

Neural-shadow quantum state tomography

Victor Wei ^{1,2,*} W. A. Coish ^{1,†} Pooya Ronagh^{2,3,4,5,‡} and Christine A. Muschik^{2,3,4,§}¹Department of Physics, McGill University, Montreal, Quebec H3A 2T8, Canada²Institute for Quantum Computing, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada³Department of Physics & Astronomy, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada⁴Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada⁵1QB Information Technologies (1QBit), Vancouver, British Columbia V6E 4B1, Canada

(Received 8 May 2023; accepted 27 March 2024; published 6 June 2024)

Quantum state tomography (QST) is the art of reconstructing an unknown quantum state through measurements. It is a key primitive for developing quantum technologies. Neural network quantum state tomography (NNQST), which aims to reconstruct the quantum state via a neural network ansatz, is often implemented via a basis-dependent cross-entropy loss function. State-of-the-art implementations of NNQST are often restricted to characterizing a particular subclass of states, to avoid an exponential growth in the number of required measurement settings. To provide a more broadly applicable method for efficient state reconstruction, we present “neural-shadow quantum state tomography” (NSQST)—an alternative neural network-based QST protocol that uses infidelity as the loss function. The infidelity is estimated using the classical shadows of the target state. Infidelity is a natural choice for training loss, benefiting from the proven measurement sample efficiency of the classical shadow formalism. Furthermore, NSQST is robust against various types of noise without any error mitigation. We numerically demonstrate the advantage of NSQST over NNQST at learning the relative phases of three target quantum states of practical interest, as well as the advantage over direct shadow estimation. NSQST greatly extends the practical reach of NNQST and provides a novel route to effective quantum state tomography.

DOI: [10.1103/PhysRevResearch.6.023250](https://doi.org/10.1103/PhysRevResearch.6.023250)

I. INTRODUCTION

Efficient methods for state reconstruction are essential in the development of advanced quantum technologies. Important applications include the efficient characterization, readout, processing, and verification of quantum systems in a variety of areas ranging from quantum computing and quantum simulation to quantum sensors and quantum networks [1–6]. However, with physical quantum platforms growing larger in recent years [7], reconstructing the target quantum state through brute-force quantum state tomography (QST) has become much more computationally demanding due to an exponentially increasing number of required measurements. To address this issue, various approaches have been proposed that are efficient in both the number of required measurement samples and in the number of parameters used to characterize the quantum state. These include classical shadows [8] and neural network quantum state tomography (NNQST) [9]. The

goal of NNQST is to produce a neural network representation of a complete physical quantum state that is close to some target state. In contrast, the classical shadows formalism does not aim to reconstruct a full quantum state, but rather to obtain a reduced classical description that allows for efficient evaluation of certain observables.

A neural network quantum state ansatz has been shown to have sufficient expressivity to represent a wide range of quantum states [10–13] using a number of model parameters that scales polynomially in the number of qubits. Furthermore, as methods for training neural networks have long been investigated in the machine learning community, many useful strategies for neural network model design and optimization have been directly adopted for NNQST [14–16]. Following the introduction of neural network quantum states [17], Torlai *et al.* proposed the first version of NNQST, an efficient QST protocol based on a restricted Boltzmann machine (RBM) neural network ansatz and a cross-entropy loss function [9]. NNQST has been applied successfully to characterize various pure states, including W states, the ground states of many-body Hamiltonians, and time-evolved many-body states [9,18,19]. Despite the promising results of NNQST in many use cases, the protocol faces a fundamental challenge: An exponentially large number of measurement settings is required to identify a general unknown quantum state (although a polynomial number is sufficient in some examples [20]). During NNQST, a series of measurements is performed in random local Pauli bases B for n qubits ($B = (P_1, P_2, \dots, P_n)$, where $P_i \in \{X, Y, Z\}$). Because this set is exponentially large,

*victor.wei203@gmail.com

†william.coish@mcgill.ca

‡pooya.ronagh@uwaterloo.ca

§christine.muschik@uwaterloo.ca

some convenient subset of all possible B must be selected for a large system, but this subset may limit the ability of NNQST to identify certain states. An important example is the phase-shifted multiqubit GHZ state, relevant to applications such as quantum sensing. In this case, the relative phase associated with nonlocal correlations cannot be captured by measurement samples from almost-diagonal local Pauli bases, i.e., bases with $m \ll n$ indices i for which $P_i = X$ or $P_i = Y$. Nonetheless, this limited set of almost-diagonal Pauli bases is widely used in NNQST implementations to avoid an exponential cost in classical postprocessing [9,18].

To address this challenge, we use the classical shadows of the target quantum state to estimate the infidelity between the model and target states. This is in contrast with approaches that use the conventional basis-dependent cross-entropy as the training loss for the neural network. This choice is motivated by two main factors. First, infidelity is a natural candidate for a loss function compared to cross-entropy; the magnitude of the basis-dependent cross-entropy loss is in general not indicative of the distance between the neural network quantum state and the target state. Additionally, infidelity is the squared Bures distance [21], a measure of the statistical distance between quantum states that enjoys metric properties such as symmetry and the triangle inequality. The infidelity is therefore a better behaved objective function for optimization. Second, the classical shadow formalism of Huang *et al.* was originally developed to address precisely the scaling issues of brute-force QST [8]. Instead of reconstructing the unknown state, shadow-based protocols, first proposed by Aaronson [22], predict certain properties of the quantum state with a polynomial number of measurements. Therefore, classical shadows provide the following two main advantages in our work: (i) they are provably efficient in the number of required measurement samples for predicting various observables (e.g., the infidelity), and (ii) there is no choice of measurement bases required and therefore no previous knowledge of the target state is assumed.

Our new pure-state QST protocol, “neural-shadow quantum state tomography” (NSQST), reconstructs the unknown quantum state in a neural network quantum state ansatz by using classical shadow estimations of the gradients of infidelity for training [Fig. 1(b)]. In our numerical experiments, NSQST demonstrates clear advantages in three example tasks: (i) reconstructing a time-evolved state in one-dimensional quantum chromodynamics, (ii) reconstructing a time-evolved state for an antiferromagnetic Heisenberg model, and (iii) reconstructing a phase-shifted multiqubit GHZ state. Moreover, the natural appearance and inversion of a depolarizing channel from randomized measurements in the classical shadow formalism makes NSQST noise-robust without any calibration or modifications to the loss function, while one of these two extra steps is required in noise-robust classical shadows [23,24]. We numerically demonstrate NSQST’s robustness against two of the most dominant sources of noise across a wide range of physical implementations: two-qubit CNOT errors and readout errors. The rest of this paper is organized as follows: In Sec. II, we summarize the methods used in our numerical simulations, including the neural network quantum state ansatz, NNQST, classical shadows, NSQST, and NSQST with pretraining. In Sec. III and Sec. IV, we provide numerical

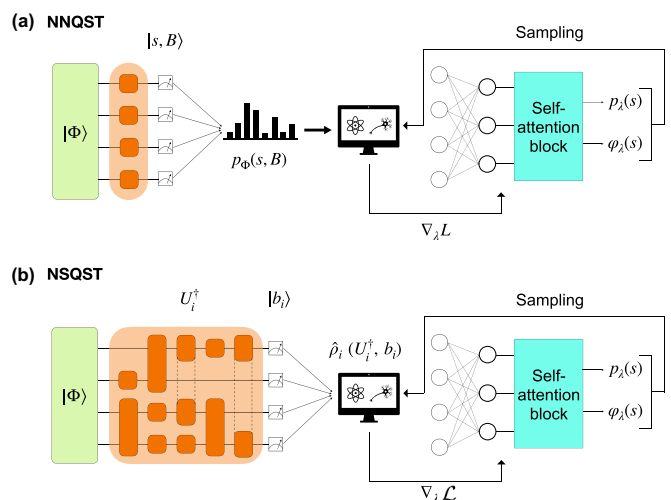


FIG. 1. Overview of the NNQST and NSQST protocols. Panel (a) shows the NNQST protocol with the cross-entropy loss function \mathcal{L}_λ . The training data determine $p_\Phi(s, B)$, the measured probability distribution of measurement outcomes s for measurements of the target state $|\Phi\rangle$ performed in the local Pauli basis B . The feedback loop on the right-hand side indicates the iterative first-order optimization for neural network training. Panel (b) displays the NSQST protocol described in Sec. II D, where the training data set consists of classical shadows only and where the network parameters λ are trained via an infidelity loss function \mathcal{L}_λ . The expression $\hat{p}_i(U_i^\dagger, b_i)$ is the stored classical shadow of the target state $|\Phi\rangle$ with the Clifford unitary U_i^\dagger and bit-string $|b_i\rangle$.

simulation results for NNQST, NSQST, and NSQST with pretraining in three useful examples, both noise-free and in the presence of noise. In particular, we provide a comparison to direct shadow estimation in Sec. III D. Finally, Sec. V summarizes the key advantages of NSQST and some possible future directions. We provide additional technical details and suggestions for further improvements to NSQST in the Appendices.

II. METHODS

In this section, we describe existing methods for characterizing quantum states and then introduce and describe two variants of NSQST. We introduce neural network quantum states in Sec. II A. State-of-the-art NNQST implementations and the classical shadow protocol are summarized in Sec. II B and Sec. II C, respectively. Our proposed NSQST protocol is described in Sec. II D. In addition, a modified NSQST protocol with pretraining is described in Sec. II E.

A. Neural network quantum state

Our pure-state neural network quantum state ansatz is adopted from Ref. [18]. The parameterized model is based on the transformer architecture [15], widely used in natural language processing and computer vision [25,26]. As compared to older architectures such as the RBM, the transformer is superior in modeling long-range interactions and allows for more efficient sampling of the encoded probability distribution due to its autoregressive property [18].

The transformer neural network quantum state ansatz takes a bit-string $s = (s_1, \dots, s_n) \in \{0, 1\}^n$ corresponding to the computational basis state $|s\rangle$, and produces a complex-valued amplitude $\langle s|\psi_\lambda\rangle = \psi_\lambda(s)$ parameterized by $\lambda = (\lambda_1, \lambda_2)$ as

$$\psi_\lambda(s) = \sqrt{p_{\lambda_1}(s)} e^{i\varphi_{\lambda_2}(s)}, \quad (1)$$

where λ_1 and λ_2 are vectors of real-valued model parameters for the normalized probability amplitudes $p_{\lambda_1}(s)$ and the phases $\varphi_{\lambda_2}(s)$ of the neural network quantum state. These amplitudes and phases may be parameterized by the neural network quantum state in various fashions. One approach is to use two completely disjoint models, independently parameterized by λ_1 and λ_2 for the amplitude and phase values, respectively [9,27]. Another approach is to use a single model parameterized by λ to encode both the amplitude and phase outputs, either via complex-valued model parameters [17,28] or by using real-valued model parameters with two disjoint layers of output neurons connected to a common preceding neural network [18]. In our numerical experiments, we use the later parametrization for NNQST and NSQST, but the modified NSQST protocol with pretraining in Sec. II E uses two separately parameterized neural networks. See Appendix A for a more detailed account of the transformer architecture.

Given a trained neural network quantum state, observables and other state properties of interest can be predicted by drawing (classical) samples from the neural network model. The number of samples required to predict the expectation value of an arbitrary Pauli string (independent of its weight) and fidelity to a computationally tractable state with bounded additive error is independent of the system size [29]. Computationally tractable states include stabilizer states and neural network quantum states. Thus, if sampling can be performed efficiently, the prediction errors from neural network quantum states are primarily due to imperfect training. We also note that not every neural network quantum state ansatz has this property of efficient observable and fidelity prediction, where an important example is a class of generative models trained on informationally complete positive-operator valued measures (IC-POVMs) [30,31].

B. Neural network quantum state tomography (NNQST)

NNQST [Fig. 1(a)] aims at obtaining a trained neural network representation that closely approximates an unknown target quantum state. The training is done by iteratively adjusting the neural network parameters along a loss gradient estimated from the measurement samples in various local Pauli bases (obtained by applying single-qubit rotations before performing measurements in the computational basis) [18]. We denote a local Pauli basis as $B = (P_1, P_2, \dots, P_n)$, with $P_i \in \{X, Y, Z\}$. If a measurement sample $s \in \{0, 1\}^n$ is obtained after performing rotations to the Pauli basis B , we store the pair (s, B) as a training sample, corresponding to a product state $|s, B\rangle$.

After choosing a subset \mathcal{B} of Pauli bases for collecting measurement samples, we estimate a loss function that represents the distance between the target state $|\Phi\rangle$ and the neural network quantum state $|\psi_\lambda\rangle$. The loss function in NNQST is based on the cross-entropy of the measurement outcome distributions for the target and neural network states in each basis

B , which is then averaged over the set of bases \mathcal{B} . Ignoring a λ -independent contribution arising from the average entropy of the target-state measurement distribution, this procedure gives the cross-entropy loss function for NNQST [18]:

$$L_\lambda = -\frac{1}{|\mathcal{B}|} \sum_{B \in \mathcal{B}} \sum_{s \in \{0, 1\}^n} p_\Phi(s, B) \ln p_{\psi_\lambda}(s, B). \quad (2)$$

Here, $p_\Phi(s, B)$ is the probability of measuring the outcome s from the target state $|\Phi\rangle$ after rotating to the Pauli basis B and $p_{\psi_\lambda}(s, B)$ is defined as

$$p_{\psi_\lambda}(s, B) = \left| \sum_{\substack{t \in \{0, 1\}^n \\ \langle s, B|t\rangle \neq 0}} \langle s, B|t\rangle \langle t|\psi_\lambda\rangle \right|^2, \quad (3)$$

where the overlap between the Pauli product state and the neural network quantum state requires a summation over the computational basis states $|t\rangle$ that satisfy $\langle s, B|t\rangle \neq 0$. Note that the number of these states $|t\rangle$ is 2^K , with K being the number of positions i where $P_i \neq Z$. This suggests that an efficient and exact calculation of $p_{\psi_\lambda}(s, B)$ requires the projective measurements to be in almost-diagonal Pauli bases for a generic neural network quantum state $|\psi_\lambda\rangle$ [18].

