

Reconstructing quantum molecular rotor ground statesIsaac J. S. De Vlucht,^{1,2} Dmitri Iouchtchenko³, Ejaaz Merali,^{1,2} Pierre-Nicholas Roy³, and Roger G. Melko^{1,2}¹*Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario N2L 3G1*²*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5*³*Department of Chemistry, University of Waterloo, Waterloo, Ontario N2L 3G1*

(Received 6 April 2020; revised 8 June 2020; accepted 9 June 2020; published 6 July 2020)

Nanomolecular assemblies of C_{60} can be synthesized to enclose dipolar molecules. The low-temperature states of such endofullerenes are described by quantum mechanical rotors, which are candidates for quantum information devices with higher-dimensional local Hilbert spaces. The experimental exploration of endofullerene arrays comes at a time when machine learning techniques are rapidly being adopted to characterize, verify, and reconstruct quantum states from measurement data. In this paper, we develop a strategy for reconstructing the ground state of chains of dipolar rotors using restricted Boltzmann machines (RBMs) adapted to train on data from higher-dimensional Hilbert spaces. We demonstrate accurate generation of energy expectation values from an RBM trained on data in the free-rotor eigenstate basis and explore the learning resources required for various chain lengths and dipolar interaction strengths. Finally, we show evidence for fundamental limitations in the accuracy achievable by RBMs due to the difficulty in imposing symmetries in the sampling procedure. We discuss possible avenues to overcome this limitation in the future, including the further development of autoregressive models such as recurrent neural networks for the purposes of quantum state reconstruction.

DOI: [10.1103/PhysRevB.102.035108](https://doi.org/10.1103/PhysRevB.102.035108)**I. INTRODUCTION**

The upcoming decade will usher in an era of increasing hybridization between two rapidly advancing technologies: artificial intelligence and near-term quantum devices. The abundance of data produced by today's quantum hardware has driven a number of machine learning innovations for reconstructing the underlying quantum state of a system from data sets comprised of projective measurements on individual qubits. A leading strategy uses generative models to learn and represent the quantum state [1–6]. This approach has culminated in the recent demonstration of the reconstruction of a Rydberg atom quantum simulator from experimental data [7]. Methods for quantum state reconstruction with generative models are advancing rapidly and are capable of learning both pure state wave functions and mixed state density matrices [6,8,9]. Reconstruction is driven by projective qubit measurement data in a variety of bases or more general positive operator valued measures (POVMs) [10].

Although the potential of quantum hardware comprised of qubits has yet to be fully harnessed, the search for other candidates as a foundation for quantum computers also continues in earnest. Generalizations from two-dimensional local Hilbert spaces (qubits) to d -dimensional ones (qudits), or to continuous variables, can yield a significantly different platform on which to devise quantum information protocols. Particularly intriguing are recent proposals to harness the infinite-dimensional Hilbert space of rotational states of a rigid body [11]. Noise-robust embeddings of finite-dimensional systems within this infinite-dimensional Hilbert space are possible, including error-correcting “molecular” codes [12].

A promising experimental platform to provide access to rotational quantum degrees of freedom is that of nanomolecular assemblies (NMAs) of endofullerenes. Endofullerenes consist of a molecule encapsulated inside a carbon cage such as the buckminsterfullerene C_{60} and can be synthesized using molecular surgery techniques [13]. It was recently suggested that NMAs composed of linear arrays of molecular endofullerenes could be used as quantum information devices [14]. These linear arrays can be synthesized using nanotechnology tools such as nanotube peapods [15]. An additional handle on the properties of the NMAs is to use their electronic degrees of freedom in order to enhance transport [16] or enable field-effect transistor behavior [17].

A more tunable medium is provided by adding the presence of dipolar molecules within the C_{60} cages, which gives discrete rotational quantum numbers appropriate for quantum information processing. One advantage of using NMAs of endofullerenes for quantum information processing is in their ability to retain quantum effects at relatively high temperatures, compared to some current qubit-based devices [18]. The presence of dipole-dipole interactions between neighboring molecules leads to coupling between the rotors and strong correlations in the system [19,20]. These correlations lead to entanglement between the rotors.

In order to study these properties, one must first calculate the many-body ground states of the NMA. This can be achieved using exact diagonalization (ED) techniques along with basis truncation [14], an approach that however scales exponentially with system size. A scalable alternative is the density matrix renormalization group (DMRG) [21–23], which has recently been extended to calculate the ground state of large chains of coupled dipolar rotors [24]. Although

these are first steps towards investigating the practicality of NMAs of endofullerenes as quantum hardware, an exploration of techniques for efficient state reconstruction using modern numerical methods is an obvious additional step in their development.

In this paper, we construct a microscopic simulation of an NMA of endofullerenes. Using a model Hamiltonian consisting of rigid rotors with dipole-dipole interactions, the ground state wave function is studied using ED and DMRG. From these simulations we observe a nontrivial sign structure of the ground state wave function in the free-rotor eigenstate basis of the rotating dipolar molecule. Despite this, we are able to reconstruct states using standard machine learning generative models solely in this computational basis. We demonstrate the procedure using a restricted Boltzmann machine (RBM), a stochastic neural network comprised of two layers connected by weight parameters. While an RBM commonly has binomial units in its visible layer, in our case it is generalized to have a visible layer consisting of multinomial units, representing the angular momentum Hilbert space of each rotor. Synthetic data is obtained through an exact sampling procedure on the DMRG wave function and is used to train the RBM parameters through a stochastic gradient descent algorithm. Through this procedure, we are able to demonstrate accurate reconstruction of the ground state energy for systems of up to eight endofullerene sites. We further characterize the reconstruction procedure by investigating the hidden layer size and input dataset size required to obtain an energy difference equal to 5% of the first excitation gap.

II. NANOMOLECULAR ASSEMBLIES OF ENDOFULLERENES

A. Hamiltonian and Hilbert space

The Hamiltonian describing a linear chain of N dipolar molecules that are equally spaced a distance r apart includes, to a first approximation, the rotational energy of each molecule with a rotational constant B and Coulombic dipole-dipole interactions with a dipole moment μ :

$$\hat{H} = \frac{B}{\hbar^2} \sum_{i=1}^N \hat{\ell}_i^2 + \frac{\mu^2}{4\pi\epsilon_0 r^3} \sum_{i<j} \hat{V}_{ij}, \quad (1)$$

where

$$\hat{\ell}_i^2 |l_j m_j\rangle = \hbar^2 \ell_i (\ell_i + 1) \delta_{ij} |l_i m_i\rangle, \quad (2)$$

and

$$\hat{V}_{ij} = \frac{\hat{x}_i \hat{x}_j + \hat{y}_i \hat{y}_j - 2\hat{z}_i \hat{z}_j}{|i-j|^3}. \quad (3)$$

Without loss of generality, we can write this Hamiltonian in dimensionless form as [24]

$$\frac{\hat{H}}{B} = \hat{K} + \frac{1}{R^3} \hat{V} = \sum_{i=1}^N \frac{\hat{\ell}_i^2}{\hbar^2} + \frac{1}{R^3} \sum_{i<j} \hat{V}_{ij}, \quad (4)$$

where,

$$R = r \left(\frac{4\pi\epsilon_0 B}{\mu^2} \right)^{1/3} \quad (5)$$

is a dimensionless quantity. When each dipolar molecule is encapsulated by a fullerene (forming an endofullerene), the resulting screening of interactions can be modeled simply by reducing μ . For example, the dipole moment of HF is reduced from 1.8265 Debye in the gas phase, to an effective 0.45 Debye when encapsulated in C_{60} [18].

