

Expressive Power and Approximation Errors of Restricted Boltzmann Machines

Guido F. Montúfar
Johannes Rauh
Nihat Ay

SFI WORKING PAPER: 2011-09-041

SFI Working Papers contain accounts of scientific work of the author(s) and do not necessarily represent the views of the Santa Fe Institute. We accept papers intended for publication in peer-reviewed journals or proceedings volumes, but not papers that have already appeared in print. Except for papers by our external faculty, papers must be based on work done at SFI, inspired by an invited visit to or collaboration at SFI, or funded by an SFI grant.

©NOTICE: This working paper is included by permission of the contributing author(s) as a means to ensure timely distribution of the scholarly and technical work on a non-commercial basis. Copyright and all rights therein are maintained by the author(s). It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may be reposted only with the explicit permission of the copyright holder.

www.santafe.edu



SANTA FE INSTITUTE

Expressive Power and Approximation Errors of Restricted Boltzmann Machines

Guido F. Montufar¹, Johannes Rauh¹, Nihat Ay^{1,2}

¹Max Planck Institute for Mathematics in the Sciences, Inselstraße 22 04103 Leipzig, Germany

²Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA

May 31, 2011

Abstract

We present explicit classes of probability distributions that can be learned by Restricted Boltzmann Machines (RBMs) depending on the number of units that they contain, and which are representative for the expressive power of the model. We use this to show that the maximal Kullback-Leibler divergence to the RBM model with n visible and m hidden units is bounded from above by $(n - 1) - \log(m + 1)$. In this way we can specify the number of hidden units that guarantees a sufficiently rich model containing different classes of distributions and respecting a given error tolerance.

1 Introduction

A Restricted Boltzmann Machine (RBM) [23, 10] is a learning system consisting of two layers of binary stochastic units, a hidden layer and a visible layer, with a complete bipartite interaction graph. It is used as a generative model for probability distributions on binary vectors. RBMs can be trained in an unsupervised way and more efficiently than general Boltzmann Machines, which are not restricted to have a bipartite interaction graph [11, 6]. Furthermore, they can be used as building blocks to progressively train and study deep learning systems [13, 4, 16, 21]. Hence, RBMs have received increasing attention in the past years.

An RBM with n visible and m hidden units generates a stationary distribution on the states of the visible units which has the following form:

$$p_{W,C,B}(v) = \frac{1}{Z_{W,C,B}} \sum_{h \in \{0,1\}^m} \exp\left(h^\top W v + C^\top h + B^\top v\right) \quad \forall v \in \{0,1\}^n,$$

where $h \in \{0, 1\}^m$ denotes a state vector of the hidden units, $W \in \mathbf{R}^{m \times n}$, $C \in \mathbf{R}^m$ and $B \in \mathbf{R}^n$ constitute the model parameters, and $Z_{W,C,B}$ is a corresponding normalization constant. In the sequel we denote by $\text{RBM}_{n,m}$ the set of all probability distributions on $\{0, 1\}^n$ which can be approximated arbitrarily well by a visible distribution generated by the RBM with m hidden and n visible units for an appropriate choice of the parameter values.

In [15] it was shown that $\text{RBM}_{n,m}$ contains any probability distribution if $m \geq 2^n + 1$. In [21] it was shown that already $m \geq 2^{n-1} - 1$ suffices. On the other hand, an RBM containing all distributions on $\{0, 1\}^n$ must have at least $2^n - 1$ parameters, and thus at least $\lceil 2^n / (n + 1) \rceil - 1$ hidden units [21]. In fact, in [8] it was shown that for most combinations of m and n the dimension of $\text{RBM}_{n,m}$ equals either the number of parameters or $2^n - 1$, whatever is smaller. However, the geometry of $\text{RBM}_{n,m}$ is intricate, and even an RBM of dimension $2^n - 1$ is not guaranteed to contain all visible distributions, see [20] for counterexamples.

In summary, an RBM that can approximate any distribution arbitrarily well must have a very large number of parameters and hidden units. In praxis, training such a large system is not desirable or even possible. However, there are at least two reasons why in many cases this is not necessary:

- An appropriate approximation of distributions is sufficient for most purposes.
- The interesting distributions the system shall simulate belong to a small class of distributions. Therefore, the model does not need to approximate all distributions.

For example, the set of optimal policies in reinforcement learning [24], the set of dynamics kernels which maximize predictive information in robotics [25], or the information flow in neural networks [3] are contained in very low dimensional manifolds. As was shown in [2], all of these are contained in 2-dimensional subsets. On the other hand, usually it is very hard to mathematically describe a set containing the optimal solutions to general problems, or a set of interesting probability distributions (for example the class of distributions generating natural images). Furthermore, although RBMs are parametric models and for any choice of the parameters we have a resulting probability distribution, in general it is difficult to explicitly specify this resulting probability distribution (or even to estimate it [18]). Due to these difficulties the number of hidden units m is often chosen on the basis of experience [12], or m is considered as a hyperparameter which is optimized by extensive search, depending on the distributions to be simulated by the RBM.

