

Unsupervised hierarchical clustering using the learning dynamics of restricted Boltzmann machines

Aurélien Decelle  and Beatriz Seoane

*Departamento de Física Teórica, Universidad Complutense de Madrid, 28040 Madrid, Spain
and Université Paris-Saclay, CNRS, INRIA Tau team, LISN, 91190 Gif-sur-Yvette, France*

Lorenzo Rosset *

Departamento de Física Teórica, Universidad Complutense de Madrid, 28040 Madrid, Spain



(Received 16 March 2023; revised 19 May 2023; accepted 9 June 2023; published 7 July 2023)

Data sets in the real world are often complex and to some degree hierarchical, with groups and subgroups of data sharing common characteristics at different levels of abstraction. Understanding and uncovering the hidden structure of these data sets is an important task that has many practical applications. To address this challenge, we present a general method for building relational data trees by exploiting the learning dynamics of the restricted Boltzmann machine. Our method is based on the mean-field approach, derived from the Plefka expansion, and developed in the context of disordered systems. It is designed to be easily interpretable. We tested our method in an artificially created hierarchical data set and on three different real-world data sets (images of digits, mutations in the human genome, and a homologous family of proteins). The method is able to automatically identify the hierarchical structure of the data. This could be useful in the study of homologous protein sequences, where the relationships between proteins are critical for understanding their function and evolution.

DOI: [10.1103/PhysRevE.108.014110](https://doi.org/10.1103/PhysRevE.108.014110)

I. INTRODUCTION

The advent of new technologies is fostering the collection of data in unprecedented quantities. The difficulty is to automatically extract information or to discover core principles from these collected data, which are frequently neither labeled nor classified. Biology provides a key illustration of this situation: modern sequencing techniques generate vast amounts of raw data (e.g., genome or protein sequences) far more quickly than researchers can connect them to biological knowledge through experimental procedures or curated annotations. Building the theoretical foundation needed to extract scientifically meaningful insights from massive data sets is more important than ever. In this study, we focus on the problem of unsupervised identification of categories and subcategories in the data without introducing interfering prior knowledge that could bias the results.

Energy-based models in unsupervised learning are particularly well suited to interpretability applications because their operating mechanism is precisely to encode the empirical distribution of the data in the Boltzmann distribution of an energy function (often called a *Hamiltonian*). If simple enough, such a Hamiltonian serves as an effective model for the data, and it can be analyzed for feature extraction. This strategy has been very popular in biology, for instance for protein-structure prediction [1–3]. In this sense, restricted Boltzmann machines (RBMs) are very well suited for this task because they are powerful enough to encode any complex data set [4], and

at the same time simple enough (formally analogous to a disordered spin Hamiltonian in physics) to be described with the tools of statistical physics. In fact, significant progress has been made in the past decade in understanding the learning mechanisms of the RBM, which has provided a theoretical basis for extracting statistical features in these models [5–9].

Today we know that learning in an RBM is triggered when the eigenvectors of its coupling matrix gradually match the first principal component analysis (PCA) directions of the data set [5–7]. As learning progresses, the RBM finds more complex nonlinear decompositions of the data, which are critical to producing increasingly reliable fake data. The goal is then to use these features to go beyond the aforementioned PCA for dimensionality reduction to identify hidden communities, or features, in the data. In addition, one can exploit the progressive nature of learning to hierarchize those features: the machine first learns global features to gradually encode more local and fine-grained features that can help us to identify subcommunities. At this point, the simplicity of the RBM architecture provides a clear advantage over more complex neural networks, as the log-likelihood of the models at different stages of the learning process can be easily explored using standard tools of statistical physics of disordered systems, such as the TAP equations to find different probability maxima [10–12]. In this work, we implement this idea in real-world data sets, exploiting current knowledge of the learning mechanisms of RBMs to build relational trees with data. We then use these hierarchies to divide the data into families and subfamilies.

Our work is structured as follows: We begin with an introduction to the RBM model and its learning mechanisms.

*lrosset@ucm.es

This is followed by a mean-field description of this model, which is used to find the multiple maxima of the model's probability distribution function that are further used to cluster the data at different stages of the learning process. Later, we explain how we exploit the dynamics of these fixed points to build a tree with a data set in which categories emerge naturally as learning encodes finer and finer features in the model. Finally, we show the implementation of this algorithm on several real-world data sets.

II. DEFINITION OF THE MODEL

The RBM is a Markov random field with pairwise interactions defined on a bipartite graph of two noninteracting layers of variables: the visible variables $\mathbf{v} = \{v_i\}_{i=1, \dots, N_v}$ represent the data, while the hidden variables $\mathbf{h} = \{h_\mu\}_{\mu=1, \dots, N_h}$ constitute a latent representation of the data that models the effective interactions among the visible variables. The joint probability distribution of visible and hidden variables is given by the Boltzmann distribution

$$p_\theta(\mathbf{v}, \mathbf{h}) = \frac{1}{Z_\theta} e^{-E(\mathbf{v}, \mathbf{h}; \theta)}, \quad \text{where } Z_\theta = \sum_{\mathbf{v}, \mathbf{h}} e^{-E(\mathbf{v}, \mathbf{h}; \theta)}. \quad (1)$$

In the previous expressions, the normalization factor Z_θ is called the *partition function*, θ refers to the parameters of the model, and E is the energy function or Hamiltonian. In this work, we focus on a setting in which the hidden units are binary variables, $h_\mu \in \{0, 1\}$, while the visible units are categorical variables that can assume values $q \in \{1, \dots, N_q\}$. We can model this kind of system by introducing a set of *Potts variables* in the visible layer; the associated Hamiltonian is

$$E(\mathbf{v}, \mathbf{h}; \theta) = - \sum_{iq} a_i^q \delta_{v_i, q} - \sum_{\mu} b_\mu h_\mu - \sum_{i\mu} \delta_{v_i, q} w_{i\mu}^q h_\mu, \quad (2)$$

where $\delta_{v_i, q}$ is the Kronecker symbol that returns 1 if the site i of the visible layer has the value q and 0 otherwise, $\mathbf{a} = \{a_i^q\}$ and $\mathbf{b} = \{b_\mu\}$ are two sets of local fields acting, respectively, on the visible and on the hidden layers, and $\mathbf{w} = \{w_{i\mu}^q\}$ is a 3-rank tensor, which we simply refer to as *weight matrix*, that models the interactions between the two layers. We sketched the Potts-RBM structure in Fig. 1. We would like to emphasize that the selection of Potts variables as visible variables for the RBM is a generalization that allows the study of sequence data within the same conceptual framework that applies to binary data. Nevertheless, binary Ising variables are sufficient for most clustering applications where the data set can be made of binary states.

The model defined by the Hamiltonian (2) is not uniquely defined, and common practice in Boltzmann learning is to fix the gauge in order to remove the flat directions caused by the overparametrization [2]. Recent works [13] suggested using the following gauge-fixing prescription:

$$\sum_q a_i^q = 0, \quad \sum_q w_{i\mu}^q = 0, \quad \forall i = 1, \dots, N_v, \quad \forall \mu = 1, \dots, N_h, \quad (3)$$

which is called *zero-sum gauge*. This choice preserves symmetry between the different categories and has the advantage




FIG. 1. Scheme of the Potts-RBM. The weights matrix is a three-dimensional tensor $\mathbf{w} = \{w_{i\mu}^q\}$. When input data $\mathbf{v}^{(k)}$ enters the machine, it selects at each site i the connection of the proper color $v_i = q$ with the hidden component μ .

of being computationally easier to implement. However, the choice of the actual gauge should not affect our clustering procedure.

The RBM is trained by maximizing the log-likelihood (LL) function of the model computed on the data set $\mathcal{D} = \{\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(M)}\}$, which is defined as

$$\begin{aligned} \mathcal{L}^{\text{uw}}(\theta|\mathcal{D}) &= \frac{1}{M} \sum_{m=1}^M \log p_\theta(\mathbf{v} = \mathbf{v}^{(m)}) \\ &= \frac{1}{M} \sum_{m=1}^M \log \sum_{\mathbf{h}} e^{-E(\mathbf{v}^{(m)}, \mathbf{h}; \theta)} - \log Z_\theta, \end{aligned} \quad (4)$$

via (stochastic) gradient ascent. For some data sets, however, it may be more convenient to look at a weighted version of the LL instead,

$$\begin{aligned} \mathcal{L}(\theta|\mathcal{D}) &= \langle \log p_\theta(\mathbf{v}) \rangle_{\mathcal{D}} = \frac{1}{M_{\text{eff}}} \sum_{m=1}^M x_m \log p_\theta(\mathbf{v} = \mathbf{v}^{(m)}) \\ &= \frac{1}{M_{\text{eff}}} \sum_{m=1}^M x_m \log \sum_{\mathbf{h}} e^{-E(\mathbf{v}^{(m)}, \mathbf{h}; \theta)} - \log Z_\theta, \end{aligned} \quad (5)$$

where $M_{\text{eff}} = \sum_m x_m$, and $\{x_m\}_{m=1, \dots, M}$ is a set of weights associated with the data samples that have to be specified for each data set. The weighted average likelihood is useful to compensate for the fact that certain categories may be over-represented in the data set compared to others. In particular, reweighting the data set can affect the process of splitting the data into categories (especially since larger groups are split first in the tree), but it has the disadvantage of requiring to know the groups in the first place. For this reason, we will generally not use this weighted version except for the protein sequence data set to partially mitigate the fact that the data entries may not be independent of each other. This is typically the case for the available protein sequence data because the genomes of some species are sequenced more frequently than others [3]. In this case, reweighting will not occur between functional groups but between sequences with similar sequence identity.

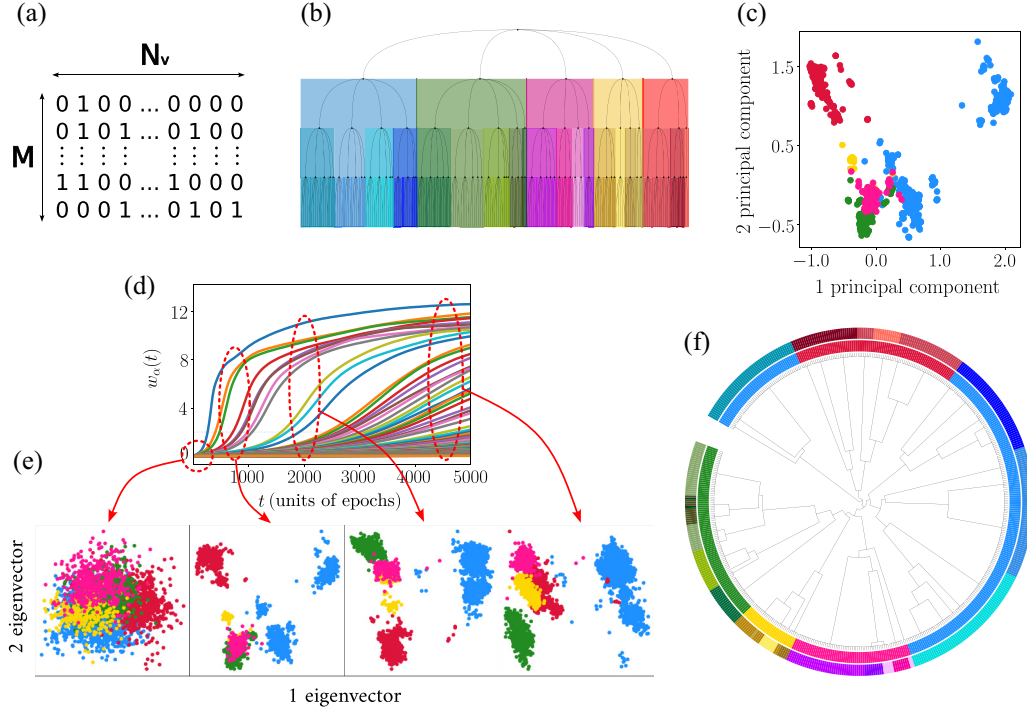


FIG. 2. Different representations of the SED data set and evolution of the features learned by the model as a function of the training time. (a) SED data set. Each row is a data vector made of N_v binary values. (b) First four levels of the real tree obtained from the SED data set. Each node of the tree corresponds to a data sample. (c) First two principal components of the SED data set's PCA. (d) Eigenvalues of the weight matrix, $\{w_\alpha\}$, as a function of the training time in terms of epochs. (e) Projections of the SED data set along the first two principal directions of the weight matrix at different epochs. (f) Tree of the SED data set reconstructed with the proposed method. The colored tags in the external rings represent the true labels (most internal ring) and sublabels (outer ring) of the data points.

