

Sample-efficient estimation of entanglement entropy through supervised learning

Maximilian Rieger ¹, Moritz Reh ^{1,*}, and Martin Gärtner ^{1,2,3,4,†}

¹*Kirchhoff-Institut für Physik, Universität Heidelberg, Im Neuenheimer Feld 227, 69120 Heidelberg, Germany*

²*Physikalisches Institut, Universität Heidelberg, Im Neuenheimer Feld 226, 69120 Heidelberg, Germany*

³*Institut für Theoretische Physik, Ruprecht-Karls-Universität Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany*

⁴*Institute of Condensed Matter Theory and Optics, Friedrich-Schiller-University Jena, Max-Wien-Platz 1, 07743 Jena, Germany*



(Received 14 September 2023; accepted 4 December 2023; published 2 January 2024)

We explore a supervised machine-learning approach to estimate the entanglement entropy of multiqubit systems from few experimental samples. We put a particular focus on estimating both aleatoric and epistemic uncertainty of the network's estimate and benchmark against the best-known conventional estimation algorithms. For states that are contained in the training distribution, we observe convergence in a regime of sample sizes in which the baseline method fails to give correct estimates, while extrapolation only seems possible for regions close to the training regime. As a further application of our method, highly relevant for quantum simulation experiments, we estimate the quantum mutual information for nonunitary evolution by training our model on different noise strengths.

DOI: [10.1103/PhysRevA.109.012403](https://doi.org/10.1103/PhysRevA.109.012403)

I. INTRODUCTION

The ultimate feat in probing the quantum nature of quantum many-body systems consists of understanding their entanglement properties [1]. Insights into phenomena such as the thermalization of closed quantum systems and many-body localization are fundamentally linked to entanglement [2–4], meaning the coherent delocalization of information among system constituents. The certification of entanglement is referred to as entanglement witnessing [5,6] and has been achieved in many systems of interest, including Bose-Einstein condensates [7], photonic systems [8], atoms in optical lattices [9], and many more. At the same time, the quantification of entanglement is disparately more challenging, as one naïvely requires the (sub-)system's density matrix ρ in order to compute the von Neumann entropy $S = -\text{tr}(\rho \ln \rho)$ or its Renyi extensions, which is of interest due to its role as an indicator of quantum phase transitions [10] or its evolution in dynamical systems [2]. However, as full tomography of ρ becomes prohibitively expensive for larger systems due to the curse of dimensionality, it becomes increasingly challenging to obtain reliable estimates of entanglement entropies at reasonable experimental and computational costs [11,12]. One option to minimize these costs consists of restricting the functional form of the state to a certain type, decreasing the number of variational parameters at the expense of introducing a bias [13–27]. Notable asymptotically unbiased estimators that have been introduced to address these scaling issues include shadow tomography [28] and distance-based approaches [29]. While being more feasible than full scale tomography of the

exponentially large density matrix ρ , the sample complexity of these methods is still too high, even for subsystems including only a small number of qubits, motivating the search for novel estimation schemes.

