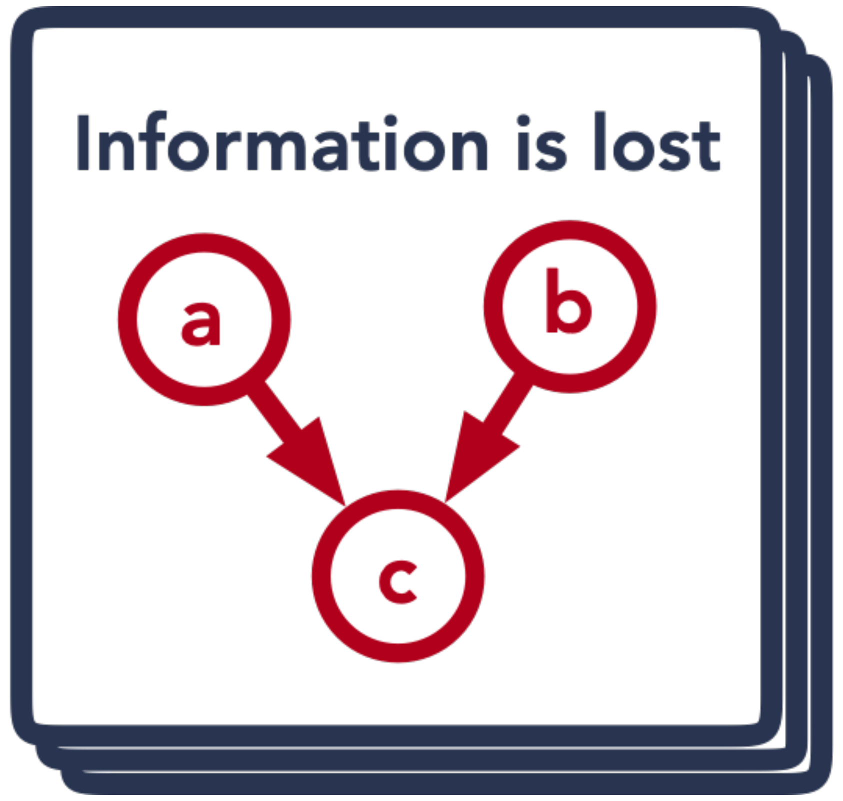

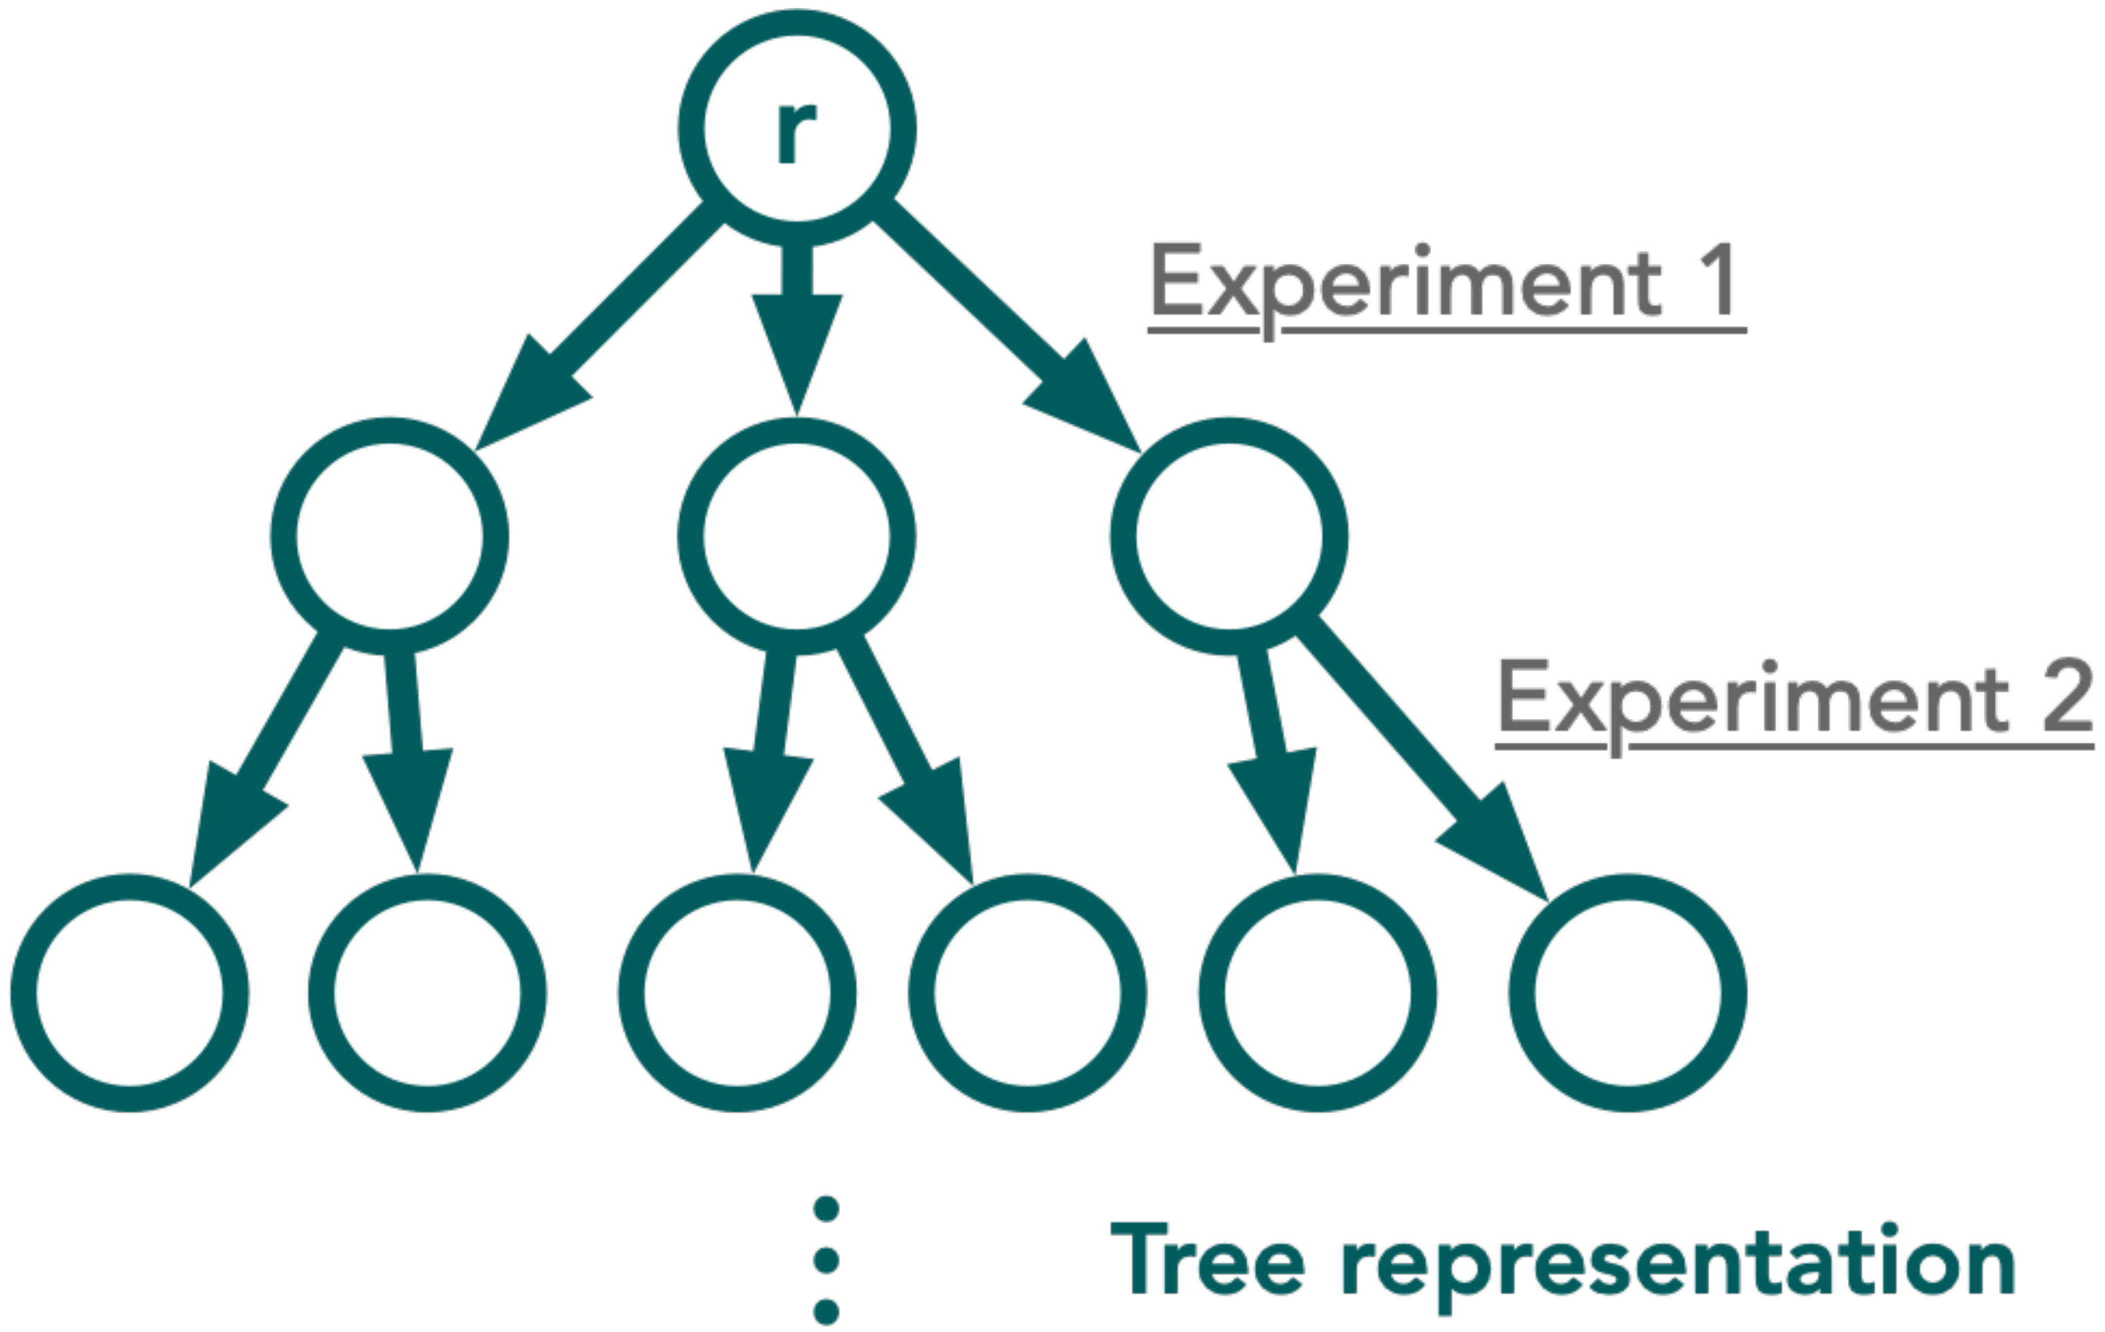

[HKP21] Huang, Kueng, Preskill. Information-theoretic bounds on quantum advantage in machine learning, PRL 2021.

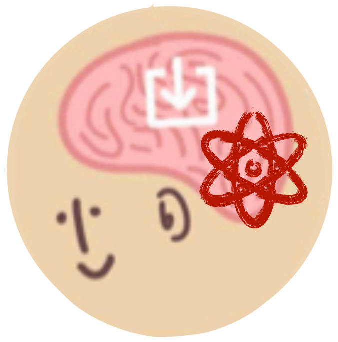

Proof Sketch: Tree representation

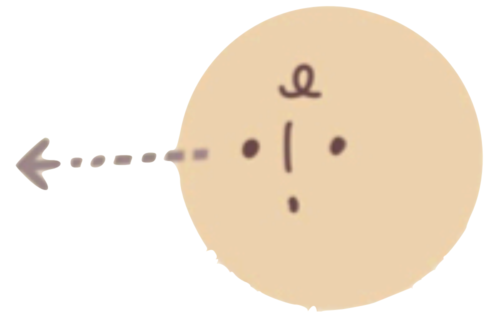

- Consider the lower bound $\Omega(2^n)$ for classical agents.
- We consider a graphical representation of the memory state of the classical agent.

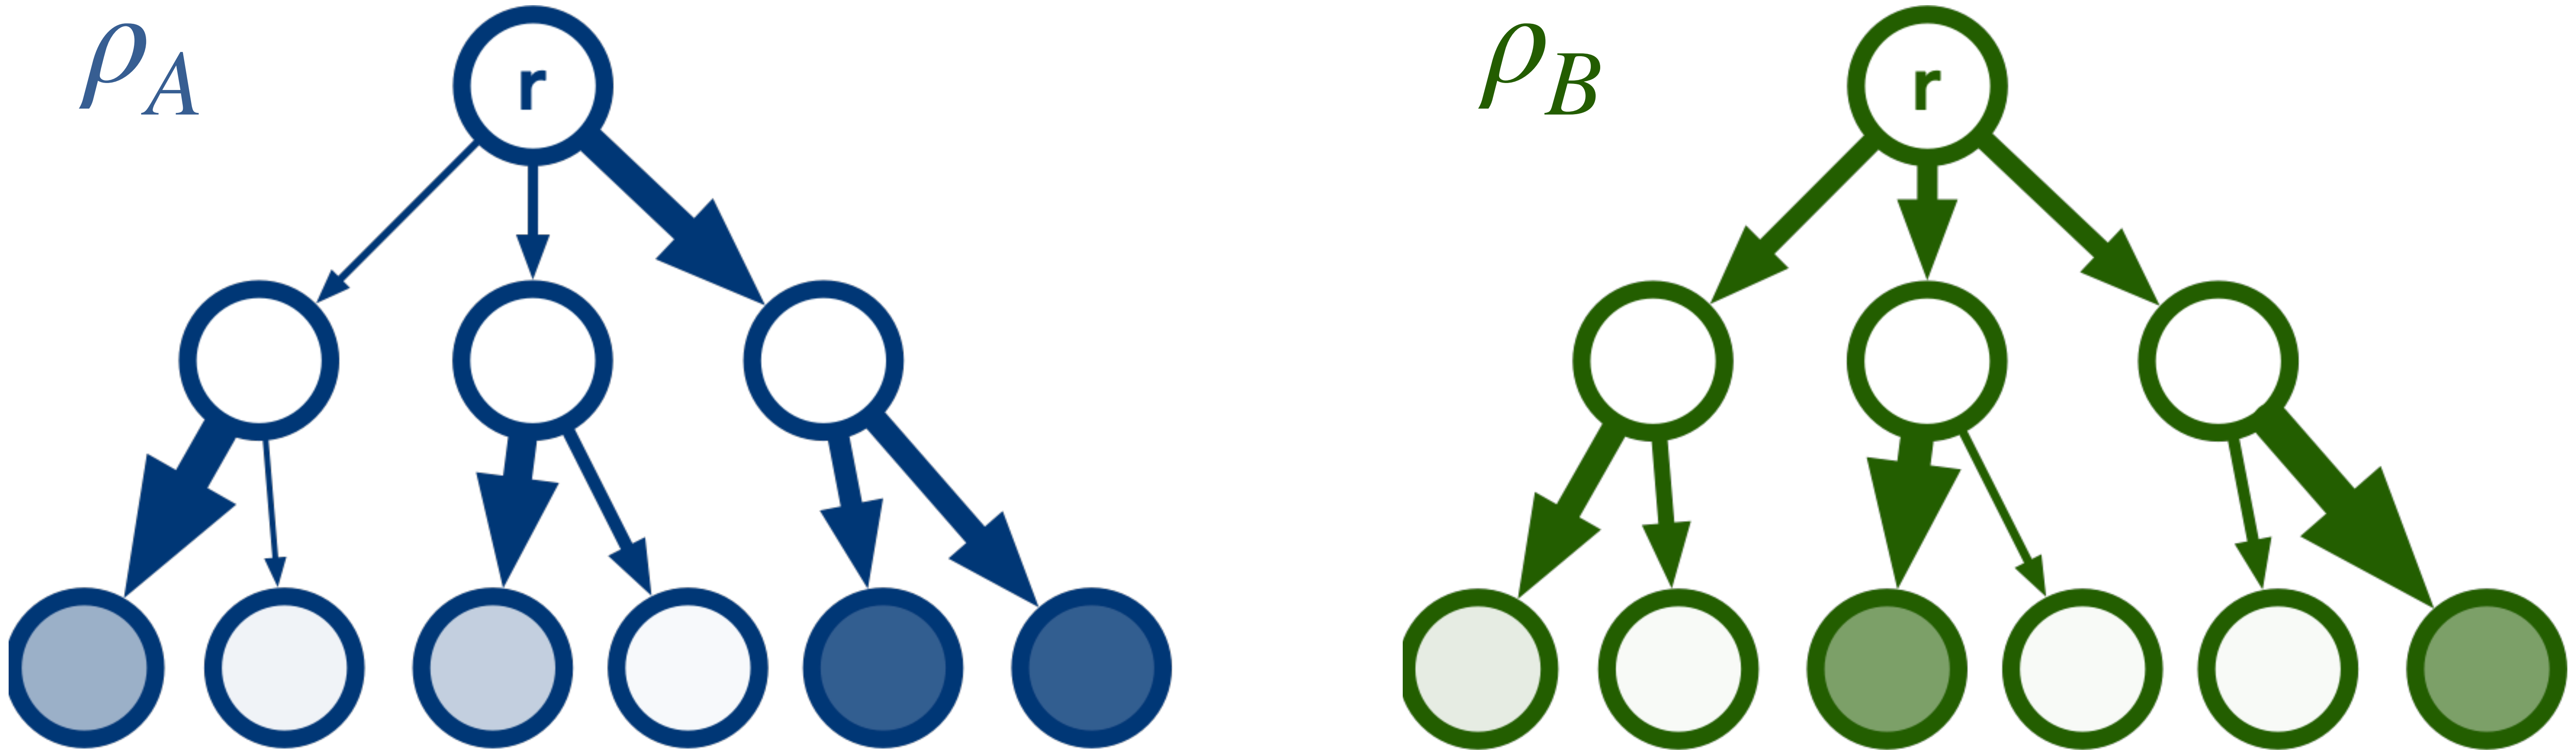

[HKP21] Huang, Kueng, Preskill. Information-theoretic bounds on quantum advantage in machine learning, PRL 2021.

Many-vs-one distinguishing task

- We reduce the prediction task to a many-vs-one distinguishing task.
- Prediction task: Estimate $\mathrm{Tr}(P\rho)$ to $1/4$ error for all $P \in \{I, X, Y, Z\}^{\otimes n} \backslash \{I^{\otimes n}\}$.
- Distinguishing task: Distinguish between null hypothesis and the alternative hypothesis.

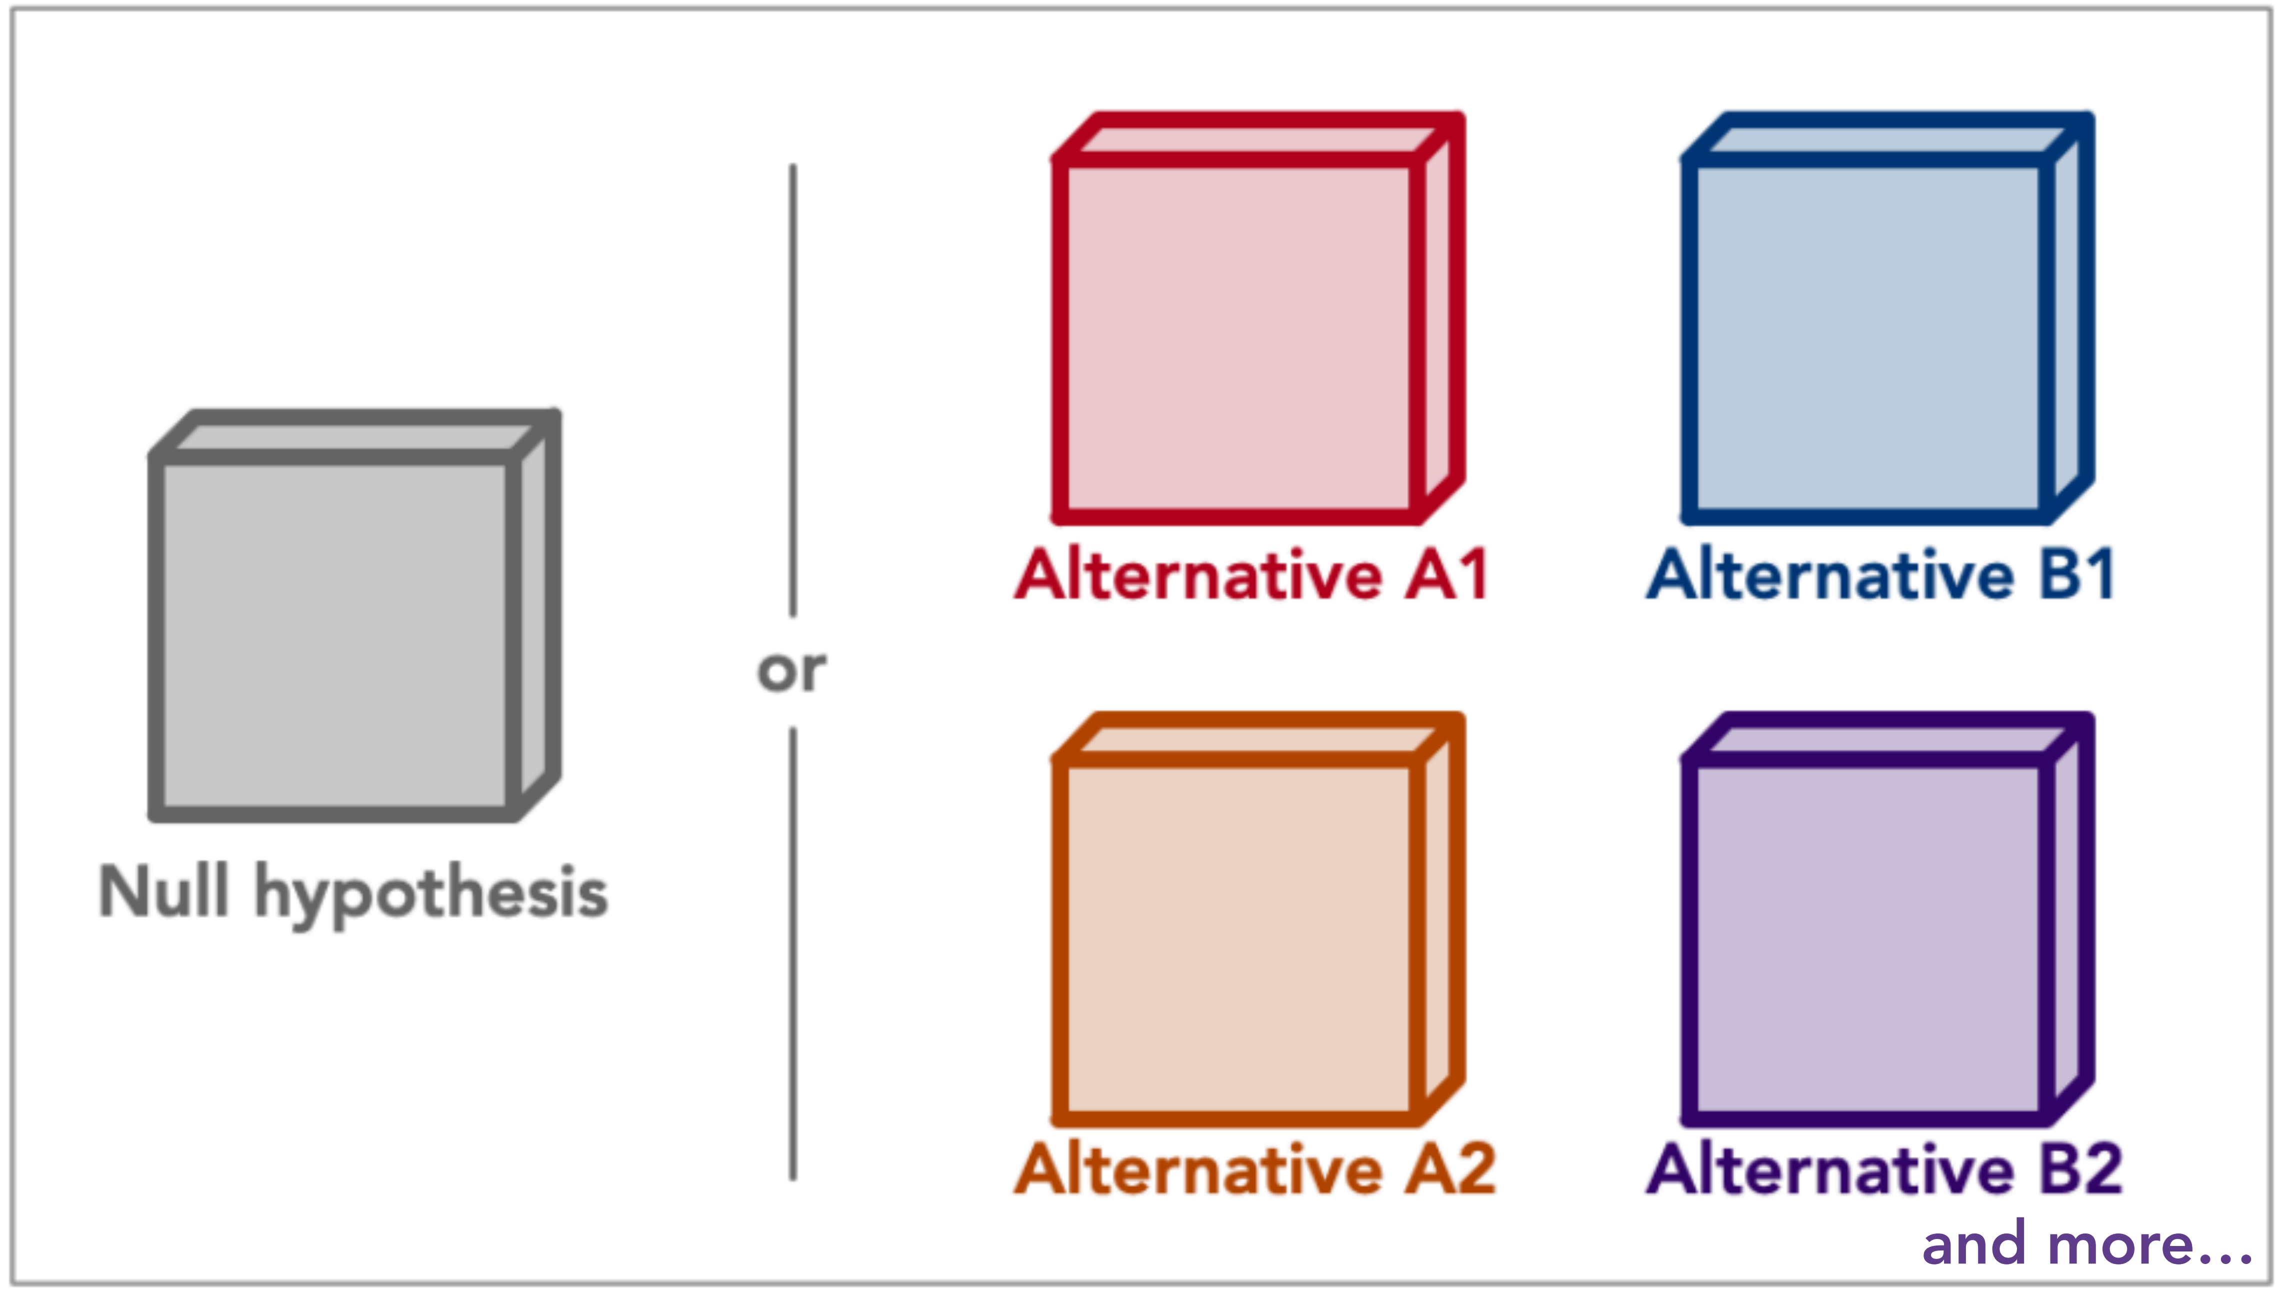

Many-versus-one distinguishing task

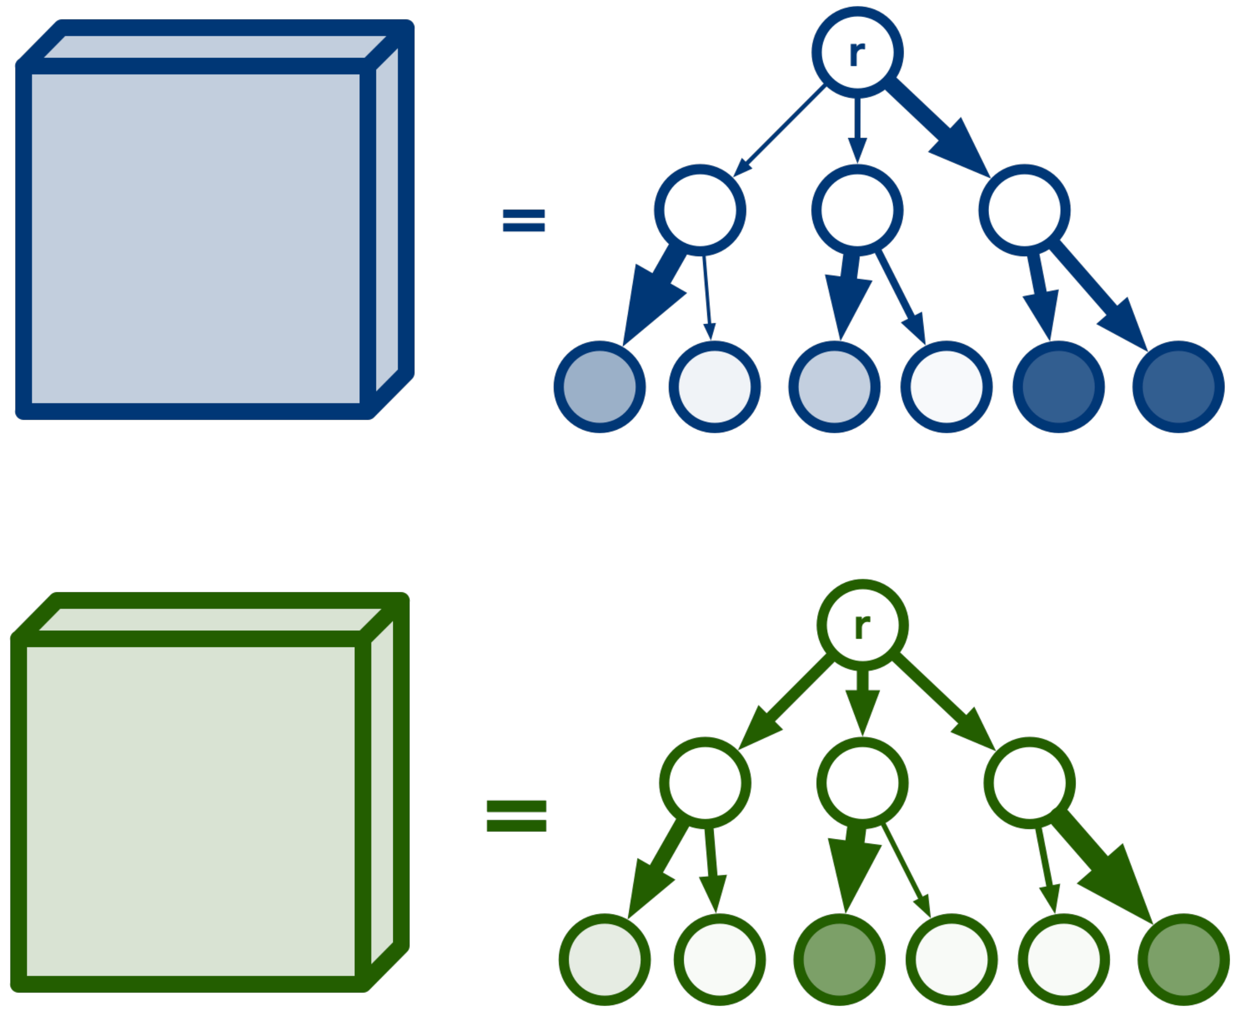

Many-vs-one distinguishing task

- We reduce the prediction task to a many-vs-one distinguishing task.
- Prediction task: Estimate $\mathrm{Tr}(P\rho)$ to $1/4$ error for all $P \in \{I, X, Y, Z\}^{\otimes n} \backslash \{I^{\otimes n}\}$.
- Distinguishing task: Distinguish between null hypothesis and the alternative hypothesis.

