

Max-Planck-Institut
für Mathematik
in den Naturwissenschaften
Leipzig

Universal Approximation of Markov Kernels by
Shallow Stochastic Feedforward Networks

by

Guido Montúfar

Preprint no.: 23

2015



Universal Approximation of Markov Kernels by Shallow Stochastic Feedforward Networks

Guido Montúfar

MONTUFAR@MIS.MPG.DE

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

Abstract

We establish upper bounds for the minimal number of hidden units for which a binary stochastic feedforward network with sigmoid activation probabilities and a single hidden layer is a universal approximator of Markov kernels. We show that each possible probabilistic assignment of the states of n output units, given the states of $k \geq 1$ input units, can be approximated arbitrarily well by a network with $2^{k-1}(2^{n-1} - 1)$ hidden units.

Keywords: universal approximation, stochastic feedforward network, Markov kernel

1. Introduction

The universal approximation capabilities of feedforward networks with one hidden layer of computational units have been studied in numerous papers and have been established under quite general conditions on the activation functions and the input-output domains (Cybenko, 1989; Hornik et al., 1989; Leshno et al., 1993; Chen and Chen, 1995; Gallant and White, 1988). Some works have also studied the minimal size of universal approximators and the quality of the approximations when the networks have only a limited number of hidden units (Hornik, 1991; Barron, 1993; Wenzel et al., 2000).

In the context of feedforward networks, the universal approximation question most commonly refers to the approximation of deterministic functions. In this paper we address a related problem that has received a bit less attention. We study the universal approximation of stochastic functions (Markov kernels) and the minimal number of hidden units in a stochastic feedforward network that is sufficient for this purpose. For a network with k input binary units and n output binary units, we are interested in maps taking inputs from $\{0, 1\}^k$ to probability distributions over outputs from $\{0, 1\}^n$. The outputs of the network are length- n binary vectors, but the outputs of the stochastic maps are length- 2^n probability vectors. We focus on shallow networks, with one single hidden layer, as the one illustrated in Figure 1, and stochastic binary units with output 1 probability given by the sigmoid of a weighted sum of their inputs. Given the number of input and output units, k and n , what is the smallest number of hidden units m that suffices to obtain a universal approximator of stochastic maps? We show that this is not more than $2^{k-1}(2^{n-1} - 1)$. We also consider the case where the weights of the output layer are fixed in advance and only the weights of the hidden layer are tunable. In that setting we show that $2^{k-1}(2^n - 1)$ hidden units are sufficient.

Some previous works have discussed compact representations of stochastic maps by feedforward networks, but focusing on the approximation of probability distributions (instead of Markov kernels) and deep networks with many hidden layers (Sutskever and Hinton, 2008; Le Roux and Bengio, 2010; Montúfar and Ay, 2011).

This paper is organized as follows. Section 2 contains basic definitions and notations, as well as a few comments about the deterministic setting. Section 3 presents our main results. Section 4 contains our analysis of the minimal number of hidden units for a network can approximate any stochastic function arbitrarily well by tuning the input weights and biases of the first layer, while the weights and biases of the second layer are kept fixed. Section 5 contains a corresponding analysis for the case when input weights and biases of both layers are tuned. Section 6 offers our conclusions and outlook.

2. Settings

This section contains definitions and basic observations.

2.1. Probability Distributions and Markov Kernels

Throughout this paper k , m , and n denote finite natural numbers. We denote by $\{0, 1\}^n$ the set of length- n binary strings. This is the set of all possible configurations of n binary units. The set of probability distributions over $\{0, 1\}^n$ is given by

$$\Delta_n := \left\{ p \in \mathbb{R}^{\{0,1\}^n} : p(x) \geq 0, \sum_{x \in \{0,1\}^n} p(x) = 1 \right\}. \quad (1)$$

This is the $(2^n - 1)$ -dimensional simplex of all vectors in 2^n -dimensional Euclidean space with non-negative entries and 1-norm equal to 1. The set of strictly positive distributions is the relative interior of Δ_n , denoted by Δ_n^+ .

A Markov kernel with source $\{0, 1\}^k$ and target $\{0, 1\}^n$ is a map from $\{0, 1\}^k$ to Δ_n . The set of all such kernels is

$$\Delta_{k,n} := \left\{ P \in \mathbb{R}^{\{0,1\}^k \times \{0,1\}^n} : P(y; x) \geq 0, \sum_{x \in \{0,1\}^n} P(y; x) = 1 \forall y \in \{0, 1\}^k \right\}.$$

Each Markov kernel is written as a matrix P with entries $P(x|y) \equiv P(y; x)$ for all $x \in \{0, 1\}^n$, for all $y \in \{0, 1\}^k$. The y -th row $P(\cdot|y)$ is a probability distribution of the output x , given the input y . The set $\Delta_{k,n}$ is the $2^k(2^n - 1)$ -dimensional polytope of $2^k \times 2^n$ row-stochastic matrices, which is the 2^k -th Cartesian power of Δ_n .

