

Advanced mean-field theory of the restricted Boltzmann machine

Haiping Huang* and Taro Toyoizumi

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

(Received 4 February 2015; published 18 May 2015)

Learning in restricted Boltzmann machine is typically hard due to the computation of gradients of log-likelihood function. To describe the network state statistics of the restricted Boltzmann machine, we develop an advanced mean-field theory based on the Bethe approximation. Our theory provides an efficient message-passing-based method that evaluates not only the partition function (free energy) but also its gradients without requiring statistical sampling. The results are compared with those obtained by the computationally expensive sampling-based method.

DOI: [10.1103/PhysRevE.91.050101](https://doi.org/10.1103/PhysRevE.91.050101)

PACS number(s): 02.50.Tt, 87.19.L-, 75.10.Nr

Restricted Boltzmann machine (RBM) forms building blocks of a deep belief network [1,2], which is able to learn complex internal representations of structured objects (such as nature image, speech, or handwriting). RBM also has wide applications in computational biology problems, for example, modeling high-dimensional neural activity data from cortical microcolumns [3].

However, learning in RBM is computationally difficult, since gradients of the log-likelihood function need to be computed at every iteration step to update the model parameters. This computation is usually accomplished by the Gibbs-sampling-based method or its variants [4,5], for which the tradeoff between accuracy and convergence speed requires careful considerations. Furthermore, an efficient way to evaluate the partition function (e.g., log-likelihood function for cross-validation analysis) remains unknown.

Here, we develop a mean-field theory for the RBM based on the cavity method (Bethe approximation) [6], which yields an efficient and fully distributed algorithm to evaluate the free energy (partition function) of a RBM of interest. The remarkable efficiency is confirmed by comparing the computation results of gradients of log-likelihood function by Gibbs sampling and the proposed mean-field theory.

A RBM [7,8] consists of one hidden layer and one visible layer without lateral connections between nodes in each layer. We assume the hidden layer has M nodes, while the visible layer has N nodes. Hidden node a with external field h_a is connected to visible node j with field ϕ_j by a symmetric coupling w_{aj} . The energy function for RBM is thus defined by $E = -\sum_{i,a} \sigma_i w_{ai} s_a - \sum_i \sigma_i \phi_i - \sum_a h_a s_a$, where σ_i and s_a are used to specify the state of visible node i and hidden node a , respectively. Due to the conditional independence of hidden nodes' state given σ , the state of the hidden nodes can be marginalized. This leads to the following probability of a visible state:

$$P(\sigma) = \frac{1}{Z} \prod_a [2 \cosh(\mathbf{w}_a \sigma + h_a)] \prod_i e^{\phi_i \sigma_i}, \quad (1)$$

where \mathbf{w}_a denotes the a th row vector of the coupling matrix \mathbf{w} . Z is a normalization constant (also called the partition function) of the model. As a model study, we assume the element of the matrix \mathbf{w} is independently and identically

distributed with a normal distribution with mean zero and variance g/N . We assume that the external field for both layers follows a normal distribution with mean zero and variance v . We denote the ratio between the number of hidden nodes and that of visible nodes by $\alpha = M/N$, where M and N can be arbitrarily large. A schematic representation of a RBM ($M = 3$, $N = 5$) is shown in Fig. 1.

An exact computation of Z requires an exponential computational complexity (2^N), which becomes impossible for a relatively large N . However, advanced mean-field approximation can be used to compute approximate values under certain condition, and its prediction should be compared with numerical simulations. Here, we propose the Bethe approximation [6] to tackle this problem. First, we transform the original model (left panel of Fig. 1) into a factor graph (right panel of Fig. 1) [9], where each square node indicates a Boltzmann factor $2 \cosh(\mathbf{w}_a \sigma + h_a)$ in Eq. (1). Then, we introduce the cavity probability $P_{i \rightarrow a}(\sigma_i)$ that the visible node i takes state σ_i in the absence of the contribution from factor node a [10], and $P_{i \rightarrow a}(\sigma_i)$ satisfies the following self-consistent equations:

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

$$\mu_{b \rightarrow i}(\sigma_i) = \sum_{\{\sigma_j | j \in \partial b \setminus i\}} 2 \cosh(\mathbf{w}_b \cdot \sigma + h_b) \prod_{j \in \partial b \setminus i} P_{j \rightarrow b}(\sigma_j), \quad (2b)$$

where the symbol \propto indicates a normalization constant, $\partial i \setminus a$ denotes the neighbors of node i except factor node a , $\partial b \setminus i$ denotes the neighbors of factor node b except visible node i , and the auxiliary quantity $\mu_{b \rightarrow i}(\sigma_i)$ represents the contribution from factor node b to visible node i given the value of σ_i [10]. With these definitions, the products in Eq. (2) are reasonable under the weak correlations assumption, whereas, the validity of this Bethe approximation should be checked by a stability analysis.

Note that the computation in Eq. (2b) is still intractable due to the summation over all possible σ except σ_i . However, because $\mathcal{U}_{b \rightarrow i} \equiv \sum_{j \in \partial b \setminus i} w_{bj} \sigma_j$ is a sum of a large number of nearly independent random variables, the central-limit theorem implies that the distribution of $\mathcal{U}_{b \rightarrow i}$ is well characterized by its mean and variance [11], i.e., $G_{b \rightarrow i} = \sum_{j \in \partial b \setminus i} w_{bj} m_{j \rightarrow b}$ and $\Xi_{b \rightarrow i}^2 \simeq \sum_{j \in \partial b \setminus i} w_{bj}^2 (1 - m_{j \rightarrow b}^2)$, respectively, where

*physhuang@gmail.com; sites.google.com/site/physhuang




FIG. 1. (Color online) A RBM is composed of one hidden layer and one visible layer. No lateral connections exist within both hidden and visible layers. Each hidden node is connected to all visible nodes with symmetric coupling weights and is responsible for capturing high-order dependence. The original RBM is shown in the left panel as an example of three hidden nodes (solid circles) and five visible nodes (empty circles). The right panel shows a transformed factor graph after marginalization of hidden states for theoretical analysis. Each factor node (square node) represents the probabilistic normalization of a hidden node given the state of all visible nodes (see the main text).

$m_{j \rightarrow b} \equiv \sum_{\sigma_j} \sigma_j P_{j \rightarrow b}(\sigma_j)$ denotes the cavity magnetization (the average of σ_j in the absence of factor node b).

Because we consider the binary spin variable $\sigma_i = \pm 1$, $P_{i \rightarrow a}(\sigma_i)$ and $\mu_{b \rightarrow i}(\sigma_i)$ can be parametrized by $m_{i \rightarrow a}$ and cavity bias $u_{b \rightarrow i}$, respectively. $u_{b \rightarrow i}$ is defined as $\frac{1}{2} \ln \frac{\mu_{b \rightarrow i}(\sigma_i=1)}{\mu_{b \rightarrow i}(\sigma_i=-1)}$. The practical recursive equations, the so-called message passing equations, are thus derived as

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

$$u_{b \rightarrow i} = \frac{1}{2} \ln \frac{\cosh(h_b + G_{b \rightarrow i} + w_{bi})}{\cosh(h_b + G_{b \rightarrow i} - w_{bi})}, \quad (3b)$$

where the $\Xi_{b \rightarrow i}$ dependency in Eq. (3b) drops because of the symmetry of cosh. The cavity magnetization can be understood as the message passing from the visible node to the factor node, while the cavity bias is interpreted as the message passing from the factor node to the visible node. This message passing based computation is much more accurate than naive mean-field approximation [12], which assumes a fully factorized distribution for Eq. (1). In contrast, Eq. (3) captures nearest neighbors' correlations.

