

Neural quantum states for entanglement-depth certification from randomized Pauli measurements

Marcin Płodzień 

Qilimanjaro Quantum Tech, Carrer de Veneçuela 74, 08019 Barcelona, Spain



(Received 22 December 2025; accepted 21 May 2026; published 8 June 2026)

Entanglement depth quantifies how many qubits share genuine multipartite entanglement, but certification typically relies on tailored witnesses or full tomography, both of which scale poorly with system size. We recast entanglement-depth and non- k -separability certification as likelihood-based model selection among neural quantum states whose architecture enforces a chosen entanglement constraint. A hierarchy of separable neural quantum states is trained on finite-shot local Pauli outcomes and compared against an unconstrained reference model trained on the same data. When all constrained models are statistically disfavored, the data certify entanglement beyond the imposed limit directly from measurement statistics, without reconstructing the density matrix. We validate the method on simulated six- and ten-qubit datasets targeting Greenberger–Horne–Zeilinger state (GHZ), Dicke, and Bell-pair states, and demonstrate robustness for mixed states under local noise. Finally, we discuss lightweight interpretability diagnostics derived from trained parameters that expose coarse entanglement patterns and qubit groupings directly from bitstring statistics.

DOI: [10.1103/29hj-7w41](https://doi.org/10.1103/29hj-7w41)

I. INTRODUCTION

Multipartite entanglement constitutes a fundamental resource for quantum information processing, metrology, and simulation. However, characterizing its structure in many-body systems remains a formidable experimental challenge. Beyond the simple dichotomy between separable and entangled states, a more granular measure is the *entanglement depth*: the minimum number of particles that must be genuinely entangled to produce the state [1–7]. This concept is formally captured by the notion of l producibility and closely related notions of k separability used in multipartite entanglement classification; a state is l producible if it can be described—even probabilistically—as a mixture of pure states that factorize into clusters of size at most l [1,8].

Conventional characterization strategies typically rely on either tailored entanglement witnesses or quantum tomography followed by postprocessing [1,4,8,9]. Witness-based methods generally necessitate carefully engineered measurement settings and often rely on prior assumptions about the target state structure, such as graph-state or symmetric Dicke-state symmetries. Conversely, tomographic techniques suffer from a scaling problem: the number of required measurements and the classical postprocessing costs grow exponentially with system size, rendering them intractable beyond a handful of qubits. Furthermore, modern quantum devices are often best

treated as black boxes generating measurement outcomes in local bases. Ideally, one would extract entanglement information directly from these statistics without committing to a costly full reconstruction of the density matrix.

Machine learning offers a complementary avenue for data-driven assessment of entanglement properties [10–15]. A common supervised setting trains classifiers to distinguish separable from entangled states using tomographic or correlator-level features [16–22], while unsupervised variants aim to cluster entangled and separable data without labeled examples [23]. Beyond binary detection, learned models have been employed to infer finer aspects of multipartite structure, including partition- or depth-related organization [24–27]. In parallel, regression-based schemes estimate entanglement measures from experimentally accessible primitives, such as collective or weak measurements [28–33]. Related research replaces explicit reconstruction with learned generative models for basis-resolved outcomes, which can be queried for downstream state properties [34,35], and kernel methods have also been explored for entanglement detection [36] (for review see Refs. [37,38]).

In this work, we pursue a data-driven strategy distinct from supervised classification. Assuming access only to bitstrings from projective measurements in local Pauli bases, we investigate whether one can certify non- k -separability and bound the entanglement depth directly from how well different generative models explain these data. Neural Quantum States (NQS) provide a natural framework for this task [39–43]. As highly expressive variational representations of many-body wave functions, NQS trained on measurement data serve as powerful generative models of the outcome distribution. This perspective has enabled neural-network approaches to quantum state tomography for pure states [44,45] and mixed states [46], with successful demonstrations on experimental platforms [47–49].

*Contact author: marcin.plodzien@qilimanjaro.tech

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. Protocol for entanglement-depth certification via generative modeling. Bitstrings obtained from randomized Pauli measurements on the target device are used to train a hierarchy of NQS models. The hierarchy consists of an unconstrained reference model ($\mathcal{A}_{\text{full}}$) and various Separable NQS architectures ($\mathcal{A}_{\mathcal{P}}$), where the bipartite connectivity is masked to enforce factorization into clusters of size at most l . By minimizing the negative log-likelihood (NLL) $\mathcal{L}^*$, we compute the likelihood gap $\Delta_{\mathcal{P}} = \mathcal{L}_{\mathcal{A}_{\mathcal{P}}}^* - \mathcal{L}_{\mathcal{A}_{\text{full}}}^*$. A persistent gap ($\Delta_{\mathcal{P}} \gg 0$) signifies that the data cannot be explained by any state within the l -producible hypothesis class, thereby certifying an entanglement depth $d_e > l$ without reconstructing the full density matrix.

We present a practical protocol for certifying non- k -separability and entanglement depth directly from finite-shot local Pauli measurement data, without making strong assumptions about the state. We train a fully entangling NQS and a hierarchy of separable neural quantum state (SNQS) ansätze—representing different multipartition structures—on the same bitstring data. The resulting hierarchy of optimized negative log-likelihoods serves as an operational witness of entanglement depth: likelihood gaps that persist for all architectures compatible with depth l rule out l producibility, thereby certifying $d_e > l$ without density-matrix reconstruction (see Fig. 1). Furthermore, we demonstrate that the parameters of the unrestricted NQS admit simple, interpretable postprocessing. Effective visible-visible couplings and affinity matrices derived from the trained restricted Boltzmann machine (RBM) weights yield informative signatures of the underlying structure—such as Bell-pair dimers or GHZ/Dicke clusters. We apply this methodology to three families of nontrivial multiqubit states, both pure and mixed, certifying entanglement depth in systems of up to ten qubits.

The paper is organized as follows. In Sec. II, we review the notions of k separability, l producibility, and entanglement depth, and introduce separable neural quantum state architectures and their combinatorial structure. In Sec. III, we describe the measurement protocol, the negative log-likelihood training objective, and the likelihood-gap criterion for non- k -separability. Section IV presents numerical results for both pure and mixed multiqubit targets, including GHZ, Bell, and Dicke clusters under local amplitude damping. In Sec. V, we analyze the trained unrestricted NQS and introduce interpretable diagnostics that recover coarse entanglement structure directly from the learned weights. We conclude in Sec. VI.

II. PRELIMINARIES

A. Multipartite separability, l producibility, and entanglement depth

We consider an N -qubit system with Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes N}$. Multipartite entanglement can be organized according to how the parties can be partitioned into smaller groups, clusters, that share quantum correlations. The language of k separability and l producibility is standard in the entanglement-depth literature [5–7].

Let $\{1, \dots, N\}$ denote the set of qubit labels and $\mathcal{P} = \{C_1, \dots, C_k\}$ a partition into k nonempty, disjoint subsets C_b with union $\{1, \dots, N\}$. A pure state $|\Psi\rangle \in \mathcal{H}$ is called k separable with respect to $\mathcal{P}$ if it factorizes as

$$|\Psi\rangle = \bigotimes_{b=1}^k |\phi_{C_b}\rangle, \quad |\phi_{C_b}\rangle \in \bigotimes_{i \in C_b} \mathbb{C}^2. \quad (1)$$

If there exists at least one partition $\mathcal{P}$ with $k \geq 2$ such that this holds, we say that $|\Psi\rangle$ is k separable. The extreme cases are $k = N$, where every cluster contains a single qubit and the state is fully separable, and $k = 1$, where the whole register forms one cluster and the state may be genuinely N -partite entangled. A pure state that is not k separable for any $k \geq 2$ is called genuinely multipartite entangled (note that $k = 1$ means no partition with two or more clusters exists under which the state factorizes) [1,4].

For mixed states, the definition is extended by convex combinations. A density operator ρ on $\mathcal{H}$ is called k separable if it admits a decomposition

$$\rho = \sum_{\lambda} p(\lambda) \bigotimes_{b=1}^k \rho_{C_b}^{(b)}(\lambda), \quad p(\lambda) \geq 0, \quad \sum_{\lambda} p(\lambda) = 1, \quad (2)$$

where for each label λ the operators $\rho_{C_b}^{(b)}(\lambda)$ act on the clusters C_b of some partition $\mathcal{P}(\lambda)$ of $\{1, \dots, N\}$ into k groups. Importantly, the partition $\mathcal{P}(\lambda)$ is in general allowed to depend on λ ; Eq. (2) thus describes classical mixtures of possibly different k -separable pure states. States that cannot be written in the form (2) are called non- k -separable. In particular, nonbiseparable states ($k = 2$) are genuinely multipartite entangled [1,4,8].

While k separability counts the number of entangled blocks, it is often more natural to characterize *how many* qubits are entangled within a block. This leads to the notion of l producibility and entanglement depth, originally introduced in the context of spin-squeezed and many-body states [1–3]. To avoid notational ambiguity, we use distinct symbols in the two definitions: k denotes the number of clusters in a k -separable partition, while l bounds the maximal cluster size in l producibility. In much of the literature, both roles are assigned to the same letter; our convention follows the usage of Refs. [1,4,7] while eliminating the resulting notational clash.

A pure state $|\Psi\rangle$ is called l producible if it can be written as a product of clusters containing at most l qubits each,

$$|\Psi\rangle = \bigotimes_{b=1}^m |\phi_{C_b}\rangle, \quad |C_b| \leq l, \quad \bigcup_b C_b = \{1, \dots, N\}. \quad (3)$$

For mixed states, ρ is l producible if it is a convex combination of l -producible pure states,

$$\rho = \sum_{\lambda} p(\lambda) |\Psi_{\lambda}\rangle\langle\Psi_{\lambda}|, \quad |\Psi_{\lambda}\rangle \text{ } l \text{ producible for every } \lambda. \quad (4)$$

Intuitively, l -producible states can be prepared by entangling at most l parties at a time. Any state that is not l producible is said to be *non- l -producible* and has entanglement extending over at least $l + 1$ qubits.

The *entanglement depth* $d_e(\rho)$ is then defined as the smallest integer l such that ρ is l producible,

$$d_e(\rho) = \min\{l \mid \rho \text{ is } l \text{ producible}\}. \quad (5)$$

Fully separable states have $d_e = 1$; states with $d_e = 2$ contain only pairwise entanglement (e.g., products of Bell pairs), and genuinely multipartite entangled states on N qubits have $d_e = N$ and are non- $(N - 1)$ -producible.

In this work, we exploit the correspondence between such partition structures and the connectivity masks of separable neural-network states (SNQS). Each SNQS architecture implements a specific partition $\mathcal{P} = \{C_1, \dots, C_B\}$ by constraining the visible-hidden connectivity of the restricted Boltzmann machine so that hidden units couple only to qubits within the same cluster. The maximal cluster size $|C_b|$ in a given architecture therefore bounds the entanglement depth of all states representable by that ansatz. By training both an unconstrained NQS and families of SNQS with different clusterings on the same measurement data and comparing their optimal negative log-likelihoods, we obtain an operational tool to falsify l -producibility hypotheses and to certify non- k -separability and entanglement depth directly from finite-shot statistics.

B. Neural quantum states and separable architectures

Neural quantum states provide a variational parametrization of many-body wave functions in terms of artificial neural networks [39,41–44]. A standard architecture for this purpose is the restricted Boltzmann machine [39]. An RBM is a bipartite undirected graphical model with two layers of binary variables: a *visible* layer encoding the physical spin configuration $\mathbf{s} \in \{\pm 1\}^N$ and a *hidden* layer of auxiliary units $\mathbf{h} \in \{\pm 1\}^M$ that induce effective many-body correlations among the visible spins. Because connections exist exclusively *between* layers (via a weight matrix W_{ij}) with no intralayer couplings [Fig. 2(a)], the hidden variables can be traced out analytically. This yields a tractable closed-form expression for the unnormalized amplitude [Eq. (7)], meaning that the Born-rule probabilities $p_{\theta}(\mathbf{s}) = |\Psi_{\theta}(\mathbf{s})|^2$ can be evaluated exactly for moderate system sizes.

