

Active learning for data-efficient modelling of quantum phase spaces

Simone Cantori¹, Gabriel Matos¹, Enrico Rinaldi¹, and Jacob Swain¹

¹Quantinuum, Partnership House, Carlisle Place, London SW1P 1BX, United Kingdom

Abstract

Neural networks provide a powerful framework for learning quantum states directly from measurement data across extended regions of parameter space. However, acquiring the data required to train these models can represent a significant cost. In this work, we introduce an active learning strategy for improving data acquisition for training conditional generative models. Based on a query-by-committee protocol, the strategy iteratively acquires data in phase space regions where an ensemble of models exhibits larger disagreement in sample distribution. The resulting trained models can identify phase transitions through susceptibility-based analyses of the learned distributions, without requiring prior knowledge of the relevant order parameters. We show that to achieve a comparable performance with uniform sampling, active learning requires approximately $3\times$ less data, representing a considerable saving in compute cost and acquisition time.

1 Introduction

Machine learning is an increasingly important tool for the study of many-body quantum systems [1]. The success of foundation models in natural language processing has motivated the development of analogous architectures for quantum many-body physics [2, 3, 4]. Such models provide a unified representation of entire families of quantum systems, though training these models often requires large datasets spanning broad regions of parameter space. Obtaining such data can be computationally expensive, especially in experimental settings, where measurements are limited by hardware resources and acquisition time.

Active learning provides a natural framework for mitigating this issue [5], by identifying data points for which acquisition, labelling or inclusion in the training set is expected to improve training performance. Prior work in active learning of quantum systems [6, 7] has approached the related task of phase diagram discovery: efficiently identifying phase boundaries. Data acquisition in this setting consists of running experiments to identify the categorical phase for a given set of input parameters. We tackle the more general problem of learning the full conditional distribution of measurement outcomes from acquired measurement samples, without the need for a priori definition of phases. This supports a broader range of analyses, including predicting arbitrary observables and characterising their behaviour along boundary-crossing trajectories.

We benchmark generative neural networks trained with active learning on a diverse set of quantum systems. Across these examples, we show that active learning consistently improves the accuracy of the learned distributions and derived observables compared to uniform data acquisition. Furthermore, we demonstrate the ability to identify quantum phase transitions through a susceptibility-based analysis of the learned probability distributions.

2 Methods

In this work we consider *families* of bitstring distributions, conditioned by a small number of parameters, $\mathbf{C} \in \mathbb{R}^D$. In particular, we look at the ground state or Trotterized dynamics of various spin Hamiltonians. The conditioning parameters comprise coefficients of Hamiltonians defining their phase space, and in the case of Trotterized dynamics, the normalised number of Trotter steps.

These parameters define a conditional probability distribution, $p(\mathbf{x}|\mathbf{C})$, namely the probability of sampling bitstring $\mathbf{x}$ from the distribution parametrised by $\mathbf{C}$. The neural network we consider is a conditional autoregressive transformer decoder model [8], using a minimal token vocabulary of $\{0, 1\}$ to represent bitstring $\mathbf{x}$. The conditional parameters $\mathbf{C}$ are each incorporated into the model by a simple




Figure 1: **(a) Overview of the active learning pipeline.** The model is first trained on a small dataset uniformly sampled across the phase space. Then, it follows an iterative active learning loop: at each round, an acquisition function determines the most informative regions of parameter space, where additional data are generated and included in the training set. **(b) Illustrative comparison of the learned distribution.** By targeting more informative regions of phase space for data acquisition, models trained with active learning produce more accurate predictions of physical quantities than models trained on a uniformly sampled dataset of equal size.

feed-forward neural network to compute a continuous embedding, which is summed with a learned token embedding $\mathbf{e}^{\text{tok}}(x_i)$, and a learned positional embedding for position i , $\mathbf{e}^{\text{pos}}(i)$.

We introduce an active learning strategy, based on *query by committee* (QBC) [9, 10], that determines at which values of $\mathbf{C}$ more data should be acquired. First, we train an ensemble of models, each initialised with different random weights, on a small dataset sampling uniformly across the domain of interest of $\mathbf{C}$. Then, we estimate the KL-divergence across the phase space between all pairs of models with Monte Carlo sampling, generate additional data in the regions of highest inter-model disagreement, and retrain the models on the expanded dataset. This process is repeated until the target dataset size S_T is reached. Fig. 1 shows an overview of the procedure.

Similar frameworks for processing different Hamiltonians with a single network have been introduced in Ref. [2, 3, 4]. A generative conditional model of this kind can be used to study the phase space of quantum systems. Ref. [2] uses peaks in the fidelity susceptibility to determine phase transitions; we use a “classical” version of the fidelity susceptibility, namely the largest eigenvalue of the Fisher information metric [11] computed from the learned distributions, which we call Fisher susceptibility.

3 Results and discussion

We first apply our active learning strategy to learning the ground state distribution in the computational basis of the periodic transverse field Ising model (TFIM) in 1D and 2D, with coupling strength fixed at $J = 1$, and using h , the transverse field strength, as the conditioning parameter (i.e. $\mathbf{C} = [h]$). We consider a 1D chain of size $N_Q = 50$ and a 7×7 2D square lattice ($N_Q = 49$). We generate data using neural quantum states (NQS) [12], employing restricted Boltzmann machines using the NetKet Python library [13] to approximate the TFIM ground states at different values of h .

