

Unsupervised feature learning from finite data by message passing: Discontinuous versus continuous phase transition

Haiping Huang* and Taro Toyoizumi†

RIKEN Brain Science Institute, Wako-shi, Saitama 351-0198, Japan

(Received 25 August 2016; revised manuscript received 1 December 2016; published 21 December 2016)

Unsupervised neural network learning extracts hidden features from unlabeled training data. This is used as a pretraining step for further supervised learning in deep networks. Hence, understanding unsupervised learning is of fundamental importance. Here, we study the unsupervised learning from a finite number of data, based on the restricted Boltzmann machine where only one hidden neuron is considered. Our study inspires an efficient message-passing algorithm to infer the hidden feature and estimate the entropy of candidate features consistent with the data. Our analysis reveals that the learning requires only a few data if the feature is salient and extensively many if the feature is weak. Moreover, the entropy of candidate features monotonically decreases with data size and becomes negative (i.e., entropy crisis) before the message passing becomes unstable, suggesting a discontinuous phase transition. In terms of convergence time of the message-passing algorithm, the unsupervised learning exhibits an easy-hard-easy phenomenon as the training data size increases. All these properties are reproduced in an approximate Hopfield model, with an exception that the entropy crisis is absent, and only continuous phase transition is observed. This key difference is also confirmed in a handwritten digits dataset. This study deepens our understanding of unsupervised learning from a finite number of data and may provide insights into its role in training deep networks.

DOI: [10.1103/PhysRevE.94.062310](https://doi.org/10.1103/PhysRevE.94.062310)

I. INTRODUCTION

Humans and other animals can learn new concepts from only a handful of training data but standard machine-learning algorithms require huge data to uncover hidden features [1]. Learning hidden features in unlabeled data is called unsupervised learning. How many data are required to learn a feature? What key factors determine the success of unsupervised learning? These fundamental questions are largely unsolved so far, and rarely explained by a physics model. Understanding how data size confines learning process is a topic of interest not only in machine learning [2] but also in cognitive neuroscience [3]. The underlying neural mechanism or inspired algorithms are still elusive, but recent progress in mean-field theory of restricted Boltzmann machine [4] allows us to develop a statistical mechanics model to understand how learning improves with data size.

Restricted Boltzmann machine (RBM) is a basic block widely used in building a deep belief network [5,6]. It consists of two layers of neurons. The visible layer receives data, and based on this, the hidden layer builds an internal representation. No lateral connections exist within each layer for computational efficiency. The symmetric connections between visible and hidden neurons represent hidden features in the data that the network learns. A common strategy to train a RBM is to apply a gradient-descent update algorithm of the connections based on stochastically sampled neural activities [6]. However, no theory was proposed to address how learning improves and how the number of candidate features decreases with data size.

Here we tackle this problem by taking a different perspective. We propose a Bayesian inference framework to uncover

the connections, i.e., a feature vector, from data by explicitly modeling its posterior distribution. Based on this framework, we develop a message-passing algorithm for inferring an optimal feature vector, while monitoring its entropy to quantify how many candidate feature vectors are consistent with given data.

II. MODEL DEFINITION

In this study, we consider a RBM [4,7] with a single hidden neuron, whose activity is generated according to $P(\sigma, h) \propto e^{-\beta E(\sigma, h)/\sqrt{N}}$, where $\sigma = \{\sigma_i | \sigma_i = \pm 1, i = 1, \dots, N\}$ represents binary activity of N visible neurons, $h = \pm 1$ is the activity of the hidden neuron, β/\sqrt{N} is the system-size and inverse-temperature dependent scaling factor, and $E(\sigma, h) = -h\xi^T\sigma$ is the energy function characterized by a feature vector ξ (T indicates vector transpose). We assume that each element of the feature vector takes a binary value $\xi_i = \pm 1$. While generalization to a case with multiple hidden neurons is possible, we do not explore it here as the analysis becomes more involved.

To perform unsupervised learning, we generate M independent samples as training data $\{\sigma^a\}_{a=1}^M$ from a RBM with randomly generated true feature vector ξ^{true} , where each element is drawn from ± 1 with equal probability. Another RBM learns this feature vector from the data. We formulate the learning process as Bayesian inference. Given the training data, the posterior distribution of the feature vector is

$$P(\xi | \{\sigma^a\}_{a=1}^M) \propto \prod_{a=1}^M P(\sigma^a | \xi) = \frac{1}{Z} \prod_{a=1}^M \cosh\left(\frac{\beta}{\sqrt{N}} \xi^T \sigma^a\right), \quad (1)$$

where Z is the partition function of the model and we assume a uniform prior about ξ . Here, the data $\{\sigma^a\}_{a=1}^M$ serves as the quenched disorder (data constraints), and the

*physhuang@gmail.com

†taro.toyoizumi@brain.riken.jp

inverse-temperature parameter β characterizes the learning difficulty on a network of dimension N . If $M > 1$, the model becomes nontrivial as the partition function cannot be computed exactly for a large N . This M can be proportional to the system size, and in this case we define a data density as $\alpha = M/N$. Hereafter, we omit the conditional dependence of $P(\xi|\{\sigma^a\}_{a=1}^M)$ on $\{\sigma^a\}_{a=1}^M$.