We consider N qubits with computational basis $\{|\mathbf{s}\rangle\}_{\mathbf{s} \in \{\pm 1\}^N}$. An RBM consists of N visible spins $s_i \in \{\pm 1\}$ and M hidden spins $h_j \in \{\pm 1\}$ with energy

$$E_{\theta}(\mathbf{s}, \mathbf{h}) = -\sum_i a_i s_i - \sum_j b_j h_j - \sum_{ij} s_i W_{ij} h_j, \quad (6)$$




FIG. 2. Restricted Boltzmann machine connectivity for $N = 6$ visible units s_i (blue, bottom row) with biases a_i and $M = 12$ hidden units h_j (red, top row) with biases b_j , connected by the weight matrix W_{ij} (gray lines). Each panel corresponds to a different partition $\mathcal{P}$ that defines a separable neural quantum state architecture: (a) unconstrained RBM ($\mathcal{A}_{\text{full}}$, $d_{\text{max}} = 6$), where every visible unit connects to every hidden unit; (b) $\mathcal{A}_{5|1}$ ($d_{\text{max}} = 5$); (c) $\mathcal{A}_{4|2}$ ($d_{\text{max}} = 4$); (d) $\mathcal{A}_{3|3}$ ($d_{\text{max}} = 3$); (e) $\mathcal{A}_{2|2|2}$ ($d_{\text{max}} = 2$); (f) $\mathcal{A}_{1|1|1|1|1|1}$ ($d_{\text{max}} = 1$, fully separable product ansatz). The absence of interblock connections (visible as spatial gaps) enforces the block-product structure of Eq. (9), bounding the entanglement depth of any representable state by $d_{\text{max}}(\mathcal{P})$.

where $\theta = \{a_i, b_j, W_{ij}\}$ are (complex) parameters. Tracing out the hidden layer yields an unnormalized amplitude

$$\begin{aligned} \tilde{\Psi}_{\theta}(\mathbf{s}) &= \sum_{\mathbf{h}} \exp[-E_{\theta}(\mathbf{s}, \mathbf{h})] \\ &= \exp\left(\sum_i a_i s_i\right) \prod_j 2 \cosh\left(b_j + \sum_i W_{ij} s_i\right), \end{aligned} \quad (7)$$

and the normalized neural quantum state is

$$|\Psi_{\theta}\rangle = \frac{1}{\sqrt{\mathcal{N}_{\theta}}} \sum_{\mathbf{s}} \tilde{\Psi}_{\theta}(\mathbf{s}) |\mathbf{s}\rangle, \quad \mathcal{N}_{\theta} = \sum_{\mathbf{s}} |\tilde{\Psi}_{\theta}(\mathbf{s})|^2. \quad (8)$$

In practice, amplitude and phase are modeled via two real RBMs [39,44], $\Psi_{\theta}(\mathbf{s}) = \exp[A_{\theta}(\mathbf{s}) + i\Phi_{\theta}(\mathbf{s})]$. The Born probabilities $p_{\theta}(\mathbf{s}) = |\Psi_{\theta}(\mathbf{s})|^2$ define a generative model for measurement outcomes in the computational basis.

The expressive power of the RBM can be constrained by masking the visible-hidden connectivity according to a chosen partition of the qubits. Let $\mathcal{P} = \{C_1, \dots, C_B\}$ be a partition of $\{1, \dots, N\}$ into disjoint clusters. We divide the hidden units into disjoint groups $H_1, \dots, H_B$ and require that units in H_b may only connect to visible qubits in C_b [50,51]. All weights W_{ij} violating this rule are set to zero [Figs. 2(b)–2(f)].

With this mask, the energy decomposes into a sum of cluster energies, $E_{\theta}(\mathbf{s}, \mathbf{h}) = \sum_{b=1}^B E^{(b)}(\mathbf{s}_{C_b}, \mathbf{h}_{H_b})$, and the

unnormalized amplitude factorizes,

$$\tilde{\Psi}_\theta(\mathbf{s}) = \prod_{b=1}^B \tilde{\psi}_b(\mathbf{s}_{C_b}), \quad (9)$$

so that the normalized neural quantum state has the form

$$|\Psi_\theta\rangle = \bigotimes_{b=1}^B |\phi_b\rangle, \quad (10)$$

with each factor $|\phi_b\rangle$ supported on the cluster C_b . Thus every SNQS implementing the partition $\mathcal{P}$ represents an B -separable state, with maximal cluster size $d_{\max}(\mathcal{P}) = \max_b |C_b|$. The entanglement depth of any state realizable by this architecture is therefore bounded by $d_{\max}(\mathcal{P})$.

Decreasing $d_{\max}(\mathcal{P})$ yields a hierarchy of increasingly constrained SNQS families: from fully separable product states ($|C_b| = 1$ for all b), through bipartite and tripartite clusterings, up to the unconstrained RBM where $\mathcal{P} = \{\{1, \dots, N\}\}$ and $d_{\max}(\mathcal{P}) = N$. This architectural hierarchy is the backbone of our non- k -separability certification protocol.

In the following, we denote by $\mathcal{A}_{\mathcal{P}}$ the SNQS family implementing partition $\mathcal{P}$. For an N -qubit system, we often use an explicit label $\mathcal{P} = n_1|n_2| \dots |n_B$ describing the block sizes in some fixed ordering of qubits. For instance, for $N = 6$, we write

$$\mathcal{A}_{2|2|2} \quad (11)$$

for the three-block SNQS where qubits are grouped as $\{0, 1\}|\{2, 3\}|\{4, 5\}$, and

$$\mathcal{A}_{1|1|1|1|1|1} \quad (12)$$

for the fully separable product ansatz. The corresponding optimized losses are written as $\mathcal{L}_{\mathcal{A}_{2|2|2}}^*$ and $\mathcal{L}_{\mathcal{A}_{1|1|1|1|1|1}}^*$, respectively. The fully connected reference NQS is denoted by $\mathcal{A}_{\text{full}}$ with optimal loss $\mathcal{L}_{\mathcal{A}_{\text{full}}}^*$.

The SNQS hierarchy relies on selecting a partition $\mathcal{P} = \{C_1, \dots, C_B\}$ of the qubit index set $\{1, \dots, N\}$ into nonempty, disjoint subsets (blocks). From the viewpoint of combinatorics, this is precisely the problem of counting set partitions, a classical topic connected to the Stirling numbers of the second kind and the Bell numbers [52–55].

The number of ways to partition an N -element set into *exactly* k nonempty, unlabeled subsets is given by a Stirling number of the second kind,

$$S(N, k) = \frac{1}{k!} \sum_{i=0}^k (-1)^{k-i} \binom{k}{i} i^N, \quad (13)$$

which counts the number of distinct equivalence relations on $\{1, \dots, N\}$ with exactly k equivalence classes. These numbers satisfy the recurrence

$$S(N+1, k) = k S(N, k) + S(N, k-1), \quad S(0, 0) = 1, \quad (14)$$

with $S(N, 0) = 0$ for $N > 0$, and have the exponential generating function

$$\sum_{N=0}^{\infty} S(N, k) \frac{z^N}{N!} = \frac{(e^z - 1)^k}{k!}. \quad (15)$$

Summing over all k yields the total number of multipartitions of an N -element set,

$$B_N = \sum_{k=0}^N S(N, k), \quad (16)$$

known as the N th *Bell number*. The Bell numbers obey

$$B_{N+1} = \sum_{k=0}^N \binom{N}{k} B_k, \quad (17)$$

and admit the exponential generating function

$$\sum_{N=0}^{\infty} B_N \frac{z^N}{N!} = e^{e^z - 1}. \quad (18)$$

Already for $N = 6$, the Stirling numbers take the values $S(6, 1) = 1$, $S(6, 2) = 31$, $S(6, 3) = 90$, $S(6, 4) = 65$, $S(6, 5) = 15$, and $S(6, 6) = 1$, so that $B_6 = 203$ distinct multipartitions exist. These include, for example, the fully separable partition $1|1|1|1|1|1$, the “dimerized” structure $2|2|2$, and the bipartition $3|3$. The combinatorial explosion of B_N with N ($B_8 = 4140$, $B_{10} = 115\,975$, etc.) shows that for even moderate system sizes an *exhaustive* scan over all partitions is neither necessary nor practical. In the SNQS setting, one therefore designs a *hierarchy* of architectures based on a carefully chosen subset of multipartitions (e.g., by increasing maximal block size or by exploiting physical symmetries), which is sufficient to probe the non- k -separability structure relevant for entanglement-depth certification.

III. MEASUREMENT PROTOCOL AND LIKELIHOOD-BASED CERTIFICATION

On a real device, the quantum state is not directly accessible; only measurement statistics in local bases are available. We therefore model the unknown target as a density operator ρ_{tgt} that can be probed only through a finite collection of measurement outcomes, summarized by empirical frequencies $f_b(\mathbf{s})$ in local measurement bases. For a given hierarchical family of ansatz states, we quantify how well each level of the hierarchy can reproduce these observed frequencies by minimizing the negative log-likelihood. The resulting likelihood values, and in particular the gaps between successive levels in the hierarchy, indicate how much expressive power is required to explain the data. Our objective is to use these likelihood gaps to certify non- k -separability and entanglement depth directly from the measurement statistics, without reconstructing ρ_{tgt} tomographically.

We consider projective measurements in local single-qubit Pauli bases. A measurement setting is specified by a basis label $b = (b_1, \dots, b_N)$, with $b_i \in \{X, Y, Z\}$ indicating whether qubit i is measured in the X , Y , or Z eigenbasis. In a given shot, we choose a basis pattern b , prepare the state ρ_{tgt} , measure all qubits in that basis, and record the classical bitstring $\mathbf{s} = (s_1, \dots, s_N) \in \{0, 1\}^N$.

Collecting N_{shots} repetitions over possibly many different bases yields a dataset

$$\mathcal{D} = \{(b_\ell, \mathbf{s}_\ell)\}_{\ell=1}^{N_{\text{shots}}}. \quad (19)$$

For each basis b and outcome $\mathbf{s}$, we define empirical frequencies

$$f_b(\mathbf{s}) = \frac{1}{N_b} \sum_{\ell: b_\ell=b} \delta_{\mathbf{s}, \mathbf{s}_\ell}, \quad (20)$$

where N_b is the total number of shots taken in basis b . In the limit $N_b \rightarrow \infty$, these converge to the Born probabilities

$$p_{\text{tgt}}(\mathbf{s} | b) = \text{Tr}[\rho_{\text{tgt}} \Pi_{\mathbf{s}}^{(b)}], \quad (21)$$

with $\Pi_{\mathbf{s}}^{(b)}$ being the projector onto the outcome $\mathbf{s}$ in basis b .

In the numerical experiments, we mimic this protocol by generating bitstrings from ideal target states, in a set of local Pauli bases (X, Y, Z) with randomly chosen patterns b .

Given a neural ansatz architecture $\mathcal{A}$ with parameters θ , the corresponding model state defines a density operator ρ_θ . For each basis b and outcome $\mathbf{s}$, the model predicts Born probabilities

$$p_\theta(\mathbf{s} | b) = \text{Tr}[\rho_\theta \Pi_{\mathbf{s}}^{(b)}]. \quad (22)$$

The natural statistical objective is the (average) negative log-likelihood of the observed dataset $\mathcal{D}$:

$$\mathcal{L}(\mathcal{A}; \theta) = -\frac{1}{N_{\text{shots}}} \sum_{\ell=1}^{N_{\text{shots}}} \ln p_\theta(\mathbf{s}_\ell | b_\ell). \quad (23)$$

Grouping identical outcomes into frequencies yields the equivalent form

$$\mathcal{L}(\mathcal{A}; \theta) = -\sum_b \sum_{\mathbf{s}} f_b(\mathbf{s}) \ln p_\theta(\mathbf{s} | b). \quad (24)$$

Minimizing (24) drives the NQS to reproduce the Born probabilities of the unknown target state as closely as the model family allows.

