Neural-Shadow Quantum State Tomography

Victor Wei,[1,2†] W. A. Coish,[1] Pooya Ronagh,[2,3,4,5] Christine A. Muschik[2,3,4]

[1]Department of Physics, McGill University, Montreal, QC, Canada. [2]Institute for Quantum Computing, University of Waterloo, Waterloo, ON, Canada.
[3]Department of Physics and Astronomy, University of Waterloo, Waterloo, ON, Canada. [4]Perimeter Institute for Theoretical Physics, Waterloo, ON, Canada.
[5]1QB Information Technologies (1QBit), Vancouver, BC, Canada.

1. Neural network quantum state tomography (NNQST)

- NNQST aims to reconstruct the target quantum state in a neural network ansatz.

$$\psi_\lambda(s) = \sqrt{p_{\lambda_1}(s)}e^{i\varphi_{\lambda_2}(s)}$$

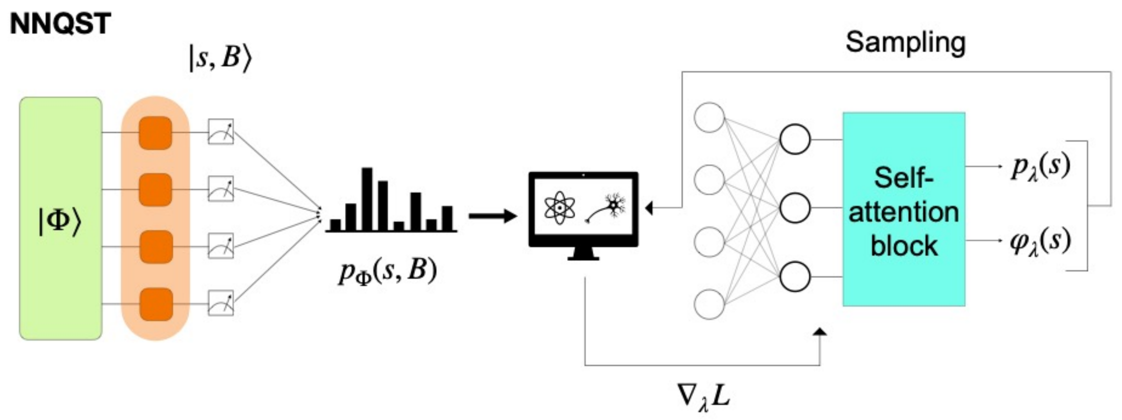

- The loss function to minimize is the cross-entropy between the two distributions (target state and neural network quantum state) averaged over a set of bases.

$$L_\lambda = -\frac{1}{|\mathcal{B}|}\sum_{B\in\mathcal{B}}\sum_{s\in\{0,1\}^n} p_\Phi(s,B)\ln p_{\psi_\lambda}(s,B).$$

2. Classical shadows

$$\hat{\rho}_i(U_i,b_i) := \mathcal{M}^{-1}(|\phi_i\rangle\langle\phi_i|) = (2^n+1)|\phi_i\rangle\langle\phi_i| - \mathbb{I}$$

$$\hat{o}^{(i)} := \mathrm{Tr}(\hat{\rho}_i O)$$

- Classical shadows (Clifford version shown above) can be used to predict observables with very few number of measurements.
- For the Clifford version, the variance of the shadow estimate is

$$\mathrm{Var}(\hat{o}) \le 3\,\mathrm{Tr}(O^2)$$

- To predict fidelity to another pure state, the number of required shadows does not grow as system size grows!

3. Neural-shadow quantum state tomography (NSQST)

- NSQST uses the collected shadows as training data and aims to reconstruct the original state in a neural network ansatz by minimizing the infidelity loss function.