2.2. Feedforward Stochastic Networks

We consider stochastic binary units of the following form. Consider the sigmoid function

$$\sigma: \mathbb{R} \rightarrow [0, 1]; a \mapsto \frac{1}{1 + \exp(-a)}.$$

For each input vector $y \in \{0, 1\}^k$, the unit computes a scalar pre-activation value by an affine function $v^\top y + b$, and outputs 1 with probability $\Pr(1) = \sigma(v^\top y + c)$ or otherwise it outputs 0 with complementary probability, $\Pr(0) = 1 - \sigma(v^\top y + c)$.

Given an input vector $y \in \{0, 1\}^k$, the probability that a feedforward layer of m stochastic units outputs $z = (z_1, \dots, z_m)^\top \in \{0, 1\}^m$ is given by the product of the output probabilities of the




Figure 1: Feedforward network with a layer of k input units, a layer of m hidden units, and a layer of n output units.

individual units in the layer,

$$\begin{aligned} \Pr(z) &= \prod_{j \in [m]} \sigma(v_j^\top y + c_j)^{z_j} (1 - \sigma(v_j^\top y + c_j))^{1-z_j} \\ &= \frac{\exp((Vy + c)^\top z)}{\sum_{z' \in \{0,1\}^m} \exp((Vy + c)^\top z')}, \end{aligned}$$

where $V = (v_1 | \dots | v_m)^\top \in \mathbb{R}^{m \times k}$ and $c = (c_1, \dots, c_m)^\top \in \mathbb{R}^m$.

Definition 1 We denote by $F_{k,m} \subseteq \Delta_{k,m}$ the set of Markov kernels that can be represented by a feedforward layer with k input units and m output units; that is, the set of kernels

$$P(z|y) = \frac{\exp((Vy + c)^\top z)}{\sum_{z'} \exp((Vy + c)^\top z')} \quad \forall z \in \{0,1\}^m, y \in \{0,1\}^k,$$

parametrized by $V \in \mathbb{R}^{m \times k}$ and $c \in \mathbb{R}^m$.

Definition 2 We denote by $F_{k,m,n} = F_{m,n} \circ F_{k,m} \subseteq \Delta_{k,n}$ the set of Markov kernels that can be represented by a feedforward network with k input units, m hidden units, and n output units; that is, the set of kernels of the form

$$P(x|y) = \sum_{z \in \{0,1\}^m} Q(z|y)R(x|z) \quad \forall y \in \{0,1\}^k, x \in \{0,1\}^n,$$

where $Q \in F_{k,m}$ and $R \in F_{m,n}$.

Fig. 1 gives a schematic illustration of a feedforward network with one input, one hidden, and one output layer.

The following set of probability distributions will be important in our analysis.

Definition 3 We denote by $\mathcal{E}_m \subseteq \Delta_m$ the set of probability distributions on $\{0,1\}^m$ of the form

$$p(z) = \frac{\exp(b^\top z)}{\sum_{z' \in \{0,1\}^m} \exp(b^\top z')} \quad \forall z \in \{0,1\}^m,$$

parametrized by $b \in \mathbb{R}^m$. This is precisely the set of probability distributions that factorize as

$$p(z) = \prod_{j \in [m]} p_j^{1-z_j} (1-p_j)^{z_j} \quad \forall z = (z_1, \dots, z_m)^\top \in \{0, 1\}^m,$$

for some $p_j \in (0, 1)$ for all $j \in [m]$.

Note that each kernel $Q \in F_{k,m}$ represented by a feedforward layer is a tuple of 2^k factorizing probability distributions, $Q(\cdot|y) \in \mathcal{E}_m$ for all $y \in \{0, 1\}^k$.

2.3. Universal Approximation

Definition 4 A set $\mathcal{M} \subseteq \Delta_{k,n}$ is a universal approximator if and only if every point from $\Delta_{k,n}$ can be approximated arbitrarily well by points from \mathcal{M} . This is the case if and only if $\overline{\mathcal{M}} = \Delta_{k,n}$, where $\overline{\mathcal{M}}$ denotes the closure of \mathcal{M} in the Euclidean topology.

We will study the universal approximation properties of $F_{k,m,n}$ in two cases. In the first case, only the first layer has free parameters, whereas the second layer has fixed parameters. This means that we consider the set $R \circ F_{k,m} \subseteq \Delta_{k,n}$ for some fixed $R \in F_{m,n}$. In the second case, all parameters are free.