For dipolar diatomic molecules such as HF trapped in C_{60} , it has been shown that computed low lying eigenstates using *ab initio* methods exhibit small coupling between molecular rotations and translations [25]. These computed eigenstates are in agreement with experimentally measured spectra [18]. Moreover, given that diatomic vibrations are at a much higher energy, molecules can be approximated as being in their ground vibrational state and behaving like rigid rotors. With regards to interactions between the dipolar molecule and the C_{60} cage, it has also been shown that HF behaves as a nearly free rotor in C_{60} (with a renormalized rotational constant of 18.523 cm^{-1}) due to the icosahedral symmetry of the cage [25].

For realistic assemblies of endofullerenes, the rotational energy is much larger than the strength of the Coulombic dipole-dipole interaction ($R \gg 1$). For example, if the adjacent rotor distance is taken to be the diameter of C_{60} , 19.0 bohr [20], and given the renormalized rotational constant and screened dipole moment of HF inside C_{60} , we find $R \approx 2.64$. It is therefore a natural choice for the single-rotor reference basis to be the free-rotor eigenstate basis $|\ell_i m_i\rangle$. For convenience, we combine the angular momentum of a single rotor ℓ_i and its projection m_i into one zero-indexed integer label, σ_i :

$$\begin{aligned} \bigotimes_{i=1}^N |\ell_i m_i\rangle &= |\ell_1 m_1, \ell_2 m_2, \dots, \ell_N m_N\rangle \\ &= |\sigma_1, \sigma_2, \dots, \sigma_N\rangle = |\sigma\rangle. \end{aligned} \quad (6)$$

Note that since the Hamiltonian in Eq. (4) commutes with the z projection of the total angular momentum, $\sum_{i=1}^N \hat{\ell}_{i,z}$, and with $\bigotimes_{i=1}^N \hat{\pi}_i$ (where $\hat{\pi}_i$ is the parity operator that inverts the z_i coordinate in the position representation: $\hat{\pi}_i |x_i y_i z_i\rangle = |x_i y_i -z_i\rangle$), it conserves the total m value $m = \sum_{i=1}^N m_i$ and the total ℓ parity $\ell_p \equiv \sum_{i=1}^N \ell_i \pmod{2}$, respectively.

Given that the local Hilbert space of each rotor is infinite dimensional, a truncation of the Hilbert space is required in order to perform numerical studies. The truncation scheme adopted in this article is to cap the available angular momentum of each rotor to a maximum value ℓ_{\max} . Thus, the local Hilbert space size \mathcal{D} becomes

$$\mathcal{D} = \sum_{\ell=0}^{\ell_{\max}} (2\ell + 1) = (\ell_{\max} + 1)^2, \quad (7)$$

and the full Hilbert space size of the system of N endofullerenes is \mathcal{D}^N . As R decreases (Coulombic interactions increase), ℓ_{\max} must be increased in order to faithfully represent the contribution of the dipole-dipole interaction and, therefore, the ground state of the system.

B. Ground state wave function sign structure

We explore the use of RBMs in reconstructing the ground state wave function of Eq. (4) under the assumption that we have access to data obtained in the free-rotor eigenstate basis. In this data-driven setting, the Hamiltonian itself is not used to determine the ground state, e.g., through exact diagonalization, variational Monte Carlo, or some other method. Rather, it is assumed that an experimental system such as an endofullerene NMA is available, from which projective measurement data is produced which enables a reconstruction. Nonetheless, any *a priori* knowledge of the experimental Hamiltonian can be taken advantage of to increase the efficiency of the reconstruction. In particular, when applying the RBM (or any generative model) for this purpose, it is crucial to be aware of the sign structure of the wave function in the computational basis, since this has a direct impact on the complexity of the machine learning architecture required. If the ground state wave function is real and positive, it can trivially be equated to the square root of the probability density, and this simplifies the generative modeling drastically [6].

The well-known Perron-Frobenius theorem gives a sufficient (but not necessary) criterion for a Hamiltonian's ground state to be sign free (i.e., have no sign structure): that all off-diagonal matrix elements are nonpositive [26]. Hamiltonian matrices that satisfy this criterion are referred to as *stochastic* [27], but Eq. (4) in the computational basis does not generally fall into this class. Thus, realizations of Eq. (4) may have a ground state with sign structure. However, we argue below that this need not necessarily be a barrier to state reconstruction with standard generative models. In order to make this argument, we briefly explore some consequences of the sign structure in the context of quantum state reconstruction.

Given a normalized state $|\psi\rangle$ with real expansion coefficients $\langle \mathbf{n} | \psi \rangle$ in some computational basis \mathcal{B} , we fix an arbitrary $|\mathbf{n}^*\rangle \in \mathcal{B}$ and partition the set of all basis states,

$$\mathcal{B} = \{|\mathbf{n}\rangle\} = \mathcal{B}^+ \cup \mathcal{B}^-, \quad (8)$$

into the equivalence classes

$$\mathcal{B}^+ = \{|\mathbf{n}\rangle \in \mathcal{B} \mid \langle \mathbf{n} | \psi \rangle \langle \mathbf{n}^* | \psi \rangle \geq 0\} \quad (9a)$$

and

$$\mathcal{B}^- = \{|\mathbf{n}\rangle \in \mathcal{B} \mid \langle \mathbf{n} | \psi \rangle \langle \mathbf{n}^* | \psi \rangle < 0\}. \quad (9b)$$

The former contains those states $|\mathbf{n}\rangle$ whose coefficients $\langle \mathbf{n} | \psi \rangle$ have the same sign as $\langle \mathbf{n}^* | \psi \rangle$ (by convention called “positive,” disregarding any global phase), while the latter states have coefficients with the opposite sign (called “negative”). Indeed our formalism could be generalized to the case where the expansion coefficients are nontrivially complex, however we do not present that result in this paper. We emphasize that in our ED and DMRG calculations of Eq. (4), we observe that all coefficients are real numbers.

We let

$$\tau^\pm = \langle \psi | \hat{\mathcal{P}}^\pm | \psi \rangle = \sum_{|\mathbf{n}\rangle \in \mathcal{B}^\pm} |\langle \mathbf{n} | \psi \rangle|^2 \quad (10)$$

be the total contributions to the wave function from positive (τ^+) and negative (τ^-) coefficients, where

$$\hat{\mathcal{P}}^\pm = \sum_{|\mathbf{n}\rangle \in \mathcal{B}^\pm} |\mathbf{n}\rangle \langle \mathbf{n}| \quad (11)$$

are projection operators satisfying $\hat{\mathcal{P}}^+ + \hat{\mathcal{P}}^- = \hat{1}$. The normalization condition demands $\tau^+ + \tau^- = 1$.

