

Deterministic and Generalized Framework for Unsupervised Learning with Restricted Boltzmann Machines

Eric W. Tramel,¹ Marylou Gabri e,² Andre Manoel,¹ Francesco Caltagirone,³ and Florent Krzakala^{2,4}

¹OWKIN, Inc., New York, New York 10022, USA

²Laboratoire de Physique Statistique,  cole Normale Sup rieure, PSL University, Paris 75005, France

³Snips, Paris 75002, France

⁴Universit  Pierre et Marie Curie, Sorbonne Universit s, Paris 75005, France
and LightOn, Paris 75005, France

 (Received 17 February 2017; revised manuscript received 24 July 2018; published 8 October 2018)

Restricted Boltzmann machines (RBMs) are energy-based neural networks which are commonly used as the building blocks for deep-architecture neural architectures. In this work, we derive a deterministic framework for the training, evaluation, and use of RBMs based upon the Thouless-Anderson-Palmer (TAP) mean-field approximation of widely connected systems with weak interactions coming from spin-glass theory. While the TAP approach has been extensively studied for fully visible binary spin systems, our construction is generalized to latent-variable models, as well as to arbitrarily distributed real-valued spin systems with bounded support. In our numerical experiments, we demonstrate the effective deterministic training of our proposed models and are able to show interesting features of unsupervised learning which could not be directly observed with sampling. Additionally, we demonstrate how to utilize our TAP-based framework for leveraging trained RBMs as joint priors in denoising problems.

DOI: [10.1103/PhysRevX.8.041006](https://doi.org/10.1103/PhysRevX.8.041006)

Subject Areas: Computational Physics,
Statistical Physics

I. INTRODUCTION

The past decade has witnessed a groundswell of research in machine learning, bolstered by the deep-learning revolution and the resurgence of neural networks [1]. Since their inception, researchers have identified the deep connection between neural networks and statistical mechanics. Perhaps the most well-known unsupervised neural models studied through the lens of statistical physics have been the Hopfield model [2,3] and the Boltzmann machine [4]. These models were proposed from a connectionist perspective of cognitive science and were studied in the context of emergent representation in unsupervised machine learning.

We can look to the Hopfield model to directly observe some of the contributions of physics to both machine learning and cognitive sciences. For example, by applying techniques from the study of spin glasses, Amit, Gutfreund, and Sompolinsky [3] were famously able to derive the memory capacity of the Hopfield model and provide a concrete understanding of the dynamics of the model via the study of its phase transitions. This fundamental

understanding of the behavior of the Hopfield model provides insight into the complexities of associative memory.

The closely related Boltzmann machine is an undirected stochastic neural network which finds its physics parallel in Ising spin-glass models [5]. Specifically, for this model, one is interested in the *inverse* problem: *learning* the couplings between spins in order to generate a particular set of configurations at equilibrium. The process of learning couplings, or *training*, is often referred to as the inverse Ising problem in the physics literature [6–8]. However, because couplings exist only between pairs of spins for the fully visible Ising spin glass, such models have limited practical application, as they cannot successfully capture higher-order correlations which might exist in a set of training configurations.

For this reason, the general Boltzmann machine introduces a set of unobserved *latent* spins. The effect of these latent spins is to abstract high-order correlations within the set of observed spins. While an optimal training of the couplings would potentially lead to a very effective general model of high-dimensional joint distributions, the intractability of this joint latent model confounds the practical application of general Boltzmann machines.

A *restricted* Boltzmann machine (RBM) is a special case of the general Boltzmann machine, where couplings exist only between latent and observed spins. This bipartite structure is key to the efficient and effective training of

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

RBM [9]. RBMs have found many applications in machine-learning problems as diverse as dimensionality reduction [10], classification [11], collaborative filtering [12], feature learning [13], and topic modeling [14]. Additionally, RBMs can be stacked into multilayer neural networks, which have played a historically fundamental role in pretraining deep-network architectures [15,16]. These constructions, known as deep-belief networks, were the first truly *deep* neural architectures, leading to the current explosion of activity in deep learning [17]. While access to vast training data sets has made such pretraining dispensable for certain tasks, RBMs remain a fundamental tool in the theory of unsupervised learning. As such, a better understanding of RBMs can be key to future developments in emergent machine intelligence.

To date, the most popular and effective approaches to training RBMs have centered on differing flavors of short-chain Monte Carlo sampling [9,18], which we cover in detail in what follows. While such techniques can yield trained RBMs which produce sampled configurations very similar to the target data set and can be used in a number of applications as detailed previously, they do not bridge the gap in understanding *what* the RBM has learned. Furthermore, understanding the modes, or internal representations, of the RBM with sampling-based frameworks has mostly consisted of subjective comparisons of sampled configurations as well as a subjective analysis of the couplings themselves, often referred to as *receptive fields* in the machine-learning literature.

Additionally, comparing two trained models, or even monitoring the training of one model, becomes problematic when using sampling-based investigative tools. For example, annealed techniques [19] can provide estimates of the log-likelihood of a model but only at a large computational cost [20,21]. At a much lower computational cost, pseudolikelihoods can be used to monitor training, but the estimates produced in this manner are inaccurate, as compared to annealed importance sampling (AIS) [19], and even AIS can fail to detect model divergence in practice [22].

In the present work, we seek to address these concerns by developing a deterministic framework to train, compare, and analyze RBMs as well as to leverage their modeling power for inference tasks. We accomplish this via statistical physics techniques through the use of the Thouless-Anderson-Palmer (TAP) formalism of the spin-glass theory [5,23–25]. In this manner, we produce a model which no longer refers to a stochastic model possessing an intractable Gibbs measure but to a *TAP machine*: an entirely self-consistent mean-field model which operates as a classical RBM but which admits deeper introspection via deterministic inference. TAP machines also naturally handle nonbinary variables as well as deep architectures. While deep Boltzmann machines’ (DBMs) [26] state-of-the-art training algorithms mix both Monte Carlo sampling and “naïve” mean-field approximation, a deep TAP machine relies entirely on the TAP mean-field approximation.

Under this interpretation, a TAP machine is not a generative probabilistic model but a deterministic model defining a set of representational magnetizations for a given training data set. Advantageously, this learning output can be computed exactly in a finite time by converging a fixed-point iteration, in contrast to the indeterminate stopping criterion of Markov-chain Monte Carlo (MCMC) sampling. This contrast is a major distinction between the TAP machine and the classical RBM, for which the true probability density function is intractable. At its core, the TAP machine training consists of arranging the minima, *solutions*, in the proposed TAP-approximated free energy so as to maximize the correlation between these solutions and the data set. In our experiments, we demonstrate how to track the growth and geometry of these solutions as a novel way to investigate the progress of unsupervised learning. We also show how to use a trained TAP machine as a prior for inference tasks.

The paper is organized as follows. In Sec. II, we formally describe the classical binary RBM and review the literature on RBM training and analysis. Subsequently, in Sec. III, we describe our proposed modification of the binary RBM to a model with arbitrary real-valued distributions with bounded support. Next, in Sec. IV, we briefly describe how to apply belief propagation to perform inference in the setting of real-valued spins. The details of this approach are pedagogically described in Appendices A and B. In Sec. V, we derive the TAP approximation of the real-valued RBM via a high-temperature expansion of a two-moment Gibbs free energy. Then, in Sec. VI, we detail how to convert this approximation to a practical training algorithm. In Sec. VII, we conduct a series of experiments on real data sets, demonstrating how to use the properties of the TAP machine interpretation to provide insight into the unsupervised learning process. We additionally show how to use a trained model for bit-flip correction as a simple example of leveraging a TAP machine for inference tasks. Last, in Appendix C, we detail the derivations of necessary distribution-specific functions.

II. RESTRICTED BOLTZMANN MACHINES

RBMs [27] are latent-variable generative models often used in the context of unsupervised learning. A set of weights and biases, the model parameters of the RBM, which correspond to the *couplings* and *local fields* present in the system, constructs an energy as a function of the data points from which follows a Gibbs-Boltzmann probability density function. In the well-known *binary* RBM, for which all visible and latent variables are in $\{0, 1\}$, the RBM distribution is

$$P(\mathbf{x}; W, \boldsymbol{\theta}) = \frac{1}{\mathcal{Z}[W, \boldsymbol{\theta}]} \sum_{\mathbf{h}} e^{\sum_{ij} W_{ij} x_i h_j + \sum_i b_i x_i + \sum_j c_j h_j}, \quad (1)$$

where $\boldsymbol{\theta} = \{\mathbf{b}, \mathbf{c}\}$ is the set of *local potentials*, i.e., the set of values which define the biases acting on each variable, and




FIG. 1. Factor-graph representation of the RBM distribution. Variables are indicated by circles, with latent variables denoted by shaded circles. The shaded rectangles indicated layer partitions within the RBM structure. Factors are represented by squares, with the right-hand side factors representing the pairwise relationships between variables and the left-hand side factors representing the influence of the localized prior distributions on the variables.

$$\mathcal{Z}[W, \boldsymbol{\theta}] = \sum_{\mathbf{x}} \sum_{\mathbf{h}} e^{\sum_{ij} W_{ij} x_i h_j + \sum_i b_i x_i + \sum_j c_j h_j}. \quad (2)$$

Here, we use the notation $\sum_{\mathbf{x}}$ and $\sum_{\mathbf{h}}$ to refer to sums over the entire space of possible configurations of visible and latent variables, respectively. When taken with respect to the parameters of the model, $\mathcal{Z}[W, \boldsymbol{\theta}]$ is known as the *partition function*. We give a factor-graph representation of the RBM distribution in Fig. 1.

As evidenced by Eq. (2), an exact computation of the normalizing partition function, and thus the probability of a given high-dimensional data point, is inaccessible in practice. Sophisticated Monte Carlo (MC) schemes relying on importance sampling [20,21] can produce estimates and bounds of the partition but at the cost of a substantial computation, running on the timescale of days or even weeks.

Thankfully, a precise estimate of the normalization is unnecessary for many RBM applications. Additionally, the bipartite structure of the RBM, which admits only couplings between the hidden and visible variables, can be leveraged to construct efficient sampling schemes. This approach is demonstrated in the contrastive divergence (CD) of Ref. [9], where very short-chain block-Gibbs sampling is shown to be sufficient for adequate RBM training. The CD approach consists of a sampling chain alternating between samples drawn from the conditional probabilities of each layer, which are dependent on the conditional expectations at the previously sampled layer. Specifically, the conditional probabilities for the hidden and visible units factorize as

$$P(\mathbf{x}|\mathbf{h}) = \prod_i \text{sigm}\left(b_i + \sum_j W_{ij} h_j\right), \quad (3)$$

$$P(\mathbf{h}|\mathbf{x}) = \prod_j \text{sigm}\left(c_j + \sum_i W_{ij} x_i\right), \quad (4)$$

where $\text{sigm}(y) = (1 + e^{-y})^{-1}$ is the logistic sigmoid function.

In order to learn the parameters of the RBM for a given training data set, one looks to maximize the following log-likelihood:

$$\ln P(\mathbf{x}; W, \boldsymbol{\theta}) = \ln \sum_{\mathbf{h}} e^{\sum_{ij} W_{ij} x_i h_j + \sum_i b_i x_i + \sum_j c_j h_j} - \ln \mathcal{Z}[W, \boldsymbol{\theta}], \quad (5)$$

via a gradient ascent on the parameters W and $\boldsymbol{\theta}$. Commonly, one does not calculate these gradients for each data point from the training set but instead calculates the gradients in average across M data points, often referred to as a *minibatch*. At each minibatch, the gradients of Eq. (5) are given as

$$\Delta W_{ij} = \langle x_i h_j \rangle_{\mathbf{x}} - \langle x_i h_j \rangle_{\text{sampled}}, \quad (6)$$

$$\Delta b_i = \langle x_i \rangle_{\mathbf{x}} - \langle x_i \rangle_{\text{sampled}}, \quad (7)$$

$$\Delta c_j = \langle h_j \rangle_{\mathbf{x}} - \langle h_j \rangle_{\text{sampled}}, \quad (8)$$

where $\langle \cdot \rangle_{\text{sampled}}$ refers to averages over particles sampled from the model and $\langle \cdot \rangle_{\mathbf{x}}$ refers to the so-called *clamped* expectations, where the values of \mathbf{x} are fixed to the training-data samples in the minibatch. In the case of the expectations involving hidden units, which are unobserved and therefore have no training data, Ref. [9] originally proposed the use of configurations sampled from $P(\mathbf{h}|\mathbf{x}; W, \mathbf{b}, \mathbf{c})$. However, one could also use the exact conditional expectations directly to calculate these clamped averages, especially in cases where sampling from these conditionals may be problematic.

Since Ref. [9], there have been a number of proposed modifications to the core sampling-based training scheme described above. The persistent trick [18] takes neatly advantage of the iterative gradient ascent over minibatches to quickly obtain thermalized Markov chains through Gibbs sampling at no extra computational cost over one step CD (CD-1). Nevertheless, the probability density function of a trained RBM is typically highly multimodal, thus making this sampling inexact. Indeed, in such glassy landscapes, mixing becomes very slow as Markov chains become stuck in metastable states, leading to over- and underrepresented, as well as missed, modes of the high-dimensional distribution, which, in turn, produces high variance estimates of means and correlations. A more accurate sampling can be achieved using parallel tempering [28,29], where particles are swapped between multiple Markov chains running at differing temperatures. This

approach, however, requires not only the additional computational burden of running more chains but also further tuning of hyperparameters, such as the number of chains and at which temperatures to run them.

As accurate sampling-based inference on RBMs can be costly, it would seem that their usefulness is limited. As learning of the RBM via a gradient ascent is dependent upon this inference, the difficulty of training a generative model with a high degree of accuracy is compounded. However, RBMs have proven to be very useful in many applications where sampling from a full-fledged generative model is unneeded. For instance, RBMs can be used as an unsupervised “feature extraction” pretraining for feed-forward networks [10,15,30]. RBMs have also been used for data-imputation tasks, e.g., image in-painting, label recovery [11], or collaborative filtering [12] by reconstructing missing data with a single visible-hidden-visible step. In truth, the CD-k training algorithm which popularized RBMs, first with binary units [9], then Gaussian units [10,31], and finally with arbitrary units [32] does not use thermalized samples to evaluate means and correlations. Instead, it focuses on the region of the configuration space nearest to the training data set [28] by using short block-Gibbs Markov chains, starting from training data points, to get fast and low variance estimates of moments. However, CD-k is prone to learn spurious minima in a configuration space far from the data, as it does not explore this region during training [28]. It also does not systematically increase the true likelihood of training data [33]. However, this training strategy has been found to be very efficient in the applications mentioned above, which consistently remain close to the data set in configuration space. One finds that CD falls short for applications which require long-chain MCMC sampling from the trained RBM, as this represents a fundamental mismatch between the training and application of the RBM. In order to address some of these shortcomings of sampling-based approaches, we now turn our attention to deterministic mean-field approximations of the RBM.

The TAP approximation [23] for disordered systems relies on the deterministic inference of approximated magnetizations, from which one can obtain estimators of all kinds of observables, starting from the log-partition or *free energy*. TAP is derived from a small-weight expansion of the variational approach and can be considered as an extension of the naïve mean-field (NMF) method (see Refs. [34,35] for the original derivation and Refs. [36,37] for pedagogical expositions). Previous works which have attempted to make use of the NMF approximation of the RBM have shown negative results [18,38].

The TAP approximation was first considered for Boltzmann machines in the context of small random models without hidden units in Ref. [39]. In the recent work of Ref. [40], this approximation is extended to a practical training algorithm for full-scale binary RBMs which is

shown to be competitive with persistent contrastive divergence [18] when applied to real-world data sets. In parallel, other works use TAP, and the related Bethe approximation, to perform inference on binary Boltzmann machines [41–44].

In the next sections, we detail how to rewrite the RBM model in the nonbinary case, for generalized distributions on the visible *and* hidden units, similar in spirit to Ref. [32]. However, unlike earlier techniques, we approach the problem of estimating the normalization of the RBM model via the tools of statistical mechanics, resulting in a fully deterministic framework for RBM inference, training, and application.

III. GENERAL DISTRIBUTIONS FOR RBMS

We now turn our attention to the case of the general RBM (GRBM), where the distributions of the hidden and visible units are not fixed. We define the distribution of interest in the following manner:

$$P(\mathbf{x}, \mathbf{h}; W, \boldsymbol{\theta}) = \frac{1}{\mathcal{Z}[W, \boldsymbol{\theta}]} e^{\sum_{i,j} x_i W_{ij} h_j} \times \prod_i P_i^v(x_i; \theta_i^v) \prod_j P_j^h(h_j; \theta_j^h), \quad (9)$$

where the sum over i, j indicates a sum over the all visible and hidden units in the model and $\boldsymbol{\theta} = \{\theta_1^v, \dots, \theta_{N_v}^v, \theta_1^h, \dots, \theta_{N_h}^h\}$ are the parameters defining the local distributions, P_i^v and P_j^h , on, respectively, the variables of \mathbf{x} and the variables of \mathbf{h} . In the case that $\mathbf{x} \in \{\pm 1\}^{N_v}$ and $\mathbf{h} \in \{\pm 1\}^{N_h}$, we can see that the distribution above reduces to a bipartite spin-glass model with $\boldsymbol{\theta}$ representing the local fields acting on the system spins; the fields \mathbf{b} and \mathbf{c} are for binary spins as described in Eq. (1). This specific case is simply a binary RBM, as described in the previous section, and which we have already considered within an extended mean-field framework in Ref. [40]. The important distinction with the model we evaluate here is that we do not assume a binary discrete distribution on the variables but instead allow for a formulation where the variables in the system can possess distributions on discrete- or real-valued bounded support. By considering this more general class of models, one can include a wide range of different models, including the Hopfield model and spike-and-slab RBMs [45], and data sets, such as images or genomic data, by varying the distributions of the hidden and visible units. The distribution of visible variables is obtained by marginalizing out the latent variables:

$$P(\mathbf{x}; W, \boldsymbol{\theta}) = \int \left(\prod_j dh_j \right) P(\mathbf{x}, \mathbf{h}; W, \boldsymbol{\theta}), \quad (10)$$

giving the log-probability

$$\begin{aligned} \ln P(\mathbf{x}; W, \boldsymbol{\theta}) &= -\ln \mathcal{Z}[W, \boldsymbol{\theta}] + \sum_i \ln P_i(x_i; \theta_i^v) \\ &+ \sum_j \ln \int dh_j P_j(h_j; \theta_j^h) e^{h_j \sum_l W_{lj} x_l}. \end{aligned} \quad (11)$$

If we take the gradients of Eq. (11) with respect to the model parameters, in the case of distribution terms $\boldsymbol{\theta}$ we find

$$\Delta \theta_i^v = \left\langle \frac{\partial}{\partial \theta_i^v} \ln P_i(x_i; \theta_i^v) \right\rangle_{\mathbf{x}} - \frac{\partial}{\partial \theta_i^v} \ln \mathcal{Z}[W, \boldsymbol{\theta}] \quad (12)$$

and

$$\Delta \theta_j^h = \left\langle \frac{\partial}{\partial \theta_j^h} \ln P_j(h_j; \theta_j^h) \right\rangle_{\mathbf{x}} - \frac{\partial}{\partial \theta_j^h} \ln \mathcal{Z}[W, \boldsymbol{\theta}], \quad (13)$$

which are generalizations of Eqs. (7) and (8). However, in the case of the gradient with respect to the couplings, we find

$$\Delta W_{ij} = \left\langle x_i \cdot f \left(\sum_l W_{lj} x_l; \theta_j^h \right) \right\rangle_{\mathbf{x}} - \frac{\partial}{\partial W_{ij}} \ln \mathcal{Z}[W, \boldsymbol{\theta}], \quad (14)$$

where the function

$$f(B; \theta) = \frac{\int dh h \cdot P(h; \theta) e^{Bh}}{\int dh P(h; \theta) e^{Bh}} \quad (15)$$

computes the conditional expectation of h_l knowing the value of the visible units. The one-dimensional integral in Eq. (15) can be computed either analytically or numerically. Note, moreover, that the data-dependent term is tractable thanks to the bipartite structure of the one-hidden layer RBM.

In contrast to the data-dependent terms, the second terms of Eqs. (12)–(14) require knowledge of the partials of the log normalization with respect to the parameter of interest. However, this term cannot be written exactly, as the explicit calculation of the normalization is intractable. Rather than resorting to sampling, we attempt to approximate the free energy $\mathbb{F} = -\ln \mathcal{Z}[W, \boldsymbol{\theta}]$ in a parametric and deterministic way, as in Ref. [40]. In the next section, we discuss how belief propagation can be used to estimate \mathbb{F} and to conduct inference on RBMs.

