



LETTER

Spectral dynamics of learning in restricted Boltzmann machines

To cite this article: A. Decelle *et al* 2017 *EPL* **119** 60001

View the [article online](#) for updates and enhancements.

Related content

- [Approximate message passing with restricted Boltzmann machine priors](#)
Eric W Tramel, Angélique Drémeau and Florent Krzakala
- [Constrained low-rank matrix estimation: phase transitions, approximate message passing and applications](#)
Thibault Lesieur, Florent Krzakala and Lenka Zdeborová
- [Statistical mechanics of unsupervised feature learning in a restricted Boltzmann machine with binary synapses](#)
Haiping Huang

Spectral dynamics of learning in restricted Boltzmann machines

A. DECELLE^{1(a)}, G. FISSORE^{1,2} and C. FURTLEHNER²

¹ *LRI, AO team, Bât 660 Université Paris Sud - Orsay Cedex 91405, France*

² *Inria Saclay - Tau team, Bât 660 Université Paris Sud - Orsay Cedex 91405, France*

received 10 August 2017; accepted in final form 30 October 2017

published online 29 November 2017

PACS 02.30.Zz – Inverse problems

PACS 02.70.Hm – Spectral methods

PACS 89.75.-k – Complex systems

Abstract – The restricted Boltzmann machine (RBM), an important tool used in machine learning in particular for unsupervised learning tasks, is investigated from the perspective of its spectral properties. Starting from empirical observations, we propose a generic statistical ensemble for the weight matrix of the RBM and characterize its mean evolution. This let us show how in the linear regime, in which the RBM is found to operate at the beginning of the training, the statistical properties of the data drive the selection of the unstable modes of the weight matrix. A set of equations characterizing the non-linear regime is then derived, unveiling in some way how the selected modes interact in later stages of the learning procedure and defining a deterministic learning curve for the RBM.

Copyright © EPLA, 2017

Introduction. – A restricted Boltzmann machine (RBM) [1] constitutes nowadays a common tool on the shelf of machine learning practitioners. It is a generative model, in the sense that it defines a probability distribution, which can be learned to approximate any distribution of data points living in some N -dimensional space, with N potentially large. It also often constitutes a building block of more complex neural network models [2,3]. The standard learning procedure called contrastive divergence [4] is well-documented [5] although being still a not so well-understood fine empirical art, with many hyperparameters to tune without much guidelines. At the same time an RBM can be regarded as a statistical physics model, being defined as a Boltzmann distribution with pairwise interactions on a bipartite graph. Similar models have been already the subject of many studies in the 1980s [6–9] which mainly concentrated on the learning capacity, *i.e.*, the number of independent patterns that could be stored in such a model. The second life of neural networks has renewed the interest of statistical physicists for such models. Recent works actually propose to exploit its statistical physics formulation to define mean-field-based learning methods using Thouless-Anderson-Palmer (TAP) equations [10–12]. Meanwhile some analysis of its static properties, assuming a given learned weight matrix W ,

have been proposed [13,14] in order to understand collective phenomena in the latent representation [15], *i.e.*, the way latent variables organize themselves to represent actual data. One common assumption made in these works is that the weights of W are i.i.d. which as we shall see is unrealistic. Concerning the learning procedure of neural networks, many recent statistical-physics-based analysis have been proposed, most of them within teacher-student setting [16] which imposes a strong assumption on the data, namely that these are generated from a model belonging to the parametric family of interest, hiding as a consequence the role played by the data themselves in the procedure. From the analysis of related models [17,18], it is already a well-established fact that a selection of the most important modes of the singular value decomposition (SVD) of the data is performed in the linear case. In fact in the simpler context of linear feed-forward models the learning dynamics can be fully characterized by means of the SVD of the data matrix [19], showing in particular the emergence of each mode by order of importance regarding singular values.

In this work we follow this guideline in the context of a general RBM. We propose to characterize both the learned RBM and the learning process itself by the SVD spectrum of the weight matrix in order to isolate the information content of an RBM. This allows us then to write a deterministic learning equation leaving aside the fluctuations.

^(a)E-mail: aurelien.decelle@lri.fr

This equation is subsequently analyzed first in the linear regime to identify the unstable deformation modes of W ; secondly at equilibrium assuming the learning is converging, in order to understand the nature of the non-linear interactions between these modes and how these are determined from the input data. In the first section we recall the RBM model and associated learning algorithm. In the second section we show how this algorithm can be described by a generic learning equation. Then we first analyze the linear regime and thereafter we describe what happens with the binary RBM. A set of dynamical parameters is shown to emerge naturally from the SVD decomposition of the weight matrix. The convergence toward equilibrium is analyzed and illustrated later with actual tests on the MNIST dataset.