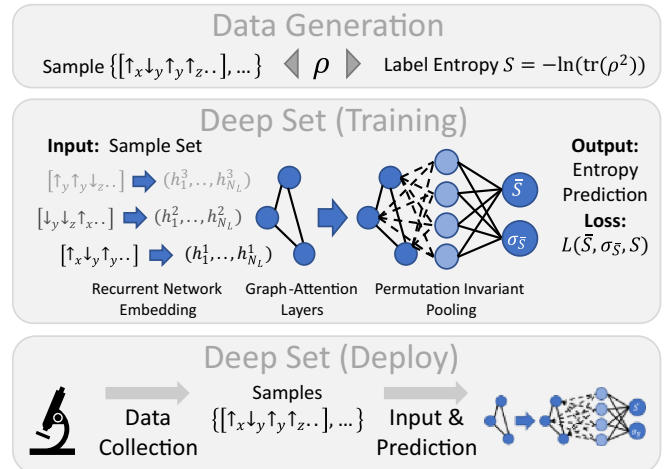


FIG. 1. Visual description of the proposed procedure. We classically simulate a system of interest described by ρ and store a quantity of interest, e.g., the entanglement entropy $S = -\ln(\text{tr}(\rho^2))$, which is difficult to estimate from samples. Simultaneously, we generate synthetic measurement data \mathcal{S} , which we use in a next step to learn a function that maps \mathcal{S} to S . The function we propose for this is a deep neural network that, in a first step, embeds all samples in a latent space, before mapping the set of latent space embeddings into another latent space, in a permutation invariant fashion. From there on, we use a feed-forward net to predict both the mean and the standard deviation of the estimate. Once trained, the model can be employed on unlabeled experimental data to give estimates for quantities with otherwise infeasible sample complexity, such as entanglement entropies.

*moritz.reh@kip.uni-heidelberg.de

†martin.gaertner@uni-jena.de

To overcome this challenge, we explore a supervised machine-learning method that aims to estimate entropic quantities of quantum states, given a small set of samples. The approach is illustrated in Fig. 1. We are motivated by recent scientific breakthroughs based on supervised learning such as AlphaFold [30], which demonstrate the capability of modern deep-learning algorithms to find seemingly intractable maps with high complexity and no apparent structure. Supervised applications in the realm of quantum physics have, for example, been explored in Refs. [31–33]. A common goal is to build a controlled estimator, meaning a reliable treatment of uncertainties of the network output, which is why we put a particular emphasis on their estimation. In contrast to related work by Koutný *et al.* [34], our aim is to use as few samples as possible, being inspired by interacting spin systems rather than photonic experiments, thereby presenting a complementary approach to the problem. Additionally, various works have studied the problem of entanglement *classification* employing machine-learning tools [35–38], while we focus on entanglement *quantification*.

We show that the proposed machine-learning method is able to outperform the more general method proposed in Ref. [29] when estimating the second-order Rényi entanglement entropy, provided that we compare them on the domain the network was trained on (in-distribution). We also examine the network’s ability to generalize to new domains and find that extrapolation is expectedly challenging. Finally, to motivate the use of the proposed method also in experimental scenarios, we study the performance in the presence of noise of varying type and strength and demonstrate correct estimation of the quantum mutual information.

II. PHYSICAL SYSTEM

While the proposed approach is applicable to any classically simulatable quantum system, we focus on the 1D transverse-field Ising model as a benchmark system.

Its Hamiltonian is given by

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

where $J \geq 0$ and we employ periodic boundary conditions. We generate entangled states by quenching the paramagnetic ground state, $|\psi_+\rangle = |+\rangle^{\otimes N}$, to the critical point $J = h = 1$. We consider both unitary time evolution, governed by the Schrödinger equation

$$i \frac{\partial}{\partial t} |\psi\rangle = H |\psi\rangle \quad |\psi(t=0)\rangle = |\psi_+\rangle, \quad (2)$$

and nonunitary time evolution, determined by the Lindblad master equation

$$\begin{aligned} \frac{\partial \rho}{\partial t} &= -i[H, \rho] + \gamma_z \sum_{i=1}^N (\sigma_z^{(i)} \rho \sigma_z^{(i)} - \rho) \\ &\quad + \gamma_- \sum_{i=1}^N \left(\sigma_-^{(i)} \rho \sigma_+^{(i)} - \frac{1}{2} \{ \sigma_+^{(i)} \sigma_-^{(i)}, \rho \} \right), \\ \rho(t=0) &= |\psi_+\rangle \langle \psi_+|. \end{aligned} \quad (3)$$

The jump operators model single-qubit noise, where σ_z corresponds to dephasing and $\sigma_- = \frac{1}{2}(\sigma_x - i\sigma_y)$ corresponds to decay. The noise strength is determined by the parameters $\gamma_z, \gamma_- \geq 0$. More details regarding the examined physical states are given in the following sections.

