Quantum Equilibrium Propagation: Gradient-Descent Training of Quantum Systems

Anonymous Author(s)
Affiliation
Address
email

Abstract

Equilibrium propagation (EP) is a training framework for physical systems that minimize an energy function. A key feature of EP is that it uses the system's intrinsic physics during both inference and training, making it a candidate for the development of energy-efficient processors for machine learning. EP has been explored in various classical physical systems, including classical Ising networks and elastic networks. We extend EP to quantum systems, where the energy function that is minimized is the mean energy functional (expectation value of the Hamiltonian), whose minimum is the Hamiltonian's ground state. As examples, we study the settings of the transverse-field Ising network and the quantum harmonic oscillator network – quantum analogues of the Ising network and elastic network.

1 Introduction

Commercial applications of machine learning (ML) are powered by classical digital computing. Meanwhile, fundamental research explores alternative computing paradigms to enhance ML capabilities. Quantum computing leverages the principles of quantum mechanics to encode and process information in ways that classical computers cannot, potentially handling exponentially larger amounts of information. In contrast, neuromorphic computing, taking inspiration from the brain's energy efficiency, aims to leverage analog physics and compute-in-memory platforms to significantly reduce the cost of inference and training in ML [Marković et al., 2020]. An emerging field of research known as 'physical learning' [Stern and Murugan, 2023] shares similar goals with neuromorphic computing, but explores the inherent physics of any physical system for computation, without necessarily mimicking neurons and synapses – see Momeni et al. [2024] for a very recent review.

A key lesson from ML research over the past decades is the effectiveness of frameworks for optimizing cost functions, e.g. the backpropagation framework. One challenge for neuromorphic computing and physical learning has been the search for such frameworks that adhere to local computation and local learning rules, which are essential features for implementation on analog compute-in-memory platforms. In recent years, several gradient-descent training frameworks for physical systems have been proposed. For instance, Lopez-Pastor and Marquardt [2023] introduced a framework applicable to arbitrary time-reversal invariant Hamiltonian systems, and Wanjura and Marquardt [2024a] developed a method for extracting weight gradients in optical systems based on linear wave scattering. The present paper focuses on the training framework known as equilibrium propagation.

Equilibrium propagation (EP), introduced in Scellier and Bengio [2017], is a training framework for energy-based systems, in which physics drives the system's state towards the minimum of an energy function (equilibrium or steady state). EP extracts the gradients of the cost function using two equilibrium states corresponding to different boundary conditions, which are then used to locally

Submitted to the Second Workshop on Machine Learning with New Compute Paradigms at NeurIPS (MLNCP 2024). Do not distribute.

adjust the trainable weights of the system. EP has been applied to various systems, including resistor networks [Kendall et al., 2020], elastic and flow networks [Stern et al., 2021], spiking networks [Martin et al., 2021], the (classical) Ising model [Laydevant et al., 2024], and coupled phase oscillators [Wang et al., 2024]. Recent experimental demonstrations have shown the applicability of EP on hardware: Dillavou et al. [2022, 2024] built two generations of self-learning resistor networks, Altman et al. [2024] built a self-learning elastic network, [Yi et al., 2023] used a variant of EP in a memristor crossbar array, and Laydevant et al. [2024] used EP on D-wave to train a classical Ising network (where, interestingly, they used quantum annealing to reach the ground state). Simulations have further underscored the potential of EP for ML applications: in particular, Laborieux and Zenke [2022] trained an energy-based convolutional network to classify a downsampled version of the ImageNet dataset. More broadly, [Zucchet and Sacramento, 2022] have highlighted EP's general applicability to any bilevel optimization problem (beyond the training of energy-based systems), including meta-learning [Zucchet et al., 2022].

We introduce Quantum Equilibrium Propagation (QEP), an extension of EP to quantum systems. In QEP, the system is brought to the ground state of its Hamiltonian, parameterized by real-valued trainable weights, to produce a prediction. The algorithm performs gradient descent on the expectation value of an observable, which serves as the cost function to optimize. Thus, in QEP, the classical EP's energy function is replaced by the system's Hamiltonian, and the equilibrium state extremizing the energy function is replaced by the ground state of the Hamiltonian. The central ingredient for translating from EP to QEP is the energy expectation value, minimized (more generally, extremized) at the Hamiltonian's ground state (more generally, eigenstates). Similar to EP, an interesting feature of QEP is the locality of the learning rule, which might be useful for the development of specialized quantum hardware with reduced classical overhead, where measurements of the weight gradients and adjustments of the trainable weights would be performed locally. To illustrate QEP, we study the settings of the transverse-field Ising network and the quantum harmonic oscillator network – quantum analogues of the Ising model and elastic network.

In parallel to this work, two other strongly related manuscripts on quantum extensions of EP have recently appeared on Arxiv [Massar and Mognetti, 2024, Wanjura and Marquardt, 2024b]. Massar and Mognetti [2024] also studied the thermal case of EP, where a thermodynamic system settles to the minimum of the free energy functional, and showed how to extract the weight gradients from thermal fluctuations alone, while Wanjura and Marquardt [2024b] established a connection between EP and Onsager's reciprocity.

