

PAPER • OPEN ACCESS

## Neural network enhanced measurement efficiency for molecular groundstates

To cite this article: Dmitri Iouchtchenko *et al* 2023 *Mach. Learn.: Sci. Technol.* **4** 015016

View the [article online](#) for updates and enhancements.

You may also like

- [Groundstate fidelity phase diagram of the fully anisotropic two-leg spin- \$\frac{1}{2}\$  XXZ ladder](#)  
Sheng-Hao Li, Qian-Qian Shi, Murray T Batchelor *et al.*
- [Bipartite fidelity for models with periodic boundary conditions](#)  
Alexi Morin-Duchesne, Gilles Perez and Jean Liénardy
- [Gapless quantum spin chains: multiple dynamics and conformal wavefunctions](#)  
Xiao Chen, Eduardo Fradkin and William Witczak-Krempa



## PAPER

## OPEN ACCESS

## RECEIVED

20 September 2022

## REVISED

8 December 2022

## ACCEPTED FOR PUBLICATION

19 January 2023

## PUBLISHED

9 February 2023

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



# Neural network enhanced measurement efficiency for molecular groundstates

Dmitri Iouchtchenko<sup>1,2,\*</sup> , Jérôme F Gonthier<sup>3</sup> , Alejandro Perdomo-Ortiz<sup>2</sup> and Roger G Melko<sup>1,4</sup> <sup>1</sup> Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada<sup>2</sup> Zapata Computing Canada Inc., 325 Front St W, Toronto, Ontario M5V 2Y1, Canada<sup>3</sup> Zapata Computing Inc., 100 Federal Street, Boston, MA 02110, United States of America<sup>4</sup> Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada

\* Author to whom any correspondence should be addressed.

E-mail: [diouchtc@uwaterloo.ca](mailto:diouchtc@uwaterloo.ca)**Keywords:** quantum tomography, classical shadow tomography, machine learning, neural networks, restricted Boltzmann machines, recurrent neural networks

## Abstract

It is believed that one of the first useful applications for a quantum computer will be the preparation of groundstates of molecular Hamiltonians. A crucial task involving state preparation and readout is obtaining physical observables of such states, which are typically estimated using projective measurements on the qubits. At present, measurement data is costly and time-consuming to obtain on any quantum computing architecture, which has significant consequences for the statistical errors of estimators. In this paper, we adapt common neural network models (restricted Boltzmann machines and recurrent neural networks) to learn complex groundstate wavefunctions for several prototypical molecular qubit Hamiltonians from typical measurement data. By relating the accuracy  $\varepsilon$  of the reconstructed groundstate energy to the number of measurements, we find that using a neural network model provides a robust improvement over using single-copy measurement outcomes alone to reconstruct observables. This enhancement yields an asymptotic scaling near  $\varepsilon^{-1}$  for the model-based approaches, as opposed to  $\varepsilon^{-2}$  in the case of classical shadow tomography.

## 1. Introduction

Finding the groundstate of a molecular Hamiltonian on a classical computer is a difficult task, but quantum phase estimation on a fault-tolerant quantum computer is known to efficiently obtain the groundstate energy under reasonable conditions [1, 2]. However, fault tolerance remains out of reach; to date, only noisy intermediate-scale quantum (NISQ) devices have been created, where the deep circuits required for phase estimation cannot be realized. Despite their shortcomings, it is believed that NISQ devices can still be used to assist classical ones in the task of calculating molecular groundstate energies [3]. Current approaches involve transforming the Hamiltonian from a second-quantized fermionic representation into qubit form [4], the groundstate of which could then be prepared on a quantum computer using strategies such as the variational quantum eigensolver (VQE) or other hybrid quantum–classical algorithms [5]. Once the groundstate—or a suitable approximation—is prepared, expectation values of physical observables can be estimated by performing measurements on the qubits.

State preparation and measurement times may vary depending on the specific experimental implementation. Regardless, given the significant hardware limitations in present-day quantum devices, it is crucial to extract information about physical observables as efficiently as possible using the available resources, including the collected measurement outcomes. The gold standard for probing an experimental state is quantum state tomography, which aims to accurately reconstruct the full wavefunction or density matrix, allowing for the general extraction of the energy and other properties. However, full state tomography comes at the cost of exponential growth of classical resources with the number of qubits, which limits its applicability to only the smallest of systems.

Beyond exact tomography, there are several promising protocols for the efficient analysis of measurement data provided by a quantum computer. In particular, it is possible to efficiently compute the expectation value of typical observables directly from samples, using single-qubit factorable positive operator valued measures (POVMs) [6] without a model-based tomographic reconstruction, or using classical shadow tomography [7–9]. In the case of typical observables like the energy, the sample complexity (defined as the number of copies  $S$  of the quantum state necessary to achieve an additive error  $\varepsilon$ ) scales as  $\varepsilon^{-2}$  for classical shadows [7].

One can also explore the accuracy of estimator reconstruction when measurement outcomes are used to train a generative model, instead of computing the expectation value directly from the samples. Motivated by techniques in machine learning, models based on neural networks can serve as powerful wavefunction ansatzes that are readily trainable from data [6, 10–12]. The goal of training a generative model is to learn the salient elements of the quantum state, while hopefully requiring fewer resources than state tomography when scaling to larger numbers of qubits [13]. Indeed, it was shown in [14] that training a restricted Boltzmann machine (RBM) [15, 16] and then sampling it (the ‘generative’ step) significantly reduced the variance of the energy estimator of some simple molecular Hamiltonians, at the possible cost of a bias due to training imperfections. Other approaches involving neural networks have recently been used for quantum control and sensing [17].