Once the iteration of Eq. (3) converges, the free energy of the model can be computed from the fixed-point solution. Under the Bethe approximation, the Bethe free energy is expressed as [10,11]:

$$F = - \sum_i \ln Z_i + (N-1) \sum_a \ln Z_a, \quad (4)$$

where $Z_i = e^{\phi_i} \prod_{b \in \partial i} \mu_{b \rightarrow i}(+1) + e^{-\phi_i} \prod_{b \in \partial i} \mu_{b \rightarrow i}(-1)$, in which $\mu_{b \rightarrow i}(\sigma_i) = 2e^{\Xi_{b \rightarrow i}^2/2} \cosh(h_b + G_{b \rightarrow i} + w_{bi}\sigma_i)$. $Z_a = 2e^{\Xi_a^2/2} \cosh(h_a + G_a)$. G_a and Ξ_a^2 are given by $\sum_{j \in \partial a} w_{aj} m_{j \rightarrow a}$ and $\sum_{j \in \partial a} w_{aj}^2 (1 - m_{j \rightarrow a}^2)$, respectively. Each stable solution of the message passing algorithm of Eq. (3) corresponds to a local minimum of the free-energy function in Eq. (4) [13].

RBM defined in Fig. 1 is basically a densely connected graphical model. Our mean-field theory provides a practical way to estimate the free energy of *single instances* (typical examples of the model). More precisely, we initialize the cavity magnetization and bias on each link of the factor graph by

random values, and then iterate Eq. (3) until it converges within a prescribed accuracy. Note that the overall time complexity is of the order $O(N^2)$; furthermore, the algorithm is fully distributed and thus amenable to large-scale applications.

In the remaining part, we demonstrate the computation of the free energy on large single instances by applying the message passing algorithm and confirm the accuracy of the results by comparing gradients of the log-likelihood with those obtained by the Gibbs sampling method. A stability analysis of the message-passing algorithm is also presented.

We run the message-passing equations on single instances of RBMs as size N , hidden-node density α , and coupling strength g are varied. As displayed in Fig. 2(a), the free-energy density decreases as α increases. Note that the density does