III. MEAN-FIELD THEORY

In the following, we compute the maximizer of the posterior marginals (MPM) estimator $\hat{\xi}_i = \arg \max_{\xi_i} P_i(\xi_i)$ [8], which maximizes the overlap $q = \frac{1}{N} \sum_{i=1}^N \xi_i^{\text{true}} \hat{\xi}_i$ between the true and estimated feature vectors. If $q = 0$, the data do not give any information about the feature vector. If $q = 1$, the feature vector is perfectly estimated. Hence, the task now is to compute marginal probabilities, e.g., $P_i(\xi_i)$, which is still a hard problem due to the interaction among data constraints. However, by mapping the model [Eq. (1)] onto a factor graph [4,9], the marginal probability can be estimated by message passing (Fig. 1) as we explain below. We first assume that elements of the feature vector on the factor graph are weakly correlated (also named Bethe approximation [10]), then by using the cavity method [9], we define a cavity probability $P_{i \rightarrow a}(\xi_i)$ of ξ_i on a modified factor graph with data node a removed. Under the weak correlation assumption, $P_{i \rightarrow a}(\xi_i)$ satisfies a recursive equation (namely belief propagation (BP) in computer science [11]):

$$P_{i \rightarrow a}(\xi_i) \propto \prod_{b \in \partial i \setminus a} \mu_{b \rightarrow i}(\xi_i), \quad (2a)$$

$$\mu_{b \rightarrow i}(\xi_i) = \sum_{\{\xi_j | j \in \partial b \setminus i\}} \cosh\left(\frac{\beta}{\sqrt{N}} \xi^T \sigma^b\right) \prod_{j \in \partial b \setminus i} P_{j \rightarrow b}(\xi_j), \quad (2b)$$

where the proportionality symbol \propto omits a normalization constant, $\partial i \setminus a$ defines the neighbors of feature node i except

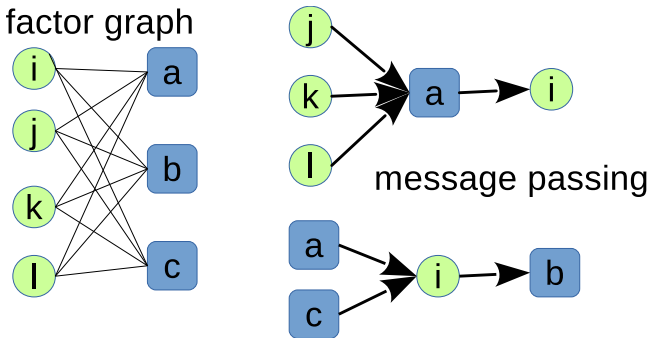


FIG. 1. Schematic illustration of factor graph representation and message passing. Left panel: circle nodes indicate features to be inferred. Square nodes indicate data constraints. A connection σ_i^a indicates how the feature ξ_i is related to a th data. Right panel: the top panel shows data node a collects messages from its neighboring features other than node i and produces an output message to node i . The bottom panel shows feature node i collects messages from its neighboring data nodes other than b and produces an output message to data node b . Iteration of these messages gives a coherent understanding of the feature learning process.

data node a , $\partial b \setminus i$ defines the neighbors of data node b except feature node i , and the auxiliary quantity $\mu_{b \rightarrow i}(\xi_i)$ represents the contribution from data node b to feature node i given the value of ξ_i [9]. An equation similar to Eq. (2) was recently derived to compute activity statistics of a RBM [4].

In the thermodynamic limit, the sum inside the hyperbolic cosine function excluding the i -dependent term in Eq. (2b) is a random variable following a normal distribution with mean $G_{b \rightarrow i}$ and variance $\Xi_{b \rightarrow i}^2$ [4], where $G_{b \rightarrow i} = \frac{1}{\sqrt{N}} \sum_{j \in \partial b \setminus i} \sigma_j^b m_{j \rightarrow b}$ and $\Xi_{b \rightarrow i}^2 \simeq \frac{1}{N} \sum_{j \in \partial b \setminus i} (1 - m_{j \rightarrow b}^2)$. The cavity magnetization is defined as $m_{j \rightarrow b} = \sum_{\xi_j} \xi_j P_{j \rightarrow b}(\xi_j)$. Thus, the intractable sum over all ξ_j ($j \neq i$) can be replaced by an integral over the normal distribution. Furthermore, because ξ_i is a binary variable, $P_{i \rightarrow a}(\xi_i)$ and $\mu_{b \rightarrow i}(\xi_i)$ are completely characterized [4] by the cavity magnetization $m_{i \rightarrow a}$ and cavity bias $u_{b \rightarrow i} = \frac{1}{2} \ln \frac{\mu_{b \rightarrow i}(\xi_i=1)}{\mu_{b \rightarrow i}(\xi_i=-1)}$, respectively. Using these expressions, the BP equation [Eq. (2)] could be reduced to the following practical recursive equations:

$$m_{i \rightarrow a} = \tanh\left(\sum_{b \in \partial i \setminus a} u_{b \rightarrow i}\right), \quad (3a)$$

$$u_{b \rightarrow i} = \tanh^{-1}\left[\tanh(\beta G_{b \rightarrow i}) \tanh(\beta \sigma_i^b / \sqrt{N})\right], \quad (3b)$$

where $m_{i \rightarrow a}$ can be interpreted as the message passing from feature i to the data constraint a , while $u_{b \rightarrow i}$ can be interpreted as the message passing from data constraint b to feature i (Fig. 1). If the weak correlation assumption is self-consistent, the BP would converge to a fixed point corresponding to a stationary point of the Bethe free-energy function with respect to the cavity messages $\{m_{i \rightarrow a}, u_{a \rightarrow i}\}$ [9].