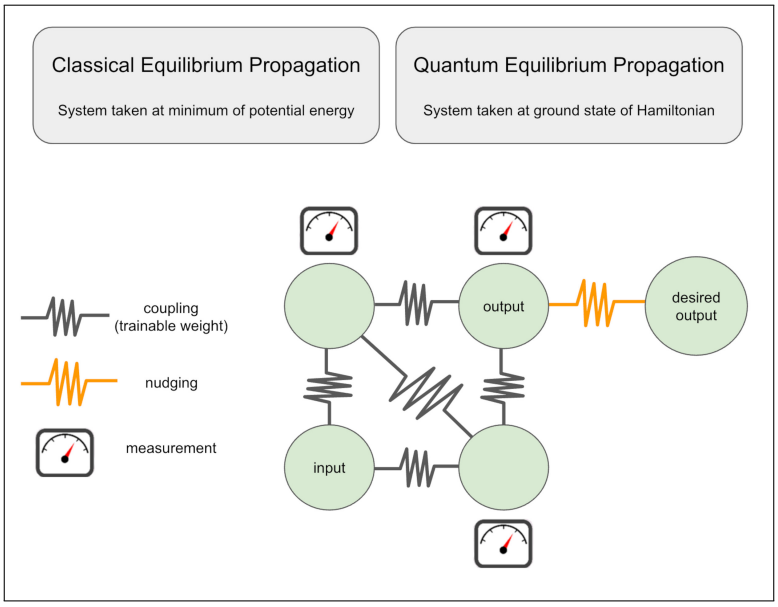

Figure 1: Example of a network, trainable via equilibrium propagation (EP). While in classical EP, the system is taken at the minimum of its energy function, in QEP the system is taken at the Hamiltonian's ground state.

2 Quantum Equilibrium Propagation

We present Quantum Equilibrium Propagation (QEP), a version of Equilibrium Propagation (EP) for quantum systems. For a brief presentation of classical EP, see Appendix A. For a brief presentation of the concepts of quantum mechanics required for QEP, see Appendix B.

We consider a quantum system, serving as a 'learning machine', whose Hamiltonian is parameterized by trainable weights $w = (w_1, w_2, \dots, w_M)$. An input x can also be supplied to the system. We denote the corresponding Hamiltonian as $\widehat{H}(w, x)$. We assume that the system can be set in its ground state $|\psi(w, x)\rangle$, characterized by

$$\widehat{H}(w,x)|\psi(w,x)\rangle = E(w,x)|\psi(w,x)\rangle, \tag{1}$$

where $E(w, x)$ is the ground state energy level, i.e. the lowest eigenvalue of $\widehat{H}(w, x)$. Similar to the classical setting, the ground state $|\psi(w, x)\rangle$ is used to encode a prediction on the desired output y, based on the supplied input x. We also assume an observable $\widehat{C}(y)$ parameterized by y, whose expectation value in the state $|\psi(w, x)\rangle$ represents the cost function to minimize,

$$\langle \widehat{C}(y)\rangle_{\psi(w,x)} = \langle\psi(w,x)|\widehat{C}(y)|\psi(w,x)\rangle. \tag{2}$$

The goal is to adjust the trainable weights of the Hamiltonian to minimize this cost function. Similar to classical EP, we assume that $\widehat{C}(y)$ is the Hamiltonian of an interaction between the system's state $|\psi\rangle$ and desired output y, and that this interaction can be integrated into the system. Specifically, we form the 'total Hamiltonian'

$$\widehat{H}^\beta = \widehat{H}(w,x) + \beta\widehat{C}(y), \tag{3}$$

where $\beta \in \mathbb{R}$ controls the strength of the new interaction. Given an input-output pair (x, y) from training data, QEP proceeds as follows.

1. Set $\beta = 0$. For each $k \in \{1, 2, \cdots, M\}$, repeat the following step T times: reach the ground state $|\psi^0_\star\rangle$ of $\widehat{H}^0$ (with associated ground state energy $E^0_\star$), characterized by
$$\widehat{H}^0|\psi^0_\star\rangle = E^0_\star|\psi^0_\star\rangle, \tag{4}$$
and measure the Hamiltonian derivative $\frac{\partial\widehat{H}^0}{\partial w_k}$. Denote the outcomes of the T measurements as $h^{(1)}_k(0), h^{(2)}_k(0), \dots, h^{(T)}_k(0)$.
2. Set $\beta > 0$ and proceed as above. For each $k \in \{1, 2, \cdots, M\}$, repeat the following step T times: reach the ground state $|\psi^\beta_\star\rangle$ of $\widehat{H}^\beta$ (with associated ground state energy $E^\beta_\star$), characterized by
$$\widehat{H}^\beta|\psi^\beta_\star\rangle = E^\beta_\star|\psi^\beta_\star\rangle, \tag{5}$$
and measure the Hamiltonian derivative $\frac{\partial\widehat{H}^\beta}{\partial w_k}$. Denote the outcomes of the T measurements as $h^{(1)}_k(\beta), h^{(2)}_k(\beta), \dots h^{(T)}_k(\beta)$.
3. Update the trainable weights $w_1, w_2, \dots, w_M$ as
$$\Delta w_k = \frac{\eta}{\beta}\left[\frac{1}{T}\sum_{t=1}^{T} h^{(t)}_k(0) - \frac{1}{T}\sum_{t=1}^{T} h^{(t)}_k(\beta)\right]. \tag{6}$$

Similar to classical EP, the central insight of QEP is that the above learning rule (Eq. (6)) approximates one step of gradient descent on the expectation value of the cost observable.

Theorem 1. *The gradient of the cost function can be approximated as*

$$\nabla_w\langle\psi(w,x)|\widehat{C}(y)|\psi(w,x)\rangle = \frac{d}{d\beta}\bigg|_{\beta=0}\langle\psi^\beta_\star|\frac{\partial\widehat{H}^\beta}{\partial w}|\psi^\beta_\star\rangle \tag{7}$$

$$\approx \frac{1}{\beta}\left[\langle\psi^\beta_\star|\frac{\partial\widehat{H}^\beta}{\partial w}|\psi^\beta_\star\rangle - \langle\psi^0_\star|\frac{\partial\widehat{H}^0}{\partial w}|\psi^0_\star\rangle\right] \tag{8}$$

$$\approx \frac{1}{\beta}\left[\frac{1}{T}\sum_{t=1}^{T} h^{(t)}_k(\beta) - \frac{1}{T}\sum_{t=1}^{T} h^{(t)}_k(0)\right]. \tag{9}$$

Theorem 1 directly follows from Theorem 2 and Lemma 3. We first comment on the differences between EP and QEP, and then discuss the properties of EP that transfer to QEP.