III. THEORETICAL CONCEPTS

We divide the system under scrutiny into the two subsystems A and B , forming a bipartition with each subsystem representing half of the spin chain. For a quantum state ρ , we define the reduced density matrix

$$\rho_A = \text{Tr}_B(\rho). \quad (4)$$

For pure states ρ , the Rényi entanglement entropy of order $n > 1$ is defined as

$$S^{(n)}(\rho_A) = \frac{1}{1-n} \ln [\text{Tr}(\rho_A^n)], \quad (5)$$

and we focus on the case $n = 2$ to allow for straightforward comparison with the conventional method proposed in Ref. [29].

However, we note that the herein developed method does not depend on this choice and can be applied to any quantity of interest. In fact, we also estimate the quantum mutual information

$$I(\rho) = S^{(1)}(\rho_A) + S^{(1)}(\rho_B) - S^{(1)}(\rho), \quad (6)$$

in the case of dissipative dynamics, generating mixed states, for which subsystem Rényi entropies no longer constitute a faithful measure of quantum correlations. Here $S^{(1)}$ is the von Neumann entropy, which can be seen by taking the limit

$$S^{(1)}(\rho) := \lim_{n \searrow 1} S^{(n)}(\rho) = -\text{Tr}[\rho \ln(\rho)]. \quad (7)$$

In the remainder of this work we abbreviate the second-order Rényi entropy of the half chain as HCE (*half chain entropy*) and the quantum mutual information as MI.

Our aim is to estimate the aforementioned quantities based on a set of projective measurements. Such projective measurements, taken in a single-basis configuration, only partly describe the quantum state at hand. To estimate entanglement entropies, one also requires phase information, which can be obtained by measuring in different basis configurations. If one is able to unambiguously infer the quantum state from the expectation values of a set of measurement operators, the set is said to be informationally complete. Positive operator valued measures (POVMs) [39] formalize this setup. A POVM \mathcal{M} is defined to be a finite set of positive self-adjoint operators,

$$\mathcal{M} = \{M_1, M_2, \dots\}, \quad M_i \geq 0, \quad M_i^\dagger = M_i, \quad (8)$$

that sum to unity, $\sum_i M_i = \mathbb{1}$. Each M_i represents one possible measurement outcome, which is observed with the probability

$$p(M_i) = \text{Tr}(\rho M_i). \quad (9)$$

Motivated by the simplicity of its experimental realization we here choose the Pauli-4-POVM, which is described by the measurement operators

$$\begin{aligned} M_1 &= \frac{1}{3} |+\rangle \langle +|, & M_2 &= \frac{1}{3} |L\rangle \langle L|, & M_3 &= \frac{1}{3} |0\rangle \langle 0|, \\ M_4 &= \mathbb{1} - M_1 - M_2 - M_3, \end{aligned} \quad (10)$$

where $|+\rangle$, $|L\rangle$, and $|0\rangle$ are the $+1$ eigenstates of the σ_x , σ_y , and σ_z operators. Operationally, this POVM can be implemented by measuring in all three Pauli bases and subsuming the outcomes with eigenvalues -1 in each basis under M_4 . The Pauli-4-POVM can be easily generalized for N qubits, by taking N -fold tensor products of the aforementioned measurement operators. A quantum state ρ of a spin chain can then be associated with a probability distribution P_ρ of generalized measurement outcomes, labeled by multi-indices $a \in \{1, 2, 3, 4\}^N$, given by

$$P_\rho(a) = \text{Tr} \left(\rho \left[\bigotimes_{i=1}^N M_{a_i} \right] \right). \quad (11)$$

IV. DATA GENERATION

Given a set of N_M measurement outcomes sampled from this probability distribution, the network will be trained to estimate quantities such as the HCE and the MI of the quantum state ρ . In order to generate the training dataset, it is necessary to classically simulate the system. Obviously, this limits the system-size regimes in which the proposed method can operate. Nevertheless, it is a meaningful extension over previous approaches as there is a gap between system sizes for which unbiased estimators give reliable results and system sizes which may be classically simulated. This statement holds irrespective of the Hamiltonian of the system and can even be extended to other quantum simulation platforms besides spin-1/2's. The dataset comprises samples from N_S states that differ in their physical parameters, for example, the evolution time t , and their associated label such as the HCE and MI. For each state we sample N_B batches of N_M measurement outcomes from the POVM probability distribution using a Markov chain Monte Carlo algorithm. The total shape of the training dataset is thus given by $(N_S, N_B, N_M, N, 4)$ where the last dimension arises from an additional one-hot encoding of the four possible single-spin measurement outcomes introduced in Eq. (10).

