

Single-layer lookup perceptrons

G.D. Tattersall, BTech, MIEE
S. Foster, BEng
R.D. Johnston, BSc, MIEE

Indexing terms: Neural networks, Perceptrons

Abstract: This paper examines the generalisation properties of various types of neural network, such as radial basis function systems and the multilayer perceptron (MLP). It is concluded that their behaviour can be explained in terms of lowpass interpolation in which discrete training examples of a function are implicitly convolved with the impulse response of a lowpass filter to produce an estimate of the function for previously unseen arguments. A different form of neural network, in the form of a single-layer lookup perceptron (SLLUP), is described, and this type of perceptron is shown to also generalise by lowpass interpolation. However, the SLLUP can learn reliably and rapidly compared to the multilayer perceptron and experiments are described which show that it compares well with the MLP on problems such as speech recognition and text-to-speech synthesis.

1 Introduction

In this paper the nature of generalisation by neural networks is related to the established signal-processing concepts of filtering, convolution and interpolation and it is argued that neural networks such as the multilayer perceptron (MLP) [1] and radial basis function (RBF) systems [2] are examples of lowpass interpolating systems which can only generalise on data which exhibit simple clustering. Although limited, this form of generalisation is often very useful when the data being processed by the net represents physical quantities, and in this paper an alternative lowpass interpolating perception is described which is able to perform this type of generalisation but with fast learning and little computation compared to the MLP and RBF.

The perceptron incorporates n -tuple pattern recognition techniques [3] in a single-layer architecture to produce a single-layer lookup perceptron (SLLUP) whose hardware realisation could be almost identical to WISARD [4] but whose *modus operandi* is very different. The SLLUP is based on a nonlinear adaptive filter proposed by Johnston [5] which has been used for echo cancellation and is functionally similar to the cerebellar model articulation controller (CMAC) [6] used in learn-

ing control systems. It is shown that this type of system can be used in most of the applications which are currently seen as the domain of neural networks such as the multilayer perceptron (MLP). The SLLUP can learn the same types of nonlinear mappings as an MLP but with a fraction of the training and computation, and an additional very desirable property of the SLLUP is that it produces a quadratic error surface and so convergence to optimal performance is assured.

It will be shown that the SLLUP is basically an interpolation system which is able to generate an estimate of a continuous mapping function from a sparse set of training examples and, as will be demonstrated, is well suited to dealing with simple nonlinear mappings such as parity detection.

The ability of the SLLUP to work on the very complex mapping problems of speech recognition and text-to-speech synthesis is also examined and compared with the performance obtainable using the multilayer perceptron. It will be seen that the SLLUP can very nearly equal the performance of the MLP in these problems, suggesting that the MLP also does little more than a straightforward sample interpolation.

2 Generalisation by neural networks

Any type of neural network can be visualised as a vector transformer which accepts an input vector or pattern X and produces an output vector Y . The functional relationship between Y and X is learned from a relatively sparse set of examples of input-output pairs which are shown to the network during a supervised learning phase. A typical arrangement for performing the learning is shown in Fig. 1, in which a training example of the

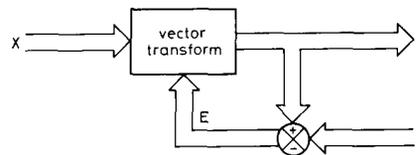


Fig. 1 Supervised learning system

required function, $f(X)$, is used as an output target for the transformer. The error between the actual output of the transformer and the target is used to adapt the internal parameters of the transformer until the error is minimised. In the context of neural networks the internal parameters are the 'synaptic weight' values and the error is used in a gradient descent algorithm such as back propagation.