Compared to the classical setting, the characteristics of quantum measurements affect the training procedure in several ways. First, since a quantum measurement only gives an unbiased estimate of the expectation value, multiple measurements of the Hamiltonian derivatives $\frac{\partial \widehat{H}}{\partial w_k}$ are required to get better estimates of the weight gradients. Second, since the state of a quantum system generally changes upon measurement of a Hamiltonian derivative, the system must be reset to its ground state after each measurement. Third, the Hamiltonian derivatives $\frac{\partial \widehat{H}}{\partial w_j}$ and $\frac{\partial \widehat{H}}{\partial w_k}$ cannot generally be measured simultaneously unless they commute (we will see, however, in Sections 2.1-2.2 that a large fraction of the Hamiltonian derivatives can typically be measured simultaneously). Fourth, QEP involves two levels of approximation in the estimate of the gradient of the cost function. In addition to the first level (also present in classical EP), which is due to the finite difference used to approximate the derivative $\frac{d}{d\beta}$ at $\beta = 0$, a second level is due to the probabilistic nature of quantum measurements.

QEP also shares multiple features of classical EP. Suppose the total Hamiltonian of the system can be expressed as the sum of Hamiltonians corresponding to individual interactions or contributions, i.e.

$$\widehat{H} = \widehat{H}_0 + \widehat{H}_1 + \widehat{H}_2 + \cdots + \widehat{H}_M \tag{10}$$

where $\widehat{H}_k$ is the Hamiltonian of an interaction parameterized by w_k (for $1 \leq k \leq M$), and $\widehat{H}_0$ does not depend on any trainable weight. Then the Hamiltonian derivatives simplify to $\frac{\partial \widehat{H}}{\partial w_k} = \frac{\partial \widehat{H}_k}{\partial w_k}$. If the trainable weight w_k is stored close to where the observable $\frac{\partial \widehat{H}_k}{\partial w_k}$ is measured, the learning rule for w_k is local. Moreover, no knowledge of $\widehat{H}_0$ is required, so the system's Hamiltonian need not be fully known. Finally, similar to the classical setting where the equilibrium state need not be stable but only a critical point (stationary state) of the energy function, QEP applies to any eigenstate of the system's Hamiltonian, not just the ground state. One condition for Eq. (8) to hold, however, is that the nudge eigenstate $|\psi_\star^\beta\rangle$ must be obtained as a smooth deformation (adiabatic transformation) of the free eigenstate $|\psi_\star^0\rangle$ when varying the nudging parameter from 0 to $\beta \neq 0$.

Next, we present for illustration the setting of the transverse-field Ising model and quantum harmonic oscillator network.

2.1 Transverse-Field Ising Model

As a first example, we consider the transverse-field Ising model, a quantum version of the classical Ising model described in Section A.1. While a classical Ising network of N classical spins can be in either of 2^N possible configurations, a quantum Ising network of N spins exists in a superposition of these 2^N configurations (i.e. a linear combination with coefficients in $\mathbb{C}$). Hence the major difference between the classical and the quantum settings: while the state of the classical model is described by a N-dimensional binary-valued vector, the state of the quantum model is described by a 2^N-dimensional complex-valued vector. We denote the $d = 2^N$ basis states as $|\sigma_1\sigma_2\cdots\sigma_N\rangle$ with $\sigma_k \in \{\uparrow, \downarrow\}$ for each $k \in \{1, 2, \ldots, N\}$, e.g. $|\uparrow\uparrow\cdots\uparrow\uparrow\rangle$, $|\uparrow\uparrow\cdots\uparrow\downarrow\rangle$ and similarly for the other $2^N - 2$ basis states.

Similar to the Ising energy function of Eq. (32), the Hamiltonian of the transverse-field Ising model has couplings between spins $J_{jk} \in \mathbb{R}$ and bias fields $h_k \in \mathbb{R}$ applied to individual spins, which we view as trainable weights. The Ising Hamitonian takes the form

$$\widehat{H}_{\text{Ising}} = -\sum_{1\leq j<k\leq N} J_{jk}\widehat{Z}_j\widehat{Z}_k - \sum_{k=1}^{N} h_k\widehat{X}_k, \tag{11}$$

where $\widehat{Z}_k$ and $\widehat{X}_k$ are the Pauli operators, defined as follows. The Pauli $\widehat{Z}_k$ operator acts as a phase-flip operator on the k-th spin, according to:

$$\widehat{Z}_k\ |\sigma_1\cdots\sigma_{k-1}\uparrow\sigma_{k+1}\cdots\sigma_N\rangle = +|\sigma_1\cdots\sigma_{k-1}\uparrow\sigma_{k+1}\cdots\sigma_N\rangle, \tag{12}$$

$$\widehat{Z}_k\ |\sigma_1\cdots\sigma_{k-1}\downarrow\sigma_{k+1}\cdots\sigma_N\rangle = -|\sigma_1\cdots\sigma_{k-1}\downarrow\sigma_{k+1}\cdots\sigma_N\rangle. \tag{13}$$

The Pauli $\widehat{X}_k$ operator acts as a bit-flip operator on the k-th spin, according to:

$$\widehat{X}_k \, |\sigma_1 \cdots \sigma_{k-1} \uparrow \sigma_{k+1} \cdots \sigma_N\rangle = |\sigma_1 \cdots \sigma_{k-1} \downarrow \sigma_{k+1} \cdots \sigma_N\rangle, \tag{14}$$

$$\widehat{X}_k \, |\sigma_1 \cdots \sigma_{k-1} \downarrow \sigma_{k+1} \cdots \sigma_N\rangle = |\sigma_1 \cdots \sigma_{k-1} \uparrow \sigma_{k+1} \cdots \sigma_N\rangle. \tag{15}$$