The gradient of the LL with respect to the model's parameters is given by

$$\frac{\partial \mathcal{L}}{\partial a_i^q} = \langle \delta_{v_i, q} \rangle_{\mathcal{D}} - \langle \delta_{v_i, q} \rangle_E, \quad \frac{\partial \mathcal{L}}{\partial b_\mu} = \langle h_\mu \rangle_{\mathcal{D}} - \langle h_\mu \rangle_E,$$

$$\text{and } \frac{\partial \mathcal{L}}{\partial w_{i\mu}^q} = \langle \delta_{v_i, q} h_\mu \rangle_{\mathcal{D}} - \langle \delta_{v_i, q} h_\mu \rangle_E, \quad (6)$$

where we denoted as $\langle \cdot \rangle_{\mathcal{D}}$ the weighted average over the data set and as $\langle \cdot \rangle_E$ the average over the model's Boltzmann measure (1). Since computing the partition function is generally intractable, this second average is typically approximated by the Block-Gibbs sampling Markov chain Monte Carlo (MCMC) method, which exploits the bipartite structure of the graph to iteratively sample the visible variables and the hidden variables conditioned on the other layer at once, hence allowing for a high degree of parallelization. In particular, in this work we have always used the persistent contrastive divergence (PCD- k) method [14]: the initial states of the parallel Markov chains used to compute the gradient are taken from the final states obtained in the previous parameter's update. In this scheme, the number of steps performed in each update is fixed at k . It has been numerically investigated that this recipe can yield quasiequilibrium models [7]. We also systematically checked that the learned models always operated in the so-called equilibrium regime [7,15], i.e., no memory of the training scheme was observed. In this paper, we will often refer to the number of epochs, N_{epochs} , performed during the learning as the *age* of the RBM, which can also be

expressed in terms of the number of the gradient updates as $N_{\text{grad}} = N_{\text{epochs}} \cdot M/N_{\text{mb}}$, where M is the number of training samples and N_{mb} is the minibatch size.

III. USING THE RBM'S LEARNING HISTORY TO CLASSIFY THE DATA

In this section, we discuss how the training trajectory of the RBM can provide useful information to reveal the hidden structure of a data set. As a proof of concept, we artificially generated a data set that mimics a completely random evolutionary process, which we call SED (synthetic evolutionary data). We started with a "mother sequence" of zeros of length $N_v = 805$ and generated a descendant set of five sequences by flipping each site from 0 to 1 with probability $p = 1/N_v$. We then assigned a label to each of the child sequences and generated several more generations with the same criterion, each sequence having a randomly chosen number of heirs between one and five. All descendants of any of the five ancestral sequences carry the same label and are referred to as a *family*. We can even go further by assigning a secondary label to each of the subtrees originating from the grandchildren of the "mother sequence." Panels (a) and (b) of Fig. 2 show, respectively, a sketch of the structure of the binary data set, and the genealogical tree corresponding to this evolutionary process.

The most naive analysis we can do to infer families from the SED data set is the PCA. We show the projections of the data onto the first two principal directions in Fig. 2(c), and

color each data point according to its family color. We can see that this analysis already does a relatively good job of separating sequences by family. For example, the red family (and perhaps the blue) can be directly isolated at this level of analysis, but the remaining points of the families remain mixed. Our goal now is to use the RBM to make a finer delineation and obtain some information about the internal structure of the families, e.g., to isolate the subfamilies of the grandchildren in Fig. 2(b). For the sake of visualization, the eigenvalues and eigenvectors of \mathbf{w} are obtained after reshaping the tensor into the form of an $(N_v \cdot N_q) \times N_h$ matrix and taking its SVD (see Appendix A 2 for details). From previous works [5,6,16] we know that, during the initial steps of learning, the RBM gradually encodes the principal directions of the data set in the SVD of the \mathbf{w} -matrix. In Fig. 2(d) we show the evolution of the eigenvalues of the weight matrix as a function of the training epochs, while in Fig. 2(e) we show the data set projected along the first two eigenvectors of \mathbf{w} at different stages of learning (different epochs). At the beginning of the learning process, the SVD eigenvectors do not contain useful information and cannot separate the families (all colored points are mixed). In a second phase, when only the first three eigenvalues have grown, the projection closely resembles that of the PCA in Fig. 2(c), thus agreeing with the previous statement that the SVD eigenvectors align with the direction of the principal directions [5]. After that, the RBM starts learning nonlinear transformations of the data. In the third frame, corresponding to the point where the RBM has about 10 expressed eigenvalues, the projections partition the different families much better even at the level of a two-dimensional projection. At much longer training times, we typically see that the clusters begin to overlap, and direct clustering of such projections using a few dimensions becomes difficult, even when the machine has learned much better representations of the data.

With this discussion, we wanted to argue that the evolution of the free energy of the model during training can be used to explore different levels of clustering, and in this way extract hierarchical information from the data. Such analysis can also give us qualitative insight into the feature extraction process in RBMs. Then, the underlying assumption is that the features learned by the RBM during the learning should have a hierarchical structure, that is, the “most relevant” properties of the data are learnt quite early during the training, while the finer substructure is detected only at later ages of the machine. The application of this method to the SED data set yields the tree reported in Fig. 2(f); we are going to discuss it in the next sections. We see that data are correctly separated into distinct families and subfamilies, and a rich structure of the data set has been detected, although the blueish data get split into two distinct branches right from the beginning of the tree. The criterion we used for categorizing the data at different ages of the learning is based on the mean-field theory applied to the RBM, and it is discussed in the next section.

IV. PROBING THE RBM LANDSCAPE: A MEAN-FIELD ANALYSIS

The method that we propose is based on the mean-field description of the RBM. In this context, we can describe the system through an approximated thermodynamic potential which depends on the system’s magnetizations. This extended free energy displays a certain number of metastable minima in the magnetization space [17], each of which corresponds to a mode (i.e., a local maximum) of the Boltzmann probability distribution. We can interpret these metastable states as *features* of the data that the model was able to capture at a certain stage of the learning. Using the mean-field theory, we can derive a set of self-consistent equations that, iterated until convergence, allow us to associate each data point in the data set to the “closest” metastable state, thereby categorizing/clustering the data into families with a similar structure to the RBM eyes. Thus, by following this assignment procedure during training for different free-energy landscapes of progressively trained machines, we have a method for hierarchically classifying the data.

A. Derivation of the self-consistent equations for the Potts RBM

In what follows, we generalize the mean-field description of Refs. [18–20], also known as Plefka expansion, to our target model defined by the Hamiltonian in Eq. (2). These high-temperature expansions have already been applied to the Boltzmann machine [21] and to RBMs with binary visible variables [10–12] and to the infinite-range Potts-glass model [22], but we have yet to generalize these results to the RBM case with Potts variables in the visible layer. In this section, we outline only the main results that will be used later to build the trees, while further details on the derivation can be found in Appendix B.

Let us consider an energy function $E_\beta(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta})$ which depends on the inverse-temperature β . By introducing two sets of temperature-dependent external fields $\mathbf{B} = \{B_i\} = (\{\phi_j^p\}, \{\psi_\mu\})$ with $i \in \{1, \dots, N_v \cdot N_q + N_h\}$, interacting, respectively, with the visible and the hidden variables, we consider the free energy of this extended system:

$$-\beta F_\beta(\mathbf{B}) = \log \sum_{\mathbf{v}, \mathbf{h}} \exp \left[-\beta E_\beta(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta}) + \sum_{iq} \phi_i^q(\beta) \delta_{v_i, q} + \sum_{\mu} \psi_\mu(\beta) h_\mu \right], \quad (7)$$

where we omit the dependence of the free energy on the systems’ parameters $\boldsymbol{\theta}$. By taking the Legendre transform of (7), we obtain a description of the system in terms of the magnetizations $\mathbf{M} = \{M_i\} = (\{f_j^p\}, \{m_\mu\})$ through the Gibbs free energy,

$$\begin{aligned} -\beta G_\beta(\mathbf{M}) &= -\beta \max_{\mathbf{B}} \left[F_\beta(\mathbf{B}) + \frac{1}{\beta} \sum_{iq} \phi_i^q(\beta) f_i^q + \frac{1}{\beta} \sum_{\mu} \psi_\mu(\beta) m_\mu \right] \\ &= \min_{\mathbf{B}} \left[\log \sum_{\mathbf{v}, \mathbf{h}} \exp \left(-\beta E_\beta(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta}) + \sum_{iq} \phi_i^q(\beta) (\delta_{v_i, q} - f_i^q) + \sum_{\mu} \psi_\mu(\beta) (h_\mu - m_\mu) \right) \right], \end{aligned} \quad (8)$$

where the maximum condition is taken because $F_\beta(\mathbf{B})$ is a concave function of the fields. Consequently, this new thermodynamic potential is a convex function of the magnetizations. For the properties of the Legendre transform, fields and magnetizations are related through

$$M_i(\mathbf{B}) = -\beta \frac{\partial}{\partial B_i} F_\beta(\mathbf{B}) \quad \text{and} \quad B_i(\mathbf{M}) = \beta \frac{\partial}{\partial M_i} G_\beta(\mathbf{M}). \quad (9)$$

The extended free energy (7) can then be recovered as the inverse Legendre transform of the Gibbs free energy (8),

$$-\beta F_\beta(\mathbf{B}) = -\beta \min_{\mathbf{M}} \left[G_\beta(\mathbf{M}) - \frac{1}{\beta} \sum_{iq} \phi_i^q(\beta) f_i^q - \frac{1}{\beta} \sum_{\mu} \psi_\mu(\beta) m_\mu \right]. \quad (10)$$

From Eq. (10), we see that the free energy of the original system—corresponding to the free energy (7) when the external fields are turned off—can be obtained as the minimum of the Gibbs free energy,

$$\begin{aligned} A_\beta &= -\frac{1}{\beta} \log \sum_{\mathbf{v}, \mathbf{h}} e^{-\beta E_\beta(\mathbf{v}, \mathbf{h}, \boldsymbol{\theta})} \\ &= F_\beta(\mathbf{B} = \mathbf{0}) = \min_{\mathbf{M}} [G_\beta(\mathbf{M})]. \end{aligned} \quad (11)$$

In this way, we derived a method to compute the free energy of the system by operating in the space of magnetizations. However, this is true only in principle, because the Gibbs free

energy is as difficult to compute as the true free energy A_β . Besides, our goal of detecting the metastable states of the system cannot be accomplished by working with the Gibbs free energy, because it displays just one global minimum. An estimation of the Gibbs free energy and the breaking of the minima degeneracy can be achieved by considering a high-temperature regime in which we can approximate the thermodynamic potential with a Taylor expansion around $\beta = 0$ truncated at a certain order [17]. In practice, the high-temperature limit corresponds to a regime in which interactions among the variables are negligible compared to thermal fluctuations. This is the typical situation that we encounter at the beginning of the learning, when the weight matrix stores only minimal information on how to couple the layers of the RBM to reproduce the data set's statistics. To formalize this, we can consider a Hamiltonian of the kind

$$-\beta E_\beta(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta}) = \sum_{iq} a_i^q \delta_{v_i, q} + \sum_{\mu} b_\mu h_\mu + \beta \sum_{iq\mu} \delta_{v_i, q} w_{i\mu}^q h_\mu, \quad (12)$$

where the inverse temperature β is just a dummy parameter that serves as the Taylor expansion and that we will eventually set to 1. Indeed, by redefining $\tilde{\mathbf{w}} = \beta \mathbf{w}$ we see that the limit $\beta \rightarrow 0$ is equivalent to the limit $\tilde{w}_\alpha \rightarrow 0$, where \tilde{w}_α are the eigenvalues of the weight matrix. That is indeed the case when the RBM has not yet learned the modes of the data set [5].