Using the law of large numbers, the cross-entropy loss can be approximated via a finite training data set \mathcal{D}_T as

$$L_\lambda \approx -\frac{1}{|\mathcal{D}_T|} \sum_{(s, B) \in \mathcal{D}_T} \ln p_{\psi_\lambda}(s, B). \quad (4)$$

An approximation for the gradient $\nabla_\lambda L$ is then directly found from Eq. (4). During training, the gradient is provided to an optimization algorithm such as stochastic gradient descent (SGD) or one of its variants (e.g., the Adam optimizer [16]). In this paper, we exclusively use the Adam optimizer.

C. Classical shadows

Shadow tomography relies on the ingenious observation that a polynomial number of measurement samples is sufficient to predict certain observables for quantum states of arbitrary size [22]. The classical shadow protocol further exploits the efficiency of the stabilizer formalism, making this procedure ready for practical experiments [8,32–34]. In this paper, we focus on estimating linear observables of the form $\text{Tr}(O\rho)$ for a pure state $\rho = |\Phi\rangle\langle\Phi|$. An important example (for $O = |\psi_\lambda\rangle\langle\psi_\lambda|$) is the fidelity between the target state $|\Phi\rangle$ and a reference state $|\psi_\lambda\rangle$. The first step in the protocol is to collect the so-called classical shadows of $|\Phi\rangle$. To obtain a single classical shadow sample, we apply a randomly sampled Clifford unitary $U_i \in \text{Cl}(2^n)$ to the quantum state and measure all n qubits in the computational basis, resulting in a single bit-string $|b_i\rangle$. The stabilizer states $|\phi_i\rangle = U_i^\dagger |b_i\rangle$ contain valuable information about ρ . Using representation theory [35], it can be shown that the density matrix obtained from an average over both the random unitaries and the measured bit-strings $\mathcal{M}(|\Phi\rangle\langle\Phi|) := \mathbb{E}_{U \sim \text{Cl}(2^n), b \sim P_\Phi(b)} [|\phi\rangle\langle\phi|]$ coincides with the outcome of a depolarizing noise channel:

$$\mathcal{M}(\rho) = \mathcal{D}_{n, 1/(2^n+1)}(\rho), \quad (5)$$

where $\mathcal{D}_{n,f}(\rho) = f\rho + (1-f)\frac{\mathbb{I}}{2^n}$ denotes the depolarizing noise channel of strength f . The original state can then be recovered as an average over classical shadows by inverting the above formula, $|\Phi\rangle\langle\Phi| = \mathbb{E}[\mathcal{M}^{-1}(|\phi\rangle\langle\phi|)]$. We emphasize that the original state can only be recovered after sampling from a prohibitively large number of Clifford unitaries, and for each of them sampling an exponentially large number of bit-strings. The classical shadows are therefore defined by [8]

$$\hat{\rho}_i(U_i, b_i) := \mathcal{M}^{-1}(|\phi_i\rangle\langle\phi_i|) = (2^n + 1)|\phi_i\rangle\langle\phi_i| - \mathbb{I}. \quad (6)$$

More generally, in the presence of a gate-independent, time-stationary, and Markovian noise channel \mathcal{E} afflicting the segment of the circuit between the preparation of the state ρ and the measurements, this definition extends to [23]

$$\hat{\rho}_i(\mathcal{E}, U_i, b_i) = \frac{1}{f(\mathcal{E})}|\phi_i\rangle\langle\phi_i| + \left(1 - \frac{1}{f(\mathcal{E})}\right)\frac{\mathbb{I}}{2^n}. \quad (7)$$

Here, $f(\mathcal{E})$ is the strength of a depolarizing noise channel $\mathcal{D}_{n,f(\mathcal{E})}$ composed of the combined effects of the channel in Eq. (5) and the twirling of the additional noise by the random Clifford unitaries, effectively imposing further depolarization. Koh and Grewal [23] derived an analytic expression for $f(\mathcal{E})$ as

$$f(\mathcal{E}) = \frac{\text{fid}(\mathcal{E}) - 1}{2^{2n} - 1}, \quad (8)$$

where

$$\text{fid}(\mathcal{E}) = \text{Tr}(\mathcal{E} \circ \text{diag}) = \sum_{s \in \{0,1\}^n} \langle s | \mathcal{E}(|s\rangle\langle s|) | s \rangle \quad (9)$$

is the sum of fidelities for the noise channel \mathcal{E} acting on each of the computational basis states $|s\rangle$ ($\text{fid}(\mathcal{E}) \in [0, 2^n]$; at the lower bound, $\text{fid}(\mathcal{E}) = 0$, the depolarizing parameter becomes negative, $f(\mathcal{E}) < 0$, but the associated depolarizing channel remains physical [23]). When the noise channel is not exactly known, extra calibration procedures are required in noise-robust classical shadow protocols [24].

Once the classical shadow samples $\{\hat{\rho}_i\}_{i=1}^N$ are collected, we calculate $\hat{\delta}^{(i)} := \text{Tr}(\hat{\rho}_i O)$ for each of the N classical shadows and obtain an estimator for the observable from an average over the N samples (alternatively, the median-of-means can be used to improve the success rate of the protocol; see Ref. [8] for more details). The key advantage of classical shadows is the bounded variance of observable estimations which, in turn, provides a bound on the number of classical shadow samples required to predict linear observables of the quantum state within a target precision. Indeed, as shown in Refs. [23,24], the number of classical shadow samples N required to estimate M arbitrary linear observables $\{O_j\}_{j=1}^M$, up to an additive error ε , scales as

$$N = \mathcal{O}\left(\frac{2^{2n} \log M}{\varepsilon^2 \text{fid}(\mathcal{E})^2} \max_{1 \leq j \leq M} \text{Tr}(O_j^2)\right). \quad (10)$$

In the case of noise-free Clifford tails, $\mathcal{E} = \mathbb{I}$, we have $\text{fid}(\mathbb{I}) = 2^n$, $f(\mathbb{I}) = (1 + 2^n)^{-1}$, and we recover the sample complexity $\mathcal{O}(\max_{1 \leq j \leq M} \text{Tr}(O_j^2) \log M / \varepsilon^2)$ presented in Ref. [8], which is indeed independent of the system size n . The variance of observable estimators is also bounded

as $\text{Var}(\hat{\delta}) \leq 3\text{Tr}(O^2)$ in this case and is independent of the system size.

D. Neural-shadow quantum state tomography (NSQST)

Given a pure target state $|\Phi\rangle$, our goal in NSQST [Fig. 1(b)] is to progressively adjust the model parameters λ , such that the associated pure state $|\psi_\lambda\rangle$ [see Eq. (1)] approaches $|\Phi\rangle$ during optimization. We approximate the fidelity between the model and target states using the classical shadow formalism described in Sec. II C, taking $O_\lambda = |\psi_\lambda\rangle\langle\psi_\lambda|$ as a linear observable, averaged with respect to $\rho = |\Phi\rangle\langle\Phi|$. The number M of observables we predict during optimization therefore coincides with the number of descent steps taken by the optimizer, as updating λ changes the observable $|\psi_\lambda\rangle\langle\psi_\lambda|$ in every iteration. By collecting N classical shadows, we can approximate our loss function (the infidelity) via

$$\begin{aligned} \mathcal{L}_\lambda(\mathcal{E}) &:= 1 - |\langle\psi_\lambda|\Phi\rangle|^2 \\ &\approx 1 - \frac{1}{N} \sum_{i=1}^N \text{Tr}(O_\lambda \hat{\rho}_i) \\ &= 1 - \frac{1}{N} \sum_{i=1}^N \langle\psi_\lambda|\hat{\rho}_i(\mathcal{E}, U_i, b_i)|\psi_\lambda\rangle \\ &= 1 - \frac{1}{2^n} \left(1 - \frac{1}{f(\mathcal{E})}\right) - \frac{1}{Nf(\mathcal{E})} \sum_{i=1}^N |\langle\phi_i|\psi_\lambda\rangle|^2. \end{aligned} \quad (11)$$

In the noise-free case $\mathcal{E} = \mathbb{I}$, this expression simplifies to

$$\mathcal{L}_\lambda(\mathbb{I}) \approx 2 - \frac{2^n + 1}{N} \sum_{i=1}^N |\langle\phi_i|\psi_\lambda\rangle|^2. \quad (12)$$

We see that, independent of the specific form of \mathcal{E} , training the model is simply equivalent to increasing the average overlap between the random stabilizer states and the model quantum state.

The next step in NSQST requires classical postprocessing to estimate the overlaps $\langle\phi_i|\psi_\lambda\rangle$. For certain states $|\psi_\lambda\rangle$ (e.g., stabilizer states), the overlap can be calculated efficiently. Many states of interest do not fall into this class, leading to a potential exponential overhead. However, we can obtain a Monte Carlo estimate of the overlap by sampling from the model quantum state $|\psi_\lambda\rangle$. In the model, we associate a probability

$$p_\lambda(s) = |\psi_\lambda(s)|^2 \quad (13)$$

to each computational basis state $|s\rangle$. Therefore,

$$\langle\phi_i|\psi_\lambda\rangle = \sum_s \frac{\phi_i^*(s)}{\psi_\lambda^*(s)} p_\lambda(s) = \left\langle \frac{\phi_i^*(s)}{\psi_\lambda^*(s)} \right\rangle_{\psi_\lambda}. \quad (14)$$

It is now straightforward to provide a Monte Carlo estimate of the above quantity [36] (see Appendix D for an alternative approach). For each sample s from the neural network, we have direct access to the exact complex-valued amplitudes $\psi_\lambda(s)$ of the neural network quantum state in the computational basis. Moreover, we can compute the stabilizer state

projections $\phi_i(s) = \langle s|U_i^\dagger|b_i\rangle$ in $O(n^2)$ time, in view of the Gottesman-Knill theorem [37–39]. Note that decomposing a randomly sampled Clifford operator into primitive unitary gates (e.g., with Hadamard, S, and CNOT gates) still takes $O(n^3)$ time [39,40]. However, this is a one-time procedure to be run for each U_i and can be done in advance of state tomography.

For first-order optimization methods (such as SGD and Adam), it is the gradient of the loss function rather than the loss function itself that must be estimated. From Eq. (11) and using the log-derivative trick, we obtain the gradient

$$\nabla_\lambda \mathcal{L}(\mathcal{E}) \approx \frac{-2}{Nf(\mathcal{E})} \sum_{i=1}^N \Re \left[\left\langle \frac{\phi_i^*(s)}{\psi_\lambda^*(s)} D_\lambda(s) \right\rangle_{\psi_\lambda} \left\langle \frac{\phi_i(s)}{\psi_\lambda(s)} \right\rangle_{\psi_\lambda} \right], \quad (15)$$

where we define the diagonal operator D_λ as

$$\begin{aligned} D_\lambda(s) &= \frac{1}{\psi_\lambda(s)} \frac{\partial \psi_\lambda(s)}{\partial \lambda} = \nabla_\lambda \ln \langle s|\psi_\lambda \rangle \\ &= \nabla_\lambda (\ln \sqrt{p_\lambda(s)} + i\varphi_\lambda(s)). \end{aligned} \quad (16)$$

A simple but important observation is that the noise enters Eq. (15) only in the overall prefactor $\propto 1/f(\mathcal{E})$. Thus, the noise may affect the learning rate, but it will not affect the direction of the gradient. This suggests that gradient-based optimization schemes can yield an accurate neural network quantum state without any noise calibration or mitigation. This is despite the fact that this same noise generally biases the estimated infidelity [see Eq. (11)].

We now discuss a possible limitation of our approach to classical postprocessing. Given N classical shadows collected experimentally, and L computational basis samples collected from the neural network quantum state, the number of Monte Carlo samples collected from the neural network quantum state must be $L \sim O(1/N^2 f(\mathcal{E})^2)$ to guarantee a bounded standard error in the approximation of the gradient from Eq. (15). Since $f(\mathcal{E}) \leq 1/(1+2^n)$, this suggests that there may be an exponential cost in performing the Monte Carlo estimations. We emphasize that this potential exponential cost in classical postprocessing does not affect the required number of classical shadows from measurements. As system sizes grow significantly larger, it will eventually become hopeless to perform an exact sum over all 2^n computational basis states and the Monte Carlo average may still lead to successful convergence with only a sub-exponential number of samples in some cases (further details and an alternative approach to performing the Monte Carlo average are discussed in Appendix D). In our numerical simulations with six qubits (see Sec. III), having $2^6 = 64$ computational basis states, we have evaluated Eq. (15) using 5000 Monte Carlo samples. With this many samples, the Monte Carlo estimation error is negligible and statistical fluctuations in the gradient are predominantly due to the finite number of classical shadows collected.

E. NSQST with pretraining

Along with the standard NSQST protocol, we also outline a modified NSQST protocol we call “NSQST with pretraining,” which combines the resources used in NNQST and the

NSQST with pre-training

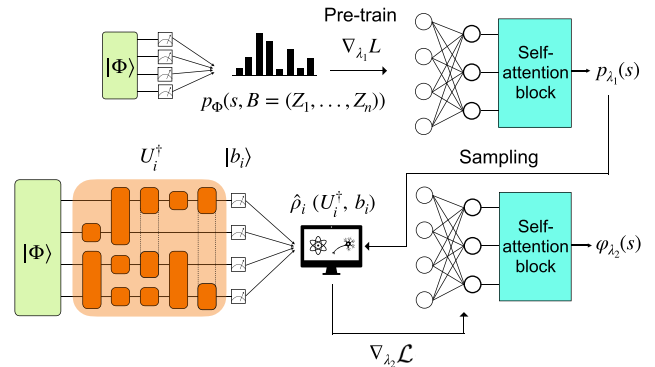


FIG. 2. Overview of the NSQST with pretraining protocol. The neural networks learning $p_{\lambda_1}(s)$ and $\varphi_{\lambda_2}(s)$ are separately parameterized, and $p_{\lambda_1}(s)$ is pretrained using NNQST with training data derived from measurements only in the computational basis.

standard NSQST protocol. NSQST with pretraining aims to find a solution with a lower infidelity than either of the other protocols alone. In this protocol we train two models with disjoint sets of parameters λ_1 and λ_2 . We call these models the *probability amplitude model* and the *phase model*. Figure 2 provides a visual overview of the protocol for NSQST with pretraining.