By comparing the number of free parameters and the dimension of $\Delta_{k,n}$, it is straightforward to obtain the following crude lower bound on the minimal number of hidden units that suffices for universal approximation:

Proposition 5 Let $k \geq 1$ and $n \geq 1$.

- If there is an $R \in \overline{F_{m,n}}$ with $R \circ \overline{F_{k,m}} = \Delta_{k,n}$, then $m \geq \lceil \frac{1}{(k+1)} 2^k (2^n - 1) \rceil$.
- If $\overline{F_{k,m,n}} = \Delta_{k,n}$, then $m \geq \lceil \frac{1}{(n+k+1)} (2^k (2^n - 1) - n) \rceil$.

2.4. Feedforward Deterministic Networks

Deterministic networks are special cases of stochastic networks. Consider a feedforward stochastic network as defined above, but where all input weights and biases, W and b , are multiplied by $r \in \mathbb{R}$. For generic choices of W and b , when $r \rightarrow \infty$ each unit outputs 0 or 1 with probability one, depending on its inputs. In this case, the kernels represented by the feedforward network are deterministic, meaning that they have the form $P(x|y) = \delta_{f(y)}(x)$ for all $x \in \{0, 1\}^n$, for all $y \in \{0, 1\}^k$, for some function $f: \{0, 1\}^k \rightarrow \{0, 1\}^n$. The feedforward networks defined by these limits are called linear threshold networks.

The representation of Boolean functions by linear threshold networks (with one binary output unit) has been studied extensively in the literature. The problem can be beautifully described as the problem of classifying subsets of vertices of the k -dimensional unit cube by an arrangement of oriented affine hyperplanes. In (Wenzel et al., 2000) it was shown that, for $k \geq 2$ and $n = 1$, the smallest number m of hidden units for which a linear threshold network with one hidden layer can compute every Boolean function satisfies $2^{k/2} - \frac{k^2}{2} < -\frac{k^2}{2} + \sqrt{\frac{k^4}{4} + 2^k} \leq m \leq \frac{3}{k+2} 2^k$.

It is important to realize that, in the deterministic setting, if a network with m hidden units can represent any Boolean function $f: \{0, 1\}^k \rightarrow \{0, 1\}$, then a network with $n \cdot m$ hidden units can

represent any function $g: \{0, 1\}^k \rightarrow \{0, 1\}^n$. This is because one can always write $g = (f_1, \dots, f_n)$ and compute the individual f_i 's in parallel groups of hidden and output units. In the stochastic setting the same is not true. In general, a joint distribution over $\{0, 1\}^n$ cannot be written in terms of n marginal distributions over $\{0, 1\}$ alone, and so, the activities of the individual output units cannot be computed independently from each other.

3. Results

We bound the minimal number of hidden units of a universal approximator from above in two cases. In the first case we consider a network whose second layer has fixed weights and biases. In the second case all weights and biases are free parameters.

3.1. Fixed Weights in the Second Layer

Theorem 6 *Let $k \geq 1$ and $n \geq 1$. There is an $R \in \overline{F_{m,n}}$ such that $R \circ \overline{F_{k,m}} = \Delta_{k,n}$, whenever $m \geq \frac{1}{2}2^k(2^n - 1)$.*

In view of the crude lower bound from Proposition 5, this upper bound is tight at least when $k = 1$.

When there are no input units, $k = 0$, we may set $\overline{F_{0,m}} = \mathcal{E}_m$ and $\Delta_{0,n} = \Delta_n$. The theorem generalizes to this case as:

Proposition 7 *Let $n \geq 2$. There is an $R \in \overline{F_{m,n}}$ with $R \circ \overline{\mathcal{E}_m} = \Delta_n$, whenever $m \geq 2^n - 1$.*

This bound is always tight, since the network uses exactly $2^n - 1$ parameters to approximate every distribution from Δ_n arbitrarily well.

3.2. Changeable Weights in the Second Layer

When both layers have changeable weights we obtain:

Theorem 8 *Let $k \geq 1$ and $n \geq 2$. Then $\overline{F_{k,m,n}} = \Delta_{k,n}$, whenever $m \geq 2^{k-1}(2^{n-1} - 1)$.*

Comparison with the crude lower bound from Proposition 5 reveals that this bound is tight at least for $(k, n) = (1, 2), (2, 2), (3, 2), (1, 3)$.

In the case of no inputs we obtain:

Proposition 9 *Let $n \geq 2$. Then $\overline{F_{m,n}} \circ \overline{\mathcal{E}_m} = \Delta_n$, whenever $m \geq 2^{n-1} - 1$.*

3.3. Outline of the Proof

Our strategy for proving Theorem 6 and Theorem 8 can be summarized as follows:

- First we show that the first layer of $\overline{F_{k,m,n}}$ can approximate Markov kernels arbitrarily well, which fix the state of some units, depending on the input, and have an arbitrary product distribution over the states of the other units. The idea is illustrated in Fig. 2.
- Then we show that the second layer can approximate deterministic kernels arbitrarily well, whose rows are copies of all point measures from Δ_n , ordered in a good way with respect to the different inputs. Note that the point measures are the vertices of the simplex Δ_n .

