

ORIGINAL CONTRIBUTION

Multilayer Feedforward Networks With a Nonpolynomial Activation Function Can Approximate Any Function

MOSHE LESHNO,¹ VLADIMIR YA. LIN,² ALLAN PINKUS,² AND SHIMON SCHOCKEN³

¹The Hebrew University, Israel, ²Technion, Israel and ³New York University

(Received 9 February 1993; revised and accepted 15 March 1993)

Abstract—Several researchers characterized the activation function under which multilayer feedforward networks can act as universal approximators. We show that most of all the characterizations that were reported thus far in the literature are special cases of the following general result: A standard multilayer feedforward network with a locally bounded piecewise continuous activation function can approximate any continuous function to any degree of accuracy if and only if the network's activation function is not a polynomial. We also emphasize the important role of the threshold, asserting that without it the last theorem does not hold.

Keywords—Multilayer feedforward networks, Activation functions, Role of threshold, Universal approximation capabilities, $L^p(\mu)$ approximation.

1. BACKGROUND

The basic building block of a neural network is a processing-unit that is linked to n input-units through a set of n directed connections. The single unit model is characterized by (1) a threshold value, denoted θ ; (2) a univariate activation function, denoted $\sigma : R \rightarrow R$; and (3) a vector of "weights," denoted $\mathbf{w} = w_1, \dots, w_n$. When an input-vector $\mathbf{x} = x_1, \dots, x_n$ is fed into the network through the input-units, the processing-unit computes the function $\sigma(\mathbf{w} \cdot \mathbf{x} - \theta)$, $\mathbf{w} \cdot \mathbf{x}$ being the standard inner-product in R^n . The value of this function is then taken to be the network's output.

A network consisting of a layer of n input-units and a layer of m processing-units can be "trained" to approximate a limited class of functions $f : R^n \rightarrow R^m$. When the network is fed with new examples of vectors $\mathbf{x} \in R^n$ and their correct mappings $f(\mathbf{x})$, a "learning algorithm" is applied to adjust the weights and the thresholds in a direction that minimizes the difference between $f(\mathbf{x})$ and the network's output. Similar back propagation learning algorithms exist for multilayer feedforward networks, and the reader is referred to Hinton (1989) for an excellent survey on the subject. This paper, however, does not concern learning. Rather,

we focus on the following fundamental question: If we are free to choose any \mathbf{w} , θ , and σ that we desire, which "real life" functions $f : R^n \rightarrow R^m$ can multilayer feedforward networks emulate?

During the last decade, multilayer feedforward networks have been shown to be quite effective in many different applications, with most papers reporting that they perform at least as well as their traditional competitors (e.g., linear discrimination models and Bayesian classifiers). This success has recently led several researchers to undertake a rigorous analysis of the mathematical properties that enable feedforward networks to perform well in the field. The motivation for this line of research was eloquently described by Hornik and his colleagues (Hornik, Stinchcombe, & White, 1989) as follows: "The apparent ability of sufficiently elaborate feedforward networks to approximate quite well nearly any function encountered in applications leads one to wonder about the ultimate capabilities of such networks. Are the successes observed to date reflective of some deep and fundamental approximation capabilities, or are they merely flukes, resulting from selective reporting and a fortuitous choice of problems?"

Previous research on the approximation capabilities of feedforward networks can be found in le Cun (1987), Cybenko (1989), Funahashi (1989), Gallant and White (1988), Hecht-Nielson (1989), Hornik et al., (1989), Irei and Miyake (1988), Lapedes and Farber (1988), Stinchcombe and White (1990), and Chui and Li (1992). These studies show that if the network's

Acknowledgement: The proof of Step 4 as given herein is due to Y. Benyamini to whom we are most appreciative.

Requests for reprints should be sent to Moshe Leshno, School of Business Administration, The Hebrew University, Mount Scopus, Jerusalem 91905 Israel.

activation functions obey an explicit set of assumptions (which vary from one paper to another), then the network can indeed be shown to be a universal approximator. For example, Gallant and White (1988) proved that a network with “cosine squasher” activation functions possess all the approximations properties of Fourier series representations. Hornik et al. (1989) extended this result and proved that a network with *arbitrary squashing* activation functions are capable of approximating any function of interest. Most recently, Hornik (1991) has proven two general results, as follows:

