Learning topological states from randomized measurements using variational tensor network tomography

Yanting Teng,¹ Rhine Samajdar,^{2,3} Katherine Van Kirk,¹ Frederik Wilde,⁴

Subir Sachdev,¹ Jens Eisert,⁴ Ryan Sweke,⁵ and Khadijeh Najafi^{6,7}

¹Department of Physics, Harvard University, Cambridge, MA 02138, USA

²Department of Physics, Princeton University, Princeton, NJ 08544, USA

³Princeton Center for Theoretical Science, Princeton University, Princeton, NJ 08544, USA

⁴Dahlem Center for Complex Quantum Systems,

Freie Universität Berlin, 14195 Berlin, Germany

⁵IBM Quantum, Almaden Research Center, San Jose, CA 95120, USA

⁶IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598 USA

⁷MIT-IBM Watson AI Lab, Cambridge, MA 02142, USA

(Dated: July 2, 2024)

Learning faithful representations of quantum states is crucial to fully characterizing the variety of many-body states created on quantum processors. While various tomographic methods such as classical shadow and MPS tomography have shown promise in characterizing a wide class of quantum states, they face unique limitations in detecting topologically ordered two-dimensional states. To address this problem, we implement and study a heuristic tomographic method that combines variational optimization on tensor networks with randomized measurement techniques. Using this approach, we demonstrate its ability to learn the ground state of the surface code Hamiltonian as well as an experimentally realizable quantum spin liquid state. In particular, we perform numerical experiments using MPS ansätze and systematically investigate the sample complexity required to achieve high fidelities for systems of sizes up to 48 qubits. In addition, we provide theoretical insights into the scaling of our learning algorithm by analyzing the statistical properties of maximum likelihood estimation. Notably, our method is sample-efficient and experimentally friendly, only requiring snapshots of the quantum state measured randomly in the X or Z bases. Using this subset of measurements, our approach can effectively learn any real pure states represented by tensor networks, and we rigorously prove that random-XZ measurements are tomographically complete for such states.

I. INTRODUCTION

In recent years, quantum simulators have proven instrumental to studying a diverse range of correlated quantum many-body systems, facilitating the investigation of phenomena that are often difficult to probe in conventional solid-state settings [1]. Today, there exist a host of candidate quantum simulation architectures, including ultracold atoms in optical lattices [2], neutral atom arrays [3], trapped ions [4], photonic devices [5], and superconducting quantum processors [6]; additionally, many of these platforms can be operated in either analog or digital modes. A common use case for such simulators is preparing the ground (or low-energy) states of some desired Hamiltonian. Importantly, going beyond simple symmetry-breaking ground states, quantum simulation has also enabled the experimental realization of highly entangled states with topological order [7, 8], characterized by fractionalized excitations and emergent gauge fields [9, 10].

Tomographic and other recovery methods for extracting information from these experiments are crucial, in particular, when they provide detailed diagnostic information about the prepared quantum state [11]. Indeed, recent approaches using randomized measurements have shown promise for measuring specific properties, such as the Renyi entropy [8, 12–14]. However, to extract Von Neumann entanglement entropy across arbitrary subsystem sizes using these methods might require multiple copies, substantially more experimental control, or more samples than what is currently available. Quantum state tomography (QST) is designed to reconstruct a quantum state from informationally complete measurements and is useful for systems where the desired observables are not known a priori [15] or cannot be measured due to the constraints of the experimental setup. Although QST is generally not scalable [16, 17], it becomes significantly more effective when applied with priors for particular classes of states. This makes OST a valuable alternative for extracting many properties from such systems, particularly when traditional methods are limited. Established state tomographic methods like matrix product state (MPS) tomography [18–20] provide a natural description with robust guarantees for onedimensional systems but fall short in accurately capturing two-dimensional quantum states, especially when they exhibit non-injective properties [21] and complex entanglement patterns [22]. While neural network quantum state tomography (NNQST) [23, 24] offers yet another option, its efficacy for topological states also remains unclear due to their intrinsic long-range entanglement and intricate phase structures across different components of the wavefunction [25].

To address the need for viable tomographic techniques for exotic two-dimensional systems [7], here, we demon-



Figure 1. Schematic illustration of variational tensor network tomography. Data collection: A target state vector $|\phi\rangle$ is measured in multiple bases for a collection of N samples defined in Eq. (1). For instance, the first panel shows measurements in a random X or Z basis for each qubit. Optimization: A tensor network model is variationally optimized from N samples using gradient descent, which outputs tensors $\hat{A}^{(N)}$ representing the state vector $|\psi(\hat{A}^{(N)})\rangle$. The middle panel illustrates an MPS model prior, as numerically investigated in this work. The loss function given by Eq. (5) includes the negative log-likelihood (NLL), typically used in maximum likelihood estimation (MLE). It is optionally regularized by physical observables estimated from the dataset (e.g., using classical shadow tomography represented by the pink dashed box). Evaluation: The algorithm outputs a tensor network or, in the current work, an MPS representation $\hat{A}^{(N)}$ as the learned state. We sometimes omit the number of samples and abbreviate the output tensor as \hat{A} . In the right panel, we evaluate the performance of our learning protocol by computing the fidelity $F = |\langle \phi | \psi(\hat{A}) \rangle|$. From the reconstructed state, we can extract arbitrary physical observables, including the von Neumann entanglement entropy and highly nonlocal Pauli strings (the inset shows an example of the nonlocal Z string operators considered in Fig. 4).

strate a heuristic state tomography method combining a tensor-network prior with maximum likelihood estimation (MLE) [26–29] and randomized measurements [25, 30]. Previous works have considered learning one-dimensional quantum states from uniformly random single-qubit measurements [26] or two-dimensional quantum channels with random Paulli measurements [29]. In this work, we focus on two-dimensional real pure states and show that they are completely characterized by only randomized XZ measurements, which randomly measure each qubit in either the X or the Z basis. The states considered here are not only potentially more complex than in previous studies but also easy to implement with high fidelity in experimental systems [31]. Moreover, we prove the tomographic completeness of randomized XZ measurements on the space of pure and real quantum states, allowing the method to be extendable to nonstoquastic real Hamiltonians' eigenstates [32]. Thus, our method provides a pragmatic pathway for characterizing two-dimensional states on near-term quantum devices, making use of only easy-to-implement measurements. We provide a detailed discussion of related works in Appendix A.

In our numerical experiments, we use an MPS ansatz as our model prior. While MPS-based representations of two-dimensional states are inherently limited, our approach sets the stage for higher-dimensional tensor networks, such as *projected entangled pair states* (PEPS), for future investigations. In addition, our method provides a scalable approach that aligns with the existing capabilities of modern simulators, thereby providing a benchmark for more advanced tomographic techniques. To supplement our numerical results, we leverage the statistical properties of MLE to provide insights into the sample complexity of our approach. More specifically, we provide a pedagogical overview of standard techniques for analyzing MLE [33–35], and discuss where they fail in this context, due to gauge freedom in MPS representations of quantum states. We then sketch how these difficulties can be circumvented by considering MLE on a manifold [36], and apply these results to the variational manifold given by our MPS prior [37].

The rest of this paper is structured as follows. In Sec. II, we present our protocol in conjunction with an in-depth discussion of our state-tomographic measurement schemes. Then, in Sec. III, we demonstrate the effectiveness of our protocol when applied to two topological states: a ground state of the surface code Hamiltonian and a \mathbb{Z}_2 quantum spin liquid state of Rydberg atoms arrayed on the ruby lattice. Our results are immediately relevant to current experiments since we only require feasible XZ measurements and consider system sizes achievable with state-of-the-art simulators. Finally, we conclude by discussing the prospects and potential directions for improvement of our methodology in Sec. IV.

II. VARIATIONAL TENSOR NETWORK TOMOGRAPHY

In this section, we introduce our variational *tensor network* (TN) tomography protocol. To begin, we briefly review the framework of randomized measurements, wherein randomly chosen unitaries are applied to a target state before measuring it in the computational

basis. Then, we utilize this data to variationally find the optimal tensor network approximation of the target state. We do so via maximum likelihood estimation, a standard technique for statistical inference; we also have the ability to regularize the MLE solution by estimations of physical quantities via classical shadows. This is done by adding a term to the MLE cost function. These steps are described below and represented by the schematic in Fig. 1.

a. Dataset and measurements. We follow the notation commonly used in randomized measurements [25, 38]. Let $|\phi\rangle$ be the unknown *n*-qubit pure state prepared by the quantum simulator, and let \mathcal{U} be the unitary ensemble with a uniform distribution over a set $\{U\}$. We apply a unitary $U \sim \mathcal{U}$ before measuring in the computational basis. Performing this process repeatedly, we collect N unitary-bit-string tuples as our dataset

$$\{(U_i, b_i)\}_{i=1}^N,\tag{1}$$

where each (U, b) is effectively sampled with probability

$$\frac{1}{|\{U\}|} \times \langle b|U|\phi\rangle\!\langle\phi|U^{\dagger}|b\rangle.$$
(2)

In this work, motivated by the available prior knowledge of the target state (which we discuss in the next subsection) as well as what measurements are experimentally feasible with high precision [12, 31], we consider two measurement schemes: (1) $\mathcal{U}_{\text{Global}}$: measuring every qubit in the X basis or every qubit in the Z basis, and (2) \mathcal{U}_{XZ} : randomly measuring each qubit in either the X or the Z basis. For the ensemble $\mathcal{U}_{\text{Global}}$, the unitaries would be uniformly sampled from

$$\left\{\prod_{i} U_i^X, \prod_{i} U_i^Z\right\}.$$
 (3)

The unitary U_i^{α} rotates the measurement basis of the *i*-th qubit to the $\alpha = X, Z$ basis. Note that the most general global controls have been considered in earlier work [15], in which the same SU(2) unitary is applied to each qubit. Here, we only focus on a discrete subset, of Pauli $X^{\otimes n}$ and $Z^{\otimes n}$ measurements, also considered in previous work [27]. For our second ensemble \mathcal{U}_{XZ} , the unitaries are uniformly sampled from

$$\left\{\prod_{i} U_{i}^{\alpha}, \, \alpha = X, Z\right\}.$$
(4)

In general, one can exploit priors on the target state to reduce the set of random unitaries required [15, 39]. For instance, consider learning the eigenstates of a real Hamiltonian. These states can be chosen to be real, and random-XZ measurements can learn any property of a state within such a space of real and pure states. This can be formally stated as follows.

Theorem 1 (Tomographic completeness). Random-XZ measurements U_{XZ} are tomographically complete on the space of real and pure states.

The intuition here is that when $|\phi\rangle$ is decomposed in terms of the computational basis states, all amplitudes can be learned by measuring every qubit in the computational (Z) basis. The (relative) signs can be learned by rotating some qubits to the X basis. We refer interested readers to Appendices C and D for details of the proof. In Appendix C, we formalize our random-XZ measurement protocol and rigorously derive the associated measurement channel in the framework of classical shadow tomography. Furthermore, in Appendix D, we motivate our restricted random-XZ measurement scheme, proving that it is tomographically complete in the space of real and pure states considered in this work. We note that this scenario is more general than that of a fully positive eigenstate of a stoquastic Hamiltonian, for which measurements in only the computational basis suffice [32].

b. Loss function. To estimate the target state from the dataset $\{(U_i, b_i)\}_{i=1}^N$, we use a tensor network ansatz $|\psi(A)\rangle$, where a sequence of local tensors is denoted by $A = (A^{[1]}, \dots, A^{[n]})$. The advantage of this formulation is that one can efficiently compute the probability of a given measurement outcome in any basis by applying the corresponding unitary to the tensor network model as shown in the middle panel of Fig. 1. This allows us to evaluate the loss function, whose minimum leads to our estimate of the target state. The loss function is

$$L(A) = -\frac{1}{N} \sum_{i=1}^{N} \log |\langle b_i | U_i | \psi(A) \rangle|^2 + \beta R(A), \quad (5)$$

where R(A) denotes the regularizer and its coefficient β describes how much we prioritize the regularization over MLE. The first term in Eq. (5) above, the *negative* log-likelihood (NLL), arises from the Kullback–Leibler divergence, which measures the difference between two probability distributions. In our case, these two probability distributions are over measurement outcomes in randomly selected bases, defined by the target and the model state, respectively. Minimizing the N-sample NLL leads to an MLE solution $\hat{A}^{(N)}$. To improve our training, we add the regularization term R(A): it penalizes the differences between our model's prediction of a physical observable $\langle \psi(A) | O | \psi(A) \rangle$ and the estimation $\langle O \rangle_N$ inferred directly from the dataset. For instance, we can choose $R(A) = \|\langle \psi(A) | O | \psi(A) \rangle - \langle O \rangle_N \|_2$ with a k-body Pauli string $O = X^{\otimes k}$, where the estimation $\langle X^{\otimes k} \rangle_N$ can be accurately determined from our global-XZ measurements. Broadly speaking, if we are confident about certain properties of the target state, our model should reflect those properties. The choices of regularization observables, therefore, depend on the systems of interest, which we discuss in Sec. III and Appendix E, where we describe the loss function and physical observables used for our numerical results.

c. MPS prior. Matrix product states, which are one of the most extensively studied families of tensor networks, have enjoyed remarkable success in accurately representing many-body quantum states. Thanks to an ef-



Figure 2. Numerical results for the surface code. (a) Schematic representation of the surface code of dimension (L_x, L_y) with Z stabilizers (red patches) and X stabilizers (blue patches). The MPS "snakes" through the system along the path shown (black arrow). (b) Scaling of the infidelity with the number of samples N for a 3×3 surface code with global-XZ (cross) or random-XZ (circle) measurements. Shades of blue indicates the regularization strength β . (c) To investigate larger system sizes, we focus on the surface code in a strip geometry $(L_y = 3)$ with a perturbation $(h_z = 0.1)$. The three lines indicate the infidelity for different numbers of qubits, n, with random-XZ measurements. Dark blue indicates a regularization $\beta = 5$. The dashed line marks the fidelity threshold for n = 9. (d) Sample complexity to reach a fixed local fidelity threshold $F_{local} = (0.99)^n$ extracted from (c). Colors and markers are the same as in (b).

fective variational algorithm known as the *density-matrix* renormalization group (DMRG) [40], MPS ansätze can be efficiently employed to describe the ground states of various many-body Hamiltonians. Strictly speaking, two-dimensional systems require the generalization of MPS, as the quasi-1D MPS description is limited to capturing short-range correlations and thus necessarily fails for sufficiently large and entangled systems. However, in practice, MPS ansätze are often used to represent ground states even in two dimensions, due to their relative advantage of efficient contractions. For the remainder of this paper, we utilize an MPS as our model prior, and we variationally find a state that is most consistent with the observed measurements, as shown in Fig. 1. The optimization of the MPS is carried out by taking gradients simultaneously with respect to all the tensor components; we provide the details in Appendix G.

III. RESULTS

In this section, we demonstrate the effectiveness of our variational tomography protocol for two topologically ordered states. While our eventual goal is to perform tomography on a state in the laboratory, for our results here, we perform numerical "experiments" using simulated target states obtained via DMRG, as detailed in Appendix F. Our protocol assumes that the laboratory state is close to a pure state that could be approximated as an MPS (albeit with a possibly large bond dimension). After obtaining our simulated measurements by sampling the target state, we initiate our protocol with random initial parameters. These parameters are iteratively adjusted until they converge to the trained state. Then, we evaluate the efficacy of this training process by computing the fidelity between the trained and target states. Here, we analyze 10 such random initializations, which are shown in our numerical results below. Of course, when using realistic experimental data, one would lack

access to the "ground truth", i.e., a classical description of the target state, which is needed to compute the fidelity exactly. In this case, we perform cross validation by comparing converged models on a test dataset [Appendix G.2], or alternatively, one could consider efficient certifications of high fidelity [41, 42].

A. Learning the surface code

In our first demonstration, we consider learning the ground state of the perturbed surface code Hamiltonian [43], $H_{sc} = -\sum_{l} S_{Z,l} - \sum_{l} S_{X,l} - h_{z} \sum_{i} Z_{i}$, defined on the square lattice. Here, $S_{P,l} = \prod_{i \in \Box_{l}} P_{i}$ denotes the *l*-th stabilizer, which is a product of Pauli P operators (P = X or P = Z) on a plaquette or boundary [Fig. 2(a)], and h_z is a perturbative field. Such a perturbation drives the ground state away from the exactly solvable limit, but it can still be well approximated by a tensor network. Let $|\phi\rangle$ be the target MPS representation of the ground state along a snake path. Although our target state is possibly perturbed away from the exact stabilizer-state limit, the expectation values of the stabilizer operators should be smoothly connected to their fixed-point values (≈ 1) as a function of the perturbation strength. This knowledge informs our choice of the set of stabilizers as observables for regularization. We estimate these stabilizers from the finite dataset either via direct measurement in the global-XZ scheme or from classical shadows for random-XZ measurements.

Our results for the surface code are presented in Fig. 2. To begin, we consider the surface code without any external perturbation $(h_z = 0)$: for n = 9 qubits arrayed on a square geometry, our protocol outputs the state vector $|\psi(\hat{A})\rangle$, and we evaluate the error of our learning protocol using the infidelity, $1 - F = 1 - |\langle \phi | \psi(\hat{A}) \rangle|$. Figure 2(b) shows that the infidelity improves with the number of samples N, for both the global-XZ and random-

XZ measurements discussed in Sec. II. For the latter, we find that regularization over the stabilizers, by increasing β , further reduces the infidelity in the sample-limited regime (N < 1000). We observe that both the global and random measurement schemes succeed in learning these states, and furthermore, the former outperforms the latter in terms of requiring fewer samples to achieve the same infidelity. However, while global-XZ measurements $\mathcal{U}_{\text{Global}}$ are sufficient for certain highly structured states such as a surface code (or GHZ) state, they may not necessarily be suited for learning generic states, wherefore random-XZ measurements \mathcal{U}_{XZ} could be beneficial. Finally, we note that the infidelity scales polynomially with the number of samples N, similarly to observations in earlier work on one-dimensional states [26]. In particular, focusing on the exact surface code ground state (for $h_z = 0$) trained without regularization ($\beta = 0$), we find an excellent fit of $1 - F \propto 1/N^{\alpha}$, as shown in Fig. 2(b), with $\alpha = 1.16$ ($\alpha = 1.05$) for random-XZ (global-XZ) measurements. To substantiate these numerical observations regarding the sample complexity needed to achieve a certain fidelity, we rigorously establish in Theorem 2 that the infidelity can be probabilistically upper bounded such that

$$\mathbb{P}\left[1 - |\langle \phi | \psi(\hat{A}^{(N)}) \rangle| \le \epsilon(N)\right] \ge 1 - \delta, \tag{6}$$

where $\epsilon(N) = O([n\chi_{\max}^2/(N\delta)]^{1/2})$ and χ_{\max} represents the maximum bond dimension. Interestingly, our numerical results, which exhibit $\alpha > 0.5$, suggest better performance than our theoretical bound for arbitrary target states. Further details and substantial analytical and rigorous insights into the sample complexity, drawing from the statistical properties of MLE, are discussed in Appendix B.

Next, we add a perturbation to the Hamiltonian, $h_z =$ 0.1, and analyze the performance of our protocol for various system sizes with up to n = 45 qubits. Given the limitations of an MPS ansatz, which is inherently quasi-onedimensional, here, we only increase the length of the system L_x along the x-axis. In Fig. 2(c), we illustrate how the infidelity scales for different system sizes (indicated by different lines) using random-XZ measurements and find that the fidelity improves with N. Finally, we analyze how the sample complexity scales with the system size n. The requirement of achieving a constant fidelity is stringent as the overlap between two many-body states can vanish exponentially with n, a phenomenon known as the orthogonality catastrophe [44]. For a finite system size, however, the overlap still serves as a useful diagnostic indicator [45]. Therefore, we relax the requirement of a fixed global fidelity by considering instead the per-site local fidelity $F_{\text{local}} \equiv F^{1/n}$. Extrapolating from the data in (c), Fig. 2(d) plots the samples required to achieve a local fidelity threshold, e.g., as indicated by the dashed line in (c) for n = 9 qubits. For the largest system size, n = 45, we find that only a few thousand measurements are necessary to attain a fidelity of $F_{\text{local}} = 0.99$.



Figure 3. Numerical results for learning quantum spin liquids. (a) The tuning parameters of the Rydberg Hamiltonian are the Rabi frequency Ω , the detuning Δ , and the blockade radius R_b . The gray arrow shows the mapping between a particular measurement bit-string on the ruby lattice and a dimer configuration on the kagome lattice: Light (dark) blue circles correspond to an atom being in the excited $|r\rangle$ (ground $|g\rangle$) state and maps to the presence (absence) of a dimer. A careful choice of the Rydberg blockade radius $R_b = 2a$, where a is the lattice spacing, ensures that the dimers do not overlap. Bottom panel: The two phases as a function of the detuning ratio Δ/Ω . (b) The ground state belongs to the trivial disordered phase at $\Delta/\Omega = 0.5$. The infidelity decreases with an increasing number of samples N. For random-XZ measurements with N < 40000, a regularization of $\beta = 1$ (red circles) further improves the fidelity compared to the case without regularization $\beta = 0$ (red squares). Global-XZ measurements (gray crosses) perform worse than their random-XZ counterparts (red). (c) Same as in (b) but for the spin-liquid state at $\Delta/\Omega = 1.7$. Relative to the case with no regularization $\beta = 0$ (light blue), regularization $\beta = 1$ (dark blue) does not improve the fidelity but reduces the number of nonconvergent outliers (above the black dashed line). Global-XZ (gray cross) measurements perform similarly to random-XZ ones (blue).

B. Learning a Rydberg quantum spin liquid

For a second application of our tomographic framework, we turn to learning ground states of strongly interacting arrays of neutral atoms. Over the last decade, these Rydberg atom arrays have evolved into mature platforms for quantum simulation and have shown great potential for probing a variety of correlated quantum phases of matter [3, 7, 46-48]. The effective Hamiltonian of this system can be written as

$$H_{\text{Ryd}} = \frac{\Omega}{2} \sum_{\ell} X_{\ell} - \Delta \sum_{\ell} \frac{1}{2} (1 + Z_{\ell}) + \frac{1}{2} \sum_{\ell, \ell'} \frac{V_{\ell, \ell'}}{4} (1 + Z_{\ell}) (1 + Z_{\ell'}), \quad (7)$$

where Ω represents the Rabi frequency, Δ is the detuning of the laser drive, and $V_{\ell,\ell'}$ is the van der Waals interaction potential. The strong Rydberg-Rydberg interactions prevent neighboring atoms lying within a "blockade radius", R_b , from being simultaneously excited to the Rydberg state, thereby engendering strong quantum correlations. For an appropriate choice of the blockade radius, configurations of Rydberg atoms on the ruby lattice can be mapped to a set of "dimers" on the kagome lattice [Fig. 3(a)]; such quantum dimer models have long been known to host topological quantum spin liquids [49, 50]. In the Rydberg system, depending on the detuning parameter Δ/Ω , the ground state (for this specifically chosen R_b) can be either a trivial disordered phase or a topologically ordered \mathbb{Z}_2 quantum spin liquid phase [7, 51– 53]; see Fig. 3(a). In the computational basis, these two phases are difficult to distinguish as they both lack any symmetry-breaking order. Here, we select two parameters in the phase diagram belonging to the spin-liquid $(\Delta/\Omega = 1.7)$ or the disordered $(\Delta/\Omega = 0.5)$ phase and perform DMRG simulations to find these ground states (see Appendix F for more details).