First, the parameters λ_1 are optimized to produce an accurate distribution $p_{\lambda_1}(s) \simeq p_\Phi(s, B)$ from measurements performed exclusively in the computational basis $[B = (Z_1, Z_2, \dots, Z_n)]$. Note that we can efficiently evaluate the loss function, Eq. (4), and its gradient in this case as they depend only on the probabilities and not on the phases. Next, we perform NSQST to train the model parameters λ_2 , learning the phases $\varphi_{\lambda_2}(s)$. However, unlike the case of standard NSQST, to perform a Monte Carlo estimate of the gradient, Eq. (15), here we select random samples from a set of computational basis states s according to the pretrained distribution $p_{\lambda_1}(s)$. Since the NSQST Monte Carlo approximations do not follow the model λ_2 in an *on policy* fashion, resampling in every iteration is no longer necessary. Nevertheless, we still resampled in our numerical experiments (described below) to reduce the sampling bias, and since the classical sampling procedure was not computationally costly in our examples.

NSQST with pretraining resembles coordinate descent optimization, with λ_1 and λ_2 being the two coordinate blocks. Optimizing λ_1 first and fixing it for the optimization of λ_2 reduces the dimension of the parameter space for the optimizers throughout the training. However, this does not guarantee convergence to a better local minimum in the loss landscape. We do not intend to demonstrate a clear advantage for NSQST with pretraining over the standard NSQST protocol, as the former uses more computational resources both experimentally (by requiring more measurements) and classically (in the form of the memory, time, and energy consumed to train the neural network quantum state). See Appendix C for a comparison of the number of model parameters and the number of measurements used for each of the two approaches. Another motivation for introducing this modified NSQST protocol is to provide new perspectives on the differences between learning

the probability amplitudes and the phases of a target quantum state, as well as to inspire other useful hybrid protocols in the future.

III. NUMERICAL SIMULATIONS WITHOUT NOISE

In this section, we first demonstrate the advantage of our NSQST protocols over the NNQST protocol in three physically relevant scenarios, then demonstrate advantages over direct shadow estimation. Specifically, we consider a model from high-energy physics (time evolution for one-dimensional quantum chromodynamics), a model from condensed-matter physics (time evolution of a Heisenberg spin chain), and a model relevant to precision measurements and quantum information science (a phase-shifted GHZ state).

For all three physical settings, we compare the performance of NNQST, NSQST, and NSQST with pretraining by measuring the exact infidelity of the trained model quantum states to the target state averaged over the last 100 iterations of training (or epochs for NNQST, see Appendix C). For NNQST's basis selection, since none of our target states is known to be the ground state of a k -local Hamiltonian (i.e., a Hamiltonian with each term acting nontrivially on at most k qubits), we simply use all of the almost-diagonal and nearest-neighbour local Pauli bases (i.e., Pauli bases with at most two neighboring terms being non- Z). The number of these bases scales linearly with the system size ($4n - 3$ bases). All NSQST protocols use only $N = 100$ resampled classical shadows per iteration for model parameter updates. We perform ten independent trials of each protocol (NNQST, NSQST, and NSQST with pretraining) in each of the three examples.

Finally, we adopt an improved pretraining strategy described in Appendix D and fix $N = 200$ Clifford shadows without resampling to demonstrate advantages of NSQST over direct shadow estimation with Clifford shadows or Pauli shadows.

A. Time-evolved state in one-dimensional quantum chromodynamics

Quantum chromodynamics (QCD) studies the fundamental strong interaction responsible for the nuclear force [41]. Lattice gauge theory, an important nonperturbative tool for studying QCD, discretizes spacetime into a lattice and the continuum results can be obtained through extrapolation [42]. Although lattice gauge theory has been extremely successful in QCD studies, simulations of many important physical phenomena such as real-time evolution are still out of reach due to the sign problem in current simulation techniques. Quantum computers are envisioned to overcome this barrier in lattice gauge theory-based QCD simulations and they may open the door to new discoveries in QCD [43–46].

We consider a Trotterized time evolution with the gauge group $SU(3)$ and aim to reconstruct the time-evolved quantum state after a given amount of time. To this end, we use the qubit formulation in Ref. [47] and study a single unit cell of the lattice. This corresponds to $n = 6$ qubits representing three quarks (red, green, blue) and three antiquarks (antired, antigreen, antiblue) as shown in Fig. 3(a). The Trotterized time evolution starts from the initial state

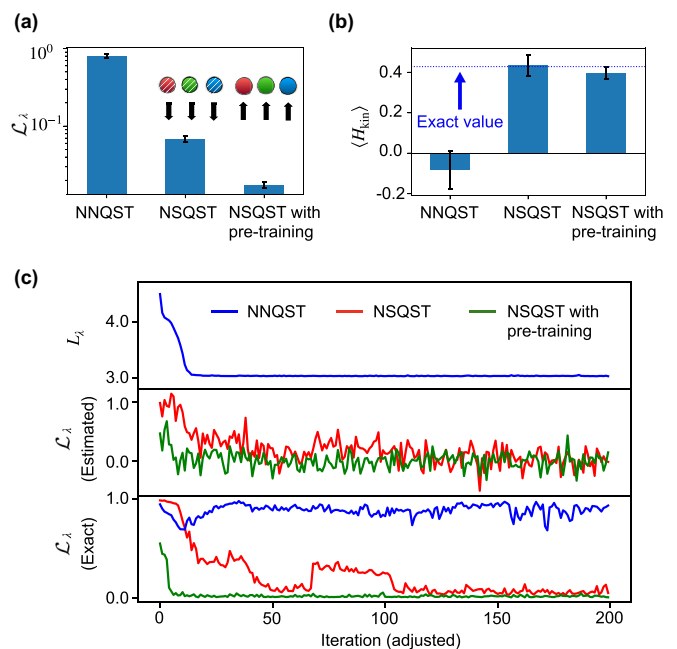


FIG. 3. Tomography of the quantum state following a one-dimensional QCD time evolution. Two Trotter steps are used for a total evolution time of $t = 1.8$. Panel (a) displays the average final-state infidelity for each of the three protocols. In each trial, we extract the (exactly calculated) average state infidelity \mathcal{L}_λ , averaged over the last 100 iterations (and further averaged over ten trials). The error bar is the standard error in the mean calculated over the ten trials. The embedded schematic shows the qubit encoding for a unit cell, containing up to three quarks (filled circles) and up to three antiquarks (striped circles). In panel (b), the plot compares the expectation value of the kinetic energy, evaluated for the neural network quantum state found in the last iteration of each trial, and averaged over ten trials for each protocol. In panel c, the optimization progress curves are displayed for a typical trial, where the adjusted iteration refers to epochs for NNQST, but rather indicates increments of ten iterations for the two NSQST protocols (a total of 2000 iterations were run in these cases). Panel (c) shows the NNQST (blue) loss L_λ in the top plot, the estimated NSQST (infidelity) loss function \mathcal{L}_λ (with and without pretraining, green and red, respectively) in the middle plot (fluctuations are dominated by the finite number $N = 100$ of classical shadows taken for each estimate), and the exact infidelity is shown in every (adjusted) iteration for all three protocols in the lower plot.

$|\Phi_0\rangle = |\downarrow\downarrow\downarrow\rangle|\uparrow\uparrow\uparrow\rangle$, which is known as the strong-coupling baryon-antibaryon state. The Hamiltonian governing the evolution is (from Ref. [47]):

$$H_{SU(3)} = H_{\text{kin}} + \tilde{m}H_m + \frac{1}{2x}H_e, \quad (17)$$

where

$$\begin{aligned} H_{\text{kin}} &= -\frac{1}{2}(\sigma_1^+ \sigma_2^z \sigma_3^z \sigma_4^- - \sigma_2^+ \sigma_3^z \sigma_4^z \sigma_5^- \\ &\quad + \sigma_3^+ \sigma_4^z \sigma_5^z \sigma_6^- + \text{H. c.}), \\ H_m &= \frac{1}{2}(6 - \sigma_1^z - \sigma_2^z - \sigma_3^z + \sigma_4^z + \sigma_5^z + \sigma_6^z), \\ H_e &= \frac{1}{3}(3 - \sigma_1^z \sigma_2^z - \sigma_1^z \sigma_3^z - \sigma_2^z \sigma_3^z), \end{aligned} \quad (18)$$

with $\tilde{m} = am$ and $x = 1/(ga)^2$ for a lattice spacing a , where m is the bare quark mass and g is the gauge coupling constant.

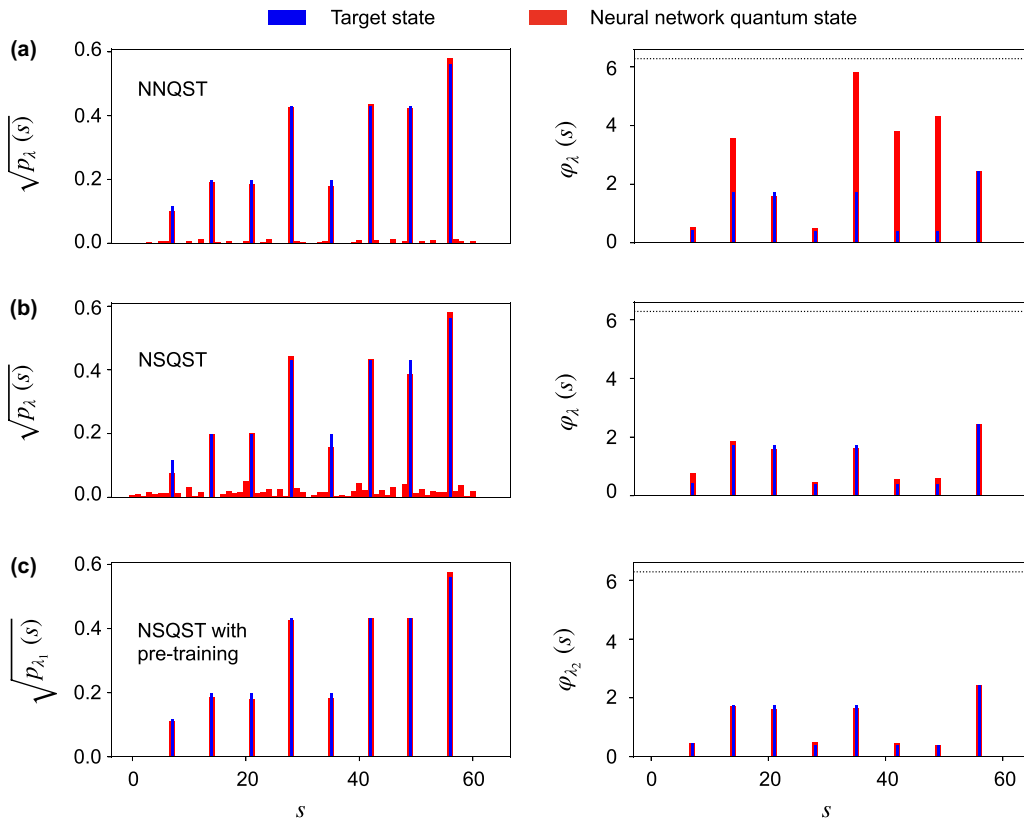


FIG. 4. Typical neural network quantum states following optimization, approximating the state after a one-dimensional QCD time evolution. Panel (a) displays a typical state found in the last iteration of NNQST training: the left plot shows the square root of the probability of the final neural network quantum state compared to the exact target state and the right plot shows the phase output of the state over the set of computational basis states s . To highlight the dominant contributions, the phase output has been truncated for states s with $\sqrt{p_{\lambda_1}(s)} < 0.1$. For the purpose of better visualization, the overall (global) phase of the neural network quantum state is chosen by aligning the phase of the most probable computational-basis state to that of the target state. The dashed line corresponds to a phase of 2π , since we choose our phase predictions to be in the range $[0, 2\pi]$. Panels (b) and (c) show typical final states from NSQST and NSQST with pretraining. We observe that both NSQST protocols succeed at learning the phase structure while NNQST fails at the same task.

We use two Trotter steps in our simulation, each for time $t = 1.8$. See Appendix B and Ref. [47] for more details on the circuit and the physical significance of this time evolution.

Figure 3 shows the results of simulating tomography on the time-evolved state $|\Phi_{SU(3)}\rangle$ using NNQST, NSQST, and NSQST with pretraining. Note that, although the NSQST protocols (with and without pretraining) are run for 2000 iterations, we use increments of ten iterations in the plot to provide a visual comparison with NNQST (which is run for 200 epochs due to faster convergence of L_λ in optimization). For NSQST with pretraining, we display the optimization progress curve only after the probabilities $p_{\lambda_1}(s)$ have been pretrained. This explains the lower initial infidelity for NSQST with pretraining. See Appendix C for further details on the simulation hyperparameters.

Based on Fig. 3, NSQST and NSQST with pretraining both result in a lower final-state infidelity, relative to NNQST, and both predict the mean kinetic energy values better than NNQST. Figure 3(c) further depicts the optimization progress curves of a typical trial. We see that for NNQST, the cross-entropy loss L_λ quickly converges with very little fluctuations, despite the continued fluctuations of the state infidelity in the lower plot near $\mathcal{L}_\lambda \simeq 1$, indicating a very small overlap with

the target state. However, standard NSQST and NSQST with pretraining both converge to a final state very close to the target state, despite fluctuations in the loss function caused by the finite number of classical shadows in each iteration. Moreover, we notice that NSQST with pretraining not only starts with a state of lower infidelity after pretraining, but also converges to a solution of lower infidelity than standard NSQST, with much more stable convergence in the end. One unexpected outcome is that NSQST with pretraining does not have a better kinetic energy prediction than the standard NSQST despite having lower infidelity. However, this can perhaps be an artifact of insufficient statistics given only ten trials. The predicted total energy and mass are also plotted in Fig. 13 (Appendix E), where NSQST and NSQST with pretraining yield significantly better predictions of total energy but not the local observable H_m .