In this paper, we implement two different generative models as wavefunction ansatzes: complex RBMs [10, 16] and complex recurrent neural networks (RNNs) [12]. We use these models to reconstruct the groundstates of standard molecular Hamiltonians for LiH, BeH<sub>2</sub>, and H<sub>2</sub> from synthetic single-copy Pauli measurement data. We perform a comparison of the errors in the energy and—when possible—the fidelity, between the models, maximum likelihood pure state tomography, and classical shadows. We find that for both the RBM and RNN, the sample complexity scales approximately as  $\varepsilon^{-1}$ , as compared to  $\varepsilon^{-2}$  for classical shadows. Our results suggest that generative models may offer a robust advantage in reconstructing physical observables in cases where data is limited, as in present-day quantum computing devices.

## 2. Models

In this section, we provide details of the generative models that we use to reconstruct the molecular groundstates from qubit measurement data. We assume that, for a given basis, each projective measurement outcome is a vector with  $N$  binary values (one per qubit), and that these outcomes are distributed according to the Born rule. However, unlike standard generative models used for classical machine learning tasks, we require that our models are capable of fully describing a pure quantum state  $|\psi\rangle$ , whose coefficients have both an amplitude and a phase. We employ two of the most well-studied generative models suited to this purpose: the complex RBM and the complex RNN. The complex RBM takes on the form described in [16]: a regular RBM, but with complex-valued parameters. In the standard basis, the complex RBM state has the coefficients

$$\langle \sigma | \varphi_{\text{RBM}}(\boldsymbol{\lambda}) \rangle = \frac{1}{\sqrt{Z(\boldsymbol{\lambda})}} \sum_{\mathbf{h}} e^{\mathbf{h}^T \mathbf{W} \sigma + \mathbf{b} \cdot \sigma + \mathbf{c} \cdot \mathbf{h}} = \frac{1}{\sqrt{Z(\boldsymbol{\lambda})}} e^{\mathbf{b} \cdot \sigma} \prod_{j=1}^{N_h} [1 + e^{(\mathbf{W} \sigma + \mathbf{c})_j}], \quad (1a)$$

where

$$Z(\boldsymbol{\lambda}) = \sum_{\sigma} \left| \sum_{\mathbf{h}} e^{\mathbf{h}^T \mathbf{W} \sigma + \mathbf{b} \cdot \sigma + \mathbf{c} \cdot \mathbf{h}} \right|^2 \quad (1b)$$

is the squared normalization,  $\boldsymbol{\lambda}$  represents the parameters ( $\mathbf{W}$ ,  $\mathbf{b}$ ,  $\mathbf{c}$ ), and  $\mathbf{h}$  is the length- $N_h$  hidden state vector.

The complex RNN is the same as in [12], but we omit the softsign function for the phases. It has the standard-basis coefficients

$$\langle \sigma | \varphi_{\text{RNN}}(\boldsymbol{\lambda}) \rangle = \prod_{j=1}^N e^{2\pi i \theta_j(\sigma)} \sqrt{p_j(\sigma)}, \quad (2a)$$

where

$$\theta_j(\sigma) = (\mathbf{u} \cdot \mathbf{h}_j(\sigma) + a) \sigma_j + \mathbf{v} \cdot \mathbf{h}_j(\sigma) + b \quad (2b)$$

is the local phase,

$$p_j(\boldsymbol{\sigma}) = \frac{e^{(\mathbf{w} \cdot \mathbf{h}_j(\boldsymbol{\sigma}) + c)\sigma_j}}{1 + e^{\mathbf{w} \cdot \mathbf{h}_j(\boldsymbol{\sigma}) + c}} \quad (2c)$$

is the conditional probability,

$$\mathbf{h}_j(\boldsymbol{\sigma}) = \tanh(\mathbf{M}\mathbf{h}_{j-1}(\boldsymbol{\sigma}) + \mathbf{p}\sigma_{j-1} + \mathbf{q}) \quad (2d)$$

is the length- $N_h$  hidden state (with  $\tanh$  as the element-wise nonlinearity),  $\boldsymbol{\lambda}$  represents the parameters ( $\mathbf{M}$ ,  $\mathbf{p}$ ,  $\mathbf{q}$ ,  $\mathbf{u}$ ,  $\mathbf{v}$ ,  $\mathbf{w}$ ,  $a$ ,  $b$ ,  $c$ ), and  $N$  is the length of the basis state vector  $\boldsymbol{\sigma}$ . We do not employ a one-hot encoding, and we initialize the hidden state recurrence with  $\mathbf{h}_0 = (0, \dots, 0)$  and  $\sigma_0 = 0$ .

In the typical case of reconstructing a single probability distribution, RBMs and RNNs are trained by minimizing the cross entropy term of the Kullback–Leibler divergence. However, for a state  $|\psi\rangle$  with both an amplitude and a phase, projective measurements must be performed in multiple orthonormal bases to capture sufficient information. This results in multiple probability distributions, and a natural approach is to combine all the cross entropies into a single loss function [14]:

$$\mathcal{L}(\boldsymbol{\lambda}) = -\frac{1}{|\mathcal{D}|} \sum_{k=1}^K \sum_{\boldsymbol{\sigma} \in \mathcal{D}_k} \log \left| \langle \boldsymbol{\sigma} | \hat{R}_k^\dagger | \varphi(\boldsymbol{\lambda}) \rangle \right|^2, \quad (3)$$

where  $\hat{R}_k$  is a unitary operator that maps the standard basis to measurement basis  $k$ ,  $\mathcal{D}_k$  is the collection of standard-basis measurement outcomes obtained from the rotated target state  $\hat{R}_k^\dagger |\psi\rangle$ ,  $K$  is the number of bases (or ‘settings’) used for projective measurements, and  $|\mathcal{D}| = \sum_{k=1}^K |\mathcal{D}_k|$ . Further model and training details are provided in appendix A.