V. NETWORK ARCHITECTURE

The network structure can be divided into three building blocks illustrated in Fig. 1. First, all POVM samples are embedded in a latent space using a long short-term memory (LSTM). The advantage of using a recurrent architecture is its ability to process inputs of arbitrary length, which allows one to utilize the same architecture for systems of different sizes, although we do not explicitly exploit this feature in the present work. The data is then be further transformed by a fully connected, permutation equivariant graph attention network (GAT). Permutation invariance is obtained by summation of all nodes of the GAT. A permutation-invariant architecture restricts the search space to the relevant domain, eliminating redundant degrees of freedom from the parameter space and facilitating an easier optimization [40].

In the last step the output of the GAT is evaluated by a dense feed-forward neural network (DFNN) with two output neurons $(\bar{S}, \sigma_{\bar{S}})$ which we interpret as the mean \bar{S} of the quantity of interest and its statistical error $\sigma_{\bar{S}}$ of the estimation. A more detailed description of the network structure is given in the Appendix.

VI. TRAINING AND UNCERTAINTY QUANTIFICATION

We train the network in a supervised fashion and generate training data using exact dynamics, as described above. If one is interested in systems of larger size, one could also generate training data using approximate methods, such as tensor networks [41,42] or neural quantum states [43] by employing the replica trick [44].

We define the loss as

$$L(\bar{S}, \sigma_{\bar{S}}; S; \theta) = \frac{(S - \bar{S})^2}{\sigma_{\bar{S}}^2} + \ln(\sigma_{\bar{S}}^2), \quad (12)$$

where S is the label, i.e., the true Rényi entropy and θ denotes the vector of network parameters. The loss takes the form of the negative log-likelihood of a Gaussian distribution, from which the standard L_2 loss is recovered by setting $\sigma_{\bar{S}}$ to 1 [45]. The uncertainty captured by $\sigma_{\bar{S}}$ is an estimate of the aleatoric uncertainty [46] and stems from the limited information on the probability distribution that can be inferred from the finite set of samples contained in the input. The aleatoric uncertainty is thus an inherent property of the dataset and cannot be reduced by means of a larger training dataset or similar.

By contrast, the reducible part of the total uncertainty which is determined by the choice of model, training scheme, etc., is usually called epistemic uncertainty. We attempt to estimate the epistemic error by individually training an ensemble of M networks with the same structure, yet different initial parameters, and averaging their outputs, after training is finished. This approach is especially effective if the network is validated on domains that it did not encounter during training, called out-of-distribution (OOD). Referring to the outputs of model m as $(\bar{S}_m, \sigma_{\bar{S}_m})$, the total uncertainty of the ensemble is given by [45]

$$\sigma_{\bar{S}}^2 = \frac{1}{M} \left(\sum_{m=1}^M \sigma_{\bar{S}_m}^2 + \bar{S}_m^2 \right) - \left(\frac{1}{M} \sum_{m=1}^M \bar{S}_m \right)^2. \quad (13)$$

The total computational cost of the approach is divided between training data generation and training of the network. In the instance of Fig. 4, the generation of the dataset took roughly 6 h on a single NVIDIA A100 GPU, whereas the training of the network took a little more than 3 h; these are representative values for all examples reported in this work.

