Efficiently Probing Quantum Systems with Compressive Measurements

Zhihui Zhu

The Ohio State University

Computer Science & Engineering

University of Hawaii

ECE Seminar

October 17, 2025

The Ohio State University

The Nobel Prize in Physics 2022

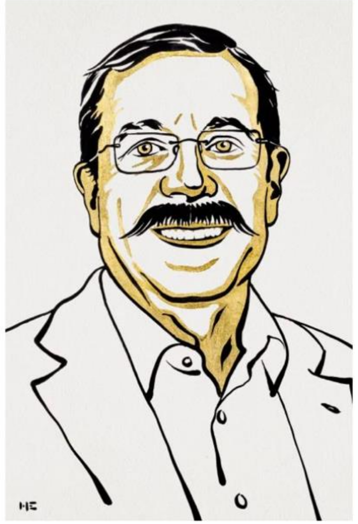

Ill. Niklas Elmehed © Nobel Prize Outreach

Alain Aspect

Prize share: 1/3

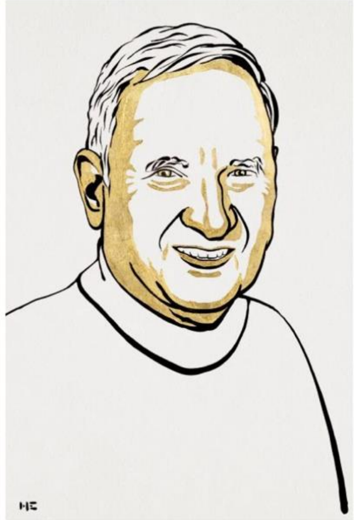

Ill. Niklas Elmehed © Nobel Prize Outreach

John F. Clauser

Prize share: 1/3

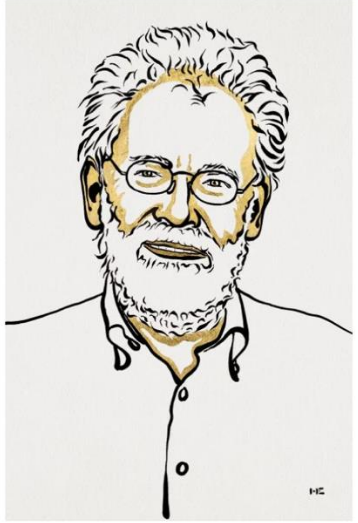

Ill. Niklas Elmehed © Nobel Prize Outreach

Anton Zeilinger

Prize share: 1/3

The Nobel Prize in Physics 2022 was awarded jointly to Alain Aspect, John F. Clauser and Anton Zeilinger "for experiments with entangled photons, establishing the violation of Bell inequalities and pioneering quantum information science"

From https://www.nobelprize.org/prizes/physics/2022/prize-announcement/

Nobel Prize in Physics 2025

Ill. Niklas Elmehed © Nobel Prize Outreach

John Clarke

Prize share: 1/3

Ill. Niklas Elmehed © Nobel Prize Outreach

Michel H. Devoret

Prize share: 1/3

Ill. Niklas Elmehed © Nobel Prize Outreach

John M. Martinis

Prize share: 1/3

to John Clarke, Michel H. Devoret and John M. Martinis "for the discovery of macroscopic quantum mechanical tunnelling and energy quantisation in an electric circuit"

https://www.nobelprize.org/prizes/physics/2025/summary/

We are in the Second Quantum Revolution

- First quantum revolution (first half of 20th):
  - discover of fundamental laws of microscopic realm
  - formulation of quantum physics

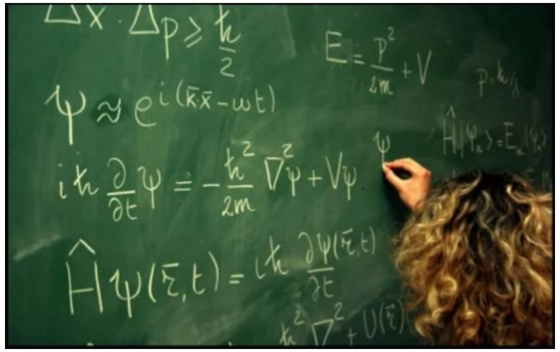

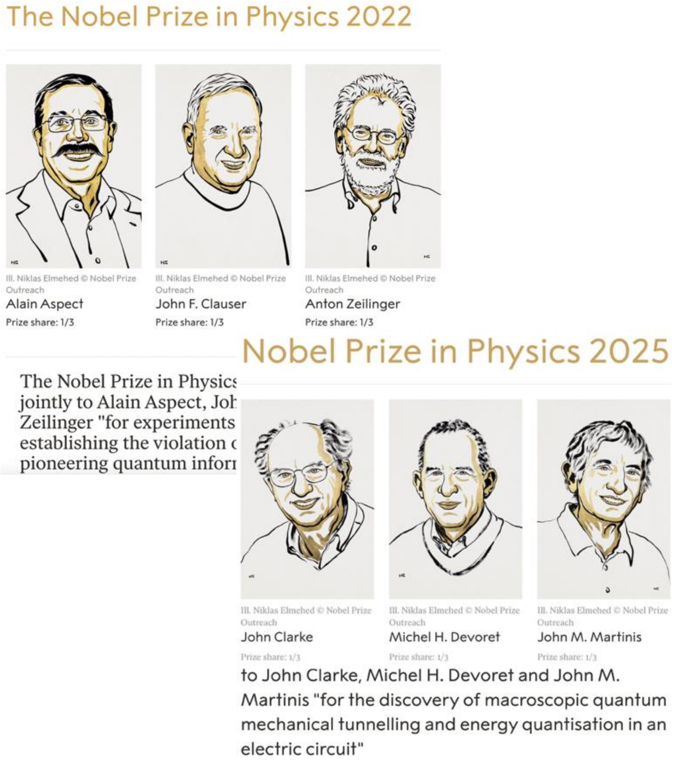

- Second quantum revolution
  - Technologies based on the manipulation of individual quantum systems
  - Use properties such as superposition and entanglement

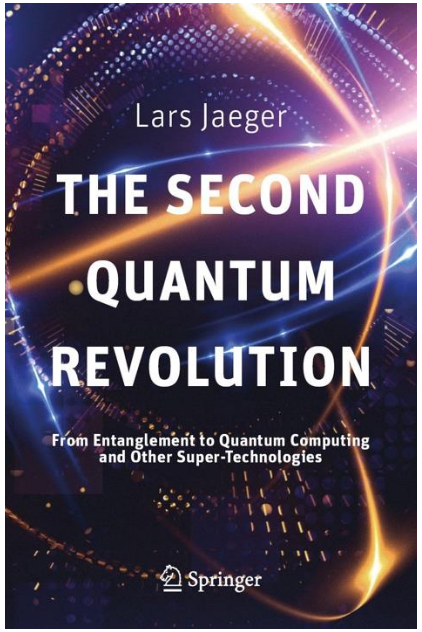

- Q sensing
- Q communication
- Q simulation
- Q computing (quantum supremacy/advantage, exponential speedup)

https://www.slideshare.net/BrunoFedrici/the-second-quantum-revolution-the-world-beyond-binary-0-and-1

QUANTUM FOR REAL-WORLD IMPACT

PHASE | Active

XPRIZE QUANTUM APPLICATIONS | Google Quantum AI

Presenting Partner: gesda

OVERVIEW ACTIVITY SPONSORS PARTNERS GUIDE

QUANTUM FOR REAL-WORLD IMPACT

$5 M

Prize Purse

Competing teams will develop new applications for quantum computers that can address complex, global challenges in climate, sustainability, health and beyond. This is a pivotal time for quantum computing; join us in shaping a future where technology meets the world's most pressing challenges head-on.

XPRIZE Quantum Applications is a 3-year, $5M global competition designed to generate quantum computing (QC) algorithms that can be put into practice to help solve real-world challenges.

https://www.xprize.org/prizes/qc-apps

What Makes Quantum Computing More Powerful?

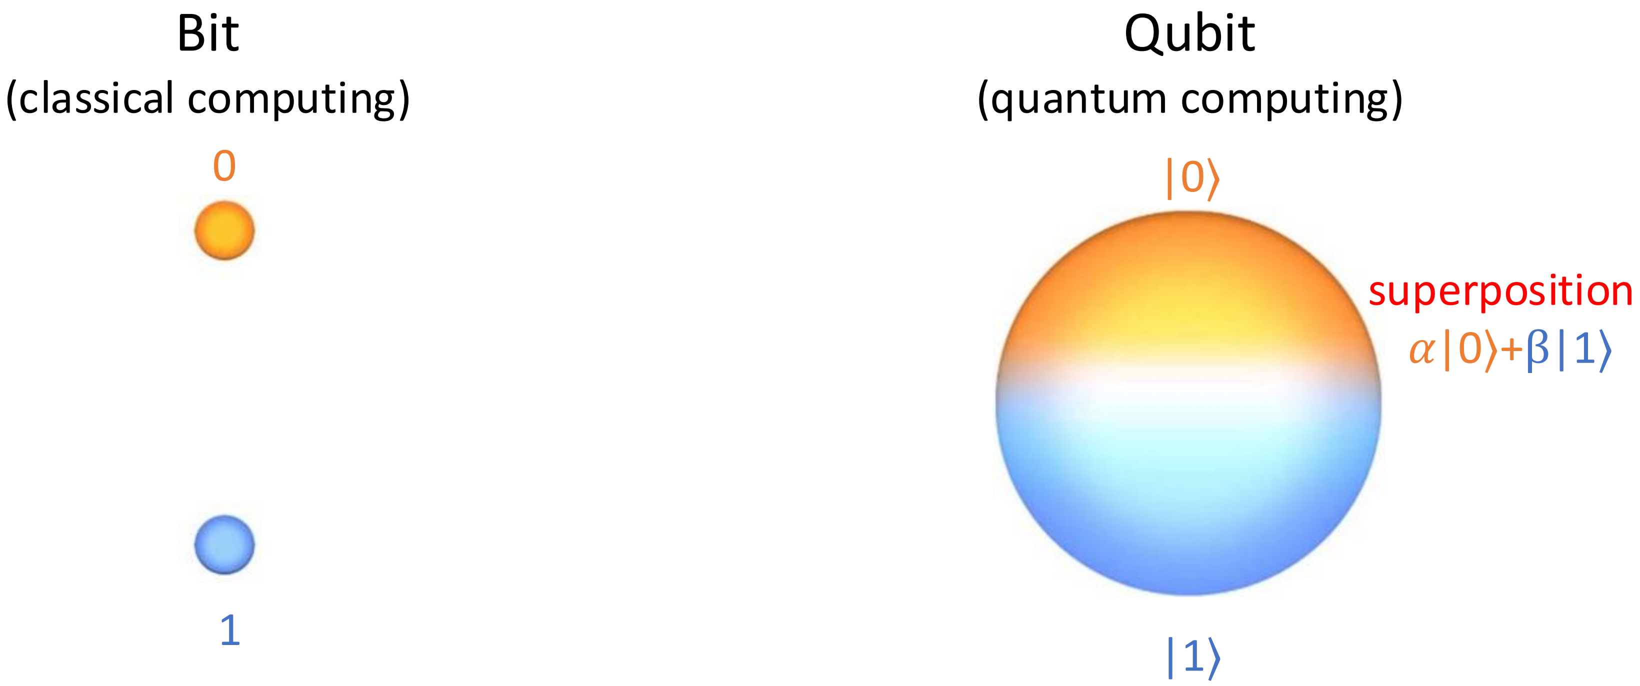

- Pros: superposition gives the ability to perform computations over classical computers
  - several qubits in superposition can crunch through a vast number of potential outcomes simultaneously
- Cons: measure the state of a quantum system is much harder

Image Credit: https://www.indiatechonline.com/viewimage.php?id=1839 https://www.youtube.com/watch?v=RCj_BJ6BddM

Learn quantum world using classical machines

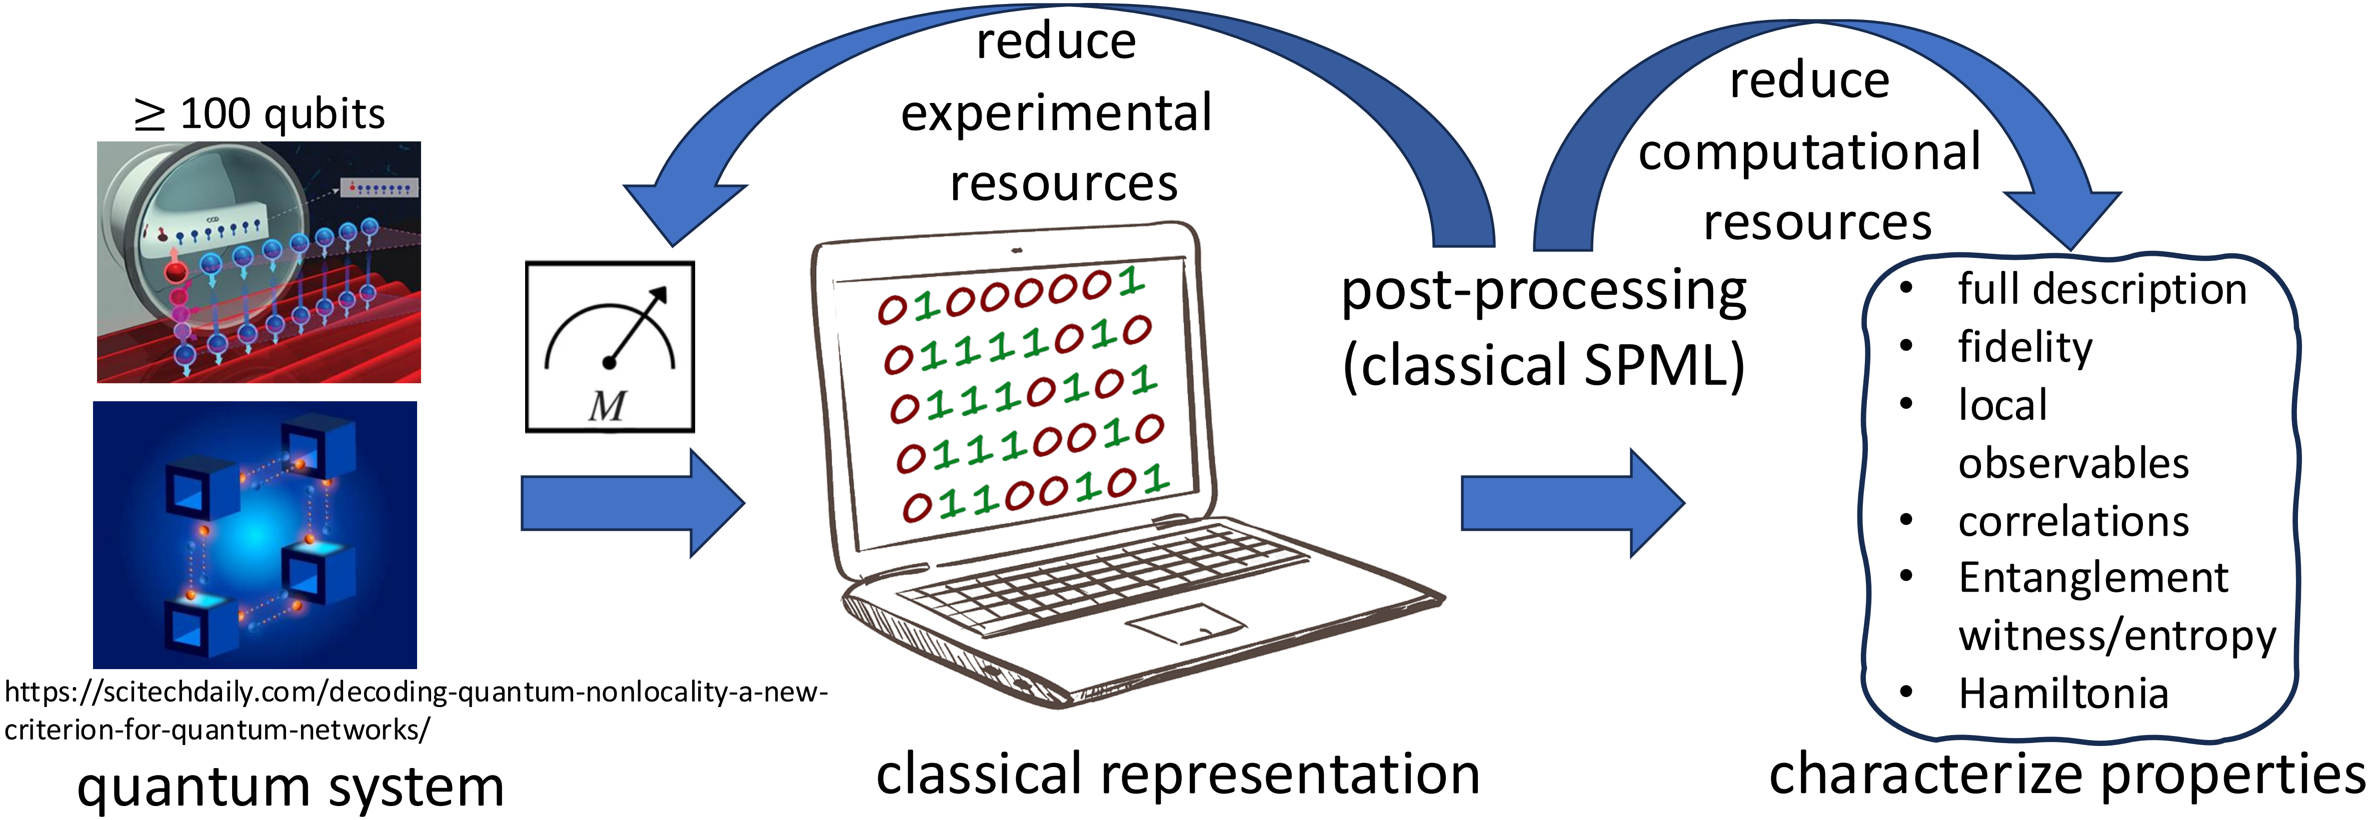

- Why do we want to construct classical representations of quantum systems?
  - Knowing what the physical system is
  - Many quantum applications rely on hybrid classical-quantum algorithms (an interface between the two)
- Exploit classical SPML for enhancing our ability to learn quantum world

Huang, H.-Y., Kueng, R. & Preskill, J. Predicting many properties of a quantum system from very few measurements. Nature Physics 2020.

