

# Quantum Equilibrium Propagation for efficient training of quantum systems based on Onsager reciprocity

Clara C. Wanjura

Max Planck Institute for the Science of Light, Staudtstraße 2, 91058 Erlangen, Germany

Florian Marquardt

Max Planck Institute for the Science of Light, Staudtstraße 2, 91058 Erlangen, Germany and  
Department of Physics, University of Erlangen-Nuremberg, 91058 Erlangen, Germany

(Dated: June 11, 2024)

The widespread adoption of machine learning and artificial intelligence in all branches of science and technology has created a need for energy-efficient, alternative hardware platforms. While such neuromorphic approaches have been proposed and realised for a wide range of platforms, physically extracting the gradients required for training remains challenging as generic approaches only exist in certain cases. Equilibrium propagation (EP) is such a procedure that has been introduced and applied to classical energy-based models which relax to an equilibrium. Here, we show a direct connection between EP and Onsager reciprocity and exploit this to derive a quantum version of EP. This can be used to optimize loss functions that depend on the expectation values of observables of an arbitrary quantum system. Specifically, we illustrate this new concept with supervised and unsupervised learning examples in which the input or the solvable task is of quantum mechanical nature, e.g., the recognition of quantum many-body ground states, quantum phase exploration, sensing and phase boundary exploration. We propose that in the future quantum EP may be used to solve tasks such as quantum phase discovery with a quantum simulator even for Hamiltonians which are numerically hard to simulate or even partially unknown. Our scheme is relevant for a variety of quantum simulation platforms such as ion chains, superconducting qubit arrays, neutral atom Rydberg tweezer arrays and strongly interacting atoms in optical lattices.

## I. INTRODUCTION

As deep learning and artificial intelligence are adopted in all branches of science and technology, the increasing complexity of neural networks has led to an exponential increase in energy consumption and training costs. This has created a need for more efficient alternatives, sparking the rapidly developing field of neuromorphic computing [1], which explores a variety of different platforms [2, 3] to design physical, analogue neural networks.

Existing backpropagation-based training strategies for neuromorphic platforms include in-silico training, requiring a faithful digital model of the system, and physics-aware backpropagation [4], combining physical inference with a simulated backward pass which relaxes these constraints. However, it is a central question whether not only inference but also training can exploit the physical dynamics [5], making full use of the energy efficiency of neuromorphic systems. For example, feedback-based parameter shifting does not require any simulation but scales unfavourably with the network size [6]. Moving towards physical implementations of efficient backpropagation, strategies for specific types of non-linearities have been developed [7–9], as well as approaches performing backpropagation only on the linear components [10, 11]. Another novel recent approach enables “forward-forward” type gradient calculation in systems which perform sequential information processing [12]. Furthermore, efficient measurements of gradients via scattering experiments can be performed in optical systems that employ a framework recently developed to produce nonlinear computation with linear wave setups [13]— such nonlinear processing was also recently demonstrated in Refs. [12, 14, 15].




FIG. 1. **The concept of quantum equilibrium propagation.** The goal is to efficiently and in a physical way obtain the gradient of some loss function (depending on expectation values measured at the “output” degrees of freedom of a quantum system) with respect to tuneable parameters. **a** Rather than shifting  $N$  parameters separately and measuring the output expectation value for each shift (parameter-shift method), Onsager reciprocity dictates that the same information can be extracted by **b** shifting, i.e. nudging, only the parameters coupling to the output observables and (in a single go) measuring the response of all  $N$  operators coupled to the training parameters (quantum equilibrium propagation). The latter procedure is more efficient as it requires only a single response experiment (or at most a small number of order 1, when some non-commuting observables have to be measured), whereas the parameter shift method requires a number of experiments scaling linearly with the number of parameters.

General approaches for physical backpropagation so far only exist in two classes of physical systems: Hamiltonian Echo Backpropagation [16], which applies to essentially lossless systems in which a time-reversal operation can be implemented, and equilibrium propagation (EP) [17, 18], which applies to energy-based, equilibrating systems.

EP stands in the tradition of contrastive learning approaches comparing measurements obtained from two different equilibria and using feedback to update parameters. Concretely,

EP consists of two phases: the free and the nudged phase. In the free phase, the input is fixed and the system relaxes into its equilibrium state. In the nudged phase, the output is “nudged” closer towards the target output and the system evolves to a new equilibrium. Comparing the system state in the free and the nudged phase, respectively, one obtains the necessary gradients which are then used to update parameters.

Since its introduction in 2017, EP has been investigated thoroughly [19–24] and a variant, coupled learning, has been developed [25, 26]. In particular, EP was proposed for training nonlinear resistor circuits [19], systems of coupled phase oscillators [24], was further adapted to spiking networks [20], to implement continual parameter updates [21] and a dynamical version was developed [22]. Experimentally, EP has been applied to train electronic systems [27, 28], elastic networks [29] and even a memristor crossbar array [30]. Furthermore, a classical Ising model has been trained using a quantum annealer to efficiently reach equilibrium [31].

Given the elegance of the EP approach, it is a natural question to ask whether the procedure can be extended to quantum systems to train a fully quantum Hamiltonian via a nudging procedure similar to classical EP. Indeed, we will show that there is a direct connection between EP and Onsager reciprocity and exploit it to derive a quantum version of EP (QEP). This new approach can be used to train efficiently arbitrary quantum systems, including specifically the highly tuneable quantum many-body systems realized nowadays in quantum simulators [32].

We illustrate this new concept with supervised and unsupervised learning examples of a quantum-mechanical nature. Specifically, we investigate the recognition of quantum many-body phases and introduce as new concepts the exploration of phase diagrams with quantum simulator platforms via efficient gradient descent optimization enabled by QEP as well as the optimisation of sensitivity (e.g. for sensing applications) and phase boundary exploration. QEP is applicable to systems which are hard or impossible to simulate classically and can be employed even in settings in which the Hamiltonian is only partially known or partially accessible.

In terms of the general question of using quantum devices for learning tasks, the area of quantum machine learning [33, 34] by now has a long history. There have been some ideas of how to learn to reproduce quantum states by adapting tuneable parameters (“quantum Boltzmann machines”, see [34]). The major research efforts in this domain are, however, spent on variational quantum circuits, which require digital quantum computing platforms for implementation (possibly even fault-tolerant), going significantly beyond the resources that we are going to assume here.