By defining $\Gamma_\beta(\mathbf{M}) = -\beta G_\beta(\mathbf{M})$ for simplicity, we obtain the following second-order expansion of the Gibbs free energy:

$$\begin{aligned} \Gamma_\beta^{(2)}(\mathbf{M}) &= \Gamma_0(\mathbf{M}) + \beta \left. \frac{\partial \Gamma_\beta(\mathbf{M})}{\partial \beta} \right|_{\beta=0} + \frac{\beta^2}{2} \left. \frac{\partial^2 \Gamma_\beta(\mathbf{M})}{\partial \beta^2} \right|_{\beta=0} \\ &= \sum_{iq} f_i^q a_i^q + \sum_{\mu} m_\mu b_\mu - \sum_{iq} f_i^q \log f_i^q - \sum_{\mu} [m_\mu \log m_\mu + (1 - m_\mu) \log(1 - m_\mu)] \\ &\quad + \beta \sum_{iq\mu} f_i^q w_{i\mu}^q m_\mu + \frac{\beta^2}{2} \sum_{\mu} (m_\mu - m_\mu^2) \left[\sum_{iq} (w_{i\mu}^q)^2 f_i^q - \sum_i \left(\sum_q w_{i\mu}^q f_i^q \right)^2 \right]. \end{aligned} \quad (13)$$

The stationary conditions (11) yield a pair of self-consistent equations that are known as the Thouless-Anderson-Palmer (TAP) equations [23], and they are the equivalent of the ones already obtained for RBM with different input types in Refs. [10,11],

$$\begin{aligned} m_\mu &= \text{sigmoid} \left\{ b_\mu + \sum_{iq} f_i^q w_{i\mu}^q - \left(m_\mu - \frac{1}{2} \right) \left[\sum_{iq} (w_{i\mu}^q)^2 f_i^q - \sum_i \left(\sum_q w_{i\mu}^q f_i^q \right)^2 \right] \right\}, \\ f_i^q &= \text{softmax} \left\{ a_i^q + \sum_{\mu} m_\mu w_{i\mu}^q + \frac{1}{2} \sum_{\mu} (m_\mu - m_\mu^2) \left[(w_{i\mu}^q)^2 - 2w_{i\mu}^q \sum_p f_i^p w_{i\mu}^p \right] \right\}. \end{aligned} \quad (14)$$

The approximated free energy (13) has many local minima, unlike the full Gibbs free energy of Eq. (8), which is convex. Although this fact could be considered a pure artefact of high-temperature expansion, in practice these various metastable states can be associated with relative maxima in the probability distribution [17]. This fact can be explicitly verified by direct visual inspection of a dimension reduction of the

generated data and the fixed points, at least at early steps in the learning history when the generated data are highly clustered. Moreover, one can also verify that this multiplication of local minima does not vanish when we include higher orders of β in the expansion. In particular, we have systematically compared the fixed points obtained with the third order of the Plefka expansion [10] with those obtained from the TAP equations in




FIG. 3. (a) Representation of the TAP equations’ dynamics projected along the two main components of the PCA. The colored circles represent the initial conditions (magnetizations) of the TAP equations associated with the data points of the SED data set. The colored lines describe the trajectories that the points follow when iterating the TAP equations, while the stars represent the final fixed points. The colors of the circles and lines match those of the corresponding fixed points. (b) Pictorial representation of the (minus) LL’s minimization trajectory during the learning on the SED data set. The scatter plots represent the data generated by the model at different ages projected along the first two components of the training data set’s PCA. The colored stars correspond to the fixed points of the TAP equations initialized with the data, and the data points’ color matches the color of the corresponding fixed point. Notice how, as the training advances, the generated data set becomes more and more structured, and the number of TAP fixed points increases accordingly. (c) ~ 500 samples from the SED training data set projected on the first two principal components of the PCA. Notice how the picture resembles the projection of the last generated data set on panel (b).

the binary RBM case for some specific data sets, and we have not found any appreciable difference that would justify the implementation of higher orders than the second in the free-energy expansion.

We want also to stress that there are other methods that allow approximating the free energy of the model through high-temperature expansions to higher orders than the second, such as belief propagation (BP) [24]. However, these methods suffer from the same main limitation of the Plefka expansion; for instance, the BP algorithm assumes the weights of the network to be uncorrelated from each other, additionally approximating the factor graph of the model as a tree. We know that this assumption is only true at the early stages of the learning trajectory when the weights are small [5,16].

One could in principle use these TAP fixed points as a way to speed up the training procedure, as suggested in previous works [10,11]. While this procedure is very effective for training models with good generative performance, there is no guarantee that the Boltzmann distribution of these models is an accurate representation of the data set [15]. In fact, for this work, we investigated this training procedure and obtained very poor results when we tried to use the RBMs trained with the TAP equations for clustering objectives. We will discuss these results later.

B. Clustering data using the TAP fixed points

Given the model’s parameters at a certain age, we can use the data set as initial conditions by mapping each data point $\mathbf{v}^{(n)}$ to a pair of magnetization sets ($\{f_i^{q,(n)}(t=0)\}$, $\{m_\mu^{(n)}(t=0)\}$) using

$$\begin{aligned} m_\mu^{(n)}(t=0) &= \text{sigmoid}\left(b_\mu + \sum_{iq} \delta_{v_i^{(n)},q} w_{i\mu}^q\right), \\ f_i^{q,(n)}(t=0) &= \text{softmax}\left(a_i^q + \sum_{\mu} m_\mu^{(n)}(t=0) w_{i\mu}^q\right) \end{aligned} \quad (15)$$

and then iterate Eqs. (14) until convergence. In this way, the data belonging to the same basin of attraction of the free energy will collapse into the same point of the magnetization space, corresponding to one of the possibly many local minima, as shown in Fig. 3(a). We can then use these multiple basins of attraction to cluster similar data together. We just need a practical way to automatically collect all the data in the same basin (or metastable state).

To this end, we can use the DBSCAN algorithm [25] to categorize the data by the fixed point at which it ends after propagating the TAP equations. Now, if we do this at different stages of learning, such a method allows us to partition the data set according to the features that the model has learned at that point (age) of training. This method is quite robust in choosing the *eps* parameter of DBSCAN, because the TAP dynamics makes the points collapse into very small regions of magnetization space that can be easily delineated by the algorithm without ambiguity. However, a relatively large value of $eps = 1$ was chosen because we sometimes want to aggregate fixed points that almost overlap but are still distinct, preferring larger categories than having extremely similar fixed points. In Fig. 3(b), we sketch how the number of fixed points and the features recognized by the model evolve during the learning. Data points that converge to the same fixed point (represented by a star) are colored accordingly. This construction allows us to visualize how the free-energy landscape changes during the training, displaying a growing number of local minima. We note that this procedure is very similar to the annealing performed on the Gaussian mixture model [26,27].

We remark that the fixed points of the TAP equations are not guaranteed to be local minima, but they might be saddle points or local maxima as well. In principle, this could be verified by computing the determinant Hessian matrix of the TAP free energy (13) at the fixed points. Yet, in practice, this check is revealed to be quite problematic since one should compute an $(N_v \cdot N_q + N_h) \times (N_v \cdot N_q + N_h)$ dimension

matrix for each one of the fixed points that are identified. Instead, we qualitatively verified that the fixed points in the early phase of learning always correspond to the local minima of free energy by visually verifying that they fall in the middle of the denser regions, i.e., identifying the modes of the probability distribution. However, in the oldest machines, we cannot guarantee that some of the fixed points may be spurious solutions. To minimize this problem, we will always take a “backward” approach to the tree construction, as will be explained in the next section.

V. GENERATING RELATIONAL TREES

In this section, we discuss how to use the aforementioned construction to produce a hierarchical description of the data set. The method we propose consists of considering multiple classifications performed using several models saved at different stages of the training (i.e., *ages*). Assuming that we saved the models corresponding to the ordered set of ages $\{t_{\text{age}}^k\}_{k=1, \dots, N_T}$, with $t_{\text{age}}^{k+1} > t_{\text{age}}^k$, we start by taking the oldest RBM at $t_{\text{age}}^{N_T}$ and we iterate the TAP equations until convergence using the dataset as initial conditions. This procedure yields a number of fixed points that is smaller than the number of data points since some data points will have collapsed into the same free-energy minimum. Then, we pick the model at $t_{\text{age}}^{N_T-1}$ and we iterate again the TAP equations, this time using as initial conditions the fixed points found at the previous step. In fact, some minima that were distinct for the oldest RBM will be degenerate for the younger one, finding themselves in the same basin of attraction of a new (larger) minimum. We can repeat this procedure backwards in time for all the selected ages until the youngest RBM, for which all the data collapse into a unique global minimum. If we keep track of the fixed points’ merging during the procedure, we can construct a tree starting from the leaves (i.e., splits made by old models) and proceeding down to the root. We depicted this procedure in Fig. 4. This procedure also allows us to efficiently avoid the spurious fixed points (the saddle points or maxima of free energy): we follow the merging of fixed points obtained in old machines into the broader basins of attraction of the fixed points of the younger machines. This procedure ensures that the tree construction will rapidly rely only on the true minima of the free energy. Yet, one could also imagine constructing a tree based on the fixed points that arise when propagating the TAP equations from all points in the data set at each age of the machine. We have found empirically that trees constructed in this way are qualitatively and quantitatively worse than those presented here (see Fig. 17 of Appendix C). We believe that a reason for this is precisely that not all fixed points are free-energy minima. Another possible explanation for this is that, in this second case, classifications performed at different ages are independent of each other. In the method we described, instead, the classification at a given age is informed by those performed at later ages. In other words, information is propagated backwards in (training) time, from the more experienced models to the rawest ones.

At this point, we have to state a criterion for selecting the ages of the RBM that we want to save during the training and that will participate in the tree construction. Ideally, one would wish to find a relation between the geometry of the

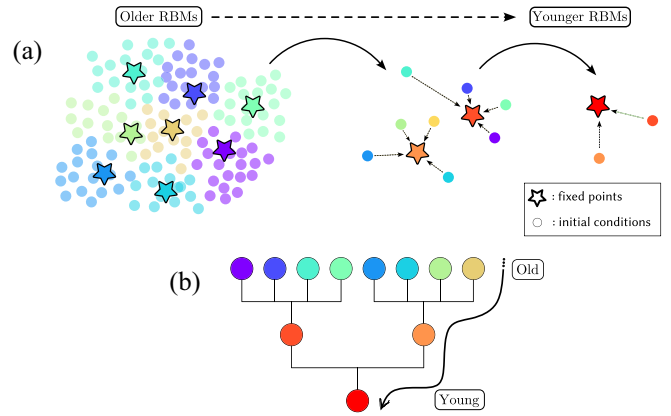


FIG. 4. (a) Sketch of the process of merging the fixed points used to build the relational tree. We start with an old RBM and let the TAP dynamics collapse the data into the fixed points of the self-consistent Eqs. (14). Then we repeat the procedure for younger and younger machines, each time propagating the TAP equations starting from the previously found fixed points, until we reach the point where all the data have collapsed into just one global minimum. (b) By following the merging of the fixed points, we obtain the relational tree of the data set.

free-energy landscape (e.g., the number of fixed points that are present) and some observable that can be monitored during the training, such as the spectra of the weight matrix, to decide optimally and automatically when to save the state of the RBM. The first splits are easily traceable by monitoring the first eigenvalues of the weight matrix. Yet, when the RBM starts to have an increasing number of “expressed” modes, we can no longer link one mode to a given split (see Fig. 11 of Appendix C). For this reason, we simply saved a certain number of models (e.g., ~ 500) equally spaced along the training history. When constructing the tree, we added a new layer only when the number of fixed points found with the TAP equations decreased (recall that we move backwards in time when using the TAP equations) from the previously added layer.