By initializing the message on each link of the factor graph (Fig. 1), we run the BP equation [Eq. (3)] until it converges within a prefixed precision. From this fixed point, one can extract useful information about the true feature vector by calculating the marginal probability as $P_i(\xi_i) = \frac{1+m_i \xi_i}{2}$ where $m_i = \tanh(\sum_{b \in \partial i} u_{b \rightarrow i})$. An alternative strategy to use passing messages to infer the true feature vector is called reinforced BP (rBP). This strategy is a kind of soft-decimation [12], which progressively enhances or weakens the current local field ($h_i^t = \sum_{b \in \partial i} u_{b \rightarrow i}$) each feature component feels by the reinforcement rule $h_i^t \leftarrow h_i^t + h_i^{t-1}$ with a probability $1 - \gamma^t$, until a solution ($\{\hat{\xi}_i = \text{sgn}(m_i)\}$) is stable over iterations. γ usually takes a value close to 1.

Another important quantity is the number of candidate feature vectors consistent with the data, characterized by the entropy per neuron $s = -\frac{1}{N} \sum_{\xi} P(\xi) \ln P(\xi)$. Under the Bethe approximation, s is evaluated as summing up contributions from single feature nodes and data nodes (Appendix A).

IV. RESULTS

We use the above mean-field theory to study unsupervised learning of the feature vector in single random realizations of the true feature vector and analyze how its performance depends on model parameters. We first explore effects of β ($N = 100$). Note that β scales the energy and hence tunes the difficulty of the learning. As shown in Fig. 2(a), it requires an extensive number of data to learn a weak feature vector

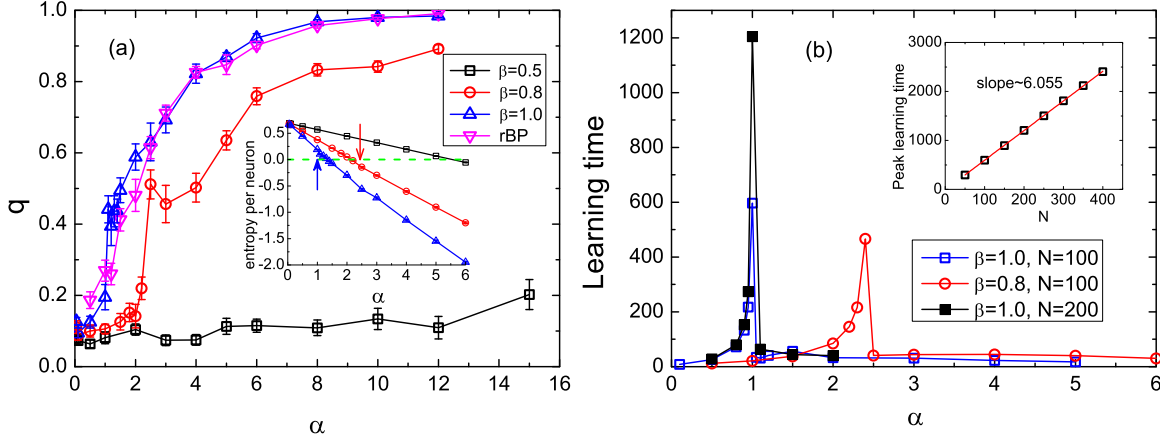


FIG. 2. Bayesian learning in the RBM model [Eq. (1)]. Thirty random realizations of the model are considered. The error bars characterize the standard deviation; $N = 100$. (a) Inference overlap q versus data density α with different values of feature saliency β . Results obtained by rBP ($\gamma = 0.95$) are also shown for comparison ($\beta = 1$). The inset shows the entropy per neuron. Arrows in the inset indicate thresholds for continuous transitions. (b) Median of convergence (learning) time defined by the number of iterations at which BP converges to a fixed point. Two different values of β are considered. For $\beta = 1$, results for a larger N are also shown. The inset shows that the peak learning time scales linearly with N ($\beta = 1$).

($\beta = 0.5$). However, as β increases, the inference becomes much better. The overlap grows more rapidly at $\beta = 1.0$. When α is above 10, one can get a nearly perfect inference of the feature vector. We also use the reinforcement strategy to infer the true feature vector, and it has nearly the same performance with reduced computer time, because the estimation of the true feature need not be carried out at the fixed point.