HORNIK THEOREM 1. *Whenever the activation function is bounded and nonconstant, then, for any finite measure μ , standard multilayer feedforward networks can approximate any function in $L^p(\mu)$ (the space of all functions on R^n such that $\int_{R^n} |f(x)|^p d\mu(x) < \infty$) arbitrarily well, provided that sufficiently many hidden units are available.*

HORNIK THEOREM 2. *Whenever the activation function is continuous, bounded and nonconstant, then, for arbitrary compact subsets $X \subseteq R^n$, standard multilayer feedforward networks can approximate any continuous function on X arbitrarily well with respect to uniform distance, provided that sufficiently many hidden units are available.*

In this paper we generalize in particular Hornik’s Theorem 2 by establishing necessary and sufficient conditions for universal approximation. In particular, we show that a standard multilayer feedforward network can approximate any continuous function to any degree of accuracy if and only if the network’s activation function is not polynomial. In addition, we emphasize and illustrate the role of the threshold value (a parameter of the activation function), without which the theorem does not hold. The theorem is intriguing because (1) the conditions that it imposes on the activation function are minimal; and (2) it embeds, as special cases, almost all the activation functions that were reported thus far in the literature.

2. MULTILAYER FEEDFORWARD NETWORKS

The general architecture of a multilayer feedforward network consists of an input layer with n input-units, an output layer with m output-units, and one or more hidden layers consisting of intermediate processing-units. Because a mapping $f : R^n \rightarrow R^m$ can be computed by m mappings $f_j : R^n \rightarrow R$, it is (theoretically) sufficient to focus on networks with one output-unit only. In addition, since our findings require only a single hidden layer, we will assume hereafter that the network consists of three layers only: input, hidden, and output. One such network is depicted in Figure 1.




FIGURE 1. Single hidden layer feedforward neural net.

In the figure, the weights-vector and the threshold value associated with the j th processing-unit are denoted w_j , and θ_j , respectively. The weights-vector associated with the single output-unit is denoted β , and the input-vector is denoted x . With this notation, we see that the function that a multilayer feedforward network computes is:

$$f(x) = \sum_{j=1}^k \beta_j \cdot \sigma(w_j \cdot x - \theta_j) \tag{1}$$

k being the number of processing-units in the hidden layer. Hence, the family of functions that can be computed by multilayer feedforward networks is characterized by four parameters, as follows:

1. The number of processing-units, denoted k ;
2. The set of weights $\{w_{i,j}\}$, one for each pair of connected units;
3. The set of threshold values $\{\theta_j\}$, one for each processing-unit;
4. An activation function $\sigma : R \rightarrow R$, same for each processing-unit.

In what follows, we denote the space of these parameters $\Lambda = \langle k, \{w_{i,j}\}, \{\theta_j\}, \sigma \rangle$, and a particular quadruple of parameters is denoted $\omega \in \Lambda$. The network with n input-units that is characterized by ω is denoted $\mathcal{N}_\omega(n)$, but for brevity we will drop the n and use the notation \mathcal{N}_ω . Finally, the function that \mathcal{N}_ω computes is denoted $f_\omega : R^n \rightarrow R^n$, and the family of all such functions is denoted $\mathcal{F} = \{f_\omega | \omega \in \Lambda\}$.

Our objective is thus to find all the functions that may be approximated by multilayer feedforward networks of the form \mathcal{N}_ω . In order to do so, we will characterize the closure $\bar{\mathcal{F}} = \text{closure}\{f_\omega | \omega \in \Lambda\}$. This closure is based on some metric defined over the set of functions from R^n to R , described in the next section.

3. DEFINITIONS

DEFINITION 1. *A metric on a set S is a function $d : S \times S \rightarrow R$ such that:*

1. $d(s, t) \geq 0$
2. $s = t$ if and only if $d(s, t) = 0$
3. $d(s, t) = d(t, s)$
4. $d(s, u) \leq d(s, t) + d(t, u)$.

If we take S to be a set of functions, the metric $d(f, g)$ will enable us to measure the distance between functions $f, g \in S$.

DEFINITION 2. The closure of a set S of a metric space (X, d) is defined as follows:

$$\text{closure}(S) = \bar{S} = \{t \mid \forall \epsilon > 0, \exists s \in S, d(s, t) < \epsilon\}.$$

DEFINITION 3. A function u defined almost everywhere with respect to Lebesgue measure ν on a measurable set Ω in R^n is said to be essentially bounded on Ω ($u \in L^\infty(\Omega)$), if $|u(x)|$ is bounded almost everywhere on Ω . We denote $u \in L^\infty(\Omega)$ with the norm

$$\begin{aligned} \|u\|_{L^\infty(\Omega)} &= \inf\{\lambda \mid \nu\{x : |u(x)| \geq \lambda\} = 0\} = \text{ess sup}_{x \in \Omega} |u(x)|. \end{aligned}$$

DEFINITION 4. A function u defined almost everywhere with respect to Lebesgue measure on a domain Ω (a domain is an open set in R^n) is said to be locally essentially bounded on Ω ($u \in L^\infty_{loc}(\Omega)$), if for every compact set $K \subset \Omega$, $u \in L^\infty(K)$.