VII. RESULTS

A. Unitary evolution

The first application we consider is to predict the HCE of states obtained via unitary time evolution of the initial state $|\psi_+\rangle$ with parameters $J = h = 1$. We set the system size to $N = 10$ and train the network on $N_S = 100$ quantum states that were taken at different points of the time interval $ht \in [0, 5]$. The batch size and sample size were chosen to be $N_B = 50$ and $N_M = 1000$. We validate the network during training on an independent dataset with the same parameters. In Fig. 2 the network's predictions and the evolution of the loss during training are shown. It is apparent that the network is capable of correctly estimating the HCE for a rather small dataset size. The HCE is shown as a function of time for convenience, but we want to emphasize that the network has information

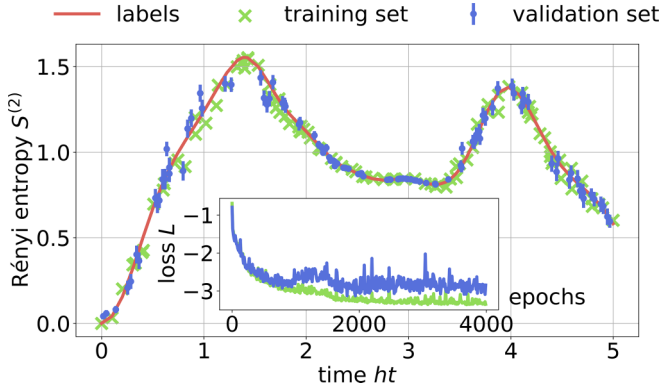


FIG. 2. Training the network for $N = 10$. Network performance on the training (green crosses) and validation (blue circles) dataset, compared to training labels (red line). After 4000 epochs of training, the parameters with minimal loss on the validation dataset are selected. Inset: Loss on training (green line, lower) and validation (blue line, upper) datasets.

on neither the particular quantum state nor the time at which the samples have been extracted from the unitary evolution. Predictions are solely based on POVM measurement results.

B. Extrapolation and baseline comparison

We have demonstrated that the network succeeds at an interpolation task, when validated on in-distribution (ID) data. We now aim to investigate the network's ability to estimate entanglement entropy on OOD data. To this end we have trained an ensemble of eight networks using the same hyperparameters and training dataset as in the previous case, but different initializations, and evaluated their performance on quantum states on the larger time interval $ht \in [0, 10]$ as shown in Fig. 3. The predictions of the network seem sensible even beyond the training region up to $ht \approx 7$. In the regime $ht \in [7, 8.5]$, however, the error bars are severely underestimated.

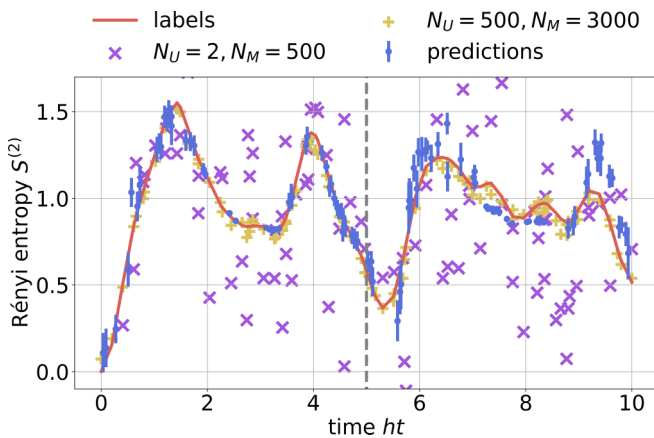


FIG. 3. Performance of the network (blue circles) trained only on the left-hand side of the time interval (dashed gray line). Violet crosses: Baseline method for $N_U = 2$ and $N_M = 500$. Yellow pluses: Baseline method for $N_U = 300$ and $N_M = 5000$.

In order to obtain a better estimate of the performance enhancement that our method gives, we here compare to the method proposed in Ref. [29], which we use as a baseline. It is based on N_M projective measurements, taken with respect to N_U randomly chosen orthonormal bases, such that a total of $N_U N_M$ measurements have to be performed. From the measurement statistics, the subsystem purity $\text{tr}(\rho_A^2)$ and thus the HCE of the system may be inferred as shown in Fig. 3 for two different choices of hyperparameters. In the first case the predictions are based on a total of 1000 measurements per state, using the same number of samples as the network. In the second case, the total amount of measurements is 1.5×10^6 . One can see that the network requires orders of magnitude fewer samples than the baseline to properly predict the HCE, if it is evaluated on ID data. For OOD data it still outperforms the baseline, if evaluated for 1000 measurements, but has less predictive power compared to the baseline at 1.5×10^6 samples, which is, however, a prohibitively large number of samples for current quantum simulation experiments.