IV. APPROXIMATION VIA BELIEF PROPAGATION

One method by which we might estimate the partition of \mathbb{F} is via belief propagation (BP) [46], which we review in Appendix A for pairwise models such as the RBM. Essentially, given a factor graph for some joint statistical

model, such as that of our RBM in Fig. 1, the BP algorithm attempts to estimate a set of marginal distributions at each variable. In the case of treelike graphs, BP provides an exact calculation of these marginals. The application of BP to factor graphs containing cycles, *loopy* BP, is not guaranteed to provide accurate estimates of the marginals. However, often these estimated marginals have a significant overlap with the true ones [47]. Additionally, it is known that the solutions of loopy BP are the fixed points of the Bethe free energy [47], which allows the construction of an approximation of \mathbb{F} . Applying this to the inverse learning problem, one can compute the gradients of this Bethe free energy in terms of the parameters of the model, allowing a gradient ascent on the Bethe-approximated log-likelihood of the training data.

One significant hurdle in the application of loopy BP to RBM learning for real-valued variables is that the messages propagated on the edges of the factor graph are continuous probability distribution functions (PDFs). In the case of discrete variables, such as Ising or Potts spins, BP messages can be written using magnetizations or the full discrete probability mass function (PMF), respectively. For binary variables, both BP and mean-field approximations of fully connected Boltzmann machines are considered in Ref. [41] in the context of inference with fixed parameters. A similar study of binary RBMs is conducted with loopy BP in Ref. [48]. It is important to note that both of these studies investigate the properties of Boltzmann machines with *independent identically distributed* random weights. While such studies permit many analytical tools for studying the behavior of the RBM, one cannot directly map these observations to RBM inference in practice, where trained weights may exhibit strong correlations both within and between receptive fields.

In order to construct a BP algorithm for PDFs over real-valued support, one requires a finite memory description of the messages. Some examples of such descriptions are given in nonparametric BP [49], moment matching [50], and relaxed BP (RBP) [51]. In Appendix B, following the example of RBP, we show how to arrive at a two-moment approximation of the continuous BP messages via a small-weight expansion on the RBM coupling parameters W . There, we also show the RBP approximated free energy of pairwise models as well as demonstrate the need for distributions with bounded support in order to preserve bounded messages.

In the next section, building upon this derivation, we consider mean-field approximations of the RBM via a high-temperature Plefka expansion.

V. TAP APPROXIMATION FOR PAIRWISE MODELS

While one could utilize the RBP approach in order to estimate the free energy of a generalized real-valued spin model, as detailed in the earlier section, such an approach

might not be desirable in practice. Specifically, if one wishes to solve the inverse learning problem, estimating model parameters from a given data set, it is necessary to estimate the gradients of the model parameters with respect to the model likelihood for each parameter update. Using a steepest-ascent approach, as detailed in Sec. III, requires one to estimate these gradients many thousands of times. For systems of large size N , the RBP scales quite poorly. Estimating a gradient requires the iteration of the RBP equations on $O(N^2)$ messages. Additionally, one must distinguish between cavity terms ($i \rightarrow j$) and marginal terms ($\rightarrow i$). If the final $O(N)$ gradients we desire can be estimated using the marginal terms alone, then requiring an iteration on the $O(N^2)$ set of messages is an extremely costly operation.

Instead, one can turn to a mean-field approach, writing the free energy, and its stationary points, in terms of the marginals alone, which can be done by including certain correction terms, up to a specified degree in the weights. In the context of RBMs, such approaches have been proposed at both the first order, the naive mean-field [38], and the second order, using the so-called Thouless-Anderson-Palmer (TAP) [23] equations for introducing an additional correction term [40,41,44,52,53]. In the case of a GRBM with arbitrary distributions on each unit, however, we must rederive the TAP approximation in terms of parameters of these distributions as well as the approximate marginalized distribution at each site, up to their first two moments. This task turns out to be closely related to the TAP approach to low-rank matrix factorization [52–56].

While it is possible to derive the stationarity conditions for the inferred marginals from the RBP messages directly by Taylor expansion, we rather focus on the free energy directly that will provide the gradients we require for training the GRBM parameters via a high-temperature expansion we present below.

Last, we point out that the TAP free-energy second-order (TAP) term depends on the statistical properties of the weight distribution. The derivation presented below assumes independent identically distributed weights, scaling as $O(1/\sqrt{N})$. This assumption is a simplification, as in practice the weight distribution cannot be known *a priori*. The distribution depends on the training data and changes throughout the learning process according to the training hyperparameters. The adaptive TAP (adaTAP) formalism [57] attempts to correct this assumption by allowing one to directly compute the “correct” second-order correction term for a realization of the W matrix without any hypothesis on how its entries are distributed. Although this algorithm is the most principled approach, its computational complexity almost rules out its implementation. Moreover, practical learning experiments indicate that training using adaTAP does not differ significantly from TAP assuming independent identically distributed weights. A more detailed discussion of the computational complexity and learning performance is described in Appendix D.

A. Derivation of the TAP free energy

We now discuss the main steps of the derivation of the TAP free energy, which was originally performed in Refs. [34,35]. We do not aim to perform it in full detail; a more pedagogical and comprehensive derivation can be found in Appendix B of Ref. [56].

In the limit $N \rightarrow \infty$, if we assume that the entries of W scale as $O(1/\sqrt{N})$ and that all sites are widely connected, on the order of the size of the system, then we can apply the TAP approximation—a high-temperature expansion of the Gibbs free energy up to second order [34,35]. In the case of a Boltzmann distribution, the global minima of the Gibbs free energy, its value at equilibrium, matches the Helmholtz free energy \mathbb{F} [58]. We derive a two-variable parameterization of the Gibbs free energy derived via the Legendre transform [59]. Additionally, we show that this two-variable Gibbs free energy both is variational and attains the Helmholtz free energy at its minima. For clarity of notation, we make our derivation in terms of a pairwise interacting Hamiltonian without enforcing any specific structure on the couplings; the bipartite structure of the RBM is reintroduced in Sec. VI.

We first introduce the inverse temperature term β to facilitate our expansion as $\beta \rightarrow 0$, $P(\mathbf{x}; \beta, W, \boldsymbol{\theta}) = e^{-\beta(\mathcal{H} - \mathbb{F}_\beta)}$, where

$$\mathcal{H} = -\sum_{(i,j)} W_{ij} x_i x_j - \frac{1}{\beta} \sum_i \ln P_i(x_i; \theta_i). \quad (16)$$

Note that β can be interpreted as the weight scaling; i.e., one can rescale the weights so that $W \leftarrow \beta W$.

We wish to derive our two-variable Gibbs free energy for this system in terms of the first two moments of the marginal distributions at each site. To accomplish this goal, we proceed as in Refs. [35,57] by first defining an augmented system under the effect of *two* auxiliary fields:

$$-\beta \mathbb{F}_\beta(\boldsymbol{\lambda}, \boldsymbol{\xi}) = \ln \int d\mathbf{x} e^{-\beta \mathcal{H} + \sum_i \lambda_i x_i + \sum_i \xi_i x_i^2}, \quad (17)$$

where we see that, as the fields disappear, $\mathbb{F}_\beta(\mathbf{0}, \mathbf{0}) = \mathbb{F}_\beta$, and we recover the true Helmholtz free energy.

We additionally note the following identities for the augmented system, namely,

$$\frac{\partial}{\partial \lambda_i} [-\beta \mathbb{F}_\beta(\boldsymbol{\lambda}, \boldsymbol{\xi})] = \langle x_i \rangle_{\boldsymbol{\lambda}, \boldsymbol{\xi}}, \quad (18)$$

$$\frac{\partial}{\partial \xi_i} [-\beta \mathbb{F}_\beta(\boldsymbol{\lambda}, \boldsymbol{\xi})] = \langle x_i^2 \rangle_{\boldsymbol{\lambda}, \boldsymbol{\xi}}, \quad (19)$$

where $\langle \cdot \rangle_{\boldsymbol{\lambda}, \boldsymbol{\xi}}$ is the average over the augmented system for the given auxiliary fields. Since the partial derivatives of the field-augmented Helmholtz free energy generate the cumulants of the Boltzmann distribution, it can be shown that the

Hessian of $-\beta\mathbb{F}_\beta(\lambda, \xi)$ is simply a covariance matrix and, subsequently, positive semidefinite. Hence, $-\beta\mathbb{F}_\beta(\lambda, \xi)$ is a convex function in terms of $\lambda \times \xi$. This convexity is shown to be true for all log partitions of exponential family distributions in Ref. [59].

We now take the Legendre transform of $-\beta\mathbb{F}_\beta(\lambda, \xi)$, introducing the conjugate variables \mathbf{a} and \mathbf{c} :

$$-\beta\mathbb{G}_\beta(\mathbf{a}, \mathbf{c}) = -\beta \sup_{\lambda, \xi} \left\{ \mathbb{F}_\beta(\lambda, \xi) + \frac{1}{\beta} \sum_i [\lambda_i a_i + \xi_i (c_i + a_i^2)] \right\}, \quad (20)$$

where we define the solution of the auxiliary fields at which $\mathbb{G}_\beta(\mathbf{a}, \mathbf{c})$ is defined as $\lambda^* \triangleq \lambda^*(\beta, \mathbf{a}, \mathbf{c})$, $\xi^* \triangleq \xi^*(\beta, \mathbf{a}, \mathbf{c})$, where we make explicit the dependence of the auxiliary field solutions on the values of the conjugate variables. Looking at the stationary points of these auxiliary fields, we find that

$$\frac{\partial}{\partial \lambda_i^*} [-\beta\mathbb{G}_\beta(\mathbf{a}, \mathbf{c})] = \frac{\partial}{\partial \lambda_i^*} [-\beta\mathbb{F}(\lambda^*, \xi^*)] - a_i, \quad (21)$$

$$\therefore a_i = \langle x_i \rangle_{\lambda^*, \xi^*}, \quad (22)$$

and

$$\frac{\partial}{\partial \xi_i^*} [-\beta\mathbb{G}_\beta(\mathbf{a}, \mathbf{c})] = \frac{\partial}{\partial \xi_i^*} [-\beta\mathbb{F}(\lambda^*, \xi^*)] - (c_i + a_i^2), \quad (23)$$

$$\therefore c_i = \langle x_i^2 \rangle_{\lambda^*, \xi^*} - \langle x_i \rangle_{\lambda^*, \xi^*}^2. \quad (24)$$

The implication of these identities is that $\mathbb{G}_\beta(\mathbf{a}, \mathbf{c})$ cannot be valid unless it meets the self-consistency constraints that \mathbf{a} and \mathbf{c} are the first and second (central) moments, respectively, of the marginal distributions of the augmented system.

Now, we wish to show the correspondence of $\mathbb{G}_\beta(\mathbf{a}, \mathbf{c})$ to the Helmholtz free energy at its unique minimum. First, let us look at the stationary points of $-\beta\mathbb{G}_\beta(\mathbf{a}, \mathbf{c})$ with respect to its parameters. We take the derivative with careful application of the chain rule to find

$$\begin{aligned} \frac{\partial}{\partial a_i} [-\beta\mathbb{G}_\beta(\mathbf{a}, \mathbf{c})] &= -\lambda_i^* - 2\xi_i^* a_i \\ &+ \sum_j \frac{\partial}{\partial \lambda_j^*} [-\beta\mathbb{G}] \frac{\partial \lambda_j^*}{\partial a_i} + \sum_j \frac{\partial}{\partial \xi_j^*} [-\beta\mathbb{G}] \frac{\partial \xi_j^*}{\partial a_i} \\ &= -\lambda_i^* - 2\xi_i^* a_i, \end{aligned} \quad (25)$$

with the terms inside the sums going to zero, as the derivatives of \mathbb{G} with respect to λ_j^* and ξ_j^* are zero by definition. Carrying through a very similar computation for the derivative with respect to c_i provides $(\partial/\partial c_i)[- \beta\mathbb{G}_\beta(\mathbf{a}, \mathbf{c})] = \xi_i^*$. This result shows that, at its solution, the Gibbs free energy must satisfy

$$\begin{aligned} \xi_i^* &= 0, \\ \lambda_i^* + 2\xi_i^* a_i &= 0, \end{aligned} \quad (26)$$

which can be true only in the event that the solutions of the auxiliary fields are truly $\lambda^* = 0$ and $\xi^* = 0$. Looking at the *inverse* Legendre transform of the Gibbs free energy for $\lambda = \mathbf{0}$, $\xi = \mathbf{0}$, we find that $\inf_{\mathbf{a}, \mathbf{c}} \mathbb{G}_\beta(\mathbf{a}, \mathbf{c}) = \mathbb{F}_\beta(\mathbf{0}, \mathbf{0}) = \mathbb{F}_\beta$, which implies that the minimum of the Gibbs free energy is equivalent to the Helmholtz free energy. This equivalence holds, since $-\beta\mathbb{G}_\beta(\mathbf{a}, \mathbf{c})$ is convex, as the Legendre transform of a convex function is itself convex. Since the Gibbs free energy can therefore possess only a single solution, then its minimum must satisfy (26) and, therefore, must be $\mathbb{F}_\beta(\mathbf{0}, \mathbf{0}) = \mathbb{F}_\beta$.

Finally, we can now rewrite the Gibbs free energy defined in Eq. (20) as a function of the moments \mathbf{a} and \mathbf{c} and parameterized by β and the GRBM parameters \mathbf{W} and $\boldsymbol{\theta}$:

$$-\beta\mathbb{G}_\beta(\mathbf{a}, \mathbf{c}; \boldsymbol{\theta}) = \ln \int d\mathbf{x} e^{-\beta\tilde{\mathcal{H}}}, \quad (27)$$

where

$$\tilde{\mathcal{H}} \triangleq \mathcal{H} - \sum_i \frac{\lambda_i^*(\beta)}{\beta} (x_i - a_i) - \frac{\xi_i^*(\beta)}{\beta} (x_i^2 - c_i - a_i^2), \quad (28)$$

where the Lagrange multipliers are given as functions of the temperature in order to make clear the order in which we apply $\beta \rightarrow 0$ later.

As this exact form of the Gibbs free energy is just as intractable as the original free energy, we apply a Taylor expansion in order to generate an approximate Gibbs free energy [34]. We make this expansion at $\beta = 0$, as in the limit of infinite temperature all interactions between sites vanish and the system can be described only in terms of individual sites and their relationship to the system average and their local potentials, allowing the Gibbs free energy to be decomposed into a sum of independent terms. Specifically, if we take the expansion up to s terms,

$$-\beta\tilde{\mathbb{G}}^{(s)}(\beta, \mathbf{a}, \mathbf{c}; \boldsymbol{\theta}) = \ln \tilde{\mathcal{Z}}_\beta|_{\beta=0} + \sum_{p=1}^s \frac{\beta^p}{p!} \frac{\partial^p}{\partial \beta^p} [\ln \tilde{\mathcal{Z}}_\beta]_{\beta=0}, \quad (29)$$

where $\tilde{\mathcal{Z}}_\beta$ is the normalization of the Boltzmann distribution defined by $\tilde{\mathcal{H}}$ at temperature β . At $\beta = 0$, we can find the first term of the expansion directly:

$$\begin{aligned} \ln \tilde{\mathcal{Z}}_0 &= -\sum_i \lambda_i^*(0) a_i - \sum_i \xi_i^*(0) (a_i^2 + c_i) \\ &+ \sum_i \ln \int dx_i P_i(x_i; \theta_i) e^{\lambda_i^*(0)x_i - \xi_i^*(0)x_i^2}, \end{aligned} \quad (30)$$

where we recognize that the last term is simply the normalization of the Gaussian-product distribution whose moments are defined in Eqs. (B9) and (B10).

We define the TAP free energy by writing the remainder of the expansions terms in the specific case of $s = 2$ [34,35]:

$$-\beta\tilde{\mathbb{G}}_{\beta}^{(2)}(\mathbf{a}, \mathbf{c}; \boldsymbol{\theta}) = \sum_i \ln Z_i(\lambda_i^*(0), \xi_i^*(0); \theta_i) - \sum_i \lambda_i^*(0)a_i - \sum_i \xi_i^*(0)(a_i^2 + c_i) + \beta \sum_{(i,j)} W_{ij} a_i a_j + \frac{\beta^2}{2} \sum_{(i,j)} W_{ij}^2 c_i c_j, \quad (31)$$

where

$$Z_i(B, A; \theta) \triangleq \int dx P_i(x; \theta) e^{(1/2)Ax^2 - Bx}. \quad (32)$$

Note that the last two terms in Eq. (31) come from the Taylor expansion in Eq. (29) and are related to the derivatives of λ^* and ξ^* evaluated at 0.

We still need to determine the values of $\lambda^*(0)$ and $\xi^*(0)$, which is done by taking the stationarity of the expanded Gibbs free energy with respect to \mathbf{a} and \mathbf{c} :

$$A_i \triangleq -2\xi_i(0)^* = -\beta^2 \sum_{j \in \partial_i} W_{ij}^2 c_j, \quad (33)$$

$$B_i \triangleq \lambda_i(0)^* = A_i a_i + \beta \sum_{j \in \partial_i} W_{ij} a_j, \quad (34)$$

where we make the definitions of \mathbf{A} and \mathbf{B} for convenience and as a direct allusion to the definitions of the cavity sums for BP inference, given in Eqs. (B5) and (B6).

Conversely, by deriving the stationarity conditions of the auxiliary fields, we obtain the self-consistency equations for \mathbf{a} and \mathbf{c} , which show us that the TAP free energy is valid only when the following self-consistencies hold:

$$a_i = f_a(B_i, A_i; \theta_i), \quad c_i = f_c(B_i, A_i; \theta_i), \quad (35)$$

where f_a and f_c are defined from Eq. (32) via $f_a \triangleq (\partial/\partial B_i) \log Z_i$ and $f_c \triangleq (\partial^2/\partial B_i^2) \log Z_i$.

Substituting these values closes the free energy on the marginal distribution moments \mathbf{a} and \mathbf{c} and completes our derivation of a free-energy approximation which is defined by $O(N)$ elements versus the $O(N^2)$ values required by RBP.

B. Solutions of the TAP free energy

As given, the TAP free energy is valid only when the self-consistency equations are met at its stationary points. Thus, only a certain set of \mathbf{a} and \mathbf{c} can have any physical

meaning. Additionally, we know that only at the minima of the exact Gibbs free energy will we have a correspondence with the original exact Helmholtz free energy.

While the exact Gibbs free energy in terms of the moments \mathbf{a} and \mathbf{c} is convex for exponential family $P_0(\mathbf{x}; \boldsymbol{\theta})$ such that $e^{-\beta\tilde{\mathcal{H}}} \geq 0$, the TAP free energy can possess multiple stationary points whose number increases rapidly as β grows [5]. Later, in Sec. VII, we show that, as GRBM training progresses, so does the number of identified TAP solutions, which can be explained due to the variance of the weights W growing with training. For fixed $\beta = 1$, as we use in our practical GRBM implementation, the variance of the weights serves as an effective inverse temperature, and its increasing magnitude has an identical effect to the system cooling as β increases.

Additionally, while the Gibbs free energy has a correspondence with the Helmholtz free energy at its minimum, this correspondence is not necessarily true for the TAP free energy. The approximate nature of the second-order expansion removes this correspondence. Thus, it may not be possible to ascertain an accurate estimate of the Helmholtz free energy from a *single* set of inferred \mathbf{a} and \mathbf{c} , as shown in Fig. 2. In the case of the naïve mean-field estimate of the Gibbs free energy, it is true that $\mathbb{F}_{\beta} \leq \tilde{\mathbb{G}}_{\beta}^{(1)}$, which implies that one should attempt to find the minima of $\tilde{\mathbb{G}}_{\beta}^{(1)}$ in order to find a more accurate estimate of \mathbb{F}_{β} , a foundational principle in variational approaches. However, while the extra expansion term in the TAP free energy should improve its accuracy in modeling \mathbb{G}_{β} over \mathbf{a} and \mathbf{c} , it does not provide a lower bound, and so an estimate of \mathbb{F}_{β} from the TAP free energy could be an under- or overestimate.

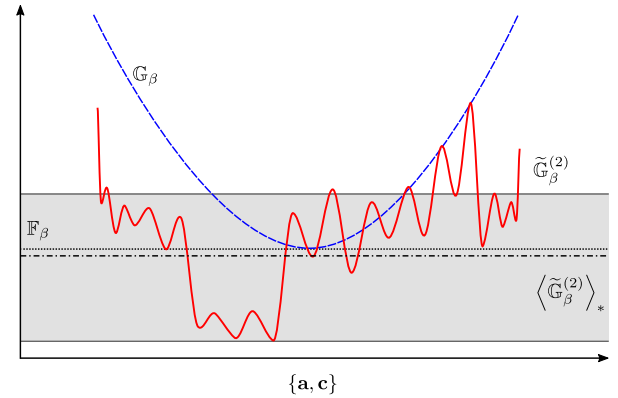


FIG. 2. Cartoon description for estimating the Helmholtz free energy (dotted line) via the Gibbs (blue dashed line) and TAP (red line) free energies. For this example of a convex Gibbs free energy, there exists one unique minimum over the moments \mathbf{a} and \mathbf{c} , and the Gibbs free energy here matches \mathbb{F}_{β} . The range of TAP free energies (gray box) gives a boundary on the location of \mathbb{F}_{β} . Averaging the TAP free energies (dash-dotted line) provides an estimate of \mathbb{F}_{β} .