DEFINITION 5. We say that a set F of functions in $L^\infty_{loc}(R^n)$ is dense in $C(R^n)$ if for every function $g \in C(R^n)$ and for every compact set $K \subset R^n$, there exists a sequence of functions $f_j \in F$ such that

$$\lim_{j \rightarrow \infty} \|g - f_j\|_{L^\infty(K)} = 0.$$

Hence, if we can show that a given set of functions F is dense in $C(R^n)$, we can conclude that for every continuous function $g \in C(R^n)$ and each compact set $K \subset R^n$, there is a function $f \in F$ such that f is a good approximation to g on K . In this paper we take $C(R^n)$ to be the family of "real world" functions that one may wish to approximate with feedforward network architectures of the form \mathcal{N}_ω . F is taken to be the family of all functions implied by the network's architecture, namely the family [eqn (1)], when ω runs over all its possible values. The key question is this: Under which necessary and sufficient conditions on σ will the family of networks \mathcal{N} be capable of approximating to any desired accuracy any given continuous function?

4. RESULTS

Let M denote the set of functions which are in $L^\infty_{loc}(R)$ and have the following property. The closure of the set of points of discontinuity of any function in M is of zero Lebesgue measure. This implies that for any $\sigma \in M$, interval $[a, b]$, and $\delta > 0$, there exists a finite number of open intervals, the union of which we denote by U , of measure δ , such that σ is uniformly

continuous on $[a, b]/U$. We will use this fact. Note that we do not demand the existence of one-sided limits at points of discontinuity.

We then have the following result:

THEOREM 1. Let $\sigma \in M$. Set

$$\Sigma_n = \text{span} \{ \sigma(\mathbf{w} \cdot \mathbf{x} + \theta) : \mathbf{w} \in R^n, \theta \in R \}.$$

Then Σ_n is dense in $C(R^n)$ if and only if σ is not an algebraic polynomial (a.e.).

PROPOSITION 1. Assume μ is a non-negative finite measure on R^n with compact support, absolutely continuous with respect to Lebesgue measure. Then Σ_n is dense in $L^p(\mu)$, $1 \leq p < \infty$, if and only if σ is not a polynomial (a.e.).

We recall that $L^p(\mu)$ is the set of all measurable functions f such that:

$$\|f\|_{L^p(\mu)} = \left(\int_{R^n} |f(x)|^p d\mu(x) \right)^{1/p} < \infty.$$

The following proposition is worth stating as it is a simple consequence of Theorem 1 and some known results.

PROPOSITION 2. If $\sigma \in M$ is not a polynomial (a.e.), then

$$\Sigma_n(\mathcal{A}) = \text{span} \{ \sigma(\lambda \mathbf{w} \cdot \mathbf{x} + \theta) : \lambda, \theta \in R, \mathbf{w} \in \mathcal{A} \}$$

is dense in $C(R^n)$ for some $\mathcal{A} \subseteq R^n$ if and only if there does not exist a nontrivial homogeneous polynomial vanishing on \mathcal{A} .

5. DISCUSSION AND CONCLUSION

First, we wish to illustrate why the threshold element is essential in the above theorems. Consider the activation function (without a threshold) $\sigma(x) = \sin(x)$. This function is not a polynomial; In addition, it is continuous, bounded, and non-constant. Now, the set $\{\sin(\mathbf{w} \cdot \mathbf{x}) \mid \mathbf{w} \in R\}$ consists of only odd functions ($\sigma(x) = -\sigma(-x)$). Thus, an even function like $\cos(x)$ cannot be approximated using this family in $[-1, 1]$, implying that $\{\sin(\mathbf{w} \cdot \mathbf{x}) \mid \mathbf{w} \in R\}$ is not dense in $C([-1, 1])$. This could be corrected by adding to the family $\sin(\cdot)$ functions with a threshold (offset) element (e.g., $\sin(x + \pi/2) = \cos(x)$). Moreover, if σ is an entire function, there exist sufficient and necessary conditions on σ under which Theorem 1 will hold without a threshold (for a more general discussion see Dahmen & Micchelli, 1987). On the other hand, the threshold may have absolutely no effect. Take for example the function $\sigma(x) = e^x$.