A. Likelihood gaps and non- k -separability

We index each SNQS architecture by the partition $\mathcal{P} = \{C_1, \dots, C_M\}$ of the qubits that it implements. We write $\mathcal{A}_{\mathcal{P}}$ for the corresponding ansatz family and $\mathcal{A}_{\text{full}}$ for the unconstrained (fully connected) NQS. Given a fixed measurement dataset and an architecture $\mathcal{A}$ with parameters θ , the NLL is $\mathcal{L}(\mathcal{A}; \theta)$ as in Eq. (24). Let $\mathcal{L}_{\mathcal{A}}^*$ denote the minimal NLL achieved by architecture $\mathcal{A}$ after training,

$$\mathcal{L}_{\mathcal{A}}^* = \min_{\theta} \mathcal{L}(\mathcal{A}; \theta). \quad (25)$$

For a partition $\mathcal{P}$, we then define the *loss gap*

$$\Delta_{\mathcal{P}} = \mathcal{L}_{\mathcal{A}_{\mathcal{P}}}^* - \mathcal{L}_{\mathcal{A}_{\text{full}}}^*. \quad (26)$$

At the level of exact global optima, $\Delta_{\mathcal{P}} \geq 0$ by model-class inclusion: the SNQS class is a strict subset of the unconstrained NQS class (obtained by zeroing specific weights), so the globally optimal NLL of the unconstrained model is guaranteed to be no worse. $\Delta_{\text{full}} = 0$, where “full” denotes the partition with a single block containing all qubits. In finite optimization runs, small violations of this bound can occur; we address this through the heuristic threshold $\overline{\Delta}_{\text{th}}$ introduced below.

Intuitively, $\Delta_{\mathcal{P}}$ measures how much worse the best $\mathcal{P}$ -separable ansatz performs compared to the best unconstrained

ansatz when both are fit to the same experimental bitstring histograms. In the ideal limit of tomographically informational data, sufficiently expressive NQS ansätze, and perfect optimization, the quantities $\mathcal{L}_{\mathcal{A}}^*$ converge to the minimal cross-entropy between the data distribution and the model family $\mathcal{A}$. If the true state ρ_{tgt} were exactly representable within $\mathcal{A}_{\mathcal{P}}$, then one would obtain $\Delta_{\mathcal{P}} \approx 0$ up to finite-shot fluctuations.

For each partition $\mathcal{P}$, the associated constrained NQS ansatz $\mathcal{A}_{\mathcal{P}}$ only represents states whose entanglement depth is bounded by

$$d_{\text{max}}(\mathcal{P}) = \max_{C \in \mathcal{P}} |C|. \quad (27)$$

If, for a fixed integer l , we find that

$$\Delta_{\mathcal{P}} > \overline{\Delta}_{\text{th}} \quad \forall \mathcal{P} \text{ such that } d_{\text{max}}(\mathcal{P}) \leq l, \quad (28)$$

then the experimental data are incompatible with all l -producible states represented in our constrained NQS hierarchy. In this proof-of-principle study, we determine $\overline{\Delta}_{\text{th}}$ heuristically based on the typical stability of the loss function after convergence. We treat likelihood differences comparable to the residual optimization fluctuations (typically $\lesssim 10^{-3}$ in our numerical experiments) as statistically insignificant, whereas structurally excluded partitions exhibit gaps orders of magnitude larger.

In this situation, we certify that ρ_{tgt} is not l producible and obtain a lower bound on the entanglement depth,

$$d_e(\rho_{\text{tgt}}) > l. \quad (29)$$

In the numerical examples below, we explicitly construct such hierarchies of SNQS architectures, compute the optimized losses $\mathcal{L}_{\mathcal{A}_{\mathcal{P}}}^*$, and report the loss gaps $\Delta_{\mathcal{P}}$. For conciseness, we summarize the resulting likelihood orderings across partitions using the compact notation

$$\mathcal{L}_{\mathcal{A}_{3|3}}^* \approx \mathcal{L}_{\mathcal{A}_{\text{full}}}^*, \quad \mathcal{L}_{\mathcal{A}_{2|2|2}}^* \gg \mathcal{L}_{\mathcal{A}_{\text{full}}}^*, \quad (30)$$

where the illustrative 3|3 and 2|2|2 partitions refer to a six-qubit system.

B. Scope and limitations

The total number of set partitions of N qubits (the Bell number B_N) grows superexponentially. In practice, however, an exhaustive scan over all architectures is neither required nor intended. When characterizing near-term devices, the hardware layout and the executed circuit dictate the relevant cuts: the user typically tests whether entanglement survives noise across specific chip boundaries. For such targeted tests, the computational cost is independent of B_N . Each training run scales as $O(E \cdot 2^N \cdot N \cdot H)$, where E is the number of epochs, N the number of qubits, and H the hidden-layer size; for the present proof-of-principle study ($N \leq 10$), this is tractable. For larger systems, sampling-based NLL estimators—as routinely used in variational Monte Carlo [39]—can replace exact enumeration, and more expressive autoregressive neural quantum states [56] can serve as drop-in replacements for RBMs, with partition constraints enforced by masking the autoregressive conditionals.

Our experiments use 200 randomly chosen local Pauli basis strings, each sampled with 2000 shots. The total budget

of 4×10^5 outcomes is comparable to that of randomized measurement protocols like classical shadows. Because the NLL objective (24) is a per-shot average, the leading statistical fluctuations decrease as $1/\sqrt{N_{\text{shots}}}$, though the prefactor may depend on system size, the gap magnitude, and the observable structure.

A large likelihood gap $\Delta_{\mathcal{P}} \gg 0$ should be interpreted as a one-sided, variational exclusion of the corresponding constrained hypothesis class: it provides evidence that the measured data are incompatible with that l -producible model family. In contrast, a small gap is inconclusive, as it may result either from true compatibility or from limited expressivity/imperfect optimization of the unconstrained ansatz. In practice, this interpretation relies on sufficiently good convergence of both constrained and unconstrained trainings; poor optimization of the constrained model can artificially enlarge the gap. For this reason, throughout the paper we treat the certification as variational with respect to the chosen model class and optimization procedure, and we use the Hilbert-Schmidt (HS) fidelity diagnostic introduced below as a sanity check on training quality. Additionally, because the unconstrained model has strictly more parameters, small likelihood improvements may partly reflect capacity to fit finite-shot noise rather than genuine entanglement; for larger systems, complexity-penalized model selection or held-out validation can sharpen this distinction. Finally, while we demonstrate the hierarchy using RBMs, the framework accommodates any generative model family with partition-constrained variants; using more expressive architectures tightens the certification lower bounds.

To confirm that the NLL minimum corresponds to an accurate generative model, we employ the following diagnostic. Whenever the exact target state ρ_{tgt} is available, as in the synthetic benchmarks of Sec. IV, we additionally compare it with the reconstructed state $\rho_{\text{est}}^{(\mathcal{A}_{\mathcal{P}})}$ obtained from a given ansatz $\mathcal{A}_{\mathcal{P}}$ using the HS overlap

$$F_{\text{HS}}(\rho_{\text{tgt}}, \rho_{\text{est}}^{(\mathcal{A}_{\mathcal{P}})}) = \frac{\text{Tr}[\rho_{\text{tgt}} \rho_{\text{est}}^{(\mathcal{A}_{\mathcal{P}})}]}{\sqrt{\text{Tr}[\rho_{\text{tgt}}^2] \text{Tr}[\rho_{\text{est}}^{(\mathcal{A}_{\mathcal{P}})^2]}} \quad (31)$$

and the associated HS distance

$$D_{\text{HS}}(\rho_{\text{tgt}}, \rho_{\text{est}}^{(\mathcal{A}_{\mathcal{P}})}) = \|\rho_{\text{tgt}} - \rho_{\text{est}}^{(\mathcal{A}_{\mathcal{P}})}\|_2 = \sqrt{\text{Tr}[(\rho_{\text{tgt}} - \rho_{\text{est}}^{(\mathcal{A}_{\mathcal{P}})})^2]}. \quad (32)$$

Both quantities are analytic and easy to evaluate numerically [57,58]. Importantly, F_{HS} and D_{HS} are used here only as *diagnostics* in numerical experiments; the actual non- k -separability and entanglement-depth certification on experimental data relies solely on the likelihood gaps $\Delta_{\mathcal{P}}$ obtained from finite-shot statistics, without any full tomographic reconstruction of ρ_{tgt} .

IV. NUMERICAL RESULTS

In this section, we illustrate the likelihood-gap certification strategy on three synthetic pure N -qubit targets of increasing structural complexity: a six-qubit GHZ state, a product of three Bell pairs, and a ten-qubit product of heterogeneous entangled clusters. In all cases, we generate finite-shot data

from the target state by measuring in a pool of 200 random local Pauli bases (with 2000 shots per basis), and we train both the unrestricted NQS ansatz $\mathcal{A}_{\text{full}}$ and a family of SNQS ansätze $\mathcal{A}_{\mathcal{P}}$ corresponding to different multipartitions $\mathcal{P}$ of the qubits.

For each architecture, we monitor the negative log-likelihood during training and record the minimal value $\mathcal{L}_{\mathcal{A}_{\mathcal{P}}}^*$ attained along the optimization trajectory. The likelihood gap $\Delta_{\mathcal{P}} = \mathcal{L}_{\mathcal{A}_{\mathcal{P}}}^* - \mathcal{L}_{\mathcal{A}_{\text{full}}}^*$, defined in Eq. (26), is then used to certify non- k -separability and entanglement depth according to the criterion (28). Whenever the exact target state is known (as here), we additionally evaluate Hilbert-Schmidt fidelity F_{HS} and the associated HS distance D_{HS} as diagnostics; however, the certification itself is phrased purely in terms of likelihood gaps and does not rely on access to ρ_{tgt} . For numerical details, see Appendix A.

Figure 3 shows the training dynamics of the NLL for a representative subset of partitions in each example, while the optimized losses and fidelities are summarized in the tables.

A. Pure states

1. Six-qubit GHZ

Our first benchmark is the pure six-qubit GHZ state in the computational basis,

$$|\psi_{\text{A}}\rangle \equiv |\psi_{\text{GHZ}}^{\text{ZZZZZZ}}\rangle = \frac{1}{\sqrt{2}}(|0_z\rangle^{\otimes 6} + |1_z\rangle^{\otimes 6}), \quad (33)$$

which is genuinely multipartite with entanglement depth $d_e(|\psi_{\text{A}}\rangle) = 6$. We train the global (unrestricted) ansatz $\mathcal{A}_{\text{full}}$ together with SNQS ansätze associated with the partitions

$$\mathcal{P} \in \{3|3, 2|4, 1|5, 2|2|2, 1|1|1|1|1|1\},$$

whose maximal block sizes are $d_{\text{max}}(\mathcal{P}) \in \{3, 4, 5, 2, 1\}$.

The global NQS converges to the lowest NLL, which we take as the reference $\mathcal{L}_{\mathcal{A}_{\text{full}}}^*$. All SNQS models exhibit sizable positive loss gaps,

$$\Delta_{1|5} \approx 0.46, \quad \Delta_{2|4} \approx 0.57, \quad \Delta_{3|3} \approx 0.60,$$

while more strongly constrained partitions lead to even larger gaps,

$$\Delta_{2|2|2} \approx 1.04, \quad \Delta_{1|1|1|1|1|1} \approx 2.01.$$

Thus every partition with $d_{\text{max}}(\mathcal{P}) \leq 5$ incurs a clear, statistically significant likelihood penalty relative to the unrestricted model (Table I).

Within our SNQS hierarchy, the measurement data are therefore incompatible with all l -producible models with $l \leq 5$, and the likelihood-gap criterion [Eq. (29)] certifies $d_e(\rho_{\text{tgt}}) > 5$, in agreement with the known entanglement depth of the GHZ state. This illustrates a pronounced “entanglement barrier”: any attempt to describe $|\psi_{\text{GHZ}}^{\text{ZZZZZZ}}\rangle$ with strictly bounded entanglement depth leads to a robust increase in NLL. For the global ansatz, the reconstructed state achieves $F_{\text{HS}} \approx 1$ and very small HS distance, confirming that the minimum of the NLL landscape corresponds to an accurate model of the target and that the observed gaps are not optimization artifacts.