We define the rectified state,

$$|\psi_{||}\rangle = \hat{\mathcal{P}}^+ |\psi\rangle - \hat{\mathcal{P}}^- |\psi\rangle, \quad (12)$$

which is normalized by construction, and whose coefficients may be written as

$$\langle \mathbf{n} | \psi_{||}\rangle = \text{sgn}(\langle \mathbf{n}^* | \psi \rangle) |\langle \mathbf{n} | \psi \rangle|, \quad (13)$$

making it clear that they are all positive. When $\mathcal{B}^- = \emptyset$ and $\tau^- = 0$, the original state is sign free (all coefficients are positive) and coincides with the rectified state: $|\psi_{||}\rangle = |\psi\rangle$. Otherwise, the rectified state is only an approximation to the original state, since all relative phase information is lost. The overlap between these states is given by

$$\langle \psi_{||} | \psi \rangle = \langle \psi | (\hat{\mathcal{P}}^+ - \hat{\mathcal{P}}^-) | \psi \rangle = 1 - 2\tau^-, \quad (14)$$

so τ^- quantifies the amount of error introduced into the state by the approximation. As shown in Appendix A, when $|\psi\rangle$ is an eigenstate of a real Hamiltonian \hat{H} with energy E , the error in the evaluation of the energy arising due to the state approximation is

$$\varepsilon = \langle \psi_{||} | \hat{H} | \psi_{||}\rangle - \langle \psi | \hat{H} | \psi \rangle \quad (15a)$$

$$= -4E\tau^- + 4\langle \psi | \hat{\mathcal{P}}^- \hat{H} \hat{\mathcal{P}}^- | \psi \rangle. \quad (15b)$$




FIG. 1. Convergence of the rectified state error measures τ^- (top panel) and $\varepsilon/\Delta E_1$ (bottom panel) with basis size truncation parameter ℓ_{\max} , for $N = 6$ and a variety of R values, computed via exact diagonalization. On this scale, crude convergence is achieved by $\ell_{\max} = 3$, and there is no visible improvement from $\ell_{\max} = 4$ to $\ell_{\max} = 5$.




FIG. 2. The rectified state error measures τ^- (top panel) and $\varepsilon/\Delta E_1$ (bottom panel) for $\ell_{\max} = 5$, computed via exact diagonalization.

We employ exact diagonalization to compute τ^- and $\varepsilon/\Delta E_1$ for the ground state of Eq. (4), using the basis truncation described in Ref. [14]. We set the energy scale to be ΔE_1 , which is the energy gap to the first excited state, and we take $|\mathbf{n}^*\rangle$ to be the noninteracting rotor ground state, $|00, \dots, 00\rangle$. From Fig. 1, it is clear that for the selected system parameters R and N , both quantities converge to small positive values with increasing basis truncation parameter ℓ_{\max} . The converged values (at $\ell_{\max} = 5$) are shown for an assortment of R and N combinations in Fig. 2. Both measures of error exhibit a peak as a function of separation distance R , but this peak shifts and shrinks with increasing system size N . Based on these results, we expect that for the parameters considered in this work ($R \geq 1$, $N \leq 8$), τ^- will be less than 10^{-4} , and $\varepsilon/\Delta E_1$ will not exceed 0.1%.

In Sec. III, we detail the reconstruction of the ground state of Eq. (4) from projective measurement data in the free-rotor basis. We consider training up to an energy threshold of 5% (see Sec. III B), which is substantially greater than $\varepsilon/\Delta E_1$, hence justifying the use of only a single basis [Eq. (6)].

III. QUANTUM STATE RECONSTRUCTION

In this section, we develop a machine learning method for the reconstruction of the ground state of a linear chain of N dipolar molecules as a model of endofullerene NMAs, with the Hamiltonian given in Eq. (4). The reconstruction proceeds by training a generative model, the restricted Boltzmann machine (RBM), on projective measurement data in the angular momentum basis only. The use of this as a computational basis is justified in Sec. II B. Quantum state reconstruction with RBMs is explained in detail in many other works [1,28–31]. However, a brief explanation of the expanded RBM scheme used in this study is offered here.




FIG. 3. Graphical representation of a multinomial RBM for learning rotor states. The visible nodes σ and the hidden nodes \mathbf{h} are represented as endofullerene cartoons and purple nodes, respectively. Connecting each visible unit to each hidden unit is a weight tensor \mathbf{W} . The external biases that couple to the visible (\mathbf{b}) and hidden (\mathbf{c}) layers are omitted from this depiction.

A. Multinomial restricted Boltzmann machine

The RBM is a two-layer, bidirectionally-connected neural network whose goal is to reconstruct a target probability distribution $q(\sigma)$ from data σ sampled from this distribution. In our case, the target distribution is taken to be the square of the ground state of Eq. (4),

$$q(\sigma) = |\psi_{\text{Exact}}(\sigma)|^2 = |\langle \sigma | \psi_{\text{Exact}} \rangle|^2. \quad (16)$$

We can calculate this ground state wave function using exact diagonalization (ED) or DMRG [24]. The data sets consist of one-hot projective measurements on each of the N sites in the reference basis [Eq. (6)] σ_{id} where i indexes the N sites and d indexes the \mathcal{D} possible projective measurement outcomes [Eq. (7)]. In other words, $\sigma_{id} = 1$ corresponds to rotor i having $(\ell_i, m_i) = \sigma_i = d$. Such a data set can be generated for small system sizes by ED. For larger system sizes, Ref. [24] presented a DMRG framework for calculating the ground state of Eq. (4), and the sampling algorithm by Ferris and Vidal can be used to generate training data in the $|\ell m\rangle$ basis [32].

Outside of machine learning in the physical sciences, multinomial data is prevalent. For instance, a dictionary has its elements (words) mapped onto a one-hot vector space in order to facilitate natural language processing. Another example is collaborative filtering, for which an extension of the traditional binary RBM to multinomial data was introduced as an application to the Netflix Prize in Ref. [30]. Such multinomial RBMs are highly suitable as generative models for the purpose of quantum state reconstruction, for example when the input datasets contain multiple outcomes, like POVMs [10]. In the following, we adapt a multinomial RBM for use on the ground state of Eq. (4) with data in the free rotor basis with a local Hilbert space size given by Eq. (7).