C. Dissipative dynamics

If applied to noisy experimental data, robustness with respect to the experimental control and noise parameters is required. To demonstrate the usefulness of the proposed method in an actual experimental application, we now investigate the ability of the network to estimate the mutual information between the subsystems A and B in a dissipative setting. We train the network on data obtained by modeling the Lindblad equation (3) with the dephasing strength γ_z and the decay strength γ_- , considering a spin chain of length $N = 8$ in the transverse-field Ising model with the initial state $\rho = |\psi_+\rangle\langle\psi_+|$. Instead of benchmarking on different evolution times, we instead limit ourselves to a fixed time $t = t_* = 0.75$ and vary the noise strengths γ_z and γ_- .

We train the network on an equidistant 5×5 grid of noise strengths $(\gamma_z, \gamma_-) \in [0, 0.5]^2$ with $N_B = 50$ batches of $N_M = 1000$ samples per noise configuration. During training we validate the network on a dataset sampled from quantum states with 20 randomly distributed noise configurations and 40 noise configurations which lie equidistantly spaced on a randomly chosen cross section in the (γ_z, γ_-) plane [Fig. 4(a), blue]. After training we selected the parameters which minimized the loss on the validation dataset and evaluated the performance of the network on a test set [Figs. 4(b) and 4(c)].

The network's predictions are overall close to the labels. We can therefore deduce that the network was indeed successful in learning to estimate the MI for any noise strength. This underlines the applicability of the proposed method to experimental situations in which one often has a rough understanding of the underlying noise model, but no precise knowledge of the noise strengths.

VIII. DISCUSSION

We have demonstrated the possibility of learning maps from informationally complete POVM measurement data to quantum entropies using deep neural networks, allowing us to significantly reduce the sample complexity compared to the baseline procedure [29]. While the estimator in Ref. [29] is asymptotically unbiased, the approach that is followed here




FIG. 4. (a) Training and validation data points in the (γ_z, γ_-) plane. Green diamonds: Training set. Blue circles: Validation set, consisting of 20 randomly distributed data points and 40 data points along a randomly chosen cross section (lower line). Red crosses: Test set, consisting of 30 randomly distributed data points and 30 validation data points along a randomly chosen cross section (lower line). (b) Performance of network (blue circles) on randomly scattered data points (red crosses). (c) Performance of the network (blue circles) on a cross section (red line) of 30 data points.

is not. Instead, we rely on training a model on labeled data and are therefore restricted to regimes that are classically simulatable. However, this still expands on the system-size regimes that one may probe using the method proposed in Ref. [29] since the regimes that allow for classical simulations comprise those for which one is able to generate sufficiently many samples for Ref. [29] to give satisfactory results.

We tested the proposed method for both ID and OOD data, observing good performance for interpolation tasks. In extrapolation tasks, however, we find regimes in which all the networks contained in the ensemble predict similar entropies, resulting in uncertainties that are too small. However, since uncertainty quantification is a highly active research field, one can expect further advances in the estimation of epistemic uncertainties, which will be directly applicable to the method proposed here. In this context, an interesting future direction is to further explore the extrapolation capabilities, particularly in system size. Here, we wish to also point out that the proposed method is not limited to the estimation of entropies and mutual information, but to any quantity and system for which both training data and labels can be generated synthetically.

The full implementation of the network architecture is available through GitHub [47].

ACKNOWLEDGMENTS

This work is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy EXC2181/1-390900948 (the Heidelberg STRUCTURES Excellence Cluster) and within the Collaborative Research Center SFB1225 (ISOQUANT). This work was partially financed by the Baden-Württemberg Stiftung gGmbH. The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through Grant No. INST 40/575-1 FUGG (JUSTUS 2 cluster). The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. [48] for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS [49] at the Jülich Supercomputing Centre (JSC).