FIG. 3. Training dynamics of the NLL for neural quantum state models fitted to finite-shot random-Pauli data. In each panel, the blue curve shows the fully entangling (“unconstrained”) ansatz, while the other curves correspond to separable neural quantum states with different partitions $\mathcal{P}$ of the qubits (see legends), whose maximal block size bounds the entanglement depth. (a) Six-qubit GHZ target [Eq. (33)]. (b) Product of three Bell pairs on six qubits [Eq. (34)]. (c) Ten-qubit product of a three-qubit GHZ state, a three-qubit Dicke state with one excitation, and a four-qubit GHZ state in local bases [Eq. (35)]. Persistent offsets between the constrained and unconstrained NLL traces signal partitions that are incompatible with the target entanglement structure. The NLL gap for partition $1|1|1|1|1|1$, $\Delta_{1|1|1|1|1|1}$ is denoted as dashed-brown vertical line.

TABLE I. Training results for the six-qubit GHZ target [Eq. (33)]. For each ansatz, we report the maximal block size $d_{\max}$, the minimal NLL gap $\Delta_{\mathcal{P}}$ with respect to the global (unrestricted) model, the Hilbert-Schmidt fidelity F_{HS} , and the Hilbert-Schmidt distance D_{HS} between the reconstructed and exact states. Rows are ordered by increasing $\Delta_{\mathcal{P}}$.

Partition $\mathcal{P}$	$d_{\max}$	$\Delta_{\mathcal{P}}$	F_{HS}	D_{HS}
Unconstrained	6	0.000	1.000	0.003
1 5	5	0.464	0.245	1.229
2 4	4	0.569	0.000	1.414
3 3	3	0.598	0.500	1.000
2 2 2	2	1.037	0.038	1.387
1 1 1 1 1 1	1	2.006	0.118	1.328

TABLE II. Training results for the product of three Bell pairs [Eq. (34)]. Columns as in Table I. Rows are ordered by increasing $\Delta_{\mathcal{P}}$.

Partition $\mathcal{P}$	$d_{\max}$	$\Delta_{\mathcal{P}}$	F_{HS}	D_{HS}
Unconstrained	6	0.000	1.000	$<10^{-2}$
2 4	4	0.000	1.000	$<10^{-2}$
2 2 2	2	0.000	1.000	$<10^{-2}$
3 3	3	0.636	0.173	1.286
1 5	5	0.649	0.499	1.001
1 1 1 1 1 1	1	1.920	0.024	1.397

2. Three Bell pairs

The second example is a six-qubit state consisting of three independent Bell pairs,

$$|\psi_{\text{B}}\rangle \equiv |\Phi^+\rangle^{\otimes 3}, \quad |\Phi^+\rangle = \frac{1}{\sqrt{2}}(|0_z 0_z\rangle + |1_z 1_z\rangle). \quad (34)$$

This state has entanglement depth $d_e(|\psi_{\text{B}}\rangle) = 2$ and exhibits only short-range two-qubit correlations. We again train the global NQS together with the same family of SNQS partitions as in the GHZ benchmark.

The likelihood hierarchy now changes qualitatively. Taking the global ansatz as reference ($\Delta_{\text{global}} = 0$), we obtain

$$\Delta_{2|2|2} \approx 0, \quad \Delta_{2|4} \approx 0,$$

within numerical precision, while partitions that break the Bell-pair structure incur substantial positive gaps,

$$\Delta_{3|3} \approx 0.64, \quad \Delta_{1|5} \approx 0.65, \quad \Delta_{1|1|1|1|1|1} \approx 1.92.$$

The ansatz with partition 2|2|2 (three two-qubit blocks) matches the performance of the global model, as visible in Fig. 3(b), reflecting the fact that the target is exactly two-producible with respect to the natural Bell-pair partition. The partition 2|4 also achieves essentially zero gap: the four-qubit block is flexible enough to internally represent two Bell pairs, so the target again lies within the corresponding model family.

From the perspective of entanglement depth, these results show that all $l = 1$ -producible (fully separable) SNQS models are strongly disfavored by the data, whereas there exist $l = 2$ -producible architectures ($d_{\max} = 2$) whose NLL is indistinguishable from the global model. Consequently, the likelihood-gap criterion certifies

$$d_e(|\psi_{\text{B}}\rangle) > 1,$$

and the smallest l for which the data admit a compatible description is $l = 2$, consistent with the exact entanglement depth. As in the GHZ case, the architectures with vanishing NLL gap yield reconstructed density matrices with $F_{\text{HS}} \approx 1$, while models with large gaps exhibit visibly reduced HS fidelities and HS distances close to unity (Table II).

3. Ten-qubit product of entangled clusters

The final benchmark is a ten-qubit pure state composed of three entangled clusters,

$$|\psi_{\text{C}}\rangle = |\psi_{\text{GHZ}}^{\text{XYZ}}\rangle \otimes |\mathbf{D}_1^{(3)}\rangle \otimes |\psi_{\text{GHZ}}^{\text{XZY}}\rangle, \quad (35)$$

TABLE III. Training results for the ten-qubit state [Eq. (35)]. Columns as in Table I. Rows are ordered by increasing $\Delta_{\mathcal{P}}$.

Partition $\mathcal{P}$	$d_{\max}$	$\Delta_{\mathcal{P}}$	F_{HS}	D_{HS}
Unconstrained	10	0.000	1.000	2.9×10^{-2}
6 4	6	0.000	1.000	1.5×10^{-2}
3 7	7	0.000	1.000	1.6×10^{-2}
3 3 4	4	0.000	1.000	1.7×10^{-2}
4 6	6	0.443	0.518	0.981
4 2 4	4	0.444	0.601	0.893
5 1 4	5	0.479	0.575	0.921
3 4 3	4	0.530	0.280	1.200
7 3	7	0.531	0.043	1.383
2 4 4	4	0.559	0.264	1.213
1 5 4	5	0.592	0.246	1.228
5 5	5	0.647	0.008	1.409
4 1 5	5	0.654	0.010	1.407
4 3 3	4	0.904	0.114	1.331
4 5 1	5	0.963	0.045	1.382
1 4 5	5	1.006	0.239	1.234
4 4 2	4	1.021	0.252	1.224
1 8 1	8	1.074	0.140	1.312
5 4 1	5	1.140	0.009	1.408
2 2 2 2 2	2	1.458	0.019	1.401
1 1 1 1 1 1 1 1 1 1	1	21.736	0.031	1.000

where the superscripts indicate local single-qubit basis rotations and $|D_n^{(N)}\rangle$ denotes a Dicke state with n excitations among N qubits,

$$|D_n^{(N)}\rangle = \binom{N}{n}^{-1/2} \sum_{\pi} P_{\pi}(|0\rangle^{\otimes(N-n)} \otimes |1\rangle^{\otimes n}), \quad (36)$$

with the sum taken over all distinct permutations P_{π} . The target thus consists of a three-qubit GHZ state in (X, Y, Z) bases, a three-qubit Dicke state with one excitation, and a four-qubit GHZ state in (X, Z, X, Y) bases. Its entanglement depth is set by the largest entangled cluster,

$$d_e(|\psi_C\rangle) = 4, \quad (37)$$

while the global structure is a tensor product across the three clusters.

To probe how this heterogeneous structure is reflected in likelihood gaps, we train the global NQS together with a broad family of SNQS partitions on random-Pauli data for $N = 10$. The partitions include two-block and three-block architectures such as 4|6, 6|4, 7|3, 3|7, 5|5, 3|3|4, 4|3|3, 3|4|3, 4|2|4, 2|4|4, 4|4|2, 5|1|4, 4|1|5, 1|5|4, 1|8|1, as well as strongly factorized models 2|2|2|2|2 and the fully product ansatz 1|1|1|1|1|1|1|1|1|1. Representative training curves are shown in Fig. 3(c), and the optimized losses and HS diagnostics are listed in Table III.

The global, fully entangling NQS sets the reference value $\mathcal{L}_{\mathcal{A}_{\text{full}}}^*$ and achieves $F_{\text{HS}} \approx 1$ with a very small HS distance. Strikingly, several structured ansätze with maximal block size $d_{\max} \geq 4$ attain essentially indistinguishable NLL and $F_{\text{HS}} \approx 1$: the partitions 6|4, 3|7, and, most importantly, the physically natural 3|3|4 all have $\Delta_{\mathcal{P}} \simeq 0$. Although these architectures group the ten qubits in different ways, none of them enforce an entanglement depth smaller than 4, and all are expressive

enough to reproduce the data at the level of finite-shot statistics. From the point of view of our protocol, they are therefore all compatible with the observed measurement outcomes.

In contrast, partitions that fragment the physical clusters or impose an incompatible factorization structure exhibit sizable positive likelihood gaps. Examples include

$$\Delta_{4|6} \approx 0.44, \quad \Delta_{4|2|4} \approx 0.44, \quad \Delta_{5|1|4} \approx 0.48, \quad \Delta_{3|4|3} \approx 0.53,$$

all accompanied by a noticeable drop in F_{HS} and an increase in HS distance. Here, the SNQS partition forces the model either to cut through at least one of the entangled clusters or to entangle qubits that are unentangled in the true tensor-product structure; in both cases, the resulting model family cannot simultaneously accommodate all empirical correlations, and the NLL gap becomes a visible “entanglement penalty.”

This trend is amplified for strongly overfragmented ansätze. Architectures with small $d_{\max}$, such as 2|2|2|2|2 and the fully product partition 1|1|1|1|1|1|1|1|1|1, show gaps of order unity and well beyond (e.g., $\Delta_{2|2|2|2|2} \approx 1.46$ and $\Delta_{1|1|\dots|1} \approx 21.7$), together with very low Hilbert-Schmidt fidelities and HS distances close to their maximal values in our normalization. These models are manifestly unable to reproduce the measurement statistics of a state with genuine four-qubit entanglement.

Two robust conclusions follow from this benchmark. First, among all tested SNQS architectures, every partition with $d_{\max} \leq 3$ (namely 2|2|2|2|2 and 1|1|1|1|1|1|1|1|1|1) exhibits a substantial positive gap $\Delta_{\mathcal{P}}$, so the data are incompatible with any l -producible model with $l \leq 3$ within our ansatz hierarchy. Second, several architectures with $d_{\max} = 4$ or larger (in particular 3|3|4) reproduce the likelihood of the global model up to finite-shot fluctuations, showing that an entanglement depth of 4 is sufficient to describe the system. Combining these observations, the likelihood-gap criterion certifies

$$d_e(|\psi_C\rangle) > 3,$$

in agreement with the known value $d_e = 4$, and it does so using only NLL comparisons between neural ansätze trained on measurement data, without full tomographic reconstruction of the ten-qubit state.

Across all three pure-state benchmarks, a coherent picture emerges: whenever the SNQS partition $\mathcal{P}$ is too restrictive to capture the true entanglement structure [i.e., it enforces a maximal block size $d_{\max}(\mathcal{P})$ smaller than the actual entanglement depth], training produces a robust positive gap $\Delta_{\mathcal{P}} \gg 0$; conversely, as soon as the hierarchy includes an ansatz with $d_{\max}(\mathcal{P})$ at least as large as the true entanglement depth, the corresponding likelihood becomes indistinguishable from that of the global model. This is precisely the behavior required for likelihood-based entanglement-depth certification without full tomography.

B. Mixed states

We now move beyond ideal pure states and address the certification of entanglement depth in more realistic scenarios where decoherence maps the target $|\psi_{\text{tgt}}\rangle$ onto a mixed density matrix ρ_{tgt} . In this regime, a useful witness must distinguish genuine quantum entanglement from classical statistical

correlations generated by incoherent noise. We demonstrate that the likelihood—gap criterion continues to work in this setting by reusing the three target states from the previous section, now subject to a noise channel.

