

Learning ground state observables from quantum computing experiments (Extended abstract submission for QTML 2026)

Ben Jaderberg¹, Freya Shah², Minjun Jeon^{1,3}, M. Emre Sahin⁴, Christa Zoufal¹, Kunal Sharma¹

¹IBM Quantum ²Department of Engineering Science, University of Oxford

³Department of Materials, University of Oxford ⁴The Hartree Centre, STFC

(Dated: June 29, 2026)

INTRODUCTION

Predicting properties of quantum many-body ground states is central to condensed-matter physics [1–3], quantum chemistry [4], and quantum information science [5]. Repeatedly evaluating ground-state properties across parameter spaces can incur substantial computational cost, creating a major bottleneck for many-body simulation. This has motivated growing interest in ML approaches that seek to amortize this cost by training models to approximate the map from Hamiltonian parameters to ground-state observables using previously generated data [6–11].

For families of gapped local Hamiltonians, classical learners have been shown to efficiently predict local ground-state properties across parameter ranges that remain within the same quantum phase [12]. Subsequent results showed that incorporating geometric locality as an inductive bias can substantially reduce the required amount of training data and improve scaling with system size [13], while more recent analyses extend these guarantees to broader classes of models, including certain deep learning architectures [14]. However, numerical experiments supporting these theoretical developments have largely used density matrix renormalization group (DMRG), even though it is not computationally efficient for 2D systems.

Meanwhile, experimental progress towards learning from quantum data has so far been limited by the difficulty of preparing accurate many-body states and estimating their observables on noisy hardware. In this work, we address this challenge for ground states of the Heisenberg XXZ model on two-dimensional heavy-hex lattices with up to 115 qubits. Figure 1 summarizes the end-to-end workflow, from quantum data generation to model training and prediction.




FIG. 1: Overview of the workflow. (Left) For a heavy-hex lattice and anisotropy J_z of the spin- $\frac{1}{2}$ Heisenberg XXZ model, we generate approximate ground states by running sample-based Krylov diagonalization (SKQD) on the *ibm_boston* quantum processor. This is followed by a basis-optimization step that learns a shallow unitary transformation of the Hamiltonian which further lowers the energy when diagonalized. We generate data by computing local observables from the resulting state, which are used to train a neural network via supervised learning. (Right) Once trained, the model predicts these observables at previously unseen values of J_z across the parameter space.

RESULTS

Verifying data quality

We study the spin- $\frac{1}{2}$ Heisenberg XXZ Hamiltonian $H(J_z) = \frac{1}{4} \sum_{\langle i,j \rangle \in \mathcal{G}} (X_i X_j + Y_i Y_j) + J_z Z_i Z_j$, where $P_i \in \{X, Y, Z\}$ are Pauli matrices, $\langle i, j \rangle$ denotes nearest neighbors on a two-dimensional heavy-hexagonal [15] graph $\mathcal{G}$ and J_z tunes the anisotropy of the model. In this work, we consider graphs of distance 5 and 7,

corresponding to $n = 57$ and $n = 115$ qubits, respectively. We generate approximate ground states of the Hamiltonian in the AFM phase using our ground-state workflow that begins with SKQD [16, 17]. Running for over 8 hours on the *ibm_boston* superconducting quantum computer, and using high-performance computing (HPC) resources for diagonalization and basis optimization, we construct ground states at 45 points in the AFM phase.

We compute expectation values of selected observables, which can be evaluated efficiently from the approximate ground state: the single-site magnetization $\langle Z \rangle$, the two-point correlator $C_{zz} = \langle Z_i Z_j \rangle$ over first- and second-nearest-neighbor edges, and the loop correlator $Z_{\text{loop}} = \langle \prod_{k \in \ell_{12}} Z_k \rangle$, where ℓ_{12} denotes a fixed loop of twelve sites on the lattice. As a reference benchmark, we compare against DMRG calculations with bond dimension $\chi = 320$, performed using the TeNPy library.

Figure 2 summarizes the resulting agreement with DMRG. In the left panel, we plot the difference between the ground state energy produced by our algorithm (E) and DMRG (E_0), expressed in units of the many-body gap $E_1 - E_0$. In the top panel of Fig. 2, we quantify the accuracy of C_{zz} and Z_{loop} . Across the phase space, the median relative errors for $n = 57$ and $n = 115$ are $M(\Delta C_{zz}) = 1.36\%$ and 1.22% , $M(\Delta Z_{\text{loop}}) = 1.57\%$ and 1.27% , respectively.