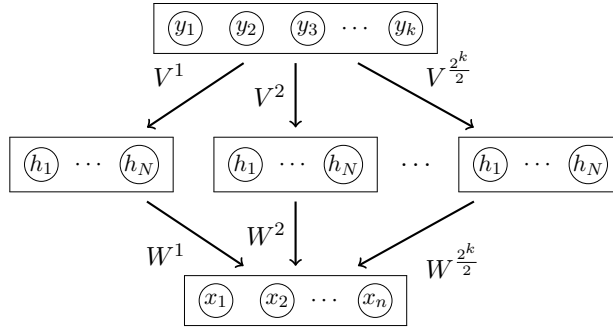


Figure 2: Illustration of the construction used in our proof. Each block of hidden units is active on a distinct subset of possible inputs. The output layer integrates the activities of the block that was activated by the input, and produces corresponding activities of the output units.

- Finally, we show that the set of product distributions of each block of hidden units is mapped to the convex hull of the rows of the kernel represented by the second layer, which is Δ_n .
- The output distributions of distinct sets of inputs is modeled individually by distinct blocks of hidden units and so we obtain the universal approximation of Markov kernels.

The goal of our analysis is to construct the individual pieces of the network as compact as possible. In particular, Lemma 10 will provide a trick that allows us to use each block of units in the hidden layer for a pair of distinct input vectors at the same time. This allows us to halve the number of hidden units that would be needed if each input had an individual block of active hidden units. Similarly, Lemma 15 will provide a trick for producing more flexible mixture components at the output layer than simply point measures. This again allows us to halve the number of hidden units of the simpler construction.

4. The Number of Hidden Units for Fixed Weights in the Second Layer

4.1. The First Layer

We start with the following lemma.

Lemma 10 *Let $y', y'' \in \{0, 1\}^k$ differ only in one entry, and let q', q'' be any two distributions on $\{0, 1\}$. Then $F_{k,1}$ can approximate the following arbitrarily well:*

$$p(\cdot|y) = \begin{cases} q', & \text{if } y = y' \\ q'', & \text{if } y = y'' \\ \delta_0, & \text{else} \end{cases} .$$

Proof Given the input weights and bias, $V \in \mathbb{R}^{1 \times k}$ and $c \in \mathbb{R}$, for each input $y \in \{0, 1\}^k$ the output probability is given by

$$p(z = 1|y) = \sigma(Vy + c). \tag{2}$$

Since the two vectors $y', y'' \in \{0, 1\}^k$ differ only in one entry, they are the vertices of an edge E of the k -dimensional unit cube. Let $l \in [k] := \{1, \dots, k\}$ be the entry in which they differ, with $y'_l = 0$ and $y''_l = 1$. Since E is a (1-dimensional) face of the cube, there is a supporting hyperplane of E . This means that there are $\tilde{V} \in \mathbb{R}^{1 \times k}$ and $\tilde{c} \in \mathbb{R}$ with $\tilde{V}y + \tilde{c} = 0$ if $y \in E$, and $\tilde{V}y + \tilde{c} < -1$ if $y \in \{0, 1\}^k \setminus E$. Let $s' = \sigma^{-1}(q'(z = 1))$ and $s'' = \sigma^{-1}(q''(z = 1))$. We define $c = \alpha\tilde{c} + s'$ and $V = \alpha\tilde{V} + (s'' - s')e_l$. Then, as $\alpha \rightarrow \infty$,

$$Vy + c = \begin{cases} s', & \text{if } y = y' \\ s'', & \text{if } y = y'' \\ -\infty, & \text{else} \end{cases}.$$

Plugging this into (2) proves the claim. ■

Given any binary vector $y \in \{0, 1\}^k$, let $\text{int}(y) := \sum_{i=1}^k 2^{i-1}y_i$ be its integer representation. Using the previous lemma, we obtain the following.

Proposition 11 *Let $N \geq 1$ and $m = 2^{k-1}N$. For each $y \in \{0, 1\}^k$, let $p(\cdot|y)$ be an arbitrary distribution from \mathcal{E}_N . The model $F_{k,m}$ can approximate the following kernel from $\Delta_{k,m}$ arbitrarily well:*

$$P(h|y) = \delta_0(h^0) \dots \delta_0(h^{\lfloor \text{int}(y)/2 \rfloor - 1}) p(h^{\lfloor \text{int}(y)/2 \rfloor} | y) \\ \times \delta_0(h^{\lfloor \text{int}(y)/2 \rfloor + 1}) \dots \delta_0(h^{2^{k-1}-1}),$$

where $h^i = (h_{Ni+1}, \dots, h_{N(i+1)})$ for all $i \in \{0, 1, \dots, 2^{k-1} - 1\}$.