In this paper we give an explicit description of classes of distributions that are contained in $\text{RBM}_{n,m}$, and which are representative for the expressive power of this model. Using this description, we estimate the maximal

Kullback-Leibler divergence between an arbitrary probability distribution and the best approximation within $\text{RBM}_{n,m}$.

This paper is organized as follows: Section 2 discusses the different kinds of errors that appear when an RBM learns. Section 3 introduces the statistical models studied in this paper. Section 4 studies submodels of $\text{RBM}_{n,m}$. An upper bound of the approximation error for RBMs is found in Section 5.

2 Approximation Error

When training an RBM to represent a distribution p there are at least three sources of errors:

1. Usually the underlying distribution p is unknown and only a set of samples generated by p is observed. These samples can be represented as an empirical distribution p^{Data} . The difference between p and p^{Data} is the generalization error.
2. The set $\text{RBM}_{n,m}$ does not contain every probability distribution, unless the number of hidden units is very large, as we outlined in the introduction. Therefore, we have an error given by the distance of p^{Data} to the best approximation $p_{\text{RBM}}^{\text{Data}}$ contained in the RBM model.
3. The learning process may yield a solution $\tilde{p}_{\text{RBM}}^{\text{Data}}$ in RBM which is not the optimal approximation $p_{\text{RBM}}^{\text{Data}}$ in terms of Kullback-Leibler divergence. This can occur if the learning algorithm gets trapped in a local optimum, or also if it optimizes an objective different from Maximum Likelihood, e.g. contrastive divergence (CD), see [6].

In this paper we study the expressive power of the RBM model and the Kullback-Leibler divergence from an arbitrary distribution to its best representation within the RBM model, this accounts for the first and second items in the above list. Estimating the approximation error is difficult because the explicit geometry of the RBM model is unknown. Our strategy is to find subsets $\mathcal{M} \subseteq \text{RBM}_{n,m}$ that are easy to describe. Then the maximal error when approximating probability distributions with an RBM is upper bounded by the maximal error when approximating with \mathcal{M} .

Consider a finite set \mathcal{X} . A real valued function on \mathcal{X} can be seen as a real vector with $|\mathcal{X}|$ entries. The set $\mathcal{P} = \mathcal{P}(\mathcal{X})$ of all probability distributions on \mathcal{X} is a $(|\mathcal{X}| - 1)$ -dimensional simplex in $\mathbf{R}^{|\mathcal{X}|}$. There are several notions of distance between probability distributions, and in turn for the error in the representation (approximation) of a probability distribution. One possibility is to use the induced distance of the Euclidian space $\mathbf{R}^{|\mathcal{X}|}$. From the point of

view of information theory, a more meaningful distance notion for probability distributions is the Kullback-Leibler divergence:

$$D(q\|p) := \sum_x q(x) \log \frac{q(x)}{p(x)}.$$

In this paper we use the basis 2 logarithm. The Kullback-Leibler (KL) divergence is non-negative and vanishes if and only if $q = p$. If the support of p does not contain the support of q it is defined as ∞ . The summands with $q(x) = 0$ are set to 0. The KL-divergence is not symmetric, but it has nice information theoretic properties [14, 7].

If $\mathcal{E} \subseteq \mathcal{P}$ is a statistical model and if $p \in \mathcal{P}$, then any probability distribution $p_{\mathcal{E}} \in \overline{\mathcal{E}}$ satisfying

$$D(p\|p_{\mathcal{E}}) = D(p\|\mathcal{E}) := \min\{D(p\|q) : q \in \overline{\mathcal{E}}\}$$

is called a (*generalized*) *reversed information projection*, or *rI*-projection. Here, $\overline{\mathcal{E}}$ denotes the closure of \mathcal{E} . If p is an empirical distribution, then one can show that any *rI*-projection is a maximum likelihood estimate.

In order to assess an RBM or some other model \mathcal{M} we use the maximal approximation error with respect to the KL-divergence when approximating arbitrary probability distributions using \mathcal{M} :

$$D_{\mathcal{M}} := \max\{D(p\|\mathcal{M}) : p \in \mathcal{P}\}.$$

For example, the maximal KL-divergence to the uniform distribution $\frac{1}{|\mathcal{X}|}$ is attained by the Dirac delta distributions δ_x , $x \in \mathcal{X}$, and amounts to:

$$D_{\left\{\frac{1}{|\mathcal{X}|}\right\}} = D(\delta_x\|\frac{1}{|\mathcal{X}|}) = \log |\mathcal{X}|. \quad (1)$$

3 Model Classes

3.1 Exponential families and product measures

In this work we only need a restricted class of exponential families, namely exponential families on a finite set with uniform reference measure. See [5] for more on exponential families.