The essential role of the threshold in our analysis is interesting in light of the biological backdrop of artificial

neural networks. Because most types of biological neurons are known to fire only when their processed inputs exceed a certain threshold value, it is intriguing to note that the same mechanism must be present in their artificial counterparts as well.

In a similar vein, our finding that activation functions need not be continuous or smooth also has an important biological interpretation, because the activation functions of real neurons may well be discontinuous, or even nonelementary. These restrictions on the activation functions have no bearing on our results, which merely require “nonpolynomiality.”

As Hornik (1991) pointed out, “Whether or not the continuity assumption can entirely be dropped is still an open and quite challenging problem.” We hope that our results solve this problem in a satisfactory way.

6. PROOFS

We use the following definition to prove our main results:

DEFINITION 6. For a function u we denote by $\text{supp}(u)$ the set $\text{supp}(u) = \{x | u(x) \neq 0\}$.

Proof of Theorem 1. We divide the proof into a series of steps.

Step 1. If σ is a polynomial, then Σ_n is not dense in $C(R^n)$.

If σ is a polynomial of degree k , then $\sigma(\mathbf{w} \cdot \mathbf{x} + \theta)$ is a polynomial of degree k for every \mathbf{w} and θ , and in fact Σ_n is exactly the set of algebraic polynomials of degree at most k . Thus Σ_n cannot be dense in $C(R^n)$. ■

In what follows we always assume that σ is not a polynomial.

Step 2. If Σ_1 is dense in $C(R)$, then Σ_n is dense in $C(R^n)$.

The space $V = \text{span}\{f(\mathbf{a} \cdot \mathbf{x}) | \mathbf{a} \in R^n, f \in C(R)\}$ is dense in $C(R^n)$. This follows in various ways, [e.g., Dahmen & Micchelli (1987), Chui & Li (1992), Vostrecov & Kreines (1961), Lin & Pinkus (in press)]. Now, let $g \in C(R^n)$ and $K \subset R^n$ be any compact subset of R^n . V is dense in $C(K)$. Thus given $\epsilon > 0$ there exist $f_i \in C(R)$ and $\mathbf{a}^i \in R^n, i = 1, \dots, k$, such that

$$|g(\mathbf{x}) - \sum_{i=1}^k f_i(\mathbf{a}^i \cdot \mathbf{x})| < \epsilon/2,$$

for all $\mathbf{x} \in K$. Now $\{\mathbf{a}^i \cdot \mathbf{x} | \mathbf{x} \in K\} \subseteq [\alpha_i, \beta_i]$ for some finite interval $[\alpha_i, \beta_i], i = 1, \dots, k$. Because Σ_1 is dense in $[\alpha_i, \beta_i], i = 1, \dots, k$, there exist constants c_{ij}, w_{ij} and $\theta_{ij}, j = 1, \dots, m_i, i = 1, \dots, k$, such that

$$|f_i(y) - \sum_{j=1}^{m_i} c_{ij} \sigma(w_{ij}y + \theta_{ij})| < \epsilon/2k,$$

for all $y \in [\alpha_i, \beta_i]$. Thus,

$$|g(\mathbf{x}) - \sum_{i=1}^k \sum_{j=1}^{m_i} c_{ij} \sigma(w_{ij}(\mathbf{a}^i \cdot \mathbf{x}) + \theta_{ij})| < \epsilon,$$

for all $\mathbf{x} \in K$. Thus Σ_1 dense in $C(R)$ implies that Σ_n is dense in $C(R^n)$. ■

Step 3. If $\sigma \in C^\infty$ (the set of all functions which have derivatives of all order), then Σ_1 is dense in $C(R)$.