## II. QUANTUM EQUILIBRIUM PROPAGATION

### A. Onsager reciprocity

Consider a parameterized Hamiltonian

$$\hat{H}(\lambda) = \sum_j \lambda_j \hat{A}_j \quad (1)$$

and its ground state  $|\Psi(\lambda)\rangle$ . A small static force coupling to  $\hat{A}_j$  (entering as a term  $\delta\lambda_j \hat{A}_j$  inside  $\hat{H}$ ) will produce a linear response in the expectation value  $\langle \hat{A}_\ell \rangle$ , given by

$$\chi_{\ell j} = \frac{\partial}{\partial \lambda_j} \langle \Psi(\lambda) | \hat{A}_\ell | \Psi(\lambda) \rangle. \quad (2)$$

Onsager reciprocity guarantees the symmetry of the susceptibility  $\chi_{j\ell} = \chi_{\ell j}$ , i.e. the same effect will be produced by a force acting on  $\hat{A}_\ell$  influencing the expectation value  $\langle \hat{A}_j \rangle$ , see Fig. 1, i.e.,

$$\frac{\partial}{\partial \lambda_j} \langle \hat{A}_\ell \rangle = \frac{\partial}{\partial \lambda_\ell} \langle \hat{A}_j \rangle. \quad (3)$$

Onsager reciprocity [35] can be derived in many ways, also for the quantum case [36]. However, the most elementary approach for static situations such as the one considered here and applied to the ground state in particular uses first-order perturbation theory for the deformation of the ground state,  $\partial_{\lambda_j} |\Psi(\lambda)\rangle = (E(\lambda) - H(\lambda))^{-1} (A_j - \langle A_j \rangle) |\Psi(\lambda)\rangle$ . Expression (3) also holds for thermal states for which the expectation values above are replaced by  $\langle \hat{A}_j \rangle = \text{Tr}(\hat{\rho} \hat{A}_j)$ , so the following results apply for arbitrary-temperature quantum equilibrium states. In the case of ground-state degeneracy, expression (3) also holds with  $\hat{\rho} = \sum_j |\Psi_j(\lambda)\rangle \langle \Psi_j(\lambda)|$  in which  $j$  sums over all degenerate states; this is approximately equivalent to a thermal state at small, but non-vanishing, temperature. Classical Onsager reciprocity is equivalent to what has been termed the “Fundamental Lemma” for classical EP [17, 18].

### B. Onsager reciprocity as a basis for quantum equilibrium propagation

We will now show that this well-known result (3) gives us access to a general version of equilibrium propagation for quantum systems. To understand that, we first consider supervised learning. Let us assume that the set of operators  $\hat{A}_j$  is split into degrees of freedom relating to the input  $j \in \mathcal{S}_{\text{in}}$ , trainable variables  $j \in \mathcal{S}_{\text{train}}$ , and the output  $j \in \mathcal{S}_{\text{out}}$ . Accordingly, the set of parameters  $\lambda$  is split into the input  $x$  containing the parameters corresponding to the input  $\lambda_j = x_j$ , the set of training parameters  $\theta$  with  $\lambda_j = \theta_j$  and the set of couplings to the output observables  $\nu$  with  $\lambda_j = \nu_j$ . Hence, a general QEP Hamiltonian is of the form  $\hat{H}(x, \theta, \nu)$  in which  $\nu = 0$  during inference. For any given training sample, the

input  $x$  is fixed by applying a field to all input degrees of freedom (we may write  $\lambda_k = x_k$  for  $k \in \mathcal{S}_{\text{in}}$ , with  $x$  representing the input vector). The output is then read off as the expectation values  $y_\ell = \langle \hat{A}_\ell \rangle$  in the operators  $\ell \in \mathcal{S}_{\text{out}}$ . Note that  $\hat{A}_\ell$  can be chosen as a projector, in which case  $y_\ell$  becomes the probability of obtaining a particular outcome in a measurement; this is useful for classification tasks.

In supervised learning, we are interested in adjusting the trainable parameters  $\theta$  in order to “nudge” the output closer to the desired target output  $y^{\text{target}}(x)$ , for any given input  $x$ . More generally, we aim to reduce the loss function  $\mathcal{L}(y, y^{\text{target}})$ , or rather its average  $\bar{\mathcal{L}}$  over many training samples  $(x, y^{\text{target}}(x))$ , via gradient descent:  $\delta\theta = -\eta\partial\bar{\mathcal{L}}/\partial\theta_j$ . To do this, we need to obtain the influence of a change in  $\theta_j$  on any of the outputs  $y_\ell$ . For a given fixed training sample, this is just the susceptibility  $\chi_{\ell j}(\lambda) = \partial\langle\hat{A}_\ell\rangle/\partial\theta_j$ . Evaluating  $\chi_{\ell j}$  for all possible trainable parameters  $j$ , see Fig. 1 a, scales unfavourably, requiring a number of different experiments that scales linearly in the number  $N$  of these parameters. Accessing the training gradient in this way amounts to the parameter-shift method which is always applicable in any neuromorphic platform but should generally be avoided whenever possible due to this unfavourable scaling. However, Onsager reciprocity, Eq. (3), tells us that we can also access the susceptibility  $\chi_{\ell j}$  by performing an alternative, much more efficient experiment that instead reveals  $\chi_{j\ell}$ : apply a small force  $\nu$  acting on the outputs and observe its influence on the expectation values of the degrees of freedom  $\langle\hat{A}_k\rangle$  connected to the trainable parameters:  $\chi_{j\ell}(\lambda) = \partial\langle\hat{A}_j\rangle/\partial\nu_\ell$ , see Fig. 1 b. This would seem to require a number of experiments that scales with the number of outputs  $N_{\text{out}}$ , already typically much smaller than the number of trainable parameters. However, by evaluating explicitly the desired gradient of the loss function, it becomes apparent that only a single experiment is in fact needed: in this experiment, a force vector  $\nu = \partial\mathcal{L}/\partial y$ , the so-called error signal, is applied to the output degrees of freedom.

In this way, Onsager reciprocity teaches us how to translate the classical equilibrium propagation approach to quantum devices. Quantum Hamiltonians of arbitrary structure can be considered.

### C. QEP procedure

We now explicitly summarize the QEP procedure. For clarity, we will from now on denote the output observables by  $\hat{O}_\ell$ . For a gradient-descent parameter update, we need to compute the derivative of the loss function,

$$\frac{\partial}{\partial\theta_j}\mathcal{L}(y, y^{\text{target}}(x)) = \sum_\ell \frac{\partial\mathcal{L}}{\partial y_\ell} \frac{\partial y_\ell}{\partial\theta_j} = \varepsilon \frac{\partial y}{\partial\theta_j} \quad (4)$$

in which the error signal vector has components  $\varepsilon_\ell = \partial\mathcal{L}/\partial y_\ell$ . For a mean-square-error loss function, we would have  $\varepsilon = 2(y - y^{\text{target}}(x))$ .

The QEP procedure for supervised learning can be summarized as follows. (i) Free phase: The nudging forces are off ( $\nu_j = 0$  for all  $j$ ) and the output expectation values  $y_\ell = \langle\hat{O}_\ell\rangle$  as well as the expectation values of all operators associated with trainable parameters  $\hat{A}_j$  are measured in the ground state of the Hamiltonian  $\hat{H}(x, \theta, 0)$ . (ii) Nudged phase: We compute the error signal  $\varepsilon$  and use it to nudge the Hamiltonian  $\hat{H}(x, \theta, \nu = \beta\varepsilon)$  by switching on the couplings to the output observables, adding a term  $\sum_\ell \nu_\ell \hat{A}_\ell$  to the Hamiltonian. The couplings are given by the vector  $\nu = \beta\varepsilon$ , in which  $\beta$  is a small parameter (keeping with the notation for classical EP [17]; this is unrelated to the inverse temperature). We again measure the expectation values of all observables  $\hat{A}_j$ . (iii) Parameter update: Using Onsager reciprocity, Eq. (3),  $\partial\langle\hat{O}_\ell\rangle/\partial\theta_j = \partial\langle\hat{A}_j\rangle/\partial\nu_\ell$ , we can approximate the gradient  $\partial\langle\hat{O}_\ell\rangle/\partial\theta_j$  and hence arrive at:

$$\frac{\partial}{\partial\theta_j}\mathcal{L}(y, y^{\text{target}}(x)) \approx \frac{\langle\hat{A}_j\rangle|_{\nu=\beta\varepsilon} - \langle\hat{A}_j\rangle|_{\nu=0}}{\beta}. \quad (5)$$

In a similar spirit as for classical Equilibrium Propagation, one may consider variants which combine positive and negative nudging [37], i.e., approximate the gradient using  $(\langle\hat{A}_j\rangle|_{\nu=\beta\varepsilon} - \langle\hat{A}_j\rangle|_{\nu=-\beta\varepsilon})/(2\beta)$ , which empirically performs better for finite nudging.