In addition, for the small system sizes considered in this work ( $N = 4, 6$ , and  $8$  qubits), it is feasible to represent the full state as a length- $2^N$  vector with complex coefficients  $\langle \boldsymbol{\sigma} | \varphi_{\text{WF}}(\boldsymbol{\lambda}) \rangle$ ; the parameters  $\boldsymbol{\lambda}$  of this ‘model’ are simply the coefficients themselves. To determine the coefficients, we begin with the observation that whenever the gradient of the loss function  $\mathcal{L}(\boldsymbol{\lambda})$  is zero (such as at a minimum), the corresponding state  $|\varphi(\boldsymbol{\lambda})\rangle$  satisfies

$$|\varphi(\boldsymbol{\lambda})\rangle = \frac{1}{|\mathcal{D}|} \sum_{k=1}^K \sum_{\boldsymbol{\sigma} \in \mathcal{D}_k} \frac{\hat{R}_k |\boldsymbol{\sigma}\rangle}{\langle \varphi(\boldsymbol{\lambda}) | \hat{R}_k |\boldsymbol{\sigma}\rangle} = T(|\varphi(\boldsymbol{\lambda})\rangle). \quad (4)$$

The nonlinear operator  $T$ —which resembles the composite physical imposition operator of [18]—has all such states as fixed points. We exploit this feature to find a complex state vector  $|\varphi_{\text{WF}}(\boldsymbol{\lambda})\rangle$  by starting from the exact state  $|\psi\rangle$  and applying  $T$  iteratively until a fixed point is reached:

$$|\varphi_{\text{WF}}(\boldsymbol{\lambda})\rangle = \tilde{T}(\tilde{T}(\dots \tilde{T}(\tilde{T}(|\psi\rangle)) \dots)), \quad (5a)$$

where

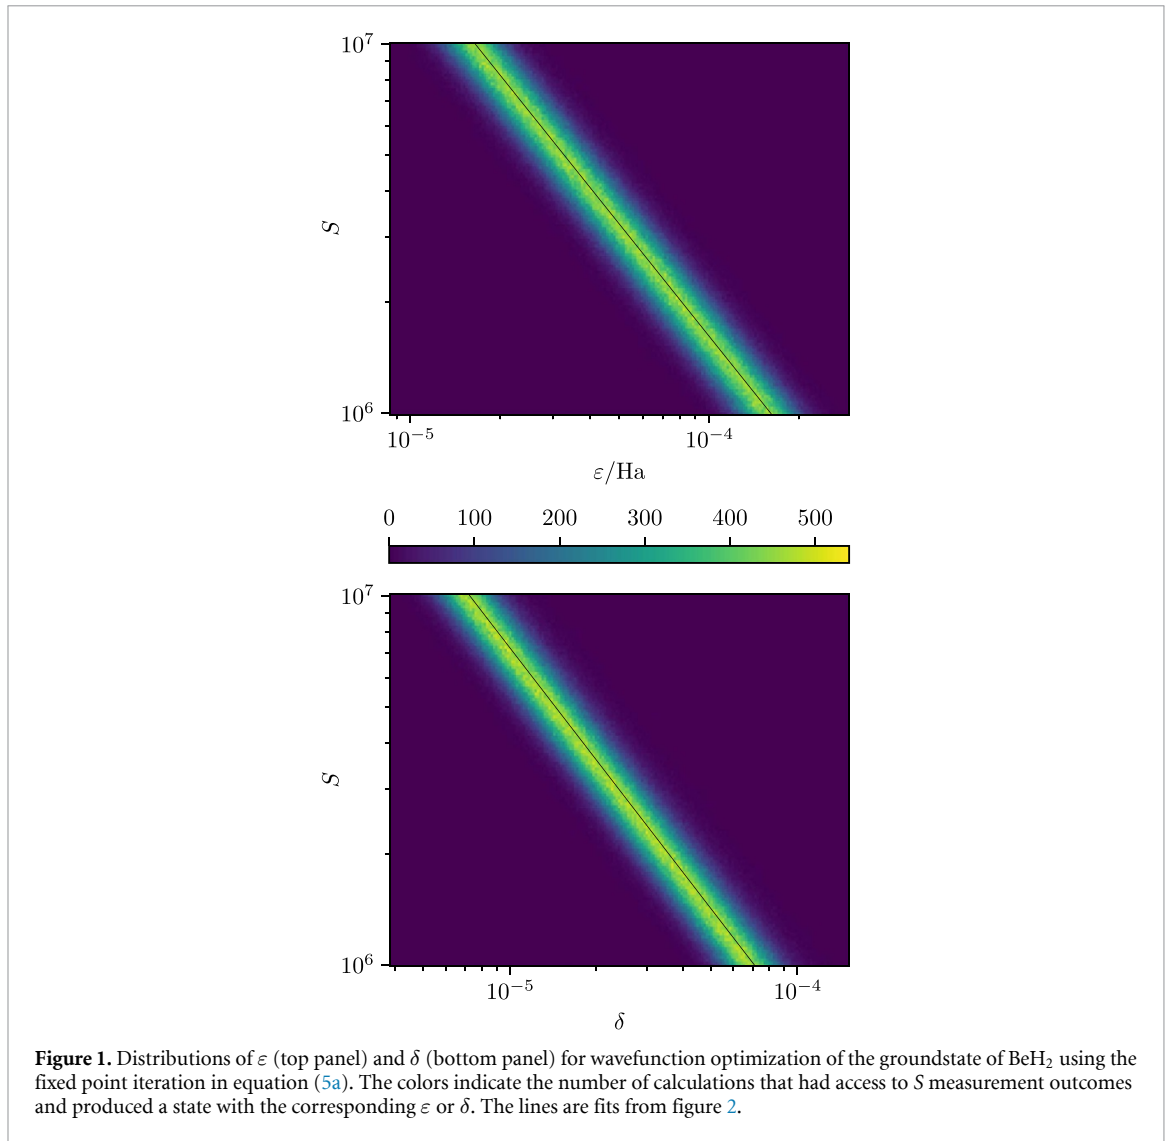
$$\tilde{T}(|\varphi\rangle) = \frac{|\varphi\rangle + T(|\varphi\rangle)}{\| |\varphi\rangle + T(|\varphi\rangle) \|} \quad (5b)$$

includes a contribution from the previous state to smooth out oscillations near convergence. Although we have no proven guarantees that this scheme will produce desirable states, we find empirically that nearly all optimized wavefunctions are found in minima of the loss function.