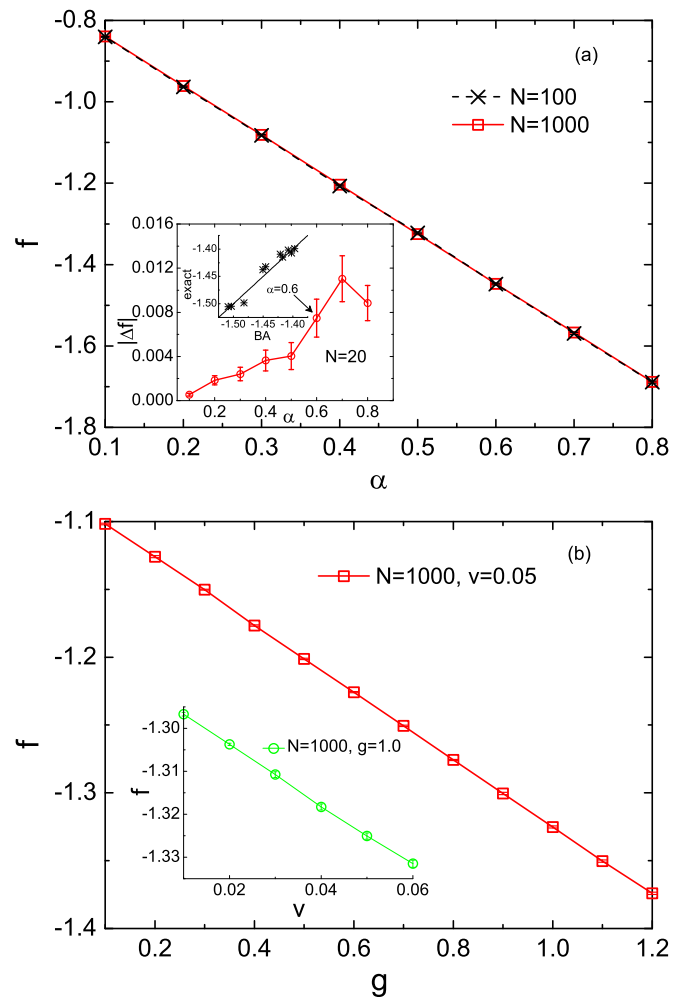


FIG. 2. (Color online) Free-energy density ($f = F/N$) of single instances of RBMs. Iterations of the message passing equations are always converged to produce the data points. The error bars give statistical errors across ten random instances. (a) Free-energy density as a function of α (density of hidden nodes); $g = 1.0$ and $v = 0.05$. The inset shows the absolute difference, $|\Delta f|$, of the free-energy density estimated by the exact enumeration and the Bethe approximation (BA) for $N = 20$ (see a comparison for ten instances with $\alpha = 0.6$ in the inset, where the line indicates equality). (b) Free-energy density as a function of weight strength g and field strength v (shown in the inset). $\alpha = 0.5$.

not change significantly at two large sizes ($N = 100$ and $N = 1000$). Furthermore, the inset of Fig. 2(a) shows that the theoretical result even matches well with the exact enumeration result for small size $N = 20$. As the variance parameter g of weights increases, the free-energy density also decreases [Fig. 2(b)]. The same property also holds when the variance v of external field increases [the inset of Fig. 2(b)]. In the explored range of g (or v) and α , Eq. (3) converges in a few steps to a single fixed point on which the free energy is calculated. Therefore, the Bethe approximation provides an accurate estimation of free energy much faster than other sampling based procedures, which are typically slow to reach an equilibrium state.

The stability of Eq. (3) can also be studied on *single instances*. Apart from the cavity magnetization, we introduce its variance as an extra message denoted by $\mathcal{V}_{i \rightarrow a}$ [14,15]. The evolution of $\mathcal{V}_{i \rightarrow a}$ follows:

$$\mathcal{V}_{i \rightarrow a} = \frac{(1 - m_{i \rightarrow a}^2)^2}{4} \sum_{b \in \partial i \setminus a} \mathcal{P}_{b \rightarrow i} \times [\tanh(\Gamma_{b \rightarrow i}) - \tanh(\Gamma_{b \rightarrow i} - 2w_{bi})]^2, \quad (5)$$

where $\Gamma_{b \rightarrow i} \equiv h_b + G_{b \rightarrow i} + w_{bi}$ and $\mathcal{P}_{b \rightarrow i} \equiv \sum_{j \in \partial b \setminus i} w_{bj}^2 \mathcal{V}_{j \rightarrow b}$. The stability is measured by the total variance $S(t) = \sum_{(i,a)} \mathcal{V}_{i \rightarrow a}(t)$ summed over all connected pairs (i,a) , where t is the iteration step. The explosion of $S(t)$ indicates the instability of the message passing equations, which is related to the divergence of the (nonlinear) spin-glass susceptibility [16], and thus the Bethe approximation becomes inconsistent. We study this effect on single instances of RBM as shown in Fig. 3. A relative strength is denoted as $\lambda = S(t+1)/S(t)$, where t denotes the step at which the iteration converges or exceeds a prefixed maximal number




FIG. 3. (Color online) Stability parameter λ as a function of model parameters. $N = 1000$, $\alpha = 0.5$, and $v = 0.05$. The error bars give statistical errors across ten random instances. The left inset gives stability versus α with $g = 1.0$ and $v = 0.05$. The right inset gives two examples (black and red) taken from the main figure at $g = 2.1$. $S(0)$ is the initial total variance.

($t_{\max} = 500$). λ grows with α and g , and the fluctuation across instances becomes strong near to the critical point ($\lambda = 1$). Note that increasing α has an equivalent effect of increasing g . In the right inset, two typical examples are shown. Near to the critical point, some instances have decaying variance strength ($\lambda < 1$), while some have growing strength ($\lambda > 1$). Equation (5) thus tells us how stable (unstable) the iteration of Eq. (3) is for a particular RBM. The algorithm converges in a few iteration steps to a solution unless the recursive process is close to the instability boundary.