A desirable property of neural networks is *generalisation*, which enables the network to produce a good

Paper 7735F (E5), first received 5th March and in revised form 3rd October 1990

Mr. Tattersall and Mr. Foster are with the School of Information Systems, University of East Anglia, Norwich NR4 7TJ, United Kingdom

Mr. Johnston is with British Telecom Research Laboratories, Martlesham Heath, Ipswich IP5 7RE, United Kingdom

estimate of an output Y given a previously unseen value of input X . In many cases, this is simply achieved by some form of interpolation. The process is illustrated in Figs. 2a and b for a one-dimensional case in which the

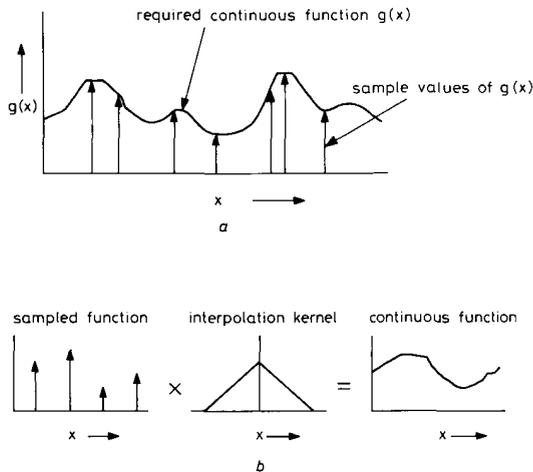


Fig. 2 Interpolation

a Training examples as samples of the underlying function
b Recovering a continuous function from its samples by filtering

function underlying a training set of discrete input-output example pairs is $y = g(x)$. Fig. 2a shows that the input-output examples are effectively samples of the function $g(x)$ and, although the samples are irregularly positioned across the pattern space, many of the ideas used in digital signal processing, such as Nyquist Sampling, aliasing, filtering and function bandwidths, are also relevant to the problem of finding this function.

In particular, if the input-output examples are samples of the underlying function, then the complete continuous function should be recoverable by passing them through a suitable lowpass interpolation filter in which they are convolved with the filter's impulse response, as illustrated in Fig. 2b. This is the essence of the radial basis function system and the single-layer lookup perceptron described in this paper.

Insight into the nature of the generalisation obtained using interpolative filtering is gained by viewing the process in the frequency domain. As an example consider a filter whose impulse response is a multivariate gaussian. The filter's frequency response is lowpass and consequently the correct continuous function will only be generated by filtering the sampled function if the spectrum of the unsampled function is also lowpass, as illustrated in Fig. 3. To emphasise this limitation, the generalisation will henceforth be called lowpass interpolation to indicate that it is only suitable if the spectrum of the function underlying the training samples is also lowpass.

In contrast to Fig. 3, Fig. 4 shows a situation in which the function underlying the samples has a bandpass spectrum. Lowpass filtering the spectrum of the sampled function will not yield the correct continuous function because a frequency shifted version of the wanted continuous function is produced.

It may be argued that real-world data will only arise from functions having a lowpass spectrum and this is often true for data which has been generated by physical processes such as speech articulation, whose generating

system contains mechanical inertia which tends to cause the spectral content of the function to be dominated by low frequencies. Such functions are appropriately gener-

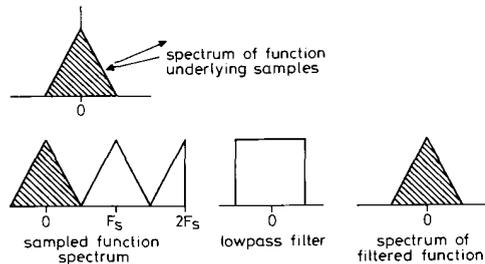


Fig. 3 Spectral view of generalisation by lowpass interpolation

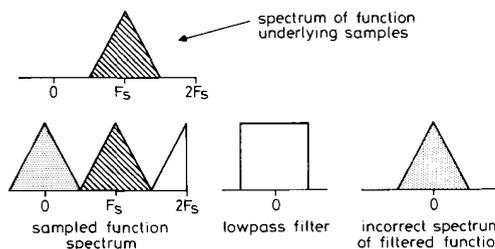


Fig. 4 Spectral view of incorrect generalisation of a function with bandpass spectrum

alised by lowpass interpolation like the radial basis function and, as will be shown shortly, sigmoidal multilayer perceptrons.

However, functions which describe logical processes which relate a set of input conditions to an output action are often characterised by bandpass spectra. A simple example is the parity function, which contains no energy at zero frequency because the function oscillates in value as each input variable changes state. Lowpass interpolation is quite inappropriate for generalising in this kind of situation and, predictably, it will usually be found that radial basis function systems or sigmoidal MLPs fail to work correctly when only trained on a subset of possible input-output examples. Correct generalisation is still possible with these types of function if a bandpass or multiband interpolation scheme is used. The Fourier perception [7] is an example of such a bandpass interpolating system, but in this paper the discussion will be limited to the case of lowpass interpolation.