Importantly, for this system, we do not have prior knowledge of which observables are important to regularize, unlike for the surface code studied above. Hence, analogously to MPS tomography, we consider estimations of all estimable Pauli strings that are supported on a local subsystem (taken to be 6 qubits here) and use them to regularize our MLE loss function. This subsystem should be chosen such that it captures the important correlations of the state and that the Paulis supported on this subsystem can be estimated accurately using randomized measurements¹. With regard to the task of learning the states of n = 48 qubits, we now show that our protocol (see Appendix \mathbf{E}) is capable of learning both regimes in the phase diagram—i.e., the spin liquid as well as the disordered state—which cannot be distinguished by a local order parameter due to the topological nature of the former.

Focusing first on the disordered state shown in Fig. 3(b), we see that the fidelity improves with the number of measurements N for all measurement schemes. Moreover, we observe a slight improvement with a regularization strength of $\beta = 1$ in the sample-limited regime



Figure 4. Prediction of spin-liquid properties. (a) A quantum spin liquid (QSL) state can be viewed as a superpostions of dimers in the computational basis. (b) Definition of measurable diagonostic quantities such as the Z string operators. Top panel: A dashed line across a bond is equivalent to applying the Pauli Z operator to the qubit on the same bond. The Z logical string P_{Logical} applies the Z string operator across the system. Bottom panel: An open string diagonal FM operator $P_{\rm FM}$, renormalized by the expectation value of a closed-loop operator; e.g., the semicircle corresponds to the product of 6 Pauli Z operators. (c) The expectation values of the $P_{\rm FM}$ (red) and $P_{\rm Logical}$ (blue) operators for different numbers of samples, N, converge to their exact values (lines). (d) Predictions (crosses) of the entanglement entropy for different sizes of the bipartition and (e) the Schmidt values (cut at site i = 21) from MPS ansätze trained with N samples converge to the exact values (open circles).

 $(N \leq 40000)$. Next, we turn to the spin liquid illustrated in 4(a). For this state as well, the infidelity generally reduces with N, as shown in Fig. 3(c), albeit with outliers that did not converge (indicated by squares above the dashed line). However, in contrast to the disordered phase, we observe that regularization, with strength $\beta = 1$ (indicated by the dark blue circles), does not improve the fidelity beyond that obtained with no regularization $\beta = 0$ (light blue squares), but the number of outliers (above the dashed line) is reduced. This lack of substantial improvement upon regularization suggests that the important operators capturing the correlations of the spin-liquid state are supported on a larger subsystem than used in our current numerical experiments.

Finally, we show that, using a successfully trained

¹ Note that we could also consider regularization from global measurements, but then we would only be able to estimate $X^{\otimes k}$ or $Z^{\otimes k}$.

MPS, we can evaluate physical properties that may be difficult to measure directly given limited experimental controls. Certain diagnostic physical observables, such as the diagonal Fredenhagen-Marcu (FM) order parameter $P_{\rm FM}$ as well as the Z logical loop operator $P_{\rm logical}$ and its X counterpart, have been previously suggested for the spin-liquid state [7, 52, 54], as illustrated in Figs. 4(a,b). These measurable observables can be used to verify the consistency of our learned states, and we see that the predictions from our trained MPS converge well to the target values [Fig. 4(c)]. Moreover, our protocol allows for the extraction of additional quantities such as the von Neumann entanglement entropy, which is essential to determining the topological entanglement entropy—a key quantity that characterizes a gapped spin liquid [55, 56]. Figure 4(d) demonstrates that the bipartite entanglement entropy converges to the target's as the bipartition varies in subsystem size for $N \approx 90000$. For a cut across roughly half of the system, the Schmidt values in Fig. 4(e)also closely match the exact values of the target state.

IV. DISCUSSION AND OUTLOOK

In this work, we combine randomized measurements with tensor networks to perform quantum state tomography via MLE, with optional regularization using classical shadows. While the randomized measurements toolbox is useful for estimating many physical properties, tensor networks are known to efficiently represent physically relevant quantum states. Here, we have considered a more restricted measurement setting, random-XZshadows, that is well motivated by both its tomographic completeness on real and pure states and its feasibility in near-term experiments. Using these XZ measurements, we can regularize our MLE loss function with observables estimated via the classical shadow framework. We showcase the performance of our protocol by learning the topological ground states of the surface code Hamiltonian as well as the Rydberg Hamiltonian on the ruby lattice. For both of these states, our protocol with random-XZmeasurements can accurately reconstruct the target state up to a fidelity of over 0.95 with under 10^5 samples. In addition, we test our shadow regularization technique on these two applications: we find that requiring the stabilizer expectations values to agree with estimations from classical shadows further improves the fidelity for the surface code, whereas for the ruby-lattice Rydberg spin liquid, this only reduces the number of outliers observed during the optimization process.

Using higher-dimensional TNs combined with randomized measurements would, in principle, allow us to learn any states of interest. In the present work, we have numerically benchmarked our protocol using MPS as our model prior, given its advantages of exact contractions and sampling. An immediate extension of our work would be to generalize this model's architecture to represent 2D area-law entangled states, such as string bond states [57], projected entangled pair states (PEPS) [58]. We leave the numerical demonstrations of our tomographic method for such tensor-network states to future work.

It constitutes an interesting perspective to extend these methods to the study of fermionic quantum states, which are the focus of attention in the study of correlated electron systems. Our analysis could be extended to quantum gas microscope measurements of optical lattice Fermi-Hubbard systems [59] which yield snapshots with the charge and spin on each site. While in this setting, quench dynamics generated by noninteracting Hamiltonians are presumably most conceivable [60], it is possible that the spin degrees of freedom can be measured using a similar unitary ensemble such as our random-XZmeasurements. The variational optimization can be carried out using fermionic neural networks [61] or auxiliary wavefunctions [62, 63], replacing the MPS we have employed in the present work. To enhance the optimization process using neural-network ansätze, it would be especially interesting to extend recent results [42] on locallyrandomized single-qubit measurements to such fermionic systems.

While our method can be generalized to quantum states realized in a wide class of quantum processors, more specifically, for Rydberg quantum simulators, it can be easily extended to learning logically encoded quantum states. In particular, our random-XZ measurements are well suited for near-term encodings, such as the color codes recently implemented in neutral atom arrays [31]. In such a Rydberg array, every logical qubit, encoded, for example, in a Steane code, can be measured in either the X or the Z basis, allowing for both error detection via post-selection and the ability to perform high-fidelity random-XZ measurements. Even if current experiments do not aim to perform full logical quantum state tomography, our characterization and derivation of the tools required for random-XZ classical shadows can be immediately applied to these logical encodings.

In settings where qubits are not logically encoded, our techniques are also experimentally amenable for physical qubits. Even as the degree of control over large quantum systems improves, single-qubit rotations will likely always remain preferable compared to many-qubit gates due to their orders-of-magnitude higher fidelity and short implementation times [64]. Therefore, the ability to accurately estimate a laboratory state from random measurements will remain attractive beyond the NISQ era. Our protocol can also be straightforwardly extended to random-Pauli measurements [25]. Moreover, this technique could be made even more attractive by incorporating adaptivity into our measurement protocol. Judiciously altering the measurement ensemble as more data is obtained allows for more efficient learning and thus reduces the sample complexity. To implement an adaptive learning scheme, one may incorporate ideas from Ref. [65] into our learning technique in order to use fewer measurements to achieve the same infidelity.

Finally, our work also lends itself to learning about experimental imperfections in state preparation. In this sense, our scheme goes beyond state verification: it learns not only the state prepared but also the unwanted deviations from the desired state, which constitute feedback to the experiment on how to better achieve a higher-fidelity preparation over many iterations. This variational state preparation is practical from a sample-complexity point of view when the underlying state has some efficient representation; in other words, an exponential number of degrees of freedom need not be learned. This is precisely the case in many condensed-matter contexts, where typical states exhibit entanglement structures that are well captured by tensor networks [22].

V. CODE AND DATA AVAILABILITY

The numerical code used for producing results and generating the figures in this paper is available at: https://github.com/teng10/tn-shadow-qst.

- J. I. Cirac and P. Zoller, Goals and opportunities in quantum simulation, Nat. Phys. 8, 264 (2012).
- [2] I. Bloch, J. Dalibard, and S. Nascimbene, Quantum simulations with ultracold quantum gases, Nat. Phys. 8, 267 (2012).
- [3] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, S. Choi, S. Sachdev, M. Greiner, V. Vuletić, and M. D. Lukin, Quantum phases of matter on a 256-atom programmable quantum simulator, Nature 595, 227 (2021).
- [4] R. Blatt and C. F. Roos, Quantum simulations with trapped ions, Nat. Phys. 8, 277 (2012).
- [5] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins, *et al.*, Quantum computational advantage with a programmable photonic processor, Nature **606**, 75 (2022).
- [6] A. A. Houck, H. E. Türeci, and J. Koch, On-chip quantum simulation with superconducting circuits, Nat. Phys. 8, 292 (2012).
- [7] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin, Probing topological spin liquids on a programmable quantum simulator, Science **374**, 1242 (2021).
- [8] K. J. Satzinger, Y. Liu, A. Smith, C. Knapp, M. Newman, C. Jones, Z. Chen, C. Quintana, X. Mi, A. Dunsworth, C. Gidney, I. Aleiner, F. Arute, K. Arya, J. Atalaya, R. Babbush, J. C. Bardin, R. Barends, J. Basso, A. Bengtsson, A. Bilmes, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. R. Derk, D. Eppens, C. Erickson, E. Farhi, L. Foaro, A. G. Fowler, B. Foxen, M. Giustina, A. Greene, J. A. Gross, M. P.

ACKNOWLEDGMENTS

We thank Jacob Barandes. Christian Bertoni. Jonathan Conrad, Zoe Holmes, Robert Huang, Dmitrii Kochkov, Nathan Leitao, Alexander Nietner, Jakob Unfried, Guifre Vidal, and Manuel Rudolph for useful discussions. We thank Hong-Ye Hu and Leandro Aolita for feedback on the manuscript. Y.T. is grateful to Dmitrii Kochkov for his inputs on the initial prototype of the numerical code base and systematic code reviews. Y.T. also thanks the groups of Zoe Holmes and Jens Eisert for their hospitality during her visits, when part of this work was completed. R.S. is supported by the Princeton Quantum Initiative Fellowship. K.V.K. acknowledges support from the Fannie and John Hertz Foundation and the National Defense Science and Engineering Graduate (ND-SEG) Fellowship. The Berlin team acknowledges support from the BMBF (MuniQCAtoms), the DFG (CRC 183, FOR 2724), the Munich Quantum Valley, and the ERC (DebuQC). Y.T. and S.S. have been supported by the U.S. Department of Energy under Grant DE-SC0019030.

Harrigan, S. D. Harrington, J. Hilton, S. Hong, T. Huang,
W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey,
Z. Jiang, D. Kafri, K. Kechedzhi, T. Khattar, S. Kim,
P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, A. Locharla, E. Lucero, O. Martin, J. R.
McClean, M. McEwen, K. C. Miao, M. Mohseni, S. Montazeri, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, T. E. O'Brien, A. Opremcak, B. Pató, A. Petukhov, N. C. Rubin, D. Sank,
V. Shvarts, D. Strain, M. Szalay, B. Villalonga, T. C.
White, Z. Yao, P. Yeh, J. Yoo, A. Zalcman, H. Neven,
S. Boixo, A. Megrant, Y. Chen, J. Kelly, V. Smelyanskiy, A. Kitaev, M. Knap, F. Pollmann, and P. Roushan,
Realizing topologically ordered states on a quantum processor, Science **374**, 1237 (2021).

- [9] X. G. Wen, Mean-field theory of spin-liquid states with finite energy gap and topological orders, Phys. Rev. B 44, 2664 (1991).
- [10] S. Sachdev, Kagomé- and triangular-lattice Heisenberg antiferromagnets: Ordering from quantum fluctuations and quantum-disordered ground states with unconfined bosonic spinons, Phys. Rev. B 45, 12377 (1992).
- [11] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, and E. Kashefi, Quantum certification and benchmarking, Nature Rev. Phys. 2, 382 (2020).
- [12] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletić, and M. D. Lukin, A quantum processor based on coherent transport of entangled atom arrays, *Nature* **604**, 451 (2022).
- [13] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, Probing Rényi entanglement entropy via randomized measurements, Science **364**, 260 (2019).

- [14] H.-Y. Hu, A. Gu, S. Majumder, H. Ren, Y. Zhang, D. S. Wang, Y.-Z. You, Z. Minev, S. F. Yelin, and A. Seif, Demonstration of robust and efficient quantum property learning with shallow shadows (2024), arXiv:2402.17911.
- [15] K. V. Kirk, J. Cotler, H.-Y. Huang, and M. D. Lukin, Hardware-efficient learning of quantum manybody states (2022), arXiv:2212.06084.
- [16] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Quantum state tomography via compressed sensing, Phys. Rev. Lett. 105, 150401 (2010).
- [17] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, Sampleoptimal tomography of quantum states, IEEE Trans. Inf. Theory 63, 5628 (2017).
- [18] 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).
- [19] T. Baumgratz, D. Gross, M. Cramer, and M. B. Plenio, Scalable reconstruction of density matrices, Phys. Rev. Lett. **111**, 020401 (2013).
- [20] R. Hübener, A. Mari, and J. Eisert, Wick's theorem for matrix product states, Phys. Rev. Lett. **110**, 040401 (2013).
- [21] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Matrix product state representations (2007), arxiv:quant-ph/0608197.
- [22] J. Eisert, M. Cramer, and M. B. Plenio, Area laws for the entanglement entropy, Rev. Mod. Phys. 82, 277 (2010).
- [23] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, Nat. Phys. 14, 447 (2018).
- [24] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, Nature Mach. Intell. 1, 155 (2019).
- [25] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. 16, 1050 (2020).
- [26] J. Wang, Z.-Y. Han, S.-B. Wang, Z. Li, L.-Z. Mu, H. Fan, and L. Wang, Scalable quantum tomography with fidelity estimation, Phys. Rev. A 101, 032321 (2020).
- [27] A. M. Gomez, S. F. Yelin, and K. Najafi, Reconstructing Quantum States Using Basis-Enhanced Born Machines (2022), arXiv:2206.01273 [quant-ph].
- [28] M. K. Kurmapu, V. Tiunova, E. Tiunov, M. Ringbauer, C. Maier, R. Blatt, T. Monz, A. K. Fedorov, and A. Lvovsky, Reconstructing complex states of a 20-qubit quantum simulator, PRX Quantum 4, 040345 (2023).
- [29] G. Torlai, C. J. Wood, A. Acharya, G. Carleo, J. Carrasquilla, and L. Aolita, Quantum process tomography with unsupervised learning and tensor networks, Nat. Commun. 14, 2858 (2023).
- [30] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, Science 377, eabk3333 (2022).
- [31] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, Logical quantum processor based on reconfigurable atom arrays, Nature **626**, 58–65 (2023).
- [32] S. Bravyi and B. Terhal, Complexity of stoquastic frustration-free Hamiltonians, SIAM J. Comput. 39, 1462 (2010).

- [33] W. K. Newey and D. McFadden, Chapter 36 Large sample estimation and hypothesis testing, in *Handbook of Econometrics*, Vol. 4 (Elsevier, 1994) pp. 2111–2245.
- [34] F. Wilde, A. Kshetrimayum, I. Roth, D. Hangleiter, R. Sweke, and J. Eisert, Scalably learning quantum many-body Hamiltonians from dynamical data (2022), arxiv:2209.14328.
- [35] T. L. Scholten and R. Blume-Kohout, Behavior of the maximum likelihood in quantum state tomography, New Journal of Physics 20, 023050 (2018).
- [36] H. Hajri, S. Said, and Y. Berthoumieu, Maximum likelihood estimators on manifolds, in *Geometric Science of Information*, Vol. 10589, edited by F. Nielsen and F. Barbaresco (Springer International Publishing, Cham, 2017) pp. 692–700, series Title: Lecture Notes in Computer Science.
- [37] J. Haegeman, M. Mariën, T. J. Osborne, and F. Verstraete, Geometry of matrix product states: Metric, parallel transport, and curvature, J. Math. Phys. 55, 021902 (2014).
- [38] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, Nature Rev. Phys. 5, 9–24 (2022).
- [39] C. Hadfield, S. Bravyi, R. Raymond, and A. Mezzacapo, Measurements of quantum Hamiltonians with locallybiased classical shadows (2020), arXiv:2006.15788.
- [40] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, Ann. Phys. 326, 96 (2011).
- [41] A. L. Shaw, Z. Chen, J. Choi, D. K. Mark, P. Scholl, R. Finkelstein, A. Elben, S. Choi, and M. Endres, Benchmarking highly entangled states on a 60-atom analogue quantum simulator, Nature 628, 71 (2024).
- [42] H.-Y. Huang, J. Preskill, and M. Soleimanifar, Certifying almost all quantum states with few single-qubit measurements (2024), arXiv:2404.07281.
- [43] A. Y. Kitaev, Fault-tolerant quantum computation by anyons, Ann. Phys. 303, 2 (2003).
- [44] P. W. Anderson, Infrared catastrophe in Fermi gases with local scattering potentials, Phys. Rev. Lett. 18, 1049 (1967).
- [45] P. Zanardi and N. Paunković, Ground state overlap and quantum phase transitions, Phys. Rev. E 74, 031123 (2006).
- [46] S. de Léséleuc, V. Lienhard, P. Scholl, D. Barredo, S. Weber, N. Lang, H. P. Büchler, T. Lahaye, and A. Browaeys, Observation of a symmetry-protected topological phase of interacting bosons with Rydberg atoms, Science 365, 775 (2019).
- [47] R. Samajdar, W. W. Ho, H. Pichler, M. D. Lukin, and S. Sachdev, Complex density wave orders and quantum phase transitions in a model of square-lattice Rydberg atom arrays, Phys. Rev. Lett. **124**, 103601 (2020).
- [48] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, A. M. Läuchli, and A. Browaeys, Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms, Nature 595, 233 (2021).
- [49] R. Moessner and K. S. Raman, Quantum dimer models, in *Introduction to Frustrated Magnetism* (Springer, 2011) pp. 437–479.
- [50] G. Misguich, D. Serban, and V. Pasquier, Quantum dimer model on the Kagome lattice: Solvable dimerliquid and Ising gauge theory, Phys. Rev. Lett. 89,

137202 (2002).

- [51] R. Samajdar, W. W. Ho, H. Pichler, M. D. Lukin, and S. Sachdev, Quantum phases of Rydberg atoms on a Kagome lattice, Proc. Natl. Acad. Sci. U.S.A. 118, e2015785118 (2021).
- [52] R. Verresen, M. D. Lukin, and A. Vishwanath, Prediction of toric code topological order from Rydberg blockade, Phys. Rev. X 11, 031005 (2021).
- [53] R. Samajdar, D. G. Joshi, Y. Teng, and S. Sachdev, Emergent Z₂ gauge theories and topological excitations in Rydberg atom arrays, Phys. Rev. Lett. **130**, 043601 (2023).
- [54] I. Cong, N. Maskara, M. C. Tran, H. Pichler, G. Semeghini, S. F. Yelin, S. Choi, and M. D. Lukin, Enhancing detection of topological order by local error correction, Nat. Commun. 15, 1527 (2024).
- [55] A. Kitaev and J. Preskill, Topological entanglement entropy, Phys. Rev. Lett. 96, 110404 (2006).
- [56] M. Levin and X.-G. Wen, Detecting topological order in a ground state wave function, Phys. Rev. Lett. 96, 110405 (2006).
- [57] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and I. Cirac, Expressive power of tensor-network factorizations for probabilistic modeling, in *Advances in Neural Information Processing Systems*, Vol. 32 (Curran Associates, Inc., 2019).
- [58] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions (2004), arXiv:cond-mat/0407066.
- [59] J. Koepsell, D. Bourgund, P. Sompet, S. Hirthe, A. Bohrdt, Y. Wang, F. Grusdt, E. Demler, G. Salomon, C. Gross, and I. Bloch, Microscopic evolution of doped Mott insulators from polaronic metal to Fermi liquid, Science **374**, 82–86 (2021).
- [60] J. Denzler, A. A. Mele, E. Derbyshire, T. Guaita, and J. Eisert, Learning fermionic correlations by evolving with random translationally invariant Hamiltonians (2023), arXiv:2309.12933 [quant-ph].
- [61] J. Robledo Moreno, G. Carleo, A. Georges, and J. Stokes, Fermionic wave functions from neural-network constrained hidden states, Proc. Natl. Acad. Sci. U.S.A 119, e2122059119 (2022).
- [62] H. Shackleton and S. Zhang, Variational wavefunctions for topologically-ordered Fermi liquids, Bull. Am. Phys. Soc., W08.00014 (March 2024).
- [63] T. Müller, R. Thomale, S. Sachdev, and Y. Iqbal, Polaronic correlations from optimized ancilla wave functions for the Fermi-Hubbard model, to appear (2024).
- [64] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin, High-fidelity parallel entangling gates on a neutral atom quantum computer, Nature 622, 268 (2023).
- [65] H. Lange, M. Kebrič, M. Buser, U. Schollwöck, F. Grusdt, and A. Bohrdt, Adaptive quantum state tomography with active learning, Quantum 7, 1129 (2023).
- [66] V. Gebhart, R. Santagati, A. A. Gentile, E. M. Gauger, D. Craig, N. Ares, L. Banchi, F. Marquardt, L. Pezzè, and C. Bonato, Learning quantum systems, Nature Rev. Phys. 5, 141 (2023).
- [67] M. Ohliger, V. Nesme, and J. Eisert, Efficient and feasible state tomography of quantum many-body systems, New J. Phys. 15, 015024 (2013).

- [68] A. Nietner, Unifying (quantum) statistical and parametrized (quantum) algorithms (2023), arXiv:2310.17716.
- [69] D. S. Kufel, J. Kemp, S. M. Linsel, C. R. Laumann, and N. Y. Yao, Approximately-symmetric neural networks for quantum spin liquids, arXiv preprint arXiv:2405.17541 (2024).
- [70] V. Wei, W. A. Coish, P. Ronagh, and C. A. Muschik, Neural-shadow quantum state tomography (2023), arXiv:2305.01078.
- [71] N.-H. Chia, C.-Y. Lai, and H.-H. Lin, Efficient learning of t-doped stabilizer states with single-copy measurements, Quantum 8, 1250 (2024).
- [72] S. Masot-Llima and A. Garcia-Saez, Stabilizer tensor networks: universal quantum simulator on a basis of stabilizer states (2024), arXiv:2403.08724.
- [73] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, A hierarchy for replica quantum advantage (2021), arXiv:2111.05874.
- [74] R. Islam, R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner, Measuring entanglement entropy in a quantum many-body system, Nature 528, 77 (2015).
- [75] A. Elben, B. Vermersch, M. Dalmonte, J. I. Cirac, and P. Zoller, Rényi entropies from random quenches in atomic Hubbard and spin models, Phys. Rev. Lett. 120, 050406 (2018).
- [76] C. Kokail, R. van Bijnen, A. Elben, B. Vermersch, and P. Zoller, Entanglement Hamiltonian tomography in quantum simulation, Nat. Phys. 17, 936 (2021).
- [77] A. Elben, J. Yu, G. Zhu, M. Hafezi, F. Pollmann, P. Zoller, and B. Vermersch, Many-body topological invariants from randomized measurements in synthetic quantum matter, Sci. Adv. 6, eaaz3666 (2020).
- [78] Z.-P. Cian, H. Dehghani, A. Elben, B. Vermersch, G. Zhu, M. Barkeshli, P. Zoller, and M. Hafezi, Manybody Chern number from statistical correlations of randomized measurements, Phys. Rev. Lett. **126**, 050501 (2021).
- [79] A. H. Werner, D. Jaschke, P. Silvi, M. Kliesch, T. Calarco, J. Eisert, and S. Montangero, Positive tensor network approach for simulating open quantum many-body systems, Physical review letters **116**, 237201 (2016).
- [80] W. Kühnel, Differentialgeometrie, 4th ed. (Vieweg, Wiesbaden, 2008).
- [81] M. Nakahara, *Geometry, Topology and Physics*, 2nd ed. (Taylor & Francis, 2003).
- [82] R. Cheng, Quantum geometric tensor (Fubini-Study metric) in simple quantum system: A pedagogical introduction (2013), arXiv:1012.1337.
- [83] I. Bengtsson and K. Zyczkowski, *Geometry of quantum states* (Cambridge University Press, Cambridge, 2020).
- [84] M. Hauru, M. Van Damme, and J. Haegeman, Riemannian optimization of isometric tensor networks, SciPost Phys. 10, 040 (2021).
- [85] I. A. Luchnikov, M. E. Krechetov, and S. N. Filippov, Riemannian geometry and automatic differentiation for optimization problems of quantum physics and quantum technologies, New J. Phys. 23, 073006 (2021).
- [86] M. Ippoliti, Y. Li, T. Rakovszky, and V. Khemani, Operator relaxation and the optimal depth of classical shadows, Phys. Rev. Lett. **130**, 230403 (2023).
- [87] H.-Y. Hu, S. Choi, and Y.-Z. You, Classical shadow tomography with locally scrambled quantum dynamics,

Phys. Rev. Research 5, 023027 (2023).