The RBM and associated learning procedure. – An RBM is a Markov random field with pairwise interactions defined on a bipartite graph formed by two layers of non-interacting variables: the visible nodes and the hidden nodes representing, respectively, data configurations and latent representations. The former noted $\mathbf{s} = \{s_i, i = 1, \dots, N_v\}$ correspond to explicit representations of the data while the latter noted $\boldsymbol{\sigma} = \{\sigma_j, j = 1, \dots, N_h\}$ are there to build arbitrary dependences among the visible units. They play the role of an interacting field among visible nodes. Usually the nodes are binary-valued (of Boolean type or Bernoulli distributed) but Gaussian distributions or more broadly arbitrary distributions on real-valued bounded support are also used [20], ultimately making RBMs adapt for more heterogeneous data sets. Here to simplify we assume that visible and hidden nodes will be taken as binary variables $s_i, \sigma_j \in \{-1, 1\}$ (using ± 1 values has the advantage of symmetrizing the equations hence avoiding to deal with “hidden” biases on the variables when considering binary $\{0, 1\}$ variables). Like the Hopfield model [6] which can actually be cast into an RBM [21] an energy function is defined for a configuration of nodes

$$E(\mathbf{s}, \boldsymbol{\sigma}) = - \sum_{i,j} s_i w_{ij} \sigma_j - \sum_{i=1}^{N_v} \eta_i s_i - \sum_{j=1}^{N_h} \theta_j \sigma_j \quad (1)$$

and this is exploited to define a joint distribution between visible and hidden units, namely the Boltzmann distribution

$$p(\mathbf{s}, \boldsymbol{\sigma}) = \frac{e^{-E(\mathbf{s}, \boldsymbol{\sigma})}}{Z}, \quad (2)$$

where W is the weight matrix and $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ are biases, or external fields on the variables. $Z = \sum_{\mathbf{s}, \boldsymbol{\sigma}} e^{-E(\mathbf{s}, \boldsymbol{\sigma})}$ is the partition function of the system. The joint distribution between visible variables is then obtained by summing over hidden ones. In this context, learning the parameters of the RBM means that, given a dataset of M samples composed of N_v variables, we ought to infer values to W , $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ such that new generated data obtained by sampling this distribution should be similar to the input data.

The general method to infer the parameters is to maximize the likelihood of the model, where the probability density function (pdf) (2) has first been summed over the hidden variables

$$\mathcal{L} = \sum_j \log \left(2 \cosh \left(\sum_i w_{ij} s_i + \theta_j \right) \right) - \log(Z). \quad (3)$$

Different methods of learning have been set up and proven to work efficiently, in particular the contrastive divergence (CD) algorithm from Hinton [4] and more recently TAP-based learning [10]. They all correspond to expressing the gradient ascent on the likelihood as

$$\Delta w_{ij} = \gamma (\langle s_i \sigma_j p(\sigma_j | \mathbf{s}) \rangle_{\text{Data}} - \langle s_i \sigma_j \rangle_{\text{RBM}}), \quad (4)$$

where γ is the learning rate. Similar equations can be derived for the biases. The main problem is the second term on the rhs of (2) which is not tractable, and various methods basically differ in their way of estimating this term (Monte Carlo chains, mean field, TAP, ...). For an efficient learning the first term also has to be approximated by making use of random mini batches of data at each step.

Deterministic dynamics of the learning. – In order to understand the dynamics of the learning we first project the CD equation (2) onto the basis defined by the SVD of W . As a generalization of eigenmodes decomposition to rectangular matrices, the SVD for a RBM is given by

$$\mathbf{W} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T, \quad (5)$$

where \mathbf{U} is an orthogonal $N_v \times N_h$ matrix whose columns are the left singular vectors \mathbf{u}^α , \mathbf{V} is an orthogonal $N_h \times N_h$ matrix whose columns are the right singular vectors \mathbf{v}^α and $\boldsymbol{\Sigma}$ is a diagonal matrix whose elements are the singular values w_α . The separation into left and right singular vectors is due to the rectangular nature of the decomposed matrix, and the similarity with eigenmodes decomposition is revealed by the following SVD equations:

$$\begin{aligned} \mathbf{W} \mathbf{v}^\alpha &= w_\alpha \mathbf{u}^\alpha, \\ \mathbf{W}^T \mathbf{u}^\alpha &= w_\alpha \mathbf{v}^\alpha. \end{aligned}$$