Let $A \in \mathbf{R}^{d \times |\mathcal{X}|}$ be a matrix. The columns A_x of A will be indexed by $x \in \mathcal{X}$. The rows of A can be interpreted as functions on \mathbf{R} . The *exponential family* \mathcal{E}_A with *sufficient statistics* A consists of all probability distributions of the form p_{λ} , $\lambda \in \mathbf{R}^d$, where

$$p_{\lambda}(x) = \frac{\exp(\lambda^{\top} A_x)}{\sum_x \exp(\lambda^{\top} A_x)}, \quad \text{for all } x \in \mathcal{X}.$$

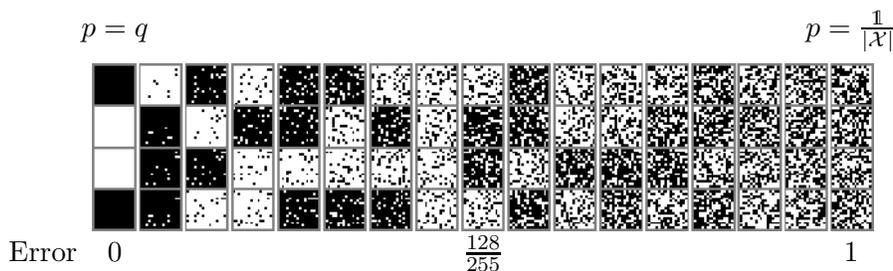


Figure 1: This figure gives an intuition on what the size of an error means. Every column shows four samples drawn from p , the best approximation of the distribution $q = \frac{1}{2}(\delta_{(1\dots 1)} + \delta_{(0\dots 0)})$ within a partition model with 2 randomly chosen cubical components, containing $(0\dots 0)$ and $(1\dots 1)$, of cardinality from 1 (first column) to $\frac{|\mathcal{X}|}{2}$ (last column). As a measure of error ranging from 0 to 1 we take $D(q\|p)/D(q\|\frac{1}{|\mathcal{X}|})$. The last column shows samples from the uniform distribution, which is, in particular, the best approximation of q within $\text{RBM}_{n,0}$. Note that an RBM with 1 hidden unit can approximate q with arbitrary accuracy, see [21] and Theorem 4.1 in the present paper.

Note that any probability distribution in \mathcal{E}_A has full support. Furthermore, \mathcal{E}_A is in general not a closed set. The closure $\overline{\mathcal{E}_A}$ (with respect to the usual topology on $\mathbf{R}^{\mathcal{X}}$) will be important in the following.

There are many links between exponential families and the KL-divergence: Let \mathcal{E} be an exponential family. For any $p \in \mathcal{P}$ there exists a unique rI -projection $p_{\mathcal{E}}$ to \mathcal{E} . The rI -projection $p_{\mathcal{E}}$ is characterized as the unique probability distribution in $\overline{\mathcal{E}} \cap (p + \ker A)$.

The most important exponential families in this work are the independence models. The *independence model* of n binary random variables consists of all probability distributions on $\{0, 1\}^n$ that factorize:

$$\overline{\mathcal{E}_n} = \left\{ p \in \mathcal{P}(\mathcal{X}) : p(x_1, \dots, x_n) = \prod_{i=1}^n p_i(x_i) \text{ for some } p_i \in \mathcal{P}(\{0, 1\}) \right\}.$$

It is the closure of an n -dimensional exponential family \mathcal{E}_n . This model corresponds to the RBM model with no hidden units. An element of the independence model is called a *product distribution*.

Lemma 3.1 (Corollary 4.1 of [1]) *Let $\overline{\mathcal{E}_n}$ be the independence model on $\{0, 1\}^n$. If $n > 0$, then $D_{\mathcal{E}_n} = (n - 1)$. The global maximizers are the distribution of the form $\frac{1}{2}(\delta_x + \delta_y)$, where $x, y \in \{0, 1\}^n$ satisfy $x_i + y_i = 1$ for all i .*

This result should be compared with (1). Although the independence model is much larger than the set $\{\frac{1}{|\mathcal{X}|}\}$, the maximal divergence decreases

only by 1. In fact, if \mathcal{E} is any exponential family of dimension k , then $D_{\mathcal{E}} \geq \log(|\mathcal{X}|/(k+1))$, see [22]. Thus, this notion of distance is rather strong.

3.2 Partition models and mixtures of products with disjoint supports

The *mixture* of m models $\mathcal{M}_1, \dots, \mathcal{M}_m \subseteq \mathcal{P}$ is the set of all convex combinations

$$p = \sum_i \alpha_i p_i, \text{ where } p_i \in \mathcal{M}_i, \alpha_i \geq 0, \sum_i \alpha_i = 1. \quad (2)$$

In general, mixture models are complicated objects. Even if all models $\mathcal{M}_1 = \dots = \mathcal{M}_m$ are equal, it is difficult to describe the mixture [17, 19]. The situation simplifies considerably if the models have disjoint supports. Note that given any partition $\xi = \{\mathcal{X}_1, \dots, \mathcal{X}_m\}$ of \mathcal{X} , any $p \in \mathcal{P}$ can be written as $p(x) = p^{\mathcal{X}_i}(x)p(\mathcal{X}_i)$ for all $x \in \mathcal{X}_i$ and $i \in \{1, \dots, m\}$, where $p^{\mathcal{X}_i}$ is a probability measure in $\mathcal{P}(\mathcal{X}_i)$ for all i .