APPENDIX: DETAILS OF NETWORK ARCHITECTURE

The input of the neural network contains N_M POVM samples, with shape $(N, 4)$ as they consist of N single-qubit measurement results with four possible outcomes [one-hot encoding, also see Eq. (11)]. As explained in the architecture section of the main text, the network architecture is divided into three parts. At the first instance each sample is transformed by a layer of LSTM cells that iterates over the single-qubit outcome, thereby mapping each sample into a latent space of dimension F . The embedded sample set therefore has dimension (N_M, F) . The embedded sample set is then treated by the GAT as a fully connected graph. During the development stages of the project we observed that a connected graph resulted in better accuracies compared to an unconnected graph, albeit incurring a much higher computational cost. In each layer of the GAT, the nodes are updated using self-attention [50]. We choose to use two layers and a feature vector of dimension F' , which corresponds to a transformation $(N_M, F) \mapsto (N_M, F')$ of the input. In order to obtain a permutation-invariant quantity from the equivariant output of the GAT we sum over all N_M nodes and receive a single invariant feature vector of dimension F' that encapsulates all information of the sample set. Finally, this feature vector is fed into a DFNN with two output neurons which correspond to the prediction of the HCE and its aleatoric error. As an optimizer for the parameters we have chosen ADAM [51]. The choice of network hyperparameters is given in Table I.

TABLE I. Network hyperparameters.

Hyperparameter	Value
RNN layers	3
RNN features (F)	20, 20, 20
GAT layers	2
GAT features (F')	10, 10
DFNN layers	2
DFNN features	4, 2
Learning rate	0.0005