To mimic energy relaxation in, e.g., superconducting qubit platforms, we model the environment by independent amplitude-damping channels acting on each qubit. For a single qubit with density matrix ρ , this channel is defined by the Kraus operators

$$K_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p} \end{pmatrix}, \quad K_1 = \begin{pmatrix} 0 & \sqrt{p} \\ 0 & 0 \end{pmatrix}, \quad (38)$$

through the completely positive trace-preserving map

$$\mathcal{E}_{\text{AD}}(\rho) = K_0 \rho K_0^\dagger + K_1 \rho K_1^\dagger. \quad (39)$$

Physically, K_0 leaves $|0\rangle$ unchanged and damps the excited state $|1\rangle$ toward $|0\rangle$, while K_1 describes a stochastic decay $|1\rangle \rightarrow |0\rangle$ with probability p . For an N -qubit state, we apply $\mathcal{E}_{\text{AD}}$ to each qubit independently,

$$\rho_{\text{tgt}} = (\mathcal{E}_{\text{AD}}^{\otimes N})[|\psi_{\text{tgt}}\rangle\langle\psi_{\text{tgt}}|], \quad (40)$$

which drives the system toward the separable ground state $|0\rangle^{\otimes N}$ and reduces the purity $\text{Tr}(\rho_{\text{tgt}}^2)$.

To accommodate mixed targets, we generalize the variational ansatz from a single pure neural quantum state to a small *ensemble* of pure NQS components. For the global (unstructured) model, we write

$$\rho_{\text{glob}} = \sum_{r=1}^R w_r |\psi_r\rangle\langle\psi_r|, \quad w_r \geq 0, \quad \sum_r w_r = 1, \quad (41)$$

where each $|\psi_r\rangle$ is a pure NQS with its own parameters and the w_r are learnable mixing weights. In all experiments below, we use a small rank $R = 4$.

For a structured mixed ansatz compatible with a partition $\mathcal{P} = \{C_1, \dots, C_B\}$, we impose the SNQS product structure inside *each* component of the ensemble. The corresponding density matrix is

$$\rho_{\mathcal{P}} = \sum_{r=1}^R w_r \bigotimes_{C \in \mathcal{P}} |\psi_C^{(r)}\rangle\langle\psi_C^{(r)}|, \quad (42)$$

where $|\psi_C^{(r)}\rangle$ is a pure NQS defined only on block $C \subset \{1, \dots, N\}$. For each component r , the global state factorizes exactly across the blocks in $\mathcal{P}$, i.e.; there is *no* entanglement between different blocks in a single component. The ensemble weights w_r introduce a hidden classical randomness: before each (virtual) experimental shot, the model “chooses” one of the product patterns r and prepares that configuration.

For a given partition $\mathcal{P}$, the ensemble thus spans a rich family of mixed states that are separable with respect to $\mathcal{P}$, in the sense that they can be written as a classical mixture of block-product pure states. Such states can exhibit complicated classical correlations between blocks, but they cannot reproduce quantum correlations across the cut. Within this framework, the likelihood gap $\Delta_{\mathcal{P}}$ retains a simple operational meaning: (1) if $\Delta_{\mathcal{P}} \approx 0$, then the measurement data are compatible with such a classical lottery over product configurations; no entanglement across $\mathcal{P}$ is needed to fit the data within our ansatz family; and (2) if $\Delta_{\mathcal{P}} \gg 0$, then

TABLE IV. Training results for the mixed six-qubit GHZ state [Eq. (33)] under amplitude damping channel ($p = 0.05$). Despite the loss of purity, all structured ensembles with $d_{\text{max}} < 6$ exhibit systematic positive likelihood gaps, showing that the data cannot be fully explained by mixtures of product-structured components with bounded entanglement depth.

Partition $\mathcal{P}$	d_{max}	$\Delta_{\mathcal{P}}$	F_{HS}	D_{HS}
Unconstrained	6	0.000	0.997	0.067
1 5	5	0.026	0.723	0.600
2 4	4	0.029	0.711	0.611
3 3	3	0.029	0.709	0.613
2 2 2	2	0.030	0.704	0.618
1 1 1 1 1 1	1	0.030	0.704	0.619

even allowing for classical mixtures of many block-product states, the model cannot reproduce the observed statistics. The remaining correlations across the cut must therefore be genuinely quantum in origin.

We start with six-qubit GHZ state of Eq. (33) under local amplitude damping. The corresponding results are summarized in Table IV. Compared to the pure-state case, the absolute sizes of the gaps are reduced by roughly an order of magnitude: all structured ensemble ansätze lie within $\Delta_{\mathcal{P}} \sim 3 \times 10^{-2}$ of the global model. This is expected, as amplitude damping drives the state partially toward the separable ground state and hence weakens the multiqubit coherences. However, a clear pattern remains: the global ensemble achieves the lowest NLL and high Hilbert-Schmidt fidelity ($F_{\text{HS}} \approx 0.997$), while *every* partition with $d_{\text{max}} < 6$ incurs a consistent positive penalty.

Within the ensemble SNQS picture, this means that, even after allowing for classical mixtures of block-product NQS, no model with maximal block size $d_{\text{max}} \leq 5$ can fully account for the noisy GHZ data. The mixed target is therefore incompatible with any l -producible ensemble SNQS with $l \leq 5$ in our hierarchy, and we again certify

$$d_e(\rho_{\text{GHZ}_6}^{(\text{mixed})}) > 5,$$

i.e., an entanglement depth of at least six qubits survives on top of the amplitude-damping noise.

The second mixed benchmark starts from three independent Bell pairs [Eq. (34)], and subjects them to the same local amplitude-damping channel. This example is particularly informative because it contains both *short-range* entanglement within each pair and potential *classical* correlations between pairs induced by noise. The results are listed in Table V.

The likelihood hierarchy follows the physical structure of the state. The partitions 2|2|2 and 2|4, which respect or coarse grain the Bell-pair grouping, achieve NLL values indistinguishable from the global ensemble and HS fidelities $F_{\text{HS}} \approx 0.998$. This shows that any correlations *between* Bell pairs created by the noise can be modeled as classical correlations between different product patterns of entangled two-qubit blocks.

In stark contrast, partitions that cut through at least one Bell pair (3|3 and 1|5) incur clear positive gaps ($\Delta_{\mathcal{P}} \approx 0.1$) and substantially reduced Hilbert-Schmidt fidelities. The fully

TABLE V. Training results for mixed state of three Bell pairs [Eq. (34)] under amplitude damping channel ($p = 0.05$).

Partition $\mathcal{P}$	$d_{\max}$	$\Delta_{\mathcal{P}}$	F_{HS}	D_{HS}
Unconstrained	6	0.000	0.998	0.061
2 4	4	0.000	0.998	0.060
2 2 2	2	0.001	0.998	0.059
3 3	3	0.099	0.806	0.523
1 5	5	0.106	0.812	0.518
1 1 1 1 1	1	0.466	0.265	0.876

product ansatz $1|1|\dots|1$ performs worst, as expected, with $\Delta \approx 0.47$ and low fidelity. Since the 3|3 ansatz already has access to a rank-4 ensemble of product states across the cut, its failure indicates that the correlations within each Bell pair cannot be decomposed into a mixture of product states, and must therefore be genuinely quantum.

Finally, we consider ten-qubit product state composed of GHZ and Dicke clusters discussed in Eq. (35), subjected to the same noise channel. Table VI reports the likelihood gaps and Hilbert-Schmidt diagnostics for the same family of partitions as in the pure-state case.

Note that the partition 7|3 shows a gap ($\Delta_{7|3} \approx 0.066$) comparable to 3|3|4, despite cutting through the four-qubit GHZ cluster. This coincidence is a noise effect: under amplitude damping ($p = 0.05$), the four-qubit GHZ coherences are partially suppressed, so the NLL penalty from fragmenting this block is comparable to the penalty incurred by the ensemble's limited rank ($R = 4$) in the 3|3|4 fit. In contrast, in the pure-state case (Table III), 7|3 exhibits a substantially larger gap ($\Delta_{7|3} \approx 0.53$) while 3|3|4 achieves $\Delta \approx 0$, clearly

TABLE VI. Training results for the ten-qubit state [Eq. (35)] under amplitude damping channel ($p = 0.05$). Columns as in Table I.

Partition $\mathcal{P}$	$d_{\max}$	$\Delta_{\mathcal{P}}$	F_{HS}	D_{HS}
Unconstrained	10	0.000	0.987	0.137
6 4	6	0.001	0.989	0.125
3 7	7	0.044	0.836	0.440
3 3 4	4	0.066	0.765	0.516
7 3	7	0.066	0.663	0.599
2 4 4	4	0.075	0.682	0.585
1 5 4	5	0.082	0.647	0.612
3 4 3	4	0.105	0.581	0.652
1 8 1	8	0.117	0.509	0.695
4 6	6	0.189	0.572	0.660
5 5	5	0.209	0.531	0.687
4 4 2	4	0.242	0.452	0.729
5 1 4	5	0.247	0.469	0.720
1 4 5	5	0.317	0.383	0.768
4 2 4	4	0.335	0.384	0.764
4 5 1	5	0.343	0.352	0.781
5 4 1	5	0.345	0.358	0.778
4 1 5	5	0.355	0.343	0.789
4 3 3	4	0.366	0.320	0.797
2 2 2 2 2	2	0.427	0.281	0.816
1 1 ... 1	1	21.55	0.039	0.799

distinguishing the two. The protocol's certification of $d_e > 3$ remains correct in both cases.

The unconstrained NQS and the 6|4 partition achieve nearly identical likelihoods and high Hilbert-Schmidt fidelities, indicating that grouping the first two clusters into a single effective block is compatible with the data. The natural tripartition 3|3|4 ($\Delta_{3|3|4} \approx 0.064$) and bipartition 3|7 ($\Delta_{3|7} \approx 0.044$), which follow the true cluster boundaries, show small gaps, attributable to the limited ensemble rank when modeling three independent mixed clusters. Scaling the ensemble rank R is necessary for highly mixed states and represents a hyperparameter trade-off in the current implementation.

More importantly, partitions that attempt to fragment the four-qubit GHZ cluster, such as 4|3|3 and 5|5, incur substantially larger gaps ($\Delta \approx 0.35$ – 0.4) and markedly reduced HS fidelities. Strongly overfragmented partitions like 2|2|2|2|2 and the fully product $1|1|\dots|1$ perform worst, as expected.

From the point of view of entanglement depth, the data are incompatible with any ensemble SNQS whose maximal block size satisfies $d_{\max} \leq 3$, while several partitions with $d_{\max} = 4$ or larger match the performance of the global ensemble within statistical uncertainty. Our likelihood-gap criterion therefore certifies $d_e(\rho_C^{(\text{mixed})}) > 3$, consistent with the four-qubit GHZ cluster present in the underlying pure state.

At the same time, the ten-qubit example illustrates an important feature of the method: when multiple partitions with $d_{\max} \geq 4$ achieve very similar likelihoods (e.g., 6|4, 3|7, 3|3|4), the NLL landscape does not single out a unique block structure. In this regime, our protocol correctly refrains from overinterpreting the data: it reliably detects that some entangled cluster of size at least four is required, but it does not artificially pin down a particular partition when the evidence is ambiguous.

For mixed states, we test separability with respect to fixed partitions $\mathcal{P}$ using low-rank mixtures of $\mathcal{P}$ -product components; allowing mixtures across different $\mathcal{P}$ would further enlarge the constrained hypothesis class and can be incorporated in future work. Our conclusions are variational with respect to the chosen SNQS model class and the tested set of partitions; large likelihood gaps rule out those hypotheses within this modeling family.