### D. Practical requirements

We now discuss the most important practical considerations for implementing quantum equilibrium propagation (QEP) in any experimental platform.

Above all, the platform needs to have tuneable couplings  $\theta$ , whose number preferably should be easy to scale up with growing system size. Such tuneable couplings have been developed for many quantum simulators [32] and quantum computing platforms by now. Examples include: (i) ion chains, where spin-spin couplings can be mediated and engineered via the vibrational modes of the chain, employing suitable Raman transitions., e.g. [38]. (ii) superconducting-qubit arrays, as employed for quantum computing, with current-tuneable couplers between neighboring qubits [39]. (iii) neutral-atom Rydberg tweezer arrays providing tuneable spin-spin couplings [40]. (iv) strongly interacting atoms in optical lattices with spatially engineered hopping and interactions, e.g. based on holographic potential shaping [41]. Other platforms, e.g. in optomechanical arrays, coupled microwave cavities, or coupled laser arrays, also demonstrate interesting tuneable coupling schemes, but they often operate out of equilibrium and are therefore not directly suitable for QEP, unless one can map them back to an equilibrium situation.

Beyond this primary requirement of a scalable number of tuneable couplings, QEP platforms also need ways to apply the output forces  $\nu$ . This demands local fields, e.g. effective magnetic fields or qubit detunings, easily available in most platforms that are flexible enough to support tuneable couplings. In addition, the expectation values of both the output

operators and of the coupling operators connected to trainable degrees of freedom should be measurable. Regarding the couplings, we note that the expectation values of interaction terms of the form  $\hat{X}_j \hat{X}_k$  or similar can easily be measured even by observing the spin operators  $\hat{X}_j$  individually (and multiplying outcomes). A Heisenberg-type coupling operator  $\sum_{\alpha=x,y,z} \hat{\alpha}_j \hat{\alpha}_k$  would need three separate measurements, for the x,y,z components, performed in independent shots of the experiment, eventually obtaining the expectation value of the complete operator. Alternatively, one could carry out a collective (two-qubit) measurement. The latter is typically performed via an ancilla, as demonstrated for syndrome extraction in quantum error correction schemes, and therefore requires more experimental effort. We note that the statistical nature of quantum physics in any case requires many runs of the experiment to measure the expectation values of output variables (needed for inference) and of coupling operators (needed for training). Each of these runs involves an equilibration step.

Finally, QEP requires efficient means to approach the equilibrium state, i.e., for the zero-temperature limit, the ground state  $|\Psi(\lambda)\rangle$  of the Hamiltonian. Efficient ground-state preparation of complex quantum many-body Hamiltonians is one of the most intensively researched questions in quantum simulation and quantum computing. For the purposes of QEP, we distinguish between two options: (i) hybrid approaches, where an external digital computer is employed during equilibration, and (ii) purely autonomous schemes. Hybrid approaches could rely on variational quantum eigensolvers [42], where an ansatz quantum circuit with continuously parametrized unitary gates is performed, the expectation value of the Hamiltonian is measured, and a classical optimization is performed to adapt parameters of the circuit. At first glance, the use of a classical optimizer to find the quantum ground state in this way might seem to contradict the basic premise of QEP or neuromorphic computing in general, i.e. using a physical system to do information processing. However, if the problem setting makes efficient use of the resulting quantum many-body ground state, even such a hybrid approach may still yield an advantage over an entirely classical digital device, in the same way that variational quantum eigensolvers are thought to be beneficial under the right circumstances as compared to numerical ground state search by classical algorithms.

Purely autonomous equilibration schemes get rid of any feedback loop. In principle, coupling to a cold environment is sufficient, but recently there has been active research into speeding up equilibration. The techniques put forward often rely in one way or another on variations of quantum reservoir engineering [43], where dissipation is introduced deliberately, and which have, e.g., been used to stabilize quantum many-body states in superconducting circuits [44]. This can happen in the form of suitable continuous driving (inspired by laser cooling), e.g. [45], or else in the form of quantum circuits that introduce gates coupling the quantum many-body system to ancilla qubits which then may be periodically reset (“digital quantum cooling”, e.g. [46]) or other schemes to reduce entropy, such as via suitable measurements [47].

### III. APPLICATIONS

#### A. Supervised learning: phase detection in a quantum many-body system

The input considered above is classical, where QEP could be used to train a quantum device to perform an essentially classical machine learning task. However, we now turn to an important class of applications where a QEP-trained system can effectively receive input that is quantum instead of classical.

In general, this setting can be realized by starting from a Hamiltonian  $\hat{H}_0(x)$ , whose quantum ground state we want to analyze with the help of a QEP setup. By tuning the classical parameters  $x$  (typically a few) we are able to realize different phases with different ground states, e.g., sweeping through some phase diagram in the case of a quantum many-body system. One task could consist in predicting, for any given  $x$ , the distinct quantum phase that the system assumes, possibly after seeing a few labeled training examples at a few parameter locations  $x$ . Another task could consist in approximating the entanglement entropy between some subsystems of  $\hat{H}_0$  or predicting any other quantity of interest that can be derived in principle by inspecting the ground state but may be hard to extract directly by simple measurements.

Any of these tasks can be addressed via QEP in the following way. We couple the system of interest, described by  $\hat{H}_0(x)$ , to a trainable physical sensor, described by  $\hat{H}_{\text{sens}}(\theta)$ , Fig. 2 a. Overall, coming back to our previous definitions, we thus have the full QEP Hamiltonian  $\hat{H}(x, \theta) = \hat{H}_0(x) + \hat{V}(\theta) + \hat{H}_{\text{sens}}(\theta)$ , where we assume that the couplings between the two systems reside inside  $\hat{V}$  (see Fig. 2).

