Quantum phase detection generalization from marginal quantum neural network models

Saverio Monaco [1, 2, *] Oriel Kiss [1, 3, *] Antonio Mandarino[4] Sofia Vallecorsa[1] and Michele Grossi[1, †]

[1]*European Organization for Nuclear Research (CERN), Geneva 1211, Switzerland*
[2]*Department of Physics, University of Padova, 35122 Padova PD, Italy*
[3]*Department of Particle and Nuclear Physics, University of Geneva, Geneva 1211, Switzerland*
[4]*International Centre for Theory of Quantum Technologies,*
University of Gdańsk, Jana Bażyńskiego 1A, 80-309 Gdańsk, Poland
(Dated: January 25, 2023)

Quantum machine learning offers a promising advantage in extracting information about quantum states, e.g. phase diagram. However, access to training labels is a major bottleneck for any supervised approach, preventing getting insights about new physics. In this Letter, using quantum convolutional neural networks, we overcome this limit by determining the phase diagram of a model where analytical solutions are lacking, by training only on marginal points of the phase diagram, where integrable models are represented. More specifically, we consider the axial next-nearest-neighbor Ising (ANNNI) Hamiltonian, which possesses a ferromagnetic, paramagnetic and antiphase, showing that the whole phase diagram can be reproduced.

Introduction. Quantum machine learning (QML) [1], where parametrized quantum circuits [2] act as statistical models, has attracted much attention recently, with applications to natural sciences [3–8] or in generative modeling [9–13]. Even if QML models benefit from high expressivity [14] and demonstrated superiority over classical models in some specific cases [15, 16], it is still unclear what kind of advantage could be obtained with quantum computers [17] in the era of deep neural networks.

Quantum data, on the other hand, could be a natural paradigm to apply QML, where quantum advantages have already been demonstrated [18]. There is hope that quantum data could be collected via quantum sensors [19], and eventually directly linked to quantum computers. In this Letter, we emulate the possibility of working with quantum data by constructing them directly on a quantum device. We use a variational ground state solver to obtain approximations of the true ground states in order to mimic noisy real world data. Specifically, this Letter addresses the computation of the phase diagram of the ground states of a Hamiltonian H using a supervised learning approach. Even if similar problems have already been explored for the binary case [20, 21], with multiple classes [22] and computed on a superconducting platform [23], all of these approaches suffer from a limitation by construction, a bottleneck. In fact, since labels are needed for the training, and because they are computed analytically or numerically, these techniques can only speed up calculations, but cannot extend beyond their validated domain. Alternatively, anomaly detection (AD), an unsupervised learning technique, has been proposed [24, 25] as a way to bypass this bottleneck, by finding structure inside the data set. However, AD can only obtain qualitative, possibly unstable, results and the

performance can therefore be difficult to assess. Instead, the proposed approach provides a clear prediction for the boundaries of the adopted model, with the possibility to evaluate the performance on a validation set.

This Letter numerically demonstrates that QML can make predictions to regions where analytical labels do not exist, after being only trained on easily computable subregions. Moreover, QML only needs very few training labels to do so, as already pointed out by Caro *et al.* [26] and Banchi *et al.* [27]. In particular, we make a step toward an out-of-distribution generalization [28], where the training and testing set do not belong to the same data distribution, which is known to be a difficult task [29]. This drastically changes the perspective, extending QML capabilities to extrapolate and eventually discover new physics when trained on well-established simpler models.

The model. We consider the axial next-nearest-neighbour Ising (ANNNI) model

$$H = J \sum_{i=1}^{N} \sigma_x^i \sigma_x^{i+1} - \kappa \sigma_x^i \sigma_x^{i+2} + h \sigma_z^i, \qquad (1)$$

where σ_a^i are the Pauli matrices acting on the ith spin, $a \in \{x, y, z\}$, and we assume open boundary conditions. The energy scale of the Hamiltonian is given by the coupling constant J (without loss of generality we set $J = 1$), while the dimensionless parameters κ and h account for the next-nearest-neighbor interaction and the transverse magnetic field, respectively. We restrict ourselves to $\kappa \geq 0$, $h \geq 0$ and even N. The difference of sign between the nearest and next-nearest interactions, leading to a ferro- or antiferro-magnetic exchange in the system, is responsible for the magnetic frustration. Thence, the ANNNI model offers the possibility to study the competing mechanism of quantum fluctuations due to the transverse magnetic field and frustration.

---

* These authors contributed equally to this work.
† michele.grossi@cern.ch

The phase diagram of the quantum model at $T = 0K$ temperature has been studied mainly by renormalization group or Monte Carlo techniques in d dimensions exploiting also the correspondence with the classical analog in $d + 1$ dimensions [30–35]. The phase diagram is quite rich and three phases have been confirmed, separated by two second-order phase transitions. The first, for low frustration ($\kappa < 0.5$) of the Ising type separates the ferromagnetic and the paramagnetic phases along the line $h_I(\kappa) \approx \frac{1-\kappa}{\kappa}\left(1 - \sqrt{\frac{1-3\kappa+4\kappa^2}{1-\kappa}}\right)$. The other one of a commensurate-incommensurate type appears between the paramagnetic phase and an antiphase for values of the field $h_C(\kappa) \approx 1.05\sqrt{(\kappa - 0.5)(\kappa - 0.1)}$, in the high frustration sector ($\kappa > 0.5$). As usual, the paramagnetic phase is the disordered one, in contrast with the two ordered phases: the ferromagnetic and the antiphase one. In particular, they are different because the former is characterized by all the spins aligned along the field direction, and the latter has a four-spin periodicity, composed of repetitions of two pairs of spins pointing in opposite directions. The point $\kappa = 0.5$ represents a multicritical point.

