How hard is it to verify a classical shadow?

Georgios Karaiskos [1] Dorian Rudolph [1] Johannes Jakob Meyer [2] Jens Eisert [2]

Sevag Gharibian [1]

[1] Institute for Photonic Quantum Systems (PhoQS) and Universität Paderborn, Germany

[2] Freie Universität Berlin, Germany

Theme

Theme

Theme

Question: How hard to check that a shadow represents what you think it represents?

Outline

1. Background: Classical Shadows
2. Our results
3. Proof sketch for local Clifford HKP protocol

Classical shadows for duck enthusiasts

1. Pick a set of duck-related properties you're interested in, $\mathcal{P} = \{P_i\}_{i=1}^m$.

Classical shadows for duck enthusiasts

1. Pick a set of duck-related properties you're interested in, $\mathcal{P} = \{P_i\}_{i=1}^m$.
2. Hire wildlife photographer (physicist) to repeat for N days:

Classical shadows for duck enthusiasts

1. Pick a set of duck-related properties you're interested in, $\mathcal{P} = \{P_i\}_{i=1}^m$.
2. Hire wildlife photographer (physicist) to repeat for N days:
   (a) Find your duck (ρ) in the wild (lab).

Classical shadows for duck enthusiasts

1. Pick a set of duck-related properties you're interested in, $\mathcal{P} = \{P_i\}_{i=1}^m$.
2. Hire wildlife photographer (physicist) to repeat for N days:
   (a) Find your duck (ρ) in the wild (lab).
   (b) Take a picture s_i (snapshot) of the duck, which scares duck away (collapse).

Classical shadows for duck enthusiasts

1. Pick a set of duck-related properties you're interested in, $\mathcal{P} = \{P_i\}_{i=1}^m$.
2. Hire wildlife photographer (physicist) to repeat for N days:
   (a) Find your duck (ρ) in the wild (lab).
   (b) Take a picture s_i (snapshot) of the duck, which scares duck away (collapse).
3. Photographer hands you stack of photos $S = \{s_1, \ldots, s_N\}$ (shadow).

Classical shadows for duck enthusiasts

1. Pick a set of duck-related properties you're interested in, $\mathcal{P} = \{P_i\}_{i=1}^m$.
2. Hire wildlife photographer (physicist) to repeat for N days:
   (a) Find your duck (ρ) in the wild (lab).
   (b) Take a picture s_i (snapshot) of the duck, which scares duck away (collapse).
3. Photographer hands you stack of photos $S = \{s_1, \ldots, s_N\}$ (shadow).
4. Using S and classical postprocessing, predict value of all P_i for your duck.

Classical shadows for duck enthusiasts

1. Pick a set of duck-related properties you're interested in, $\mathcal{P} = \{P_i\}_{i=1}^{m}$.
2. Hire wildlife photographer (physicist) to repeat for N days:
   (a) Find your duck (ρ) in the wild (lab).
   (b) Take a picture s_i (snapshot) of the duck, which scares duck away (collapse).
3. Photographer hands you stack of photos $S = \{s_1, \ldots, s_N\}$ (shadow).
4. Using S and classical postprocessing, predict value of all P_i for your duck.

Question: How many snapshots N do you need to capture m properties?

Classical shadows for duck enthusiasts

1. Pick a set of duck-related properties you're interested in, $\mathcal{P} = \{P_i\}_{i=1}^m$.
2. Hire wildlife photographer (physicist) to repeat for N days:
   (a) Find your duck (ρ) in the wild (lab).
   (b) Take a picture s_i (snapshot) of the duck, which scares duck away (collapse).
3. Photographer hands you stack of photos $S = \{s_1, \ldots, s_N\}$ (shadow).
4. Using S and classical postprocessing, predict value of all P_i for your duck.

Question: How many snapshots N do you need to capture m properties?

Answer: $N \sim \text{polylog}(m)$ [Aaronson 2018], i.e. sample-efficient.

Caveat: Not time efficient.

Huang-Kueng-Preskill (HKP) classical shadows [HKP20]

General framework:

1. Pick set of measurement operators, $\mathcal{P} = \{P_i\}_{i=1}^m$, and set of unitaries $\mathcal{U}$.
2. Repeat N times:
   (a) Produce n-qubit state ρ in lab.
   (b) Pick random $U \in \mathcal{U}$, measure $U\rho U^\dagger$ in standard basis to obtain snapshot $s_i \in \{0,1\}^n$.
3. Using $S = \{s_i\}$ and classical median-of-means, predict value of all $\mathrm{Tr}(\rho P_i)$ within additive error ϵ.