Proof We divide the set $\{0, 1\}^k$ of all possible inputs into 2^{k-1} disjoint pairs with successive decimal values. The i -th pair consists of the two vectors y with $\lfloor \text{int}(y)/2 \rfloor = i$, for all $i \in \{0, \dots, 2^{k-1} - 1\}$. The kernel P has the property that, for the i -th input pair, all output units are inactive with probability one, except those with index $Ni + 1, \dots, N(i + 1)$. Given a joint distribution q on the states of I units, let q_j denote the corresponding marginal distribution on the states of the j -th of these units. By Lemma 10, we can set

$$P_{Ni+j}(\cdot|y) = \begin{cases} p_j(\cdot|y), & \text{if } \text{int}(y) = 2i \\ p_j(\cdot|y), & \text{if } \text{int}(y) = 2i + 1 \\ \delta_0, & \text{else} \end{cases}$$

for all $i \in \{0, \dots, 2^{k-1} - 1\}$ and $j \in [N]$. ■

4.2. The Second Layer

For the second layer we will consider deterministic kernels. Given a binary vector z , let $l(z) := \lceil \log_2(\text{int}(z) + 1) \rceil$ denote the largest j with $z_j = 1$. Here we set $l(0, \dots, 0) = 0$. Given an integer $l \in \{0, \dots, 2^n - 1\}$, let $\text{bin}_n(l)$ denote the n -bit binary representation of l ; that is, the vector with $\text{int}(\text{bin}_n(l)) = l$.




Figure 3: Illustration of Lemma 13 for $n = 2$, $N = 2^n - 1 = 3$. There is an arrangement of n hyperplanes which divides the vertices of a $(2^n - 1)$ -dimensional cube as $0|1|2, 3|4, 5, 6, 7|\dots|2^{2^n-2}, \dots, 2^{2^n-1} - 1$.

Lemma 12 *Let $N = 2^n - 1$. The set $F_{N,n}$ can approximate the following deterministic kernel arbitrarily well:*

$$Q(\cdot|z) = \delta_{\text{bin}_n l(z)}(\cdot) \quad \forall z \in \{0, 1\}^N.$$

In words, the z -th row of Q indicates the largest non-zero entry of the binary vector z . For example, for $n = 2$ we have $N = 3$ and

$$Q = \begin{matrix} & \begin{matrix} 00 & 01 & 10 & 11 \end{matrix} \\ \begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{pmatrix} & \begin{matrix} 000 \\ 001 \\ 010 \\ 011 \\ 100 \\ 101 \\ 110 \\ 111 \end{matrix} \end{matrix}.$$

Proof of Lemma 12 Given the input and bias weights, $W \in \mathbb{R}^{n \times N}$ and $b \in \mathbb{R}^n$, for each input $z \in \{0, 1\}^N$ the output distribution is the product distribution $p(\cdot|z) \in \mathcal{E}_n$ with exponential parameters $Wz + b$. If $\text{sgn}(Wz + b) = \text{sgn}(x - \frac{1}{2})$ for some $x \in \{0, 1\}^n$, then the product distribution with parameters $\alpha(Wz + b)$, $\alpha \rightarrow \infty$ tends to δ_x . We only need to show that there is a choice of W and b with $\text{sgn}(Wz + b) = \text{sgn}(f(z) - \frac{1}{2})$, $f(z) = \text{bin}_n \lceil \log_2(\text{int}(z) + 1) \rceil$, for all $z \in \{0, 1\}^N$. That is precisely the statement of Lemma 13. \blacksquare

We used the following lemma in the proof of Lemma 12. For $l = 0, 1, \dots, 2^n - 1$, the l -th orthant of \mathbb{R}^n is the set of all vectors $r \in \mathbb{R}^n$ with strictly positive or negative entries and $\text{int}(\text{hs}(r)) = l$, where hs denotes the entry-wise Heaviside function, assigning value 0 to negative entries and value 1 to positive entries.