Unlike the complex RBM and RNN, which have limited expressivity, these wavefunctions can represent all states in the Hilbert space. The optimized wavefunctions are determined solely by the loss function and the initial state, without any inductive bias due to the choice of model. Thus, the process of wavefunction optimization through fixed point iteration yields a reference state with which to compare the other data-driven reconstructions using the generative models.

### 3. Results and discussion

We are interested in reconstructing observables for groundstates of three molecular Hamiltonians  $\hat{H}$  for LiH, BeH<sub>2</sub>, and H<sub>2</sub>. In order to implement these in a form suitable for a quantum computer, each fermionic Hamiltonian is decomposed into a qubit Hamiltonian as a sum of Pauli strings. This is obtained via the parity transformation implemented in Qiskit [19]; see appendix B for more details.




Each of the qubit Hamiltonians is sufficiently small that its groundstate  $|\psi\rangle$  can be found using exact diagonalization. This allows us to calculate the exact energies (or any generic observable), as well as to produce a synthetic measurement dataset of arbitrary size. For each synthetic measurement outcome from a given molecular groundstate, an orthonormal basis  $k$  is chosen uniformly at random from the  $K$  bases determined by the Pauli string expansion of the corresponding Hamiltonian, and an outcome label  $\sigma$  is selected with the Born rule probability

$$p_k(\sigma) = \left| \langle \sigma | \hat{R}_k^\dagger | \psi \rangle \right|^2. \tag{6}$$

The resulting datasets are used to train the RBM and RNN generative models. By way of comparison, we use the same datasets to implement maximum likelihood tomography on all  $2 \times 2^N$  real-valued parameters of the full wavefunction via fixed point iteration. We also compare to uniform classical shadows by implementing algorithm 1 of [8].

For a given groundstate, the sample complexity  $S$  is defined as the number of measurements (i.e. the number of copies of the state) required to achieve a target accuracy of the reconstructed state or observable. This formulation implies that  $S$  is a function of the accuracy parameters  $\varepsilon$  and  $\delta$ . However, in practice, one must fix  $S$  and obtain a state with some energy and fidelity, which will be random variables that depend on both the observed measurement outcomes and the details of the reconstruction method. Figure 1 reveals that the distributions of

$$\varepsilon = \langle \varphi(\boldsymbol{\lambda}) | \hat{H} | \varphi(\boldsymbol{\lambda}) \rangle - \langle \psi | \hat{H} | \psi \rangle \tag{7a}$$



and

$$\delta = 1 - |\langle \varphi(\boldsymbol{\lambda}) | \psi \rangle|^2 \quad (7b)$$

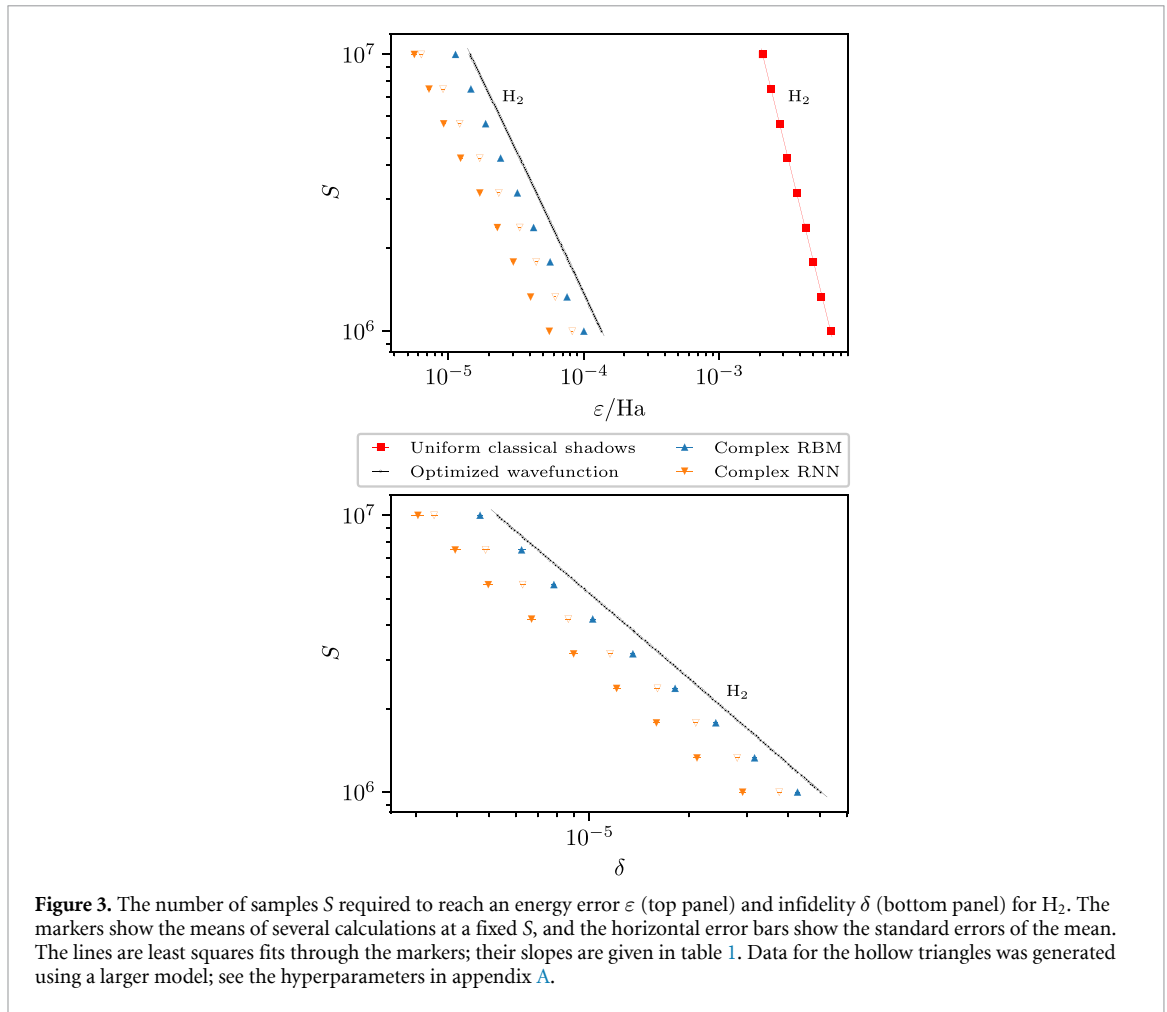
for maximum likelihood tomography of the BeH<sub>2</sub> groundstate vary smoothly with  $S$ . Therefore, we examine averages over multiple independent datasets with fixed  $S$  in order to draw conclusions about the scaling of  $S$ .