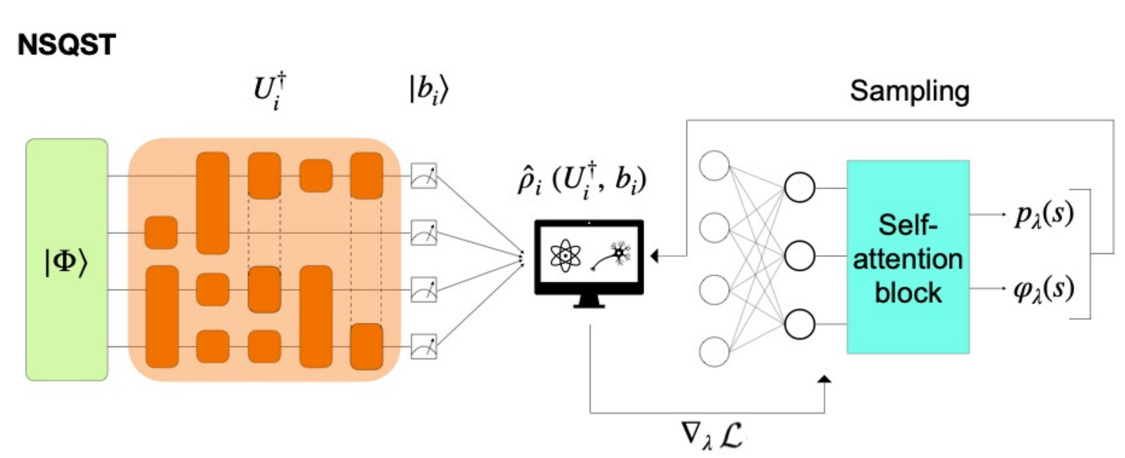

- By the proven property of (Clifford) classical shadows, the number of shadows needed to predict infidelity to a pure state with small additive error is independent of the system size.

$$\begin{aligned}\mathcal{L}_\lambda(\mathcal{E}) &:= 1 - |\langle\psi_\lambda|\Phi\rangle|^2\\ &\approx 1 - \frac{1}{N}\sum_{i=1}^{N}\mathrm{Tr}(O_\lambda\hat{\rho}_i)\\ &= 1 - \frac{1}{N}\sum_{i=1}^{N}\langle\psi_\lambda|\,\hat{\rho}_i(\mathcal{E},U_i,b_i)\,|\psi_\lambda\rangle\\ &= 1 - \frac{1}{2^n}\left(1-\frac{1}{f(\mathcal{E})}\right) - \frac{1}{Nf(\mathcal{E})}\sum_{i=1}^{N}|\,\langle\phi_i|\psi_\lambda\rangle\,|^2\end{aligned}$$

- NSQST with pre-training combines the resources used in NNQST and the standard NSQST protocol.

NSQST with pre-training

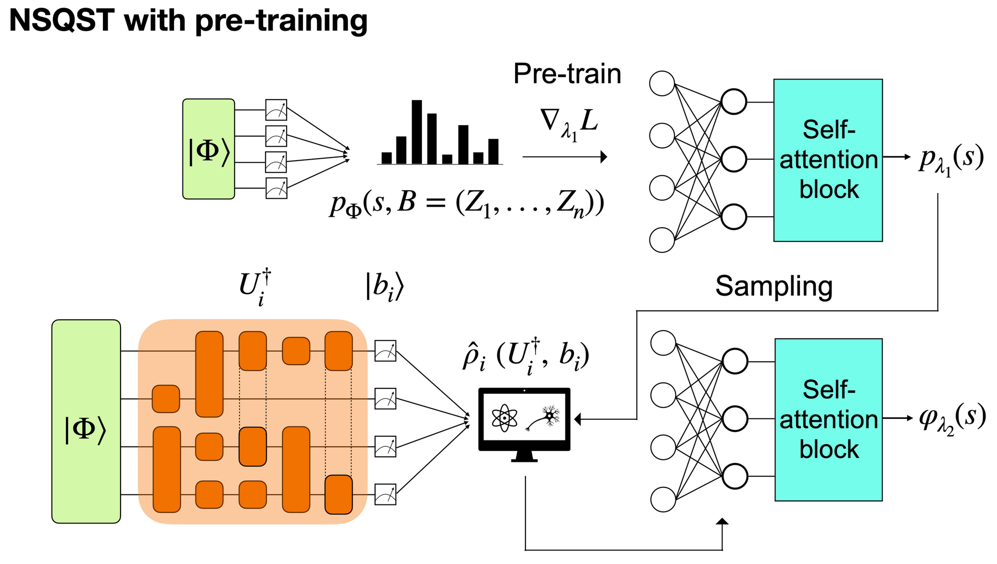

4. Trotterized time evolution with gauge group SU(3)

- To have a qubit formulation, we apply gauge transformation to eliminate the gauge degrees of freedom from the Hamiltonian, followed by a Jordan-Wigner transformation.