Lemma 13 *Let $N = 2^n - 1$. There is an affine map $\{0, 1\}^N \rightarrow \mathbb{R}^n$; $z \mapsto Wz + b$, sending $\{z \in \{0, 1\}^N : l(z) = l\}$ to the l -th orthant of \mathbb{R}^n , for all $l \in \{0, 1, \dots, N\}$.*

Proof Consider the affine map $z \mapsto Wz + b$, where $b = -(1, \dots, 1)^\top$ and the l -th column of W is $2^{l+1}(\text{bin}_n(l) - \frac{1}{2})$ for all $l \in \{1, \dots, N\}$. For this choice, $\text{sgn}(Wz + b) = \text{sgn}(\text{bin}_n(l(z)) - \frac{1}{2})$ lies in the l -th orthant of \mathbb{R}^n . ■

Lemma 13 is illustrated in Fig. 3 for $n = 2$. As another example, for $n = 3$ the affine map can be defined as $z \mapsto Wz + b$, where

$$b = \begin{pmatrix} -1 \\ -1 \\ -1 \end{pmatrix}, \quad \text{and} \quad W = \begin{pmatrix} 2 & -4 & 8 & -16 & 32 & -64 & 128 \\ -2 & 4 & 8 & -16 & -32 & 64 & 128 \\ -2 & -4 & -8 & 16 & 32 & 64 & 128 \end{pmatrix}.$$

Proposition 14 Let $N = 2^n - 1$ and let Q be defined as in Lemma 12. Then $Q \circ \mathcal{E}_N = \Delta_n^+$.

Proof Consider a strictly positive product distribution $p \in \mathcal{E}_N$ with $p(z) = \prod_i p_i^{1-z_i} (1-p_i)^{z_i}$ for all $z \in \{0, 1\}^N$. Then $p^\top Q \in \Delta_n$ is the vector $q = (q_0, q_1, \dots, q_N)$ with entries

$$\begin{aligned} q_i &= \sum_{z: l(z)=i} p(z) \\ &= \sum_{z_k, k < i} \left(\prod_{k < i} p_k^{1-z_k} (1-p_k)^{z_k} \right) (1-p_i) \left(\prod_{j > i} p_j \right) \\ &= (1-p_i) \prod_{j > i} p_j \end{aligned}$$

for all $i = 1, \dots, N$, and $q_0 = \prod_{j > 0} p_j$. Therefore,

$$\frac{q_i}{q_0} = \frac{1-p_i}{p_i} \frac{1}{\prod_{j=1}^{i-1} p_j} \quad \forall i = 1, \dots, N. \quad (3)$$

Since $\frac{1-p_i}{p_i}$ can be made arbitrary in $(0, \infty)$ by choosing an appropriate p_i , independently of p_j , for $j < i$, the quotient $\frac{q_i}{q_0}$ can be made arbitrary in $(0, \infty)$ for all $i \in \{1, \dots, N\}$. This implies that q can be made arbitrary in Δ_n^+ . ■

Proof of Theorem 6 This follows from Proposition 11 and Proposition 14. ■

5. The Number of Hidden Units for Changeable Weights in the Second Layer

In order to prove Theorem 8 we will use the same construction of the first layer as in the previous section. For the second layer we will use the following refinement of Lemma 12.

Lemma 15 Let $n \geq 2$ and $N = 2^{n-1} - 1$. The set $F_{N,n}$ can approximate the following kernels arbitrarily well:

$$Q(\cdot|z) = \lambda_z \delta_{\text{bin}_n 2l(z)}(\cdot) + (1 - \lambda_z) \delta_{\text{bin}_n 2l(z)+1}(\cdot) \quad \forall z \in \{0, 1\}^N,$$

where λ_z are certain (not mutually independent) weights in $[0, 1]$. Given any $r_l \in \mathbb{R}_+$ for all $l \in \{0, 1, \dots, N\}$, it is possible to choose the λ_z 's such that

$$\frac{\sum_{z: l(z)=l} \lambda_z}{\sum_{z: l(z)=l} (1 - \lambda_z)} = r_l \quad \forall l \in \{0, 1, \dots, N\}.$$

In words, the z -th row of Q is a convex combination of the indicators of $2l(z)$ and $2l(z) + 1$, and, furthermore, the total weight assigned to $2l$ relative to the total weight assigned to $2l + 1$ can be made arbitrary for each l . For example, for $n = 2$ we have $N = 1$ and

$$Q = \begin{pmatrix} 00 & 01 & 10 & 11 \\ \lambda_0 & (1 - \lambda_0) & & \\ & & \lambda_1 & (1 - \lambda_1) \end{pmatrix} \begin{matrix} 0 \\ 1 \\ 0 \\ 1 \end{matrix}.$$

The sum of all weights in a given even column can be made arbitrary, relative to the sum of all weights in the column right next to it, for all $N + 1$ such pairs of columns simultaneously.

Proof of Lemma 15 Consider the sets $Z_l = \{z \in \{0, 1\}^N : l(z) = l\}$, for $l = 0, 1, \dots, N$. Let $W' \in \mathbb{R}^{(n-1) \times N}$ and $b' \in \mathbb{R}^{n-1}$ be the input weights and biases defined in Lemma 12. We define W and b by appending a row (μ_1, \dots, μ_N) on top of W' and an entry μ_0 on top of b' .