Applying this procedure to the artificial data set SED, we create the tree we showed as teasing in Fig. 2(f). The leaves of the graph (real data points) were colored according to the labels of each entry, as shown in the genealogical tree in Fig. 2(b). We consider two possible levels of relationships: Ancestral families are shown in five colors (blue, green, pink, yellow, and red) and next-generation entries are shown in different shades of the same color family. We see that with this tree we can perfectly delineate all five families based on the main branches of the constructed tree, and the same is true for the subfamilies initiated by the grandchildren (shown in the outer crowns of the diagram). Since the nodes of the tree correspond to points in the magnetization space, we can visualize the evolution of the learned features during learning, as shown in Fig. 5(a). Looking at the innermost nodes of the tree, we can see that the RBM first learns the characteristic patterns of the first generation of sequences by recognizing the mutation sites, and then gradually distinguishes the second-generation mutations, and so on. Interestingly, the magnetizations of the various fixed points tell us exactly where the mutations (which led to the creation of each family and subfamily) occurred

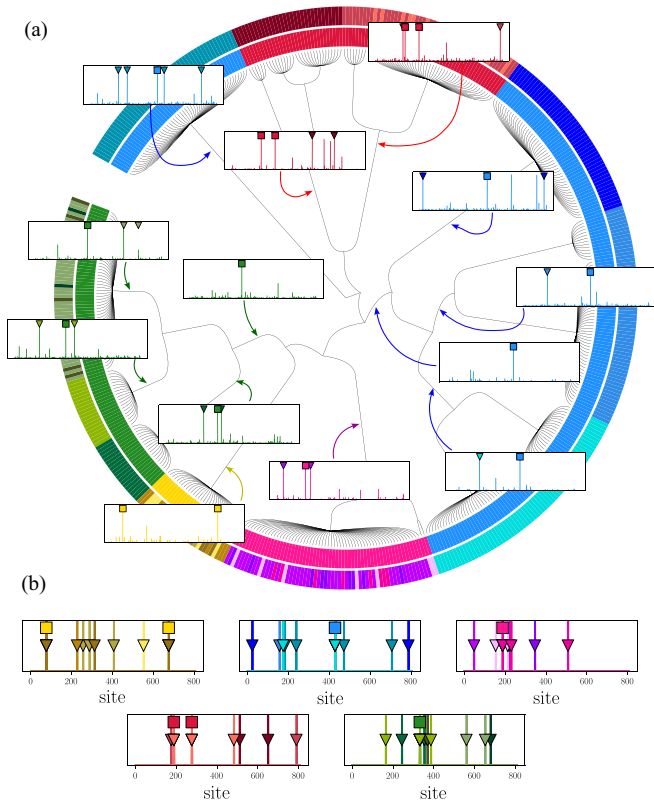


FIG. 5. Representation of the features learned by the RBM during training on the SED data set. (a) The same tree as in Fig. 2(f), truncated to 15 layers. The plots in the insets represent the comparison between the magnetization profiles (solid colored lines) associated with some internal nodes of the tree and the “patterns” of the first and second generations of the genealogical tree in Fig. 2(b). The x -axis of the boxes denotes the index site along the sequence, whereas the y -axis represents the value of the magnetization. The colored squares indicate the locations where mutations occurred in the first generation of progeny (passing from the ancestral mother sequence with all 0s to a sequence of 0s plus a few 1s), and the triangles indicate the mutations in the second generations corresponding to this subset (secondary coloring labels). (b) Representation of the location of all first- and second-generation mutations. Each box represents one of the first five families. The squares mark the mutations in the first generation, while the triangles and the vertical lines highlight the mutation sites in the second generation (in both cases compared to the ancestral mother sequence). Each color represents a different second-generation family [retaining the color scheme of Fig. 2(b)].

during the evolutionary process [see Fig. 5(b)]. Furthermore, from this diagram, it is easy to understand why the main blue family is already divided into two main groups at the root of the tree. The magnetization profile of the blue family in the upper left part of Fig. 5, indeed, shows a burst of mutations that occurred very early in the evolutionary process and caused this subfamily to diverge from the rest of the sequences.

VI. RESULTS

So far, we have shown the performance of our method only on a synthetic data set whose hierarchical relational structure

was given by design. Now, we apply our method to three different real-world data sets with different categories of data points. First, we analyze a simple image data set where all labels are known and where features can be visualized directly. In particular, we analyze a relational structure of the MNIST data set [28], which contains images of handwritten digits whose label is the actual digit. Second, we analyze the human genome data set (HGD) [29], which captures the genetic mutations of 5008 individuals from around the world. This data set is even more interesting because the sequences contain two levels of annotation (continental origin and population group). Third, we conclude with the last and most challenging data set: amino acid sequences of a homologous family of proteins (the CPF protein) with a large functional diversity. The biological function of these sequences is known for only a very small fraction, namely those that have been studied experimentally. The challenge, then, is to assign a probable biological function to each of the unclassified sequences based on limited experimental knowledge. For a detailed description of the data sets considered, see Appendix A.

We show the projection of all three of these data sets along the first two principal components of the PCA in Fig. 6. In particular, we point out the different degree of clusterization displayed in each of these three data sets: while the distribution of the MNIST data points is quite broad and connected, the HGD and the CPF protein are rather clustered, which is known to sensibly damage the quality of the training process due to sampling convergence problems [30].

A. MNIST data set

The MNIST data set consists of binary images (pixels can be either 0 or 1) of handwritten digits. Each example is associated with a numeric label—from 0 to 9—indicating the digit represented in the image. We show in Fig. 7 the relational tree obtained by our method for MNIST. As we can see, the RBM is mostly able to separate the different digits, although there are some cases in which two categories are mixed together (e.g., the digit 4 with the 9). Dealing with pictures has the advantage that we can visualize the features learned by the RBM during learning and try to explain them qualitatively. The first few layers of the tree obtained using a subset of the MNIST data set are shown in Fig. 8, where we also highlighted the images of some data that were misclassified. Since the internal nodes of the tree correspond to actual points in the magnetization space, we can visualize the evolution of the features learned from the RBM as they unfold along the tree and see where and when each cluster detail appears. For example, the first split separates digit 0 from the rest. As training time (or tree depth) increases, the 0 appears to remain completely isolated while the rest of the digits are gradually resolved as the RBM learns finer details. We can also observe the proximity of some digits; in addition to the 4 and the 9, we see that the 7 is also quite close to the 9. In principle, it would be possible to further analyze the obtained results and find interesting features that may have been unknown before the analysis.

B. Human genome data set

In the HGD, each data point is a binary vector corresponding to a human individual, where a 1 is placed in

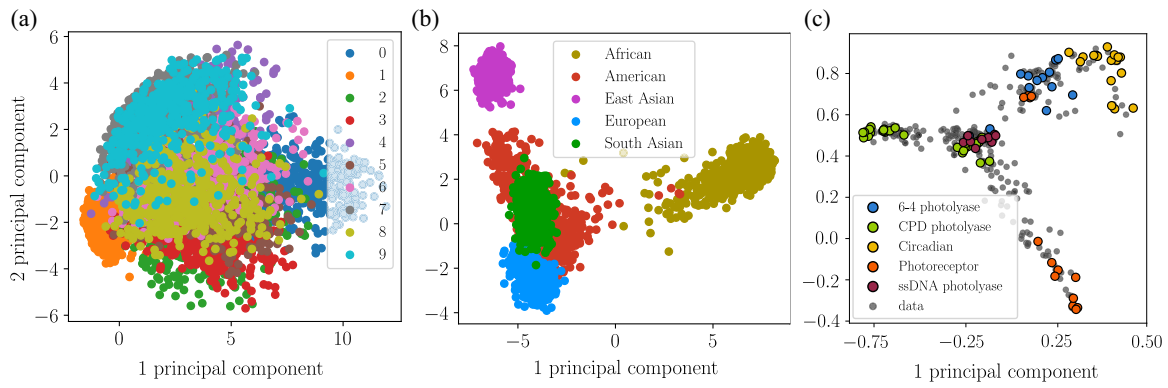


FIG. 6. Projections of the data sets along the first two principal components of the PCA. The colors correspond to the known labels of the data points. (a) MNIST dataset. (b) Human Genome data set. (c) CPF protein. All the data are shown in black and the colored dots correspond to the experimentally annotated sequences.

correspondence to the sites that present a mutated gene with respect to a reference sequence, and a 0 is placed where the gene is maintained unaltered. For this data set, we have two sets of labels since each individual is associated with a certain continental area (superpopulation) and population. This two-level characterization of the data set allows us to test the discriminative power of the RBM at different resolutions. The tree obtained for this data set is shown in Fig. 9. Over-

all, the separation of the major labels, corresponding to the continental area, is performed quite well, although sometimes the tree blends together European and American individuals. On the other hand, the substructure corresponding to the different populations is almost never detected by the model at this stage of the training, apart from the isolation of Peruvian and Finnish people. By training an RBM separately on each continental area, we verified that, in some cases, it would be




FIG. 7. Complete tree generated on 500 samples from the MNIST data set. The external coloured tags represent the true labels, while the colored regions inside the tree have been manually assigned for the sake of visualization.

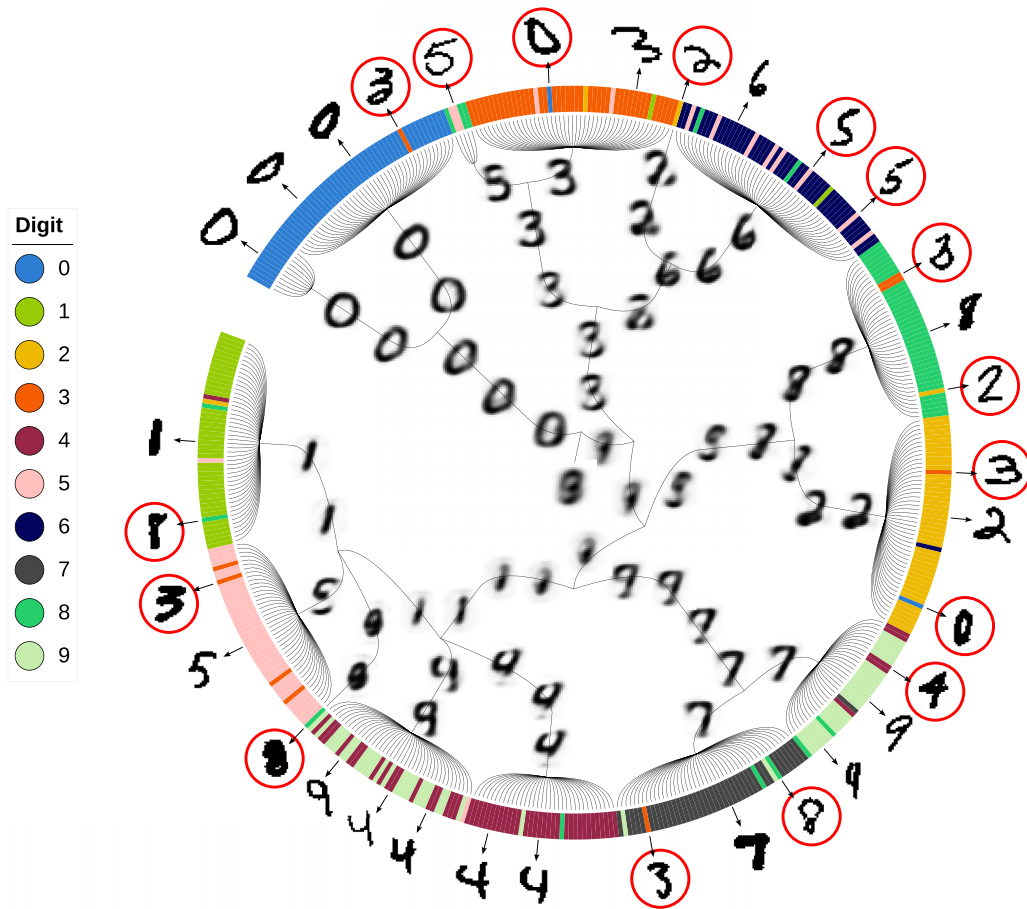


FIG. 8. First 10 layers of the tree obtained on 500 samples from the MNIST data set. The colored tags in the external ring correspond to the true labels of the data. Inside the tree, we superimposed the images of the RBM’s fixed points corresponding to the nodes along the branches. The intensity of the pixels represents the probability that the machine assigns to that pixel to be turned on. We also report the images from the data set corresponding to some of the leaves. In particular, we put a red circle around the digits that have been wrongly classified by the algorithm in order to highlight the configurations that fooled the machine.

possible to discriminate a subcategory from the others (see, for instance, Fig. 12 of Appendix C), while in other cases the data set in itself and the RBM cannot resolve the different populations, therefore appearing indistinguishable among each other. We should keep in mind here that, when applying the RBM on populations, we are dealing with a reduced number of samples in the training set and that the impossibility of separating the different populations might be a consequence of the lack of a statistical signal. We remark that, since only a subpart of the genome has been used here, we do not claim at all that these populations would be indistinguishable in general.