- [88] C. Bertoni, J. Haferkamp, M. Hinsche, M. Ioannou, J. Eisert, and H. Pashayan, Shallow shadows: Expectation estimation using low-depth random Clifford circuits, Phys. Rev. Lett. **132**, "in press" (2024).
- [89] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar, X.-Z. Luo, B. Nash, X. Gao, B. Barak, E. Farhi, S. Sachdev, N. Gemelke, L. Zhou, S. Choi, H. Pichler, S.-T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin, Quantum optimization of maximum independent set using Rydberg atom arrays, Science **376**, 1209 (2022).
- [90] B. F. Schiffer, D. S. Wild, N. Maskara, M. Cain, M. D. Lukin, and R. Samajdar, Circumventing superexponential runtimes for hard instances of quantum adiabatic

optimization, Phys. Rev. Res. 6, 013271 (2024).

- [91] M. Tikhanovskaya, S. Sachdev, and R. Samajdar, Equilibrium dynamics of infinite-range quantum spin glasses in a field, PRX Quantum 5, 020313 (2024).
- [92] M. Troyer and U.-J. Wiese, Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations, Phys. Rev. Lett. 94, 170201 (2005).
- [93] D. Hangleiter, I. Roth, D. Nagaj, and J. Eisert, Easing the Monte Carlo sign problem, Sci. Adv. 6, eabb8341 (2020).
- [94] J. Gray, Quimb: A Python Package for Quantum Information and Many-Body Calculations, J. Open Source Softw. 3, 819 (2018).

Appendix

CONTENTS

I.	Introduction	1
II.	Variational tensor network tomography	2
III.	Results A. Learning the surface code B. Learning a Rydberg quantum spin liquid	4 4 5
IV.	Discussion and outlook	7
V.	Code and data availability	8
	Acknowledgments	8
	References	8
А.	Relation to previous works A.1. Potential advantages and limitations	13 14
В.	 Asymptotic MLE error via MPS manifold B.1. Matrix product states B.2. Differential geometry B.3. Projective Hilbert space as a manifold B.4. Manifold of matrix product states B.5. Maximum likelihood estimation on manifold B.6. Main result: A concentration inequality for infidelity B.7. Discussions and interpretations 	15 17 18 21 22 23 25 31
C.	Random- XZ classical shadow tomography C.1. Performing random- XZ classical shadow tomography C.2. Derivation of tools for random- XZ classical shadows	32 32 33
D.	Tomographic completeness of randomized XZ measurements D.1. Discussion of results D.2. Proof of Theorem 1	37 37 38
E.	Loss function E.1. Maximum likelihood estimation: Randomized Kullback–Leibler divergence E.2. Regularization	$\begin{array}{c} 40\\ 40\\ 41 \end{array}$
F.	Dataset and state generation with DMRG F.1. Perturbed surface code F.2. Ruby-lattice quantum spin liquid	43 43 43
G.	Training details G.1. Optimization scheme via gradient descent G.2. Model selection	45 45 45

APPENDIX A: Relation to previous works

This appendix aims to provide a comparison of our method to existing quantum state tomography techniques. We highlight the unique challenges and opportunities in characterizing complex quantum systems, particularly twodimensional topological states. In addition to the references summarized below, we direct interested readers to Ref. [66] for a comprehensive review of quantum state tomography. Finally, having discussed advantages of our method for learning quasi-2d quantum states with topological order, we summarize the limitations faced by our method discussed in the main text.

- 1. Matrix product state (MPS) tomography [18–20, 67] is an effective method for reconstructing injective onedimensional quantum states with short-range correlations associated with a local parent Hamiltonian, by finding the best MPS consistent with local reduced density matrices from randomized measurements. In principle, these reduced density matrices can be estimated from Pauli measurements using the techniques of classical shadows [68]. Its application to 2d topological states, however, is challenging: first of all, the size of the reduced density matrices required for uniquely determining a quantum state in 2d is unclear, and secondly, due to nontrivial ground-state degeneracies of topological states, they cannot be uniquely determined by local operators and thus, cannot be efficiently determined with this method.
- 2. Variational MPS tomography [26–28] extends MPS tomography to one-dimensional states without the injectivity constraint. It uses a variational approach to optimize the MPS representation through maximum likelihood estimation (MLE). Prior work also considered using matrix product operator (MPO) for quantum process tomography [29]. Our work extends this protocol to 2d states using a "snake path", commonly adopted in cylindrical DMRG calculations for 2d systems. We further enhance our method by incorporating regularization techniques inspired by classical shadow tomography, combining the benefits of both approaches to better characterize two-dimensional quantum systems. We refer the reader to Sec. II for the details of our protocol.
- 3. Neural network quantum state tomography (NNQST) [23] utilizes neural networks as variational ansätze within the MLE framework to model quantum states. Although particularly successful for states with fully positive amplitudes (for example, ground states of frustration-free stoquastic Hamiltonians [32]), this method requires careful consideration when the target states have nontrivial signs and phases. Using measurement outcomes from different bases is therefore essential—a requirement directly addressed in MPS-based approaches, which can inherently relate measurements post local unitary transformations to learn relative phases. In this context, it may be useful to explore a most recently proposed architecture that can accurately represent quantum spin liquids by incorporating physical symmetries [69].
- 4. Generative modeling for density matrix reconstruction [24] extends the use of neural networks to mixed-state tomography. This approach uses single-qubit positive operator-valued measure (POVM) measurements for information-complete reconstructions. While generative models can theoretically model measurement probability distributions across an exponentially large outcome space, practical challenges arise in ensuring convergence to the accurate density matrix, particularly given the sparse nature of such distributions in large systems. It is unclear how effective this approach is [25] for learning entangled quantum states such as the topological ordered states considered in this work.
- 5. Classical shadow tomography [25, 30] efficiently estimates properties of quantum states by measuring them post application of random unitary rotations, followed by post-processing. This technique has demonstrated that a number of samples independent of the system size is sufficient for estimating the fidelity of specific projectors, such as ground states of the toric code Hamiltonian [25], given prior knowledge of the target projector and the utilization of global Clifford measurements. A significant extension of this method [15, 30] employed a dimensional reduction technique along with a nonlinear kernel for post-processing to distinguish toric code states from trivial states using only random Pauli measurements. This unsupervised learning task, while innovative, requires the assembly of numerous distinct quantum states for shadow tomography, presenting a considerable experimental demand. Our method, utilizing random-XZ measurements, sidesteps these requirements, offering a more experimentally feasible approach for a broader range of quantum systems.
- 6. Shadow-NNQST [70] combines a neural-network ansatz used in NNQST with "regularization" of fidelity estimations from classical shadow tomography. This regularization step necessitates global Clifford measurements for efficient fidelity estimation [25], and an arbitrary Clifford unitary on n qubits has nontrivial depth. In contrast, our approach leverages random-XZ measurements, avoiding the need for such specific, experimentally costly measurement protocols. More recently, a quantity termed the "shadow overlap" has demonstrated faithful certification of fidelity requiring only polynomial sample complexity for a large set of quantum states [42],

which could significantly improve the training process of NNQST or our variational tensor network tomography. Notably, this technique also only uses single-qubit measurements that are experimentally feasible.

- 7. Learning an unknown stabilizer state can be carried out efficiently using single-copy measurements [71]. Efficient learning of stabilizer states through single-copy measurements offers a direct method for those topological states that are also stabilizer states. Our protocol expands its applicability to nonstabilizer states amenable to representation as an MPS, which allows for broader quantum state characterization beyond the stabilizer formalism. A recent work has considered combining MPS with the stabilizer formalism [72], which allows for the simulations of states beyond the scope of either method individually—this framework could be also explored in the current context of quantum state tomography.
- 8. *Predicting specific properties of quantum states.* If one knows the targeted properties that one wants to learn, there are a plethora of alternative methods, of which we list the most relevant approaches below. We will not try to compare our protocol to all of these directly, since such methods differ in their flexibility with regard to experimental measurements as well as the number of samples required to achieve accurate predictions.
 - (a) Rényi entropy:
 - A multicopy approach [73], based on measuring two copies of a quantum state, has been demonstrated experimentally [12, 74]. However, this requires the preparation of multiple identical copies of a state while demanding advanced control over the quantum system to ensure the fidelity of the copies.
 - Randomized measurements [13, 14, 75] leverage the statistical properties of the outcomes to estimate the Rényi entropy without needing multiple copies of a system. This is both similar to and encompasses classical shadow tomography.
 - (b) Von Neumann entropy:
 - Entanglement Hamiltonian tomography [76] assumes that the reduced density matrices of a subsystem can be approximated by an ansatz involving local operators in the Hamiltonian. By learning the coefficients of this ansatz, the technique aims to reconstruct the entanglement Hamiltonian and thereby predict the von Neumann entropy.
 - (c) Topological invariants:
 - Direct tomography can be applied to small subsystem sizes, from which the topological entanglement entropy can be extracted [55, 56]; this approach was demonstrated experimentally for a surface code [8]. Nevertheless, the applicability of this technique for general systems remains unclear, particularly with regard to the required size of the subsystem for accurate measurement of the entanglement entropy.
 - Randomized measurements can also be used to predict topological invariants of systems such as *symmetry protected topological* (SPT) phases [77] or fractional quantum Hall states [78].

A.1. Potential advantages and limitations

Our variational tensor network tomography, aimed at diagnostics and recovery of quantum states, offers potential for extracting various properties such as the Von Neumann entanglement entropy, difficult to learn using other techniques. This method combines maximum likelihood estimation of the quantum state with randomized measurements and is detailed in previous works such as [26, 28, 29]. We demonstrate this approach using simplified XZmeasurements, which proves sufficient for the two-dimensional topological quantum states considered, with complex long-range entanglement. Such quantum states are represented by non-injective MPS, and can not be learned via MPS tomography [18, 19].

One significant limitation is that the target state must admit an efficient tensor network representation. In numerical experiments, we use an MPS ansatz, recognizing that while it effectively approximates one-dimensional states, it may not capture the full complexity of two-dimensional systems. This works sets the stage for future exploration of more sophisticated tensor networks. We also discuss the adaptation of statistical properties of MLE for quantum tomography, outlining both its potential and the challenges posed by the gauge freedom in MPS representations.

Another limitation is that real laboratory states can be mixed, which requires using generalizations of tensor networks such as locally purified matrix product operators [79]. Additionally, one generally do not have access to the ground truth fidelity with the lab state. This necessitates methods like cross-validation as discussed in Appendix G or efficient fidelity certifications [42] to validate the trained model against empirical data.

APPENDIX B: Asymptotic MLE error via MPS manifold

Our variational algorithm in the framework of maximum likelihood estimation (MLE) aims to find the MPS representation $|\psi(A)\rangle$ that is most consistent with measurements of the target state vector $|\phi\rangle$. Intuitively, we expect that the accuracy of the optimal MLE solution improves with more samples. The goal of this appendix is to analyze the accuracy of such a solution as a function of the number of samples N. More precisely, we provide a probabilistic upper bound of the infidelity 1 - F by deriving a concentration inequality using the statistical properties of MLE. First, we summarize our main finding of the probabilistic bound in the theorem below:

Theorem 2 (Probabilistic bound). Let $|\psi(\hat{A}^{(N)})\rangle$ be the MLE solution with a maximum bond dimension χ_{\max} for an n-qubit target state vector $|\phi\rangle$. With probability greater than $1 - \delta$, the infidelity is upper bounded asymptotically in the limit of large number of samples N as

$$\mathbb{P}\left[1 - |\langle \phi | \psi(\hat{A}^{(N)}) \rangle| \le \epsilon(N)\right] \ge 1 - \delta,\tag{B1}$$

where $\epsilon(N) = O\left(\sqrt{\frac{n\chi_{\max}^2}{N\delta}}\right).$

We provide a detailed and rigorous proof of Theorem 2 concerning the *optimal* MLE solution, focusing on conditions under which it represents a global minimum. Achieving such a global minimum may not always be practical, and our theorem does not address the complexity of reaching such a solution. Moreover, the theorem is specifically applicable in the *asymptotic limit* as $N \to \infty$. For the validity of our theorem, we assume that the target state is within the variational space used for MLE, i.e., it is an MPS of known bond dimension. In practice this condition might not be satisfied.

Ideally, one would like to prove a theorem that applies to our variational MPS tomography *in practice*. There are a wide variety of technical obstacles to achieving this; thus, our goal in this work is to take the first steps toward this direction and provide a thorough introduction of the required mathematical tools. In this process, we (a) present an initial theorem under a specific set of assumptions, and (b) offer pedagogical insights into the technical tools and conceptual challenges involved in proving our theorem and potentially more comprehensive statements. The success of our numerical results, combined with these theoretical explorations, provides both motivation and foundational insights for developing more specialized analyses.

To build some intuition, we first review the maximum likelihood estimation (MLE) within traditional settings. It turns out that these results cannot be directly applied to our context due to the gauge freedom of MPS. To generalize the MLE analysis to our MPS-estimation setting, we first elaborate on the gaps that we are filling in and how our statement differs from conventional results.

Let us begin with a standard MLE problem, where the goal is to estimate a single parameter. Considering a sequence of N independent and identically distributed (i.i.d.) samples $(x_i)_{i=1}^N$ drawn from a target Gaussian distribution \mathcal{P}_{θ_0} with mean θ_0 and a known variance σ^2 ,

$$\mathcal{P}_{\theta_0}(x) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left[-\frac{(x-\theta_0)^2}{2\sigma^2}\right].$$
(B2)

Since the variance is known, our goal is to find the estimator $\hat{\theta}^{(N)}$ of the mean such that the likelihood of the observed samples is maximized. The *negative log-likelihood* (NLL) of observing the N samples is

$$\mathcal{L}(\theta) = -\frac{1}{N} \sum_{i=1}^{N} \log\left(\exp\left[-\frac{(x_i - \theta)^2}{2\sigma^2}\right]\right) = \frac{1}{N} \sum_{i=1}^{N} \frac{(x_i - \theta)^2}{2\sigma^2}.$$
(B3)

For this simple case of using a Gaussian distribution as our model prior, the minimization of the NLL \mathcal{L} is the same as minimizing the least-squared error. Differentiating \mathcal{L} with respect to θ we readily find the extremum $\hat{\theta}^{(N)} = \frac{1}{N} \sum_{i} x_{i}$, which is the well-known empirical mean.

Note that the estimator $\hat{\theta}^{(N)}$ is a random variable, as it depends on the random samples $(x_i)_{i=1}^N$. Hence, we can describe the probability distribution of $\hat{\theta}^{(N)}$. Under certain conditions this distribution tends to the normal distribution as the number of samples goes to infinity. This property is known as *asymptotic normality* and has been studied rigorously [33]. We will state the precise technical conditions for asymptotic normality later (in the context of MLE on manifolds) in Lemma 3, but emphasize here already that, in general, they entail two crucial requirements: 1) the MLE solution is consistent such that θ_0 is the unique solution—the unique minimizer of \mathcal{L} in the infinite-sample

$$\sqrt{N}(\hat{\theta}^{(N)} - \theta_0) \xrightarrow{\mathrm{d.}} \mathcal{N}(0, \sigma^2).$$
(B4)

Here, $\stackrel{d}{\to}$ denotes convergence in distribution, i.e., as N increases, the sequence of distributions of the random variables $\sqrt{N}(\hat{\theta}^{(N)} - \theta_0)$ for growing N approximates the normal distribution increasingly well. With the normality condition, we can write down a concentration inequality to probabilistically upper bound the error $|\hat{\theta}^{(N)} - \theta_0|$. As N increases, the error reduces with a rate of $1/\sqrt{N}$ so that in the infinite-N limit, the convergence $\hat{\theta}^{(N)} \xrightarrow{P} \theta_0$ is guaranteed by the consistency condition. To illustrate how this condition could fail, let us consider a modified Gaussian distribution periodic in the parameter

$$\mathcal{P}_{\theta_0} = \frac{1}{\sqrt{2\pi\sigma}} \exp\left[-\frac{(x-\sin(\theta_0))^2}{2\sigma^2}\right].$$
(B5)

Now, a parameter shifted by a multiples of 2π still represents the same distribution $\mathcal{P}_{\theta_0} = \mathcal{P}_{\theta_0+2\pi}$. Therefore, $\hat{\theta}^{(N)}$ need not converge to θ_0 . In physics, the phenomenon where multiple parameters describe the same function is known as *gauge redundancy*. In this simple scenario, we can explicitly account for the periodicity in the optimal estimator, or equivalently impose "gauge fixing" conditions to select a unique representative, by restricting the parameter domain to $[0, 2\pi)$. However, for distributions parametrized by MPS, it is challenging in practice to fix the gauge degrees of freedom consistently. Specifically, transforming MPS to canonical forms using singular value decomposition only partially fixes the gauge. Additional unitary-gauge degrees of freedom remain in such a decomposition. Moreover, degeneracies in the singular values introduce redundancy within the subspace spanned by degenerate eigenvectors.

Consequently, the MPS tensors that minimize our loss function do not satisfy the consistency condition. Nevertheless, if our measurement scheme is tomographically complete [see Appendix D], the physical state itself meets the consistency condition. This implies that even though multiple MPS tensor configurations may solve the MLE problem, they all correspond to the same physical state. Thus, our goal is to investigate the asymptotic properties of MLE estimators across the set of physical states represented by MPS. These states, represented by a well-defined class of MPS described further in this section, form a manifold. Therefore, our analysis aims to understand the asymptotic properties of MLE estimators on this manifold. Such a generalized asymptotic normality on a manifold has been established in previous work [36]. What remains is to relate the normality condition to infidelity, which we discuss by lifting the distance on a manifold to the state distance in the embedding Hilbert space via Lemma 4.

In the rest of this self-contained appendix, we provide the necessary background on MLE and differential geometry for readers interested in technical details and a proof of the theorem, provided in Appendix B.6. The first five sections review existing results, which we then use in our proof. First, we briefly review notations for MPS commonly used for the rest of the section in Appendix B.1. We provide a minimal list of technical definitions in differential geometry in Appendix B.2, which can be skipped for readers already familiar with the topic. Equipped with this knowledge of differential geometry, we review how the projective Hilbert space can be viewed as a manifold in Appendix B.3. Then, given a subspace of the projective Hilbert space defined by an MPS, we review properties of the projective MPS manifold in Appendix B.4. Finally, our main result relies on a theorem establishing asymptotic normality of the MLE estimator on a manifold, which we restate in Appendix B.5. We conclude this section with a summary of our main result as well as our interpretation of its relations to our numerical results in Appendix B.7.

B.1. Matrix product states

In this subsection, we introduce the notations and definitions used for *matrix product states* (MPS) throughout the rest of the subsection, following Ref. [37]. To represent a quantum state across n sites with physical dimension d, we define the MPS tensors over complex numbers as

$$\mathbb{A}_{\chi} = \bigoplus_{i=1}^{n} \mathbb{C}^{\chi_i \times d \times \chi_{i+1}} \cong \mathbb{C}^l.$$
(B6)

Here, $\chi = (\chi_1, \ldots, \chi_{n+1})$ represents the bond dimensions and sums to a total tensor dimension of $l = d \sum_{i=1}^n \chi_i \chi_{i+1}$. Specifically, the tensor at the *i*-th site is denoted by $A^{[i]}$. The tensor elements are $A_{\alpha,\beta}^{[i]\delta}$, with $1 \le \alpha \le \chi_i$, $1 \le \beta \le \chi_{i+1}$, and $1 \le \delta \le d$. For this work, we limit the physical dimension to d = 2 (representing qubits) and consider open boundary conditions such that $\chi_1 = \chi_{n+1} = 1$. We define an MPS via a map from the space of MPS tensors \mathbb{A}_{χ} to the Hilbert space $\mathcal{H} \cong \mathbb{C}^{d^n}$ of *n* sites.

$$\psi_{\chi} : \mathbb{A}_{\chi} \longrightarrow \mathcal{H}, \tag{B7}$$

$$A \longmapsto \sum_{b \in \{0,1\}^n} A^{[1]b_1} \cdots A^{[n]b_n} |b\rangle.$$
(B8)

The set $\{|b\rangle\}$ represents the computational basis, where the bit-string $b \in \{0,1\}^n$, unless specified otherwise. Note that in the main text, we also use the notation $|\psi(A)\rangle$ to refer to the physical state, which is inferred from the context.

Physically distinct quantum states are identified up to a global phase and normalization factor, and are unique elements in the *projective Hilbert space* P(H). The elements in this space are equivalence classes, often called "rays" of vectors in H.

Definition 1 (Projective Hilbert space [37]). Associated with a complex Hilbert space \mathcal{H} , the projective Hilbert space is defined as

$$\mathsf{P}(\mathcal{H}) := \mathcal{H}/\operatorname{GL}(1,\mathbb{C}),\tag{B9}$$

which contains all rays represented by the equivalence classes $[|\psi\rangle]$. Two nonzero vectors $|\psi\rangle, |\phi\rangle \in \mathcal{H}$ are equivalent if and only if $|\psi\rangle = \lambda |\phi\rangle$ for some $\lambda \in GL(1, \mathbb{C})$, the group of nonzero complex numbers. For simplicity, we use the notation $[\psi]$ to denote $[|\psi\rangle]$.

We can restrict MPS to the projective Hilbert space $P(\mathcal{H})$ by defining a new mapping $\tilde{\psi}_{\chi}$. This projection maps an MPS tensor onto a ray, effectively associating each MPS with a unique physical state representation. While it is possible to consider the manifold of MPS without this projection [37], using the projective MPS simplifies our MLE analysis by removing ambiguious normalization or phase factors that are not physically observable.

Definition 2 (Projective MPS [37] (Definition 13)). The map $\tilde{\psi}_{\chi}$ from the space of MPS tensors \mathbb{A}_{χ} to the projective Hilbert space $\mathsf{P}(\mathcal{H})$ is

$$\widetilde{\psi}_{\chi} : \mathbb{A}_{\chi} \longrightarrow \mathsf{P}(\mathcal{H}),$$
(B10)

$$A \longmapsto [\psi_{\chi}(A)], \tag{B11}$$

where $[\psi_{\chi}(A)]$ denotes the ray representing the equivalence class of the state vector $\psi_{\chi}(A)$ under rescaling by nonzero complex numbers.