In Fig. 4, the amplitudes and phases are displayed for typical final neural network quantum states for each protocol. We see that the NNQST protocol fails at learning the phase structure of the target state, despite accurately learning the probability distribution. This observation is consistent with NNQST's convergence to a poor local minimum in the lower plot of Fig. 3(c), with the infidelity values stuck at around 1.0.

However, standard NSQST and NSQST with pretraining are both successful at learning the phase structure of the target state, while NSQST with pretraining also learns the probability distribution better.

B. Time-evolved state for a one-dimensional Heisenberg antiferromagnet

The Heisenberg model describes magnetic systems quantum mechanically. Understanding the properties of the quantum Heisenberg model is crucial in many fields, including condensed matter physics, material science, and quantum information theory [48–50]. In this example, we perform tomography on a state that has evolved in time under the action of the one-dimensional antiferromagnetic Heisenberg (AFH) Hamiltonian. We use four Trotter time steps to approximate the time evolution.

The one-dimensional AFH model Hamiltonian is

$$H_{\text{AFH}} = \sum_{i=1}^{n-1} (\sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y + \sigma_i^z \sigma_{i+1}^z). \quad (19)$$

We choose $n = 6$ for our simulation and we take open boundary conditions. The 6-qubit initial state is set to the classical Néel-ordered state $|\Phi_0\rangle = |\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle$ and our target state occurs after evolving under the Heisenberg Hamiltonian up to time $t = 0.8$. The circuit describing the Trotterized time evolution is given in Appendix B.

Figure 5 shows the simulation results for performing tomography on the time-evolved AFH state using NNQST, NSQST, and NSQST with pretraining. We see that NSQST and NSQST with pretraining both reach a lower final state infidelity than NNQST in Fig. 5(a). Figure 5(b) displays the mean staggered magnetization (along x) at each site. We observe that NSQST with pretraining results in a tighter spread of values about the exact result, relative to NNQST or NSQST across all sites. Comparing NNQST and standard NSQST, we observe that the standard NSQST protocol has significantly worse predictions at sites 3 and 4 than NNQST, with a mean more than one standard error away from the exact value, despite having a significantly lower final state infidelity. This is likely due to the fact that NNQST is trained using nearly diagonal measurement data, providing direct access to the staggered magnetization observable of interest, whereas NSQST was trained using the infidelity loss. This result demonstrates that reaching a lower infidelity does not necessarily imply a better prediction of local observables, although we can improve the standard NSQST protocol by using more classical shadows per iteration. Finally, the typical optimization progress curves in Fig. 5(c) are consistent with the statistical results shown in Fig. 5(a). For NNQST, the infidelity does not converge stably, despite the convergence of its loss function.

The probability amplitudes and phases obtained after training on a time-evolved AFH state are shown in Fig. 6. Here, the two highest peaks in $p_\lambda(s)$ correspond to the two Néel states $|\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\downarrow\uparrow\downarrow\uparrow\rangle$ [51]. Figure 6 further confirms the advantage of NSQST with pretraining. Not only does NSQST with pretraining find more accurate phases, it also finds a better description of the probability distribution

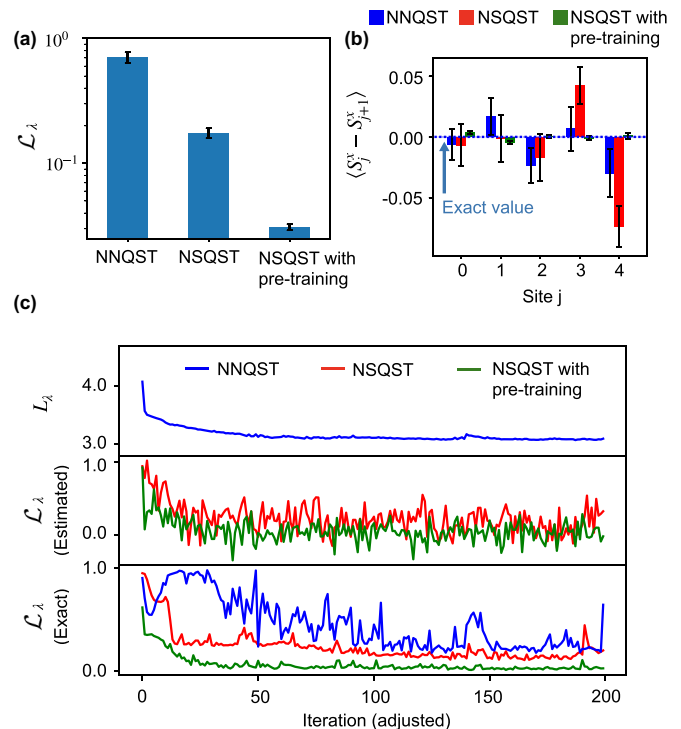


FIG. 5. Training and results for a simulation of tomography on the time-evolved state for a one-dimensional AFH model. Four Trotter steps are used for a total evolution time of $t = 0.8$. Panel (a) compares the final state infidelity, averaged over ten trials for the three protocols, following the same procedure used for one-dimensional QCD time evolution. Panel (b) compares the predicted mean staggered magnetization in the x direction (where $S_j^x = \frac{1}{2}\sigma_j^x$), following a Trotterized time evolution under the AFH model, for all the three protocols. In panel (c), we show optimization progress curves for a typical run, with the NNQST (blue) loss L_λ in the top plot, NSQST loss functions (with and without pretraining, green and red, respectively) in the middle plot, and the exact infidelity in every adjusted iteration for all three protocols in the lower plot.

since the pretraining involves many measurement samples from the all- Z basis (whereas NNQST splits the same number of measurement samples over multiple bases).

C. The phase-shifted GHZ state

In this last example, we consider the tomography of a phase-shifted GHZ state. Here, our target is a 6-qubit GHZ state with a relative phase of $\frac{\pi}{2}$. A GHZ state is a maximally entangled state that is highly relevant to quantum information science due to its nonclassical correlations [52]. Moreover, the GHZ state is the only n -qubit pure state that cannot be uniquely determined from its associated $(n - 1)$ -qubit reduced density matrices [53], indicating genuine multipartite entanglement.

As shown in Fig. 7(a), both NSQST and NSQST with pretraining result in a significantly lower average final state infidelity than NNQST. The optimization progress curves displayed in Fig. 7(b) confirm this result, as we see that the infidelity of the NNQST state is rapidly fluctuating during training, quite distinctly from the previous two examples. Given

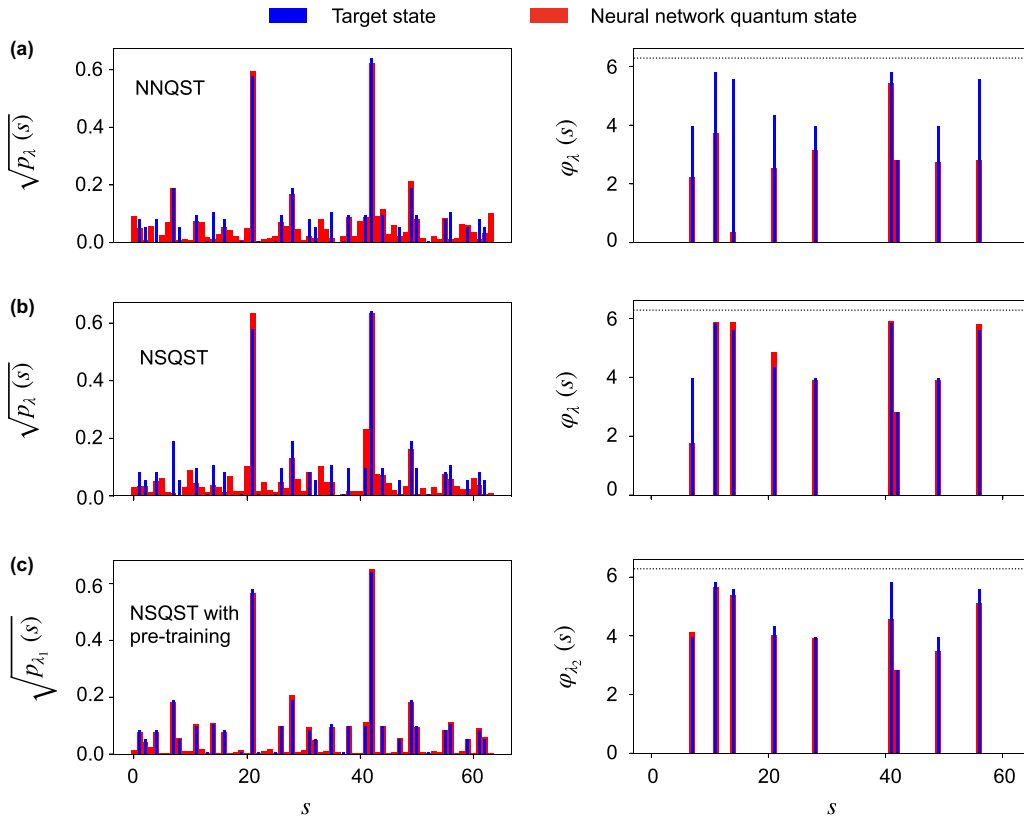


FIG. 6. Typical final neural network quantum states, trained on the time-evolved state of a one-dimensional AFH model. Panel (a) displays a typical final state in the last iteration of NNQST optimization, generated using the same procedure from Fig. 4. Panels (b) and (c) show the typical final states from NSQST and from NSQST with pretraining, respectively.

that we have employed a widely used adaptive optimizer and that we have chosen a reasonable initial learning rate (5×10^{-3}), the occurrence of such a divergence is likely due to the fact that the NNQST loss function does not incorporate tomographically complete information about the target GHZ state.

D. Comparison with direct shadow estimation

So far we have only compared the performance of NNQST, NSQST, and NSQST with pretraining, but an important question is whether any of the above methods has an advantage over direct shadow estimation. In this subsection, we compare the performance of NSQST with pretraining and direct shadow estimation for a one-dimensional QCD time-evolved state from Sec. III A. In addition, we perform a scalability study of the phase-shifted GHZ state with up to 40 qubits, comparing NSQST to direct shadow estimation. To minimize the number of Clifford shadows used for training, we fix 200 Clifford shadows as training data and do not resample in every iteration. In addition to the original pretraining protocol from Sec. II E, we adopt the improved pretraining strategy described in Appendix D. A typical optimization progress curve with the improved pretraining strategy is shown in Fig. 14 from Appendix E, where very few iterations are required for convergence. Once training is completed, we compare the prediction errors of NSQST with pretraining, Clifford shadows, and Pauli shadows.

As shown in Fig. 8, we have compared the absolute prediction error of NSQST with pretraining (with and without an improved strategy) and two types of direct shadow estimation methods. For a fair comparison, we have randomly sampled the same number (1200) measurements for each method. For shadow reconstruction, 1200 Clifford shadows or 1200 Pauli shadows were used. For NSQST with pretraining, 1000 computational-basis measurements were used for pretraining and 200 shadows were used. In Fig. 8(a), we observe that using an improved strategy, NSQST with pretraining achieves a significantly smaller prediction error than either Pauli shadows or Clifford shadows. This is expected, as the kinetic-energy Hamiltonian H_{kin} in Eq. (18) contains high-weight Pauli strings, and Pauli shadows are provably efficient at predicting only local observables [8]. However, Clifford shadows have an exponentially growing variance bound for any Pauli observables irrespective of locality [since $\text{Tr}(O^2) = 2^n$ in Eq. (10)], which explains the large prediction error for kinetic energy. In Fig. 8(b), we report the predicted error in the fidelity to the ideal time-evolved state. The Pauli shadows yield the largest prediction error in fidelity estimation due to the exponentially growing variance bound for nonlocal observables. Finally, in Fig. 8(c), we apply the four methods to the problem of predicting a single Pauli string with increasing weight, where we change the identity matrix to Pauli-X at each site as the weight increases. The nonlocal observable of interest ($X \dots$) corresponds to the Wilson loop operator in lattice gauge theory with \mathbb{Z}_2 symmetry [54]. Since predicting

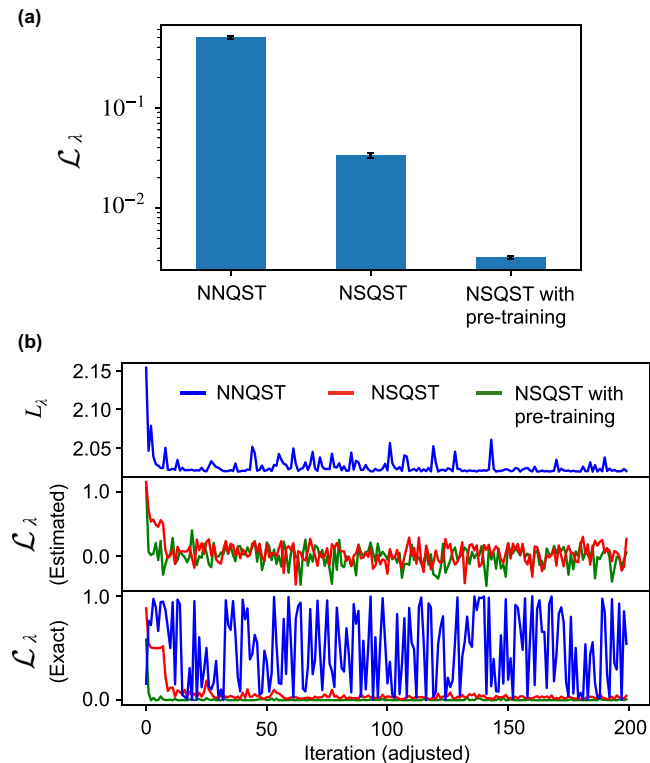


FIG. 7. Simulation of tomography on a six-qubit phase-shifted GHZ state. Panel (a) compares the final state infidelity, averaged over ten trials for each of the three protocols. Panel (b) shows typical optimization progress curves for NNQST (blue), NSQST (red), and NSQST with pretraining (green).

high-weight observables is a hard task for both shadow protocols, the prediction error from NSQST with pretraining is much lower than for either Clifford or Pauli shadows for most of the observables. We also observe a consistent increasing prediction error for the Pauli shadows as the weight increases. A natural question that arises is the scalability of NSQST's advantages over direct shadow estimation. While a trained neural network quantum state closely approximating the target state has more predictive power than classical shadows alone, there is no guarantee of successful convergence during training. For instance, learning a general multiqubit probability distribution without any prior knowledge in the pretraining step is hard [55] and would eventually require exponentially growing resources. The key to having a scalable advantage is to leverage prior knowledge of the prepared target state and to find ways to impose these known constraints in the neural network ansatz and the loss function [36,56,57].