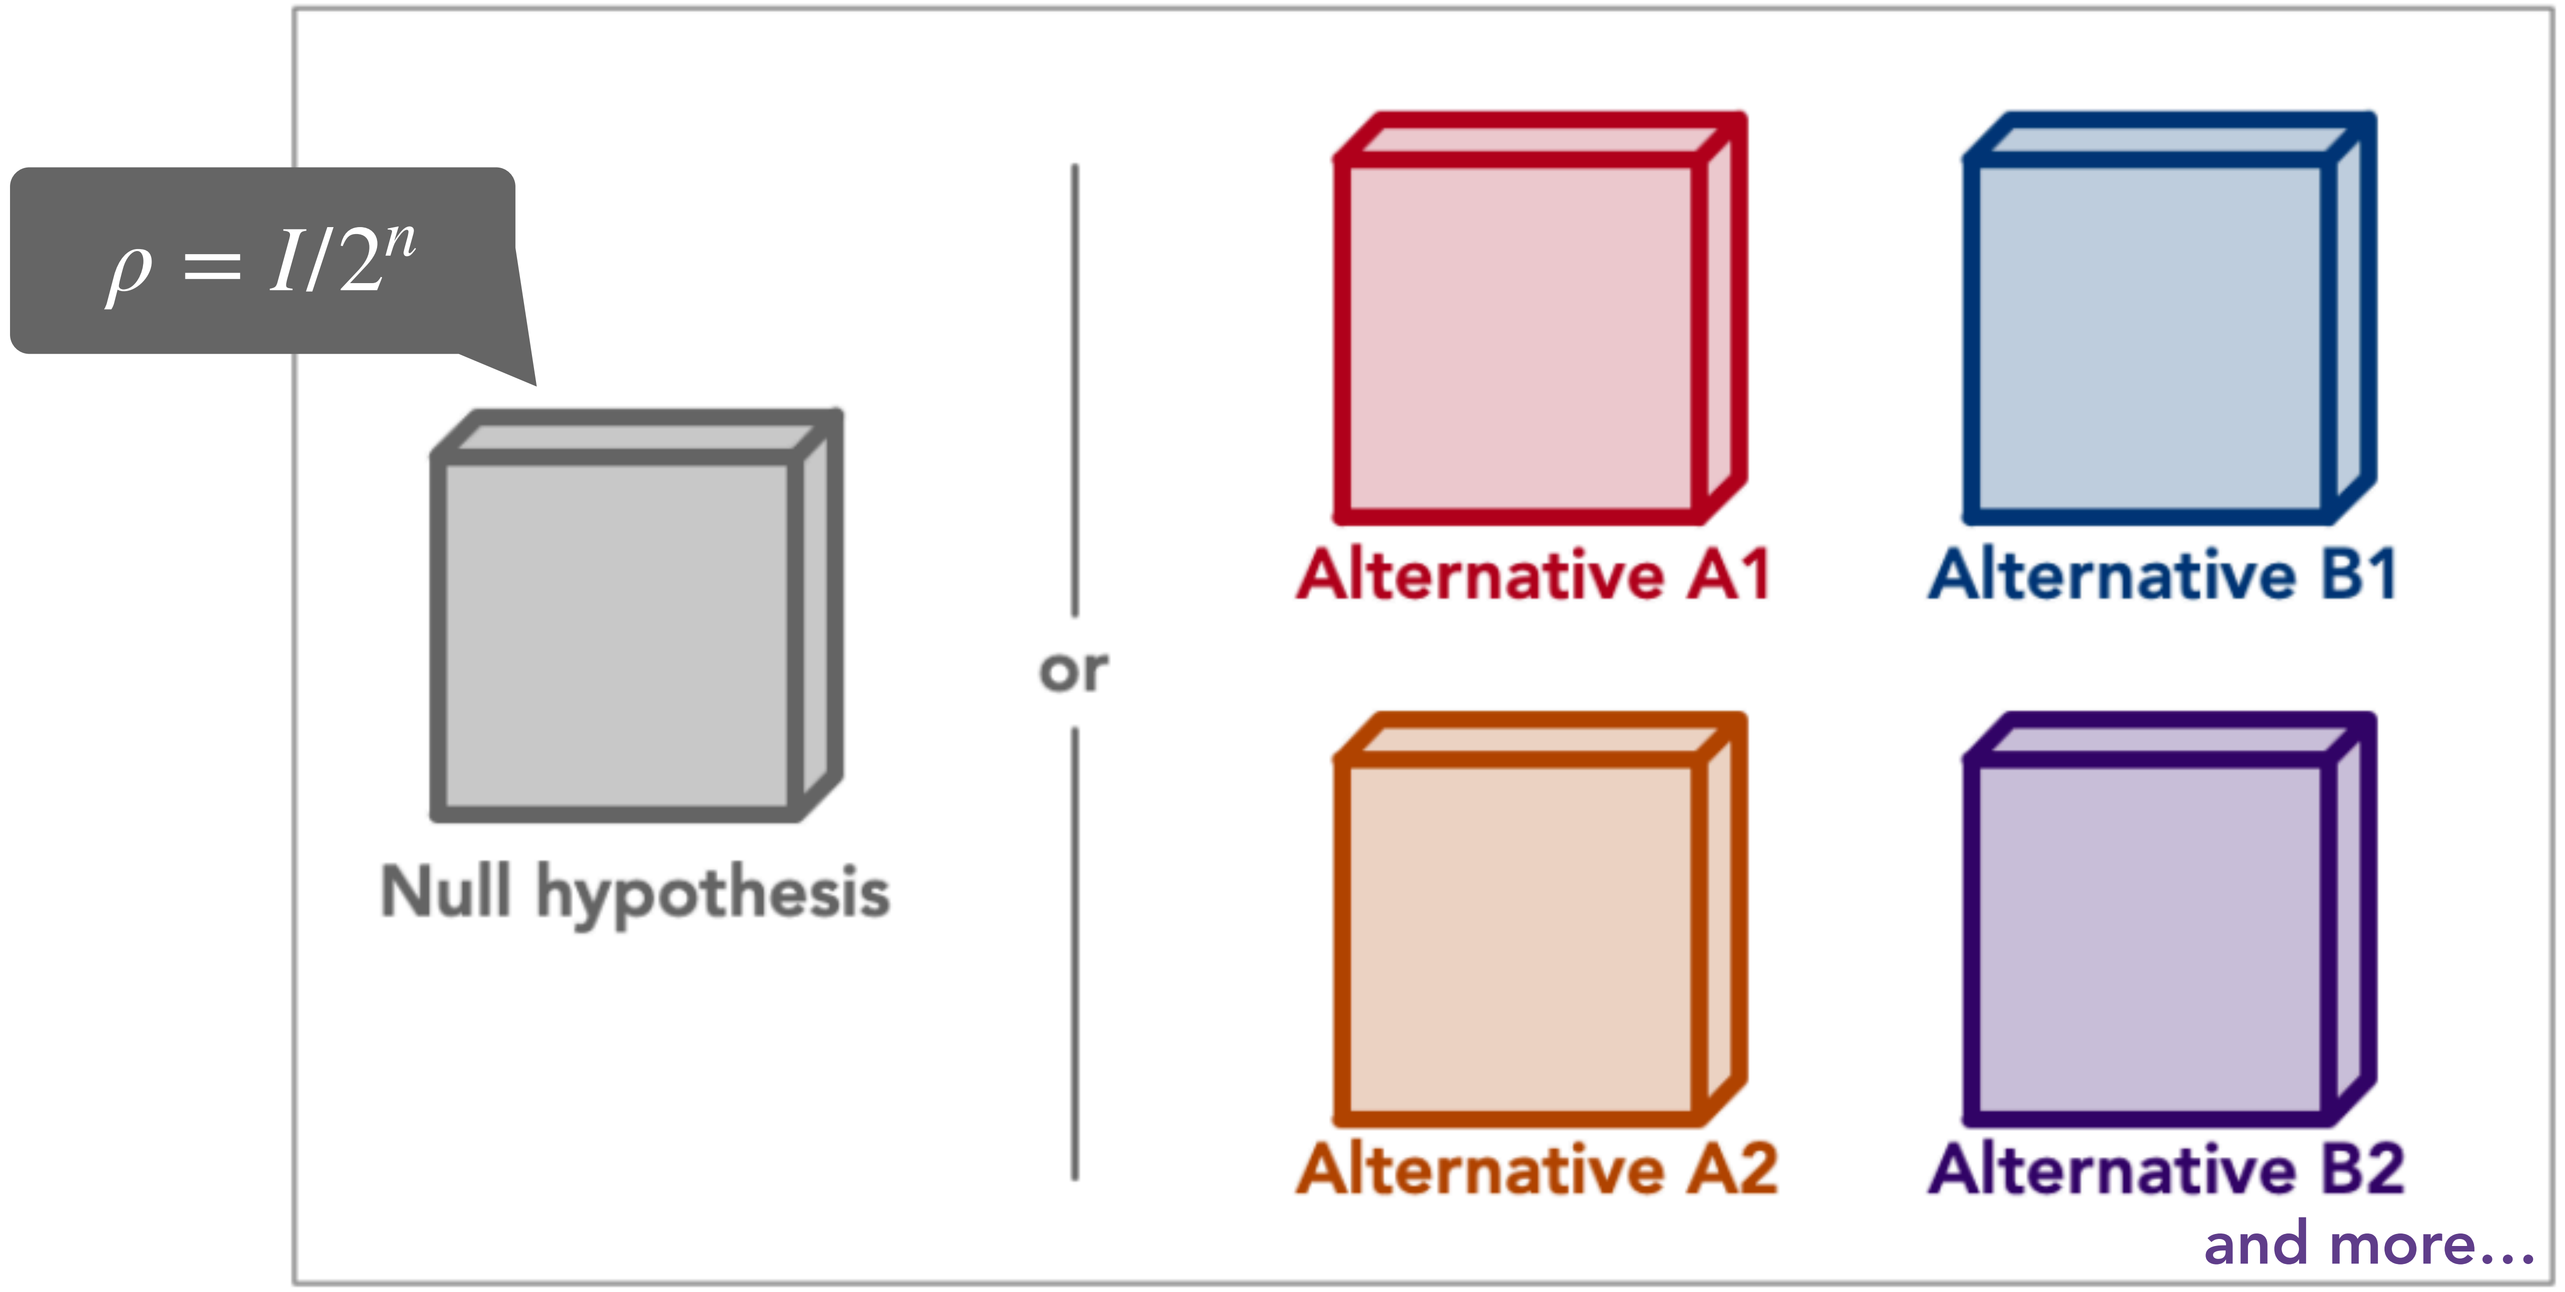

Many-versus-one distinguishing task

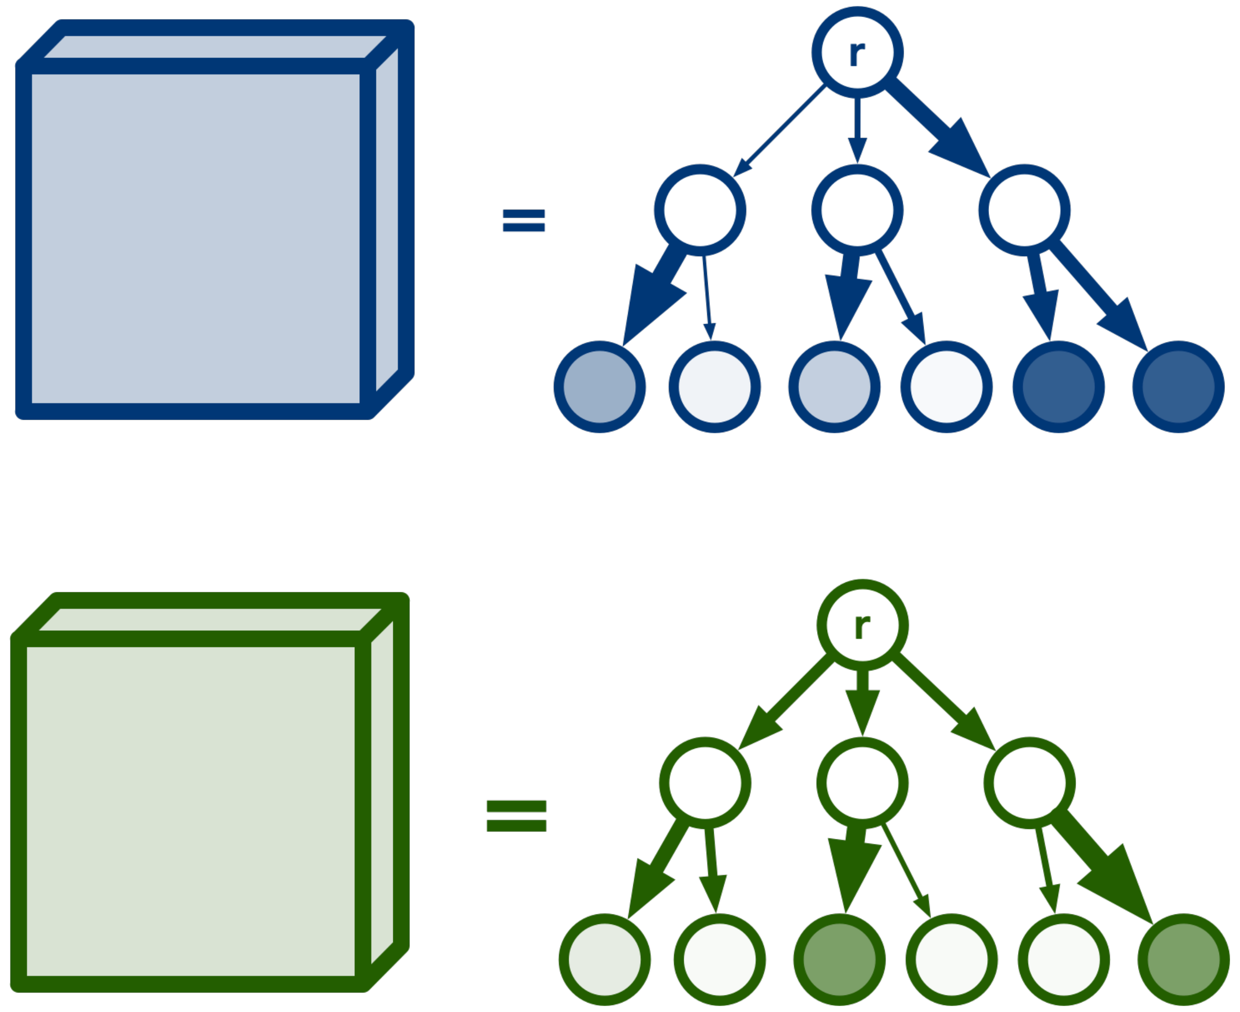

Many-vs-one distinguishing task

- We reduce the prediction task to a many-vs-one distinguishing task.
- Prediction task: Estimate $\mathrm{Tr}(P\rho)$ to $1/4$ error for all $P \in \{I, X, Y, Z\}^{\otimes n} \backslash \{I^{\otimes n}\}$.
- Distinguishing task: Distinguish between null hypothesis and the alternative hypothesis.

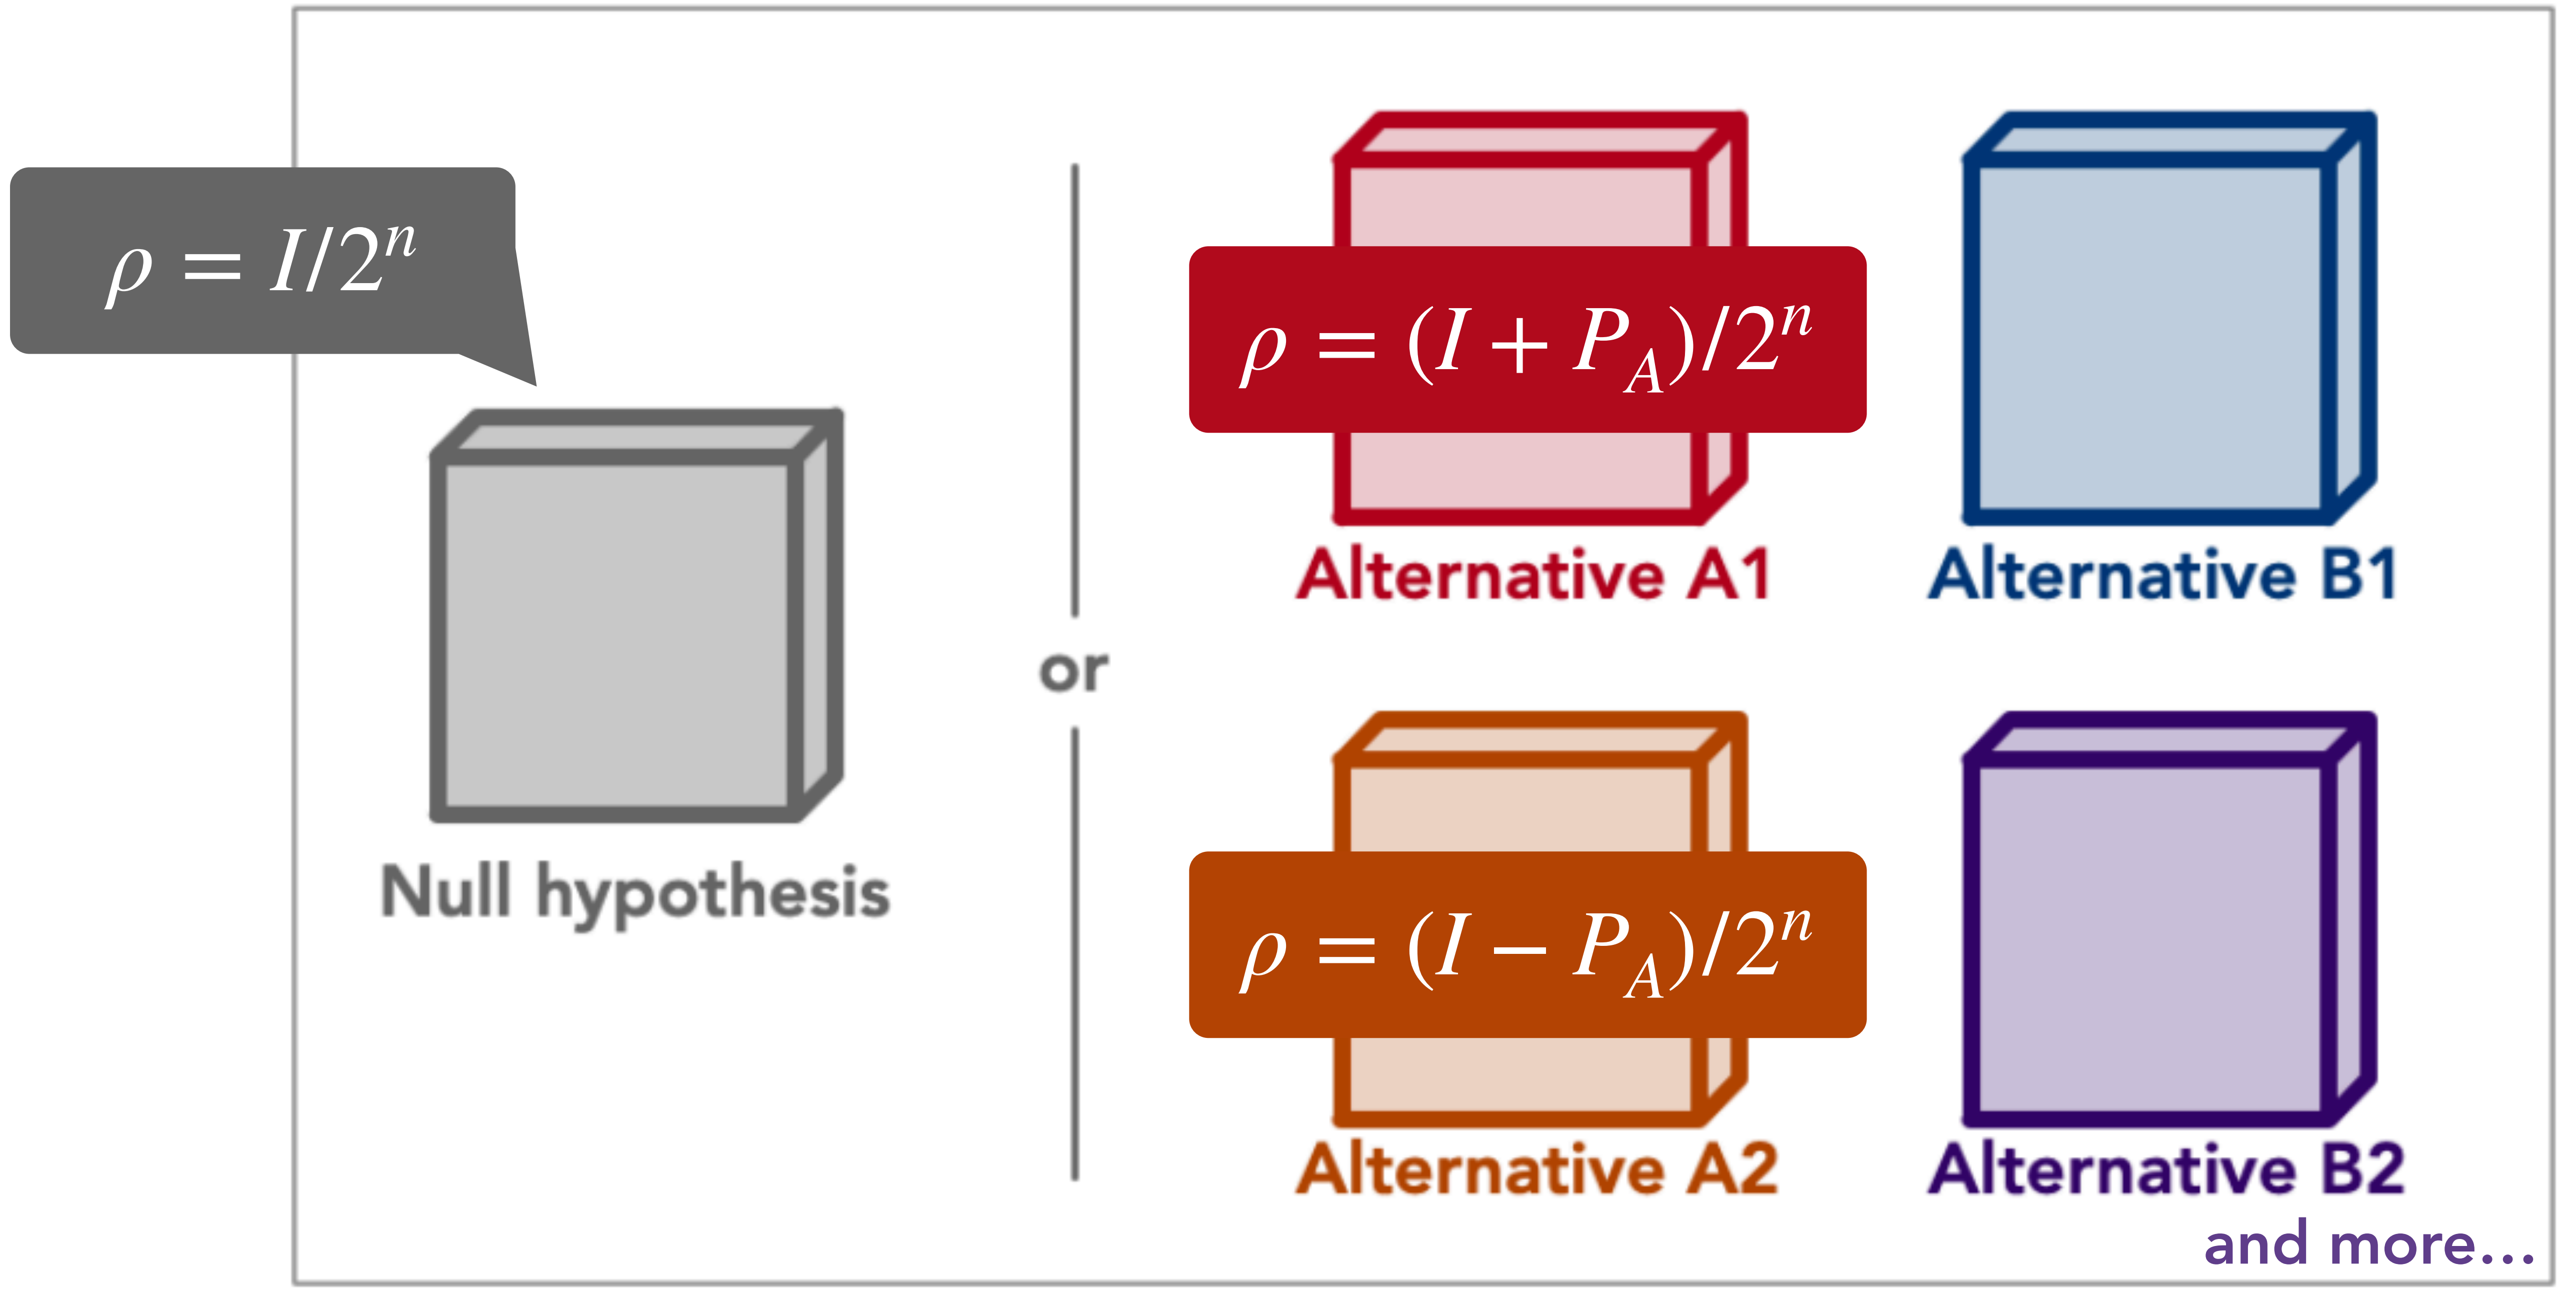

Many-versus-one distinguishing task

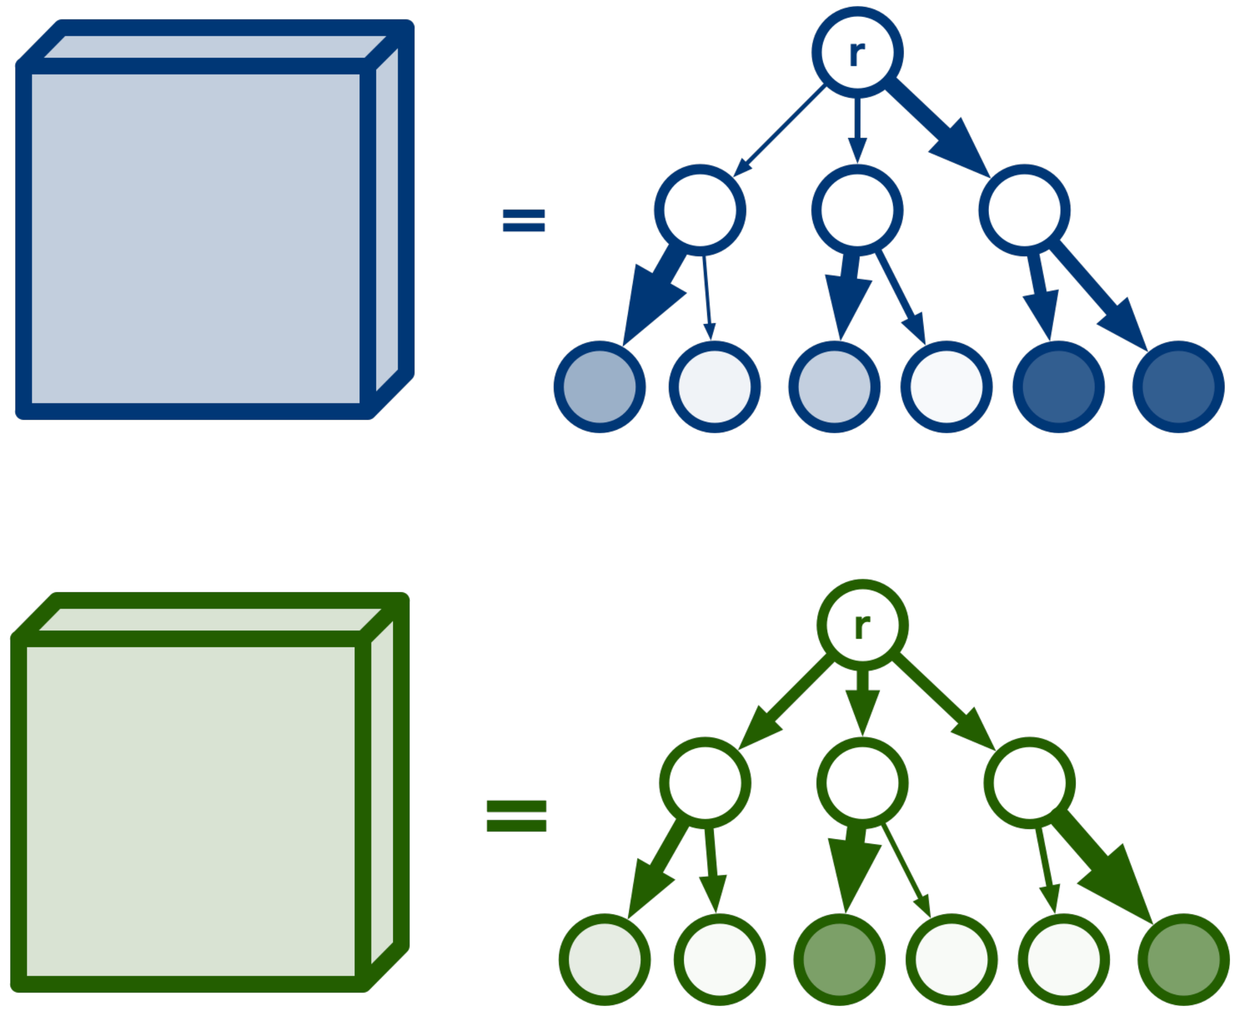

Many-vs-one distinguishing task

- We reduce the prediction task to a many-vs-one distinguishing task.
- Prediction task: Estimate $\mathrm{Tr}(P\rho)$ to $1/4$ error for all $P \in \{I, X, Y, Z\}^{\otimes n} \backslash \{I^{\otimes n}\}$.
- Distinguishing task: Distinguish between null hypothesis and the alternative hypothesis.