We mention here that other relevant lines have been numerically addressed but not confirmed. One signaling an infinite-order phase transition of the Berezinskii–Kosterlitz–Thouless (BKT) type for $h_{BKT}(\kappa) \approx 1.05(\kappa - 0.5)$, delimiting a floating phase between the paramagnetic and the antiphase [34], and a disorder line where the model is exactly solvable known as the Peschel-Emery (PE) line $h_{PE}(\kappa) \approx \frac{1}{4\kappa} - \kappa$ [33, 36].

Variational state preparation. The purpose of the variational quantum eigensolver (VQE) [37] is to calculate the ground state energy of a Hamiltonian $H(\kappa, h)$ on a quantum computer. Using the Rayleigh-Ritz variational principle, the VQE minimizes the energy expectation value of a parametrized wavefunction and has been successfully applied in quantum chemistry [38–40], in nuclear physics [41–43] or in frustrated magnetic systems [44, 45]. Here, we are interested in the final eigenstates, represented by an ansatz $|\psi(\theta; \kappa, h)\rangle$, to be used as quantum data. Typically, the ansatz is chosen as a hardware-efficient (HEA) quantum circuit [38, 46], which is built with low connectivity and gates that can be easily run on noisy intermediate-scale quantum (NISQ) [47] devices. For instance, we use $D = 6(9)$ repetitions of a layer consisting of free rotations around the y axis $R_y(\theta) = e^{-i\theta\sigma_y/2}$ and controlled-NOT (CNOT) gates with linear connectivity $CX_{i,i+1}$ for $0 \leq i < N$ [48], for $N = 6(12)$ spin systems. The optimization is performed using the gradient-descent-based ADAM algorithm [49], with an initial learning rate of 0.3 and a parameter recycling scheme to improve the convergence [50]. Moreover, we note that the VQE can also be used to recursively compute excited states [51], which we used to show that the ground states of the ANNNI model are only degenerate at the boundaries in the phase diagram,

where the ground states corresponding to the different phases are competing, excluding the bit flip symmetry at $h = 0$. Finally, we asses the accuracy of the VQE states by comparing with the exact energy and observe that the relative error ratio is always below 1%. Moreover, it seems that the energy accuracy distribution is able to reveal the Peschel-Emery line, since the predicted energy values are more accurate along it. More details about the implementation, optimization, degeneracy and accuracy can be found in Appendix A.

Quantum convolutional neural networks (QCNNs). QCNNs are a class of quantum circuits, inspired by classical convolutional neural networks (CNN) [52], originally proposed in [20]. The QCNN is trained to detect quantum phase transitions, effectively learning an observable $O(\theta)$ that linearly separates two states $|\psi_A\rangle$ and $|\psi_B\rangle$ from two different phases A and B, such that $\langle\psi_A|O(\theta)|\psi_A\rangle < 0 < \langle\psi_B|O(\theta)|\psi_B\rangle$ [53], which exist since the phases in the ANNNI model are not topological. Intuitively, non topological phases of matter exhibit macroscopic differences, which can be captured by the variational observable $O(\theta)$. In principle, quantum phase detection could be performed by measuring different string order parameters (SOPs) [20]. However, the SOP vanishes near the phase transition, thus requiring exponentially many samples for the classification. On the other hand, the QCNN output is much sharper, therefore reducing the sample complexity. This changes quantum phase detection to the task of designing and training an appropriate ansatz.

In our implementation, the QCNN starts with a free rotation layer around the y axis, followed by blocks consisting of convolutions, free rotations, and pooling layers that halve the number of qubits to k until $k = \lceil\log_2(K)\rceil$, where K it the total number of quantum phases. Finally, a fully connected layer and measurement are performed in the computational basis. An example with $N = 6$ qubits is shown in Figure 1 where we have free y axis rotations (yellow), $R(\vec{\theta}) = \bigotimes_{i=1}^{N} R_y(\vec{\theta}_i)$, two-qubit convolutions (light green) $C(\theta) = \bigotimes_{i=1}^{2} R_y(\theta)$, pooling (red) $P(\vec{\theta}, \phi, b) = R_y(\vec{\theta}_b)R_x(\phi)$ with $b \in \{0, 1\}$ the value of the measured qubit, and a two-qubit fully connected (dark green) gate $F(\vec{\theta}^{(1)}, \vec{\theta}^{(2)}) = \left(\bigotimes_{i=1}^{2} R_y(\vec{\theta}_1^{(i)})R_x(\vec{\theta}_2^{(i)})R_y(\vec{\theta}_3^{(i)})\right)CX_{1,2}$.

QCNNs have been shown to be resistant to barren plateaus [54] due to their distance from low $T2$ design and are therefore good candidates for any quantum learning tasks. The analogy with CNN holds in the quantum settings since convolution and pooling layers are functions of shared parameters and the reduction of the circuit's dimension is guaranteed by the intermediate measurement. The whole algorithm flow starts with the QCNN taking as input ground states $|\psi(\theta; \kappa, h)\rangle$ from the Hamiltonian family $H(\kappa, h)$, obtained through the VQE. The quantum network then outputs the
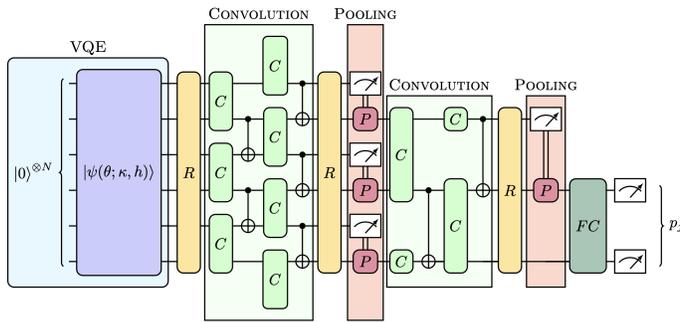
Figure 1: Circuit architecture: VQE states (blue) are the input of the quantum convolutional neural network composed of free rotations R (yellow), convolutions C (light green), pooling P (red), and a fully connected layer F (dark green).
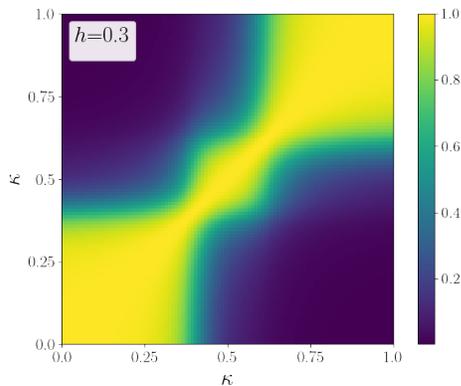
Figure 2: Fidelity between the ground states of the ANNNI model at $h = 0.3$ and $N = 12$. We observe three different cluster, corresponding to the ferromagnetic, paramagnetic, and antiphase, respectively.