In this setting, the gradients of the Ising Hamiltonian with respect to the trainable weights, required in the learning rule of Eq. (6), are given by

$$\frac{\partial \widehat{H}_{\text{Ising}}}{\partial J_{jk}} = -\widehat{Z}_j \widehat{Z}_k, \qquad \frac{\partial \widehat{H}_{\text{Ising}}}{\partial h_k} = -\widehat{X}_k. \tag{16}$$

Importantly, the Pauli $\widehat{Z}_k$ operators ($1 \leq k \leq N$) commute with one another, allowing them to be measured simultaneously (when the system is in the ground state). Similarly, the Pauli $\widehat{X}_k$ operators ($1 \leq k \leq N$) also commute and can be measured simultaneously. However, $\widehat{Z}_j$ and $\widehat{X}_k$ do not commute, so they cannot be measured simultaneously.

2.2 Quantum Harmonic Oscillator Network

As a second example, we consider the quantum harmonic oscillator network, a quantum analogue of the elastic network model presented in Section A.2. It consists of N quantum particles, such as atoms, interacting via harmonic potentials. We assume for simplicity that each atom is described by a one-dimensional (rather than three-dimensional) position. While the classical elastic network model is described by the N-dimensional vector of positions of the atoms $(r_1, r_2, \cdots, r_N) \in \mathbb{R}^N$, the state of the quantum network is a superposition of all these configurations. In this setting, the state vector is a function $\psi : \mathbb{R}^N \to \mathbb{C}$ (the wave function), that assigns a complex number $\psi(r_1, r_2, \cdots, r_N)$ to each possible configuration $(r_1, r_2, \cdots, r_N)$. Contrary to the quantum Ising network, the corresponding Hilbert space is infinite-dimensional.

The position and momentum operators of the i-th atom, denoted as $\widehat{r}_i$ and $\widehat{p}_i$, are defined by their action on the wavefunction as follows:

$$(\widehat{r}_i \psi)(r_1, r_2, \cdots, r_N) = r_i \psi(r_1, r_2, \cdots, r_N), \tag{17}$$

$$(\widehat{p}_i \psi)(r_1, r_2, \cdots, r_N) = -\mathbf{i}\hbar \frac{\partial \psi}{\partial r_i}(r_1, r_2, \cdots, r_N), \tag{18}$$

where $\mathbf{i}$ is the imaginary unit ($\mathbf{i}^2 = -1$) and $\hbar$ is the reduced Planck constant. Similar to the elastic energy function, the Hamiltonian of the quantum harmonic oscillator network is given by:

$$\widehat{H}_{\text{QHO}} = \sum_{i=1}^{N} \frac{\widehat{p}_i^2}{2m_i} + \frac{1}{2} \sum_{1 \leq i,j \leq N} k_{ij} \left(\widehat{r}_i - \widehat{r}_j\right)^2, \tag{19}$$

where $\frac{\widehat{p}_i^2}{2m_i}$ is the kinetic energy operator of the i-th atom, and $\frac{1}{2} k_{ij} \left(\widehat{r}_i - \widehat{r}_j\right)^2$ is the harmonic potential operator between the i-th and j-th atoms. In these expressions, m_i are the masses of the atoms and k_{ij} are the spring constants, which serve as trainable weights. The partial derivatives of the Hamiltonian with respect to the spring constants are given by

$$\frac{\partial \widehat{H}_{\text{QHO}}}{\partial k_{ij}} = \frac{1}{2} \left(\widehat{r}_i - \widehat{r}_j\right)^2. \tag{20}$$

It is straightforward to verify that all $\widehat{r}_i$ operators commute with one another, which allows us to measure these observables simultaneously to obtain the gradients of the cost function.

3 Discussion

Equilibrium Propagation (EP) has been studied in various classical physical systems, and has been implemented experimentally in resistor networks, classical Ising networks and elastic networks. EP is generally applicable to systems that extremize an energy functional. Quantum Equilibrium Propagation (QEP) extends EP to quantum systems, where the extremized functional is the mean energy, achieving its extrema at the eigenstates of the system's Hamiltonian. QEP can serve both as a

normative framework for training quantum systems, and to justify the use of EP in physical systems where quantum effects may arise and where the classical EP framework is not directly applicable. For instance, QEP could provide insights into the work by Laydevant et al. [2024], which employed quantum annealing in the context of classical EP.

Similar to EP, one attractive feature of QEP is the locality of the learning rule. This locality might be advantageous for building specialized quantum computers with reduced classical computation overhead, where the physical quantities serving as trainable weights would be located near the location where the Hamiltonian derivatives are measured. Another attractive feature is that the Hamitonian terms of interactions that do not depend on trainable weights need not be known explicitly, so that QEP is partially agnostic to the system's Hamiltonian. A difference in the quantum setting is the probabilistic nature of measurements. Multiple measurements are generally required to obtain accurate gradient estimates of the cost function, which could necessitate additional memory to store the outcomes. To address this, studying the effect of single measurements of single eigenstates (free or nudge) on the variance of the gradient estimator would be useful, similar to the study by Williams et al. [2023] in the classical setting. Another difference with the classical setting is that Hamiltonian derivatives generally cannot be measured simultaneously, unless they commute. We discussed the transverse-field Ising model and the quantum harmonic oscillator network and we have seen that measurements in these models can largely be parallelized.