$$H_{SU(3)} = H_{kin} + \tilde{m}H_m + \frac{1}{2x}H_e,$$

where

$$\begin{aligned}H_{kin} &= -\frac{1}{2}(\sigma_1^+\sigma_2^z\sigma_3^z\sigma_4^- - \sigma_2^+\sigma_3^z\sigma_4^z\sigma_5^-\\ &\qquad + \sigma_3^+\sigma_4^z\sigma_5^z\sigma_6^- + \mathrm{H.c.}),\\ H_m &= \frac{1}{2}\left(6 - \sigma_1^z - \sigma_2^z - \sigma_3^z + \sigma_4^z + \sigma_5^z + \sigma_6^z\right),\\ H_e &= \frac{1}{3}\left(3 - \sigma_1^z\sigma_2^z - \sigma_1^z\sigma_3^z - \sigma_2^z\sigma_3^z\right),\end{aligned}$$

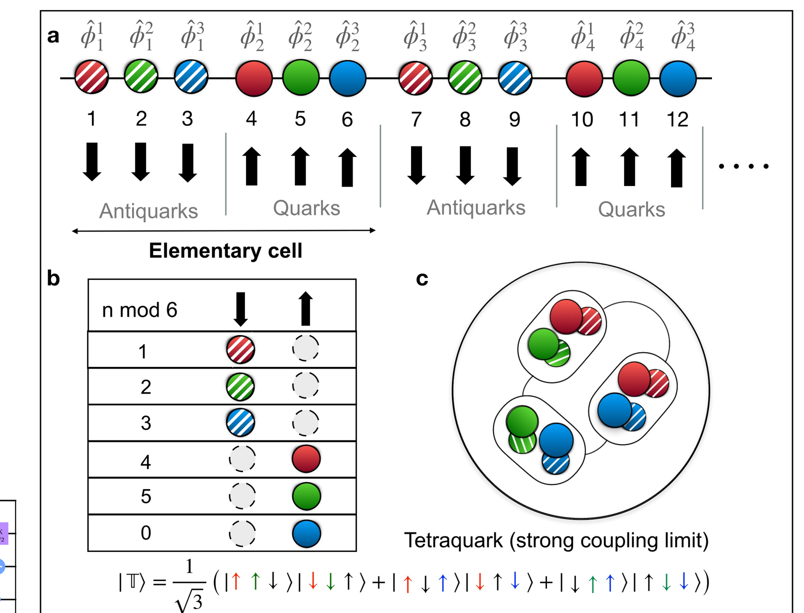

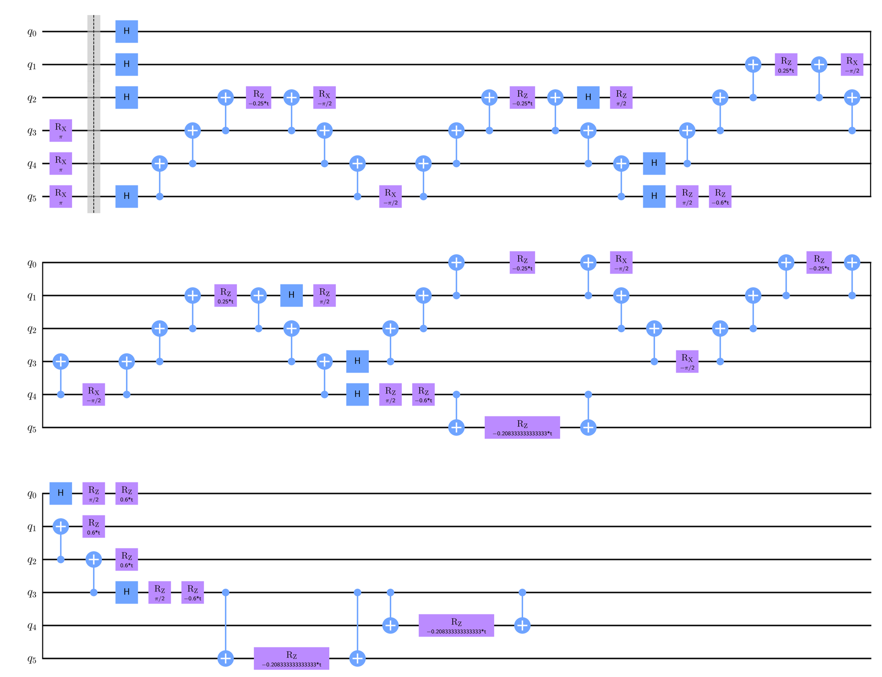

- To implement time evolution on a quantum computer, we use the Trotterized time evolution operator.
- A single 6-qubit Trotter step quantum circuit is shown on the left.

5. Advantages over NNQST

- We perform tomography of a 6-qubit time-evolved state (with four Trotter steps) under the SU(3) Hamiltonian.
- NNQST fails to learn the relative phases of the target state and predicts kinetic energy incorrectly.
- NSQST with pre-training combines the best of two worlds: learning the probability distribution from cross-entropy loss and the relative phases from shadow-based infidelity loss.