Finally, the bottom panel of Fig. 2 shows that the loop correlator data preserve the correct spatial variations across the lattice. At $J_z = 4.0$ for $n = 115$, the DMRG loop pattern exhibits two boundary-adjacent loops with larger values than those in the bulk. The approximate ground state produced by our workflow reproduces this pattern despite the high precision required to resolve it. This shows that the constructed dataset preserves local structure rather than only matching spatial averages, motivating its use as training targets for learning ground-state properties.




FIG. 2: Accuracy of data generated from quantum experiments. (Left) Error in the approximate ground state energy E relative to the gap to the first excited state $(|E - E_0|)/(E_1 - E_0)$ as calculated by DMRG. (Top) Mean percentage error between the quantum data and DMRG for C_{zz} and Z_{loop} . (Bottom) Spatial map of $\langle Z_{\text{loop}} \rangle$ for $n = 115$ at $J_z = 4.0$.

In-distribution generalization

We next show that neural network models trained on this data can generalize to accurately predict local observables at unseen points within the interior of the AFM phase, providing a practical realization of the in-phase learning setting analyzed theoretically in Refs. [12–14]. We use the architecture introduced in Ref. [14] which combines the outputs of several local multi-layer perceptrons to predict a single observable $\langle \hat{O}_1 \rangle$. To evaluate generalization within the phase, we use training data from the parameter range $1.6 \leq J_z \leq 6.0$ and partition this interval into three strata: 4.1–6.0, 2.1–4.0, and 1.6–2.0. We construct a held-out test set by removing 20% of points from each stratum. We train three separate models to predict expectation values across the lattice for the observable families Z , C_{zz} and Z_{loop} , respectively.

For $n = 57$, the models achieve $\leq 2\%$ average error for all three observable types, where here we take the mean of the relative errors across the lattice (i.e., across 57 values for Z). For $n = 115$, we observe a similar trend, but with a lower average error, remaining below 1% across all observable families and points in the phase space. Overall, we find that models trained on a limited set of quantum data points can predict spatially resolved local observables at unseen J_z with percent-level accuracy on both system sizes.

Out-of-distribution generalization

We next evaluate the ability of the models to generalize outside of the training interval $1.6 \leq J_z \leq 6.0$ and towards the phase boundary at $J_z = 1.0$. In general, computing ground states near a phase boundary is difficult for both classical and quantum methods. For the XXZ model, the transition at $J_z = 1.0$ from the gapped AFM phase to the gapless critical XY phase is marked by the closing of the spectral gap. The gap decreases rapidly as J_z approaches the phase boundary, with the most pronounced change occurring below the training cutoff at $J_z = 1.6$. This identifies the region $1.1 \leq J_z \leq 1.4$ as one in which the ground-state physics changes over a short parameter interval, making it an important testbed for the generalization capabilities of the trained models. To do this, we run an additional set of experiments for $1.1 \leq J_z \leq 1.4$ to generate test data.

Subsequently, we consider the relative error between the ML prediction and quantum experiments for the staggered magnetization m_s . Despite the absence of training data below $J_z = 1.6$, the models achieve $< 3\%$ error in this region across both system sizes. Similarly, for the two-point correlators C_{zz} and the loop correlators, the spatially averaged error remains below 4% across the extrapolation region. This shows that the models remain accurate when applied toward the phase boundary, where the closing spectral gap makes the ground-state observables vary more rapidly. Finally, the right panel of Fig. 3 shows that the models capture spatial variations of the observables at $J_z = 1.1$, the closest point to the phase boundary considered here. Taken together, Fig. 3 shows that the models generalize beyond the training distribution not only in terms of averaged observables, but also in their specific spatial structure.




FIG. 3: Model generalization towards the phase boundary. (Left) Spectral gap $E_1 - E_0$ as a function of J_z on approaching the phase boundary at $J_z = 1.0$. (Middle) Relative error between local observables predicted by ML and data from new quantum experiments outside the training regime. (Right) Representative slice through the lattice at $J_z = 1.1$, comparing the spatially resolved observables from ML and the quantum experiments.