If $\sigma \in C^\infty(R)$ then because $[\sigma((w+h)x + \theta) - \sigma(wx + \theta)]/h \in \Sigma_1$ for every $w, \theta \in R$ and $h \neq 0$, it follows that $(d/dw)\sigma(wx + \theta) \in \overline{\Sigma_1}$. By the same argument $(d^k/dw^k)\sigma(wx + \theta) \in \overline{\Sigma_1}$ for all $k \in N$ (and all $w, \theta \in R$). Now $(d^k/dw^k)\sigma(wx + \theta) = x^k \sigma^{(k)}(wx + \theta)$, where $\sigma^{(k)}$ denotes the k th derivative of σ , and since σ is not a polynomial there exists a $\theta_k \in R$ such that $\sigma^{(k)}(\theta_k) \neq 0$. Thus,

$$x^k \sigma^{(k)}(\theta_k) = \left. \frac{d^k}{dw^k} \sigma(wx + \theta) \right|_{w=0, \theta=\theta_k} \in \overline{\Sigma_1}.$$

This implies that $\overline{\Sigma_1}$ contains all polynomials. By Weierstrass's Theorem it follows that $\overline{\Sigma_1}$ contains $C(K)$ for each $K \subset R$. That is, Σ_1 is dense in $C(R)$. ■

Step 4. For each $\varphi \in C_0$, (C^∞ function with compact support), $\sigma * \varphi \in \overline{\Sigma_1}$.

We first recall that

$$(\sigma * \varphi)(x) = \int \sigma(x - y)\varphi(y)dy,$$

is the convolution of σ and φ , and is well-defined. We prove Step 4 constructively. (If σ were continuous this could easily be proven using a soft analysis approach.)

Without loss of generality, assume that $\text{supp } \varphi \subseteq [-\alpha, \alpha]$, and that we wish to prove that we can uniformly approximate $\sigma * \varphi$ from Σ_1 on $[-\alpha, \alpha]$. We will prove that

$$\sum_{i=1}^m \sigma(x - y_i)\varphi(y_i)\Delta y_i$$

uniformly converges to $\sigma * \varphi$ on $[-\alpha, \alpha]$, where

$$y_i = -\alpha + \frac{2i\alpha}{m}, \quad i = 1, \dots, m,$$

and $\Delta y_i = 2\alpha/m, i = 1, \dots, m$.

Given $\epsilon > 0$, we first choose $\delta > 0$ so that:

$$10\delta \|\sigma\|_{L^\infty[-2\alpha, 2\alpha]} \|\varphi\|_{L^\infty} \leq \epsilon. \tag{2}$$

For this given $\delta > 0$, we know that there exists a finite

number $r(\delta)$ of intervals, the measure of whose union U is δ , such that σ is uniformly continuous on $[-2\alpha, 2\alpha]/U$. We now choose m sufficiently large so that $m\delta > \alpha r(\delta)$, and:

If $|s - t| \leq 2\alpha/m$, then

$$|\varphi(s) - \varphi(t)| \leq \frac{\epsilon}{2\alpha\|\sigma\|_{L^\infty[-2\alpha, 2\alpha]}}. \tag{3}$$

If $s, t \in [-2\alpha, 2\alpha]/U$, and $|s - t| \leq 2\alpha/m$, then

$$|\sigma(s) - \sigma(t)| \leq \frac{\epsilon}{\|\varphi\|_{L^1}}. \tag{4}$$

All these conditions can be satisfied. Equation (3) follows from the uniform continuity of φ . By assumption σ is uniformly continuous on $[-2\alpha, 2\alpha]/U$ and thus eqn (4) holds.

Fix $x \in [-\alpha, \alpha]$. Set $\Delta_i = [y_{i-1}, y_i]$, ($y_0 = \alpha$). Now

$$\begin{aligned} & \left| \int \sigma(x - y)\varphi(y)dy - \sum_{i=1}^m \int_{\Delta_i} \sigma(x - y_i)\varphi(y)dy \right| \\ & \leq \sum_{i=1}^m \int_{\Delta_i} |\sigma(x - y) - \sigma(x - y_i)| |\varphi(y)| dy, \end{aligned}$$

because $\text{supp } \varphi \subset [-\alpha, \alpha]$. If $x - \Delta_i$ does not intersect the U , then from eqn (4),

$$\int_{\Delta_i} |\sigma(x - y) - \sigma(x - y_i)| |\varphi(y)| dy \leq \frac{\epsilon}{\|\varphi\|_{L^1}} \int_{\Delta_i} |\varphi(y)| dy.$$

Thus if we sum over those Δ_i for which this holds we get an error of at most ϵ .