Throughout the rest of this section, what we refer to by $|\psi\rangle$ or $[\psi]$ should be understood from the context as denoting either a normalized state vector or a ray in the projective Hilbert space, in distinction to the projective mapping $\tilde{\psi}_{\chi}$. Note that the image $\tilde{\psi}_{\chi}(\mathbb{A}_{\chi})$ represents only a subspace of $\mathsf{P}(\mathcal{H})$ and does not possess the structure of a vector space ². In the following sections, we will provide constraints on the domain of MPS tensors such that their image under $\tilde{\psi}_{\chi}$ (and ψ_{χ}) forms a differentiable manifold within $\mathsf{P}(\mathcal{H})$. Equipping the set of MPS with such an additional structure, we will be able to rigorously analyze the accuracy of the MLE solution on the space of a manifold, as detailed in Appendix B.5.

² For instance, consider the sum of two MPS of bond dimension χ as $(|\psi_{\chi}\rangle + |\phi_{\chi}\rangle)/\sqrt{2}$. This new state has bond dimension up to 2χ , which is beyond the original space defined for bond dimension χ .

B.2. Differential geometry

This subsection introduces the essential concepts of differential geometry required to navigate through the subsequent discussions in Appendix B.3, B.4, and B.5. Assuming no prior expertise in differential geometry, we aim to outline here the key definitions and provide intuitive explanations for these concepts that can be found in standard textbooks [80]. While this section is in principle self-consistent, we highly encourage readers who would like to fully understand this material to consult Refs. [80, 81] for more context. Readers who are well familiar with differential geometry of Riemannian manifolds may choose to skip most of this section and only briefly review Def. 8 and Fact 1, which introduces terminologies used specifically in the context of our theorem.

Let us begin by briefly recalling the concept of an r-dimensional differentiable manifold M. We use $\theta \in M$ to denote elements on the manifold, where, in the subsequent sections, θ parametrizes some underlying probability distribution. Formally, any local neighbourhood $M_i \subset M$ can be related to Euclidean space \mathbb{R}^r . This relationship is established through an injective map $\varphi_i : M_i \to \mathbb{R}^r$, known as a *chart*. A chart provides a local coordinate system near each point in its neighbourhood. Intuitively, the manifold is covered by a collection of neighbourhoods M_i such that the union $\bigcup_{i \in I} M_i = M$ represents the entire manifold.

The collection of charts, known as an *atlas*, allows us to study functions defined on M using the usual machinery from calculus defined on \mathbb{R}^r . Moreover, for M to be called differentiable, any composition $\varphi_i^{-1} \circ \varphi_j$ defined on any two overlapping patches $M_i \cap M_j \neq \emptyset$ must be differentiable. This allows us to define the tangent vectors X at any point $\theta \in M$. In alignment with Ref. [37], we adopt a geometric perspective of tangent vectors by representing them as column vectors $(X_1, \dots, X_r)^{\top}$. The components of the vector X depend on the local coordinates in \mathbb{R}^r , defined by the chosen chart φ .³ The space of all tangent vectors at θ is called the tangent space $T_{\theta}M$ and it is isomorphic to \mathbb{R}^r . Consequently, each tangent vector can be expressed as $X = \sum_{a=1}^r X_a e_a$ where $\{e_a\}$ is a basis of $T_{\theta}M$.

To compute the lengths and angles of the tangent vectors, it is essential to define an inner product for each tangent space, called *metric*. In general, a metric provides an inner product for each tangent space $T_{\theta}M$ which explicitly depends on the point θ , i.e., as one moves across the manifold, the inner product *changes*. When the inner product is real-valued and its change is gradual, such that it is differentiable (as precisely defined below), the metric is called a *Riemannian* metric. This allows us to track the change of lengths and angles between tangent vectors as one moves across the manifold so that we can define distances between points on the manifold.

Definition 3 (Riemannian metric [80]). A Riemannian metric g is a smooth family of inner products $\{g_{\theta} | \theta \in M\}$ on the tangent spaces of M. The metric g_{θ} at each point θ , defined by

$$g_{\theta}: \mathrm{T}_{\theta}M \times \mathrm{T}_{\theta}M \longrightarrow \mathbb{R}, \tag{B12}$$

is bilinear and positive definite. Represented in the basis $\{e_a\}$, its matrix elements $(g_{\theta})_{a,b} = g_{\theta}(e_a, e_b)$ vary smoothly with respect to θ .

In the discussion that follows, we will use g to denote the metric g_{θ} , as it is understood that the metric is defined for the tangent space at each point $\theta \in M$. Additionally, we adopt the notation $\langle \bullet | \bullet \rangle_g$ to represent the inner product of any two tangent vectors, which is computed according to the metric g.

Definition 4 (Riemannian manifold). A differentiable manifold M together with a Riemannian metric g forms a Riemannian manifold (M, g).

To be able to discuss curves connecting two points on a manifold (Def. 8), we introduce the concept of a Riemannian connection, which generalizes the notion of the covariant derivative.

Definition 5 (Riemannian connection [80]). Let $\Gamma(TM)$ be the set of differentiable vector fields $\theta \mapsto X \in T_{\theta}M$ on a Riemannian manifold (M,g). A map $\nabla : \Gamma(TM) \times \Gamma(TM) \to \Gamma(TM)$ is called a Riemannian connection, typically expressed as $\nabla_X Y$ representing $\nabla(X,Y)$. It holds that

- 1. $\nabla_{fX+hY}Z = f\nabla_X Z + h\nabla_Y Z$,
- 2. $\nabla_X(Y+Z) = \nabla_X Y + \nabla_X Z$,

³ Tangent vectors to M at points θ can be naturally seen as directional derivatives. For a smooth function $f : M \to \mathbb{R}$ defined on the manifold, if X is a tangent vector to the manifold M at θ , then the directional derivative (defined in Def. 6) is given by $D_X(f) = X(f)$. For what follows, it is convenient to directly refer to tangent vectors.

3.
$$\nabla_X f Z = f \nabla_X Z + (\mathbf{D}_X f) Z$$
,

4.
$$D_X g(Y, Z) = g(\nabla_X Y, Z) + g(Y, \nabla_X Z),$$

5.
$$\nabla_X Y - \nabla_Y X - [X, Y] = 0,$$

for all $X, Y, Z \in \Gamma(TM)$ and all smooth functions $f, h : M \to \mathbb{R}$ defined on the manifold. $[\bullet, \bullet]$ is the Lie bracket defined by $[X, Y](f) := D_X D_Y f - D_Y D_X f$.

Here we use $D_X f$ to denote the directional derivative of a scalar-valued function f, also commonly denoted by X(f) in the literature. The directional derivative has a direct correspondence with the tangent vectors, and can be formally defined as follows:

Definition 6 (Directional derivative). Let M be a differentiable manifold, $\theta \in M$ and $X \in T_{\theta}M$. Given a differentiable function $f: M \to \mathbb{R}$, the directional derivative $D_X f$ (also denoted X(f)) at θ is defined as

$$D_X f \Big|_{\theta} = \frac{d}{dt} f(\gamma(t)) \Big|_{t=0}, \tag{B13}$$

where $\gamma : \mathbb{R} \to M$ is a differentiable curve such that $\gamma(0) = \theta$ and the velocity $\dot{\gamma}(0) = X$.

The Riemannian connection is sometimes also referred to as the *Levi-Civita connection* and for each Riemannian manifold, there exists exactly one Riemannian connection [80]. As mentioned above and as the notation suggests, the Riemannian connection provides us with a notion of directional (or covariant, to be more precise) derivatives on a manifold. This allows us to describe how vectors change as they are moved along a curve on the manifold. Specifically, we can now define a notion of *parallelism* along curves, a concept that is very intuitive in Euclidean space but much less so in curved spaces.

Definition 7 (Parallel transport). Let $\alpha : [0,1] \to M$ be a smooth curve on the Riemannian manifold (M,g) and $\dot{\alpha}(t)$ be its velocity vector in the tangent space $T_{\alpha(t)}M$. The parallel transport of a tangent vector $X_0 \in T_{\alpha(0)}M$ along the curve α under a Riemannian connection ∇ is defined by the equation

$$\nabla_{\dot{\alpha}(t)}X(t) = 0,\tag{B14}$$

where $X(t) \in T_{\alpha(t)}M$ for all $t \in [0,1]$ and $X(0) = X_0$. The vector field X is said to be parallel along α if it satisfies the above condition.

A special case of parallel vector fields is given by *geodesic curves*. The velocity vector $\dot{\gamma}$ of a geodesic curve $\gamma : I \to M$ (with $I \subseteq \mathbb{R}$ some parameter interval) is parallel along γ itself. For simplicity, we will take the unit interval I = [0, 1]because a curve with any other interval can always be reparameterized. The notion of a geodesic extends the concept of a straight line from Euclidean space to curved spaces like manifolds. In general, there can be multiple geodesic curves connecting two points on a manifold⁴. For our purpose, we only consider the curves that represent the *shortest* path between two points on the manifold, parametrized by arc length on the unit interval.

Definition 8 (Shortest connecting geodesic). Let (M, g) be a Riemannian manifold and $\nabla : \Gamma(TM) \times \Gamma(TM) \rightarrow \Gamma(TM)$ the associated Riemannian connection. Given two points $\theta_1, \theta_2 \in W$ where $W \subset M$ is a sufficiently small neighborhood, the curve

$$\gamma: [0,1] \longrightarrow W, \tag{B15}$$

$$t \mapsto \gamma(t),$$
 (B16)

is the shortest connecting geodesic of θ_1 and θ_2 if $\gamma(0) = \theta_1$ and $\gamma(1) = \theta_2$, and for all $t \in [0, 1]$ the tangent vectors $\dot{\gamma}$ satisfy

$$\nabla_{\dot{\gamma}}\dot{\gamma} = 0. \tag{B17}$$

This implies that $\dot{\gamma}$ is parallel transported along the geodesic curve γ .

⁴ For example, the great circle on the sphere S^2 is a geodesic curve. Two points on the great circle separate the circle into two arcs. The shorter arc is the shortest connecting geodesic, and the longer arc is the longer connecting geodesic.

Note, that in general, one can loosen the condition for geodesics to $\nabla_{\dot{\gamma}}\dot{\gamma} = \lambda(t)\dot{\gamma}$ where λ is some scalar function and γ is required to be regular ($\dot{\gamma}(t) \neq 0$ for all t). However, by setting $\nabla_{\dot{\gamma}}\dot{\gamma} = 0$, we enforce that the velocity vector $\dot{\gamma}$ of the geodesic cannot change in the direction of the curve, so that the velocity vector has a constant length throughout the curve. This is formally stated in Fact 1 below.

Fact 1 (Velocity vector). Consider a geodesic γ (defined in Def. 8) on a Riemannian manifold (M, g) of dimension r. If $\{e_a(t)\}_{a=1}^r$ represents the set of basis vectors parallel transported along γ , the velocity vector $\dot{\gamma}$ can be expressed as

$$\dot{\gamma} = \sum_{a=1}^{r} \Delta_a e_a(t), \tag{B18}$$

where the decomposition coefficients $\{\Delta_a\}_{a=1}^r$ are constant due to the constraint $\nabla_{\dot{\gamma}}\dot{\gamma} = 0$. Moreover, the length of the velocity vector is given by

$$\begin{aligned} \|\dot{\gamma}(t)\| &:= \sqrt{\langle \dot{\gamma} | \dot{\gamma} \rangle_g} \\ &= \sqrt{\sum_{a,b} \Delta_a \, g_{a,b} \, \Delta_b}, \end{aligned} \tag{B19}$$

which is a constant for all $t \in [0, 1]$, due to the aforementioned constraint $\nabla_{\dot{\gamma}}\dot{\gamma} = 0$ representing the constant speed at which γ traverses the manifold. The metric tensor components $g_{a,b}(t) = \langle e_a(t)|e_b(t)\rangle_g$ in the basis $\{e_a(t)\}_{a=1}^r$ along the curve are constant as well.

Proof. We begin by expanding the condition for the geodesics for a given coordinate system

$$0 = \nabla_{\dot{\gamma}}\dot{\gamma}$$
(B20)
$$= \sum_{a} \nabla_{\dot{\gamma}} \Delta_{a} e_{a}$$
$$= \sum_{a} \Delta_{a} \nabla_{\dot{\gamma}} e_{a} + \sum_{a} (D_{\dot{\gamma}} \Delta_{a}) e_{a}.$$

Using the definition of the parallel transport of the basis vectors, $\nabla_{\dot{\gamma}} e_a(t) = 0$, we conclude that $\sum_a (D_{\dot{\gamma}} \Delta_i) e_a = 0$. Since the basis vectors $\{e_a\}$ are linearly independent, the coefficients $D_{\dot{\gamma}} \Delta_a$ must all be zero individually, hence Δ_a is a constant in t for all a. Moreover, any Riemannian connection preserves the metric (Property 4 in Def. 5) and thus, we have

$$D_{\dot{\gamma}}g_{a,b}(t) = \langle \nabla_{\dot{\gamma}}e_a(t)|e_b(t)\rangle + \langle e_a(t)|\nabla_{\dot{\gamma}}e_b(t)\rangle, \tag{B21}$$

which is zero since $\nabla_{\dot{\gamma}} e_a(t) = 0$ for all $a = 1, \ldots, r$, by definition.

Given a connecting geodesic curve on a Riemannian manifold, the length of the curve can serve as a distance measure between the two end points θ_1 and θ_2 . The length of the curve depends on the metric along the path. Let γ be a connecting geodesic as in Def. 8. The arc length $L_{\gamma}(0, 1)$ of any segment of the curve between $\gamma(0)$ and $\gamma(1)$ is given by

$$L_{\gamma}(0,1) = \int_{0}^{1} \mathrm{d}t \sqrt{\langle \dot{\gamma} | \dot{\gamma} \rangle_{g}}$$

$$= \sqrt{\langle \dot{\gamma} | \dot{\gamma} \rangle_{g}},$$
(B22)

which is the constant length of the velocity vector along γ .

B.3. Projective Hilbert space as a manifold

In this subsection, we apply the previously discussed concepts of differential geometry to describe quantum systems. In particular, we review how the projective Hilbert space $P(\mathcal{H})$ can be understood as a manifold [37]. These descriptions allow us to relate the length of geodesics within this manifold to the quantum state overlap—the fidelity—that we aim to analyze.

Definition 9 (Projective Hilbert space as a manifold [37]). The projective Hilbert space $P(\mathcal{H}) \simeq \mathbb{CP}^{2^n-1}$ is a differentiable manifold, whose elements correspond to vectors in the Hilbert space modulo a global phase and normalization.

Note that $\mathsf{P}(\mathcal{H})$ can be naturally seen as a complex manifold of complex dimension $2^n - 1$, although it could, at the same time, also be defined as a real manifold of dimension $2(2^n - 1)$. Following Ref. [81] (Sec. 8.2) and Ref. [37] (Sec. IIA), we can take the tangent space of the *complex* manifold $\mathsf{P}(\mathcal{H})$ at some point $[\psi] \in \mathsf{P}(\mathcal{H})$ as the real vector space $\mathsf{T}_{[\psi]}\mathsf{P}(\mathcal{H})$ of dimension $2(2^n - 1)$. To avoid confusion, we remark that it can be useful to consider the *complexification* $\mathsf{T}_{[\psi]}^{\mathbb{C}}\mathsf{P}(\mathcal{H})$ (i.e., the *complex* span of any basis of $\mathsf{T}_{[\psi]}\mathsf{P}(\mathcal{H})$) of this vector space to define a set of so-called holomorphic tangent vectors. Note, however, that $\mathsf{T}_{[\psi]}^{\mathbb{C}}\mathsf{P}(\mathcal{H})$ has complex dimension $2(2^n - 1)$ and is therefore over-parametrized.

For our result, we only rely on the manifold $\mathsf{P}(\mathcal{H})$ having the structure of a Riemannian manifold and therefore, we will only work with $\mathsf{T}_{[\psi]}\mathsf{P}(\mathcal{H})$. However, as we will see, $\mathsf{T}_{[\psi]}\mathsf{P}(\mathcal{H})$ is closely related to the Hilbert space \mathcal{H} itself. Hence, it is more natural to view the $2(2^n - 1)$ -dimensional real tangent vectors as $2^n - 1$ -dimensional complex vectors. Those correspond to vectors in \mathcal{H} that are orthogonal to $[\psi]$. We stress that these vectors are *not* elements of the complexified tangent space (which would be complex vectors of dimension $2(2^n - 1)$).

Definition 10 (Tangent space of projective Hilbert space [37]). The tangent space $T_{[\psi]}\mathsf{P}(\mathcal{H}) = \mathcal{H}/\sim$ at a base point denoted by $[\psi]$ is the space of all tangent vectors that satisfy the following equivalence relation. Let $|\psi_1\rangle, |\psi_2\rangle \in \mathcal{H}$ be two state vectors in the Hilbert space,

$$|\psi_1\rangle \sim |\psi_2\rangle :\iff |\psi_1\rangle - |\psi_2\rangle = \alpha |\psi\rangle, \tag{B23}$$

for some constant $\alpha \in \mathbb{C}$.

The equivalence relation removes one dimension from the Hilbert space so that the tangent space is isomorphic to the Euclidean complex space and the Eucliean real space of twice the dimension, i.e., $T_{[\psi]}\mathsf{P}(\mathcal{H}) \simeq \mathbb{C}^{2^n-1} \simeq \mathbb{R}^{2(2^n-1)}$.

For the purpose of our result, we only rely on the minimal structure of a Riemannian manifold. $P(\mathcal{H})$ can be equipped with the Fubini-Study metric in the context of quantum geometry. Such a choice of the Fubini-Study metric is natural because it represents the inner product between two state vectors in the Hilbert space.

Definition 11 (Fubini-Study metric [82]). Let $|\psi_1\rangle$ and $|\psi_2\rangle$ be two tangent vectors in $T_{[\psi]}\mathsf{P}(\mathcal{H})$ and $|\psi\rangle$ be the normalized representation of $[\psi]$. The Fubini-Study metric is defined as

 $g: \mathrm{T}_{[\psi]}\mathsf{P}(\mathcal{H}) \times \mathrm{T}_{[\psi]}\mathsf{P}(\mathcal{H}) \longrightarrow \mathbb{R}, \tag{B24}$

$$(|\psi_1\rangle, |\psi_2\rangle) \longmapsto \operatorname{Re}\left[\langle \psi_1 | \psi_2 \rangle - \langle \psi_1 | \psi \rangle \langle \psi | \psi_2 \rangle\right]. \tag{B25}$$

Note that here we have defined the Fubini-Study metric by only taking the real part of what is defined in Ref. [37]. This is a more common definition used in other literature [82], and suffices for our purpose of defining the length of a geodesic.

In the following, we denote the Fubini-Study metric by g, since the choice of metric for $\mathsf{P}(\mathcal{H})$ is unambiguous. Using the Fubini-Study metric, we can calculate the distance between two rays in projective Hilbert space.

Definition 12 (State distance defined by Fubini-Study metric in projective Hilbert space). Let $\gamma : [0,1] \rightarrow P(\mathcal{H})$ be the shortest geodesic (see Def. 8) connecting two rays $\gamma(0) = [\phi]$ and $\gamma(1) = [\psi]$ on the Riemannian manifold $(P(\mathcal{H}), g)$. The Fubini-Study distance between these two quantum states is defined as the geodesic curve length

$$d_{\rm FS}([\phi], [\psi]) := \int_0^1 \mathrm{d}t \sqrt{\langle \dot{\gamma} | \dot{\gamma} \rangle_g}.$$
 (B26)

Note that this Fubini-Study distance is known to be related to the angle between the two states $[\phi]$ and $[\psi]$. This is often stated without any proof, but rather the intuitive motivation is that the shortest path in projective space corresponds to the arc length of the enclosed segment of the unit circle containing $[\phi]$ and $[\psi]$. Since we define the state distance via the curve length of the *shortest* connecting geodesic, we state an explicit expression of the Fubini-Study distance below.

Lemma 1 (Fubini-Study distance). The Fubini-Study distance can be calculated as

$$d_{\rm FS}([\phi], [\psi]) = \arccos |\langle \phi | \psi \rangle|. \tag{B27}$$

Ref. [83] (Sec. 4.5) contains a proof by defining the Fubini-Study distance as given by Eq. (B27) and showing that the Fubini-Study metric in in Eq. (B24) is consistent.

B.4. Manifold of matrix product states

Building on the notations established for MPS in Appendix B.1 and the concepts from differential geometry outlined in Appendix B.2 and B.3, we now proceed to introduce the variational set of projective MPS, as described in Ref. [37]. It turns out that the variational set $\psi_{\chi}(\mathbb{A}_{\chi})$ (with \mathbb{A}_{χ} defined in Eq. (B6)) does not have the structure of a manifold. In essence, points in \mathbb{A}_{χ} where one of the Schmidt ranks is lower than specified by the bond dimensions $\chi_1, \ldots, \chi_{n+1}$ constitute self intersections⁵ in $\psi_{\chi}(\mathbb{A}_{\chi})$. We refer to Ref. [37] for more details on this issue and only mention here that one needs to introduce additional constraints in order to define a variational set which forms a valid Riemannian manifold.

Let us begin with the MPS tensors $A^{[1]}, \ldots, A^{[n]}$ as defined in Eq. (B7) and denote the transfer matrices by $T^{[i]}_{(\alpha,\alpha'),(\beta,\beta')} = \sum_{\delta} A^{[i]\delta}_{\alpha,\beta} \bar{A}^{[i]\delta}_{\alpha',\beta'}$, where \bar{A} represents the complex conjugation of A. Furthermore, we denote the left and right virtual density matrices by

$$l^{[i]}(A) = T^{[1]} \cdots T^{[i-1]}, \tag{B30}$$

$$r^{[i]}(A) = T^{[i-1]} \cdots T^{[n]},$$
(B31)

while we define $l^{[0]}(A) = r^{[n]}(A) = 1$. As mentioned above, we need to exclude MPS tensors for which the Schmidt rank between sites *i* and *i* + 1 is not equal to χ_{i+1} , i.e., where any of the left or right density matrices are not positive definite.

Definition 13 (Full-rank MPS tensors ([37] Sec III.B Definition 10)). The subset $\mathcal{A}_{\chi} \subset \mathbb{A}_{\chi}$ of full-rank MPS tensors contains all tensors characterized by having positive-definite virtual left and right density matrices $l^{[i]}(A)$ and $r^{[i]}(A)$, at every site i:

$$\mathcal{A}_{\chi} = \{ A \in \mathbb{A}_{\chi} \mid l^{[i]}(A) > 0, \ r^{[i]}(A) > 0, \forall i \}.$$
(B32)

We can now define the set of projective full-rank MPS

$$\tilde{M} := \tilde{\psi}_{\chi}(\mathcal{A}_{\chi}) \tag{B33}$$

to represent physical states within the projective Hilbert space that are described by full-rank MPS tensors. Having removed the self intersections, this variational set has the structure of a Riemannian manifold.

$$\lim(M) = 8,\tag{B28}$$

$$\gamma_{i,j}(t) = |0,0\rangle + t|i,j\rangle,\tag{B29}$$

for any i = 0, 1, 2 and any j = 0, 1, 2. For any $t \in \mathbb{R}$, any i = 0, 1, 2, and any j = 0, 1, 2, the Schmidt rank of $\gamma_{i,j}(t)$ is at most 2 and indeed, all velocity vectors $\dot{\gamma}_{i,j}$ are linearly independent. This would require the tangent space of M to be 9-dimensional at $|0, 0\rangle$, which contradicts the fact that the manifold is 8-dimensional, as discussed above.

⁵ Technically, a self intersection at a point $\theta \in M$ is a point which has two neighborhoods M_1 and M_2 such that the images under the charts $\varphi_1(M_1)$ and $\varphi_2(M_2)$ are open (required by the definition of a differentiable manifold), but the images $\varphi_1(M_1 \cap M_2)$ and $\varphi_2(M_1 \cap M_2)$ are not open. A simple example is the figure eight, which is the set of points in \mathbb{R}^2 satisfying $y^2 = x^2 - x^4$. This is a 1-dimensional manifold if one excludes the self-intersection point (x, y) = (0, 0).