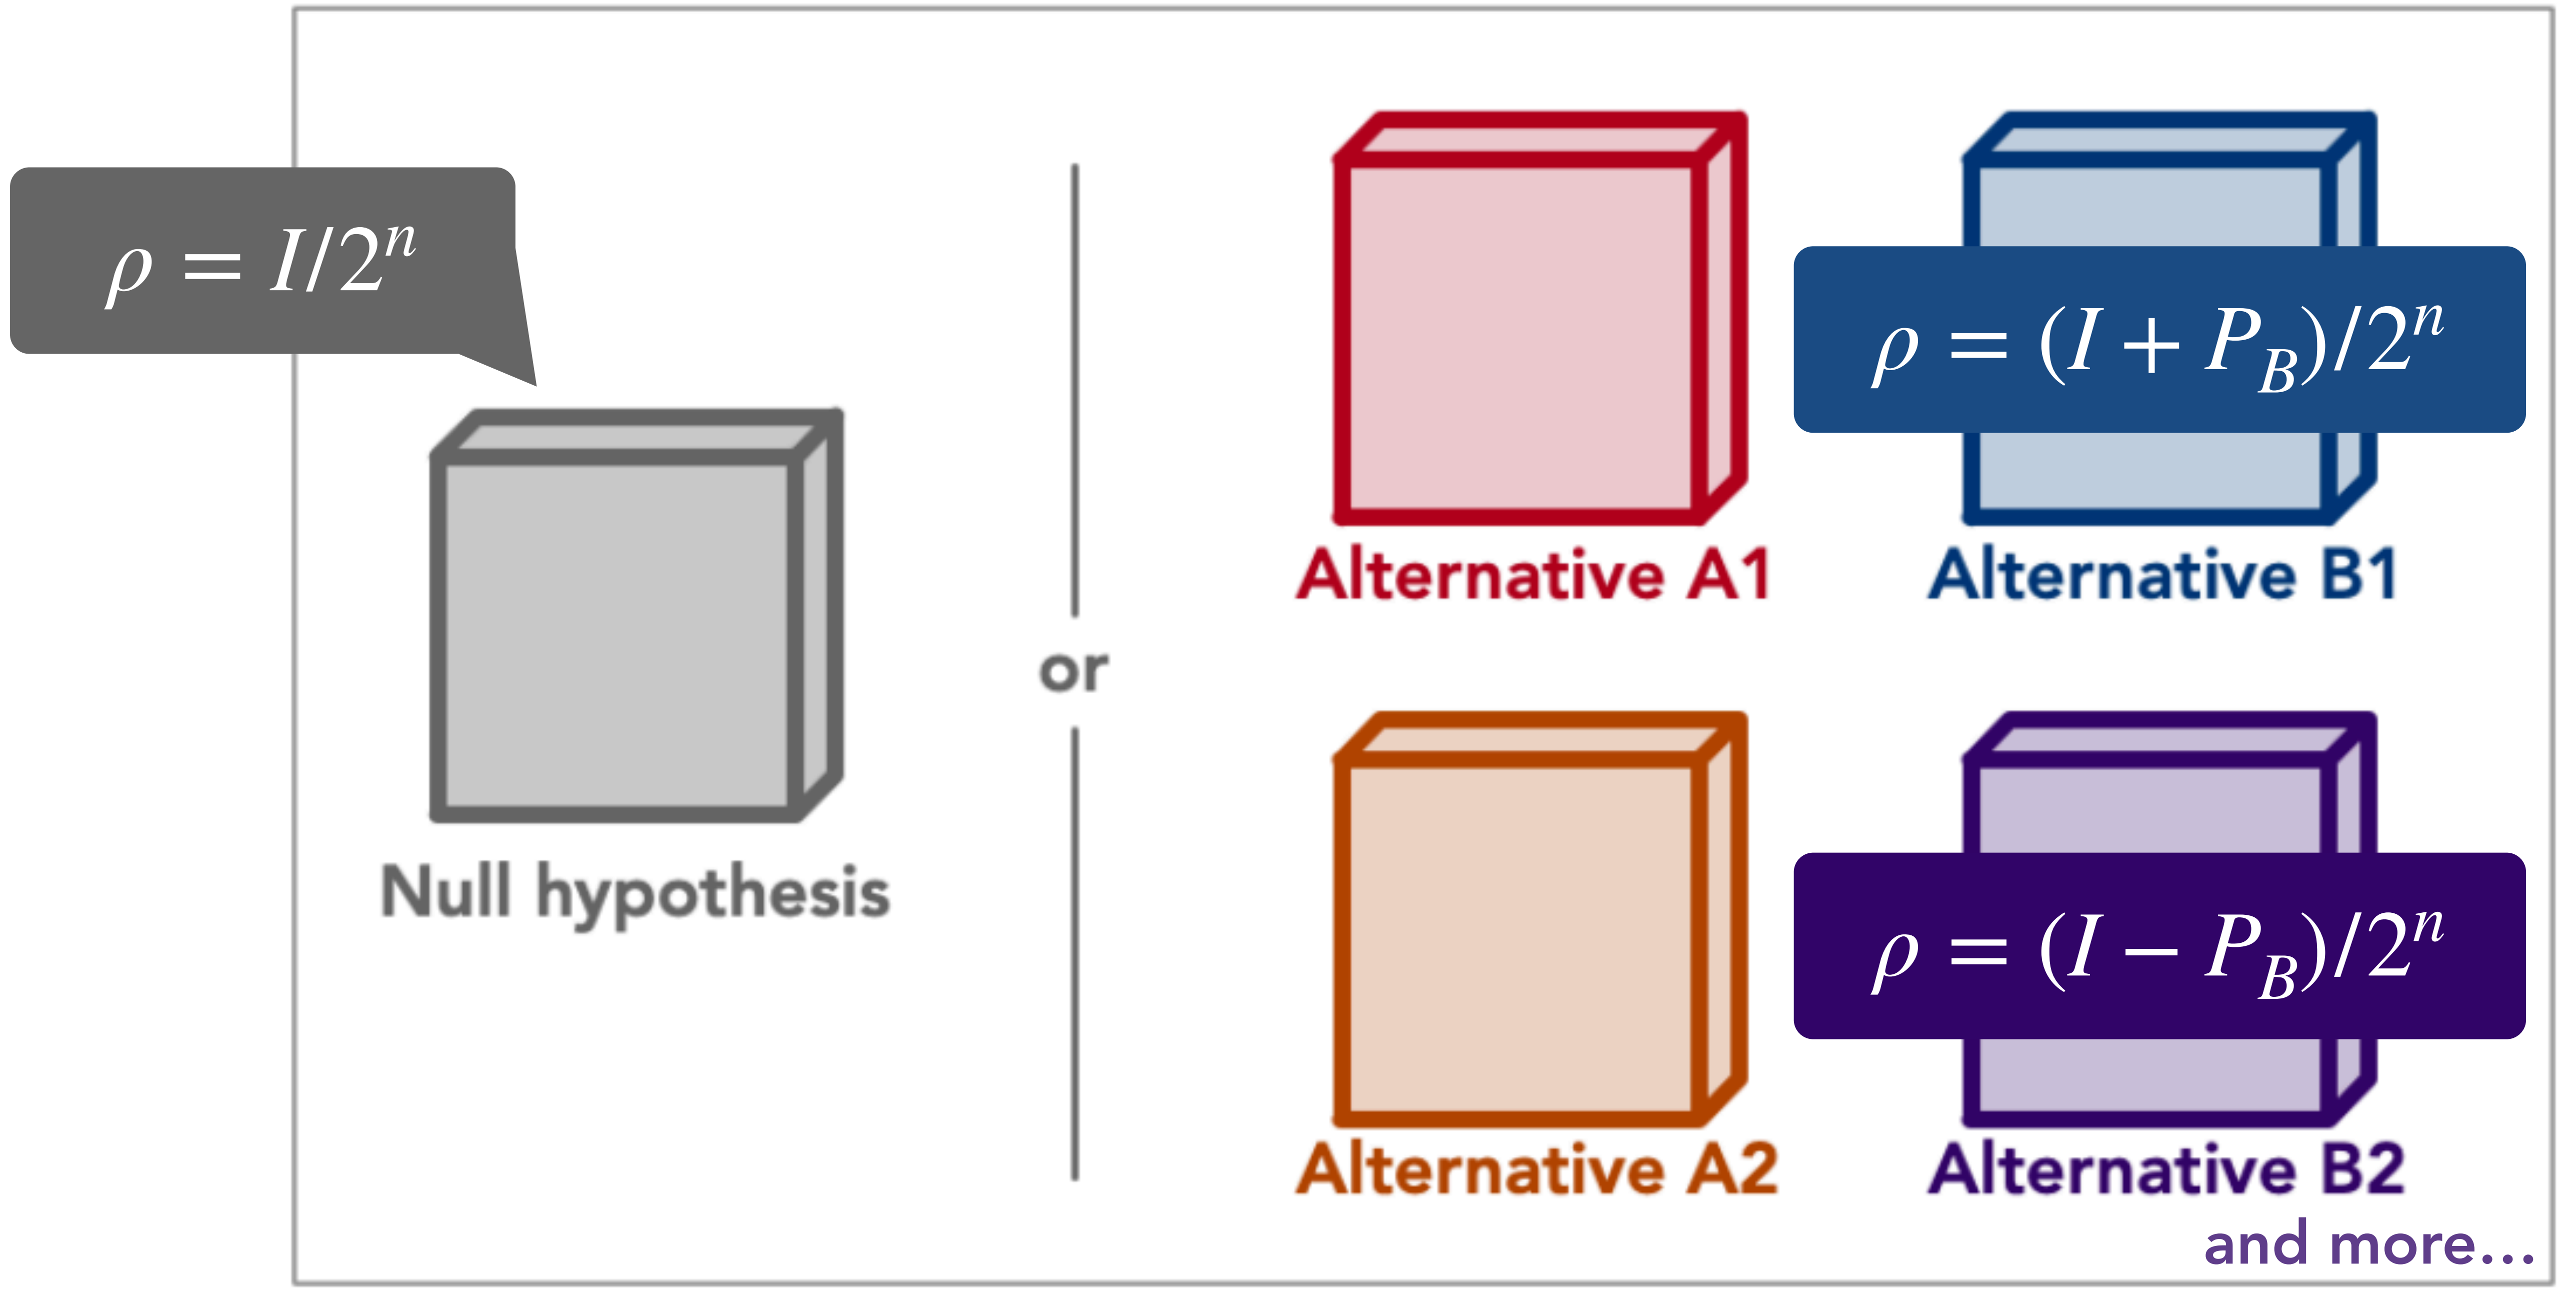

Many-versus-one distinguishing task

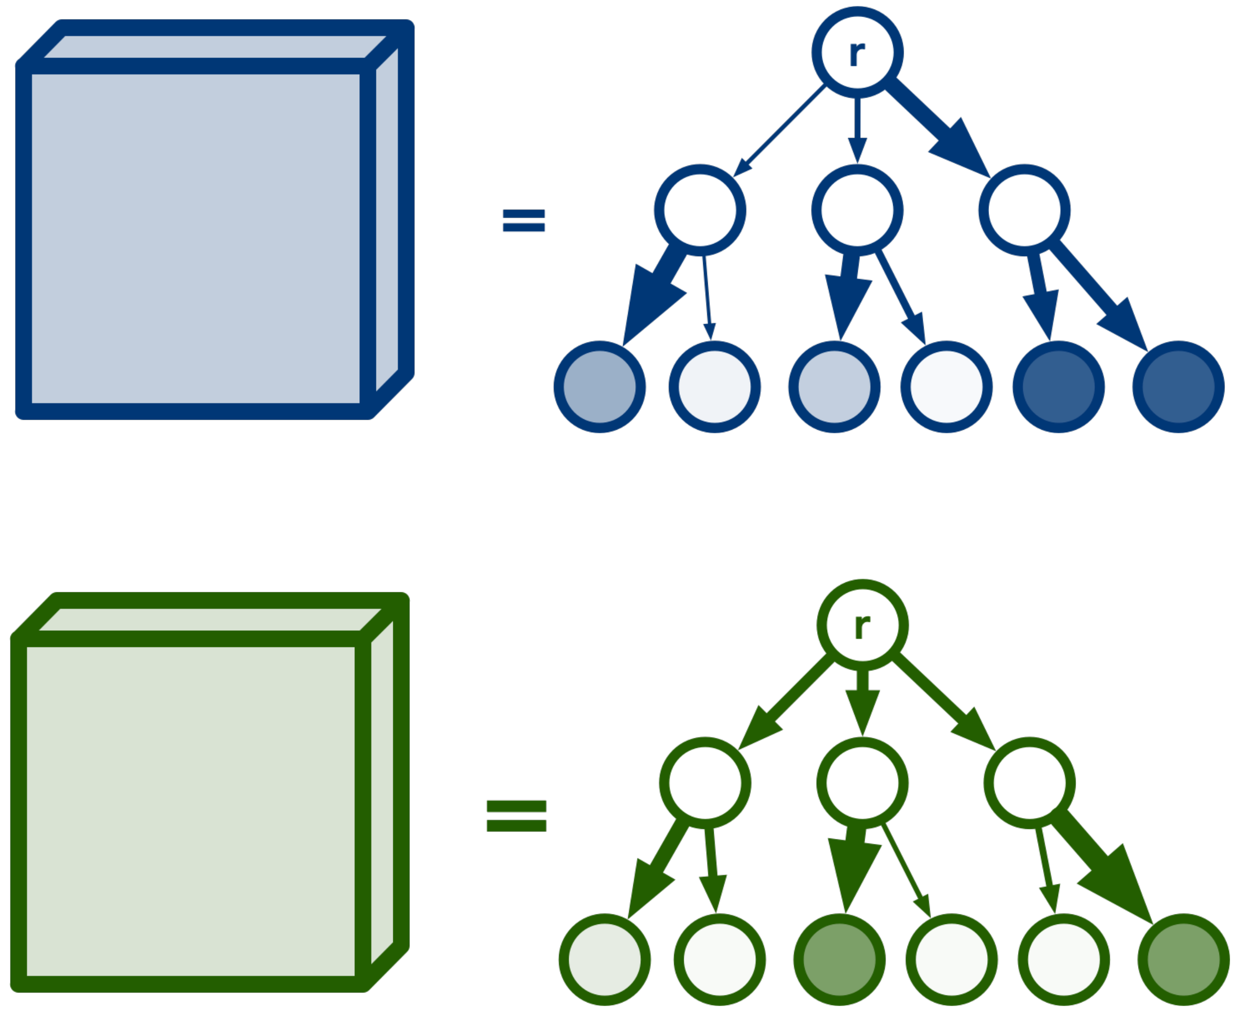

Information-theoretic bound

- Succeeding in the many-vs-one distinguishing task implies a lower bound on the TV in the probability over leaf nodes.
- We then upper bound TV with a function of the number of experiments (i.e., samples).

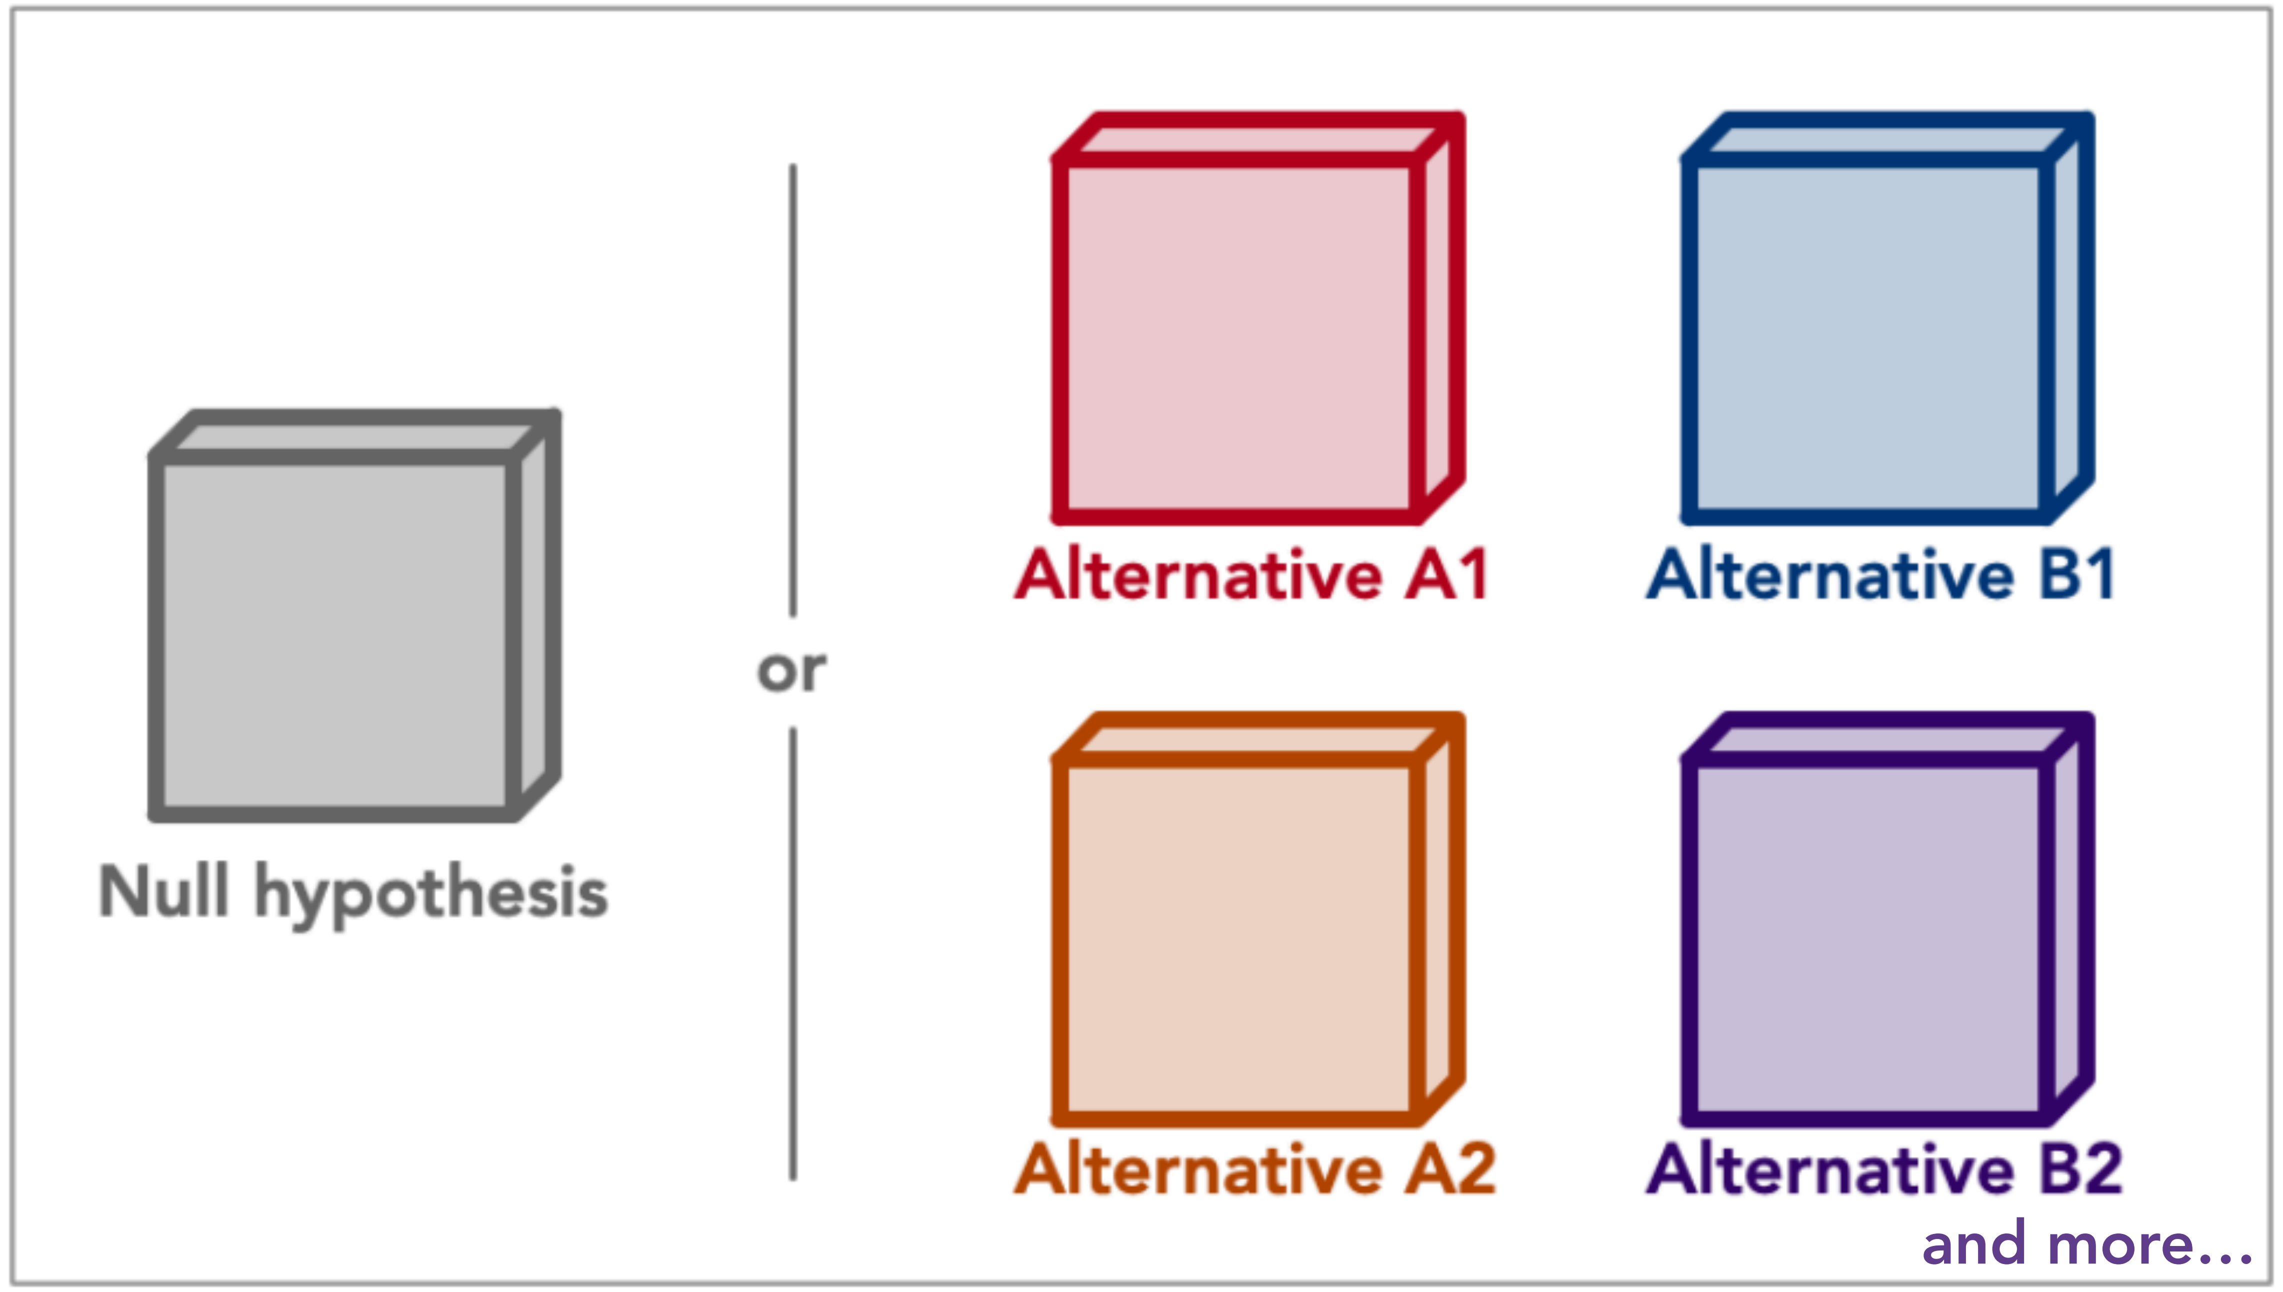

Many-versus-one distinguishing task

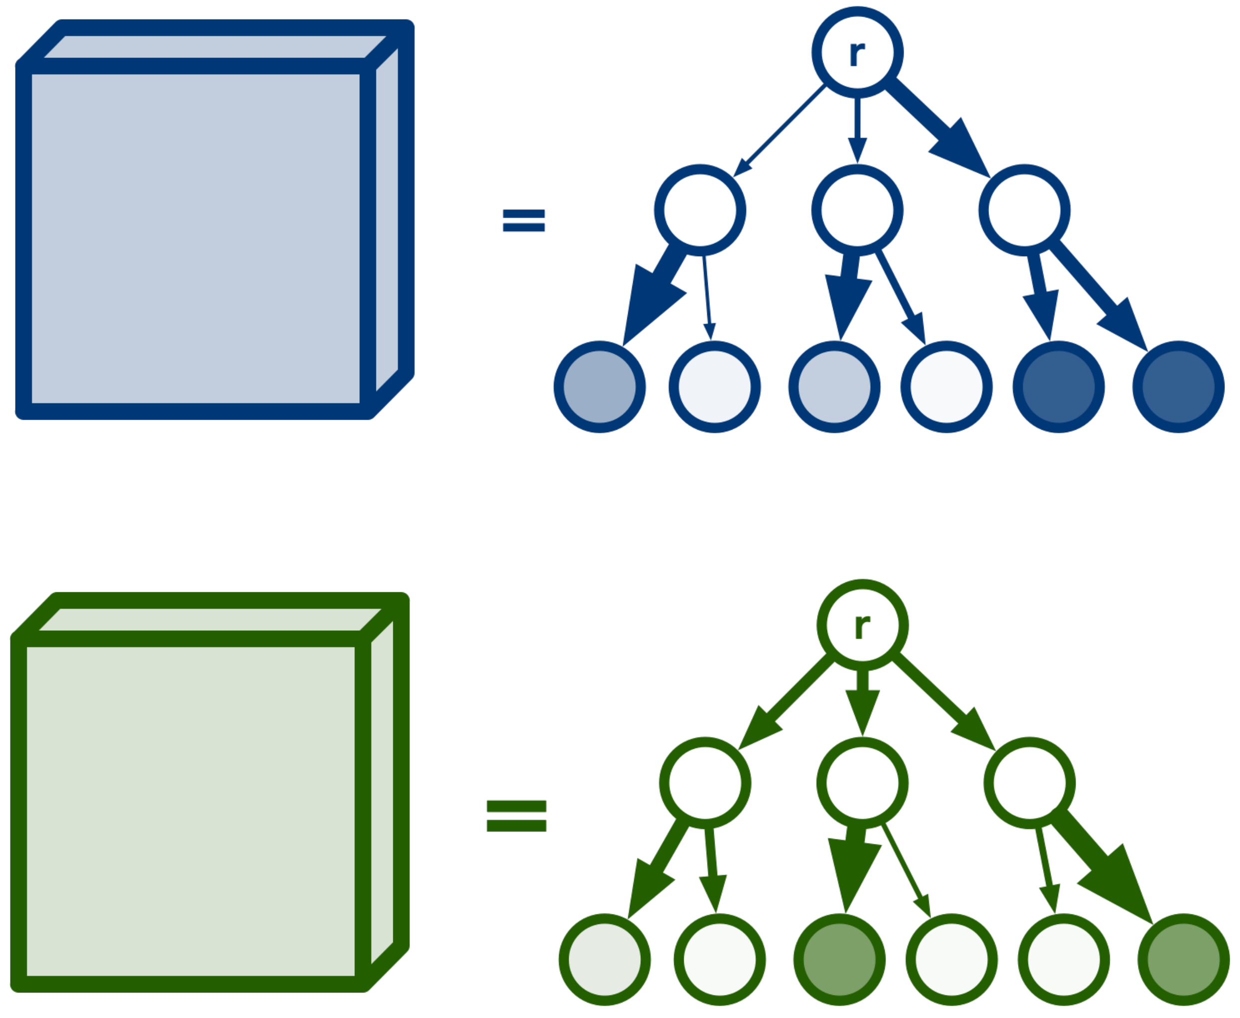

Quantum advantage in predicting Pauli observables

- 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.

Quantum advantage in predicting general observables

- The classical/quantum agent learns about the unknown n-qubit state ρ.
- Subsequently, the agent predicts $\mathrm{Tr}(O_i \rho)$ from a known set $O_1, \ldots, O_M$.

Theorem

Classical agent needs $\tilde{\Omega}(\min(M, 2^n))$ experiments to predict the M observables,
but quantum agent only needs $\mathcal{O}(n \log^2 M)$ experiments to predict the observables.

Quantum agent uses the truly-quantum shadow tomography [Badescu, O'Donnel]:
"Online learning" + "Quantum threshold search."

Exponential quantum advantage

Predicting many incompatible observables

To predict all Pauli observables $\{I, X, Y, Z\}^{\otimes n}$,
classical agent needs $\Omega(2^n)$ experiments,
quantum agent only needs $\mathcal{O}(n)$ experiments.

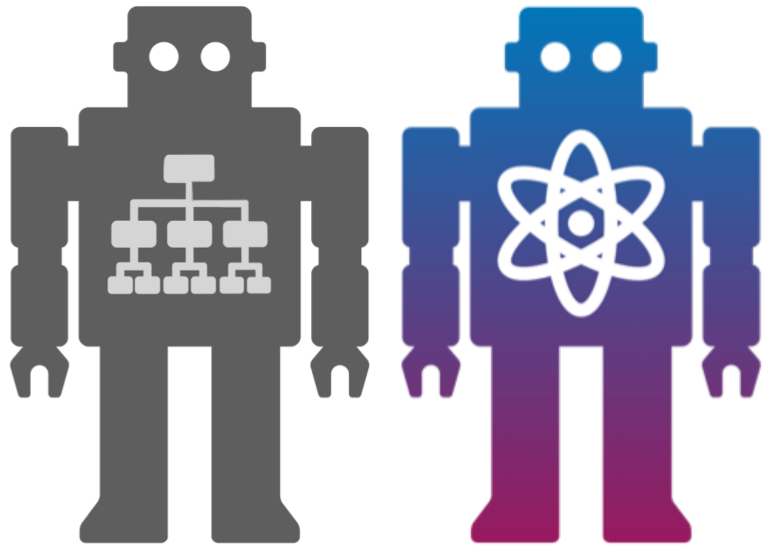

[HKP21] Huang, Kueng, Preskill. Information-theoretic bounds on quantum advantage in machine learning, *Physical Review Letters*, 2021.
[CCHL21] Chen, Cotler, Huang, Li. Exponential separations in learning with and without quantum memory, *FOCS*, 2021.
[HBC+] Huang, et all. Quantum advantage in learning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

Exponential quantum advantage

Predicting many incompatible observables

To predict all Pauli observables $\{I, X, Y, Z\}^{\otimes n}$,
classical agent needs $\Omega(2^n)$ experiments,
quantum agent only needs $\mathcal{O}(n)$ experiments.

Performing quantum PCA

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

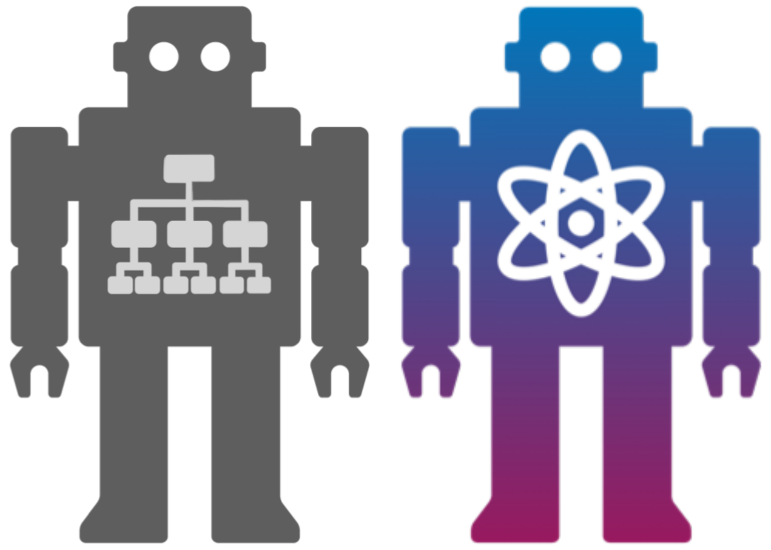

[HKP21] Huang, Kueng, Preskill. Information-theoretic bounds on quantum advantage in machine learning, *Physical Review Letters*, 2021.
[CCHL21] Chen, Cotler, Huang, Li. Exponential separations in learning with and without quantum memory, *FOCS*, 2021.
[HBC+] Huang, et all. Quantum advantage in learning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

Exponential quantum advantage

Predicting many incompatible observables

To predict all Pauli observables $\{I, X, Y, Z\}^{\otimes n}$, classical agent needs $\Omega(2^n)$ experiments, quantum agent only needs $\mathcal{O}(n)$ experiments.

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 gives exponential separations

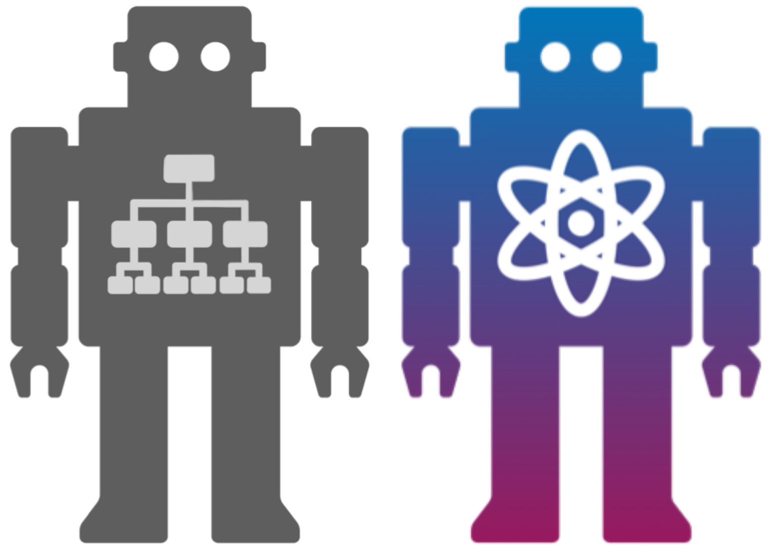

[HKP21] Huang, Kueng, Preskill. Information-theoretic bounds on quantum advantage in machine learning, *Physical Review Letters*, 2021.
[CCHL21] Chen, Cotler, Huang, Li. Exponential separations in learning with and without quantum memory, *FOCS*, 2021.
[HBC+] Huang, et all. Quantum advantage in learning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

Exponential quantum advantage

Predicting many incompatible observables

To predict all Pauli observables $\{I, X, Y, Z\}^{\otimes n}$, classical agent needs $\Omega(2^n)$ experiments, quantum agent only needs $\mathcal{O}(n)$ experiments.

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 gives exponential separations

Learning physical dynamics

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

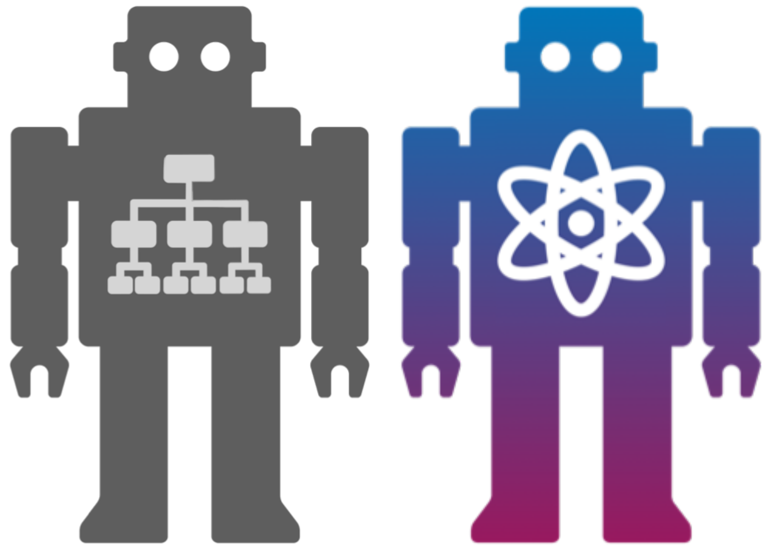

[HKP21] Huang, Kueng, Preskill. Information-theoretic bounds on quantum advantage in machine learning, *Physical Review Letters*, 2021.
[CCHL21] Chen, Cotler, Huang, Li. Exponential separations in learning with and without quantum memory, *FOCS*, 2021.
[HBC+] Huang, et all. Quantum advantage in learning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

Quantum advantage in NISQ

- Do these quantum advantages persist in noisy quantum computers?
  Yes! Rigorous analysis in [HFP22], Experiments in [HBC+].

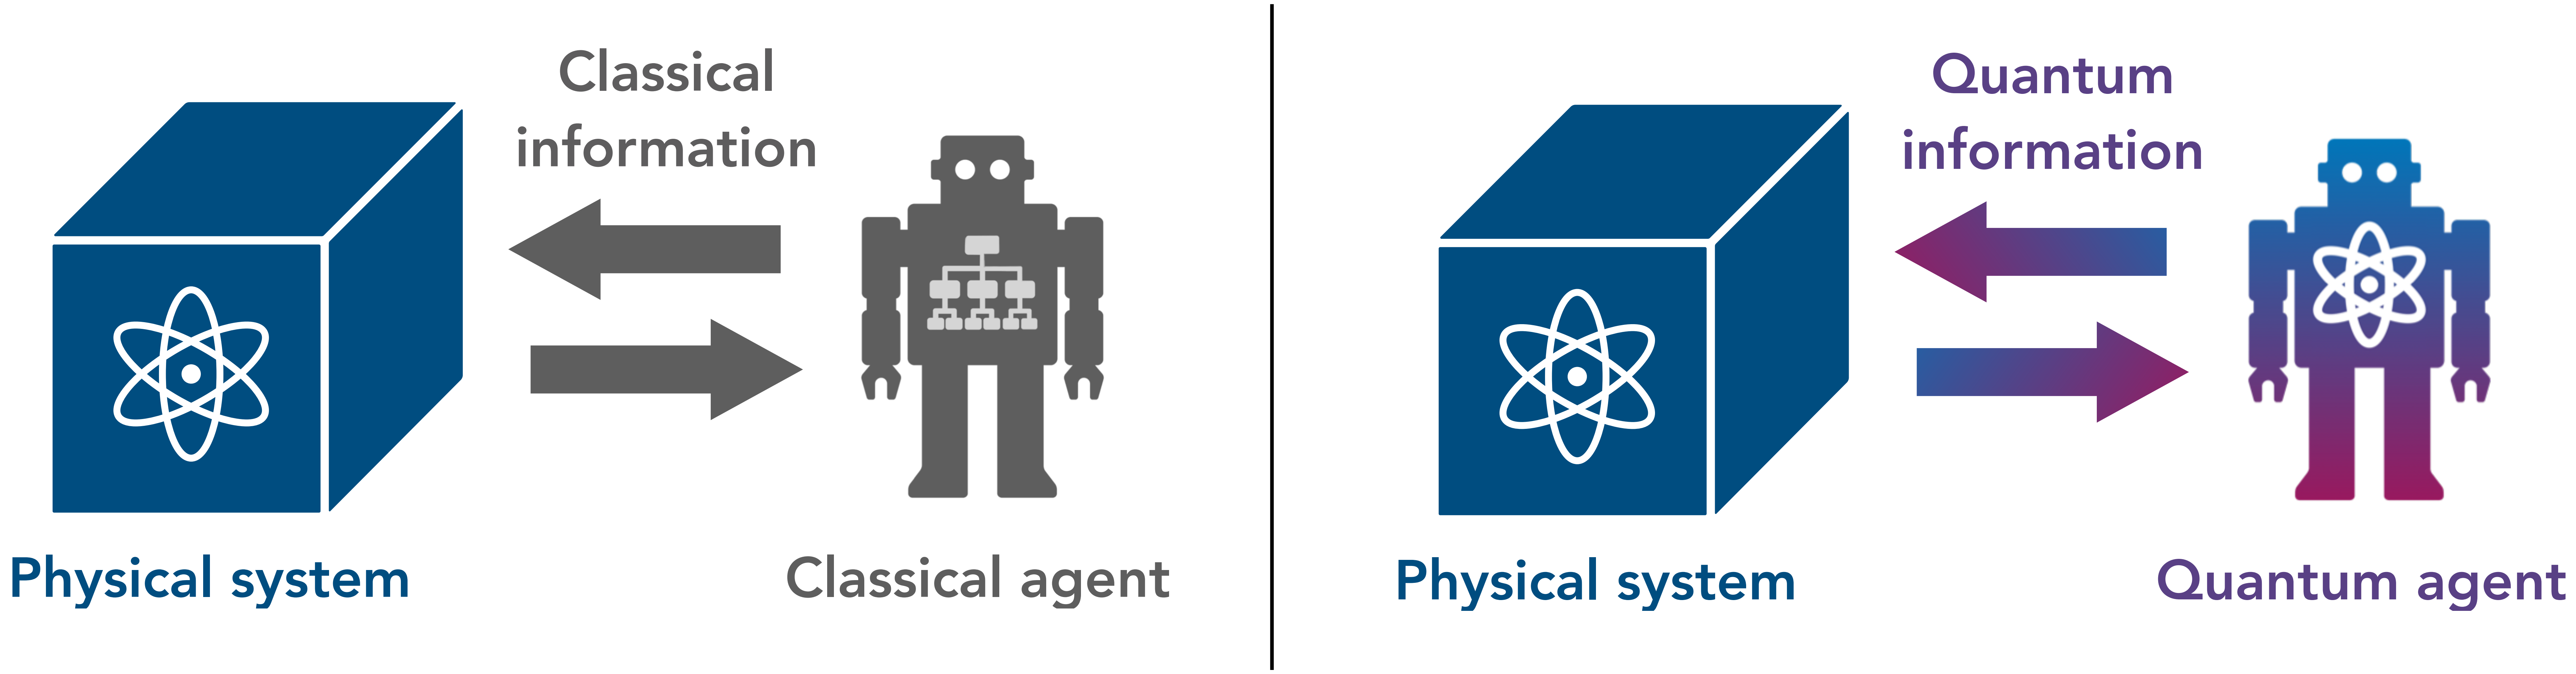

[HFP22] Huang, Flammia, Preskill. Foundations for learning from noisy quantum experiments, *QIP*, 2022.
[HBC+] Huang, Broughton, Cotler, Chen, Li, Mohseni, Neven, Babbush, Kueng, Preskill, McClean. Quantum advantage in learning from experiments, *Science*, 2022.