probability $p_j(\kappa, h)$ of being in one of the $K = 3$ phases (ferromagnetic, paramagnetic or antiphase), where $p_j(\kappa, h)$ is computed as the probability of measuring the state $|01\rangle, |10\rangle, |11\rangle$ on the two output qubits. Since the phase diagram of the ANNNI model only contains three phases, the state $|00\rangle$ is interpreted as a *garbage* class.

Generalization.

The main contribution of this Letter is to demonstrate the ability of QCNN to work in a partial supervised approach and thus get closer to an out-of-distribution generalization by training on a set of easily available labels. We first argue that this generalization is expected to hold according to [27] if the ground states of the ANNNI model are clustered, i.e., if the fidelity between states in the same phase is high [55–57], while being low between different phases. This is indeed the case as shown in Figure 2 along the line $h = 0.3$ for the $N = 12$ spin case. Even if the requirements of the generalisation results from [26] do not hold since the training data

are only located on the boundaries, and specifically not independent and identically distributed (i.i.d.), we observe a numerical agreement with the generalization error's scaling behavior predicted in Ref. [26], i.e., $\mathcal{O}\left(\sqrt{\frac{T}{n}}\right)$, where T is the number of parameters and n the number of training points. Since the QCNN is composed of $T = \mathcal{O}[\log(N)]$ parameters [20], we can control the expected risk by training on $n = \mathcal{O}[\log(N)]$ points.

Training set. The training data set consists of the composition of points from two analytical models derived from the simplification of the physical model used. Specifically, we consider the integrable Ising model in transverse field in the case $\kappa = 0$ and the *quasi classical* model when $h = 0$, where quantum fluctuations no longer exist. We demonstrate that QCNNs extend their prediction to the all phase diagram when only trained on the marginal model given by $\mathcal{S}_X^n \subseteq \{(\kappa, h) \in \{0\} \times [0, 2]\} \cup \{(\kappa, h) \in [0, 1] \times \{0\}\}$. We consider three types of subsets $X \in \{GC, G2, U\}$, $\mathcal{S}_{GC}^n$ where n training points are sampled normally around each critical point $\{(0,1), (0.5,0)\}$, $\mathcal{S}_{G2}^n$ where n training points are sampled normally at the middle of each phase $\{(0,1.5), (0,0.5), (0.25,0), (0.75,0)\}$ and $\mathcal{S}_U^n$ where n data points are drawn uniformly on both axes. The QCNN is trained using the cross entropy $\mathcal{L}$ loss,

$$\mathcal{L} = -\frac{1}{|\mathcal{S}_X^n|} \sum_{(\kappa, h) \in \mathcal{S}_X^n} \sum_{j=1}^{K} y_j(\kappa, h) \log(p_j(\kappa, h)) \quad (2)$$

between the one-hot classical labels $y_j(\kappa, h)$ and the predictions on the training region $\mathcal{S}_X^n$ of the phase space.

Results.

Once we have introduced the problem and defined the techniques used, we can analyze the quality of the results obtained under ideal conditions with a quantum simulator.