C. CPF protein data set

In the two cases previously discussed, we had the label information for all the data points. For practical purposes, a more interesting scenario is one in which only a fraction of the data has been annotated. To this end, we considered a biologically relevant data set consisting of homologous sequences representing the CPF (Cryptochrome/Photolyase) protein. For this data set, only 86 out of 397 sequences have been experimentally classified, based on their biological function, into

five different categories [31–34]. Since the amino acid chains are not the same length, the sequences have to be previously aligned in the form of a multisequence alignment (MSA) matrix before being fed to the RBM. Each site of the sequences can take one over $N_q = 21$ possible values (20 amino acids + a gap symbol), hence making the data set particularly suitable to be studied with the Potts-RBM model. Using our tree construction, we may therefore infer a classification for the sequences not previously annotated by looking at the label of the annotated sequences that fall in the same branch.

The algorithm’s results on the CPF protein data set are shown in Fig. 10. Coherently with the phylogenetic tree (see Fig. 14 of Appendix C), the sequences classified as CPD photolyase are found in two distinct main branches of the tree, and a subtree containing photoreceptor sequences is found within the (6-4) photolyase main branch. It is indeed a known fact in biology that the concept of a biological function is related, but not identifiable, to the phylogeny of the protein family. In other words, proteins that diverged in their amino acid sequences may still carry the same function. For instance, in Ref. [35] the authors propose a hierarchical classification method, called ProfileView, that is able to correctly identify 11 of the known subgroups of the CPF protein beyond the




FIG. 10. Tree generated from the CPF protein family. The external colored tags, when available, represent the experimental classification [31–34], while the color of the leaves represents the classification into 11 (plus one unidentified) subgroups classified by ProfileView in Ref. [35]. Since several ProfileView categories may correspond to the same experimental classification, to avoid confusion we kept the colors of the two classifications independent of each other.

age, but only those that find some earlier fixed points in their basin of attraction.

Given that our method is based upon the mean-field approximation, which assumes the model to be in the early stage of the learning, one may allegedly wonder how far we can go with the training and still obtain a reliable approximation of the model’s free energy. Since we do not have access to the true free energy of the model, we cannot directly assess the goodness of our second-order approximation. Even so, we can have a clue on the validity of the approximation by comparing the fixed points of the TAP equations with those obtained with the self-consistent equations corresponding to the Gibbs free energy (8) truncated at the first order in β , which we call *naive* mean-field (nMF). The point of the training at which the two sets of solutions start to diverge can be considered a safe upper bound below which we are guaranteed that our method is working properly. Performing this check for the various data sets reveals that, at the point in which the main features of interest (corresponding to the labeled categories) are separated inside the tree, the fixed points of the first- and second-order approximations of the free energy are still very similar. This is particularly true when the data set is very clustered. Then, when the two sets of solutions start to change significantly,

we observe that the number of nMF fixed points grows much more than that of the TAP solutions, as we can see in Fig. 13 of Appendix C.

Another interesting question concerns the variability of the trees obtained with repeated training of the same RBM model. Since training the RBM is a stochastic process, we might expect the tree structure to change slightly with different training runs. We verified that there is indeed some degree of variability, as we can see in Fig. 15 in Appendix C for the MNIST data set. However, we also note that some key features of the tree are preserved across experiments, such as the easy separation of digit 0 and the difficult distinction between digits 4 and 9. The same behavior was observed for other data sets, such as the CPF protein data set. In this case, we observe that the position of the branches within the tree may vary with repeated experiments, but the hierarchical organization of the subtrees is preserved.

It is worth making a final comment on the possibility of using the training trajectory of the RBM to construct the tree. Since training the RBM can be viewed as an effective annealing of the model, one might want to try to construct the tree by simply using the most trained model and gradually increasing the temperature (i.e., decreasing β) while iterating

the equations from TAP repeatedly until convergence. Although this method gave quite similar results to our original version on the HGD and SED data sets, it proved very inadequate when applied to the MNIST data set (see Fig. 15 in Appendix C). A good reason for using the learning trajectory instead of a trajectory in temperature is that, in the first case, the system undergoes continuous phase transitions [5] and the evolution of the probability distribution of the model is smooth. In the second case, on the other hand, the system undergoes first-order phase transitions in temperature, which determine very different free-energy landscapes at different temperatures. This could explain why the use of the learning trajectory leads to better results in the hierarchical classification.

Finally, one can ask what the result of inference would be if one were to use an RBM trained with a different method from the one we used. For example, one can train the RBM by computing the negative term of the gradient using the TAP equations (14) as proposed in Refs. [10,11]. We have explored this possibility, but we have found empirically that models trained in this way tend to lose the hierarchical properties as the fixed points emerge during learning. We show in Fig. 18 of Appendix C an example of a tree obtained with an RBM trained with the TAP equations. Similarly, trees obtained with nonequilibrium machines, i.e., RBMs trained with nonpersistent, nonconvergent chains [7,15], e.g., with chains initialized with random distributions or the data set, show some degree of hierarchization in clustering, but they are not specific enough to correctly delineate the different families.

VIII. CONCLUSIONS

In this work, we derived the TAP equations for the Potts RBM defined by the Hamiltonian (2), and we used this mean-field description of the model to study the evolution of the RBM free energy during training. We then developed a method to infer the hierarchical structures of the data sets from the training trajectory. Using labeled data, we showed that our method is able to generate meaningful trees when applied to four different data sets. In addition, we were able to examine the internal nodes of the tree generated on the SED and MNIST data sets to see how the internal representations of the model change during learning. Although the tree construction is inherently stochastic, we demonstrated that some relevant features are conserved across different realizations. Finally, we tested our method on a protein data set and showed that it can provide useful functional information by only capturing the statistics of the input MSA in a purely unsupervised manner, i.e., without any prior knowledge about the data. These results suggest potential promising applications in biology to detect functional signals in homologous protein data sets.

ACKNOWLEDGMENTS

We would like to thank Giovanni Catania for his help in the early stages of this project in deriving the TAP equations for the Potts RBMs, Alessandra Carbone for useful discussions about the applications of our methods to protein sequences, and Edoardo Sarti for discussions and technical help. We acknowledge financial support by the Comunidad de Madrid and the Complutense University of Madrid (Spain) through the

Atracción de Talento programs (Refs. 2019-T1/TIC-13298 for A.D. and L.R. and 2019-T1/TIC-12776 for B.S.), the Banco Santander and the UCM (Grant No. PR44/21-29937), and Ministerio de Economía y Competitividad, Agencia Estatal de Investigación, and Fondo Europeo de Desarrollo Regional (FEDER) (Spain and European Union) through Grant No. PID2021-125506NA-I00.

APPENDIX A: STUDIED DATA SETS AND TECHNICAL DETAILS

We tested our algorithm on four different data sets:

(i) The MNIST data set [28] consists of gray-scale images representing handwritten digits, from 0 to 9, of dimension 28×28 pixels. The data set is thoroughly annotated, namely, we have a label for each sample indicating the represented digit. For our experiments, we extracted 10 000 samples from the training set and we applied the tree algorithm to a subset of 500 data samples. The images have been flattened and converted in binary format by setting to 1 the pixels with a normalized value above 0.3 and to 0 otherwise. In this way, the data set we used was made of 784-dimensional binary vectors.

(ii) The human genome data set (HGD) [29] represents human genetic variations of a population of 5008 individuals sampled from 26 populations in Africa, east Asia, south Asia, Europe, and the Americas. Each sample is a sequence of 805 binary variables, $v_i \in \{0, 1\}$, representing the alteration or not of a gene with respect to a reference genetic sequence. For these data, we have a complete two-level hierarchical labeling. At a higher level, the sequences are classified based on the continental origin of the individuals, while at a finer level, we also have a label for the particular population from which the data were sequenced. This allows us to test the discriminative capabilities of the RBM at different resolutions. We trained the RBM on 4508 samples.

(iii) The synthetic evolutionary data (SED) data set has been artificially created by us with the aim of simulating a controlled evolutionary process fully driven by random mutations. We started with an ancestor sequence of length $N_v = 805$ and we produced five heirs by flipping at least one site of the sequence. Then, we associated a label with each one of the heirs, and we iterated the following evolutionary process: given a sequence, we generated a random number between 1 and 5, which corresponds to the number of children. Each child is obtained by flipping each site of the mother sequence with probability $p = 1/N_v$, and it is added to the data set if the sequence was not already present. Each child inherits the label of the mother sequence. We repeated this process at each generation until we obtained a data set of 4508 sequences.

(iv) Finally, we tested the algorithm on a data set of homologous protein sequences. The data set we consider is made of 397 amino acid sequences representing the cryptochrome/photolyase (CPF) protein. Proteins belonging to the CPF family share the same fold, yet they perform several different functions. Cryptochromes are divided into photoreceptors that activate specific pathways in the presence of light and light-independent circadian clock transcriptional regulators, while photolyases are light-activated enzymes encharged with repairing UV-damaged DNA. In particular, 86 of the available sequences have been experimentally

TABLE I. Technical specifics of the RBMs used for generating the trees reported in the main text. All machines have been trained in PCD mode [7]. Only for the tree of Fig. 2(f) has a binary RBM been used in place of the Potts version described in the main text.

Tree	Data-set	epochs	Minibatch total gradient		Learning		
			size	updates	N_{Gibbs}	rate	N_h
Fig. 2(f)	SED	10 000	512	9×10^4	100	10^{-2}	512
Figs. 8,7	MNIST	10 000	512	2×10^5	100	10^{-3}	512
Fig. 9	HGD	10 000	512	9×10^4	100	10^{-3}	512
Fig. 10	CPF	10 000	128	3×10^4	100	10^{-3}	512

classified into five categories according to their biological function [31–34]: circadian, photoreceptor, (6-4) photolyase, CPD photolyase, and ssDNA photolyase. Finally, we validated our model by also considering the finer classification in 11 categories determined in [4].

To be processed by the RBM, the sequences must be previously aligned in a matrix called multisequence alignment (MSA). The alignment is performed by inserting a special symbol, the *gap* ‘–’, a proper number of times within adjacent sites of the sequences. Each site can therefore assume 1 over 21 possible values (20 amino acids + the gap). After the alignment, the length of the obtained sequences was $N_v = 7403$. To perform the MSA we used the software MAFFT with a large ep parameter of 1 to penalize long alignments.

Concerning the weights used for computing the weighted average in Eq. (5) in this data set, we first had to define a set of clusters by gathering all sequences with sequence identity greater than 0.8, and then we applied the same criterion described above. The clustering has been done using the software MMseqs2.

All the trees shown in the article have been generated using the iTOL software.

1. Technical details of the RBM models used

In Table I we report the hyperparameters of the RBMs used for generating the trees reported in the main article.

2. Projecting the data set along the weight matrix principal directions

To project a data set made of categorical variables along the first principal directions of the weight matrix, we first have to transpose the tensor \mathbf{w} along the first two dimensions and then reshape it from an $N_v \times N_q \times N_h$ tensor to an $(N_q \cdot N_v) \times N_h$ matrix:

$$\mathbf{w} \in \mathbb{R}^{N_v \times N_q \times N_h} \xrightarrow{\text{transpose}} \hat{\mathbf{w}} \in \mathbb{R}^{N_q \times N_v \times N_h} \xrightarrow{\text{reshape}} \tilde{\mathbf{w}} \in \mathbb{R}^{(N_q \cdot N_v) \times N_h}.$$

Then, we consider the singular value decomposition (SVD) of $\tilde{\mathbf{w}}$,

$$\tilde{\mathbf{w}} = USV^T, \quad (\text{A1})$$

where $U \in \mathbb{R}^{(N_q \cdot N_v) \times (N_q \cdot N_v)}$, S is an $(N_q \cdot N_v) \times N_h$ matrix with the singular values of \mathbf{w} in the diagonal, and $V \in \mathbb{R}^{N_h \times N_h}$. The

next step is to compute the one-hot representation of the data set, which consists in associating with each categorical variable q an N_q -dimensional vector that has value 1 in position q and 0 otherwise:

$$\mathcal{D} \in \mathbb{R}^{M \times N_v} \xrightarrow{\text{one-hot}} \hat{\mathcal{D}} \in \mathbb{R}^{M \times N_v \times N_q} \xrightarrow{\text{reshape}} \tilde{\mathcal{D}} \in \mathbb{R}^{M \times (N_v \cdot N_q)}.$$

Finally, the projected data set is given by

$$\mathcal{D}_{\text{proj}} = \tilde{\mathcal{D}}U \in \mathbb{R}^{M \times (N_v \cdot N_q)}, \quad (\text{A2})$$

where each row is a vector associated with a data sample whose components correspond to the projections of the one-hot representation of the data point along the principal directions of the weight matrix.