Caveat

- The rest of this talk is to describe the basic ideas of quantum state learning or estimation to folks with EECS background, but maybe don't know quantum
- There will be ≈ 0 physical intuition provided in the talk
- Instead, we will illustrate from mathematical perspective (particularly, linear algebra + probability)

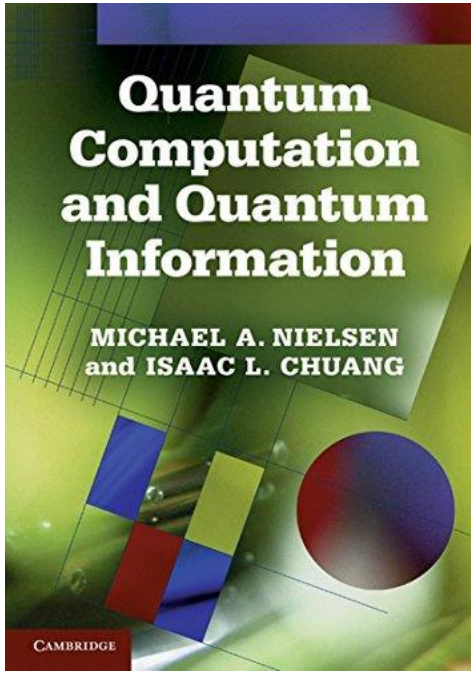

Quantum state learning/estimation

- Quantum state learning: how can we learn about quantum objects
  - called learning in quantum and TCS, e.g., FOCS 2024 Workshop: Recent Advances in Quantum Learning
  - called estimation or inverse problem in signal processing

- Set up an experiment that can produce copies of a quantum state ρ at a time

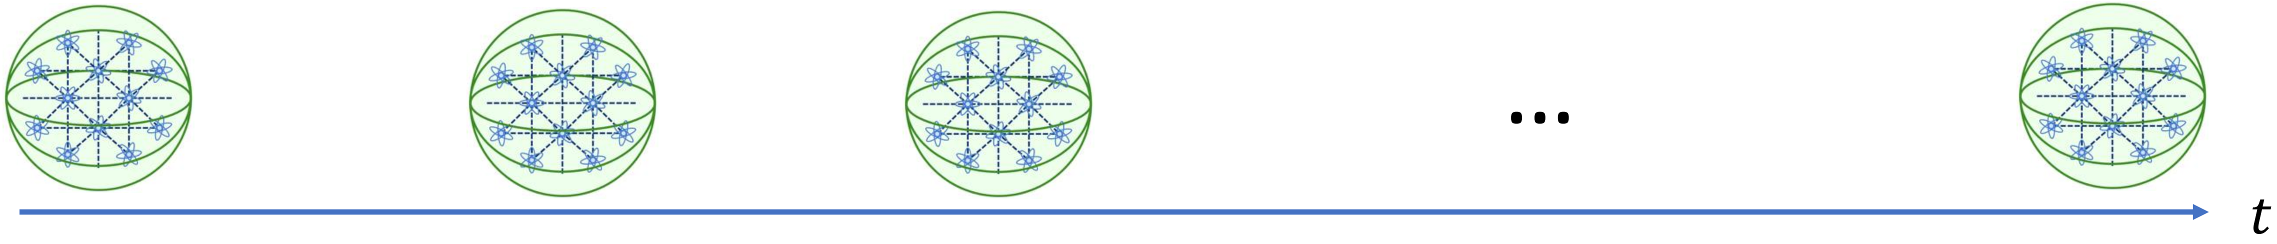

- Sample complexity: how many copies of ρ needed to learn about it?

Quantum measurement

- An n-qudit quantum system is described by a density matrix $\boldsymbol{\rho} \in \mathbb{C}^{d^n \times d^n}$

$$\boldsymbol{\rho} \succcurlyeq 0, \qquad \mathrm{trace}(\boldsymbol{\rho}) = 1$$

- One does not simply sample from a quantum state
- To interact with a quantum state, one must specify a measurement $\{\boldsymbol{A}_1, \boldsymbol{A}_2, \ldots, \boldsymbol{A}_K\}$ such that
  – $\boldsymbol{A}_K \succcurlyeq 0$
  – $\boldsymbol{A}_1 + \cdots + \boldsymbol{A}_K = \mathbf{I}$
  – positive operator-valued measure (POVM)

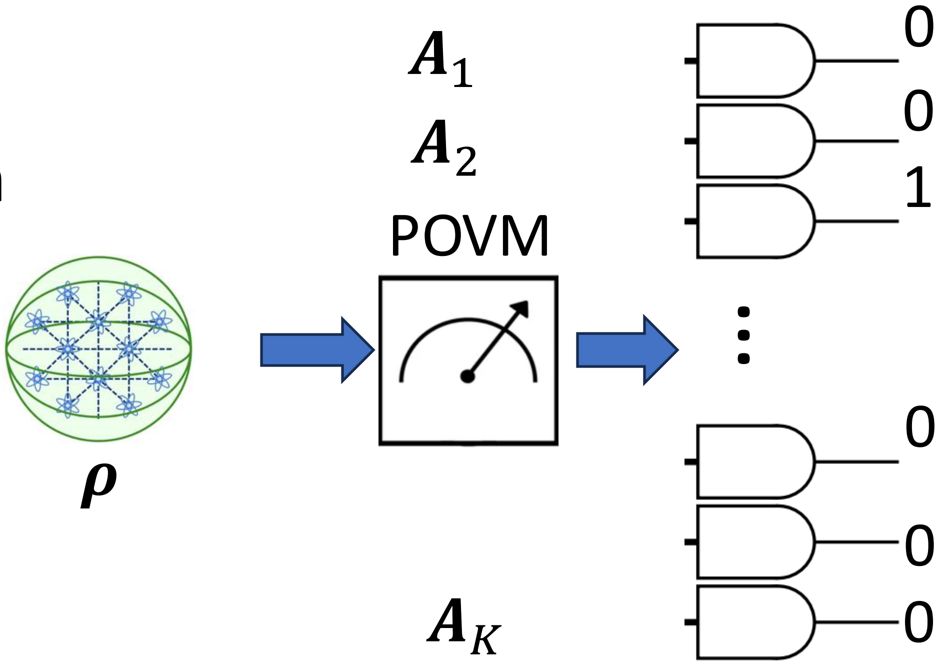

- Born's rule: The outcome of measuring $\boldsymbol{\rho}$ using this POVM is a random one-hot vector; we observe $\boldsymbol{A}_k$ with probability $p_k = \mathrm{trace}(\boldsymbol{A}_k \boldsymbol{\rho})$.
- Additionally, measuring a state collapses the state

Quantum measurement

- Often repeat this process many times on different copies of $\boldsymbol{\rho}$

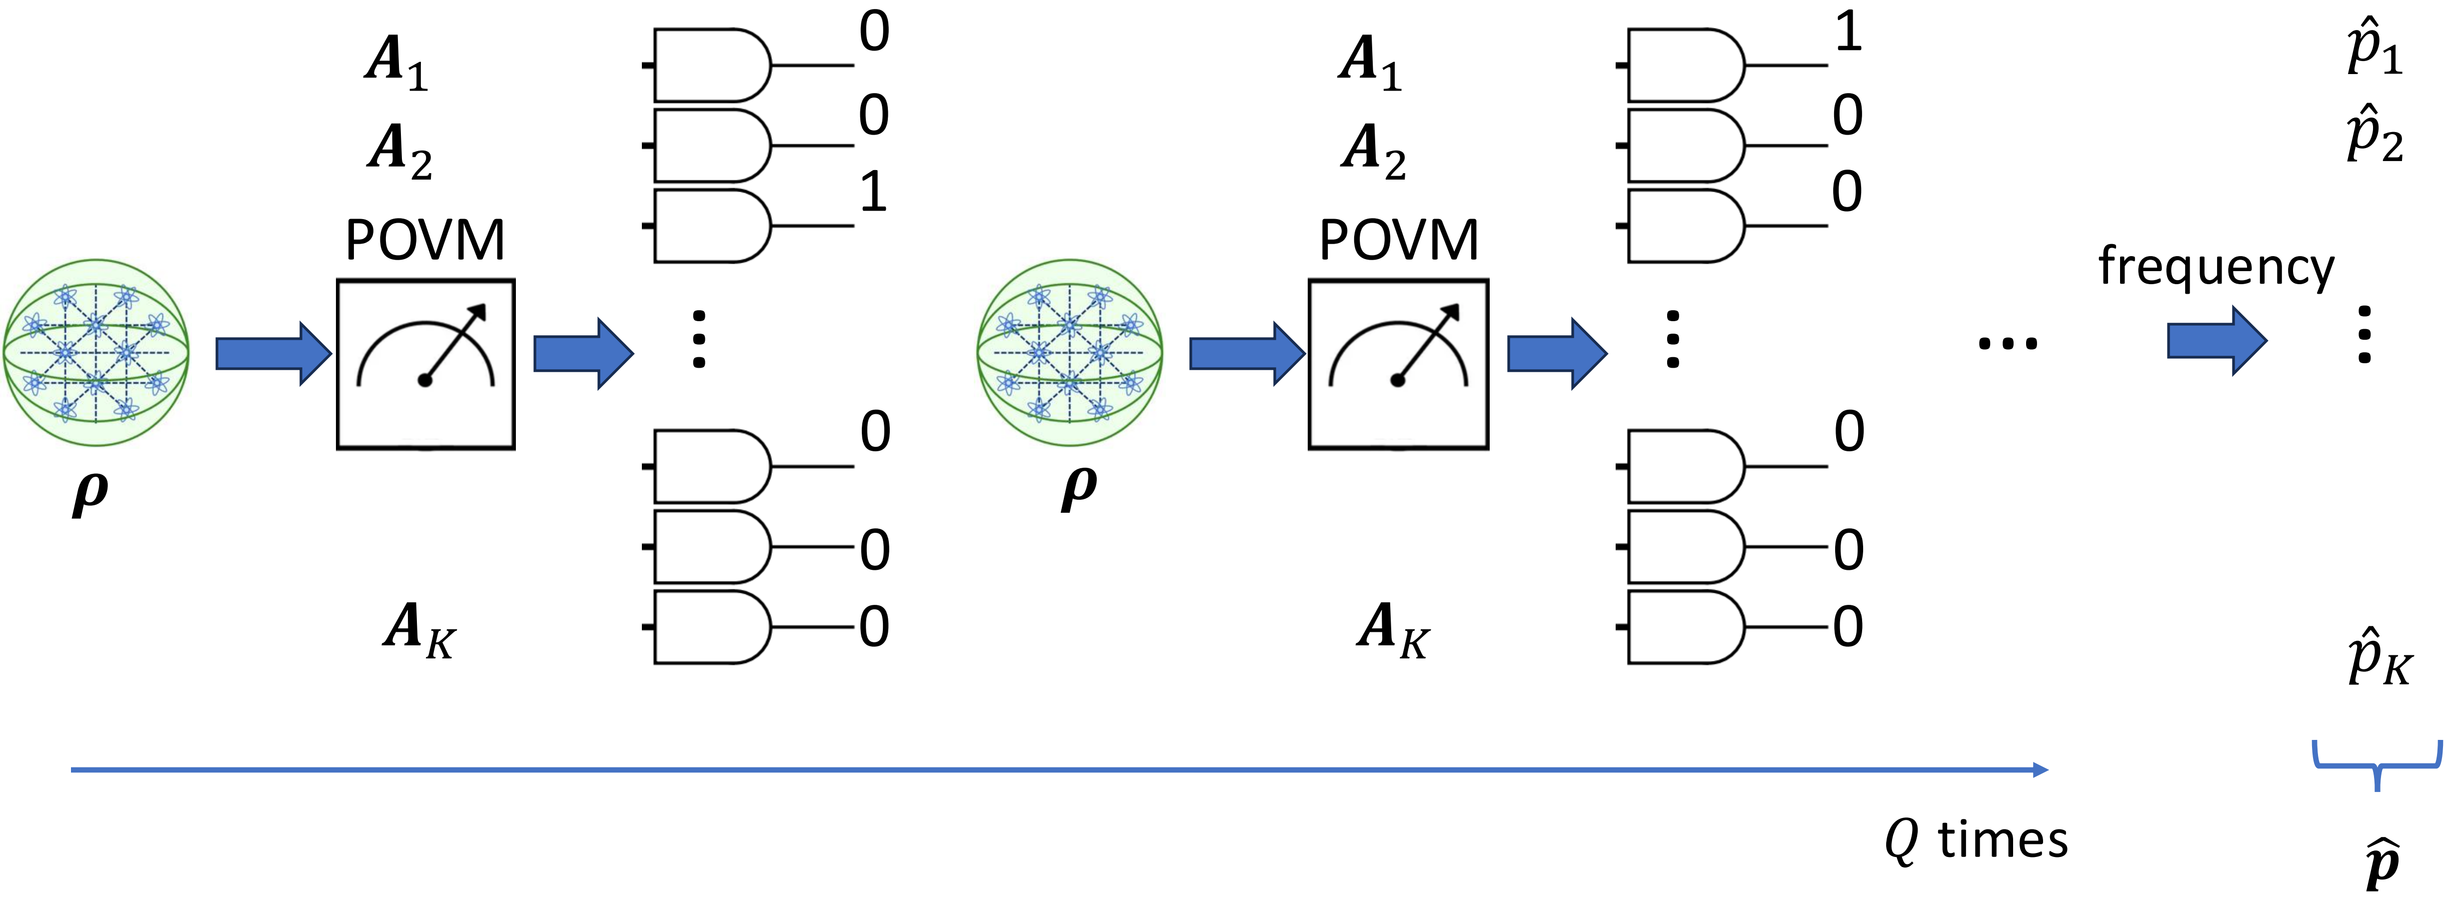

- Goal: learn the properties of the state $\boldsymbol{\rho}$ from $\widehat{\boldsymbol{p}}$

$$\mathbb{E}[\hat{p}_k] = p_k = \mathrm{trace}(\boldsymbol{A}_k \boldsymbol{\rho})$$

Curse of dimensionality

- An n-qudit state $\boldsymbol{\rho}$ lives in the space $\mathbb{C}^{d^n \times d^n}$ satisfies

$$\boldsymbol{\rho} \succcurlyeq 0, \qquad \text{trace}(\boldsymbol{\rho}) = 1$$

- The outcome of measuring $\boldsymbol{\rho}$ using the POVM is a random variable, i.e., we observe $\boldsymbol{A}_k$ with probability $p_k = \text{trace}(\boldsymbol{A}_k \boldsymbol{\rho})$. So the entropy of the outcome is at most $\log_2(K)$
- Thus, $Q = \text{poly}(d^n)$ sample complexity is required for learning general state

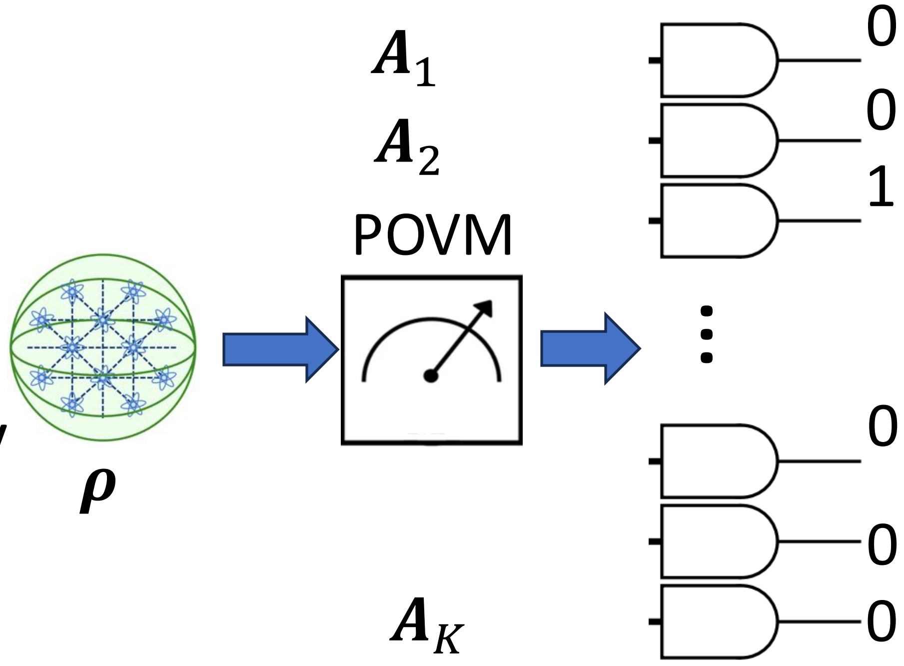

Haah, J., Harrow, A., Ji, Z., Wu, X. & Yu, N. Sample-optimal tomography of quantum states. IEEE TIT, 2017.