Huang-Kueng-Preskill (HKP) classical shadows [HKP20]

General framework:

1. Pick set of measurement operators, $\mathcal{P} = \{P_i\}_{i=1}^m$, and set of unitaries $\mathcal{U}$.
2. Repeat N times:
   (a) Produce n-qubit state ρ in lab.
   (b) Pick random $U \in \mathcal{U}$, measure $U\rho U^\dagger$ in standard basis to obtain snapshot $s_i \in \{0,1\}^n$.
3. Using $S = \{s_i\}$ and classical median-of-means, predict value of all $\mathrm{Tr}(\rho P_i)$ within additive error ϵ.

Sample complexity:

$$O\left(\frac{\log(m)}{\epsilon^2}\max_{1\leq i\leq m}\left\|P_i - \frac{\mathrm{Tr}(P_i)}{2^n}I\right\|_{\text{shadow}}^2\right),$$

where $\|\cdot\|_{\text{shadow}}$ depends on choice of $\mathcal{P}$ and $\mathcal{U}$.

Time complexity: Depends on complexity of implementing $\mathcal{U}$ and corresponding postprocessing.

Two efficient instantiations of HKP classical shadows [HKP20]

Local Clifford:

- Properties to predict: $\mathcal{P} = \{P_i\}_{i=1}^m$ are k-local observables.
- Random measurements: Independent single-qubit Clifford measurements on each qubit.
- Sample complexity: $O(\log(m) 4^k \max_{1 \leq i \leq m} \|P_i\|_\infty^2 / \epsilon^2)$.

Two efficient instantiations of HKP classical shadows [HKP20]

Local Clifford:

- Properties to predict: $\mathcal{P} = \{P_i\}_{i=1}^m$ are k-local observables.
- Random measurements: Independent single-qubit Clifford measurements on each qubit.
- Sample complexity: $O(\log(m)4^k \max_{1\leq i\leq m} \|P_i\|_\infty^2/\epsilon^2)$.

Global Clifford:

- Properties to predict: $\mathcal{P} = \{P_i\}_{i=1}^m$.
- Random measurements: Global Clifford measurements on all n qubits.
- Sample complexity: $O(\log(m) \max_{1\leq i\leq m} \|P_i\|_F^2/\epsilon^2)$.

Central question of this work

Given $S = \{s_i\}$, how hard to verify if S consistent with some n-qubit ρ?

Central question of this work

Given $S = \{s_i\}$, how hard to verify if S consistent with some n-qubit ρ?

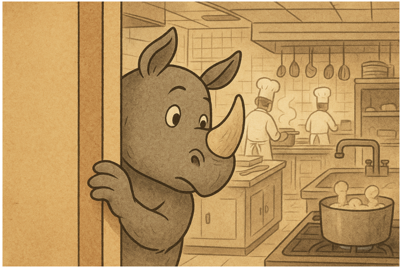

Problem: First need general formal definition of "classical shadow".

Formalization

Classical shadow

A shadow on n qubits is a 3-tuple (S, O, A), where

- (Shadow) Multi-set $S = \{s_i\}_{i=1}^N$ of $\mathrm{poly}(n)$-bit strings, i.e. *snapshots*,
- (Observables) Set $O = \{O_i\}_{i=1}^m$ of observables s.t. $\|O_i\|_\infty \leq 1$ and $1 \leq m \leq 4^n$. O is poly-time uniformly generated, i.e. given index i, poly-size quantum circuit for O_i can be efficiently produced.
- (Recovery) Efficient classical algorithm A which, given S and $i \in [m]$, produces $A(S, i) \in [-1, 1]$.

Formalization

Classical shadow

A shadow on n qubits is a 3-tuple (S, O, A), where

- (Shadow) Multi-set $S = \{s_i\}_{i=1}^{N}$ of $\text{poly}(n)$-bit strings, i.e. *snapshots*,
- (Observables) Set $O = \{O_i\}_{i=1}^{m}$ of observables s.t. $\|O_i\|_\infty \leq 1$ and $1 \leq m \leq 4^n$. O is poly-time uniformly generated, i.e. given index i, poly-size quantum circuit for O_i can be efficiently produced.
- (Recovery) Efficient classical algorithm A which, given S and $i \in [m]$, produces $A(S, i) \in [-1, 1]$.