To see how this phenomenon manifests itself in the variational set of MPS $\psi_{\chi}(\mathbb{A}_{\chi})$, we can consider the example of n = 2 qudits of local dimension d = 3 and bond dimension $\chi = 2$. The (complex) dimension of the corresponding manifold $M = \psi_{\chi=2}(\mathcal{A}_{\chi=2})$ is

since the dimension of $\mathbb{A}_{\chi=2}$ (and therefore also $\mathcal{A}_{\chi=2}$) is 12 and the dimension of the structure group is 4 (see Theorem 14 in Ref. [37]). Therefore, the tangent space of M at any point has complex dimension 8.

However, if the state vector $|0,0\rangle$ —which has Schmidt rank 1—was an element of M, there would be 9 linearly independent "tangent vectors" at $|0,0\rangle$. This can be easily seen by considering the curves

Lemma 2 (Projective MPS manifold ([37] Theorem 15)). The set \tilde{M} , defined as the image of the projective MPS mapping $\tilde{\psi}_{\chi}$ over the full-rank MPS tensors \mathcal{A}_{χ} , forms a differentiable manifold within the projective Hilbert space $\mathsf{P}(\mathcal{H})$, known as the projective MPS manifold.

Note that in Ref. [37], it is demonstrated that M is actually a Kähler manifold when one endows it with the complexvalued Fubini-Study metric, which is shown to be a Hermitian metric. This introduces additional structure—studied in complex geometry—to the manifold. Since we only rely on the variational set being a well-defined Riemannian manifold, we ignore the additional structure here and only define the real-valued Fubini-Study metric in alignment with Def. 11 by inducing the Fubini-Study metric from the embedding space $P(\mathcal{H})$. The projective MPS manifold is a submanifold $\tilde{M} \subset P(\mathcal{H})$ of the projective Hilbert space, which is equipped with the Fubini-Study metric defined in Def. 11. For each point $[\psi] \in \tilde{M}$, the associated tangent space $T_{[\psi]}\tilde{M}$ is a subspace of $T_{[\psi]}P(\mathcal{H})$. As a result, one can define the induced Fubini-Study metric \tilde{g} on \tilde{M} such that for any $[\psi] \in \tilde{M}$ and any $|\phi_1\rangle, |\phi_2\rangle \in T_{[\psi]}\tilde{M}$, one has

$$\langle \phi_1 | \phi_2 \rangle_{\tilde{q}} = \langle \phi_1 | \phi_2 \rangle_q, \tag{B34}$$

as done in Sec. IIIE of Ref. [37].

B.5. Maximum likelihood estimation on manifold

In this subsection, we review and provide intuitions for existing results on the asymptotic normality of *maximum likelihood estimation* (MLE) on a manifold [36] in Lemma 3. More concretely, Ref. [36] defines an appropriate notion of asymptotic normality of the distance between the estimator and the target on a manifold. Our main result of Theorem 2 follows directly by applying Lemma 3 to our analysis of MLE on the manifold of projective MPS, as defined in Appendix B.4.

Let (M, g) be a Riemannian manifold of dimension $r := \dim(M)$ with a Riemannian connection ∇ and let \mathcal{P}_{θ} be a family of distributions parametrized by elements $\theta \in M$. Let \mathcal{S} represent the sample space of the distributions, i.e., $\mathcal{P}_{\theta} : \mathcal{S} \to [0, 1]$. Our goal is to identify the target point $\theta_0 \in M$ on the manifold. To achieve this, we assume access to a *dataset*, which is formally defined as a random variable $D^{(N)}$ consisting of N samples drawn i.i.d. from the target distribution \mathcal{P}_{θ_0} . Given a realization of the dataset $D = (x_i \sim \mathcal{P}_{\theta_0})_{i=1}^N$, MLE aims to identify the element on the manifold θ^D for which the associated distribution $\mathcal{P}_{\theta D}$ is most likely to have generated the dataset D. This is achieved by minimizing the negative log-likelihood (NLL). More specifically, we define the NLL for the dataset instance D, denoted by $\mathcal{L}_{\text{NLL}}^D$, as a map from the manifold M to the real numbers,

$$\mathcal{L}^{D}_{\mathrm{NLL}}: M \longrightarrow \mathbb{R}, \tag{B35}$$

$$\theta \mapsto -\frac{1}{N} \sum_{i=1}^{N} \log \mathcal{P}_{\theta}[x_i],$$
(B36)

which measures how likely the observed samples are under the distribution \mathcal{P}_{θ} . By minimizing the NLL for the dataset D we find the estimator

λT

$$\theta^{D} = \operatorname{argmin}_{\theta \in M} \mathcal{L}_{\mathrm{NLL}}^{D}(\theta).$$
(B37)

For clarity, we use $\mathcal{L}_{\text{NLL}}^{D}(\theta)$ and θ^{D} to denote the NLL and MLE estimator for the specific dataset D consisting of N samples, whereas we use the notation $\hat{\theta}^{(N)}$ for the random variable corresponding to the minimizer of NLL associated to the N-sample dataset $D^{(N)}$ drawn randomly. In Euclidean space, minimization can be achieved by finding points where the gradient of the loss function vanishes. On a manifold M, we can define an estimating form ω , which plays the role of the gradient.

Definition 14 (Estimating form ([36] Definition 1)). A map

$$\omega: \mathcal{S} \times M \to \mathrm{T}M,\tag{B38}$$

is an estimating form if $\mathbb{E}_{x \sim \mathcal{P}_{\theta}}[\omega(x, \theta)] = 0$ for all $\theta \in M$. TM is the tangent bundle, which is a collection of tangent spaces over all points in the manifold.

A more concrete expression of the estimating form follows from the covariant derivative of the NLL.

$$\omega(x,\theta) = -\left(\mathcal{D}_{e_1}\log(\mathcal{P}_{\theta}[x]), \dots, \mathcal{D}_{e_r}\log(\mathcal{P}_{\theta}[x])\right)^{\top},\tag{B39}$$

where $\{e_a\}_{a=1}^r$ is a basis of the tangent space $T_{\theta}M$, and D_X denotes the directional derivative with respect to a tangent vector $X \in T_{\theta}M$ as introduced in Def. 6.

Proof. To show $\mathbb{E}_{x \sim \mathcal{P}_{\theta}}[\omega(x, \theta)] = 0$, we write down its definition

$$\mathbb{E}_{x \sim \mathcal{P}_{\theta}}[\omega(x,\theta)] = -\frac{1}{|\mathcal{S}|} \sum_{x \in \mathcal{S}} \mathcal{P}_{\theta}[x] \Big(D_{e_1} \log(\mathcal{P}_{\theta}[x]), \dots, D_{e_r} \log(\mathcal{P}_{\theta}[x]) \Big)^{\top},$$

$$= -\frac{1}{|\mathcal{S}|} \sum_{x \in \mathcal{S}} \Big(D_{e_1} \mathcal{P}_{\theta}[x], \dots, D_{e_r} \mathcal{P}_{\theta}[x] \Big)^{\top},$$

$$= -\Big(D_{e_1} \underbrace{\frac{1}{|\mathcal{S}|} \sum_{x \in \mathcal{S}} \mathcal{P}_{\theta}[x]}_{=1}, \dots, D_{e_r} \frac{1}{|\mathcal{S}|} \sum_{x \in \mathcal{S}} \mathcal{P}_{\theta}[x] \Big)^{\top},$$

$$= 0.$$
(B40)

In this definition, the directional derivative is well defined as long as \mathcal{P}_{θ} is differentiable, while terms with probability $\mathcal{P}_{\theta}[x] = 0$ are to be omitted in the sum, by the limit $y \log y \to 0$ as $y \to 0$. Consequently, the set of estimating forms $\{\omega(x, \theta), \theta \in M\}$ are differentiable vector fields.

Additionally, the asymptotic normality of an MLE estimator in Euclidean space requires the Hessian of the loss function to be well behaved around the target parameters. In the manifold setting, we need a similar condition and define the Hessian matrix below.

Definition 15 (Hessian matrix [36]). Let ω be a differentiable estimating form. Choose a basis set $\{e_a\}_{a=1}^r \in T_{\theta}M$ in the tangent space of the manifold at θ . The Hessian matrix, in this basis, is defined through the Riemannian connection of ω , and is given by its matrix elements:

$$\mathbf{H}_{a,b}(x,\theta) = \langle \nabla_{e_a} \omega(x,\theta) | e_b \rangle_a. \tag{B41}$$

In subsequent discussions, we often consider the expectation of the Hessian matrix elements over a distribution \mathcal{P}_{θ} given by

$$\mathbf{H}_{a,b}(\theta) = \mathbb{E}_{x \sim \mathcal{P}_{\theta}} \Big[\langle \nabla_{e_a} \omega(x, \theta) | e_b \rangle_a \Big]. \tag{B42}$$

Having reviewed the relevant concepts, we are now prepared to state the theorem concerning asymptotic normality of MLE on a manifold following Ref. [36]. Let us adopt the following notation: given an estimate $\theta \in M$, we denote by $\gamma_{\theta}(t)$ the shortest connecting geodesic (see Def. 8) such that $\gamma(0) = \theta_0$ and $\gamma(1) = \theta$, i.e., γ_{θ} is the shortest path connecting the MLE estimate to the target. Moreover, let $\{e_a\}_{a=1}^r$ be a basis of the target tangent space $T_{\theta_0}M$ and given a shortest connecting curve γ_{θ} , denote the parallel transport (see Def. 7) of the basis vectors by $\{e_a(t)\}_{a=1}^r$, i.e., $e_a(0) = e_a$ for all a. This allows us to decompose tangent vectors with respect to well-defined basis vectors along the *entire* curve γ_{θ} .

Lemma 3 (Asymptotic normality ([36] Theorem 2)). Let ω be the differentiable estimating form as defined in Def. 14 and H be the Hessian matrix defined in Def. 15. Let $(x_i)_{i=1}^N$ be a sequence of *i.i.d.* samples from \mathcal{P}_{θ_0} , where $\theta_0 \in M$ is the target element. Denote by $\hat{\theta}^{(N)}$ the MLE estimator, which is the minimizer of the NLL, and let γ be the connecting geodesic between $\gamma(0) = \theta_0$ and $\gamma(1) = \hat{\theta}^{(N)}$. Furthermore, consider the following conditions:

- 1. The estimator $\hat{\theta}^{(N)}$ satisfies $\sum_{i=1}^{N} \omega(x_i, \hat{\theta}^{(N)}) = 0$ (first-order condition) for all N and converges in distribution to the target $\hat{\theta}^{(N)} \xrightarrow{\mathbf{p}} \theta_0$ as $N \to \infty$.
- 2. The expected Hessian matrix $H_{a,b}(\theta_0)$ at the target point, defined in a basis $\{e_a\}_{a=1}^r$, is finite and invertible.

3. The Hessian $H(x,\theta)$ in a neighborhood of θ_0 converges uniformly to $H(x,\theta_0)$ as $\theta \to \theta_0$ in the following sense:

$$\mathbb{E}_{x \sim \mathcal{P}_{\theta_0}} \left[\sup_{\theta \in B_{\theta_0}(\delta)} \left| \langle \nabla_{e_a(1)} \omega(x, \theta) | e_b(1) \rangle - \langle \nabla_{e_a} \omega(x, \theta_0) | e_b \rangle \right| \right] \xrightarrow{\mathrm{p.}} 0$$
(B43)

as $\delta \to 0$, where $B_{\theta_0}(\delta)$ is the set of all points with distance at most δ from θ_0 (the " δ -Ball" around θ_0).

When these conditions are satisfied the components $\Delta_1, \ldots, \Delta_r$ of the velocity vector $\dot{\gamma}(t) = \sum_a \Delta_a e_a(t)$ converge in distribution to a multivariate normal distribution

$$\sqrt{N}(\Delta_1, \dots, \Delta_r)^\top \xrightarrow{\mathrm{d.}} \mathcal{N}(0, \Sigma), \tag{B44}$$

where $\Sigma = (\mathrm{H}^{\dagger})^{-1}\Gamma\mathrm{H}$ and the Gram matrix $\Gamma_{a,b} = \mathbb{E}_{x \sim \mathcal{P}_{\theta_0}} \left[\langle \omega(x,\theta_0) | e_a \rangle \langle \omega(x,\theta_0) | e_b \rangle \right].$

B.6. Main result: A concentration inequality for infidelity

In this subsection, we prove our main result that the infidelity between an MLE estimator and the target state can be probabilistically upper bounded when the number of samples N is asymptotically large. Theorem 2 states that such an upper bound depends on the properties of the target state and converges to zero with a rate depending on the number of samples as $1/\sqrt{N}$. A central component of our proof relies on the asymptotic normality of the MLE estimator on a manifold discussed in Lemma 3. Before rigorously proving our theorem, we first integrate the concepts from previous sections and restate the normality condition in the context of the projective MPS manifold \tilde{M} , as discussed in Appendix B.4 and B.5.

Consider the target state $[\phi] \in \tilde{M}$, with a normalized representative $|\phi\rangle \in \mathcal{H}$. Given a unitary ensemble \mathcal{U} , the sample space \mathcal{S} consists of all unitary-bit-string pairs $\mathcal{S} = \mathcal{U} \times \{0,1\}^n$. Let \mathcal{D} denote the data probability distribution corresponding to the target state $[\phi]$ given by

$$\mathcal{D}[(U,b),[\phi]] = \frac{1}{|\mathcal{U}|} \times \langle b|U|\phi\rangle\!\langle\phi|U^{\dagger}|b\rangle.$$
(B45)

For a dataset $D = \{(U_i, b_i)\}_{i=1}^N$ drawn i.i.d. from \mathcal{D} , its corresponding *negative log-likelihood* (NLL) is defined as

$$\mathcal{L}^{D}_{\mathrm{NLL}}: \tilde{M} \longrightarrow \mathbb{R}, \tag{B46}$$

$$[\psi] \longmapsto -\frac{1}{N} \sum_{i=1}^{N} \log |\langle b_i | U_i | \psi \rangle|^2, \tag{B47}$$

Notice that by introducing a scaling factor c > 1, the loss associated with any vector can be arbitrarily reduced through the transformation $|\psi\rangle \mapsto c |\psi\rangle$. However, such rescaling does not affect physical observables and is effectively managed by constraining the analysis to the projective Hilbert space. This highlights the practical value of defining the projective MPS and its associated projective manifold. The MLE estimate is given by the minimizer of the NLL

$$[\psi^{D}] = \operatorname{argmin}_{[\psi] \in \tilde{M}} \mathcal{L}_{\mathrm{NLL}}^{D} ([\psi]).$$
(B48)

Using the notation established in Appendix B.5, we use $[\hat{\psi}^{(N)}]$ to denote the random variable of the MLE estimator, which is the minimizer of the N-sample NLL for a dataset random variable $D^{(N)}$. Using Fact 2 we find a concrete expression of the estimating form by taking covariant derivative of the NLL as follows.

Corollary 1 (Differentiable estimating form). $\omega((U,b), [\psi]) = -(D_{e_1} \log |\langle b|U|\psi\rangle|^2, \dots, D_{e_r} \log |\langle b|U|\psi\rangle|^2)^\top$ is a differentiable estimating form.

In order to establish the essential property of asymptotic normality of the projective MPS estimator, as required to prove our main theorem, we first adapt the necessary conditions in Lemma 3 to our MPS manifold setting as follows.

Proof. In the infinite-sample limit, he NLL is given by

and the estimator converges to the target state in probability $[\hat{\psi}^{(N)}] \xrightarrow{\mathbf{p}} [\phi]$.

$$\lim_{N \to \infty} \mathcal{L}_{\text{NLL}}([\psi]) = -\frac{1}{|\mathcal{U}|} \sum_{(U,b) \in \mathcal{U} \times \{0,1\}^n} |\langle b|U|\phi \rangle|^2 \log\left(|\langle b|U|\psi \rangle|^2 / |\mathcal{U}|\right), \tag{B49}$$

where $|\phi\rangle$ and $|\psi\rangle$ are normalized representations of their rays. The NLL loss function achieves the minimum denoted by $[\hat{\psi}]^{-6}$ if the probabilities agree for all measurement outcomes

$$|\langle b|U|\hat{\psi}\rangle|^2 = |\langle b|U|\phi\rangle|^2, \quad \forall (b,U).$$
(B50)

Since the unitary ensemble \mathcal{U} is tomographically complete, Eq. (B50) implies $|\hat{\psi}\rangle = c|\phi\rangle$ up to a global phase c, so that $[\hat{\psi}] = [\phi]$.

Assumption 1. The expected Hessian at the target state $[\phi]$ is invertible and finite.

Assumption 2. The Hessian at $[\hat{\psi}^{(N)}]$ in a small neighbourhood near the target converges uniformly to the Hessian at $[\phi]$ in expectation with respect to \mathcal{D} .

Having established or assumed the conditions outlined in Lemma 3, we are now prepared to formally apply the normality result to the MPS-manifold setting. This application is detailed in Corollary 2, which extends the lemma to our specific context.

Corollary 2 (Asymptotic normality on the projective MPS manifold). Let \mathcal{D} be the distribution corresponding to the target state $[\phi]$ (defined in Eq. (B45)), and $D^{(N)} = \{(U_i, b_i)\}_{i=1}^N$ be a sequence of i.i.d. samples from \mathcal{D} . Let $[\hat{\psi}^{(N)}]$ be the minimizer of the negative log-likelihood function and ω its associated estimating form as defined in Corollary 1. Define γ as the shortest connecting geodesic such that $\gamma(0) = [\phi]$, a point $\gamma(1) = [\hat{\psi}^{(N)}]$, and $\dot{\gamma}$ as the velocity vector along the curve. Let the velocity vector be decomposed in the parallel transporte—along γ —basis vectors $\{e_a(t)\}$ as $\dot{\gamma} = \sum_a \Delta_a e_a(t)$. Suppose the conditions of Proposition 1, Assumption 1, and Assumption 2 are satisfied. By Lemma 3, the velocity vector's components $\{\Delta_a\}_{a=1}^r$ converge in distribution to a multivariate normal distribution

$$\sqrt{N}(\Delta_1, \dots, \Delta_r)^\top \xrightarrow{\mathrm{d.}} \mathcal{N}(0, \Sigma), \tag{B51}$$

where $\Sigma = (\mathrm{H}^{\dagger})^{-1}\Gamma\mathrm{H}$ and the Gram matrix $\Gamma_{a,b} = \mathbb{E}_{(U,b)\sim\mathcal{D}}[\langle \omega((U,b), [\phi])|e_a\rangle\langle \omega((U,b), [\phi])|e_b\rangle].$

Corollary 2 establishes the asymptotic normality of the velocity vector's components along the shortest connecting geodesic γ in the projective MPS manifold \tilde{M} , which is crucial for determining the geodesic length (see Fact 1 and Eq. (B22)). Since \tilde{M} is a submanifold of the projective Hilbert space $\mathsf{P}(\mathcal{H})$ with an induced metric \tilde{g} , lifting $\gamma : [0,1] \to \tilde{M}$ to this embedding space to $\gamma' : [0,1] \to \mathsf{P}(\mathcal{H})$ reveals that its length is always longer than or equal to that of a shortest connecting geodesic in $\mathsf{P}(\mathcal{H})$; equality occurs only when the geodesic curve in $\mathsf{P}(\mathcal{H})$ is exactly γ' . As the geodesic curve in $\mathsf{P}(\mathcal{H})$ defines the Fubini-Study distance, the length of the geodesic thus relates to the state overlap using Lemma 1. This relation allows us to upper bound the infidelity between quantum states by the length of γ within \tilde{M} . Using the properties and asymptotic behavior of these geodesic paths, we provide a bounded estimation of the state infidelity in Lemma 4 below.

⁶ The minimum is given by the Shannon entropy of the target distribution \mathcal{D} , calculated as $-\sum_{(U,b)} \mathcal{D}[(U,b),[\phi]] \log \mathcal{D}[(U,b),[\phi]]$. This follows from Gibbs' inequality, which, for two discrete probability distributions $\{p_1, \dots, p_N\}$ and $\{q_1, \dots, q_N\}$, states that $-\sum_{i=1}^N p_i \log p_i \leq -\sum_{i=1}^N p_i \log q_i$. The equality is achieved when two distributions are the same $p_i = q_i$.

27

Lemma 4 (Infidelity bounded by geodesic curve length). Given two states $[\psi], [\phi]$ in \tilde{M} , the projective space of MPS (see Lemma 2), denote the shortest connecting geodesic by $\gamma : [0,1] \rightarrow \tilde{M}$, such that $\gamma(0) = [\psi]$ and $\gamma(1) = [\phi]$, as defined in Def. 8. The infidelity of $[\psi]$ and $[\phi]$ is bounded by the length of the geodesic:

$$1 - |\langle \psi | \phi \rangle| \le \int_0^1 \mathrm{d}t \sqrt{\langle \dot{\gamma} | \dot{\gamma} \rangle_{\tilde{g}}},\tag{B52}$$

where $\langle \bullet | \bullet \rangle_{\tilde{a}}$ denotes the induced Fubini-Study metric on \tilde{M} .

Proof. By Lemma 1, the infidelity between two state vectors $|\psi\rangle$ and $|\phi\rangle$ is upper bounded by their Fubini-Study distance in the projective Hilbert space as

$$1 - |\langle \psi | \phi \rangle| = 1 - \cos d_{\rm FS}([\phi], [\psi]), \qquad (B53)$$

$$\leq d_{\rm FS}([\phi], [\psi]).$$

Recall that the Fubini-Study distance is defined as the length of the shortest curve between two points in projective Hilbert space (see Def. 12). Hence, for any differentiable curve $\alpha : [0, 1] \to \mathsf{P}(\mathcal{H})$, with $\alpha(0) = [\psi]$ and $\alpha(1) = [\phi]$, the length of α also upper bounds the infidelity as

$$1 - |\langle \psi | \phi \rangle| \le \int_0^1 \mathrm{d}t \sqrt{\langle \dot{\alpha} | \dot{\alpha} \rangle_g},\tag{B54}$$

where $\langle \bullet | \bullet \rangle_g$ denotes the Fubini-Study metric on $\mathsf{P}(\mathcal{H})$. Note that the choice of the parameter interval [0, 1] is without loss of generality because we can reparametrize any connecting curve to an arbitrary closed interval of our choice. Since \tilde{M} is a submanifold of $\mathsf{P}(\mathcal{H})$, we can lift the shortest connecting geodesic γ from the projective MPS manifold \tilde{M} into projective Hilbert space $\mathsf{P}(\mathcal{H})$. To be precise, we can formally distinguish the curve γ in the two abstract manifolds \tilde{M} and $\mathsf{P}(\mathcal{H})$. For this, we formally identify each point $\gamma(t) \in \tilde{M}$ with its corresponding point in $\gamma'(t) \in \mathsf{P}(\mathcal{H})$. Since \tilde{g} is induced by g, the lengths of any tangent vectors are the same in both the submanifold and the embedding manifold $\langle \dot{\gamma} | \dot{\gamma} \rangle_{\tilde{g}} = \langle \dot{\gamma}' | \dot{\gamma}' \rangle_g$ (see Eqs. (103) and (108) in Ref. [37]). Therefore, the lengths of γ and γ' are equal. Choosing the curve to be $\alpha = \gamma'$ in Eq. (B54), we have proved the result stated in Eq. (B52).

Finally, using Corollary 2 and Lemma 4, we provide an upper bound on the infidelity, which is our main result below.