Curse of dimensionality

- $\mathrm{poly}(d^n)$ sample complexity is needed for learning a general quantum state
- But we may not care about learning anything about any quantum states

- **Practical states** have very low-dimensional structures
  - low-rank
  - matrix product state
  - neural quantum state

- Learning restricted properties
  - fidelity
  - correlations
  - entanglement witness
  - etc.

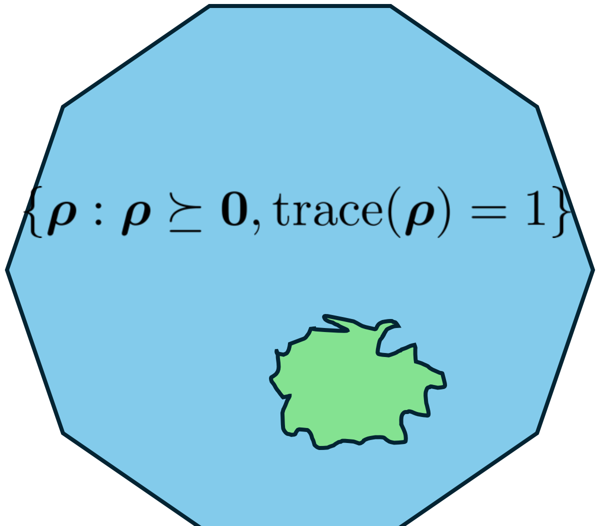

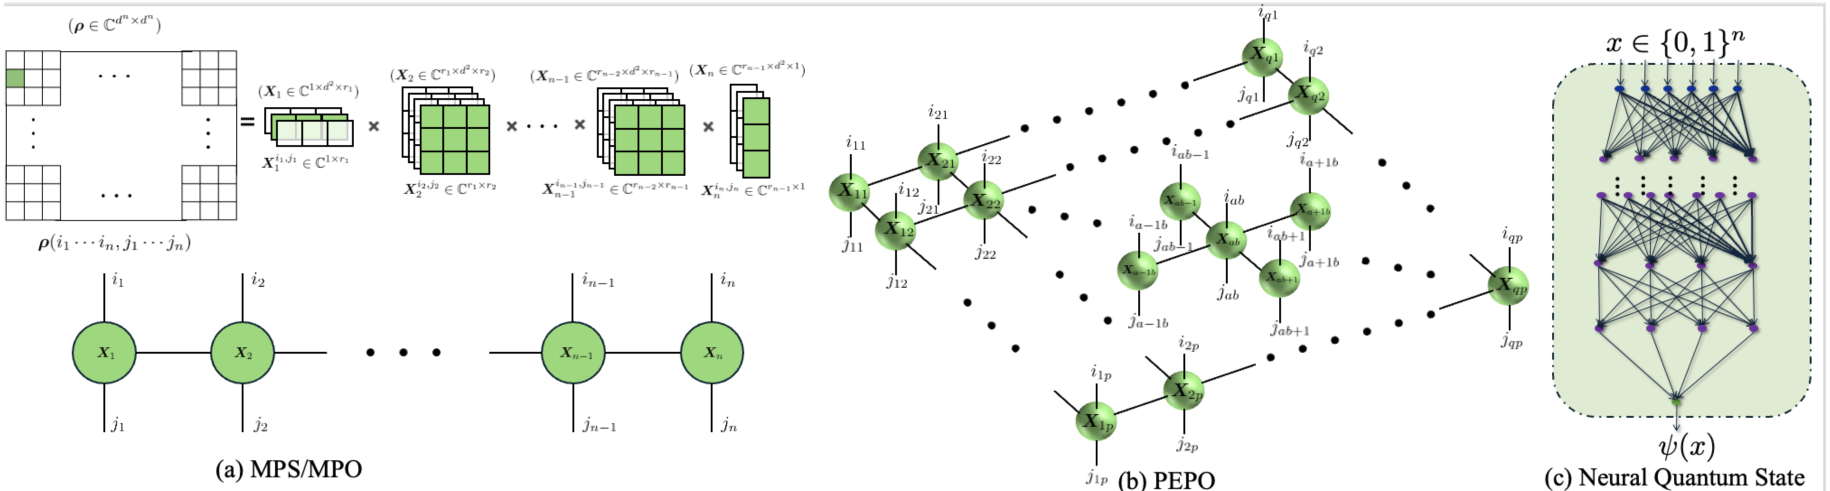

(a) MPS/MPO (b) PEPO (c) Neural Quantum State

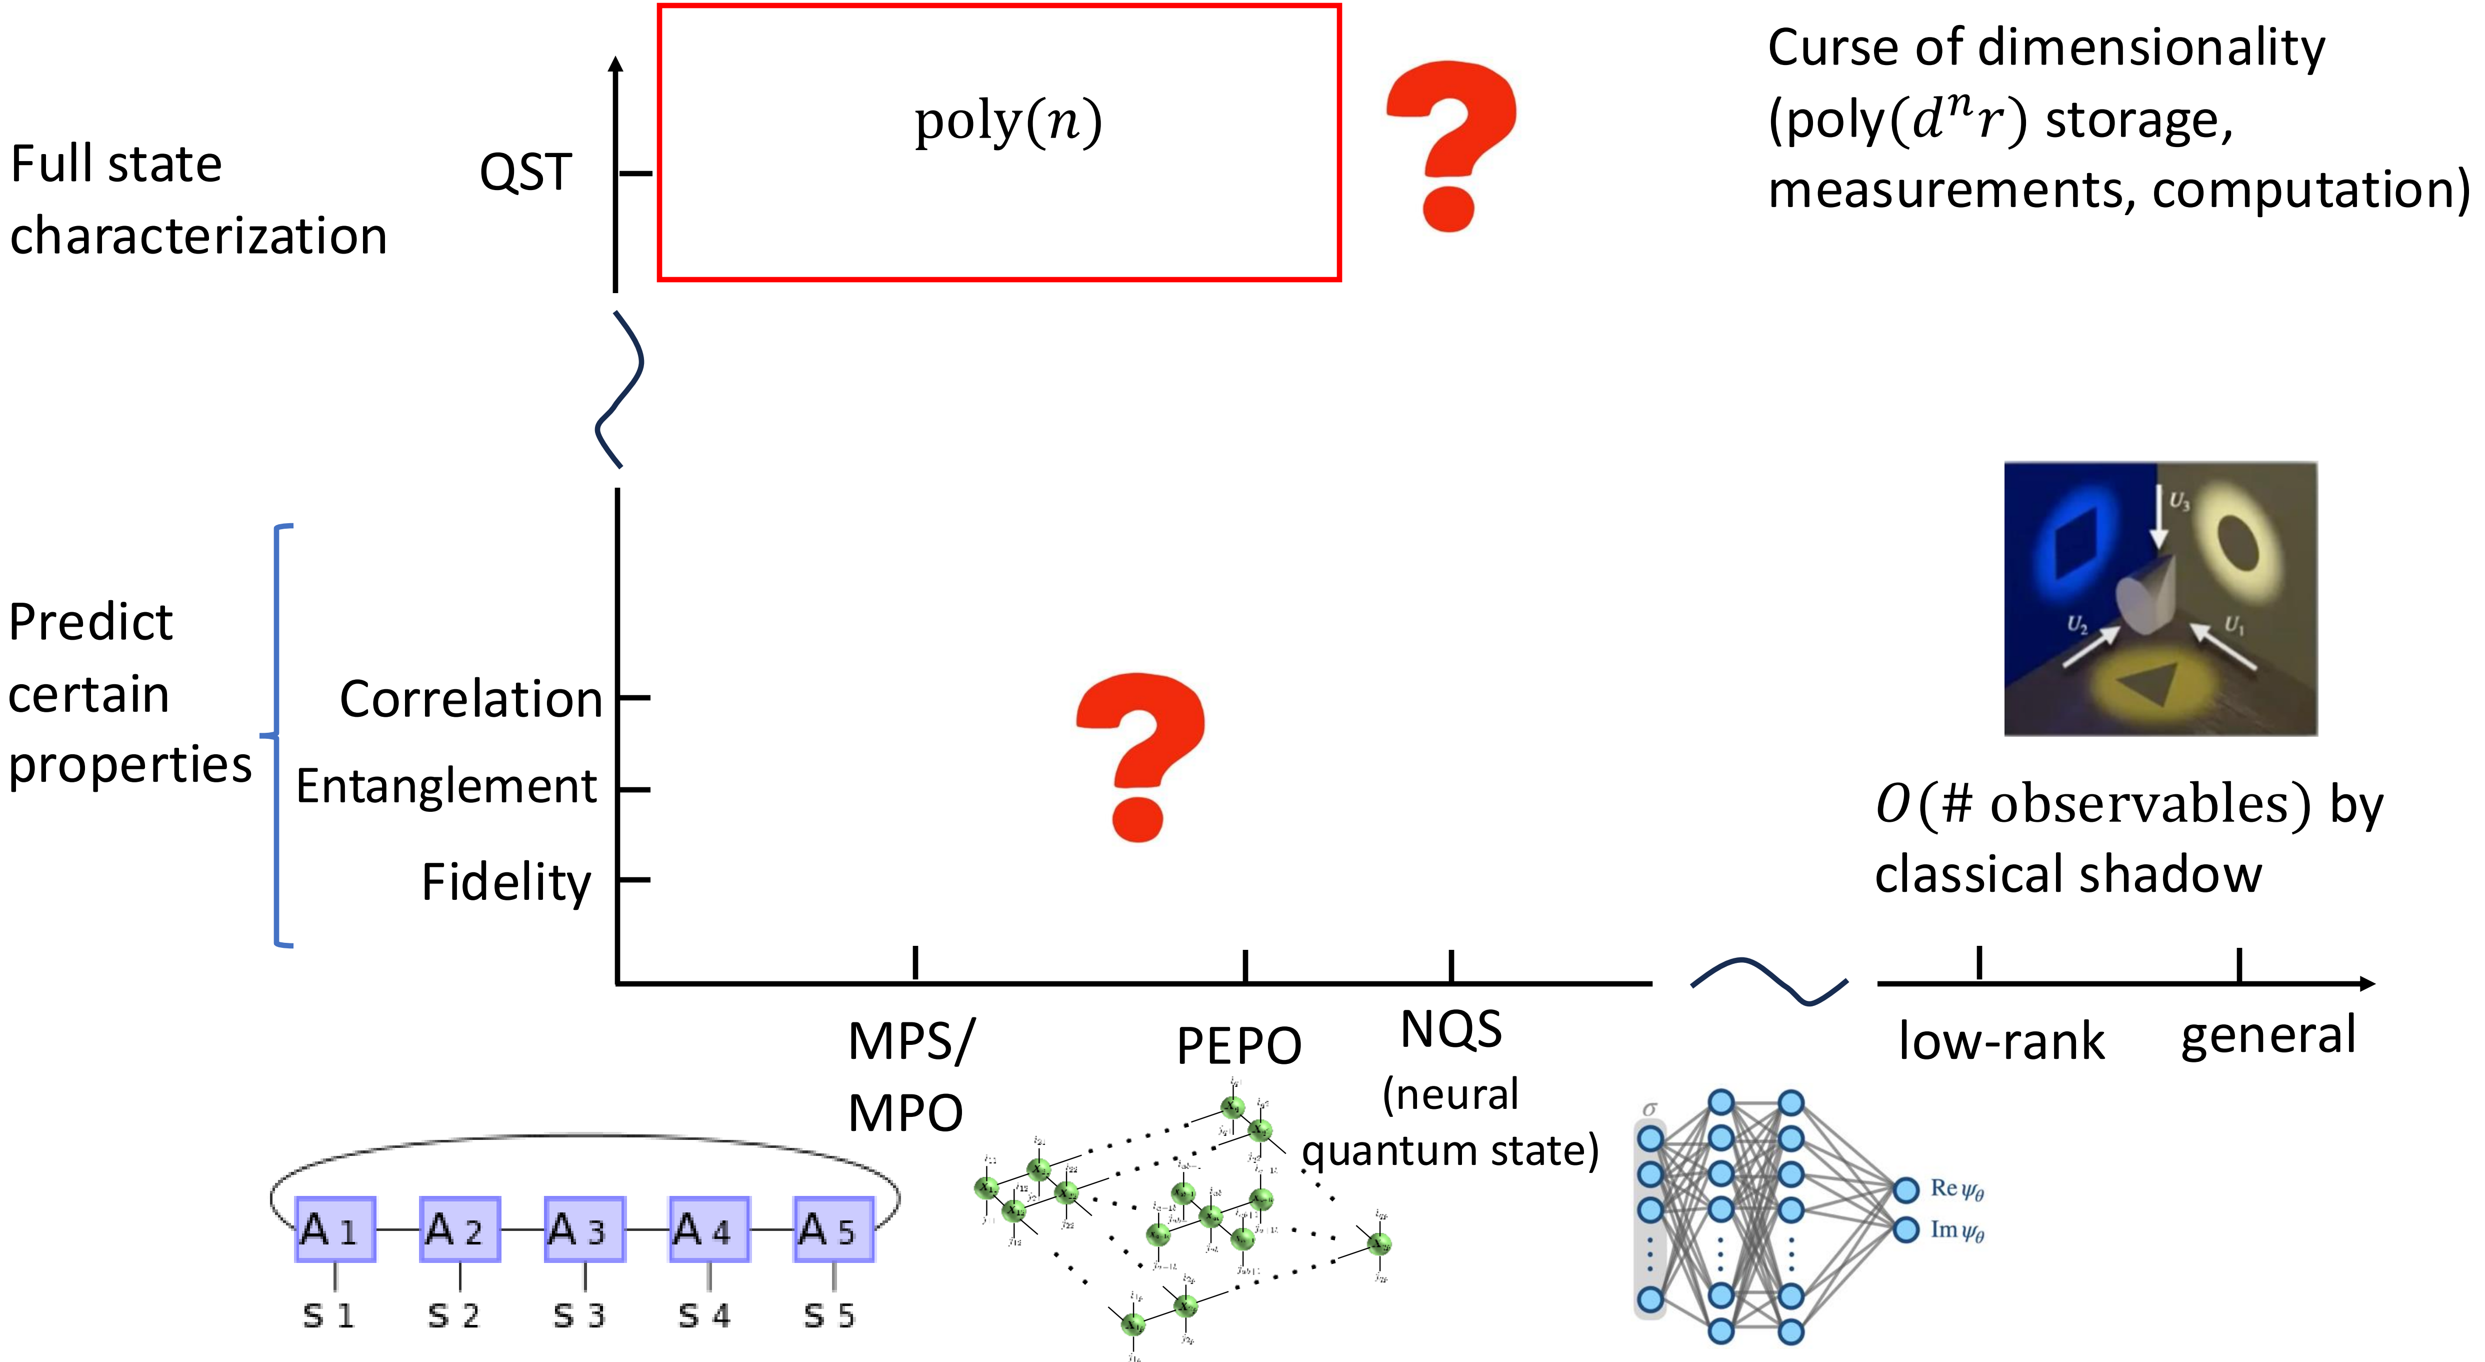

Full state characterization
QST
poly(n)
?
Curse of dimensionality
(poly($d^n r$) storage,
measurements, computation)
Predict certain properties
Correlation
Entanglement
Fidelity
?
U_3
U_2
U_1
O(# observables) by
classical shadow
MPS/
MPO
PEPO
NQS
(neural
quantum state)
low-rank
general
A1 A2 A3 A4 A5
s1 s2 s3 s4 s5
σ
$\mathrm{Re}\,\psi_\theta$
$\mathrm{Im}\,\psi_\theta$