All N sites form the *visible* layer of the network, while interdependencies between sites in the visible layer are captured through n_h binary *hidden* nodes, taking on the values h_j (see Fig. 3). Each site in the visible layer is connected to each node in the hidden layer through a set of weights, \mathbf{W} , and each site and hidden node is coupled to an external bias field, \mathbf{b} and \mathbf{c} , respectively; we refer to all the parameters simultaneously as $\lambda = (\mathbf{W}, \mathbf{b}, \mathbf{c})$. There are no intralayer connections. At the core of this model is a joint probability distribution over both layers defined by a Boltzmann distribution,

$$p_\lambda(\sigma, \mathbf{h}) = \frac{e^{-E_\lambda(\sigma, \mathbf{h})}}{Z_\lambda}, \quad (17)$$

with an energy

$$\begin{aligned}
 E_\lambda(\boldsymbol{\sigma}, \mathbf{h}) = & - \sum_{i=1}^N \sum_{j=1}^{n_h} \sum_{d=1}^{\mathcal{D}} W_{ijd} h_j \sigma_{id} \\
 & - \sum_{i=1}^N \sum_{d=1}^{\mathcal{D}} \sigma_{id} b_{id} \\
 & - \sum_{j=1}^{n_h} h_j c_j,
 \end{aligned} \quad (18)$$

and a partition function (normalization)

$$Z_\lambda = \sum_{\boldsymbol{\sigma}, \mathbf{h}} e^{-E_\lambda(\boldsymbol{\sigma}, \mathbf{h})}. \quad (19)$$

The marginal distribution over the visible layer, $p_\lambda(\boldsymbol{\sigma})$, can be obtained by summing over the hidden layer:

$$p_\lambda(\boldsymbol{\sigma}) = \sum_{\mathbf{h}} p_\lambda(\boldsymbol{\sigma}, \mathbf{h}) = \frac{e^{-\mathcal{E}_\lambda(\boldsymbol{\sigma})}}{Z_\lambda}, \quad (20)$$

where,

$$\begin{aligned}
 \mathcal{E}_\lambda(\boldsymbol{\sigma}) = & - \sum_{i=1}^N \sum_{d=1}^{\mathcal{D}} b_{id} \sigma_{id} \\
 & - \sum_{j=1}^{n_h} \ln \left(1 + \exp \left(c_j + \sum_{i=1}^N \sum_{d=1}^{\mathcal{D}} W_{ijd} \sigma_{id} \right) \right).
 \end{aligned} \quad (21)$$

With the multinomial RBM thus defined, training proceeds as usual through minimization of the Kullback-Leibler divergence. Many of the details of training of the multinomial RBM are similar to a conventional RBM with binary hidden and visible layers. We provide these remaining details in Appendix B for completeness.

B. Training evaluators

We are interested in training an RBM from synthetic data generated from the exact ground state of Eq. (4), obtained from ED or DMRG. A central question is the quality of the quantum state that is reconstructed by the RBM. Here, we develop a criterion to assess the quality of reconstruction, which can be used on endofullerene systems of various sizes N . For example, one common criterion is the fidelity between the RBM state and the exact ground state. However, the fidelity is expensive to compute for large N , since the full marginal distribution p_λ must be calculated, including the normalization Z_λ . As noted previously [33], it is much more efficient to take advantage of the simplicity of Gibbs sampling in the RBM and to calculate the expectation value of a local observable such as the energy.

In order to treat the RBM as a wave-function ansatz, we must take the square root of the marginal distribution:

$$\psi_\lambda(\boldsymbol{\sigma}) = \sqrt{p_\lambda(\boldsymbol{\sigma})}. \quad (22)$$

The probability values are strictly non-negative, so the resulting state $|\psi_\lambda\rangle$ is inherently real and sign free in this basis. When the RBM is sufficiently well trained [i.e., $p_\lambda(\boldsymbol{\sigma}) \approx$

$q(\boldsymbol{\sigma})$], the state that it encodes should be the rectified version of $|\psi_{\text{Exact}}\rangle$. It is then straightforward (at least in principle) to compute the RBM expectation value of an arbitrary operator \hat{A} as

$$\langle \hat{A} \rangle_{\psi_\lambda} = \langle \psi_\lambda | \hat{A} | \psi_\lambda \rangle = \sum_{\boldsymbol{\sigma} \boldsymbol{\sigma}'} \psi_\lambda(\boldsymbol{\sigma}) \psi_\lambda(\boldsymbol{\sigma}') A_{\boldsymbol{\sigma} \boldsymbol{\sigma}'}, \quad (23)$$

where $A_{\boldsymbol{\sigma} \boldsymbol{\sigma}'} = \langle \boldsymbol{\sigma} | \hat{A} | \boldsymbol{\sigma}' \rangle$. In practice, this explicit summation is not tractable, so it is replaced by Monte Carlo sampling.

To evaluate the expectation value of Eq. (4), it is useful to separate it into its diagonal and off-diagonal terms; i.e. the rotational energy \hat{K} and dipole-dipole potential \hat{V} , respectively. Given a set of samples, $\mathbf{G} = \{\boldsymbol{\sigma}^g\}$, produced from p_λ by Gibbs sampling, one can trivially approximate the expectation value of the rotational energy term as a diagonal Monte Carlo estimator:

$$\langle \hat{K} \rangle_{\psi_\lambda} = \sum_{\boldsymbol{\sigma}} |\psi_\lambda(\boldsymbol{\sigma})|^2 \ell_{\boldsymbol{\sigma}}^2 \approx \frac{1}{|\mathbf{G}|} \sum_{g=1}^{|\mathbf{G}|} \ell_{\boldsymbol{\sigma}^g}^2, \quad (24a)$$

where

$$\ell_{\boldsymbol{\sigma}}^2 = \langle \boldsymbol{\sigma} | \hat{K} | \boldsymbol{\sigma} \rangle = \sum_{i=1}^N \ell_i(\ell_i + 1). \quad (24b)$$

To avoid calculating $\psi_\lambda(\boldsymbol{\sigma})$ for the potential energy term, the general expression in Eq. (23) can be approximated as an off-diagonal Monte Carlo estimator [2,6]:

$$\langle \hat{V} \rangle_{\psi_\lambda} \approx \frac{1}{|\mathbf{G}|} \sum_{g=1}^{|\mathbf{G}|} \sum_{\boldsymbol{\sigma}'} \frac{\tilde{\psi}_\lambda(\boldsymbol{\sigma}')}{\tilde{\psi}_\lambda(\boldsymbol{\sigma}^g)} V_{\boldsymbol{\sigma}^g \boldsymbol{\sigma}'}, \quad (25)$$

where

$$\tilde{\psi}_\lambda(\boldsymbol{\sigma}) = \psi_\lambda(\boldsymbol{\sigma}) \sqrt{Z_\lambda} = e^{-\mathcal{E}(\boldsymbol{\sigma})/2} \quad (26)$$

is the unnormalized RBM wave function, which may be evaluated without knowledge of the normalization Z_λ . As $V_{\boldsymbol{\sigma} \boldsymbol{\sigma}'}$ is very sparse [24], most of the terms in the sum over $\boldsymbol{\sigma}'$ in Eq. (25) may be skipped, avoiding the unfavorable scaling of the basis size. The total energy can therefore be efficiently approximated by adding Eqs. (24a) and (25):

$$E_{\text{RBM}} \approx \frac{1}{|\mathbf{G}|} \sum_{g=1}^{|\mathbf{G}|} \ell_{\boldsymbol{\sigma}^g}^2 + \frac{1}{R^3} \sum_{\boldsymbol{\sigma}'} \frac{\tilde{\psi}_\lambda(\boldsymbol{\sigma}')}{\tilde{\psi}_\lambda(\boldsymbol{\sigma}^g)} V_{\boldsymbol{\sigma}^g \boldsymbol{\sigma}'}. \quad (27)$$

Then, the training of the RBM can be evaluated by periodically generating samples from the RBM, and calculating the energy from those samples, along with the relative error in this energy as compared to the exact value. As explained in Appendix C, we use

$$\delta = \left| \frac{E_{\text{RBM}} - E_{\text{Exact}}}{\Delta E_1} \right| \quad (28)$$

as a measure of the quality of the state encoded in the RBM, with the exact energy and excitation gap calculated by ED or DMRG. In Sec. III C, we will use $\delta = 5\%$ as a *learning criterion*, in analogy to that defined in the scaling studies of Ref. [33].