RBM can be used to model the real data, and the parameters are fitted to maximize the probability of observing the training data [2]. This would lead to computation of the following quantities, $m_i \equiv \langle \sigma_i \rangle$, $\hat{m}_a \equiv \langle \tanh(\mathbf{w}_a \boldsymbol{\sigma} + h_a) \rangle$, and $C_{aj} \equiv \langle \tanh(\mathbf{w}_a \boldsymbol{\sigma} + h_a) \sigma_j \rangle$, where the average, $\langle \cdot \rangle$, is taken over the distribution defined in Eq. (1), which is intractable without approximations. Here, m_i is the average of visible state σ_i , \hat{m}_a is the average of hidden state s_a , and C_{aj} is the correlation between s_a and σ_j . Although an accurate evaluation of the above quantities requires a sufficiently long Gibbs sampling (so-called the most difficult negative phase in the machine-learning community [2]), we can compute them by the message passing equations and then compare the results with those obtained by Gibbs sampling to check the consistency of the theory.

Following the same spirit, our theory gives the theoretical evaluation of the above quantities as

$$m_i = \tanh \left(\phi_i + \sum_{b \in \partial i} u_{b \rightarrow i} \right), \quad (6a)$$

$$\hat{m}_a = \int Dx \tanh(\tilde{\Xi}_a x + \tilde{G}_a), \quad (6b)$$

$$C_{aj} \simeq \hat{m}_a m_j + w_{aj} (1 - m_j^2) A_a, \quad (6c)$$

where $Dx \equiv e^{-x^2/2}/\sqrt{2\pi} dx$ is a Gaussian measure. $A_a \equiv 1 - \int Dx \tanh^2(\tilde{\Xi}_a x + \tilde{G}_a)$, $\tilde{G}_a = \sum_{k \in \partial a} w_{ak} m_k + h_a$, and $\tilde{\Xi}_a^2 \simeq \sum_{k \in \partial a} w_{ak}^2 (1 - m_k^2)$ [17]. Equation (6) is computed based on the fixed point of the iterative algorithm [Eq. (3)].

We used alternating Gibbs sampling [5] to evaluate the equilibrium average of the gradients. More precisely, the hidden nodes are updated in parallel according to $P(s_a = 1 | \boldsymbol{\sigma}) = e^{\mathbf{w}_a \boldsymbol{\sigma} + h_a} / [2 \cosh(\mathbf{w}_a \boldsymbol{\sigma} + h_a)]$, while the visible nodes are then all updated in parallel according to $P(\sigma_i = 1 | \mathbf{s}) = e^{\mathbf{w}_i^T \mathbf{s} + \phi_i} / [2 \cosh(\mathbf{w}_i^T \mathbf{s} + \phi_i)]$, where \mathbf{w}_i is the i th column of the weight matrix. Note that the visible nodes are conditionally independent given the hidden states and vice versa [5]. These two steps of updates form one full step of the alternating Gibbs sampling. If this Markov chain is run for a sufficiently long time, the stationary (equilibrium) distribution is expected to be reached, from which the averages can be estimated. We test our theory in a system with $N = 100$ visible nodes and run the Markov chain with 10^6 steps for thermal equilibration and the other 4×10^6 steps to collect a total number of 10^5 samples to calculate the average. We measure the performance by the root-mean-square (RMS) error between the Gibbs sampling (GS) result and the Bethe approximation (BA) result, which is shown in Fig. 4. The RMS error is