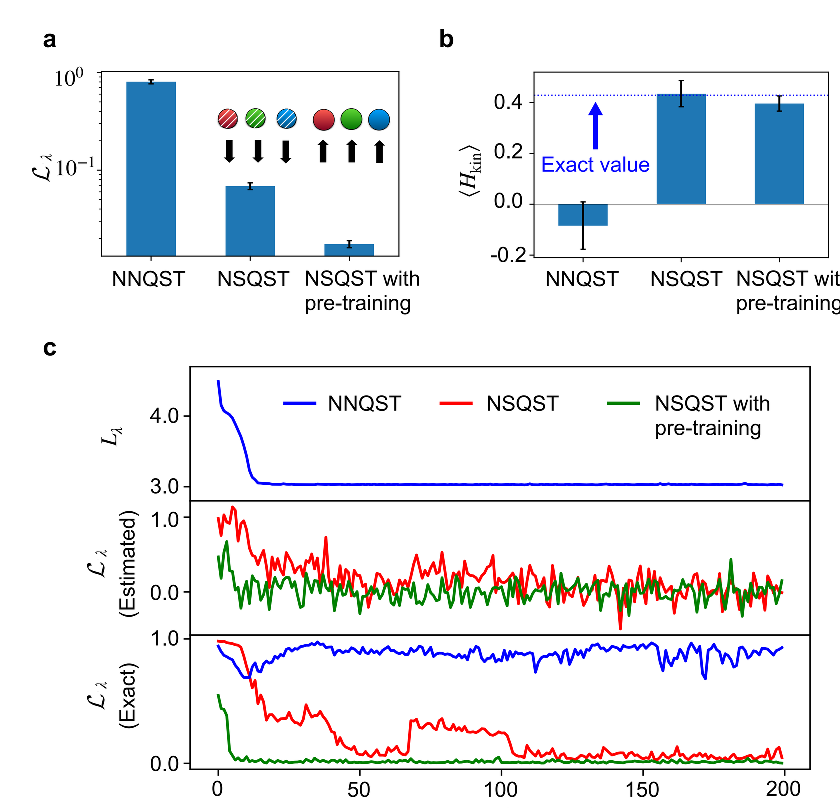

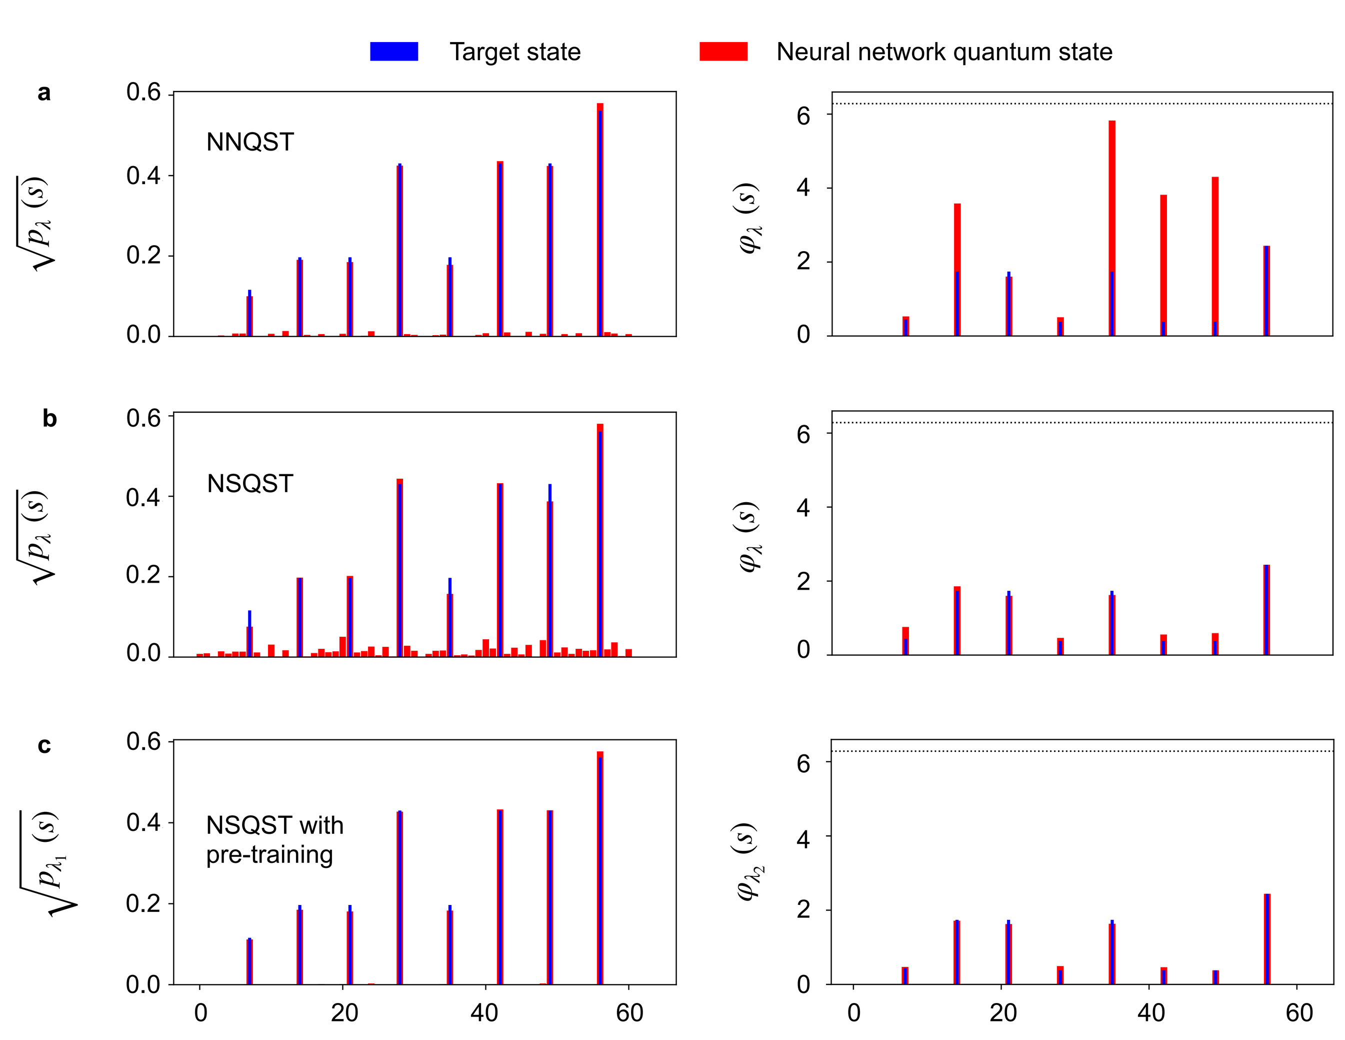

6. Advantages over direct shadow estimation

- It's important to investigate NSQST's advantages over direct shadow estimation. In other words, does the trained neural network quantum state have more predictive power than the shadow training data alone?
- Theoretically, Pauli shadows are known to be sample-efficient for predicting local observables but not global observables, whereas Clifford shadows are known to be sample-efficient at predicting low-rank observables (such as fidelity) but not high-rank observables (such as Pauli observables) irrespective of locality.
- For 6-qubit time-evolved state from Trotterized SU(3) time evolution, we compare the prediction error of various protocols.
- The same number of measurements is used in each protocol.