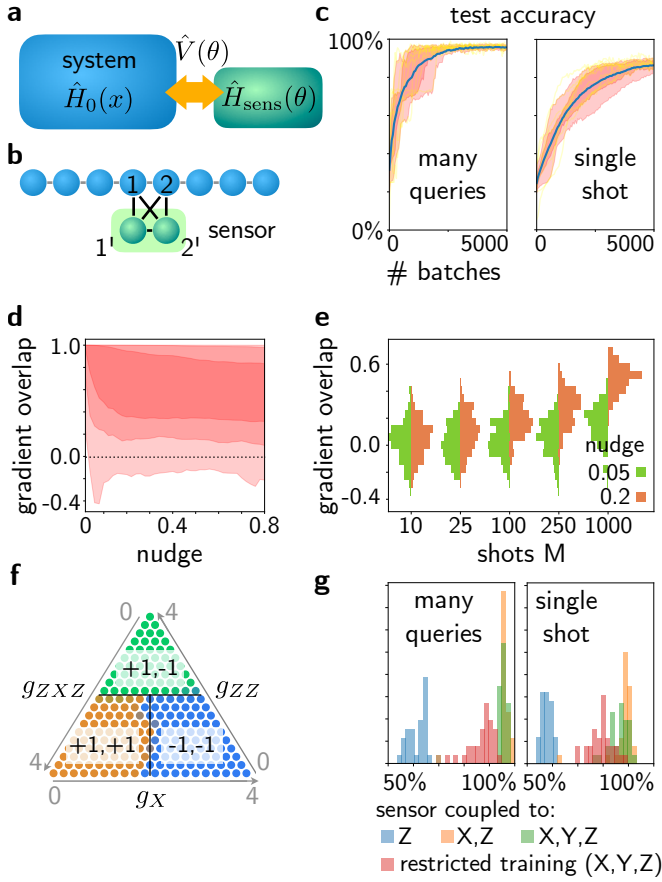
Ideally, the couplings inside the system of interest,  $\hat{H}_0$ , should be stronger than the couplings to the sensor and within that sensor. This will ensure that the system of interest is only weakly perturbed, while the recognition model can still react strongly to the features of the ground state  $|\psi_0\rangle$  of  $\hat{H}_0$ . However, in our numerical experiments (below) we have seen that even outside this scenario successful learning is possible.

Distinguishing quantum phases can serve as an important application, and it has been considered before in the context of quantum machine learning based on gate-based quantum computers, where a unitary circuit acts on a given ground state encoded in a multi-qubit register [48–50]. In contrast to that, QEP relies on equilibration and moreover is far more general in the choice of systems – e.g. it does not require qubits as degrees of freedom nor the ability to perform gates nor any detailed knowledge of all aspects of the Hamiltonian.

Both for this task as well as others analyzed below, we choose the cluster Ising Hamiltonian (see e.g. [50]) as an illustrative quantum many body system,

$$\hat{H}_0 = g_{\text{ZXZ}} \sum_j \hat{Z}_{j-1} \hat{X}_j \hat{Z}_{j+1} - g_{\text{ZZ}} \sum_j \hat{Z}_j \hat{Z}_{j+1} - g_{\text{X}} \sum_j \hat{X}_j, \quad (6)$$

in which  $\hat{X}, \hat{Y}, \hat{Z}$  are the Pauli matrices. This model has three phases, including a topologically nontrivial one. We



**FIG. 2. Supervised learning: Learning recognition of quantum phases using QEP.** (a) Schematic of a trainable quantum sensor coupled to a system. (b) Specific example of a two-qubit sensor coupled to a 1D cluster transverse Ising Hamiltonian at two locations, where readout of  $\hat{Z}_{1'}$  and  $\hat{Z}_{2'}$  is supposed to indicate the phase. The 51 tuneable couplings are learned using QEP. (c) Evolution of test accuracy during supervised training on the whole phase diagram (for a chain of length  $N = 8$ ). Multiple training runs (yellow), confidence intervals as areas (red; at 50% and 80%), and average accuracy (blue). Accuracies for “many queries” (asking whether the maximum-probability detector outcome matches correct phase) and “single shot” (probability to indicate correct phase in single quantum measurement) [batches of 10 training samples; projection noise for  $M = 10$  measurement shots per sample is accounted for; nudge parameter  $\beta = 0.4$ ]. (d) Overlap of batch-averaged gradient estimate with true gradient direction, vs. nudge parameter. Confidence intervals (red, 95%, 80%, 50%) depict distribution over many batches (batch size 10, no measurement shot noise). (e) Gradient overlap histograms vs. measurement shots  $M$ , for two different nudge parameters (batch size 10). (f) Test of phase recognition: probabilities of measuring the trained sensor in one of the three different combinations of  $Z_1, Z_2$ , shown in orange/green/blue; true phase boundaries in black ( $g_z + g_{xx} + g_{zx} = 4$ ). (g) Histogram of final test accuracies for repeated training runs (parameters as above), for a sensor that only couples to  $\hat{Z}$  operators in the chain (or only to  $\hat{X}$  and  $\hat{Z}$ ), and for a sensor trained only on a small patch in the middle of each phase but tested throughout (“restricted”).

will now regard  $x = (g_{zxz}, g_{zz}, g_x)$  as the input parameters. We add a sensor made of qubits that are coupled in all possible ways among each other (2-local, with terms like  $\hat{X}_\alpha \hat{Y}_\beta$  or  $Z_\alpha$  etc) and to a limited region in the chain (couplings  $\hat{Z}_\alpha \hat{X}_j$  etc, where  $j$  is a spin in the chain). The sensor qubits are measured in the  $Z$ -basis, and the resulting configuration is supposed to announce the detected phase. Here we used a mean-square-error loss function, although a categorical cross-entropy would have been suitable as well. We find that even a small-scale sensor of only two qubits, coupled to two spins in the chain, has sufficient expressivity to properly learn the known phase diagram of the cluster Ising model (Fig. 2 b,c), when trained in a supervised fashion using QEP. Inspecting the solution also reveals a surprise: the sensor-system couplings are not weak, and the spin configuration nearby the sensor is visibly perturbed, without however deteriorating the functionality. We also confirm that the quantum nature of the coupling is important. This can be ascertained by comparing to a sensor that is only allowed to couple to  $\hat{Z}$  operators in the chain, which performs much more poorly (Fig. 2 g). By contrast, allowing coupling to non-commuting observables yields the observed good accuracy.

An important general aspect of QEP training is the unavoidable projection shot noise encountered in any quantum experiment, where expectation values are obtained by repeated measurements with individually discrete outcomes – in contrast to the classical situation. This hampers training, since the estimate of the gradient is noisy, with fluctuations  $\sim 1/\sqrt{M}$ , where each of the  $M$  shots requires a renewed equilibration. Fortunately, we find that this effect can be counteracted by employing an increased finite nudging strength  $\mu$ , effectively boosting the contrast in estimating the response of expectation values. Finite nudging, however, leads to a deviation from the linear response that would yield the ideal gradients. Therefore ultimately there is a sweet spot, balancing shot noise vs. nonlinearity of the response, to obtain optimized training convergence (Fig. 2 d,e). In our numerical experiments, we find that the goal of minimizing the total number of experimental runs (i.e. number of training samples multiplied by number of shots per sample) is best achieved by keeping  $M$  small and simply taking more batches. Apparently, this leads to more variety in the observed training data and better training performance. In experiments, the number of shots  $M$  could also be reduced by coupling multiple sensors at different locations to the system, having them share their coupling values. Finally, in principle, there is an alternative to averaging over projective measurements, namely performing weak continuous measurements of the expectation values (a weakly coupled sensor is able to realize this). This could be performed without rethermalizing to the ground state.

## B. Unsupervised learning

We now discuss two applications for unsupervised learning for which the gradients are used for optimization tasks.

