

Many-body localized hidden generative models

Weishun Zhong,^{1,2,*} Xun Gao,³ Susanne F. Yelin,³ and Khadijeh Najafi⁴

¹*Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

²*IBM Quantum, MIT-IBM Watson AI Laboratory, Cambridge, Massachusetts 02142, USA*

³*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

⁴*IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, New York 10598, USA*



(Received 18 October 2022; revised 28 June 2023; accepted 17 September 2024; published 16 October 2024)

Quantum generative models hold the promise of accelerating or improving machine learning tasks by leveraging the probabilistic nature of quantum states, but the successful optimization of these models remains a difficult challenge. To tackle this challenge, we present a new architecture for quantum generative modeling that combines insights from classical machine learning and quantum phases of matter. In particular, our model utilizes both many-body localized (MBL) dynamics and hidden units to improve the optimization of the model. We demonstrate the applicability of our model on a diverse set of classical and quantum tasks, including a toy version of MNIST handwritten digits, quantum data obtained from quantum many-body states, and nonlocal parity data. Our architecture and algorithm provide novel strategies of utilizing quantum many-body systems as learning resources and reveal a powerful connection between disorder, interaction, and learning in quantum many-body systems.

DOI: [10.1103/PhysRevResearch.6.043041](https://doi.org/10.1103/PhysRevResearch.6.043041)

I. INTRODUCTION

The computational power of quantum processors is the subject of considerable amount of recent research, in particular with regard to scaling and a potential quantum advantage [1–7]. While the advent of a fully error corrected quantum computer requires yet another milestone, the immediate application of noisy quantum hardware with a clear advantage over classical computation becomes even more crucial. In this regard, the interface of quantum computing and machine learning has been increasingly brought into focus. For example, the rise of hybrid variational algorithms, such as variational quantum eigensolvers (VQE) [8] and the quantum approximate optimization algorithm (QAOA) [9], which use a parametrized quantum circuit as variational ansatz and optimize the parameters classically, has been considered particularly promising as they aim to obtain heuristic and approximate solutions.

However, the exponential dimension of the Hilbert space and the random characteristics of parametrized quantum circuits makes their training very challenging due to the existence of barren plateaus [10]. More recently, yet another approach to quantum machine learning has emerged, which is known as brain-inspired [11–15]. One interesting category consists of quantum reservoir computing (QRC) where a fixed

reservoir geometry scrutinizing the unitary dynamics of an interacting quantum system allows versatile machine learning tasks [16–19]. While QRCs have shown many advantages, they are mainly appropriate for discriminative tasks such as classification or regression.

The goal of generative models, however, is to learn an unknown data probability distribution p_{data} in order to subsequently sample from p_{data} and thus generate new and previously unseen data. Such tasks can, for example, be performed by the recently introduced quantum generative models known as Born machines [20,21]. Born machines for many-body problems have early on shown to be successful in conjunction with tensor network state ansätze. The elements of these matrix product states or tree tensor networks and their bond dimensions can be optimized during training to effectively approximate p_{data} [20–22]. While Born machines have also been tested with parameterized quantum circuits [23], we address here the question of whether there are other quantum many-body states that can be used as ansatz for Born machine to any advantage.

Quantum many-body systems display many phases in the presence of disorder, in particular, the break-down of thermalization and thus localization of the wave function in the so-called *many-body localized* (MBL) phase. Here, emergent integrals of motion can be utilized as quantum memories [24]. The failure of such systems to anneal [25] has inspired their use in QRC [19] for learning tasks, with particular enhancement close to the phase transition [26].

Here, we extend quantum-inspired generative models into the MBL phase, and introduce a hidden architecture to increase the representation power of our generative model. While many recent work has also quantum generative models in the MBL phase [27], using a similar quenched approach, our work differs in the hidden architecture and the characterization of learnability and expressibility.

*Present address: School of Natural Sciences, Institute for Advanced Study, Princeton, NJ, 08540, USA; Contact author: wszhong@ias.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. Illustration of the MBL hidden Born machine. (a) XXZ spin chain in 1D with periodic boundary condition. The faded color spins are the hidden units h_i , and the solid color spins are the visible units v_i . (b) The probability distribution of finding individual states in the z -basis represents the model distribution for the generative model, which is coded as normalized pixel values of an image. (c) An illustration of the loss landscape defined by our hidden MBL generative model. The training is done by optimizing disorder configurations in the Hamiltonian during each quantum quench, which is then used to evolve the initial state $|\psi_0\rangle$ over successive layers of quenches toward a final state $|\psi_f\rangle$ which gives rise to the desired model distribution.

In this paper, we introduce the new quantum generative model, the hidden Born machine in Sec. II and describe our training algorithm in Sec. IV. To illustrate the effect of hidden units, we introduce the randomly driven Born machine in Sec. IV B and compare its performance with the hidden Born machine by learning patterns of MNIST hand written digits. Next, in Sec. V, we investigate the learning power of the hidden Born machine both in the thermal phase and the MBL phase, and numerically show that the thermal phase fails to learn data obtained from quantum systems either in MBL or in thermal phase. Tracking von Neumann entanglement entropy and Hamming distance during training suggests that localization is crucial to learning. In Sec. V C we further show that while the hidden Born machine trained in the MBL phase is able to capture the underlying structure of the parity data, a hidden Born machine trained in the Anderson localized phase fails to do so, shedding light on the fact that the interplay between interaction and disorder plays an important role in learning. Finally, we conclude and discuss possible direction for future works.

II. HIDDEN BORN MACHINES

Born machine [28–32] is a generative model that parameterized the probability distribution of observing a given configuration $\mathbf{z}$ of the system according to the probabilistic interpretation of its associated quantum wave function $\psi(\mathbf{z})$,

$$p_{\text{Born}}(\mathbf{z}) = \frac{|\psi(\mathbf{z})|^2}{\mathcal{N}}, \quad (1)$$

where $\mathcal{N} = \sum_{\mathbf{z}} |\psi(\mathbf{z})|^2$ is the overall normalization of the wave function. Note that $\mathcal{N}$ is only required in tensor network

ansätze but not in physical systems. Training of the Born machine is done by minimizing the discrepancy between the model distribution $p_{\text{Born}}(\mathbf{z})$ and the data distribution $q_{\text{data}}(\mathbf{z})$.

In the language of Boltzmann machine [33–35], the units that are used for generating configurations are called “visible.” Meanwhile, adding “hidden” units prove to be a powerful architecture for the Boltzmann machine as it provides a way to decouple the complex interaction among the visible units at the expense of introducing interaction between the hidden and the visible units [36,37]. In Eq. (1), all units of the system are used to generate configurations that are compared against data and therefore all units are visible. In a similar spirit, we introduce hidden units to the Born machine by defining the probability distribution of observing a given visible spin configuration $\mathbf{z}$ to be its expectation value in z basis after tracing out the hidden units,

$$p_{\text{hidden}}(\mathbf{z}) = \text{Tr} \rho_{\text{vis}} \Pi_{\mathbf{z}}, \quad (2)$$

where

$$\rho_{\text{vis}} = \text{Tr}_h |\psi\rangle\langle\psi|, \quad (3)$$

is the reduced density matrix for the visible units, and $\Pi_{\mathbf{z}} = |\mathbf{z}\rangle\langle\mathbf{z}|$ is the projection operator onto the z basis of the visible part of the system [see Fig. 1(a) for an illustration of our model]. Note that normalization is implicit in Eq. (3) for ρ_{vis} to be a density matrix.

The hidden units offer two main advantages:

(1) when traced out, the hidden units provide an optimization advantage by acting as an effective heat bath for the visible units, in such a way that the systems are less prone to get trapped in local minima;

(2) the hidden Born machine in Eq. (2) offers expressive power advantage over the basic Born machine in Eq. (1).

We numerically confirm the mechanism in (1) by comparing the performance between the hidden Born machine and a regular Born machine with an artificially introduced heat bath in Sec. IV B. For (2), the hidden units provide additional degrees of freedom to parametrize generic probability distributions in an efficient way (see more in Appendix A).