If $\mu_j < 0$ for all $j = 0, 1, \dots, N$, then $z \mapsto Wz + b$ maps Z_l to the $2l$ -th orthant of \mathbb{R}^n , for each $l = 0, 1, \dots, N$.

Consider now some arbitrary fixed choice of μ_j , $j < l$. Choosing $\mu_l < 0$ with $|\mu_l| > \sum_{j < l} |\mu_j|$, Z_l is mapped to the $2l$ -th orthant. If $\mu_l \rightarrow -\infty$, then $\lambda_z \rightarrow 1$ for all z with $l(z) = l$. As we increase μ_l to a sufficiently large positive value, the elements of Z_l gradually are mapped to the $(2l + 1)$ -th orthant. If $\mu_l \rightarrow \infty$, then $(1 - \lambda_z) \rightarrow 1$ for all z with $l(z) = l$. By continuity, there is a choice of μ_l such that $\frac{\sum_{z: l(z)=l} \lambda_z}{\sum_{z: l(z)=l} (1 - \lambda_z)} = r_l$.

Note that the images of Z_j , $j < l$, are independent of the i -th rows of W for all $i = l, \dots, N$. Hence changing μ_l does not have any influence on the images of Z_l nor on λ_z for $z: l(z) < l$. Tuning μ_i sequentially, starting with $i = 0$, we obtain a kernel that approximates any Q of the claimed form arbitrarily well. \blacksquare

Let Q_n^N be the collection of kernels described in Lemma 15.

Proposition 16 Let $n \geq 2$ and $N = 2^{n-1} - 1$. Then $Q_n^N \circ \mathcal{E}_N = \Delta_n^+$.

Proof Consider a strictly positive product distribution $p \in \mathcal{E}_N$ with $p(z) = \prod_{i=1}^N p_i^{z_i-1} (1 - p_i)^{z_i}$ for all $z \in \{0, 1\}^N$. Then $p^\top Q \in \Delta_n$ is a vector $(q_0, q_1, \dots, q_{2N+1})$ whose entries satisfy

$$q_{2i} + q_{2i+1} = (1 - p_i) \prod_{j>i} p_j,$$

for all $i = 1, \dots, N$ and $q_0 + q_1 = \prod_{j>0} p_j$. As in the proof of Proposition 14, this implies that the vector $(q_0 + q_1, q_2 + q_3, \dots, q_{2N} + q_{2N+1})$ can be made arbitrary in Δ_{n-1}^+ . This is irrespective of the coefficients $\lambda_0, \dots, \lambda_N$. Now all we need to show is that we can make q_{2i} arbitrary relative to q_{2i+1} for all $i = 0, \dots, N$.

We have

$$\begin{aligned}
 q_{2i} &= \sum_{z: l(z)=i} \lambda_z p(z) \\
 &= \left(\sum_{z: l(z)=i} \lambda_z \left(\prod_{k<i} p_k^{1-z_k} (1-p_k)^{z_k} \right) \right) \\
 &\quad \times (1-p_i) \left(\prod_{j>i} p_j \right)
 \end{aligned}$$

and

$$\begin{aligned}
 q_{2i+1} &= \sum_{z: l(z)=i} (1-\lambda_z) p(z) \\
 &= \left(\sum_{z: l(z)=i} (1-\lambda_z) \left(\prod_{k<i} p_k^{1-z_k} (1-p_k)^{z_k} \right) \right) \\
 &\quad \times (1-p_i) \left(\prod_{j>i} p_j \right).
 \end{aligned}$$

Therefore,

$$\frac{q_{2i}}{q_{2i+1}} = \frac{\sum_{z: l(z)=i} \lambda_z \left(\prod_{k<i} p_k^{1-z_k} (1-p_k)^{z_k} \right)}{\sum_{z: l(z)=i} (1-\lambda_z) \left(\prod_{k<i} p_k^{1-z_k} (1-p_k)^{z_k} \right)}.$$

By Lemma 15 it is possible to choose all λ_z arbitrarily close to zero for all z with $l(z) = i$ and have them transition continuously to values arbitrarily close to one (independently of the values of λ_z , $z: l(z) \neq i$). Since all p_k are strictly positive, this implies that the quotient $\frac{q_{2i}}{q_{2i+1}}$ takes all values in $(0, \infty)$ as the λ_z , $z: l(z) = i$ transition from zero to one. ■

Proof of Theorem 8 This follows from Proposition 11 and Proposition 16. ■

6. Conclusions

This article proves upper bounds on the minimal size of binary shallow stochastic feedforward networks with sigmoid activation probabilities that can approximate any stochastic function with a given number of binary inputs and outputs arbitrarily well. By our analysis, if all parameters of the network are free, $2^{k-1}(2^{n-1} - 1)$ hidden units suffice, and, if only the parameters of the first layer are free, $2^{k-1}(2^n - 1)$ hidden units suffice.