Classical Shadow Validity (CSV)

Given shadow (S, O, A), α, β s.t. $\beta - \alpha \geq 1/\text{poly}(n)$, decide:

- **Yes**: $\exists$ n-qubit state ρ s.t. $\forall$ $i \in [m]$, $|\mathrm{Tr}(O_i\rho) - A(S, i)| \leq \alpha$.
- **No**: $\forall$ n-qubit states ρ $\exists$ some $i \in [m]$ s.t. $|\mathrm{Tr}(O_i\rho) - A(S, i)| \geq \beta$.

Outline

Our results 1: Polynomially many observables O_i

Hardness results:

Theorem 1

CSV for local Clifford HKP shadows is QMA-complete, even for 6-local O_i on a spatially sparse hypergraph.

Interpretation:

- Even for protocols as simple as local Clifford HKP, shadow verification is hard!

Our results 1: Polynomially many observables O_i

Hardness results:

Theorem 1

CSV for local Clifford HKP shadows is QMA-complete, even for 6-local O_i on a spatially sparse hypergraph.

Interpretation:

- Even for protocols as simple as local Clifford HKP, shadow verification is hard!

Theorem 2 (High-dimensional generalization of HKP)

CSV for local [Mao, Yi, Zhu 2025] shadows is QMA-complete for odd prime local dimension $d \geq 11$, even for 2-local nearest-neighbor O_i on a line.

Our results 1: Polynomially many observables O_i

"Dequantization" results:

Theorem 3

CSV for global Clifford HKP shadows is classically efficiently solvable if (a) $\|O_i\|_F \leq \text{poly}(n)$ for all i, and (b) we are given query and sampling access to each O_i.

Recall:

- Sample complexity: $O(\log(m)\max_{1\leq i\leq m}\|P_i\|_F^2/\epsilon^2)$.
- Dequantization covers precisely regime where HKP is efficient!

Our results 1: Polynomially many observables O_i

"Dequantization" results:

Theorem 3

CSV for global Clifford HKP shadows is classically efficiently solvable if (a) $\|O_i\|_F \leq \text{poly}(n)$ for all i, and (b) we are given query and sampling access to each O_i.

Recall:

- Sample complexity: $O(\log(m) \max_{1\leq i\leq m} \|P_i\|_F^2/\epsilon^2)$.
- Dequantization covers precisely regime where HKP is efficient!

Theorem 4

CSV for global [Mao, Yi, Zhu 2025] is classically efficiently solvable if (a) $\|O_i\|_F \leq \text{poly}(n)$ for all i, and (b) we are given query and sampling access to each O_i.

Our results 2: Exponentially many observables O_i

Question: Since sample complexity is polylog(m), what about CSV for exponentially many observables m?

Our results 2: Exponentially many observables O_i

Question: Since sample complexity is polylog(m), what about CSV for exponentially many observables m?

[King, Gosset, Kothari and Babbush 2025]

- Classical shadow protocol for $O = \{I, X, Y, Z\}^n$.
- Preparing shadow is poly sample complexity, exponential time complexity.
  - Preparation time not relevant for CSV, since given shadow as input.
- "Rapid recovery": Poly-time recovery algorithm A for $\epsilon \in \Theta(1)$ precision.

Our results 2: Exponentially many observables O_i

Question: Since sample complexity is polylog(m), what about CSV for exponentially many observables m?

[King, Gosset, Kothari and Babbush 2025]

- Classical shadow protocol for $O = \{I, X, Y, Z\}^n$.
- Preparing shadow is poly sample complexity, exponential time complexity.
  - Preparation time not relevant for CSV, since given shadow as input.
- "Rapid recovery": Poly-time recovery algorithm A for $\epsilon \in \Theta(1)$ precision.

Theorem 5

CSV for exponentially many observables and constant error ϵ recovery precision is qc-Σ_2-complete.

Our results 2: Exponentially many observables O_i

Question: Since sample complexity is polylog(m), what about CSV for exponentially many observables m?