Remarkably, the overlap improves at some α , revealing that the RBM could extract the hidden feature vector only after sufficient data are shown. This critical α decreases with the saliency β of the hidden feature. A statistical analysis of Eq. (3) reveals that $\alpha_c = \frac{1}{\beta^2}$ (Appendix C). Next, we show the entropy per neuron in the inset of Fig. 2(a). This quantity that describes how many feature vectors are consistent with the data becomes negative (i.e., entropy crisis [13,14]) at a zero-entropy α_{ZE} . However, the BP equation is still stable (convergent), and thus the instability occurs after the entropy crisis. This suggests the existence of a discontinuous glass transition at a value of α less than or equal to α_{ZE} , as commonly observed in some spin glass models of combinatorial satisfaction problems [14–17]. As shown in the inset of Fig. 2(a), α_c can be either larger or smaller than α_{ZE} , depending on β . If $\alpha_c < \alpha_{ZE}$, a continuous transition is followed by a discontinuous transition, where intrastate and interstate overlaps [18] bifurcate discontinuously. If $\alpha_c > \alpha_{ZE}$, the predicted continuous transition is inaccurate at least under the replica symmetric assumption. Detailed discussions about this novel property will be provided in a forthcoming extended work [19]. Interestingly, despite the likely glass transition and entropy crisis, the fixed point of BP still yields good inference of the feature vector, which may be related to the Nishimori condition (Bayes-optimal inference) [19,20], since the temperature parameter used in inference is the same as that used to generate the data.

Next, we study the median of learning time. The learning time is measured as the number of iterations when the message passing converges. Figure 2(b) shows that learning is fast at small α , slow around the critical α , and becomes fast again at

large α . This easy-hard-easy phenomenon can be understood as follows. When a few data are presented, the inference is less constrained, and thus there exist many candidate feature vectors consistent with the data, the BP converges fast to estimate the marginals. Once relatively many data are presented, the number of candidate feature vectors reduces [Fig. 2(a)], and the BP requires more iterations to find a candidate feature vector. We also observe that the peak learning time at α_c scales linearly with N within the measured range [the inset of Fig. 2(b)]. This intermediate region is thus termed hard phase. Large β moves the hard phase to the small α region. Once the number of data is sufficiently large, the inference becomes easy once again as the number of necessary iterations drops drastically. This may be because the feature space around the true feature vector dominates the posterior probability in the presence of sufficient data.

V. CONNECTION TO AN APPROXIMATE HOPFIELD MODEL

It is interesting to show that one can also perform the same unsupervised learning by using an associative memory (Hopfield) model defined by

$$\tilde{P}(\xi) \propto \prod_a e^{\frac{\tilde{\beta}}{2N} (\xi^T \sigma^a)^2}, \quad (4)$$

where $\tilde{\beta} = \beta^2$. This posterior distribution about a feature vector ξ given data $\{\sigma^a\}_{a=1}^M$ can be obtained by a small- β expansion of Eq. (1) [21]. This relationship indicates that one can infer the feature vector of a RBM using this Hopfield model if β is small enough. In this case, the true feature vector is also interpreted as a unique memory pattern in the Hopfield model. By a small- β expansion, we interpret the unsupervised learning in a RBM model as recovering stored pattern in a Hopfield model from noisy data [22]. In the following, we generate data by a RBM with true feature ξ^{true} and compute

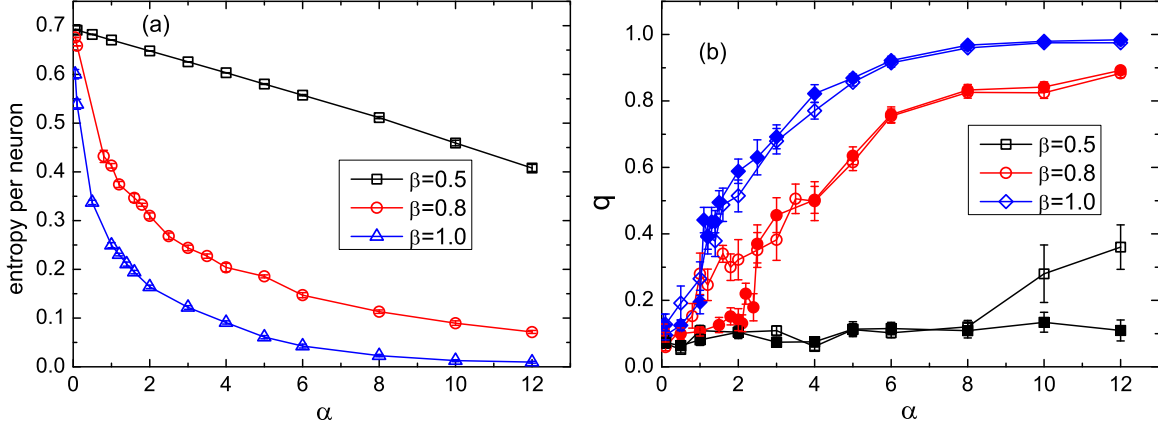


FIG. 3. Bayesian learning in the approximate Hopfield model [Eq. (4)]. Thirty random realizations of the model are considered. The error bars characterize the standard deviation; $N = 100$. (a) Entropy per neuron versus the data density α . (b) Overlap versus the data density. Results for RBM (solid symbols) are shown for comparison.

the MPM estimator $\hat{\xi}$ of the memory pattern in Eq. (4) based on the given data.