Theorem 2 (Probabilistic bound). Let \tilde{M} be the manifold of full-rank MPS of a maximum bond dimensions χ_{\max} , as per Eq. (B33), and let $[\phi] \in \tilde{M}$ be some target state. Let $D^{(N)} = \{(U_i, b_i)\}_{i=1}^N$ be the random variable representing a dataset of N samples drawn from the distribution $\mathcal{D}[(U, b), [\phi]]$ defined in Eq. (B45). Denote by $[\hat{\psi}^{(N)}]$ the random variable representing the minimizer, in \tilde{M} , of the N-sample negative log-likelihood defined by $D^{(N)}$, and let ω be its associated estimating form as defined in Corollary 1. Assume the conditions of Proposition 1, Assumption 1, and Assumption 2 are satisfied. Then, for any $\delta \in (0, 1]$, there exists N_0 such that for all $N > N_0$, one has

$$\mathbb{P}\left[1 - |\langle \hat{\psi}^{(N)} | \phi \rangle| \le \epsilon(N)\right] \ge 1 - \delta, \tag{B55}$$

where

$$\epsilon(N) = O\left(\sqrt{\frac{n\chi_{\max}^2}{N\delta}}\right). \tag{B56}$$

Proof. For simplicity, we assume that $|\hat{\psi}^{(N)}\rangle$ and $|\phi\rangle$ are normalized elements of the equivalence classes $[\hat{\psi}^{(N)}]$ and $[\phi]$, respectively. Let \tilde{g} be the Fubini-Study metric of the projective MPS manifold \tilde{M} induced by the Fubini-Study metric of the projective Hilbert space via Eq. (B34). On this manifold, γ is the shortest connecting geodesic between $\gamma(0) = [\phi]$ and $\gamma(1) = [\hat{\psi}^{(N)}]$. By Proposition. 1, the estimator converges to the target state in the infinite-sample limit

$$[\hat{\psi}^{(N)}] \xrightarrow{\mathbf{p}} [\phi],$$
 (B57)

as $N \to \infty$. To quantify the error between the N-sample estimator and the target state, we use Lemma. 4 to upper bound the N-sample infidelity by the length of the geodesic curve. More specifically, according to Fact 1, this geodesic length is the same as the length of the velocity vector, which is constant throughout the curve, so that

$$1 - |\langle \hat{\psi}^{(N)} | \phi \rangle| \le \sqrt{\langle \dot{\gamma} | \dot{\gamma} \rangle_{\tilde{g}}}.$$
(B58)

Therefore, to obtain an upper bound for the infidelity, it suffices to analyze the velocity vector defined at an arbitrary point along the curve, which we choose to be $\gamma(0) = [\phi]$. The velocity vector, in the basis $\{e_a\}_{a=1}^r$ of the tangent space $T_{[\phi]}\tilde{M}$, can be written as

$$\dot{\gamma} = \sum_{a=1}^{\prime} \Delta_a^{(N)} e_a. \tag{B59}$$

Expressing the Fubini-Study metric in the same basis as $\tilde{g}_{a,b}$, the squared length of the velocity vector is explicitly given by

$$\langle \dot{\gamma} | \dot{\gamma} \rangle_{\tilde{g}} = \sum_{a,b} \Delta_a^{(N)} \, \tilde{g}_{a,b} \, \Delta_b^{(N)}. \tag{B60}$$

For simplicity, we represent the velocity vector's components as a column vector $\Delta^{(N)} = (\Delta_1^{(N)}, \dots, \Delta_r^{(N)})^{\top}$, so that its squared length is

$$\langle \dot{\gamma} | \dot{\gamma} \rangle_{\tilde{g}} = \Delta^{(N)^{\top}} \tilde{g} \, \Delta^{(N)}, \tag{B61}$$

where, in slight abuse of notation, we write $\tilde{g} = (\tilde{g}_{a,b})_{a,b=1}^r$. By virtue of Assumptions 1 and 2, together with the consistency of the estimator in Eq. (B57), Corollary 2 states that the *N*-sample estimator $\sqrt{N}\Delta^{(N)}$ converges in distribution to the normal distribution as

$$\sqrt{N}\Delta^{(N)} \stackrel{\mathrm{d}}{\to} \Delta^{(\infty)} \sim \mathcal{N}(0, \Sigma). \tag{B62}$$

Therefore, in the infinite-sample limit $N \to \infty$, we can upper bound the probability of the norm of $\Delta^{(\infty)}$ (with respect to the metric \tilde{g}) greater than an error ε using Markov's inequality as

$$\mathbb{P}\left[\Delta^{(\infty)} \tilde{g} \, \Delta^{(\infty)} \ge \varepsilon\right] \le \frac{\mathbb{E}\left(\Delta^{(\infty)} \tilde{g} \, \Delta^{(\infty)}\right)}{\varepsilon}.$$
(B63)

To compute the expectation of the $\Delta^{(\infty)}$ -inner product on the r.h.s., we use the following identity

$$\mathbb{E}\left(\Delta^{(\infty)^{\top}}\tilde{g}\,\Delta^{(\infty)}\right) = \mathbb{E}\left(\operatorname{Tr}(\tilde{g}\,\Delta^{(\infty)}\Delta^{(\infty)^{\top}})\right), \qquad (B64)$$
$$= \operatorname{Tr}\left(\tilde{g}\,\mathbb{E}\left(\Delta^{(\infty)}\Delta^{(\infty)^{\top}}\right)\right), \qquad (B7)$$
$$= \operatorname{Tr}\left(\tilde{g}\,\Sigma\right).$$

Going from the second to the third line, first we have used the definition of the covariance matrix

$$\Sigma = \mathbb{E}\left(\Delta^{(\infty)}\Delta^{(\infty)}\right) - \mathbb{E}\left(\Delta^{(\infty)}\right)\mathbb{E}\left(\Delta^{(\infty)}\right).$$
(B65)

Second, using the zero-mean property of $\mathbb{E}(\Delta^{(\infty)}) = 0$, we arrive at an explicit expression for the expectation of the $\Delta^{(\infty)}$ -inner product as

$$\mathbb{E}\left(\Delta^{(\infty)}{}^{\top}\tilde{g}\,\Delta^{(\infty)}\right) = \operatorname{Tr}\left(\tilde{g}\,\Sigma\right). \tag{B66}$$

Plugging the expectation values derived in Eq. (B66) into the r.h.s. of Eq. (B63), we have the following bound

$$\mathbb{P}\left[\Delta^{(\infty)}{}^{\top}\tilde{g}\,\Delta^{(\infty)} \ge \varepsilon\right] \le \frac{\operatorname{Tr}\left(\tilde{g}\,\Sigma\right)}{\varepsilon}.\tag{B67}$$

To analyze the finite- $N \Delta^{(N)}$ -inner product estimator, we first notice that the estimator converges in probability to its infinite-N limit as

$$\mathbb{P}\left[N\Delta^{(N)}{}^{\top}\tilde{g}\,\Delta^{(N)} \ge \varepsilon\right] \xrightarrow{N \to \infty} \mathbb{P}\left[\Delta^{(\infty)}{}^{\top}\tilde{g}\,\Delta^{(\infty)} \ge \varepsilon\right].$$
(B68)

Once we choose a convergence error tolerance η , there exists an N_0 such that the finite-sample estimator is close to the infinite-sample limit upper bounded by η for all $N > N_0$:

$$\mathbb{P}\left[N\Delta^{(N)}{}^{\top}\tilde{g}\,\Delta^{(N)} \ge \varepsilon\right] - \mathbb{P}\left[\Delta^{(\infty)}{}^{\top}\tilde{g}\,\Delta^{(\infty)} \ge \varepsilon\right] \le \eta.$$
(B69)

Therefore, we can bound the finite-sample estimator by

$$\mathbb{P}\left[N\Delta^{(N)}{}^{\top}\tilde{g}\,\Delta^{(N)} \ge \varepsilon\right] \le \frac{\operatorname{Tr}(\tilde{g}\,\Sigma)}{\varepsilon} + \eta.$$
(B70)

Notice that in the l.h.s., there is a factor of N; in order to bound the squared length, we can rescale $\varepsilon' = \varepsilon/N$, so that the following inequality is satisfied for all $N \ge N_0$

$$\mathbb{P}\left[\Delta^{(N)}{}^{\top}\tilde{g}\,\Delta^{(N)} \ge \varepsilon'\right] \le \frac{\mathrm{Tr}(g\Sigma)}{N\varepsilon'} + \eta. \tag{B71}$$

By choosing a tolerance probability δ such that the lower bound for the squared length is satisfied,

$$\frac{\operatorname{Tr}\left(\tilde{g}\,\Sigma\right)}{N\varepsilon'(N)} + \eta = \delta,\tag{B72}$$

we can solve for the N-sample error $\varepsilon'(N)$ as

$$\varepsilon'(N) = \frac{\operatorname{Tr}(\tilde{g}\,\Sigma)}{N(\delta - \eta)}.\tag{B73}$$

Finally, we can upper bound the squared length by taking the complement of the inequality. With probability $1 - \delta$, the squared-length estimator is upper bounded by $\varepsilon'(N)$, i.e.

$$\mathbb{P}\left[\Delta^{(N)}{}^{\top}\tilde{g}\,\Delta^{(N)} \le \varepsilon'(N)\right] \ge 1 - \delta. \tag{B74}$$

To achieve a bound for the infidelity, we take the square root of both sides of the inequality and arrive at an upper bound of the length of the velocity vector, which bounds the infidelity by Eq. (B58) and Eq. (B61), so that

$$\mathbb{P}\left[1 - |\langle \hat{\psi}^{(N)} | \phi \rangle| \le \sqrt{\varepsilon'(N)}\right] \ge 1 - \delta.$$
(B75)

Finally, we can upper bound the N-sample error threshold as

$$\varepsilon'(N) \le \frac{\dim(M) \|\tilde{g}\,\Sigma\|_{\infty}}{N(\delta - \eta)},\tag{B76}$$

where $\| \bullet \|_{\infty}$ is the ∞ -Schatten norm. The dimension of the projective MPS manifold can be upper bounded as $\dim(\tilde{M}) = O(n\chi_{\max}^2)$ where χ_{\max} is the maximum bond dimension $\chi_{\max} = \max(\chi)$.

The η -term in the denominator is a consequence of the fact that the estimator $[\hat{\psi}^{(N)}]$ is only asymptotically normal. The smaller one wants to make η , the larger N_0 becomes so that we have to move "closer" to the asymptotic regime. Asymptotically, i.e., in the limit $N \to \infty$, we can choose $\eta \to 0$ to arrive at our main result in Theorem 2

$$\mathbb{P}\left[1 - |\langle \hat{\psi}^{(N)} | \phi \rangle| \le \epsilon(N)\right] \ge 1 - \delta, \tag{B77}$$

where the infidelity error threshold is given by

$$\epsilon(N) = O\left(\sqrt{\frac{n\chi_{\max}^2}{N\delta}}\right). \tag{B78}$$



Supplementary Figure 1. Infidelity scaling for interpolation between a random and GHZ 3-qubit state. The numerical results illustrate the polynomial scaling of the infidelity with the number of samples N across a family of states, parametrized by x. Here, x ranges from 0 (darkest color, representing the GHZ state) to 1 (lightest color, representing the random state). Dashed lines indicate polynomial fits of the form $1 - F = c \times N^{-\alpha}$, with values for c and α provided in the legend.

B.7. Discussions and interpretations

In this subsection, we discuss Theorem 2 and its connection to the numerical results presented in the main text. This theorem demonstrates that the MPS solution of maximum likelihood estimate (MLE) approaches the target state when the number of samples N is sufficiently large, with the infidelity 1 - F converging at a rate proportional to $1/\sqrt{N}$. This provides a theoretical foundation for understanding the scaling behavior observed in our numerical experiments.

First, we comment on how this theoretical result relates to our numerics in practice. In the derivation of our theorem, we have assumed that our variational space forms a projective MPS manifold. However, in our numerical simulations, we fix the bond dimension across sites, including variational states that are not always full rank, although they do include a projective MPS manifold as a subspace. In addition, several assumptions underpinning our result could be compromised: 1) The Hessian of the *negative log-likelihood* (NLL) loss function may not behave as expected. 2) The unitary ensemble is not tomographically complete in our variational space. Notably, while Theorem 1 confirms that our random-XZ measurements are tomographically complete for real and pure states, our variational approach involves complex numbers for enhanced stability in optimization, leading to a space of complex pure states. To overcome this discrepancy, incorporating random-Pauli measurements, which are known to be tomographically complete, could provide a solution. Alternatively, our analysis of the complex projective MPS manifold could be adapted to the setting of a real projective MPS, and one can verify that the necessary conditions on the manifold structure are satisfied in the subspace of real states.

Focusing on our numerical observations, we empirically find that the infidelity decreases polynomially with the number of samples N as $1 - F \propto \frac{1}{N^{\alpha}}$. Theorem 2 generally predicts that $\alpha = 0.5$. Specifically, a strong alignment with the polynomial functional form indicates that N is sufficiently large to approach the asymptotic limit. In practice, we see that some of our numerical results for n = 9 qubits exhibit such a polynomial scaling of the infidelity with $\alpha \approx 1$ in Fig. 2(b), while others do not (see Fig. 3(b,c)). The violation of our prediction could be explained by the invalidity of one or more of our assumptions. Intriguingly, $\alpha = 1$ is actually better than the upper bound provided in our theorem, pointing to additional complexities.

To address discrepancies between our numerical observations and our theoretical framework, we have performed additional numerical experiments using random states to validate our theorem's predictions. In Supplementary Fig. 1, we illustrate how the infidelity scales with the number of samples N for states that interpolate between a GHZ state $|\text{GHZ}\rangle$ and a random state vector $|\psi_{\chi=2}\rangle$ with bond dimension 2. The level of randomness is controlled by the parameter x, as defined by

$$|\phi\rangle = \sqrt{1 - x} |\text{GHZ}\rangle + \sqrt{x} |\psi_{\chi=2}\rangle. \tag{B79}$$

Our findings reveal that for a purely random state (x = 1), the observed scaling exponent $\alpha = 0.53$ aligns perfectly with our theoretical prediction. Conversely, at x = 0, corresponding to the GHZ state, an unexpected deviation with $\alpha = 1$ mirrors the anomaly previously observed in Fig. 2(b). This suggests that the inherent structures of stabilizer states, like the GHZ and the surface code state, might facilitate a better scaling than our current upper bound.

Finally, we note that in our practical numerics, we perform MLE in the space of MPS tensors, and not directly on the manifold. Theoretically, the trajectory towards the minimum would differ between these two approaches. Nonetheless, our theorem focuses on the properties of the *optimal* MLE estimator, acknowledging that practically achieving this solution may not always be efficient. Note, however, that once we have found MPS tensors corresponding to the minimizer of the loss function, they correspond to a *unique* physical state in Hilbert space. Thus, the theorem still makes a meaningful statement about the practical optimizer with respect to the MPS tensors, under the condition that the optimizer finds a global minimizer of the loss function. We leave potential improvement of the optimization by incorporating insights from the manifold structure [84, 85] for future work.

APPENDIX C: Random-XZ classical shadow tomography

This self-contained appendix fully characterizes random-XZ classical shadow tomography and derives the necessary tools to implement it. We use the results outlined in this section when estimating observables in the regularization step of Eq. (5); see also Appendix E.2 for the details of regularization using random-XZ classical shadows. We are also motivated to characterize random-XZ classical shadows because measurements of this form allow one to fully characterize real, pure states, which are the target states in this work (refer to Appendix D for a proof). Random-XZ classical shadows may also be of interest more broadly due to the ease of their experimental implementation across different platforms. For instance, with superconducting qubits, measurement in different bases only requires single-qubit rotation which are cheap to implement. In the context of Rydberg atom arrays, this holds for both physical and logical qubits. Single-qubit rotations on physical qubits have high fidelities [64]. For logical qubits, transversal Hadamard rotates the logical to the X basis by using single-qubit physical rotations, and even without these transversal Hadamard gates, logical qubits can still be effectively measured in the X basis [31]. In this latter case, performing XZ measurements is not only doable but also provides the ability to perform error detection (e.g., consider the [[15,1,3]] Reed-Muller 3D color code).

Random-XZ classical shadows involve preparing some (unknown) state, randomly measuring some *n*-qubit Pauli string composed of only X and Z single-site Paulis, and then post-processing to predict properties of interest. We assume that the chosen Pauli is uniformly sampled from the set of all *n*-qubit strings composed of only Xs and Zs. As this set of measurements is not informationally complete, random-XZ classical shadows do not allow us to predict any arbitrary observable. However, we can predict some observables of interest, namely, those that contribute to our regularization step. In this appendix, we will discuss which observables can be estimated with this technique, discuss how to post-process the data to estimate observables, and how to calculate the sample complexity required for learning Pauli strings. We specifically discuss the sample complexity for learning Pauli strings because these are the observables that we estimate in our regularization step. In the first section of this appendix, we introduce these ideas pedagogically, with the rigorous backings for all our claims provided in the second section.

C.1. Performing random-XZ classical shadow tomography

Here, we derive an expression for the classical shadow of the quantum state measured by the randomized XZ measurements discussed in Appendix D. Random-XZ classical shadows involve taking many random measurements on some prepared state. Each measurement involves two steps. First, the state is rotated under some unitary U which is sampled from $U \sim \mathcal{U}_{XZ}$, and then, in the second step, the state is measured in the computational basis states $\{|b\rangle\langle b|\}_{b\in\{0,1\}^n}$. This ensemble of unitaries \mathcal{U}_{XZ} exhibits a uniform probability distribution over unitaries U which take the form

$$U = \bigotimes_{i=1}^{n} U^{\alpha_i} \text{ where } \alpha_i \in \{X, Z\}.$$
(C1)

Here, U^{α_i} is the single-qubit unitary which rotates from the computational basis to the (different) basis of the Pauli operator σ_i^{α} acting on site *i*. Therefore, acting on the prepared state, the ensemble \mathcal{U}_{XZ} allows one to randomly measure any *n*-qubit Pauli string composed of only *X* and *Z* single-site Paulis. Again, this set of measurements is not informationally complete, and as such, we cannot predict any arbitrary observable with random-*XZ* classical shadows.

In this appendix, we will use the nomenclature developed in Ref. [15]. In particular, we define the visible space of our random-XZ measurement scheme as the space of all operators O for which we *can* estimate the expectation value $Tr(\rho O)$; we denote this as VisibleSpace(\mathcal{U}). Similarly, the invisible space is the space of operators that are invisible to us under the implementable unitaries \mathcal{U} , and we label this as InvisibleSpace(\mathcal{U}). Note that the full operator space $\mathcal{B}(\mathcal{H})$ associated to the Hilbert space \mathcal{H} can be expressed with the following tensor sum:

$$\mathcal{B}(\mathcal{H}) \simeq \mathsf{VisibleSpace}(\mathcal{U}) \oplus \mathsf{InvisibleSpace}(\mathcal{U}), \tag{C2}$$

where the symbol \simeq stands for "isomorphic to". As a general rule, a more limited set of implementable unitaries \mathcal{U} corresponds to a smaller VisibleSpace(\mathcal{U}). In Corollary 4, we find that the random-XZ visible space is spanned by Pauli strings constructed from single-site 1, X, and Z Paulis. We estimate these Paulis and then plug them into our regularization step.

We can estimate observables $O \in \mathsf{VisibleSpace}(\mathcal{U}_{XZ})$ by post-processing the data to construct an estimate $\hat{\rho}$ of the true density matrix ρ . This estimate is called the "classical shadow" of the true density matrix, and we can use it to compute $\mathrm{Tr}(\hat{\rho}O) \approx \mathrm{Tr}(\rho O)$. We construct the classical shadow using the data from our measurements. Each

measurement snapshot can be characterized by the unitary applied, $U = \bigotimes_i U^{\alpha_i}$, and the computational basis state measured, $|b\rangle = \bigotimes_i |b_i\rangle$. Here the *i* indexes the qubits. We can construct a single-shot estimate of our density matrix, $\hat{\rho}_1$, from random-XZ data as

$$\hat{\rho}_{1} = \mathcal{M}_{XZ}^{-1}(U^{\dagger}|b\rangle\!\langle b|U)$$

$$= \otimes_{i=1}^{n} \left(\operatorname{Tr}(U^{\alpha_{i}\dagger}|b_{i}\rangle\!\langle b_{i}|U^{\alpha_{i}}X)X + \operatorname{Tr}(U^{\alpha_{i}\dagger}|b_{i}\rangle\!\langle b_{i}|U^{\alpha_{i}}Z)Z + \frac{1}{2}\mathbb{I} \right)$$

$$= \otimes_{i=1}^{n} \left(2U^{\alpha_{i}\dagger}|b_{i}\rangle\!\langle b_{i}|U^{\alpha_{i}} - \frac{\mathbb{I}}{2} \right).$$
(C3)

The subscript 1 above indicates that the classical shadow was constructed with data from 1 measurement. The channel \mathcal{M}_{XZ}^{-1} is the Moore-Penrose pseudoinverse of the measurement channel \mathcal{M}_{XZ} derived in Proposition 2 in the next subsection. This channel models the measurement step of the randomized-measurement procedure, and for an infinite set of informationally complete measurements, it guarantees that we will recover our original density matrix:

$$\rho = \lim_{N \to \infty} \hat{\rho}_N,\tag{C4}$$

where $\hat{\rho}_N$ is the classical shadow constructed from N measurements. One example, of an informationally complete measurement scheme is that of random Pauli measurements, where each qubit is randomly measured in either the X, Y, or Z basis. Intuitively, since the Pauli strings form a basis of the operator space $\mathcal{B}(\mathcal{H})$, this allows one to predict any observable. However, in our particular setting, the random-XZ measurements are not informationally complete, and as such, the infinite-measurement XZ classical shadow will not exactly recover every ρ . Instead, with an infinite number of XZ measurements, we will recover the exact density matrix projected onto the visible space. Therefore, we can only learn the observables in the visible space because we only learn this portion of ρ .

The sample complexity of classical shadow protocols is defined in terms of the shadow norm $||O||_{\text{shadow}}^2$, an upper bound on the variance of predicting $\text{Tr}(\rho O)$ with this measurement-and-reconstruction technique. Of course, to get an actual number of samples needed to predict $\text{Tr}(\rho O)$ to some precision ϵ , one would need to employ a concentration inequality. While the original classical shadows proposal demonstrated a median-of-means estimation technique [25], this concentration inequality has a large prefactor, which increases the number of samples by almost two orders of magnitude. Instead, we recommend that the interested reader employs Bernstein's inequality, as has been used in Ref. [15], which uses an empirical average over the data from N shots

$$\hat{\rho}_N = \frac{1}{N} \sum_{j=1}^N \mathcal{M}_{XZ}^{-1} \left(U^{(j)\dagger} | b^{(j)} \rangle \langle b^{(j)} | U^{(j)} \right).$$
(C5)

Here, we represent the unitary applied in the *j*th random measurement and the outcome bit-string of that measurement as $U^{(j)}$ and $b^{(j)}$, respectively.

Just as is the case with random Pauli classical shadows (measuring any Pauli string in a given shot), random-XZ classical shadows are most effective at learning local observables, the intuition being that we more often measure in bases that tell us something about that observable. Consider, for example, the observables $X \otimes \mathbb{1}^{\otimes n-1}$ and $X^{\otimes n}$. As long as we measure the first qubit in the X basis, the measurement will provide information about $\text{Tr}(\rho X \otimes \mathbb{1}^{\otimes n-1})$. This happens with probability 1/2: there is a 1/2 chance of measuring the first qubit in the X basis and a 1/2 chance of measuring it in the Z basis. Now, let us consider the second observable: $X^{\otimes n}$. In order to learn about the expectation value of $X^{\otimes n}$ with our state, every qubit must be measured in the X basis, and this happens with probability $1/2^n$ since there are n qubits and each has a 1/2 chance of being measured in X. Therefore, it is much easier to learn $X \otimes \mathbb{1}^{\otimes n-1}$ than $X^{\otimes n}$. In the next subsection of this appendix, we provide sample complexity upper bounds on Pauli strings in the form of the shadow norm of Pauli strings. This result is consistent with our intuition that random-XZ classical shadows can more easily learn local Pauli strings than nonlocal Pauli strings. In particular, a k-local Pauli string P_k (k non-identity single-site Paulis) exhibits a shadow norm of $||P_k||^2_{\text{shadow}} = 2^k$.