APPENDIX B: TAP EQUATIONS FOR THE POTTS RBM

We start by considering a Hamiltonian of the kind

$$-\beta E_\beta(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta}) = \sum_{iq} a_i^q \delta_{v_i, q} + \sum_{\mu} b_{\mu} h_{\mu} + \beta \sum_{iq\mu} \delta_{v_i, q} w_{i\mu}^q h_{\mu}, \quad (\text{B1})$$

where the inverse-temperature parameter, β , regulates the interactions among the variables. Our goal is to study the metastable minima of the free energy in the high-temperature regime $\beta \rightarrow 0$, which is equivalent to considering the limit in which the eigenvalues of the weight matrix are vanishing.

Let us introduce two sets of external fields, $\{\phi_j^p\}$ and $\{\psi_{\mu}\}$, which act, respectively, on the visible and hidden variables and depend on the inverse-temperature β . We denote the total external field vector as $\mathbf{B} = \{B_i\} = (\{\phi_j^p\}, \{\psi_{\mu}\})$, with $i \in \{1, \dots, N_v \cdot N_q + N_h\}$. The free energy of this extended system is

$$-\beta F_\beta(\mathbf{B}) = \log \sum_{\mathbf{v}, \mathbf{h}} \exp \left[-\beta E_\beta(\mathbf{v}, \mathbf{h}; \boldsymbol{\theta}) + \sum_{iq} \phi_i^q(\beta) \delta_{v_i, q} + \sum_{\mu} \psi_{\mu}(\beta) h_{\mu} \right], \quad (\text{B2})$$

where we omitted the free energy’s dependence from the model’s parameters $\boldsymbol{\theta}$. The local magnetizations in the presence of these fields, $\mathbf{M} = \{M_i\} = (\{f_j^p\}, \{m_{\mu}\})$, are

$$f_i^q = \langle \delta_{v_i, q} \rangle_{\mathbf{B}} = -\beta \frac{\partial}{\partial \phi_i^q} F_\beta(\mathbf{B}) \quad \text{and} \quad m_{\mu} = \langle h_{\mu} \rangle_{\mathbf{B}} = -\beta \frac{\partial}{\partial \psi_{\mu}} F_\beta(\mathbf{B}). \quad (\text{B3})$$

Notice that the free energy in Eq. (B2) is a concave function of the external fields. To see this, we can represent a state of the system with the vector $\mathbf{x} = (\boldsymbol{\delta}^{\text{oh}}, \mathbf{h}) \in \mathbb{R}^{N_v \cdot N_q + N_h}$, where $\boldsymbol{\delta}^{\text{oh}} \in \mathbb{R}^{N_v \cdot N_q}$ is the one-hot representation of the Potts states. Then, the Hessian of the free energy with respect to the fields is

$$\frac{\partial^2 F_\beta(\mathbf{B})}{\partial B_j \partial B_i} = -\frac{1}{\beta} \frac{\partial M_i}{\partial B_j} = -\frac{1}{\beta} \langle (x_i - M_i)(x_j - M_j) \rangle_{\mathbf{B}} \quad \text{for } i, j \in \{1, \dots, N_v \cdot N_q + N_h\}, \quad (\text{B4})$$

which is seminegative defined.

We can describe the system in terms of the magnetizations by considering the Gibbs free energy, which is obtained by taking the Legendre transform of (B2):

$$\begin{aligned} -\beta G_\beta(\mathbf{M}) &= -\beta \max_{\mathbf{B}} \left[F_\beta(\mathbf{B}) + \frac{1}{\beta} \sum_{iq} \phi_i^q(\beta) f_i^q + \frac{1}{\beta} \sum_{\mu} \psi_\mu(\beta) m_\mu \right] \\ &= \min_{\mathbf{B}} \left[\log \sum_{\mathbf{v}, \mathbf{h}} \exp \left(\sum_{iq} a_i^q \delta_{v_i, q} + \sum_{\mu} b_\mu h_\mu + \beta \sum_{iq\mu} \delta_{v_i, q} w_{i\mu}^q h_\mu + \sum_{iq} \phi_i^q(\beta) (\delta_{v_i, q} - f_i^q) + \sum_{\mu} \psi_\mu(\beta) (h_\mu - m_\mu) \right) \right]; \end{aligned} \quad (\text{B5})$$

in this way the fields $\mathbf{B} = \mathbf{B}(\mathbf{M})$ satisfy Eqs. (B3), namely they enforce the magnetizations \mathbf{M} . The solution $\mathbf{B}(\mathbf{M})$ must satisfy

$$\begin{aligned} B_i &= \beta \frac{\partial}{\partial M_i} G_\beta(\mathbf{M}), \quad \frac{\partial^2 G_\beta(\mathbf{M})}{\partial M_j \partial M_i} = \frac{1}{\beta} \frac{\partial}{\partial M_j} B_i(\mathbf{M}) \\ &= \frac{1}{\beta} \left(\frac{\partial M_j}{\partial B_i} \right)^{-1} = - \left(\frac{\partial^2 F_\beta(\mathbf{B})}{\partial B_i \partial B_j} \right)^{-1}, \end{aligned} \quad (\text{B6})$$

meaning that the Gibbs free energy is a convex function of the magnetizations. The free energy is the inverse Legendre transform of $G_\beta(\mathbf{M})$,

$$\begin{aligned} -\beta F_\beta(\mathbf{B}) &= -\beta \min_{\mathbf{M}} \left[G_\beta(\mathbf{M}) - \frac{1}{\beta} \sum_{iq} \phi_i^q(\beta) f_i^q \right. \\ &\quad \left. - \frac{1}{\beta} \sum_{\mu} \psi_\mu(\beta) m_\mu \right], \end{aligned} \quad (\text{B7})$$

and the solution $\mathbf{M}(\mathbf{B})$ satisfies Eq. (B6) and it is the global minimum of the Gibbs free energy. From Eq. (B7) we can see that, by turning off the external fields \mathbf{B} , we retrieve the free energy of the original system at the minimum of the Gibbs

free energy:

$$A_\beta = -\frac{1}{\beta} \sum_{\mathbf{v}, \mathbf{h}} e^{-\beta E_\beta(\mathbf{v}, \mathbf{h}, \boldsymbol{\theta})} = F_\beta(\mathbf{B} = \mathbf{0}) = \min_{\mathbf{M}} [G_\beta(\mathbf{M})]. \quad (\text{B8})$$

At this point, we have two problems. The first problem is that the Gibbs free energy is a convex function of the magnetizations, meaning that we cannot use it directly to identify the metastable minima of the free energy. Moreover, we are prevented from computing it in the first place because it would be as difficult to evaluate as the original free energy. The two issues can be overcome if we restrict ourselves to the high-temperature limit, corresponding to the early stage of the learning when the RBM has not yet learnt the modes of the data set. In this regime, we can approximate the Gibbs free energy with a Taylor expansion around $\beta = 0$. For ease of notation, let us define the Gibbs free entropy as $\Gamma_\beta(\mathbf{M}) = -\beta G_\beta(\mathbf{M})$. We want to compute the second-order expansion:

$$\begin{aligned} \Gamma_\beta^{(2)}(\mathbf{M}) &= \Gamma_0(\mathbf{M}) + \beta \left. \frac{\partial \Gamma_\beta(\mathbf{M})}{\partial \beta} \right|_{\beta=0} \\ &\quad + \frac{\beta^2}{2} \left. \frac{\partial^2 \Gamma_\beta(\mathbf{M})}{\partial \beta^2} \right|_{\beta=0}. \end{aligned} \quad (\text{B9})$$

The high-temperature approximation has the effect of breaking the degeneracy of the Gibbs free energy, hence allowing us to detect the metastable pure states. Moreover, the expansion (B9) can be directly computed, because at $\beta = 0$ the interactions among the variables disappear, and the probability measure factorizes. Given an observable O , we can define the high-temperature average as

$$\langle O \rangle = \frac{\sum_{\mathbf{v}, \mathbf{h}} O(\mathbf{v}, \mathbf{h}) \exp(\sum_{iq} [a_i^q + \phi_i^q(0)] \delta_{v_i, q} + \sum_{\mu} [b_\mu + \psi_\mu(0)] h_\mu)}{\sum_{\mathbf{v}, \mathbf{h}} \exp(\sum_{iq} [a_i^q + \phi_i^q(0)] \delta_{v_i, q} + \sum_{\mu} [b_\mu + \psi_\mu(0)] h_\mu)}. \quad (\text{B10})$$

In particular, we can compute the magnetizations in the high-temperature regime as

$$\begin{aligned} m_\mu &= \langle h_\mu \rangle = \frac{\sum_{h_\mu=0,1} h_\mu \exp(\sum_{\mu} h_\mu [\psi_\mu(0) + b_\mu])}{\sum_{h_\mu=0,1} \exp(\sum_{\mu} h_\mu [\psi_\mu(0) + b_\mu])} \\ &= \frac{1}{1 + e^{-b_\mu - \psi_\mu(0)}} = \text{sigmoid}[b_\mu + \psi_\mu(0)], \quad (\text{B11}) \\ f_i^q &= \langle \delta_{v_i, q} \rangle = \frac{\sum_{\mathbf{v}} \delta_{v_i, q} \exp(\sum_p \delta_{v_i, p} [\phi_i^p(0) + a_i^p])}{\sum_{\mathbf{v}} \exp(\sum_p \delta_{v_i, p} [\phi_i^p(0) + a_i^p])} \end{aligned}$$

$$= \frac{e^{a_i^q + \phi_i^q(0)}}{\sum_p e^{a_i^p + \phi_i^p(0)}} = \text{softmax}[a_i^q + \phi_i^q(0)]. \quad (\text{B12})$$

We now proceed by computing the various terms of Eq. (B9). For ease of notation, let us define

$$\mathcal{I} = \sum_{iq\mu} \delta_{v_i, q} w_{i\mu}^q h_\mu, \quad (\text{B13})$$

$$\mathcal{F}(\beta) = \sum_{iq} \phi_i^q(\beta) (\delta_{v_i, q} - f_i^q) + \sum_{\mu} \psi_\mu(\beta) (h_\mu - m_\mu), \quad (\text{B14})$$

so that the Gibbs free entropy can be written as

$$\Gamma_\beta(\mathbf{M}) = \min_{\mathbf{B}} \left[\log \sum_{v,h} \exp \left(\beta \mathcal{I} + \mathcal{F}(\beta) + \sum_{iq} a_i^q \delta_{v_i,q} + \sum_{\mu} b_{\mu} h_{\mu} \right) \right]. \quad (\text{B15})$$

1. Entropic term

The zero-order term corresponds to the mean-field entropy of the model, and it is given by

$$\begin{aligned} \Gamma_0(\mathbf{M}) = & - \sum_{iq} \phi_i^q(0) f_i^q - \sum_{\mu} \psi_{\mu}(0) m_{\mu} \\ & + \sum_i \log \left[\sum_v \exp(a_i^v + \phi_i^v(0)) \right] \\ & + \sum_{\mu} \log[1 + \exp(b_{\mu} + \psi_{\mu}(0))]. \end{aligned} \quad (\text{B16})$$

To make explicit the dependence on the magnetizations $\{m_{\mu}\}$ and $\{f_i^q\}$ into Eq. (B16), we need to invert Eqs. (B11) and (B12). The first equation leads straightforwardly to

$$\psi_{\mu}(0) = \log \left(\frac{m_{\mu}}{1 - m_{\mu}} \right) - b_{\mu}, \quad (\text{B17})$$

while from the second equation we get

$$\phi_i^q(0) = \log f_i^q - a_i^q + c_i, \quad (\text{B18})$$

where $c_i = \log[\sum_p \exp(a_i^p + \phi_i^p(0))]$. This last arbitrary constant comes from the fact that the Potts model displays a gauge invariance [3]. A convenient choice to fix it is the *lattice-gas gauge*, for which we impose

$$a_i^{N_q} = \phi_i^{N_q} = w_{i\mu}^{N_q} = 0. \quad (\text{B19})$$

In this way, we obtain the inverse relation of the softmax function (B12):

$$\phi_i^q(0) = \log \left(\frac{f_i^q}{f_i^{N_q}} \right) - a_i^q. \quad (\text{B20})$$

By inserting these relations into Eq. (B16) and remembering the normalization condition $\sum_q f_i^q = 1$, we obtain the familiar expression of the mean-field Shannon entropy of a set of binary and categorical variables plus two interaction terms between the magnetizations and the local fields of the model,

$$\begin{aligned} \Gamma_0(\mathbf{M}) = & \sum_{iq} f_i^q a_i^q + \sum_{\mu} m_{\mu} b_{\mu} - \sum_{iq} f_i^q \log f_i^q \\ & - \sum_{\mu} [m_{\mu} \log m_{\mu} + (1 - m_{\mu}) \log(1 - m_{\mu})]. \end{aligned} \quad (\text{B21})$$

2. First-order term

The first derivative is readily obtained as

$$\left. \frac{\partial \Gamma_{\beta}}{\partial \beta} \right|_{\beta=0} = \langle \mathcal{I} \rangle + \underbrace{\langle \partial_{\beta} \mathcal{F}(0) \rangle}_{=0} = \sum_{i\mu} f_i^q w_{i\mu}^q m_{\mu}, \quad (\text{B22})$$

which is merely the naive mean-field interaction energy.