Outline

- Statistical analysis: sampling complexity for estimating structured states
  – matrix product state/operator

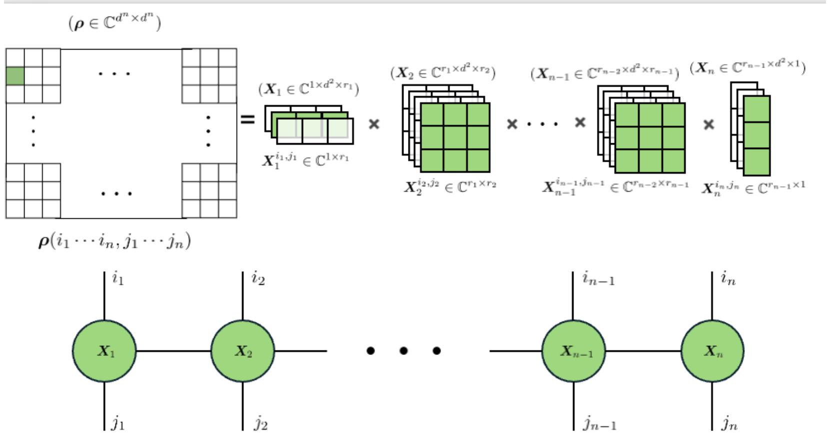

- Algorithm design: projected classical shadow for estimating structured states
  – efficient algorithm with nearly optimal sampling complexity

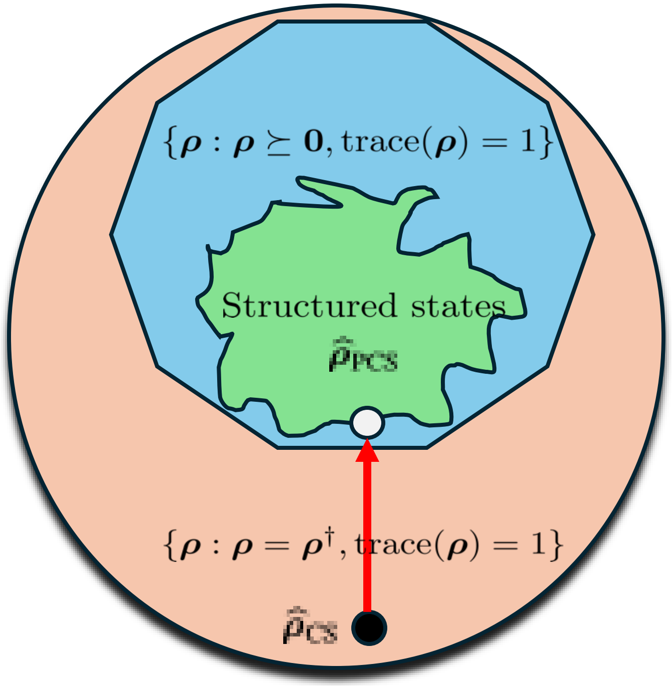

Matrix product state/operator

- A quantum state is a matrix product operator if its entries can be expressed as

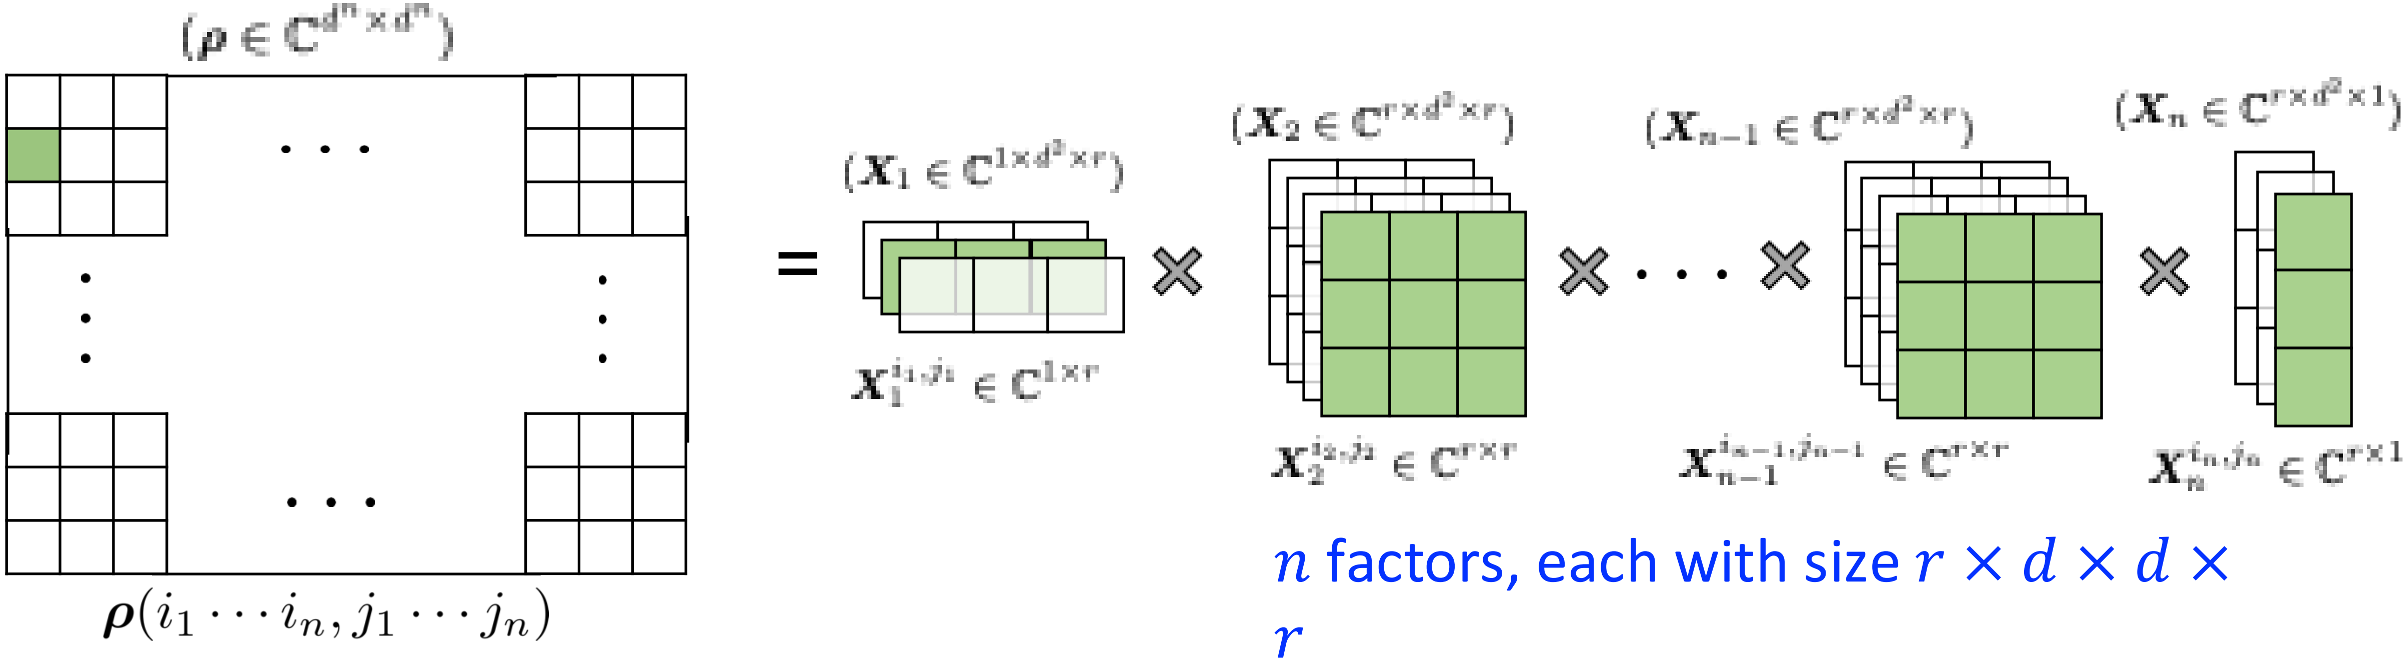

n factors, each with size $r \times d \times d \times r$

  - # parameters: $O(nd^2r^2)$, r is called the bond dimension, or MPO rank
  - The manifold of all MPO_r has dimension $(n-1)d^2r^2 + 2d^2r - (n-1)r^2$
  - Known as tensor train decomposition if reshape the state as a tensor
  - Many quantum systems with short-range interactions obey such structures

A. H. Werner et al., “Positive tensor network approach for simulating open quantum many-body systems,” PRL 2026.
Holtz, Rohwedder, Schneider. "On manifolds of tensors of fixed TT-rank." Numerische Mathematik 120.4 (2012): 701-731.

Informationally complete POVMs

- A POVM is informationally complete (IC) if it consists of $K \geq d^{2n}$ matrices and can form d^{2n} linearly independent matrices by linear combination.
- The induced linear mapping $\mathcal{A}$ is invertible
- With a good estimation of the probability distribution $\widehat{\boldsymbol{p}}$, we get $\boldsymbol{\rho} = \mathcal{A}^{-1}(\widehat{\boldsymbol{p}})$
- It could be hard to implement physically
- Often it is useful to study the sample complexity

ideal case: measure the state ∞ times

$\boldsymbol{A}_1$ $\quad p_1 = \mathrm{trace}(\boldsymbol{A}_1\boldsymbol{\rho})$

$\boldsymbol{A}_2$ $\quad p_2 = \mathrm{trace}(\boldsymbol{A}_2\boldsymbol{\rho})$

POVM

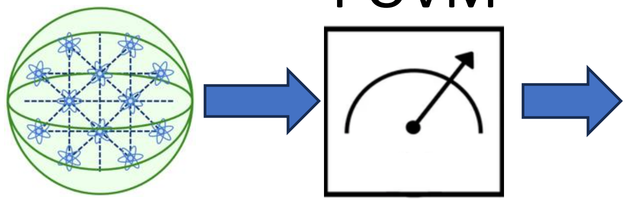

$\boldsymbol{\rho}$

$\vdots$

$\boldsymbol{A}_K$ $\quad p_K = \mathrm{trace}(\boldsymbol{A}_K\boldsymbol{\rho})$

$\boldsymbol{p} = \mathcal{A}(\boldsymbol{\rho})$

Informationally complete POVMs

A finite set of normalized vectors $\boldsymbol{w}_1, \dots, \boldsymbol{w}_K \in \mathbb{C}^{d^n}$ is called a spherical quantum t-design if

$$\frac{1}{K}\sum_{k=1}^{K}\left(\boldsymbol{w}_k\boldsymbol{w}_k^{\mathrm{H}}\right)^{\otimes s} = \int\left(\boldsymbol{w}\boldsymbol{w}^{\mathrm{H}}\right)^{\otimes s} d\boldsymbol{w}, \forall s \leq t,$$

where the integral is taken with respect to the Haar measure on the complex unit sphere.

- The induced t-design POVM $\{\boldsymbol{A}_k = \boldsymbol{w}_k\boldsymbol{w}_k^{\mathrm{H}}\}$ is IC and satisfies

$$||\mathcal{A}(\boldsymbol{\rho})||_2^2 \approx \frac{||\boldsymbol{\rho}||_F^2 + \left(\mathrm{trace}(\boldsymbol{\rho})\right)^2}{K}$$

for any Hermitian matrices $\boldsymbol{\rho}$.

Sample complexity with IC POVMs

Theorem (informal) Suppose $\{\boldsymbol{A}_k\}$ forms a t-design POVM ($t \geq 3$). For an MPO state $\boldsymbol{\rho}^\star$ in MPO_r, we measure the state Q times using the POVM. Then any global solution of the constrained least-squares estimator

$$\widehat{\boldsymbol{\rho}} = \underset{\boldsymbol{\rho}:\,\mathrm{MPO}_r}{\mathrm{argmin}} ||\mathcal{A}(\boldsymbol{\rho}) - \widehat{\boldsymbol{p}}||_2^2$$

satisfies

$$||\widehat{\boldsymbol{\rho}} - \boldsymbol{\rho}^\star||_F \leq \epsilon$$

with high probability as long as

$$Q \gtrsim \frac{nd^2r^2(\text{log factors})}{\epsilon^2}$$

- It matches the DOF in MPOs and is nearly optimal (if ignore the log factors).
- This supports the use of low-Q measurement schemes.

Haar random projective measurements

- Moving toward more practical measurements, consider a randomly generated Haar-distributed unitary matrix $\boldsymbol{U} = [\boldsymbol{u}_1 \ \cdots \ \boldsymbol{u}_{d^n}]$. With this, we can form projective measurement with POVM
$$\boldsymbol{A}_1 = \boldsymbol{u}_1\boldsymbol{u}_1^{\mathrm{H}},\ \boldsymbol{A}_2 = \boldsymbol{u}_2\boldsymbol{u}_2^{\mathrm{H}}, \dots,\ \boldsymbol{A}_{d^n} = \boldsymbol{u}_{d^n}\boldsymbol{u}_{d^n}^{\mathrm{H}}$$

- The probability of observing k-th outcome is
$$p_k = \mathrm{trace}(\boldsymbol{A}_k\boldsymbol{\rho}) = \boldsymbol{u}_k^{\mathrm{H}}\boldsymbol{\rho}\boldsymbol{u}_k = \boldsymbol{e}_k^{\mathrm{T}}\left(\boldsymbol{U}^{\mathrm{H}}\boldsymbol{\rho}\boldsymbol{U}\right)\boldsymbol{e}_k$$
- Apply a random unitary matrix $\boldsymbol{U}$ to rotate the the state and then perform a computational-basis measurement with $\{\boldsymbol{e}_k\}$
  - A universal quantum computer can approximately generate such random unitary to any given precision, though the complexity varies

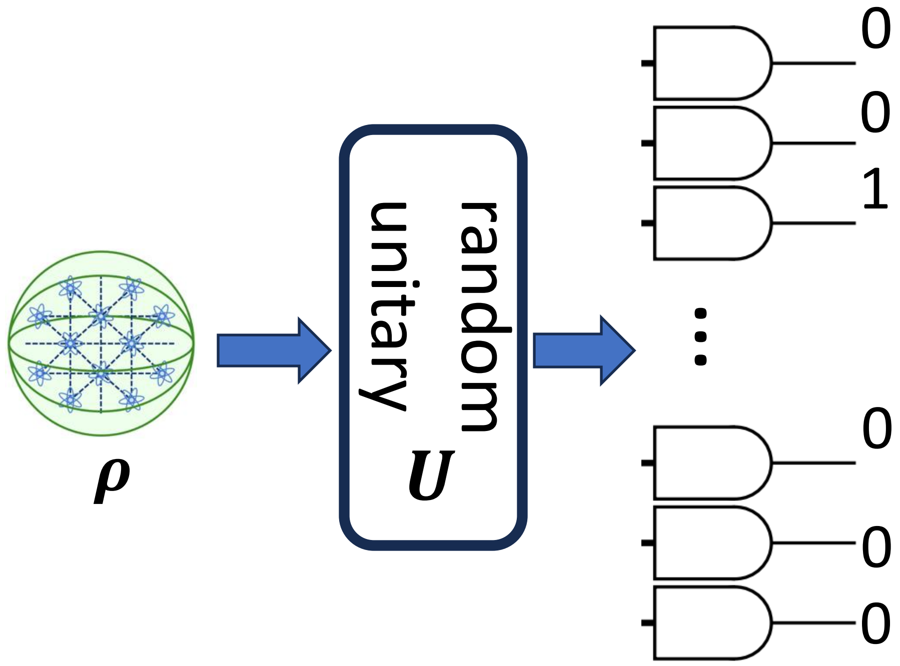

Haar random projective measurements

- Moving toward more practical measurements, consider a randomly generated Haar-distributed unitary matrix $\boldsymbol{U} = [\boldsymbol{u}_1 \ \cdots \ \boldsymbol{u}_{d^n}]$. With this, we can form projective measurement with POVM

$$\boldsymbol{A}_1 = \boldsymbol{u}_1 \boldsymbol{u}_1^{\mathrm{H}}, \ \boldsymbol{A}_2 = \boldsymbol{u}_2 \boldsymbol{u}_2^{\mathrm{H}}, \dots, \ \boldsymbol{A}_{d^n} = \boldsymbol{u}_{d^n} \boldsymbol{u}_{d^n}^{\mathrm{H}}$$

- This POVM is not IC. So in general, multiple such POVMs are needed.