[King, Gosset, Kothari and Babbush 2025]

- Classical shadow protocol for $O = \{I, X, Y, Z\}^n$.
- Preparing shadow is poly sample complexity, exponential time complexity.
  - Preparation time not relevant for CSV, since given shadow as input.
- "Rapid recovery": Poly-time recovery algorithm A for $\epsilon \in \Theta(1)$ precision.

Theorem 5

CSV for exponentially many observables and constant error ϵ recovery precision is qc-Σ_2-complete.

- Pro: First natural complete problem for qc-Σ_2, quantum generalization of Σ_2^p
  - qc-Σ_2: $\exists$ quantum proof $|\psi\rangle$ s.t. $\forall$ classical proofs y, quantum verifier V accepts $(|\psi\rangle, y)$.

Our results 2: Exponentially many observables O_i

Question: Since sample complexity is polylog(m), what about CSV for exponentially many observables m?

[King, Gosset, Kothari and Babbush 2025]

- Classical shadow protocol for $O = \{I, X, Y, Z\}^n$.
- Preparing shadow is poly sample complexity, exponential time complexity.
  - Preparation time not relevant for CSV, since given shadow as input.
- "Rapid recovery": Poly-time recovery algorithm A for $\epsilon \in \Theta(1)$ precision.

Theorem 5

CSV for exponentially many observables and constant error ϵ recovery precision is qc-Σ_2-complete.

- Pro: First natural complete problem for qc-Σ_2, quantum generalization of Σ_2^p
  - qc-Σ_2: $\exists$ quantum proof $|\psi\rangle$ s.t. $\forall$ classical proofs y, quantum verifier V accepts $(|\psi\rangle, y)$.
- Con/open question: Cannot prove it for observable set $O = \{I, X, Y, Z\}^n$.

Our results 3: Variants of CSV

Theorem 6

CSV where consistent state must be product, i.e., $\rho = \rho_A \otimes \rho_B$, is:

- $\mathrm{QMA}(2)$-complete for polynomially many observables, and
- qcq-Σ_3-complete for exponentially many observables.

Here:

- $\mathrm{QMA}(2)$ is QMA but with tensor product proof $\rho = \rho_A \otimes \rho_B$.
- qcq-Σ_3: $\exists$ quantum proof $|\psi\rangle$ s.t. $\forall$ classical proofs y $\exists$ quantum proof $|\phi\rangle$, s.t. V accepts $(|\psi\rangle, y, |\phi\rangle)$.

Our results 3: Variants of CSV

Theorem 6

CSV where consistent state must be product, i.e., $\rho = \rho_A \otimes \rho_B$, is:

- $\mathrm{QMA}(2)$-complete for polynomially many observables, and
- qcq-Σ_3-complete for exponentially many observables.

Here:

- $\mathrm{QMA}(2)$ is QMA but with tensor product proof $\rho = \rho_A \otimes \rho_B$.
- qcq-Σ_3: $\exists$ quantum proof $|\psi\rangle$ s.t. $\forall$ classical proofs y $\exists$ quantum proof $|\phi\rangle$, s.t. V accepts $(|\psi\rangle, y, |\phi\rangle)$.

Theorem 7

Verifying if set of shadows, each possibly with different observables, all correspond to the same ρ is:

- QMA-complete for polynomially many observables, and
- qc-Σ_2-complete for exponentially many observables.

Outline

Theorem 1

CSV for local Clifford HKP shadows is QMA-complete, even for 6-local O_i on a spatially sparse hypergraph.

Theorem 1

CSV for local Clifford HKP shadows is QMA-complete, even for 6-local O_i on a spatially sparse hypergraph.

Naive idea: CSV sounds suspiciously like QMA-complete Consistency of Local Density Matrices problem...

CONSISTENCY [Liu 2006]

Given poly many k-local density matrices ρ_i, is there global n-qubit ρ consistent with all ρ_i?

Theorem 1

CSV for local Clifford HKP shadows is QMA-complete, even for 6-local O_i on a spatially sparse hypergraph.

Naive idea: CSV sounds suspiciously like QMA-complete Consistency of Local Density Matrices problem...

CONSISTENCY [Liu 2006]

Given poly many k-local density matrices ρ_i, is there global n-qubit ρ consistent with all ρ_i?

Indeed: CONSISTENCY is special case of CSV $\implies$ get QMA-hardness of general CSV problem for free!