Instead, one might attempt to obtain an estimate of the Helmholtz free energy by utilizing either all or a subset of the equilibrium solutions of the TAP free energy. Since there is no manner by which we might distinguish the equilibrium moments by their proximity to the unknown \mathbb{F}_β , averaging the TAP free energy across its solutions, denoted as $\langle \mathbb{G}_\beta^{(2)} \rangle_*$, can serve as a simple estimator of \mathbb{F}_β [5]. In Ref. [60], a weighting is introduced to the average, correcting the Helmholtz free-energy estimate at a low temperature by removing the overinfluence of the exponential number of high-energy solutions. The weights in this approach are proportional to the exponents of each solution's TAP free energy, placing a much stronger emphasis on low-energy solutions.

However, such an approach is not well justified in the our general setting of P_i , where we expect large deviations from the expectations derived for the Sherrington-Kirkpatrick (SK) model. Additionally, while this weighting scheme is shown across the *entire set* of solutions for a particular random SK model, in our case, we are interested in the solution space centered on the particular data set we wish to model. Since the solutions are computed by iterating the TAP self-consistency equations, we can easily probe this region by initializing the iteration according to the training data. Subsequently, we do not encounter a band of high-energy solutions that we must weight against. Instead, we obtain a set of solutions over a small region of the support of the TAP free energy. Because of the uniformity of these solutions, unweighted averaging across the solutions seems the best approach in terms of efficiency. In the subsequent section, we explore some of these properties numerically for trained RBMs.

VI. RBMS AS TAP MACHINES

To utilize the TAP inference of Sec. V, we need to write the TAP free energy in terms of the variables of the RBM. To clarify the bipartite structure of the GRBM, we rewrite the TAP free energy in terms of the hidden and visible variables at fixed temperature $\beta = 1$:

$$\begin{aligned}
& -\mathbb{F}_{\text{RBM}}(\mathbf{a}^v, \mathbf{c}^v, \mathbf{a}^h, \mathbf{c}^h; \boldsymbol{\theta}) \\
&= \sum_i \left\{ \ln Z_i^v(B_i^v, A_i^v; \theta_i^v) - B_i^v a_i^v + \frac{1}{2} A_i^v [(a_i^v)^2 + c_i^v] \right\} \\
&+ \sum_j \left\{ \ln Z_j^h(B_j^h, A_j^h; \theta_j^h) - B_j^h a_j^h + \frac{1}{2} A_j^h [(a_j^h)^2 + c_j^h] \right\} \\
&+ \sum_{ij} \left\{ W_{ij} a_i^v a_j^h + \frac{1}{2} W_{ij}^2 c_i^v c_j^h \right\}, \quad (36)
\end{aligned}$$

where $\{\mathbf{a}^v, \mathbf{c}^v\}$ and $\{\mathbf{a}^h, \mathbf{c}^h\}$ are the means and variances of the visible and hidden variables, respectively.

As in Sec. V, solutions of the TAP GRBM free energy can be found by a fixed-point iteration, as shown in Algorithm 1, which bears much resemblance to the

Algorithm 1. TAP inference for GRBMs.

Input: $W, \boldsymbol{\theta}$

Initialize: $t = 0, \mathbf{a}^{v,(0)}, \mathbf{c}^{v,(0)}$

repeat

Hidden side updates

$$A_j^{h,(t+1)} = -\sum_i W_{ij}^2 c_i^{v,(t)}$$

$$B_j^{h,(t+1)} = A_j^{h,(t+1)} a_j^{h,(t)} + \sum_i W_{ij} a_i^{v,(t)}$$

$$a_j^{h,(t+1)} = f_a^h(B_j^{h,(t+1)}, A_j^{h,(t+1)}; \theta_j^h)$$

$$c_j^{h,(t+1)} = f_c^h(B_j^{h,(t+1)}, A_j^{h,(t+1)}; \theta_j^h)$$

Visible side updates

$$A_i^{v,(t+1)} = -\sum_j W_{ij}^2 c_j^{h,(t+1)}$$

$$B_i^{v,(t+1)} = A_i^{v,(t+1)} a_i^{v,(t)} + \sum_j W_{ij} a_j^{h,(t+1)}$$

$$a_i^{v,(t+1)} = f_a^v(B_i^{v,(t+1)}, A_i^{v,(t+1)}; \theta_i^v)$$

$$c_i^{v,(t+1)} = f_c^v(B_i^{v,(t+1)}, A_i^{v,(t+1)}; \theta_i^v)$$

$t = t + 1$

until convergence

AMP iteration derived in the context of compressed sensing [61,62] and matrix factorization [53,55,56]. We note that, rather than updating over the entire system at each time step, fixing one side at a time has the effect of stabilizing the fixed-point iteration. For clarity, Algorithm 1 is written for a single initialization of the visible marginals. However, as noted in Sec. VB, there exist a large number of initialization-dependent solutions to the TAP free energy. Thus, in order to capture the plurality of modes present in the TAP free-energy landscape, one should run this inference independently for many different initializations.

If the use case of the GRBM requires that we only train the GRBM tightly to the data space (e.g., data imputation), it makes sense to fix the initializations of the inference to points drawn from the data set:

$$\mathbf{a}^{v,(0)} = \mathbf{x}^{(m)}, \quad \text{where } m \in \{1, \dots, M\}, \quad (37)$$

$$\mathbf{c}^{v,(0)} = \mathbf{0}. \quad (38)$$

In order to train the GRBM more holistically, structured random initializations can help probe modes outside of the data space. In this work, we do not employ this strategy, restricting ourselves to a deterministic initialization.

For a set of TAP solutions $\{\mathbf{a}_k, \mathbf{c}_k, \mathbf{B}_k, \mathbf{A}_k\}$ for $k \in \{1, \dots, K\}$ at fixed GRBM parameters $\{W, \boldsymbol{\theta}\}$, the TAP-approximated log-likelihood can be written as

$$\begin{aligned}
\ln P(\mathbf{x}; W, \boldsymbol{\theta}) &\approx \sum_i \ln P_i(x_i; \theta_i^v) \\
&+ \sum_j \ln Z_j^h \left(\sum_i W_{ij} x_i, 0; \theta_j^h \right) \\
&+ \frac{1}{K} \sum_k \mathbb{F}_{\text{RBM}}(\mathbf{a}_k^v, \mathbf{c}_k^v, \mathbf{a}_k^h, \mathbf{c}_k^h; \boldsymbol{\theta}), \quad (39)
\end{aligned}$$

where $Z_j^h(B, 0; \theta)$ is the normalization of the conditional expectation of Eq. (15), since $f_a(B, 0; \theta) = f(B; \theta)$.

After reintroducing an averaging of the log-likelihood over the samples in the minibatch, the gradients of the TAP-approximated GRBM log-likelihood with respect to the model parameters are given by

$$\Delta W_{ij} \approx \frac{1}{M} \sum_m x_i^{(m)} f_a^h \left(\sum_i W_{ij} x_i^{(m)}, 0; \theta_j^h \right) - \frac{1}{K} \sum_k \{ a_{i,k}^v a_{j,k}^h + W_{ij} c_{i,k}^v c_{j,k}^h \}, \quad (40)$$

$$\Delta \theta_j^h \approx \frac{1}{M} \sum_m \frac{\partial}{\partial \theta_j^h} \left[\ln Z_j^h \left(\sum_i W_{ij} x_i^{(m)}, 0; \theta_j^h \right) \right] - \frac{1}{K} \sum_k \frac{\partial}{\partial \theta_j^h} [\ln Z_j^h(B_{j,k}^h, A_{j,k}^h; \theta_j^h)], \quad (41)$$

$$\Delta \theta_i^v \approx \frac{1}{M} \sum_m \frac{\partial}{\partial \theta_i^v} [\ln P_i^v(x_i^{(m)}; \theta_i^v)] - \frac{1}{K} \sum_k \frac{\partial}{\partial \theta_i^v} [\ln Z_i^v(B_{i,k}^v, A_{i,k}^v; \theta_i^v)]. \quad (42)$$

In the presented gradients, we make the point that the set of data samples and the set of TAP solutions can have different cardinality. For example, one might employ a minibatch strategy to training, where the set of data samples used in the gradient calculation might be on the order of 10^2 . However, depending on the application of the GRBM, one might desire to probe a very large number of TAP solutions in order to have a more accurate picture of the representations learned by the GRBM. In this case, one might start with a very large number of initializations, resulting in a very large number, $K \gg M$, of unique TAP solutions. Or, on the contrary, while one might start with a number of initializations equal to M , the number of unique solutions might be $K \ll M$, especially early in training or when the number of hidden units is small.

Using these gradients, a simple gradient ascent with a fixed or monotonically decreasing step size γ can be used to update these GRBM parameters. We present the final GRBM training algorithm in Algorithm 2.

Besides considering nonbinary units, another natural extension of traditional RBMs is to consider additional hidden layers, as in DBMs. It is possible to define and train *deep* TAP machines, as well. Probabilistic DBMs are substantially harder to train than RBMs as the data-dependent (or clamped) terms of the gradient updates (40)–(42) become intractable with depth. Interestingly, state-of-the-art training algorithms retain a Monte Carlo evaluation of other intractable terms while introducing a naïve mean-field approximation of these data-dependent terms. For deep TAP machines, we consistently utilize the TAP equations. The explicit definition and training algorithm are fully described in Appendix E.

Algorithm 2. GRBM training.

Input: $\mathbf{X}, T, M, K, R(\cdot)$

Initialize: $W_{ij}^{(0)} \sim \mathcal{N}(0, \sigma), \theta^{(0)} \leftarrow \mathbf{X}$

repeat

for all minibatches of size M **do**

$\mathbf{a}^{(t+1)}, \mathbf{c}^{(t+1)}, \mathbf{B}^{(t+1)}, \mathbf{A}^{(t+1)} \leftarrow$ Algorithm 1 ($W^{(t)}, \theta^{(t)} \times K$)

$W_{ij}^{(t+1)} \leftarrow W_{ij}^{(t)} + \gamma \Delta W_{ij}^{(t)} + \gamma \epsilon R(W_{ij}^{(t)}) + \eta \Delta W_{ij}^{(t-1)}$

$\theta_j^{h,(t+1)} \leftarrow \theta_j^{h,(t)} + \gamma \Delta \theta_j^{h,(t)}$

$\theta_i^{v,(t+1)} \leftarrow \theta_i^{v,(t)} + \gamma \Delta \theta_i^{v,(t)}$

end for

$t \leftarrow t + 1$

until $t > T$

VII. EXPERIMENTS

A. Data sets

MNIST.—The MNIST handwritten digit data set [63] consists of both a training and a testing set, each with 60 000 and 10 000 samples, respectively. The data samples are real-valued 28×28 pixel 8-bit gray-scale images which we normalize to the dynamic range of $[0, 1]$. The images themselves are centered crops of the digits “0” through “9” in roughly balanced proportion. We construct two separate versions of the MNIST data set. The first, which we refer to as *binary MNIST*, applies a thresholding such that pixel values > 0.5 are set to 1 and all others to 0. The second, *real MNIST*, simply refers to the normalized data set introduced above.

CBCL.—The CBCL face database [64] consists of both face and nonface 8-bit gray-scale 19×19 pixel images. For our experiments, we utilize only the face images. The database contains 2429 training and 472 testing samples of face images. For our experiments, we normalize the samples to the dynamic range of $[0, 1]$.

B. Learning dynamics

We now investigate the behavior of the GRBM over the course of the learning procedure, looking at a few metrics of interest: the TAP-approximated log-likelihood of the training data set, the TAP free energy, and the number of discovered TAP solutions. We note that each of these metrics is unique to the TAP-based model of the GRBM.

While it is empirically shown in Ref. [40] that CD does indeed increase the TAP log-likelihood in the case of binary RBMs, the specific construction of CD is entirely independent from the TAP model of the GRBM. Thus, it is hard to say that a CD or TAP-trained GRBM is “better” in a general case. At present, we present comparisons between TAP GRBMs of varying complexity trained under fixed hyperparameter settings, as indicated in Table I.

In Fig. 3, we see a comparison of the TAP log-likelihood as a function of training *epochs* for binary MNIST for binary RBMs consisting of differing numbers of hidden units. As the

TABLE I. Parameter settings for GRBM training.

	Binary MNIST	Real MNIST	CBCL
N_v	784	784	361
N_h	{25, 50, 100, 500}	{100, 500}	256
M	100	100	20
K	100	100	20
Prior vis.	B.	Tr. Gauss.-B.	Tr. Gauss.
Prior hid.	B.	B.	B.
$R(\cdot)$	ℓ_2	ℓ_2	ℓ_2
γ	0.005	$[10^{-2}, 10^{-5}]$	0.005
ϵ	0.001	0.001	0.01
η	0.5	0.5	0.5

gradient ascent on the log-likelihood is performed batch by batch over the training data, we define one *epoch* to be a single pass over the training data: Every example is presented to the gradient ascent once. The specifics of this particular experiment are given in the caption. We note that, for equal comparison across varying model complexity, this log-likelihood is normalized over the number of visible and hidden units present in the model. In this way, we observe a “per-unit” TAP log-likelihood, which gives us a measure of the concentration of representational power encapsulated in each unit of the model. Increasing values of the normalized TAP log-likelihood indicate that the evaluated training samples are becoming more likely given the state of the GRBM model parameters.

It can be observed that, at each level of complexity, the TAP log-likelihood of the data rapidly increases as the values of W quickly adjust from random initializations to receptive fields correlated with the training data. However, across each of the tested models, by about the 20th epoch the rate of increase of the TAP log-likelihood tapers off to a constant rate of improvement.

For reference, we also show a subset of the trained receptive fields, i.e., the rows of W , for each of the tested experiments. Since the full set of receptive fields would be too large to display, we attempt to show some representative samples in Fig. 4 by looking at the extreme samples in terms of spatial spread or localization and activity over the training set. We observe that the trained GRBMs, in the case of both binary MNIST and real MNIST, are able to learn both the localized and stroke features commonly observed in the literature for binary RBMs trained on the MNIST data set [9,65]. It is interesting to note that even in the case of real MNIST, where we are using the novel implementation of truncated Gauss-Bernoulli visible units (see Appendix C 2), we are able to observe similar learned features as in the case of binary MNIST. We take this as an empirical indication that the proposed framework of GRBM learning is truly learning correlations present in the data set as intended. Finally, we see feature localization increase with the number of hidden units.

To date, understanding “what” an RBM learns from the unlabeled data has mostly been a purely subjective exercise

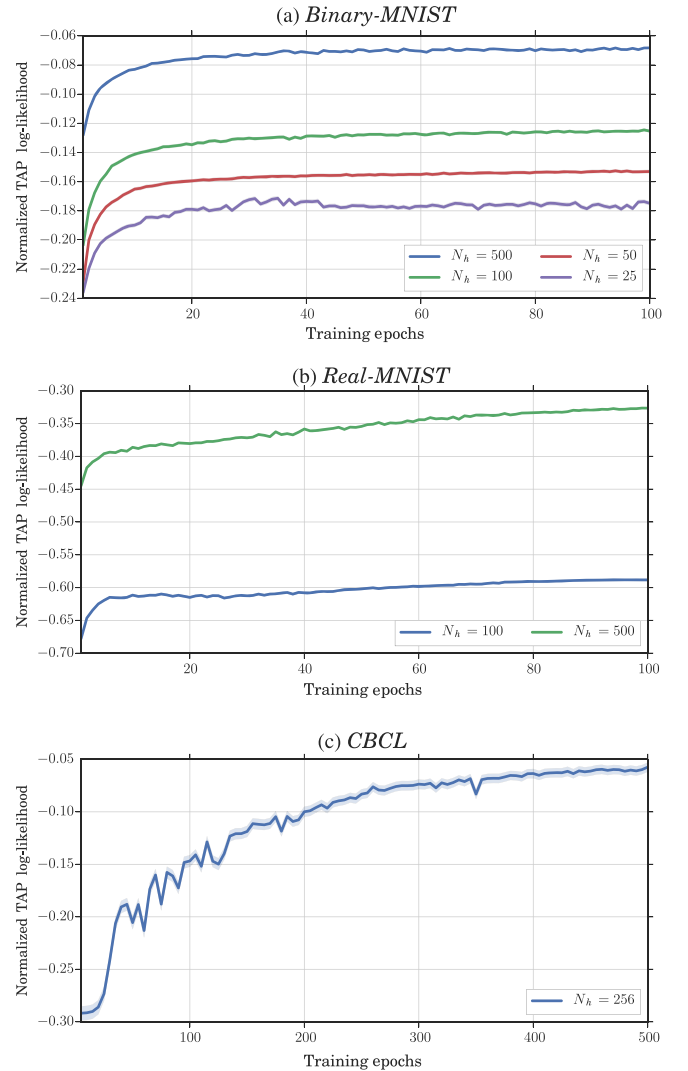


FIG. 3. Training performance over 100 epochs for the tested data sets over varying numbers of hidden units, N_h . Performance is measured in terms of the normalized (per-unit) TAP log-likelihood estimate computed for 10 000 training data samples. The TAP free energy is estimated using the unique TAP solutions thermalized from initial conditions drawn from the data samples, as in Eqs. (37) and (38). Thermalization is determined by the convergence of the magnetizations up to a difference of 10^{-8} in MSE between iterations. Solid lines indicate the average normalized TAP log-likelihood over the tested training samples, and shaded regions indicate standard error.

in studying the receptive fields, as shown in Fig. 4. However, the interpretation of the GRBM as a TAP machine can provide us with a novel insight in the nature and dynamics of GRBM learning via the stationary points of the TAP free energy, which we detail in the next section.

C. Probing the GRBM

Given the deterministic nature of the TAP framework, it is possible to investigate the structure of the modes which a given set of GRBM parameters produces in the free-energy

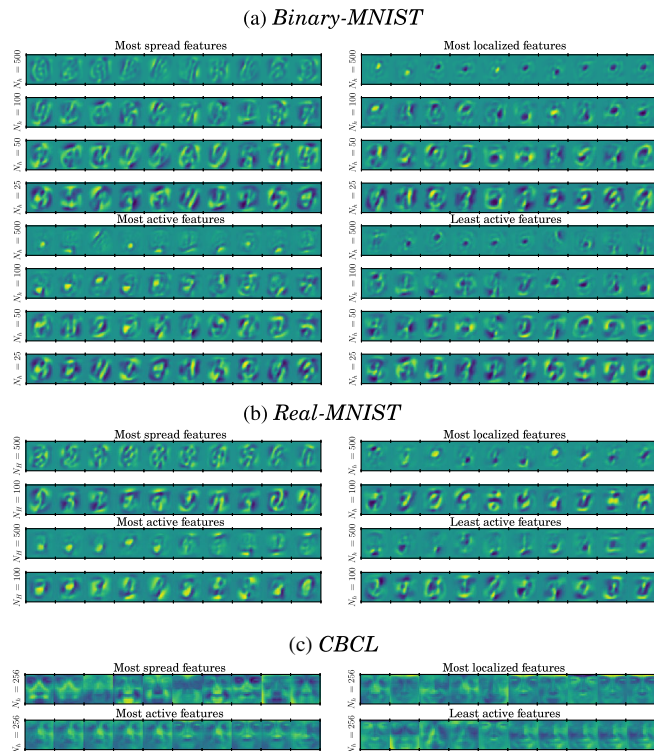


FIG. 4. Subsets of the final receptive fields, i.e., the columns of W , obtained by TAP training of GRBM models with varying numbers of hidden units, N_h . For the receptive fields, dark blue and yellow are mapped to -1 and $+1$, respectively, and green indicates a value of 0. Receptive fields are ranked according to two criteria: first, spread, and conversely localization, as measured by the ℓ_p -norm of the receptive field, for $p = 0.1$, and second, by activation level, as measured by the mean activation of each receptive field’s corresponding hidden unit averaged across the training data set.

landscape. Understanding the nature and concentration of these modes gives us an intuition on the representational power of the GRBM.

To date, observing the modes of a given GRBM model could be approached only via long-chain sampling. Given enough sampling chains from a diverse set of initial conditions, thermalizing these chains produces a set of samples from which one could attempt to derive statistics, such as concentrations of the samples in their high-dimensional space, to attempt to pinpoint the likely modes in the model. However, the number of required chains to resolve these features increases with the dimensionality of the space and the number of potential modes which might exist in the space. Because of this, the numerical evaluation we carry out here would be impractical with sampling techniques.

The RBP and mean-field models of the RBM allow us to directly obtain the modes of the model by running inference to solve the direct problem. Given a diverse set of initial conditions, such as a given training data set, running RBP or TAP provides a deterministic mapping between the initial conditions drawn from the data, as in Eqs. (37) and (38),

and the “nearest” solution of the TAP free energy. If the initial point is drawn from the data set, then this solution can be interpreted as the RBM’s best-matching internal representation for the data point.