We consider the usual situation where $N_h < N_v$, which means that the rank of W is at most N_h . $W(t)$ represents the learned weight matrix at time t . Let $\{w_\alpha(t) \in [0, +\infty[\}$, $\{u_\alpha(t) \in \mathbb{R}^{N_v}\}$ and $\{v_\alpha(t) \in \mathbb{R}^{N_h}\}$ such that the following decomposition $w_{ij}(t) = \sum_\alpha u_i^\alpha(t) w_\alpha(t) v_j^\alpha(t)$ holds. Discarding stochastic fluctuations usually inherent to the learning procedure and letting the learning rate $\gamma \rightarrow 0$, the continuous version of (2) can be recast as follows:

$$\begin{aligned} \left(\frac{dw}{dt} \right)_{\alpha\beta} &= \delta_{\alpha,\beta} \frac{dw_\alpha}{dt}(t) \\ &+ (1 - \delta_{\alpha,\beta}) (w_\beta(t) \Omega_{\beta\alpha}^v(t) + w_\alpha(t) \Omega_{\alpha\beta}^h(t)) = \\ &\langle s_\alpha \sigma_\beta \rangle_{\text{Data}} - \langle s_\alpha \sigma_\beta \rangle_{\text{RBM}}, \end{aligned} \quad (6)$$

$$\Omega_{\alpha\beta}^v(t) = -\Omega_{\beta\alpha}^v \stackrel{\text{def}}{=} \frac{d\mathbf{u}^{\alpha,T}}{dt} \mathbf{u}^\beta = \frac{-1}{w_\alpha + w_\beta} \left(\frac{dw}{dt} \right)_{\alpha\beta}^A + \frac{1}{w_\alpha - w_\beta} \left(\frac{dw}{dt} \right)_{\alpha\beta}^S \quad (7)$$

$$\Omega_{\alpha\beta}^h(t) = -\Omega_{\beta\alpha}^h \stackrel{\text{def}}{=} \frac{d\mathbf{v}^{\alpha,T}}{dt} \mathbf{v}^\beta = \frac{1}{w_\alpha + w_\beta} \left(\frac{dw}{dt} \right)_{\alpha\beta}^A + \frac{1}{w_\alpha - w_\beta} \left(\frac{dw}{dt} \right)_{\alpha\beta}^S. \quad (8)$$

Here everything is expressed in the reference frame defined by singular vectors of W . $s_\alpha = \sum_i u_i^\alpha s_i$ and $\sigma_\alpha = \sum_j v_j^\alpha \sigma_j$ represent spin configurations in this frame. Note that one has to keep track of the original reference frame to be able to evaluate the data and RBM average in particular when the basic variables are discrete. We have introduced the skew-symmetric rotation generators $\Omega_{\alpha\beta}^{v,h}(t)$ of the basis vectors induced by the dynamics. These tell us how the data rotate relatively to this frame. The superscripts S, A indicate the symmetric (respectively, anti-symmetric) part of the matrix. Note that these equations become singular when some degeneracy occurs in W because then the SVD is not uniquely defined. This is not really a problem since we are interested in rotations among non-degenerate modes, the rest corresponding to gauge degrees of freedom. Similar equations can be derived for the fields $\eta_\alpha(t) \stackrel{\text{def}}{=} \sum_i \eta_i(t) u_i^\alpha(t)$ and $\theta_\alpha(t) \stackrel{\text{def}}{=} \sum_j v_j^\alpha(t) \theta_j(t)$ projected onto the SVD modes. At this point we make the assumption that the learning dynamics is represented by a trajectory of $(\{w_\alpha(t), \eta_\alpha(t), \theta_\alpha(t), \Omega_{\alpha\beta}^{v,h}(t)\})$, while the specific realization of the u_i^α and v_j^α is considered to be irrelevant, and can be averaged out with respect to some simple distributions, as long as this average is correlated with the data. This means that the decomposition $\hat{s}_\alpha = \sum_i u_i^\alpha \hat{s}_i$ of any given sample configuration is assumed also to be kept fixed while averaging. What matters mainly is the strength given by $w_\alpha(t)$ and the rotation given by $\Omega_{\alpha\beta}^{v,h}(t)$ of these SVD modes. Assuming for example i.i.d. centered normal distribution with respective variance $1/N_v$ and $1/N_h$ for u_i^α and v_j^α , the empirical term takes the simple form:

$$\langle s_\alpha \sigma_\beta \rangle_{\text{Data}} = \frac{1}{N_h} \left\langle s_\alpha (s_\beta w_\beta - \theta_\beta) V \left(\frac{1}{N_h} \sum_\gamma (w_\gamma s_\gamma - \theta_\gamma)^2 \right) \right\rangle_{\text{Data}},$$

where $V(x) = \int dy \frac{e^{-y^2/2}}{\sqrt{2\pi}} \text{sech}^2(\sqrt{x}y)$, (9)

which actually depends on the activation function (an hyperbolic tangent in this case). The main point here is that the empirical term defines an operator whose decomposition onto the SVD modes of W functionally depends solely on w_α, θ_α and on the projection of the data on the SVD modes of W . This term is precisely driving the dynamics. The adaptation of the RBM to this driving force is

given by the second term which can be as well estimated in the thermodynamic limit, as a function of w_α, θ_α and η_α alone.