Theorem 1

CSV for local Clifford HKP shadows is QMA-complete, even for 6-local O_i on a spatially sparse hypergraph.

Naive idea: CSV sounds suspiciously like QMA-complete Consistency of Local Density Matrices problem...

CONSISTENCY [Liu 2006]

Given poly many k-local density matrices ρ_i, is there global n-qubit ρ consistent with all ρ_i?

Indeed: CONSISTENCY is special case of CSV $\implies$ get QMA-hardness of general CSV problem for free!

Challenge:

- Does not suffice for QMA-hardness of classical shadows used in practice, e.g. HKP shadows
- Unlike local reduced states ρ_i, HKP shadows are highly non-local objects

Theorem 1

CSV for local Clifford HKP shadows is QMA-complete, even for 6-local O_i on a spatially sparse hypergraph.

Naive idea: CSV sounds suspiciously like QMA-complete Consistency of Local Density Matrices problem...

CONSISTENCY [Liu 2006]

Given poly many k-local density matrices ρ_i, is there global n-qubit ρ consistent with all ρ_i?

Indeed: CONSISTENCY is special case of CSV $\implies$ get QMA-hardness of general CSV problem for free!

Challenge:

- Does not suffice for QMA-hardness of classical shadows used in practice, e.g. HKP shadows
- Unlike local reduced states ρ_i, HKP shadows are highly non-local objects
  - Reducing CONSISTENCY to HKP $=$ stitching together overlapping local data to build global state

Theorem 1

CSV for local Clifford HKP shadows is QMA-complete, even for 6-local O_i on a spatially sparse hypergraph.

Naive idea: CSV sounds suspiciously like QMA-complete Consistency of Local Density Matrices problem...

CONSISTENCY [Liu 2006]

Given poly many k-local density matrices ρ_i, is there global n-qubit ρ consistent with all ρ_i?

Indeed: CONSISTENCY is special case of CSV $\implies$ get QMA-hardness of general CSV problem for free!

Challenge:

- Does not suffice for QMA-hardness of classical shadows used in practice, e.g. HKP shadows
- Unlike local reduced states ρ_i, HKP shadows are highly non-local objects
  - Reducing CONSISTENCY to HKP $=$ stitching together overlapping local data to build global state
  - Even worse, this has to be done while simulating measurement statistics of HKP shadows...

Local Clifford HKP shadows:

- Properties to predict: $\mathcal{P} = \{P_i\}_{i=1}^m$ are k-local observables.
- Random measurements: Independent single-qubit Clifford measurements on each qubit.
  - Given ρ, measure each qubit at random in eigenbasis of X, Y, or Z.
  - Let $|\psi_i\rangle$ be obtained result for qubit i.

Local Clifford HKP shadows:

- Properties to predict: $\mathcal{P} = \{P_i\}_{i=1}^m$ are k-local observables.
- Random measurements: Independent single-qubit Clifford measurements on each qubit.
  - Given ρ, measure each qubit at random in eigenbasis of X, Y, or Z.
  - Let $|\psi_i\rangle$ be obtained result for qubit i.
  - Defining $\hat{\eta}_j = 3|\psi_i\rangle\langle\psi_i| - I$, obtain snapshot $s = \hat{\eta}_1 \otimes \cdots \otimes \hat{\eta}_n$.

Local Clifford HKP shadows:

- Properties to predict: $\mathcal{P} = \{P_i\}_{i=1}^m$ are k-local observables.
- Random measurements: Independent single-qubit Clifford measurements on each qubit.
  - Given ρ, measure each qubit at random in eigenbasis of X, Y, or Z.
  - Let $|\psi_i\rangle$ be obtained result for qubit i.
  - Defining $\hat{\eta}_j = 3|\psi_i\rangle\langle\psi_i| - I$, obtain snapshot $s = \hat{\eta}_1 \otimes \cdots \otimes \hat{\eta}_n$.
  - Shadow $S = \{s_1, s_2, \ldots, s_m\}$ is set of all snapshots.