If a large number of structurally diverse data points map to a single solution, then this mapping may be an indicator that the GRBM parameters are not sufficient to model the diverse nature of the data and, perhaps, further changes to the model parameters or hyperparameters are required. Conversely, if the number of solutions explodes, being roughly equivalent to the number of initial data points, then this explosion indicates a potential spin-glass phase, that the specific RBM is overtrained, perhaps memorizing the original data samples during training. Additionally, when in such a phase, the large set of TAP solutions may be replete with spurious solutions which convey very little structural information about the data set. In this case, hyperparameters of the model may need to be tuned in order to ensure that the model possesses a meaningful generalization over the data space.

To observe these effects, we obtain a subset of the TAP solutions by initializing the TAP iteration with initial conditions drawn from the data set, running the iteration until convergence, and then counting the unique TAP solutions. We present some measures on these solutions in Fig. 2. Here, we count both the number of unique TAP solutions as well as the distribution of the TAP free energy over these solutions, across training epochs. There are a few common features across the tested data sets. First, the early phase of training shows a marked increase of the TAP free energy, which then gradually declines as training continues. Comparing the point of inflection in the TAP free energy against the normalized TAP log-likelihood shown in Fig. 3 shows that the early phase of GRBM training is dominated by the reinforcement of the empirical moments of the training data, with the GRBM model correlations playing a small role in the gradient of Eq. (14), which makes sense, as the random initialization of $W \sim \mathcal{N}(0, \sigma)$ for $\sigma \approx 10^{-3}$ implies that the hidden units are almost independent of the training data. Thus, the TAP solutions at the early stage of learning are driven by, and correlated with, the local potentials on the hidden and visible variables.

The effect of this influence is that the TAP free-energy landscape possesses very few modes in the data space. Figure 5 shows this effect very clearly, as the number of TAP solutions starts at 1 and then steadily increases with training. Because the positive data term of Eq. (14) is dominant, the GRBM parameters do not appear to minimize the TAP free energy, as we would expect. However, as more TAP solutions appear, the data and model terms of the gradient become balanced, and the TAP free energy is minimized. It is at this point of inflection that we see a leveling off of the normalized TAP log-likelihood.

Second, we observe free-energy bands in the TAP solutions. This feature is especially pronounced in the

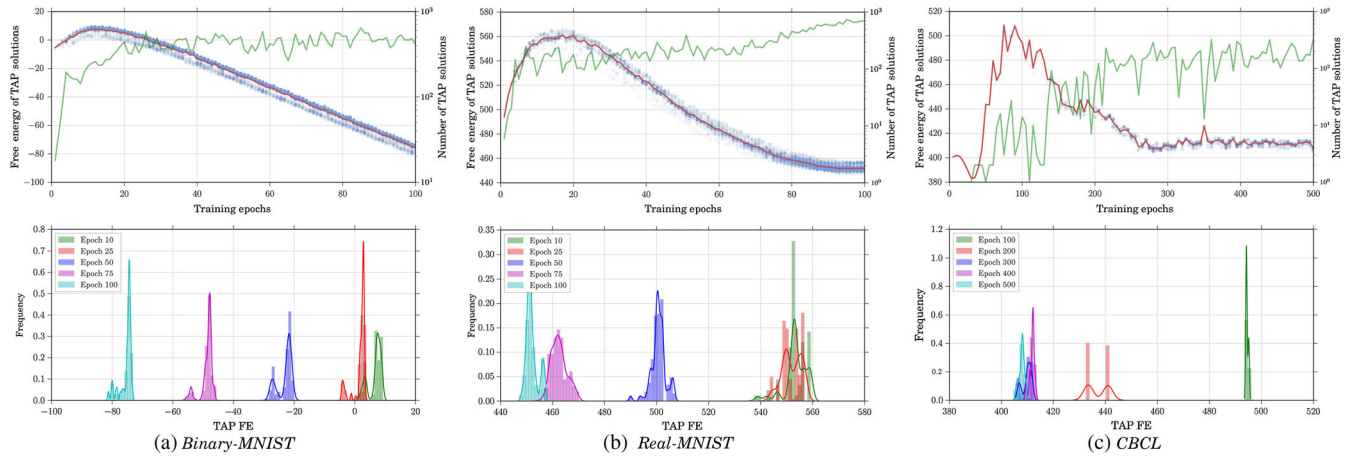


FIG. 5. Distribution of free-energy estimates of TAP solutions as a function of training epochs for the three different data sets. In the case of the two MNIST experiments, the number of hidden units is the same. $N_h = 100$ and 10 000 samples drawn from the training data are used as initial conditions. For CBCL, 2400 training samples are used. Top row: TAP free energy for all unique TAP solutions (transparent blue dots) and the Helmholtz free energy estimate via uniform averaging (red line). The number of unique TAP solutions is also given (green line). Bottom row: Detail of TAP free-energy distributions for slices of training. Histograms are given as bars, while kernel density estimates of the TAP free-energy distribution are given as curves.

case of the binary-MNIST experiment. Here, at all training epochs, there exist two significant modes in the free-energy distribution over the TAP solutions. We see this effect more clearly in the training-slice histograms shown in the bottom row in Fig. 5(a). In the case of the real-MNIST experiment, we see that the free-energy distributions do not exhibit such tight banding, but they do show the presence of some high- and low-energy solutions which persist across training. The main feature across experiments is the multimodal structure of the free-energy distribution. Finally, we note that for both real MNIST and binary MNIST, in the case of $N_h = 100$, we do not empirically observe an explosion of TAP solutions, a potential indicator of a spin-glass phase, since the proportion of unique TAP solutions to the initial data points remains less than 10%.

In order to investigate whether the modes in the TAP free-energy distributions are randomly assigned over configuration space or exist in separate continuous partitions of the configuration space, we need to look at the proximity of the solutions in the configuration space. Because this space cannot be observed in its ambient dimensionality, we project the configuration space into a two-dimensional embedding in Fig. 6. Here, we utilize the well-known Isomap [66] algorithm for calculating a two-dimensional manifold which approximately preserves local neighborhoods present in the original space. Using this visualization, we observe that as training progresses the assignment of high and low free energy to TAP solutions does not appear random in nature but seems to be inherent to the structure of the solutions themselves, that is, their location in the configuration space. Additionally, in Fig. 6, we can see the progression from few TAP solutions to many and how they spread across the configuration space. It is interesting to note how the solutions start from a highly correlated state and then proceed to diversify.

We can also observe the TAP solutions with respect to the initializations which produced them, as shown in Fig. 7. In these charts, we use a similar approach as in Fig. 6, mapping all high-dimensional data points, as well as TAP magnetizations, into a 2D embedding using Isomap. This approach allows us to see, in an approximate way, how the TAP solutions distribute themselves over the data space. We also show how the number of TAP solutions grows from few to many over training and how they maintain a spread distribution over the data space, which demonstrates how the training procedure is altering the parameters of the model so as to place TAP solutions within dense regions of the data space. For the sake of clarity, we have not included lines indicating the attribution of an initial data point to its resultant TAP solution. However, as training progresses, one sees that the TAP solutions act as attractors over the data space, clustering together data points which the TAP machine recognizes as similar.

D. Inference for denoising

Serving as a prior for inference is one particular use case for the TAP machine interpretation of the GRBM. As a simple demonstration, we turn to the common signal processing task of *denoising*. Specifically, given a planted signal, one observes a set of noisy observations which are measures of the true signal corrupted by some stochastic process. Denoising tasks are ubiquitous in signal processing, both at an analog level (e.g., additive and shot noise) and at the level of digital communications (e.g., binary symmetric and erasure channels). The goal of this task is to produce the most accurate estimate of the unknown signal. In the analog case, this estimate may be a measure of the mean-square error (MSE) between the estimate and the true signal. In the binary case, this estimate may be a measure of

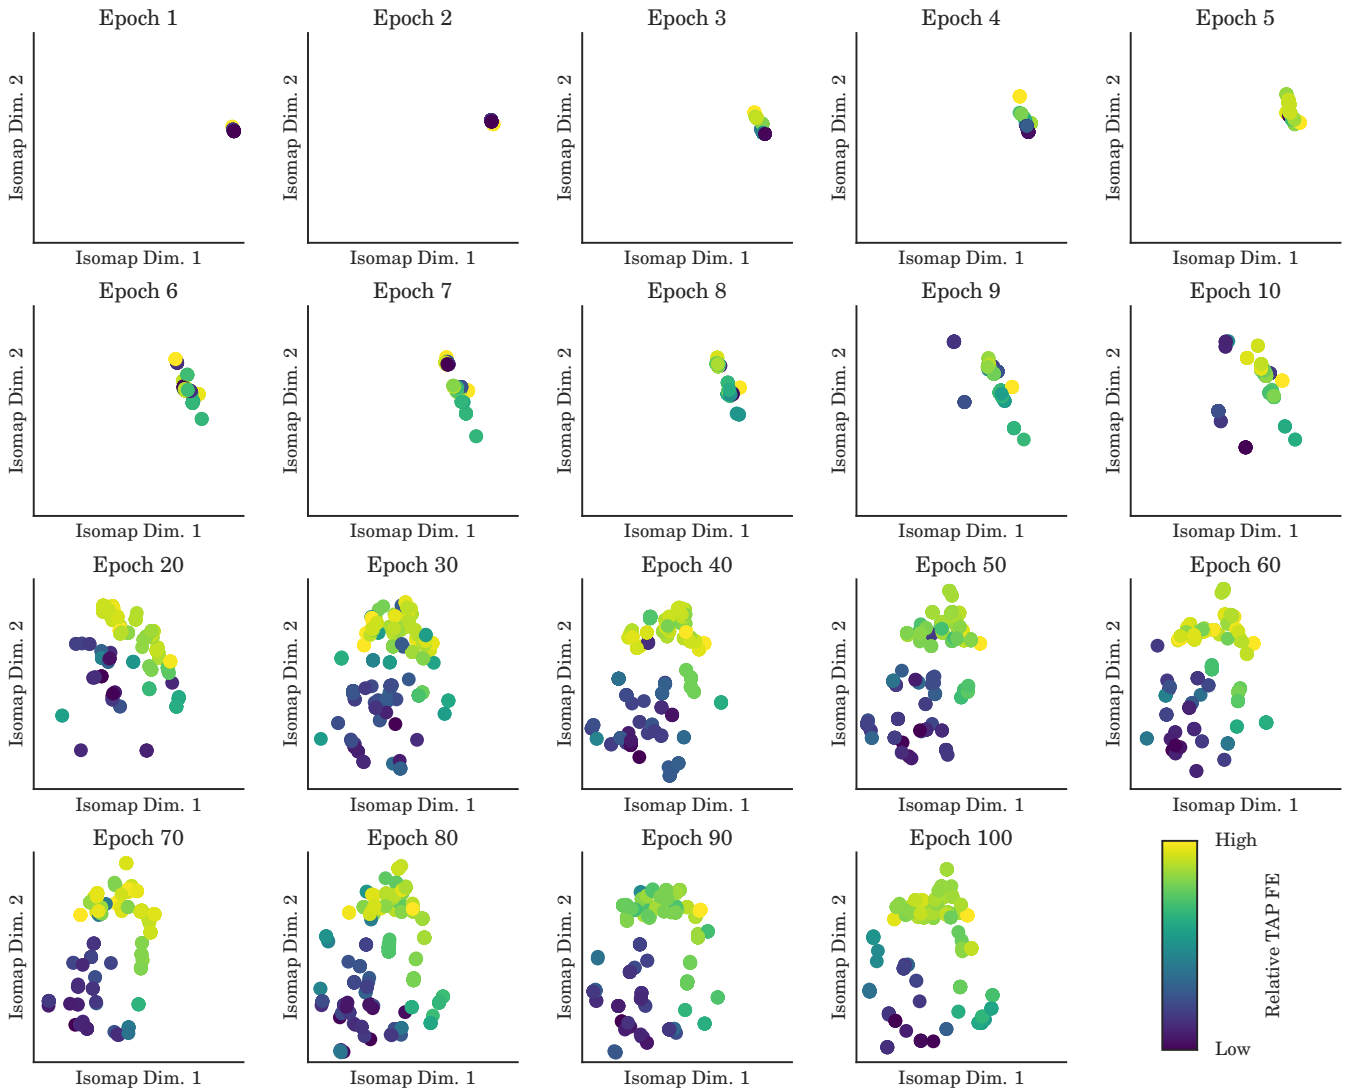


FIG. 6. Isomap visualization of TAP solutions for binary MNIST over training epochs for $N_h = 100$. All TAP solutions are mapped to the same two-dimensional embedding via an Isomap transform fitted to the TAP solutions of epoch 100. The embedding is performed on both the hidden and visible inferred expectations \mathbf{a}^v and \mathbf{a}^h . The color mapping corresponds to the TAP free-energy values of each solution, with the range of colors normalized between the minimum and maximum free energies of the solution at each training epoch.

accuracy, counting the number of incorrect estimates, or some other function of the binary confusion matrix, such as the F1 score or Matthews correlation coefficient (MCC).

For a fixed set of observations and channel parameters, if we assume that the original signal is drawn from some unknown and intractable generating distribution, then, as we construct more and more accurate tractable approximate priors, the more accurately we can construct an estimate of the original signal.

In other words, the more we know about the structure and content of the unknown signal *a priori*, the closer our estimate can be. Often, as in the case of wavelet-based image denoising, statistics are gathered on the transform coefficients of particular image classes, and heuristic denoising approaches are designed by hand accordingly [67]. By-hand derivation of denoising

algorithms works well in practice owing to its generality. Specific *a priori* information about the original signal is not required, beyond its signal class (e.g., natural images, human speech, radar return timings). However, meaningful features must be assumed or investigated by practitioners before successful inference can take place.

1. Denoising the binary symmetric channel

For binary denoising problems, we assume a binary symmetric channel (BSC) defined in the following manner. Given some binary signal $\mathbf{x} \in \{0, 1\}^N$, we observe the signal $\mathbf{y} \in \{0, 1\}^N$ as \mathbf{x} with independent bit flips occurring with probability p , which gives the following likelihood at each observation:

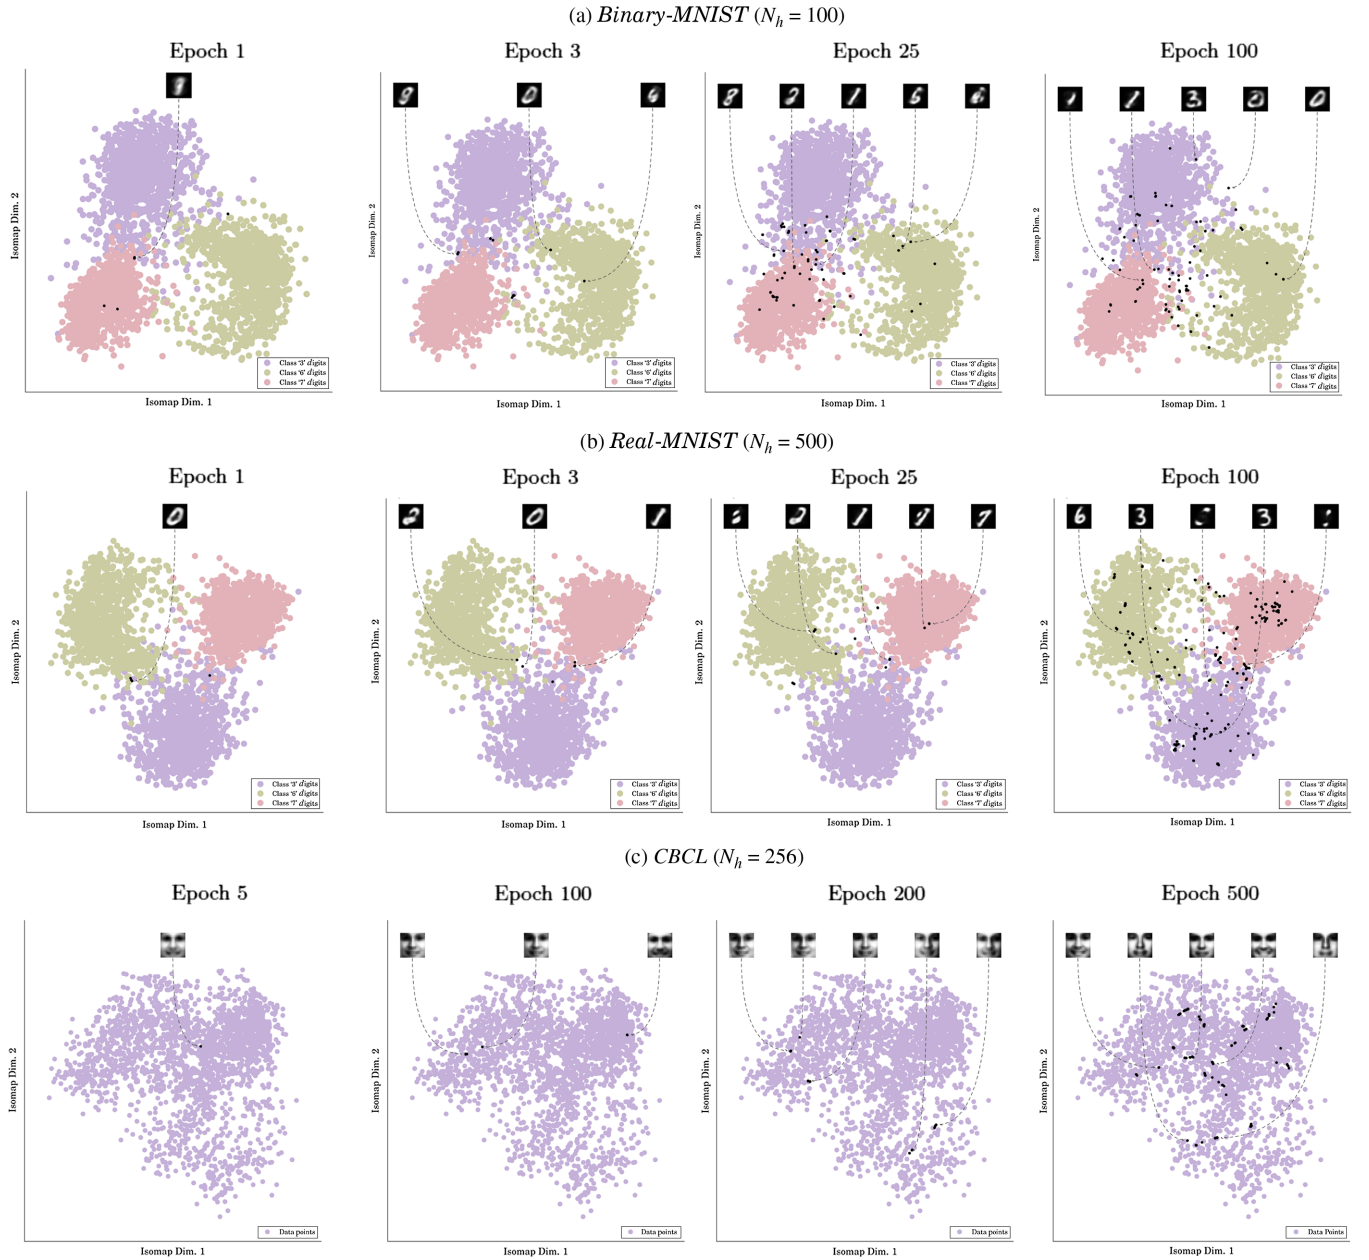


FIG. 7. Comparison of initial conditions for TAP equilibration (colored dots) compared to converged TAP solutions (black dots) for the tested data sets at different stages of training. For each data set, a two-dimensional Isomap embedding is calculated over the initialization data. Subsequently, the magnetizations of the TAP solutions are embedded in the same space. In each case, all initial variances are set to 0, as in Eq. (38). Also, a random selection of TAP solution magnetizations are chosen to provide some context for the representations the RBM is learning. (a) Binary MNIST: Here, the approximately 3100 digits corresponding to the classes “3,” “6,” and “7” are drawn from the first 10 000 training samples of the binarized MNIST data set as initializations. A reduced set of labels is used for readability. (b) Real MNIST: The same initializations are used as in (a); however, the initializations are not binarized. (c) CBCL: All available training face images are used as initializations.

$$P(\mathbf{y}|\mathbf{x}) = (1-p) \prod_i \left(\frac{p}{1-p} \right)^{\delta_{x_i, y_i}}, \quad (43)$$

$$P(\mathbf{y}|\mathbf{x}) = \frac{1}{Z(\mathbf{y})} e^{\sum_i D_i x_i}, \quad (44)$$

which can be shown to have the equivalent representation as a Boltzmann distribution:

where $D_i \triangleq \ln(p/1-p)(2y_i - 1)$. For a given prior distribution $P(\mathbf{x})$, the posterior distribution is given by Bayes' rule:

$$P(\mathbf{x}|\mathbf{y}) = \frac{e^{\sum_i D_i x_i} P(\mathbf{x})}{\sum_{\mathbf{x}} e^{\sum_i D_i x_i} P(\mathbf{x})}. \quad (45)$$

By assuming a factorized $P(\mathbf{x}) = \prod_i m_i^{x_i} (1 - m_i)^{1-x_i}$, where m_i might be per-site empirical averages obtained from available training data, the posterior factorizes, and we can construct the Bayes-optimal pointwise estimator (OPE) as the average $\langle x_i \rangle_{P(x_i|y_i)}$, which is just $P(x_i = 1|y_i)$ for our binary problem. Thus, the OPE at each site x_i is given as

$$\begin{aligned} P(x_i = 1|y_i) &= \frac{1}{1 + \left(\frac{1-m_i}{m_i}\right) \times \left(\frac{1-p}{p}\right)^{2y_i-1}} \\ &= \text{sigm}\left(\ln \frac{m}{1-m} + (2y_i - 1) \ln \frac{p}{1-p}\right). \end{aligned} \quad (46)$$

For a given data set, the OPE gives us the best-case performance using only pointwise statistics from the data set, namely, empirical estimates of the magnetizations m_i . We can see from Eq. (46) that the OPE returns either the observations, in the case of $p = 0$, or the prior magnetizations m_i , in the case $p = 0.5$. In this case of complete information loss, the worst-case performance is bounded according to the deviation of the data set from its mean. We present the performance of the OPE in Fig. 8 for the binary-MNIST data set, which makes the OPE a valuable baseline comparison and sanity check for the GRBM approximation of $P(\mathbf{x})$. As the GRBM model takes into account both pointwise and pairwise relationships in the data, a properly trained GRBM should provide estimates *at least as good as* the OPE.