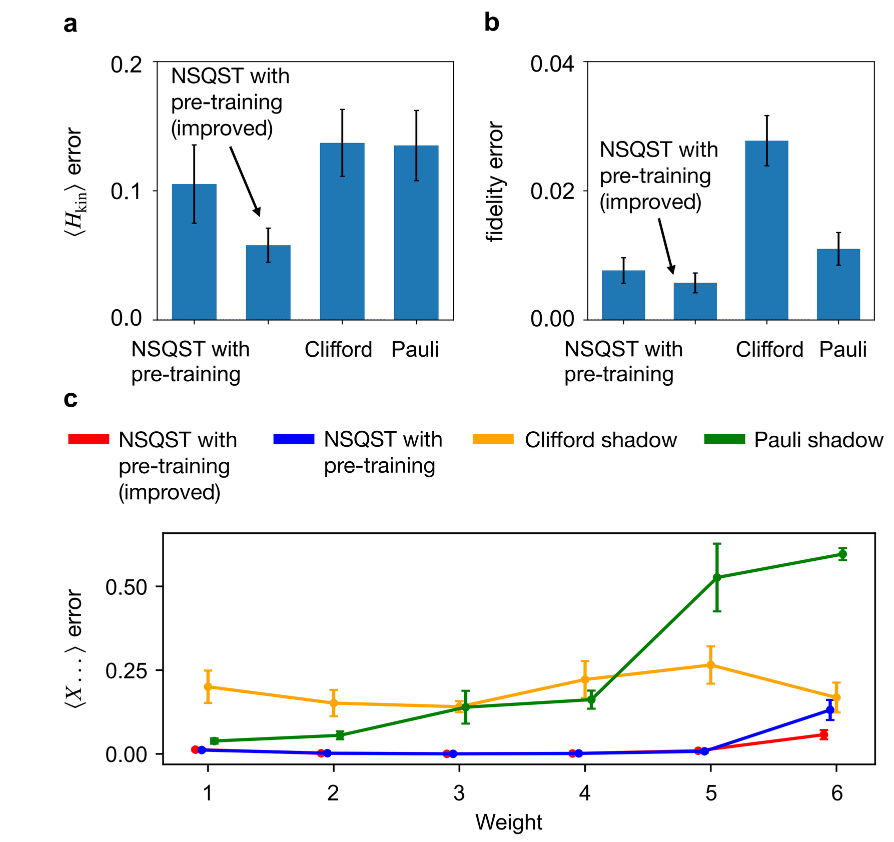

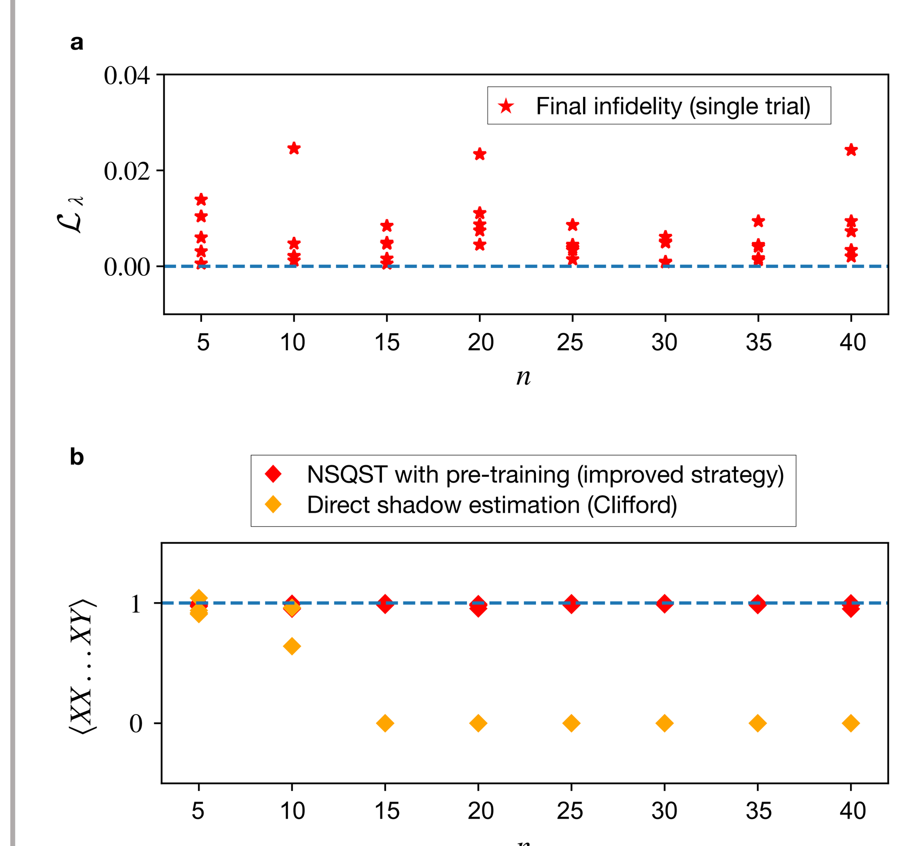

- For phase-shifted GHZ state, we compare NSQST with pretraining and direct (Clifford) shadow estimation up to 40 qubits, while using the same number of measurements in each protocol.
- The target state is a stabilizer state, and the measured Pauli string is one of its stabilizers.

7. Noise robustness

- We perform tomography of a 6-qubit phase-shifted GHZ state in the presence of a noise channel.