V. INTERPRETABILITY OF UNRESTRICTED NEURAL QUANTUM STATES

While the likelihood-gap protocol treats the unrestricted ansatz $\mathcal{A}_{\text{full}}$ primarily as a flexible reference for model comparison, it is natural to ask what physical information is encoded in its trained parameters. In physics applications, accurate likelihood fits are most useful when they can be accompanied by diagnostics that provide sanity checks on learned structure and yield compact, human-interpretable summaries that can guide subsequent modeling choices. A number of approaches along these lines have been developed, including tools for interpreting neural decision functions and assessing reliability [59,60], interpretable-by-construction architectures and order-parameter extractors that expose correlation structure [61–65], and sparsification or weight-space analyses that connect salient parameters and generalization behavior to underlying physical organization

[66–69]. In the present setting, we show that simple post-processing of the unrestricted NQS weights yields qualitative signatures of subsystem and entanglement structure directly at the level of learned parameters.

We begin with the characterization of pure target states with the help of two standard tools. First, we consider a two-body correlator

$$c_{ij}^{\alpha\beta} = \langle \sigma_i^\alpha \sigma_j^\beta \rangle - \langle \sigma_i^\alpha \rangle \langle \sigma_j^\beta \rangle, \quad i \neq j, \quad (43)$$

$\alpha, \beta \in \{X, Y, Z\}$, and define its basis aggregation

$$C_{ij} = \left(\sum_{\alpha, \beta \in \{X, Y, Z\}} (c_{ij}^{\alpha\beta})^2 \right)^{1/2}. \quad (44)$$

The C_{ij} correlator can be extracted either from the quantum state, directly from raw measurement bitstrings, or from trained NQS. In randomized-basis datasets, a direct estimator of each component $c_{ij}^{\alpha\beta}$ uses only those experimental shots in which qubits i and j were measured in the corresponding local bases (α, β) . For finite measurement budgets, this results in an estimator whose variance is both basis dependent and uneven across pairs (i, j) and settings (α, β) , because different basis patterns occur with different frequencies in the dataset. A trained generative model provides an alternative, model-based estimate of the same correlators: once the NQS is fitted by maximum likelihood to the full dataset, expectation values can be evaluated under the learned state using all measurement settings simultaneously, which can reduce estimator variance at the cost of model bias, analogous to standard regularized or model-based estimators. We additionally quantify two-qubit entanglement via the negativity. For each qubit pair (i, j) , we consider the reduced density matrix $\rho_{ij} = \text{Tr}_{\bar{i}\bar{j}}(\rho)$ and take its partial transpose with respect to one subsystem (here j), denoted $\rho_{ij}^{T_j}$. The negativity is then defined as

$$\mathcal{N}_{ij} = \sum_{\lambda_k < 0} |\lambda_k|, \quad (45)$$

where $\{\lambda_k\}$ are the eigenvalues of $\rho_{ij}^{T_j}$. We consider negativity as a reference entanglement diagnostic. In the synthetic benchmarks here, it is computed from the exact target state; in an experimental setting, it could only be accessed through a reconstructed model state via the trained NQS.

Figure 4 shows C_{ij} (top row) and $\mathcal{N}_{ij}$ (bottom row) for considered pure states $|\psi_{A,B,C}\rangle$ (columns from left to right, respectively). For six-qubit GHZ state, $|\psi_A\rangle$ [panel (a)], C_{ij} displays broadly distributed two-body correlations, reflecting its global symmetry, while negativity [panel (d)] vanishes, as the GHZ state is genuinely entangled state, and after tracing out any qubit, the reduced density matrix has no entanglement. For the product of three Bell pairs $|\psi_B\rangle$ [panel (b)], C_{ij} consists of three isolated strong links, sharply revealing the 2|2|2 cluster structure. The negativity [panel (e)] displays similar structure, showing strong entanglement between qubits in each of the Bell pairs. For the heterogeneous ten-qubit state, $|\psi_C\rangle$, C_{ij} reveals a block pattern consistent with the intended 3|3|4 factorization [panel (c)]. The negativity [panel (f)] shows strong entanglement within central three-qubit block of the $|D_1^{(3)}\rangle$ Dicke state and vanishes for blocks of GHZ states.




FIG. 4. Aggregated two-body correlator C_{ij} [Eq. (44) (top row)] and negativity [Eq. (45) (bottom row)] for the three benchmark states: six-qubit GHZ state $|\psi_A\rangle$ [Eq. (33)] product of three Bell pairs $|\psi_B\rangle$ [Eq. (34)] and ten-qubit product of three entangled clusters $|\psi_C\rangle$ [Eq. (35)].

A. Weight-space diagnostics of the unrestricted NQS

We now turn to structure encoded directly in the trained RBM weights. For a spin configuration $\mathbf{s} \in \{\pm 1\}^N$, we write the NQS wave function as

$$\Psi_\theta(\mathbf{s}) = \exp[A_\theta(\mathbf{s}) + i\Phi_\theta(\mathbf{s})], \quad (46)$$

where A_θ and Φ_θ are represented by RBMs of the form

$$X_\theta(\mathbf{s}) = \sum_i x_i s_i + \sum_h \ln \left[2 \cosh \left(y_h + \sum_i W_{ih}^{(X)} s_i \right) \right], \quad (47)$$

where $X \in \{A, \Phi\}$. In what follows, we focus on the visible-hidden weight matrices $W^{(A)}$ and $W^{(\Phi)}$, and construct from them two $N \times N$ summaries that can be compared to the correlator C_{ij} .

1. Effective qubits couplings

We first relate the RBM representation of Eq. (47) to an effective Ising-like model on the visible spins [70–72]. For definiteness, we consider the amplitude RBM and suppress the superscript (A) on the weights. Writing

$$A_\theta(\mathbf{s}) = \sum_i a_i s_i + \sum_h \ln \left[2 \cosh \left(b_h + \sum_i W_{ih} s_i \right) \right], \quad (48)$$

we specialize to the amplitude network, using the notation (a_i, b_h) for the visible and hidden biases in place of the generic labels (x_i, y_h) of Eq. (47), while $s_i \in \{\pm 1\}$ remains the spin configuration. We introduce $x_h(\mathbf{s}) = b_h + \sum_i W_{ih} s_i$. For moderate values of x_h , one may expand $\ln(2 \cosh x)$ around $x = 0$,

$$\ln(2 \cosh x) = \ln 2 + \frac{x^2}{2} - \frac{x^4}{12} + \mathcal{O}(x^6), \quad (49)$$

which inserted into Eq. (48) yields a quadratic approximation

$$A_\theta(\mathbf{s}) \approx A_0 + \sum_i h_i s_i + \frac{1}{2} \sum_{i,j} J_{ij} s_i s_j, \quad (50)$$




FIG. 5. Effective visible nodes (qubits) couplings extracted from the trained weight W of an unrestricted NQS. Each panel shows the absolute value of the normalized coupling matrix $|J_{ij}^{(X)}|$, with $X = A$ (amplitude, top row) or $X = \Phi$ (phase, bottom row), computed from the Gram matrices $J^{(X)} = W^{(X)}W^{(X)T}$ of the trained unconstrained NQS. Columns correspond to the three pure targets: (a), (d) six-qubit GHZ $|\psi_A\rangle$ [Eq. (33)] (b), (e) product of three Bell pairs $|\psi_B\rangle$ [Eq. (34)] and (c), (f) ten-qubit product state with GHZ and Dicke clusters $|\psi_C\rangle$ [Eq. (35)]. For $|\psi_A\rangle$, the couplings are comparatively homogeneous, reflecting the genuinely global character of the GHZ entanglement. For $|\psi_B\rangle$, the amplitude couplings display three localized 2×2 blocks, directly exposing the underlying $2|2|2$ structure. For $|\psi_C\rangle$, the dominant amplitude couplings concentrate on the three-qubit Dicke cluster, consistent with the fact that pairwise entanglement is present only in that block [Fig. 4(f)]. The rotated bases GHZ states are not simply visible, reflecting complicated structure in the computational basis, natural for $J_{ij}^{(X)}$ characteristic. The phase couplings $|J_{ij}^{(\Phi)}|$ (bottom row) show a more diffuse structure. For the GHZ state [panel (d)], the relative phases between computational basis states are simple (a single phase between $|0\rangle^{\otimes N}$ and $|1\rangle^{\otimes N}$), so the phase RBM learns a homogeneous pattern. For the Bell-pair state [panel (e)], the network shows weaker but visible two-qubit modularity. For the heterogeneous state [panel (f)], both $|J^{(A)}|$ and $|J^{(\Phi)}|$ show complex patterns reflecting the nontrivial local basis transformations. Overall, $|J_{ij}^{(A)}|$ captures the dominant outcome probabilities, whereas $|J_{ij}^{(\Phi)}|$ controls the relative phases.

where A_0 is a constant. Using $x_h^2(\mathbf{s}) = b_h^2 + 2b_h \sum_i W_{ih}s_i + \sum_{i,j} W_{ih}W_{jh}s_i s_j$, one finds

$$J_{ij} = \sum_h W_{ih}W_{jh} + \text{higher-order terms.} \quad (51)$$

An analogous construction applied to the phase RBM yields a matrix $J^{(\Phi)}$.

Figure 5 presents $|J^{(X)}|$ ($X \in \{A, \Phi\}$; amplitude: top row, phase: bottom row). The coupling element is nonvanishing when qubits i and j couple strongly to overlapping sets of hidden units, and small when their hidden-layer connectivity is largely disjoint. The visible structures in $|J^{(A)}|$ align with the basis-aggregated two-body correlations [Eq. (44)] and negativity, [Eq. (45)] (see Fig. 4). In particular, the Bell-pair state $|\psi_B\rangle$ produces three isolated 2×2 blocks in $|J^{(A)}|$, mirroring the three strong links in C_{ij} . The GHZ state $|\psi_A\rangle$ leads to a more homogeneous pattern, consistent with the absence of a small-cluster decomposition at the two-body level. For the heterogeneous state $|\psi_C\rangle$, written in locally rotated




FIG. 6. Row-based affinity matrices for the unrestricted NQS. Each panel shows the normalized affinity $A_{ij}^{(X)}$ defined in Eq. (54), with $X = A$ (amplitude, top row) or $X = \Phi$ (phase, bottom row), computed from the row vectors of the weight matrices $W^{(X)}$ of the trained unconstrained NQS. Columns correspond to the three pure target states $|\psi_A\rangle$, $|\psi_B\rangle$, and $|\psi_C\rangle$ as in Fig. 5. Large values of $A_{ij}^{(X)}$ indicate that qubits i and j couple to the hidden layer in a similar way. For the product of three Bell pairs $|\psi_B\rangle$, the amplitude affinity exhibits three nearly independent 2×2 blocks, while for the ten-qubit state $|\psi_C\rangle$ it reveals three clusters of sizes 3, 3, and 4, consistent with the intended tensor-product structure. The GHZ state $|\psi_A\rangle$ does not exhibit a pronounced block-diagonal pattern, consistent with the absence of a natural small-cluster decomposition.

bases, the weight-space structure of amplitude RBM, $|J_{ij}^{(A)}|$, reveals structure captured by negativity [see Fig. 4(f)], while the blocks of GHZ states are not visible in amplitude RBM, reflecting the increased complexity of its representation in the computational basis. The effectiveness of this diagnostic depends on how simply the target state is represented in the computational basis, since $A_\theta(\mathbf{s})$ and $\Phi_\theta(\mathbf{s})$ are defined in that basis.

2. Row-based affinity between visible units

Whereas $J^{(X)}$ emphasizes pairwise interaction strength mediated by shared hidden units, a complementary question is whether different qubits play similar roles within the RBM. This is captured by comparing the rows of $W^{(X)}$. Let

$$\mathbf{v}_i^{(X)} = (W_{i1}^{(X)}, \dots, W_{iM}^{(X)}) \in \mathbb{R}^M \quad (52)$$

denote the row vector associated with qubit i . We define squared row distances

$$d_{ij}^2 = \|\mathbf{v}_i^{(X)} - \mathbf{v}_j^{(X)}\|_2^2, \quad (53)$$

and rescale them into a normalized affinity

$$A_{ij}^{(X)} = 1 - \frac{d_{ij}^2}{\max_{k,\ell} d_{k\ell}^2}, \quad A_{ii}^{(X)} = 1, \quad (54)$$

with $X \in \{A, \Phi\}$. Large $A_{ij}^{(X)}$ indicates that qubits i and j couple to the hidden layer in a similar manner.