As a proof of concept, we numerically study the sample complexity scaling of learning a phase-shifted GHZ state using NSQST with pretraining and using the improved strategy from Appendix D. With 3000 measurements in the computational basis, 200 Clifford shadows, and 5000 Monte Carlo samples for each system size, we investigate the scaling of the final infidelity. As shown in Fig. 9(a), the final infidelity does not grow as the system size increases. This is expected, as the multiqubit GHZ state is sparse in the computational basis and only a single relative phase needs to be determined.

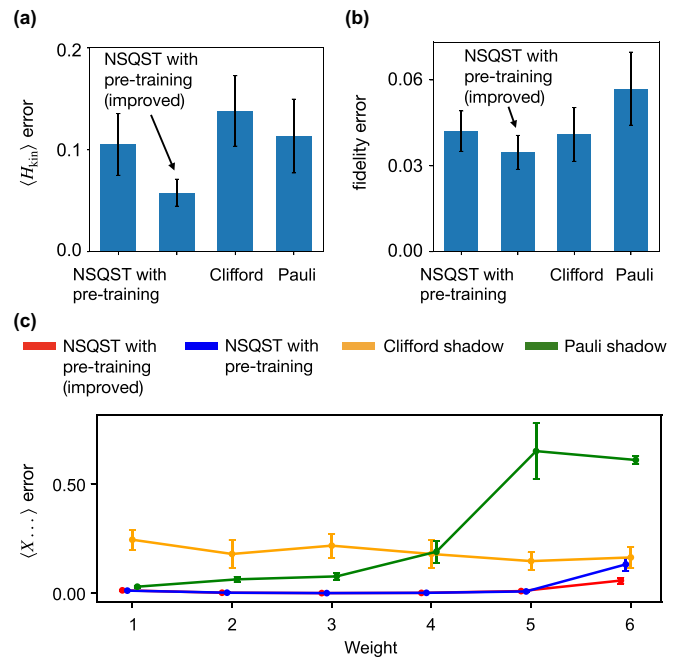


FIG. 8. Comparison of NSQST with pretraining to direct shadow estimation. In each trial of NSQST with pretraining, 200 Clifford shadows are used as training data without resampling in every iteration and 1000 measurements in the computational basis are used in pretraining. For direct shadow estimation, 1200 Clifford shadows and 1200 Pauli shadows are used. Panel (a) compares the absolute error in the predicted kinetic energy, averaged over ten trials for each of the four protocols. Panel (b) compares the absolute error in the predicted fidelity to the ideal time-evolved state, averaged over ten trials for each of the four protocols. Panel (c) compares the absolute error in the predicted expectation value of a Pauli string observable of various weight in the Pauli-X basis, averaged over ten trials for each of the four protocols. The data points are slightly shifted relative to the ticks of the x axis for a better display of error bars.

Although learning the GHZ state using a neural network ansatz is a trivial example, the sparseness property of the GHZ state is not exploited by direct shadow estimation and the collected Clifford shadows as training data would not be sample-efficient at predicting Pauli observables. In Fig. 9(b), we observe that direct shadow estimation fails to predict expectation values $\langle XX \dots XY \rangle$ accurately and yields only values of zero as the system size increases. The result presented in Fig. 9 suggests that there is hope to reconstruct sufficiently sparse states with sub-exponentially growing resources using NSQST, and potentially nonsparse states as well with enough known constraints imposed. Finally, we emphasize that, as compared to the neural network quantum state ansatz trained on IC-POVMs data in Ref. [30], our chosen state ansatz explained in Sec. II A is sample-efficient at predicting Pauli string observables of arbitrary weight and fidelity to any classically tractable states.

We make a final remark on the predictive power of randomized measurements alone versus a trained variational pure state. While one may generally expect the trained pure state to inherit the features of the training data, this may not be true in specific cases for NSQST and Clifford shadows, where

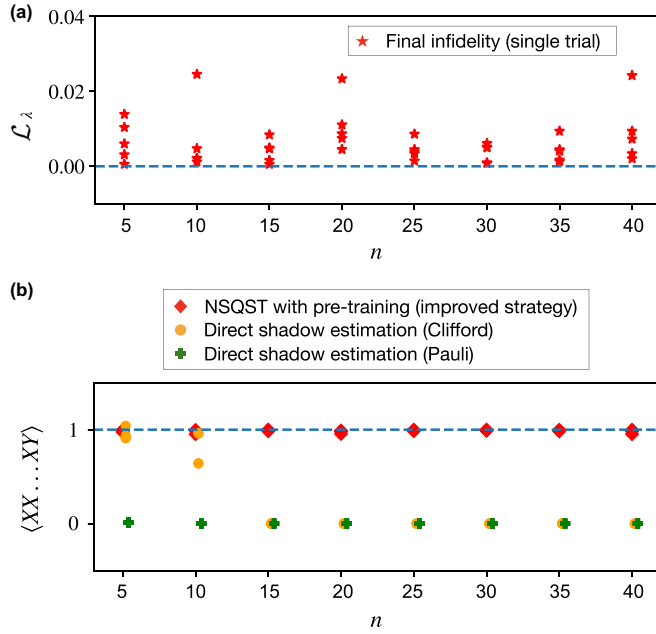


FIG. 9. Infidelity and predicted expectation value of the phase-shifted GHZ state as the system size grows. For each system size, we generate 3000 computational-basis measurements and 200 Clifford shadows. In panel (a), with independently sampled measurement data and 5000 Monte Carlo samples, we run NSQST with pretraining using the improved strategy for five trials and report the individual final infidelities. Note that the number of Monte Carlo samples used during training is much less than the number of basis states, which is not an issue if the target state is sufficiently sparse, as in the case of the multiqubit GHZ state. In panel (b), we plot the predicted expectation value of the Pauli string $XX\dots XY$, which is one of the target state’s stabilizers. In comparison, the expectation values predicted from five trials of direct shadow estimation are plotted, each with 3200 independently sampled Clifford or Pauli shadows. The data points are slightly shifted relative to the ticks of the x axis for a better display.

the trained pure state’s predictive power mainly depends on the global reconstruction error (quantum infidelity), rather than an estimator for a particular observable. Moreover, the variational training framework of NSQST is not limited to a neural network quantum state with Clifford shadows. Other variational ansatzes such as matrix product states and other randomized measurement schemes such as Hamiltonian-driven shadows should be explored with proper locality adjustments, for practical advantages such as hardware-aware measurements and scalable classical postprocessing [58–61].

IV. NUMERICAL SIMULATIONS WITH NOISE

We now numerically investigate the noise robustness of our NSQST protocol, focusing on the same phase-shifted GHZ state as in Sec. III C. We consider two different sources of noise affecting the Clifford circuit used to evaluate our infidelity-based loss function using classical shadows (see Sec. II D). The first noise model, (a particular case of the model already introduced in Sec. II C), describes either measurement (readout) errors or gate-independent time-stationary Markovian (GTM) noise. The second noise model describes

imperfect two-qubit entangling gates in the Clifford circuit. In the following, we introduce both noise models and discuss their effects on the fidelity of the reconstructed state.

Our first model is an amplitude damping channel applied before measurements. The amplitude damping channel is suitable for investigating the effect of measurement errors in the computational basis. The n -qubit amplitude damping noise channel $AD_{n,p}$ with channel parameter p is defined as

$$AD_{n,p} = AD_{1,p}^{\otimes n}, \quad (20)$$

where

$$AD_{1,p} : \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{10} & \rho_{11} \end{pmatrix} \mapsto \begin{pmatrix} \rho_{00} + (1-p)\rho_{11} & \sqrt{p}\rho_{01} \\ \sqrt{p}\rho_{10} & p\rho_{11} \end{pmatrix}. \quad (21)$$

Apart from modeling measurement noise, this noise channel also serves as a suitable model for studying gate-independent, time-stationary, and Markovian (GTM) noise [24]. In this case, each gate that appears in the Clifford circuit U_i is subject to the same noise map. The resulting noisy random Clifford circuit \tilde{U}_i can be decomposed into $\mathcal{E}U_i$ with \mathcal{E} being a noise channel applied after the ideal Clifford unitary.

To demonstrate the noise robustness of our NSQST protocol, we first perform tomography on a phase-shifted GHZ state (having a relative phase of $\frac{\pi}{2}$). Despite the presence of the amplitude damping noise $\mathcal{E} = AD_{n,p}$, we simulate NSQST using the noise-free gradient expression $\nabla_\lambda \mathcal{L}(\mathbb{I})$ in Eq. (15). As discussed in Sec. II D, the noise-free gradient expression in NSQST will still yield an estimate that is directed along the true gradient, being modified only with an overall prefactor. In contrast, the noise-free loss function $\mathcal{L}(\mathbb{I})$ and the true loss $\mathcal{L}(\mathcal{E})$ are related nontrivially in the presence of noise:

$$\mathcal{L}(\mathcal{E}) = \frac{1}{(2^n + 1)f(\mathcal{E})} \mathcal{L}(\mathbb{I}) + \frac{(4^n - 1)f(\mathcal{E}) - 2^n + 1}{2^n(2^n + 1)f(\mathcal{E})}. \quad (22)$$

This means that our estimated loss function no longer converges to zero, while the infidelity between the neural network quantum state and the target state approaches zero during training.

Figure 10(a) shows the simulation results for the effects of an amplitude damping noise channel applied before measurement. First, we observe that the average exact infidelity of the last 100 iterations remains small despite the growing noise channel strength. The increasing loss function value (red curves) is evidence of the growing variance in our gradient estimations, and will eventually lead to failure of the optimizer to converge to a state close to the target state. Intuitively, since the classical shadows method only uses the measured bit-string, but not the phase for postprocessing, only the diagonal bit-flip errors in Eq. (21) contribute to the noise model and these are twirled into depolarizing noise by random Clifford circuits. Finally, the agreement between the exact infidelity (blue curve) and the transformed loss function [the right-hand side of Eq. (22) represented by the orange curve] validates our theoretical account. Here, we have used Eq. (8) to find the

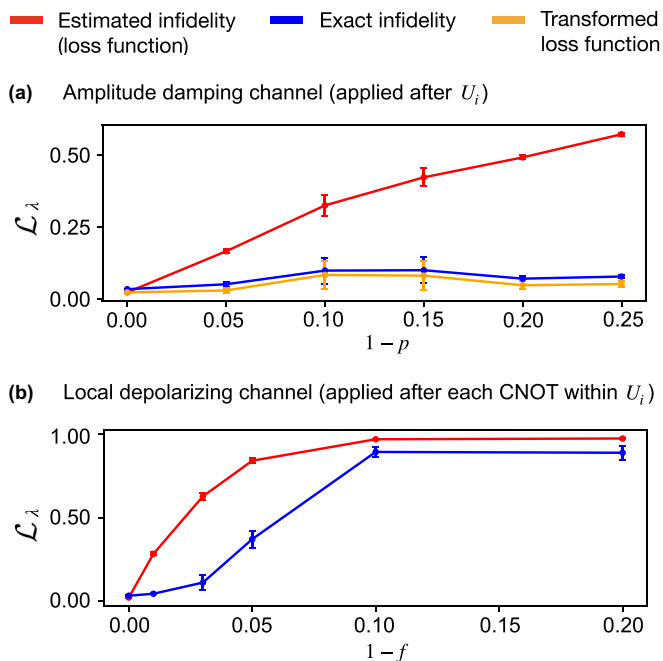


FIG. 10. Simulation of tomography for a phase-shifted GHZ state in the presence of noise. Panel (a) displays the average loss function (red) defined in Eq. (23) and the exact infidelity (blue) for the amplitude damping channel $\text{AD}_{n,p}$ (the loss function is averaged over the last 100 iterations for each trial and then the average is taken over ten trials). The strength of the noise increases with increasing $1 - p$. The noiseless infidelity loss function $\mathcal{L}(\mathbb{I})$ is then transformed into an estimated infidelity for the noisy case $\mathcal{L}(\mathcal{E})$ using Eq. (22) for the amplitude damping channel $\mathcal{E} = \text{AD}_{n,p}$ to obtain the transformed cost function. The error bars represent the standard error in the mean over ten trials. In panel (b), we show the average loss function and exact infidelity for the local depolarizing noise model with a two-qubit depolarizing channel applied after every CNOT gate in the appended random Clifford circuit U_i . The channel parameter $1 - f$ characterizes the growing strength of the noise. We do not plot the transformed loss function in this case because the CNOT-dependent local depolarizing noise model does not have an analytic noisy shadow expression.

depolarizing noise channel strength

$$f(\text{AD}_{n,p}) = \frac{(1+p)^n - 1}{2^{2n} - 1}. \quad (23)$$

Note that $f(\mathcal{E})$ may be hard to estimate in practice. However, since this parameter does not affect the direction of the estimated gradient, we expect training to converge to the same optimal parameters λ with or without noise. It is therefore not necessary to compensate for noise by computing the linear transformation in Eq. (22) as long as we can verify the successful convergence of training.