Similar to the classical setting where EP can be applied to any stationary state (critical point) of the system's energy function, QEP can be applied to any eigenstate of the system's Hamiltonian, not just the ground state. While this feature provides greater flexibility, a caveat is that for Eq. (8) to hold, the nudge eigenstate must in principle be obtained as a smooth deformation (adiabatic transformation) of the free eigenstate. It remains to be seen whether this condition is necessary or can be further relaxed in practice.

The next step would be to simulate QEP numerically, but we leave this for future works. For quantum Ising networks, the exact diagonalization method becomes impractical when the number of spins N exceeds a few dozen due to the exponential growth of the state space ($d = 2^N$). For larger systems, approximate methods such as the Density Matrix Renormalization Group (DMRG) and Variational Monte Carlo (VMC) could be employed.

References

L. E. Altman, M. Stern, A. J. Liu, and D. J. Durian. Experimental demonstration of coupled learning in elastic networks. *Physical Review Applied*, 22(2):024053, 2024.

S. Dillavou, M. Stern, A. J. Liu, and D. J. Durian. Demonstration of decentralized physics-driven learning. *Physical Review Applied*, 18(1):014040, 2022.

S. Dillavou, B. D. Beyer, M. Stern, A. J. Liu, M. Z. Miskin, and D. J. Durian. Machine learning without a processor: Emergent learning in a nonlinear analog network. *Proceedings of the National Academy of Sciences*, 121(28):e2319718121, 2024.

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.

A. Laborieux and F. Zenke. Holomorphic equilibrium propagation computes exact gradients through finite size oscillations. *Advances in Neural Information Processing Systems*, 35:12950–12963, 2022.

J. Laydevant, D. Marković, and J. Grollier. Training an ising machine with equilibrium propagation. *Nature Communications*, 15(1):3671, 2024.

V. Lopez-Pastor and F. Marquardt. Self-learning machines based on hamiltonian echo backpropagation. *Physical Review X*, 13(3):031020, 2023.

D. Marković, A. Mizrahi, D. Querlioz, and J. Grollier. Physics for neuromorphic computing. *Nature Reviews Physics*, 2(9):499–510, 2020.

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(3), 2021.

S. Massar and B. M. Mognetti. Equilibrium propagation: the quantum and the thermal cases. *arXiv preprint arXiv:2405.08467*, 2024.

A. Momeni, B. Rahmani, B. Scellier, L. G. Wright, P. L. McMahon, C. C. Wanjura, Y. Li, A. Skalli, N. G. Berloff, T. Onodera, et al. Training of physical neural networks. *arXiv preprint arXiv:2406.03372*, 2024.

B. Scellier. *A deep learning theory for neural networks grounded in physics*. PhD thesis, Université de Montréal, 2021.

B. Scellier and Y. Bengio. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in computational neuroscience*, 11:24, 2017.

M. Stern and A. Murugan. Learning without neurons in physical systems. *Annual Review of Condensed Matter Physics*, 14:417–441, 2023.

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(2):021045, 2021.

Q. Wang, C. C. Wanjura, and F. Marquardt. Training coupled phase oscillators as a neuromorphic platform using equilibrium propagation. *arXiv preprint arXiv:2402.08579*, 2024.

C. C. Wanjura and F. Marquardt. Fully nonlinear neuromorphic computing with linear wave scattering. *Nature Physics*, pages 1–7, 2024a.

C. C. Wanjura and F. Marquardt. Quantum equilibrium propagation for efficient training of quantum systems based on onsager reciprocity. *arXiv e-prints*, pages arXiv–2406, 2024b.

E. Williams, C. Bredenberg, and G. Lajoie. Flexible phase dynamics for bio-plausible contrastive learning. In *International Conference on Machine Learning*, pages 37042–37065. PMLR, 2023.

S.-i. Yi, J. D. Kendall, R. S. Williams, and S. Kumar. Activity-difference training of deep neural networks using memristor crossbars. *Nature Electronics*, 6(1):45–51, 2023.

N. Zucchet and J. Sacramento. Beyond backpropagation: bilevel optimization through implicit differentiation and equilibrium propagation. *Neural Computation*, 34(12):2309–2346, 2022.

N. Zucchet, S. Schug, J. Von Oswald, D. Zhao, and J. Sacramento. A contrastive rule for meta-learning. *Advances in Neural Information Processing Systems*, 35:25921–25936, 2022.

A Equilibrium Propagation

In this appendix, we review the Equilibrium Propagation (EP) training framework [Scellier and Bengio, 2017] and the classical Ising network and elastic network where it has been used.

EP applies in physical systems governed by dynamics that drive their state s towards the minimum of an energy function $\mathcal{E}$. These systems contain trainable weights $w = (w_1, w_2, \dots, w_M)$ and can take an input x supplied as a boundary condition. We denote the corresponding energy function as $\mathcal{E}(w, x, s)$. During inference, given an input x, the system evolves towards its equilibrium or steady state, characterized by

$$s(w, x) = \arg\min_s \mathcal{E}(w, x, s). \tag{21}$$

The system thus implements a function $x \mapsto s(w, x)$. Training the system consists in adjusting the weights w so that $s(w, \cdot)$ matches a desired input-output function. Mathematically, we use a cost function $\mathcal{C}(s(w, x), y)$ which, given an input x and its associated desired output y, measures the accuracy of the prediction $s(w, x)$ by comparing it with y. Training the system can be formulated as a bilevel optimization problem:

$$\text{minimize } \mathcal{J}(w) = \mathbb{E}_{(x,y)} \left[\mathcal{C}(s(w, x), y) \right], \tag{22}$$