We study our ability to reconstruct the phase diagram of the ANNNI model, characterized by a non trivial disordered paramagnetic phase, the ordered ferromagnetic phase, and antiphase one. To test the stability of the proposed approach, we consider the model with an increasing number of spins $N = 6, 12$ and sampling a different number of points $0 < n \leq 100$ used for the training. By virtue of the quality of the results, we evaluated the influence of different sampling of the training points corresponding to the two physical models that could affect the quality of the classification. A summary of the results can be qualitatively seen in Figure 3, while more quantitative results for the QCNN are displayed in Appendix B. In the first row, we have the phase diagram reconstruction for the ANNNI model with six spins, where the black lines represent the analytical transition explained above in the model section. The second line in the figure shows the same for a system with $N = 12$ spins.
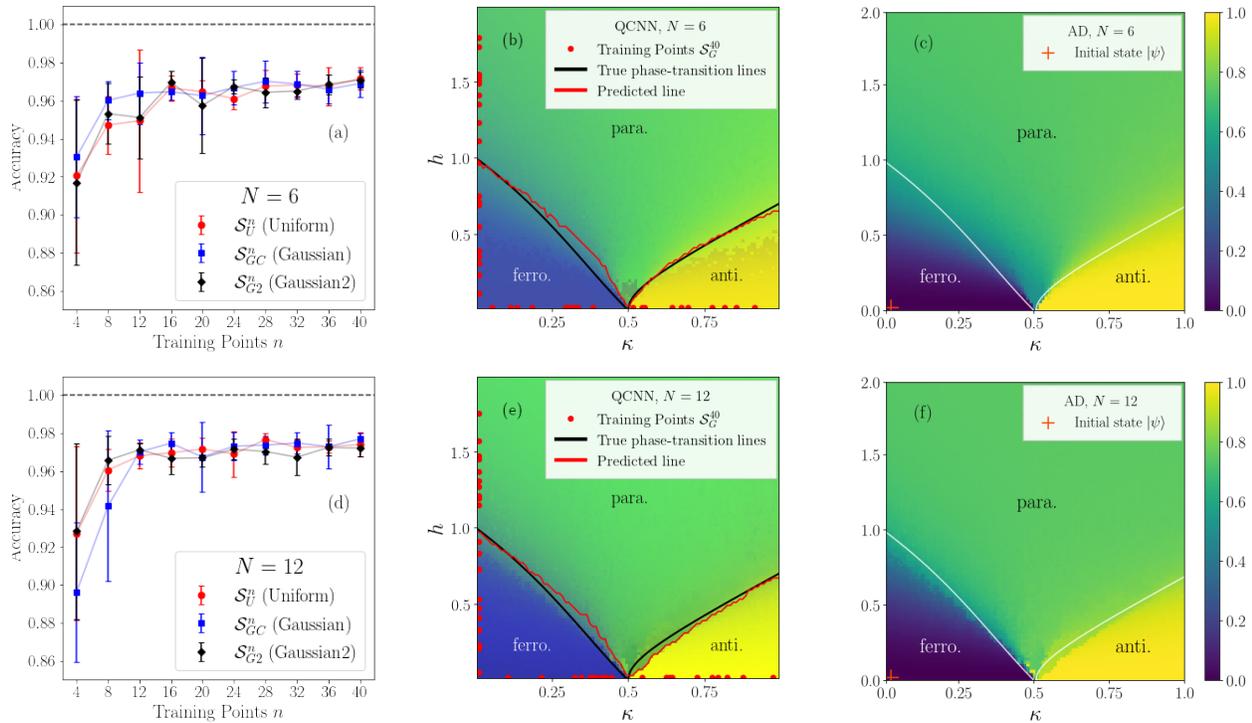
Figure 3: Quantum phases classification. (a) shows the classification accuracy of the QCNN as a function of the number of training points n, for the Gaussian centered around the critical points (GC blue), Gaussian centered around the middle of each phase ($G2$ black) and uniform sampling (red). (b) displays the phase diagram predicted by the QCNN trained on $\mathcal{S}_{G2}^{40}$ (red dots) where the color represents the probability mixture of being in one of the three phases and the red lines the predicted boundaries, while (c) shows the anomaly score for a $N = 6$ spins systems trained on the initial state $|\psi\rangle$ (red cross). (d)-(f) are similar but for $N = 12$ spins. The black lines are $h_I(\kappa)$ for $\kappa < 0.5$ and $h_C(\kappa)$ for $\kappa > 0.5$

The first column shows the accuracy, computed on the whole phase space, as a function of the number of training points n, for the Gaussian centered around the critical points $X = GC$ (blue), around the middle of each phase $X = G2$ (black), and the uniform $X = U$ (red) sampling scheme, where the error bars correspond to one standard deviation from ten independent runs. We observe that the accuracy quickly increases with n, before saturating for $n \geq 20$, as argued in Ref. [26] and that the sampling strategy does not play a major role. More importantly, sampling away from the critical points is enough. The second column displays the phase diagram obtained with the QCNN trained on $n = 40$ points. Colour shades represent the continuous probability distribution of the QCNN for our multiclass classifier as a probability mixture (blue, green and yellow times the relevant probability), while the red lines represent the predicted boundaries. The individual probabilities of each phases predicted by the QCNN are shown in Appendix B. The last column instead shows the comparison to the unsupervised learning approach inspired from [25] where the autoencoder is trained to compress the single red cross $|\psi\rangle$, and tested on the remaining points. In a nutshell, the

autoencoder is expected to perform poorly if the states are far away in the Hilbert space, i.e., if they belong to different phases, thus leading to a high compression score. The color scale shows the compression loss of each state. Additional details about the implementation of the anomaly detection can be found in Appendix C. It is worth noting that although only one training point is sufficient to obtain a qualitatively good phase diagram, only QCNNs allow a quantitative prediction for the phase. Moreover, while the QCNNs are stable when changing the training set, it is easy to find initial states where AD performs poorly, for instance by starting in the paramagnetic phase. The relatively good performance of AD can be explained by the product state nature of the training point. Hence, the product state can be easily compressed with the autoencoder [58], while states corresponding to a high magnetic field cannot.

Conclusions. This Letter addresses the computation of the phase diagram of a non integrable model, by training a QCNN on the limiting integrable regions of the considered ANNNI model. We provide numerical evidence that the QCNNs are able to generalize from non-i.i.d. training data, which is a challenging task in

general. The numerical simulations suggest that QCNNs can carry this task with more than 97% accuracy, using only $n = 20$ quantum data points distributed on the two integrable axes of the phase space. Moreover, the data points do not need to be close to the critical points. The accuracy of the QCNN quickly increases to reach its maximum as a function of the number of training points, suggesting that QCNNs can generalize from a few data points. Being a supervised method, the QCNN is not able to detect phases that are not present in the training set $\mathcal{S}_X^n$, i.e., the boundaries, such as the BKT phase transition and the PE line. Nevertheless, AD is also not able to reveal them and is limited to qualitative predictions, while a supervised approach gives quantitative results whose quality can be easily evaluated on the validation set. Moreover, by approaching out-of-distribution generalization, we propose a solution to the bottleneck of needing training labels, that are generally challenging to obtain. Consequently, we make a step into extending the reach of QML to useful applications in physics. Future work should be performed to detect phases not present in the training set, such as the floating phase or the PE line, by either affording $\mathcal{O}(1)$ training points inside these unrepresented phases or mixing the QCNN with the unsupervised approach.

CODE AVAILABILITY

The code used to generate the data set and the figures of the present Letter is publicly available [59].

[1] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd, "Quantum machine learning," Nature **549**, 195–202 (2017).