We proceed now with a discussion of the second noise model, which assumes that entangling gates are the dominant source of error. For numerical simulations, we decompose each random Clifford unitary U_i into CNOT, Hadamard, and phase gate operations. Subsequently, a local two-qubit noise map is applied after each CNOT gate in U_i . We consider the

depolarizing noise channel

$$\mathcal{D}_{n,f}(\rho) = f\rho + (1-f)\frac{\mathbb{I}}{2^n}, \quad (24)$$

with $n = 2$ and a fixed noise strength $1 - f$. This noise model is not GTM and no longer has an analytic noisy shadow expression. However, we still expect NSQST to be fairly noise robust based on the numerically demonstrated robustness of classical shadows against many non-GTM errors, such as pulse miscalibration noise [24].

As shown in Fig. 10(b), NSQST exhibits some measure of noise robustness even in the presence of a more realistic non-GTM noise model. This is reflected in the positive curvature of the blue curve for decreasing noise, $1 - f \rightarrow 0$, leading to a weak-noise limit where the exact infidelity (blue curve) is small relative to the estimated infidelity (red curve). Our randomly sampled six-qubit unitary U_i has an average of 21 CNOT gates (see Appendix C), leading to a substantial accumulation of errors. Thus, the noise parameter $1 - p$ controlling one-time measurement errors is not comparable to the parameter $1 - f$ controlling the noise on the individual CNOT gates. A transformed loss function curve is not presented in Fig. 10(b) because our local (two-qubit) depolarizing noise model does not yield an analytic $f(\mathcal{E})$ expression. The robustness of the classical shadows formalism against many other non-GTM noise models (with an extra calibration step) has been well studied in Ref. [24], while our NSQST protocol holds a similar noise robustness without any extra calibration steps.

V. CONCLUSIONS AND OUTLOOK

In this work, we have proposed a new QST protocol, neural-shadow quantum state tomography (NSQST). We have demonstrated its clear advantages over state-of-the-art implementations of neural network quantum state tomography (NNQST) in three relevant settings, as well as advantages over direct shadow estimation. We have further shown that NSQST is noise robust. Our study of the benefits of NSQST suggests that the choice of infidelity as a loss function has great potential to broaden the applicability of neural network-based tomography methods to a wider range of quantum states.

In Appendix D we describe technical developments (reuse of classical shadows and alternative Monte Carlo methods) that can be pursued to further enhance the performance of NSQST. Another direction for future work would be to tailor NSQST more closely to emerging quantum hardware platforms. This can be done by exploring NSQST with alternative shadow protocols. In particular, it would be interesting to investigate hardware-aware classical shadows that use the native interactions of the quantum device [58,59,62]. In addition, future work should extend NSQST to mixed-state protocols [31].

Relative to classical shadow protocols, which only allow for efficient fidelity and local observable predictions, but no efficient state reconstruction, NSQST achieves the goal of reconstructing a physical state that approximates a target quantum state via a variational ansatz. The variational ansatz in NSQST comes with the convenience of a quantum state and can be used to predict many global observables of interest

beyond the reach of direct shadow estimation [63]. Moreover, NSQST inherits the advantages of NNQST. For example, we can incorporate symmetry constraints of the target state, reducing the computational resources needed [36,56,64]. Once trained, relative to the large number of classical shadows that must be collected, the variational ansatz in NSQST may yield a more efficient classical representation of the state. Finally, as demonstrated in NSQST with pretraining, the trained variational ansatz approximating the target state can be fed into a second round of optimization, performed with respect to a new loss function. This possibility provides great flexibility in addressing a variety of tasks, including, for example, error mitigation in classical postprocessing [18].

NSQST is an efficient state reconstruction method. It will be useful as a benchmarking tool, an important element for testing the performance of near-term quantum devices as they scale up. In particular, NSQST can be used to construct a “digital twin” [65] of the prepared target state, where a neural network quantum state can be used for experimentally relevant simulation [66,67], cross-platform verification [68], error mitigation [18] and other uses. Having access to a digital twin of the target quantum state will become increasingly relevant for accelerating the development of quantum technologies. Further down the road, we also foresee great potential for NSQST as a stepping stone for interfacing classical probabilistic graphical models and quantum circuits, where data stored in quantum circuits can be transferred to classical memory and vice versa, leading to new hybrid computing approaches.

ACKNOWLEDGMENTS

We thank Abhijit Chakraborty, Luca Dellantonio, David Gosset, Arsalan Motamedi, Jinglei Zhang, Jan Friedrich Haase, Yasar Atas, and Randy Lewis for useful discussions. This research was undertaken thanks in part to funding from the Canada First Research Excellence Fund (CFREF), the Natural Sciences and Engineering Research Council (NSERC), the New Frontiers in Research Fund (NFRF), and the Fonds de Recherche-Nature et Technologies (FRQNT). C.M. acknowledges the Alfred P. Sloan foundation for a Sloan Research Fellowship and an Ontario Early Researcher Award. P.R. further acknowledges the support of NSERC Discovery Grant No. RGPIN-2022-03339, Mike and Ophelia Lazaridis, Innovation, Science and Economic Development Canada (ISED), IQBit, Perimeter Institute for Theoretical Physics (PI), and subsequently, PI’s partial support by the Government of Canada through ISED, and by the Province of Ontario through the Ministry of Colleges and Universities.

The numerical implementation of NNQST from Ref. [18] can be found at Ref. [71]. The numerical implementation of NSQST can be found at Ref. [72].

APPENDIX A: NEURAL NETWORK QUANTUM STATE ARCHITECTURE

For NNQST and NSQST, we use the transformer-based neural network quantum state ansatz directly adopted from Ref. [18]. A central component in the ansatz is the transformer layer, which has a self-attention block followed by a linear

layer. With a bit-string $s = (s_1, \dots, s_n) \in \{0, 1\}^n$ as input, s is extended to $\tilde{s} = (0, s)$ by prefixing a zero bit. Then, each bit \tilde{s}_j is encoded into a D -dimensional representation space using a learned embedding governed by f_{jd} , yielding the encoded bit $e_{jd}^{(0)}$ with $j \in \{0, \dots, n\}$ and $d \in \{1, \dots, D\}$. The encoded input is then processed using K transformer layers.

We outline the parameters involved in a single transformer layer indexed by k in the ansatz, and refer the reader to Ref. [18] for more details:

(1) The query, key, value matrices $Q_{\text{hid}}^{(k)}$, $K_{\text{hid}}^{(k)}$, and $V_{\text{hid}}^{(k)}$, where $h \in \{1, \dots, H\}$, $d \in \{1, \dots, D\}$, $i \in \{1, \dots, D/H\}$. We require D/H to be an integer.

(2) A matrix to process the output of the self-attention heads, $O_{de}^{(k)}$, with $d, e \in \{1, \dots, D\}$.

(3) A weight matrix and a bias vector of the linear layer, $W_{de}^{(k)}$ and $b_d^{(k)}$, with $d, e \in \{1, \dots, D\}$.

Once we have passed the final transformer layer, scalar-valued logits ℓ_j are obtained by using an extra linear layer. The conditional probabilities directly used in sampling are then given by

$$p_\lambda(s_j = 1 | s_1, \dots, s_{j-1}) = \sigma(\ell_{j-1}), \quad (\text{A1})$$

where $\sigma(\ell_{j-1}) = \frac{1}{1+e^{-\ell_{j-1}}}$ is the logistic sigmoid function. Since the outcome at index j is conditional on the preceding indices $\tilde{j} \leq j$, we can efficiently draw unbiased samples from the probability distribution $p_\lambda(s)$ by proceeding one bit at a time. The phase output $\varphi_\lambda(s)$ is obtained by first concatenating the output of the final transformer layer to a vector of length n , then projecting the vector to a single scalar value using a linear layer [separate from the linear layer used in obtaining $p_\lambda(s)$].

For NSQST with pretraining, our $p_{\lambda_1}(s)$ is parameterized in the same way as in standard NSQST, except that we remove the phase output layer. The phase outcome $\varphi_{\lambda_2}(s)$ is encoded in a separate transformer-based neural network ansatz, where we remove the other linear layer (the one producing scalar-valued logits representing the probability amplitudes). Thus, the encoded quantum state in NSQST with pretraining has its probability distribution and phase output separately parameterized by model parameters λ_1 and λ_2 , respectively.

APPENDIX B: TROTTERIZED TIME EVOLUTION CIRCUITS

In this section, we elaborate on the details of Trotterized time evolution circuits for the two noiseless numerical experiments (Secs. III A and III B).

1. One-dimensional QCD time evolution

For the one-dimensional QCD Hamiltonian in Eq. (17), we choose $\tilde{m} = 1.2$ and $x = 0.8$. Two Trotter steps are used for a total evolution time of $t = 1.8$. With the initial state chosen as $|\Phi_0\rangle = |\downarrow\downarrow\downarrow\rangle|\uparrow\uparrow\uparrow\rangle$, Fig. 11 shows the circuit (from Ref. [47]) used to generate the time-evolved target state $|\Phi_{SU(3)}\rangle$.

Global Phase: 42.4115008234622 - 2.1125* π

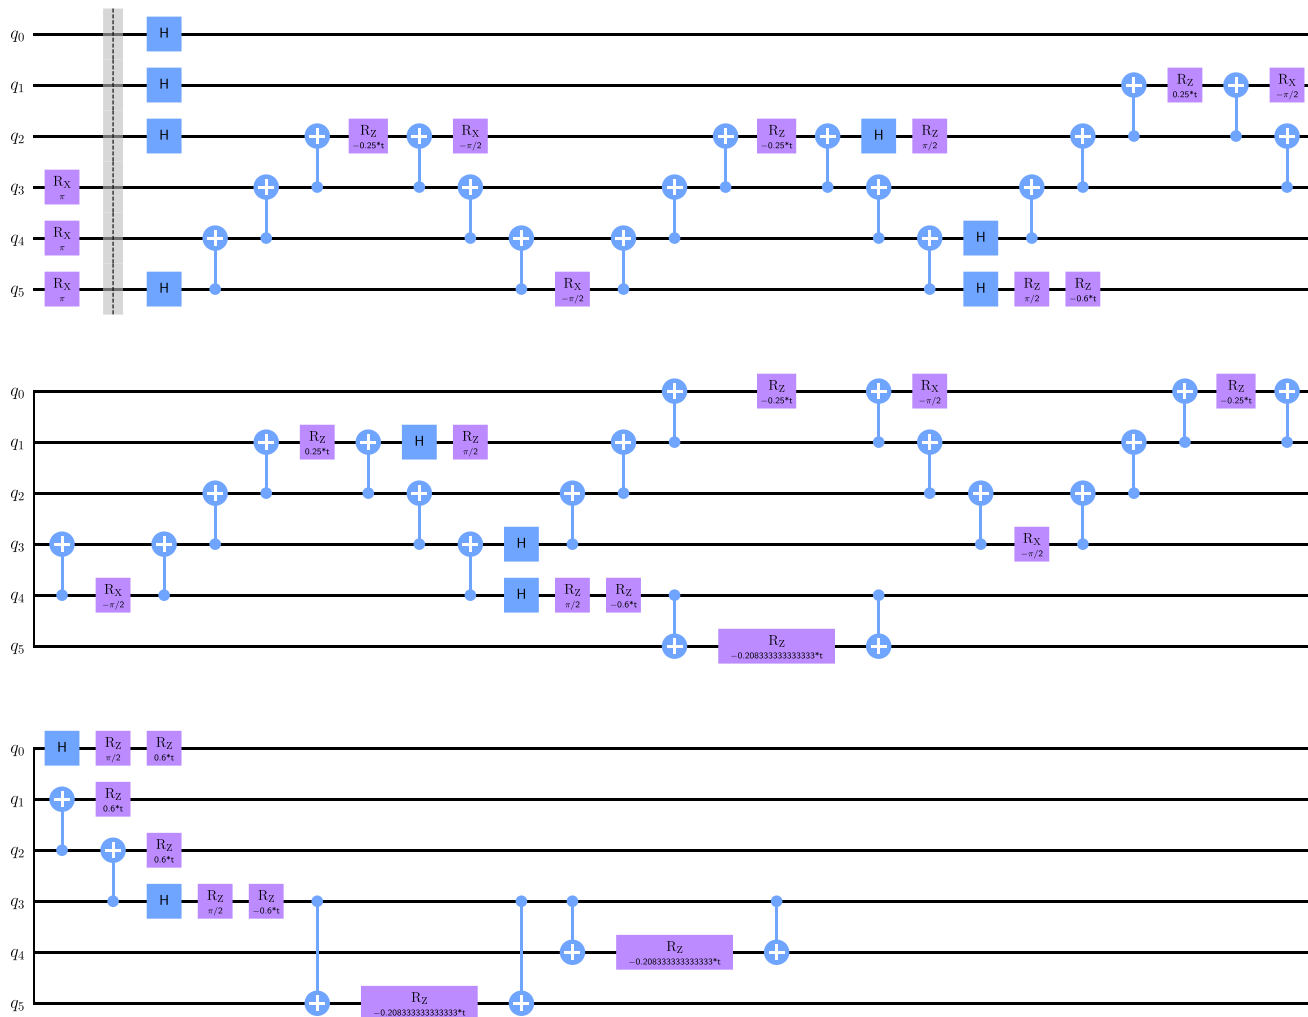


FIG. 11. Simulation circuit for the one-dimensional QCD model. The initial state preparation circuit (before the barrier) and a single Trotter step (after the barrier) are drawn using Qiskit [69]. In our numerical experiments an evolution for time $t = 1.8$ is decomposed into two Trotter steps.

2. One-dimensional AFH time evolution

For the one-dimensional AFH model Hamiltonian in Eq. (19), we choose the initial state $|\Phi_0\rangle = |\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle$ and set the total evolution time as $t = 0.8$, decomposed into four Trotter steps to generate the target state. Figure 12 shows the initial state preparation and the circuit of one Trotter step.

APPENDIX C: SIMULATION HYPERPARAMETERS

For all $n = 6$ transformer-based neural network quantum states, we have used two transformer layers ($K = 2$), four attention heads ($H = 4$), and eight internal dimensions ($D = 8$). The neural network quantum states for NNQST and the standard NSQST have 858 model parameters. NSQST with pretraining uses two neural networks for the same quantum state, where the probability distribution ansatz $p_{\lambda_1}(s)$ has 801 model parameters and the phase output ansatz $\varphi_{\lambda_2}(s)$ has 849 model parameters, a total of 1650 model parameters. It is likely that we have used more model parameters

than necessary for a 6-qubit pure state, as the focus of our demonstrations is studying the new loss function. For $n > 6$ transformer-based neural network quantum states used for NSQST with pretraining, we have used three transformer layers ($K = 3$), four attention heads ($H = 4$), and eight internal dimensions ($D = 8$), corresponding to a number of model parameters from 2466 to 3186 for system size from $n = 10$ to $n = 40$.