III. MANY-BODY LOCALIZATION PHASE

Previously, different ansätze for $|\psi\rangle$ has been introduced for the Born machine, notably tensor networks states and states prepared by both digital quantum circuits and analog quantum many-body systems [20,21,23,27]. In this paper, we will be adopting the latter approach, and focus on a specific type of quantum many-body systems that admits a many-body localization (MBL) phase. In the following, we show that the MBL phase has two prominent effects on the training dynamics.

(1) Stabilizing training trajectories. On one hand, the effective heat bath from the hidden units assists the system in escaping local minima. On the other hand, this noise might cause the wave function to wander wildly around the Hilbert space, making it difficult for the system to find optimal solutions. Localization prevents such chaotic behavior and enhances the trainability of the model, as illustrated in Fig. 5.

(2) Resemblance to associative memory. Our quantum generative model exhibits memory characteristics such as pattern recognition ability, similar to those observed in associative memory architectures inspired by physical systems [38–41]. We illustrate this in Fig. 6.

Many-body localized ansätze

It is generally believed that, thermalization in quantum system wipes out the microscopic information associated with the initial state. Even in the case of closed quantum system, the information of initial state quickly spreads throughout the entire system, implying that no local measurements can retrieve those information [42,43]. However, it is known that strong disorder leads to localization, preventing the system to thermalize. Furthermore, the localization manifests itself in the form of memory associated with the lack of transport. While the localization in the presence of strong disorder was first introduced in noninteracting systems by Anderson [44], more recently, it was shown that the localization and break down of thermalization can also happen in strongly interacting systems, leading to new dynamical of phase of matter known as many-body localization (MBL) [45,46].

In the MBL phase, eigenstates of the system do not satisfy eigenstate thermalization hypothesis (ETH) and the wave functions become localized in the Hilbert space. Such ergodicity breaking renders the system to retain memory of its initial state, and offers advantage in controlling and preparing desired quantum many-body states and has been also realized experimentally [47]. The XXZ model of quantum spin chain is well-known to develop a MBL phase when the disorder strength exceeds the MBL mobility edge [48].

We perform numerical simulation with the XXZ-Hamiltonian $\hat{\mathcal{H}}_{\text{XXZ}}$ defined as

$$\hat{\mathcal{H}}_{\text{XXZ}} = \sum_i^{L-1} J_{xy} (\hat{S}_i^x \hat{S}_{i+1}^x + \hat{S}_i^y \hat{S}_{i+1}^y) + \sum_i^{L-1} J_{zz} \hat{S}_i^z \hat{S}_{i+1}^z, \quad (4)$$

where $\hat{S}_i^\alpha$ ($\alpha \in \{x, y, z\}$) are Pauli spin 1/2 operators acting on spins $i \in 1, \dots, L$, and $L = L_v + L_h$ consists of L_v visible units and L_h hidden units. J_{xy} and J_{zz} are couplings in the xy plane and z direction, respectively. Then, we consider a series of M quenches $\hat{\mathcal{H}}_{\text{quench}}(\Theta_m)$ in the z direction:

$$\hat{\mathcal{H}}_{\text{total}} = \hat{\mathcal{H}}_{\text{XXZ}} + \hat{\mathcal{H}}_{\text{quench}}(\Theta_m), \quad (5)$$

where $\hat{\mathcal{H}}_{\text{quench}}(\Theta_m) = \sum_i h_i^m \hat{S}_i^z$ and we have denoted the tunable parameters in the system collectively as $\Theta_m = \{h_i^m\}$. During each quench m , h_i^m are drawn i.i.d. from the uniform distribution over the interval $[-h_d, h_d]$, where h_d is the disorder strength. Notice that when $J_{zz} = 0$, the model reduces to noninteracting XY model with random transverse field exhibiting single particle localization. Once we turn on the J_{zz} interaction, the spins couple via Heisenberg interaction and MBL phase emerges when $h_c \sim 3.5$ (for $J_{zz} = J_{xy} = 1$) [48–50]. See more details in Appendix B.

In Sec. IV, we will explain the training algorithm under the time evolution implied by series of quenches in $\hat{\mathcal{H}}_{\text{total}}$, and learning through optimizing the values of disordered field h_i^m at each site.

IV. TRAINING OF HIDDEN MBL BORN MACHINE

A. Learning algorithm

The basic idea behind the training of Hidden Born machine is the following: given target distribution q_{data} , and a loss function $\mathcal{L}(p_{\text{model}}, q_{\text{data}})$ that measures the discrepancy between model distribution and data distribution, training of the MBL Born machine is achieved through time-evolving the system with the Hamiltonian in Eq. (4), then optimizing Θ_m over N different disorder realizations for each quench m . After obtaining the final state at the M th quench, we evaluate the model distribution of the MBL hidden Born machine from Eq. (2) and use it as our generative model [see Fig. 1(c) for an illustration of the learning process]. We use maximum mean discrepancy (MMD) loss as our loss function:

$$\mathcal{L}_{\text{MMD}} = \left\| \sum_x p(x) \phi(x) - \sum_x q(x) \phi(x) \right\|^2, \quad (6)$$

where x is samples for estimating the MMD loss, and $\phi(x)$ are kernel functions that one can choose (see more details in Appendix C).

The learning algorithm is summarized by the pseudocode in algorithm1 and illustrated in Fig. 2.

Given the reduced density matrix ρ_M of the L_v visible spins at the final layer of the quench $m = M$, we compute the model distribution from Eq. (2), which gives the probability $p_{\text{model}}(\mathbf{z})$ of finding each of the 2^{L_v} basis states $\mathbf{z}$ in the visible part of the system. For learning image data, we then interpret the probabilities as pixel values (normalized to be within 0 and 1), and reshape it into an image of size $2^{L_v/2} \times 2^{L_v/2}$ [see Fig. 1(b)]. For quantum data, we interpret these probabilities

ALGORITHM 1. Training of MBL hidden Born machine.

Initialize the system in some initial state $|\psi(\Theta_{m=0})\rangle \equiv |\psi_0\rangle$
 and choose $\Theta_0 = \mathbf{0}$;
while $m < M$ **do**
 while $n < N$ **do**
 Sample $\Theta_m^{(n)}$ uniformly from the interval $[-h_d, h_d]$;
 Time-evolve the state $|\psi_{m+1}^{(n)}\rangle = \hat{U}(\Theta_m^{(n)})|\psi_m\rangle$ with
 $\hat{U} = \hat{T} \exp(-i \int_0^T dt \hat{H}_{\text{total}})$;
 Trace out the hidden units
 $\rho_{m+1}^{(n)} = \text{Tr}_h |\psi_{m+1}^{(n)}\rangle \langle \psi_{m+1}^{(n)}|$;
 Compute $\mathcal{L}(\Theta_m^{(n)})$ from
 $p_{\text{hidden}}^{(n)}(\mathbf{z}) = \text{Tr} \rho_{m+1}^{(n)} \Pi_Z$;
 $n \leftarrow n + 1$
 end while
 $\Theta_m = \text{argmax}_{\Theta_m^{(n)}} \mathcal{L}(\Theta_m^{(n)})$;
 $|\psi_{m+1}\rangle = \hat{U}(\Theta_m)|\psi_m\rangle$
 $m \leftarrow m + 1$
end while
 Denote the training outcome as $p_{\text{model}}(\mathbf{z}) = \text{Tr} \rho_M \Pi_Z$.

as measurement outcomes obtained from the quantum state. For more details, see Appendix D.

B. Randomly driven MBL Born machine

In classical machine learning, stochasticity is found to have the effect of smoothing out loss landscape and helps to avoid local minima [36,51,52]. When introducing the hidden Born machine in Eq. (2), the hidden units are traced out and effectively act as a heat bath for the remaining visible units and provide a source for stochasticity. To understand the source of the learnability advantage provided by the hidden units, we construct a Born machine with an artificially introduced heat bath in this section. We then compare its performance with that of the hidden Born machine, which operates within an effective heat bath induced by the hidden units.