FIG. 4. The error in the energy δ during the training process for $N = 4$ and $R = 1.0, 1.1$. For $R = 1.0, 1.1$, we used $n_h = 5, 4$, respectively, along with 10 000 input samples. The shaded region around each line is the standard error in δ .

C. Results

Endofullerene NMAs defined by Eq. (4) are characterized solely by the parameter R , which determines the relative strengths of the kinetic and potential energy terms. Previous Monte Carlo, path integral ground state, and DMRG simulations on finite-size systems have provided evidence that a continuous phase transition exists near $R = 1$ [24,34,35]. We refer to this as the *critical* value of R , for simplicity. In the following, we consider two values of R : $R = 1.0$ and 1.1 , both with $\ell_{\max} = 3$. For all RBMs, $\delta = 0.05$ was used as the learning criterion, i.e., the threshold at which training was considered to be complete. Note that since the endofullerene NMAs that we study are for finite sizes N , systems at criticality will not display gap closing. To calculate the energy from samples generated by the RBM during training (E_{RBM} from Eq. (27)), 7500 to 10 000 independent samples were generated using 7500 to 10 000 Gibbs steps for $N = 2, 4, 6$, and 8 . An example of the training process for $N = 4$ is given in Fig. 4.

The hidden layer serves the purpose of acting as a feature detector for the visible layer of the RBM. It is therefore expected that for systems that have long correlation lengths or high entanglement, a large number of hidden units may be required to obtain an accurate reconstruction. Similarly, quantum states with a larger number of nontrivial amplitudes in the computational basis will require more training data (i.e., higher sample complexity) in order to obtain an accurate reconstruction. In the language of one-dimensional chains of dipolar rotors governed by Eq. (4), one expects $R = 1.1$ to possess fewer nontrivial configurations outside of the free-rotor ground state ($|00, \dots, 00\rangle$) than at the critical value of $R = 1.0$. The relative contribution of the free-rotor ground state as a function of N and R is given in Table I. Evidently, for larger R and smaller N , the free-rotor ground state's probability amplitude vastly outweighs that of all other probability amplitudes.

From the above arguments, we expect that the ground state near the critical point ($R = 1.0$) should be harder to learn, both in terms of hidden layer size and sample complexity, than off-critical systems (such as $R = 1.1$). This hypothesis can be tested as a function of N , which can be used to elucidate the scaling behavior as in Ref. [33]. To do so, we train a

TABLE I. The ratio between the free-rotor ground state probability amplitude and the sum of all other probability amplitudes, $|\psi_{\text{Exact}}(0)|^2 / \sum_{i \neq 0} |\psi_{\text{Exact}}(i)|^2$, as calculated via DMRG.

	$N = 2$	$N = 3$	$N = 4$
$R = 1.0$	24.7	10.3	5.9
$R = 1.1$	274.1	132.2	86.6

variety of system sizes N , systematically increasing resources (hidden units, or data set size) until the RBM energy satisfies our learning criterion $\delta = 5\%$. This defines the minimum resources required for the reconstruction. The results of this procedure are shown in Fig. 5. Plot (a) shows the scaling of the ground state reconstruction with minimum hidden layer size (n_h). Similar to the procedure defined in Ref. [33], to make this scaling plot we assume access to unlimited data. In practice we find that $|\mathbf{D}| = \mathcal{O}(10^4)$ completely saturates the training data sample complexity for the system sizes studied. Figure 5(a) empirically defines $n_{h,\min}$ as a function of N . As evident from the plot, not only does the critical point $R = 1.0$ require a larger number of hidden units than the off-critical $R = 1.1$ for system sizes $N > 2$, the functional form of the scaling curve also appears to be different. Next, in Fig. 5(b), we fix the number of hidden units in the RBM to $n_{h,\min} + 1$ and systematically increase the size of the training data set $|\mathbf{D}|$ until the learning criterion is satisfied. As hypothesized, the critical value of R requires a larger sample complexity for every $N > 2$, although it is harder to discern a simple scaling curve for $R = 1.0$. In both cases, one may question whether these scaling results lie in the asymptotic limit; clearly, future work on larger system sizes will be required to answer this.

Finally, we comment on efficiency issues inherent to the sampling procedure of the RBM, called *block Gibbs* sampling (see Appendix B). In any procedure that uses Gibbs

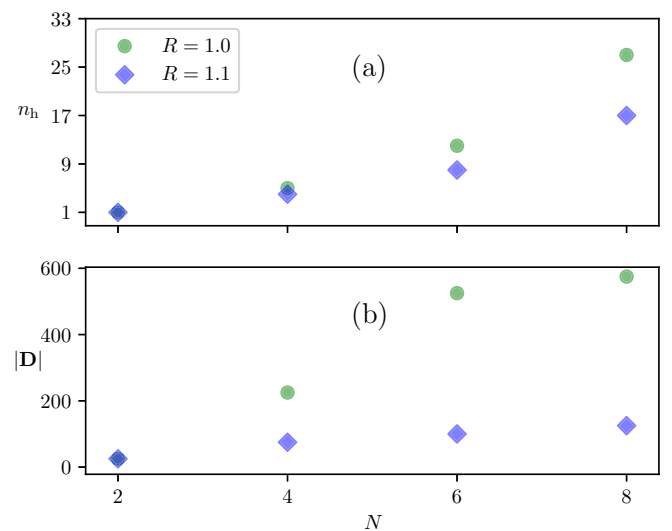


FIG. 5. (a) The minimum hidden layer size ($n_{h,\min}$) required to train an RBM to $\delta = 5\%$ given a saturated input data size. (b) The minimum input data size ($|\mathbf{D}|_{\min}$) required to train an RBM to $\delta = 5\%$ given a saturated hidden layer size.