Figure 6 shows that $A^{(X)}$ provides a clear clustering signal that is consistent with the correlation C_{ij} and the coupling matrices $J^{(X)}$. For the Bell-pair state $|\psi_B\rangle$, the affinity reveals three nearly independent 2×2 blocks, while for the heterogeneous state $|\psi_C\rangle$ it groups qubits into three clusters of sizes

3, 3, and 4, recovering the intended tensor-product structure directly at the level of learned parameters. The $A^{(X)}$ serves as a data-driven proxy for subsystem structure and can be used to propose candidate partitions $\mathcal{P}$ for the SNQS hierarchy.

The metrics considered here, C_{ij} , $J^{(X)}$, and $A^{(X)}$, provide a compact tomography-free representation of pairwise correlations that can be extracted from raw measurement bitstrings and is faithfully reproduced by the trained unrestricted NQS. These provide a principled route from an unrestricted fit to candidate separability partitions for structured NQS models.

VI. CONCLUSIONS

We have recast the problem of certifying non- k -separability and entanglement depth as an experimentally driven, likelihood-based model-comparison task between neural quantum states. Instead of reconstructing a density matrix and subsequently applying entanglement criteria, we compare the optimized negative log-likelihoods of a fully entangling neural ansatz and a hierarchy of partition-constrained structured neural quantum states. Within this hierarchy, each multipartition $\mathcal{P}$ defines a l -producible model family whose maximal block size bounds the entanglement depth, and robust positive likelihood gaps provide an operational witness that the data are incompatible with that separability class.

A key ingredient is the use of structured neural quantum states for both pure and mixed targets. For pure states, separability is built into the architecture via masked connectivity that enforces a given product structure. For mixed states, we promote these structured states to low-rank ensembles so that classical correlations between blocks can be captured without generating entanglement across the cuts. In this way, the same likelihood-gap criterion applies uniformly in the presence or absence of noise and distinguishes classical noise-induced correlations from genuinely quantum ones.

Beyond certification, we considered an interpretability scheme for neural quantum states that extracts structural information about the underlying many-body state directly from the trained NQS parameters. Rather than treating the unrestricted ansatz as an uninterpretable one, we show that its learned weights encode coarse entanglement features such as natural qubit groupings and cluster structure. This provides qualitative insight into the organization of multipartite entanglement directly from measurement bitstrings, without explicit state reconstruction, and can be used to guide the design of structured ansatz families for further hypothesis testing.

Our numerical experiments indicate that useful bounds on non- k -separability, together with qualitative information on the entanglement structure, can be extracted from finite-shot local Pauli data and relatively modest neural architectures, without full tomography or an exhaustive scan over partitions. In unconstrained scenarios, the interpretability diagnostics of Sec. V (effective couplings $J_{ij}^{(X)}$ and affinities $A_{ij}^{(X)}$) can guide partition selection directly from a single unrestricted training run, reducing the search overhead. The dominant computational cost in the current implementation is the exact enumeration of 2^N configurations; for larger systems, sampling-based NLL estimators and more expressive autoregressive neural quantum states [56] can replace exact

enumeration while preserving the partition-constrained structure that underpins the certification hierarchy. This generative capability distinguishes our method from randomized measurement protocols [73], such as classical shadows [74], which have successfully addressed the scalability of estimating scalar observables and fidelities. While highly efficient for specific properties, shadow-based estimators do not yield a generative model of the system's statistics, nor do they directly expose the granular structure of multipartite correlations. In contrast, our results demonstrate that structured NQS provide a best-fit generative approximation of the state itself, allowing one to not only certify entanglement depth but also inspect the learned weights to recover physical insights—such as effective coupling graphs and qubit affinity—that remain accessible only through a model-based approach.

Several research directions remain open, including a systematic treatment of statistical uncertainty in likelihood gaps. Beyond the systematic treatment of statistical uncertainty, the interpretable weight diagnostics introduced here pave the way for automated model selection. Future work could employ Graph Neural Networks (GNNs) to process the effective interaction graphs derived from the unconstrained NQS. By mapping these topological features to qubit clusterings, GNNs could serve as a data-driven proposal mechanism, automatically identifying the most relevant candidate partitions to test within the structured ansatz hierarchy [75].

ACKNOWLEDGMENTS

We thank Grzegorz Rajchel-Mieldzioć, Arnau Riera, and Anna Dawid for useful comments. We acknowledge RES resources provided by Barcelona Supercomputing Center in Marenostrum 5 to NNO-2025-3-0004, and funds from MICIU/AEI/10.13039/501100011033/FEDER, UE.

DATA AVAILABILITY

The data that support the findings of this article are openly available [76].

APPENDIX: RBM PARAMETRIZATION OF THE NQS AND SNNS ANSÄTZE

This Appendix records the precise RBM architectures used in the numerical experiment [76]. We focus on the concrete parametrization implemented in the code (parameter tensors and their dimensions) and avoid rederiving expressions already introduced in the main text.

1. Visible variables and notation

The visible layer is an Ising encoding of computational-basis bitstrings, $s \in \{-1, +1\}^N$, obtained from $x \in \{0, 1\}^N$ by $s_i = 2x_i - 1$. All model probabilities required for training are evaluated by explicit enumeration over all 2^N configurations.

The pure-state NQS wave function is represented as a product of an amplitude RBM and an independent phase RBM,

$$\psi_\theta(s) = \exp[\mathcal{A}_\theta(s)] \exp[i 2\pi \Phi_\theta(s)]. \quad (\text{A1})$$

Both $\mathcal{A}_\theta$ and Φ_θ are computed from RBM-like hidden layers with the visible variables marginalized analytically (i.e., in “free-energy” form). The amplitude RBM uses H_{amp} hidden units with parameters

$$a \in \mathbb{R}^N, \quad b \in \mathbb{R}^{H_{\text{amp}}}, \quad W \in \mathbb{R}^{N \times H_{\text{amp}}}. \quad (\text{A2})$$

The phase RBM uses H_{phase} hidden units with parameters

$$c \in \mathbb{R}^N, \quad d \in \mathbb{R}^{H_{\text{phase}}}, \quad U \in \mathbb{R}^{N \times H_{\text{phase}}}. \quad (\text{A3})$$

In the implementation, the phase hidden contributions are passed through the bounded nonlinearity

$$g(z) = \frac{2}{\pi} \arctan[\tanh(z)] + \frac{1}{2}, \quad (\text{A4})$$

and the overall complex phase factor is $\exp(i2\pi \Phi_\theta)$. Throughout training and evaluation, the state vector ψ_θ is normalized explicitly, $\psi_\theta \rightarrow \psi_\theta / \|\psi_\theta\|_2$.

The ensemble model represents a density matrix as a convex mixture of pure NQS states,

$$\rho_\theta = \sum_{k=1}^K w_k |\psi_{\theta_k}\rangle \langle \psi_{\theta_k}|, \quad w_k = \frac{e^{\ell_k}}{\sum_{j=1}^K e^{\ell_j}}. \quad (\text{A5})$$

Each component $|\psi_{\theta_k}\rangle$ is parametrized by its own set of RBM weights and biases $\theta_k = \{a^{(k)}, b^{(k)}, W^{(k)}, c^{(k)}, d^{(k)}, U^{(k)}\}$, with the same dimensions as in the pure-state ansatz. The mixing weights are determined by unconstrained logits $\ell \in \mathbb{R}^K$.

For constrained NQS, the qubits are partitioned into disjoint clusters $\{C_b\}_b$, with $N_b = |C_b|$ and $\sum_b N_b = N$. Each block is assigned an RBM-based NQS of the same functional form as the global ansatz, but restricted to the visible spins s_{S_b} . A pure constrained NQS wave function is the product of wave functions, $\psi_{\text{SQNS}}(s) = \prod_b \psi_{\theta_b}^{(b)}(s_{C_b})$, with independent parameters $\theta_b = \{a_b, b_b, W_b, c_b, d_b, U_b\}$ of dimensions $a_b, c_b \in \mathbb{R}^{N_b}$, $b_b \in \mathbb{R}^{H_{\text{amp}}^{(b)}}$, $d_b \in \mathbb{R}^{H_{\text{phase}}^{(b)}}$, $W_b \in \mathbb{R}^{N_b \times H_{\text{amp}}^{(b)}}$, $U_b \in \mathbb{R}^{N_b \times H_{\text{phase}}^{(b)}}$. In the simulations, if block hidden sizes are not set explicitly, one uses $H_{\text{amp}}^{(b)} = H_{\text{amp}}$ and $H_{\text{phase}}^{(b)} = H_{\text{phase}}$.

We fixed the RBM hidden-layer sizes to $H_{\text{amp}} = 64$ and $H_{\text{phase}} = 64$, and initialize all trainable parameters (including ensemble logits) from an i.i.d. normal distribution with standard deviation 10^{-2} . Calculations are performed in double precision. Optimization is carried out by minimizing the negative log-likelihood with Adam optimizer using a cosine-decay learning-rate schedule (from $\eta_0 = 5 \times 10^{-3}$ to $10^{-2}\eta_0$ over 6000 steps) implemented in JAX/Optax [77].

-
- [1] O. Gühne and G. Tóth, Entanglement detection, *Phys. Rep.* **474**, 1 (2009).
 - [2] A. Sørensen and K. Mølmer, Entanglement and extreme spin squeezing, *Phys. Rev. Lett.* **86**, 4431 (2001).
 - [3] G. Tóth and O. Gühne, Detecting genuine multipartite entanglement with two local measurements, *Phys. Rev. Lett.* **94**, 060501 (2005).
 - [4] O. Gühne and M. Seevinck, Separability criteria for genuine multiparticle entanglement, *New J. Phys.* **12**, 053002 (2010).
 - [5] D. Bruß and G. Leuchs (eds.), *Quantum Information: From Foundations to Quantum Technology Applications* (Wiley, 2016).
 - [6] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, Cambridge, 2018).
 - [7] M. M. Wilde, *Quantum Information Theory* (Cambridge University Press, 2017).
 - [8] M. Huber and J. I. de Vicente, Structure of multidimensional entanglement in multipartite systems, *Phys. Rev. Lett.* **110**, 030501 (2013).
 - [9] M. Płodzień, J. Chwedeńczuk, M. Lewenstein, and G. Rajchel-Mieldzioć, Entanglement classification and non- k -separability certification via Greenberger-Horne-Zeilinger-class fidelity, *Phys. Rev. A* **110**, 032428 (2024).
 - [10] C. B. D. Goes, A. Canabarro, E. I. Duzzioni, and T. O. Maciel, Automated machine learning can classify bound entangled states with tomograms, *Quantum Inf. Process.* **20**, 99 (2021).
 - [11] B. C. Hiesmayr, Free versus bound entanglement, a NP-hard problem tackled by machine learning, *Sci. Rep.* **11**, 19739 (2021).
 - [12] M. Krawczyk, J. Pawłowski, M. M. Maška, and K. Roszak, Data-driven criteria for quantum correlations, *Phys. Rev. A* **109**, 022405 (2024).
 - [13] J. Pawłowski and M. Krawczyk, Identification of quantum entanglement with Siamese convolutional neural networks and semisupervised learning, *Phys. Rev. Appl.* **22**, 014068 (2024).
 - [14] N. Taghaddomi, A. Mani, A. Fahim, and A. Bakouei, Effective detection of quantum discord by using convolutional neural networks, *Quantum Mach. Intell.* **7**, 40 (2025).
 - [15] A. Dawid, J. Arnold, B. Requena, A. Gresch, M. Płodzień, K. Donatella, K. A. Nicoli, P. Stornati, R. Koch, M. Büttner, *et al.*, *Machine Learning in Quantum Sciences* (Cambridge University Press, 2025).
 - [16] S. Lu, S. Huang, K. Li, J. Li, J. Chen, D. Lu, Z. Ji, Y. Shen, D. Zhou, and B. Zeng, Separability-entanglement classifier via machine learning, *Phys. Rev. A* **98**, 012315 (2018).
 - [17] N. Asif, U. Khalid, A. Khan, T. Q. Duong, and H. Shin, Entanglement detection with artificial neural networks, *Sci. Rep.* **13**, 1562 (2023).
 - [18] J. Ureña, A. Sojo, J. Bermejo-Vega, and D. Manzano, Entanglement detection with classical deep neural networks, *Sci. Rep.* **14**, 18109 (2024).
 - [19] M. Wieśniak, Performance comparison of Gilbert’s algorithm and machine learning in classifying Bell-diagonal two-qutrit entanglement, *Sci. Rep.* **13**, 19500 (2023).
 - [20] D. Koutný, L. Motka, Z. Hradil, J. Řeháček, and L. L. Sánchez-Soto, Neural-network quantum state tomography, *Phys. Rev. A* **106**, 012409 (2022).
 - [21] A. M. Palmieri, G. Müller-Rigat, A. K. Srivastava, M. Lewenstein, G. Rajchel-Mieldzioć, and M. Płodzień, Enhancing quantum state tomography via resource-efficient attention-based neural networks, *Phys. Rev. Res.* **6**, 033248 (2024).
 - [22] P. Baláž, M. Krawczyk, J. Pawłowski, and K. Roszak, Learning entanglement from tomography data: Contradictory measure-