$$\text{subject to } s(w, x) = \arg\min_s \mathcal{E}(w, x, s), \tag{23}$$

where $\mathbb{E}_{(x,y)}$ represents the expectation value over input-output pairs (x, y) from the training data. The conventional method to solve this problem is gradient descent on the upper-level objective: at each step of training, an input-output pair (x, y) is picked from the training data, and the trainable weights are adjusted in proportion to the gradient of the cost function: $\Delta w = -\eta \nabla_w \mathcal{C}(s(w, x), y)$, where $\eta > 0$ is a learning rate. The remaining task is to obtain or estimate the weight gradients, $\nabla_w \mathcal{C}(s(w, x), y)$, using the system's physics. This is what EP enables us to do. The central idea of EP is to view the cost function $\mathcal{C}(s, y)$ as the energy of an interaction between the state variables (s) and desired outputs (y), which can be incorporated into the system's energy function to form the 'total energy function',

$$\mathcal{E}^\beta(w, x, s, y) = \mathcal{E}(w, x, s) + \beta \mathcal{C}(s, y), \tag{24}$$

where $\beta \in \mathbb{R}$ is a parameter termed the 'nudging parameter' that controls the strength of this new interaction. See Figure 1. EP proceeds in three steps:

1. Set $\beta = 0$ and let the system settle to an equilibrium state $s^0_\star$, called the 'free state', characterized by
$$s^0_\star = \arg\min_s \mathcal{E}^0(w, x, s, y) = s(w, x). \tag{25}$$
For each $k \in \{1, 2, \cdots, M\}$, measure $\frac{\partial \mathcal{E}^0}{\partial w_k}(w_1, \dots, w_M, x, s^0_\star, y)$, i.e. the partial derivative of the energy function with respect to w_k.

2. Set $\beta > 0$ and let the system reach a new equilibrium state $s^\beta_\star$, called the 'nudge state', characterized by
$$s^\beta_\star = \arg\min_s \mathcal{E}^\beta(w, x, s, y). \tag{26}$$
Measure again the partial derivative $\frac{\partial \mathcal{E}^\beta}{\partial w_k}(w_1, \dots, w_M, x, s^\beta_\star, y)$ for each $k \in \{1, 2, \cdots, M\}$.

3. Update the trainable weights $w_1, w_2, \dots, w_M$ as
$$\Delta w_k = \frac{\eta}{\beta} \left[\frac{\partial \mathcal{E}^0}{\partial w_k}(w, x, s^0_\star, y) - \frac{\partial \mathcal{E}^\beta}{\partial w_k}(w, x, s^\beta_\star, y) \right], \tag{27}$$
where $\eta > 0$ is a learning rate.

The main theoretical result of EP is that the above contrastive learning rule (27) approximates one step of gradient descent on the cost function.

Theorem 2 (Equilibrium Propagation)**.** *The gradient of the cost function with respect to the trainable weights can be approximated as*

$$\nabla_w \mathcal{C}(s(w,x),y) = \frac{d}{d\beta}\frac{\partial \mathcal{E}^\beta}{\partial w}(w,x,s^\beta_\star,y)\bigg|_{\beta=0} \tag{28}$$

$$\approx \frac{1}{\beta}\left[\frac{\partial \mathcal{E}^\beta}{\partial w}(w,x,s^\beta_\star,y) - \frac{\partial \mathcal{E}^0}{\partial w}(w,x,s^0_\star,y)\right]. \tag{29}$$

Theorem 2 is proved in Scellier and Bengio [2017]. One difference with traditional machine learning methods which use automatic differentiation (i.e. backpropagation) is that EP does not perform exact gradient descent on the cost function, but rather approximates it. We review improved versions of EP that mitigate this problem in Appendix **??**. Next, we discuss some important features of EP.

First, one of the primary interests of EP is that, in a wide range of physical systems, the contrastive learning rule of Eq. (27) is local for each trainable weight. To see this, let $w = (w_1, w_2, \dots, w_M)$ be the set of trainable weights, and assume that the energy function is separable,

$$\mathcal{E} = \mathcal{E}_0 + \mathcal{E}_1 + \mathcal{E}_2 + \cdots + \mathcal{E}_M, \tag{30}$$

where each $\mathcal{E}_k$ (with $1 \leq k \leq M$) is the energy term of an interaction parameterized by w_k (and w_k only), while $\mathcal{E}_0$ is an energy term that does not depend on any trainable weight. The energy derivatives arising in the learning rule simplify as $\frac{\partial \mathcal{E}}{\partial w_k} = \frac{\partial \mathcal{E}_k}{\partial w_k}$. If the energy term $\mathcal{E}_k$ involves only state variables spatially close to w_k, the learning rule for w_k is local in space. Below we illustrate this property in the classical Ising model [Laydevant et al., 2024] and elastic network model [Stern et al., 2021].

Second, the system's energy function may be partially unknown. Specifically, in the above decomposition, while knowledge of the energy derivatives $\frac{\partial \mathcal{E}_k}{\partial w_k}$ is required, the energy term $\mathcal{E}_0$ need not be analytically known. Below we illustrate this property in the elastic network model.

Last, the equilibrium states $s(w,x)$ and $s^\beta_\star$ of Eq. (21) and Eq. (26) need not be stable (i.e. minima of their repective energy functions). Theorem 2 is more generally valid when $s(w,x)$ and $s^\beta_\star$ are critical points (i.e. saddle points) of their respective energy functions, satisfying the stationary conditions

$$\frac{\partial \mathcal{E}}{\partial s}(w,x,s(w,x)) = 0, \qquad \frac{\partial \mathcal{E}^\beta}{\partial s}(w,x,s^\beta_\star,y) = 0. \tag{31}$$

See Scellier [2021, Chapter 2] for a brief discussion on this matter. A condition for Theorem 2 to hold is that the nudge state $s^\beta_\star$ must be the stationary state obtained as a smooth deformation of $s^0_\star$ as we gradually vary the nudging parameter from 0 to $\beta \neq 0$.

Next, we present two examples of physical systems where EP is applicable. We illustrate the locality of the learning rule and the fact that the energy function may be partially unknown.