Local Clifford HKP shadows:

- Properties to predict: $\mathcal{P} = \{P_i\}_{i=1}^m$ are k-local observables.
- Random measurements: Independent single-qubit Clifford measurements on each qubit.
  - Given ρ, measure each qubit at random in eigenbasis of X, Y, or Z.
  - Let $|\psi_i\rangle$ be obtained result for qubit i.
  - Defining $\hat{\eta}_j = 3|\psi_i\rangle\langle\psi_i| - I$, obtain snapshot $s = \hat{\eta}_1 \otimes \cdots \otimes \hat{\eta}_n$.
  - Shadow $S = \{s_1, s_2, \ldots, s_m\}$ is set of all snapshots.

Goal

Given k-local states e.g. $\rho_{1,2,5}$, $\rho_{2,4,9}$, etc, map to n-local snapshots e.g. $s_1 = \hat{\eta}_{1,1} \otimes \cdots \otimes \hat{\eta}_n$, $s_2 = \hat{\eta}_{2,1} \otimes \cdots \otimes \hat{\eta}_{2,n}$, etc, such that:

Local Clifford HKP shadows:

- Properties to predict: $\mathcal{P} = \{P_i\}_{i=1}^m$ are k-local observables.
- Random measurements: Independent single-qubit Clifford measurements on each qubit.
  - Given ρ, measure each qubit at random in eigenbasis of X, Y, or Z.
  - Let $|\psi_i\rangle$ be obtained result for qubit i.
  - Defining $\hat{\eta}_j = 3|\psi_i\rangle\langle\psi_i| - I$, obtain snapshot $s = \hat{\eta}_1 \otimes \cdots \otimes \hat{\eta}_n$.
  - Shadow $S = \{s_1, s_2, \ldots, s_m\}$ is set of all snapshots.

Goal

Given k-local states e.g. $\rho_{1,2,5}$, $\rho_{2,4,9}$, etc, map to n-local snapshots e.g. $s_1 = \hat{\eta}_{1,1} \otimes \cdots \otimes \hat{\eta}_n$, $s_2 = \hat{\eta}_{2,1} \otimes \cdots \otimes \hat{\eta}_{2,n}$, etc, such that:

- $\rho_{i,j,k}$ globally consistent iff frequency of s_i consistent with measuring some global ρ with random local Cliffords, and

Local Clifford HKP shadows:

- Properties to predict: $\mathcal{P} = \{P_i\}_{i=1}^m$ are k-local observables.
- Random measurements: Independent single-qubit Clifford measurements on each qubit.
  - Given ρ, measure each qubit at random in eigenbasis of X, Y, or Z.
  - Let $|\psi_i\rangle$ be obtained result for qubit i.
  - Defining $\hat{\eta}_j = 3|\psi_i\rangle\langle\psi_i| - I$, obtain snapshot $s = \hat{\eta}_1 \otimes \cdots \otimes \hat{\eta}_n$.
  - Shadow $S = \{s_1, s_2, \ldots, s_m\}$ is set of all snapshots.

Goal

Given k-local states e.g. $\rho_{1,2,5}$, $\rho_{2,4,9}$, etc, map to n-local snapshots e.g. $s_1 = \hat{\eta}_{1,1} \otimes \cdots \otimes \hat{\eta}_n$, $s_2 = \hat{\eta}_{2,1} \otimes \cdots \otimes \hat{\eta}_{2,n}$, etc, such that:

- $\rho_{i,j,k}$ globally consistent iff frequency of s_i consistent with measuring some global ρ with random local Cliffords, and
- $E[s_i] = \rho$, where expectation with respect to randomness in HKP protocol.

Proof steps - Reduction from 1D CONSISTENCY

1. Prove 1D CONSISTENCY problem on qudits with $d = 8$ is QMA-hard under many-one reductions.
   - Easier to "stitch together" nearest neighbor local states $\rho_{1,2}$, $\rho_{2,3}$, etc on the line.

Proof steps - Reduction from 1D CONSISTENCY

1. Prove 1D CONSISTENCY problem on qudits with $d = 8$ is QMA-hard under many-one reductions.
   - Easier to "stitch together" nearest neighbor local states $\rho_{1,2}$, $\rho_{2,3}$, etc on the line.
   - Combines locally simulatable codes [Broadbent, Grilo 2022] + 1D QMA-hardness of local Hamiltonian [Hallgren, Nagaj, Narayanaswami 2013].