All NNQST protocol trials use 21 ($4n - 3$ with $n = 6$) nearly diagonal and nearest-neighbour Pauli bases, where 512 measurement samples are used for each basis. The batch sizes are 128 and all NNQST simulations are run for 200 epochs, where one epoch refers to one sweep over the entire training data set. In all simulations the Adam optimizer is used. For the one-dimensional QCD and AFH time evolution examples, the learning rate is set to 10^{-3} . For the GHZ state tomography, a learning rate of 5×10^{-3} is used.

All $n = 6$ NSQST protocol trials (including the ones with pretraining) use 100 classical shadow samples ($N = 100$) and 5000 neural network quantum state samples (both resampled

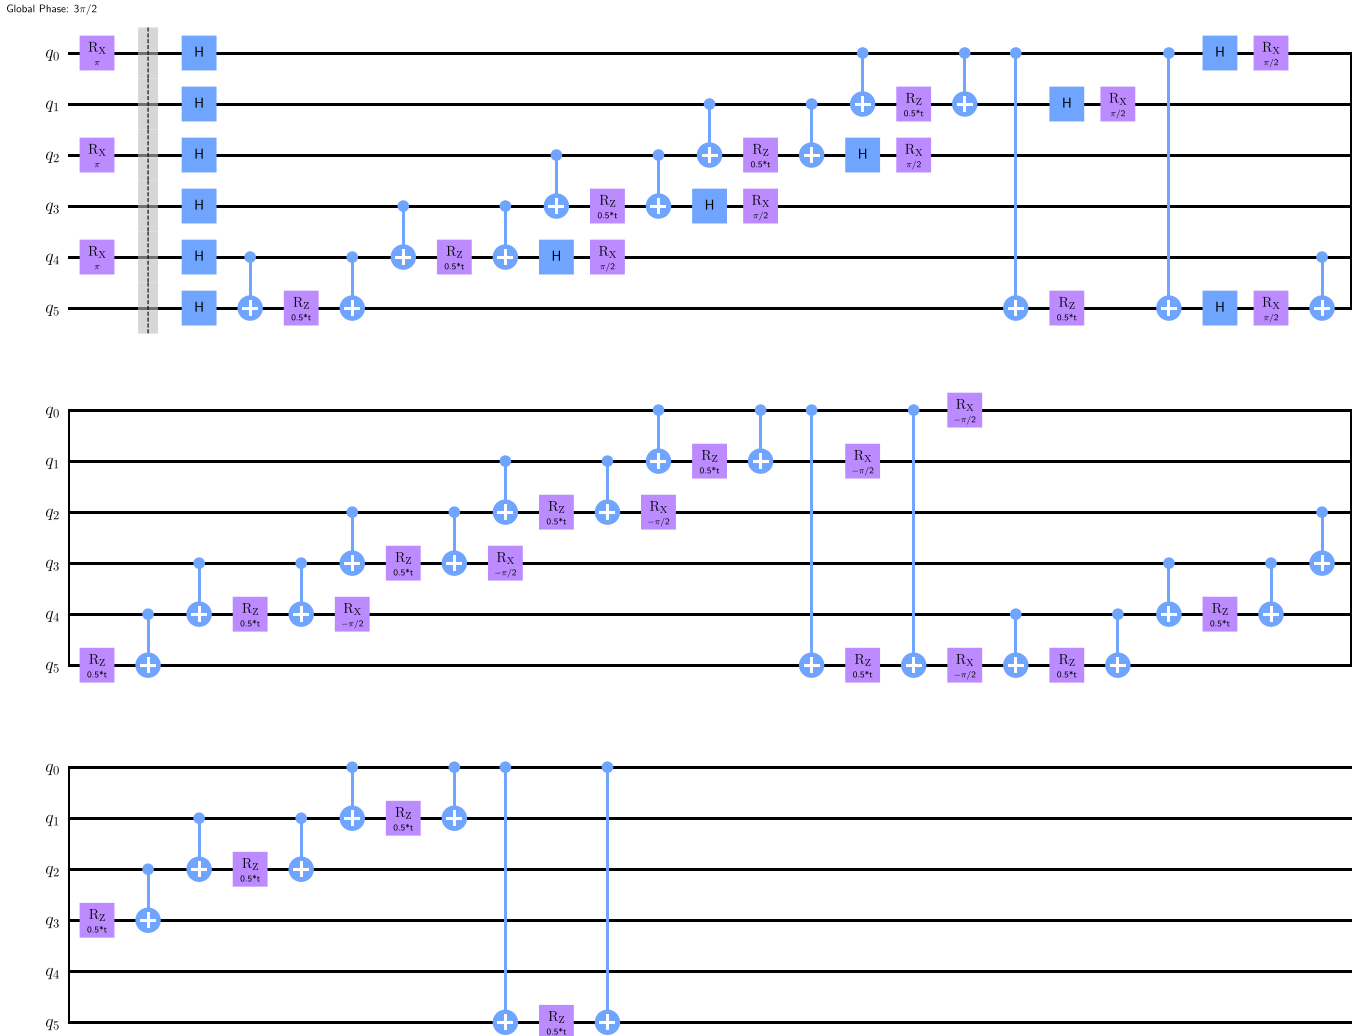


FIG. 12. Simulation circuit for the one-dimensional AFH model. The initial state preparation circuit (before the barrier) and a single Trotter step (after the barrier) are drawn using Qiskit [69]. In our numerical experiments an evolution for time $t = 0.8$ is decomposed into four Trotter steps.

in every iteration) per iteration. All simulations are run for 2000 iterations with the learning rate 10^{-2} for the Adam optimizer. Note that we have not leveraged the reusability of the classical shadows, which can bring a significant reduction in the number of measurements. The randomly sampled 6-qubit Clifford circuits U_i^\dagger have 20.8 ± 0.4 CNOT gates, assuming all-to-all connectivity.

The $n = 6$ modified NSQST experiment with pretraining uses 10752 (512×21) computational basis measurement samples in the pretraining stage using NNQST’s optimization framework, with identical hyperparameters to that of the NNQST protocol. NSQST with pretraining (with or without improved strategy) in Sec. III D uses 1000 computational basis measurements for $n = 6$ and 3000 computational basis measurements for $n > 6$ in the pretraining stage, and a fixed set of 200 Clifford classical shadows are used for training the relative phases.

For predicting observable values given a trained neural network quantum state, exact calculations are done for $n = 6$ and 20 000 Monte Carlo samples are used for $n > 6$.

APPENDIX D: METHOD IMPROVEMENT

For simplicity, we have naively resampled the 100 classical shadows in every iteration, which is not necessary as the previously sampled classical shadows could be reused. We expect many possible improvements could be made to reduce the number of sampled Clifford circuits and the number of measurements, such as measuring more bit-strings for each U_i . Hybrid training protocols beyond NSQST with pretraining should also be explored.

We recognize that an exponential classical computational cost may be induced for convergence due to the growing variance of the gradient in Eq. (15). To intuitively see this, we note that estimating the overlap in Eq. (14) to within $\tilde{\epsilon}$ additive error with failure probability at most δ requires $L \sim O(\frac{1}{\tilde{\epsilon}^2} \log \frac{1}{\delta})$ classical samples drawn from $|\psi_\lambda\rangle$ [70]. However, the prefactor $\frac{1}{Nf(\mathcal{E})}$ in Eq. (15) will make the error ϵ in estimating the gradient exponentially large. Since we want ϵ to be bounded, we demand $\tilde{\epsilon} \sim Nf(\mathcal{E})$ to be exponentially small and this directly translates to $L \sim O(1/N^2 f(\mathcal{E})^2)$ as discussed

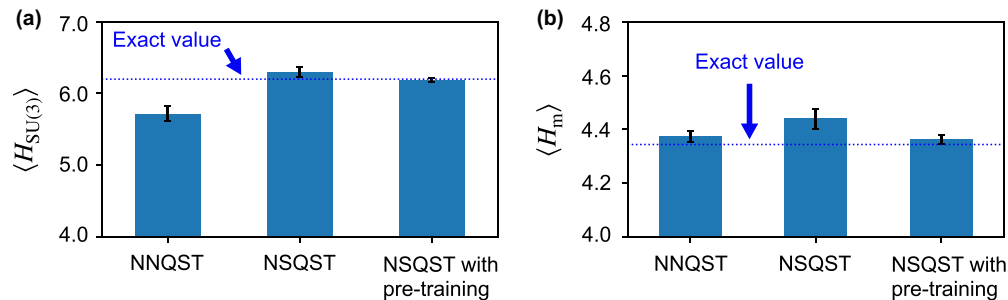


FIG. 13. Additional plots for the quantum state following a one-dimensional QCD time evolution. Panel (a) displays the expectation values of the total energy, evaluated for the neural network quantum state found in the last iteration of each trial, and averaged over ten trials for each protocol. Panel (b) displays the expectation values of the mass Hamiltonian.

in Sec. II D. Imposing harder constraints on the quantum state ansatz may remove or alleviate this issue, and numerical experiments of larger system sizes should be done to explore NSQST's limitations in the future.

An additional complication can arise when estimating the inner product shown in Eq. (14) from a finite number of Monte Carlo samples. In practice, the overlap is estimated in terms of a subset \mathcal{S} of distinct bit-strings s using

$$\langle \phi_i | \psi_\lambda \rangle = \left\langle \frac{\phi_i^*(s)}{\psi_\lambda^*(s)} \right\rangle_{\psi_\lambda} \simeq \sum_{s \in \mathcal{S}} P(s) \frac{\phi_i^*(s)}{\psi_\lambda^*(s)}. \quad (\text{D1})$$

Here, $P(s) = f_s/N_s \simeq p_\lambda(s) = |\psi_\lambda(s)|^2$ is determined from the frequency f_s of the bit-string s found from N_s samples drawn according to the probability distribution $p_\lambda(s)$. For a transformer-based neural network architecture, the samples can be generated efficiently bit-by-bit using the procedure described in Appendix A. Up to a constant factor, the right-hand side of Eq. (D1) can be interpreted as the exact overlap between a stabilizer state $|\phi_i\rangle$ and a fictitious state $|\Psi_\lambda\rangle$ with wave function:

$$\Psi_\lambda(s) := \frac{P(s)}{\psi_\lambda^*(s)A_S}; \quad A_S^2 = \sum_{s \in \mathcal{S}} \frac{P^2(s)}{|\psi_\lambda(s)|^2}. \quad (\text{D2})$$

The normalization constant approaches $A_S = 1$ when $P(s) = p_\lambda(s) = |\psi_\lambda(s)|^2$ (e.g., when the sample set \mathcal{S} includes all s). However, a problem arises when we sample only over a subset of possible bit-strings s . In this case, it may be that $|\psi_\lambda^*(s)| \ll \sqrt{P(s)}$ for some s , leading to $A_S \gg 1$. Estimating the infi-




FIG. 14. Typical optimization progress curve from NSQST with pretraining and fixed Clifford shadows. Unlike the other plots, the iteration number on the x axis is not adjusted and corresponds to every gradient update during optimization.

delity from classical shadows to obtain the NSQST loss function [Eq. (11)] through Monte Carlo samples as in Eq. (D1) requires the estimated overlaps $\langle \phi_i | \psi_\lambda \rangle \simeq A_S \langle \phi_i | \Psi_\lambda \rangle$. When an incomplete sample set is taken, the factor A_S can become very large, leading to an unphysical blow-up, potentially leading to estimated overlaps $\gg 1$. In this limit, the Monte Carlo estimate is meaningless. A simple solution to this problem could be to truncate the set $\mathcal{S} \rightarrow \mathcal{S}'$, allowing only for bit-strings s for which $|\psi_\lambda(s)|/\sqrt{P(s)}$ exceeds some threshold value, then we replace the normalization constant $A_S \rightarrow A_{\mathcal{S}'}$. For a given task, it may be difficult to establish truncation thresholds that maintain convergence to an accurate state. In the rest of this Appendix, we give an alternative procedure that does not show the ill-conditioned “blow-up” from a finite Monte Carlo sample size, while avoiding predetermined truncation thresholds.

To avoid the pitfalls of representing a Monte Carlo average as in Eq. (D1), we consider a hybrid NSQST protocol. In this hybrid protocol, the classical shadows are only used to learn the phases $\varphi_{\lambda_2}(s)$. The probability amplitudes $p_{\lambda_1}(s)$ are learned using NNQST from measurements performed in the computational basis (similar to NSQST with pretraining): $p_{\lambda_1}(s) \simeq p_\Phi(s, B)$ with $B = (Z_1, Z_2, \dots, Z_n)$. The difference between this new hybrid protocol and NSQST with pretraining is in learning the phases. To train the phase model, we calculate the gradient of the loss function with an alternative approximation for the inner product:

$$\langle \phi_i | \psi_\lambda \rangle \approx \langle \phi_i | \tilde{\psi}_\lambda \rangle := \sum_{s \in \mathcal{S}} \phi_i^*(s) e^{i\varphi_{\lambda_2}(s)} \sqrt{P(s)}, \quad (\text{D3})$$

where we have introduced an alternative normalized state:

$$\tilde{\psi}_\lambda(s) := e^{i\varphi_{\lambda_2}(s)} \sqrt{P(s)}. \quad (\text{D4})$$

The estimated loss function in this new hybrid protocol is the shadow-estimated infidelity between the target state $|\Phi\rangle$ and the state $|\tilde{\psi}_\lambda\rangle$. The gradient of this infidelity can be calculated to optimize the phases from variations in the parameters λ_2 :

$$\nabla_{\lambda_2} \mathcal{L}(\mathcal{E}) \approx \frac{-2}{Nf(\mathcal{E})} \sum_{i=1}^N \Re[\partial_{\lambda_2} \langle \phi_i | \tilde{\psi}_\lambda \rangle \cdot \langle \tilde{\psi}_\lambda | \phi_i \rangle]. \quad (\text{D5})$$

We can efficiently evaluate the right-hand side of Eq. (D5) exactly for a sub-exponential number of distinct bit-strings $s \in \mathcal{S}$. The optimization procedure is then limited only by the expressivity and accuracy of the sparse approximation $\psi_\lambda(s)$ for the neural network quantum state $\psi_\lambda(s)$, arising from a finite number of samples. This new hybrid NSQST protocol does not suffer from the “blow-up” described above and it may converge with a sub-exponential number of samples, especially when $|\psi_\lambda\rangle$ is sufficiently sparse in the computational basis. This alternative strategy was unnecessary in most of our numerical experiments given the very small system size $n = 6$ and the very large number of Monte Carlo samples $N_s = 5000$.