*a. Phase exploration:* In the first example, we would like to explore the phase diagram of a quantum many-body

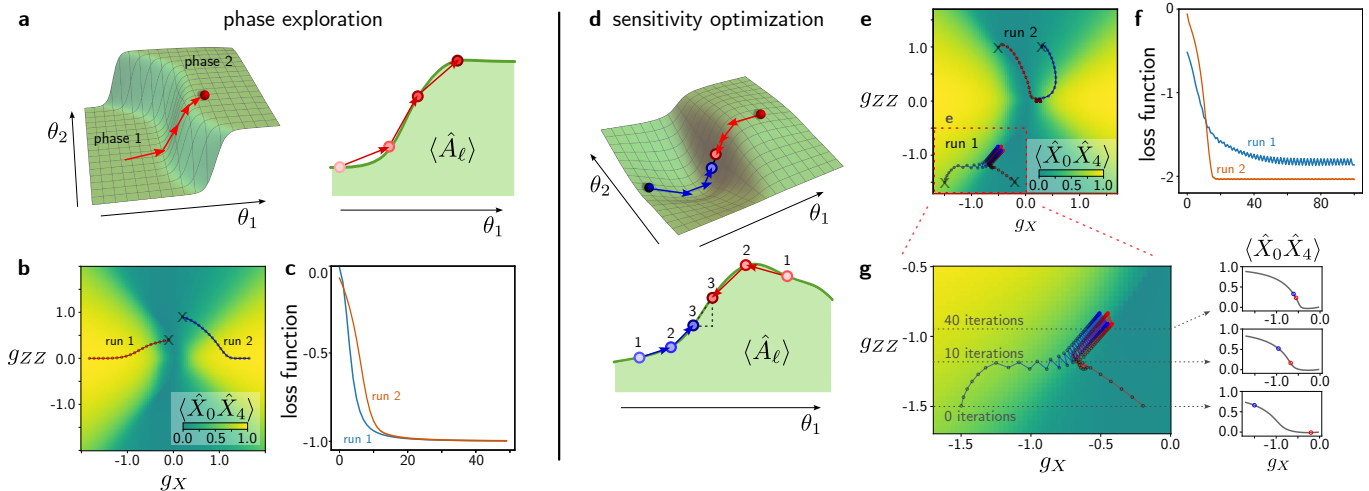


FIG. 3. **Unsupervised learning applications: phase exploration and sensitivity optimization for sensing.** **a** Sketch of phase space exploration using QEP. An expectation value of interest,  $\langle \hat{A}_\ell \rangle$ , is maximised. **b** Two example trajectories in the phase diagram of the cluster Ising Hamiltonian (6) each starting at the position in the phase diagram marked with  $x$  and **c** the corresponding loss functions. The colours in **b** show  $\langle \hat{X}_0 \hat{X}_4 \rangle$ . **d** Sketch of the concept of sensitivity optimization. The derivative of an observable of interest,  $\langle \hat{A}_\ell \rangle$ , w.r.t. a certain parameter (or w.r.t. a vector of parameters) is optimised. Concretely, this scheme can be used to devise optimal sensors or to find phase boundaries in a phase diagram. **e-g** Optimization of  $\partial \langle \hat{X}_0 \hat{X}_4 \rangle / \partial g_X$  (magnetic field sensor). **e** Two example runs, each showing the trajectories of two points in the phase diagram of the cluster Ising Hamiltonian (6) and **f** the corresponding loss functions. The two points are initialized at the position marked with  $\times$  and converge towards the phase boundary where the derivative w.r.t.  $g_X$  is maximal. In the second run, **g** the trajectories first converge towards a phase boundary and then follow it, suggesting that the technique may be employed to trace out phase boundaries. The plots on the right show cuts through the phase diagram at various steps during the training.

system. In practice, this could be an interesting task for characterizing the capabilities of a quantum simulator when we would like to explore the phase diagram of a system which is partially unknown (e.g. some Hamiltonian terms are not known or cannot be tuned) and computationally hard to simulate classically. We can nevertheless ask whether phases exist that maximize (or minimize) certain expectation values and use QEP to explore the phase diagram and find such regimes. In this scenario, the relevant QEP-Hamiltonian of this unsupervised learning task is  $\hat{H}(\theta, \nu)$ , not containing any inputs  $x$ .

Fig. 3 **a** illustrates the procedure: starting from an initial set of parameters  $\theta$  (a point in the phase diagram), we optimize  $\theta$  by computing the gradients with QEP to find the (potentially local) maximum of the expectation value of interest. The only difference in the QEP procedure here is that the derivative of the simple loss function  $\mathcal{L} = -y$  w.r.t. the output variable yields a trivial error signal of  $-1$  which can then be inserted in the gradient calculation above according to Eq. (5).

Here, we exemplify the procedure by examining a slice through the phase diagram of the cluster Ising Hamiltonian (6) when we fix  $g_{ZZ} = -0.5 \equiv \text{const}$ . We optimize the Néel order parameter  $y \equiv \langle \hat{X}_0 \hat{X}_4 \rangle$ , such that the loss function is simply  $\mathcal{L}(y) \equiv -\langle \hat{X}_0 \hat{X}_4 \rangle$ . We show two example trajectories in Fig. 3 **b**. In both cases, the trajectories quickly converge to the  $g_{ZZ} = 0$  line and then move along it, going into the paramagnetic phase. Accordingly, the loss function, Fig. 3 **c**, rapidly decreases, with the decrease becoming slower as the trajectory moves along the  $g_{ZZ} = 0$  line.

We suggest that, in general, this approach could be an ef-

ficient technique for exploring higher dimensional phase diagrams which cannot simply be mapped out without considerable effort by simply sweeping all of the parameters.

**b. Sensitivity optimization:** In the second unsupervised learning application, the aim is to maximize the *derivative* of some expectation value of interest w.r.t. a certain parameter  $\theta_j$ . As Fig. 3 **d** illustrates, in contrast to the previous examples, we now start with two points,  $\theta^{(1)}$  and  $\theta^{(2)}$ , in the phase diagram at which we compute expectation values  $y_{1,2} \equiv \langle \hat{A}_\ell \rangle_{|\theta^{(1,2)}}$  and maximize the slope calculated from the difference quotient of the output expectation values at these two points. Concretely, the corresponding loss function has the form

$$\mathcal{L}(y_1, y_2) = - \left| \frac{\langle \hat{A}_\ell \rangle_{|\theta^{(1)}} - \langle \hat{A}_\ell \rangle_{|\theta^{(2)}}}{\theta_j^{(1)} - \theta_j^{(2)}} \right|. \quad (7)$$

An appealing application of this procedure could be to find the optimal working point of sensors, such as magnetic field sensors.

To illustrate this scheme, we again consider the cluster Ising Hamiltonian (6) and optimize the slope of  $\langle \hat{X}_0 \hat{X}_4 \rangle$  w.r.t. the parameter  $g_X$ , which is proportional to the magnetic field. Fig. 3 **e** shows two sets of example trajectories and the corresponding loss functions, Fig. 3 **f**. We observe that the trajectories converge to the phase boundary, where the slope is largest. For the first run, Fig. 3 **g**, we see that the trajectory moves along the phase boundary, since the slope is larger for smaller  $g_{ZZ}$ . We suggest that, in the future, this feature could be exploited to more generally map out phase boundaries.

#### IV. CONCLUSION AND OUTLOOK

We have exploited Onsager reciprocity to derive a quantum version of equilibrium propagation, which in its classical form is one of the major general training techniques for neuromorphic platforms. We have shown that this can be used successfully even for situations where the input is effectively a quantum state (as in classifying quantum phases via supervised learning), as well as for unsupervised learning tasks re-

lated to exploring the phase diagrams of quantum simulators. In all of these cases, QEP can be applied even when classically, the Hamiltonian is hard or impossible to simulate. A large variety of experimental platforms should be amenable to implementations of quantum equilibrium propagation.