FIG. 2. Schematics of the learning algorithm as in algorithm 1. At the m th quench, we independently evolve N copies of the state $|\psi_m\rangle$ with different time-evolution operators $\Theta_m^{(n)}$ sampled from the same distribution. At the $(m+1)$ th quench, we pick the $|\psi_m^{(n)}\rangle$ with the lowest loss value [based on the loss function Eq. (6)] from the previous quench as our new starting point and evolve again. As we repeat this process, the learning resembles a directed random walk in the Hilbert space.




FIG. 3. Learning toy MNIST digit patterns. The top two rows are different data instances q_{data} in our toy MNIST digit patterns dataset. The bottom two rows are the corresponding learning outcome p_{model} from our MBL hidden Born machine (each digit trained separately).

Let's consider the Hamiltonian Eq. (4) with applied external random drives $\hat{H}_{\text{RD}}$ in the x direction (we can also apply random drives in the xy plane and the result will be similar),

$$\hat{H}_{\text{RD}}(t) = \sum_i d_i^m(t) \hat{S}_i^x. \quad (7)$$

To model the heat bath, we would like $\{d_i^m(t)\}$ to be like a white noise,

$$\langle d_i^m(t) d_i^m(0) \rangle = 2D\delta(t), \quad (8)$$

where D is the amplitude of the white noise and is proportional to the temperature of the bath. In the simulation, we split the driven interval T into intervals of auto-correlation time τ , and require that Eq. (8) holds for $t > \tau$. Outside of this correlation time, $d_i^m(t)$ is drawn i.i.d. from $\mathcal{N}(0, \sqrt{2D})$. We refer to this model as randomly driven Born machine (RDBM).

To illuminate on the learning power of the hidden Born machine, here, we compare the three models: the basic Born machine (BM), the RDBM, and the hidden Born machine (hBM). We task all three models with a toy dataset constructed from the images of MNIST dataset [53] (downsampled to 2^{L_v} pixels). Our toy dataset consists of mean pixel values across all different styles within a single type of MNIST digit, see 'target patterns' in Fig. 3 (also see Appendix D).

We perform the training of the hidden Born machine using the algorithm described in algorithm 1, and show the corresponding learning outcomes in Fig. 3. Our results indicate our hidden model can learn different patterns of MNIST digits accurately (the results of basic BM and RDBM are omitted). In Fig. 4, we compare the performance of three models by plotting the MMD loss on the toy MNIST dataset as a function of the number of quenches, m . We find that both the RDBM and hBM outperform the BM, with the RDBM and hBM converging to a similar terminal loss value (although the RDBM slightly underperforms the hBM). This observation suggests that the stochasticity introduced by the heat bath indeed improves learning, and the learnability advantage of hidden units is of a thermal nature. However, the RDBM is more computationally intensive. Modeling the white noise as




FIG. 4. Model comparisons. (a) Basic Born machine (BM), randomly driven Born machine (RDBM), and hidden Born machine (hBM). Log-MMD loss as a function of quench layer number m . The solid curves are averaged over 100 different realizations, with one standard deviation included as the shades. The hidden Born machine achieves the lowest MMD loss throughout and at the end of the training.

expressed in Eq. (8) requires discretizing the quench time T into small intervals of τ , which generally increases the total runtime by a factor of T/τ . Therefore the hBM is more efficient than the RDBM while being more trainable than the basic BM.

V. LEARNABILITY IN DIFFERENT PHASES

We have already seen that the hidden Born machine in the MBL phase can properly learn the toy MNIST dataset (Fig. 3). An important question arises as to whether learning can happen in the thermal phase. In the thermal phase, information spreads throughout the system, which makes it difficult to extract. In the quenched approach as in Eq. (5), the state of the system in the thermal phase changes wildly between successive quenches and effectively only parameters in the last layer of the quench would be trained. In contrast, as the system becomes more disordered and enters the MBL phase, the breakdown of thermalization and emergence of lo-

cal integrals of motion leads to local memory, which is useful for directing the state toward a target corner of the Hilbert space [Fig. 1(c)]. We aim to understand the effect of disorder on learning by comparing the learning ability of the hidden Born machine in the MBL and thermal phases.

In Fig. 5(a), we show a log-log plot of the final quench layer MMD loss on the toy MNIST dataset as a function of disorder strength h_d . By varying the disorder strength, the system in Eq. (4) can exhibit both a thermal phase (denoted as ETH) and an MBL phase depending on whether the disorder strength exceeds the critical value $h_c \sim 3.5$ (for $J_{zz} = J_{xy} = 1$). We observe that the loss value has a significant change at the transition from the thermal phase (corresponding to $h_d < 3.5$ indicated by pink shade) into the MBL phase (indicated by the blue shade). The relatively high value of MMD loss in the thermal phase indicates that the hidden Born machine fails to learn. As we increase the disorder, the MMD loss deep in the MBL phase decreases significantly, indicating better learning power of the MBL phase. We can attribute the better learning power in the MBL phase to the quantum memory and the emergent local integral of motions. In contrast to the thermal phase, the thermalization mechanism wipes out all the information from the initial conditions, as observed similarly in the case of quantum reservoir computing in the MBL phase [19].

To better quantify the learning mechanism in the MBL phase, we investigate the time evolution of quantities underlying MBL physics during the quenched steps. First, we investigate the Hamming distance (HD) defined as

$$\mathcal{D}(t) = \frac{1}{2} - \frac{1}{2L} \sum_i \langle \psi_0 | \sigma_i^z(t) \sigma_i^z(0) | \psi_0 \rangle, \quad (9)$$

which gives a measure of a number of spin flips concerning the initial state ψ_0 normalized by the length of chain L . It's expected that in a long time the HD approaches 0.5 in the thermal phase and decreases as one increases the disorder [54]. In Fig. 5(b), we show the trajectory of HD at the end of each quench $\mathcal{D}^m(t = T)$. We observe that, evolving in the thermal phase the HD fluctuates around the value of 0.5 as expected, while in the MBL phase, the HD reaches a lower value of




FIG. 5. Training hidden Born machine in thermal and MBL phases. (a) The terminal (at the final layer of quench) MMD loss of the hidden Born machine on the toy MNIST task is plotted as a function of disorder strength h_d . The results are averaged over 100 realizations and one standard deviation is included as shade. (b) Hamming distance concerning the initial state (normalized by L) as a function of quenches m . In the thermal phase, states change discontinuously over successive quenches, whereas in the MBL phase, states change gradually toward the target state that gives rise to the desired distribution. (c) Entanglement entropy per site as a function of quenches m , confirming that our system evolves under dynamics distinctive in the thermal/MBL phases.

about 0.33. The more significant fluctuations in the thermal phase indicate that the system retains little information about the most recent quench, and therefore is difficult to manipulate toward a target state that gives desired probability distribution. In contrast, the relatively small fluctuations in the MBL phase suggest that the system changes gradually between successive quenches and is more amenable to directed evolution by quenches.

One hallmark of MBL phase is the logarithmically slow growth of von Neumann entanglement entropy ($S_{\text{ent}}^m = -\text{Tr} \rho_m \ln \rho_m$) due to the presence of strong interaction. Notice that ρ_m is the reduced density matrix at quench m , obtained by tracing over the complementary part of the system with respect to the subsystem of interest. This can be considered as a slow dephasing mechanism implying that not all information of initial state survives [55–57]. To confirm that our system indeed evolves under MBL/thermal dynamics when trained in these two phases, in Fig. 5(c), we track the value of S_{ent}^m over different quenches. In the MBL phase, S_{ent}^m shows a quick saturation, while in the thermal phase the entanglement entropy changes significantly from successive quenches, a behavior expected from the thermal phase.