- ment importance for neural networks and random forests, *Phys. Rev. A* **112**, 022431 (2025).
- [23] Y. Chen, Y. Pan, G. Zhang, and S. Cheng, Detecting quantum entanglement with unsupervised learning, *Quantum Sci. Technol.* **7**, 015005 (2022).
- [24] C. Chen, C. Ren, H. Lin, and H. Lu, Entanglement structure detection via machine learning, *Quantum Sci. Technol.* **6**, 035017 (2021).
- [25] J. Y. Khoo and M. Heyl, Quantum entanglement recognition, *Phys. Rev. Res.* **3**, 033135 (2021).
- [26] R. Li, J. Du, Z. Qin, S. Zhang, C. Du, Y. Zhou, and Z. Xiao, Entanglement structure detection via computer vision, *Phys. Rev. A* **110**, 012448 (2024).
- [27] R. Li, S. Zhang, Z. Qin, C. Du, Y. Zhou, and Z. Xiao, Low-cost detection of high-dimensional multipartite entanglement structures, *Phys. Rev. Appl.* **23**, 044033 (2025).
- [28] J. Gray, L. Banchi, A. Bayat, and S. Bose, Machine-learning-assisted many-body entanglement measurement, *Phys. Rev. Lett.* **121**, 150503 (2018).
- [29] J. Roik, K. Bartkiewicz, A. Černoch, and K. Lemr, Entanglement quantification from collective measurements processed by machine learning, *Phys. Lett. A* **446**, 128270 (2022).
- [30] X. Lin, Z. Chen, and Z. Wei, Quantifying unknown entanglement by neural networks, *Quantum Info. Proc.* **22**, 341 (2023).
- [31] C. Feng and L. Chen, Quantifying quantum entanglement via machine learning models, *Commun. Theor. Phys.* **76**, 075104 (2024).
- [32] Y. Wang, S. Wang, J. Xing, Y. Du, and X. Wu, Quantifying unknown multiqubit entanglement using machine learning, *Entropy* **27**, 185 (2025).
- [33] M. Yang, Y. Xiao, Z.-Y. Hao, Y.-W. Liao, J.-H. Cao, K. Sun, E.-H. Wang, Z.-H. Liu, Y. Shikano, J.-S. Xu, C.-F. Li, and G.-C. Guo, Entanglement quantification via weak measurements assisted by deep learning, *Photonics Res.* **12**, 712 (2024).
- [34] A. W. R. Smith, J. Gray, and M. S. Kim, Efficient quantum state sample tomography with basis-dependent neural networks, *PRX Quantum* **2**, 020348 (2021).
- [35] V. Wei, W. A. Coish, P. Ronagh, and C. A. Muschik, Neural-shadow quantum state tomography, *Phys. Rev. Res.* **6**, 023250 (2024).
- [36] A. Martínez-Sabiote, M. Skotiniotis, J. J. Bermejo-Vega, D. Manzano, and C. Cano, Entanglement detection with quantum-inspired kernels and SVMs, *J. Supercomput.* **82**, 100 (2026).
- [37] J. M. Varela, A. B. de Palhares, and D. H. G. Duarte, Entanglement detection and quantification through machine learning: A comprehensive review, *Braz. J. Phys.* **56**, 25 (2026).
- [38] L. Jiao, X. Song, C. You, X. Liu, L. Li, P. Chen, X. Tang, Z. Feng, F. Liu, Y. Guo, S. Yang, Y. Li, X. Zhang, W. Ma, S. Wang, J. Bai, and B. Hou, AI meets physics: A comprehensive survey, *Artif. Intell. Rev.* **57**, 256 (2024).
- [39] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [40] D.-L. Deng, X. Li, and S. Das Sarma, Quantum entanglement in neural network states, *Phys. Rev. X* **7**, 021021 (2017).
- [41] R. G. Melko, G. Carleo, J. Carrasquilla, and J. I. Cirac, Restricted Boltzmann machines in quantum physics, *Nat. Phys.* **15**, 887 (2019).
- [42] M. Medvidović and J. R. Moreno, Neural-network quantum states for many-body physics, *Eur. Phys. J. Plus* **139**, 631 (2024).
- [43] H. Lange, A. Van de Walle, A. Abedinnia, and A. Bohrdt, From architectures to applications: A review of neural quantum states, *Quantum Sci. Technol.* **9**, 040501 (2024).
- [44] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. G. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [45] M. J. S. Beach, I. De Vlucht, A. Golubeva, P. Huembeli, B. Kulchytskyy, X. Luo, R. G. Melko, E. Merali, and G. Torlai, Qucumber: Wavefunction reconstruction with neural networks, *SciPost Phys.* **7**, 009 (2019).
- [46] G. Torlai and R. G. Melko, Latent space purification via neural density operators, *Phys. Rev. Lett.* **120**, 240503 (2018).
- [47] G. Torlai, B. Timar, E. P. L. van Nieuwenburg, H. Levine, A. Omran, A. Keesling, H. Bernien, M. Greiner, V. Vuletić, M. D. Lukin, R. G. Melko, and M. Endres, Integrating neural networks with a quantum simulator for state reconstruction, *Phys. Rev. Lett.* **123**, 230504 (2019).
- [48] M. Neugebauer, L. Fischer, A. Jäger, S. Czischek, S. Jochim, M. Weidemüller, and M. Gärtner, Neural-network quantum state tomography in a two-qubit experiment, *Phys. Rev. A* **102**, 042604 (2020).
- [49] E. S. Tiunov, V. V. Tiunova (Vyborova), A. E. Ulanov, A. I. Lvovsky, and A. K. Fedorov, Experimental quantum homodyne tomography via machine learning, *Optica* **7**, 448 (2020).
- [50] C. Harney, S. Pirandola, A. Ferraro, and M. Paternostro, Entanglement classification via neural network quantum states, *New J. Phys.* **22**, 045001 (2020).
- [51] C. Harney, M. Paternostro, and S. Pirandola, Mixed state entanglement classification using artificial neural networks, *New J. Phys.* **23**, 063033 (2021).
- [52] L. Comtet, *Advanced Combinatorics: The Art of Finite and Infinite Expansions*, translated by J. W. Nienhuys (D. Reidel Publishing Company, Dordrecht, 1974), See Chap. 5 on Partitions of a Set.
- [53] R. P. Stanley, *Enumerative Combinatorics* (Cambridge University Press, Cambridge, 2012), Vol. 1, 2nd ed. see Chap. 1 for set partitions and Bell numbers.
- [54] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete Mathematics: A Foundation for Computer Science*, 2nd ed. (Addison-Wesley, Reading, MA, USA, 1994).
- [55] P. J. Cameron, *Combinatorics: Topics, Techniques, Algorithms* (Cambridge University Press, Cambridge, 1994).
- [56] O. Sharir, Y. Levine, N. Wiesner, G. Carleo, and A. Shashua, Deep autoregressive models for the efficient variational simulation of many-body quantum systems, *Phys. Rev. Lett.* **124**, 020503 (2020).
- [57] J. A. Miszczak, Z. Puchala, P. Horodecki, A. Uhlmann, and K. Życzkowski, Sub- and super-fidelity as bounds for quantum fidelity, *Quantum Inf. Comput.* **9**, 103 (2009).
- [58] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University Press, 2010).
- [59] A. Dawid, P. Huembeli, M. Tomza, M. Lewenstein, and A. Dauphin, Phase detection with neural networks: Interpreting the black box, *New J. Phys.* **22**, 115001 (2020).
- [60] A. Dawid, P. Huembeli, M. Tomza, M. Lewenstein, and A. Dauphin, Hessian-based toolbox for reliable and interpretable

- machine learning in physics, *Mach. Learn.: Sci. Technol.* **3**, 015002 (2022).
- [61] S. J. Wetzel and M. Scherzer, Machine learning of explicit order parameters: From the Ising model to SU(2) lattice gauge theory, *Phys. Rev. B* **96**, 184410 (2017).
- [62] C. Miles, A. Bohrdt, R. Wu, C. Chiu, M. Xu, G. Ji, M. Greiner, K. Q. Weinberger, E. Demler, and E.-A. Kim, Correlator convolutional neural networks as an interpretable architecture for image-like quantum matter data, *Nat. Commun.* **12**, 3905 (2021).
- [63] A. Valenti, E. Greplova, N. H. Lindner, and S. D. Huber, Correlation-enhanced neural networks as interpretable variational quantum states, *Phys. Rev. Res.* **4**, L012010 (2022).
- [64] K. Cybiński, J. Enouen, A. Georges, and A. Dawid, Speak so a physicist can understand you! TetrisCNN for detecting phase transitions and order parameters, *arXiv:2411.02237*.
- [65] F. Döschl and A. Bohrdt, Importance of correlations for neural quantum states, *arXiv:2508.14152*.
- [66] A. Golubeva and R. G. Melko, Pruning a restricted Boltzmann machine for quantum state reconstruction, *Phys. Rev. B* **105**, 125124 (2022).
- [67] B. Barton, J. Carrasquilla, C. Roth, and A. Valenti, Connectivity determines the capability of sparse neural network quantum states, *Mach. Learn.: Sci. Technol.* **7**, 035025 (2026).
- [68] V. Hernandez, T. Spriggs, S. Khaleefah, and E. Greplova, Adiabatic fine-tuning of neural quantum states enables detection of phase transitions in weight space, *arXiv:2503.17140*.
- [69] M. S. Moss, A. Orfi, C. Roth, A. M. Sengupta, A. Georges, D. Sels, A. Dawid, and A. Valenti, Double descent: When do neural quantum states generalize? *Phys. Rev. E* **113**, 045303 (2026).
- [70] E. Rrapaj and A. Roggero, Exact representations of many-body interactions with restricted-Boltzmann-machine neural networks, *Phys. Rev. E* **103**, 013302 (2021).
- [71] A. Decelle, C. Furtlehner, A. D. J. N. Gómez, and B. Seoane, Inferring effective couplings with restricted Boltzmann machines, *SciPost Phys.* **16**, 095 (2024).
- [72] A. Decelle, A. d. J. Navas Gómez, and B. Seoane, Inferring higher-order couplings with neural networks, *Phys. Rev. Lett.* **135**, 207301 (2025).
- [73] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, *Nat. Rev. Phys.* **5**, 9 (2023).
- [74] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [75] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, *arXiv:2104.13478*.
- [76] M. Płodzień, Neural-Quantum-States-for-entanglement-depth-certification, Github, 2025, <https://github.com/MarcinPlodzien/Neural-Quantum-States-for-entanglement-depth-certification>.
- [77] DeepMind, I. Babuschkin, K. Baumli, A. Bell, S. Bhupatiraju, J. Bruce, P. Buchlovsky, D. Budden, T. Cai, A. Clark, I. Danihelka, *et al.*, The DeepMind JAX ecosystem, Github, 2020, <https://github.com/google-deeppmind>.