Demonstration on Sycamore: Quantum advantage in learning states

Utilizing a total of 40 qubits

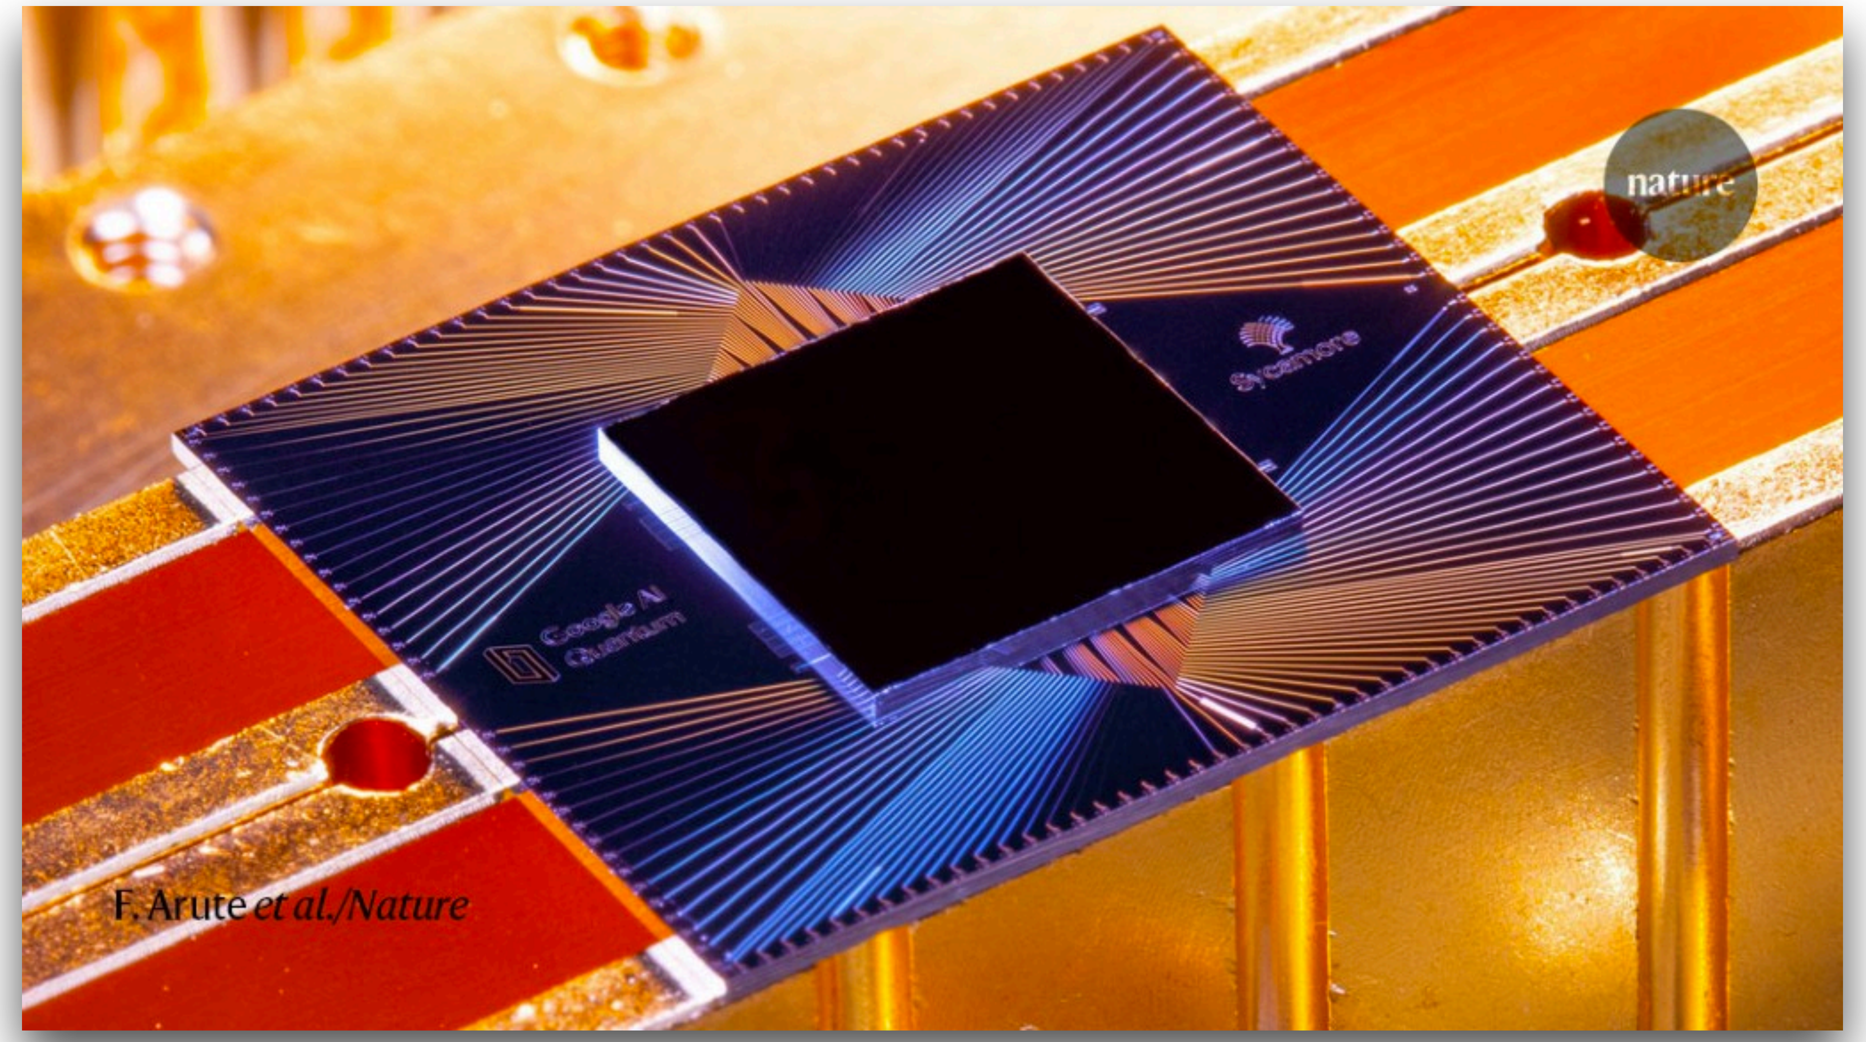

Sycamore Processor

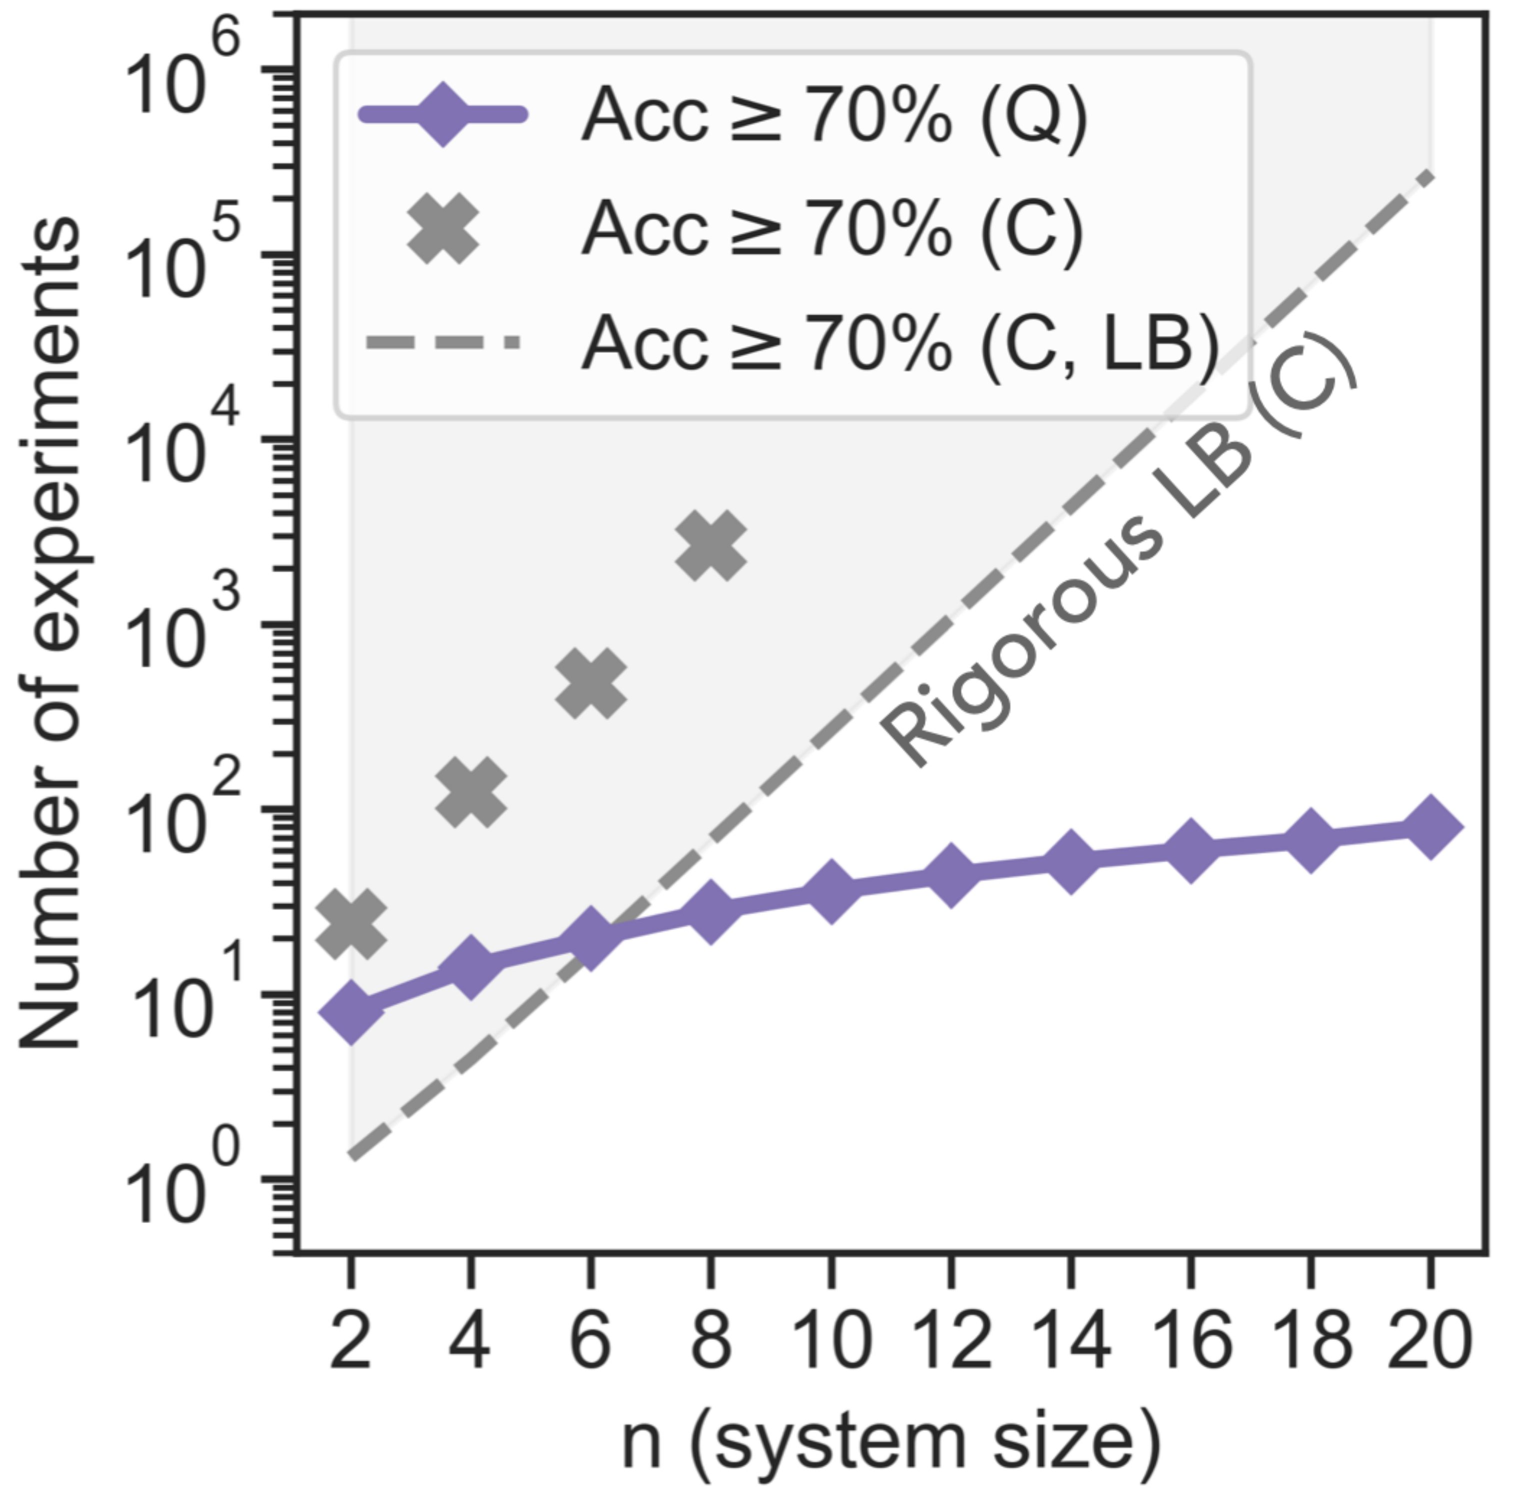

[HFP22] Huang, Flammia, Preskill. Foundations for learning from noisy quantum experiments, *QIP*, 2022.
[HBC+] Huang, Broughton, Cotler, Chen, Li, Mohseni, Neven, Babbush, Kueng, Preskill, McClean. Quantum advantage in learning from experiments, *Science*, 2022.

Demonstration on Sycamore: Quantum advantage in learning states

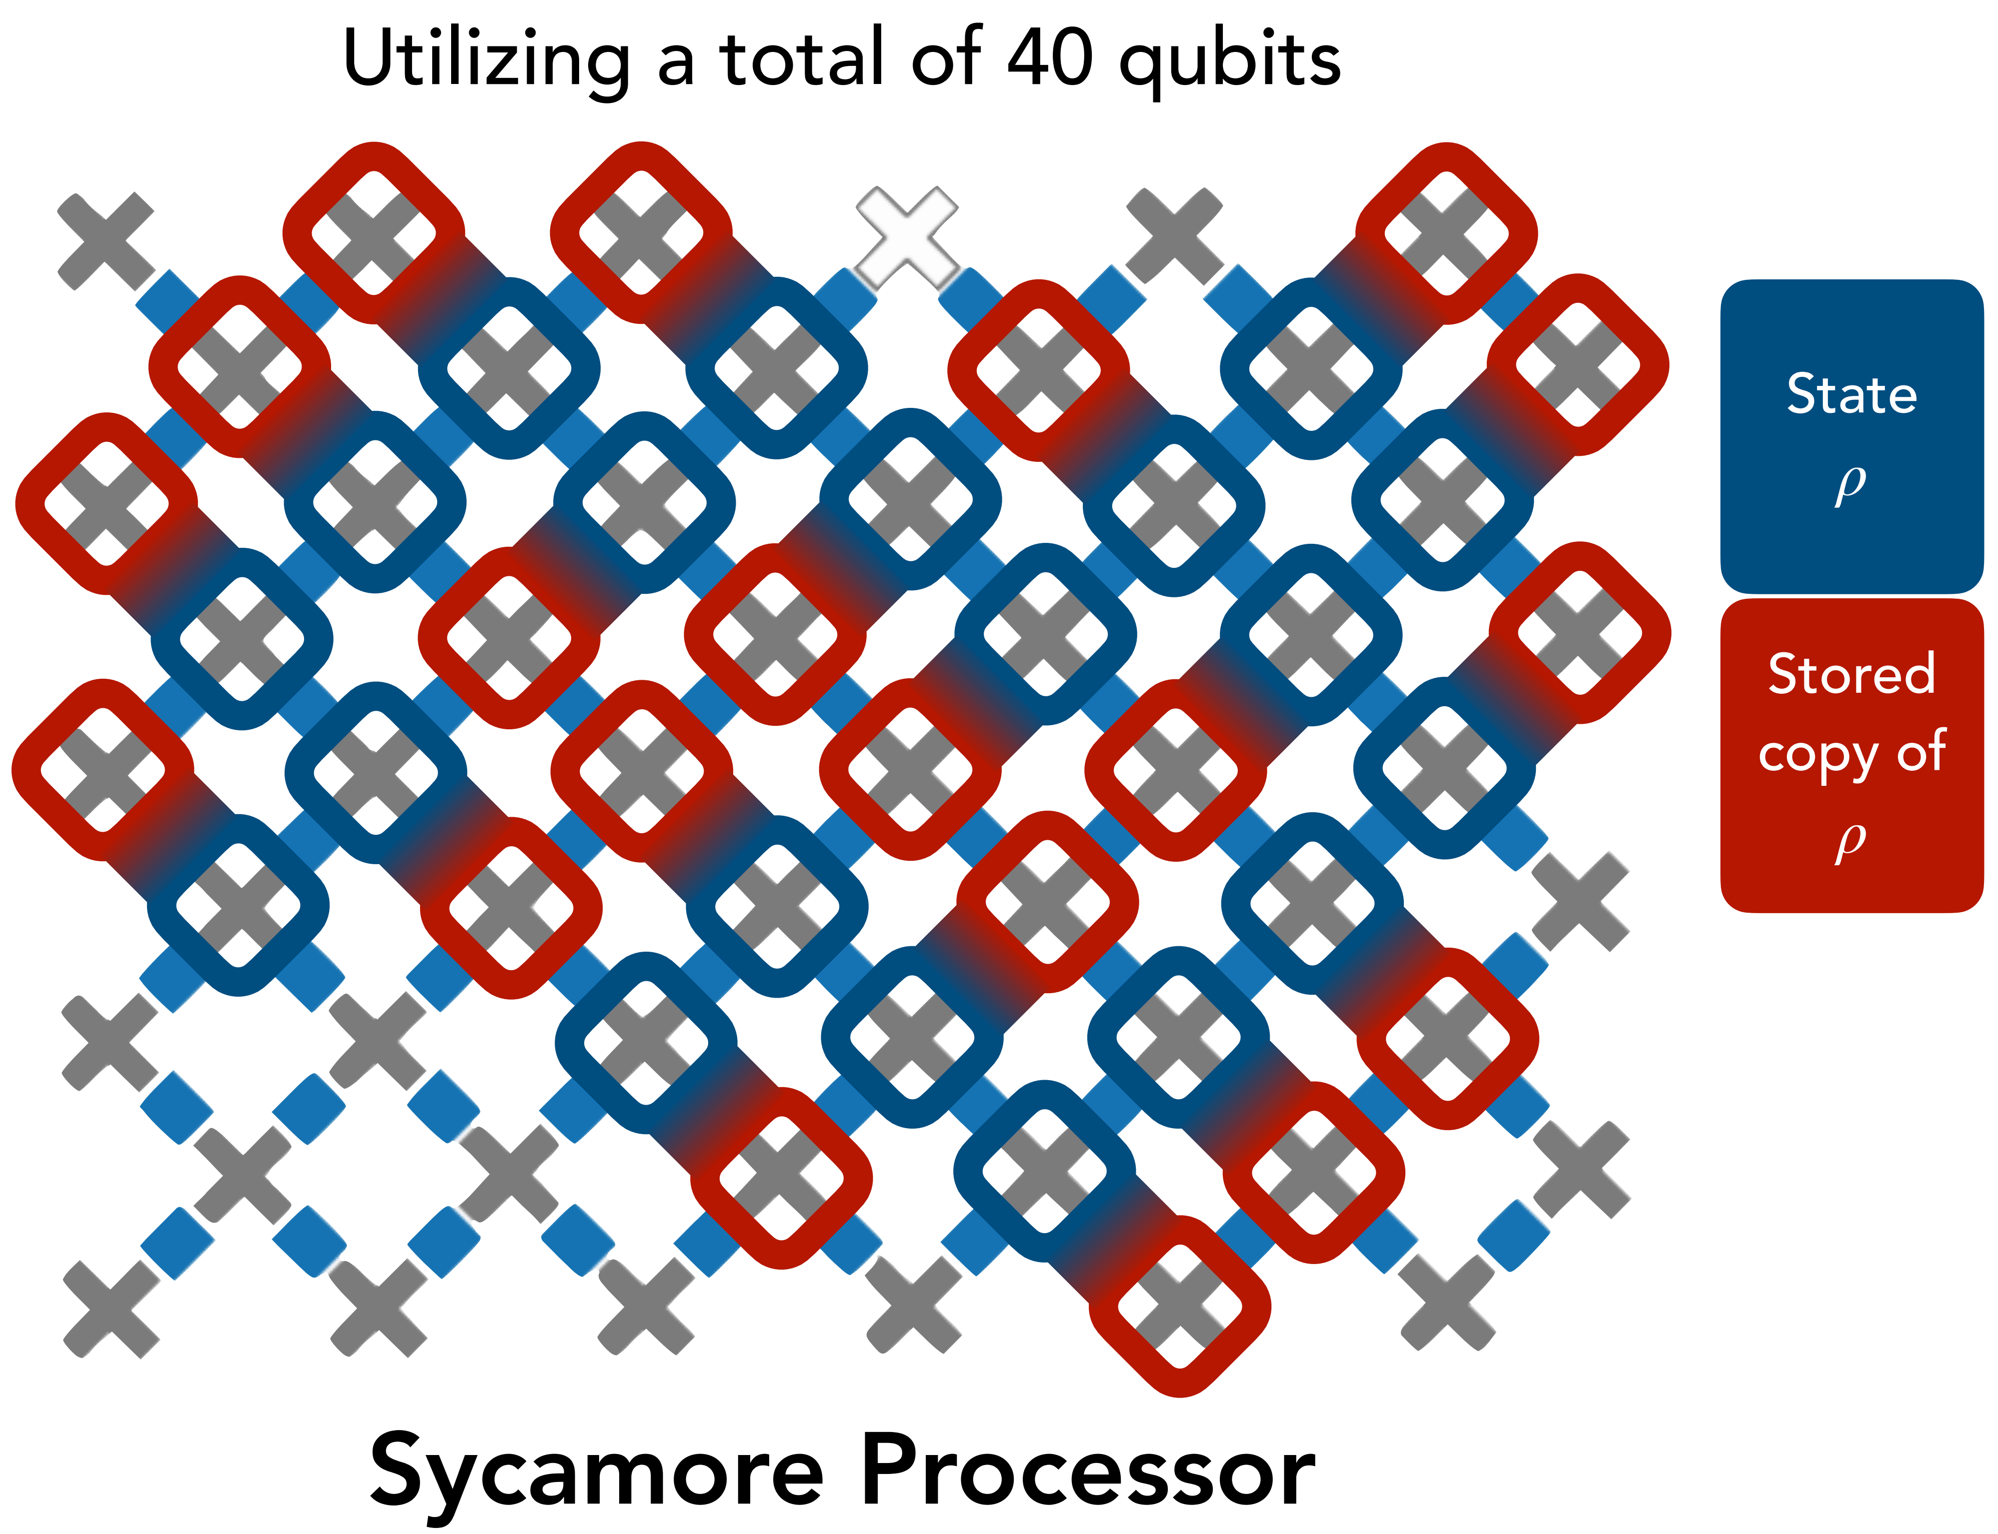

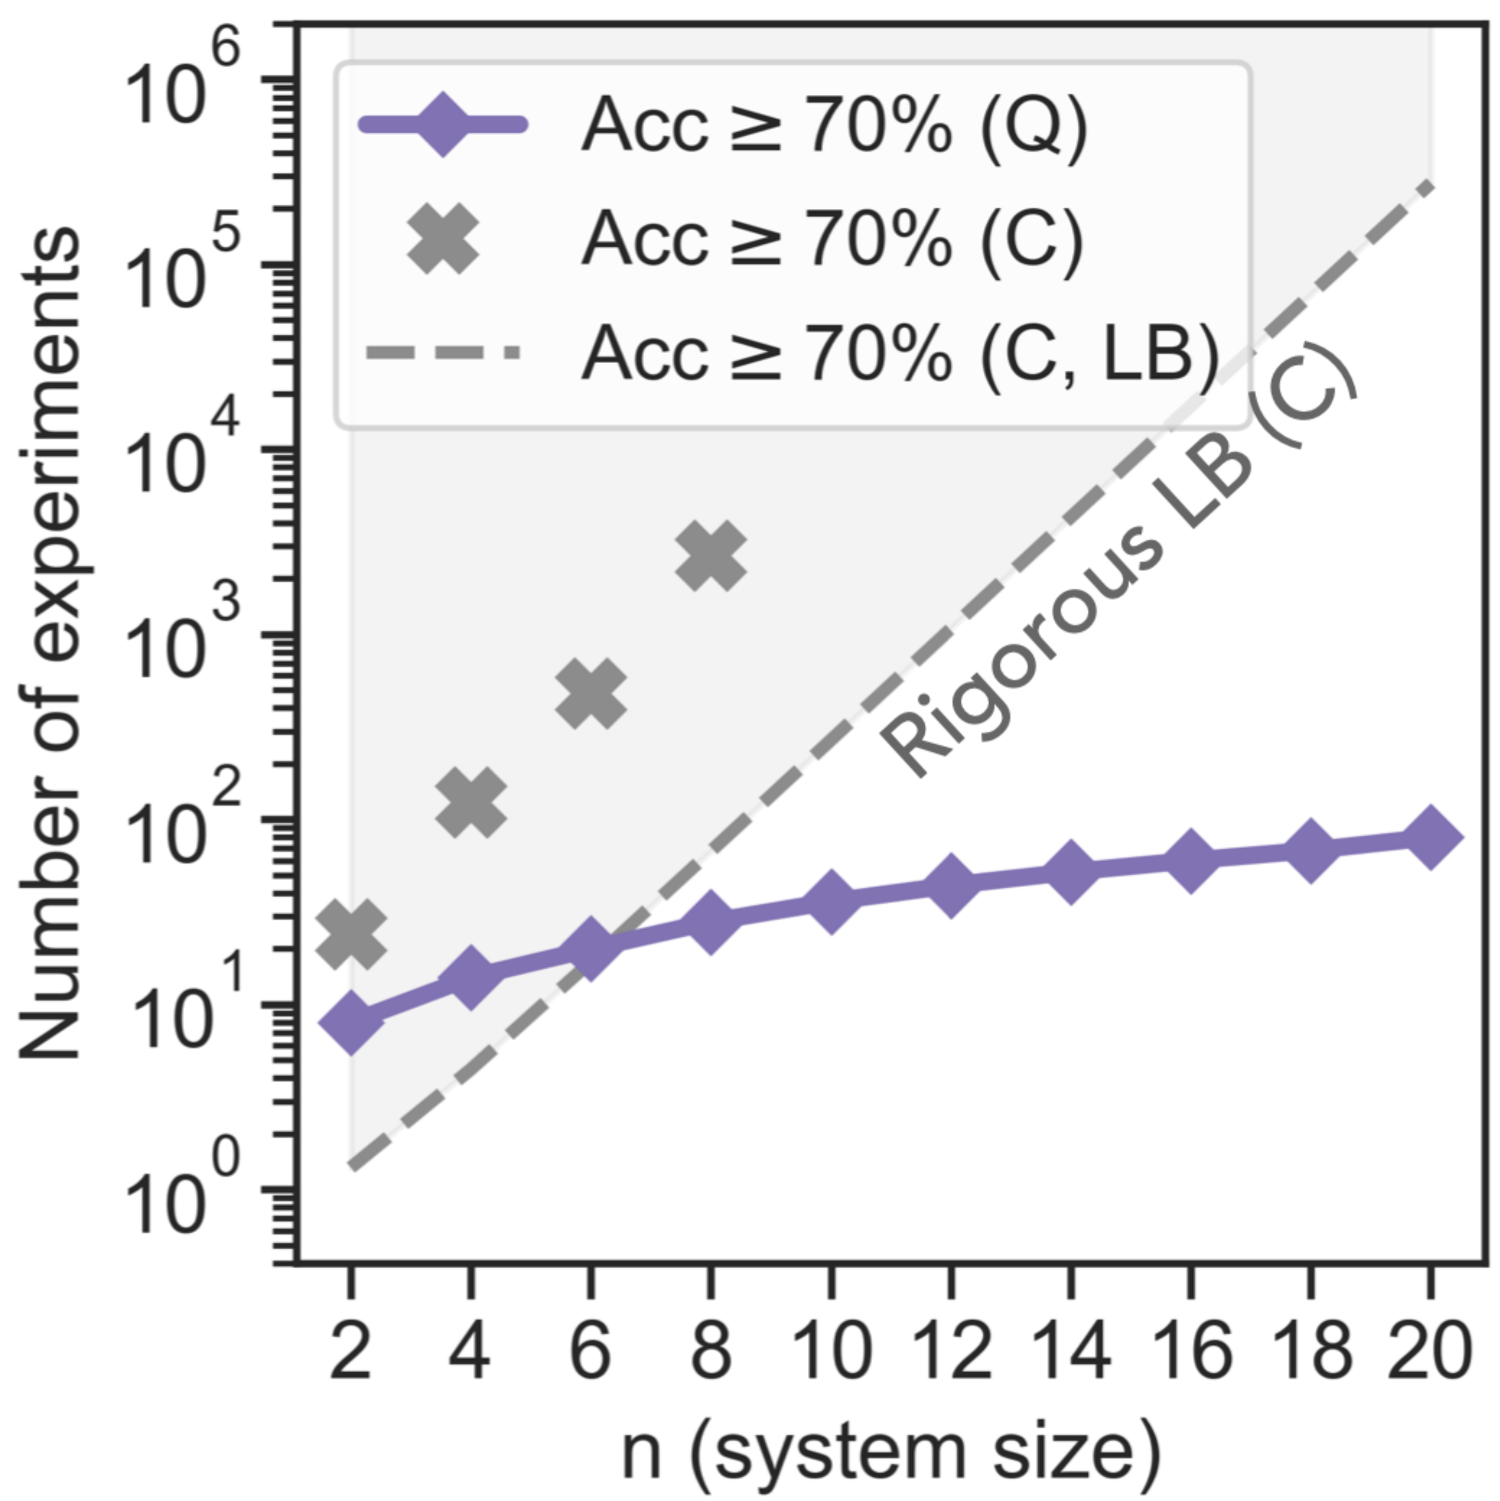

[HFP22] Huang, Flammia, Preskill. Foundations for learning from noisy quantum experiments, *QIP*, 2022.
[HBC+] Huang, Broughton, Cotler, Chen, Li, Mohseni, Neven, Babbush, Kueng, Preskill, McClean. Quantum advantage in learning from experiments, *Science*, 2022.