Analogous to the derivation of Eq. (3), we can derive the practical BP corresponding to the posterior probability [Eq. (4)]:

$$m_{i \rightarrow a} = \tanh \left(\frac{\tilde{\beta}}{\sqrt{N}} \sum_{b \in \partial i \setminus a} \sigma_i^b \tilde{G}_{b \rightarrow i} F_{b \rightarrow i} \right), \quad (5)$$

where $\tilde{G}_{b \rightarrow i} = \frac{1}{\sqrt{N}} \sum_{k \in \partial b \setminus i} \sigma_k^b m_{k \rightarrow b}$, $F_{b \rightarrow i} = 1 + \frac{\tilde{\beta} C_{b \rightarrow i}}{1 - \tilde{\beta} C_{b \rightarrow i}}$ in which $C_{b \rightarrow i} = \frac{1}{N} \sum_{k \in \partial b \setminus i} (1 - m_{k \rightarrow b}^2)$. The entropy of candidate feature vectors can also be evaluated from the fixed point of this iterative equation (Appendix B).

Bayesian learning performance of the Hopfield model is shown in Fig. 3. This model does not show an entropy crisis in the explored range of α . As α increases, the entropy decreases much more slowly for weak features than for strong ones. For $\beta = 0.5$, the overlap stays slightly above zero for a wide range of α . At a sufficiently large α (~ 8), the overlap starts to increase continuously. It is impossible to predict underlying structures if the number of data is insufficient. This phenomenon is named retarded learning first observed in unsupervised learning based on Gaussian or mixture-of-Gaussian data [23,24]. At large α where the entropy value approaches zero, the overlap approaches one. All the properties except the entropy crisis are qualitatively similar in both the approximate Hopfield and RBM model. The absence of entropy crisis may be related to the absence of p -spin ($p > 2$) interactions in the approximate model, which has thus only the continuous glass transition at $\alpha_c = (\frac{1}{\beta} - 1)^2$, predicted by a statistical analysis of Eq. (5) (Appendix C). Note that the spin glass transition in a standard Hopfield model where many random patterns are stored is of second order [25]. The current analysis sheds light on understanding the relationship between RBM and associative memory networks [21,26] within an unsupervised learning framework.

VI. APPLICATIONS OF OUR THEORY ON MNIST DATASET

Finally, we test our theory on the MNIST handwritten digit dataset [27]. For simplicity, we consider some combinations of two different digits (e.g., 0 and 1, 3 and 7). Each digit is represented by a 28×28 gray-scale image. Results show that the real dataset shares common properties with our model (Fig. 4), which does not change qualitatively when different combinations even the whole dataset are used. Importantly, the presence and absence of the entropy crisis in a RBM and a Hopfield model, respectively, are also confirmed in this real dataset.

Given the dataset [e.g., a set of digit 0 and 1 images (see Fig. 5)], our theory could estimate the corresponding hidden feature vector consistent with these images. The feature vector is also called the receptive field of the hidden neuron, because it determines the firing response of the hidden neuron. We organize the inferred feature vector into a 28×28 matrix (the same size as the input image), and plot the matrix as the gray-scale image (black indicates $\xi_i = -1$, and white indicates $\xi_i = 1$). As shown in Fig. 6, the inferred feature

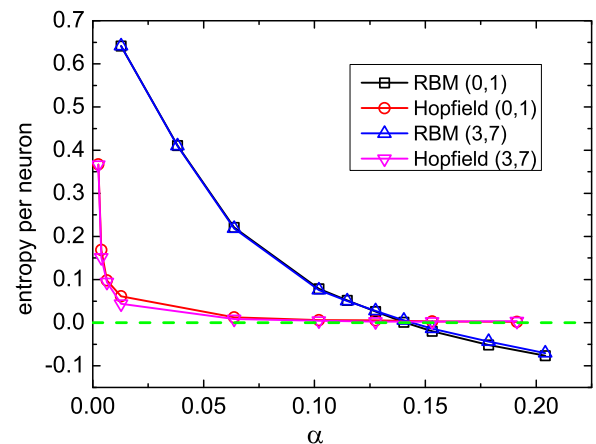


FIG. 4. Entropy per neuron versus α for MNIST dataset [handwritten digits (0, 1) and (3, 7)] at $\beta = 1$.




FIG. 5. Two examples of digits (0 and 1) from the MNIST dataset are shown.

vector gets much better as the number of input images grows, serving as a local structure detector. This is indicated by a few active synaptic weights in the center of the image, where local characteristics of handwritten digits are approximately captured when sufficient data samples are shown to the

machine. In a deep network composed of many layers of stacked RBM, the low-level features (e.g., edges or contours) detected by the hidden neuron could be passed to deeper layers, where high-level information (e.g., object identity) could be extracted [29].