3. Second-order term

Given an observable $O(\beta)$, the derivative of its average at inverse-temperature β can be computed as

$$\begin{aligned} \frac{\partial}{\partial \beta} \langle O(\beta) \rangle_{\beta} = & \langle \partial_{\beta} O(\beta) \rangle_{\beta} + \langle O(\beta) \mathcal{I} \rangle_{\beta} + \langle O(\beta) \partial_{\beta} \mathcal{F}(\beta) \rangle_{\beta} \\ & - \langle O(\beta) \rangle_{\beta} \langle \mathcal{I} \rangle_{\beta} - \langle O(\beta) \rangle_{\beta} \langle \partial_{\beta} \mathcal{F}(\beta) \rangle_{\beta}. \end{aligned} \quad (\text{B23})$$

For $\beta = 0$, the expression simplifies into

$$\begin{aligned} \left. \frac{\partial}{\partial \beta} \langle O(0) \rangle \right|_{\beta=0} = & \langle \partial_{\beta} O(0) \rangle + \langle O(0) \mathcal{I} \rangle + \langle O(0) \partial_{\beta} \mathcal{F}(0) \rangle \\ & - \langle O(0) \rangle_{\beta} \langle \mathcal{I} \rangle, \end{aligned} \quad (\text{B24})$$

so that the second derivative of the Gibbs free entropy is

$$\left. \frac{\partial^2 \Gamma_{\beta}}{\partial \beta^2} \right|_{\beta=0} = \frac{\partial}{\partial \beta} \langle \mathcal{I} \rangle = \langle \mathcal{I}^2 \rangle - \langle \mathcal{I} \rangle^2 + \langle \mathcal{I} \partial_{\beta} \mathcal{F}(0) \rangle. \quad (\text{B25})$$

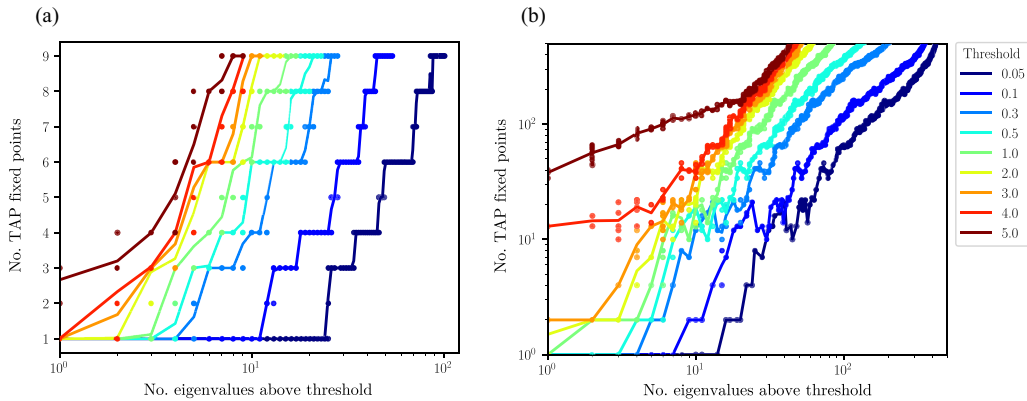


FIG. 11. For each model saved during the training, we recorded the number of TAP fixed points and the number of eigenvalues of the weight matrix above different thresholds. The different colors correspond to different thresholds, and the continuous lines represent the average number of TAP fixed points expected when a given number of eigenvalues are above the threshold. Panel (a) shows the results obtained on the SED data set, whereas panel (b) shows the results on the MNIST data set.

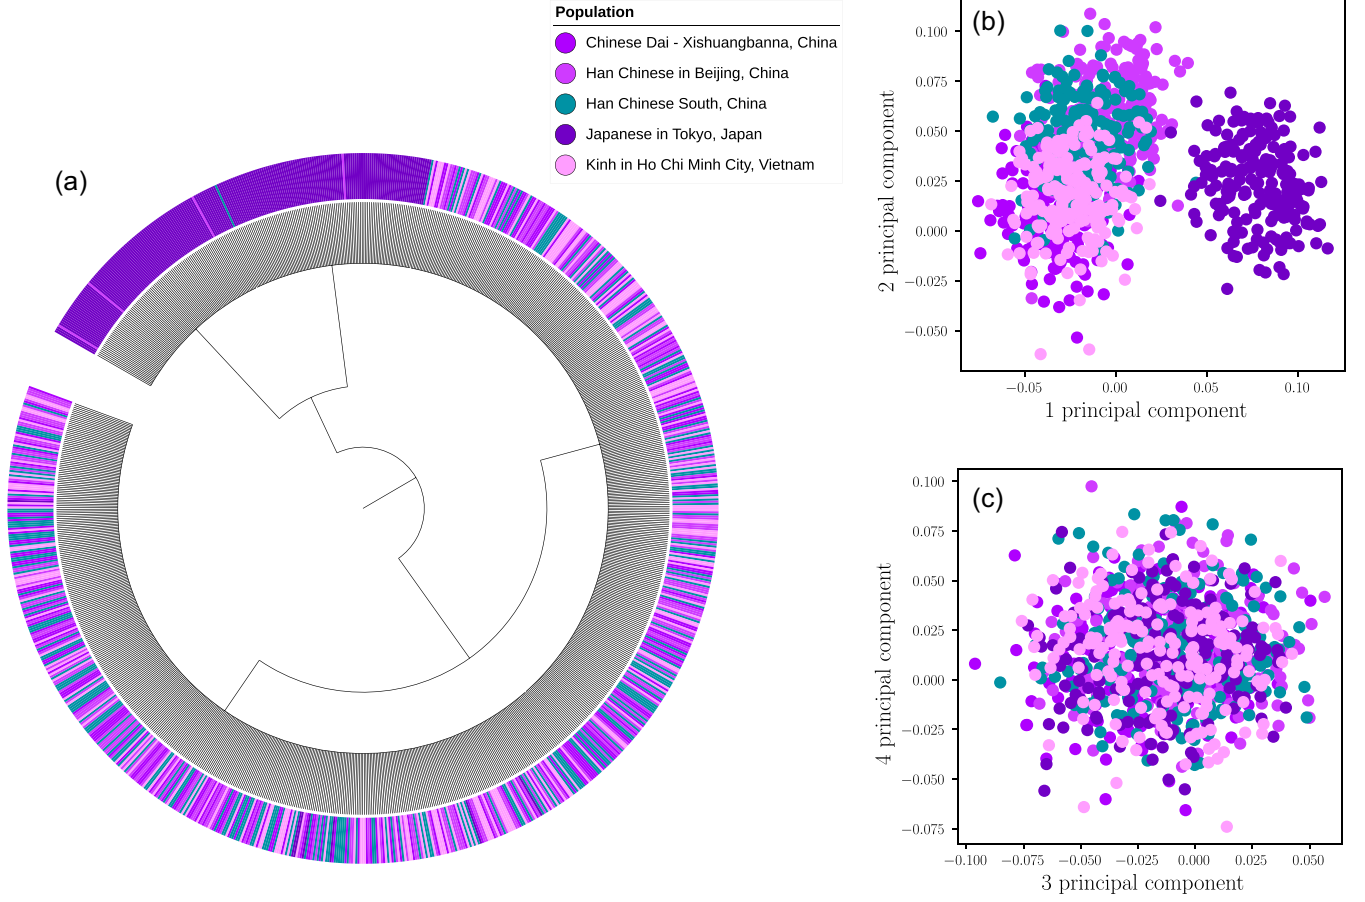


FIG. 12. (a) Tree generated on all the HGD data denoted as “East Asian.” The tree is able to only distinguish the data labeled as “Japanese in Tokyo, Japan” from the others. (b), (c) Projection of the data set along the first and second principal components [panel (b)] and along the third and fourth principal components [panel (c)].

To proceed further with the computation, we need to use the following pair correlation functions computed at $\beta = 0$:

$$\begin{aligned} \langle h_\mu h_\nu \rangle &= m_\mu \delta_{\mu\nu} + (1 - \delta_{\mu\nu}) m_\mu m_\nu, \\ \langle \delta_{v_i, q} \delta_{v_j, p} \rangle &= f_i^q \delta_{ij} \delta_{qp} + (1 - \delta_{ij}) f_i^q f_j^p. \end{aligned} \quad (\text{B26})$$

Also, we need to compute the first derivative of the external fields with respect to β at $\beta = 0$. These can be easily obtained

using the Schwarz theorem:

$$\begin{aligned} \partial_\beta \phi_i^q(0) &= - \left. \frac{\partial^2 \Gamma_\beta}{\partial \beta \partial f_i^q} \right|_{\beta=0} = - \left. \frac{\partial^2 \Gamma_\beta}{\partial f_i^q \partial \beta} \right|_{\beta=0} = - \sum_\mu w_{i\mu}^q m_\mu, \\ \partial_\beta \psi_\mu(0) &= - \left. \frac{\partial^2 \Gamma_\beta}{\partial \beta \partial m_\mu} \right|_{\beta=0} = - \left. \frac{\partial^2 \Gamma_\beta}{\partial m_\mu \partial \beta} \right|_{\beta=0} = - \sum_{iq} f_i^q w_{i\mu}^q. \end{aligned} \quad (\text{B27})$$

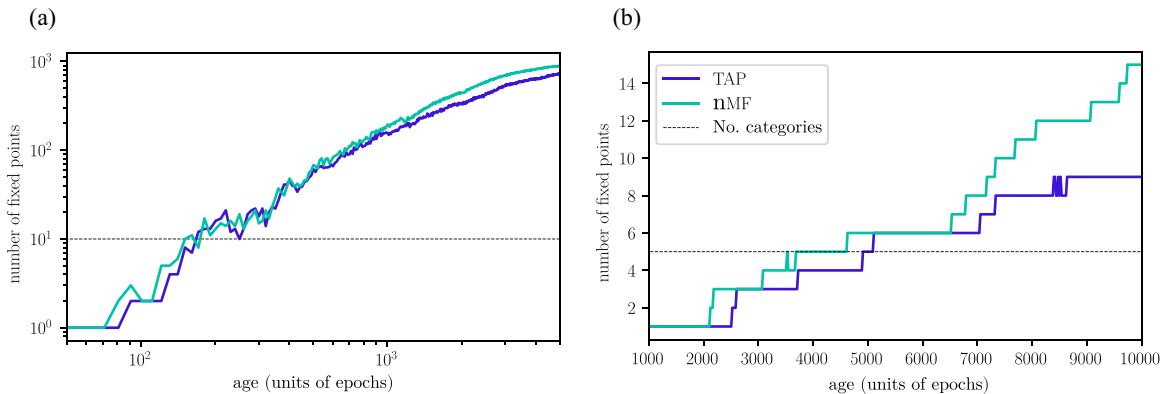


FIG. 13. Number of fixed points of the approximated free energy on the MNIST data set [panel (a)] and on the SED data set [panel (b)] as a function of the RBM’s age. Notice that plot (a) is in log-log scale, whereas plot (b) is in linear-linear scale.




FIG. 14. Phylogenetic tree obtained from the CPF protein data set using the neighbor-joining method and the software FastTree with default parameters. The external colored tags correspond to the experimental classification, whereas the leaf’s color represents the classification inferred by ProfileView in Ref. [35].