A. Pattern recognition

Pattern recognition has been implemented in a variety of analog classical systems ranging from molecular self-assembly to elastic networks [40,41,58–63]. It is interesting to ask whether quantum systems possesses similar power. In this section, we demonstrate the pattern recognition ability of the MBL hidden Born machine. Here, we take the same toy dataset of MNIST digit patterns as in Fig. 3. Each pattern $\xi^\mu \in [0, 1]^{2^{L_v}}$ is a (normalized) vector in the pixel space, where L_v is the length of the visible units, and $\mu = 1, 2, \dots, P$ denotes the pattern index. We encode the patterns into the hidden Born machine by setting $p_{\text{data}} = \sum_\mu \xi^\mu$ [64]. Again, we perform the training of the hidden Born machine using the algorithm in algorithm 1 (see first column of Fig. 6 for the learned patterns from p_{model}). After training, we obtain the target final state $|\psi_M\rangle$, along with a series of unitaries $\{\hat{U}(\Theta_m) \equiv \hat{U}_m\}_{m=0}^M$ that defines the entire history of intermediate states during successive quenches, $|\psi_m\rangle = \prod_{i=0}^m \hat{U}_{m-i} |\psi_0\rangle$, which upon tracing out hidden units becomes intermediate model distributions, $p_m = \text{Tr} \text{Tr}_h |\psi_m\rangle \langle \psi_m| \Pi_Z$. Now given a partially corrupted pattern $\tilde{\xi}^\mu$ and the state $|\tilde{\psi}^\mu\rangle$ that gives rise to this corrupted pattern, $|\tilde{\psi}^\mu(z)|^2/\mathcal{N} = \tilde{\xi}^\mu$ (see the second column of Fig. 6 for examples of corrupted patterns), we can identify the “closest” intermediate state $|\psi_{m^*}\rangle$ where $m^* = \text{argmax}_m \text{MMD}(\tilde{\xi}^\mu, p_m)$. Then we apply unitary time evolution to the corrupted state $|\tilde{\psi}^\mu\rangle$ using the series of learned unitaries starting from m^* and obtain the ‘retrieved’ state $|\hat{\psi}^\mu\rangle \equiv \prod_{i=0}^{M-m^*} \hat{U}_{M-i} |\tilde{\psi}^\mu\rangle$. We can then compute the corresponding retrieved pattern as $\hat{\xi}^\mu \equiv \text{Tr} \text{Tr}_h |\hat{\psi}^\mu\rangle \langle \hat{\psi}^\mu| \Pi_Z$ (see last column of Fig. 6 for the retrieved patterns). As shown in the top row of Fig. 6, in the case of a single pattern (a digit 0), the MBL hidden Born machine can retrieve a complete pattern from a corrupted pattern (a partially corrupted digit 0). As shown in the bottom row of Fig. 6, in the case of multiple patterns (a superposition of 0 and 1), the MBL hidden Born




FIG. 6. Pattern recognition task by the MBL hidden Born machine. Given a corrupted pattern $\tilde{\xi}^\mu$ and its corresponding corrupted state $|\tilde{\psi}^\mu\rangle$, we find the quench layer number m^* such that the intermediate model distribution p_{m^*} resembles the corrupted pattern most. Then we time-evolve $|\tilde{\psi}^\mu\rangle$ with the series of learned unitaries $\hat{U}_i$ starting from m^* to obtain the retrieved state $|\hat{\psi}^\mu\rangle$, from which we can then obtain the retrieved pattern $\hat{\xi}^\mu$. Top row: after learning a single pattern (digit 0), a complete 0 can be retrieved from a partially corrupted 0. Bottom row: after learning multiple patterns (superposition of digit 0 and 1), complete 0 or 1 can be selectively retrieved from partially corrupted 0 and 1, respectively.

machine can selectively retrieve complete patterns (0 or 1) based on the input corrupted pattern [65].

B. Learning quantum dataset

We have demonstrated the power of MBL Born machine in learning classical data of the toy MNIST digit patterns, now we explore the ability of the MBL Born machine in learning data obtained from measurements of quantum states. While quantum state tomography is the standard method for state reconstruction, it becomes a daunting task as the system size increases. In this respect, quantum machine learning has shown great success in learning quantum states from a limited amount of data [22,66–71]. In this section, we use the hidden Born machine to learn data obtained from quantum many-body states prepared by Eq. (5) subject to a single layer of quench, but with disorder strengths h_d different from the phases that the hidden Born machine is trained in.

In Fig. 7, we demonstrate the learning ability of Born machine in the thermal and MBL phase. In Fig. 7 left/right, we compare the measurement outcome sampled from the exact simulation q_{data} in MBL/thermal phase (denoted as ETH) (shown in purple), with those learned via hidden Born machine trained in MBL phase (shown in blue). In the insets, we show the classical fidelity between the model distribution p and data distribution q , $F(p, q) = (\sum_i \sqrt{p_i q_i})^2$. We see that the hidden Born machine trained in MBL phase can capture the underlying probability distribution obtained from both the MBL and thermal phases with high fidelity (~ 0.98), while the hidden Born machine trained in the thermal phase fails to learn either. Notice that to learn the quantum state, one needs




FIG. 7. Learning quantum dataset. Left/right: MBL hidden Born machine trained in MBL phase learns the probability distribution corresponding to an MBL/thermal (denoted as ETH) target state. Insets: classical fidelities between the model and the data distributions. Model trained in the MBL phase has better learning capability than model trained in the thermal phase.

to measure the informationally complete basis as reported in Ref. [22].

C. Learning parity dataset

In the previous sections, we have discussed the role of localization and emergent memory in learning various datasets, however, the role played by interaction in the many-body localized phase remains unclear. To shed light on the role of interaction and its interplay with the disorder, here, we investigate the power of MBL phase in the learning parity dataset and compare it with both thermal and Anderson localized phase which can be obtained by setting $J_{zz} = 0$ in Eq. (4). Here, we consider the even parity dataset, which is defined as set of bit-string $(b_1, b_2, \dots, b_L)$ of length L with $b_i \in \{0, 1\}$, such that the parity function $\Pi(b_1, b_2, \dots, b_L) := \sum_{i=1}^L b_i \mod 2$ is equal to 0. While this is a classical bitstring, it appears as a measurement outcome of particular quantum observables in a certain basis such as the measurement outcome of GHZ state in the x basis.

Previous studies have indicated challenging learning on this dataset, in particular, training the Born machine based on MPS with gradient descent optimization schemes has encountered failures [72], while quantum-inspired optimization schemes such as density matrix learning has shown great success with the caveat in their scaling [73]. Here, we investigate the power of our hidden Born machine across various phases in learning the parity dataset. Our numerical results (Fig. 8) demonstrate the interesting fact that both the MBL phase and Anderson localized phase show better performance compared to the thermal phase. The better learning performance in these two phases suggests that the emergence of the integral of motion and memory plays an important role in learning. We further notice that the MBL phase has a better performance even though the Anderson localized phase is known to have better memory. In the latter, the strong localization prevents the transport of information across the system, leading to a lesser learning power. While the value of fidelity around $F_{MBL} = 0.75$ is not too high, reflecting the hardness of learning the parity dataset, our MBL hidden Born machine still shows a better performance compared to MPS Born machine which was reported a fidelity of $F_{MPS} = 0.48$ [72]. Our numerical results indicate that both disorder and in-

teraction are crucial for the successful learning of our model. The MBL phase enhances the learnability of the hidden Born Machine because it incorporates these two properties. However, given that the XXZ model is not the only one with such characteristics, we anticipate that the hidden Born machine, when evolved under other Hamiltonians with disorder and interaction—such as spin-glass Hamiltonians [74]—would exhibit similar learnability. Intuitively, one can think of the Hidden Born Machine as carving a path through the Hilbert space, capitalizing on both local memory and interaction to reach the target state.