During the final stage of completion of this manuscript and shortly before submission, two strongly related preprints appeared on the arXiv [51, 52], also introducing a quantum version of equilibrium propagation, with somewhat different use cases.

- 
- [1] D. Marković, A. Mizrahi, D. Querlioz, and J. Grollier, Physics for neuromorphic computing, *Nature Reviews Physics* **2**, 499 (2020).
- [2] G. Wetzstein, A. Ozcan, S. Gigan, S. Fan, D. Englund, M. Soljačić, C. Denz, D. A. Miller, and D. Psaltis, Inference in artificial intelligence with deep optics and photonics, *Nature* **588**, 39 (2020).
- [3] B. J. Shastri, A. N. Tait, T. Ferreira de Lima, W. H. Pernice, H. Bhaskaran, C. D. Wright, and P. R. Prucnal, Photonics for artificial intelligence and neuromorphic computing, *Nature Photonics* **15**, 102 (2021).
- [4] L. G. Wright, T. Onodera, M. M. Stein, T. Wang, D. T. Schachter, Z. Hu, and P. L. McMahon, Deep physical neural networks trained with backpropagation, *Nature* **601**, 549 (2022).
- [5] A. Momeni, B. Rahmani, B. Scellier, L. G. Wright, P. L. McMahon, C. C. Wanjura, Y. Li, A. Skalli, N. G. Berloff, T. Onodera, I. Oguz, F. Morichetti, P. del Hougne, M. L. Gallo, A. Sebastian, A. Mirhoseini, C. Zhang, D. Marković, D. Brunner, C. Moser, S. Gigan, F. Marquardt, A. Ozcan, J. Grollier, A. J. Liu, D. Psaltis, A. Alù, and R. Fleury, Training of physical neural networks (2024), [arXiv:2406.03372 \[physics.app-ph\]](https://arxiv.org/abs/2406.03372).
- [6] S. Bartunov, A. Santoro, B. Richards, L. Marris, G. E. Hinton, and T. Lillicrap, Assessing the scalability of biologically-motivated deep learning algorithms and architectures, in *Advances in Neural Information Processing Systems*, Vol. 31, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., 2018).
- [7] D. Psaltis, D. Brady, X.-G. Gu, and S. Lin, Holography in artificial neural networks, *Nature* **343**, 325 (1990).
- [8] X. Guo, T. D. Barrett, Z. M. Wang, and A. Lvovsky, Backpropagation through nonlinear units for the all-optical training of neural networks, *Photonics Research* **9**, B71 (2021).
- [9] J. Spall, X. Guo, and A. I. Lvovsky, Training neural networks with end-to-end optical backpropagation (2023), [arXiv:2308.05226 \[physics.optics\]](https://arxiv.org/abs/2308.05226).
- [10] T. W. Hughes, M. Minkov, Y. Shi, and S. Fan, Training of photonic neural networks through in situ backpropagation and gradient measurement, *Optica* **5**, 864 (2018).
- [11] S. Pai, Z. Sun, T. W. Hughes, T. Park, B. Bartlett, I. A. Williamson, M. Minkov, M. Milanizadeh, N. Abebe, F. Morichetti, *et al.*, Experimentally realized in situ backpropagation for deep learning in photonic neural networks, *Science* **380**, 398 (2023).
- [12] A. Momeni, B. Rahmani, M. Malléjac, P. del Hougne, and R. Fleury, Backpropagation-free training of deep physical neural networks, *Science* **382**, 1297 (2023), <https://www.science.org/doi/pdf/10.1126/science.adi8474>.
- [13] C. C. Wanjura and F. Marquardt, Fully non-linear neuromorphic computing with linear wave scattering (2023), [arXiv:2308.16181 \[physics.optics\]](https://arxiv.org/abs/2308.16181).
- [14] M. Yildirim, N. U. Dinc, I. Oguz, D. Psaltis, and C. Moser, Nonlinear processing with linear optics (2023), [arXiv:2307.08533 \[physics.optics\]](https://arxiv.org/abs/2307.08533).
- [15] F. Xia, K. Kim, Y. Eliezer, L. Shaughnessy, S. Gigan, and H. Cao, Deep learning with passive optical nonlinear mapping (2023), [arXiv:2307.08558 \[physics.optics\]](https://arxiv.org/abs/2307.08558).
- [16] V. López-Pastor and F. Marquardt, Self-learning machines based on hamiltonian echo backpropagation, *Phys. Rev. X* **13**, 031020 (2023).
- [17] B. Scellier and Y. Bengio, Equilibrium propagation: Bridging the gap between energy-based models and backpropagation, *Frontiers in computational neuroscience* **11**, 24 (2017).
- [18] B. Scellier, A deep learning theory for neural networks grounded in physics, arXiv preprint arXiv:2103.09985 (2021).
- [19] J. Kendall, R. Pantone, K. Manickavasagam, Y. Bengio, and B. Scellier, Training end-to-end analog neural networks with equilibrium propagation, arXiv preprint arXiv:2006.01981 (2020).
- [20] E. Martin, M. Ernoult, J. Laydevant, S. Li, D. Querlioz, T. Petrisor, and J. Grollier, Eqspike: spike-driven equilibrium propagation for neuromorphic implementations, *Iscience* **24** (2021).
- [21] M. Ernoult, J. Grollier, D. Querlioz, Y. Bengio, and B. Scellier, Equilibrium propagation with continual weight updates, arXiv preprint arXiv:2005.04168 (2020).
- [22] B. Scellier, S. Mishra, Y. Bengio, and Y. Ollivier, Agnostic physics-driven deep learning (2022), [arXiv:2205.15021 \[cs.LG\]](https://arxiv.org/abs/2205.15021).
- [23] M. Falk, A. Strupp, B. Scellier, and A. Murugan, Contrastive learning through non-equilibrium memory (2023), [arXiv:2312.17723 \[cond-mat.dis-nn\]](https://arxiv.org/abs/2312.17723).
- [24] Q. Wang, C. C. Wanjura, and F. Marquardt, Training coupled phase oscillators as a neuromorphic platform using equilibrium propagation (2024), [arXiv:2402.08579 \[cs.ET\]](https://arxiv.org/abs/2402.08579).
- [25] M. Stern, D. Hexner, J. W. Rocks, and A. J. Liu, Supervised learning in physical networks: From machine learning to learning machines, *Physical Review X* **11**, 021045 (2021).
- [26] M. Stern, A. J. Liu, and V. Balasubramanian, The physical effects of learning (2023), [arXiv:2306.12928 \[cond-mat.dis-nn\]](https://arxiv.org/abs/2306.12928).
- [27] S. Dillavou, M. Stern, A. J. Liu, and D. J. Durian, Demonstration of decentralized physics-driven learning, *Physical Review Applied* **18**, 014040 (2022).
- [28] S. Dillavou, B. D. Beyer, M. Stern, M. Z. Miskin, A. J. Liu, and D. J. Durian, Machine learning without a processor: Emergent