It is interesting to compare these results with what is known about universal approximation of Markov kernels by shallow undirected stochastic networks, called conditional restricted Boltzmann machines. For those networks previous work (Montúfar et al., 2014) has shown that $2^{k-1}(2^n - 1)$ hidden units suffice, whereby, if the number k of input units is large enough, $\frac{1}{4}2^k(2^n - 1 + 1/30)$

suffice. These bounds are sandwiched between our bounds for feedforward networks. In the case of no input units, our bound $2^{n-1} - 1$ equals the tightest known bound for universal approximation of probability distributions by restricted Boltzmann machines (Montúfar and Ay, 2011). Hence, given the current state of knowledge, if we were to specify a smallest possible universal approximator of Markov kernels or probability distributions, feedforward networks would seem preferable. However, verifying the tightness of the bounds appears to be a very challenging problem. Recently it has been observed that undirected networks can represent many kernels that can be represented by feedforward networks, especially when these are not too stochastic (Montúfar et al., 2014; Montúfar, 2014). In future work it would be interesting to compare the representational power of both network architectures in more detail.

We think that it is possible to adapt our analysis to cover deep architectures as well. This should allow us to conclude that a multilayer feedforward stochastic network with k input units and n output units is a universal approximator of Markov kernels if it has about 2^{n-1} hidden layers, each containing about $n2^{k-1}$ units. The verification of this claim is left for future work. In relation with this, the results presented in this paper should be helpful for analyzing the relative representational power of shallow vs. deep stochastic feedforward networks, a topic that has attracted much interest in recent years and that still poses a great many questions.

Acknowledgments

I would like to thank the Santa Fe Institute, Santa Fe, NM, USA, and the RIKEN Brain Science Institute, Hirosawa, Saitama, Japan, for hosting me during the work on this article.

References

- Andrew R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. *Information Theory, IEEE Transactions on*, 39(3):930–945, May 1993. ISSN 0018-9448. doi: 10.1109/18.256500.
- Tianping Chen and Hong Chen. Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems. *Neural Networks, IEEE Transactions on*, 6(4):911–917, Jul 1995. ISSN 1045-9227. doi: 10.1109/72.392253.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989. doi: 10.1007/BF02551274. URL <http://dx.doi.org/10.1007/BF02551274>.
- A. Ronald Gallant and Halbert White. There exists a neural network that does not make avoidable mistakes. In *Neural Networks, 1988., IEEE International Conference on*, pages 657–664 vol.1, July 1988. doi: 10.1109/ICNN.1988.23903.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, March 1991. ISSN 0893-6080. doi: 10.1016/0893-6080(91)90009-T. URL [http://dx.doi.org/10.1016/0893-6080\(91\)90009-T](http://dx.doi.org/10.1016/0893-6080(91)90009-T).
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359 – 366, 1989. ISSN 0893-6080. doi: 10.

1016/0893-6080(89)90020-8. URL [http://dx.doi.org/10.1016/0893-6080\(89\)90020-8](http://dx.doi.org/10.1016/0893-6080(89)90020-8).

Nicolas Le Roux and Yoshua Bengio. Deep belief networks are compact universal approximators. *Neural Computation*, 22(8):2192–2207, August 2010. ISSN 0899-7667. doi: 10.1162/neco.2010.08-09-1081. URL <http://dx.doi.org/10.1162/neco.2010.08-09-1081>.

Moshe Leshno, Vladimir Ya Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*, 6(6):861–867, 1993.

Guido Montúfar. Deep narrow boltzmann machines are universal approximators. *arXiv preprint arXiv:1411.3784*, 2014.

Guido Montúfar and Nihat Ay. Refinements of universal approximation results for deep belief networks and restricted Boltzmann machines. *Neural Computation*, 23(5):1306–1319, May 2011. ISSN 0899-7667. doi: 10.1162/NECO_a_00113. URL http://dx.doi.org/10.1162/NECO_a_00113.

Guido Montúfar, Nihat Ay, and Keyan Zahedi. Geometry and expressive power of conditional restricted boltzmann machines for sensorimotor control. *arXiv preprint arXiv:1402.3346*, 2014.

Ilya Sutskever and Geoffrey E. Hinton. Deep, narrow sigmoid belief networks are universal approximators. *Neural Computation*, 20(11):2629–2636, November 2008. ISSN 0899-7667. doi: 10.1162/neco.2008.12-07-661. URL <http://dx.doi.org/10.1162/neco.2008.12-07-661>.

Walter Wenzel, Nihat Ay, and Frank Pasemann. Hyperplane arrangements separating arbitrary vertex classes in n-cubes. *Adv. Appl. Math.*, 25(3):284–306, 2000.