VI. CONCLUSION AND OUTLOOK

In this work, we have introduced a novel quantum generative model, the hidden MBL Born machine as a powerful quantum-inspired generative model. Although parameterized quantum circuit has become one of the focal point in the realm of quantum machine learning, their training scheme poses many challenges as one is required to search in an exponential Hilbert space, which resembles finding a needle in a haystack [10]. While other variational algorithms such as QAOA offer a different scheme of finding solutions in Hilbert space which is akin to adiabatic computing, here, by utilizing unique properties of MBL phase such as localization and memory, we




FIG. 8. Learning parity dataset. Different bars in the horizontal axis correspond to the model trained in the MBL, thermal, and Anderson localized phases, respectively. The vertical-axis shows the classical fidelity of the model. A model trained in the MBL phase exhibits the highest fidelity despite the dataset is highly nonlocal. Comparing model performances in three phases suggests that both disorder and interaction are important for learning.

develop a Born machine evolving under MBL dynamics such that by optimizing over values of disorder at each site we can reach a desired target state in the Hilbert space.

In this work, we aimed to answer two key questions, namely, whether MBL phase and hidden units can be used as resource for learning, and what is the underlying mechanism of learning. By performing various numerical experiments in different phases, with and without hidden units, we show that successful learning relies on three key factors, hidden units, disorder, and interaction.

Our work opens up a new horizon in utilizing exotic quantum phases of matter as quantum-inspired generative models. While we have explored the role of disorder in the MBL phase, an immediate question that follows is whether another disordered quantum phase would be capable of learning, which is left for future work. In addition, although we show that our model contains states that cannot be simulated in polynomial time by classical computer, it remains an important question of designing efficient quantum algorithms to prepare such states. Similarly, it remains an open question whether the MBL hidden Born machine can learn other classically intractable distributions [32,75,76], such as the distribution generated in quantum supremacy experiments [2–6]. Furthermore, our quenched Born machine resembles specific adiabatic schedule, and whether we can utilize our model as quantum variational algorithm awaits further investigation. Although we have quantified the learning mechanism during the training by tracking both local and nonlocal quantities such as Hamming distance and entanglement entropy, more quantitative studies such as the existence of barren plateau and over-parametrization in the context of quantum kernel learning remain an important questions for future study [10,77,78].

ACKNOWLEDGMENTS

The authors would like to thank Nishad Maskara and Nicole Yunger Halpern for useful discussions. W.Z. acknowledges supports from IBM Research (Quantum) summer internship program. SFY would like to acknowledge funding from NSF (through CUA and Q-SeNSE) and AFOSR.

APPENDIX A: EXPRESSIBILITY OF MBL HIDDEN BORN MACHINE

1. Expressibility of hidden units

The hidden Born machine Eq. (2) generalizes the basic Born machine (BM) defined by Eq. (1), in the sense that the class of probability distributions expressible by the basic Born machine is a subset of that of the hidden Born machine.

Proposition 1. For the same set of visible spins $\mathbf{v}$, let $p_{\text{BM}}(\mathbf{z})$ denote the model distribution realized by the basic Born machine, and $p_{\text{hBM}}(\mathbf{z})$ denote the model distribution realized by the hidden Born machine, then $\{p_{\text{BM}}(\mathbf{z})\} \subseteq \{p_{\text{hBM}}(\mathbf{z})\}$.

Proposition 1 further suggests that the minimum achievable training loss for the hidden Born machine is less than or equal to that of the basic Born machine, a property that we will confirm numerically in Sec. IV.

In the following, we assume only that the visible and hidden part couple through an interaction term in the Hamiltonian.

Let's consider a basic Born machine consisting of only visible units $\mathbf{v} = \{v_i\}$, with Hamiltonian $\hat{\mathcal{H}}_v$. Now consider adding hidden units $\mathbf{h} = \{h_j\}$ to the system with Hamiltonian $\hat{\mathcal{H}}_h$ and the hidden units couple with the visible ones through an interaction Hamiltonian $\hat{\mathcal{H}}_{\text{int}}$. The full Hamiltonian can be written as

$$\hat{\mathcal{H}}_{vh}[\mathbf{v}, \mathbf{h}] = \hat{\mathcal{H}}_v[\mathbf{v}] + \hat{\mathcal{H}}_h[\mathbf{h}] + \hat{\mathcal{H}}_{\text{int}}[\mathbf{v}, \mathbf{h}], \quad (\text{A1})$$

where all the $\hat{\mathcal{H}}$'s in general can be time-dependent. Let's assume that the basic Born machine model is described by just the visible part of Hamiltonian in Eq. (A1), $\hat{\mathcal{H}}_{\text{BM}} = \hat{\mathcal{H}}_v(\Theta^{\text{BM}})$, and the hidden Born machine is described by the full Hamiltonian, $\hat{\mathcal{H}}_{\text{hBM}} = \hat{\mathcal{H}}_{vh}(\Theta^{\text{hBM}})$, where Θ^{BM} and Θ^{hBM} denotes the parameters in the Hamiltonian to be optimized during learning.

Proof. Let's denote the initial state for the BM as $|\psi_0^v\rangle \in \mathcal{H}_v$. Let $\hat{\mathcal{U}}_v = \hat{\mathcal{T}} \exp(-i \int_0^T dt \hat{\mathcal{H}}_v)$. Then, the final state of BM is $|\psi_f^v\rangle = \hat{\mathcal{U}}_v |\psi_0^v\rangle$. Choose an initial product state for the hBM, $|\psi_0^{vh}\rangle = |\psi_0^v\rangle \otimes |\psi_0^h\rangle \in \mathcal{H}_v \otimes \mathcal{H}_h$ for some $|\psi_0^h\rangle \in \mathcal{H}_h$. Choose Θ^{hBM} to be such that $\hat{\mathcal{H}}_v^{\text{hBM}} = \hat{\mathcal{H}}_v^{\text{BM}}$ and $\|\hat{\mathcal{H}}_{\text{int}}^{\text{hBM}}\| \gg \|\hat{\mathcal{H}}_{\text{int}}^{\text{BM}}\|$.

Then, we have

$$|\psi_f^{vh}\rangle \approx \hat{\mathcal{U}}_v |\psi_0^v\rangle \otimes \hat{\mathcal{U}}_h |\psi_0^h\rangle = |\psi_f^v\rangle \otimes |\psi_f^h\rangle \quad (\text{A2})$$

where we have defined $|\psi_f^h\rangle \equiv \hat{\mathcal{U}}_h |\psi_0^h\rangle$. With this choice, now $\rho_{\text{vis}} = \text{Tr}_h |\psi_f^{vh}\rangle \langle \psi_f^{vh}| = |\psi_f^v\rangle \langle \psi_f^v|$, and $p_{\text{hBM}}(\mathbf{z}) = \text{Tr} \rho_{\text{vis}} \Pi_{\mathbf{z}} = |\psi_f^v(\mathbf{z})|^2 = p_{\text{BM}}(\mathbf{z})$, where p_{BM} is automatically normalized ($\mathcal{N} = 1$) for physical systems as in our case. Therefore the class of probability distributions described by BM is contained in hBM. ■

2. Mapping XXZ chain into Ising model

There has been extensive studies on the expressive power of quantum models. In particular, quantum computational advantage for sampling problem has been proved (based on standard computational complexity assumptions) in a translation-invariant Ising model [79]. In particular, we show that the XXZ model in 2-dimension, with proper choice of disorder parameters, can be reduced to an Ising model that contains brickwork state that is classically intractable Ref. [79].

Proposition 2. The XXZ model in 2D subject to quench in z direction can be reduced to an Ising model.

We have recovered the case in Ref. [79]. Note that for the proof in Ref. [79] to work, one also need to initialize the system in all $|+\rangle$ states and subsequently perform all measurements in the x – basis. While our numeric are mostly restricted to the 1-dimension case as it can be studied by exact diagonalization, the XXZ model can be realized in any dimensions. This classically hard instance implies that our model cannot be simulated in polynomial time by a classical computer and therefore offers an advantage in its expressive power.




FIG. 9. Checkerboard partition of the 2D lattice. We apply π pulses at the $\mathcal{K}$ sites (shown in blue) to effectively turn off the interaction in the x and y directions upon time averaging.