Linear instabilities. – First let us consider the linear regime which can be analyzed thoroughly. It can be obtained by rescaling all the weights and fields by a common “inverse temperature” β factor and let this go to zero in eqs. (6). This limit can be understood by keeping up to quadratic terms in the mean-field free energy and should correspond to the first stages of the learning. In this limit, magnetizations (μ_v, μ_h) of visible and hidden variables have Gaussian fluctuations with covariance matrix

$$C(\mu_v, \mu_h) \stackrel{\text{def}}{=} \begin{bmatrix} \sigma_v^{-2} & -W \\ -W^T & \sigma_h^{-2} \end{bmatrix}^{-1}$$

with $\sigma_v^2 = \sigma_h^2 = 1$ introduced for the sake of generality when considering general linear RBM. To simplify the exposition, we discard the biases of the data and related fields $(\theta_\alpha, \eta_\alpha)$ of the RBM. In that case the empirical term in (7) involves directly the covariance matrix of the data expressed in the frame defined by the SVD modes of W :

$$\langle s_\alpha \sigma_\beta \rangle_{\text{Data}} = \sigma_h^2 w_\beta \langle s_\alpha s_\beta \rangle_{\text{Data}}.$$

From $C(\mu_v, \mu_h)$ we get the other terms yielding the following equations:

$$\frac{dw_\alpha}{dt} = w_\alpha \sigma_h^2 \left(\langle s_\alpha^2 \rangle_{\text{Data}} - \frac{\sigma_v^2}{1 - \sigma_v^2 \sigma_h^2 w_\alpha^2} \right),$$

$$\Omega_{\alpha\beta}^{v,h} = (1 - \delta_{\alpha\beta}) \sigma_h^2 \left(\frac{w_\beta - w_\alpha}{w_\alpha + w_\beta} \mp \frac{w_\beta + w_\alpha}{w_\alpha - w_\beta} \right) \langle s_\alpha s_\beta \rangle_{\text{Data}}.$$

Note that these equations are exact for a linear RBM, since they can be derived without any reference to the coordinates of u_α and v_α over which we average in the non-linear regime. These equations tell us that, during the learning the vectors \mathbf{u}^α (and also \mathbf{v}^α) will rotate until being aligned to the principal components of the data, *i.e.*, until $\langle s_\alpha s_\beta \rangle_{\text{Data}}$ becomes diagonal. Then calling \hat{w}_α^2 the corresponding empirical variance given by the data, the system reaches the following equilibrium values:

$$w_\alpha^2 = \begin{cases} \frac{\hat{w}_\alpha^2 - \sigma_v^2}{\sigma_v^2 \sigma_h^2 \hat{w}_\alpha^2}, & \text{if } \hat{w}_\alpha^2 > \sigma_v^2, \\ 0, & \text{if } \hat{w}_\alpha^2 \leq \sigma_v^2. \end{cases}$$

From this we see that the RBM selects the strongest SVD modes in the data. The linear instabilities correspond to directions for which the variance of the data is above the threshold σ_v^2 . This determines the deformations of the weight matrix which can develop during the learning and will eventually interact, following the usual mechanism of non-linear pattern formation like, *e.g.*, in reaction-diffusion processes [22]. Other possible deformations are damped to zero. The linear RBM will therefore learn all (up to N_h) principal components that passed the




Fig. 1: (Colour online) Time evolution of the eigenvalues in the linear model and of the likelihood. We observe very clearly how the different modes emerge from the bulk and how the likelihood increases at each eigenvalue learned. In the inset, the scalar product of the vectors \mathbf{u} obtained from the SVD of the data and of \mathbf{w} . The \mathbf{u} 's of \mathbf{w} are aligned with the SVD of the data at the end of the learning.

threshold but it is important to remember that the resulting distribution will still be unimodal. Note that this selection mechanism is already known to occur for linear auto-encoders [18] or some other similar linear Boltzmann machines [17]. In fig. 1 we can see the eigenvalues being learned one by one in a linear RBM. For non-linear RBM when the system escapes the linear regime, a well-suited mean-field theory is required to understand the dynamics and the steady-state regime.