Lemma 3.2 *Let $\xi = \{\mathcal{X}_1, \dots, \mathcal{X}_m\}$ be a partition of \mathcal{X} and let $\mathcal{M}_1, \dots, \mathcal{M}_m$ be statistical models such that $\mathcal{M}_i \subseteq \mathcal{P}(\mathcal{X}_i)$. Consider any $p \in \mathcal{P}$ and corresponding $p^{\mathcal{X}_i}$ such that $p(x) = p^{\mathcal{X}_i}(x)p(\mathcal{X}_i)$ for $x \in \mathcal{X}_i$. Let p_i be an rI -projection of $p^{\mathcal{X}_i}$ to \mathcal{M}_i . Then the rI -projection $p_{\mathcal{M}}$ of P to the mixture \mathcal{M} of $\mathcal{M}_1, \dots, \mathcal{M}_m$ satisfies*

$$p_{\mathcal{M}}(x) = p(\mathcal{X}_i)p_i(x), \quad \text{whenever } x \in \mathcal{X}_i.$$

Therefore, $D_{\mathcal{M}} = \max_{i=1, \dots, m} D_{\mathcal{M}_i}$.

Proof Let $p \in \mathcal{M}$ be as in (2). Then $D(q||p) = \sum_{i=1}^m q(\mathcal{X}_i)D(q^{\mathcal{X}_i}||p_i)$ for all $q \in \mathcal{P}$. For fixed q this sum is minimal if and only if each term is minimal. \square

If each \mathcal{M}_i is an exponential family, then the mixture is also an exponential family (this is not true if the supports of the models \mathcal{M}_i are not disjoint). In the rest of this section we discuss two examples.

If each \mathcal{M}_i equals the set containing just the uniform distribution on \mathcal{X}_i , then \mathcal{M} is called the *partition model* of ξ , denoted with \mathcal{P}_{ξ} . The partition model \mathcal{P}_{ξ} is given by all distributions with constant value on each block \mathcal{X}_i , i.e. those that satisfy $p(x) = p(y)$ for all $x, y \in \mathcal{X}_i$. This is the closure of the exponential family with sufficient statistics

$$A_x = (\chi_1(x), \chi_2(x), \dots, \chi_d(x))^{\top},$$

where $\chi_i := \chi_{\mathcal{X}_i}$ is 1 on $x \in \mathcal{X}_i$, and 0 everywhere else. See [22] for interesting properties of partition models.

The partition models include for example the set of finite exchangeable distributions, (see e.g. [9]), where the partition components are given by the

sets of binary vectors which have the same number of entries equal to one. The probability of a vector v depends only on the number of ‘successes’ 1, but not on their position.

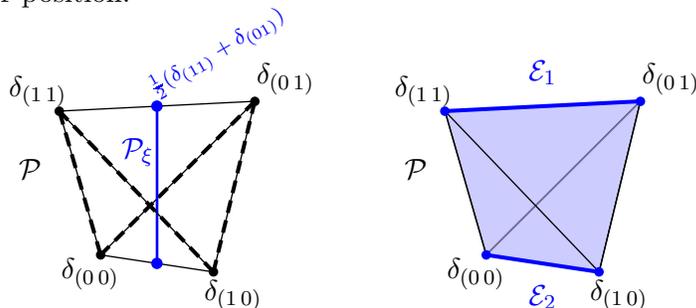


Figure 2: Left: The simplex of distributions on two binary variables. The blue line represents the partition model \mathcal{P}_ξ with partition $\xi = \{(11), (01)\} \cup \{(00), (10)\}$. The dashed lines represent the set of KL-divergence maximizers for \mathcal{P}_ξ . Right: The mixture of the product distributions \mathcal{E}_1 and \mathcal{E}_2 with disjoint supports on $\{(11), (01)\}$ and $\{(00), (10)\}$ corresponding to the same partition ξ equals the whole simplex \mathcal{P} .

Corollary 3.1 *Let $\xi = \{\mathcal{X}_1, \dots, \mathcal{X}_m\}$ be a partition of \mathcal{X} . Then*

$$D_{\mathcal{P}_\xi} = \max_{i=1, \dots, m} \log |\mathcal{X}_i|.$$

Now assume that $\mathcal{X} = \{0, 1\}^n$ is the set of binary vectors of length n . As a subset of \mathbf{R}^n it consists of the vertices (extreme points) of the n -dimensional hypercube. The vertices of a k -dimensional face of the n -cube are given by fixing the values of x in $n - k$ positions:

$$\{x \in \{0, 1\}^n : x_i = \tilde{x}_i, \forall i \in I, \text{ for some } I \subseteq \{1, \dots, n\}, |I| = n - k\}$$

We call such a subset $\mathcal{Y} \subseteq \mathcal{X}$ *cubical* or a *face* of the n -cube. A cubical subset of cardinality 2^k can be naturally identified with $\{0, 1\}^k$. This identification allows to define independence models and product measures on $\mathcal{P}(\mathcal{Y}) \subseteq \mathcal{P}(\mathcal{X})$. Note that product measures on \mathcal{Y} are also product measures on \mathcal{X} , and the independence model on \mathcal{Y} is a subset of the independence model on \mathcal{X} .