APPENDIX E: ADDITIONAL PLOTS

In this section, we provide additional plots relevant to the numerical simulation results presented in Sec. III.

In Fig. 13, the predicted total energy and mass from the three protocols are plotted, where we see that NNQST fails to yield a better prediction of total energy than NSQST in Fig. 13(a). However, as shown in Fig. 13(b), NNQST predicts mass H_m more accurately than NSQST, which is a local observable from Eq. (18).

In Fig. 14, a typical optimization progress curve is plotted for the numerical results presented in Sec. III D. The iteration number is not adjusted and pretraining is repeated for every trial.

-
- [1] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, and E. Kashefi, Quantum certification and benchmarking, *Nat. Rev. Phys.* **2**, 382 (2020).
- [2] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, *Phys. Rev. A* **77**, 012307 (2008).
- [3] J. Carrasco, A. Elben, C. Kokail, B. Kraus, and P. Zoller, Theoretical and experimental perspectives of quantum verification, *PRX Quantum* **2**, 010102 (2021).
- [4] J. Choi, A. L. Shaw, I. S. Madjarov, X. Xie, R. Finkelstein, J. P. Covey, J. S. Cotler, D. K. Mark, H.-Y. Huang, A. Kale *et al.*, Preparing random states and benchmarking with many-body quantum chaos, *Nature (London)* **613**, 468 (2023).
- [5] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, *Nat. Rev. Phys.* **5**, 9 (2023).
- [6] R. Stricker, M. Meth, L. Postler, C. Edmunds, C. Ferrie, R. Blatt, P. Schindler, T. Monz, R. Kueng, and M. Ringbauer, Experimental single-setting quantum state tomography, *PRX Quantum* **3**, 040310 (2022).
- [7] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [8] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [9] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [10] Y. Huang, J. E. Moore *et al.*, Neural network representation of tensor network and chiral states, *Phys. Rev. Lett.* **127**, 170601 (2021).
- [11] O. Sharir, A. Shashua, and G. Carleo, Neural tensor contractions and the expressive power of deep neural quantum states, *Phys. Rev. B* **106**, 205136 (2022).
- [12] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, Equivalence of restricted Boltzmann machines and tensor network states, *Phys. Rev. B* **97**, 085104 (2018).
- [13] D.-L. Deng, X. Li, and S. Das Sarma, Quantum entanglement in neural network states, *Phys. Rev. X* **7**, 021021 (2017).
- [14] G. E. Hinton, A practical guide to training restricted Boltzmann machines, in *Neural Networks: Tricks of the Trade* (Springer, Berlin, 2012), pp. 599–619.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, Attention is all you need, *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2017), Vol. 30..
- [16] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, in *International Conference on Learning Representations (ICLR)* (San Diego, CA, USA, 2015).
- [17] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [18] E. R. Bennewitz, F. Hopfmueller, B. Kulchitsky, J. Carrasquilla, and P. Ronagh, Neural error mitigation of near-term quantum simulations, *Nat. Mach. Intell.* **4**, 618 (2022).
- [19] G. Torlai, B. Timar, E. P. L. van Nieuwenburg, H. Levine, A. Omran, A. Keesling, H. Bernien, M. Greiner, V. Vuletić, M. D. Lukin, R. G. Lukin, M. Lukin *et al.*, Integrating neural networks with a quantum simulator for state reconstruction, *Phys. Rev. Lett.* **123**, 230504 (2019).
- [20] J. Chen, H. Dawkins, Z. Ji, N. Johnston, D. Kribs, F. Shultz, and B. Zeng, Uniqueness of quantum states compatible with given measurement results, *Phys. Rev. A* **88**, 012109 (2013).
- [21] S. Luo and Q. Zhang, Informational distance on quantum-state space, *Phys. Rev. A* **69**, 032106 (2004).
- [22] S. Aaronson, Shadow tomography of quantum states, in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing* (ACM, 2018), pp. 325–338.
- [23] D. E. Koh and S. Grewal, Classical shadows with noise, *Quantum* **6**, 776 (2022).
- [24] S. Chen, W. Yu, P. Zeng, and S. T. Flammia, Robust shadow estimation, *PRX Quantum* **2**, 030348 (2021).
- [25] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, Transformers: State-of-the-art natural language processing, in *Proceedings of the Conference on Empirical Methods in Natural Language Processing: System Demonstrations* (Association for Computational Linguistics, 2020), pp. 38–45.
- [26] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah, Transformers in vision: A survey, *ACM Comput. Surv. (CSUR)* **54**, 1 (2022).
- [27] G. Torlai and R. G. Melko, Latent space purification via neural density operators, *Phys. Rev. Lett.* **120**, 240503 (2018).
- [28] F. Vicentini, D. Hofmann, A. Szabó, D. Wu, C. Roth, C. Giuliani, G. Pescia, J. Nys, V. Vargas-Calderón,

- N. Astrakhantsev *et al.*, Netket 3: Machine learning toolbox for many-body quantum systems, *SciPost Phys. Codebases*, **7** (2022).
- [29] V. Havlicek, Amplitude ratios and neural network quantum states, *Quantum* **7**, 938 (2023).
- [30] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, *Nat. Mach. Intell.* **1**, 155 (2019).
- [31] P. Cha, P. Ginsparg, F. Wu, J. Carrasquilla, P. L. McMahon, and E.-A. Kim, Attention-based quantum tomography, *Mach. Learn.: Sci. Technol.* **3**, 01LT01 (2022).
- [32] G. I. Struchalin, Y. A. Zagorovskii, E. V. Kovlakov, S. S. Straupe, and S. P. Kulik, Experimental estimation of quantum state properties from classical shadows, *PRX Quantum* **2**, 010307 (2021).
- [33] S. Becker, N. Datta, L. Lami, and C. Rouzé, Classical shadow tomography for continuous variables quantum systems, *IEEE Transactions on Information Theory* (IEEE, 2024).
- [34] T. Zhang, J. Sun, X.-X. Fang, X.-M. Zhang, X. Yuan, and H. Lu, Experimental quantum state measurement with classical shadows, *Phys. Rev. Lett.* **127**, 200501 (2021).
- [35] Z. Webb, The clifford group forms a unitary 3-design, *Quantum Inf. Comput.* **16**, 1379 (2015).
- [36] K. Choo, G. Carleo, N. Regnault, and T. Neupert, Symmetries and many-body excitations with neural-network quantum states, *Phys. Rev. Lett.* **121**, 167204 (2018).
- [37] D. Gottesman, *Stabilizer Codes and Quantum Error Correction* (California Institute of Technology, Pasadena, CA, 1997).
- [38] C. Gidney, Stim: A fast stabilizer circuit simulator, *Quantum* **5**, 497 (2021).
- [39] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
- [40] S. Bravyi and D. Maslov, Hadamard-free circuits expose the structure of the clifford group, *IEEE Trans. Inf. Theory* **67**, 4546 (2021).
- [41] W. Marciano and H. Pagels, Quantum chromodynamics, *Phys. Rep.* **36**, 137 (1978).
- [42] J. B. Kogut, The lattice gauge theory approach to quantum chromodynamics, *Rev. Mod. Phys.* **55**, 775 (1983).
- [43] M. C. Banuls, R. Blatt, J. Catani, A. Celi, J. I. Cirac, M. Dalmonte, L. Fallani, K. Jansen, M. Lewenstein, S. Montangero *et al.*, Simulating lattice gauge theories within quantum technologies, *Eur. Phys. J. D* **74**, 165 (2020).
- [44] B. Nachman, D. Provasoli, W. A. de Jong, and C. W. Bauer, Quantum algorithm for high energy physics simulations, *Phys. Rev. Lett.* **126**, 062001 (2021).
- [45] M. Dalmonte and S. Montangero, Lattice gauge theory simulations in the quantum information era, *Contemp. Phys.* **57**, 388 (2016).
- [46] E. A. Martinez, C. A. Muschik, P. Schindler, D. Nigg, A. Erhard, M. Heyl, P. Hauke, M. Dalmonte, T. Monz, P. Zoller *et al.*, Real-time dynamics of lattice gauge theories with a few-qubit quantum computer, *Nature (London)* **534**, 516 (2016).
- [47] Y. Y. Atas, J. F. Haase, J. Zhang, V. Wei, S. M.-L. Pfaendler, R. Lewis, and C. A. Muschik, Simulating one-dimensional quantum chromodynamics on a quantum computer: Real-time evolutions of tetra-and pentaquarks, *Phys. Rev. Res.* **5**, 033184 (2023).
- [48] M. Motta, C. Sun, A. T. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. Brandão, and G. K. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nat. Phys.* **16**, 205 (2020).
- [49] S.-S. Gong, W. Zhu, and D. Sheng, Emergent chiral spin liquid: Fractional quantum Hall effect in a kagome Heisenberg model, *Sci. Rep.* **4**, 1 (2014).
- [50] D. P. Arovas and A. Auerbach, Functional integral theories of low-dimensional quantum Heisenberg models, *Phys. Rev. B* **38**, 316 (1988).
- [51] T. Kennedy, E. H. Lieb, and B. S. Shastry, Existence of Néel order in some spin-1/2 Heisenberg antiferromagnets, *J. Stat. Phys.* **53**, 1019 (1988).
- [52] M. Żukowski, A. Zeilinger, M. Horne, and H. Weinfurter, Quest for GHZ states, *Acta Phys. Pol. A* **93**, 187 (1998).
- [53] S. N. Walck and D. W. Lyons, Only n-qubit Greenberger-Horne-Zeilinger states are undetermined by their reduced density matrices, *Phys. Rev. Lett.* **100**, 050501 (2008).
- [54] E. Zohar, Local manipulation and measurement of nonlocal many-body operators in lattice gauge theory quantum simulators, *Phys. Rev. D* **101**, 034518 (2020).
- [55] A. Blum, A. Kalai, and H. Wasserman, Noise-tolerant learning, the parity problem, and the statistical query model, *J. ACM (JACM)* **50**, 506 (2003).
- [56] S. Morawetz, I. J. S. De Vlucht, J. Carrasquilla, and R. G. Melko, U(1)-symmetric recurrent neural networks for quantum state reconstruction, *Phys. Rev. A* **104**, 012401 (2021).
- [57] A. Haldar, O. Tavakol, and T. Scaffidi, Variational wave functions for Sachdev-Ye-Kitaev models, *Phys. Rev. Res.* **3**, 023020 (2021).
- [58] H.-Y. Hu and Y.-Z. You, Hamiltonian-driven shadow tomography of quantum states, *Phys. Rev. Res.* **4**, 013054 (2022).
- [59] H.-Y. Hu, S. Choi, and Y.-Z. You, Classical shadow tomography with locally scrambled quantum dynamics, *Phys. Rev. Res.* **5**, 023027 (2023).
- [60] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, Efficient quantum state tomography, *Nat. Commun.* **1**, 149 (2010).
- [61] B. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. Buyskikh, A. Daley, M. Cramer *et al.*, Efficient tomography of a quantum many-body system, *Nat. Phys.* **13**, 1158 (2017).
- [62] K. Van Kirk, J. Cotler, H.-Y. Huang, and M. D. Lukin, Hardware-efficient learning of quantum many-body states, [arXiv:2212.06084](https://arxiv.org/abs/2212.06084).
- [63] M. Hibat-Allah, R. G. Melko, and J. Carrasquilla, Investigating topological order using recurrent neural networks, *Phys. Rev. B* **108**, 075152 (2023).
- [64] D. Luo, G. Carleo, B. K. Clark, and J. Stokes, Gauge equivariant neural networks for quantum lattice gauge theories, *Phys. Rev. Lett.* **127**, 276402 (2021).
- [65] F. Tao, H. Zhang, A. Liu, and A. Y. Nee, Digital twin in industry: State-of-the-art, *IEEE Trans. Ind. Inf.* **15**, 2405 (2018).
- [66] I. L. Gutiérrez and C. B. Mendl, Real time evolution with neural-network quantum states, *Quantum* **6**, 627 (2022).
- [67] M. Medvidović and G. Carleo, Classical variational simulation of the quantum approximate optimization algorithm, *npj Quantum Inf.* **7**, 101 (2021).
- [68] A. Elben, B. Vermersch, R. van Bijnen, C. Kokail, T. Brydges, C. Maier, M. K. Joshi, R. Blatt, C. F. Roos, and P. Zoller,

- Cross-platform verification of intermediate scale quantum devices, *Phys. Rev. Lett.* **124**, 010504 (2020).
- [69] G. Aleksandrowicz, T. Alexander, P. Barkoutsos, L. Bello, Y. Ben-Haim, D. Bucher, F. J. Cabrera-Hernández, J. Carballo-Franquis, A. Chen, C.-F. Chen *et al.*, Qiskit: An open-source framework for quantum computing, <https://zenodo.org/records/2562111> (2019), accessed Mar. 16.
- [70] E. Tang, A quantum-inspired classical algorithm for recommendation systems, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (ACM, 2019), pp. 217–228.
- [71] <https://github.com/IQB-Information-Technologies/NEM>.
- [72] <https://github.com/victor11235/Neural-Shadow-QST>.