Theorem (informal) Aggregating measurements from M such unitary bases, then with

$$M \gtrsim O(nd^2r^2 \log n)$$

the linear map $\mathcal{A}$ satisfies one-side RIP $||\mathcal{A}(\boldsymbol{\rho})||_2^2 \gtrsim \sqrt{\frac{M}{d^n}}$ for all MPO_r.

- Ensure distinct measurements ($\mathcal{A}(\boldsymbol{\rho}_1) \neq \mathcal{A}(\boldsymbol{\rho}_2)$) as long as $\boldsymbol{\rho}_1 \neq \boldsymbol{\rho}_2$

Sample complexity with Haar random measurements

Theorem (informal) Aggregating measurements from M such Haar-distributed rank-one POVMs, with Q state copies per POVM, any global solution of the constrained least-squares estimator

$$\widehat{\boldsymbol{\rho}} = \underset{\boldsymbol{\rho}:\,\mathrm{MPO}_r}{\operatorname{argmin}} \|\mathcal{A}(\boldsymbol{\rho}) - \widehat{\boldsymbol{p}}\|_2^2$$

satisfies

$$\|\widehat{\boldsymbol{\rho}} - \boldsymbol{\rho}^\star\|_F \leq \epsilon$$

with high probability as long as

$$QM \gtrsim \frac{n^3 d^2 r^2 (\text{log factors})}{\epsilon^2}.$$

- n^3 could be further improved.
- This supports the use of low-QM measurement schemes.

Nonconvex Algorithms

- The constrained least-squares estimator involves a nonconvex opt problem

$$\widehat{\boldsymbol{\rho}} = \operatorname*{argmin}_{\boldsymbol{\rho}:\,\mathrm{MPO}_r} ||\mathcal{A}(\boldsymbol{\rho}) - \widehat{\boldsymbol{p}}||_2^2$$

- Projected gradient descent or iterative hard thresholding
  - no tractable algorithm for exact projection, only quasi-optimal approximation
  - we provide convergence guarantee with additional conditions for IC-PVOM
- Factorization approach: optimize over the factors $\boldsymbol{X}_1, \dots, \boldsymbol{X}_n$

$$\min_{\boldsymbol{X}_1,\dots,\boldsymbol{X}_n} ||\mathcal{A}([\boldsymbol{X}_1, \dots, \boldsymbol{X}_n]) - \widehat{\boldsymbol{p}}||_2^2$$

  - transfer the nonconvex set to nonconvex optimization over the factors
  - reduced memory
  - we establish local linear convergence for Gaussian measurement ensembles $\mathcal{A}$

Z. Qin, C. Jameson, A. Goldar, M. Wakin, Z. Gong, Z. Zhu, ‘‘Sample-Optimal Quantum State Tomography for Structured Quantum States in One Dimension," 2024.
Z. Qin, M. B. Wakin, and Z. Zhu, ‘‘Guaranteed Nonconvex Factorization Approach for Tensor Train Recovery," JMLR 2025.

Experiments

- Pure state MPS; $M = 1$; iterative hard thresholding

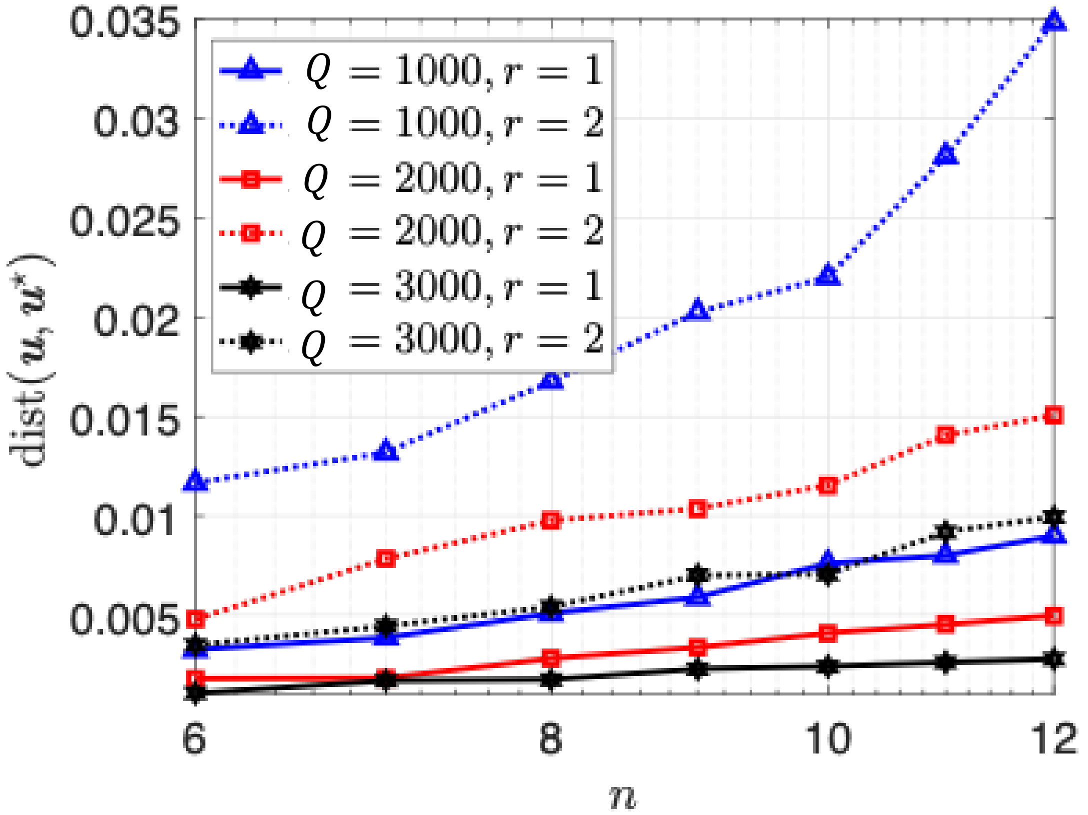

- Error increases with r, decreases with Q, increases only polynomially with n

Outline

- Statistical analysis: sampling complexity for estimating structured states
  – matrix product state/operator

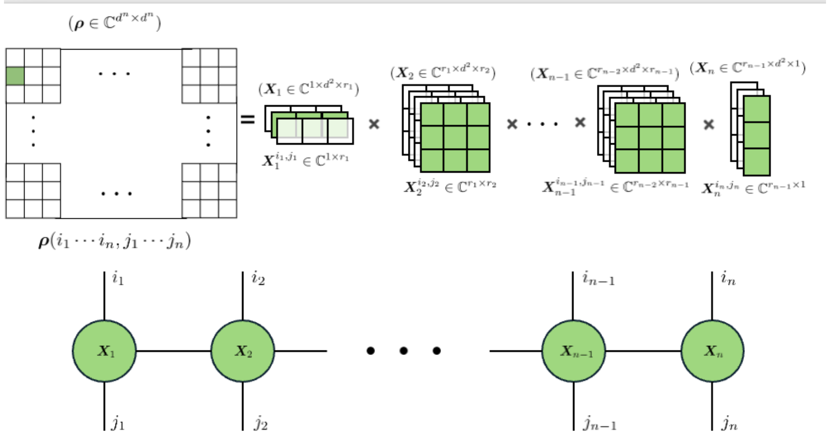

- Algorithm design: projected classical shadow for estimating structured states
  – efficient algorithm with nearly optimal sampling complexity

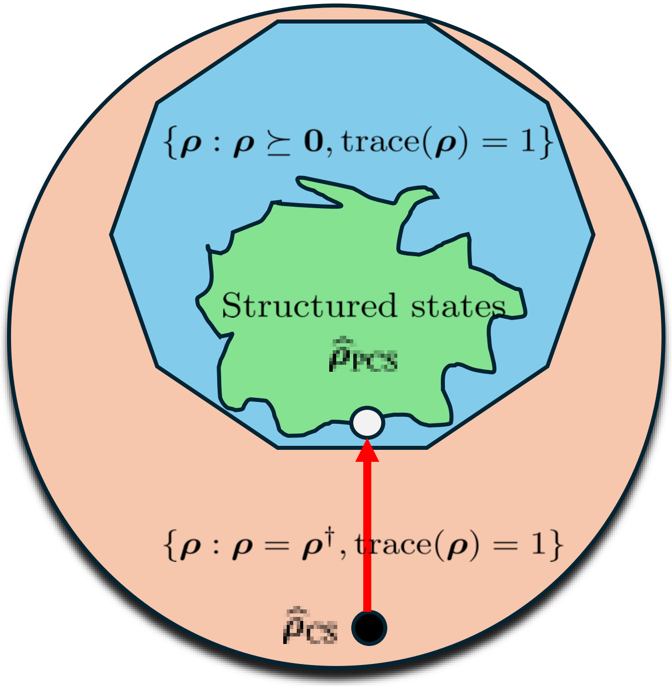

Classical shadow

Repeat the following single-shot measurement M times [HKP'20]:

- Generate a random unitary matrix $\boldsymbol{U_i}$ to rotate the quantum system
- Measure the system in the computational basis once to get $\widehat{\boldsymbol{p}}_i \in \{0,1\}^n$
- Store the "classical shadow": $\boldsymbol{\rho}_i = (d^n + 1)\underbrace{\boldsymbol{U_i}\widehat{\boldsymbol{p}}_i}_{\text{one column of } \boldsymbol{U_i}}(\boldsymbol{U_i}\widehat{\boldsymbol{p}}_i)^{\mathrm{H}} - \mathbf{I}_{d^n}$

$$\begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}$$

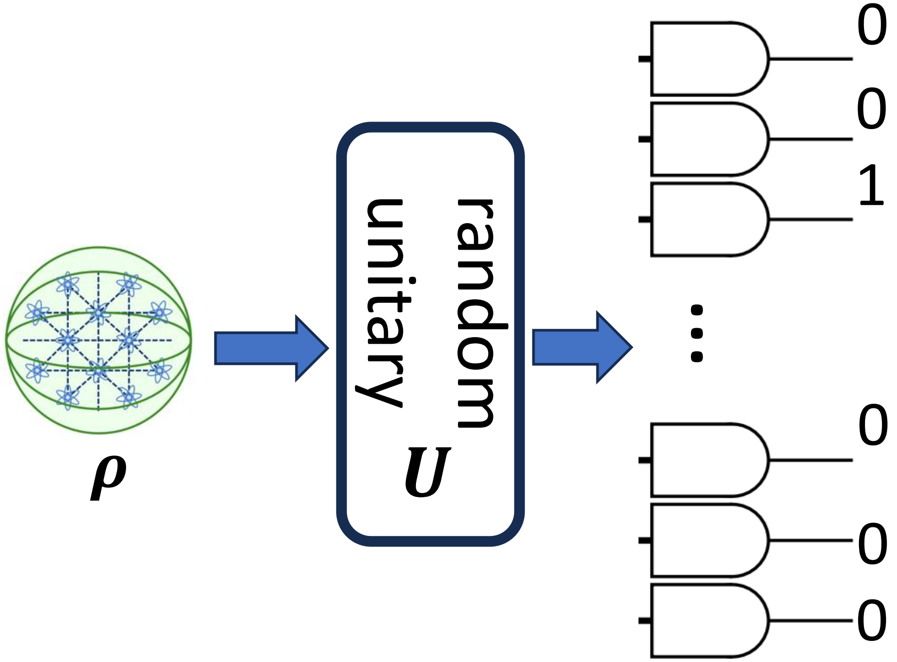

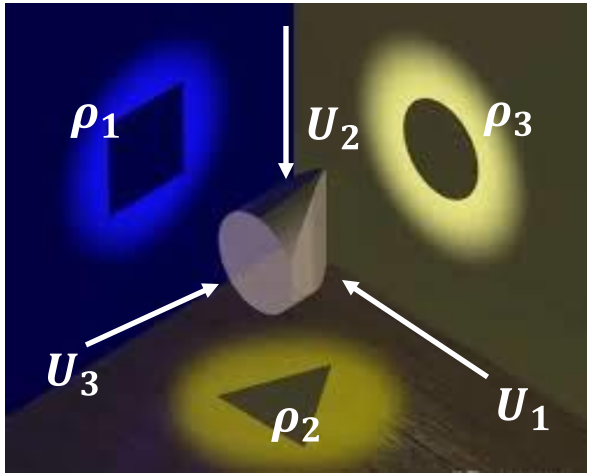

Huang, H.-Y., Kueng, R. & Preskill, J. Predicting many properties of a quantum system from very few measurements. Nature Physics 2020.

Classical shadow

Repeat the following single-shot measurement M times:

- Generate a random unitary matrix $\boldsymbol{U_i}$ to rotate the quantum system
- Measure the system in the computational basis once to get $\widehat{\boldsymbol{p}}_i \in \{0,1\}^n$
- Store the "classical shadow": $\boldsymbol{\rho}_i = (d^n + 1)\underbrace{\boldsymbol{U_i}\widehat{\boldsymbol{p}}_i}_{\text{one column of } \boldsymbol{U_i}}(\boldsymbol{U_i}\widehat{\boldsymbol{p}}_i)^{\mathrm{H}} - \mathbf{I}_{d^n}$

$$\begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}$$

- $\boldsymbol{\rho}_i$ is not a physical state as it has eigenvalues $\sigma_1 = d^n, \sigma_2 = \cdots = \sigma_{d^n} = -1$
- full state, regardless of the ground-truth state $\boldsymbol{\rho}$
- it is random in $\boldsymbol{U_i}$ and $\widehat{\boldsymbol{p}}_i | \boldsymbol{U_i}$
- it is an unbiased estimator of $\boldsymbol{\rho}$

$$\mathbb{E}_{\boldsymbol{U_i},\widehat{\boldsymbol{p}}_i}[\boldsymbol{\rho}_i] = \boldsymbol{\rho}$$

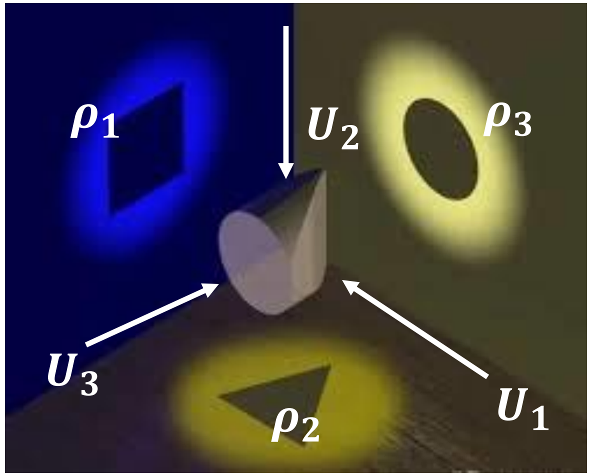

Classical shadow for regularizing LS

- Consider a least-squares estimator
  $$\widehat{\boldsymbol{\rho}}_{\mathrm{LS}} = \underset{\boldsymbol{\rho}' \in \mathbb{C}^{d^n \times d^n}}{\operatorname{argmin}} \, ||\widehat{\boldsymbol{p}} - \mathcal{A}(\boldsymbol{\rho}')||^2$$
  – without constraints, leads to a simple closed-form solution
  $$\widehat{\boldsymbol{\rho}}_{\mathrm{LS}} = (\mathcal{A}^*\mathcal{A})^{+}(\mathcal{A}^*(\widehat{\boldsymbol{p}}))$$
  $$= \frac{1}{M}\sum_i \left(\frac{1}{Q}\mathcal{A}^*\mathcal{A}\right)^{+} (\boldsymbol{U}_i\widehat{\boldsymbol{p}}_i)(\boldsymbol{U}_i\widehat{\boldsymbol{p}}_i)^{\mathrm{H}}$$
  – $(\cdot)^+$ is the pseudo-inverse
  – biases to zero when $M \ll d^n$