Proof steps - Reduction from 1D CONSISTENCY

1. Prove 1D CONSISTENCY problem on qudits with $d = 8$ is QMA-hard under many-one reductions.
   - Easier to "stitch together" nearest neighbor local states $\rho_{1,2}$, $\rho_{2,3}$, etc on the line.
   - Combines locally simulatable codes [Broadbent, Grilo 2022] + 1D QMA-hardness of local Hamiltonian [Hallgren, Nagaj, Narayanaswami 2013].
2. Write each qudit q_i as qubit triples T_i. Consider all possible 6-local HKP snapshots on pairs (T_i, T_{i+1}).

Proof steps - Reduction from 1D CONSISTENCY

1. Prove 1D CONSISTENCY problem on qudits with $d = 8$ is QMA-hard under many-one reductions.
   - Easier to "stitch together" nearest neighbor local states $\rho_{1,2}$, $\rho_{2,3}$, etc on the line.
   - Combines locally simulatable codes [Broadbent, Grilo 2022] + 1D QMA-hardness of local Hamiltonian [Hallgren, Nagaj, Narayanaswami 2013].
2. Write each qudit q_i as qubit triples T_i. Consider all possible 6-local HKP snapshots on pairs (T_i, T_{i+1}).
3. Write linear program (LP) to compute "how much probability" to put onto each local snapshot, s.t.:
   - "Local probabilities" consistent with global HKP shadow iff 1D CONSISTENCY was YES instance.
   - Leverages HKP requirement that $E[s_i] = \rho$.

Proof steps - Reduction from 1D CONSISTENCY

1. Prove 1D CONSISTENCY problem on qudits with $d = 8$ is QMA-hard under many-one reductions.
   - Easier to "stitch together" nearest neighbor local states $\rho_{1,2}$, $\rho_{2,3}$, etc on the line.
   - Combines locally simulatable codes [Broadbent, Grilo 2022] + 1D QMA-hardness of local Hamiltonian [Hallgren, Nagaj, Narayanaswami 2013].
2. Write each qudit q_i as qubit triples T_i. Consider all possible 6-local HKP snapshots on pairs (T_i, T_{i+1}).
3. Write linear program (LP) to compute "how much probability" to put onto each local snapshot, s.t.:
   - "Local probabilities" consistent with global HKP shadow iff 1D CONSISTENCY was YES instance.
   - Leverages HKP requirement that $E[s_i] = \rho$.
   - Problem: LP returns real numbers, but each local snapshot can only occur integer number of times in shadow.

Proof steps - Reduction from 1D CONSISTENCY

1. Prove 1D CONSISTENCY problem on qudits with $d = 8$ is QMA-hard under many-one reductions.
   - Easier to "stitch together" nearest neighbor local states $\rho_{1,2}$, $\rho_{2,3}$, etc on the line.
   - Combines locally simulatable codes [Broadbent, Grilo 2022] + 1D QMA-hardness of local Hamiltonian [Hallgren, Nagaj, Narayanaswami 2013].
2. Write each qudit q_i as qubit triples T_i. Consider all possible 6-local HKP snapshots on pairs (T_i, T_{i+1}).
3. Write linear program (LP) to compute "how much probability" to put onto each local snapshot, s.t.:
   - "Local probabilities" consistent with global HKP shadow iff 1D CONSISTENCY was YES instance.
   - Leverages HKP requirement that $E[s_i] = \rho$.
   - Problem: LP returns real numbers, but each local snapshot can only occur integer number of times in shadow.
4. "Round" LP into integer program (IP) to give integer weights on local snapshots.

Proof steps - Reduction from 1D CONSISTENCY

1. Prove 1D CONSISTENCY problem on qudits with $d = 8$ is QMA-hard under many-one reductions.
   - Easier to "stitch together" nearest neighbor local states $\rho_{1,2}$, $\rho_{2,3}$, etc on the line.
   - Combines locally simulatable codes [Broadbent, Grilo 2022] + 1D QMA-hardness of local Hamiltonian [Hallgren, Nagaj, Narayanaswami 2013].
2. Write each qudit q_i as qubit triples T_i. Consider all possible 6-local HKP snapshots on pairs (T_i, T_{i+1}).
3. Write linear program (LP) to compute "how much probability" to put onto each local snapshot, s.t.:
   - "Local probabilities" consistent with global HKP shadow iff 1D CONSISTENCY was YES instance.
   - Leverages HKP requirement that $E[s_i] = \rho$.
   - Problem: LP returns real numbers, but each local snapshot can only occur integer number of times in shadow.
4. "Round" LP into integer program (IP) to give integer weights on local snapshots.
   - Problem: IPs generally NP-hard to solve...