A.1 Ising Network

As a first example, we present the (classical) Ising model of coupled spins. This widely studied model has been explored in particular as a computing platform for machine learning, and recently studied in the context of EP [Laydevant et al., 2024]. The model consists of N classical spins, characterized by their state $\sigma_k \in \{+1, -1\}$ for $1 \leq k \leq N$, representing "up" or "down" states. The state of the system is represented by the N-dimensional vector of spin states, $s = (\sigma_1, \sigma_2, \dots, \sigma_N)$, so the state space is discrete and finite, consisting of 2^N possible configurations. The Ising energy function that the system seeks to minimize is defined as

$$\mathcal{E}_{\text{Ising}}(\sigma_1, \sigma_2, \dots, \sigma_N) = -\sum_{1 \leq j < k \leq N} J_{jk}\sigma_j\sigma_k - \sum_{k=1}^{N} h_k\sigma_k, \tag{32}$$

where J_{jk} represents the couplings between spins, and h_k represents the bias fields applied to individual spins. These parameters serve as trainable weights in the model. The partial derivatives of the energy function with respect to these trainable weights, given by

$$\frac{\partial \mathcal{E}_{\text{Ising}}}{\partial J_{jk}} = -\sigma_j\sigma_k, \qquad \frac{\partial \mathcal{E}_{\text{Ising}}}{\partial h_k} = -\sigma_k, \tag{33}$$

involve only information that is locally available to J_{jk} and h_k, respectively.

In terms of hardware demonstration, Laydevant et al. [2024] implemented an Ising network on the D-Wave Ising machine, employing the quantum annealing procedure of D-Wave to reach the ground state. They trained it to classify the MNIST handwritten digits using EP. They also emulated a small convolutional Ising network, using the Chimera architecture of D-Wave's chips to implement the necessary convolutional operations.

A.2 Elastic Network

As a second example, we consider the elastic network model, studied by Stern et al. [2021] in the context of Coupled Learning (CL), a variant of EP. We follow the methodology of Kendall et al. [2020] to treat the case of nonlinear networks and illustrate that EP is partly agnostic to the system's energy function.

We consider a network of N masses $m_1, m_2, \dots, m_N$ interconnected by springs. We denote the position of mass m_i as $\vec{r}_i$, and we denote $\vec{r}_{ij} := \vec{r}_i - \vec{r}_j$. We assume that the force $\vec{F}_{ij}$ exerted by the spring between masses m_i and m_j on mass m_j is central, and we denote $\mathcal{E}_{ij}$ the resulting elastic potential energy stored in the spring. These are defined as

$$\vec{F}_{ij} = -f_{ij}(\|\vec{r}_{ij}\|)\frac{\vec{r}_{ij}}{\|\vec{r}_{ij}\|}, \qquad \mathcal{E}_{ij} = \int_0^{\|\vec{r}_{ij}\|} f_{ij}(u)du, \tag{34}$$

where $f_{ij}(\cdot)$ is a linear or nonlinear characteristic. We also assume that some of the network springs follow Hooke's law, $f_{ij}(\|\vec{r}_{ij}\|) = k_{ij}\left(\|\vec{r}_{ij}\| - \ell_{ij}\right)$, where k_{ij} is the spring constant and ℓ_j is the spring's rest length, and we view these k_{ij}'s and ℓ_j's as the trainable weights of the system. The corresponding energy term is $\mathcal{E}_{ij} = \frac{1}{2}k_{ij}\left(\|\vec{r}_{ij}\| - \ell_{ij}\right)^2$. The state of the system is the vector of mass positions, $s = (\vec{r}_1, \vec{r}_2, \dots, \vec{r}_N)$, and the total elastic energy stored in the network is given by

$$\mathcal{E}_{\text{elastic}}(\vec{r}_1, \vec{r}_2, \dots, \vec{r}_N) = \underbrace{\sum_{\text{untrained } (i,j)} \mathcal{E}_{ij}}_{=\mathcal{E}_0} + \sum_{\text{trainable } (i,j)} \frac{1}{2}k_{ij}\left(\|\vec{r}_i - \vec{r}_j\| - \ell_{ij}\right)^2. \tag{35}$$

The partial derivatives of the energy function with respect to these weights are given by

$$\frac{\partial \mathcal{E}_{\text{elastic}}}{\partial k_{ij}} = \frac{1}{2}\left(\|\vec{r}_i - \vec{r}_j\| - \ell_{ij}\right)^2, \qquad \frac{\partial \mathcal{E}_{\text{elastic}}}{\partial \ell_{ij}} = k_{ij}\left(\ell_{ij} - \|\vec{r}_i - \vec{r}_j\|\right). \tag{36}$$

As in the Ising model, the energy function is separable, so the learning rule is local. This example also illustrates that EP is agnostic to the characteristics f_{ij} of untrained springs: no knowledge of $\mathcal{E}_0$ is required. A difference with the Ising model is that the space of possible network configurations, $\mathbb{R}^{3N}$, is continuous and infinite, while the space of configurations of the Ising model is discrete and finite.

An experimental realization of an elastic network that learns using CL was performed by Altman et al. [2024]. In their implementation, they used the spring rest lengths ℓ_{ij} as trainable weights, while keeping the spring constants fixed (untrained).

B Concepts of Quantum Mechanics

Next, we present the concepts of quantum mechanics that are necessary to extend EP to quantum systems. In particular, we present the variational principle of quantum mechanics that enables this extension (Lemma 3).