Proof. In 2D,

$$\hat{\mathcal{H}}_{\text{XXZ}} = \sum_{\langle i,j \rangle} J_{xy} (\hat{S}_i^x \hat{S}_j^x + \hat{S}_i^y \hat{S}_j^y) + \sum_{\langle i,j \rangle} J_{zz} \hat{S}_i^z \hat{S}_j^z, \quad (\text{A3})$$

where the interactions are between nearest neighbors. During a quench $\hat{\mathcal{H}}_{\text{quench}}$ of duration t_m , we can divide the disorder into a time-dependent and a time-independent part,

$$h_i^m(t) = J_i^m(t) + B_i^m. \quad (\text{A4})$$

In the case of bipartite lattice, we can partition the vertices into two partitions, and denote the sites in one partition as $\mathcal{K} = \{k_1, k_2, \dots, k_{L/2}\}$ and another partition as $\mathcal{N} = \{n_1, n_2, \dots, n_{L/2}\}$ (assuming N even). For example, in the case of a square lattice, $\mathcal{K}$ and $\mathcal{N}$ correspond to the blue and white sites of the checkerboard coloring (Fig. 9).

For the set of $\mathcal{K}$ spins, we turn on a π pulse in the middle of the quench ($k \in \mathcal{K}$), such that:

$$J_k^m(t) = \begin{cases} 0 & \text{if } 0 \leq t < \frac{t_m}{2} \\ \frac{\pi}{2\Delta t} & \text{if } \frac{t_m}{2} \leq t < \frac{t_m}{2} + \Delta t \\ 0 & \text{if } \frac{t_m}{2} + \Delta t \leq t < t_m \end{cases}, \quad (\text{A5})$$

where $\Delta t \ll t_m$ is a short duration of time. With this choice of disorder, the time evolution operator reduces to

$$\hat{\mathcal{U}} = e^{-i\hat{\mathcal{H}}_{\text{XXZ}}\Delta t} (\prod_k i\hat{Z}_k) e^{-i\sum_k B_k^m \hat{Z}_k}. \quad (\text{A6})$$

Now the Pauli Z 's effectively flip the signs of the $\hat{S}^x \hat{S}^x$ and $\hat{S}^y \hat{S}^y$ terms in the XXZ-Hamiltonian, and upon integrating over the duration of a quench cancels out with the corresponding terms in first half of the quench. Therefore, after time evolution of a quench, the effective Hamiltonian is left with only Ising interactions,

$$\hat{\mathcal{H}}_{\text{eff}}^m = \sum_{\langle i,j \rangle} J_{zz} \hat{S}_i^z \hat{S}_j^z + \sum_i B_i^m \hat{S}_i^z. \quad (\text{A7})$$

■

APPENDIX B: MBL PHASE TRANSITION

In this section, we present the details of the numerical simulation of XXZ model [Eq. (4)] and confirm the thermal to MBL phase transition. We simulate the XXZ model using exact diagonalization methods provided by the QUSPIN package [80,81]. Throughout the paper, we use parameters $J_{xy} = J_{zz} = 1$.

One hallmark of the MBL phase is the Poisson distribution of level spacings in the eigenspectrum of the Hamiltonian




FIG. 10. Level statistics of $L = 16$ XXZ model subject to quenches in the z direction. The thermal phase (denoted as ETH) is simulated with $h_d = 0.1$, with resulting $\langle r_\alpha \rangle \approx 0.391$; the MBL phase is simulated with $h_d = 3.9$, with resulting $\langle r_\alpha \rangle \approx 0.529$. Results are averaged over 1000 different realizations.

[82–84]. The level statistics $Pr(r_\alpha)$ is defined as the normalized distribution of

$$r_\alpha = \frac{\min(\Delta_{\alpha+1}, \Delta_\alpha)}{\max(\Delta_{\alpha+1}, \Delta_\alpha)}, \quad (\text{B1})$$

where $\Delta_\alpha = E_{\alpha+1} - E_\alpha$ are the level spacings in the eigenspectrum. In Fig. 10, we show the level statistics in a simulation of $L = 16$ spins described by Eq. (5) subject to a single quench $M = 1$, at two different disorder strengths: $h_d = 0.1$ and 3.9 (the critical disorder strength is $h_c \sim 3.5$ for $J_{zz} = J_{xy} = 1$). We see that indeed the level statistics in the thermal phase obeys Wigner-Dyson statistics ($\langle r_\alpha \rangle \approx 0.391$), and in the MBL phase obeys Poisson statistics ($\langle r_\alpha \rangle \approx 0.529$), confirming the existence of thermal-MBL phase transition.

Another hallmark of MBL phase is the area law scaling of von Neumann entanglement entropy ($S_{\text{ent}} = -\text{Tr} \rho \ln \rho$), compared to the volume law scaling in the thermal phase. We numerically calculate the half-system entanglement entropy in the middle of the spectrum for the Hamiltonian in Eq. (5), and perform a scaling analysis for different L and different disorder strengths h (see Fig. 11). Our numerical results agrees with those reported in Ref. [48].

APPENDIX C: TRAINING MBL HIDDEN BORN MACHINE

1. MMD loss

Previously, KL divergence has been suggested for training MBL Born machine as a generative model [27]. However, KL divergence does not capture correlations within data, and suffers from infinities outside the support of data distribution. Moreover, KL divergence requires full knowledge of data distribution which is often unrealistic. To remedy these situations, the maximum mean discrepancy (MMD) loss has been proposed for training Born machines [28]. The MMD




FIG. 11. Scaling analysis of entanglement entropy. We plot the entanglement entropy per site S_{ent}/L as a function of system size L for different disorder strengths h . Volume law scaling in the thermal phase (small h) leads to constant S_{ent}/L , while area law scaling in the MBL phase (for large h) leads to decreasing S_{ent}/L .

loss measures the distance between model distribution p and target distribution q , by comparing their mean embeddings in the feature space, and one can use samples to estimate the loss. The (squared) MMD loss can be written as

$$\begin{aligned} \mathcal{L}_{\text{MMD}} &= \left\| \sum_x p(x)\phi(x) - \sum_x q(x)\phi(x) \right\|^2 \\ &= \mathbb{E}_{x,x' \sim p} k(x, x') + \mathbb{E}_{y,y' \sim q} k(y, y') \\ &\quad - 2\mathbb{E}_{x \sim p, y \sim q} k(x, y), \end{aligned} \quad (\text{C1})$$

where we have employed the kernel trick and write $k(x, y) = \phi(x)^T \phi(y)$. In our model, we use a Gaussian mixture kernel $k(x, y) = \frac{1}{c} \sum_{i=1}^c \exp(-\frac{1}{2\sigma_i^2} |x - y|^2)$ of four channels $c = 4$, with corresponding bandwidths $\sigma_i^2 = [0.1, 0.25, 4, 10]$. The bandwidths are chosen such that our Gaussian kernels are able to capture both the local features and the global features in the target distribution.

2. Parameters

In the training of our MBL hidden Born machine, we use $N = 6 + 2$ (6 visible spins and 2 hidden spins), and $M = 100$ quenches and search over $N = 500$ different disorder realizations.

Generally, we found that more hidden units lead to better learning outcome. However, for the tasks considered in this

study, further increasing the number of hidden units beyond two does not result in substantial improvements in the results. Nonetheless, we expect that increasing the number of hidden units may prove beneficial for more complicated tasks, as suggested in the future directions.

3. Runtime scaling

The runtime complexity of our learning algorithm is primarily determined by its use of Monte Carlo searches, resulting in an overall complexity of $\mathcal{O}(TNM)$. Here, T represents the evolution time of the system during each quench, N is the task-dependent number of Monte Carlo samples, and M denotes the constant number of quenches. Notably, the specific influence of T depends on the simulation method employed. In the case of exact diagonalization, T scales exponentially with system size, leading to substantial computational requirements for larger systems. However, it is anticipated that near-term analog or digital quantum computers will significantly enhance the efficiency of these operations, mitigating the impact of this scaling behavior.