Demonstration on Sycamore: Quantum advantage in learning states

Conventional experiments

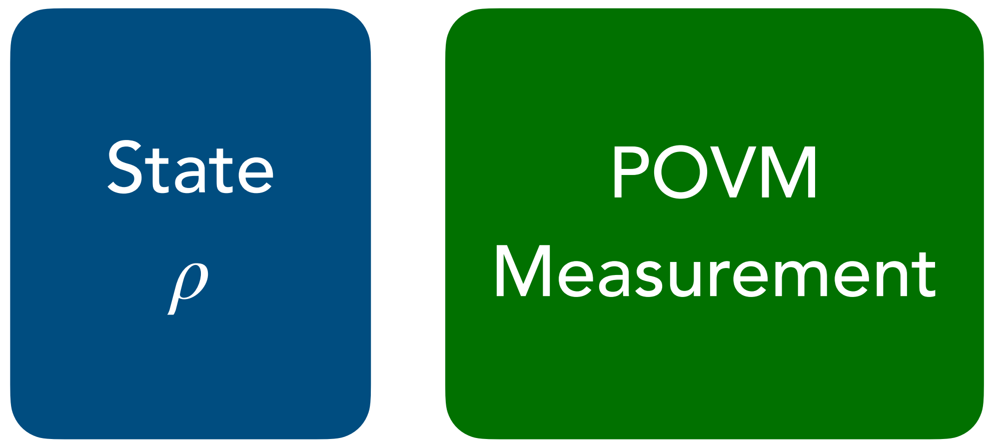

Quantum-enhanced experiments

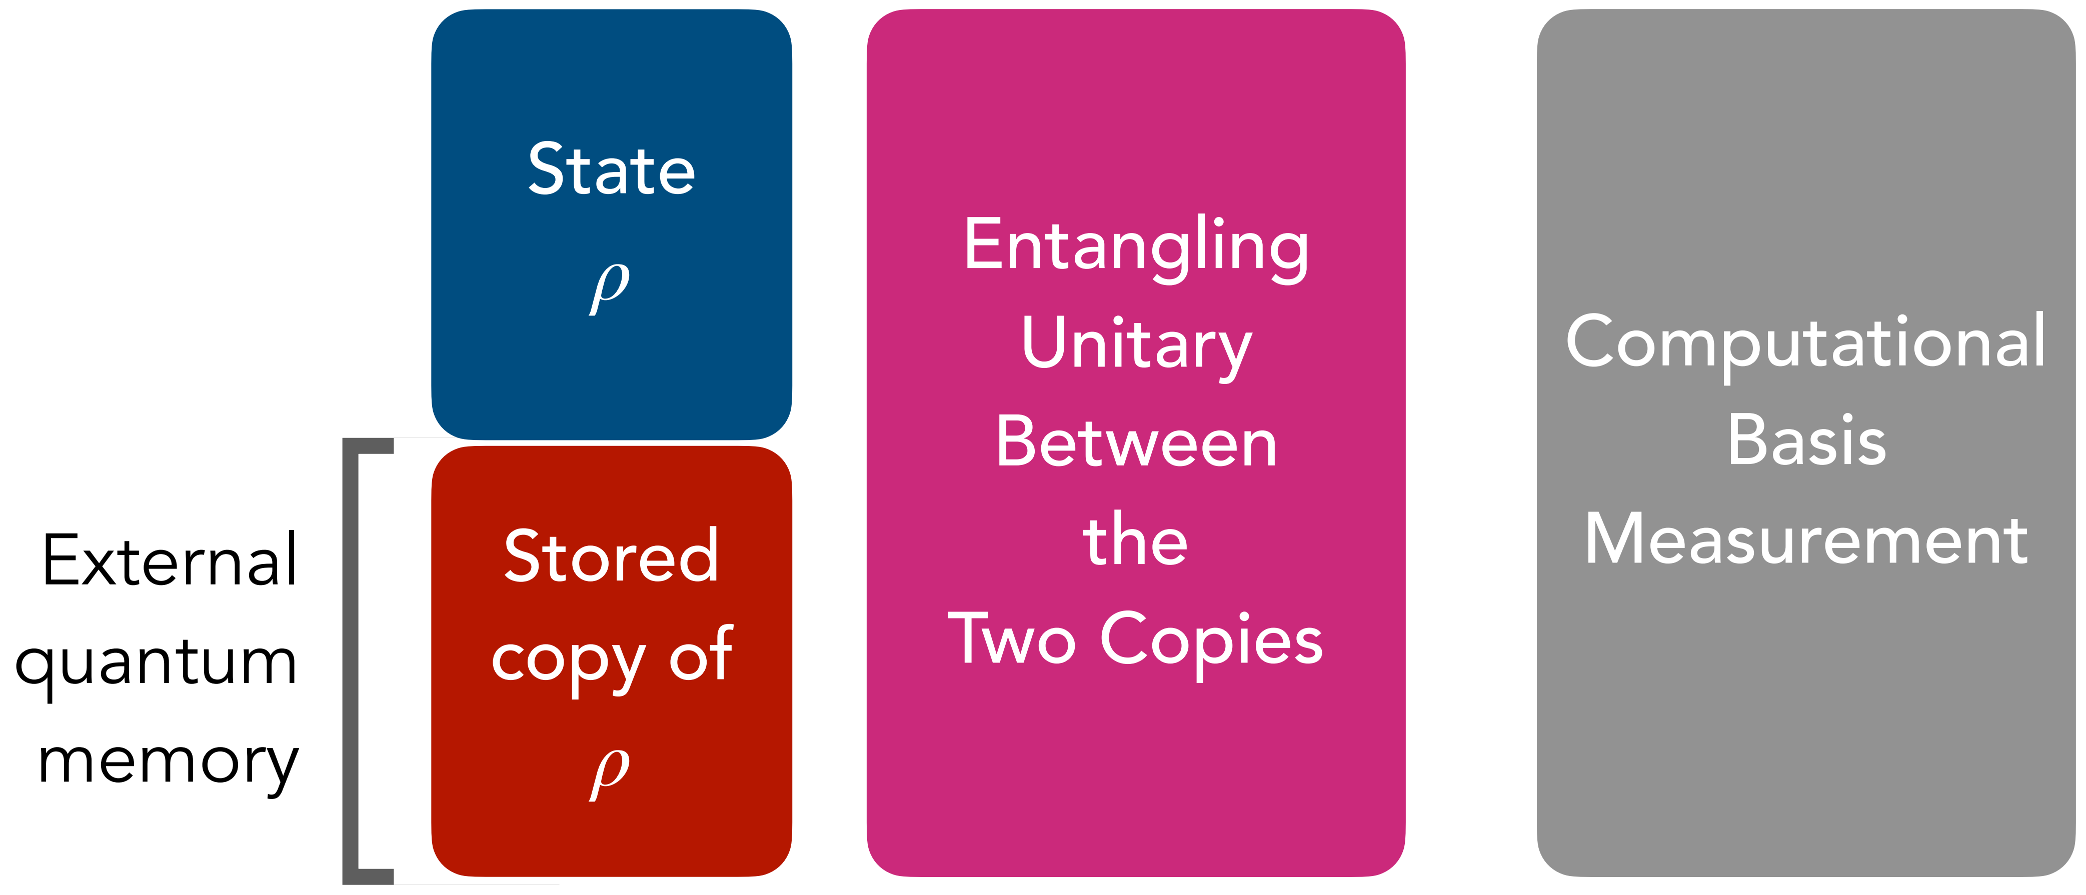

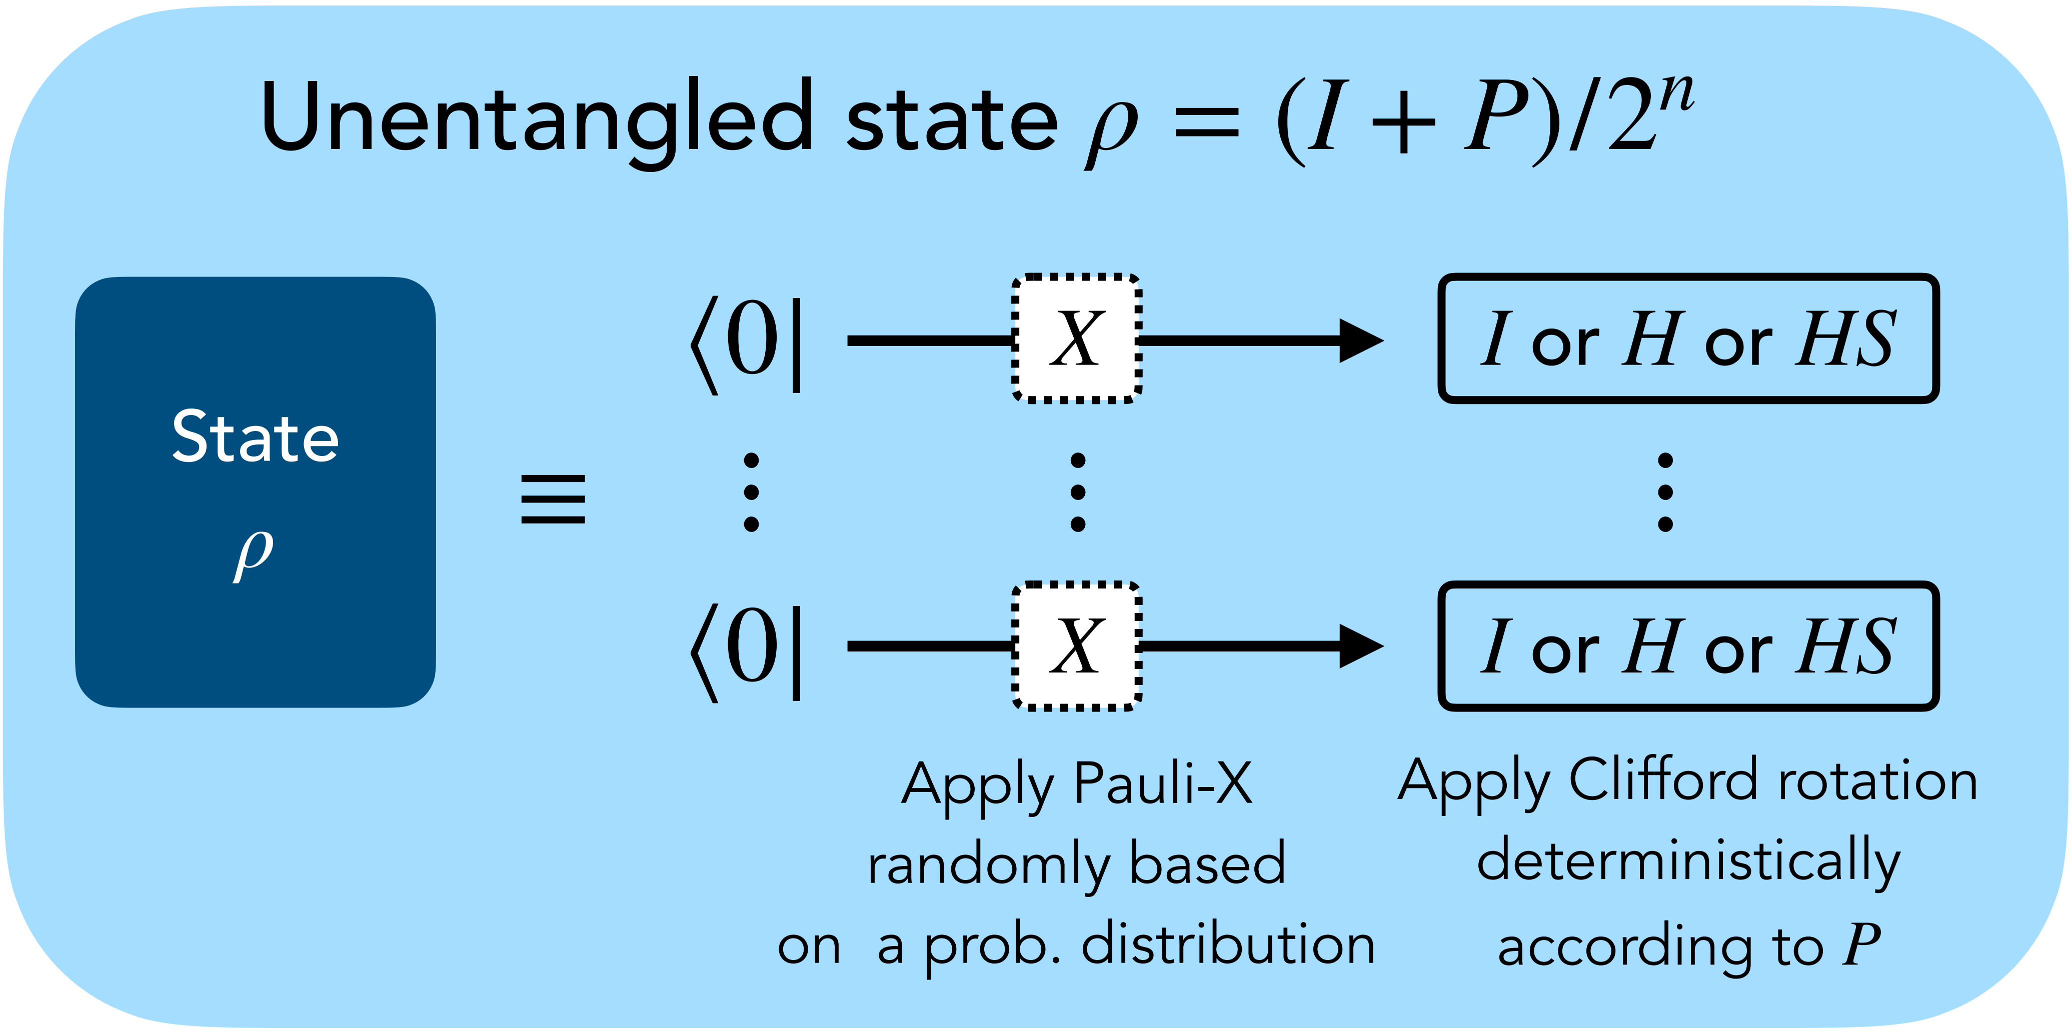

Demonstration on Sycamore:
Quantum advantage in learning states

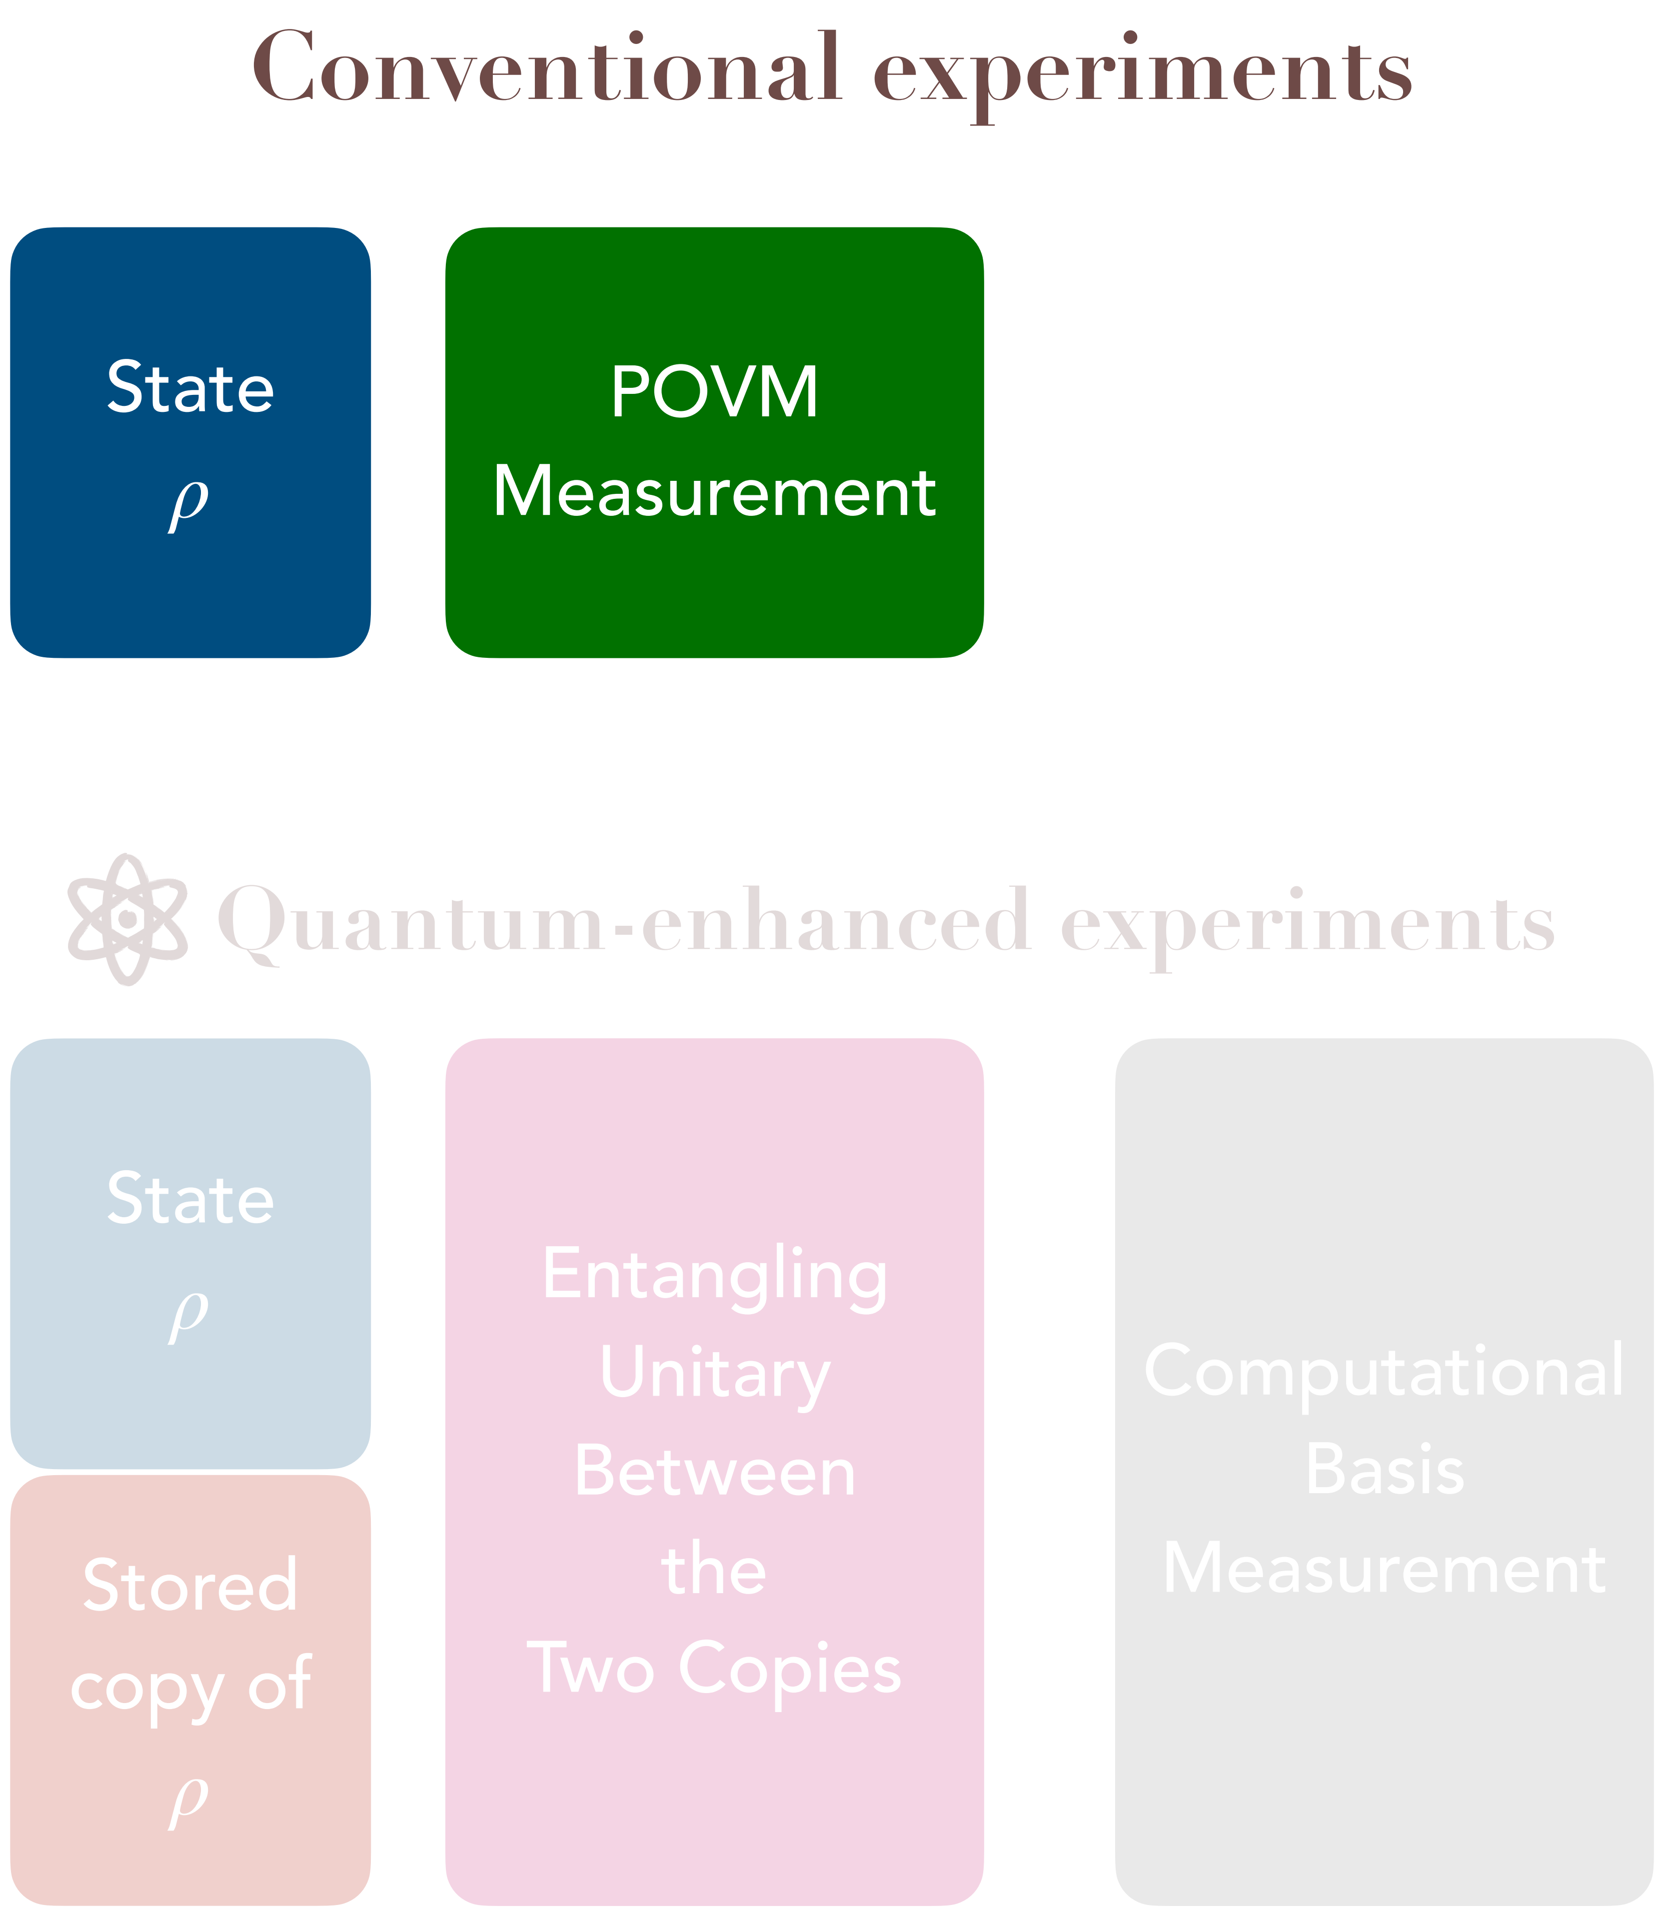

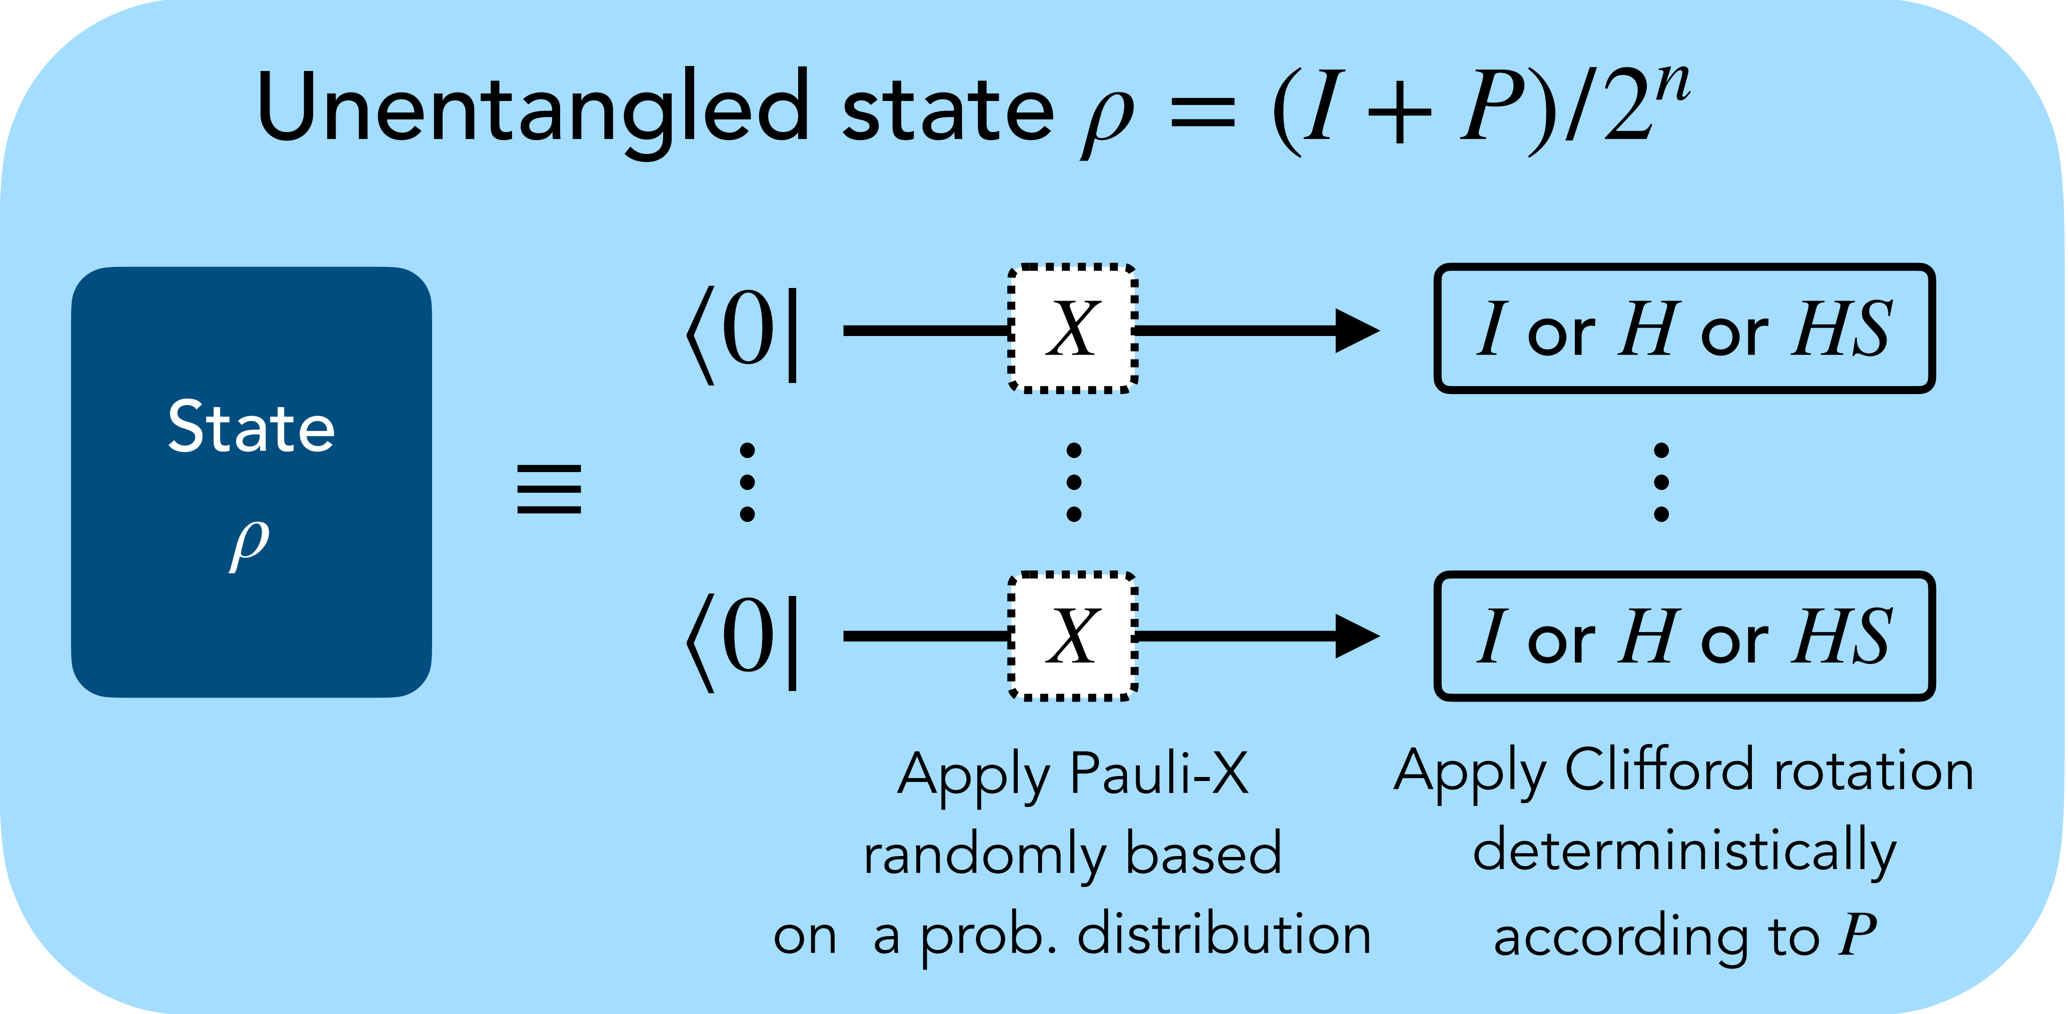

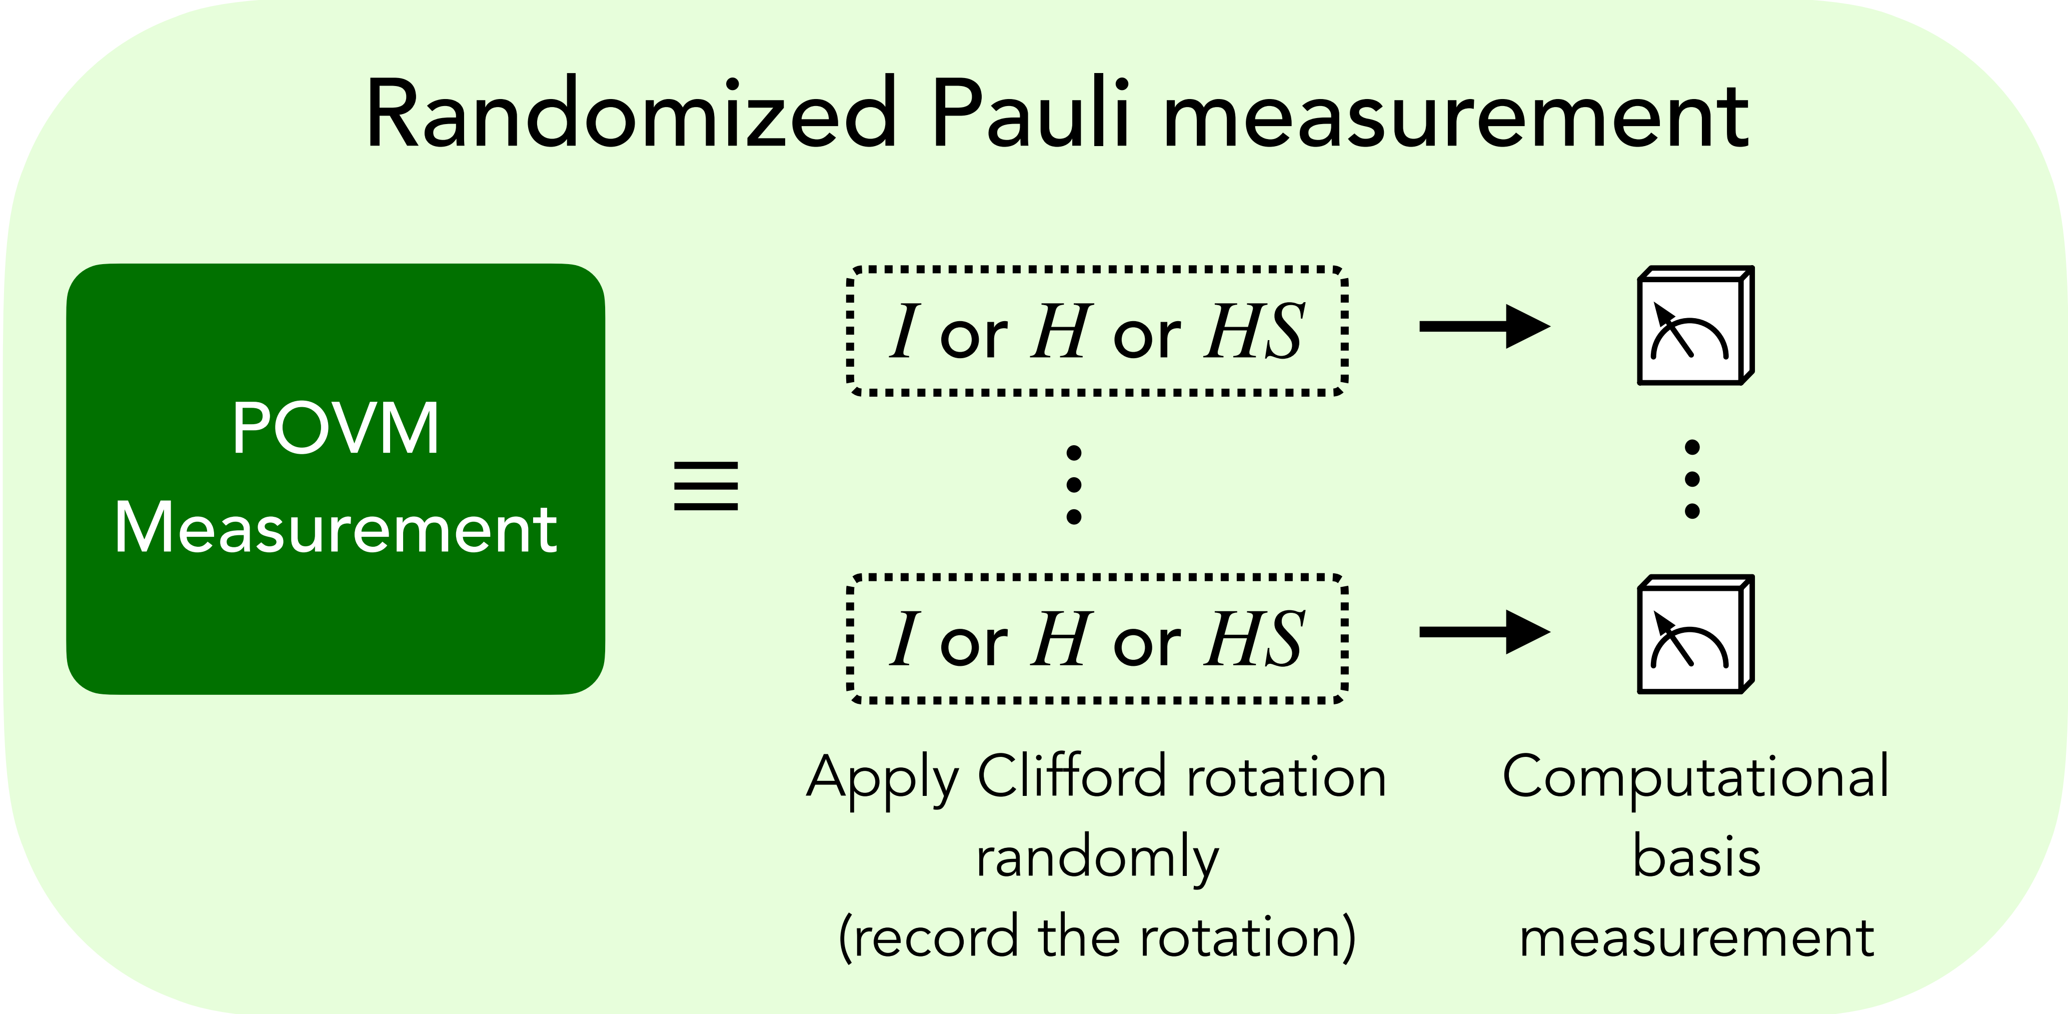

Demonstration on Sycamore:
Quantum advantage in learning states

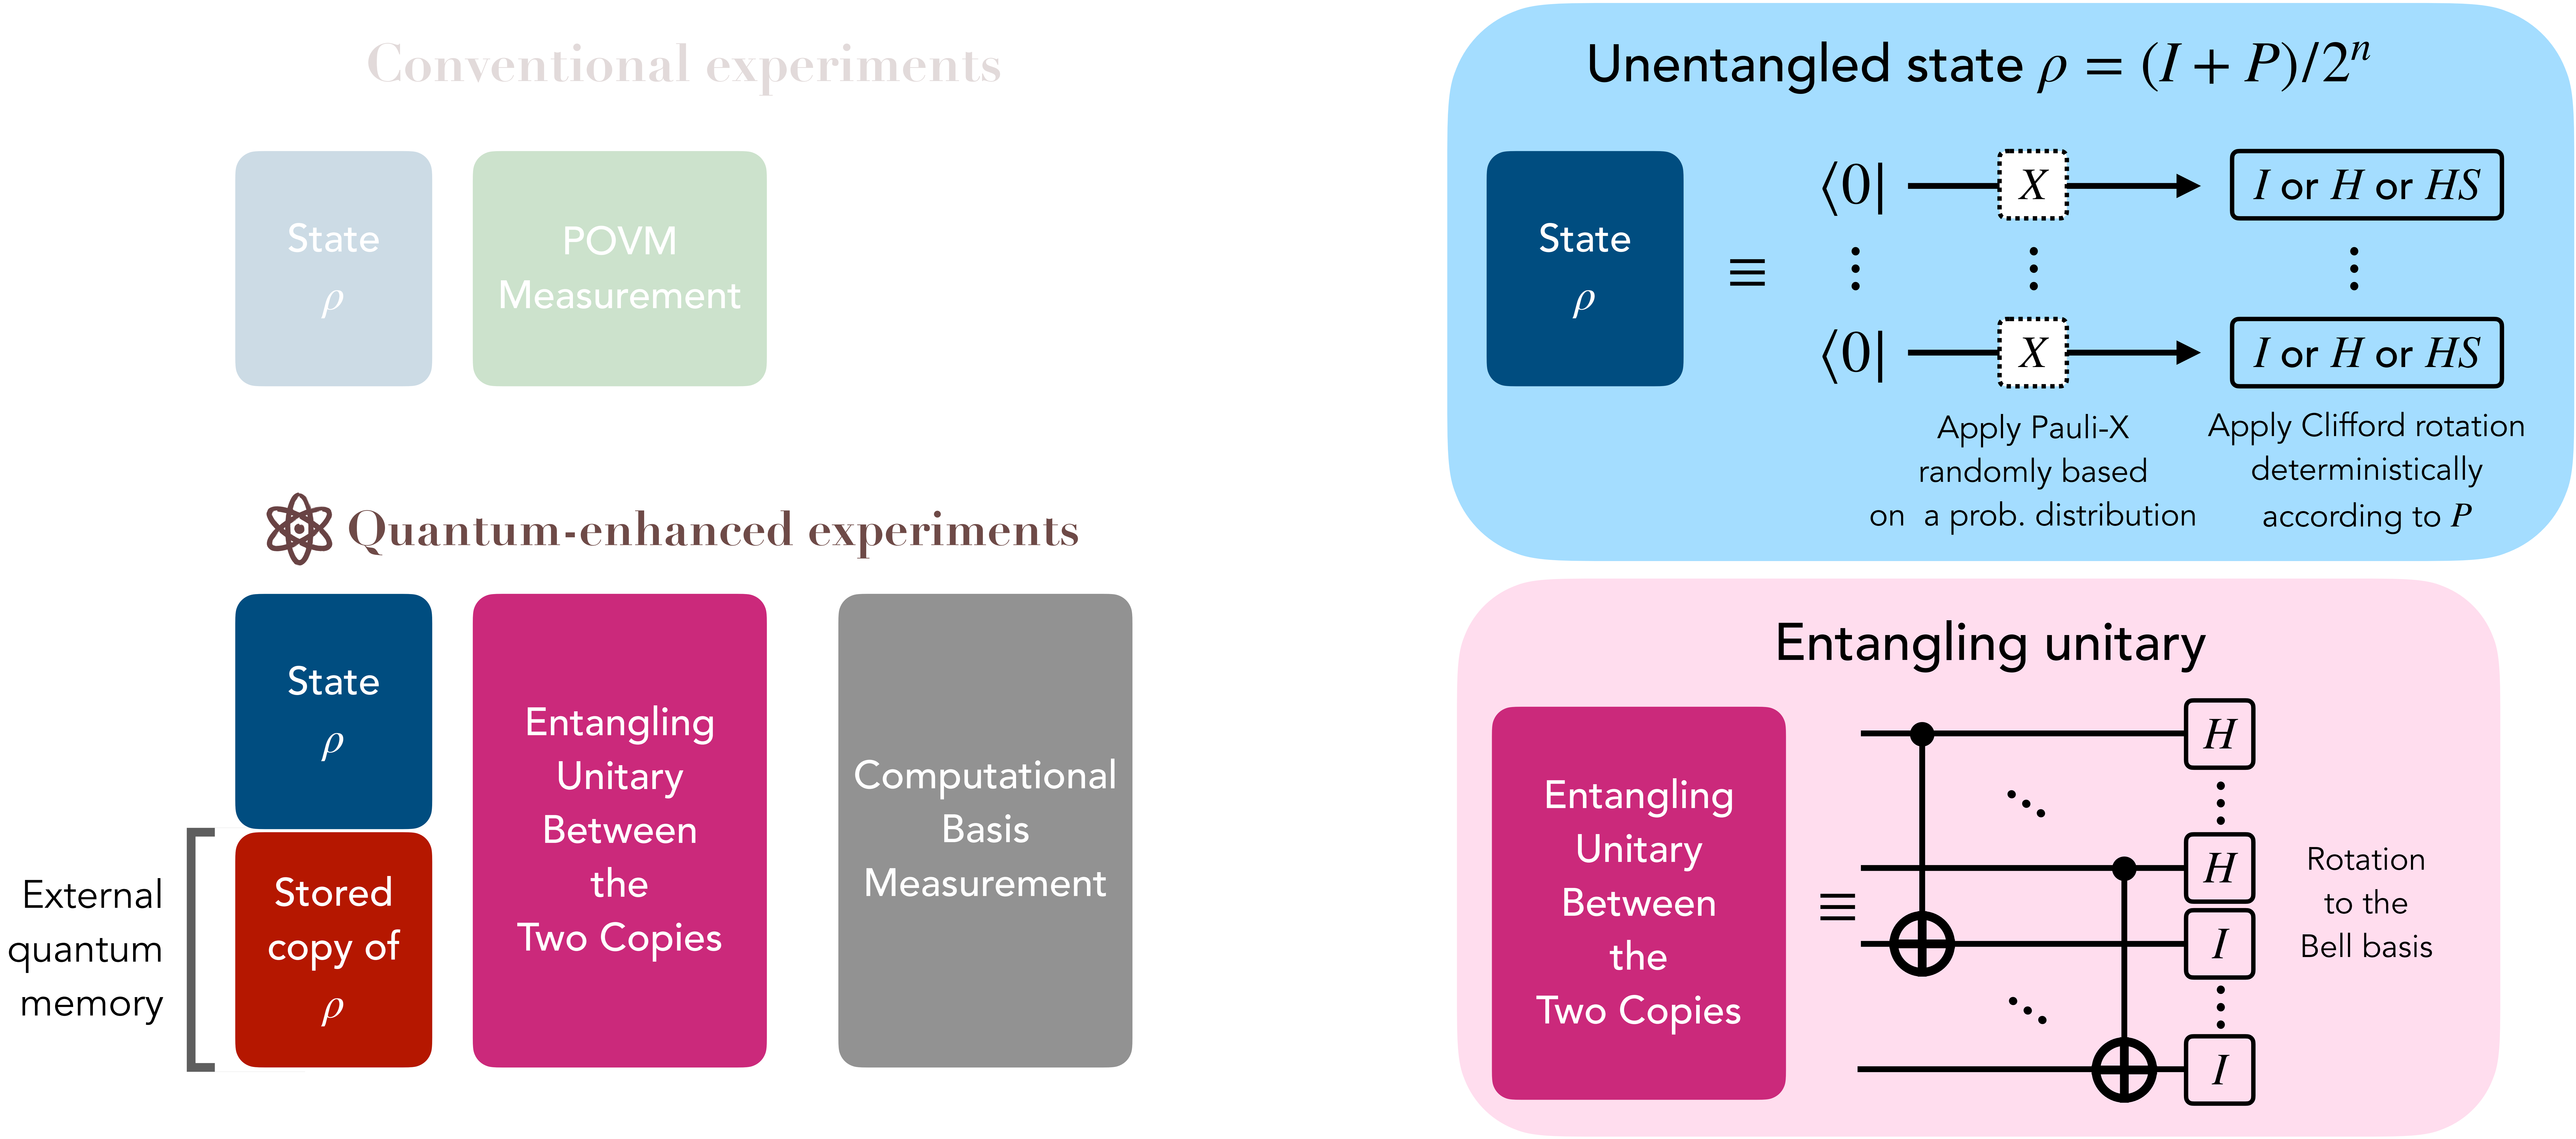