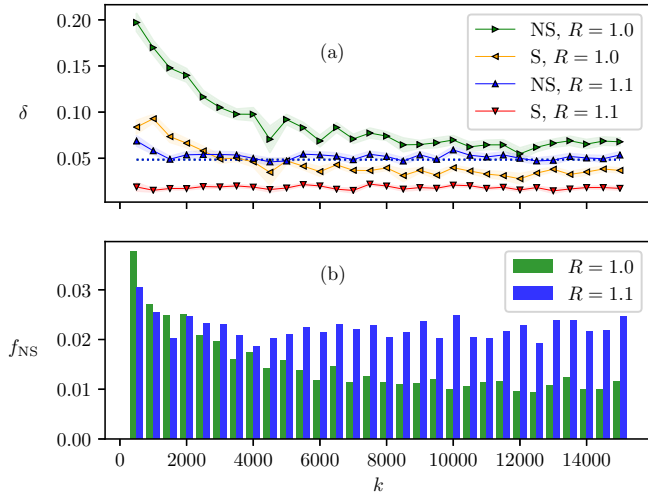


FIG. 6. (a) δ as a function of the number of Gibbs sampling steps k for trained RBMs for $N = 6$ using $n_{h,\min}$. Data denoted by S, $R = 1.0, 1.1$ correspond to δ as calculated by only counting samples satisfying the symmetries of Eq. (4), while NS, $R = 1.0, 1.1$ included all samples generated. The dashed line represents the value of δ that the RBMs are trained to. At each Gibbs sampling step, 10 000 independent samples were generated to calculate δ . (b) The fraction of samples f_{NS} generated at a given number of Gibbs sampling steps k that violate the symmetries of Eq. (4).

sampling, particularly training, or estimator calculation, the state of each visible node is produced independently from the others based on a conditional probability involving the state of all hidden nodes. The “restricted” nature of the RBM ensures this independence, which also means that conditional relationships between different visible units cannot be directly enforced. In other words, it is difficult to implement symmetry restrictions on visible states in a standard RBM that manifest as null probability amplitudes in the true ground state. Recall, as discussed in Sec. II A, the Hamiltonian in Eq. (4) conserves the total m value $m = \sum_{i=1}^N m_i$ and the total ℓ parity $\ell_p \equiv \sum_{i=1}^N \ell_i \pmod{2}$. In the free-rotor basis [Eq. (6)], this manifests in the ground state wave function as null probability amplitudes for configurations that violate these symmetries.

Figure 6(a) shows δ [Eq. (28)] as a function of the number of Gibbs steps k as calculated from trained RBMs for $N = 6$ (when initialized to an all-zero configuration). Even for our threshold of $\delta \leq 5\%$, a non-negligible amount of samples generated from the RBM violate the total m and ℓ_p symmetries as evident in Fig. 6(b). If these symmetry-violating samples are removed, Fig. 6(a) shows that the value of δ improves significantly, as expected. We note that the Gibbs sampling chain equilibration time here is much larger than that required for similar studies on the transverse-field Ising model [33]. Such shortcomings in generating symmetry-violating samples via Gibbs sampling can be overcome by other classes of generative models, in particular wave functions such as recurrent neural networks, where recent work has demonstrated the enforcement of symmetries in the autoregressive architecture [36].

IV. DISCUSSION

In this paper we study NMAs of endofullerenes through a model of one-dimensional quantum rotors with Coulombic dipole-dipole interactions. As part of an exploration of their suitability for quantum information processing purposes, we ask whether efficient quantum state reconstruction is possible for the ground states of such systems, using recently developed machine learning techniques. In the case of the dipolar quantum rotor chains, the model Hamiltonian is nonstoquastic in the angular momentum basis, implying that a trivial (positive) sign structure is not guaranteed for the ground state wave function. This fact holds significance for the prospects of reconstruction with generative models, which are more efficient for real and positive wave functions, with measurement outcomes that can be mapped directly to classical probability distributions. We study the ground state wave function numerically on finite system sizes using ED and DMRG. By analyzing the relative error associated with neglecting the negative wave-function amplitudes, we determine that their contribution is negligible in the expectation value of the energy up to typical uncertainties in the machine learning reconstruction procedure.

In light of this, we demonstrate here a machine learning reconstruction technique of using an RBM with multinomial visible units to capture the local Hilbert space of the dipolar rotor. Using the free-rotor angular momentum states as a computational basis, we generate synthetic measurement data using DMRG and train a multinomial RBM on a variety of one-dimensional chains of length N . Examining the energy difference between the RBM and the DMRG, we find that typical accuracies obtained by the machine learning procedure are well above the errors associated with neglecting negative amplitudes. Defining a *learning criterion* as 5% difference, in units of the first excitation gap, we examine the scaling of the reconstruction for two values parametrizing the strength of the dipolar interaction: one near a quantum critical point and one farther away. We observe that for the critical system, the size of the RBM hidden layer required to achieve the learning criterion is larger than for the off-critical system, assuming the availability of a sufficiently large measurement dataset. Conversely, holding the number of hidden units fixed near the minimum required to achieve the learning criterion, critical dipolar chains require significantly more data to train.

In the future, it would be interesting to continue this investigation for much larger system sizes to attempt to extract asymptotic scaling functions, as was done previously at the critical point of the transverse-field Ising model [33]. Scaling functions could be investigated for multiple criteria on different local and nonlocal estimators, such as correlation functions or entanglement entropies. From the observation in this work of significant fluctuations in the energy estimator of trained RBMs, it is possible that other types of generative models that do not rely on Gibbs sampling may be more suitable for such investigations. In particular, autoregressive models such as recurrent neural networks allow various symmetries to be taken into account when generating samples [36]. Their use in quantum state reconstruction is being explored in earnest.

Finally, we note that there are in general two classes of quantum wave functions from the perspective of machine learning state reconstruction. First are those that can be efficiently learned with standard generative models, which are designed to represent classical probability distributions. This class requires measurement data in only the computational basis. Second are those wave functions that do not have strictly real and positive amplitudes in the computational basis. These require generative models that are modified to account for the sign or phase structure, as well as requiring measurement data in a significant number of different bases to satisfy informational completeness. However, as we have shown in the case of dipolar quantum rotors chains, some quantum systems that do not rigorously fall in the first class may have a sufficiently small number of nonpositive amplitudes to allow for an accurate reconstruction when one is interested in reproducing local observables. It is an interesting theoretical question to ask how many physical systems fall into the class of states that are efficient to reconstruct with standard machine learning techniques, in spite of their sign structure.

ACKNOWLEDGMENTS

We thank J. Carrasquilla, A. Golubeva, and G. Torlai for many enlightening discussions. R.G.M. is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), Canada Research Chair (CRC) program, and the Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities. P.-N.R. (RGPIN-2016-04403) acknowledges NSERC, the Ontario Ministry of Research and Innovation (MRI), the CRC program (950-231024), the Canada Foundation for Innovation (CFI) (project No. 35232), and the Canada First Research Excellence Fund (CFREF).