- learning in a nonlinear electronic metamaterial, arXiv preprint arXiv:2311.00537 (2023).
- [29] L. E. Altman, M. Stern, A. J. Liu, and D. J. Durian, Experimental demonstration of coupled learning in elastic networks, arXiv preprint arXiv:2311.00170 (2023).
- [30] S. Oh, J. An, S. Cho, R. Yoon, and K.-S. Min, Memristor crossbar circuits implementing equilibrium propagation for on-device learning, *Micromachines* **14**, 1367 (2023).
- [31] J. Laydevant, D. Markovic, and J. Grollier, Training an ising machine with equilibrium propagation (2023), arXiv:2305.18321 [cs.NE].
- [32] E. Altman, K. R. Brown, G. Carleo, L. D. Carr, E. Demler, C. Chin, B. DeMarco, S. E. Economou, M. A. Eriksson, K.-M. C. Fu, *et al.*, Quantum simulators: Architectures and opportunities, *PRX quantum* **2**, 017003 (2021).
- [33] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, *Contemporary Physics* **56**, 172 (2015).
- [34] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
- [35] L. Onsager, Reciprocal relations in irreversible processes. i., *Phys. Rev.* **37**, 405 (1931).
- [36] R. Kubo, The fluctuation-dissipation theorem, *Reports on Progress in Physics* **29**, 255 (1966).
- [37] B. Scellier, M. Ernoult, J. Kendall, and S. Kumar, Energy-based learning algorithms for analog computing: a comparative study (2023), arXiv:2312.15103 [cs.LG].
- [38] K. Kim, M. Chang, R. Islam, S. Korenblit, L. Duan, and C. Monroe, Entanglement and tunable spin-spin couplings between trapped ions using multiple transverse modes., *Physical Review Letters* **103**, 120502 (2009).
- [39] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, *et al.*, Qubit architecture with high coherence and fast tunable coupling, *Physical review letters* **113**, 220502 (2014).
- [40] L.-M. Steinert, P. Osterholz, R. Eberhard, L. Festa, N. Lorenz, Z. Chen, A. Trautmann, and C. Gross, Spatially tunable spin interactions in neutral atom arrays, *Physical Review Letters* **130**, 243001 (2023).
- [41] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice, *Nature* **462**, 74 (2009).
- [42] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, *et al.*, The variational quantum eigensolver: a review of methods and best practices, *Physics Reports* **986**, 1 (2022).
- [43] J. Poyatos, J. I. Cirac, and P. Zoller, Quantum reservoir engineering with laser cooled trapped ions, *Physical review letters* **77**, 4728 (1996).
- [44] R. Ma, B. Saxberg, C. Owens, N. Leung, Y. Lu, J. Simon, and D. I. Schuster, A dissipatively stabilized mott insulator of photons, *Nature* **566**, 51 (2019).
- [45] M. Raghunandan, F. Wolf, C. Ospelkaus, P. O. Schmidt, and H. Weimer, Initialization of quantum simulators by sympathetic cooling, *Science Advances* **6**, eaaw9268 (2020).
- [46] S. Polla, Y. Herasymenko, and T. E. O’Brien, Quantum digital cooling, *Physical Review A* **104**, 012414 (2021).
- [47] J. Cotler, S. Choi, A. Lukin, H. Gharibyan, T. Grover, M. E. Tai, M. Rispoli, R. Schittko, P. M. Preiss, A. M. Kaufman, *et al.*, Quantum virtual cooling, *Physical Review X* **9**, 031013 (2019).
- [48] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nature Physics* **15**, 1273 (2019).
- [49] J. Herrmann, S. M. Lima, A. Remm, P. Zapletal, N. A. McMahon, C. Scarato, F. Swiadek, C. K. Andersen, C. Hellings, S. Krinner, *et al.*, Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases, *Nature communications* **13**, 4144 (2022).
- [50] Y.-J. Liu, A. Smith, M. Knap, and F. Pollmann, Model-independent learning of quantum phases of matter with quantum convolutional neural networks, *Physical Review Letters* **130**, 220603 (2023).
- [51] S. Massar and B. M. Mognetti, Equilibrium propagation: the quantum and the thermal cases (2024), arXiv:2405.08467 [quant-ph].
- [52] B. Scellier, Quantum equilibrium propagation: Gradient-descent training of quantum systems (2024), arXiv:2406.00879 [quant-ph].

## APPENDIX

### Appendix: Some details for the supervised learning example

We considered a cluster Ising chain of  $N = 8$  spins, and a sensor of 2 qubits, such that the total Hilbert space is 1024-dimensional. This allows efficient exact diagonalization using the Lanczos algorithm applied to sparse matrices, to find the ground state of the coupled system. The three considered output operators are projectors onto three states, each with a definite combination of the sensor operators  $\hat{Z}_{1'}$  and  $\hat{Z}_{2'}$ , as shown in the figure; e.g.  $\hat{P}_{1,-1} = (1 + \hat{Z}_{1'})(1 - \hat{Z}_{2'})/4$  projects onto the combination  $(+1, -1)$  and would be used to indicate the ferromagnetic phase, which is reached when  $g_{ZZ}$  dominates. The expectation value is correspondingly the probability to observe this particular combination in a projective measurement of these two operators. The training samples are drawn uniformly from the triangular phase diagram shown in the figure, where (following convention in the literature) we set the sum of all three coupling parameters to 4. The phase boundaries are known for this benchmark quantum many-body model [50], which allows us to provide the correct labels: in each phase, one of the projectors is 1, while the other two are 0 (“one-hot-encoding”), and the assignment of the three combinations of  $(Z_{1'}, Z_{2'})$  to the three phases is defined in a fixed arbitrary way.

Gradients are obtained using the symmetric nudging procedure, by adding the nudged output operators  $\sum_j \nu_\ell \hat{A}_\ell$  to the Hamiltonian, where in our case  $\hat{A}_\ell = \hat{P}_\ell$  and  $l \in (+1, +1), (+1, -1), (-1, -1)$ . The ground state for the nudged Hamiltonian (for both signs of nudging) is recalculated using sparse Lanczos diagonalization. Interestingly, our experiments have shown that this is numerically more efficient than attempting to obtain the exact linear response using first-order perturbation theory, which involves solving a linear system of equations when applying  $(E - \hat{H})^{-1}$  to the perturbed ground state. After the approximate gradient has been obtained using QEP, we use it inside an Adam adaptive gradient descent optimizer to update the parameters—which would be possible also for the QEP gradients obtained in real experiments. The learning rate employed in the numerical examples was set to 0.01.



Test accuracies are measured on a test set of 200 points that are also uniformly randomly distributed across the triangular space and which are fixed before the training run. We distinguish two measures of accuracy: In the ‘‘many queries’’ version, we imagine that one would obtain the expectation values of the three projectors, which requires multiple shots, and assign as official outcome the phase whose associated projector has the largest expectation value (largest measurement probability). This is similar to how accuracy would be assessed for classical machine-learning classification models. In the ‘‘single shot’’ version, we imagine to run inference only once (equilibrating to the ground state once) and performing a single measurement of the two sensor operators. The outcome will be declared correct if the combination matches the correct label of the true underlying phase for this parameter combination. Single-shot performance is more difficult, but even so the training results show that single-shot accuracy can also reach relatively high values. Repetition, e.g. using three shots and taking a majority vote when possible, will quickly boost the accuracy (until it reaches the ‘many queries’ result in the limit of many shots).