3 Nyquist sampling, filtering and generalisation

Three important points are raised by viewing the process of function generalisation as interpolative filtering of a sparse set of samples of the function. These points are equally applicable to lowpass and bandpass interpolation and provide insight into the limitations of any generalising machine.

3.1 Nyquist sampling criteria

Sufficient training examples must be given so that there are minimum of two samples per cycle of the highest frequency in the mapping function. This suggests that the complexity of a mapping function be specified in terms of its bandwidth B and that the maximum interval between training examples should be no greater than $1/(2B)$.

3.2 Bandwidth of interpolation filters

To obtain a perfect, continuous function from the training examples, the interpolation filter should have a rectangular frequency response of B Hz bandwidth. This would require an infinite length sinc (x) impulse response, which is impossible. Practical interpolation filters should have a bandwidth of as near as possible B Hz, but it must be recognised that finite impulse response filters will always produce an error in the estimate of the continuous function.

3.3 Uniformity of function sampling interval

It is very unlikely that the training examples supplied to a supervised learning machine will be uniformly distributed across the pattern space. This means that the distances between samples of the required mapping function are non-uniform. The Nyquist sampling theorem requires at least two samples per cycle of the function if it is to be recovered without loss of information. However, a uniform sample interval is not specified and so the irregularity of the training points does not necessarily mean that the continuous mapping function cannot be recovered. Unfortunately, a simple interpolation filter is unable to recover a continuous function from a set of irregular samples because the function will be non-uniformly scaled in proportion to the density of the samples. This is not a problem if the interpolation is learned by iteratively adjusting the scale factor of the impulse response of the interpolation filter depending on which region of the pattern space it is operating in, and an example of this process is shown by the radial basis function system described in the next Section.

4 Radial basis functions as lowpass interpolators

The operation of the radial basis function network is illustrated in Fig. 5, in which a set of basis functions are

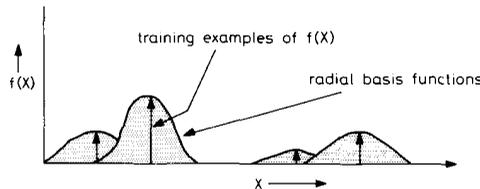


Fig. 5 Synthesis of a function using radial basis functions

added together so that their sum closely fits discrete training examples of a continuous function which is to be learned by the system. Typically the basis functions are multivariate gaussians whose amplitude, mean value, and variance can be scaled to make the network output fit the given examples of the function.

The lowpass filter action of the network is easily understood by initially assuming that the training samples of the function are regularly spaced and that one basis function is positioned over every sample in the function domain. In this situation the output of the network is the convolution of the training samples with the radial basis function. Spectrally, this is lowpass filtering because the Fourier transform of the multivariate gaussian is a lowpass frequency response.

In general the samples are not positioned regularly and the bandwidth and amplitude of each of the basis functions must be individually adjusted to match the sample rate in its locality. Thus, if the samples are very sparse in one region of the pattern space, the bandwidth

of the basis function has to be reduced so that proper interpolation can occur. This corresponds to increasing the variance of the basis function.

Very often it is not possible to place a basis function over every single training sample because of computational load, and in this case, the function is subsampled by using a relatively small number of basis functions for its synthesis. The positions of the basis functions in the pattern space are adapted iteratively to optimise the accuracy of the synthesised function and the bandwidth of each radial basis function is reduced to reflect the lower effective sample rate of the function.

5 MLPs as lowpass interpolators

Sigmoidal MLPs also perform as lowpass interpolators and are therefore only capable of generalising correctly on functions whose spectrum is lowpass. This behaviour is easily understood by looking at the three-layer scalar-to-scalar mapping network with one layer of hidden units shown in Fig. 6. The output of the network is the

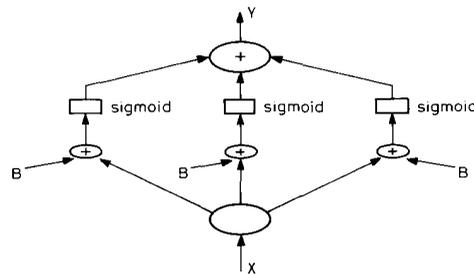


Fig. 6 MLP as a synthesiser of a mapping function

weighted summation of the sigmoid functions produced by each of the hidden units and the slope and offset of each of the functions is set by the weights connecting each of the hidden units to the input units. All the weights are adjusted during learning until the desired output function is synthesised as illustrated in Fig. 7.