Such averages are shown for all the methods in figures 2 and 3. For uniform classical shadows, the estimated energy is used directly to compute the energy error:

$$\varepsilon = \left| \bar{E} - \langle \psi | \hat{H} | \psi \rangle \right|. \quad (8)$$

We chose to work in the regime of  $10^6$  to  $10^7$  measurements, as this is currently tractable on some experimental hardware. For example, on the order of  $10^7$  measurement outcomes can be obtained per hour using a superconducting quantum computer [20] (although for other hardware, e.g. neutral atom arrays, a realistic number may be several orders of magnitude less [21]). With this many measurement outcomes, we found that linear fits are appropriate for model-based reconstructions ( $R^2 > 0.998$ ), optimized wavefunctions ( $R^2 > 0.999$ ), and uniform classical shadows ( $R^2 > 0.999$ ); their slopes are summarized in table 1. Since the data is log-transformed for plotting, the slopes of the fits are the exponents of the power law expressions relating the sample complexity  $S$  to the estimates of reconstruction quality ( $\varepsilon$  and  $\delta$ ).

Assuming that results in this region of  $S$  are in the asymptotic limit, we may compare the scaling between the methods mentioned above. As listed in table 1, uniform classical shadows have a sample complexity of approximately  $S \propto \varepsilon^{-2}$ , as expected from the method's structure as a direct average over random samples. On the other hand, the model-based approaches get very near to scaling as  $S \propto \varepsilon^{-1}$ . This implies that—in the small error limit—significantly fewer measurements are needed to achieve the same quality of approximation. In the present case of small  $N$ , the fidelity  $1 - \delta$  between the model and target wavefunctions can be computed directly; as shown in table 1, the model-based approaches are close to  $S \propto \delta^{-1}$  using only single-copy measurements.



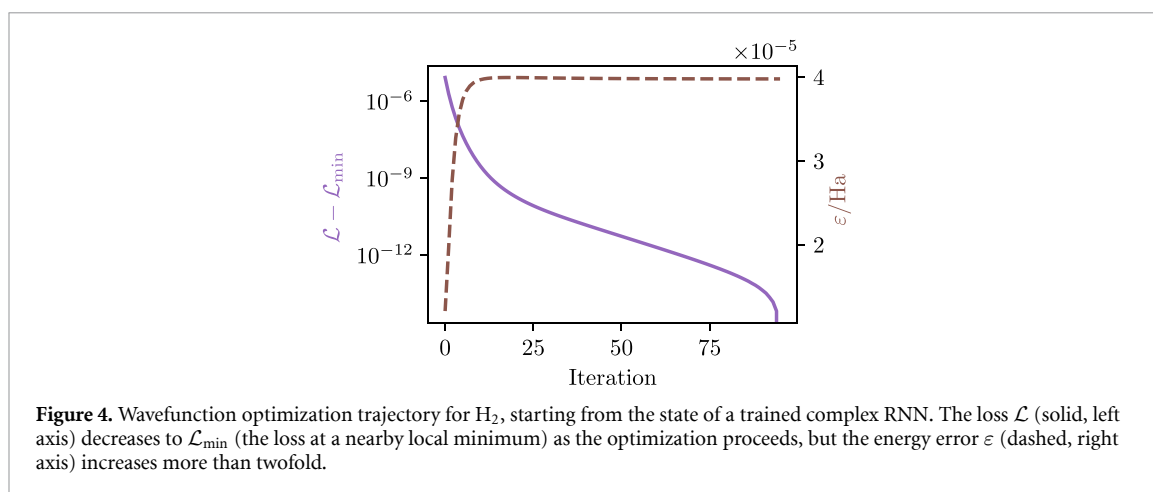
**Figure 3.** The number of samples  $S$  required to reach an energy error  $\varepsilon$  (top panel) and infidelity  $\delta$  (bottom panel) for  $\text{H}_2$ . The markers show the means of several calculations at a fixed  $S$ , and the horizontal error bars show the standard errors of the mean. The lines are least squares fits through the markers; their slopes are given in table 1. Data for the hollow triangles was generated using a larger model; see the hyperparameters in appendix A.

**Table 1.** Slopes of the linear fits in figures 2 and 3 (for visual clarity, fits for the complex RBM and RNN are not drawn).