The effect of the data size on the feature learning performance can be quantitatively measured by the precision-recall curve. First, we computed the weighted sum (local field H) of the hidden neuron given the input image. These fields are then ranked. Digits 1 and 0 are discriminated by introducing a threshold H_{th} for corresponding local fields. The true positive (TP) event is identified when digit 1 is predicted by a local field above the threshold. The false-positive (FP) event is identified when digit 0 is wrongly predicted by a local field above the threshold. The false-negative event (FN) is identified when digit 1 is wrongly predicted by a local field below the threshold. The recall (RC) is defined as $RC = \frac{P_{TP}}{P_{TP} + P_{FN}}$, and the precision (PR) is defined as $PR = \frac{P_{TP}}{P_{TP} + P_{FP}}$, where P_{TP} is defined as the number of TP events in all presented samples. Thus the precision-recall curve is the parametric curve of $PR(H_{th})$ and $RC(H_{th})$. The closer we can get to (1, 1) in the plot, the better the unsupervised feature learning understands the embedded feature structure of the data. As shown in Fig. 7, as the data size increases, the performance improves,




FIG. 6. Inferred features evolve with the number of input images (M) on the MNIST dataset. The top panel displays the result for the RBM model, and the bottom panel displays the result for the approximate Hopfield model.

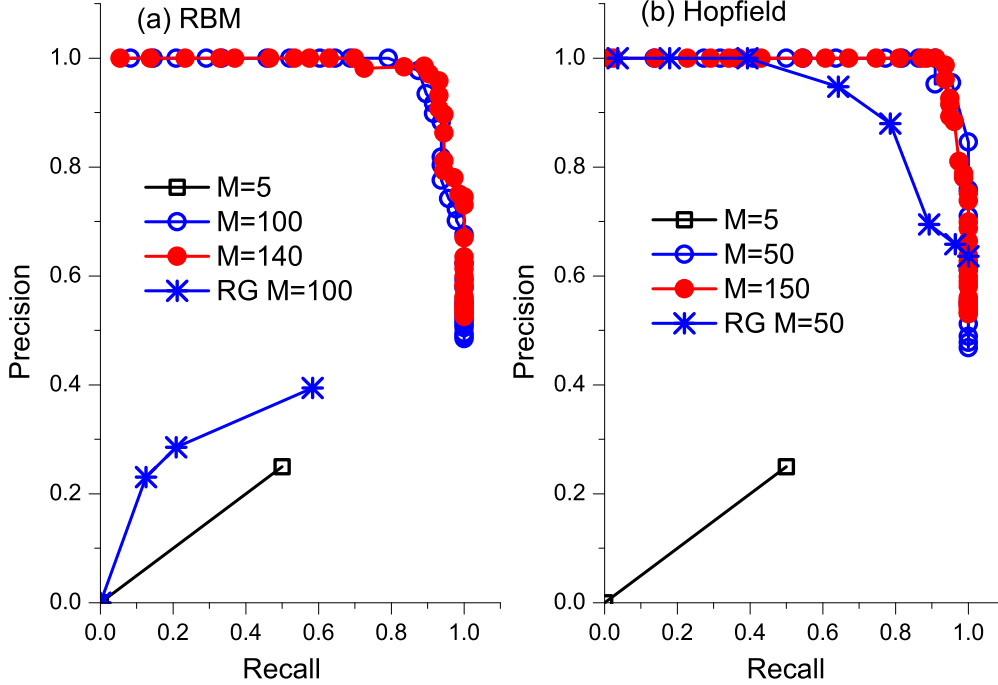


FIG. 7. The precision-recall curve corresponding to Fig. 6. The unsupervised learning performance is also compared with the random guess (RG) case.

and it behaves much better than a random guess of synaptic weights.

VII. CONCLUSION

In conclusion, we build a physics model of unsupervised learning based on the RBM and propose a Bayesian inference framework to extract a hidden feature vector from a finite number of data. The mean-field theory inspires an efficient message passing procedure to infer the hidden feature. Unlike previous approaches, each data in this work is treated as a constraint on the factor graph, and the message passing carries out a probabilistic inference of the hidden feature without sampling activities of neurons. We show that salient features can be recovered by even a few data. Conversely, it is impossible to recover the weak features by a finite amount of data. Interestingly, the entropy of candidate feature vectors becomes negative before the message passing algorithm becomes unstable, suggesting a discontinuous glass transition to resolve the entropy crisis, a typical statistical mechanics phenomenon revealed in studies of spin glass models [14–17]. In terms of the

convergence time of the message passing algorithm, we reveal the easy-hard-easy phenomenon for the current unsupervised learning, which is explained by our theory. All these properties except the entropy crisis are also observed in an approximate Hopfield model, where we infer from data a hidden feature of a RBM. Remarkably, these phenomena are also confirmed in a real dataset (Fig. 4). This work provides a theoretical basis to understand efficient neuromorphic implementation of RBM with simple binary feature elements [28]. We also hope our study will provide important insights into a physics understanding of unsupervised learning, especially its important role in pretraining deep neural networks for superior performances [6].

ACKNOWLEDGMENTS

H.H. thanks Pan Zhang for useful discussions. This work was supported by the program for Brain Mapping by Integrated Neurotechnologies for Disease Studies (Brain/MINDS) from Japan Agency for Medical Research and development, AMED, and by RIKEN Brain Science Institute.