The **state vector** of a quantum system, denoted $|\psi\rangle$, belongs to a complex vector space $\mathcal{H}$ equipped with an inner product $\langle\cdot|\cdot\rangle$ (specifically, a Hilbert space). For simplicity of presentation, we assume here that $\mathcal{H}$ is finite-dimensional with dimension d. The system's **Hamiltonian**, $\widehat{H}$, is a linear operator acting on the Hilbert space, $\widehat{H} : \mathcal{H} \to \mathcal{H}$, with the property of being self-adjoint. Due to $\widehat{H}$ being self-adjoint, its eigenvalues are real. We denote the eigenvectors of $\widehat{H}$ as $|\psi_0\rangle$, $|\psi_1\rangle$, ..., $|\psi_{d-1}\rangle$, and the associated eigenvalues as $E_0 \leq E_1 \leq \ldots \leq E_{d-1}$, such that:

$$\widehat{H}|\psi_k\rangle = E_k|\psi_k\rangle, \qquad 0 \leq k \leq d-1. \tag{37}$$

Eq. (37) is known as the time-independent Schrödinger equation. The eigenvectors $|\psi_k\rangle$ are also called the **eigenstates** of the Hamiltonian, and their associated eigenvalues E_k are the energy levels. The eigenstate $|\psi_0\rangle$ with the lowest energy level is the **ground state**.

In quantum mechanics, a measurable physical quantity is represented by a self-adjoint operator, $\widehat{O} : \mathcal{H} \to \mathcal{H}$, called an **observable**. The set of possible outcomes of measuring $\widehat{O}$ is the set of eigenvalues of $\widehat{O}$, denoted o_0, o_1, ..., o_{d-1}, which are real due to the self-adjoint property. A peculiar aspect of quantum mechanics is that measurement outcomes are inherently probabilistic. When the system is in state $|\psi\rangle$, the probability of obtaining outcome o_k upon measuring $\widehat{O}$ is given by the Born rule, $p_k = |\langle o_k|\psi\rangle|^2$, where $|o_k\rangle$ is the eigenstate associated with o_k, i.e. such that $\widehat{O}|o_k\rangle = o_k|o_k\rangle$. The **expectation value** of a measurement of $\widehat{O}$ when the system is in state $|\psi\rangle$ is denoted $\langle\widehat{O}\rangle_\psi$ and calculated as $\langle\widehat{O}\rangle_\psi = \sum_{k=1}^{d} p_k o_k$. Using the spectral theorem for self-adjoint operators, it can be shown that this expectation value rewrites

$$\langle\widehat{O}\rangle_\psi = \langle\psi|\widehat{O}|\psi\rangle. \tag{38}$$

In statistical terms, the expectation value represents the average result of a large number of measurements of the observable $\widehat{O}$ performed on the system in state $|\psi\rangle$.

The Hamiltonian $\widehat{H}$ is an example of an observable, with possible measurement outcomes being the energy levels E_0, E_1, ..., E_{d-1}. The central result that allows us to transpose EP to quantum systems is the following variational formulation of the Hamiltonian's eigenstates (Lemma 3, proved in Appendix **??**). It tells us that the Hamiltonian's expectation value and the ground state can be viewed as EP's 'energy function' and 'equilibrium state', respectively.

Lemma 3. *The ground state $|\psi_0\rangle$ achieves the minimum of the Hamiltonian's expectation value:*

$$|\psi_0\rangle = \underset{\psi\in\mathcal{H}, \|\psi\|=1}{\arg\min} \langle\psi|\widehat{H}|\psi\rangle. \tag{39}$$

More generally, the eigenstates of the Hamiltonian $\widehat{H}$ are the critical points of the Rayleigh quotient $\psi \mapsto \frac{\langle\psi|\widehat{H}|\psi\rangle}{\langle\psi|\psi\rangle}$.

Another peculiar aspect of quantum mechanics is that the act of measuring an observable usually changes the system's state. Specifically, upon measurement of an observable $\widehat{O}$, if the outcome is o_k, then the system's state $|\psi\rangle$ instantaneously "collapses" to the eigenstate $|o_k\rangle$ corresponding to eigenvalue o_k. This principle, known as **state collapse**, implies that measuring $\widehat{O}$ a second time immediately after the first measurement will yield the same outcome o_k and leave the state $|\psi\rangle = |o_k\rangle$ unchanged, in accordance with the Born rule ($p_k = \langle\psi|o_k\rangle = \langle o_k|o_k\rangle = 1$). However, state collapse has another consequence that does not exist in classical mechanics. Suppose we want to measure two observables $\widehat{O}$ and $\widehat{P}$ in state $|\psi\rangle$, aiming to obtain (unbiased) estimates of both $\langle\widehat{O}\rangle_\psi$ and $\langle\widehat{P}\rangle_\psi$. Since, in general, the system is no longer in state $|\psi\rangle$ after measuring $\widehat{O}$, we must first reset the system to state $|\psi\rangle$ before measuring $\widehat{P}$. Nonetheless, there is one notable case where it is legitimate to measure $\widehat{O}$ and $\widehat{P}$ successively without resetting the state of the system between the two measurements: this is when the two observables **commute**, i.e.

$$\widehat{O}\widehat{P} = \widehat{P}\widehat{O}. \tag{40}$$

In this case, the two operators $\widehat{P}$ and $\widehat{O}$ are simultaneous diagonalizable. Assuming for simplicity that the eigenvalues of $\widehat{O}$ are all distinct, this means that the eigenstates $|o_0\rangle$, ..., $|o_{d-1}\rangle$ of $\widehat{O}$ are also eigenstates of $\widehat{P}$. Therefore, the probability of collapsing to any eigenstate $|o_k\rangle$ is the same under each observable (given by the Born rule, $p_k = |\langle o_k|\psi\rangle|^2$), and subsequent measurements of either $\widehat{O}$ or $\widehat{P}$ will leave the state unchanged. This allows for successive measurements of $\widehat{O}$ and $\widehat{P}$ to obtain unbiased estimates of their expectation values in the initial state $|\psi\rangle$, without resetting the system between the measurements.