Let us now consider those intervals Δ_i for which $(x - \Delta_i) \cap U \neq \emptyset$. We denote such intervals by $\tilde{\Delta}_i$. Because U has measure δ and is composed of $r(\delta)$ intervals, the total length of the $\tilde{\Delta}_i$ intervals is at most $\delta + (4\alpha/m)r(\delta)$. By our choice of m , we have that $\delta + (4\alpha/m)r(\delta) \leq 5\delta$. Thus from eqn (2),

$$\begin{aligned} & \sum \int_{\tilde{\Delta}_i} |\sigma(x - y) - \sigma(x - y_i)| |\varphi(y)| dy \\ & \leq 2\|\sigma\|_{L^\infty[-2\alpha, 2\alpha]} \|\varphi\|_{L^1} 5\delta \leq \epsilon. \end{aligned}$$

Finally,

$$\begin{aligned} & \left| \sum_{i=1}^m \int_{\Delta_i} \sigma(x - y_i)\varphi(y)dy - \sum_{i=1}^m \sigma(x - y_i)\varphi(y_i)\Delta y_i \right| \\ & = \left| \sum_{i=1}^m \int_{\Delta_i} \sigma(x - y_i)[\varphi(y) - \varphi(y_i)]dy \right| \\ & \leq \sum_{i=1}^m \int_{\Delta_i} |\sigma(x - y_i)| |\varphi(y) - \varphi(y_i)| dy, \end{aligned}$$

and from eqn (3)

$$\leq \sum_{i=1}^m \int_{\Delta_i} |\sigma(x - y_i)| dy \left[\frac{\epsilon}{2\alpha\|\sigma\|_{L^\infty[-2\alpha, 2\alpha]}} \right] \leq \epsilon.$$

Thus, we obtain

$$\left| \int \sigma(x - y)\varphi(y)dy - \sum_{i=1}^m \sigma(x - y_i)\varphi(y_i)\Delta y_i \right| \leq 3\epsilon,$$

for all $x \in [-\alpha, \alpha]$. ■

Step 5. If for some $\varphi \in C_0^\infty$ we have that $\sigma*\varphi$ is not a polynomial, then Σ_1 is dense in $C(R)$.

From Step 4, $\sigma*\varphi \in \overline{\Sigma_1}$. It thus follows that $(\sigma*\varphi)(wx + \theta)$ is also in $\overline{\Sigma_1}$, for each $w, \theta \in R$. Now for σ and any $\varphi \in C_0^\infty$, we have $\sigma*\varphi \in C^\infty$, see Adams (1975, pp. 29-31). Thus from Step 3, if $\sigma*\varphi$ is not a polynomial then Σ_1 is dense in $C(R)$. ■

We therefore now assume that $\sigma*\varphi$ is a polynomial for all $\varphi \in C_0^\infty$. We will conclude from this fact that σ is itself a polynomial (a.e.).

Step 6. If for all $\varphi \in C_0^\infty$, $\sigma*\varphi$ is a polynomial, then there exists an $m \in N$ such that $\sigma*\varphi$ is a polynomial of degree at most m for all $\varphi \in C_0^\infty$.

For any $a < b$, define the set of functions $C_0^\infty[a, b]$ to be the set of all C_0^∞ functions with support in $[a, b]$. We first prove the claim in the case of $\varphi \in C_0^\infty[a, b]$. We define a metric ρ on $C_0^\infty[a, b]$ by:

$$\rho(\varphi_1, \varphi_2) = \sum_{n=0}^{\infty} 2^{-n} \frac{\|\varphi_1 - \varphi_2\|_n}{1 + \|\varphi_1 - \varphi_2\|_n},$$

where $\|\varphi\|_n = \sum_{j=0}^n \sup_{x \in [a, b]} |\varphi^{(j)}(x)|$. $C_0^\infty[a, b]$ with the ρ metric is a complete metric vector space (Fréchet space).

By assumption $\sigma*\varphi$ is a polynomial for any $\varphi \in C_0^\infty[a, b]$.

Define:

$$V_k = \{ \varphi \in C_0^\infty[a, b] \mid \text{degree}(\sigma*\varphi) \leq k \}.$$

We have that V_k is a closed subspace, $V_k \subseteq V_{k+1}$, and

$$\bigcup_{k=0}^{\infty} V_k = C_0^\infty[a, b].$$

As $C_0^\infty[a, b]$ is a complete metric space, by Baire's Category Theorem [Bachman and Narici (1972, p. 77)] there exists an integer m such that $V_m = C_0^\infty[a, b]$ ($C_0^\infty[a, b]$ is of the second category and therefore some V_m contains a non-void open set. Because V_m is a vector space thus $V_m = C_0^\infty[a, b]$). This completes the proof for the $C_0^\infty[a, b]$ case. For the general case we note that the number m does not depend on the interval $[a, b]$. This can be seen as follows. By translation m depends at most of the length of the interval. Let $[A, B]$ be any interval. For $\varphi \in C_0^\infty[A, B]$ we can find $\varphi_i \in C_0^\infty[a_i, b_i]$, $i = 1, \dots, k$, such that $[A, B] \subseteq \bigcup_{i=1}^k [a_i, b_i]$, $b_i -$