C.2. Derivation of tools for random-XZ classical shadows

In the remainder of this appendix, we derive the tools and claims discussed above for performing random-XZ classical shadow tomography. First, we will derive the classical shadow measurement channel \mathcal{M}_{XZ} for the learning procedure described in the previous subsection. Once a set of implementable unitaries has been specified, the measurement scheme can be defined in the language of probability density functions. For the random-XZ measurement

scheme considered above, we define our ensemble \mathcal{U}_{XZ} to have a uniform probability density function in the space of all possible unitaries rotating into X and Z single-qubit bases. This allows us to use Observation 4 from Ref. [15] in order to define the XZ visible space. Specifically, the image of the XZ measurement channel defines the visible space. While the visible space is technically independent of the underlying probability density function (being an indicator of what can be learned with the measurement unitaries), the uniform distribution allows us to characterize the visible space since each unitary has a nonzero probability of being sampled [15]. In addition to characterizing the visible space, we will also use the measurement channel to define the shadow norm of Pauli strings. As the unitaries in the ensemble are Cliffords, the eigenoperators of the measurement channel are Pauli strings, and the eigenvalues are the inverse shadow norms of the corresponding eigenoperators [86, 87].

Proposition 2 (Random-XZ classical shadow measurement channel). Assume a measurement primitive U_{XZ} such that the associated probability density function is uniform over random-XZ unitaries. The random-XZ measurement channel for n qubits can be written as

$$\mathcal{M}_{\mathrm{XZ}}(A) = \sum_{P \in \mathrm{Pauli}_{1,X,Z}(n)} \frac{1}{2^{k(P)}} \mathrm{tr}(AP)P, \qquad (C6)$$

where $\operatorname{Pauli}_{1,X,Z}(n)$ is the set of n-qubit Pauli strings constructed from only $\mathbb{1}s$, Xs, and Zs, and k(P) is the locality of Pauli string P. Note that we assume all Pauli strings P above are normalized with respect to the Hilbert-Schmidt inner product: $\operatorname{tr}(PP^{\dagger}) = 1$.

Proof. For the measurement primitive [25] assumed above, the random unitaries are uniformly sampled from \mathcal{U}_{XZ} , which is a subset of the Clifford group. As such, by Lemma 1 of Ref. [88], the eigenoperators of the corresponding measurement channel \mathcal{M}_{XZ} are Pauli strings. Therefore, we can express the action of the measurement channel as $M_{XZ}(P) = \lambda_P P$, where λ_P denotes the eigenvalue corresponding to the Pauli string eigenoperator P. It suffices to show that a k(P)-local Pauli string $P \in \text{Pauli}_{1,X,Z}(n)$ has eigenvalue $\lambda_P = 1/2^{k(P)}$, and all other Pauli strings $P \notin \text{Pauli}_{1,X,Z}(n)$ have eigenvalue 0.

To show this, we will derive an expression for λ_P . Let us start by rearranging the measurement channel. We can start with its original definition from Ref. [25], to get

$$\mathcal{M}(P) = \mathbb{E}_{U \sim \mathcal{U}_{XZ}} \sum_{b \in \{0,1\}^n} U^{\dagger} |b\rangle \langle b| U \operatorname{Tr} \left(U^{\dagger} |b\rangle \langle b| UP \right)$$

$$= \mathbb{E}_{U \sim \mathcal{U}_{XZ}} \operatorname{Tr}_B \left[\sum_{b \in \{0,1\}^n} U^{\dagger} |b\rangle \langle b| U \otimes (U^{\dagger} |b\rangle \langle b| UP)_B \right]$$

$$= \mathbb{E}_{U \sim \mathcal{U}_{XZ}} \operatorname{Tr}_B \left[U^{\dagger \otimes 2} \sum_{b} |bb\rangle \langle b, b| U^{\otimes 2} (\mathbb{1} \otimes P) \right]$$

$$= \mathbb{E}_{U \sim \mathcal{U}_{XZ}} \sum_{\tilde{P} \in \operatorname{Pauli}_{1,Z}(n)} U^{\dagger} \tilde{P} U \operatorname{Tr}(P U^{\dagger} \tilde{P} U).$$
(C7)

In the second line, we take the trace over subsystem "B", which we define to be the second subsystem (i.e., the one with P in the term $\mathbb{1} \otimes P$). In the third line, Pauli_{1,Z}(n) is the set of all normalized Pauli strings containing only $\mathbb{1}$ s and Zs. One can begin to see the transition from line two to line three by considering the single-qubit measurement channel. For a single qubit, we have $\sum_{b \in \{0,1\}} |b,b\rangle \langle b,b| = \frac{1}{2}(\mathbb{1}\mathbb{1} + ZZ)$. This can be extended to n qubits, where the n-qubit version $\sum_{b \in \{0,1\}^n} |b,b\rangle \langle b,b|$ becomes a sum of all length-n $\mathbb{1},Z$ strings tensored twice, normalized by a factor of $1/2^n$ [15]. This normalization does not appear in the final expression above because we absorb it into the two \tilde{P} s such that they are all normalized with respect to the Hilbert-Schmidt inner product.

We can further simplify this expression by noticing that $\text{Tr}(P \ U^{\dagger} \tilde{P} U)$ will be nonzero only when $P = U^{\dagger} \tilde{P} U$. With this observation, the expression becomes

$$\mathcal{M}(P) = \left(\mathbb{E}_{U \sim \mathcal{U}_{XZ}} \sum_{\substack{\tilde{P} \in \text{Pauli}_{1,Z}(n), \\ \text{s.t.}\tilde{P} = UPU^{\dagger}}} 1 \right) P.$$
(C8)

The quantity in parenthesis is simply an expression for λ_P , the eigenvalue corresponding to Pauli string P. It is the probability that one of our random unitaries $U \in \mathcal{U}_{XZ}$ rotates P into a Pauli string UPU^{\dagger} composed of only 1s and

Zs. We will separate the 4^n Pauli strings into two sets and consider their eigenvalues separately. First, let's consider the Pauli strings P in Pauli_{1,X,Z}(n). All sites in P which have a 1 will be trivially rotated back to 1 because U in UPU^{\dagger} separates into $U = \bigotimes_i U_i$. All that we must consider are the non-identity sites of P, and there are k(P) of them. Regardless of whether each site has an X or a Z, there is a 1/2 probability that each site is measured in the corresponding basis. Therefore, the eigenvalue of a k(P)-local Pauli string $P \in \text{Pauli}_{1,X,Z}(n)$ is $\lambda_P = 1/2^{k(P)}$.

Next, let us consider Pauli strings that are *not* in $\operatorname{Pauli}_{1,X,Z}(n)$. In other words, such Pauli strings P contain at least one single-site Y Pauli by definition. Since the random-XZ unitaries measure each qubit in either the X or the Z basis, there exists no U in \mathcal{U}_{XZ} that would rotate P to a string in $\operatorname{Pauli}_{1,Z}(n)$. Therefore, there is no $\tilde{P} \in \operatorname{Pauli}_{1,Z}(n)$ such that $\tilde{P} = UPU^{\dagger}$, so the eigenvalue of any Pauli string P containing at least one single-site Y Pauli is $\lambda_P = 0$.

Corollary 3 (Random-XZ shadow norm of Pauli strings). Assume a measurement primitive \mathcal{U}_{XZ} such that the probability density function p(U) is uniform over random-XZ unitaries. Then, the shadow norm of a Pauli string $P \in \text{Pauli}_{1,X,Z}(n)$ is

$$\|P\|_{shadow}^2 = 2^{k(P)},\tag{C9}$$

where P is normalized with respect to the Hilbert-Schmidt inner product and k(P) is the locality of P. Otherwise, for all Pauli strings $P \notin \text{Pauli}_{1,X,Z}(n)$, the shadow norm is $||P||_{shadow}^2 = \infty$, and, therefore, the observable cannot be learned.

Proof. The shadow norm of an eigenoperator O of the measurement channel \mathcal{M} is $||O||^2_{\text{shadow}} = 1/\lambda_O$, where λ_O is the eigenvalue of \mathcal{M} corresponding to O: $\mathcal{M}(O) = \lambda_O O$ (see Ref. [86] for a proof). The eigenoperators of our XZ measurement channel \mathcal{M}_{XZ} are the Pauli strings (by Lemma 1 of Ref. [88]). Therefore, the result follows from Proposition 2.

The intuition behind this result is as follows. Our measurment channel \mathcal{M}_{XZ} has a nontrivial null space, and as such, there are eigenoperator(s) of this channel whose eigenvalue is zero. By definition, these eigenoperators can *never* be learned and are invisible under these measurement settings [15]. In the language of probabilities, we do not have access to information about these eigenoperators because the probability of measuring in the eigenbases of these operators is zero. This is further elucidated by considering the shadow norm: the shadow norm of an operator in the null space of the measurement channel will be infinite. Therefore, no matter how many measurements are taken, it can never be learned.

Corollary 4 (Visible space of random-XZ measurements). For the set of unitaries in the ensemble \mathcal{U}_{XZ} defined above, the VisibleSpace(\mathcal{U}_{XZ}) on n qubits is the span of Pauli_{1,X,Z}(n), where Pauli_{1,X,Z}(n) is the set of n-qubit Pauli strings constructed from only 1s, Xs, and Zs.

Proof. Observation 4 of Ref. [15] states that if unitaries are sampled uniformly at random in the ensemble \mathcal{U} , then VisibleSpace(\mathcal{U}) is the image of the associated measurement channel \mathcal{M} . Following Proposition 2, the image of the random-XZ measurement channel is the span of Pauli_{1,X,Z}(n): all Pauli strings composed of single-site X, Z and 1s. Therefore, the VisibleSpace(\mathcal{U}_{XZ}) is the span of Pauli_{1,X,Z}(n).

This result also informs our intuition about the XZ visible space, the space of observables that we can estimate using the unitaries \mathcal{U}_{XZ} . Recall that each classical shadow takes the form $\hat{\rho} = \mathcal{M}_{XZ}^{-1}(U^{\dagger}|b\rangle\langle b|U)$. Definition 3 in Ref. [15] states that $U^{\dagger}|b\rangle\langle b|U \in \mathsf{VisibleSpace}(\mathcal{U}_{XZ})$, and per the observations above, the \mathcal{M} channel is effectively block diagonal with respect to the visible and invisible spaces. Therefore, $\hat{\rho} \in \mathsf{VisibleSpace}(\mathcal{U}_{XZ})$, and in the infinitemeasurement limit, the classical shadow will exactly be the projection of ρ onto $\mathsf{VisibleSpace}(\mathcal{U}_{XZ})$. As such, we can define the projected density matrix ρ_{XZ} , as the portion of the true density matrix ρ that lives in the $\mathsf{VisibleSpace}(\mathcal{U}_{XZ})$

$$\rho_{XZ} = \sum_{P \in \text{Pauli}_{1,X,Z}(n)} \text{Tr}(P\rho)P.$$
(C10)

This projected density matrix represents the part of ρ that is learnable using random-XZ measurements. Consequently, when we use classical shadows generated by the XZ unitaries \mathcal{U}_{XZ} to estimate $\langle O \rangle$ used in our regularization step in Appendix E, we find that we can only estimate operators O that live in the visible space.

Finally, we conclude this appendix by using the results of the derivations above to find the expression for the inverse of the measurement channel \mathcal{M}_{XZ} . This inverse measurement channel is a crucial tool because, as discussed at the start of this appendix, it is used to compute the classical shadow. Here we will show that this inverse measurement

$$\mathcal{M}_{XZ}(A) = \left(\mathcal{M}_{XZ}^{(1)}\right)^{\otimes n}(A), \quad \text{where} \quad \mathcal{M}_{XZ}^{(1)}(B) = \frac{1}{4}\left(\operatorname{Tr}(BX)X + \operatorname{Tr}(BZ)Z\right) + \frac{1}{2}\operatorname{Tr}(B)\mathbb{I}$$
(C11)

for all A and B in the *n*-qubit and single-qubit visible spaces, respectively. Explicitly, an operator in the single-qubit visible subspace can be decomposed as follows:

$$B = \frac{1}{2} \left(\operatorname{Tr}(BX)X + \operatorname{Tr}(BZ)Z + \operatorname{Tr}(B)\mathbb{I} \right).$$
(C12)

Next, we notice that each single-site channel $\mathcal{M}_{XZ}^{(1)}$ is not only independent, but also the same on every site. Thus, all that is left is to take the Moore-Penrose pseudoinverse of this single qubit channel. This yields the following result:

$$\mathcal{M}^{-1}{}^{(1)}_{XZ}(B) = \left(\operatorname{Tr}(BX)X + \operatorname{Tr}(BZ)Z + \frac{1}{2}\operatorname{Tr}(B)\mathbb{I} \right),$$
(C13)

which matches our result in equation Eq. (C3).

APPENDIX D: Tomographic completeness of randomized XZ measurements

The target states considered in this work are real pure states. We are interested in such states because they are eigenstates of a Hamiltonian whose matrix elements are purely real, e.g., a Hamiltonian that is equivalent to a composition of X and Z operators up to a global rotation. For example, if we have a (complex) Hamiltonian $H = \sum_{i,j} Z_i Z_j + \sum_i Y_i$, then we can perform a rotation identically at each site such that $Z_i \to Z_i$ and $Y_i \to X_i$. These Hamiltonians, which are interesting from the perspective of both condensed matter physics and quantum optimization [89–91], can suffer from a sign problem [92], so finding their ground states is a highly nontrivial task [93].

In this appendix, we will explore how and when we can fully characterize real pure states. These states can be expressed as a pure-state vector $|\psi\rangle = \sum_b c_b |b\rangle$ with real coefficients $c_b^* = c_b$. For simplicity, we will assume that the basis $\{|b\rangle\}_{b\in\{0,1\}^n}$ is the computational basis, unless indicated otherwise. We will demonstrate that we can fully characterize these states by only measuring each qubit in either the X or the Z basis (Lemma 5). While our numerical demonstrations are not always limited to measurements only in the X or Z bases, our motivation for specifically pursuing measurement schemes of this type is that randomized-XZ measurements are tomographically complete on the subspace of real and pure states (Theorem 1).

Below, this appendix is separated into two parts. First, we will provide an overarching discussion of our formal statements in order to develop our intuition and contextualize them in our broader work. We will then provide rigorous proofs of our claims.

D.1. Discussion of results

As elaborated on in the previous appendix, we can parametrize the learning power of a measurement protocol by looking at the visible space. Recall that the visible space is the space of all operators O for which we can estimate the expectation value $\text{Tr}(\rho O)$. In particular, in Corollary 4, we defined the visible space corresponding to the unitaries in the random-XZ ensemble \mathcal{U}_{XZ} . We found that the visible space is *not* the entire operator space $\mathcal{B}(\mathcal{H})$. In other words, there are some observables that are invisible to us—they can never be learned. With random-XZ measurements, we can measure any Pauli string in $\text{Pauli}_{1,X,Z}(n)$, the set of *n*-qubit Pauli strings constructed from only single-site 1, X, and Z Paulis. Therefore, the visible space is the span of all such Pauli strings. It turns out that this XZ measurement scheme is sufficient for learning states that live in the visible space as well as those with some components outside the invisible space, as long as certain assumptions are satisfied. One such example of this is a major result of this work: randomized-XZ measurements are tomographically complete on the subspace of real and pure states (Theorem 1).

Real, pure states do not live entirely in VisibleSpace(\mathcal{U}_{XZ}). For example, the GHZ states have nontrivial support on Pauli strings with an even number of single-site Y Paulis. Even though these states do not live entirely in VisibleSpace(\mathcal{U}_{XZ}), we can still fully characterize them with access to the visible XZ observables. This is because we can rely on the assumption that the state is both real and pure. Therefore, even though real, pure states are not fully encapsulated in the visible space, the XZ measurement scheme is sufficient for fully characterizing these states. To gain some intuition, let us consider a real, pure two-qubit state vector $|\psi\rangle$. We can write this state as a density matrix and decompose it in terms of its support on all Pauli strings $P \in \text{Pauli}_{1,X,Y,Z}(n)$ as

$$\begin{aligned} |\psi\rangle\langle\psi| &= \frac{1}{2^n} \left(\sum_{P \in \text{Pauli}_{1,X,Y,Z}(n)} c_P P\right) \\ &= \frac{1}{2^n} \left(\sum_{P' \in \text{Pauli}_{1,X,Z}(n)} c_{P'} P' + c_{YY} YY\right). \end{aligned}$$
(D1)

Since the state is real, its support must be zero on strings with an odd number of Ys, and therefore, there is only one term remaining that contains single-site Y Paulis: the two-qubit string $Y \otimes Y$. As a result, this two-qubit state does not live entirely in the visible space because $Y \otimes Y \notin \text{VisibleSpace}(\mathcal{U}_{XZ})$. However, we can still estimate the coefficient c_{YY} using the purity condition: $\text{Tr}(|\psi\rangle\langle\psi|^2) = 1$. In other words, since our XZ measurements allow us to learn all coefficients except c_{YY} , we can solve for c_{YY} using this extra constraint. While, so far, we have been considering a real, pure n = 2 qubit state, it turns out that this holds true for any n-qubit real, pure state. The information from the XZ measurements allows us to extract the support of $|\psi\rangle\langle\psi|$ on Pauli strings containing an even number of Ys. This will be shown rigorously in the following subsection.

D.2. Proof of Theorem 1

In this subsection, we prove that XZ measurements, where each qubit is measured in either the X or the Z basis, are tomographically complete on the space of real pure states. Tomographic completeness is similar to informational completeness but defined with respect to state characterization. A measurement primitive, which samples $U \sim \mathcal{U}$ and subsequently measures the rotated state in the computational basis $\{|b\rangle : b \in \{0,1\}^n\}$, is tomographically complete if and only if for each $\rho \neq \sigma$, there exist $U \in \mathcal{U}$ and b such that $\langle b|U\rho U^{\dagger}|b\rangle \neq \langle b|U\sigma U^{\dagger}|b\rangle$ [25]. For our purposes, we are only interested in establishing tomographic completeness on a subspace of the density matrices, namely, all states that are both real and pure.

Theorem 1 states that real, pure states are tomographically complete under XZ measurements. While our theorem might at first be surprising, it demonstrates that there is always a unitary $U \sim \mathcal{U}_{XZ}$ such that its measurement will allow us to differentiate between real, pure states ρ and σ . Furthermore, it elucidates that the visible space does not immediately render an understanding of tomographic completeness in the presence of extra assumptions.

Before providing the rigorous details, we first outline a sketch of the proof. Theorem 1 follows immediately from Lemma 5, which shows that we can fully characterize a real state vector $|\psi\rangle$ up to a global phase using only XZ measurements. We prove Lemma 5 by induction on the number of qubits k. Our base case is a real, pure state on k = 1 qubits. This can be immediately learned by measuring in the X and Z bases. The inductive step (k qubits) involves first learning all the amplitudes $|c_b|$ of $|\psi\rangle$, which can again be expressed as

$$|\psi\rangle = \sum_{b \in \{0,1\}^n} c_b |b\rangle.$$
(D2)

We learn the amplitudes by measuring every qubit in the Z basis. Then, we use our inductive assumption to argue that we can always characterize all relative signs using XZ measurements. In particular, our technique is more powerful than simply proving learnability, because we explicitly provide an algorithm by which one can use these measurements to fully characterize a real, pure state up to a global sign. Therefore, we show that XZ measurements are tomographically complete on the space of real pure states. We leave the details of this algorithm to the interested reader (see below).

Lemma 5 (Random-XZ measurements are sufficient for characterizing a real, pure state). Up to a global sign, we can uniquely characterize any real state vector $|\psi\rangle$ using XZ measurements, formally defined as measuring any qubit in either the X or the Z basis. We represent our real pure state in the computational basis $\{|b\rangle\}_{b\in\{0,1\}^n}$ as $|\psi\rangle = \sum_b c_b |b\rangle$, where $c_b^* = c_b$ and $\sum_b |c_b|^2 = 1$. We uniquely characterize the state by learning each c_b up to a global sign.

Proof. To fully characterize the state vector $|\psi\rangle$, we must learn all coefficients $\{c_b\}$ up to a global sign. Each coefficient is characterized by its amplitude and sign: $c_b = \operatorname{sgn}(c_b) \times |c_b|$. We can extract the amplitudes by measuring all qubits in the Z basis – whereafter, it remains to learn the relative signs. The remainder of this proof will show how we can learn the relative signs with XZ measurements. We will proceed by induction on the number of qubits k.

Base case (k = 1 qubits). We extract the relative sign between $|0\rangle$ and $|1\rangle$ by learning the sign of the expectation value $\langle \psi | X | \psi \rangle$: sgn $(\langle \psi | X | \psi \rangle)$.

Inductive step (k qubits). Assume we can learn all the relative signs of a real, pure state on k-1 qubits with XZ measurements. We have a k-qubit state. If we measure our kth qubit in the computational Z basis, then, by the inductive assumption, we can learn all the relative signs in the $\{|0\rangle \otimes |\tilde{b}\rangle\}_{\tilde{b} \in \{0,1\}^{k-1}}$ subspace as well as in the $\{|1\rangle \otimes |\tilde{b}\rangle\}_{\tilde{b} \in \{0,1\}^{k-1}}$ subspace. Here, the tilde on \tilde{b} indicates that the bit-string is on k-1 qubits instead of k. Also note that when we project the kth qubit of our state vector $|\psi\rangle$ onto $|0\rangle$ or $|1\rangle$, the state on the remaining k-1 qubits is projected onto a real pure state.