The k -nearest-neighbor (k -NN) algorithm represents a different heuristic approach to the same problem [68]. In this case, the noisy measurements are compared to a set of exemplars: the training data set. Then, according to some distance metric such as MSE or correlation, one finds the k exemplars with a minimal distance to the noisy observations to serve as a basis for recovering the original binary signal. One can use some arbitrary approach for fusing these exemplars together into the final estimate, but the simplest case is a simple average. In the case that $k \rightarrow \infty$, the performance when using averaging is again bounded by the empirical magnetizations. In the other limit of $k = 1$, the estimate is simply the nearest exemplar. It is hard to show the limiting performance of this approach, as it is dependent on the distances, in the chosen metric, between the exemplars and the observations, as well as the interplay between the noise channel and the distance metric.

However, it can be seen directly that this approach is nonoptimal, as this approach will not yield the true signal at $p = 0$ unless the true signal is itself contained within the training data. We show the performance for $k = 1$ in Fig. 8. The advantage of this approach is that it successfully regularizes against noise as $p \rightarrow 0.5$, as the nearest




FIG. 8. Average denoising performance for reconstruction from bit-flip errors on binary MNIST over the probability of a bit to be flipped. Denoising via inference on the binary-binary RBM is denoised by the varying numbers of hidden units ($N_h = \{25, 50, 100, 500\}$). Also shown as baseline comparisons are the OPE given the empirical factorized magnetizations at each site (dashed black line) and a 1-NN matching from the training set (solid black line). All experiments are run over the same 1000 data samples drawn from the held-out test set and compared using the MCC. Binary estimates are obtained for the OPE and TAP inferred estimates by rounding the resulting magnetizations.

exemplar is always noise-free and at least marginally correlated with the original signal, up to the distance metric. Additionally, we see that it performs better than the OPE in the regime $p > 0.2$, which can be explained since we can think of the k -NN approach as implicitly, though indirectly, taking into account higher-order correlations in the data set by naively returning data exemplars; all the estimates trivially possess the same arbitrarily complex structure as the unknown signal.

Using the GRBM, we can hope to capture the best points of both approaches. First, we hope to perfectly estimate the original signal in the case $p = 0$. Second, we hope to leverage the pairwise correlations present in the data set, returning estimates which retain the structure of the data even as $p \rightarrow 0.5$. For GRBM denoising of the BSC, we no longer have a factorized posterior. Instead, we have the GRBM likelihood given in Eq. (9) summed over the hidden units. Using the definition of the binary prior given in Appendix C 3,

$$P(\mathbf{x}; W, U) \propto \sum_{\mathbf{h}} e^{\sum_{ij} x_i W_{ij} h_j + \sum_i U_i x_i + \sum_j U_j h_j}. \quad (47)$$

Since both the GRBM and the BSC channel likelihood are written as exponential family distributions, $P(\mathbf{y}|\mathbf{x}) \times P(\mathbf{x}; \mathbf{W}, \mathbf{U}) \propto \sum_{\mathbf{h}} e^{\sum_{ij} x_i W_{ij} h_j + \sum_i (D_i + U_i) x_i + \sum_j U_j h_j}$. Finding the averages $\langle x_i \rangle$ for this model simply consists in running the TAP-based inference of Algorithm 1 for the modified visible binary prior $\mathcal{B}(x_i; U_i + D_i)$. One heuristic caveat of this approach is that we must take into account the multimodal nature of the TAP free energy. Since we must initialize somewhere, and the resulting inference estimate is dependent upon this initialization, we initialize the inference with the OPE result.

We can see that, as $p \rightarrow 0$, $D_i \rightarrow \pm\infty$ and the highest probability configuration becomes observations. So, in the limit, we are able to obtain the true signal, just as the OPE, especially since we initialize within the well of this potential, which is shown for binary RBMs trained with varying numbers of hidden units in Fig. 8. In every case, for $p = 0$, the true signal is recovered. In the case of $N_h = \{25, 50, 100\}$, we see that the TAP inference on the binary RBM always outperforms the OPE. Additionally, we see that, in each of these cases, the performance closely mirrors that of the 1-NN as $p \rightarrow 0.5$. In the limit $p = 0.5$, we see that the result of the TAP inference is, essentially, uncorrelated with the original signal, as in this case there is no extra potential present to bias the inference, and the resulting estimate is simply an arbitrary solution of the TAP free energy. As this closely mirrors the exemplar selection in 1-NN, the MCC curves for the two approaches are similar.




FIG. 9. Subjective comparison of denoising estimates for a single-digit image for $p = [0, 0.5]$. For the OPE and RBM approaches, the inferred posterior averages $\langle x_i \rangle$ are shown, rather than the final configurations, where white and black represent the values 0 and 1, respectively. At each tested value of p , the same noise realization is used for each method. As p increases, for $N_h = \{25, 50, 100\}$, the TAP inference for the RBM provides estimates which still possess digit structure. In the case $N_h = 500$, the TAP inference gets caught in spurious and undesirable minima as p increases.

In the case of $N_h = 500$, we can see that an over-training effect occurs. Essentially, at low values of p , the TAP inference over the binary RBM is able to more accurately identify the original signal. However, at a certain point, owing to the increased number of solutions in the TAP free energy, there exist many undesirable minima around the noisy solutions, leading to poor denoising estimates. One can observe this subjectively in Fig. 9, where in the case of $N_h = 500$, the TAP inference results in either nearly zero modes or in very localized ones. This result would seem to indicate that the landscape of the TAP free energy around the initializations is becoming more unstable as the density of solutions increases around it. Additionally, since the TAP free-energy landscape is probed using only data points during training, the clustering of solutions around noisy samples remains ambiguous. Augmenting the initializations used when calculating the TAP solutions for the gradient estimate with noisy data samples could help alleviate this problem and regularize the TAP free-energy landscape in the space of noisy data samples.

VIII. DISCUSSION

In this paper, we propose a novel interpretation of the RBM within a fully tractable and deterministic framework of learning and inference via TAP approximation. This deterministic construction allows novel tools for scoring unsupervised models and investigation of the memory of trained models, as well as allowing their efficient use as structured joint priors for inverse problems. While deterministic methods based on NMF for RBM training are shown to be inferior to CD-k in Ref. [38], the level of approximation accuracy afforded by TAP finally makes the deterministic approach to RBMs effective, as shown in the case of binary RBMs in Ref. [40].

Additionally, our construction is generalized over the distribution of both the hidden and visible units. This generalization is unique to our work, as other works propose unique training methods and models when changing the distribution of the visible units. For example, this construction can be seen in the modified Hamiltonians used for real-valued data [69,70]. This construction allows us to consider binary, real-valued, and sparse real-valued data sets within the same framework. Additionally, one can also consider other architectures by changing the distributions imposed on the hidden unit. Here, we present experiments using only binary hidden units, but one could also use our proposed framework for Gaussian-distributed hidden units, thus mimicking a Hopfield network [2]. Or, also, sparse Gauss-Bernoulli distributed hidden units could mimic the same functionality as that proposed by the spike-and-slab RBM [45]. We leave these investigations to further works on this topic.

Our proposed framework also offers a possibility to explore the statistical mechanics of these latent variable models at the level of TAP approximation. Specifically, for a given statistical model of the weights W , both the cavity method and replica can begin to make predictions about these unsupervised models. An analytical understanding of the complexity of the free-energy landscape, and its transitions as a function of model hyperparameters, can allow for a richer understanding of statistically optimal network construction for learning tasks. In the case of random networks, there has already been some progress in this area, as shown in Refs. [71,72]. However, similar comprehensive studies conducted on learning in a realistic setting are still yet to be realized. Finally, as our framework can be applied to deep Boltzmann machines with minimal alteration, it can also potentially lead to a richer understanding of deep networks and the role of hierarchy in regularizing the learning problem in high dimensionality.

ACKNOWLEDGMENTS

This work has been supported by the ERC under the European Union's FP7 Grant Agreement 307087-SPARCS as well as by the French Agence Nationale de la Recherche under Grant No. ANR-17-CE23-0023-01 PAIL. Additional funding is acknowledged from "Chaire de recherche sur les modèles et sciences des données," Fondation CFM pour la Recherche-ENS.

APPENDIX A: BELIEF PROPAGATION FOR PAIRWISE MODELS

In order to estimate the derivatives of \mathbb{F} , we must first construct a BP algorithm on the factor-graph representation given in Fig. 1. We note that this graph, in terms of the variables \mathbf{x} , does not make an explicit distinction between the latent and visible variables. We instead treat this graph in full generality so as to clarify the derivation and notation. This graph corresponds to the following joint distribution over \mathbf{x} :

$$P(\mathbf{x}; W, \boldsymbol{\theta}) \propto e^{\sum_{(i,j)} \phi(x_i, x_j; W_{ij}) + \sum_i \phi(x_i; \theta_i)}, \quad (\text{A1})$$

where $\sum_{(i,j)}$ is a sum over the edges in the graph. In the case of a Boltzmann machine, any two variables are connected via pairwise factors,

$$\phi(x_i, x_j; W_{ij}) = e^{W_{ij} x_i x_j}, \quad (\text{A2})$$

and all variables are also influenced by univariate factors written trivially as $\phi(x_i; \theta_i) = P_i(x_i; \theta_i)$.

A message passing can be constructed on this factor graph by writing messages from variables to factors and also from factors to variables. Since all factors are at most degree 2, we can write the messages for this system as variable-to-variable messages [61]:

$$\nu_{i \rightarrow j}^{(t+1)}(x_i) = \frac{\phi(x_i; \theta_i)}{Z_{i \rightarrow j}} \prod_{l \in \partial_i / j} \int dx_l \phi(x_i, x_l; W_{il}) \nu_{l \rightarrow i}^{(t)}(x_l). \quad (\text{A3})$$

Here, the notation $i \rightarrow j$ represents a message *from* variable index i *to* variable index j , and ∂_i / j refers to all neighbors of variable index i *except* variable index j . We denote neighboring variables as those which share a pairwise factor. Finally, the superscripts on the messages refer to the time index of the BP iteration, which implies the successive application of Eq. (A3) until convergence on the set of messages $\underline{\nu} = \{\nu_{i \rightarrow j}(x_i) : (i, j) \in E\}$, where E is the set of all pairs of neighboring variables. We also note the inclusion of the message normalization term $Z_{i \rightarrow j}$, which ensures that all messages are valid PDFs. Additionally, it is possible to write the marginal beliefs at each variable by collecting the messages from all their neighbors:

$$\nu_{\rightarrow i}^{(t)}(x_i) = \frac{\phi(x_i; \theta_i)}{Z_{\rightarrow i}} \prod_{j \in \partial_i} \int dx_j \phi(x_i, x_j; W_{ij}) \nu_{j \rightarrow i}^{(t)}(x_j). \quad (\text{A4})$$

Subsequently, the Bethe free energy can be written for a converged set of messages, $\underline{\nu}^*$, according to Ref. [61] as

$$\begin{aligned} \mathbb{F}_B[\underline{\nu}^*] &= \sum_i \mathbb{F}_i[\underline{\nu}^*] - \sum_{(i,j) \in E} \mathbb{F}_{(i,j)}[\underline{\nu}^*] \\ &= \sum_i \ln Z_{\rightarrow i}^* - \sum_{(i,j)} \ln \left[\int dx_i dx_j \phi(x_i, x_j; W_{ij}) \nu_{j \rightarrow i}^*(x_j) \right. \\ &\quad \left. \times \phi(x_i, x_j; W_{ij}) \nu_{i \rightarrow j}^*(x_i) \right], \end{aligned} \quad (\text{A5})$$

where $Z_{\rightarrow i}^*$ refers to the normalization of the set of the marginal belief at site i derived from $\underline{\nu}^*$.

Unfortunately, the message passing of Eq. (A3) cannot be written as a computable algorithm due to the continuous nature of the PDFs. Instead, we must find some manner by which to parameterize the messages. In the case of binary variables, as in Ref. [40], each message PDF can be exactly parameterized by its expectation. However, for this general case formulation, we cannot make the same assumption. Instead, we turn to *relaxed* BP (RBP) [51], described in the next section, which assumes a two-moment parameterization of the messages.

APPENDIX B: RBP FOR PAIRWISE MODELS

We now consider one possible parametric approximation of the message set via RBP [51]. This approach has also gone by a number of different names in parallel rediscoveries of the approach, e.g., moment matching [50] and nonparametric BP [49]. In essence, we are assuming that all messages $\underline{\nu}$ can be well approximated by their mean and variance, a Gaussian assumption. This approximation arises from a second-order

expansion assuming small weights W_{ij} . By making this assumption, we ultimately are able to close an approximation of the messages on their two first moments, $a_{i \rightarrow j} \triangleq \langle x_i \rangle_{\nu_{i \rightarrow j}}$ and $c_{i \rightarrow j} \triangleq \langle x_i^2 \rangle_{\nu_{i \rightarrow j}} - \langle x_i \rangle_{\nu_{i \rightarrow j}}^2$.

1. Derivation via small-weight expansion

Considering the marginalization taking place in Eq. (A3), we perform a second-order expansion assuming that $W_{il} \rightarrow 0$. We start by taking the Taylor series of the incoming message marginal for negligible weights:

$$\begin{aligned} & \int dx_l e^{W_{il} x_i x_l} \cdot \nu_{l \rightarrow i}^{(t)}(x_l) \\ &= 1 + W_{il} \int dx_l \frac{\partial}{\partial W_{il}} [e^{W_{il} x_i x_l}]_{W_{il}=0} \nu_{l \rightarrow i}^{(t)}(x_l) \\ & \quad + \frac{1}{2} W_{il}^2 \int dx_l \frac{\partial^2}{\partial W_{il}^2} [e^{W_{il} x_i x_l}]_{W_{il}=0} \nu_{l \rightarrow i}^{(t)}(x_l) \\ & \quad + O(W_{il}^3). \end{aligned} \quad (\text{B1})$$

Now, we approximate the series by dropping the terms less than $O(W_{il}^3)$. This approximation can be justified in the event that all weight values satisfy $|W_{il}| < 1$. Identifying the integrals from the expansion as moments, we see the following approximation:

$$\approx 1 + x_i W_{il} a_{l \rightarrow i}^{(t)} + \frac{1}{2} x_i^2 W_{il}^2 \langle x_l^2 \rangle_{\nu_{l \rightarrow i}^{(t)}}. \quad (\text{B2})$$

However, we would like to write this approximation in terms of the central second moment. Through a second approximation that neglects $O(W_{il}^3)$ terms, we arrive at our desired parameterization of the incoming message marginalization in terms of the message's two first *central* moments:

$$\begin{aligned} &= \exp \left\{ \log \left[1 + x_i W_{il} a_{l \rightarrow i}^{(t)} + \frac{1}{2} x_i^2 W_{il}^2 \langle x_l^2 \rangle_{\nu_{l \rightarrow i}^{(t)}} \right] \right\} \\ &\approx e^{x_i W_{il} a_{l \rightarrow i}^{(t)} + (1/2) x_i^2 W_{il}^2 c_{l \rightarrow i}^{(t)} + O(W_{il}^3)} \\ &\approx e^{x_i W_{il} a_{l \rightarrow i}^{(t)} + (1/2) x_i^2 W_{il}^2 c_{l \rightarrow i}^{(t)}}. \end{aligned} \quad (\text{B3})$$

We now substitute this approximation back into Eq. (A3) to get

$$\nu_{i \rightarrow j}^{(t+1)}(x_i) \approx \tilde{\nu}_{i \rightarrow j}^{(t+1)}(x_i) = \frac{\phi(x_i; \theta_i)}{Z_{i \rightarrow j}} e^{x_i B_{i \rightarrow j}^{(t)} - (1/2) x_i^2 A_{i \rightarrow j}^{(t)}}, \quad (\text{B4})$$

where

$$B_{i \rightarrow j}^{(t)} \triangleq \sum_{l \in \partial_i / j} W_{il} a_{l \rightarrow i}^{(t)}, \quad (\text{B5})$$

$$A_{i \rightarrow j}^{(t)} \triangleq - \sum_{l \in \partial_i / j} W_{il}^2 c_{l \rightarrow i}^{(t)}. \quad (\text{B6})$$

From here, we can see that we have now a set of closed equations due to the dependence of $\mathbf{A}^{(t)}$ and $\mathbf{B}^{(t)}$ on the moments $\mathbf{a}^{(t)}$ and $\mathbf{c}^{(t)}$, and vice versa. The values of these moments can be written as a function of $\mathbf{A}^{(t)}$ and $\mathbf{B}^{(t)}$ which is dependent upon the form of the local potentials $\phi(x_i; \theta_i)$, i.e., the prior distribution we assign to the variables themselves,

$$a_{i \rightarrow j}^{(t)} = f_a(B_{i \rightarrow j}^{(t-1)}, A_{i \rightarrow j}^{(t-1)}; \theta_i), \quad (\text{B7})$$

$$c_{i \rightarrow j}^{(t)} = f_c(B_{i \rightarrow j}^{(t-1)}, A_{i \rightarrow j}^{(t-1)}; \theta_i), \quad (\text{B8})$$

where

$$f_a(B, A; \theta) \triangleq \int dx x \frac{\phi(x; \theta)}{Z} e^{xB - (1/2)x^2 A}, \quad (\text{B9})$$

$$\begin{aligned} f_c(B, A; \theta) &\triangleq -f_a(B, A; \theta)^2 \\ &\quad + \int dx x^2 \frac{\phi(x; \theta)}{Z} e^{xB - (1/2)x^2 A}, \end{aligned} \quad (\text{B10})$$

and Z is simply the normalization $\int dx \phi(x; \theta) e^{xB - (1/2)x^2 A}$. The inferred marginal distributions at each site can be calculated via the same functions but instead using all of the incoming messages, i.e., $a_i^{(t)} = f_a(B_{\rightarrow i}^{(t)}, A_{\rightarrow i}^{(t)}; \theta_i)$ and $c_i^{(t)} = f_c(B_{\rightarrow i}^{(t)}, A_{\rightarrow i}^{(t)}; \theta_i)$. In Appendix C 1, we give the closed forms of these moment calculations for a few different choices of $\phi(x; \theta)$.

If one wants to obtain an estimate of the free energy for a given set of parameters θ , it is possible to iterate between Eqs. (B6) and (B5) and Eqs. (B7) (B8) until, ideally, convergence. It is important to note, however, due to both the potentially loopy nature of the network as well as small-weight expansion, that the BP iteration is not guaranteed to converge [41,61]. Additionally, while we retain the time indices in our derivation, it is not clear whether one should attempt to iterate these message in a fully sequential or parallel fashion or if some clustering and partitioning of the variables should be applied to determine the update order dynamically.