[2] Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini, "Parameterized quantum circuits as machine learning models," Quantum Science and Technology **4**, 4 (2019).

[3] Oriel Kiss, Francesco Tacchino, Sofia Vallecorsa, and Ivano Tavernelli, "Quantum neural networks force field generation," Mach. Learn.: Sci. Technol. **3** (2022), https://doi.org/10.1088/2632-2153/ac7d3c.

[4] Vishal S. Ngairangbam, Michael Spannowsky, and Michihisa Takeuchi, "Anomaly detection in high-energy physics using a quantum autoencoder," Phys. Rev. D **105**, 095004 (2022).

[5] Chen Zhou, Jay Chan, Wen Guan, Shaojun Sun, Alex Zeng Wang, Sau Lan Wu, Miron Livny, Federico Carminati, Alberto Di Meglio, Andy C. Y. Li, Joseph Lykken, Panagiotis Spentzouris, Samuel Yen-Chi Chen, Shinjae Yoo, and Tzu-Chieh Wei, "Application of Quantum Machine Learning to High Energy Physics Analysis at LHC using IBM Quantum Computer Simulators and IBM Quantum Computer Hardware," PoS **ICHEP2020**, 930 (2021).

[6] Sau Lan Wu, Jay Chan, Wen Guan, Shaojun Sun, Alex Wang, Chen Zhou, Miron Livny, Federico Carminati, Alberto Di Meglio, Andy C Y Li, Joseph Lykken, Panagiotis Spentzouris, Samuel Yen-Chi Chen, Shinjae Yoo, and Tzu-Chieh Wei, "Application of quantum machine learning using the quantum variational classifier method to high energy physics analysis at the LHC on IBM quantum computer simulator and hardware with 10 qubits," Journal of Physics G: Nuclear and Particle Physics **48**, 125003 (2021).

[7] YaoChong Li, Ri-Gui Zhou, RuQing Xu, Jia Luo, and WenWen Hu, "A quantum deep convolutional neural network for image recognition," Quantum Science and Technology **5**, 044003 (2020).

[8] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," Phys. Rev. A **98**, 032309 (2018).

[9] Daniel Mills Brian Coyle, Vincent Danos, and Elham Kashefi, "The born supremacy: quantum advantage and training of an ising born machine," npj Quantum Inf **6**, 60 (2020).

[10] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner, "Quantum generative adversarial networks for learning and loading random distributions," npj Quantum Inf **5** (2019), https://doi.org/10.1038/s41534-019-0223-2.

[11] Manuel S. Rudolph, Ntwali Bashige Toussaint, Amara Katabarwa, Sonika Johri, Borja Peropadre, and Alejandro Perdomo-Ortiz, "Generation of high-resolution handwritten digits with an ion-trap quantum computer," Phys. Rev. X **12**, 031010 (2022).

[12] Oriel Kiss, Michele Grossi, Enrique Kajomovitz, and Sofia Vallecorsa, "Conditional born machine for monte carlo event generation," Phys. Rev. A **106**, 022612 (2022).

[13] Andrea Delgado and Kathleen E. Hamilton, "Unsupervised quantum circuit learning in high energy physics," Phys. Rev. D **106**, 096006 (2022).

[14] Amira Abbas, David Sutter, Christa Zoufal, Aurelien Lucchi, Alessio Figalli, and Stefan Woerner, "The power of quantum neural networks," Nature Computational Science **1**, 403–409 (2021).

[15] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R. McClean, "Power of data in quan-

[15] tum machine learning," Nat Commun **2631** (2021), https://doi.org/10.1038/s41467-021-22539-9.

[16] Jennifer R. Glick, Tanvi P. Gujarati, Antonio D. Corcoles, Youngseok Kim, Abhinav Kandala, Jay M. Gambetta, and Kristan Temme, "Covariant quantum kernels for data with group structure," ArXiv e-prints (2021), arXiv:2105.03406 [quant-ph].

[17] Maria Schuld and Nathan Killoran, "Is quantum advantage the right goal for quantum machine learning?" PRX Quantum **3**, 030101 (2022).

[18] Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, and Jarrod R. McClean, "Quantum advantage in learning from experiments," Science **376**, 1182–1186 (2022).

[19] Christian D. Marciniak, Thomas Feldker, Ivan Pogorelov, Raphael Kaubruegger, Denis V. Vasilyev, Rick van Bijnen, Philipp Schindler, Peter Zoller, Rainer Blatt, and Thomas Monz, "Optimal metrology with programmable quantum sensors," Nature **603**, 604–609 (2022).

[20] Iris Cong, Soonwon Choi, and Mikhail D. Lukin, "Quantum convolutional neural networks," Nat. Phys. **15**, 1273–1278 (2019).

[21] A. V. Uvarov, A. S. Kardashin, and J. D. Biamonte, "Machine learning phase transitions with a quantum processor," Phys. Rev. A **102**, 012415 (2020).

[22] Marco Lazzarin, Davide Emilio Galli, and Enrico Prati, "Multi-class quantum classifiers with tensor network circuits for quantum phase recognition," Physics Letters A **434**, 128056 (2022).

[23] Johannes Herrmann, Sergi Masot Llima, Ants Remm, Petr Zapletal, Nathan A. McMahon, Colin Scarato, François Swiadek, Christian Kraglund Andersen, Christoph Hellings, Sebastian Krinner, Nathan Lacroix, Stefania Lazar, Michael Kerschbaum, Dante Colao Zanuz, Graham J. Norris, Michael J. Hartmann, Andreas Wallraff, and Christopher Eichler, "Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases," Nat Commun **13** (2022), https://doi.org/10.1038/s41467-022-31679-5.

[24] Korbinian Kottmann, Patrick Huembeli, Maciej Lewenstein, and Antonio Acín, "Unsupervised phase discovery with deep anomaly detection," Phys. Rev. Lett. **125**, 170603 (2020).