For assessing the influence of shot noise, we replace the exact expectation values by numerical values that are drawn from a Gaussian distribution centered around that value, with the correct variance  $\text{Var} \hat{A}_\ell / M$ , where  $M$  is the number of measurement shots (recall each of those will usually require a new equilibration, unless one adopts some of the strategies mentioned in the main text). Replacing the true distribution by a Gaussian is a reasonable approximation unless  $M$  is very small.

### Appendix: Different approaches for nudging

In classical EP, the idea is to add the loss function  $\mathcal{L}(y, y^{\text{target}}(x))$  to the energy, multiplied by  $\beta$ , where  $\beta$  is a small constant. Instead, in the main text we advocated simply adding  $\sum_\ell \nu_\ell \hat{A}_\ell$  to the Hamiltonian, where  $\hat{A}_\ell$  are the output operators and  $\nu_\ell = \beta \partial \mathcal{L} / \partial y_\ell$  is the nudge force, since for small  $\beta$  this produces the force needed to elicit the linear response required for the gradient. If we were instead to translate directly the classical EP prescription to the quantum Hamiltonian, we could use  $\beta \mathcal{L}(\hat{y}, y^{\text{target}}(x))$ , with the operator version of the outputs,  $\hat{y}_\ell = \hat{A}_\ell$  for  $\ell \in \mathcal{S}_{\text{out}}$ , replacing the expectation values  $y_\ell$ . Expanding this to linear order in  $\hat{A}_\ell$ , we would get the same result as our ansatz. The higher-order terms in  $\beta \mathcal{L}$  would lead to further corrections to the Hamiltonian, and depending on the quantum fluctuations in  $\hat{A}_\ell$  these might be as large as the linear-order term itself (e.g. for mean-squared error loss functions, these could correspond to a stiffening of the potential acting on  $\hat{A}_\ell$ ). Since these higher-order terms might also be more difficult to implement experimentally, depending on the shape of  $\hat{A}_\ell$ , and since they are, at the same time, not needed to evaluate the gradient, we opted for the procedure explained in the main text, adding only linear terms to the Hamiltonian. It would be interesting in the future to compare the various approaches for finite nudging, where

$\beta$  is not small (which is actually the case for the numerical experiments).

### Appendix: Further details for unsupervised learning examples

#### Phase exploration

We again consider a cluster Ising Hamiltonian, of 10 spins, and search for the points in the phase diagram that maximize  $\langle \hat{X}_0 \hat{X}_4 \rangle$ , hence,  $\hat{X}_0 \hat{X}_4$  is the output operator. This is expected to be maximized in the paramagnetic phase. Initial parameters will typically be chosen randomly. For illustrative purposes, we manually select the initial parameters, i.e. the starting points of the trajectories in Fig. 3 b. Concretely, the two starting points are  $g_X = -0.1, g_{ZZ} = 0.4$  (run 1) and  $g_X = 0.9, g_{ZZ} = 0.9$  (run 2). We obtain gradients according to the QEP procedure by switching on the coupling to the output operator  $\nu \hat{X}_0 \hat{X}_4$ . As in the supervised learning example, the ground state is computed using sparse Lanczos diagonalization. For the results shown in Fig. 3 b, both the learning rate and the nudge parameter  $\nu$  were set to 0.1

#### Sensitivity optimization

In the second unsupervised learning example, we search for the largest slope of  $\langle \hat{X}_0 \hat{X}_4 \rangle$  w.r.t.  $g_X$  in the phase diagram of the cluster Ising Hamiltonian (6) with 10 spins by optimizing the loss function (7). Concretely, we consider two different values of  $g_X^{(1,2)}$  while the other parameters are the same ( $g_{ZZ}$  is trained while  $g_{ZZZ} = -0.5 \equiv \text{const.}$ ).

During one step of the training, we update  $g_X^{(1)}, g_X^{(2)}$  and  $g_{ZZ}$ . To obtain the necessary gradients, we need to compute the following derivatives (for brevity we denote  $\partial / \partial \theta_j$  by  $\partial_{\theta_j}$ )

$$\begin{aligned} \partial_{g_X^{(1)}} \mathcal{L} = & \varepsilon \frac{\partial_{g_X^{(1)}} \langle \hat{X}_0 \hat{X}_4 \rangle |_{g_X^{(1)}}}{|g_X^{(1)} - g_X^{(2)}|} \\ & + \text{sgn}(g_X^{(1)} - g_X^{(2)}) \frac{\mathcal{L}}{|g_X^{(1)} - g_X^{(2)}|} \end{aligned} \quad (8)$$

$$\begin{aligned} \partial_{g_X^{(2)}} \mathcal{L} = & -\varepsilon \frac{\partial_{g_X^{(2)}} \langle \hat{X}_0 \hat{X}_4 \rangle |_{g_X^{(2)}}}{|g_X^{(1)} - g_X^{(2)}|} \\ & - \text{sgn}(g_X^{(1)} - g_X^{(2)}) \frac{\mathcal{L}}{|g_X^{(1)} - g_X^{(2)}|} \end{aligned} \quad (9)$$

$$\partial_{g_{ZZ}} \mathcal{L} = \varepsilon \partial_{g_{ZZ}} \frac{\left( \langle \hat{X}_0 \hat{X}_4 \rangle |_{g_X^{(1)}} - \langle \hat{X}_0 \hat{X}_4 \rangle |_{g_X^{(2)}} \right)}{|g_X^{(1)} - g_X^{(2)}|} \quad (10)$$

with  $\varepsilon \equiv \text{sgn}[\langle \hat{X}_0 \hat{X}_4 \rangle |_{g_X^{(1)}} - \langle \hat{X}_0 \hat{X}_4 \rangle |_{g_X^{(2)}}]$ . In all of the expressions, we use QEP to extract  $\partial_{\theta_j} \langle \hat{X}_0 \hat{X}_4 \rangle |_{g_X^{(1,2)}}$  with  $\theta_j \in \{g_X^{(1)}, g_X^{(2)}, g_{ZZ}\}$ .

To that end, as outlined in the main text, we approximate the derivative by comparing the nudged and the free expectation

value

$$\frac{\partial}{\partial \theta_j} \langle \hat{X}_0 \hat{X}_4 \rangle \approx \frac{\langle \hat{X}_0 \hat{X}_4 \rangle|_{\nu=\beta\varepsilon} - \langle \hat{X}_0 \hat{X}_4 \rangle|_{\nu=0}}{\beta}. \quad (11)$$

As before, to evaluate the expectation value in the nudged phase, we couple to the output operator by adding  $\nu \hat{X}_0 \hat{X}_4$  to the Hamiltonian. Ground states in the free and the nudged phase are again computed using sparse Lanczos diagonaliza-

tion.

To illustrate the procedure, we consider two different starting points in the phase diagram:  $g_X^{(1)} = -0.2$ ,  $g_X^{(2)} = -1.5$ ,  $g_{ZZ} = -1.5$  (run 1; each points starts in a different phase) and  $g_X^{(1)} = -0.5$ ,  $g_X^{(2)} = 0.3$ ,  $g_{ZZ} = 1.0$  (run 2; both points are in the same phase). Both the learning rate and the nudging parameter were set to 0.1.