Non-linear regime. – During the linear regime some specific modes are selected and at some point these modes start to interact in a non-trivial manner. The empirical term in (8) involves higher-order statistics of the data as exemplified by (9) and the Gaussian estimation with $\sigma_v^2 = \sigma_h^2 = 1$ of the RBM response term $\langle s_\alpha \sigma_\beta \rangle_{\text{RBM}}$ is no longer valid. In order to estimate this term in the thermodynamic limit, some assumptions on the form of the weight matrix are needed. A common assumption consists in considering i.i.d. random variables for the weights w_{ij} and this, like for example in [13–15], generally leads to a Marchenko-Pastur distribution of the singular values of W , which, as we shall see in the next section, is unrealistic. Instead, based on our experiments such distribution corresponds to the noise of the weight matrix, while its information content is better expressed by the presence of SVD modes outside of the bulk. This leads us to write the weight matrix as

$$w_{ij} = \sum_{\alpha=1}^K w_\alpha u_i^\alpha v_j^\alpha + r_{ij}, \quad (10)$$

where the $w_\alpha = O(1)$ are isolated singular values (describing a rank K matrix), the \mathbf{u}^α and \mathbf{v}^α are the eigenvectors of the SVD decomposition and the $r_{ij} = \mathcal{N}(0, \sigma^2/L)$ where $L = \sqrt{N_h N_v}$ are i.i.d. corresponding to noise. To be consistent with the linear analysis, these modes are

assumed to span the (left) subspace corresponding to the part of the empirical SVD above threshold while r spans the complementary space of empirical modes below threshold. We limit the analysis here to the case where K is finite. This then allows us to assume simple distributions p_u and p_v for the components of \mathbf{u}^α and \mathbf{v}^α considered i.i.d. for instance. This altogether defines our statistical ensemble of RBM to which we restrict ourselves to study the learning procedure. For K extensive we should instead average over the orthogonal group which would lead to a slightly different mean-field theory [23,24]. In the present form our model of RBM is similar to the Hopfield model and recent generalizations [25], the patterns being represented by the SVD modes outside the bulk. The main difference, in addition to the bipartite structure of the graph, is the non-degeneracy of the singular values w_α . Still the analysis in the thermodynamic limit follows classical treatments like [7,26] for the Hopfield model or [14] for bipartite models. The starting point is to express the average over u, v and weights r_{ij} of the log partition function Z in (2) with the help of the replica trick:

$$\mathbb{E}_{u,v,r}[\log(Z)] = \lim_{p \rightarrow 0} \frac{d}{dp} \mathbb{E}_{u,v,r}[Z^p].$$

After averaging over the i.i.d. weights, 4 sets of order parameters $\{(m_\alpha^a, \bar{m}_\alpha^a), a = 1, \dots, p; \alpha = 1, \dots, K\}$ and $\{(Q_{ab}, \bar{Q}_{ab}), a, b = 1, \dots, p; a \neq b\}$ are introduced with the help of two distinct Hubbard-Stratonovich transformations. These variables represent the following quantities:

$$m_\alpha^a \sim \frac{1}{\sqrt{L}} \mathbb{E}_{u,v,r}(\langle \sigma_\alpha^a \rangle), \quad \bar{m}_\alpha^a \sim \frac{1}{\sqrt{L}} \mathbb{E}_{u,v,r}(\langle s_\alpha^a \rangle), \\ Q_{ab} \sim \mathbb{E}_{u,v,r}(\langle \sigma_i^a \sigma_i^b \rangle), \quad \bar{Q}_{ab} \sim \mathbb{E}_{u,v,r}(\langle s_j^a s_j^b \rangle),$$

namely the correlations of the hidden (respectively, visible) states with the left (respectively, right) singular vectors and the Edward-Anderson order parameters measuring the correlation between replicas of hidden or visible states. \mathbb{E}_u and \mathbb{E}_v denote an average with respect to the rescaled components $u \simeq \sqrt{N_v} u_i^\alpha$ and $v \simeq \sqrt{N_h} v_j^\alpha$ of the SVD modes. The transformations involve pairs of complex integration variables because of the asymmetry introduced by the two-layers structure by contrast to fully connected models. They lead to the following representation:

see equation on top of the next page

$\kappa = N_h/N_v$ and $B[\bar{m}, \bar{Q}]$ obtained from $A[m, Q]$ by replacing u by v , η by θ and κ by $1/\kappa$. The thermodynamic properties are obtained by first letting $L \rightarrow \infty$ allowing for a saddle point approximation and then the limit $p \rightarrow 0$ is taken. We restrict here the discussion to replica symmetric (RS) saddle points [27]. The breakdown of RS can actually be determined by computing the so-called AT line [28] and will be detailed somewhere else [29]. In the RS case the set $\{(Q_{ab}, \bar{Q}_{ab})\}$ reduces to a pair (q, \bar{q}) of spin glass parameters, while quenched magnetization towards the SVD directions are now represented by $\{(m_\alpha, \bar{m}_\alpha), \alpha = 1, \dots, K\}$.