Demonstration on Sycamore: Quantum advantage in learning states

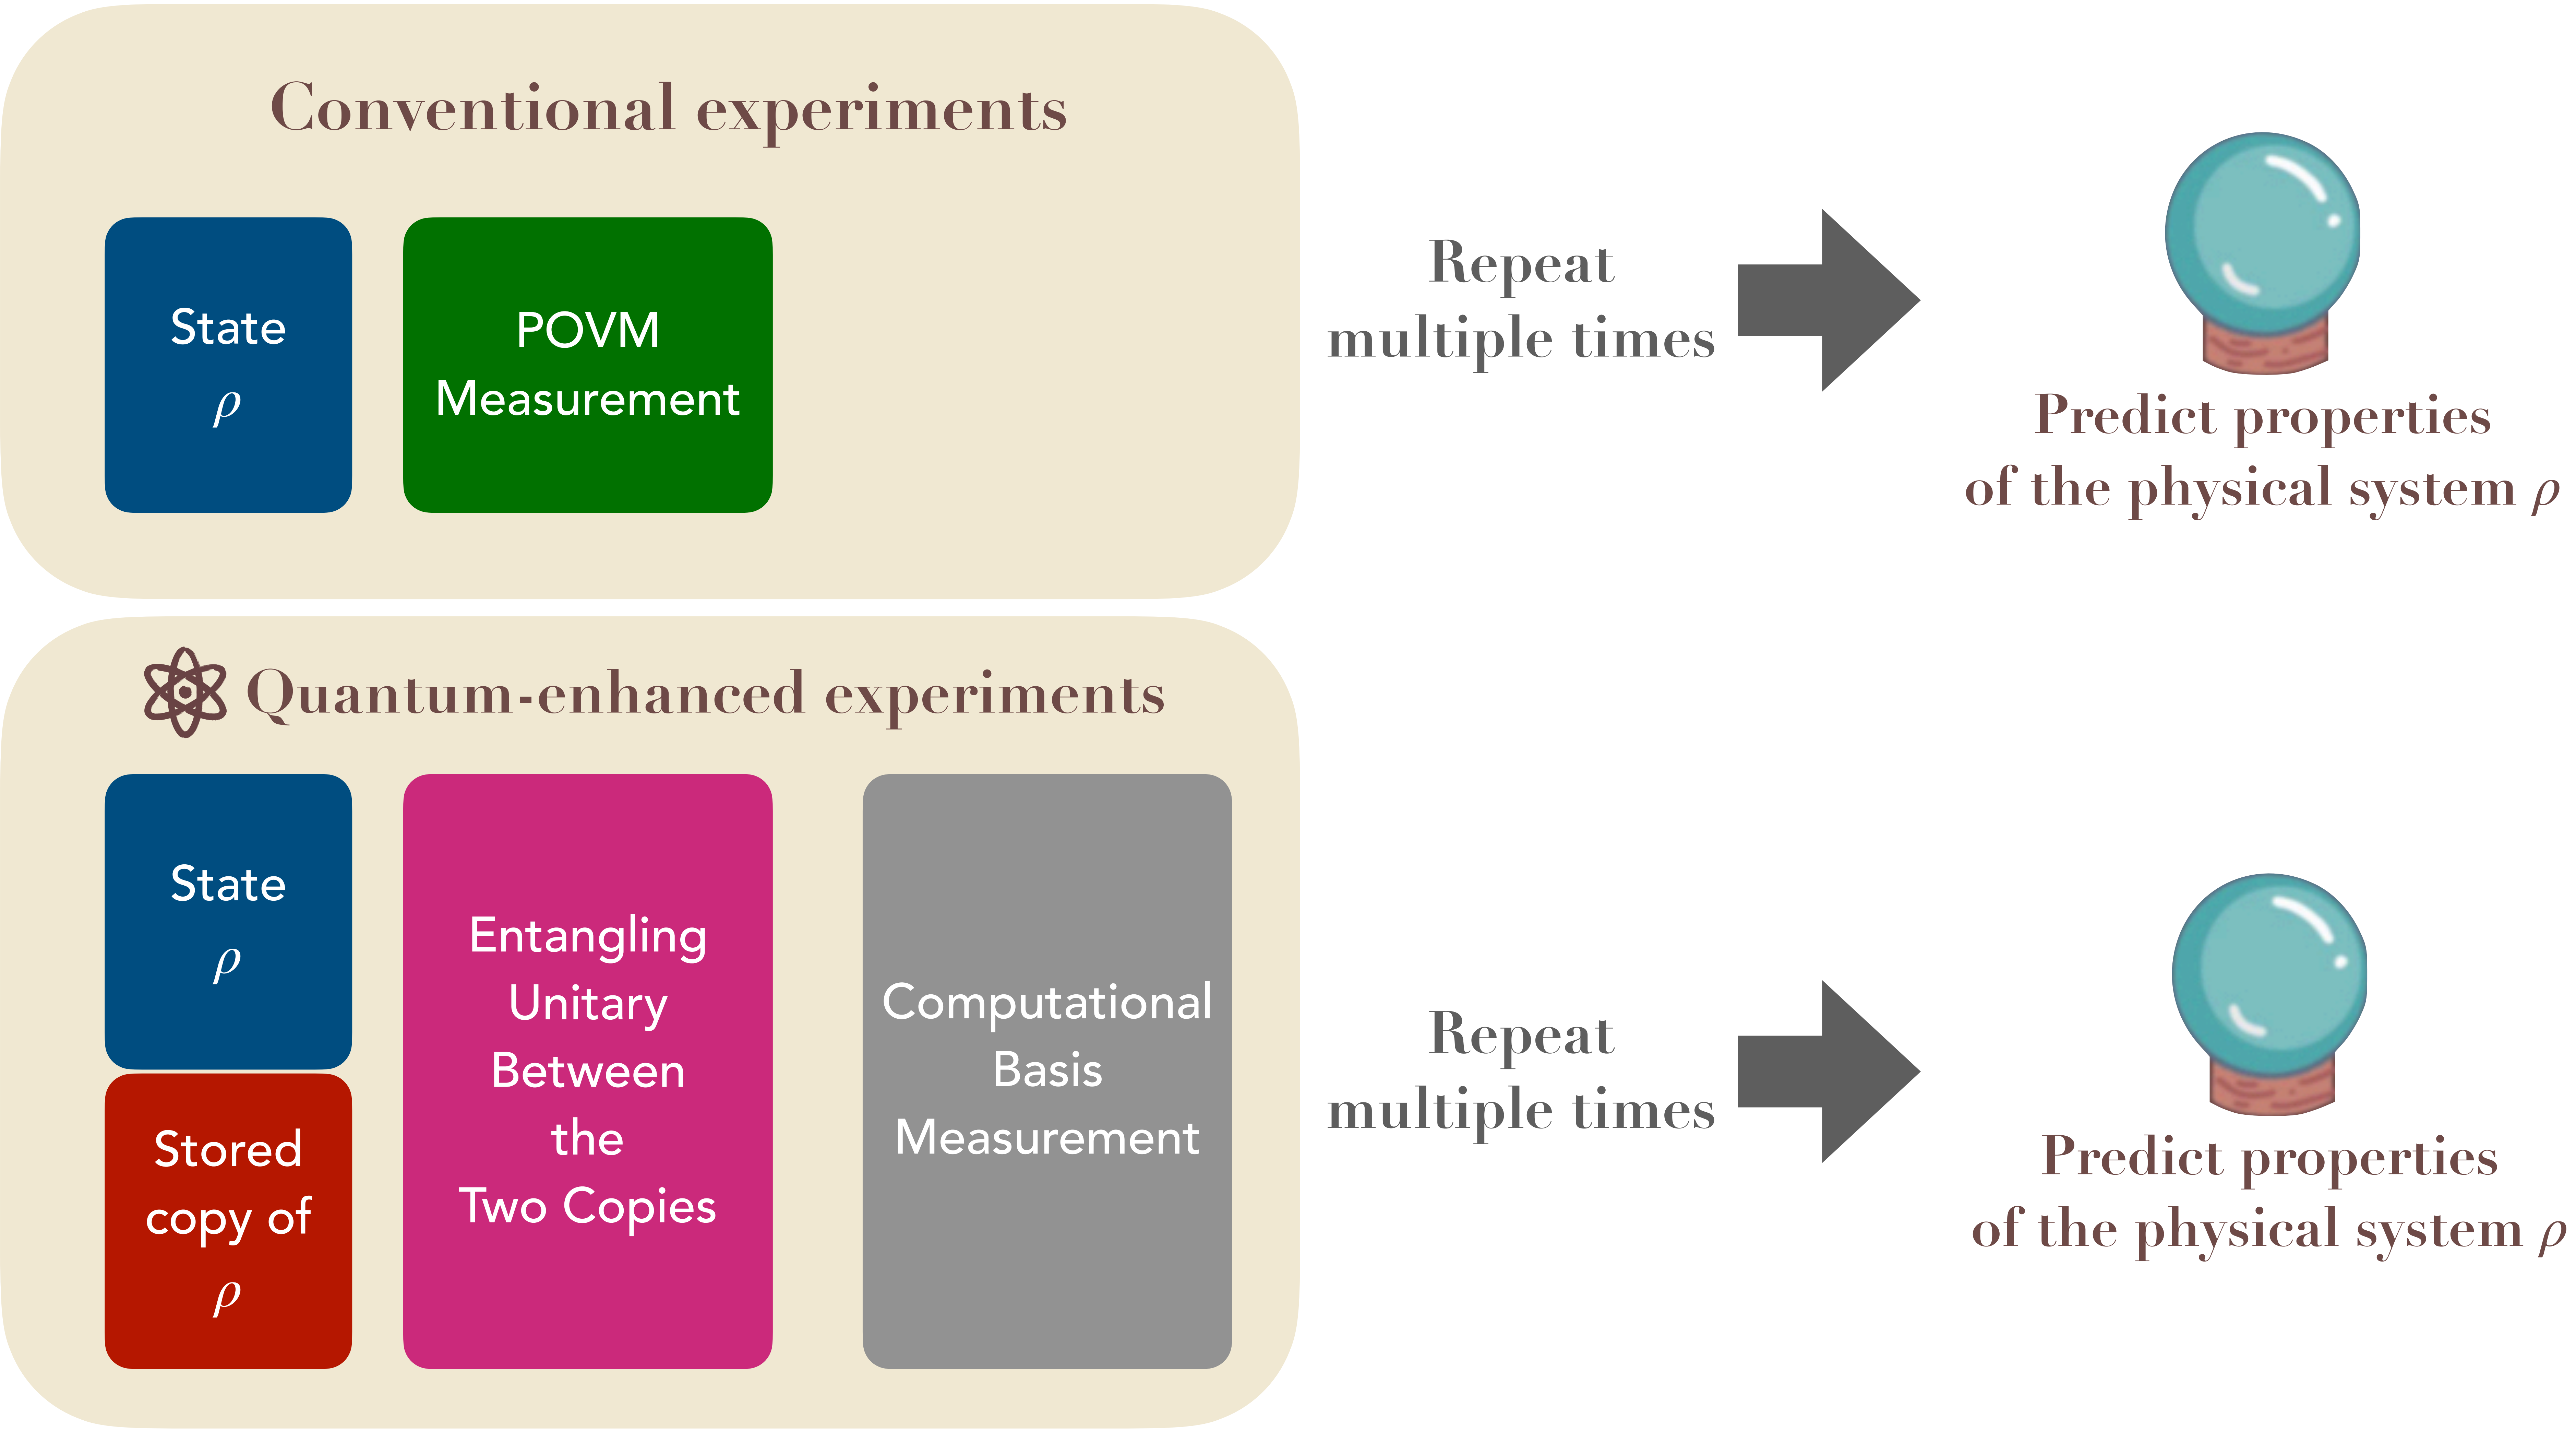

Demonstration on Sycamore:
Quantum advantage in learning states

Utilizing a total of 40 qubits

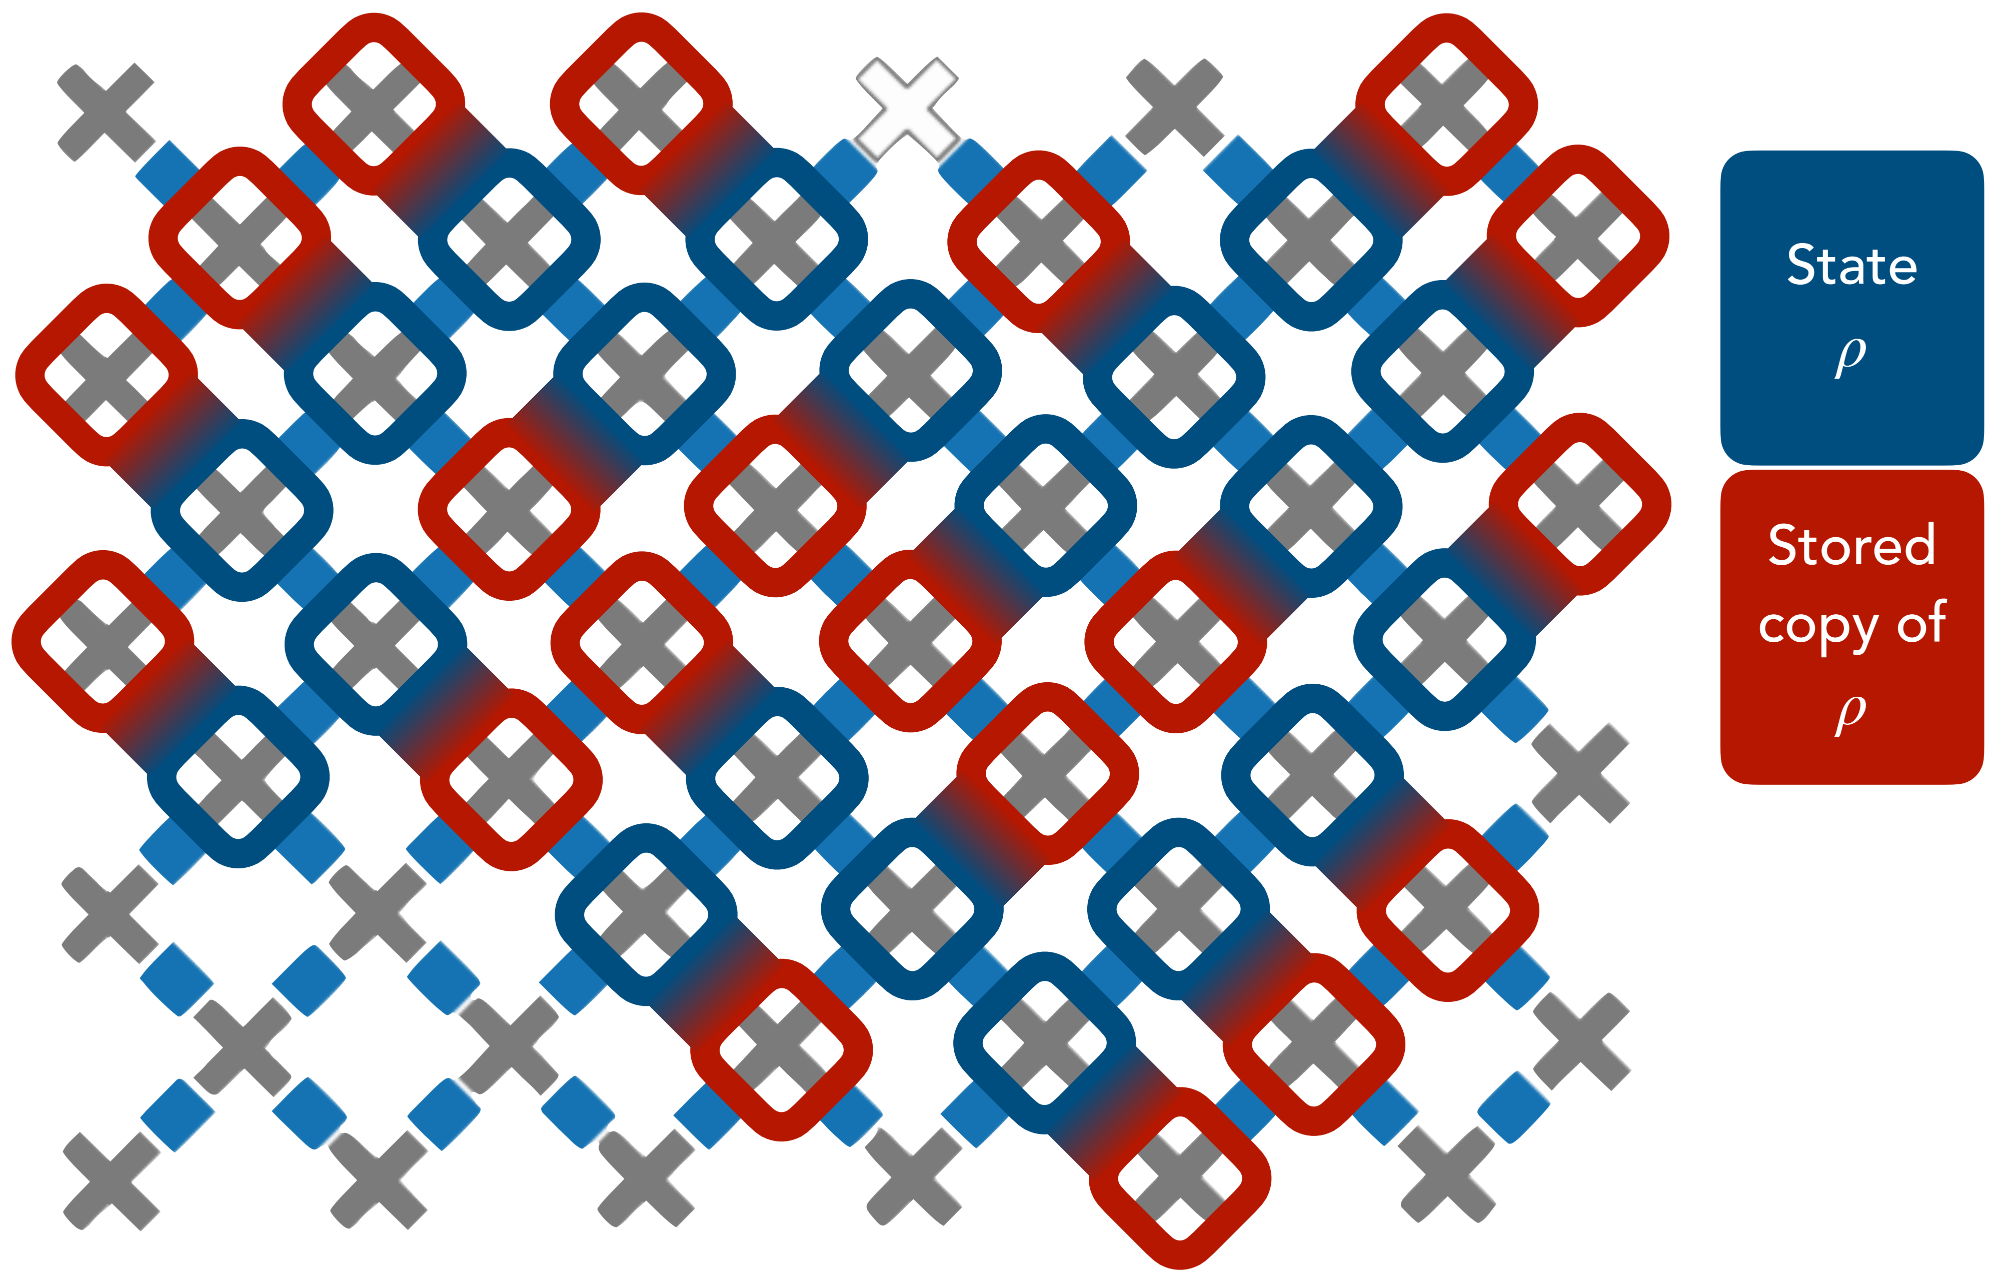

Sycamore Processor

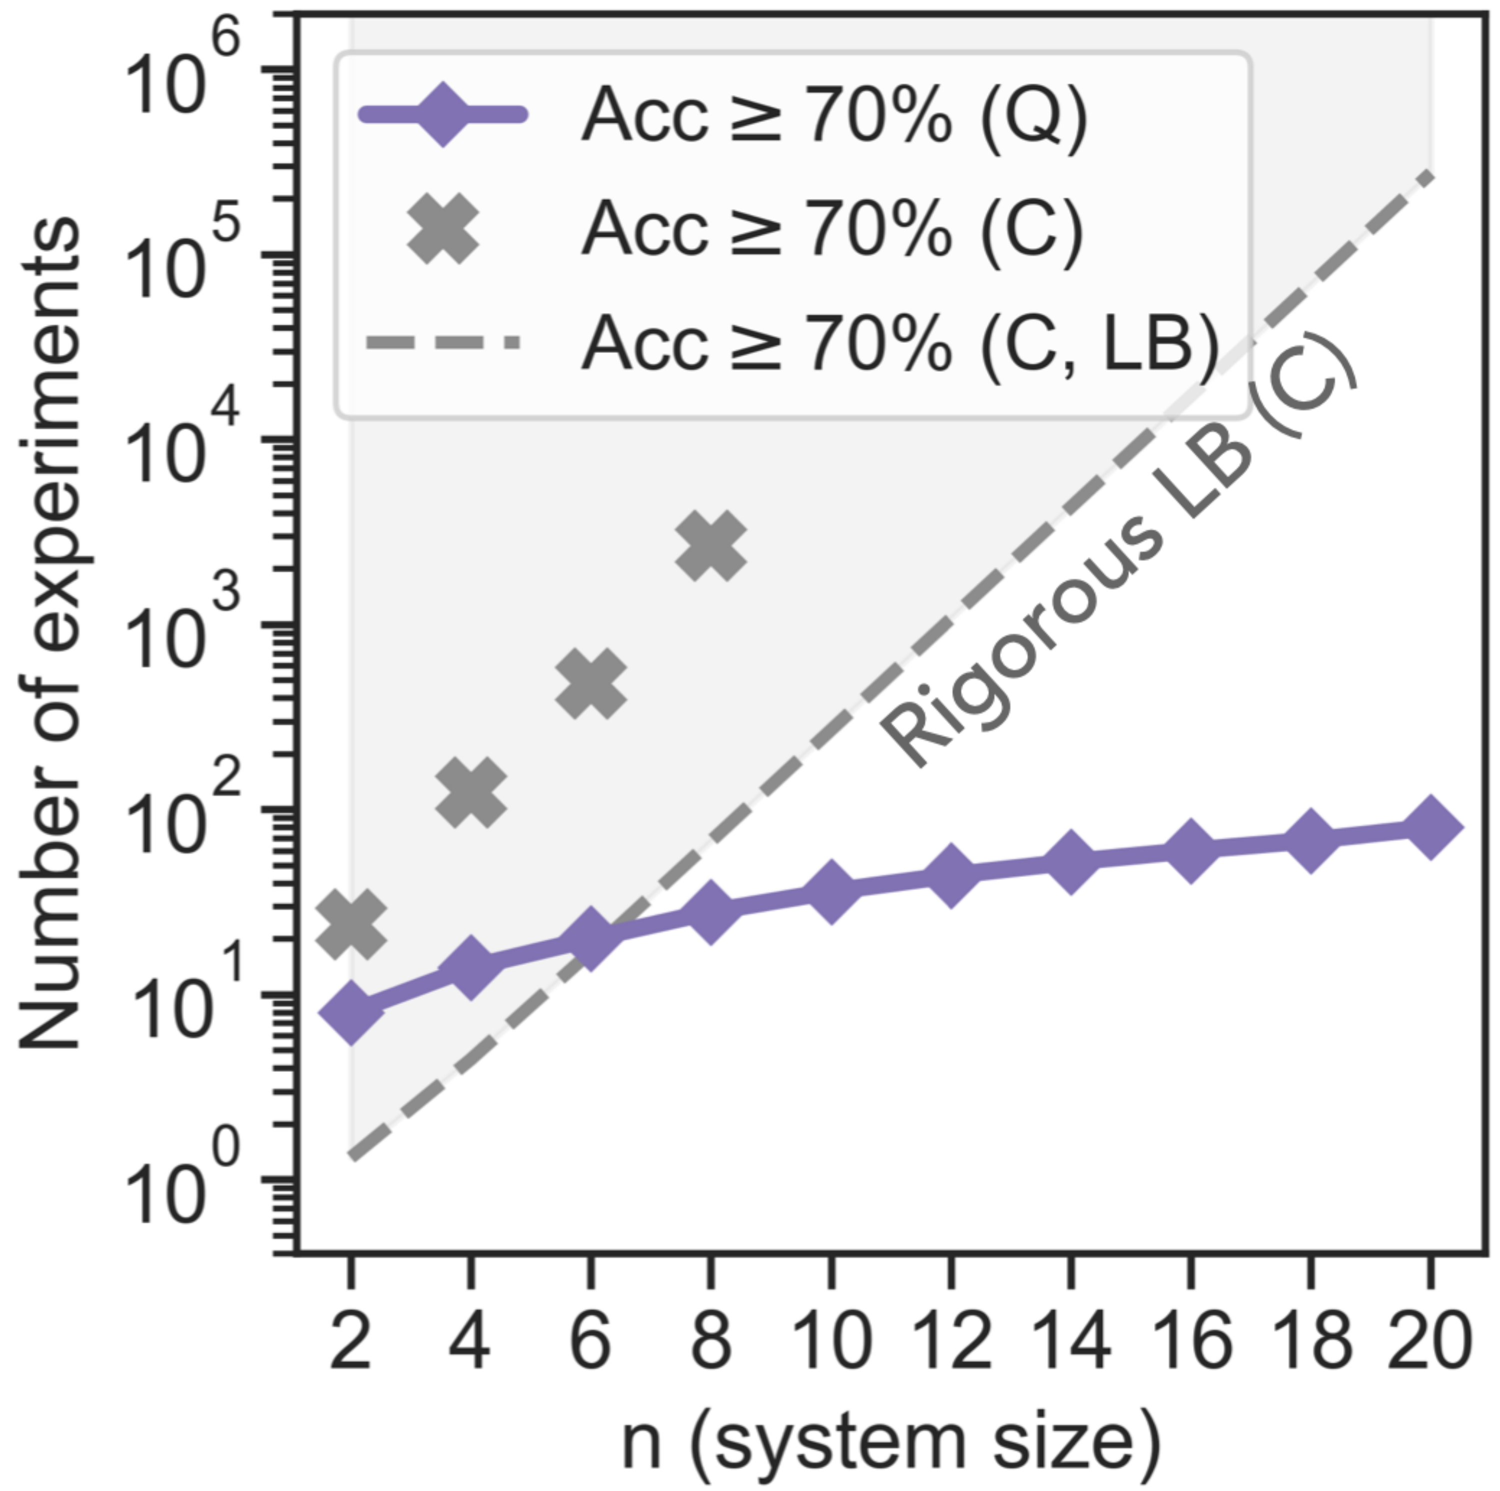

[HFP22] Huang, Flammia, Preskill. Foundations for learning from noisy quantum experiments, *QIP*, 2022.
[HBC+] Huang, Broughton, Cotler, Chen, Li, Mohseni, Neven, Babbush, Kueng, Preskill, McClean. Quantum advantage in learning from experiments, *Science*, 2022.