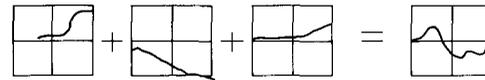


Fig. 7 Synthesis of a mapping function from the weighted sum of hidden-unit outputs

Assuming the weight values feeding into the hidden units are not allowed to become very high, each hidden unit produces a smoothly changing basis function which contributes to the overall function.

Any smooth, non-periodic function, such as the sigmoid, will have a lowpass spectrum and consequently the MLP will perform generalisation by lowpass interpolation and is incapable of correctly generalising functions whose spectrum is bandpass.

6 The single-layer lookup perceptron (SLLUP)

The SLLUP is another example of a lowpass interpolating system. However, it has clear advantages in terms of learning speed and computation when compared to the MLP and RBF, and it will be described in detail in

this paper. The SLLUP is based upon a nonlinear adaptive filter proposed by Johnston [5] and is, in part, inspired by the WISARD [4, 15, 16] system and the technique of n -tuple sampling proposed by Bledsoe and Browning [3]. The operation of the SLLUP is best understood in terms of the WISARD style architecture shown in Fig. 8, although it will become apparent that

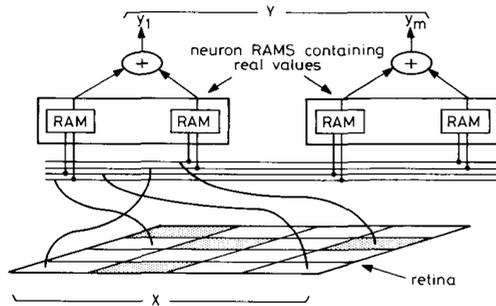


Fig. 8 WISARD architecture used as a SLLUP

the way in which the SLLUP uses the architecture and, indeed, the practical realisation of the SLLUP are different from WISARD.

The WISARD architecture consists of a retina in which an input pattern X is encoded as an image of black and white pixels formed by bits of the code representing the scalar elements of X . Random connections are made onto the pixels in the image and groups of n connections are formed into n -tuples which are used to address a large number of RAMs. The RAMs themselves are grouped into 'neuron' blocks called discriminators and the outputs of all the RAMs in the i th block are added to form y_i , the value of the i th element of the output vector Y .

WISARD is a *pattern classifier* whose function is to produce a high score at the output of a single-class discriminator when a pattern belonging to the class is applied to the retina. This is achieved by associating a single discriminator with each of the pattern classes to be recognised and then applying example patterns of each class. A value of '1' is stored in each addressed location of the RAMs in the discriminator and it can be shown [8] that, after training with many examples, the output value produced by the discriminator for class C_i is approximately proportional to class conditional probability $P(C_i|X)$ of the input pattern X . The conditional probabilities are then used to perform Bayesian classification of input patterns of unknown class. Various modifications can be made to improve the estimate of the class conditional probabilities, such as making the value of a RAM location equal to the number of times it has been addressed during training, rather than '1' [9]. However, the fundamental function is to provide a set of probability estimates which can be used for Bayesian classification.

In contrast to the WISARD pattern classifier, most supervised neural nets are designed to learn arbitrary mapping functions between input and output vectors. The SLLUP uses n -tuple sampling techniques to implement arbitrary mapping functions in the same way as a conventional neural net but with the computational simplicity and learning reliability of the WISARD system.

The SLLUP can be implemented using the same hardware as the WISARD with the exceptions that each

RAM location must be several bits wide so that a wide range of values can be represented, no class is associated with each of the discriminators, and, most importantly, a different training algorithm is used. In reality, it is not practical to implement the SLLUP using a WISARD style architecture because of its inefficient partitioning of memory, and it is usually better to use a single, large multiplexed memory in conjunction with hash addressing rather than small individual RAMs with n -tuple connections.