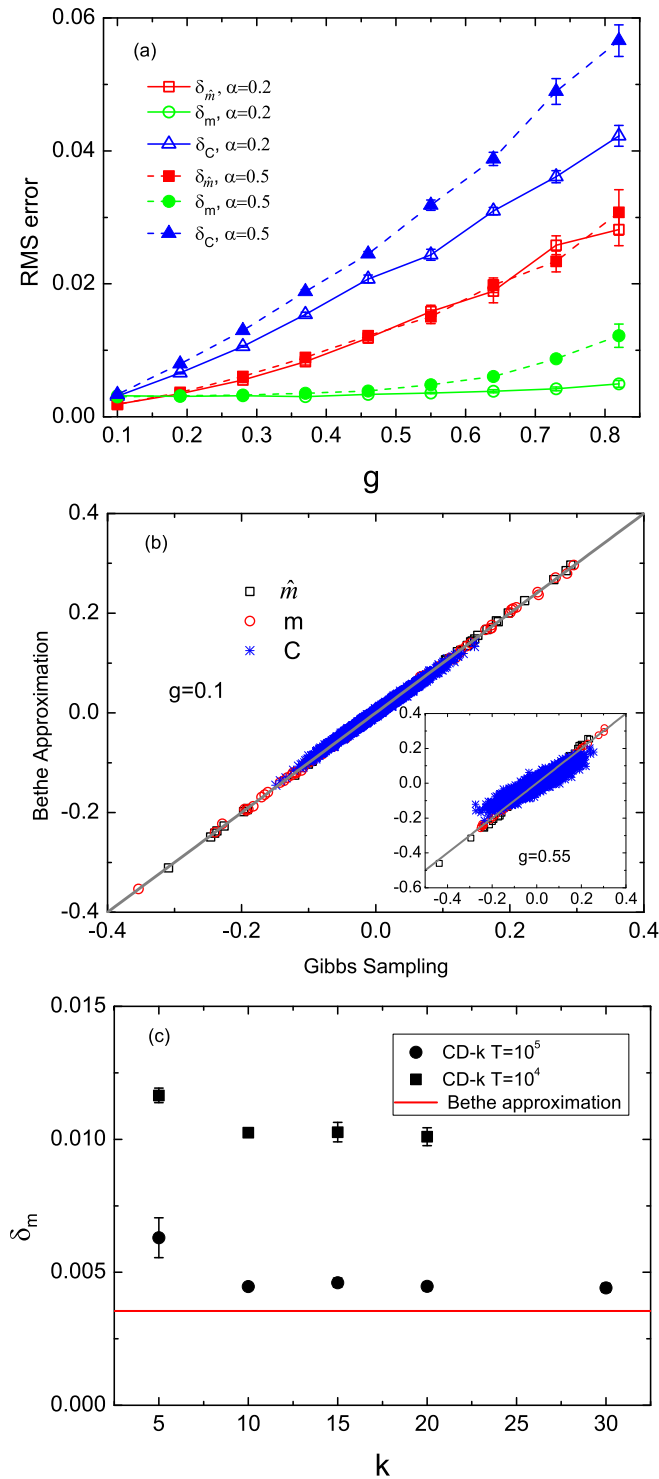


FIG. 4. (Color online) Evaluation performance of mean-field theory in comparison with the Gibbs sampling. Iterations of the recursive equations always converged to produce the data points. The error bars give statistical errors across ten random instances. (a) RMS error as a function of g ; $N = 100$ and $v = 0.02$. (b) Scatter plot for a typical example of $N = 100$, $\alpha = 0.5$, $g = 0.1$, and $v = 0.02$. The inset shows an example of $g = 0.55$ (other parameters do not change). The line indicates equality. (c) RMS error δ_m reached by CD- k as a function of k in comparison to the Bethe approximation; $N = 100$, $\alpha = 0.2$, $g = 0.55$, and $v = 0.02$. The result is averaged over five random instances.

defined as $\delta_Y \equiv \sqrt{\frac{1}{|Y|} \sum_{i=1}^{|Y|} (Y_i^{GS} - Y_i^{BA})^2}$, where Y takes m , \hat{m} , or C , and $|Y|$ indicates the number of these parameters. Small RMS error indicates that the Bethe approximation is accurate.

As shown in Fig. 4(a), all evaluation errors grow with the weight strength g , which is reasonable since our mean-field theory will break down when the network enters a strongly correlated state, as already shown by the stability analysis. In a similar manner, the error grows with the hidden node density α , because each hidden node puts a constraint to the network and all constraints compete with each other to give an equilibrium state, resulting in strong correlations with high α . However, the magnitude of all errors is small, implying that one can acquire accurate estimation of gradients of log-likelihood function by passing messages on a factor graph as well. We show this point more clearly with a scatter plot in Fig. 4(b) (Bethe approximation result versus Gibbs sampling result). This accuracy is obtained by requiring much fewer computational costs compared with the Gibbs sampling. The comparison further confirms the efficiency of the proposed mean-field method across a wide range of model parameters.