[25] Korbinian Kottmann, Friederike Metz, Joana Fraxanet, and Niccolò Baldelli, "Variational quantum anomaly detection: Unsupervised mapping of phase diagrams on a physical quantum computer," Phys. Rev. Research **3**, 043184 (2021).

[26] Matthias C. Caro, Hsin-Yuan Huang, M. Cerezo, Kunal Sharma, Andrew Sornborger, Lukasz Cincio, and Patrick J. Coles, "Generalization in quantum machine learning from few training data," Nat Commun **13**, 4919 (2022).

[27] Leonardo Banchi, Jason Pereira, and Stefano Pirandola, "Generalization in quantum machine learning: A quantum information standpoint," PRX Quantum **2**, 040321 (2021).

[28] Matthias C. Caro, Hsin-Yuan Huang, Nicholas Ezzell, Joe Gibbs, Andrew T. Sornborger, Lukasz Cincio, Patrick J. Coles, and Zoë Holmes, "Out-of-distribution generalization for learning quantum dynamics," ArXiv e-prints (2022), arXiv:2204.10268 [quant-ph].

[29] Haotian Ye, Chuanlong Xie, Tianle Cai, Ruichen Li, Zhenguo Li, and Liwei Wang, "Towards a theoretical framework of out-of-distribution generalization," in Advances in Neural Information Processing Systems, edited by A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (2021).

[30] Walter Selke, "The annni model — theoretical analysis and experimental application," Physics Reports **170**, 213–264 (1988).

[31] Paulo R. Colares Guimarães, João A. Plascak, Francisco C. Sá Barreto, and João Florencio, "Quantum phase transitions in the one-dimensional transverse ising model with second-neighbor interactions," Phys. Rev. B **66**, 064413 (2002).

[32] Anjan Kumar Chandra and Subinay Dasgupta, "Floating phase in the one-dimensional transverse axial next-nearest-neighbor ising model," Phys. Rev. E **75**, 021105 (2007).

[33] Matteo Beccaria, Massimo Campostrini, and Alessandra Feo, "Density-matrix renormalization-group study of the disorder line in the quantum axial next-nearest-neighbor ising model," Phys. Rev. B **73**, 052402 (2006).

[34] Matteo Beccaria, Massimo Campostrini, and Alessandra Feo, "Evidence for a floating phase of the transverse annni model at high frustration," Phys. Rev. B **76**, 094410 (2007).

[35] Askery Canabarro, Felipe Fernandes Fanchini, André Luiz Malvezzi, Rodrigo Pereira, and Rafael Chaves, "Unveiling phase transitions with machine learning," Phys. Rev. B **100**, 045129 (2019).

[36] I Peschel and VJ Emery, "Calculation of spin correlations in two-dimensional ising systems from one-dimensional kinetic models," Zeitschrift für Physik B Condensed Matter **43**, 241–249 (1981).

[37] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," Nature Communications **5**, 4123 (2014).

[38] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," Nature **549**, 242–246 (2017).

[39] Panagiotis Kl. Barkoutsos, Jerome F. Gonthier, Igor Sokolov, Nikolaj Moll, Gian Salis, Andreas Fuhrer, Marc Ganzhorn, Daniel J. Egger, Matthias Troyer, Antonio Mezzacapo, Stefan Filipp, and Ivano Tavernelli, "Quantum algorithms for electronic structure calculations: Particle-hole hamiltonian and optimized wave-function expansions," Phys. Rev. A **98**, 022322 (2018).

[40] Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik, "Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz," Quantum Sci. Technol. **4**, 014008 (2019).

[41] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, "Cloud quantum computing of an atomic nucleus," Phys. Rev. Lett. **120**, 210501 (2018).

[42] I. Stetcu, A. Baroni, and J. Carlson, "Variational approaches to constructing the many-body nuclear ground state for quantum computing," Phys. Rev. C **105**, 064308 (2022).

[43] Oriel Kiss, Michele Grossi, Pavel Lougovski, Federico Sanchez, Sofia Vallecorsa, and Thomas Papenbrock, "Quantum computing of the ^{6}Li nucleus via ordered unitary coupled clusters," Phys. Rev. C **106**, 034325 (2022).

[44] Alexey Uvarov, Jacob D. Biamonte, and Dmitry Yudin, "Variational quantum eigensolver for frustrated quantum systems," Phys. Rev. B **102**, 075104 (2020).

[45] Michele Grossi, Oriel Kiss, Francesco De Luca, Carlo Zollo, Ian Gremese, and Antonio Mandarino, "Finite-size criticality in fully connected spin models on superconducting quantum hardware," ArXiv e-prints (2022), arXiv:2208.02731 [quant-ph].

[46] Sukin Sim, Peter D. Johnson, and Alán Aspuru-Guzik, "Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms," Advanced Quantum Technologies **2**, 1900070 (2019).

[47] John Preskill, "Quantum Computing in the NISQ era and beyond," Quantum **2**, 79 (2018).

[48] Michael A. Nielsen and Isaac L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, 2010).

[49] Diederik P. Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, edited by Yoshua Bengio and Yann LeCun (2015).

[50] Stuart M. Harwood, Dimitar Trenev, Spencer T. Stober, Panagiotis Barkoutsos, Tanvi P. Gujarati, Sarah Mostame, and Donny Greenberg, "Improving the variational quantum eigensolver using variational adiabatic quantum computing," ACM Transactions on Quantum Computing **3** (2022), 10.1145/3479197.

[51] Oscar Higgott, Daochen Wang, and Stephen Brierley, "Variational quantum computation of excited states," Quantum **3**, 156 (2019).

[52] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," Neural Computation **1**, 541–551 (1989).

[53] Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V. Albert, and John Preskill, "Provably efficient machine learning for quantum many-body problems," Science **377**, eabk3333 (2022), https://www.science.org/doi/pdf/10.1126/science.abk3333.