2. RBP approximate Bethe free energy

Additionally, we can write the specific form of the Bethe free energy under the RBP two-moment parameterization of the messages. In this case, we can simply apply the small-weight expansion of Eq. (B3) to the Bethe free energy for pairwise models given in Eq. (A5):

$$\begin{aligned} \tilde{\mathbb{F}}_B[\tilde{\mathcal{Z}}^*] &= \sum_i \ln \tilde{Z}_{\rightarrow i}^* \\ &\quad - \sum_{(i,j)} \ln \left[\int dx_i \tilde{\nu}_{i \rightarrow j}^*(x_i) e^{x_i W_{ij} a_{j \rightarrow i}^* + (1/2) x_i^2 W_{ij}^2 c_{j \rightarrow i}^*} \right]. \end{aligned} \quad (\text{B11})$$

Subsequently, using the final message definition given in Eq. (B4), we can see that

$$\tilde{\mathbb{F}}_B[\tilde{\mathcal{Z}}^*] = \sum_i \ln \tilde{Z}_{i \rightarrow}^* - \sum_{(i,j)} \{\ln \tilde{Z}_{i \rightarrow}^* - \ln \tilde{Z}_{i \rightarrow j}^*\}, \quad (\text{B12})$$

which, correcting for double counting, can also be written as

$$\tilde{\mathbb{F}}_B[\tilde{\mathcal{Z}}^*] = \sum_i \left(1 - \frac{1}{2}d_i\right) \ln \tilde{Z}_{i \rightarrow}^* + \frac{1}{2} \sum_{i,j} \ln \tilde{Z}_{i \rightarrow j}^*, \quad (\text{B13})$$

where d_i is the degree at site i , $|\partial_i|$.

3. Enforcing bounded messages

While we write the RBP messages (B4) as though they are Gaussian distributions, this is a slight, since $A_{i \rightarrow j}^{(t)} \leq 0$ as $W_{il}^2 c_{l \rightarrow i}^{(t)} \geq 0 \forall i, l$. The implication of the expansion is that, in general, the messages are in fact unbounded. This unboundedness is a direct result of the form of the conventional RBM pairwise factor $e^{x_i W_{ij} x_j}$.

There are a few avenues available to us to address these unbounded messages and produce a meaningful message passing for generalized RBMs. Let us consider the cases for which the messages are unbounded *given a specific variable distribution*. Assume that site x_i is assigned a Gaussian prior, $\phi(x_i; \theta_i = \{V_i, U_i\}) \propto e^{x_i U_i - (1/2)x_i^2 V_i}$. In this case, the RBP message reads

$$\tilde{v}_{i \rightarrow j}^{(t+1)}(x_i) = \frac{1}{\tilde{Z}_{i \rightarrow j}} e^{x_i (U_i + B_{i \rightarrow j}^{(t)}) - (1/2)x_i^2 (V_i + A_{i \rightarrow j}^{(t)})}. \quad (\text{B14})$$

In this case, the message is unbounded in the event that the weighted sum of all incoming neighbor variances at i exceeds the inverse variance of Gaussian prior on x_i ,

$$V_i + A_{i \rightarrow j}^{(t)} < 0, \quad (\text{B15})$$

$$V_i < \sum_{l \in \partial_i / j} W_{il}^2 c_{l \rightarrow i}^{(t)}, \quad (\text{B16})$$

$$\sigma_i^2 > \left[\sum_{l \in \partial_i / j} W_{il}^2 c_{l \rightarrow i}^{(t)} \right]^{-1}, \quad (\text{B17})$$

where σ_i^2 is the variance of the Gaussian prior. Said another way, this condition is telling us that, when the message passing starts to tell us that if the variance at x_i is smaller than that of its prior, the messages become unbounded and fail to be meaningful probability distributions, and our expansion fails. The implication is that the RBP message passing should be utilized in contexts where there exists some, preferably strong, evidence at each site, or the weights in W should be sufficiently small. The stronger this local potential, or the smaller the weights, the more favorable the model is to the

RBP inference. This observation mirrors those made in Ref. [41]; however, here the authors make the observation that, in this setting, BP based on a small-weight expansion for *binary* variables fails to converge. In our case, without taking some form of regularization, the inference fails entirely. Thus, large magnitude couplings W_{il} must be backed with a high degree of evidence at site i and l . This property could be utilized for the inverse learning problem, where one must learn the couplings W given a data set, in order to constrain the learning to parameters which are amenable to the RBP inference.

One direct manner to create probability distributions from otherwise unbounded continuous functions is via truncation. Specifically, we enforce a noninfinite normalization factor by restricting the support of the distribution to some subset of \mathbb{R} . In this case, just slightly violating the bounded condition above induces a uniform message distribution over the distribution support, while a strong violation causes the distribution to concentrate on the boundaries of the support. Another approach might simply be to fix a hard boundary constraint on $A_{i \rightarrow j}$, thus never permitting unbounded messages to occur.

APPENDIX C: CALCULATIONS FOR SPECIFIC VARIABLE DISTRIBUTIONS

1. Truncated Gaussian units

In general, the truncated Gaussian is defined in the following manner:

$$\mathcal{TG}(x; \mu, \sigma^2, [\alpha, \omega]) = \frac{1}{\sqrt{2\pi\sigma^2}} \cdot \frac{1}{\Phi\left[\frac{\omega-\mu}{\sigma}\right] - \Phi\left[\frac{\alpha-\mu}{\sigma}\right]} \cdot e^{-(x-\mu)^2/2\sigma^2}, \quad (\text{C1})$$

where μ and σ^2 are the mean and variance, respectively, of the original Gaussian prior to truncation, the range $[\alpha, \omega]$ defines the lower and upper bounds of the truncation, $-\infty \leq \alpha < \omega \leq \infty$, and $\Phi[\cdot]$ is the CDF for the normal distribution. To make things easier for us later, we define the prior in a little bit of a different manner by making the following definitions:

$$V \triangleq \frac{1}{\sigma^2}, \quad U \triangleq \frac{\mu}{\sigma^2}, \quad (\text{C2})$$

and writing the distribution for the parameters $\theta = \{U, V, [\alpha, \omega]\}$ as

$$\mathcal{TG}_+(x; \theta) = \frac{\frac{2}{e^{U^2/2V}} \sqrt{\frac{V}{2\pi}} e^{-(1/2)Vx^2 + Ux}}{\text{Erf}\left[\sqrt{\frac{V}{2}}\left(\omega - \frac{U}{V}\right)\right] - \text{Erf}\left[\sqrt{\frac{V}{2}}\left(\alpha - \frac{U}{V}\right)\right]}, \quad (\text{C3})$$

where $\text{Erf}[\cdot]$ is the error function and the last step follows from the identity $\Phi[x] = \frac{1}{2} + \frac{1}{2}\text{Erf}(x/\sqrt{2})$. Here, the subscript $+$ is used to indicate $V > 0$.

In the case that $V < 0$, we write \mathcal{TG} in terms of the imaginary error function, $\text{Erfi}[x] \triangleq -i\text{Erf}[ix]$, and, noting that for $V < 0$, $V = -1 \cdot |V|$,

$$\mathcal{TG}_-(x; \theta) = \frac{\frac{2}{e^{U^2/2V}} \sqrt{\frac{|V|}{2\pi}} e^{-(1/2)Vx^2+Ux}}{\text{Erfi}\left[\sqrt{\frac{|V|}{2}}(\omega - \frac{U}{V})\right] - \text{Erfi}\left[\sqrt{\frac{|V|}{2}}(\alpha - \frac{U}{V})\right]}. \quad (\text{C4})$$

We use this negative variance version of the truncated Gaussian to handle the special case of $A + V < 0$ first detailed in Appendix B 3.

We now detail the computation of the partition and first two moments of the Gaussian-product distribution of \mathcal{TG} , $Q(x; \theta, A, B) = (1/Z)\mathcal{TG}(x; \theta)e^{-(1/2)Ax^2+Bx}$. The calculation of the moments as a function of A and B provides the definitions of f_a and f_c , while the calculation of the normalization Z provides terms necessary for both the computation of the TAP free energy as well as the gradients necessary for learning θ during RBM training.

First, we calculate the normalization of $Q(x; \theta, A, B)$ in terms of all free parameters. To do this, consider the truncated normalization of the following product of Gaussians:

$$Z_Q = \frac{1}{Z_{\mathcal{TG}}} \int_{\alpha}^{\omega} dx e^{-(1/2)(A+V)x^2+(B+U)x}, \quad (\text{C5})$$

where $Z_{\mathcal{TG}}$ is defined such that $\mathcal{TG}(x; \theta) = (1/Z_{\mathcal{TG}})e^{-(1/2)Vx^2+Ux}$. We need to make note of the special case of $A + V < 0$; thus,

$$Z_Q = \frac{1}{Z_{\mathcal{TG}}} \sqrt{\frac{\pi}{2|A+V|}} e^{(B+U)^2/2(A+V)} \times \begin{cases} d_+, & A + V > 0, \\ d_-, & A + V < 0, \end{cases} \quad (\text{C6})$$

where $d_+ \triangleq \text{Erf}[h_{\omega}] - \text{Erf}[h_{\alpha}]$, $d_- \triangleq \text{Erfi}[h_{\omega}] - \text{Erfi}[h_{\alpha}]$, and $h_x \triangleq \sqrt{(|A+V|/2)}[x - (B+U/A+V)]$.

Since $Q(x; \theta, A, B)$ is simply a truncated Gaussian with updated parameters, the first moment is given according to the well-known truncated Gaussian expectation. While this expectation is usually written in terms of a mean and variance of the *untruncated* Gaussian distribution and for the case of positive mean, we instead write the expectation in terms of the exponential polynomial coefficients and note the special case $A + V < 0$:

$$f_a(B, A; \theta) = \frac{B+U}{A+V} + \sqrt{\frac{2}{\pi|A+V|}} \times \begin{cases} \frac{e^{-h_{\alpha}^2} - e^{-h_{\omega}^2}}{d_+}, & A + V > 0, \\ \frac{e^{h_{\alpha}^2} - e^{h_{\omega}^2}}{d_-}, & A + V < 0. \end{cases} \quad (\text{C7})$$

Next, we write the variance of $Q(x; \theta, A, B)$ as a function of A and B . As earlier, since Q has the specific form of a truncated Gaussian distribution, we can utilize the well-known variance formula for such a distribution. As in the case of f_a , we modify this function for the special case of $A + V < 0$. Specifically,

$$f_c(B, A; \theta) = \frac{1}{A+V} - \left(f_a(B, A; \theta) - \frac{B+U}{A+V} \right)^2 + \frac{2}{\sqrt{\pi}(A+V)} \times \begin{cases} \frac{h_{\alpha}e^{-h_{\alpha}^2} - h_{\omega}e^{-h_{\omega}^2}}{d_+}, & A + V > 0, \\ \frac{h_{\alpha}e^{h_{\alpha}^2} - h_{\omega}e^{h_{\omega}^2}}{d_-}, & A + V < 0. \end{cases} \quad (\text{C8})$$

a. Gradients of the log-likelihood

To determine the gradients of the log-likelihood with respect to the model parameters, it is necessary to calculate the gradients of both $\ln \mathcal{TG}_+(x; \theta)$ and $\ln Z_Q$ in terms of the distribution parameters U and V . We assume that the boundary terms α and ω remain fixed. Since both of these distributions are truncated Gaussians, we can treat them both in terms of the derivatives of the log normalization of a general-case truncated Gaussian:

$$\ln Z_{\mathcal{TG}} = -\frac{U^2}{2V} + \frac{1}{2} \ln |V| - \ln d_{\pm}. \quad (\text{C9})$$

For Boltzmann measures of the given quadratic form, we know that $(\partial/\partial U) \ln Z_{\mathcal{TG}} = \langle x \rangle_{\mathcal{TG}}$ and $(\partial/\partial V) \ln Z_{\mathcal{TG}} = -\frac{1}{2} \langle x^2 \rangle_{\mathcal{TG}}$. From these relations, we can write the necessary derivatives. For the case of $\ln \mathcal{TG}_+(x; \theta)$, we have

$$\frac{\partial}{\partial U} \ln \mathcal{TG}_+(x; \theta) = x - \frac{\partial}{\partial U} \ln Z_{\mathcal{TG}} = x - \langle x \rangle_{\mathcal{TG}} \quad (\text{C10})$$

and

$$\begin{aligned} \frac{\partial}{\partial V} \ln \mathcal{TG}_+(x; \theta) &= -\frac{1}{2} x^2 - \frac{\partial}{\partial V} \ln Z_{\mathcal{TG}} \\ &= -\frac{1}{2} (x^2 - \langle x^2 \rangle_{\mathcal{TG}}), \end{aligned} \quad (\text{C11})$$

which we can see as just the difference between the data and the moments of the truncated Gaussian distribution.

Next, the derivatives of Z_Q , for fixed A and B , can be written in terms of f_a and f_c , the moments of Q :

$$\frac{\partial}{\partial U} \ln Z_Q = f_a(B, A; \theta) - \langle x \rangle_{\mathcal{TG}}, \quad (\text{C12})$$

$$\frac{\partial}{\partial V} \ln Z_Q = -\frac{1}{2} [f_c(B, A; \theta) + f_a(B, A; \theta)^2 - \langle x^2 \rangle_{\mathcal{TG}}]. \quad (\text{C13})$$

Finally, we are ready to write the gradients of the log-likelihood to be used for both hidden and visible updates with the truncated Gaussian distribution. For a given set of minibatch data indexed by $m \in \{1, \dots, M\}$ and a number of TAP solutions indexed by $p \in \{1, \dots, P\}$, the gradients of a visible variable are given by

$$\begin{aligned} \Delta U_i &= \frac{1}{M} \sum_m (x_i^{(m)} - \langle x_i \rangle_{\mathcal{TG}}) \\ &\quad - \frac{1}{P} \sum_p (a_i^{(p)} - \langle x_i \rangle_{\mathcal{TG}}) \\ &= \langle x_i \rangle_M - \langle a_i \rangle_P \end{aligned} \quad (\text{C14})$$

and

$$\Delta V_i = -\frac{1}{2} [\langle x_i^2 \rangle_M - \langle c_i + a_i^2 \rangle_P], \quad (\text{C15})$$

where $\langle \cdot \rangle_M$ and $\langle \cdot \rangle_P$ are averages over the minibatch and TAP solutions, respectively.

For updates of hidden side variables using the truncated Gaussian distribution, we have the following gradients for updating their parameters:

$$\Delta U_j = \langle \tilde{a}_j \rangle_M - \langle a_j \rangle_P \quad (\text{C16})$$

and

$$\Delta V_j = -\frac{1}{2} [\langle \tilde{c}_j + \tilde{a}_j^2 \rangle_M - \langle c_j + a_j^2 \rangle_P], \quad (\text{C17})$$

where the moments $\tilde{a}_j^{(m)} \triangleq f_a(\sum_i W_{ij} x_i^{(m)}, 0; \theta_j)$ and $\tilde{c}_j^{(m)} \triangleq f_c(\sum_i W_{ij} x_i^{(m)}, 0; \theta_j)$ are defined for convenience. Using these gradients, we can now update the local biasing distributions for truncated Gaussian variables.

b. Numerical considerations

Using the truncated Gaussian prior comes with a few numerical issues which must be carefully considered. While we have already addressed the cases when $A + V < 0$, we have not addressed the case where

$$\left| \frac{B + U}{A + V} \right| \gg |\omega|, |\alpha|. \quad (\text{C18})$$

One can see how this case complicates matters by observing the term $(1/d_{\pm})$, which occurs in both the first and second moment computations. When the magnitude of the limits ω and α become vanishingly small in comparison to the scaled joint mean term, then we see that $d_{\pm} \rightarrow 0$. However, the numerators of both f_a and f_c also go to zero, which implies that we may be able to find some method of approximation to find estimates of these moments without, up to numerical precision, dividing by zero.

In order to handle this eventuality in our implementation, we make a Taylor series expansion of the $(1/d_{\pm})$ term in

the following way. First, we note that d_{\pm} has the form $\text{Erf}[\eta(\omega - \mu)] - \text{Erf}[\eta(\alpha - \mu)]$, which can be rewritten as $\text{Erf}[z] - \text{Erf}[z - \epsilon]$ for $z \triangleq \eta(\omega - \mu)$, and $\epsilon \triangleq \eta(\alpha - \omega)$ is a multiple of the difference between the two boundaries of the truncated Gaussian distribution. Since we wish to consider the case that $\mu \rightarrow \infty$, we take the Taylor expansion centered at $z = \infty$, since the value of μ dominates. From here, we find that the following approximation works well in practice:

$$\begin{aligned} \frac{e^{h_a^2} - e^{h_{\omega}^2}}{d_+} &\approx \frac{\mathcal{S}^{(n)}(h_{\omega}^2)}{(\alpha - \frac{B+U}{A+V})(1 - e^{h_a^2 - h_{\omega}^2})} \\ &\quad + \frac{\mathcal{S}^{(n)}(h_a^2)}{(\omega - \frac{B+U}{A+V})(1 - e^{h_{\omega}^2 - h_a^2})}, \end{aligned} \quad (\text{C19})$$

where $\mathcal{S}^{(n)}(\cdot)$ is the n -term power-series representation of the error function. In our experiments, we use $n = 11$. A similar approximation can be used for the variances in the same situation. While this approximation could potentially be computationally costly for large n , we note that it is used only for updates on variables for which a small value of d_{\pm} has been detected.

2. Truncated Gauss-Bernoulli units

For truncated Gauss-Bernoulli distributed units, we form the distribution as a mixture between a delta function and \mathcal{TG} , with an extra term $(1 - \rho)$ which controls the density at $x = 0$; thus, for $\theta = \{\rho, U, V, [\alpha, \omega]\}$ we have

$$\mathcal{TGB}(x; \theta) \triangleq (1 - \rho)\delta(x) + \rho\mathcal{TG}(x; \theta) \quad (\text{C20})$$

$$= \frac{1}{Z_{\mathcal{TGB}}} [(1 - \rho)\delta(x) + \rho e^{-(1/2)Vx^2 + Ux}] \quad (\text{C21})$$

where $Z_{\mathcal{TGB}} \triangleq (1 - \rho) + \rho Z_{\mathcal{TG}}$ and $\delta(x)$ is the Dirac delta function such that $\delta(0) = 1$ and 0 everywhere else. By using a construction such that the truncation is done on the Gaussian mode alone and not across the entire distribution, we can easily write the necessary functions of this distribution in terms of the values we have already calculated in Appendix C 1 as long as $0 \in [\alpha, \omega]$. Additionally, it is useful for our calculations to define the probability of x to be nonzero according to \mathcal{TGB} :

$$P[x \neq 0] = \frac{\rho Z_{\mathcal{TG}}}{Z_{\mathcal{TGB}}}. \quad (\text{C22})$$

We now continue as in the previous Appendices and write the normalization and first two moments of the Gaussian product distribution $Q(x; \theta, A, B) = (1/Z_Q)\mathcal{TGB}(x; \theta)e^{-(1/2)Ax^2 + Bx}$. First, the normalization can be written simply as a function of the truncated Gaussian normalization modified by A and B , as in Eq. (C6):

$$Z_Q = (1 - \rho) + \rho Z_{Q, \mathcal{TG}}. \quad (\text{C23})$$

Next, the first moment of Q can be found by recalling the relation $(\partial/\partial U) \ln Z_Q = \langle x \rangle_Q$ and, consequently, that $(\partial/\partial U) Z_Q = Z_Q \langle x \rangle_Q$. Thus,

$$f_a(B, A; \theta) = f_a^{\mathcal{TG}}(B, A; \theta) \cdot P[x \neq 0], \quad (\text{C24})$$

where the nonzero probability is calculated according to A and B . For the second moment, we note that $(\partial/\partial V) Z_Q = -\frac{1}{2} Z_Q \langle x^2 \rangle_{\mathcal{TG}}$ to find

$$\begin{aligned} f_c(B, A; \theta) &= P[x \neq 0] \cdot [f_a^{\mathcal{TG}}(B, A; \theta)^2 + f_c^{\mathcal{TG}}(B, A; \theta)] \\ &\quad - f_a(B, A; \theta)^2 \\ &= (P[x \neq 0] - P[x \neq 0]^2) f_a^{\mathcal{TG}}(B, A; \theta) \\ &\quad + P[x \neq 0] f_c^{\mathcal{TG}}(B, A; \theta). \end{aligned} \quad (\text{C25})$$

Next, we turn our attention to the log-likelihood gradients necessary for updating the parameters U , V , and ρ during training. First, we look at the derivatives of $\ln \mathcal{TGB}$ required for updates on visible units. In order to calculate these derivatives, we split the log probability into two cases:

$$\ln \mathcal{TGB}(x; \theta) = \begin{cases} -\ln Z_{\mathcal{TGB}}, & x = 0, \\ \ln \rho - \frac{1}{2} V x^2 + Ux - \ln Z_{\mathcal{TGB}}, & x \neq 0. \end{cases} \quad (\text{C26})$$

Consequently, the derivatives of the log probability are written as the following:

$$\frac{\partial}{\partial U} [\ln \mathcal{TGB}(x; \theta)] = x - \langle x \rangle_{\mathcal{TGB}} \quad (\text{C27})$$

and

$$\frac{\partial}{\partial V} [\ln \mathcal{TGB}(x; \theta)] = -\frac{1}{2} (x^2 - \langle x^2 \rangle_{\mathcal{TGB}}). \quad (\text{C28})$$

The derivative with respect to ρ is a bit more complicated, as we cannot use the same identities. Additionally, we must also consider the two cases of $x = 0$ and $x \neq 0$ separately. Thus,

$$\begin{aligned} \frac{\partial}{\partial \rho} [\ln \mathcal{TGB}(x; \theta)] &= \begin{cases} \rho Z_{\mathcal{TG}} - \frac{1}{\rho} P[x \neq 0], & x = 0, \\ \rho Z_{\mathcal{TG}} - \frac{1}{\rho} (P[x \neq 0] - 1), & x \neq 0 \end{cases} \\ &= \rho Z_{\mathcal{TG}} - \frac{1}{\rho} \{1 - \delta(x) - P[x \neq 0]\}, \end{aligned} \quad (\text{C29})$$

which can be rewritten in the more concise form

$$\frac{\partial}{\partial \rho} [\ln \mathcal{TGB}(x; \theta)] = \frac{\delta(x) - \rho}{\rho(1 - \rho)} \quad (\text{C30})$$

by noting the complement of the support probability $P[x=0] = 1 - P[x \neq 0] = (1 - \rho/Z_{\mathcal{TGB}}) = 1 - \rho$ and making the appropriate substitution.