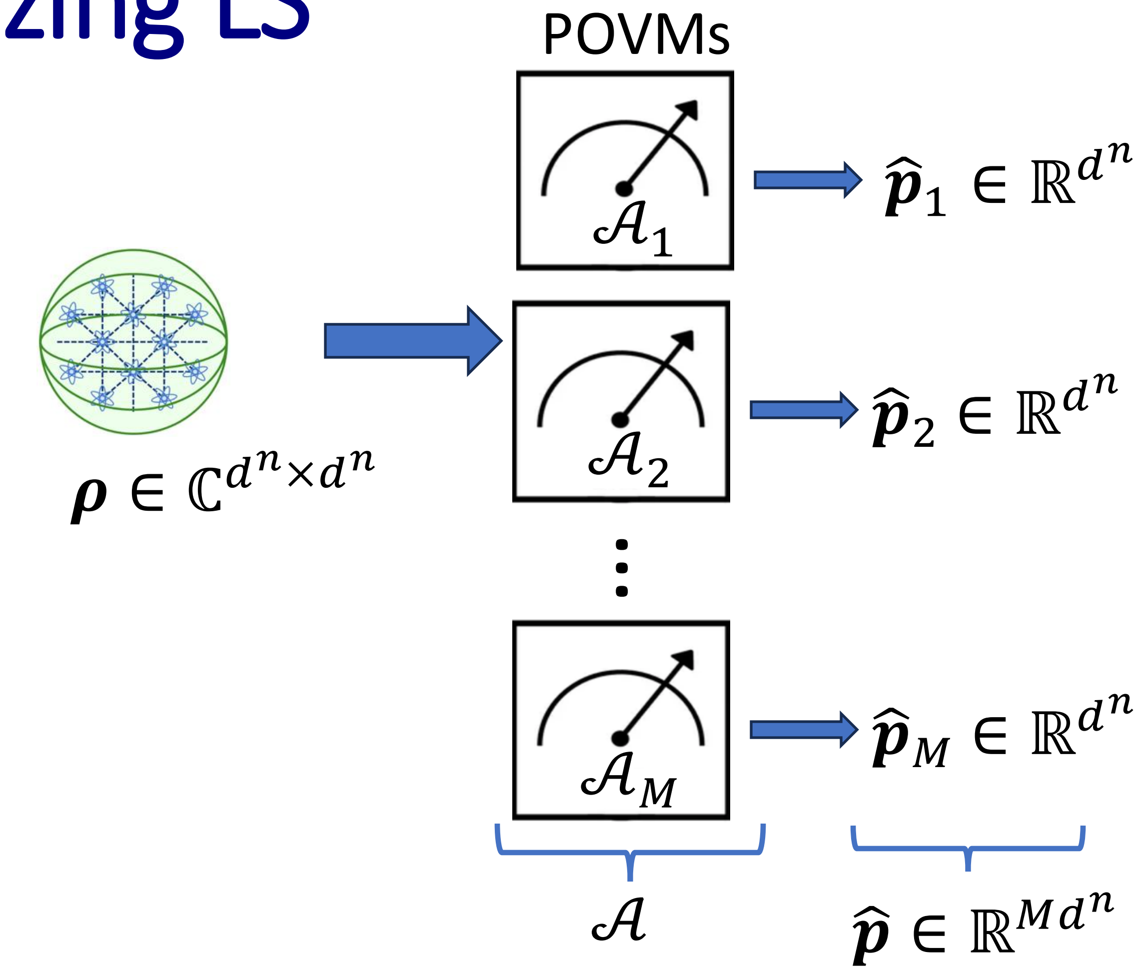

- By CLT, $\frac{1}{M}\mathcal{A}^*\mathcal{A} \to \mathbb{E}\underbrace{\left[\frac{1}{M}\mathcal{A}^*\mathcal{A}\right]}_{\text{quantum channel } \mathcal{M}}$

$$\widehat{\boldsymbol{\rho}}_{\mathrm{CS}} = \frac{1}{M}\sum_i \underbrace{\mathcal{M}^{-1}(\boldsymbol{U}_i\widehat{\boldsymbol{p}}_i(\boldsymbol{U}_i\widehat{\boldsymbol{p}}_i)^{\mathrm{H}})}_{\text{classical shadow } \boldsymbol{\rho}_i}$$

Shadows all the way down

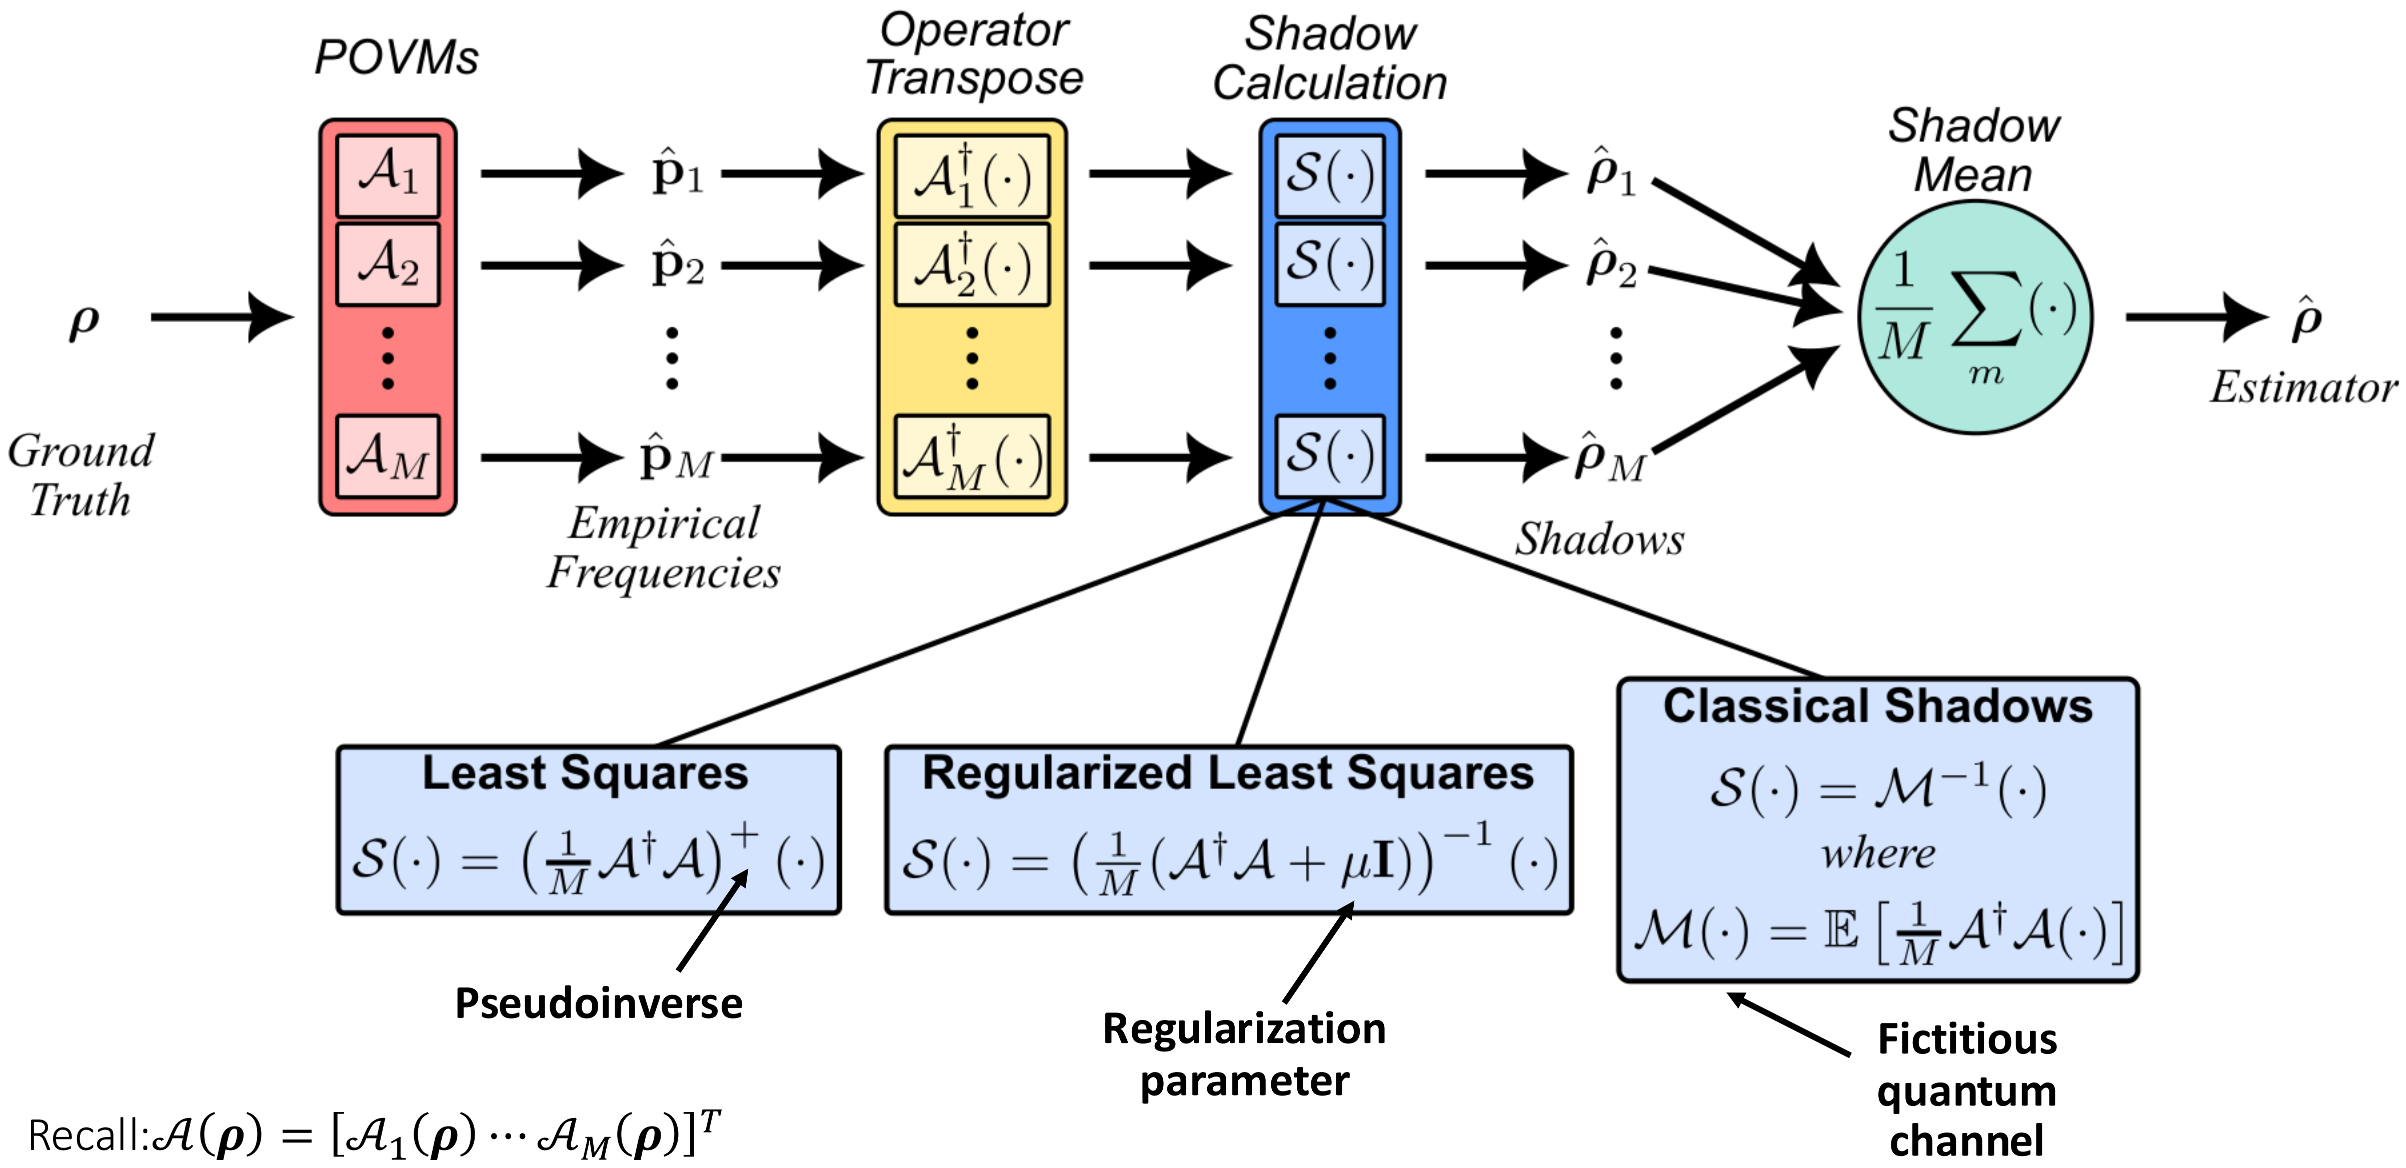

Recall: $\mathcal{A}(\boldsymbol{\rho}) = [\mathcal{A}_1(\boldsymbol{\rho}) \cdots \mathcal{A}_M(\boldsymbol{\rho})]^T$

Classical shadow vs RLS

	Bias	Variance	Distribution dependent?
LS	high	high	N
CS	low (zero)	high	Y
RLS	high	low	N

- RLS and CS trade off bias and variance differently
- Use them to estimate properties of the state $\boldsymbol{\rho}$
  – e.g., estimate the linear observable $\lambda = \mathrm{trace}(\boldsymbol{\Lambda\rho})$ by

$$\hat{\lambda} = \mathrm{trace}(\boldsymbol{\Lambda}\hat{\boldsymbol{\rho}}_{\mathrm{CS}}) = \frac{1}{M}\sum_i \mathrm{trace}(\boldsymbol{\Lambda}\boldsymbol{\rho}_{\boldsymbol{i}})$$

  – media of mean can reduce the variance

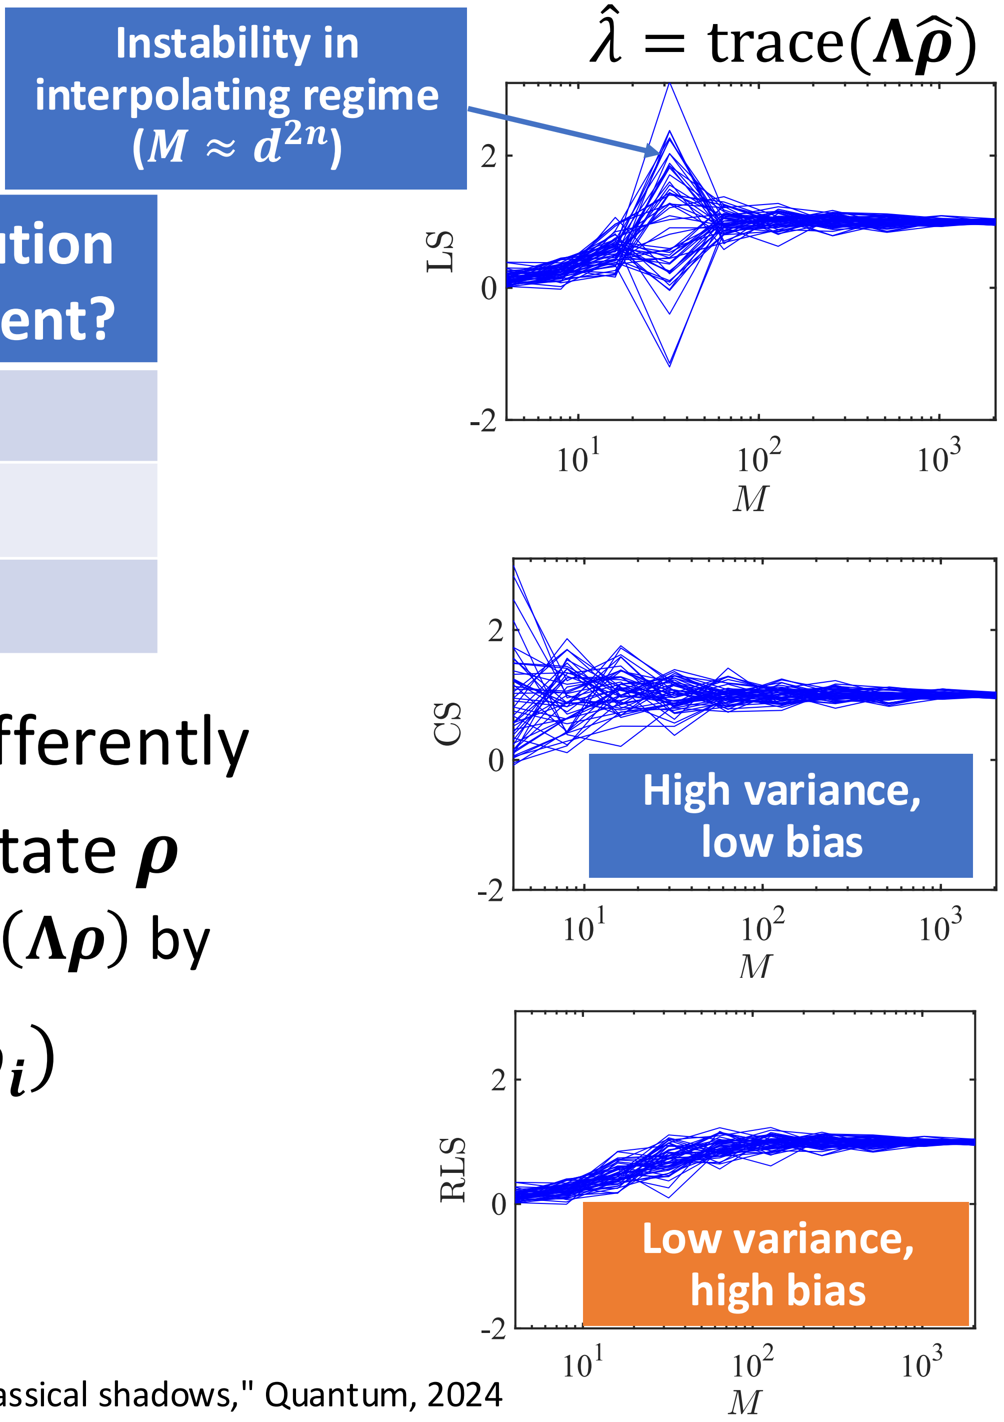

Z. Zhu, J. Lukens, B. Kirby, ''On the connection between least squares, regularization, and classical shadows," Quantum, 2024

Classical shadow for tomograph?