APPENDIX A: THE NUMBER OF CANDIDATE FEATURE VECTORS IN THE RBM MODEL

For the unsupervised learning in the RBM model, one important quantity is the number of candidate feature vectors consistent with the input noisy data, characterized by the entropy per neuron $s = -\frac{1}{N} \sum_{\xi} P(\xi) \ln P(\xi)$. Under the Bethe approximation [10], s is evaluated as summing up contributions from single feature nodes and data nodes:

$$Ns = \sum_i \Delta S_i - (N - 1) \sum_a \Delta S_a, \tag{A1}$$

where single feature node contribution is expressed as $\Delta S_i = \sum_{a \in \partial i} [\beta^2 \Xi_{a \rightarrow i}^2 / 2 + \ln \cosh(\beta G_{a \rightarrow i} + \beta \sigma_i^a / \sqrt{N})] + \ln(1 + \prod_{a \in \partial i} \mathcal{G}_{a \rightarrow i}) - [\sum_{a \in \partial i} \mathcal{H}_{a \rightarrow i}(+1) + \prod_{a \in \partial i} \mathcal{G}_{a \rightarrow i} \sum_{a \in \partial i} \mathcal{H}_{a \rightarrow i}(-1)] / (1 + \prod_{a \in \partial i} \mathcal{G}_{a \rightarrow i})$, and single data node

contribution $\Delta S_a = \ln \cosh(\beta G_a) - \beta^2 \Xi_a^2/2 - \beta G_a \tanh(\beta G_a)$. We define $\mathcal{G}_{a \rightarrow i} = e^{-2u_{a \rightarrow i}}$, $\mathcal{H}_{a \rightarrow i}(\xi_i) = \beta^2 \Xi_{a \rightarrow i}^2 + (\beta G_{a \rightarrow i} + \beta \sigma_i^a \xi_i / \sqrt{N}) \tanh(\beta G_{a \rightarrow i} + \beta \sigma_i^a \xi_i / \sqrt{N})$, $G_a = \frac{1}{\sqrt{N}} \sum_{i \in \partial a} \sigma_i^a m_{i \rightarrow a}$, and $\Xi_a^2 = \frac{1}{N} \sum_{i \in \partial a} (1 - m_{i \rightarrow a}^2)$.

APPENDIX B: THE NUMBER OF CANDIDATE FEATURE VECTORS IN THE APPROXIMATE HOPFIELD MODEL

For the approximate Hopfield model, the entropy can be evaluated as $Ns = \sum_i \Delta S_i - (N-1) \sum_a \Delta S_a$, where single feature node contribution reads $\Delta S_i = - \sum_{a \in \partial i} [\frac{1}{2} \ln(1 - \tilde{\beta} C_{a \rightarrow i}) + \frac{\tilde{\beta} C_{a \rightarrow i}}{2(1 - \tilde{\beta} C_{a \rightarrow i})} + \frac{\tilde{\beta}}{2} (1/N + \tilde{G}_{a \rightarrow i}) F'_{a \rightarrow i}] + \ln(2 \cosh(\tilde{\beta} H_i)) - (\tilde{\beta} H_i + \tilde{\beta} H'_i) \tanh(\tilde{\beta} H_i)$, and single data node contribution $\Delta S_a = -\frac{1}{2} \ln(1 - \tilde{\beta} C_a) - \frac{\tilde{\beta} C_a}{2(1 - \tilde{\beta} C_a)} - \frac{\tilde{\beta}}{2} \tilde{G}_a^2 F'_a$, where $\tilde{G}_a = \frac{1}{\sqrt{N}} \sum_{k \in \partial a} \sigma_k^a m_{k \rightarrow a}$, $F'_{a \rightarrow i} = \frac{\tilde{\beta} C_{a \rightarrow i}}{(1 - \tilde{\beta} C_{a \rightarrow i})^2}$, $F'_a = \frac{\tilde{\beta} C_a}{(1 - \tilde{\beta} C_a)^2}$, $C_a = \frac{1}{N} \sum_{k \in \partial a} (1 - m_{k \rightarrow a}^2)$, $H_i = \frac{1}{\sqrt{N}} \sum_{b \in \partial i} \sigma_i^b \tilde{G}_{b \rightarrow i} F_{b \rightarrow i}$, and $H'_i = \frac{1}{\sqrt{N}} \sum_{b \in \partial i} \sigma_i^b \tilde{G}_{b \rightarrow i} F'_{b \rightarrow i}$.

APPENDIX C: A STATISTICAL ANALYSIS OF PRACTICAL BP EQUATIONS

We first statistically analyze the practical BP equations [Eq. (3) in the main text] for the RBM. In a large N limit, the cavity bias can be approximated as $u_{b \rightarrow i} \simeq \frac{\beta \sigma_i^b}{\sqrt{N}} \tanh \beta G_{b \rightarrow i}$. We then define a cavity field as $h_{i \rightarrow a} = \sum_{b \in \partial i \setminus a} u_{b \rightarrow i}$. The sum in the cavity field involves an order of $\mathcal{O}(N)$ terms, which are assumed to be nearly independent under the replica symmetric assumption. Therefore, the cavity field follows a normal distribution with mean zero and variance $\alpha \beta^2 \hat{Q}$, where $\hat{Q} \equiv \langle \tanh^2 \beta G_{b \rightarrow i} \rangle$. Note that $G_{b \rightarrow i}$ is also a random variable subject to a normal distribution with zero mean and variance Q . Q is defined by $Q = \frac{1}{N} \sum_i m_i^2$. To derive the variance of $G_{b \rightarrow i}$, we used $\frac{1}{N} \sum_{k \in \partial b \setminus i} m_{k \rightarrow b}^2 \simeq Q$, which is reasonable in the large N limit (thermodynamics limit). Finally, we arrive at the following thermodynamic recursive equation:

$$Q = \int Dz \tanh^2 \beta \sqrt{\alpha \hat{Q} z}, \quad (\text{C1a})$$

$$\hat{Q} = \int Dz \tanh^2 \beta \sqrt{Q z}, \quad (\text{C1b})$$

where $Dz = \frac{dz e^{-z^2/2}}{\sqrt{2\pi}}$. Note that $Q = 0$ is a stable solution of Eq. (C1) only when $\alpha \leq \alpha_c = \frac{1}{\beta^4}$. The threshold can be obtained by a Taylor expansion of Eq. (C1) around $Q = 0$.

Next, we perform a statistical analysis of the practical BP equation [Eq. (5) in the main text] for the approximate Hopfield model. Similarly, a cavity field defined by $h_{i \rightarrow a} = \frac{1}{\sqrt{N}} \sum_{b \in \partial i \setminus a} \sigma_i^b \tilde{G}_{b \rightarrow i} F_{b \rightarrow i}$ can be approximated by a random variable following a normal distribution with mean zero and variance $\frac{\alpha Q}{(1 - \tilde{\beta}(1 - Q))^2}$, where $Q \equiv \frac{1}{N} \sum_i m_i^2$. Consequently, we obtain the final thermodynamic equation as

$$Q = \int Dz \tanh^2 \left(\frac{\tilde{\beta}}{1 - \tilde{\beta}(1 - Q)} \sqrt{\alpha Q z} \right). \quad (\text{C2})$$

Obviously, $Q = 0$ is a solution of Eq. (C2), which is stable up to $\alpha_c = (\frac{1}{\tilde{\beta}} - 1)^2$. The threshold can be analogously derived by a linear stability analysis.

- [1] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, *Science* **350**, 1332 (2015).
- [2] G. E. Hinton, *Trends Cogn. Sci.* **11**, 428 (2007).
- [3] D. Kersten, P. Mamassian, and A. Yuille, *Annu. Rev. Psychol.* **55**, 271 (2004).
- [4] H. Huang and T. Toyozumi, *Phys. Rev. E* **91**, 050101 (2015).
- [5] G. E. Hinton and R. R. Salakhutdinov, *Science* **313**, 504 (2006).
- [6] Y. Bengio, A. Courville, and P. Vincent, *IEEE Trans. Pattern Anal. Machine Intell.* **35**, 1798 (2013).
- [7] G. Hinton, S. Osindero, and Y. Teh, *Neural Comput.* **18**, 1527 (2006).
- [8] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press, Oxford, 2001).
- [9] M. Mézard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, Oxford, 2009).
- [10] M. Mézard and G. Parisi, *Eur. Phys. J. B* **20**, 217 (2001).
- [11] J. S. Yedidia, W. T. Freeman, and Y. Weiss, *IEEE Trans. Inf. Theory* **51**, 2282 (2005).
- [12] A. Braunstein and R. Zecchina, *Phys. Rev. Lett.* **96**, 030201 (2006).
- [13] B. Derrida, *Phys. Rev. Lett.* **45**, 79 (1980).
- [14] O. C. Martin, M. Mézard, and O. Rivoire, *Phys. Rev. Lett.* **93**, 217205 (2004).
- [15] A. Montanari, *Eur. Phys. J. B* **23**, 121 (2001).
- [16] H. Hai-Ping, *Commun. Theor. Phys.* **63**, 115 (2015).
- [17] H. Huang and Y. Kabashima, *Phys. Rev. E* **90**, 052813 (2014).
- [18] E. Barkai, D. Hansel, and I. Kanter, *Phys. Rev. Lett.* **65**, 2312 (1990).
- [19] H. Huang, [arXiv:1612.01717](https://arxiv.org/abs/1612.01717).
- [20] Y. Iba, *J. Phys. A: Math. Gen.* **32**, 3875 (1999).
- [21] E. Agliari, A. Barra, A. Galluzzi, F. Guerra, and F. Moauro, *Phys. Rev. Lett.* **109**, 268101 (2012).
- [22] S. Cocco, R. Monasson, and V. Sessak, *Phys. Rev. E* **83**, 051123 (2011).
- [23] M. Biehl and A. Mietzner, *J. Phys. A: Math. Gen.* **27**, 1885 (1994).

- [24] T. L. H. Watkin and J. P. Nadal, *J. Phys. A: Math. Gen.* **27**, 1899 (1994).
- [25] D. J. Amit, H. Gutfreund, and H. Sompolinsky, *Phys. Rev. Lett.* **55**, 1530 (1985).
- [26] M. Mézard, [arXiv:1608.01558](https://arxiv.org/abs/1608.01558).
- [27] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, *Proc. IEEE* **86**, 2278 (1998).
- [28] M. Courbariaux, Y. Bengio, and J.-P. David, *Advances in Neural Information Processing Systems 28*, edited by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2015), pp. 3105–3113.
- [29] James J. DiCarlo, D. Zoccolan, and N. C. Rust, *Neuron* **73**, 415 (2012).