APPENDIX A: RECTIFIED STATE EXPECTATION VALUES

We wish to determine the error

$$\varepsilon = \langle \psi_{||} | \hat{A} | \psi_{||} \rangle - \langle \psi | \hat{A} | \psi \rangle \quad (\text{A1})$$

made in the expectation value of an operator \hat{A} by using the rectified state $|\psi_{||}\rangle$ as an approximation for $|\psi\rangle$. By making use of the identity $\hat{P}^+ + \hat{P}^- = \hat{1}$, the first term may be expressed as

$$\langle \psi_{||} | \hat{A} | \psi_{||} \rangle = \langle \psi | \hat{A} | \psi \rangle - 2\langle \psi | \hat{P}^+ \hat{A} \hat{P}^- + \hat{P}^- \hat{A} \hat{P}^+ | \psi \rangle, \quad (\text{A2})$$

so we immediately obtain the general expression

$$\varepsilon = -2\langle \psi | \hat{P}^+ \hat{A} \hat{P}^- + \hat{P}^- \hat{A} \hat{P}^+ | \psi \rangle \quad (\text{A3a})$$

or, equivalently,

$$\varepsilon = 2\langle \psi_{||} | \hat{P}^+ \hat{A} \hat{P}^- + \hat{P}^- \hat{A} \hat{P}^+ | \psi_{||} \rangle. \quad (\text{A3b})$$




FIG. 7. Correlation between the rectification error measures in the top and bottom panels of Fig. 2. Multiple points for a single system size N correspond to different separation distances R .

The second form follows from $\hat{P}^\pm |\psi_{||}\rangle = \pm \hat{P}^\pm |\psi\rangle$. When \hat{A} is Hermitian ($\hat{A}^\dagger = \hat{A}$), this simplifies to

$$\varepsilon = -4\Re\langle \psi | \hat{P}^+ \hat{A} \hat{P}^- | \psi \rangle, \quad (\text{A4})$$

while for diagonal \hat{A} we get

$$\varepsilon = 0. \quad (\text{A5})$$

Of particular interest is the case of a real Hamiltonian \hat{H} having $|\psi\rangle$ as an eigenstate with eigenvalue E . The energy error due to the coefficient rectification is

$$\varepsilon = -4E\tau^- + 4\langle \psi | \hat{P}^- \hat{H} \hat{P}^- | \psi \rangle, \quad (\text{A6})$$

which suggests that τ^- and ε are intimately connected. This relationship is demonstrated for small systems in Fig. 7, where the comparison of τ^- and $\varepsilon/\Delta E_1$ values from Fig. 2 across several combinations of the R and N system parameters reveals a positive correlation.

APPENDIX B: TRAINING THE RESTRICTED BOLTZMANN MACHINE

To train the RBM [i.e., tune the parameters λ , such that $p_\lambda(\sigma)$ approximates Eq. (16)], the *Kullback-Leibler (KL) divergence* is used as the RBM's cost function C_λ , where

$$C_\lambda = \sum_\sigma q(\sigma) \ln \left(\frac{q(\sigma)}{p_\lambda(\sigma)} \right). \quad (\text{B1})$$

The optimization is carried out through stochastic gradient descent with the update rule

$$\lambda \leftarrow \lambda - \eta \nabla_\lambda C_\lambda, \quad (\text{B2})$$

$$\nabla_\lambda = \left[\frac{\partial}{\partial W_{ij}}, \frac{\partial}{\partial b_{id}}, \frac{\partial}{\partial c_j} \right], \quad (\text{B3})$$

where η is referred to as the *learning rate* and $\nabla_\lambda C_\lambda = -\nabla_\lambda \langle \ln(p_\lambda(\sigma)) \rangle_q$. The expectation value over the target distribution $q(\sigma)$ is approximated to be the empirical distribution of a set of training data samples \mathbf{D} (many sets of local projective measurements of ℓ and m for each site). $\nabla_\lambda C_\lambda$ can then be simplified to

$$\nabla_\lambda C_\lambda \approx \langle \nabla_\lambda \mathcal{E}_\lambda(\sigma) \rangle_{\mathbf{D}} - \langle \nabla_\lambda \mathcal{E}_\lambda(\sigma) \rangle_{p_\lambda}, \quad (\text{B4})$$

where

$$\langle \nabla_{\lambda} \mathcal{E}_{\lambda}(\boldsymbol{\sigma}) \rangle_{\mathbf{D}} = \frac{1}{|\mathbf{D}|} \sum_{\boldsymbol{\sigma} \in \mathbf{D}} \nabla_{\lambda} \mathcal{E}_{\lambda}(\boldsymbol{\sigma}). \quad (\text{B5})$$

The expectation value over p_{λ} poses problems, as one must calculate Z_{λ} which is computationally costly. This can be avoided by approximating this expectation value as an estimator over samples from p_{λ} using Gibbs sampling [29], wherein the variables in a given layer are sampled from its conditional distribution given the current configurations of the other layer: $p_{\lambda}(\sigma_{id} | \mathbf{h})$ and $p_{\lambda}(h_j = 1 | \boldsymbol{\sigma})$. Starting from an initial configuration of the visible layer (typically an all-zero configuration), $\boldsymbol{\sigma}^{(0)}$, one alternates between sampling a new set of h and $\boldsymbol{\sigma}$ via $p_{\lambda}(\sigma_{id} | \mathbf{h})$ and $p_{\lambda}(h_j = 1 | \boldsymbol{\sigma})$. The number of times one updates the original configuration is referred to as the number of Gibbs sampling steps k . If one generates a set of samples, $\boldsymbol{\Gamma}$, in this fashion, the expectation value over the marginal distribution approximates to

$$\langle \nabla_{\lambda} \mathcal{E}_{\lambda}(\boldsymbol{\sigma}) \rangle_{p_{\lambda}} \approx \frac{1}{|\boldsymbol{\Gamma}|} \sum_{\boldsymbol{\sigma}^{(k)} \in \boldsymbol{\Gamma}} \nabla_{\lambda} \mathcal{E}_{\lambda}(\boldsymbol{\sigma}^{(k)}). \quad (\text{B6})$$

To introduce some noise into the gradient updates so as to not get stuck in a local minimum of C_{λ} , stochastic gradient decent, calculating $\langle \nabla_{\lambda} \mathcal{E}_{\lambda}(\boldsymbol{\sigma}) \rangle_{\mathbf{D}}$ by way of many averaged updates over mini-batches, \mathbf{P} , of \mathbf{D} (equally sized disjoint subsets of \mathbf{D}) is performed instead [29]. The individual gradients of \mathcal{E}_{λ} with respect to the RBM parameters are

$$\frac{\partial \mathcal{E}_{\lambda}(\boldsymbol{\sigma})}{\partial W_{ijd}} = -p_{\lambda}(h_j = 1 | \boldsymbol{\sigma}) \sigma_{id}, \quad (\text{B7a})$$

$$\frac{\partial \mathcal{E}_{\lambda}(\boldsymbol{\sigma})}{\partial c_j} = -p_{\lambda}(h_j = 1 | \boldsymbol{\sigma})_j, \quad (\text{B7b})$$

$$\frac{\partial \mathcal{E}_{\lambda}(\boldsymbol{\sigma})}{\partial b_{id}} = -\sigma_{id}. \quad (\text{B7c})$$

TABLE II. Hyperparameters used for training.