	LiH		BeH <sub>2</sub>		H <sub>2</sub>	
	$\varepsilon$	$\delta$	$\varepsilon$	$\delta$	$\varepsilon$	$\delta$
Shadows	-2.01	—	-2.00	—	-1.99	—
Wavefunction	-1.09	-1.13	-1.01	-1.01	-1.03	-1.02
Complex RBM	-1.17	-1.18	-1.00	-1.01	-1.05	-1.05
Complex RNN	-1.12	-1.16	-0.97	-1.00	-0.99	-1.01

One concern with models that are optimized to reproduce a finite collection of measurement outcomes is the possibility of overfitting to the training data. Indeed, we consistently observe this with the optimized wavefunctions, the vast majority (>99.998%) of which end up at a minimum of the loss function after several rounds of fixed point iteration. Although it is not always possible to find a state in the Hilbert space that perfectly reproduces the observations, being in a good loss minimum suggests some level of memorization of the training data, an intuition that has been around since the inception of generative models [22]. In contrast, the states produced by the trained complex RBMs and RNNs never completely minimize the loss; states with a smaller loss always exist, but either the models cannot express them or the training procedure cannot find them. To highlight the difference in the final states achieved by these two approaches, we show in figure 4 an optimization trajectory that begins from the state of a trained complex RNN, but is further optimized by repeated application of the nonlinear operator  $\tilde{T}$  from equation (5b). The state that is attained by the end of the fixed point iteration has a loss that is lower (better memorization) at the cost of a larger energy error (worse generalization, as described in appendix C).

The effect of this trade-off can be observed in figure 3 for  $\text{H}_2$ , where the models invariably generalize better than the optimized wavefunctions. Using larger models (which are expected to be more expressive and possibly more trainable) can also negatively impact the ability to generalize, as shown by the hollow triangles



in figure 3. It is likely that a description of the actual ground state of  $H_2$  is much simpler (in the sense of Kolmogorov complexity) than a description of any particular realization of measurement data, so a more constrained model might stumble upon the former while searching for the latter. Thus, the limited flexibility of the models can be a boon for groundstate energy estimation in practice. Additionally, if one is willing to consider a variational Monte Carlo approach, overfitting can be avoided altogether by switching to an energy-based loss after the initial portion of the training [23, 24].

#### 4. Conclusions and outlook

In this paper, we explore the use of generative models adapted from modern machine learning for the efficient reconstruction of observables from prototypical qubit Hamiltonians in quantum chemistry: LiH,  $BeH_2$ , and  $H_2$ . The neural network ansatzes we use are well-capable of representing groundstate wavefunctions, can be made systematically more expressive (by increasing e.g. the number of hidden units), and have well-studied and efficient training heuristics. Using synthetic datasets of typical noiseless experimental measurements on the groundstate wavefunction, we calculate the sample complexity  $S$ : the number of independent measurements required to reproduce the energy or fidelity up to a given accuracy.

We find that the energy reconstruction obtained by training a generative model results in an energy error  $\varepsilon$  that is orders of magnitude smaller than that achieved by classical shadow tomography using the same amount of measurement data. More generally, these findings suggest that  $S$  for the model-based reconstruction should scale asymptotically as  $\varepsilon^{-1}$ , as opposed to  $\varepsilon^{-2}$  for classical shadows.

Of course, in practice, this type of improvement would depend on the model training heuristics scaling favorably. While very few results relevant to this scaling exist in the machine learning literature, efficient scaling of the training procedure has been observed previously in reconstructions of many-qubit groundstates relevant for condensed matter physics [13]. It remains to be seen whether the asymptotic advantage for the sample complexity in the model-based case persists for larger molecules.

Our results illustrate that some applications of present-day quantum computers can benefit from the use of generative models adapted from machine learning. For the small molecules studied here, we can conclude that a model-based approach is better able to make use of hardware with a data rate limited by state preparation and measurement than alternatives such as classical shadow tomography. More generally, hybrid quantum–classical algorithms such as VQE can benefit from modeling quantum states using neural networks. These approaches require multiple states to be prepared on a quantum computer, with very many measurements performed over the course of the algorithm [25, 26]. Insertion of neural network models into the hybrid workflow has the potential to drastically reduce the amount of quantum computation involved, enabling the study of systems which would otherwise be prohibitively expensive.

#### Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.



## Acknowledgment

The authors would like to thank Yi Hong Teoh, Marta Mauri, Juan Carrasquilla, Alexander A Kunitsa, Peter D Johnson, and Jhonathan Romero for insightful discussions. This work was supported by Mitacs through the Mitacs Elevate program. RGM acknowledges support from the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada Research Chair (CRC) program, the New Frontiers in Research Fund, and the Perimeter Institute for Theoretical Physics. Research at the 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 Economic Development, Job Creation and Trade. The numerical calculations in this work were possible thanks to the computational resources provided by Compute Canada.

## Appendix A. Models and training

The complex RBMs and RNNs were trained using the Adam optimizer [27] in the Flux machine learning library for Julia [28, 29], with the hyperparameters given in table 2. The hyperparameters were chosen to provide a reasonable level of convergence of the energy simultaneously for small and large  $S$ . Since the  $S$  measurements are allocated to the  $K$  bases uniformly at random, approximately  $S/K$  measurements are performed in each basis.

To simplify the training procedure and reduce the number of hyperparameters, no minibatching was performed during training. For both optimization methods (Adam and fixed point iteration), all  $S$  measurement outcomes were used for each update of the model parameters. For larger models, it may be necessary to split the training data into random minibatches, which will have the effect of stochastically perturbing the loss landscape at each optimization step. This ‘blurring’ of the loss minima should act as an implicit regularizer and improve the generalization of the trained models [30–32].

**Table 2.** Model and training hyperparameters. The number of real-valued parameters is  $N_P = 2(NN_h + N + N_h)$  for complex RBMs and  $N_P = N_h^2 + 5N_h + 3$  for complex RNNs. Hyperparameters for the large  $H_2$  RNN ( $N_h = 64$ , hollow triangles in figure 3) are also included.