Demonstration on Sycamore: Quantum advantage in learning dynamics

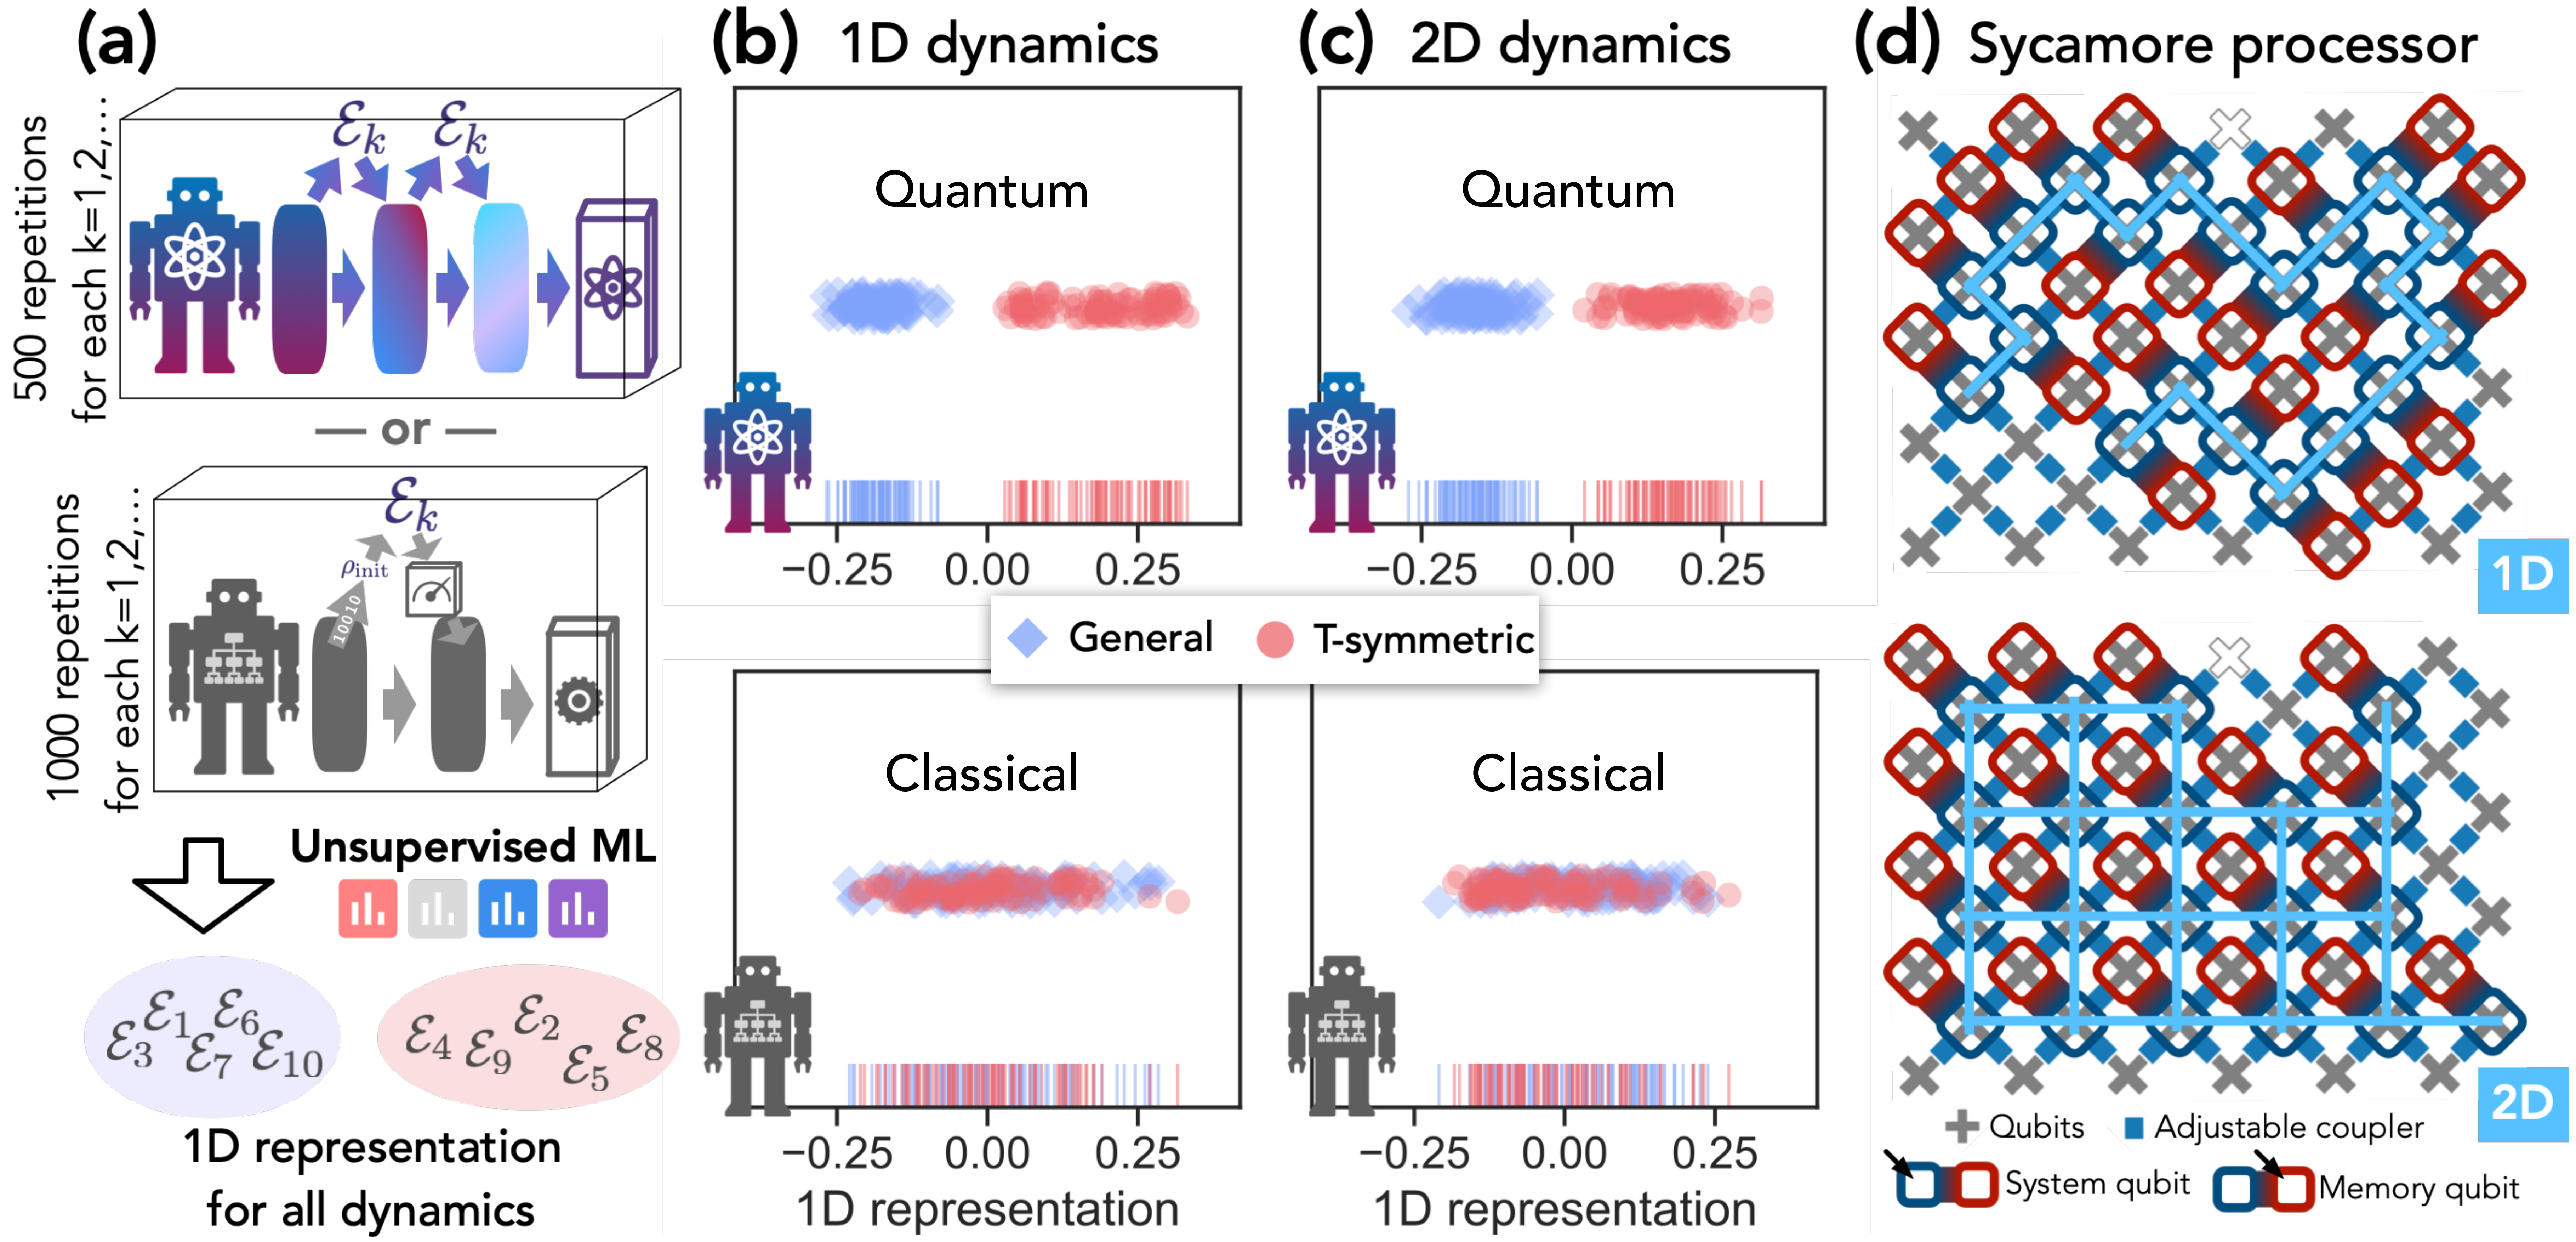

[HFP22] Huang, Flammia, Preskill. Foundations for learning from noisy quantum experiments, *QIP*, 2022.

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

Demonstration on Sycamore:
Quantum advantage in learning dynamics

Conventional experiments

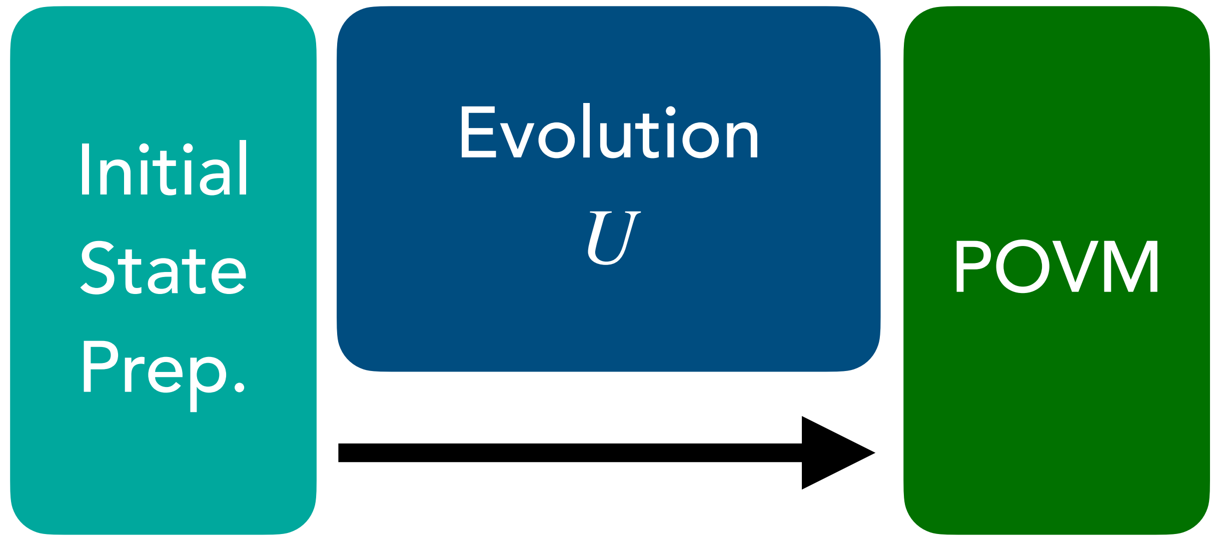

Quantum-enhanced experiments

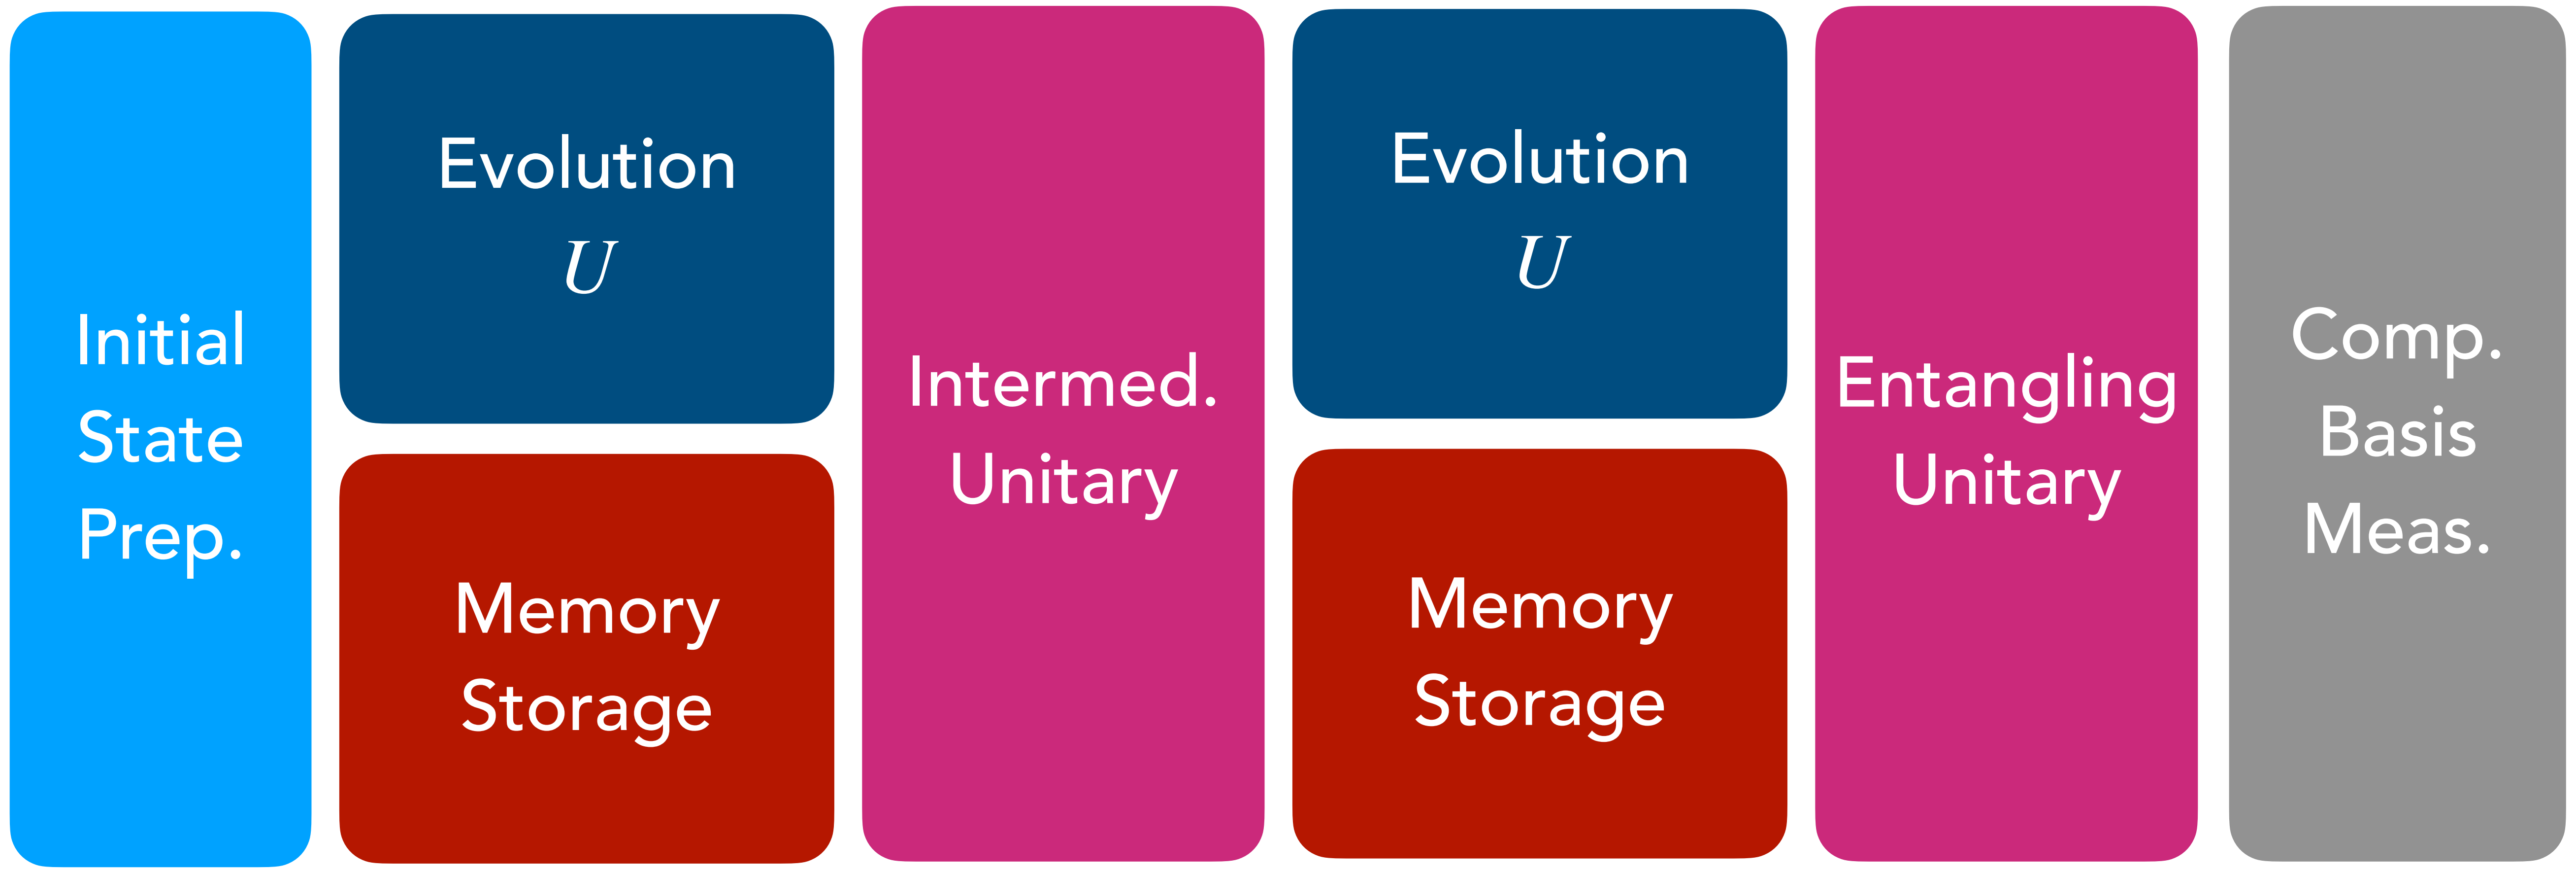

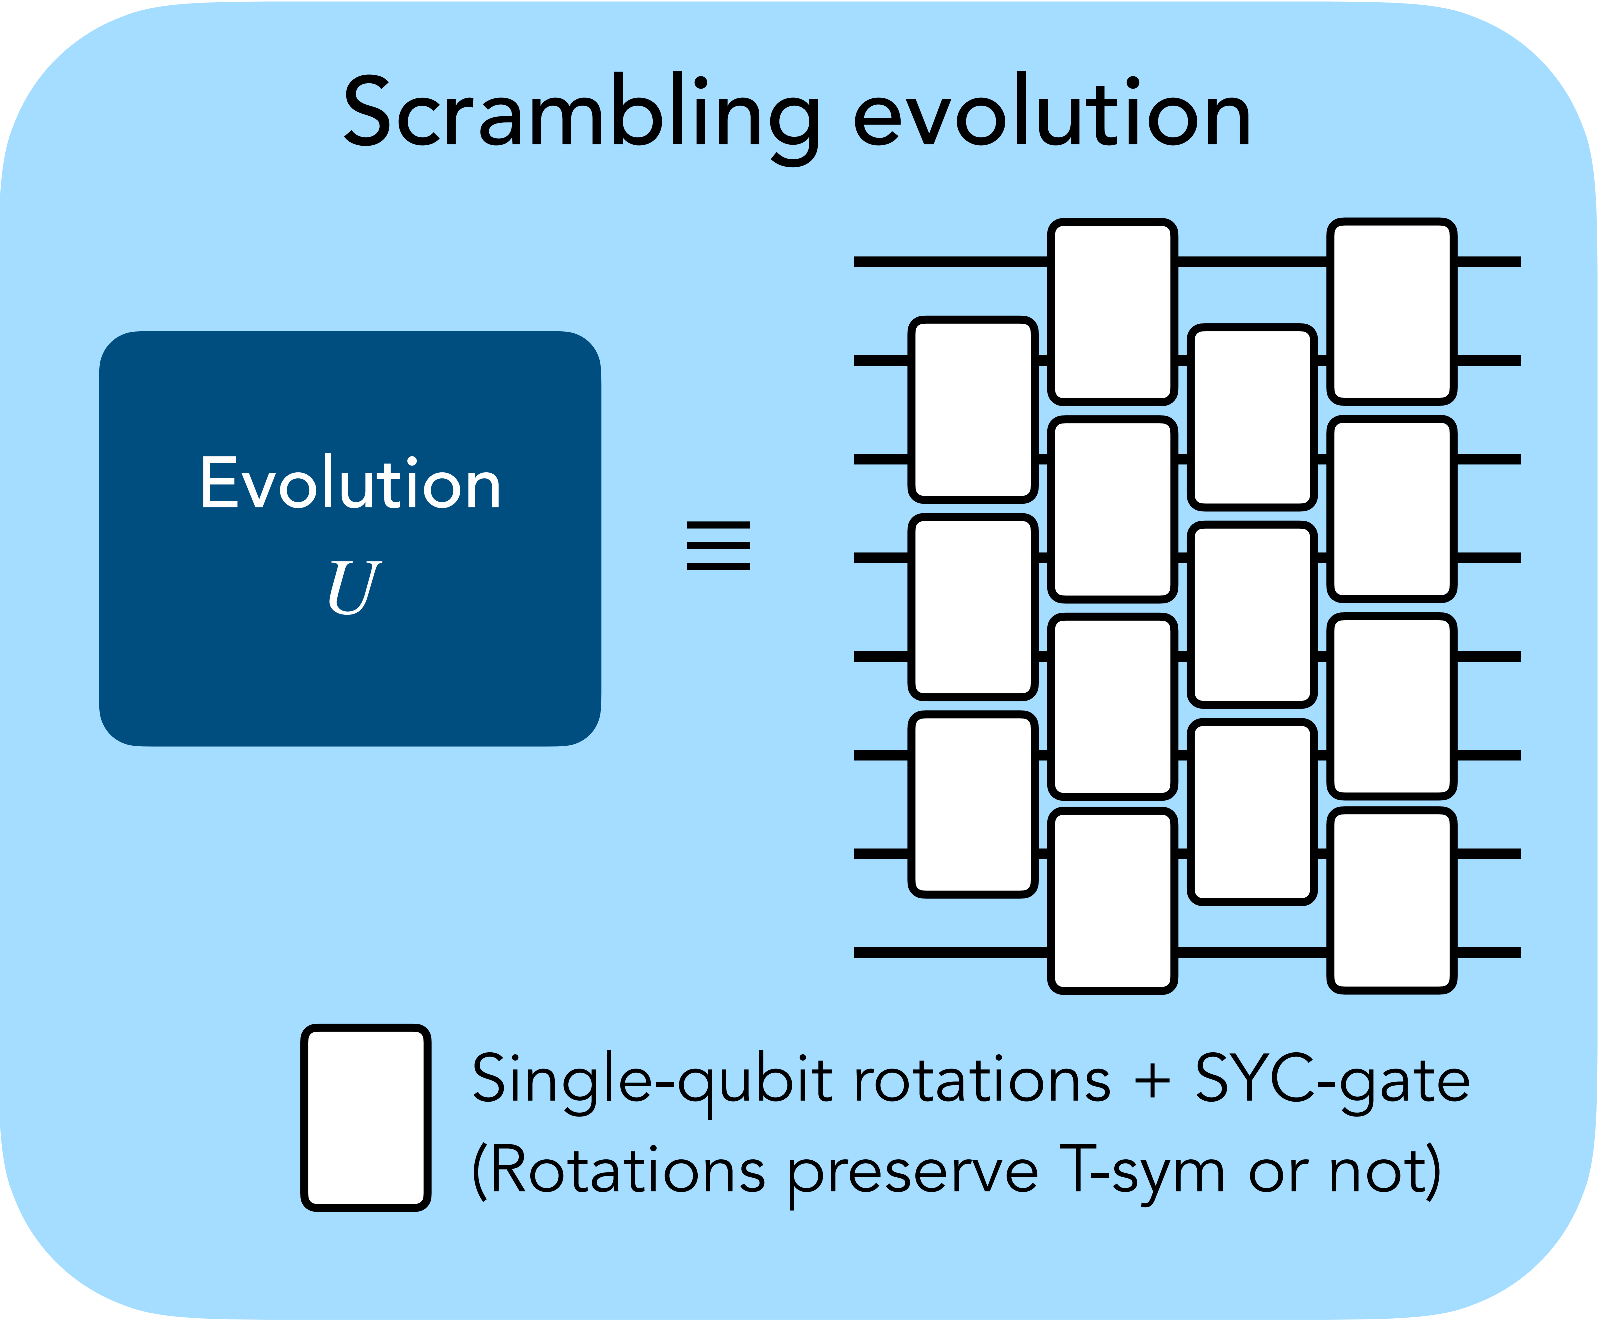

Demonstration on Sycamore: Quantum advantage in learning dynamics

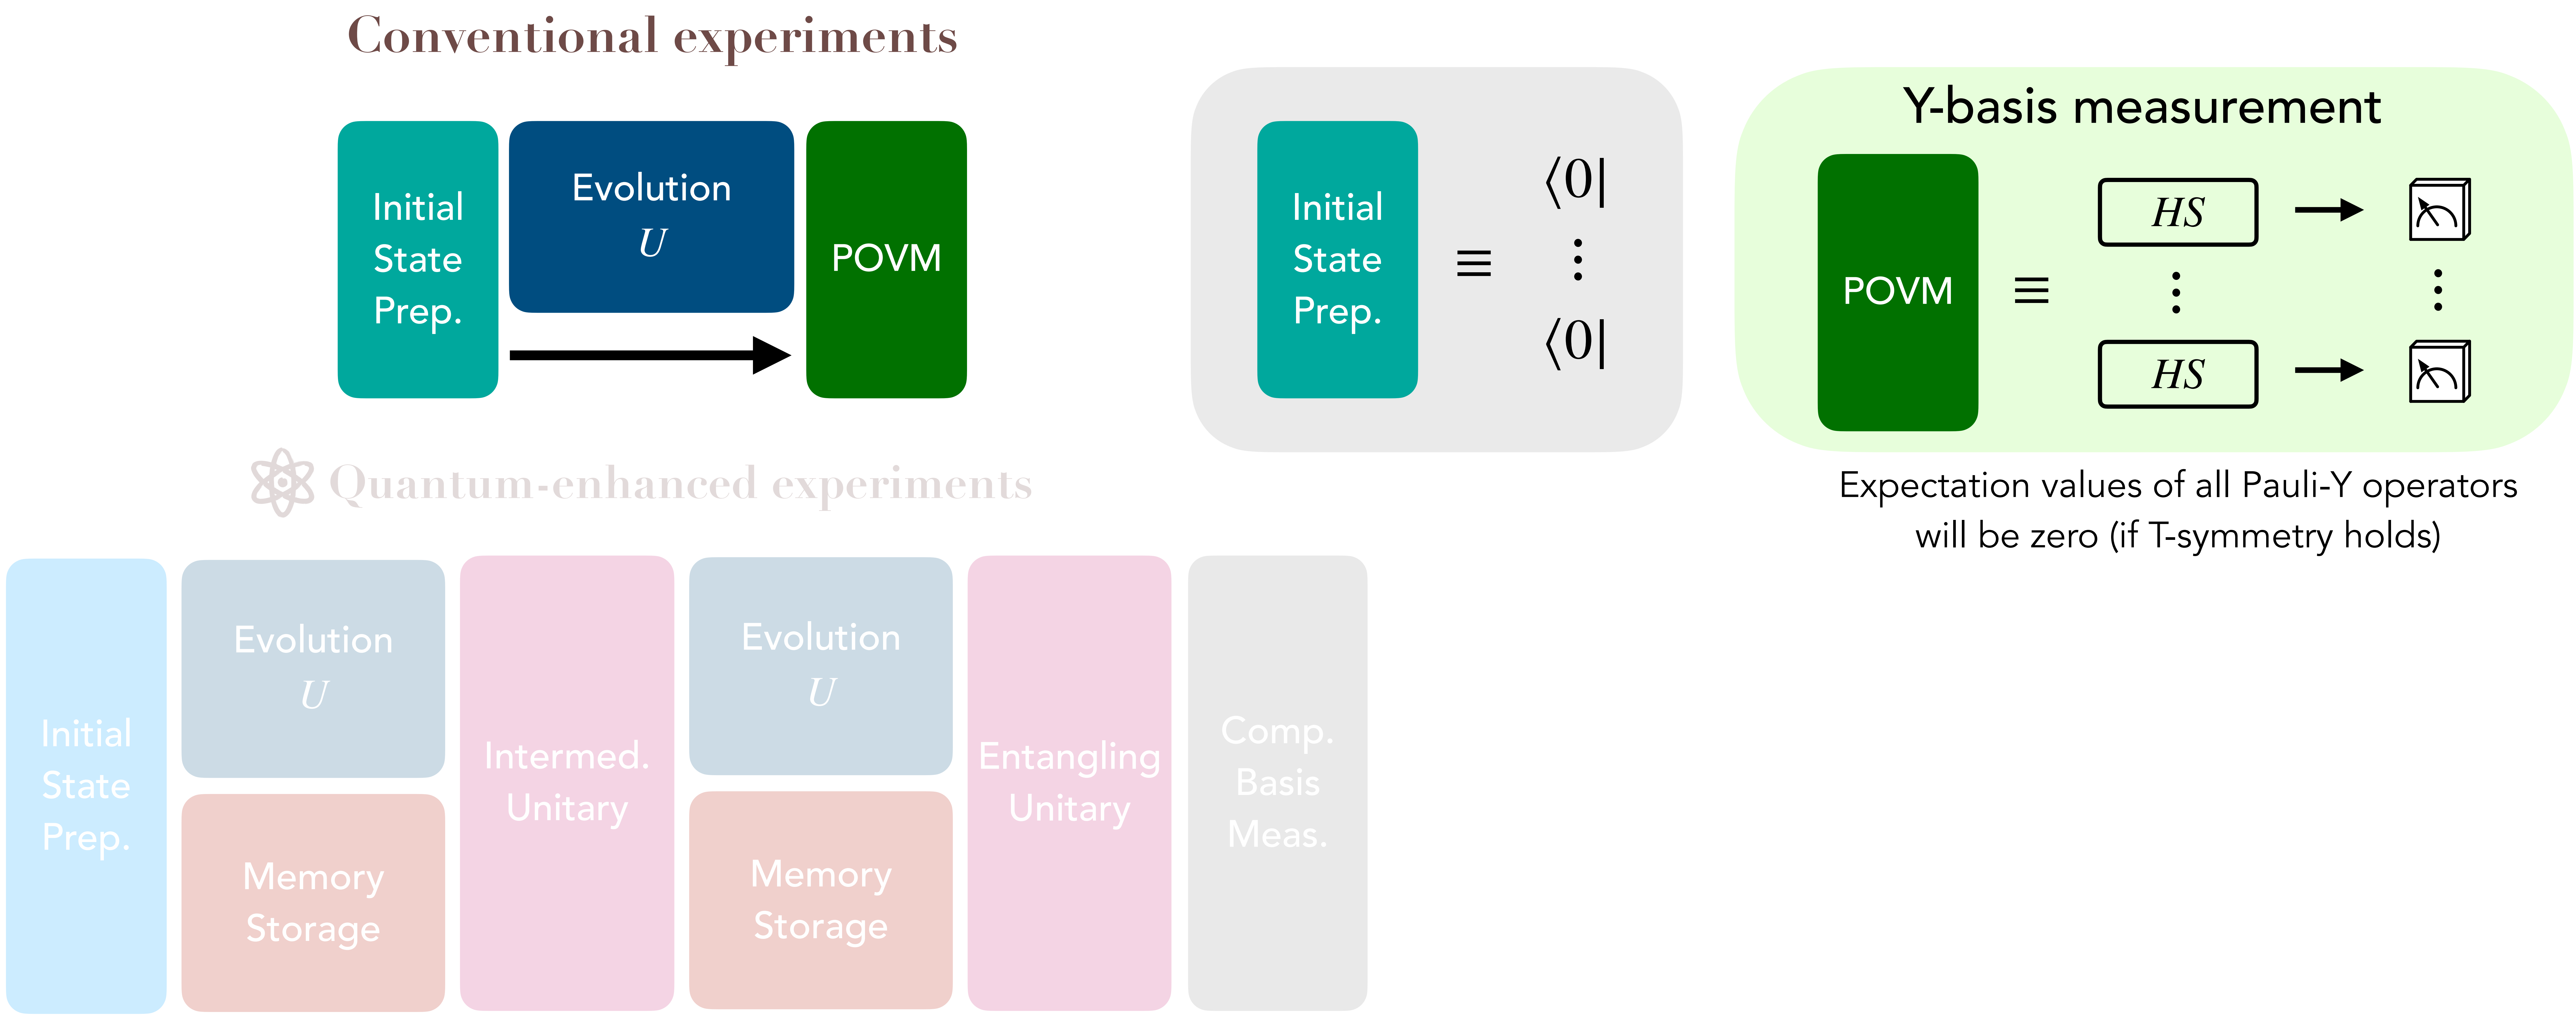

Expectation values of all Pauli-Y operators will be zero (if T-symmetry holds)