$$\mathbb{E}_{u,v,r}[Z^p] = \int \prod_{a,\alpha} \frac{d\bar{m}_\alpha^a}{2\pi} \prod_{a \neq b} \frac{dQ_{ab} d\bar{Q}_{ab}}{2\pi} \exp \left\{ -L \left(\sum_{a,\alpha} w_\alpha m_\alpha \bar{m}_\alpha + \frac{\sigma^2}{2} \sum_{a \neq b} Q_{ab} \bar{Q}_{ab} - \frac{1}{\sqrt{\kappa}} A[m, Q] - \sqrt{\kappa} B[\bar{m}, \bar{Q}] \right) \right\}$$

with $A[m, Q] \stackrel{\text{def}}{=} \log \left[\sum_{S^a \in \{-1, 1\}} \mathbb{E}_u \left(e^{\frac{\sqrt{\kappa}\sigma^2}{2} \sum_{a \neq b} Q_{ab} S^a S^b + \kappa^{\frac{1}{4}} \sum_{a,\alpha} (m_\alpha^a w_\alpha - \eta_\alpha) u^\alpha S^a} \right) \right],$

Letting $x = \mathcal{N}(0, 1)$ and skipping some details, the saddle point equations are given by

$$(m_\alpha, \bar{m}_\alpha) = \mathbb{E} \left(\kappa^{\frac{1}{4}} v^\alpha \tanh(\bar{h}(x, v)), \kappa^{-\frac{1}{4}} u^\alpha \tanh(h(x, u)) \right), \quad (11)$$

$$(q, \bar{q}) = \mathbb{E} \left(\tanh^2(\bar{h}(x, v)), \tanh^2(h(x, u)) \right), \quad (12)$$

with \mathbb{E} denoting the average over (u, v, x) and

$$h(x, u) \stackrel{\text{def}}{=} \kappa^{\frac{1}{4}} \left(\sigma \sqrt{q} x + \sum_{\gamma} (w_\gamma m_\gamma - \eta_\gamma) u^\gamma \right),$$

$$\bar{h}(x, v) \stackrel{\text{def}}{=} \kappa^{-\frac{1}{4}} \left(\sigma \sqrt{\bar{q}} x + \sum_{\gamma} (w_\gamma \bar{m}_\gamma - \theta_\gamma) v^\gamma \right).$$

These fixed point equations can be solved numerically to tell us how the variables condensate on the SVD modes within each equilibrium state of the distribution and whether a spin glass phase is present or not. The important point here is that with K finite and a non-degenerate spectrum the mode with highest singular value dominates the ferromagnetic phase. The phase diagram looks in fact similar to the one of the SK model with ferromagnetic coupling, when $1/\sigma$ is interpreted as temperature and w_{max}/σ as the ferromagnetic coupling. Some subtleties arise when considering various ways of averaging over singular vectors components [29]. In [15,30] it is underlined the importance of the capability of networks to produce compositional states structured by combination of hidden variables. In our representation, we do not have direct access to this property, but to the dual one in some sense, namely states corresponding to a combination of modes. Their presence and their structure are rather sensitive to the way the average over u and v is performed. In this respect the case where \mathbf{u}^α and \mathbf{v}^α are Gaussian i.i.d. distributed is very special: all other fixed points associated to lower modes can be shown to be unstable as well as fixed points associated to combinations of modes. Instead, for other distributions with smaller kurtosis, like uniform or Bernoulli, stable fixed points associated to many different single modes or combinations of modes can exist and contribute to the thermodynamics.

Coming back to the learning dynamics, the first thing which is expected, already from the linear analysis, is that the noise term in (10) vanishes by condensing into a delta function of zero modes. Then the term corresponding

to the response of the RBM in (6) is estimated (in absence of bias) in the thermodynamic by means of the order parameters defined previously:

$$\langle s_\alpha \sigma_\beta \rangle_{\text{RBM}} = \frac{L}{Z_{\text{MF}}} \sum_{q=1}^C e^{-F_q} \bar{m}_\alpha^{(q)} m_\beta^{(q)}, \quad Z_{\text{MF}} \stackrel{\text{def}}{=} \sum_q e^{-F_q},$$

where the index q runs over all stable fixed point solutions of (11), (12) weighted accordingly to their free energy. These are the dominant contributions as long as free energy differences are $O(1)$, internal fluctuations given by each fixed point are comparatively of order $O(1/L)$. Note that this is the reason why the RBM needs to reach a ferromagnetic phase with many states to be able to match the empirical term in (7) in order to converge. For instance, in the case of a multimodal data distribution with many well-separated clusters, the SVD modes of W which will develop are the one pointing in the direction of the magnetizations defined by these clusters. In this simple case the RBM will evolve as in the linear case to a state such that the empirical term becomes diagonal, while the singular values adjust themselves until matching the proper magnetization in each fixed point. More precise statements about the phase diagram of the RBM and the behaviour of our dynamical equations including the dynamics of the external fields η_α and θ_α will be given in [29].