$a_i = b - a$ and $\varphi = \sum_{i=1}^k \varphi_i$. Thus $\sigma * \varphi = \sum_{i=1}^k \sigma * \varphi_i$, and for every $i = 1, \dots, k$, $\sigma * \varphi_i$ is a polynomial of degree less than or equal to m . Therefore $\text{degree}(\sigma * \varphi) \leq m$. ■

Step 7. If $\sigma * \varphi$ is a polynomial of degree at most m for all $\varphi \in C_0^\infty$, then σ is a polynomial of degree at most m (a.e.).

From Step 6,

$$\int \sigma(x - y)\varphi^{(m+1)}(y)dy = 0$$

for all $\varphi \in C_0^\infty$. From standard results in Distribution Theory, [e.g., Friedman (1963, 57-59)], σ is itself a polynomial of degree at most m (a.e.). ■

REMARK 1. Step 6 is one of those folklore results we were rather surprised not to have succeeded in finding in the literature. There are other proofs thereof.

REMARK 2. A reading of the proof of Theorem 1 shows that the problem of approximating a function g on some compact K of R^n from Σ_n can almost be divided into two parts. One part is the approximation of $g(x)$ by functions of the form $\sum_i f_i(\mathbf{a}^i \cdot \mathbf{x})$ where the f_i are functions in $C(R)$. The other is the approximation of f_i on the appropriate set from Σ_1 . Because $C(R)$ is separable, one can choose $\sigma \in C(R)$ so that for each and every $f \in C(R)$ and any interval $[a, b]$,

$$0 = \inf_{c, w, \theta} \max_{a \leq x \leq b} |f(x) - c\sigma(wx + \theta)|.$$

That is, only one "processing unit" is needed. However there remains the problem of approximating $g(x)$ by $\sum_i f_i(\mathbf{a}^i \cdot \mathbf{x})$ (these latter are called ridge functions or plane waves), which seems to be the more difficult problem.

REMARK 3. If σ has a jump discontinuity, say at 0, and is continuous in $[-\eta, 0)$ and $(0, \eta]$ (some $\eta > 0$) with $\lim_{x \rightarrow 0^-} \sigma(x)$ and $\lim_{x \rightarrow 0^+} \sigma(x)$ existing and unequal, then one can obtain Theorem 1 almost directly (from after Step 2). That is, given any $f \in C(R)$ and any K compact in R , it is possible to approximate f from Σ_1 on K . Constants are in $\Sigma_1(c\sigma(\theta))$, and thus choosing $w \in \{-1, 1\}$ and multiplying by a constant we can assume that

$$\lim_{x \rightarrow 0^-} \sigma(x) = 0, \quad \lim_{x \rightarrow 0^+} \sigma(x) = 1.$$

Letting $w \rightarrow 0$ in $\sigma(wx)$, we can then prove that the function $\chi \in \Sigma_1$, where $\chi(x) = 0$ for $x < 0$, and $\chi(x) = 1$ for $x > 0$. It is now easy to see how linear combinations of χ and its translates can uniformly approximate any continuous function on any finite interval (and thus any compact subset of R).

REMARK 4. There is another method of proof of Theorem 1 for continuous functions, which is simple but

based on a deep result. It comes from the theory of mean-periodic functions introduced by Schwartz (1947). One consequence of that theory is that if $\sigma \in C(R)$ is not a polynomial, then

$$\text{span}\{\sigma(x + \theta) : \theta \in R\},$$

contains in its closure a function of the form $e^{\lambda x} \cos vx$ for some real $(\lambda, v) \neq (0, 0)$. Because any such function is in C^∞ , Theorem 1 then follows without the necessity of proving Steps 4-7.

Proof of Proposition 2. If σ is a polynomial of degree m , then Σ_n is contained in the set of polynomial of total degree $\leq m$, and thus cannot be dense in $L^p(\mu)$, $1 \leq p < \infty$.