Demonstration on Sycamore: Quantum advantage in learning dynamics

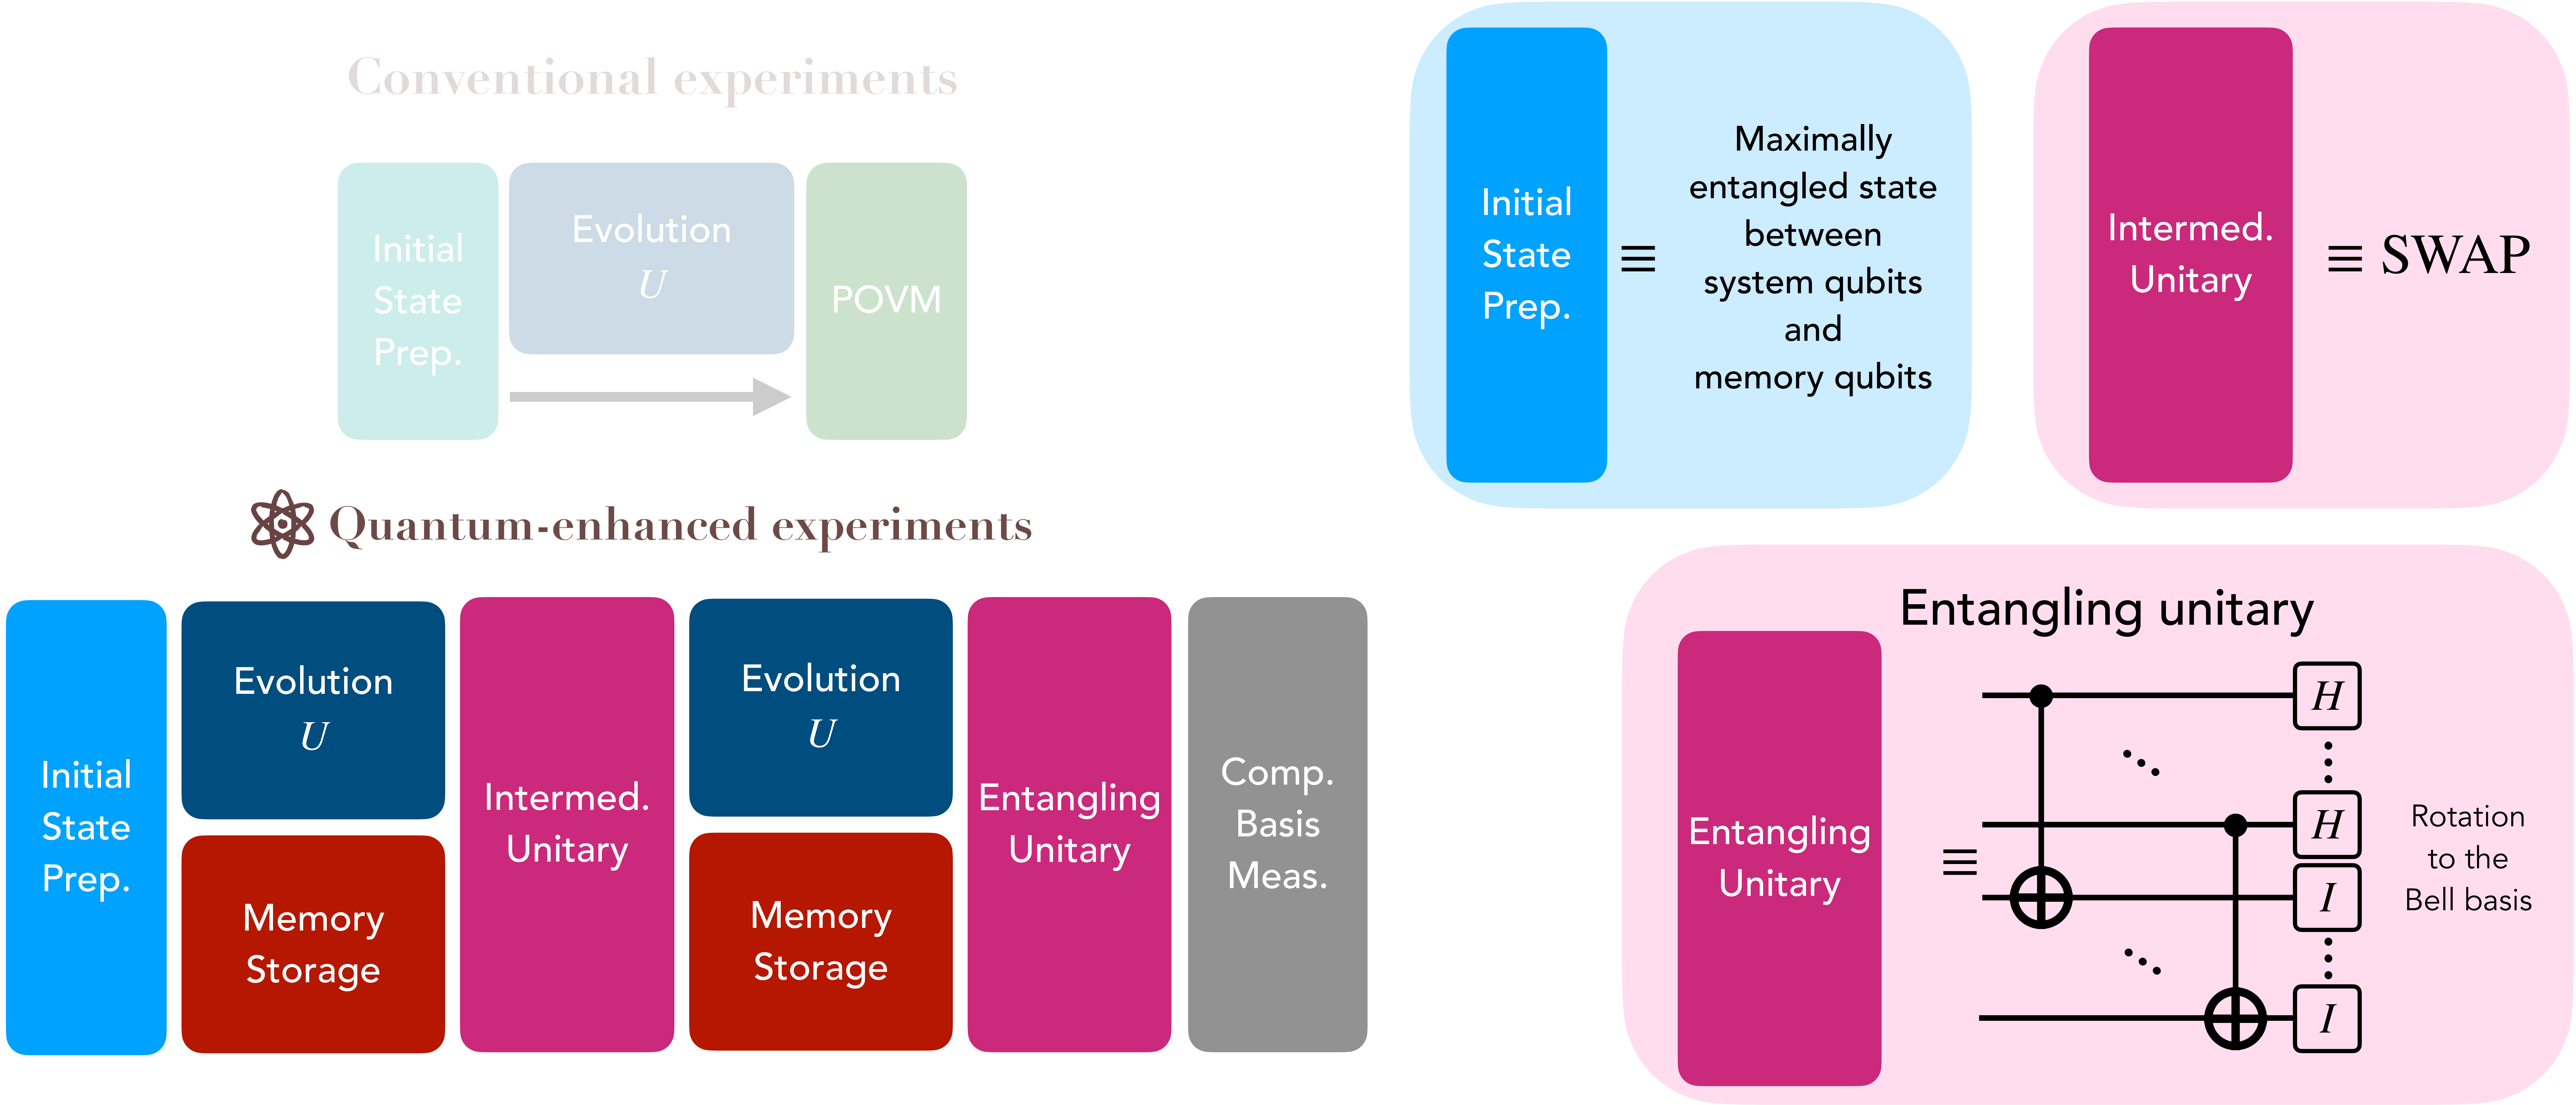

Demonstration on Sycamore: Quantum advantage in learning dynamics

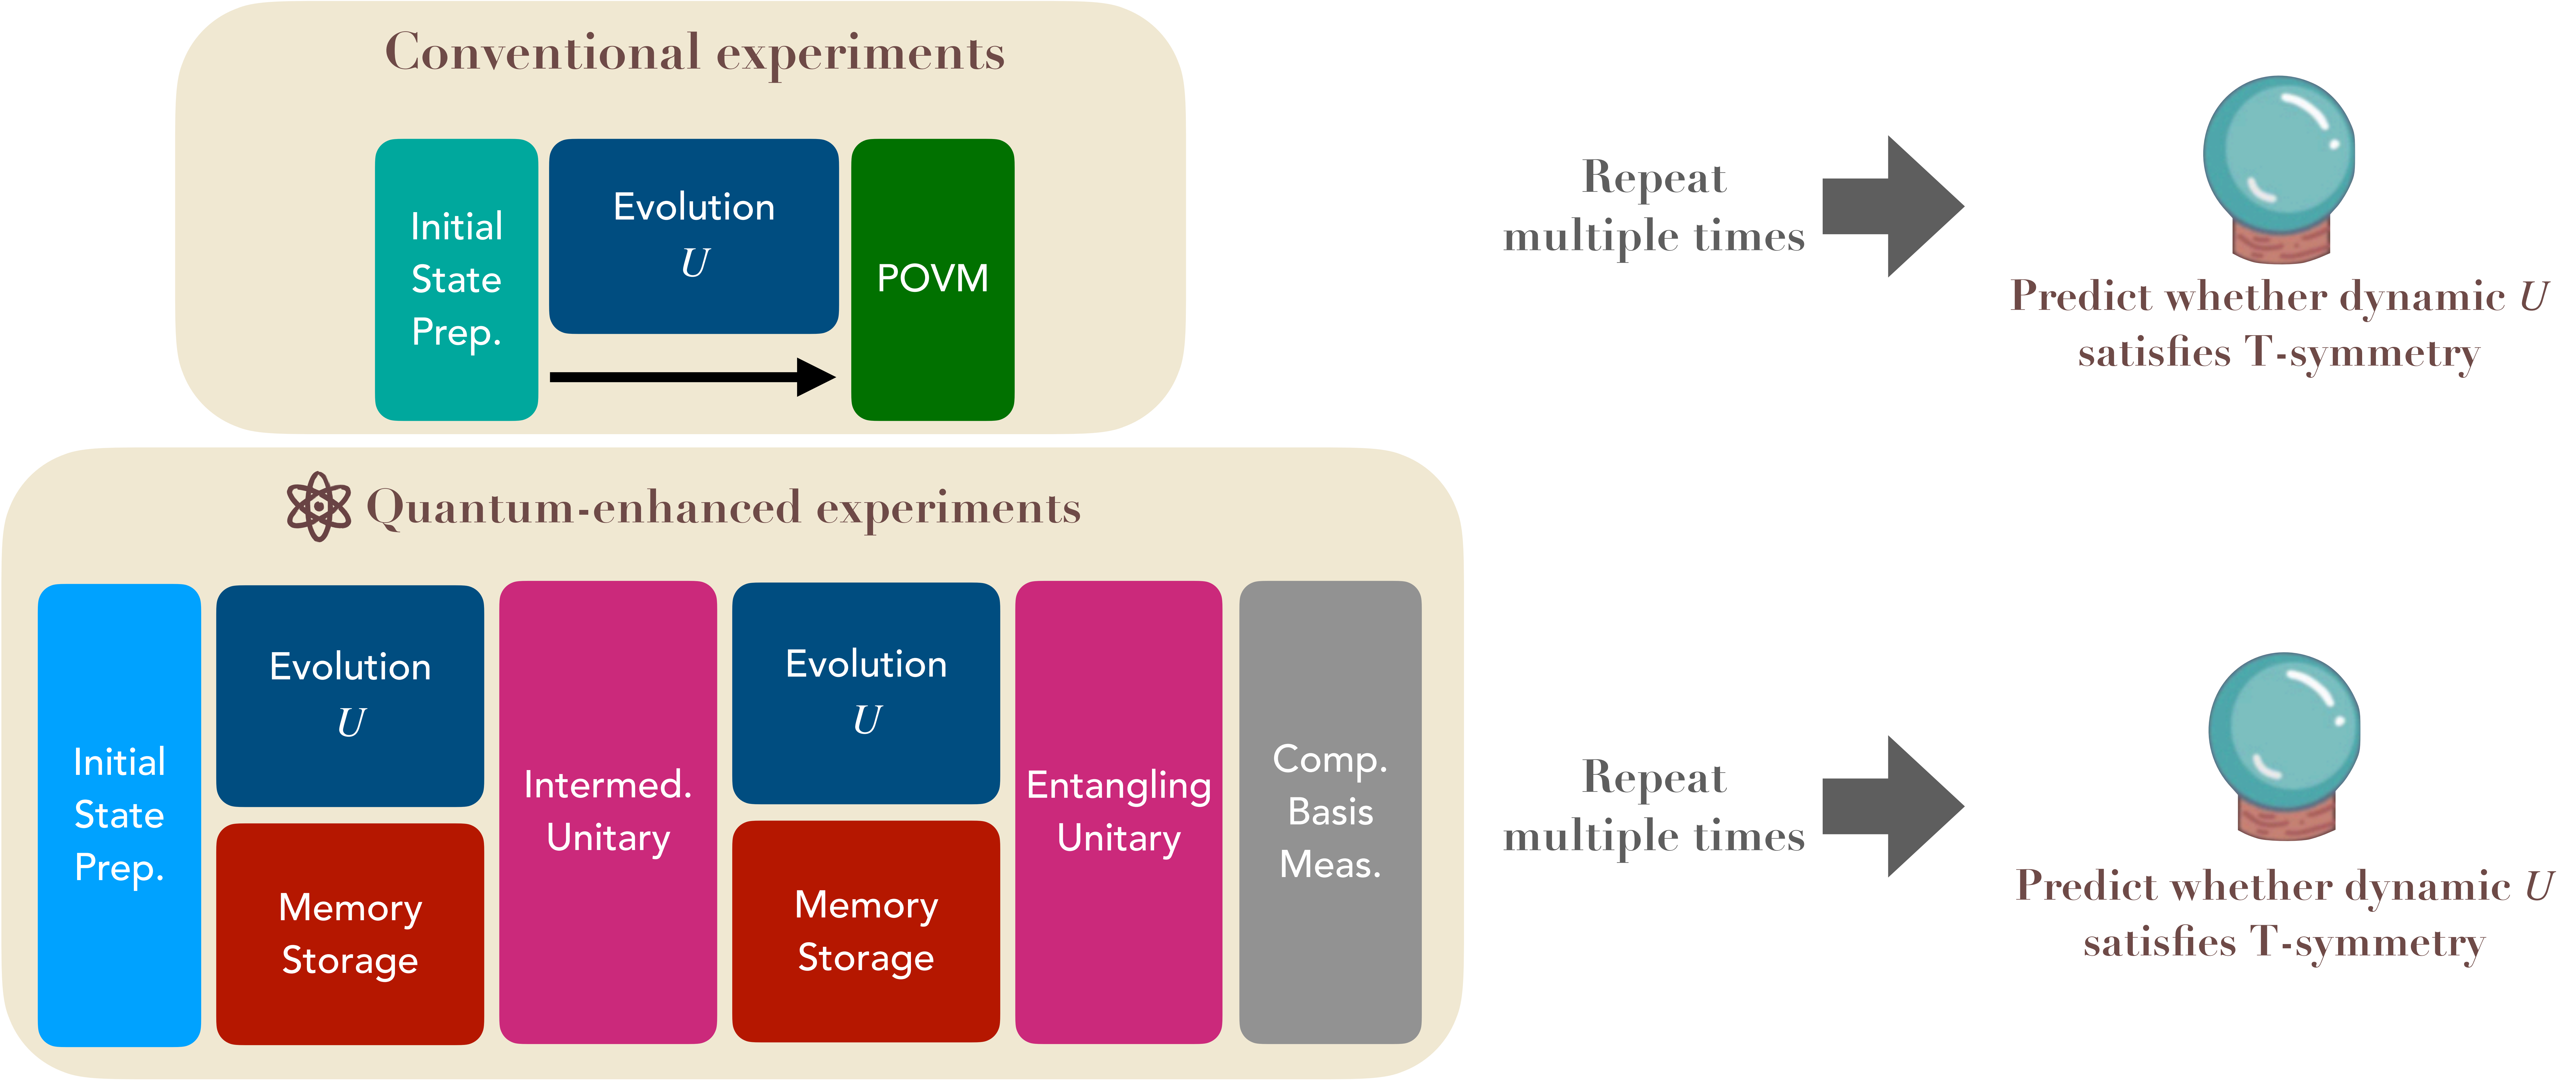

Demonstration on Sycamore: Quantum advantage in learning dynamics

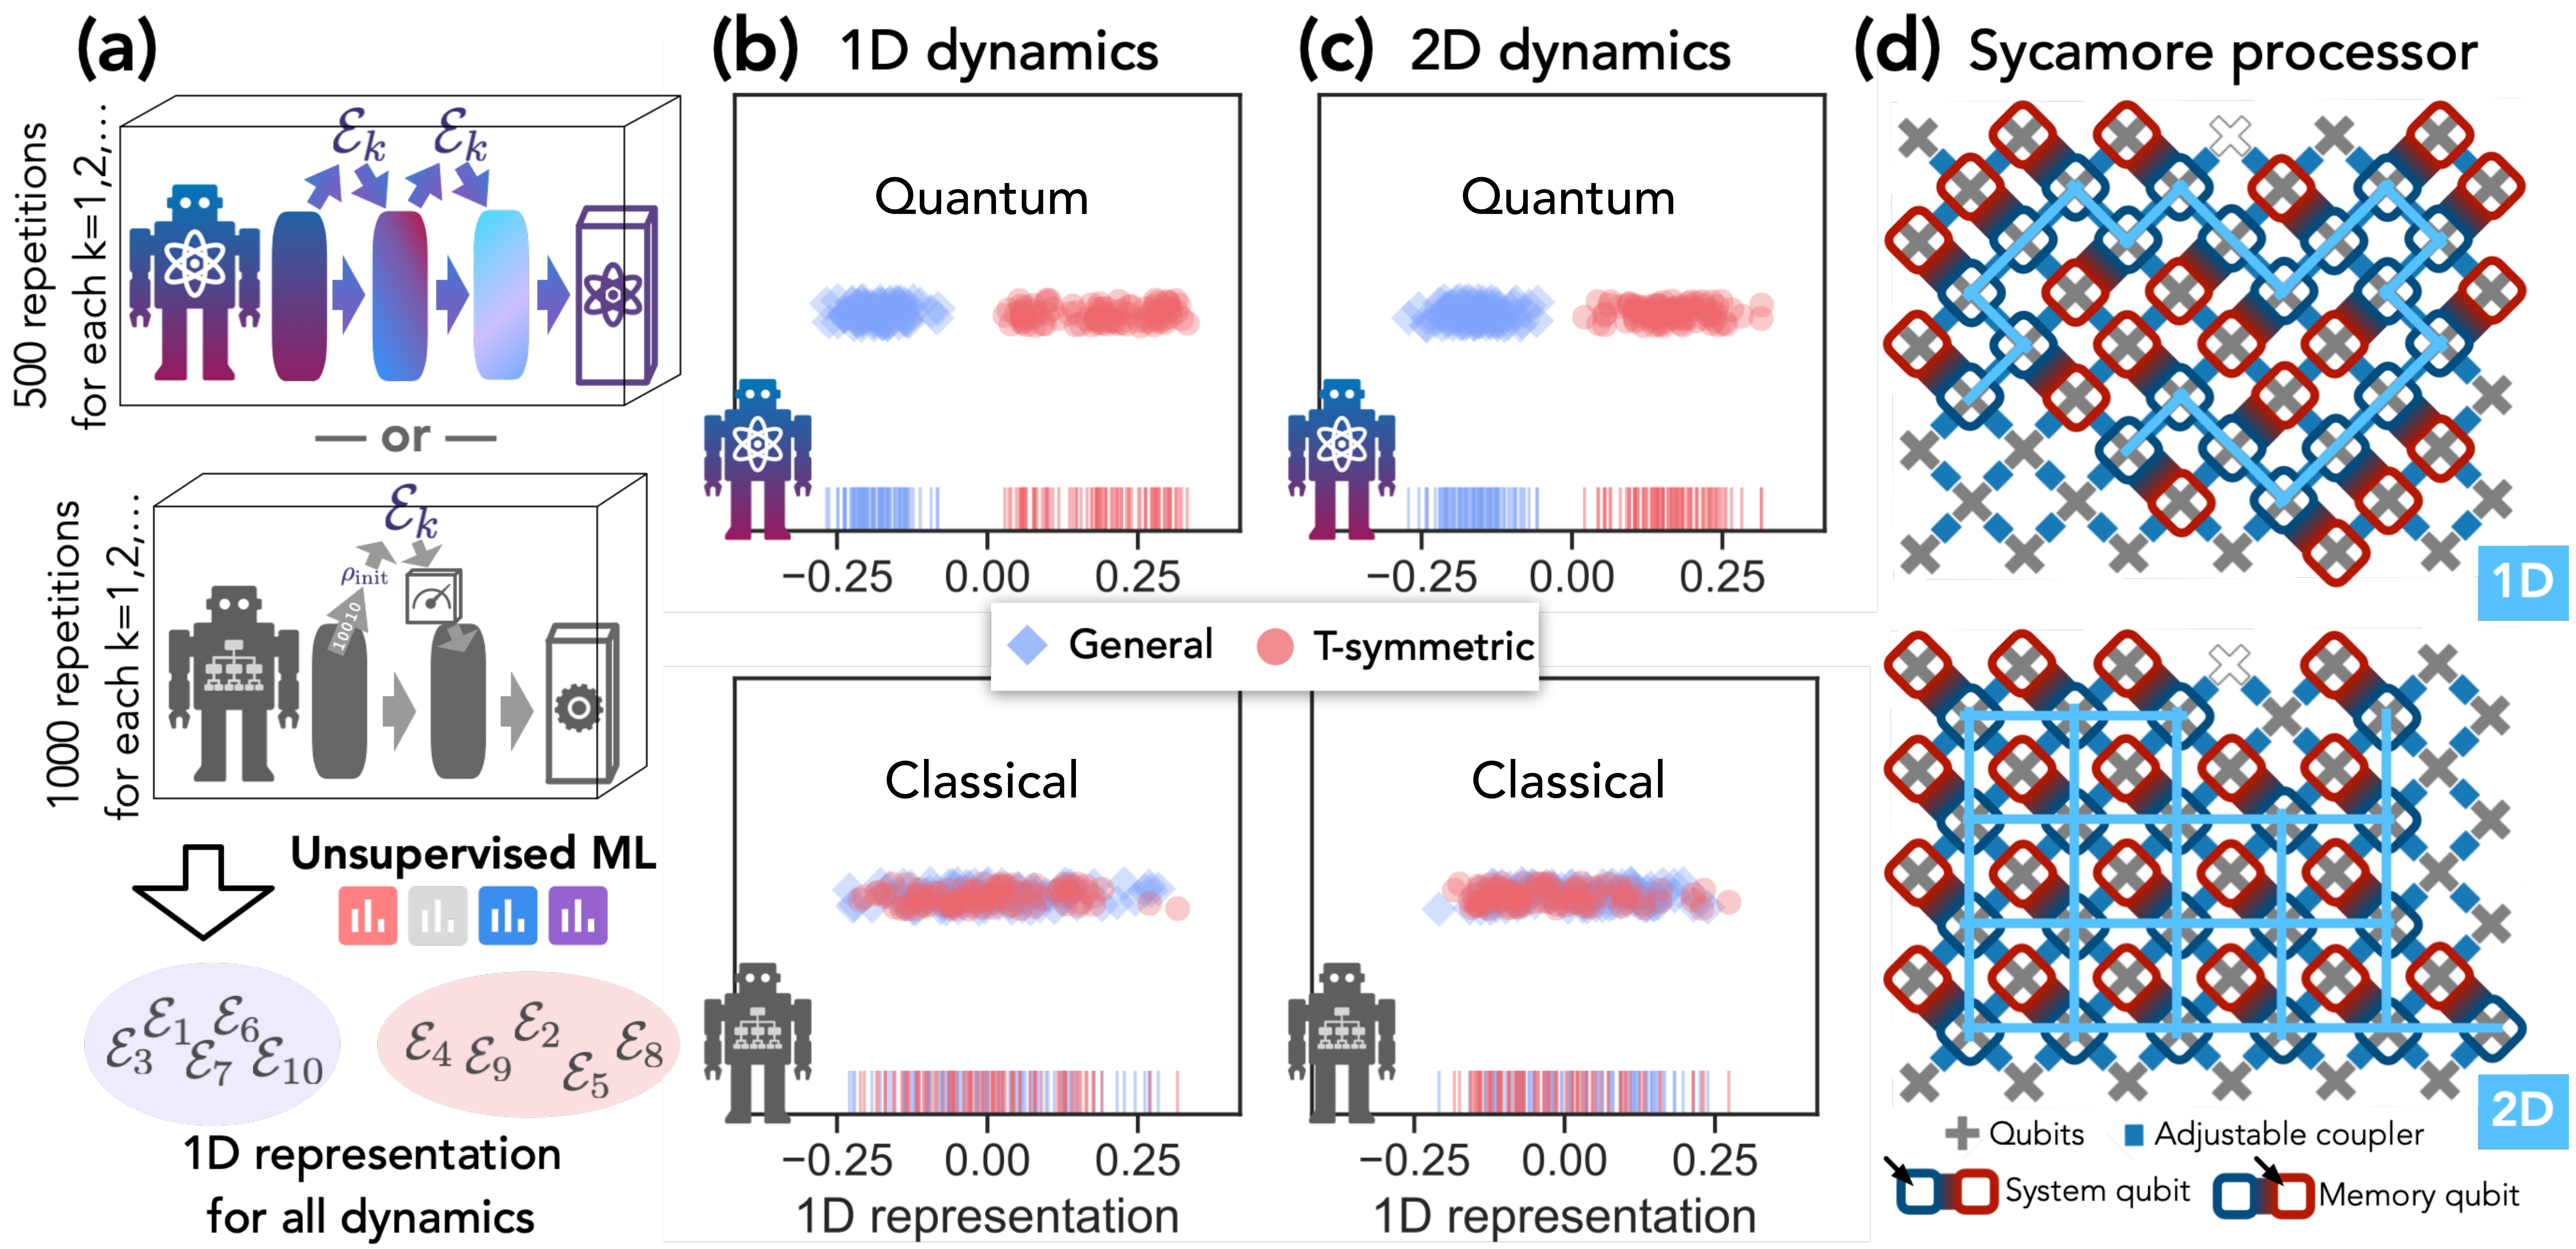

[HFP22] Huang, Flammia, Preskill. Foundations for learning from noisy quantum experiments, *QIP*, 2022.

[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?

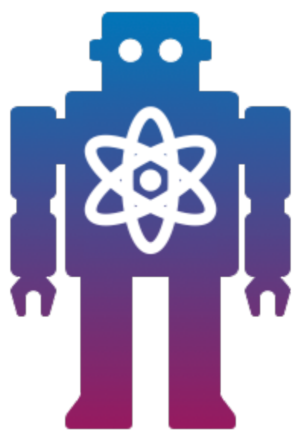

Overview

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?

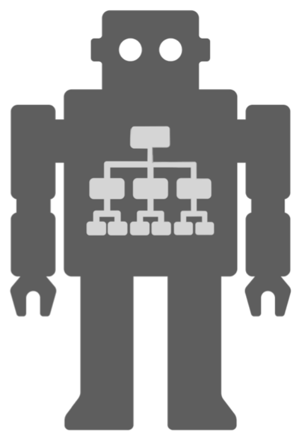

Learning with quantum machines

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

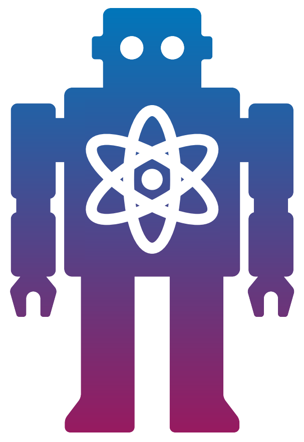

Classical machines $\ll$ Classical learning machines $\ll$ Quantum learning machines

(predicting ground states) (uncovering symmetry, …)

Conclusion

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

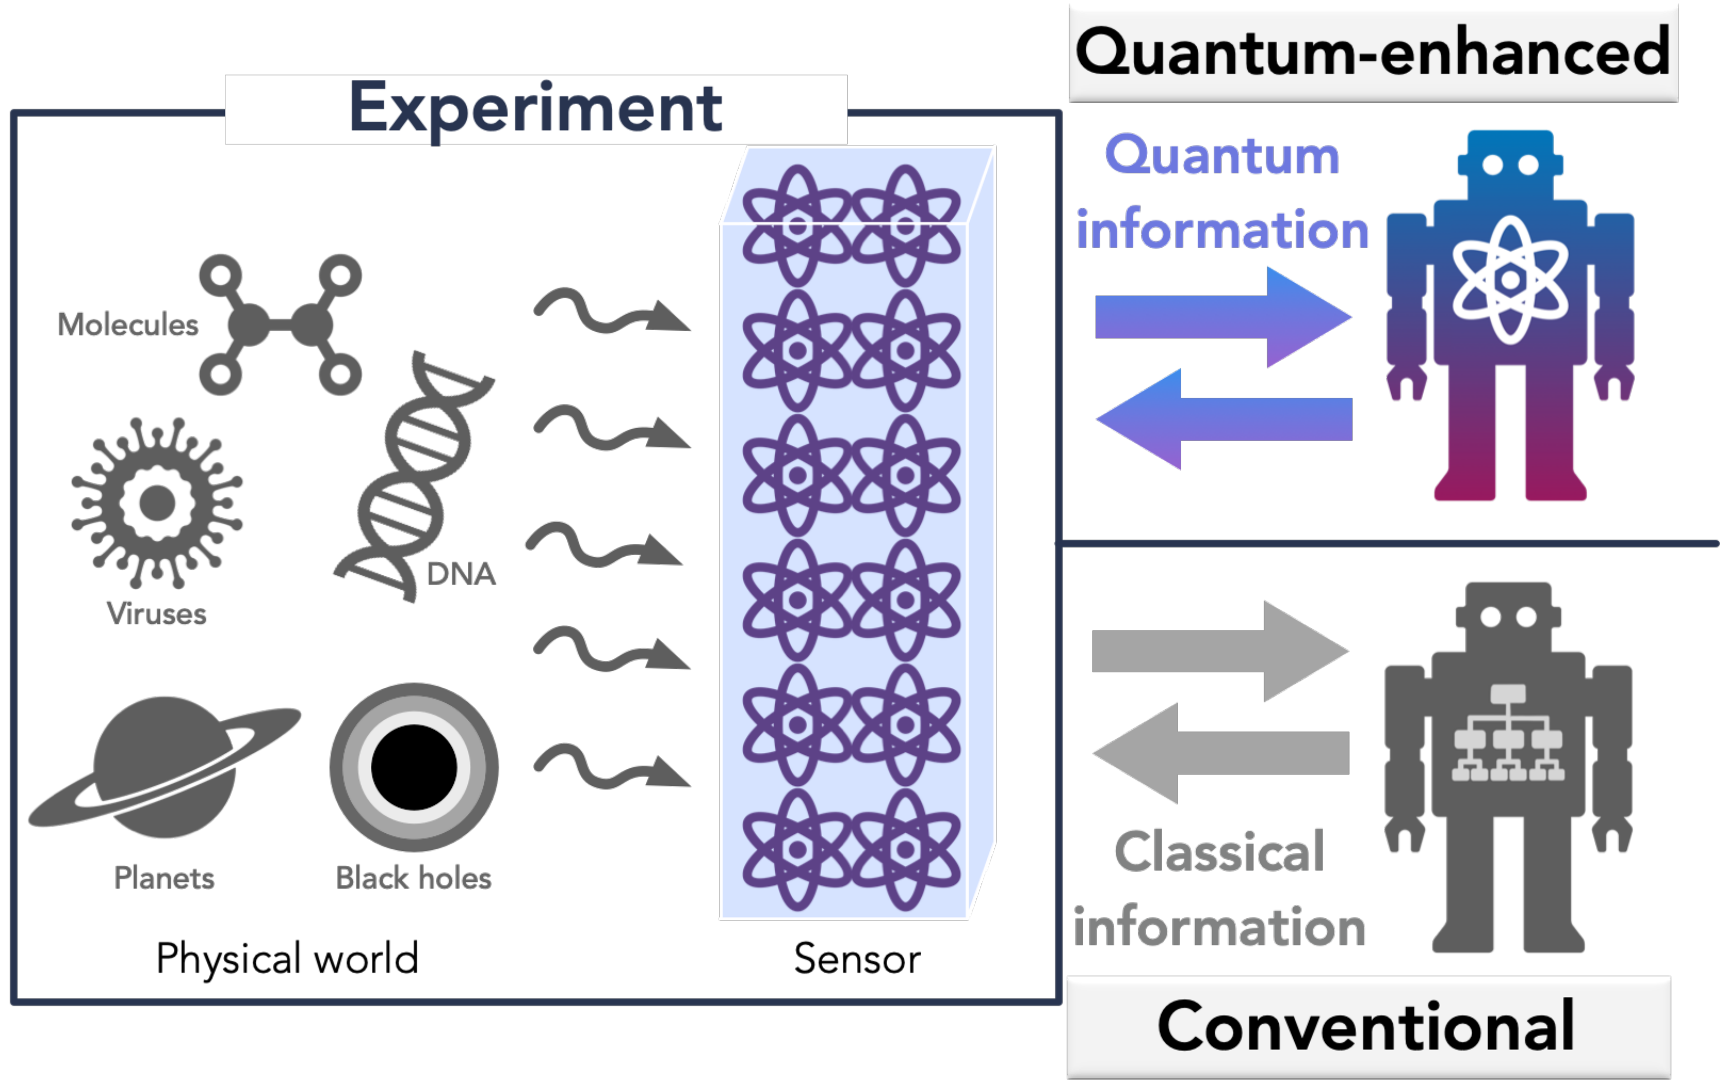

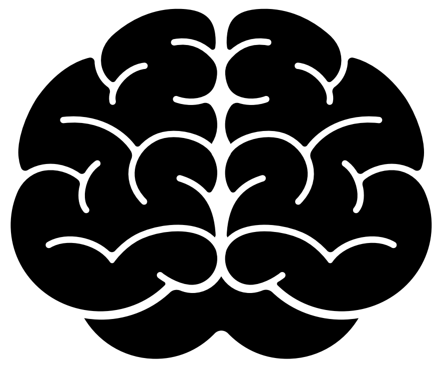

Long-term ambitions

- Could we develop an algorithmic theory to accelerate/automate (quantum) science and the discovery of new physical phenomena?
- Could we 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)