DISCUSSION

In this work, we demonstrate an end-to-end workflow for learning ground-state observables from experimental quantum data in a large, two-dimensional interacting many-body system. Combining SKQD with an optimized entangled-basis transformation of the projected Hamiltonian, we obtain approximate ground states of the Heisenberg XXZ model on heavy-hex lattices with up to 115 qubits. Once trained on this dataset, the neural networks accurately predict ground-state observables at unseen Hamiltonian parameters.

Looking ahead, a broader motivation for this workflow is the use of quantum processors to generate training data for states inaccessible to classical methods. Although gapped two-dimensional local Hamiltonians generally satisfy an entanglement area law [1, 18], this scaling does not guarantee an efficient PEPS representation [2, 19]. Given the known hardness of contracting PEPS [20], quantum computing provides a scalable alternative for generating ground-state data. As hardware improves, these workflows will enable classical models to learn from quantum data when classical approximations become prohibitively expensive.

[1] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, Reviews of modern physics **82**, 277 (2010).

- [2] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions, arXiv preprint cond-mat/0407066 (2004).
- [3] L. Amico, R. Fazio, A. Osterloh, and V. Vedral, Entanglement in many-body systems, *Reviews of modern physics* **80**, 517 (2008).
- [4] W. Kohn and L. J. Sham, Self-consistent equations including exchange and correlation effects, *Physical review* **140**, A1133 (1965).
- [5] J. I. Latorre, E. Rico, and G. Vidal, Ground state entanglement in quantum spin chains, arXiv preprint quant-ph/0304098 (2003).
- [6] J. Behler and M. Parrinello, Generalized neural-network representation of high-dimensional potential-energy surfaces, *Physical review letters* **98**, 146401 (2007).
- [7] R. Ramakrishnan, P. O. Dral, M. Rupp, and O. A. Von Lilienfeld, Big data meets quantum chemistry approximations: the δ -machine learning approach, *Journal of chemical theory and computation* **11**, 2087 (2015).
- [8] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nature Physics* **13**, 431 (2017).
- [9] F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, Bypassing the kohn-sham equations with machine learning, *Nature communications* **8**, 872 (2017).
- [10] J. Schmidt, J. Shi, P. Borlido, L. Chen, S. Botti, and M. A. Marques, Predicting the thermodynamic stability of solids combining density functional theory and machine learning, *Chemistry of Materials* **29**, 5090 (2017).
- [11] K. Ch'Ng, J. Carrasquilla, R. G. Melko, and E. Khatami, Machine learning phases of strongly correlated fermions, *Physical Review X* **7**, 031038 (2017).
- [12] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, *Science* **377**, eabk3333 (2022).
- [13] L. Lewis, H.-Y. Huang, V. T. Tran, S. Lehner, R. Kueng, and J. Preskill, Improved machine learning algorithm for predicting ground state properties, *nature communications* **15**, 895 (2024).
- [14] M. Wanner, L. Lewis, C. Bhattacharyya, D. Dubhashi, and A. Gheorghiu, Predicting ground state properties: Constant sample complexity and deep learning algorithms, *Advances in Neural Information Processing Systems* **37**, 33962 (2024).
- [15] C. Chamberland, G. Zhu, T. J. Yoder, J. B. Hertzberg, and A. W. Cross, Topological and subsystem codes on low-degree graphs with flag qubits, *Physical Review X* **10**, 011022 (2020).
- [16] J. Yu, J. R. Moreno, J. T. Iosue, L. Bertels, D. Claudino, B. Fuller, P. Groszkowski, T. S. Humble, P. Jurcevic, W. Kirby, *et al.*, Quantum-centric algorithm for sample-based krylov diagonalization, arXiv preprint arXiv:2501.09702 (2025).
- [17] W. Kirby, B. Pokharel, J. R. Moreno, K. C. Smith, S. Bravyi, A. Deshpande, C. Evangelinos, B. Fuller, J. R. Garrison, B. Jaderberg, *et al.*, Observation of improved accuracy over classical sparse ground-state solvers using a quantum computer, arXiv preprint arXiv:2603.03496 (2026).
- [18] A. Anshu, I. Arad, and D. Gosset, An area law for 2d frustration-free spin systems, in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing* (2022) pp. 12–18.
- [19] Y. Ge and J. Eisert, Area laws and efficient descriptions of quantum many-body states, *New Journal of Physics* **18**, 083026 (2016).
- [20] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Computational complexity of projected entangled pair states, *Physical review letters* **98**, 140506 (2007).