Let K denote the support of μ . $C(K)$ is dense in $L^p(\mu)$ [e.g., Adams (1975, p. 31)] and Σ_n is dense in $C(K)$ in the uniform norm. Thus given $f \in L^p(\mu)$ and $\epsilon > 0$ there exists a $g \in C(K)$ such that

$$\|f - g\|_{L^p(\mu)} \leq \epsilon/2,$$

and for this given $g \in C(K)$ there exists an $h \in \Sigma_n$ such that

$$\|g - h\|_{L^\infty(K)} \leq \frac{\epsilon}{2c}.$$

where $c = \mu^{1/p}(K)$. Thus $\|g - h\|_{L^p(\mu)} \leq \epsilon/2$, and

$$\|f - h\|_{L^p(\mu)} \leq \|f - g\|_{L^p(\mu)} + \|g - h\|_{L^p(\mu)} \leq \epsilon. \quad \blacksquare$$

Proof of Proposition 3. In Vostrecov and Kreines (1961) [see also Lin and Pinkus (in press)] can be found the fact that for given $\mathcal{A} \subset R^n$

$$M(\mathcal{A}) = \text{span}\{f(\mathbf{w} \cdot \mathbf{x}) | f \in C(R), \mathbf{w} \in \mathcal{A}\},$$

is dense in $C(R^n)$ if and only if there does not exist a non-trivial homogeneous polynomial vanishing on \mathcal{A} . Now $\text{span}\{\sigma(\lambda \mathbf{w} \cdot \mathbf{x} + \theta) | \lambda, \theta \in R\} \subseteq \text{span}\{f(\mathbf{w} \cdot \mathbf{x}) | f \in C(R)\}$ for every $\mathbf{w} \in \mathcal{A}$. This proves the necessity. To prove the sufficiency assume $M(\mathcal{A})$ is dense in $C(R^n)$ and use the argument as given in Step 2 of the proof of Theorem 1 to show that if Σ_1 is dense in $C(R)$ then $\Sigma_n(\mathcal{A})$ is dense in $C(R^n)$. ■

REFERENCES

Adams, R. A. (1975). *Sobolov spaces*. New York: Academic Press.
 Bachman, G., & Narici, L. (1972). *Functional analysis* (5th ed.). New York: Academic Press.
 Chui, C. K., & Xin Li. (1992). Approximation by ridge functions and neural networks with one hidden layer. *Journal of Approximation Theory*, **70**, 131-141.
 Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, **2**, 303-314.
 Dahmen, W., & Micchelli, C. A. (1987). Some remarks on ridge functions. *Approximation Theory and its Applications*, **3**, 2-3.
 Friedman, A. (1963). *Generalized functions and partial differential equations*. Englewood Cliffs, NJ: Prentice-Hall.

- Funahashi, K. (1989). On the approximate realization of continuous mappings by neural networks. *Neural Networks*, **2**, 183–192.
- Gallant, A. R. & White, H. (1988). There exists a neural network that does not make avoidable mistakes. *Proceedings of the IEEE Second International Conference on Neural Networks* (I, 657–664). San Diego: SOS Printing.
- Hecht-Nielsen, R. (1989). Theory of the back propagation neural network. *Proceedings of the International Joint Conference on Neural Networks*. (I, 593–606). San Diego: SOS Printing.
- Hinton, G. E. (1989). Connectionist learning procedure. *Artificial Intelligence*, **40**, 185–234.
- Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks*, **4**, 251–257.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, **2**, 359–366.
- Irie, B., & Miyake, S. (1988). Capabilities of three layer perceptrons. *Proceedings of the IEEE Second International Conference on Neural Networks* (I, 641–648). San Diego: SOS Printing.
- Lapedes, A., & Farber, R. (1988). *How neural networks work* (Tech. Rep. LA-UR-88-418). Los Alamos, NM: Los Alamos National Library.
- LeCun, Y. (1987). *Models connexionistes de l'apprentissage*. Master's thesis, Université Pierre et Marie Curie, Paris.
- Lin, V. Ya., & Pinkus, A. (in press). Fundamentality of ridge functions. *Journal of Approximation Theory*.
- Schwartz, L. (1947). Théorie générale des fonctions moyenne-périodiques. *Annals of Mathematics*, **48**, 857–929.
- Stinchcombe, M., & White, H. (1990). *Approximating and learning unknown mappings using multilayer feedforward networks with bounded weights* (Tech. Rep.). San Diego: University of California, Dept. of Economics.
- Vostrecov, B. A., & Kreines, M. A. (1961). Approximation of continuous functions by superposition of plane waves. *Doklady Akademia Nauk SSSR* (Soviet Math. Dokl.), **140**, 2.