Corollary 3.2 *Let $\xi = \{\mathcal{X}_1, \dots, \mathcal{X}_m\}$ be a partition of $\mathcal{X} = \{0, 1\}^n$ into cubical sets. For any i let \mathcal{E}_i be the independence model on \mathcal{X}_i , and let \mathcal{M} be the mixture of $\mathcal{E}_1, \dots, \mathcal{E}_m$. Then*

$$D_{\mathcal{M}} = \max_{i=1, \dots, m} \log(|\mathcal{X}_i|) - 1.$$

See Figure 1 for an intuition on the approximation error of partition models, and see Figure 2 for small examples of a partition model and of a mixture of products with disjoint support.

4 Classes of distributions that RBMs can learn

Consider a set $\xi = \{\mathcal{X}_i\}_{i=1}^m$ of m disjoint cubical sets \mathcal{X}_i in \mathcal{X} . Such a ξ is a partition of some subset $\cup \xi = \cup_i \mathcal{X}_i$ of \mathcal{X} into m disjoint cubical sets. We write G_m for the collection of all such partitions. We have the following result:

Theorem 4.1 *RBM $_{n,m}$ contains the following distributions:*

- *Any mixture of one arbitrary product distribution, $m - k$ product distributions with support on arbitrary but disjoint faces of the n -cube, and k arbitrary distributions with support on any edges of the n -cube, for any $0 \leq k \leq m$. In particular:*
- *Any mixture of $m+1$ product distributions with disjoint cubical supports. In consequence also: Any distribution belonging to any partition model in G_{m+1} .*

Restricting the cubical sets of the second item to edges, i.e. pairs of vectors differing in one entry, we see that the above theorem implies the following previously known result, which was shown in [21] based on [15]:

Corollary 4.1 *RBM $_{n,m}$ contains the following distributions:*

- *Any distribution with a support set that can be covered by $m+1$ pairs of vectors differing in one entry. In particular, this includes:*
- *Any distribution in \mathcal{P} with a support of cardinality smaller than or equal to $m+1$.*

Corollary 4.1 implies that an RBM with $m \geq 2^{n-1} - 1$ hidden units is a universal approximator of distributions on $\{0,1\}^n$, i.e. can approximate any distribution to an arbitrarily good accuracy.

Assume $m+1 = 2^k$ and let ξ be a partition of \mathcal{X} into $m+1$ disjoint cubical sets of equal size. Let us denote by $\mathcal{P}_{\xi,1}$ the set of all distributions which can be written as a mixture of $m+1$ product distributions with support on the elements of ξ . The dimension of $\mathcal{P}_{\xi,1}$ is given by

$$\begin{aligned} \dim \mathcal{P}_{\xi,1} &= (m+1) \log \left(\frac{2^n}{m+1} \right) + m+1 + n \\ &= (m+1) \cdot n + (m+1) + n - (m+1) \log(m+1) . \end{aligned}$$

The dimension of the set of visible distribution represented by an RBM is at most equal to the number of paramters, see [21], this is $m \cdot n + m + n$. This means that the class given above has roughly the same dimension of the set of distributions that can be represented. In fact,

$$\dim \mathcal{P}_{\xi,1} - \dim \text{RBM}_{m-1} = n + 1 - (m+1) \log(m+1) .$$

This means that the class of distributions $\mathcal{P}_{\xi,1}$ which by Theorem 4.1 can be represented by $\text{RBM}_{n,m}$ is not contained in $\text{RBM}_{n,m-1}$ when $(m+1)^{m+1} \leq 2^{n+1}$.

Proof of Theorem 4.1 The proof draws on ideas from [15] and [21]. An RBM with no hidden units can represent precisely the independence model, i.e. all product distributions, and in particular any uniform distribution on a face of the n -cube.

Consider an RBM with $m-1$ hidden units. For any choice of the parameters $W \in \mathbf{R}^{m-1 \times n}$, $B \in \mathbf{R}^n$, $C \in \mathbf{R}^{m-1}$ we can write the resulting distribution on the visible units as:

$$p(v) = \frac{\sum_h z(v, h)}{\sum_{v', h'} z(v', h')}, \quad (3)$$

where $z(v, h) = \exp(hWv + Bv + Ch)$. Appending one additional hidden unit, with connection weights w to the visible units and bias c , produces a new distribution which can be written as follows:

$$p_{w,c}(v) = \frac{(1 + \exp(wv + c)) \sum_h z(v, h)}{\sum_{v', h'} (1 + \exp(wv' + c)) z(v', h')}.$$

Consider now any set $I \subseteq [n] := \{1, \dots, n\}$ and an arbitrary visible vector $u \in \mathcal{X}$. The values of u in the positions $[n] \setminus I$ define a face $F := \{v \in \mathcal{X} : v_i = u_i, \forall i \notin I\}$ of the n -cube of dimension $|I|$. Let $\mathbf{1} := (1, \dots, 1) \in \mathbf{R}^n$ and denote by $u^{I,0}$ the vector with entries $u_i^{I,0} = u_i, \forall i \notin I$ and $u_i^{I,0} = 0, \forall i \in I$. Let $\lambda^I \in \mathbf{R}^n$ with $\lambda_i^I = 0, \forall i \notin I$ and let $\lambda_c, a \in \mathbf{R}$. Define the connection weights w and c as follows:

$$\begin{aligned} w &= a(u^{I,0} - \frac{1}{2}\mathbf{1}^{I,0}) + \lambda^I, \\ c &= -a(u^{I,0} - \frac{1}{2}\mathbf{1}^{I,0})^\top u + \lambda_c. \end{aligned}$$