We now write the derivatives of the $\ln Z_Q$ in terms of U , V , and ρ . These take the same form as those written in Appendix C 1. Subsequently, the equations for ΔU_i , ΔV_i , ΔU_j , and ΔV_j all remain consistent, just under the modification of all moments being taken with respect to the \mathcal{TGB} . We need only to write the gradients $\Delta \rho_i$ and $\Delta \rho_j$. Starting from the log partition and applying the same identity used to write Eq. (C30),

$$\frac{\partial}{\partial \rho} [\ln Z_Q] = \frac{P_Q[x \neq 0] - \rho}{\rho(1 - \rho)}, \quad (\text{C31})$$

which gives us the final gradients

$$\Delta \rho_i = \frac{1}{\rho_i(1 - \rho_i)} [\langle \delta(x_i) \rangle_M - \langle P_Q[x_i \neq 0] \rangle_P] \quad (\text{C32})$$

for visible units and

$$\Delta \rho_j = \frac{1}{\rho_j(1 - \rho_j)} [\langle \tilde{P}_Q[x_j \neq 0] \rangle_M - \langle P_Q[x_j \neq 0] \rangle_P] \quad (\text{C33})$$

for hidden units, where \tilde{P}_Q is calculated in the naive-mean-field manner described in Appendix C 1, where $\tilde{A}_j^{(m)} = 0$ and $\tilde{B}_j^{(m)} = \sum_i W_{ij} x_i^{(m)}$.

3. Binary units

We define the distribution for binary units to be the Bernoulli distribution such that $x \in \{0, 1\}$:

$$\mathcal{B}(x; m) = (1 - m)^{1-x} m^x, \quad (\text{C34})$$

where $m \triangleq \text{Prob}[x = 1] = \langle x \rangle_{\mathcal{B}}$. We can also write \mathcal{B} as the Boltzmann distribution

$$\mathcal{B}(x; U) = \frac{1}{Z_{\mathcal{B}}} e^{Ux}, \quad (\text{C35})$$

where $U \triangleq \ln(m/1 - m)$ and $Z_{\mathcal{B}} \triangleq 1 + e^U$. Next, we calculate the normalization and moments of the distribution $Q(x; A, B, U) = (1/Z_Q Z_{\mathcal{B}}) e^{-(1/2)Ax^2 + (B+U)x}$. For the normalization, we have

$$\begin{aligned} Z_Q &= \frac{1}{Z_{\mathcal{B}}} \sum_{x \in \{0,1\}} e^{-(1/2)Ax^2 + (B+U)x} \\ &= \frac{1}{Z_{\mathcal{B}}} [1 + e^{-(1/2)A + B + U}]. \end{aligned} \quad (\text{C36})$$

Subsequently, for the moments of Q , we have

$$f_a(B, A; U) = \frac{\partial}{\partial U} [\ln Z_{\mathcal{B}}] = \text{sigm} \left(U + B - \frac{1}{2}A \right), \quad (\text{C37})$$

where sigm is the logistic sigmoid function. Subsequently, the variance for the binary unit can be calculated directly as

$$f_c(B, A; U) = f_a(B, A; U) - f_a(B, A; U)^2. \quad (\text{C38})$$

Next, if we wish to define the learning gradients on U , we write the derivatives of $\ln B(x; U)$ and $\ln Z_Q$ with respect to U ,

$$\frac{\partial}{\partial U} \ln B(x; U) = \frac{\partial}{\partial U} [Ux - \ln Z_B] = x - \langle x \rangle_B, \quad (\text{C39})$$

and, for the log normalization,

$$\frac{\partial}{\partial U} \ln Z_B = f_a(B, A; U) - \langle x \rangle_B. \quad (\text{C40})$$

The resulting gradients for the distribution terms U are

$$\Delta U_i = \langle x_i \rangle_M - \langle a_i \rangle_P \quad (\text{C41})$$

for visible units and

$$\Delta U_j = \langle \tilde{a}_j \rangle_M - \langle a_j \rangle_P \quad (\text{C42})$$

for hidden units, where $\tilde{a}_j \triangleq f_a(\sum_i W_{ij} x_i^{(m)}, 0; U)$.

APPENDIX D: ADAPTIVE TAP

When performing inference, one could employ instead of TAP a variant known as adaptive TAP (or adaTAP) [57], which gives, in general, more accurate results, albeit being slower to iterate. We briefly investigate here the performance of this method in the binary case.

The adaTAP algorithm is more generally presented without a distinction between visible and hidden variables. We thus write the algorithm for generic weight matrix J and bias vector \mathbf{H} , in practice here defined by blocks

$$J = \begin{pmatrix} 0 & W \\ W^T & 0 \end{pmatrix} \in \mathbb{R}^{(N_v+N_h) \times (N_v+N_h)},$$

$$\mathbf{H} = \begin{pmatrix} \mathbf{b} \\ \mathbf{c} \end{pmatrix} \in \mathbb{R}^{(N_v+N_h)},$$

The proposed implementation Algorithm 3 uses the recently introduced vector approximate message passing (VAMP) [73] to find the adaTAP fixed points. After convergence, quantities with subscripts 1 and 2 are equal and identify with the outputs of the TAP inference Algorithm 1,

$$\mathbf{A}_1 = \mathbf{A}_2 = \begin{pmatrix} \mathbf{A}^v \\ \mathbf{A}^h \end{pmatrix}, \quad \mathbf{B}_1 = \mathbf{B}_2 = \begin{pmatrix} \mathbf{B}^v \\ \mathbf{B}^h \end{pmatrix},$$

$$\mathbf{a}_1 = \mathbf{a}_2 = \begin{pmatrix} \mathbf{a}^v \\ \mathbf{a}^h \end{pmatrix}, \quad \mathbf{c}_1 = \mathbf{c}_2 = \begin{pmatrix} \mathbf{c}^v \\ \mathbf{c}^h \end{pmatrix},$$

again compactly defined by blocks over the visible and hidden units. Here, \mathbf{c}_2 is defined as the diagonal of the C_2 matrix, which gives an estimate of the correlation between different units and must be computed at each step of the algorithm. These quantities are then incorporated to our training Algorithm 2.

The computational burden of Algorithm 3 lies in the matrix inversion needed to evaluate C_2 , which needs to be performed at each iteration. In Fig. 10 (right), we compare the time needed to perform one iteration of both the algorithms under identical experimental conditions.

This larger cost per iteration should be compensated, in principle, by a more accurate inference procedure. However, that does not seem to translate to improvements in the training performance. Figure 10 (left) presents a

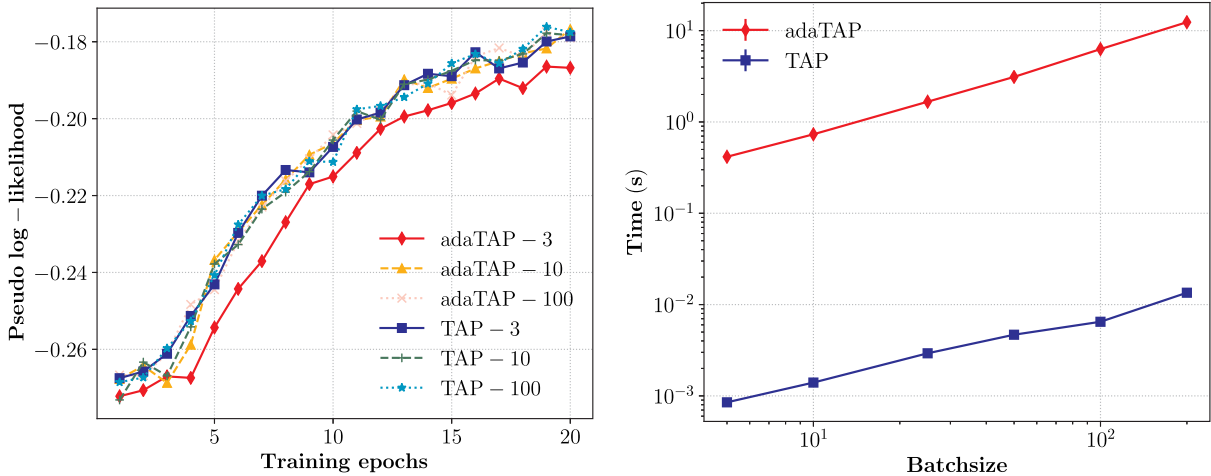


FIG. 10. Left: Evolution of the pseudolikelihood along the training of an RBM with 784 binary visible units and 500 binary hidden units. Training is performed using the 5000 first images of binarized MNIST, with a learning rate of 0.001 and batches of size 100. The different curves correspond to different strategies of estimation of the likelihood gradients, with either TAP or adaTAP. Both algorithms are iterated for a fixed number of times (3, 10, and 100). In all cases, a damping of 0.5 is used. All methods yield comparable results in terms of training performance, except for adaTAP with only three iterations, which shows a poorer performance. Right: Computation time for one iteration of the inference algorithm, as a function of the batch size. Time is reported in seconds for identical experimental settings. The need for a matrix inversion for each batch element makes VAMP 3 orders of magnitude slower than TAP.

minimal test on 5000 MNIST training samples, where performances are reported in terms of the pseudolikelihood. We evaluate both algorithms for different numbers of iterations. Results suggest that all strategies are roughly equivalent, except for running adaTAP for a very small number of iterations, which always leads to a poorer result.

We thus conclude that, as far as proposing a tractable and efficient training algorithm for RBMs, the TAP inference seems to serve the purpose more appropriately.

APPENDIX E: DEEP BOLTZMANN MACHINES

1. Model and inference

It is possible to define as well deep models of Boltzmann machines by considering several stacked hidden layers. These deep Boltzmann machines (DBMs) [26] consist in a straightforward extension of RBMs. The distribution corresponding to a DBM with L hidden layers indexed by l is

$$P(\mathbf{x}, \mathbf{h}^{(1)}, \dots, \mathbf{h}^{(L)}; W^{(1)}, \dots, W^{(L)}, \boldsymbol{\theta}) = \frac{1}{\mathcal{Z}[W^{(1)}, \dots, W^{(L)}, \boldsymbol{\theta}]} e^{\sum_{i,j} x_i W_{ij} h_j^{(1)} + \sum_l \sum_{i,j} h_i^{(l)} W_{ij}^{(l)} h_j^{(l+1)}} \times \prod_i P_i^v(x_i; \theta_i^v) \prod_l \prod_j P_j^{h^{(l)}}(h_j^{(l)}; \theta_j^{h^{(l)}}). \quad (\text{E1})$$

Similarly to RBMs, the distribution of visible variables is obtained by marginalizing out the latent variables:

$$P(\mathbf{x}; W^{(1)}, \dots, W^{(L)}, \boldsymbol{\theta}) = \int \left(\prod_l \prod_j dh_j^{(l)} \right) P(\mathbf{x}, \mathbf{h}^{(1)}, \dots, \mathbf{h}^{(L)}; W^{(1)}, \dots, W^{(L)}, \boldsymbol{\theta}), \quad (\text{E2})$$

yielding the following log-likelihood:

$$\begin{aligned} \ln P(\mathbf{x}; W^{(1)}, \dots, W^{(L)}, \boldsymbol{\theta}) &= -\ln \mathcal{Z}[W^{(1)}, \dots, W^{(L)}, \boldsymbol{\theta}] + \sum_i \ln P_i(x_i; \theta_i^v) \\ &+ \ln \int \left(\prod_l \prod_j dh_j^{(l)} \right) e^{\sum_{i,j} x_i W_{ij} h_j^{(1)} + \sum_l \sum_{i,j} h_i^{(l)} W_{ij}^{(l)} h_j^{(l+1)}} \\ &\times \prod_l \prod_j P_j^{h^{(l)}}(h_j^{(l)}; \theta_j^{h^{(l)}}). \end{aligned} \quad (\text{E3})$$

The major difference between RBMs and DBMs lies in the complexity of evaluating the above expression. Whereas for RBMs only the log-partition features a problematic multi-dimensional integral Eq. (11), here both the log-partition and the last term are intractable. This additional complication carries through to the computation of the gradients necessary for training, since the data-dependent term deriving from the last term of Eq. (E3) is no longer tractable.

This intractability follows from the fact that hidden units in neighboring layers are now connected to each other and are thus no longer conditionally independent. Interestingly, the

first proposal to deal with the data-dependent terms of DBMs consisted in using a naive mean-field approximation [26], while keeping a Monte Carlo-based strategy to compute gradients deriving from the log-partition. In this work, we propose instead to use the TAP approximation for both of them, hence improving on the NMF approximation and avoiding any sampling of the rather complicated RBMs.

The TAP equations related to the log-partition $\mathcal{Z}[W^{(1)}, \dots, W^{(L)}, \boldsymbol{\theta}]$ follow directly from the general derivation of Sec. V for fully connected models, with, however, a different weight matrix being used. For instance, the effective weight matrix of a DBM with two hidden layers is defined by blocks as

$$W = \begin{pmatrix} 0 & W^{(1)} & 0 \\ W^{(1)T} & 0 & W^{(2)} \\ 0 & W^{(2)T} & 0 \end{pmatrix}. \quad (\text{E4})$$

Thus, implementing the GRBM inference Algorithm 1 with the proper weights outputs TAP solutions $\{\mathbf{a}, \mathbf{c}, \mathbf{B}, \mathbf{A}\}$, each of them a vector with components corresponding to the different units in the DBM.

For the last term of Eq. (E3), we recognize the log-partition of a model closely related to the considered DBM, where visible units are not anymore variable but fixed, or *clamped*, at values x_i , and the original interaction between visible and first hidden layer units is replaced by an additional local field on each $h_j^{(1)}$ equal to $\sum_i W_{ij}^{(1)} x_i$. Finally, under this simple modification of the Hamiltonian, TAP equations follow again from the general derivation in Sec. V. The resultant TAP solutions depending on data points \mathbf{x} are said to be clamped and denoted as $\{\bar{\mathbf{a}}(\mathbf{x}), \bar{\mathbf{c}}(\mathbf{x}), \bar{\mathbf{B}}(\mathbf{x}), \bar{\mathbf{A}}(\mathbf{x})\}$.

2. Training algorithm and experiments

The gradients of the log-likelihood with respect to the model parameters $\boldsymbol{\theta}$ are similar to the RBM ones, given by Eqs. (12)–(14). However, the first data-dependent term cannot be analytically computed anymore, and we use the clamped TAP solutions to approximate it. The second term is evaluated using the data-independent TAP solutions, similarly to our strategy for RBMs. The corresponding expressions of the gradients are

$$\begin{aligned} \Delta \theta_i^v &\approx \frac{1}{M} \sum_m \frac{\partial}{\partial \theta_i^v} [\ln P_i^v(x_i^{(m)}; \theta_i^v)] \\ &- \frac{1}{K} \sum_k \frac{\partial}{\partial \theta_i^v} [\ln Z_i^v(\mathbf{B}_{i,k}^v, A_{i,k}^v; \theta_i^v)], \end{aligned} \quad (\text{E5})$$

$$\begin{aligned} \Delta \theta_j^{h^{(l)}} &\approx \frac{1}{M} \sum_m \frac{\partial}{\partial \theta_j^{h^{(l)}}} [\ln Z_j^h(\bar{\mathbf{B}}_j^h(\mathbf{x}^{(m)}), \bar{A}_j^h(\mathbf{x}^{(m)}); \theta_j^{h^{(l)}})] \\ &- \frac{1}{K} \sum_k \frac{\partial}{\partial \theta_j^{h^{(l)}}} [\ln Z_j^h(\mathbf{B}_{j,k}^h, A_{j,k}^h; \theta_j^{h^{(l)}})], \end{aligned} \quad (\text{E6})$$

Algorithm 3. AdaTAP inference for binary-binary RBMs.*Input:* J, \mathbf{H} *Initialize:* $t = 0, \mathbf{A}_1^{(0)}, \mathbf{B}_1^{(0)}$ **repeat***Prior updates*

$$a_{1,i}^{(t+1)} = \text{sigm}(B_{1,i}^{(t)} - A_{1,i}^{(t)}/2)$$

$$c_{1,i}^{(t+1)} = a_{1,i}^{(t+1)}(1 - a_{1,i}^{(t+1)})$$

$$A_{2,i}^{(t+1)} = 1/c_{1,i}^{(t+1)} - A_{1,i}^{(t+1)}$$

$$B_{2,i}^{(t+1)} = 1/(1 - a_{1,i}^{(t+1)}) - B_{1,i}^{(t+1)}$$

Interaction updates

$$C_2^{(t+1)} = (\text{diag}(A_2^{(t+1)}) - J)^{-1}$$

$$a_{2,i}^{(t+1)} = \sum_j (C_2^{(t+1)})_{ij} (B_{2,j}^{(t+1)} + H_j)$$

$$A_{1,i}^{(t+1)} = 1/(C_2^{(t+1)})_{ii} - A_{2,i}^{(t+1)}$$

$$B_{1,i}^{(t+1)} = a_{2,i}^{(t+1)} / (C_2^{(t+1)})_{ii} - B_{2,i}^{(t+1)}$$

 $t = t + 1$ **until** convergence

$$\Delta W_{ij}^{(1)} \approx \frac{1}{M} \sum_m x_i^{(m)} \bar{a}_j^{h^{(1)}}(\mathbf{x}^{(m)}) - \frac{1}{K} \sum_k \{a_{i,k}^v a_{j,k}^{h^{(1)}} + W_{ij}^{(1)} c_{i,k}^v c_{j,k}^{h^{(1)}}\}, \quad (\text{E7})$$

$$\Delta W_{ij}^{(l)} \approx \frac{1}{M} \sum_m \{\bar{a}_i^{h^{(l-1)}}(\mathbf{x}^{(m)}) \bar{a}_j^{h^{(l)}}(\mathbf{x}^{(m)}) + W_{ij}^{(l)} \bar{c}_i^{h^{(l-1)}}(\mathbf{x}^{(m)}) \bar{c}_j^{h^{(l)}}(\mathbf{x}^{(m)})\} - \frac{1}{K} \sum_k \{a_{i,k}^{h^{(l-1)}} a_{j,k}^{h^{(l)}} + W_{ij}^{(l)} c_{i,k}^{h^{(l-1)}} c_{j,k}^{h^{(l)}}\} \quad \text{for } l \geq 2. \quad (\text{E8})$$

Algorithm 4. GDBM training.*Input:* $\mathbf{X}, T_{\text{pre train}}, T_{\text{joint train}}, M, K, R(\cdot)$ *Pretraining*

$$W^{(1)}, \boldsymbol{\theta}^v, \boldsymbol{\theta}^{h^{(1)}} \leftarrow \text{Algorithm 2}(\mathbf{X}, T_{\text{pre train}}, M, K, R(\cdot))$$

for all hidden layers $l \geq 2$ **do**

$$\mathbf{H}^{(l-1)} \sim P(\mathbf{h}^{(l)} | \mathbf{H}^{(l-1)}, \dots, \mathbf{H}^1, \mathbf{X})$$

$$W^{(l)}, \boldsymbol{\theta}^{h^{(l)}} \leftarrow \text{Algorithm 2}(\mathbf{H}^{(l-1)}, T_{\text{pre train}}, M, K, R(\cdot))$$

end for*Joint training**Initialize:* $t = 0$ **repeat****for** all minibatches \mathbf{X}^B of size M **do**

$$\mathbf{a}^{(t+1)}, \mathbf{c}^{(t+1)}, \mathbf{B}^{(t+1)}, \mathbf{A}^{(t+1)} \leftarrow \text{Algorithm 1}(W^{(t)}, \boldsymbol{\theta}^{(t)}) \times K$$

$$\bar{\mathbf{a}}^{(t+1)}, \bar{\mathbf{c}}^{(t+1)}, \bar{\mathbf{B}}^{(t+1)}, \bar{\mathbf{A}}^{(t+1)} \leftarrow \text{Algorithm 1}(\mathbf{X}^B, W^{(t)}, \boldsymbol{\theta}^{(t)})$$

$$W_{ij}^{(l),(t+1)} \leftarrow W_{ij}^{(l),(t)} + \gamma \Delta W_{ij}^{(l),(t)}$$

$$\theta_j^{h^{(l),(t+1)}} \leftarrow \theta_j^{h^{(l),(t)}} + \gamma \Delta \theta_j^{h^{(l),(t)}}$$

$$\theta_i^{v,(t+1)} \leftarrow \theta_i^{v,(t)} + \gamma \Delta \theta_i^{v,(t)}$$

end for $t \leftarrow t + 1$ **until** $t > T_{\text{joint train}}$

These expressions can be plugged to a gradient ascent algorithm, as in the RBM training Algorithm 2. Nevertheless, this simple strategy of simultaneous training of all the parameters of the model (*joint training*) usually fails, as the magnitude of weights of deep layers typically remains very small and the model eventually resembles a mere RBM. Several regularizations have been proposed to tackle this well-known problem of DBM training [16,26,74–76]. In our experiments, we used a greedy layerwise *pretraining* [26], which consists in computing a meaningful initialization of the weights by training the