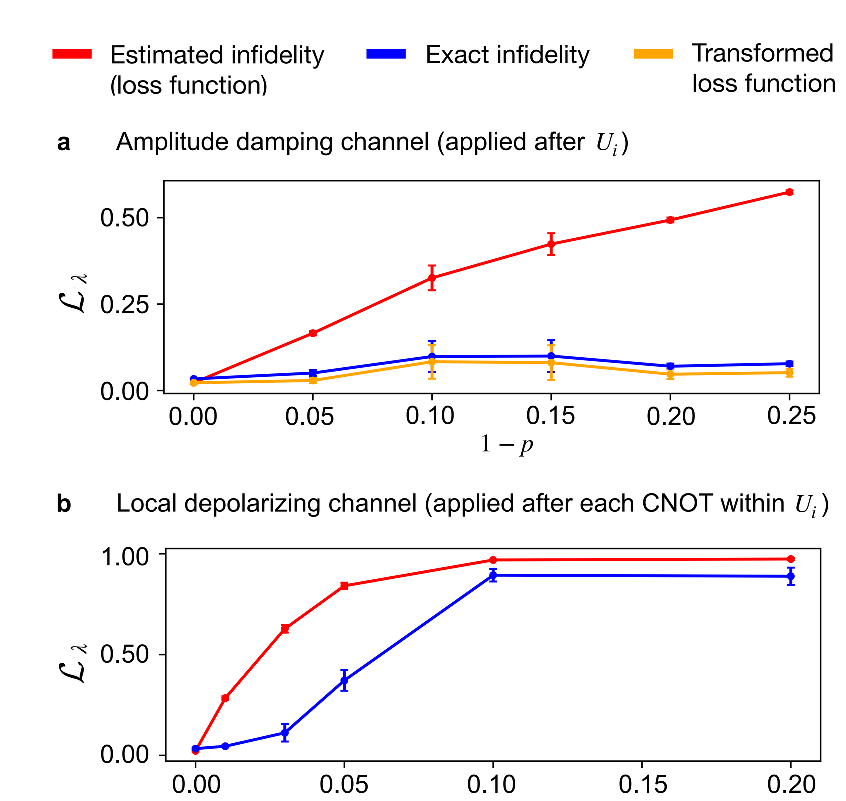

References:

1. Victor Wei, W. A. Coish, Pooya Ronagh, and Christine A. Muschik, "Neural-Shadow Quantum State Tomography", arXiv preprint arXiv:2305.01078 (2023).
2. Y. Y. Atas, J. F. Haase, J. Zhang, V. Wei, S. M.-L. Pfaendler, R. Lewis, and C. A. Muschik, "Real-time Evolution of SU(3) Hadrons on a Quantum Computer", arXiv preprint arXiv:2207.03473 (2022).