For this choice and $a \rightarrow \infty$ equation (4) yields:

$$p_{w,c}(v) = \begin{cases} \frac{p(v)}{1 + \sum_{v' \in F} \exp(\lambda^I \cdot v' + \lambda_c) p(v')}, & \forall v \notin F \\ \frac{(1 + \exp(\lambda^I \cdot v + \lambda_c)) p(v)}{1 + \sum_{v' \in F} \exp(\lambda^I \cdot v' + \lambda_c) p(v')}, & \forall v \in F \end{cases}. \quad (4)$$

If the initial p from equation (3) is such that its restriction to F is a product distribution, then $p(v) = K \exp(\eta^I \cdot v), \forall v \in F$, where K is a constant and η^I is a vector with $\eta_i^I = 0, \forall i \notin I$. We can choose $\lambda^I = \beta^I - \eta^I$, and $\exp(\lambda_c) = \alpha \frac{1}{K \sum_{v \in F} \exp(\beta^I \cdot v)}$. For this choice, equation (4) yields:

$$p_{w,c} = (\alpha - 1)p + \alpha \hat{p},$$

where \hat{p} is a product distribution with support in F and arbitrary natural parameters β^f , and α is an arbitrary mixture weight in $[0, 1]$. Finally, the product distributions on edges of the cube are arbitrary, see [19] or [21] for details, and hence the restriction of any p to any edge is a product distribution. \square

5 Maximal Approximation Errors of RBMs

By Theorem 4.1, $\text{RBM}_{n,m}$ contains the union of many exponential families. The results from Section 3 give immediate bounds on $D_{\text{RBM}_{n,m}}$: Let $m < 2^{n-1} - 1$. All partition models for partitions of $\{0, 1\}^n$ into $m + 1$ cubical sets are contained in $\text{RBM}_{n,m}$. Applying Corollary 3.1 to such a partition where the cardinality of all blocks is $2^{n - \lfloor \log(m+1) \rfloor}$ or less yields the bound $D_{\text{RBM}_{n,m}} \leq n - \lfloor \log(m+1) \rfloor$. Similarly, Theorem 4.1 and Corollary 3.2 imply $D_{\text{RBM}_{n,m}} \leq n - 1 - \lfloor \log(m+1) \rfloor$. In this section we derive an improved bound which strictly decreases, as m increases, until 0 is reached.

Theorem 5.1 *Let $m \leq 2^{n-1} - 1$. Then the maximal Kullback-Leibler divergence from any distribution on $\{0, 1\}^n$ to $\text{RBM}_{n,m}$ is not larger than*

$$\max_{p \in \mathcal{P}} D(p \| \text{RBM}_{n,m}) \leq (n - 1) - \log(m + 1).$$

Conversely, given an error tolerance $0 \leq \epsilon \leq 1$, the choice $m \geq 2^{(n-1)(1-\epsilon)} - 1$ ensures a sufficiently rich RBM model that satisfies $D_{\text{RBM}_{n,m}} \leq \epsilon D_{\text{RBM}_{n,0}}$.

For $m = 2^{n-1} - 1$ the error vanishes, corresponding to the fact that an RBM with that many hidden units is a universal approximator. In Figure 3 we use computer experiments to illustrate Theorem 5.1. The proof makes use of the following lemma:

Lemma 5.1 *Let $n_1, \dots, n_m \geq 0$ such that $2^{n_1} + \dots + 2^{n_m} = 2^n$. Let \mathcal{M} be the union of all mixtures of independent models corresponding to all cubical partitions of \mathcal{X} into blocks of cardinalities $2^{n_1}, \dots, 2^{n_m}$. Then $D_{\mathcal{M}} \leq \sum_{i:n_i > 1} \frac{n_i - 1}{2^{n - n_i}}$.*

Proof of Lemma 5.1 The proof is by induction on n . If $n = 1$, then $m = 1$ or $m = 2$, and in both cases it is easy to see that the inequality holds (both sides vanish). If $n > 1$, then order the n_i such that $n_1 \geq n_2 \geq \dots \geq n_m \geq 0$. Without loss of generality assume $m > 1$.

Let $p \in \mathcal{P}(\mathcal{X})$, and let \mathcal{Y} be a cubical subset of \mathcal{X} of cardinality 2^{n-1} such that $p(\mathcal{Y}) \leq \frac{1}{2}$. Since the numbers $2^{n_1} + \dots + 2^{n_i}$ for $i = 1, \dots, m$ contain all multiples of 2^{n_1} up to 2^n and $2^n/2^{n_1}$ is even, there exists k such that $2^{n_1} + \dots + 2^{n_k} = 2^{n-1} = 2^{n_{k+1}} + \dots + 2^{n_m}$.