In the SLLUP mode of operation, the output of each 'discriminator' forms an element of the output vector of the SLLUP and the system is trained by applying a vector X to its input which causes a specific set of n -tuple addresses to be generated that access corresponding contents in each of the RAMs. The summation of the outputs of each group of RAMs produces the elements of the output vector Y . This vector is compared with the desired output T and the resultant error vector R is used to modify the values of the currently addressed RAM locations so that the next time the same input vector is applied the output Y is nearer to the desired output T .

Repeated application of different training vectors allows the system to learn the required input-output mapping $Y = f(X)$. It is important to notice that with appropriate choice of n -tuple order and number of RAMs in each neuron block, the system can estimate the best function $f(X)$ to fit a rather sparse training set. That is, it is not necessary to expose the machine to all possible input-output vector pairs because it is able to interpolate the required function between training points. This property will be analysed in more detail later in the paper.

6.1 Analysis of the SLLUP learning procedure

The adaptation of the RAM contents to develop the required mapping function is performed using the steepest descent algorithm [10, 17] to minimise the mean square error between the actual outputs and target outputs of the system. It is therefore necessary to obtain a value for the derivative of the average error power with respect to the values stored in each location of each RAM so that the required change in each RAM location value can be determined.

Each RAM in the system is addressed by an n -tuple whose value depends on the vector X contained in the input image. Thus, the output of each RAM in the system depends in some complex way on X , such that the output of the j th RAM in the i th 'neuron' block can be expressed as $C_{ij}(X)$. The output of the i th neuron is then given by eqn. 1 where Q is the number of RAMs per neuron:

$$y_i = \sum_{j=1}^Q C_{ij}(\bar{x}) \quad (1)$$

If the target output vector when X is input is $T = [t_1, \dots, t_N]$, then eqn. 1 can be used to express the mean square output error of the system as

$$E^2 = \frac{1}{N} \sum_{i=1}^N \left\{ \sum_{k=1}^Q C_{ik}(\bar{x}) - t_i \right\}^2 \quad (2)$$

Eqn. 2 shows that the highest power term involving $C_{ij}(X)$ in the expression for mean square error is two. This indicates a single minimum, quadratic surface, and so convergence to a global optimum is guaranteed using a gradient descent algorithm. The gradient term required is simply calculated from eqn. 2 as

$$\frac{\partial E^2}{\partial C_{ij}^n(X)} = 2(y_i - t_i) = 2e_i \quad (3)$$

where e_i is the difference between the output of the i th neuron and its target value. So, the algorithm for modifying the RAM contents becomes

$$C_{ij}^{n+1}(X) = C_{ij}^n(X) + k(y_i - t_i) \quad (4)$$

6.2 The SLLUP interpolation kernel

It can be shown that the SLLUP performs generalisation by lowpass filtering the training samples in the same way as a radial basis function network. The basis functions or interpolation kernels are generated implicitly by the architecture of the SLLUP and their form depends strongly on the way in which the input vector X is encoded, but it will be seen that typically they have a double-sided decaying exponential shape. The frequency response of a lowpass filter having this kind of impulse response is far from an ideal 'brick wall' interpolation filter, but the poor cutoff response is not a problem if the mapping function is sampled at a sufficiently high rate such that insignificant amounts of aliasing occur. This means using a sufficient number of examples in training the network.

The simplest coding, from an analytical point of view, is bar chart coding in which the value of each element of X is represented by the number of 'on' bits in a bar, and the operation of the SLLUP using this code will be analysed before going on to consider other input-vector encoding techniques using binary and Gray codes as well as other codes with controlled redundancy.

6.3 Bar chart coding

A bar chart encoding of a two-dimensional pattern $[x_1, x_2]$ is shown in Fig. 9. In this example each vector

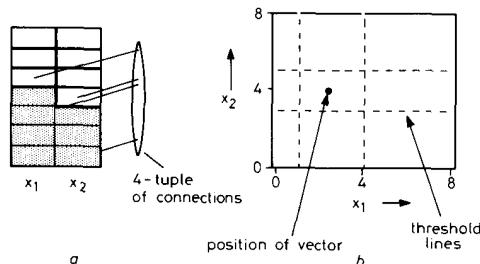


Fig. 9 Bar chart coding of input vector

a Bar chart coding of 2D vector [4, 2]
b Pattern space

element has been quantised to one of eight values and the vector value shown is [3, 4].