Note that the Gibbs sampling result serves as the ground truth here, since we run the Markov chain for a long time. More practically, one can estimate the statistics by k -steps contrastive divergence (CD- k) algorithm [5] that requires a time complexity of $O(kTMN)$, where T denotes the number of sample particles. However, to reach a similar accuracy as the Bethe approximation, it typically requires $k \geq 10$ and $T \sim 10^5$ under the current setup [Fig. 4(c)]. In contrast, the Bethe approximation yields a time complexity of $O(nMN)$ with $n < 100$, where n is the number of iterations and one iteration involves the update of MN cavity messages.

In conclusion, we propose a mean-field theory for the RBM, a widely used model in the machine-learning community and biological data analysis. The theory captures nearest neighbors' correlations by operating on the cavity factor graph (by removing factor nodes), leading to an approximate estimation of the free-energy function (log-likelihood function) for single instances of large-size networks, for which the standard Gibbs sampling procedure becomes prohibitively slow to get a reliable result (e.g., for evaluating the log-likelihood function for cross-validation analysis). Moreover, we replace the normal Gibbs sampling with a mean-field computation based on message passing algorithm, to estimate the gradients of log-likelihood function and show its efficiency by extensive numerical simulations on single instances. The nature of this fast inference lies in the fact that the information is exchanged locally between factor nodes and visible nodes, to reach a coherent fixed point, which may provide a computation paradigm for probabilistic inference in neural networks. We expect the mean-field theory inspired calculation will be useful in practical applications and bring more insights to understand the RBM and its role in deep learning [18].

This work was supported by RIKEN Brain Science Institute and the Brain Mapping by Integrated Neurotechnologies for Disease Studies (Brain/MINDS) by the Ministry of Education, Culture, Sports, Science and Technology of Japan (MEXT).

- [1] G. E. Hinton and R. R. Salakhutdinov, *Science* **313**, 504 (2006).
- [2] Y. Bengio, A. Courville, and P. Vincent, *IEEE Trans. Pattern Anal. Machine Intell.* **35**, 1798 (2013).
- [3] U. Köster, J. Sohl-Dickstein, C. M. Gray, and B. A. Olshausen, *PLoS Comput. Biol.* **10**, e1003684 (2014).
- [4] G. Hinton, *Neural Comput.* **14**, 1771 (2002).
- [5] G. Hinton, S. Osindero, and Y. Teh, *Neural Comput.* **18**, 1527 (2006).
- [6] M. Mézard and G. Parisi, *Eur. Phys. J. B* **20**, 217 (2001).
- [7] P. Smolensky, *Information Processing in Dynamical Systems: Foundations of Harmony Theory* (MIT Press, Cambridge, MA, 1986), pp. 194–281.
- [8] Y. Freund and D. Haussler, Technical Report (Santa Cruz, CA, 1994).
- [9] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, *IEEE Trans. Inf. Theory* **47**, 498 (2001).
- [10] M. Mézard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, Oxford, 2009).
- [11] H. Huang, K. Y. M. Wong, and Y. Kabashima, *J. Phys. A: Math. Theor.* **46**, 375002 (2013).
- [12] H. J. Kappen and F. B. Rodriguez, *Neural Comput.* **10**, 1137 (1998).
- [13] T. Heskes, *Neural Comput.* **16**, 2379 (2004).
- [14] Y. Kabashima, *J. Phys. A* **36**, 11111 (2003).
- [15] A. Montanari and F. Ricci-Tersenghi, *Phys. Rev. B* **70**, 134406 (2004).
- [16] O. Rivoire, G. Biroli, O. Martin, and M. Mézard, *Eur. Phys. J. B* **37**, 55 (2004).
- [17] H. Huang and Y. Kabashima, *J. Stat. Mech.: Theory Exp.* (2014) P05020.
- [18] P. Mehta and D. J. Schwab, [arXiv:1410.3831](https://arxiv.org/abs/1410.3831).