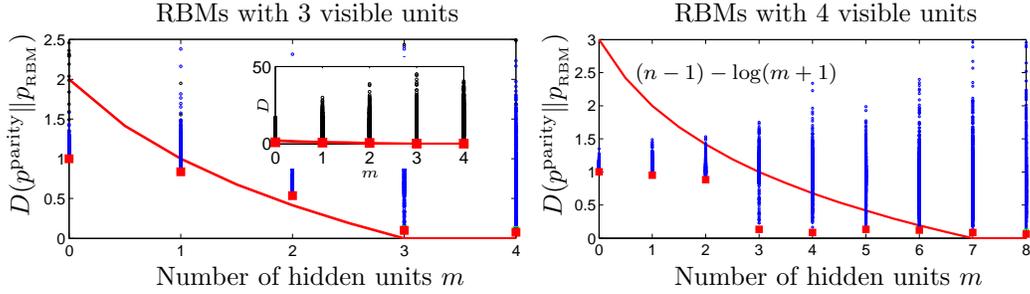


Figure 3: This figure demonstrates our results for $n = 3$ and $n = 4$ visible units. The red curves represent the bounds from Theorem 5.1. We fixed p^{parity} as target distribution, the uniform distribution on binary length n vectors with an even number of ones. The distribution p^{parity} is not the KL-maximizer from $\text{RBM}_{n,m}$, but it is in general difficult to represent. Qualitatively, samples from p^{parity} look like uniformly distributed, and representing p^{parity} requires the maximal number of product mixture components [20, 19]. For both values of n and each $m = 0, \dots, 2^n/2$ we initialized 500 resp. 1000 RBMs at parameter values chosen uniformly at random in the range $[-10, 10]$. The inset of the left figure shows the resulting KL-divergence $D(p^{\text{parity}} \| p_{\text{RBM}}^{\text{rand}})$ (for $n = 4$ the resulting KL-divergence was larger). Randomly chosen distributions in $\text{RBM}_{n,m}$ are likely to be very far from the target distribution. We trained these randomly initialized RBMs using CD for 500 training epochs, learning rate 1 and a list of even parity vectors as training data. The result after training is given by the blue circles. After training the RBMs the result is often not better than the uniform distribution, for which $D(p^{\text{parity}} \| \frac{\mathbb{1}}{|\{0,1\}^n|}) = 1$. For each m , the best set of parameters after training was used to initialize a further CD training with a smaller learning rate (green squares, mostly covered) followed by a short maximum likelihood gradient ascent (red filled squares).

Let \mathcal{M}' be the union of all mixtures of independence models corresponding to all cubical partitions Ξ of \mathcal{X} into blocks of cardinalities n_1, \dots, n_m such that $\Xi^1 \cup \dots \cup \Xi^k = \mathcal{Y}$. In the following, the symbol \sum'_i shall denote summation over all indices i such that $n_i > 1$. By induction

$$D(p \| \mathcal{M}) \leq D(p \| \mathcal{M}') \leq p(\mathcal{Y}) \sum'_{i=1}^k \frac{n_i - 1}{2^{n-1-n_i}} + p(\mathcal{X} \setminus \mathcal{Y}) \sum'_{j=k+1}^m \frac{n_j - 1}{2^{n-1-n_j}}. \quad (5)$$

There exist $j_1 = k + 1 < j_2 < \dots < j_k \leq j_{k+1} - 1 = m$ such that $2^{n_i} = 2^{n_{j_i}} + \dots + 2^{n_{j_{i+1}-1}}$ for all $i \leq k$. Note that

$$\sum'_{j=j_i}^{j_{i+1}} \frac{n_j - 1}{2^{n-1-n_j}} \leq \frac{n_i - 1}{2^{n-1}} (2^{n_{j_i}} + \dots + 2^{n_{j_{i+1}-1}}) = \frac{n_i - 1}{2^{n-1-n_i}},$$

and therefore

$$\left(\frac{1}{2} - p(\mathcal{Y})\right) \frac{n_i - 1}{2^{n-1-n_i}} + \left(\frac{1}{2} - p(\mathcal{X} \setminus \mathcal{Y})\right) \sum_{j=j_i}^{j_{i+1}-1} \frac{n_j - 1}{2^{n-1-n_j}} \geq 0.$$

Adding these terms for $i = 1, \dots, k$ to the right hand side of equation (5) yields

$$D(p \parallel \mathcal{M}) \leq \frac{1}{2} \sum_{i=1}^k \frac{n_i - 1}{2^{n-1-n_i}} + \frac{1}{2} \sum_{j=k+1}^m \frac{n_j - 1}{2^{n-1-n_j}},$$

from which the assertions follow. \square

Proof of Theorem 5.1 From Theorem 4.1 we know that $\text{RBM}_{n,m}$ contains the union \mathcal{M} of all mixtures of independent models corresponding to all partitions with up to $m + 1$ cubical blocks. Hence, $D_{\text{RBM}_{n,m}} \leq D_{\mathcal{M}}$. Let $k = n - \lceil \log(m+1) \rceil$ and $l = 2m + 2 - 2^{n-k+1} \geq 0$; then $l2^{k-1} + (m+1)2^k = 2^n$. Lemma 5.1 with $n_1 = \dots = n_l = k - 1$ and $n_{l+1} = \dots = n_{m+1} = k$ implies

$$D_{\mathcal{M}} \leq \frac{l(k-2)}{2^{n-k+1}} + \frac{(m+1-l)(k-1)}{2^{n-k}} = k - \frac{m+1}{2^{n-k}}.$$