Tests on the MNIST dataset. – We illustrate our results on the MNIST dataset. The MNIST dataset is composed of 60000 images of handwritten digits of 28×28 pixels. It is known that RBMs perform reasonably well on this dataset and therefore we can now interpret in the light of the preceding sections how the learning goes. For the training of the MNIST dataset we use the following parameters. The weights of the matrix W were initiated randomly from a centered Gaussian distribution with a variance of 0.01 such that the Marchenko-Pastur (MP) bulk does not pass the threshold. The visible fields are initialized to reproduce the empirical mean of the data for each visible variable. The hidden field is put to zero. The learning rate is chosen to be ≈ 0.01 . With these parameters we verified that our machine was able to sample digits in a satisfactory way after 20 epochs. Now we can investigate the value of some observables introduced previously. First, we look at the SVD modes of the matrix w during the learning in fig. 2. We see that, after seeing only few updates the system has already learned many SVD modes from the data.




Fig. 2: (Colour online) (a) Singular values distribution of the initial random matrix compared to the Marchenko-Pastur law. (b) With the training we can see some singular values strengthening and overcoming the threshold set by the Marchenko-Pastur law. (c) Distribution of the singular values after a long training: we can see many outliers spread above threshold and a spike of below-threshold singular values near zero.




Fig. 3: (a) First mode learnt by the RBM with the external visible field initialized as a null vector. (b) External visible field initialized on the empirical mean. (c) First principal components extracted from the training set. (d) Principal components extracted from the training set (starting from the second). (e) The first 10 modes of a RBM trained for 1 epoch. (f) Same as (e) but after a 10 epochs training.

In fig. 2(a)–(c), we observe what is expected from the linear regime. Some modes escape from the MP bulk of the eigenvalues while other condense down to zero. In particular, we can see that the modes at the beginning of the learning correspond exactly to the SVD modes of the data, see fig. 3. In this figure, we notice that the modes of the W matrix are the same as the ones of the data at the beginning of the learning as predicted by the linear theory.

After many epochs, we observe in fig. 3(f) that non-linear effects have deformed the SVD modes of W with respect to the beginning of the learning. We can also look at the evolution of the eigenvalues of W . In fig. 4 we observe their evolution and when they start to be amplified (or dumped). In the inset, we see how the strongest modes get out of the bulk and increase while the lowest ones are dumped after many epochs. We also observe that the top part of the spectrum of W appears flattened as compared




Fig. 4: (Colour online) Log-log plot of the singular values represented as discrete abscissas (in decreasing order) with their magnitude reported on the ordinates. The RBM contained 400 hidden variables. A cutoff is highlighted by the onset of the linear behaviour and the SVD modes of the data in black. We qualitatively observe that beyond some α_{tresh} the modes are dumped while before they are amplified. In the inset, the time evolution of the modes 1, 2, 10, 100, 350, 400 during the learning as a function of the number of epochs, we see that for large value of α , the modes are decreasing. We observe that the linear cutoff (around $\alpha \approx 50$ seems different from the one observed when going deep into the non-linear regime ($\alpha \approx 250$).

to the empirical SVD spectrum. This presumably favors the expression of many states of similar free energy related to various digit configurations, able to contribute to the RBM response term in (6).

Discussion. – The equations obtained for the dynamics and the mean-field theory that allows us to compute them constitute a phenomenological description of the learning of an RBM. This is assumed to represent a typical learning trajectory in the limit of infinite batch size. These equations have been obtained by averaging over the components of left and right SVD vectors of the weight matrix, keeping fixed a certain number of quantities considered to be the relevant ones, fully characterizing a typical RBM during the learning process. This averaging corresponds actually to a standard self-averaging assumption in a RS phase. The singular values spectrum $\{w_\alpha\}$ is playing the main role. The projections $(\eta_\alpha, \theta_\alpha)$ of the bias onto the eigenmodes of W are also considered as intrinsic quantities. Finally the rotation vectors $\{\Omega_{\alpha,\beta}^{v,h}\}$ give us the relative motion of the data with respect to the time-dependent frame given by the singular vectors of W . In our phenomenological description the learning dynamics is represented by a trajectory of $\{w_\alpha(t), \eta_\alpha(t), \theta_\alpha(t), \Omega_{\alpha,\beta}^{v,h}(t)\}$ which is uniquely determined by our equations once an initial condition specified by the decomposition of the data on the singular vectors of W is given. By contrast to usual approaches which rely on the teacher-student scenario, we may obtain generic learning curves of non-linear neural networks, which are driven by intrinsic properties of the data. The point is to give insights into the relationship between model and data. This allows us to give

some elements for understanding of which properties of the data drive the learning and how they are represented in the model. Eventually this will lead us to identify and correct some flaws of the present learning methods.