Hyperparameter	Value	Description
η	0.001, 0.01	Learning rate
$ \mathbf{P} $	20	Positive-batch size
$ \boldsymbol{\Gamma} $	10	Negative-batch size
k	10	# of Gibbs steps

The mini-batch size, $|\mathbf{P}|$, is typically called the *positive-batch size*, while $|\boldsymbol{\Gamma}|$ is typically called the *negative-batch size*. One full pass through the entire data set \mathbf{D} is called an epoch. The training hyperparameters used throughout this study for $R = 1.0$ and 1.1 are given in Table II. A learning rate of 0.001 was employed for the hidden layer scaling, while 0.01 was used for the data size scaling.

APPENDIX C: TRAINING THRESHOLD

Let $|\psi\rangle$ be a normalized state that is an arbitrary superposition of the eigenstates, $|n\rangle$, of a Hamiltonian, \hat{H} : $|\psi\rangle = \sum_n c_n |n\rangle$. The deviation of the energy of this state, E_{ψ} , from the ground state energy of \hat{H} , E_0 , is

$$E_{\psi} - E_0 = \langle \psi | \hat{H} | \psi \rangle - \langle 0 | \hat{H} | 0 \rangle = \sum_{n=1}^{\infty} |c_n|^2 \Delta E_n, \quad (\text{C1})$$

where $\Delta E_n = E_n - E_0$. Relative to the first excited state gap, ΔE_1 , the above expression may be written as

$$\frac{E_{\psi} - E_0}{\Delta E_1} = |c_1|^2 + \sum_{n=2}^{\infty} |c_n|^2 \frac{\Delta E_n}{\Delta E_1}. \quad (\text{C2})$$

When E_{ψ} approaches E_0 , we expect the dominant contribution to the difference to be from the first excited state. Therefore, neglecting all terms beyond the first on the right-hand side of Eq. (C2), we see that the left-hand side should act as a proxy for the amount of excited state contamination in $|\psi\rangle$.

- [1] G. Torlai and R. G. Melko, *Phys. Rev. B* **94**, 165134 (2016).
 [2] M. Kieferová and N. Wiebe, *Phys. Rev. A* **96**, 062327 (2017).
 [3] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. G. Melko, and G. Carleo, *Nat. Phys.* **14**, 447 (2018).
 [4] H. J. Kappen, *J. Phys. A: Math. Theor.* **53**, 214001 (2020).
 [5] G. Torlai and R. G. Melko, *Annu. Rev. Condens. Matter Phys.* **11**, 325 (2020).
 [6] M. Beach, I. De Vlucht, A. Golubeva, P. Huembeli, B. Kulchytskyy, X. Luo, R. Melko, E. Merali, and G. Torlai, *SciPost Phys.* **7**, 009 (2019).
 [7] 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, *Phys. Rev. Lett.* **123**, 230504 (2019).
 [8] G. Torlai and R. G. Melko, *Phys. Rev. Lett.* **120**, 240503 (2018).
 [9] A. Melkani, C. Gneiting, and F. Nori, [arXiv:1911.07506](https://arxiv.org/abs/1911.07506).
 [10] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, *Nat. Mach. Intellig.* **1**, 155 (2019).
 [11] D. Gottesman, A. Kitaev, and J. Preskill, *Phys. Rev. A* **64**, 012310 (2001).
 [12] V. V. Albert, J. P. Covey, and J. Preskill, [arXiv:1911.00099v2](https://arxiv.org/abs/1911.00099v2).
 [13] K. Komatsu, M. Murata, and Y. Murata, *Science* **307**, 238 (2005).
 [14] T. Halverson, D. Iouchtchenko, and P. N. Roy, *J. Chem. Phys.* **148**, 074112 (2018).
 [15] S. Berber, Y.-K. Kwon, and D. Tománek, *Phys. Rev. Lett.* **88**, 185502 (2002).
 [16] B. Xu and X. Chen, *Phys. Rev. Lett.* **110**, 252001 (2013).
 [17] C. H. Lee, K. T. Kang, K. S. Park, M. S. Kim, H. S. Kim, H. G. Kim, J. E. Fischer, and A. T. Johnson, *Jpn. J. Appl. Phys.* **42**, 5392 (2003).
 [18] A. Krachmalnicoff, R. Bounds, S. Mamone, S. Alom, M. Concistrè, B. Meier, K. Kouřil, M. E. Light, M. R. Johnson, S. Rols, A. J. Horsewill, A. Shugai, U. Nagel, T. Rõm, M. Carravetta, M. H. Levitt, and R. J. Whitby, *Nat. Chem.* **8**, 953 (2016).

- [19] M. Xu, S. Ye, and Z. Bačić, *J. Phys. Chem. Lett.* **6**, 3721 (2015).
- [20] P. M. Felker and Z. Bačić, *J. Chem. Phys.* **146**, 084303 (2017).
- [21] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [22] S. R. White, *Phys. Rev. B* **48**, 10345 (1993).
- [23] U. Schollwoeck, *Ann. Phys.* **326**, 96 (2011).
- [24] D. Iouchtchenko and P. N. Roy, *J. Chem. Phys.* **148**, 134115 (2018).
- [25] Y. N. Kalugina and P.-N. Roy, *J. Chem. Phys.* **147**, 244303 (2017).
- [26] P. Tarazaga, M. Raydan, and A. Hurman, *Lin. Alg. Appl.* **328**, 57 (2001).
- [27] S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, and B. M. Terhal, [arXiv:quant-ph/0606140v4](https://arxiv.org/abs/quant-ph/0606140v4).
- [28] G. Torlai, Augmenting Quantum Mechanics with Artificial Intelligence, Ph.D. Thesis, UWSpace (2018).
- [29] G. E. Hinton, in *Neural Networks: Tricks of the Trade*, edited by G. Montavon, G. B. Orr, and K.-R. Müller (Springer, Berlin, Heidelberg, 2012), Vol. 7700, p. 599.
- [30] R. Salakhutdinov, A. Mnih, and G. Hinton, in *Proceedings of the 24th International Conference on Machine Learning - ICML '07* (ACM Press, Corvalis, Oregon, 2007), p. 791.
- [31] K. Murphy, *Machine Learning: A Probabilistic Perspective* (The MIT Press, Cambridge, MA, 2012).
- [32] A. J. Ferris and G. Vidal, *Phys. Rev. B* **85**, 165146 (2012).
- [33] D. Sehayek, A. Golubeva, M. S. Albergo, B. Kulchytskyy, G. Torlai, and R. Melko, *Phys. Rev. B* **100**, 195125 (2019).
- [34] B. P. Abolins, R. E. Zillich, and K. B. Whaley, *J. Low Temp. Phys.* **165**, 249 (2011).
- [35] T. Sahoo, D. Iouchtchenko, C. M. Herdman, and P.-N. Roy, *J. Chem. Phys.* **152**, 184113 (2020).
- [36] M. Hibat-Allah, M. Ganahl, L. E. Hayward, R. G. Melko, and J. Carrasquilla, *Phys. Rev. Research* **2**, 023358 (2020).