[54] Arthur Pesah, M. Cerezo, Samson Wang, Tyler Volkoff, Andrew T. Sornborger, and Patrick J. Coles, "Absence of barren plateaus in quantum convolutional neural networks," Phys. Rev. X **11**, 041011 (2021).

[55] Matteo Bina, Antonio Mandarino, Stefano Olivares, and Matteo G. A. Paris, "Drawbacks of the use of fidelity to assess quantum resources," Phys. Rev. A **89**, 012305 (2014).

[56] Antonio Mandarino, Matteo Bina, Stefano Olivares, and Matteo GA Paris, "About the use of fidelity in continuous variable systems," International Journal of Quantum Information **12**, 1461015 (2014).

[57] Antonio Mandarino, Matteo Bina, Carmen Porto, Simone Cialdi, Stefano Olivares, and Matteo G. A. Paris, "Assessing the significance of fidelity as a figure of merit in quantum state reconstruction of discrete and continuous-variable systems," Phys. Rev. A **93**, 062118 (2016).

[58] Jonathan Romero, Jonathan P Olson, and Alan Aspuru-Guzik, "Quantum autoencoders for efficient compression of quantum data," Quantum Science and Technology **2**, 045001 (2017).

[59] https://phaseestimation.readthedocs.io/en/hotfixes-1.0.0a/index.html.

[60] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, Shahnawaz Ahmed M. Sohaib Alam, Juan Miguel Arrazola, Carsten Blank, Alain Delgado, Soran Jahangiri, Keri McKiernan, Johannes Jakob Meyer, Zeyue Niu, Antal Száva, and Nathan Killoran, "Pennylane: Automatic differentiation of hybrid quantum-classical computations," ArXiv e-prints (2018), arXiv:1811.04968 [quant-ph].

[61] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran, "Evaluating analytic gradients on quantum hardware," Phys. Rev. A **99**, 032331 (2019).

Appendix A: Quantum data set

The variational quantum eigensolver (VQE) is used to construct a quantum data set in the form of quantum circuits representing the ground states of the ANNNI model. In a nutshell, a parameterized wavefunction is prepared on a quantum computer, and its parameters are iteratively updated to minimize the expected value of the energy. Hence, using the variational principle, this parameterized wavefunction shall be a good approximation of the ground state, assuming that the ansatz is expressive enough and that the optimization is successful. We recall that the ansatz is constructed with $D = 6(9)$ repetitions of a layer consisting of free rotations around the y-axis $R_y(\theta) = e^{-i\theta\sigma_y/2}$ and CNOT gates with linear connectivity $CX_{i,i+1}$ for $0 \leq i < N$, for $N = 6(12)$ spins. The optimization is performed with the ADAM algorithm, which uses first-order gradient descent with momentum, where the gradients are obtained with the automatic differentiation framework provided by Pennylane [60]. We note that the gradients can be obtained on quantum hardware using the parameter-shift rule [61]. To improve the convergence and accuracy of the quantum data set, we perform five optimization rounds composed of 1000 update steps by reducing the initial learning rate from 0.3 to 0.1 between each round. Moreover, we use a parameters recycling scheme [50], where the initial parameters at the point (h, κ) are chosen to be the converged parameters from the previously computed neighboring site. This strategy improves the speed and accuracy of the optimization while keeping a high fidelity between neighboring states, as expected from a physical point of view. Figure 4 shows the accuracy of the VQE ground states energy with respect to the exact one for the $N = 12$ spins cases. We observe that the relative error ratio

$$\Delta E = \left| \frac{E_{\text{VQE}} - E_{\text{exact}}}{E_{\text{exact}}} \right| \tag{A1}$$

is always below 1%. Moreover, we notice that the accuracy is significantly better on the Peschel-Emery disorder
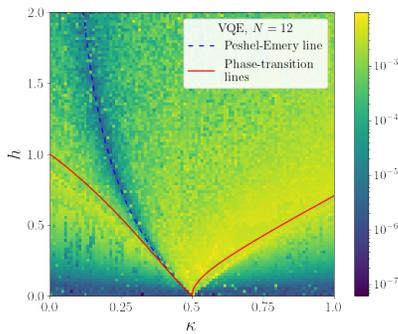
Figure 4: Normalised difference between the energy obtained with exact diagonalization and those obtained with the VQE on a log scale. Red lines correspond to the numerical reference boundaries obtained through Monte Carlo, while the dashed blue line represents the second-order PE line.
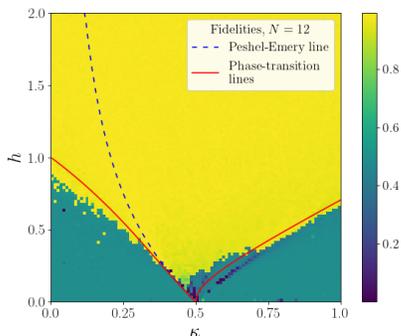
Figure 5: Fidelity between the VQE and exact ground states for $N = 12$ spins. Red lines correspond to the numerical reference boundaries obtained through Monte Carlo, while the dashed blue line represents the second-order PE line.

line, indicating that the accuracy of the VQE algorithm could be used to detect interesting regions of the phase diagram. For instance, it could be used in conjunction with the DMRG algorithm to expose such disorder lines. We also report the fidelity between the VQE and exact ground states in Figure 5. We observe that the paramagnetic states are almost exactly reproduced, while the fidelity for low magnetic field values is about 50%, and lower around the critical points.