By expanding the expression (B25), after many simplifications we end up with the following second-order Taylor expansion of the Gibbs free entropy:

$$\begin{aligned}
 \Gamma_{\beta}^{(2)}(\mathbf{M}) &= \Gamma_0(\mathbf{M}) + \beta \left. \frac{\partial \Gamma_{\beta}(\mathbf{M})}{\partial \beta} \right|_{\beta=0} + \frac{\beta^2}{2} \left. \frac{\partial^2 \Gamma_{\beta}(\mathbf{M})}{\partial \beta^2} \right|_{\beta=0} \\
 &= \sum_{iq} f_i^q a_i^q + \sum_{\mu} m_{\mu} b_{\mu} - \sum_{iq} f_i^q \log f_i^q - \sum_{\mu} [m_{\mu} \log m_{\mu} + (1 - m_{\mu}) \log(1 - m_{\mu})] + \beta \sum_{iq\mu} f_i^q w_{i\mu}^q m_{\mu} \\
 &\quad + \frac{\beta^2}{2} \sum_{\mu} (m_{\mu} - m_{\mu}^2) \left[\sum_{iq} (w_{i\mu}^q)^2 f_i^q - \sum_i \left(\sum_q w_{i\mu}^q f_i^q \right)^2 \right].
 \end{aligned}
 \tag{B28}$$




FIG. 15. Comparison between three trees of the MNIST data set generated with three models having the same parameters but trained using different random seeds.




FIG. 16. Tree of the MNIST data set generated through an inverse-annealing of the oldest RBM model. We applied the tree construction procedure to a set of ~ 100 models at linearly spaced inverse temperatures in the interval $\beta \in [0.65, 1]$, starting from the coldest model at $\beta = 1$ and progressively increasing the temperature.

The previous expression is a valid approximation of the true free energy of the model only when the magnetizations fulfill the stationary conditions (B8):

$$\frac{\partial \Gamma_{\beta}^{(2)}}{\partial m_{\mu}} = 0, \quad \frac{\partial \Gamma_{\beta}^{(2)}}{\partial f_i^q} = 0. \quad (\text{B29})$$

When computing the stationary condition with respect to the visible magnetization, one has to take care of the fact that the normalization condition on $\{f_i^q\}$ leaves us with only $N_q - 1$ independent colors, so that $f_i^{N_q} = 1 - \sum_{q=1}^{N_q-1} f_i^q$. This has to be considered only in the entropic term, since elsewhere the magnetization couples with the weight matrix that is zero along the color N_q because of the lattice-gas gauge.

The full computation of the stationary conditions (B29) at $\beta = 1$ yields

$$\begin{aligned} \log\left(\frac{m_{\mu}}{1-m_{\mu}}\right) - b_{\mu} &= \sum_{iq} f_i^q w_{i\mu}^q - \left(m_{\mu} - \frac{1}{2}\right) \left[\sum_{iq} (w_{i\mu}^q)^2 f_i^q - \sum_i \left(\sum_q w_{i\mu}^q f_i^q \right)^2 \right], \\ \log\left(\frac{f_i^q}{f_i^{N_q}}\right) - a_i^q &= \sum_{\mu} m_{\mu} w_{i\mu}^q + \frac{1}{2} \sum_{\mu} (m_{\mu} - m_{\mu}^2) \left[(w_{i\mu}^q)^2 - 2w_{i\mu}^q \sum_p f_i^p w_{i\mu}^p \right], \end{aligned} \quad (\text{B30})$$




FIG. 17. Tree obtained on the MNIST data set by iterating the TAP equations starting from the data at each age.

where at the left-hand side we recognize, respectively, the inverse functions of the sigmoid (B11) and of the softmax (B12). Eventually, we end up with the following self-consistent TAP equations:

$$\begin{aligned}
 m_\mu &= \text{sigmoid} \left\{ b_\mu + \sum_{iq} f_i^q w_{i\mu}^q - \left(m_\mu - \frac{1}{2} \right) \left[\sum_{iq} (w_{i\mu}^q)^2 f_i^q - \sum_i \left(\sum_q w_{i\mu}^q f_i^q \right)^2 \right] \right\}, \\
 f_i^q &= \text{softmax} \left\{ a_i^q + \sum_\mu m_\mu w_{i\mu}^q + \frac{1}{2} \sum_\mu (m_\mu - m_\mu^2) \left[(w_{i\mu}^q)^2 - 2w_{i\mu}^q \sum_p f_i^p w_{i\mu}^p \right] \right\}. \quad (\text{B31})
 \end{aligned}$$

APPENDIX C: ADDITIONAL FIGURES

In this section, we show additional experiments that we performed to better understand the limitations of the presented method and the possibility of considering other tree construction schemes.

In Figs. 11 and 13 we study how the number of solutions of TAP equations evolves with the training time of the RBM (in Fig. 11 we study the correspondence between the number of points and the number of eigenvalues of the weight matrix above a certain threshold and in Fig. 13 we compare the number of fixed points obtained with naive mean-field and TAP).

Figure 12 shows the result of applying our method to a subset of the data to test the resolution power of the machine when trained with only a subcategory of the Human Genome dataset.

For comparison with our method, we show in Fig. 14 the phylogenetic tree of the protein family presented in the main text.

In Fig. 15, we show three alternative trees obtained by applying our method to three independent trainings of the MNIST dataset.

In Fig. 16, we show the tree obtained for the MNIST dataset by constructing the tree using a temperature annealing of the older machine instead of following the learning trajectory as proposed in our method.




FIG. 18. Tree obtained on the MNIST data set using an RBM trained with PCD-100 using the TAP equations.

In Fig. 17, we show the tree obtained by constructing a hierarchy with fixed points for each age instead of following the fusion of fixed points along the training trajectory and proposed in the main text.

In Fig. 18, we show the tree obtained when training the RBM with the TAP fixed points to approximate the gradient instead as with MCMC sampling and the PCD scheme as done in the main text.

-
- [1] M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, and T. Hwa, Identification of direct residue contacts in protein–protein interaction by message passing, *Proc. Natl. Acad. Sci. USA* **106**, 67 (2009).
 - [2] F. Morcos, A. Pagnani, B. Lunt, A. Bertolino, D. S. Marks, C. Sander, R. Zecchina, J. N. Onuchic, T. Hwa, and M. Weigt, Direct-coupling analysis of residue coevolution captures native contacts across many protein families, *Proc. Natl. Acad. Sci. USA* **108**, E1293 (2011).
 - [3] S. Cocco, C. Feinauer, M. Figliuzzi, R. Monasson, and M. Weigt, Inverse statistical physics of protein sequences: A key issues review, *Rep. Prog. Phys.* **81**, 032601 (2018).
 - [4] N. Le Roux and Y. Bengio, Representational power of restricted boltzmann machines and deep belief networks, *Neural Comput.* **20**, 1631 (2008).
 - [5] A. Decelle, G. Fissore, and C. Furtlehner, Spectral dynamics of learning in restricted boltzmann machines, *Europhys. Lett.* **119**, 60001 (2017).
 - [6] A. Decelle, G. Fissore, and C. Furtlehner, Thermodynamics of restricted boltzmann machines and related learning dynamics, *J. Stat. Phys.* **172**, 1576 (2018).
 - [7] A. Decelle, C. Furtlehner, and B. Seoane, Equilibrium and non-equilibrium regimes in the learning of restricted Boltzmann machines, in *Advances in Neural Information Processing Systems*, edited by M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, and J. W. Vaughan (Curran Associates, Red Hook, NY, 2021), Vol. 34, pp. 5345–5359.
 - [8] J. Tubiana and R. Monasson, Emergence of Compositional Representations in Restricted Boltzmann Machines, *Phys. Rev. Lett.* **118**, 138301 (2017).
 - [9] J. Tubiana, S. Cocco, and R. Monasson, Learning protein constitutive motifs from sequence data, *Elife* **8**, e39397 (2019).
 - [10] M. Gabrié, E. W. Tramel, and F. Krzakala, Training restricted Boltzmann machine via the Thouless-Anderson-Palmer free energy, in *Advances in Neural Information Processing Systems*,

- edited by C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, Red Hook, NY, 2015), Vol. 28.
- [11] E. W. Tramel, M. Gabrié, A. Manoel, F. Caltagirone, and F. Krzakala, Deterministic and Generalized Framework for Unsupervised Learning with Restricted Boltzmann Machines, *Phys. Rev. X* **8**, 041006 (2018).
- [12] A. Maillard, L. Foini, A. L. Castellanos, F. Krzakala, M. Mézard, and L. Zdeborová, High-temperature expansions and message passing algorithms, *J. Stat. Mech.* (2019) 113301.
- [13] J. Tubiana, Restricted Boltzmann machines: From compositional representations to protein sequence analysis, Ph.D. theses, Université Paris Sciences et Lettres (2018).
- [14] T. Tieleman, Training restricted Boltzmann machines using approximations to the likelihood gradient, in *Proceedings of the 25th International Conference on Machine Learning*, ICML'08 (ACM Press, New York, NY, 2008), pp. 1064–1071.
- [15] E. Agoritsas, G. Catania, A. Decelle, and B. Seoane, Explaining the effects of non-convergent sampling in the training of energy-based models, [arXiv:2301.09428](https://arxiv.org/abs/2301.09428).
- [16] A. Decelle and C. Furtlehner, Restricted boltzmann machine: Recent advances and mean-field theory, *Chin. Phys. B* **30**, 040202 (2021).
- [17] F. Zamponi, Mean field theory of spin glasses, [arXiv:1008.4844](https://arxiv.org/abs/1008.4844) (2010).
- [18] T. Plefka, Convergence condition of the tap equation for the infinite-ranged ising spin glass model, *J. Phys. A* **15**, 1971 (1982).
- [19] A. Georges and J. S. Yedidia, How to expand around mean-field theory using high-temperature expansions, *J. Phys. A* **24**, 2173 (1991).
- [20] T. Tanaka, Information geometry of mean-field approximation, *Neural Comput.* **12**, 1951 (2000).
- [21] T. Tanaka, Mean-field theory of boltzmann machine learning, *Phys. Rev. E* **58**, 2302 (1998).
- [22] T. Yokota, Ordered phase for the infinite-range potts-glass model, *Phys. Rev. B* **51**, 962 (1995).
- [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] H. Huang, *Statistical Mechanics of Neural Networks* (Springer, Singapore, 2021).
- [25] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96 (AAAI Press, Portland, Oregon, 1996), pp. 226–231.
- [26] K. Rose, E. Gurewitz, and G. C. Fox, Statistical Mechanics and Phase Transitions in Clustering, *Phys. Rev. Lett.* **65**, 945 (1990).
- [27] T. Bonnaire, A. Decelle, and N. Aghanim, Cascade of phase transitions for multiscale clustering, *Phys. Rev. E* **103**, 012105 (2021).
- [28] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* **86**, 2278 (1998).
- [29] G. P. Consortium *et al.*, A global reference for human genetic variation, *Nature (London)* **526**, 68 (2015).
- [30] N. Béreux, A. Decelle, C. Furtlehner, and B. Seoane, Learning a restricted Boltzmann machine using biased Monte Carlo sampling, *SciPost Phys.* **14**, 032 (2023).
- [31] A. Sancar, Structure and function of dna photolyase and cryptochrome blue-light photoreceptors, *Chem. Rev.* **103**, 2203 (2003).
- [32] K. Brettel and M. Byrdin, Reaction mechanisms of dna photolyase, *Curr. Opin. Struct. Biol.* **20**, 693 (2010).
- [33] I. Chaves, R. Pokorny, M. Byrdin, N. Hoang, T. Ritz, K. Brettel, L.-O. Essen, G. T. J. van der Horst, A. Batschauer, and M. Ahmad, The cryptochromes: Blue light photoreceptors in plants and animals, *Ann. Rev. Plant Biol.* **62**, 335 (2011).
- [34] M. Jaubert, J.-P. Bouly, M. R. d'Alcalà, and A. Falciatore, Light sensing and responses in marine microalgae, *Curr. Opin. Plant Biol.* **37**, 70 (2017).
- [35] R. Vicedomini, J. Bouly, E. Laine, A. Falciatore, and A. Carbone, Multiple profile models extract features from protein sequence data and resolve functional diversity of very different protein families, *Mol. Biol. Evol.* **39**, msac070 (2022).