- [1] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Rev. Mod. Phys.* **81**, 865 (2009).
- [2] P. Calabrese and J. Cardy, *J. Stat. Mech.* (2005) P04010.
- [3] D. Basko, I. Aleiner, and B. Altshuler, *Ann. Phys. (Amsterdam, Neth.)* **321**, 1126 (2006).
- [4] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [5] O. Gühne and G. Tóth, *Phys. Rep.* **474**, 1 (2009).
- [6] B. M. Terhal, *Phys. Lett. A* **271**, 319 (2000).
- [7] J. Estève, C. Gross, A. Weller, S. Giovanazzi, and M. K. Oberthaler, *Nature (London)* **455**, 1216 (2008).
- [8] A. Mair, A. Vaziri, G. Weihs, and A. Zeilinger, *Nature (London)* **412**, 313 (2001).
- [9] H.-N. Dai, B. Yang, A. Reingruber, X.-F. Xu, X. Jiang, Y.-A. Chen, Z.-S. Yuan, and J.-W. Pan, *Nat. Phys.* **12**, 783 (2016).
- [10] T. J. Osborne and M. A. Nielsen, *Phys. Rev. A* **66**, 032110 (2002).
- [11] R. Blume-Kohout, *New J. Phys.* **12**, 043034 (2010).
- [12] A. I. Lvovsky, *J. Opt. B: Quantum Semiclassical Opt.* **6**, S556 (2004).
- [13] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, *Nat. Commun.* **1**, 149 (2010).
- [14] T. Baumgratz, D. Gross, M. Cramer, and M. B. Plenio, *Phys. Rev. Lett.* **111**, 020401 (2013).
- [15] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, A. J. Daley, M. Cramer, M. B. Plenio, R. Blatt, and C. F. Roos, *Nat. Phys.* **13**, 1158 (2017).
- [16] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, *Phys. Rev. Lett.* **105**, 150401 (2010).
- [17] C. Schwemmer, G. Tóth, A. Niggebaum, T. Moroder, D. Gross, O. Gühne, and H. Weinfurter, *Phys. Rev. Lett.* **113**, 040503 (2014).
- [18] C. A. Riofrío, D. Gross, S. T. Flammia, T. Monz, D. Nigg, R. Blatt, and J. Eisert, *Nat. Commun.* **8**, 15305 (2017).
- [19] G. Tóth, W. Wieczorek, D. Gross, R. Krischek, C. Schwemmer, and H. Weinfurter, *Phys. Rev. Lett.* **105**, 250403 (2010).
- [20] T. Moroder, P. Hyllus, G. Tóth, C. Schwemmer, A. Niggebaum, S. Gaile, O. Gühne, and H. Weinfurter, *New J. Phys.* **14**, 105001 (2012).
- [21] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Nat. Phys.* **14**, 447 (2018).
- [22] P. Cha, P. Ginsparg, F. Wu, J. Carrasquilla, P. L. McMahon, and E.-A. Kim, *Mach. Learn.: Sci. Technol.* **3**, 01LT01 (2022).
- [23] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, *Nat. Mach. Intell.* **1**, 155 (2019).
- [24] T. V. Zache, C. Kokail, B. Sundar, and P. Zoller, *Quantum* **6**, 702 (2022).
- [25] M. K. Joshi, C. Kokail, R. van Bijnen, F. Kranzl, T. V. Zache, R. Blatt, C. F. Roos, and P. Zoller, *Nature* **624**, 539 (2023).
- [26] C. Kokail, R. van Bijnen, A. Elben, B. Vermersch, and P. Zoller, *Nat. Phys.* **17**, 936 (2021).
- [27] A. Anshu, S. Arunachalam, T. Kuwahara, and M. Soleimanifar, *Nat. Phys.* **17**, 931 (2021).
- [28] H.-Y. Huang, R. Kueng, and J. Preskill, *Nature (London)* **16**, (2020).
- [29] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, *Science* **364**, 260 (2019).
- [30] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Židek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back *et al.*, *Nature (London)* **596**, 583 (2021).
- [31] F. Møller, T. Schweigler, M. Tajik, J. Sabino, F. Cataldini, S.-C. Ji, and J. Schmiedmayer, *Phys. Rev. A* **104**, 043305 (2021).
- [32] E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber, *Nat. Phys.* **13**, 435 (2017).
- [33] J. Carrasquilla and R. G. Melko, *Nat. Phys.* **13**, 431 (2017).
- [34] D. Koutný, L. Ginés, M. Moczala-Dusanowska, S. Höfling, C. Schneider, A. Predojević, and M. Ježek, *Sci. Adv.* **9**, eadd7131 (2023).
- [35] S. Lu, S. Huang, K. Li, J. Li, J. Chen, D. Lu, Z. Ji, Y. Shen, D. Zhou, and B. Zeng, *Phys. Rev. A* **98**, 012315 (2018).
- [36] C. Harney, S. Pirandola, A. Ferraro, and M. Paternostro, *New J. Phys.* **22**, 045001 (2020).
- [37] C. Harney, M. Paternostro, and S. Pirandola, *New J. Phys.* **23**, 063033 (2021).
- [38] N. Asif, U. Khalid, A. Khan, T. Q. Duong, and H. Shin, *Sci. Rep.* **13**, (2023).
- [39] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University, Cambridge, England, 2010).
- [40] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, [arXiv:2104.13478](https://arxiv.org/abs/2104.13478).
- [41] U. Schollwöck, *Ann. Phys. (Amsterdam, Neth.)* **326**, 96 (2011).
- [42] R. Orús, *Ann. Phys. (Amsterdam, Neth.)* **349**, 117 (2014).
- [43] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
- [44] M. B. Hastings, I. González, A. B. Kallin, and R. G. Melko, *Phys. Rev. Lett.* **104**, 157201 (2010).
- [45] B. Lakshminarayanan, A. Pritzel, and C. Blundell, [arXiv:1612.01474](https://arxiv.org/abs/1612.01474).
- [46] E. Hüllermeier and W. Waegeman, *Mach. Learn.* **110**, 457 (2021).
- [47] GitHub at <https://github.com/MaximilianRie/EntropyEstimation>.
- [48] www.gauss-centre.eu.
- [49] D. Krause, Forschungszentrum Jülich, Jülich Supercomputing Centre, *JLSRF* **5**, A171 (2019).
- [50] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, in *Proceedings of the International Conference on Learning Representations (ICLR)*, (2018).
- [51] D. P. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).