* * *

We thank B. SEOANE helping us improving the quality of the letter.

REFERENCES

- [1] SMOLENSKY P., in *Parallel Distributed Processing*, Vol. 1, edited by RUMELHART D. and MCLELLAND J. (MIT Press) 1986, Chapt. 6: *Information Processing in Dynamical Systems: Foundations of Harmony Theory*, pp. 194–281.
- [2] HINTON G. E. and SALAKHUTDINOV R. R., *Science*, **313** (2006) 504.
- [3] SALAKHUTDINOV R. and HINTON G., *PMLR*, **5** (2009) 448.
- [4] HINTON G. E., *Neural Comput.*, **14** (2002) 1771.
- [5] HINTON G. E., *A Practical Guide to Training Restricted Boltzmann Machines* (Springer, Berlin, Heidelberg) 2012, pp. 599–619.
- [6] HOPFIELD J. J., *Proc. Nat. Acad. Sci. U.S.A.*, **79** (1982) 2554.
- [7] AMIT D. J., GUTFREUND H. and SOMPOLINSKY H., *Ann. Phys.*, **173** (1987) 30.
- [8] GARDNER E., *Europhys. Lett.*, **4** (1987) 481.
- [9] GARDNER E. and DERRIDA B., *J. Phys. A: Math. Gen.*, **21** (1988) 271.
- [10] GABRIÉ M., TRAMEL E. W. and KRZAKALA F., *Training restricted Boltzmann machines via the Thouless-Anderson-Palmer free energy*, in *Proceedings of the 28th International Conference on Neural Information Processing Systems, NIPS'15* (Curran Associates, Inc.) 2015, pp. 640–648.
- [11] HUANG H. and TOYOIZUMI T., *Phys. Rev. E*, **91** (2015) 050101.
- [12] TAKAHASHI C. and YASUDA M., *J. Phys. Soc. Jpn.*, **85** (2016) 034001.
- [13] HUANG H., *J. Stat. Mech.: Theory Exp.* (2017) 053302.
- [14] BARRA A., GENOVESE G., SOLLICH P. and TANTARI D., *Phase diagram of restricted Boltzmann machines and generalized Hopfield networks with arbitrary priors*, arXiv:1702.05882 (2017).
- [15] MONASSON R. and TUBIANA J., *Phys. Rev. Lett.*, **118** (2017) 138301.
- [16] ZDEBOROVÁ L. and KRZAKALA F., *Adv. Phys.*, **65** (2016) 453.
- [17] TIPPING M. E. and BISHOP C. M., *Neural Comput.*, **11** (1999) 443.
- [18] BOURLARD H. and KAMP Y., *Biol. Cybern.*, **59** (1988) 291.
- [19] SAXE A. M., MCCLELLAND J. L. and GANGULI S., *Exact solutions to the nonlinear dynamics of learning in deep linear neural networks*, arXiv:1312.6120 (2014).
- [20] TRAMEL E. W., GABRIÉ M., MANOEL A., CALTAGIRONE F. and KRZAKALA F., *A Deterministic and Generalized Framework for Unsupervised Learning with Restricted Boltzmann Machines*, arXiv:1702.03260 (2017).
- [21] BARRA A., BERNACCHIA A., SANTUCCI E. and CONTUCCI P., *Neural Netw.*, **34** (2012) 1.
- [22] HOHENBERG P. C. and CROSS M. C., *An Introduction to Pattern Formation in Nonequilibrium Systems* (Springer, Berlin, Heidelberg) 1987, pp. 55–92.
- [23] PARISI G. and POTTERS M., *J. Phys. A: Math. Gen.*, **28** (1995) 5267.
- [24] OPPER M. and WINTHER O., *Phys. Rev. E*, **64** (2001) 056131.
- [25] MÉZARD M., *Phys. Rev. E*, **95** (2017) 022117.
- [26] AMIT D. J., GUTFREUND H. and SOMPOLINSKY H., *Phys. Rev. A*, **32** (1985) 1007.
- [27] MÉZARD M., PARISI G. and VIRASORO M. A., *Spin Glass Theory and Beyond* (World Scientific, Singapore) 1987.
- [28] ALMEIDA J. R. L. and THOULESS D. J., *J. Phys. A: Math. Gen.*, **11** (1978) 983.
- [29] DECELLE A., FISSORE G. and FURTELEHNER C., in preparation.
- [30] AGLIARI E., BARRA A., GALLUZZI A., GUERRA F. and MOAURO F., *Phys. Rev. Lett.*, **109** (2012) 268101.