- Can we use classical shadow for full characterization?

$$\mathbb{E}[||\widehat{\boldsymbol{\rho}}_{\mathrm{CS}} - \boldsymbol{\rho}^\star||_F^2] \approx \frac{d^{2n}}{M}$$

- Still needs d^{2n} sample complexity to get a stable recovery of the full state
- Projected classical shadow (PCS) to incorporate prior information about the state structure into CS

$$\widehat{\boldsymbol{\rho}}_{\mathrm{PCS}} = \mathcal{P}_{\mathbb{X}}(\widehat{\boldsymbol{\rho}}_{\mathrm{CS}}) = \operatorname*{argmin}_{\boldsymbol{\rho}' \in \mathbb{X}} ||\boldsymbol{\rho}' - \widehat{\boldsymbol{\rho}}_{\mathrm{CS}}||_F^2$$

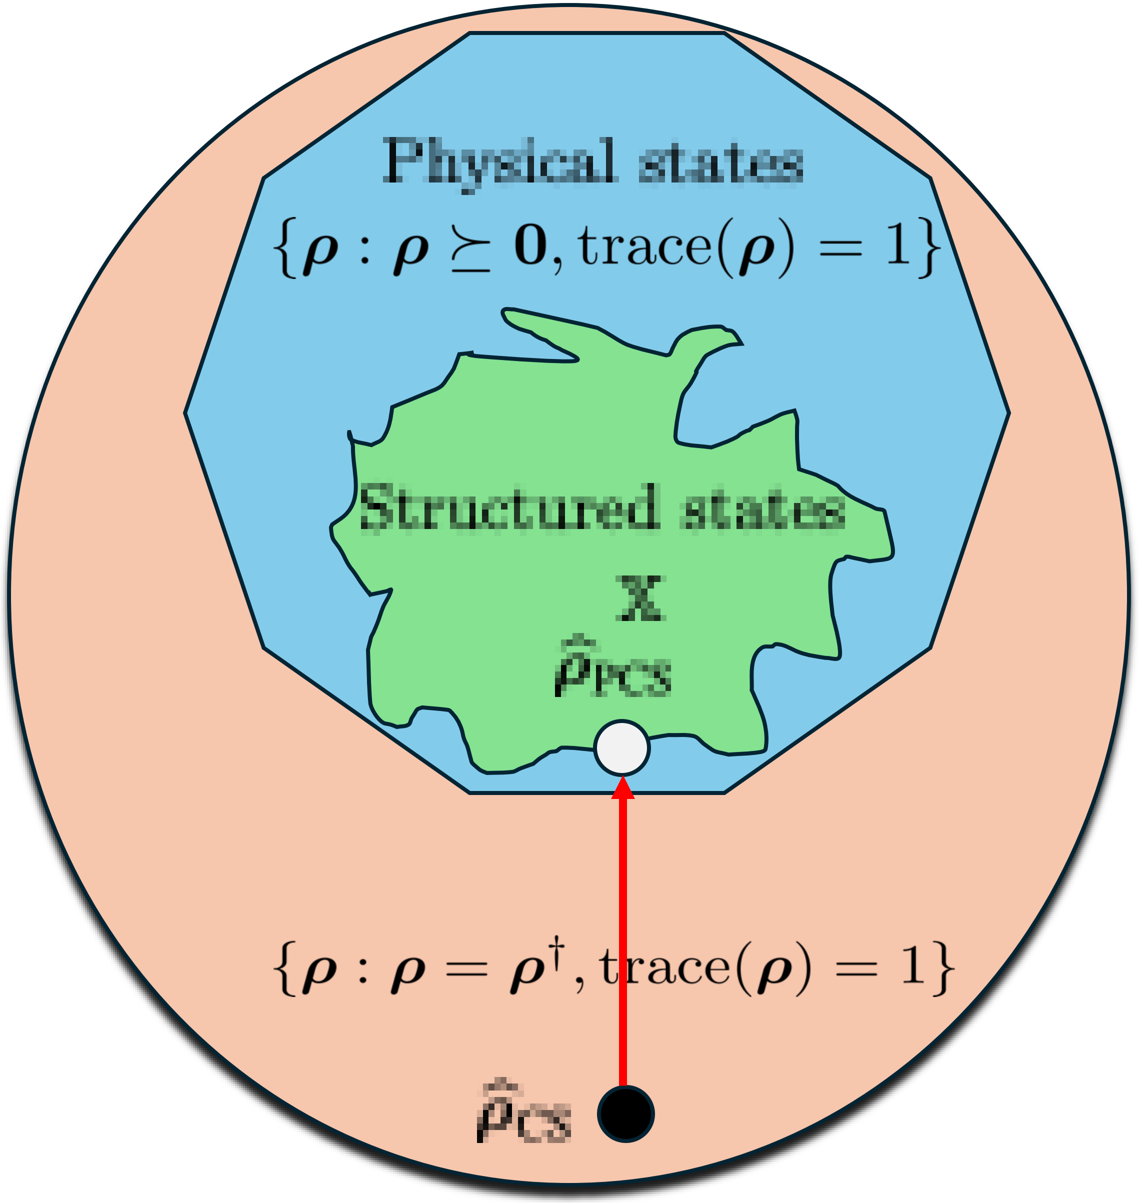

Projected classical shadow

Theorem (PCS; informal) For a given state $\boldsymbol{\rho}^\star \in \mathbb{X}$, $\widehat{\boldsymbol{\rho}}_{\text{PCS}}$ satisfies

$$\|\boldsymbol{\rho}^\star - \widehat{\boldsymbol{\rho}}_{\text{PCS}}\|_F \leq \epsilon$$

with high probability as long as

$$M \gtrsim \frac{\log N(\mathbb{X})}{\epsilon^2}$$

- Covering number $N(\mathbb{X})$ captures how "big" the space
- $\log N(\mathbb{X})$ often captures the degrees of freedom, so PCS achieves nearly optimal sample complexity

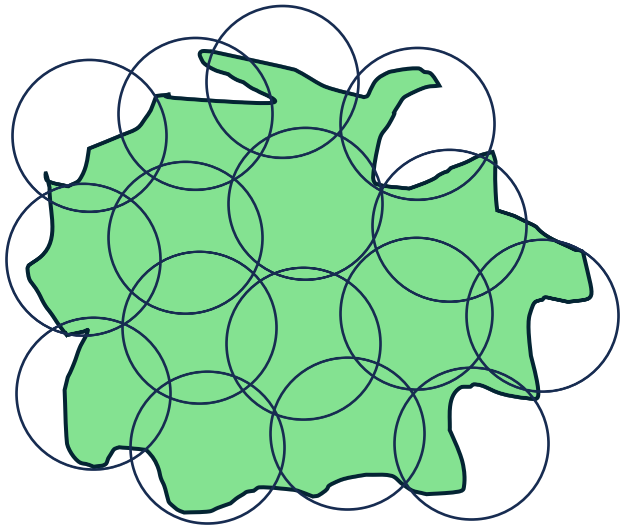

Z. Qin, J. Lukens, B. Kirby, Z. Zhu, Enhancing Quantum State Reconstruction with Structured Classical Shadows, npj Quantum Information, 2025.

Projected classical shadow

- Approximate projection $\tilde{\mathcal{P}}_{\mathbb{X}}$ satisfies $||\tilde{\mathcal{P}}_{\mathbb{X}}(\boldsymbol{\rho}) - \boldsymbol{\rho}||_{\boldsymbol{F}} \leq \alpha||\mathcal{P}_{\mathbb{X}}(\boldsymbol{\rho}) - \boldsymbol{\rho}||_{\boldsymbol{F}}$

Theorem (Approximate PCS; informal) $\tilde{\boldsymbol{\rho}}_{\mathrm{PCS}} = \tilde{\mathcal{P}}_{\mathbb{X}}(\hat{\boldsymbol{\rho}}_{\mathrm{CS}})$ satisfies

$$||\boldsymbol{\rho}^{\star} - \tilde{\boldsymbol{\rho}}_{\mathrm{PCS}}||_F \leq \epsilon$$

with high probability as long as

$$M \gtrsim \frac{\alpha \log N(\mathbb{X})}{\epsilon^2}$$

- Enable the use of efficient methods for approximate projection computation
- Sample complexity grows only proportionally

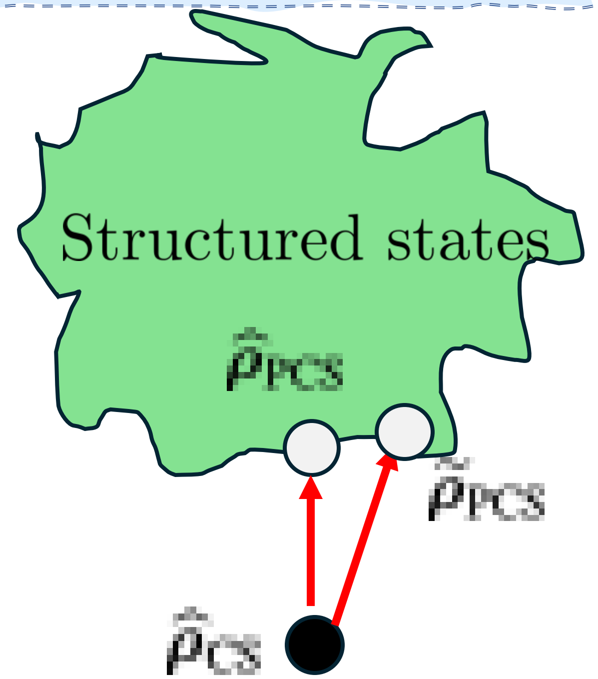

Projected classical shadow for low-rank states

Theorem (informal) For low (matrix) rank states $\mathbb{X}$ with rank γ, $\widehat{\boldsymbol{\rho}}_{\mathrm{PCS}}$ satisfies

$$||\boldsymbol{\rho}^{\star} - \widehat{\boldsymbol{\rho}}_{\mathrm{PCS}}||_F \leq \epsilon$$

with high probability as long as

$$M \gtrsim \frac{d^n \gamma}{\epsilon^2}$$

- Projection can be computed by eigen-decomposition
- $\Omega(d^n \gamma^2 / \epsilon^2)$ sample complexity for $||\boldsymbol{\rho}^{\star} - \widehat{\boldsymbol{\rho}}_{\mathrm{PCS}}||_1 \leq \epsilon$
- Match the optimal guarantee (up to log terms) for independent measurements

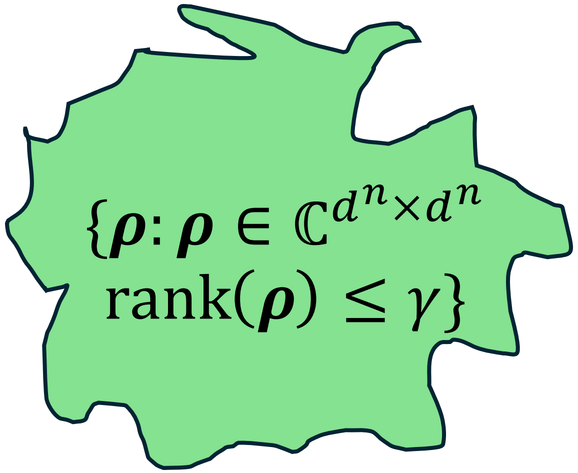

$\log N(\mathbb{X}) \lesssim d^n \gamma$

Haah, J., Harrow, A., Ji, Z., Wu, X. & Yu, N. Sample-optimal tomography of quantum states. IEEE Trans. Inf. Theory 2017.

Projected classical shadow for MPOs

Theorem (informal) For a given MPO $\boldsymbol{\rho}^\star \in \mathbb{X}$, $\widehat{\boldsymbol{\rho}}_{\mathrm{PCS}}$ satisfies

$$||\boldsymbol{\rho}^\star - \widehat{\boldsymbol{\rho}}_{\mathrm{PCS}}||_F \leq \epsilon$$

with high probability as long as $M \gtrsim \frac{nd^2r^2}{\epsilon^2}$ (matches the DOF).

The approximate $\widetilde{\boldsymbol{\rho}}_{\mathrm{PCS}}$ with $\tilde{\mathcal{P}}_{\mathbb{X}}$ given by TT-SVD satisfies the same guarantee with $M \gtrsim \frac{n^2d^2r^2}{\epsilon^2}$.

- However, no known algorithm to compute the exact projection. TT-SVD is a sequential SVD-based algorithm with quasi-optimal guarantee

$$||\mathrm{SVD}^{\mathrm{TT}}(\boldsymbol{\rho}) - \boldsymbol{\rho}||_F^2 \leq (n-1) \min_{\boldsymbol{\rho}':\,\mathrm{MPO}_r} ||\boldsymbol{\rho}' - \boldsymbol{\rho}||_2^2$$

- PCS with TT-SVD is quasi-optimal for MPOs ($O(n)$ larger than the DOF)

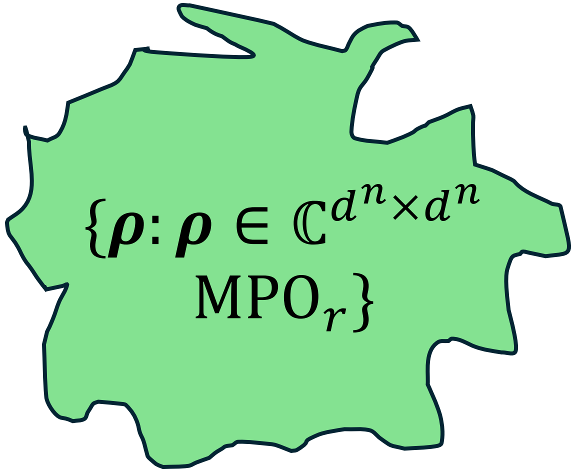

$\log N(\mathbb{X}) \lesssim nd^2r^2$

IV Oseledets, "Tensor-Train Decomposition", SIAM Journal on Scientific Computing, 2021.