Molecule	Model	$N_h$	$N_P$	Learning rate	Epochs
LiH	RBM	3	38	$3 \times 10^{-3}$	$1.6 \times 10^5$
LiH	RNN	4	39	$1 \times 10^{-3}$	$1.0 \times 10^5$
BeH <sub>2</sub>	RBM	9	138	$3 \times 10^{-4}$	$1.0 \times 10^5$
BeH <sub>2</sub>	RNN	15	303	$1 \times 10^{-4}$	$4.5 \times 10^5$
H <sub>2</sub>	RBM	5	106	$3 \times 10^{-3}$	$1.2 \times 10^5$
H <sub>2</sub>	RNN	9	129	$3 \times 10^{-4}$	$1.2 \times 10^5$
H <sub>2</sub>	RNN	64	4419	$1 \times 10^{-5}$	$1.2 \times 10^5$

**Table 3.** Details of the molecular Hamiltonians. The transformation parameters were chosen to produce the same system sizes  $N$  as in [14].

	LiH	BeH <sub>2</sub>	H <sub>2</sub>
Bond length	1.55 Å	1.32 Å	0.75 Å
Basis set	STO-3G	STO-3G	6-31G
Spin orbitals	12	14	8
freeze_core	True	True	False
orbital_reduction	[-3, -2]	[-3]	None
two_qubit_reduction	True	True	False
z2symmetry_reduction	None	[+1, +1]	None
Qubits ( $N$ )	4	6	8
Pauli terms	100	216	185
Measurement bases ( $K$ )	25	83	54

## Appendix B. Molecular Hamiltonians

We studied three molecules near their equilibrium bond lengths: LiH, BeH<sub>2</sub>, and H<sub>2</sub>. In the case of BeH<sub>2</sub>, we only considered linear geometries with equal bond lengths. The electronic Hamiltonians were transformed into qubit form using Qiskit's `FermionicTransformation` class with `Psi4` [33] as the driver and the mapping set to `FermionicQubitMappingType.PARITY`. Energies are reported in hartrees (Ha). Per-molecule information is provided in table 3.

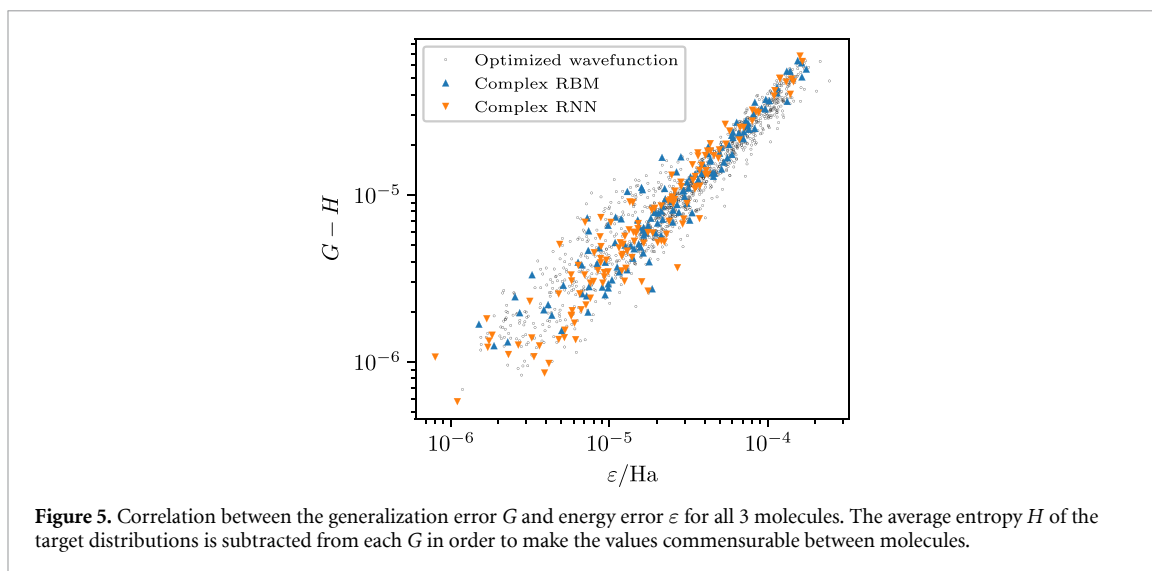
## Appendix C. Energy error as a proxy for generalization

One might intuitively expect that a smaller additive energy error  $\varepsilon$  indicates that the generative model is better able to generalize to unseen configurations. For example, if the model's energy is computed by Monte Carlo, evaluating the estimator on samples that are more representative of the true distribution should provide a result that is closer to the true energy.

The complex RBM and RNN are trained to reproduce a particular finite collection of data using the empirical risk (the cross entropy relative to the histogram of the data). Hence, a suitable measure of generalization error is the true risk (the cross entropy relative to the target distribution) [34],

$$G(\lambda) = -\frac{1}{K} \sum_{k=1}^K \sum_{\sigma} p_k(\sigma) \log \left| \langle \sigma | \hat{R}_k^\dagger | \varphi(\lambda) \rangle \right|^2, \quad (C1)$$

which is obtained in the limit of infinite data. As can be seen in figure 5, for the wavefunctions obtained in this work, there is a strong correlation between the generalization error  $G$  and the energy error  $\varepsilon$ , suggesting that the latter is representative of the former.