What is, therefore, left to be shown is that we can learn the relative sign between these two subspaces. If the amplitudes of all states in either of the subspaces are zero, then the relative sign is inconsequential because we already know the state up to a global phase. Otherwise, there must exist at least one state in each subspace that has nonzero amplitude. Let us choose two computational basis states: $b^{(0)}$ from the $\{|0\rangle \otimes |\tilde{b}\rangle\}_{\tilde{b}}$ subspace and $b^{(1)}$ from the $\{|1\rangle \otimes |\tilde{b}\rangle\}_{\tilde{b}}$ subspace, respectively. We choose them such that they that have nonzero amplitude and, among all possible choices, they have the smallest Hamming weight $\mathbf{d}(b^{(0)}, b^{(1)})$; this Hamming weight must be greater than or equal to 1 because the *k*th bit will always differ. We can extract the relative sign of the two spaces by estimating the

observable $O \in \mathsf{VisibleSpace}(\mathcal{U}_{XZ})$ given by

$$O = \bigotimes_{i=1}^{k} O_i; \quad O_i = \begin{cases} \left| + \right\rangle \!\! \left\langle + \right| & \text{if } i = k, \\ \left| b_i^{(0)} \right\rangle \! \left\langle b_i^{(0)} \right| & \text{if } b_i^{(0)} = b_i^{(1)}, \\ X & \text{otherwise,} \end{cases}$$
(D3)

where $b_i^{(0)}$ is the *i*th bit of bit-string $b^{(0)}$, and likewise $b_i^{(1)}$ is the *i*th bit of $b^{(1)}$. Consider first the case when the Hamming weight is $\mathbf{d}(b^{(0)}, b^{(1)}) = 1$. Then, since the *k*th qubit must always differ between the two states, all of the remaining bits are the same: $\forall i \neq k, b_i^{(0)} = b_i^{(1)}$. Therefore, $\langle \psi | O | \psi \rangle$ takes the form $|c_{b^{(0)}} + c_{b^{(1)}}|^2$. Since we already know $|c_{b^{(0)}}|$ and $|c_{b^{(1)}}|$, and by assumption they are nonzero, we can uniquely identify the relative sign between the subspaces. Next, we consider the case when the Hamming weight is $\mathbf{d}(b^{(0)}, b^{(1)}) > 1$ and learn $\langle \psi | O | \psi \rangle$. Notice that O's projection of $|\psi\rangle$ onto the matching bits in $b^{(0)}$ and $b^{(1)}$ leaves only two (nonzero amplitude) basis state components:

$$\left(\otimes_{i \text{ s.t. } b_{i}^{(0)}=b_{i}^{(1)}}\left\langle b_{i}^{(0)}\right|\right)|\psi\rangle = c_{b^{(0)}}\left(\otimes_{i \text{ s.t. } b_{i}^{(0)}=b_{i}^{(1)}}\left\langle b_{i}^{(0)}\right|\right)|b^{(0)}\rangle + c_{b^{(1)}}\left(\otimes_{i \text{ s.t. } b_{i}^{(0)}=b_{i}^{(1)}}\left\langle b_{i}^{(0)}\right|\right)|b^{(1)}\rangle. \tag{D4}$$

This is because we chose the pair $b^{(0)}$ and $b^{(1)}$ such that they have the smallest Hamming weight among all possible pairings. Then, the X operator flips any remaining bits at $i \neq k$, and we obtain a state with precisely swapped amplitudes

$$\left(\otimes_{i \text{ s.t. } b_{i}^{(0)}=b_{i}^{(1)}}\left\langle b_{i}^{(0)}\right|\right)\left(c_{b^{(0)}}|b^{(1)}\rangle+c_{b^{(1)}}|b^{(0)}\rangle\right).$$
(D5)

Consequently, evaluating the expectation value yields $\langle \psi | O | \psi \rangle = c_{h^{(0)}} \times c_{h^{(1)}}$. Since $|c_{h^{(0)}}|$ and $|c_{h^{(1)}}|$ are known and (by assumption) nonzero, we can uniquely identify the relative sign between the subspaces.

Theorem 1 (Tomographic completeness of random-XZ measurements on the space of real, pure states). A measurement primitive which samples $U \sim \mathcal{U}_{XZ}$ and subsequently measures the rotated state in the computational basis $\{|b\rangle : b \in \{0,1\}^n\}$ is tomographically complete [25] on the space of real, pure states. Namely, for all states $|\psi\rangle\langle\psi|\neq|\phi\rangle\langle\phi|$, there exist $U\in\mathcal{U}_{XZ}$ and b such that $\langle b|U\rho U^{\dagger}|b\rangle\neq\langle b|U\sigma U^{\dagger}|b\rangle$.

Proof. We proceed by contradiction. Assume that this measurement primitive is not tomographically complete on the space of real, pure states. Then, there exist real, pure states $|\psi\rangle\langle\psi|\neq|\phi\rangle\langle\phi|$ such that

$$\forall U, b \ \langle b|U|\psi\rangle\langle\psi|U^{\dagger}|b\rangle = \langle b|U|\phi\rangle\langle\phi|U^{\dagger}|b\rangle. \tag{D6}$$

In other words, after the application of any U, the resulting computational basis state probabilities will always be the same: we cannot differentiate between $|\psi\rangle\langle\psi|$ and $|\phi\rangle\langle\phi|$ with this measurement primitive. However, Lemma 5 shows that this same measurement primitive can uniquely characterize, up to a global phase, both $|\psi\rangle$ and $|\phi\rangle$. Therefore, since these states differ by more than just some global phase $(|\psi\rangle\langle\psi|\neq|\phi\rangle\langle\phi|)$ by definition), we arrive at a contradiction.

APPENDIX E: Loss function

This appendix presents the details of the loss function in Eq. (5). In the first subsection, we provide some background for maximum likelihood estimation and derive its contribution as the first term in our loss function. In the next subsection, we explicitly state the physical observables that we use for regularization in our numerical results. While keeping most of the section self-contained, some of the notations and terminologies used have been set up in Appendix C.

E.1. Maximum likelihood estimation: Randomized Kullback-Leibler divergence

First, we review the standard framework of maximum likelihood estimation (MLE), before applying such a loss function to our specific setting of quantum measurements. In general, we observe data from an underlying probability distribution $\mathcal{D}: X \to \mathbb{R}$ on the dataset domain X, and our goal is to find the most likely model with parameters A denoted as $\mathcal{Q}_A: X \to \mathbb{R}$ that is consistent with our observations. To find such a model, we would want to variationally optimize for the model's parameters A by minimizing the Kullback–Leibler (KL) divergence

$$K(\mathcal{D}|\mathcal{Q}_A) = -\sum_{x \in X} \mathcal{D}(x) \log \frac{\mathcal{Q}_A(x)}{\mathcal{D}(x)}.$$
(E1)

Since the data distribution \mathcal{D} is a constant, we just need to minimize the cross-entropy contribution to the KL divergence

$$\mathbb{H}(\mathcal{D}|\mathcal{Q}_A) = -\sum_{x \in X} \mathcal{D}(x) \log \mathcal{Q}_A(x).$$
(E2)

More specifically, for our purposes, the dataset consists of bit-string measurements of a target state $|\phi\rangle\langle\phi|$ in the computational basis after applying a unitary U. In a randomized measurement scheme, such as the XZ-shadow measurements discussed in Appendix C, we first draw a unitary from an ensemble \mathcal{U} with a uniform probability over a set $\{U\}$. Then, we measure the bit-string b in the computational basis. Taken together, we sample a data point, which is a tuple of the sampled unitary and the measured bit-string $x_i = (U_i, b_i)$, with probability

$$\mathcal{D}[(U,b)] = \frac{1}{|\{U\}|} \times \langle b|U|\phi\rangle\!\langle\phi|U^{\dagger}|b\rangle.$$
(E3)

Similarly, we define the probability of observing this data point from a model parametrized by A as Q_A . Then, the *randomized* KL divergence is

$$K(\mathcal{D}|\mathcal{Q}_A) = -\mathbb{E}_{U \sim \mathcal{U}} \sum_{b \in \{0,1\}^n} \mathcal{D}(U,b) \log \frac{\mathcal{Q}_A(U,b)}{\mathcal{D}(U,b)}.$$
 (E4a)

As in Eq. (E2), we can drop the constant factor from the data distribution and obtain the randomized cross entropy or the infinite-sample loss function

$$\mathbb{H}(\mathcal{D}|\mathcal{Q}_A) = -\mathbb{E}_{U \sim \mathcal{U}} \sum_{b \in \{0,1\}^n} \mathcal{D}(U,b) \log \left[\mathcal{Q}_A(U,b)\right].$$
(E5)

For N-samples i.i.d. drawn from \mathcal{D} as our dataset denoted by $D = \{(U_i, b_i)\}_{i=1}^N$, this leads to the negative log-likelihood (NLL) contribution to our loss function which is

$$L_{\text{NLL}}^D(A) = -\frac{1}{N} \sum_{i=1}^N \log \mathcal{Q}_A(b_i, U_i).$$
(E6)

In this work, we focus on a model distribution given by a MPS $|\psi_{\chi}(A)\rangle$ with tensors A and bond dimension χ such that

$$\mathcal{Q}_A(U,b) = \frac{1}{|\{U\}|} \times |\langle b|U|\psi_{\chi}(A)\rangle|^2.$$
(E7)

Dropping the constant normalization factor from the uniform distribution over the unitary ensemble gives us the MLE contribution described in Eq. (5) of the main text.

E.2. Regularization

In this section, we discuss the second contribution to our loss function, which additionally regularizes the optimization landscape of our training process. Note that these estimations from finite-N samples have deviations from the true physical expectation values; in general, we do not want our regularization strength to be dominant.

In the main text, we numerically explore different regularization strengths. Moreover, we have focused on two measurement schemes: global-XZ measurements, and random-XZ measurements. For the former, we measure bitstrings in two fixed bases $X^{\otimes n}$ and $Z^{\otimes n}$, and we can estimate certain physical observables using empirical means. For the latter, using the toolbox of randomized measurements, we apply random unitaries from an ensemble and measure bit-strings in the computational basis. Depending on the choice of the unitary architectures, different physical properties can be estimated. Performing classical shadow tomography using random-XZ measurements—as detailed in Appendix C—allows us to estimate any observables in the visible space of our measurement. We now discuss our choices of regularizers for the two systems considered in the main text:

Perturbed surface code. Here, we use the stabilizers S of the exact surface code defined in Appendix F.1. The intuition behind such a choice is that we know that the stabilizer expectations will change continuously with the perturbation and that these operators would have large weight in the density matrix. The regularization contribution to the loss function is

$$R(A) = \sum_{S \in \mathcal{S}} |\langle S \rangle_N - \langle \psi(A) | S | \psi(A) \rangle|^2,$$
(E8)

where $\langle S \rangle_N$ are estimates from the measurements and the second term is the model's prediction. The regularizer R(A) would be modulated by β , which is a hyperparameter controlling its strength. The knowledge of these specific observables allows us to estimate their expectation values in two measurement schemes.

With global-XZ measurements in the bases $X^{\otimes n}$ and $Z^{\otimes n}$, we estimate the stabilizer expectations using empirical means. For instance, to estimate $S_{Z,3} = Z_1 Z_2$ from fixed $Z^{\otimes n}$ basis measurements $\{Z_1^{(i)} Z_2^{(i)} \cdots Z_n^{(i)}\}_{i=1}^N$, we average over the measurement outcomes

$$\langle Z_1 Z_2 \rangle_{N, \text{Global}} = -\frac{1}{N} \sum_{i=1}^N Z_1^{(i)} Z_2^{(i)}.$$
 (E9)

In the scheme of random-XZ measurements, we use classical shadow tomography for such estimation. In particular, we can estimate any k-local observable $S_{Z,l}$ using the classical shadows $\hat{\rho}_N$ defined in Eq. (C5) as

$$\langle S_{Z,l} \rangle_{N, \text{Random}} = \text{Tr}(S_{Z,l} \hat{\rho}_N).$$
 (E10)

As discussed in Appendix C, the number of samples required to estimate such observables accurately scales exponentially as $\mathcal{O}(2^k)$. Here, a key surprise of our protocol is that even though we only regularize with physically local observables $k \ll n$, we are able to see an improvement in global quantities as the infidelity is reduced.

Rydberg atom arrays. More generally, if we do not know a priori which operators are important—as is indeed the case for non-fixed-point states realizable in quantum simulators such as Rydberg atom arrays—we would need to consider all the operators that one can estimate using a finite number of samples. The intuition behind the regularization contribution is that capturing more physical observables accurately is consistent with improvement of the state fidelity. In fact, describing a subsystem with a size comparable to the correlation length guarantees an efficient learning of a one-dimensional injective quantum state using MPS tomography [18, 19]. Here, using random-XZ measurements, the space of the operators we can estimate is spanned by all combinations of 1, X, Z strings normalized with respect to the Hilbert-Schmidt inner product Pauli_{1,X,Z}(n). This motivates us to define the projected density matrix

$$\rho_{XZ} = \sum_{P \in \text{Pauli}_{1,X,Z}(n)} \text{Tr}(P\rho)P.$$
(E11)

Now, the classical shadows learnable using random-XZ measurements [cf. Eq. (C3)] converge to this projection in the infinite-data limit, i.e.,

$$\rho_{XZ} = \lim_{N \to \infty} \hat{\rho}_m. \tag{E12}$$

In the finite-sample regime, we regularize on observables that are supported on subsystems B of 6 qubits, described by a classical shadow

$$\hat{\rho}_N^B = \otimes_{i \in B} \left(2U^{\alpha_i \dagger} |b_i\rangle \langle b_i | U^{\alpha_i} - \frac{\mathbb{I}}{2} \right).$$
(E13)

Such a projection can also be computed from our MPS model $|\psi(A)\rangle$. First, we trace out the complement \bar{B}

$$\tilde{\rho}^B(A) = \Pr_B(|\psi(A)\rangle\!\langle\psi(A)|). \tag{E14}$$

Next, we project the reduced density matrix onto $\mathsf{VisibleSpace}(\mathcal{U}_{XZ})$ as

$$\tilde{\rho}_{XZ}^B(A) = \sum_{P \in \text{Pauli}_{1,X,Z}(n)} \text{Tr}(P \tilde{\rho}^B(A)) P.$$
(E15)

Putting all the pieces together, we regularize over these projected density matrices—supported on all the unit cells that cover the lattice—using the second term in our loss function

$$R(A) = \|\hat{\rho}_N^B - \tilde{\rho}_{XZ}^B(A)\|_2.$$
(E16)

Here, the matrix norm is given by the Schatten-2 norm (or Frobenius norm)

$$\|O\|_{2} = \left[\sum_{i,j} O_{i,j}^{2}\right]^{\frac{1}{2}}.$$
(E17)

APPENDIX F: Dataset and state generation with DMRG

In this appendix, we discuss the generation of our numerical dataset using Gibbs sampling from the target MPS. The target MPS is variationally optimized using the *density-matrix renormalization group* (DMRG) algorithm. Here, we outline a snippet of the numerical code written in the Python language. The calculations are performed using an open source library quimb [94]. All of the code used to generate our numerical results is available at: https://github.com/teng10/tn-shadow-qst.

F.1. Perturbed surface code

In this section, we discuss the details of our surface code numerics and provide the relevant code snippets used to generate our dataset. The perturbed surface code Hamiltonian for a system of n qubits is defined as

$$H_{\rm sc} = -\sum_{l} S_{Z,l} - \sum_{l} S_{X,l} - h_z \sum_{i} Z_i, \quad S_{Z,l} = \prod_{i \in \Box_l} Z_i, \quad S_{X,l} = \prod_{i \in \Box_l} X_i, \tag{F1}$$

where $S = \{S_{Z,l}, S_{X,l}, l = 1, \dots, \frac{n-1}{2}\}$ is the set of n-1 stabilizers around a square, as illustrated in Fig. 2(a). To construct the Hamiltonian for an $(L_x, L_y) = (3, 3)$ system, the following lines create an object called *PhysicalSystem* and generate its matrix product operator representation.

```
surface_code = physical_systems.SurfaceCode(3, 3)
surface_code_mpo = surface_code.get_ham()
```

To run the DMRG algorithm, first, a random MPS is selected as an initial state. Then, the MPO and this initial state are used to solve for an MPS approximation of the ground state with a given bond dimension.

```
import quimb.tensor as qtn
size = 9
bond_dim = 10
mps = qtn.MPS_rand_state(size, bond_dim)
dmrg = qtn.DMRG2(surface_code_mpo, bond_dims=bond_dim, p0=mps)
dmrg.solve()
result_mps = dmrg.state
```

Finally, to generate the dataset for training, we first select a measurement scheme (e.g., random-XZ [see Appendix D]), and specify the number of samples required.

```
1 init_seed = 42 # for reproducibility
2 sampling_method = 'x_or_z_basis_sampler' # random XZ measurements
3 num_samples = 1000
4 ds = run_data_generation._run_data_generation(init_seed, num_samples, sampling_method, mps)
```

F.2. Ruby-lattice quantum spin liquid

In this section, we include the details of our DMRG calculation for preparing ground states of the Rydberg Hamiltonian on a ruby lattice [7]. Specifically, we consider a system of atoms arrayed on a ruby lattice with an aspect ratio of $\rho = \sqrt{3}$ [Supplementary Fig. 2] and cylindrical boundary conditions. The Hamiltonian is given by

$$H_{\text{Ryd}} = \frac{\Omega}{2} \sum_{\ell} X_{\ell} - \Delta \sum_{\ell} \frac{1+Z_{\ell}}{2} + \frac{1}{2} \sum_{\ell,\ell'} \frac{V_{\ell,\ell'}}{4} (1+Z_{\ell})(1+Z_{\ell'}).$$
(F2)

Denoting by $|b\rangle = |\pm\rangle$ the eigenstate of the Pauli Z operator in the computational basis (i.e., $Z|b\rangle = b|b\rangle$), we work with a sign convention such that b = +1 (-1) corresponds to an atom being in the Rydberg (ground) state.

Depending on the ratio of the detuning to the Rabi frequency, Δ/Ω , different quantum phases can be prepared [51, 52]. The native van der Waals interaction potential decays rapidly with the distance as $V_{\ell,\ell'} = V_0/|r_\ell - r_{\ell'}|^6$. Choosing the interaction strength such that $\Omega \ll V_{\ell,\ell'}$, the strong repulsion between neighboring atoms implements the Rydberg blockade effect, such that no two atoms are allowed to be simultaneously excited within a blockade radius of $R_b = (V_0/\Omega)^{\frac{1}{6}}$. While, in principle, the interactions act between all pairs of atoms, in practice, an approximated potential

obtained by truncating the long-range tails of the interaction is known to capture the essential physics of the spinliquid phase [52]. In our DMRG simulations, we only retain interactions up to third-nearest neighbors (r = 2a, where a is the lattice spacing), as illustrated in Fig. 3(a). Setting $\Omega = 1$, we approximate the potential as a step function

$$V_{\ell,\ell'} = \begin{cases} 47\Omega, & |r_{\ell} - r_{\ell'}| \le 2a, \\ 0, & \text{otherwise.} \end{cases}$$
(F3)

Note that unlike the so-called PXP model [52], in which the Hilbert space is constrained to forbid simultaneous excitations of neighboring atoms, here, the blockade is only imposed as an energetic penalty. The interaction scale in Eq. (F3) follows from the choice of the blockade radius, $R_b = 3.8a$. Moreover, the Rabi frequency Ω is positive, so the present Hamiltonian is nonstoquastic [32], without guarantees of positive wavefunction amplitudes in the computational basis.

In the present work, we focus on the ruby-lattice Rydberg Hamiltonian with periodic boundary conditions along the y-direction and open boundary conditions along the x-axis. We perform numerical simulations for a lattice of size $(L_x, L_y) = (4, 2)$; the unit cell has 6 basis atoms, so the total number of qubits is n = 48. To mitigate finite-size effects, we compensate for the fewer neighboring atoms of any site along the boundary by adding a local field $h_{bd} = -0.6\Omega$ to the regions shaded by the blue boxes in Supplementary Fig. 2:

$$H_{\text{Ryd, bd}} = H_{\text{Ryd}} - h_{\text{bd}} \sum_{\ell \in \partial} \frac{1 + Z_{\ell}}{2}.$$
 (F4)

```
1 delta = 1.7 # detuning frequency for spin liquid state
```

- 2 boundary_z_field = -0.6 # adding an onsite field to mitigate boundary effects
- 3 rydberg = physical_systems.RubyRydbergPXP(4, 2)
- 4 rydberg_mpo = rydberg.get_ham(delta, boundary_z_field)



Supplementary Figure 2. Rydberg atom array on a cylinder. The ruby lattice is placed on a cylinder, imposing periodic (open) boundary conditions along the y(x)-axis. The numbers label the snake-like path for the tensors in the MPS representation of the ground state. An onsite field applied to the sites in the blue boxes compensates for the reduced coordination number of atoms situated along the open boundaries, thereby preventing them from getting pinned.

APPENDIX G: Training details

In this appendix, we discuss our approach to optimizing *matrix product state* (MPS) tensors through simultaneous updates of all tensor components by gradient descent. We then highlight the selection process for models likely to converge in the absence of direct access to the infidelity. This can be done by partitioning datasets into training and test datasets and post selecting models that successfully generalize on the test dataset. This method's efficacy is illustrated through the analysis of training and test *negative log-likelihood* (NLL) metrics. As evidenced in the surface code example, we evaluate the model performance's with sample size to ensure robust generalization.

G.1. Optimization scheme via gradient descent

We approach the optimization of *matrix product state* (MPS) tensors by updating all tensors simultaneously. While the target states we are learning are real states, we use MPS tensors with complex variables because we have found that optimization over complex tensors is more robust, as also observed in Ref. [27]. The minimization of our loss function involves the logarithm of sums over expectation values of $U_i^{\dagger}|b_i\rangle\langle b_i|U_i$ defined by the measured unitary-bit-strings $D = \{(U_i, b_i)\}_{i=1}^N$,

$$L_{\rm NLL}^D(A) = -\frac{1}{N} \sum_{i=1}^N \log\left(\langle \psi(A^*) | U_i^{\dagger} | b_i \rangle \langle b_i | U_i | \psi(A) \rangle\right). \tag{G1}$$

For simplicity, we have considered the loss function without a regularization with $\beta = 0$. While $U_i^{\dagger} |b_i\rangle \langle b_i | U_i$ can be represented as a *matrix product operator* (MPO), this loss function is different from the traditional eigenvalue problem of finding the ground state of a local Hamiltonian typically considered in DMRG calculations. Given the presence of the logarithmic term in our cost function, previous work has considered a "DMRG-like" sweeping algorithm [26]. However, to simplify the training process, we have adopted global updates of MPS tensors over local sweeping updates and find that this strategy suffices for our purpose. Future improvements are possible by adopting optimization techniques on the MPS manifold [84, 85] as discussed in B.7, particularly focusing on the projection of gauge degrees of freedom to refine tensor updates.

Note that in order to minimize the loss function, we calculate gradients for all MPS tensors in the direction opposite to that dictated by the derivative of the complex conjugate 7

$$A' = A - \delta \frac{\mathrm{d}L}{\mathrm{d}\bar{A}}.\tag{G2}$$

To integrate with machine-learning optimization techniques, our procedure employs a two-step strategy: initial descent towards an approximate global minimum using first-order *stochastic gradient descent* (SGD), followed by refinement through a second-order quasi-Newton's method, the *limited-memory Broyden-Fletcher-Goldfarb-Shanno* (L-BFGS) algorithm. This combination allows for efficient preliminary optimization followed by detailed adjustments.

G.2. Model selection

In applying our protocol to a laboratory state, in general, we would not have access to the infidelity. Here, we outline the steps to select models that are likely to converge without access to the fidelity. In the context of machine learning, a common approach is to divide the full dataset into the training and test datasets. First, we train the MPS model on only the training dataset partition and then evaluate the *negative log-likelihood* (NLL) in Eq. (E6) using the test dataset. For the model to not be "overfitting" the training dataset, we would want both the training and test NLLs to be small. Specifically, we consider the surface code numerics as an example below. In Supplementary Fig. 3, we plot both the training NLL and the test NLL for different numbers of samples. We observe that for a small number of samples, while we achieve a lower training NLL, the test NLL is higher, indicating overfitting of the MPS to the training measurements. On the other hand, an MPS trained with a large number of samples achieves a lower infidelity and hence, better generalization.

⁷ To see this, let $f(z, \bar{z}) : \mathbb{C} \to \mathbb{R}$ be a real-valued function of variables z = x + iy and $\bar{z} = x - iy$. Note that the loss function could be nonholomorphic. An update to minimize the loss function is given by $(x' + iy') = x - \delta \frac{\partial f}{\partial x} + i(y - \delta \frac{\partial f}{\partial x})$. This is the same as $z' = z - \frac{\delta}{2} \frac{\partial f}{\partial \bar{z}}$. Numerically, we compute $\frac{\partial L}{\partial A}$ using JAX and take the direction given by its complex conjugate $(\frac{\partial L}{\partial A})^*$.

A different approach is to estimate the fidelity of the learned MPS model with the laboratory state [26] by using classical shadow tomography with global random unitary ensembles, which could be hard to implement in practice. Finally, we remark that a recent protocol would also allow us to efficiently certify the fidelity using single-qubit measurements given a query model [42], in our case, an MPS.



Supplementary Figure 3. Model selection by comparing training and test losses for a surface code with different system sizes and measurement schemes. On the dashed lines, the training NLL = test NLL; ideally, the model should be selected along the dashed line *and* with a small training NLL. The models which have small NLLs and are close to the dashed line yield high fidelities (darker colors).