Experiments on low-rank state

- Estimating $n = 4$ qubit low (matrix) rank states

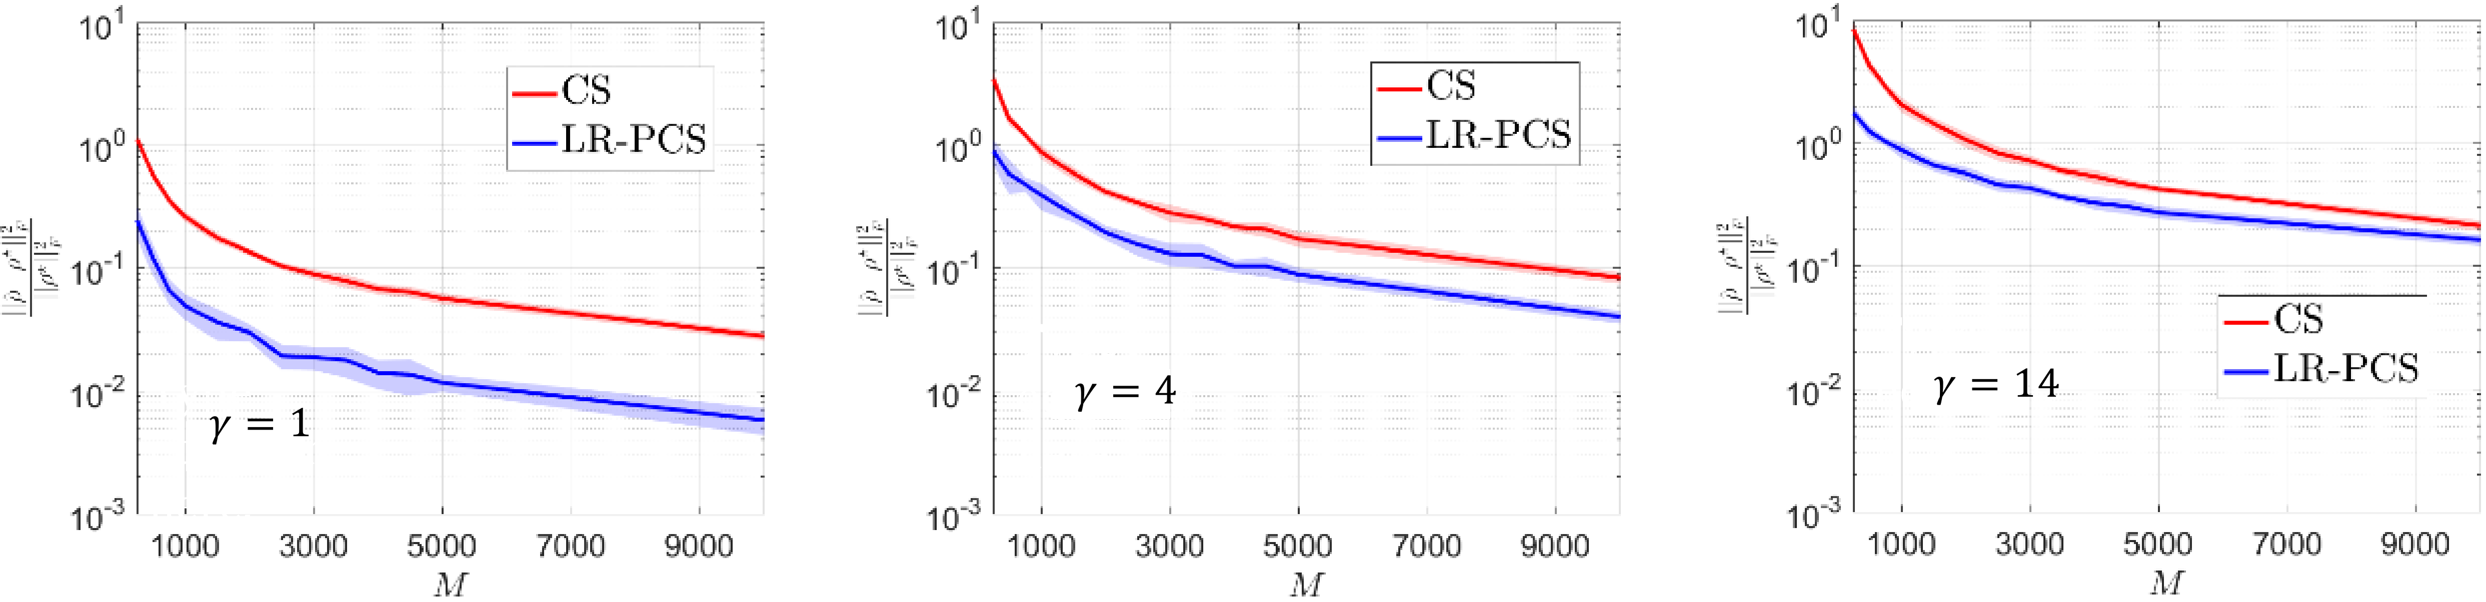

- PCS outperforms CS, particularly when the rank is low

Experiments on MPOs

- Estimating $n = 7$ qubit MPOs

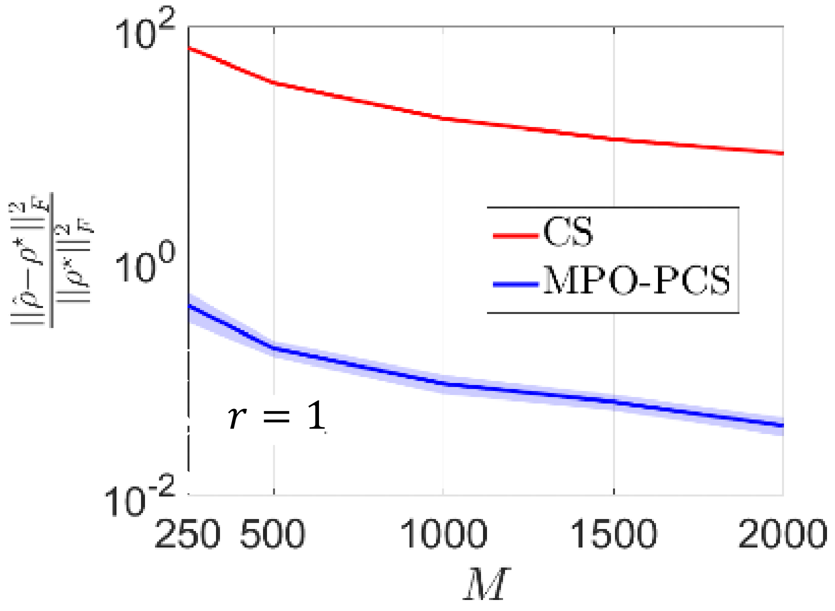

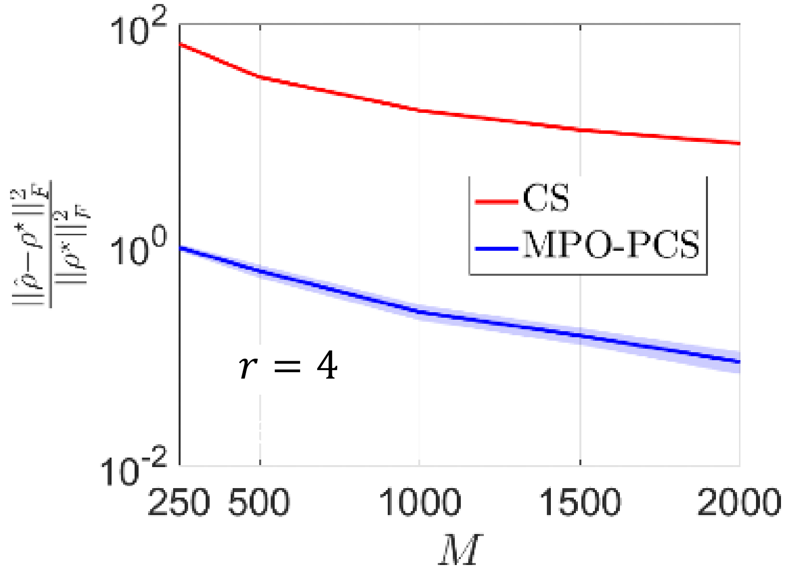

- PCS significantly outperforms CS

Experiments on practical states

- Estimating $n = 7$ qubit thermal and GHZ states

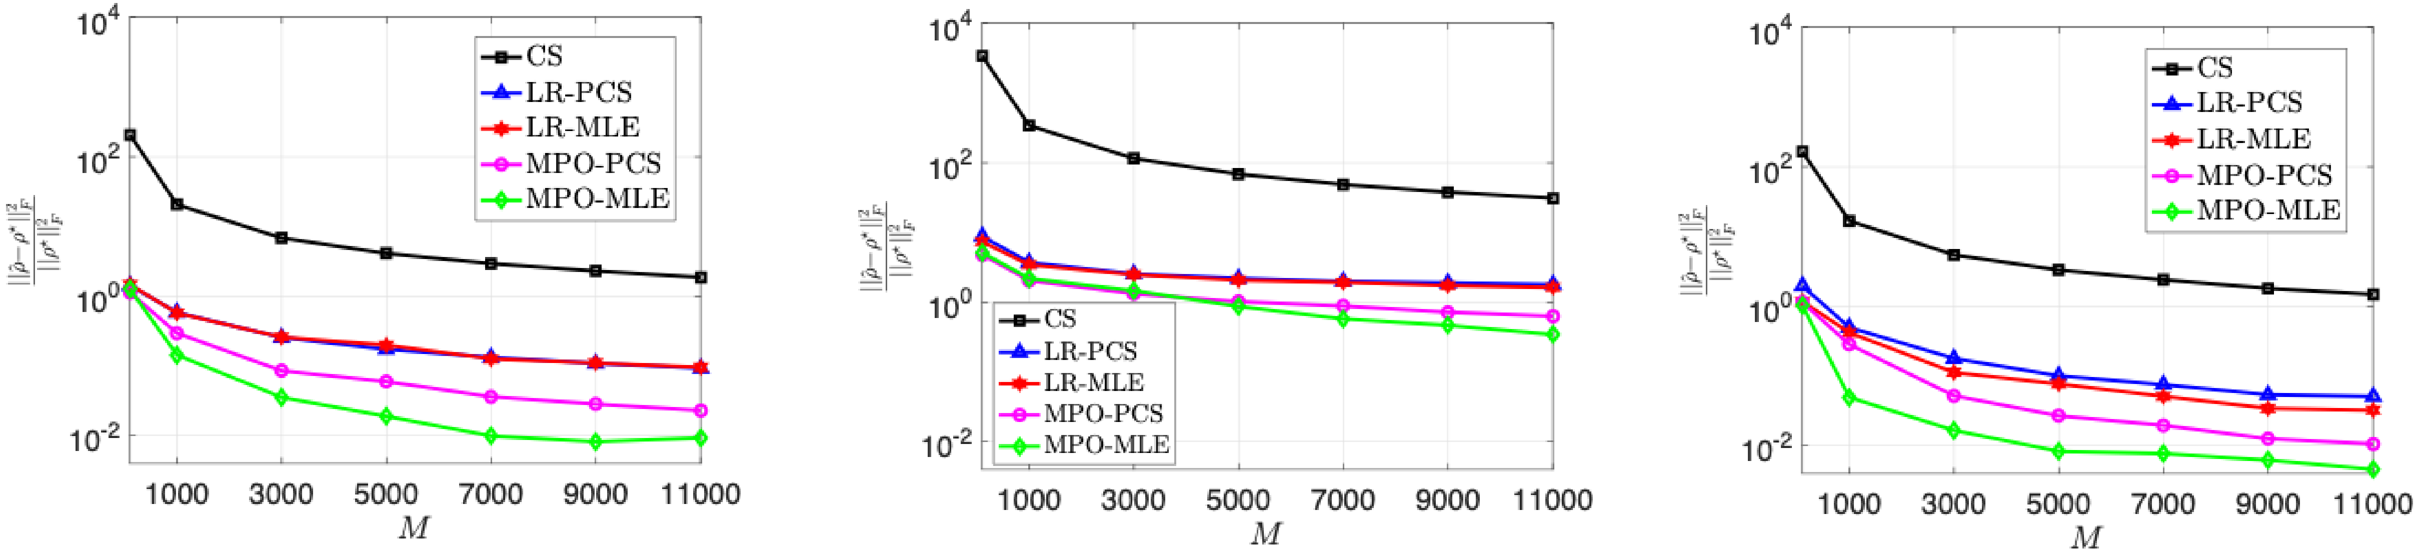

thermal state ($T = 0.2$) thermal state ($T = 2$) GHZ state

- Both LR-PCS and MPO-PCS outperform original CS
- PCS matches the performances of maximum likelihood estimation (MLE)

One more thing: experimental cost for POVMs

- We shift from IC-POVM to Haar random projective measurement because the former is challenging to implement. But Haar-random unitaries
  - are uniformly distributed over $\mathrm{U}(d^n)$, require $\exp(n)$ time to generate

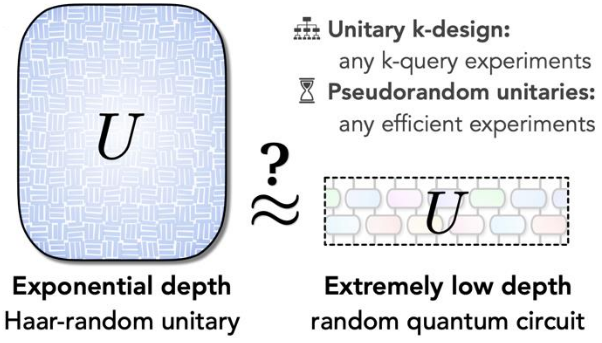

- Pseudorandom unitaries
- Local measurements
  - Extreme case: randomly rotate each qudit separately

$$\underset{d^n \times d^n}{\boldsymbol{U}} = \underset{d \times d}{\boldsymbol{U}_1} \otimes \underset{d \times d}{\boldsymbol{U}_2} \otimes \cdots \otimes \underset{d \times d}{\boldsymbol{U}_n}$$

  - Statistical analysis will be more challenging

T Schuster, J Haferkamp, HY Huang, Random unitaries in extremely low depth, Science 2025.

Conclusion

Quantum states have very low-dim structures. We can exploit them for efficient quantum state learning

- reduce sampling complexity
- design efficient post-processing algorithm

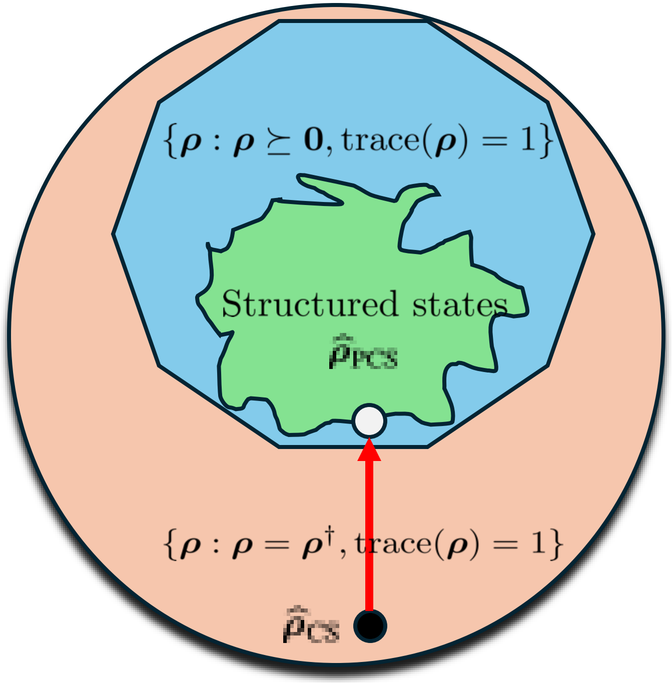

- We are in the noisy intermediate scale quantum (NISQ) ear. We don't have general, large scale, fault-tolerant quantum computation yet.
  - extend the analysis to local measurements
  - deal with State Preparation and Measurement (SPAM) errors

Z. Qin, J. Lukens, B. Kirby, Z. Zhu, Enhancing Quantum State Reconstruction with Structured Classical Shadows, npj Quantum Information, 2025.
Z. Qin, C. Jameson, Z. Gong, M. B. Wakin, Z. Zhu, Quantum State Tomography for Matrix Product Density Operators, IEEE TIT 2024.
Z. Qin, et al., Z. Zhu, Sample-Efficient Quantum State Tomography for Structured Quantum States in One Dimension, 2024.
Z. Zhu, J. Lukens, B. Kirby, "On the connection between least squares, regularization, and classical shadows," Quantum, 2024

Acknowledgements

Brian Kirby
ARL

Joseph Lukens
Purdue

Zhen Qin
OSU/UMich

Alireza Goldar
CSM

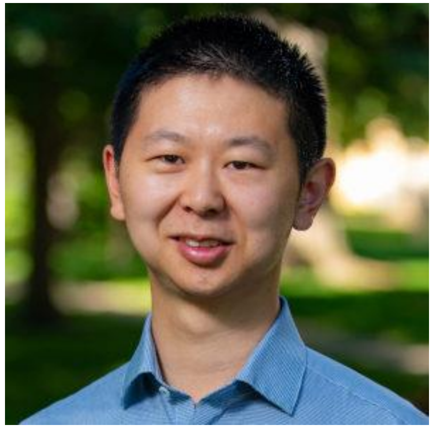

Zhexuan Gong
CSM

Casey Jameson
CSM

Mike Wakin
CSM

Research supported by NSF grants CCF-2241298 and ECCS-2409701.

Thank you for your attention!