## ORCID iDs

Dmitri Iouchtchenko  <https://orcid.org/0000-0001-9308-3682>

Jérôme F Gonthier  <https://orcid.org/0000-0002-2933-4085>

Alejandro Perdomo-Ortiz  <https://orcid.org/0000-0001-7176-4719>

Roger G Melko  <https://orcid.org/0000-0002-5505-8176>

## References

- [1] Kitaev A Y 1995 Quantum measurements and the Abelian stabilizer problem (arXiv:quant-ph/9511026)
- [2] Abrams D S and Lloyd S 1999 Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors *Phys. Rev. Lett.* **83** 5162
- [3] Kandala A, Mezzacapo A, Temme K, Takita M, Brink M, Chow J M and Gambetta J M 2017 Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets *Nature* **549** 242
- [4] Tranter A, Love P J, Mintert F and Coveney P V 2018 A comparison of the Bravyi–Kitaev and Jordan–Wigner transformations for the quantum simulation of quantum chemistry *J. Chem. Theory Comput.* **14** 5617
- [5] McClean J R, Romero J, Babbush R and Aspuru-Guzik A 2016 The theory of variational hybrid quantum-classical algorithms *New J. Phys.* **18** 023023
- [6] Carrasquilla J, Torlai G, Melko R G and Aolita L 2019 Reconstructing quantum states with generative models *Nat. Mach. Intell.* **1** 155
- [7] Huang H-Y, Kueng R and Preskill J 2020 Predicting many properties of a quantum system from very few measurements *Nat. Phys.* **16** 1050
- [8] Hadfield C, Bravyi S, Raymond R and Mezzacapo A 2020 Measurements of quantum Hamiltonians with locally-biased classical shadows (arXiv:2006.15788)
- [9] Stricker R, Meth M, Postler L, Edmunds C, Ferrie C, Blatt R, Schindler P, Monz T, Kueng R and Ringbauer M 2022 Experimental single-setting quantum state tomography *PRX Quantum* **3** 040310
- [10] Torlai G, Mazzola G, Carrasquilla J, Troyer M, Melko R and Carleo G 2018 Neural-network quantum state tomography *Nat. Phys.* **14** 447
- [11] Melko R G, Carleo G, Carrasquilla J and Cirac J I 2019 Restricted Boltzmann machines in quantum physics *Nat. Phys.* **15** 887
- [12] Hibat-Allah M, Ganahl M, Hayward L E, Melko R G and Carrasquilla J 2020 Recurrent neural network wave functions *Phys. Rev. Res.* **2** 023358
- [13] Schayek D, Golubeva A, Albergo M S, Kulchytskyy B, Torlai G and Melko R G 2019 Learnability scaling of quantum states: restricted Boltzmann machines *Phys. Rev. B* **100** 195125
- [14] Torlai G, Mazzola G, Carleo G and Mezzacapo A 2020 Precise measurement of quantum observables with neural-network estimators *Phys. Rev. Res.* **2** 022060(R)
- [15] Torlai G and Melko R G 2016 Learning thermodynamics with Boltzmann machines *Phys. Rev. B* **94** 165134
- [16] Carleo G and Troyer M 2017 Solving the quantum many-body problem with artificial neural networks *Science* **355** 602
- [17] Xiao T, Fan J and Zeng G 2022 Parameter estimation in quantum sensing based on deep reinforcement learning *npj Quantum Inf.* **8** 2
- [18] Goyeneche D and de la Torre A C 2014 Quantum tomography meets dynamical systems and bifurcations theory *J. Math. Phys.* **55** 062103
- [19] Anis M S *et al* 2021 Qiskit: an open-source framework for quantum computing (<https://doi.org/10.5281/zenodo.2573505>)
- [20] Arute F *et al* 2019 Quantum supremacy using a programmable superconducting processor *Nature* **574** 505
- [21] Endres M, Bernien H, Keesling A, Levine H, Anschuetz E R, Krajenbrink A, Senko C, Vuletic V, Greiner M and Lukin M D 2016 Atom-by-atom assembly of defect-free one-dimensional cold atom arrays *Science* **354** 1024
- [22] Hopfield J J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci.* **79** 2554

- [23] Bennewitz E R, Hopfmueller F, Kulchytskyy B, Carrasquilla J and Ronagh P 2021 Neural error mitigation of near-term quantum simulations (arXiv:[2105.08086](#))
- [24] Czischek S, Moss M S, Radzihovsky M, Merali E and Melko R G 2022 Data-enhanced variational Monte Carlo for Rydberg atom arrays *Phys. Rev. B* **105** 205108
- [25] Wecker D, Hastings M B and Troyer M 2015 Progress towards practical quantum variational algorithms *Phys. Rev. A* **92** 042303
- [26] Gonthier J F, Radin M D, Buda C, Doskocil E J, Abuan C M and Romero J 2020 Identifying challenges towards practical quantum advantage through resource estimation: the measurement roadblock in the variational quantum eigensolver (arXiv:[2012.04001](#))
- [27] Kingma D P and Ba J 2017 Adam: a method for stochastic optimization (arXiv:[1412.6980](#))
- [28] Innes M 2018 Flux: Elegant machine learning with Julia *J. Open Source Softw.* **3** 602
- [29] Innes M, Saba E, Fischer K, Gandhi D, Rudilosso M C, Joy N M, Karmali T, Pal A and Shah V 2018 Fashionable modelling with Flux (arXiv:[1811.01457](#))
- [30] Keskar N S, Mudigere D, Nocedal J, Smelyanskiy M and Tang P T P 2017 On large-batch training for deep learning: generalization gap and sharp minima *ICLR 2017 Conf. Track (OpenReview.net, Palais des Congrès Neptune, Toulon, France)*
- [31] Smith S, Elsen E and De S 2020 On the generalization benefit of noise in stochastic gradient descent *Proc. 37th Int. Conf. on Machine Learning (PMLR)* pp 9058–67
- [32] Smith S L, Dherin B, Barrett D G T and De S 2021 On the origin of implicit regularization in stochastic gradient descent *Int. Conf. on Learning Representations (ICLR 2021)* OpenReview.net
- [33] Smith D G A *et al* 2020 Psi4 1.4: Open-source software for high-throughput quantum chemistry *J. Chem. Phys.* **152** 184108
- [34] Hansen L and Larsen J 1996 Unsupervised learning and generalization *Proc. IEEE Int. Conf. on Neural Networks (ICNN'96)* vol 1 (Washington, DC: IEEE) pp 25–30