The assertion follows from $\log(m+1) \leq (n-k) + \frac{m+1}{2^{n-k}} - 1$, where $\log(1+x) \leq x$ for all $x > 0$ was used. \square

6 Conclusion

In this paper we studied the expressive power of the Restricted Boltzmann Machine model depending on the number of visible and hidden units that it contains. We presented a hierarchy of explicit classes of probability distributions that an RBM can represent according to the number n of visible units and the number m of hidden units. These classes include any mixture of an arbitrary product distribution and m further product distributions with disjoint supports. The geometry of these submodels is easier to study than that of the RBM models, while these subsets still capture many of the distributions contained in the RBM models. Using these results we derived bounds for the approximation errors of RBMs. We showed that it is always possible to reduce the error to $(n - 1) - \log(m + 1)$ or less. That is, given any target distribution, there is a distribution within the RBM model for which the Kullback-Leibler divergence between both is not larger than that number. Our results give a theoretical basis for selecting the size of an RBM which accounts for a desired error tolerance.

Computer experiments showed that the bound captures the order of magnitude of the true approximation error, at least for small examples. However, learning may not always find the best approximation, resulting in an error that may well exceed our bound.

References

- [1] N. Ay and A. Knauf. Maximizing multi-information. *Kybernetika*, 42:517–538, 2006.
- [2] N. Ay, G. Montufar, and J. Rauh. Selection criteria for neuromanifolds of stochastic dynamics. *International Conference on Cognitive Neurodynamics*, 2011.
- [3] N. Ay and T. Wennekers. Dynamical properties of strongly interacting Markov chains. *Neural Networks*, 16:1483–1497, 2003.
- [4] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy layer-wise training of deep networks. *NIPS*, 2007.
- [5] L. Brown. *Fundamentals of Statistical Exponential Families: With Applications in Statistical Decision Theory*. Institute of Mathematical Statistics, Hayworth, CA, USA, 1986.
- [6] M. A. Carreira-Perpiñan and G. E. Hinton. On contrastive divergence learning. In *Proceedings of the 10-th International Workshop on Artificial Intelligence and Statistics*, 2005.
- [7] T. M. Cover and J. A. Thomas. *Elements of Information Theory, 2nd Edition*. John Wiley & Sons, 2006.
- [8] M. A. Cueto, J. Morton, and B. Sturmfels. Geometry of the restricted Boltzmann machine. In M. A. G. Viana and H. P. Wynn, editors, *Algebraic methods in statistics and probability II, AMS Special Session*, volume 2. American Mathematical Society, 2010.
- [9] P. Diaconis and D. Freedman. Finite exchangeable sequences. *The Annals of Probability*, 8:745–764, 1980.
- [10] Y. Freund and D. Haussler. Unsupervised learning of distributions on binary vectors using 2-layer networks. *NIPS*, pages 912–919, 1992.
- [11] G. E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14:1771–1800, 2002.
- [12] G. E. Hinton. A practical guide to training restricted Boltzmann machines, version 1. Technical report, UTML2010-003, University of Toronto, 2010.
- [13] G. E. Hinton, S. Osindero, and Y. Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554, 2006.

- [14] S. Kullback and R. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.
- [15] N. Le Roux and Y. Bengio. Representational power of restricted Boltzmann machines and deep belief networks. *Neural Computation*, 20(6):1631–1649, 2008.
- [16] N. Le Roux and Y. Bengio. Deep belief networks are compact universal approximators. *Neural Computation*, 22:2192–2207, 2010.
- [17] B. Lindsay. *Mixture models: theory, geometry, and applications*. NSF-CBMS regional conference series in probability and statistics. Institute of Mathematical Statistics, 1995.
- [18] P. M. Long and R. A. Servedio. Restricted Boltzmann machines are hard to approximately evaluate or simulate. In *Proceedings of the 27-th ICML*, pages 703–710, 2010.
- [19] G. Montufar. Mixture decompositions using a decomposition of the sample space. *ArXiv 1008.0204*, 2010.
- [20] G. Montufar. Mixture models and representational power of RBMs, DBNs and DBMs. *NIPS Deep Learning and Unsupervised Feature Learning Workshop*, 2010.
- [21] G. Montufar and N. Ay. Refinements of universal approximation results for deep belief networks and restricted Boltzmann machines. *Neural Computation*, 23(5):1306–1319, 2011.
- [22] J. Rauh. *Finding the maximizers of the information divergence from an exponential family*. PhD thesis, Universität Leipzig, 2011. Submitted.
- [23] P. Smolensky. Information processing in dynamical systems: foundations of harmony theory. In *Symposium on Parallel and Distributed Processing*, 1986.
- [24] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning)*. The MIT Press, March 1998.
- [25] K. G. Zahedi, N. Ay, and R. Der. Higher coordination with less control – a result of information maximization in the sensomotor loop. *Adaptive Behavior*, 18(3-4):338–355, 2010.