APPENDIX D: DATA ENCODING

Here, we describe the detailed data encoding scheme and our toy dataset of MNIST digit patterns in this section. Given a reduced density matrix ρ_{vis} of L visible spins, we compute the distribution of finding each of the 2^L basis states in our computational basis, and interpret the result as pixel values. We then reshape this probability vector into an image of size $2^{L/2} \times 2^{L/2}$.

On the other hand, given an image $\tilde{x}_\mu \in \mathbb{R}^{n \times n}$, where $n \times n$ is the number of original pixels in the image, we first down-sample it to $2^{L/2} \times 2^{L/2}$ pixels, then normalized the pixel values to be within 0 and 1.

Our toy dataset of MNIST digit patterns are constructed as follows: we take all the training images $\tilde{x}^\mu$ from a digit class, downsample to $2^{L/2} \times 2^{L/2}$ pixels, and compute each pixel as the average value $\bar{x}_i = 1/P \sum_{\mu=1}^P x_i^\mu$ across different styles within this digit class, where $i = 1, \dots, 2^L$. We then normalized the pixels to $\bar{x}_i \rightarrow \bar{x}_i / \sum_i \bar{x}_i$ and interpret the result as q_{data} . We take caution that this is different from learning the MNIST distribution in generative models. The latter refers to learning the joint probability distribution over all pixels in the image, and our toy data set corresponds to taking the mean-field limit of this joint probability distribution, which ignores the complicated correlations among pixels. This is akin to learning a single pattern (the averaged MNIST images shown in Fig. 3), and the reason for taking the average pixel value is such that we will be able to perform pattern recognition with imperfect initial states.

[1] P. Ball *et al.*, First quantum computer to pack 100 qubits enters crowded race, *Nature (London)* **599**, 542 (2021).
 [2] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature (London)* **574**, 505 (2019).

[3] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, and R. Wisnieff, Leveraging secondary storage to simulate deep 54-qubit sycamore circuits, [arXiv:1910.09534](https://arxiv.org/abs/1910.09534).
 [4] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z.

- Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, Quantum computational advantage using photons, *Science* **370**, 1460 (2020).
- [5] M. Gong, S. Wang, C. Zha, M.-C. Chen, H.-L. Huang, Y. Wu, Q. Zhu, Y. Zhao, S. Li, S. Guo, H. Qian, Y. Ye, F. Chen, C. Ying, J. Yu, D. Fan, D. Wu, H. Su, H. Deng, H. Rong, K. Zhang, S. Cao, J. Lin, Y. Xu, L. Sun, C. Guo, N. Li, F. Liang, V. M. Bastidas, K. Nemoto, W. J. Munro, Y.-H. Huo, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, Quantum walks on a programmable two-dimensional 62-qubit superconducting processor, *Science* **372**, 948 (2021).
 - [6] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan *et al.*, Strong quantum computational advantage using a superconducting quantum processor, *Phys. Rev. Lett.* **127**, 180501 (2021).
 - [7] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho *et al.*, Quantum phases of matter on a 256-atom programmable quantum simulator, *Nature (London)* **595**, 227 (2021).
 - [8] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
 - [9] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
 - [10] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
 - [11] D. Marković and J. Grollier, Quantum neuromorphic computing, *Appl. Phys. Lett.* **117**, 150501 (2020).
 - [12] P. Pfeiffer, I. Egusquiza, M. Di Ventra, M. Sanz, and E. Solano, Quantum memristors, *Sci. Rep.* **6**, 29507 (2016).
 - [13] T. Gonzalez-Raya, X.-H. Cheng, I. L. Egusquiza, X. Chen, M. Sanz, and E. Solano, Quantized single-ion-channel Hodgkin-Huxley model for quantum neurons, *Phys. Rev. Appl.* **12**, 014037 (2019).
 - [14] T. Gonzalez-Raya, E. Solano, and M. Sanz, Quantized three-ion-channel neuron model for neural action potentials, *Quantum* **4**, 224 (2020).
 - [15] E. Torrontegui and J. J. García-Ripoll, Unitary quantum perceptron as efficient universal approximator, *Europhys. Lett.* **125**, 30004 (2019).
 - [16] K. Fujii and K. Nakajima, Quantum reservoir computing: a reservoir approach toward quantum machine learning on near-term quantum devices, in *Reservoir Computing: Theory, Physical Implementations, and Applications* (Springer, 2021), pp. 423–450.
 - [17] K. Nakajima, K. Fujii, M. Negoro, K. Mitarai, and M. Kitagawa, Boosting computational power through spatial multiplexing in quantum reservoir computing, *Phys. Rev. Appl.* **11**, 034021 (2019).
 - [18] R. A. Bravo, K. Najafi, X. Gao, and S. F. Yelin, Quantum reservoir computing using arrays of Rydberg atoms, *PRX Quantum* **3**, 030325 (2022).
 - [19] W. Xia, J. Zou, X. Qiu, and X. Li, The reservoir learning power across quantum many-body localization transition, *Front. Phys.* **17**, 33506 (2022).
 - [20] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, Unsupervised generative modeling using matrix product states, *Phys. Rev. X* **8**, 031012 (2018).
 - [21] S. Cheng, L. Wang, T. Xiang, and P. Zhang, Tree tensor networks for generative modeling, *Phys. Rev. B* **99**, 155131 (2019).
 - [22] A. M. Gomez, S. F. Yelin, and K. Najafi, Reconstructing quantum states using basis-enhanced born machines, [arXiv:2206.01273](https://arxiv.org/abs/2206.01273) (2022).
 - [23] H.-J. Liao, J.-G. Liu, L. Wang, and T. Xiang, Differentiable programming tensor networks, *Phys. Rev. X* **9**, 031041 (2019).
 - [24] D. A. Huse, R. Nandkishore, V. Oganesyan, A. Pal, and S. L. Sondhi, Localization-protected quantum order, *Phys. Rev. B* **88**, 014206 (2013).
 - [25] B. Altshuler, H. Krovi, and J. Roland, Anderson localization makes adiabatic quantum optimization fail, *Proc. Natl. Acad. Sci.* **107**, 12446 (2010).
 - [26] R. Martínez-Peña, G. L. Giorgi, J. Nokkala, M. C. Soriano, and R. Zambrini, Dynamical phase transitions in quantum reservoir computing, *Phys. Rev. Lett.* **127**, 100502 (2021).
 - [27] J. Tangpanitanon, S. Thanasilp, N. Dangniam, M.-A. Lemonde, and D. G. Angelakis, Expressibility and trainability of parametrized analog quantum systems for machine learning applications, *Phys. Rev. Res.* **2**, 043364 (2020).
 - [28] J.-G. Liu and L. Wang, Differentiable learning of quantum circuit born machines, *Phys. Rev. A* **98**, 062324 (2018).
 - [29] S. Cheng, J. Chen, and L. Wang, Information perspective to probabilistic modeling: Boltzmann machines versus born machines, *Entropy* **20**, 583 (2018).
 - [30] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, A generative modeling approach for benchmarking and training shallow quantum circuits, *npj Quantum Inf.* **5**, 45 (2019).
 - [31] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).
 - [32] B. Coyle, D. Mills, V. Danos, and E. Kashefi, The born supremacy: Quantum advantage and training of an ising born machine, *npj Quantum Inf.* **6**, 60 (2020).
 - [33] P. Smolensky, *Information Processing in Dynamical Systems: Foundations of Harmony Theory* (Colorado University at Boulder Department of Computer Science, 1986), pp. 194–281.
 - [34] G. E. Hinton and R. R. Salakhutdinov, Reducing the dimensionality of data with neural networks, *Science* **313**, 504 (2006).
 - [35] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, A learning algorithm for boltzmann machines, *Cognitive science* **9**, 147 (1985).
 - [36] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to machine learning for physicists, *Phys. Rep.* **810**, 1 (2019).
 - [37] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, *Nat. Commun.* **8**, 662 (2017).
 - [38] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities., *Proc. Natl. Acad. Sci.* **79**, 2554 (1982).
 - [39] D. Krotov and J. J. Hopfield, Dense associative memory for pattern recognition, *Advances in Neural Information Processing Systems*, Vol. 29 (Curran Associates, Inc., 2016).
 - [40] A. Murugan, Z. Zeravcic, M. P. Brenner, and S. Leibler, Multifarious assembly mixtures: Systems allowing retrieval of diverse stored structures, *Proc. Natl. Acad. Sci.* **112**, 54 (2015).