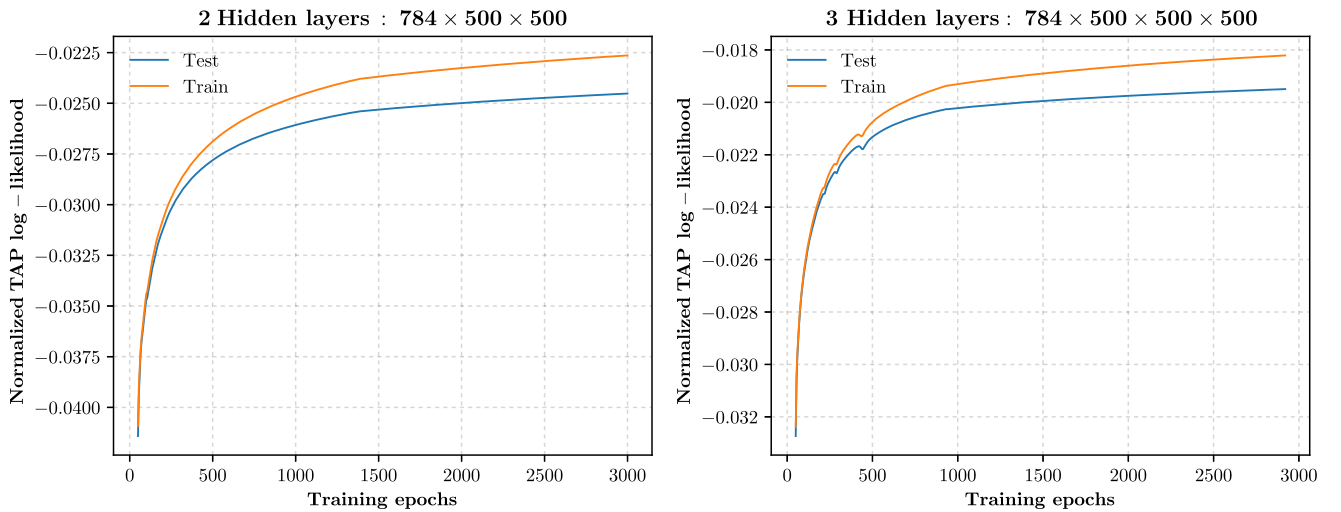


FIG. 11. Training performances over 3000 training epochs for a two-hidden-layer (left) and a three-hidden-layer (right) DBM on the binarized MNIST data sets. Both models are pretrained for 50 epochs with the same learning rate of 0.001. The training performance is measured as the normalized (per unit) TAP log-likelihood for the test images (blue curve) and the train images (orange curve).

RBM layer by layer, before performing the joint training. The complete algorithm is described in Algorithm 4.

Figure 11 shows the evolution of the TAP log-likelihood for a two-hidden-layer and a three-hidden-layer DBM, trained with the above-described algorithm.

APPENDIX F: COMPARISON OF THE TAP LOG-LIKELIHOOD WITH OTHER SURROGATES

As detailed in the main text, we consider the TAP log-likelihood as a *surrogate* to the true and intractable log-likelihood, as there are no guarantees to how close it lies to the true value, nor does it provide any bound on the value of the true log-likelihood—as discussed in Sec. V B, the TAP estimate may fall either above or below. Surrogates, such as the pseudolikelihood [77], are widely used in Boltzmann machine learning [18]. As a surrogate for the true log-likelihood, besides its fast convergence and deterministic calculation, the TAP log-likelihood possesses many interesting properties which we study in our experiments. For instance, TAP machines can be used for denoising, as demonstrated in Sec. VII D, or as priors for other applications of statistical inference.

Another piece of supporting evidence for the quality of the TAP log-likelihood as a surrogate comes from studying how consistent it is with other estimators and surrogates. We present two of them here: the pseudolikelihood (PL)

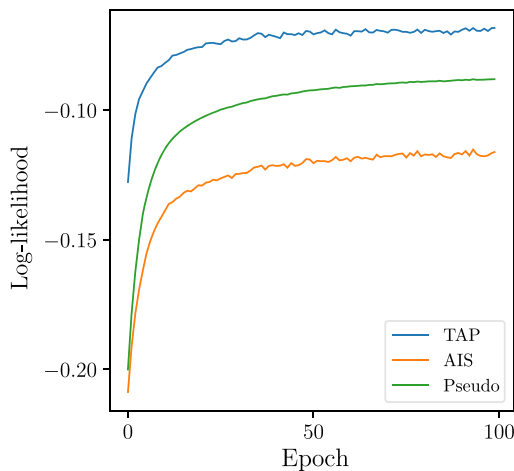


FIG. 12. Comparison of the TAP log-likelihood to two other measures: the estimate provided by AIS and the (log) PL. All three measures are evaluated for the same sets of parameters, which correspond to the training of a binary-binary RBM using the TAP log-likelihood [see Fig. 3(a)]. The AIS results are averaged over 100 runs. For each run, 14 500 intermediate distributions are generated between the initial and target distribution using the same schedule as in Ref. [78]. The pseudolikelihood is computed using a stochastic approximation over 100 of the 784 pixels. Even though the training is performed so as to maximize the TAP log-likelihood, the AIS and PL estimates also increase, indicating the measures to be consistent.

and the log-likelihood estimate provided by annealed importance sampling (AIS) [19]. For this experiment, we compare log-likelihood estimates from TAP, PL, and AIS over the course of 100 epochs of training performed on a single binary-binary RBM with 500 hidden units using the binary-MNIST data set as training data. The training is performed by maximizing the TAP log-likelihood, and each of the three approaches produces either an estimate or a surrogate to the true log-likelihood at the end of each epoch [as in Fig. 3(a)]. We show the comparison of these surrogates over training in Fig. 12.

All three measures have the same qualitative behavior as a function of epochs, indicating they are consistent to each other in this particular experiment. Notably, over training, one does not need to accurately estimate the *value* of the underlying true log-likelihood but rather provide the same quality as a metric, that is, if one set of model parameters has a larger log-likelihood than another according to the true log-likelihood, then it should also be larger in the surrogate. As long as the “landscape” of the log-likelihood is preserved by the surrogate, then it should be sufficient for training accurate models. While the true log-likelihood is unknowable in the context of this experiment, comparing to the AIS estimate shows similar growth in the log-likelihood between the TAP surrogate and AIS. Thus, we can observe that maximizing the TAP log-likelihood does indeed appear to produce models which improve the true log-likelihood overtraining, indicating some correspondence between the TAP log-likelihood and the true log-likelihood, up to the AIS estimate.

- [1] Y. LeCun, Y. Bengio, and G. Hinton, *Deep Learning*, *Nature (London)* **521**, 436 (2015).
- [2] J. J. Hopfield, *Neural Networks and Physical Systems with Emergent Collective Computational Abilities*, *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554 (1982).
- [3] D. J. Amit, H. Gutfreund, and H. Sompolinsky, *Numbers of Patterns in a Spin-Glass Model of Neural Networks*, *Phys. Rev. Lett.* **55**, 1530 (1985).
- [4] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *A Learning Algorithm for Boltzmann Machines*, *Cogn. Sci.* **9**, 147 (1985).
- [5] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond*, *Lect. Notes Phys.* (World Scientific, Singapore, 1987), Vol. 9.
- [6] V. Sessak and R. Monasson, *Small-Correlation Expansions for the Inverse Ising Problem*, *J. Phys. A* **42**, 055001 (2009).
- [7] F. Ricci-Tersenghi, *The Bethe Approximation for Solving the Inverse Ising Problem: A Comparison with Other Inference Methods*, *J. Stat. Mech.* (2012) P08015.
- [8] M. Ekeberg, C. Lökvist, Y. Lan, M. Weigt, and E. Aurell, *Improved Contact Prediction in Proteins: Using Pseudolikelihoods to Infer Potts Models*, *Phys. Rev. E* **87**, 012707 (2013).
- [9] G. Hinton, *Training Products of Experts by Minimizing Contrastive Divergence*, *Neural Comput.* **14**, 1771 (2002).

- [10] G. Hinton and R. Salakhutdinov, *Reducing the Dimensionality of Data with Neural Networks*, *Science* **313**, 504 (2006).
- [11] H. Larochelle and Y. Bengio, in *Proceedings of the International Conference on Machine Learning* (2008), pp. 536–543.
- [12] R. Salakhutdinov, A. Mnih, and G. Hinton, in *Proceedings of the International Conference on Machine Learning* (2007), pp. 791–798.
- [13] A. Coates, A. Y. Ng, and H. Lee, in *Proceedings of the International Conference on Artificial Intelligence and Statistics* (2011), pp. 215–223.
- [14] G. Hinton and R. Salakhutdinov, in *Proceedings of Advances in Neural Information Processing Systems* (2009), pp. 1607–1614.
- [15] G. Hinton, S. Osindero, and Y.-W. Teh, *A Fast Learning Algorithm for Deep Belief Nets*, *Neural Comput.*, **18**, 1527 (2006).
- [16] R. Salakhutdinov and H. Larochelle, in *Proceedings of the International Conference on Artificial Intelligence and Statistics* (2010), pp. 693–700.
- [17] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, 2016), <http://www.deeplearningbook.org>.
- [18] T. Tieleman, in *Proceedings of the International Conference on Machine Learning* (2008), pp. 1064–1071.
- [19] R. M. Neal, *Annealed Importance Sampling*, *Stat. Comput.* **11**, 125 (2001).
- [20] R. Salakhutdinov and I. Murray, in *Proceedings of the International Conference on Machine Learning* (2008), pp. 872–879.
- [21] Y. Burda, R. B. Grosse, and R. Salakhutdinov, in *Proceedings of the International Conference on Artificial Intelligence and Statistics* (2015), Vol. 38, pp. 102–110.
- [22] H. Schulz, A. Müller, and S. Behnke, in *Proceedings of the NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, Whistler, CA, 2010 (2010).
- [23] D. J. Thouless, P. W. Anderson, and R. G. Palmer, *Solution of “Solvable Model of a Spin Glass”*, *Philos. Mag.* **35**, 593 (1977).
- [24] M. Opper and O. Winther, *A Mean Field Algorithm for Bayes Learning in Large Feedforward Neural Networks*, *Proceedings of Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, 1997), pp. 225–231.
- [25] L. Bachschmid-Roman, C. Battistin, M. Opper, and Y. Roudi, *Variational Perturbation and Extended Plefka Approaches to Dynamics on Random Networks: The Case of the Kinetic Ising Model*, *J. Phys. A* **49**, 434003 (2016).
- [26] R. Salakhutdinov and G. Hinton, in *Proceedings of the International Conference on Artificial Intelligence and Statistics* (2009), pp. 448–455.
- [27] P. Smolensky, *Information Processing in Dynamical Systems: Foundations of Harmony Theory* (MIT, Cambridge, MA, 1986).
- [28] G. Desjardins, A. Courville, Y. Bengio, P. Vincent, and O. Delalleau, *Tempered Markov Chain Monte Carlo for training of Restricted Boltzmann Machines*, *Proceedings of the International Conference on Artificial Intelligence and Statistics* (2010), pp. 145–152.
- [29] R. Salakhutdinov, *Learning in Markov Random Fields Using Tempered Transitions*, *Proceedings of Advances in Neural Information Processing Systems* (2010), Vol. 23, p. 1.
- [30] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle, *Greedy Layer-wise Training of Deep Networks*, *Proceedings of Advances in Neural Information Processing Systems* (2007), Vol. 19, p. 153.
- [31] K. Cho, *Improved Learning Algorithms for Restricted Boltzmann Machines*, Master’s Thesis, Aalto University School of Science, 2011.
- [32] M. Welling, M. Rosen-Zvi, and G. E. Hinton, *Exponential Family Harmoniums with an Application to Information Retrieval*, *Proceedings of Advances in Neural Information Processing Systems* (2004), p. 1481.
- [33] A. Fischer and C. Igel, *Empirical Analysis of the Divergence of Gibbs Sampling Based Learning Algorithms for Restricted Boltzmann Machines*, *Proceedings of the International Conference on Artificial Neural Networks* (2010), pp. 1–10.
- [34] T. Plefka, *Convergence Condition of the TAP Equation for the Infinite-Ranged Ising Spin Glass Model*, *J. Phys. A* **15**, 1971 (1982).
- [35] A. Georges and J. S. Yedidia, *How to Expand around Mean-Field Theory Using High-Temperature Expansions*, *J. Phys. A* **24**, 2173 (1991).
- [36] J. Yedidia, in *Advanced Mean Field Methods: Theory and Practice* (MIT, Cambridge, MA, 2001), pp. 21–36.
- [37] F. Zamponi, *Mean Field Theory of Spin Glasses* (to be published).
- [38] M. Welling and G. Hinton, in *Proceedings of the International Conference on Artificial Neural Networks* (2002), pp. 351–357.
- [39] H. J. Kappen and F. B. Rodríguez, *Efficient Learning in Boltzmann Machines Using Linear Response Theory*, *Neural Comput.* **10**, 1137 (1998).
- [40] M. Gabrié, E. W. Tramel, and F. Krzakala, *Training Restricted Boltzmann Machines via the Thouless-Anderson-Palmer Free Energy*, *Proceedings of Advances in Neural Information Processing Systems* (2015), p. 640.
- [41] M. Welling and Y. W. Teh, *Approximate Inference in Boltzmann Machines*, *Artif. Intell.* **143**, 19 (2003).
- [42] E. W. Tramel, A. Drémeau, and F. Krzakala, *Approximate Message Passing with Restricted Boltzmann Machine Priors*, *J. Stat. Mech.* (2016) 073401.
- [43] F. Ricci-Tersenghi, *The Bethe Approximation for Solving the Inverse Ising Problem: A Comparison with Other Inference Methods*, *J. Stat. Mech.* (2012) P08015.
- [44] M. Mézard, *Mean-Field Message-Passing Equations in the Hopfield Model and Its Generalizations*, *Phys. Rev. E* **95**, 022117 (2017).
- [45] A. Courville, J. Bergstra, and Y. Bengio, in *Proceedings of the International Conference on Artificial Intelligence and Statistics* (2011), Vol. 15.
- [46] J. Pearl, in *Proceedings of the AAAI Conference on Artificial Intelligence* (1982), pp. 133–136.
- [47] J. S. Yedidia, W. T. Freeman, and Y. Weiss, *Bethe Free Energy, Kikuchi Approximations, and Belief Propagation Algorithms*, *Proceedings of Advances in Neural Information Processing Systems* (2001), Vol. 13.
- [48] H. Huang and T. Toyozumi, *Advanced Mean-Field Theory of the Restricted Boltzmann Machine*, *Phys. Rev. E* **91**, 050101 (2015).

- [49] E. B. Sudderth, A. T. Ihler, M. Isard, W. T. Freeman, and A. S. Willsky, *Nonparametric Belief Propagation*, *Commun. ACM* **53**, 95 (2010).
- [50] M. Opper and O. Winther, *Expectation Consistent Approximate Inference*, *J. Mach. Learn. Res.* **6**, 2177 (2005).
- [51] S. Rangan, in *Proceedings of the Annual Conference on Information Science and Systems* (2010), pp. 1–6.
- [52] T. Lesieur, F. Krzakala, and L. Zdeborová, in *Proceedings of the IEEE International Symposium on Information Theory* (IEEE, New York, 2015), pp. 1635–1639.
- [53] T. Lesieur, F. Krzakala, and L. Zdeborová, in *Proceedings of the Allerton Conference on Communication, Control, and Computing* (2015), pp. 680–687.
- [54] Y. Deshpande and A. Montanari, in *Proceedings of the IEEE International Symposium on Information Theory* (IEEE, New York, 2014), pp. 2197–2201.
- [55] S. Rangan and A. K. Fletcher, in *Proceedings of the IEEE International Symposium on Information Theory* (IEEE, New York, 2012), pp. 1246–1250.
- [56] T. Lesieur, F. Krzakala, and L. Zdeborová, *Constrained Low-Rank Matrix Estimation: Phase Transitions, Approximate Message Passing and Applications*, *J. Stat. Mech.* (2017) 073403.
- [57] M. Opper and O. Winther, *Tractable Approximations for Probabilistic Models: The Adaptive Thouless-Anderson-Palmer Mean Field Approach*, *Phys. Rev. Lett.* **86**, 3695 (2001).
- [58] J. S. Yedidia, *An Idiosyncratic Journey beyond Mean Field Theory* (MIT, Cambridge, MA, 2000).
- [59] M. J. Wainwright and M. I. Jordan, *Graphical Models, Exponential Families, and Variational Inference*, Foundations and Trends in Machine Learning (Now Publishers Inc., Hannover, 2008), Vol. 1.
- [60] C. D. Dominicis and A. P. Young, *Weighted Averages and Order Parameters for the Infinite Range Ising Spin Glass*, *J. Phys. A* **16**, 2063 (1983).
- [61] M. Mézard and A. Montanari, *Information, Physics, and Computation* (Oxford University, New York, 2009).
- [62] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, *Probabilistic Reconstruction in Compressed Sensing: Algorithms, Phase Diagrams, and Threshold Achieving Matrices*, *J. Stat. Mech.* (2012) P08009.
- [63] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, *Gradient-Based Learning Applied to Document Recognition*, *Proc. IEEE* **86**, 2278 (1998).
- [64] MIT Center for Biological and Computation Learning, CBCL face database No. 1, 2000.
- [65] K. Cho, T. Raiko, and A. Ilin, *Enhanced Gradient for Training Restricted Boltzmann Machines*, *Neural Comput.* **25**, 805 (2013).
- [66] J. B. Tenenbaum, V. de Silva, and J. C. Langford, *A Global Geometric Framework for Nonlinear Dimensionality Reduction*, *Science* **290**, 2319 (2000).
- [67] L. Şendur and I. W. Selesnick, *Bivariate Shrinkage Functions for Wavelet-based Denoising Exploiting Inter-scale Dependency*, *IEEE Trans. Signal Process.* **50**, 2744 (2002).
- [68] B. van Ginneken and A. Mendrik, in *Proceedings of the International Conference on Pattern Recognition* (2006).
- [69] J. Melchior, N. Wang, and L. Wiskott, *Gaussian-Binary Restricted Boltzmann Machines on Modeling Natural Image Statistics*, *PLoS One* **12**, e0171015 (2017).
- [70] K. Cho, A. Ilin, and T. Raiko, in *Proceedings of the International Conference on Artificial Neural Networks* (2011), pp. 10–17.
- [71] J. Tubiana and R. Monasson, *Emergence of Compositional Representations in Restricted Boltzmann Machines*, *Phys. Rev. Lett.* **118**, 138301 (2017).
- [72] A. Barra, G. Genovese, P. Sollich, and D. Tantari, *Phase Transitions in Restricted Boltzmann Machines with Generic Priors*, *Phys. Rev. E* **96**, 042156 (2017).
- [73] S. Rangan, P. Schniter, and A. K. Fletcher, in *Proceedings of the 2017 IEEE International Symposium on Information Theory (ISIT)* (IEEE, New York, 2017), pp. 1588–1592.
- [74] G. Montavon and K. R. Müller, in *Neural Networks: Tricks of the Trade* (Springer, New York, 2012), pp. 621–637.
- [75] G. Desjardins, A. Courville, and Y. Bengio, *On Training Deep Boltzmann Machines*, arXiv:1203.4416.
- [76] J. Melchior, A. Fischer, and L. Wiskott, *How to Center Deep Boltzmann Machines*, *J. Mach. Learn. Res.* **17**, 1 (2016).
- [77] J. Besag, *Statistical Analysis of Non-Lattice Data*, *Statistician* **24**179 (1975).
- [78] R. Salakhutdinov, MATLAB code for AIS, http://www.utstat.toronto.edu/~rsalakhu/rbm_ais.html.