Proof steps - Reduction from 1D CONSISTENCY

1. Prove 1D CONSISTENCY problem on qudits with $d = 8$ is QMA-hard under many-one reductions.
   - Easier to "stitch together" nearest neighbor local states $\rho_{1,2}$, $\rho_{2,3}$, etc on the line.
   - Combines locally simulatable codes [Broadbent, Grilo 2022] + 1D QMA-hardness of local Hamiltonian [Hallgren, Nagaj, Narayanaswami 2013].
2. Write each qudit q_i as qubit triples T_i. Consider all possible 6-local HKP snapshots on pairs (T_i, T_{i+1}).
3. Write linear program (LP) to compute "how much probability" to put onto each local snapshot, s.t.:
   - "Local probabilities" consistent with global HKP shadow iff 1D CONSISTENCY was YES instance.
   - Leverages HKP requirement that $E[s_i] = \rho$.
   - Problem: LP returns real numbers, but each local snapshot can only occur integer number of times in shadow.
4. "Round" LP into integer program (IP) to give integer weights on local snapshots.
   - Problem: IPs generally NP-hard to solve...
   - Solution: Leverage 1D structure to solve via dynamic programming in poly-time.

Proof steps - Reduction from 1D CONSISTENCY

1. Prove 1D CONSISTENCY problem on qudits with $d = 8$ is QMA-hard under many-one reductions.
   - Easier to "stitch together" nearest neighbor local states $\rho_{1,2}$, $\rho_{2,3}$, etc on the line.
   - Combines locally simulatable codes [Broadbent, Grilo 2022] + 1D QMA-hardness of local Hamiltonian [Hallgren, Nagaj, Narayanaswami 2013].
2. Write each qudit q_i as qubit triples T_i. Consider all possible 6-local HKP snapshots on pairs (T_i, T_{i+1}).
3. Write linear program (LP) to compute "how much probability" to put onto each local snapshot, s.t.:
   - "Local probabilities" consistent with global HKP shadow iff 1D CONSISTENCY was YES instance.
   - Leverages HKP requirement that $E[s_i] = \rho$.
   - Problem: LP returns real numbers, but each local snapshot can only occur integer number of times in shadow.
4. "Round" LP into integer program (IP) to give integer weights on local snapshots.
   - Problem: IPs generally NP-hard to solve...
   - Solution: Leverage 1D structure to solve via dynamic programming in poly-time.
5. Construct global snapshots by stitching together local snapshots under appropriate permutations given by perfect matching.

Summary

- Formal definition of classical shadows and their verification
- QMA-hardness for verifying local Clifford HKP shadows
- “Dequantization” of global Clifford HKP shadows
- qc-Σ_2-completeness of shadow verification of exponentially many observables

Summary

- Formal definition of classical shadows and their verification
- QMA-hardness for verifying local Clifford HKP shadows
- "Dequantization" of global Clifford HKP shadows
- qc-Σ_2-completeness of shadow verification of exponentially many observables

Open questions

- qc-Σ_2-hardness of verifying King-Gosset-Kothari-Babbush shadows for observable set $\{I, X, Y, Z\}^n$?
- Are there cases where shadow verification be done efficiently without sampling assumptions?

Summary

- Formal definition of classical shadows and their verification
- QMA-hardness for verifying local Clifford HKP shadows
- "Dequantization" of global Clifford HKP shadows
- qc-Σ_2-completeness of shadow verification of exponentially many observables

Open questions

- qc-Σ_2-hardness of verifying King-Gosset-Kothari-Babbush shadows for observable set $\{I, X, Y, Z\}^n$?
- Are there cases where shadow verification be done efficiently without sampling assumptions?

Summary

- Formal definition of classical shadows and their verification
- QMA-hardness for verifying local Clifford HKP shadows
- "Dequantization" of global Clifford HKP shadows
- qc-Σ_2-completeness of shadow verification of exponentially many observables

Open questions

- qc-Σ_2-hardness of verifying King-Gosset-Kothari-Babbush shadows for observable set $\{I, X, Y, Z\}^n$?
- Are there cases where shadow verification be done efficiently without sampling assumptions?