In Fig. 2, we show the improvement obtained with the QBC active learning strategy over the standard uniform acquisition strategy. Specifically, we report the relative error $\epsilon = \left| 1 - \hat{E}_0 / E_0^{\text{NQS}} \right|$ where E_0^{NQS} denotes the approximate ground-state energy obtained with the NQS, and $\hat{E}_0$ is the corresponding estimate produced by our transformer model.

We also tested our strategy on the axial next-nearest neighbour Ising model (ANNNI) [14], with coupling strength $J_1 = 1$, using κ , the frustration parameter, and h , the transverse field strength, as conditioning parameters. To evaluate its effectiveness, we consider a system with $N_Q = 10$ spins and compute the KL divergence between the probability distribution learned by the transformer and the exact ground-state distribution. In Fig. 3, we show the improvement of QBC with respect to the standard uniform acquisition strategy. This model features a 2D phase space defined by κ and h . The ferromagnetic, paramagnetic and antiphase regions can be detected using $\langle m_Z^2 \rangle$ and the squared staggered magnetization [14], $\langle s_Z^2 \rangle$. In Fig. 4, we compare the order parameters obtained from the exact ground state and from the QBC training with $S_T = 10^4$. We also show the Fisher susceptibility detecting both phase transitions.

Similar results were observed for Trotterized quantum dynamics (not shown). Across quantum spin




Figure 2: Relative error ϵ as a function of the transverse field h for the 1D-TFIM (a) and the 2D-TFIM (b). Comparison between uniform acquisition at different values of S_T and the QBC protocol. At constant dataset size, relative error through the phase transition improves by a factor of $\sim 2.2 - 3$. To attain this performance with the uniform sampling strategy would require an increase in dataset size of factor ~ 3 .




Figure 3: KL divergence with respect to the ground truth as a function of κ and h for the ANNNI model with $N_Q = 10$ and $S_T = 10^4$. Panel (a) and (b) show the KL divergence for the uniform baseline training and the QBC training, respectively. Panel (c) shows the difference between the uniform and QBC strategies.




Figure 4: Phase-space analysis of the ANNNI model. Panel (a) shows the Fisher susceptibility as a function of κ and h . Panels (b) and (c) report the order parameter $\langle m_Z^2 \rangle$ obtained from the exact ground state and the transformer model, respectively, while panels (d) and (e) show the corresponding results for $\langle s_Z^2 \rangle$.

models and quantum circuit dynamics, active learning consistently improves the accuracy of learned probability distributions and derived physical observables compared to uniform data acquisition. The resulting models characterise phase diagrams and detect phase transitions through the Fisher susceptibility computed from the learned distributions. Our results indicate that active learning can substantially improve the data efficiency of generative conditional models for quantum many-body systems, reducing the amount of training data required to achieve a given level of accuracy. This is particularly relevant in settings where data generation is computationally expensive or where measurements are obtained from quantum hardware.

References

- [1] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4), December 2019.
- [2] Riccardo Rende, Luciano Loris Viteritti, Federico Becca, Antonello Scardicchio, Alessandro Laio, and Giuseppe Carleo. Foundation neural-networks quantum states as a unified ansatz for multiple hamiltonians. *Nature Communications*, 16(1), August 2025.
- [3] Zihao Qi, Christopher Earls, and Yang Peng. Neural operator quantum state: A foundation model for quantum dynamics, 2026.
- [4] Timothy Zaklama, Daniele Guerri, and Liang Fu. Attention-based foundation model for quantum states, 2025.
- [5] Dongyuan Li, Zhen Wang, Yankai Chen, Renhe Jiang, Weiping Ding, and Manabu Okumura. A survey on deep active learning: Recent advances and new frontiers, 2024.
- [6] Chengyu Dai and Sharon C. Glotzer. Efficient phase diagram sampling by active learning. *Journal of Physical Chemistry B*, 124(7):1275–1284, 2020.
- [7] Max Zhu, Jian Yao, Marcus Mynatt, Hubert Pugzlys, Shuyi Li, Sergio Bacallado, Qingyuan Zhao, and Chunjing Jia. Active learning for discovering complex phase diagrams with gaussian processes. *arXiv preprint arXiv:2409.07042*, 2024.
- [8] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023.
- [9] Hideitsu Hino and Shinto Eguchi. Active learning by query by committee with robust divergences, 2022.
- [10] H. S. Seung, M. Oppen, and H. Sompolinsky. Query by committee. In *Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory*, Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory, pages 287–294. Publ by ACM, 1992. Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory ; Conference date: 27-07-1992 Through 29-07-1992.
- [11] Victor Kasatkin, Evgeny Mozgunov, Nicholas Ezzell, and Daniel Lidar. Detecting quantum and classical phase transitions via unsupervised machine learning of the fisher information metric, 2024.
- [12] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, February 2017.
- [13] Filippo Vicentini, Damian Hofmann, Attila Szabó, Dian Wu, Christopher Roth, Clemens Giuliani, Gabriel Pescia, Jannes Nys, Vladimir Vargas-Calderón, Nikita Astrakhantsev, and Giuseppe Carleo. Netket 3: Machine learning toolbox for many-body quantum systems. *SciPost Phys. Codebases*, page 7, 2022.
- [14] M. Cea, M. Grossi, S. Monaco, E. Rico, L. Tagliacozzo, and S. Vallecorsa. Exploring the phase diagram of the quantum one-dimensional annni model, 2024.