Appendix B: Details on the predicted phase diagram

As explained in the result section, the QCNN outputs the probability $p_j(\kappa, h)$ of obtaining the four two-qubit state $|01\rangle, |10\rangle, |11\rangle$ and $|00\rangle$. We associate the first three to the physical ferro-, para-magnetic, and antiphase, while the last one is treated as a garbage class. In the main text, we plot the probability mixture of the three physical phases with a blue-green-yellow color chan-
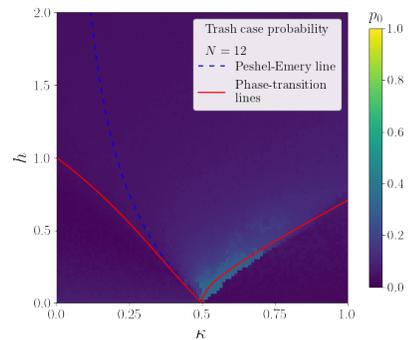
(a)
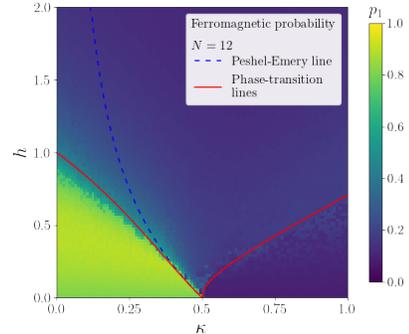
(b)
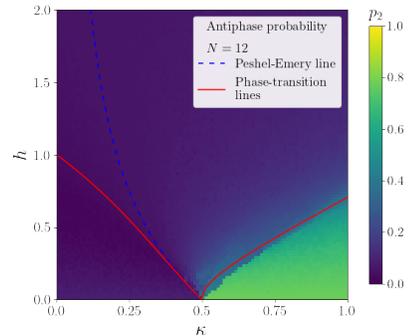
(c)
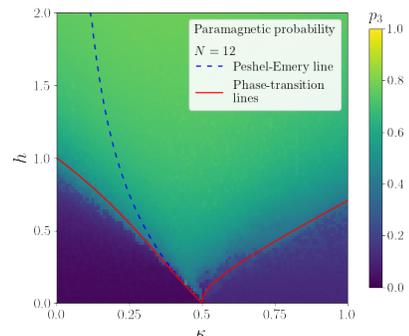
(d)

Figure 6: Individual probability $p_i(\kappa, h)$, $0 \le i < 4$ (a-d), as a function of the Hamiltonian parameters. Red lines correspond to the numerical reference boundaries obtained through Monte Carlo, while the dashed blue line represents the second-order PE line.
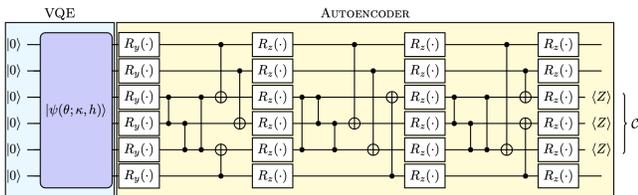
Figure 7: Compression circuit (yellow) and anomaly score measurement ($\mathcal{C}$) of the ground states of $H(\kappa, h)$ obtained through a VQE (blue). The $\cdot$ represents independent parameters.

nel. For clarity and completeness, the individual probability $p_i(\kappa, h)$ with $N = 12$ spins are shown in Figure 6, where the color indicates the magnitude of the corresponding probability.

We observe that the probability of being in the garbage class is almost zero, except near the triple point. Moreover, all predicted phase boundaries are sharp, indicating the classifier's confidence.

Appendix C: Anomaly Detection

For the reader's convenience, we will recall the unsupervised anomaly detection (AD) scheme, initially proposed by Kottmann *et al.* [25], to draw the phase diagram of the Bose-Hubbard model. Since it is an unsupervised learning technique, it bypasses the bottleneck of needing classical training labels and is, therefore, an alternative to the approach taken in this letter.

As a first step, an initial state $|\psi\rangle$ is chosen in the data set composed of the ground states of H. Although there is no formal restriction, it should lie far from any critical points. A quantum encoder [58] is then trained to learn to compress $|\psi\rangle$ on a k-qubit state $|\phi\rangle$, with

quantum register q_C and $k < N$, i.e., to write $|\psi\rangle = |\phi\rangle \otimes |T\rangle$, where the latter is a $(N-k)$-qubit trash state with register q_T. The registers here refer to the indices of the qubits composing the states. In practice, an anomaly score based on the Hamming distance between the trash state $|T\rangle$ to $|0\rangle^{\otimes(N-k)}$, written as

$$\mathcal{C} = \frac{1}{2} \sum_{j \in q_T} (1 - \langle Z_j \rangle), \tag{C1}$$

and we make the choice $k = \lfloor N/2 \rfloor$. Intuitively, the encoder compresses similar states, i.e., states in the same phase, with success but will fail to compress states in a different phase, leading to a high anomaly score. The encoder, as proposed in [25], is composed of D layers of independent $R_y(\theta)$ rotations on all qubits and $CZ_{i,j}$ gates for $i \in q_C, j \in q_T$ and $i, j \in q_T$. We use a slightly modified version, with a first layer of $R_y(\cdot)$ individual rotations, followed by $D = 3$ layers composed of $CX_{i,j}$ gates for $i \in q_C$ and $j \in q_T$, $CZ_{i,j}$ gates with $i, j \in q_T$ and independent $R_z(\cdot)$ rotations as displayed in Figure 7 for $N = 6$.

We highlight a few differences with the supervised approach. First, the anomaly score measurement is highly dependent on the choice of the initial state $|\psi\rangle$, and can often lead to phase diagrams without any clear phase separation. Moreover, there is no quantitative way to assess the validity of the predicted phase diagram, while with the QCNN, we may evaluate the accuracy on the validation set. Finally, the anomaly score only provided qualitative results. Hence, only a continuous number (the anomaly syndrome) is associated with each point, and there is no canonical way to assign it to a particular phase. On the other hand, the QCNN outputs the probability of being in each phase and therefore, a unique predicted phase can be assigned to each quantum state.