- [41] W. Zhong, D. J. Schwab, and A. Murugan, Associative pattern recognition through macro-molecular self-assembly, *J. Stat. Phys.* **167**, 806 (2017).
- [42] J. M. Deutsch, Quantum statistical mechanics in a closed system, *Phys. Rev. A* **43**, 2046 (1991).
- [43] M. Srednicki, Chaos and quantum thermalization, *Phys. Rev. E* **50**, 888 (1994).
- [44] P. W. Anderson, Absence of diffusion in certain random lattices, *Phys. Rev.* **109**, 1492 (1958).
- [45] D. Basko, I. Aleiner, and B. Altshuler, Metal-insulator transition in a weakly interacting many-electron system with localized single-particle states, *Ann. Phys.* **321**, 1126 (2006).
- [46] V. Oganesyan and D. A. Huse, Localization of interacting fermions at high temperature, *Phys. Rev. B* **75**, 155111 (2007).
- [47] J. Smith, A. Lee, P. Richerme, B. Neyenhuis, P. W. Hess, P. Hauke, M. Heyl, D. A. Huse, and C. Monroe, Many-body localization in a quantum simulator with programmable random disorder, *Nat. Phys.* **12**, 907 (2016).
- [48] D. J. Luitz, N. Laflorencie, and F. Alet, Many-body localization edge in the random-field heisenberg chain, *Phys. Rev. B* **91**, 081103(R) (2015).
- [49] A. Pal and D. A. Huse, Many-body localization phase transition, *Phys. Rev. B* **82**, 174411 (2010).
- [50] T. Devakul and R. R. P. Singh, Early breakdown of area-law entanglement at the many-body delocalization transition, *Phys. Rev. Lett.* **115**, 187201 (2015).
- [51] L. Bottou, Stochastic gradient learning in neural networks, *Proc. Neuro-Nimes* **91**, 12 (1991).
- [52] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, 2016).
- [53] L. Deng, The mnist database of handwritten digit images for machine learning research, *IEEE Signal Process Mag.* **29**, 141 (2012).
- [54] P. Hauke and M. Heyl, Many-body localization and quantum ergodicity in disordered long-range ising models, *Phys. Rev. B* **92**, 134204 (2015).
- [55] J. H. Bardarson, F. Pollmann, and J. E. Moore, Unbounded growth of entanglement in models of many-body localization, *Phys. Rev. Lett.* **109**, 017202 (2012).
- [56] M. Serbyn, Z. Papić, and D. A. Abanin, Quantum quenches in the many-body localized phase, *Phys. Rev. B* **90**, 174302 (2014).
- [57] D. A. Huse, R. Nandkishore, and V. Oganesyan, Phenomenology of fully many-body-localized systems, *Phys. Rev. B* **90**, 174202 (2014).
- [58] E. Winfree, *Algorithmic Self-Assembly of DNA* (California Institute of Technology, 1998).
- [59] L. Qian, D. Soloveichik, and E. Winfree, Efficient Turing-universal computation with DNA polymers, in *DNA Computing and Molecular Programming: 16th International Conference, DNA 16, Hong Kong, China, June 14-17, 2010*, Revised Selected Papers 16 (Springer, Berlin, Heidelberg, 2011).
- [60] J. O'Brien and A. Murugan, Temporal pattern recognition through analog molecular computation, *ACS Synth. Biol.* **8**, 826 (2019).
- [61] M. Stern, M. B. Pinson, and A. Murugan, Continual learning of multiple memories in mechanical networks, *Phys. Rev. X* **10**, 031044 (2020).
- [62] W. Zhong, J. M. Gold, S. Marzen, J. L. England, and N. Y. Halpern, Machine learning outperforms thermodynamics in measuring how well a many-body system learns a drive, *Sci. Rep.* **11**, 9333 (2021).
- [63] M. Stern and A. Murugan, Learning without neurons in physical systems, *Ann. Rev. Condensed Matter Phys.* **14**, 417 (2023).
- [64] While there exists other more sophisticated encoding schemes, here we choose the simplest one for illustration.
- [65] However, one should note that just like in classical pattern recognition [38], if the input pattern gets too corrupted and does not resemble any of the encoded patterns, this procedure will fail.
- [66] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, *Nat Mach Intell* **1**, 155 (2019).
- [67] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [68] J. Wang, Z.-Y. Han, S.-B. Wang, Z. Li, L.-Z. Mu, H. Fan, and L. Wang, Scalable quantum tomography with fidelity estimation, *Phys. Rev. A* **101**, 032321 (2020).
- [69] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, *Science* **377**, 682 (2022).
- [70] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [71] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, *Nat. Phys. Rev.* **5**, 9 (2023).
- [72] K. Najafi, M. E. Stoudenmire, S. Yelin, A. Azizi, X. Gao, M. D. Lukin, and M. Mohseni, Limitations of gradient-based Born Machines over tensor networks on learning quantum nonlocality, in *34th Conference on Neural Information Processing Systems*, First Workshop on Quantum Tensor Networks in Machine Learning (JMLR, 2020), p. 10.
- [73] T.-D. Bradley, E. M. Stoudenmire, and J. Terilla, Modeling sequences with quantum states: A look under the hood, *Mach. Learn.: Sci. Technol.* **1**, 035008 (2020).
- [74] J. A. Kjäll, J. H. Bardarson, and F. Pollmann, Many-body localization in a disordered quantum ising chain, *Phys. Rev. Lett.* **113**, 107204 (2014).
- [75] B. Fefferman, M. Foss-Feig, and A. V. Gorshkov, Exact sampling hardness of ising spin models, *Phys. Rev. A* **96**, 032324 (2017).
- [76] R. Movassagh, The hardness of random quantum circuits, *Nat. Phys.* **19**, 1719 (2023).
- [77] T. L. Patti, K. Najafi, X. Gao, and S. F. Yelin, Entanglement devised barren plateau mitigation, *Phys. Rev. Res.* **3**, 033090 (2021).
- [78] J. Liu, K. Najafi, K. Sharma, F. Tacchino, L. Jiang, and A. Mezzacapo, An analytic theory for the dynamics of wide quantum neural networks, *Phys. Rev. Lett.* **130**, 150601 (2023).
- [79] X. Gao, S.-T. Wang, and L.-M. Duan, Quantum supremacy for simulating a translation-invariant ising spin model, *Phys. Rev. Lett.* **118**, 040502 (2017).
- [80] P. Weinberg and M. Bukov, Quspin: a python package for dynamics and exact diagonalisation of quantum many body systems part I: spin chains, *SciPost Phys.* **2**, 003 (2017).

- [81] P. Weinberg and M. Bukov, Quspin: A python package for dynamics and exact diagonalisation of quantum many body systems. part ii: bosons, fermions and higher spins, [SciPost Phys. 7, 020 \(2019\)](#).
- [82] R. Nandkishore and D. A. Huse, Many body localization and thermalization in quantum statistical mechanics, [Annu. Rev. Condens. Matter Phys. 6, 15 \(2015\)](#).
- [83] F. Alet and N. Laflorencie, Many-body localization: An introduction and selected topics, [C. R. Phys. 19, 498 \(2018\)](#).
- [84] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, Colloquium: Many-body localization, thermalization, and entanglement, [Rev. Mod. Phys. 91, 021001 \(2019\)](#).