Consider a single 1-tuple connection made onto the q th pixel of the x_1 bar. If $x_1 > q$ the 1-tuple value will be one, whereas if $x_1 < q$ then its value is zero. In pattern space the value of '1' or '0' on this connection changes as an imaginary threshold line at $x_1 = q$ is crossed. This idea can be extended to the n connections forming an n -tuple. The value on each connection tells on which side of the associated threshold line the current input pattern lies. The intersections of the lines delineate particular regions of pattern space which are associated with particular n -tuple values as shown in Fig. 9 for the case of a 4-tuple. If many n -tuples are connected onto the input retina, the combination of n -tuple values delineate smaller and smaller regions of pattern space, with each region becoming a regular square as the number of n -tuples becomes very high.

The idea of a pattern space dissected by threshold lines can be used to predict the form of the basis function produced by a SLLUP. Assume that the RAMs have initially zero contents: a single training pattern X_1 is applied to the system and the contents of the locations in the RAMs addressed by each n -tuple's value are iteratively modified until they produce the specified target output vector T_1 . If another pattern is now applied to the system, starting at X_1 and gradually moving away from X_1 , an increasing number of the threshold lines will be crossed.

Each time a threshold line is crossed, the address formed by an n -tuple of connections on the retina changes value, and one of the RAMs which was contributing to the output value T_1 switches to another location and produces zero output. Thus the output from the SLLUP slowly drops towards zero as the input pattern is progressively displaced from the training pattern and the shape of the fall in output value defined the form of the implicit basis function.

The shape of the basis function can be derived analytically if it is assumed that large numbers of n -tuples are connected to the input pattern retina. Let there be Q RAMs in each 'neuron' and D dimensions of length W in the pattern space. The output from each 'neuron' which is caused by an input pattern displaced by distance x from the position of the training pattern is given by $s(x)$:

$$s(x) = t - r(x)C_{av} \quad (5)$$

where t is the output from the neuron when the input is the training pattern, $r(x)$ is the number of RAMs whose addresses change when the input pattern is moved x units away in pattern space, and C_{av} is the average contribution of each RAM to the output value at the training point. However,

$$r(x) = \int_0^x p(x) dx \quad (6)$$

where $p(x)$ is the probability density of crossing a threshold line at position x in the pattern space. Assuming large numbers of randomly connected n -tuples, this density is just the number of active threshold lines $q(x)$ cutting any axis of the pattern space divided by the pattern space width over which the lines are distributed, that is

$$p(x) = \frac{q(x)}{W} \quad (7)$$

To solve for $p(x)$, $r(x)$ and $s(x)$ we investigate the variation of $q(x)$ as x increases by δx :

$$q(x + \delta x) = q(x) - \frac{n}{D} \frac{q(x)}{W} \delta x \quad (8)$$

The n/D term in the equation arises because crossing a single threshold line renders the other $n - 1$ lines within the n -tuple ineffective, i.e. subsequent crossing of any one of these other lines cannot affect the output from the RAM any more because it is already switched to produce zero output. The n lines which are effectively removed from play are spread over D dimensions and so the average number of lines along a particular axis which are deactivated by crossing just one threshold line is n/D . Rearranging eqn. 8 gives

$$\frac{\delta q(x)}{\delta x} = -\frac{n}{WD} q(x) \quad (9)$$

This first-order differential equation has a solution of

$$q(x) = \frac{Q^n}{D} e^{-nx/W D} \quad (10)$$

and can be used in conjunction with eqns. 5 and 7 to provide the required expression for the interpolation kernel.

$$s(x) = t(1 - e^{-n|x|/W D}) \quad (11)$$

The approximate validity of this expression has been confirmed by a computer simulation in which a SLLUP has been trained on a single input pattern and the output caused by other patterns covering a 2-D input space recorded. The results are shown in Fig. 10a for a system using 24 quantisation levels for each dimension with 32 RAMs addressed by 4-tuples of connections onto the input retina.

6.4 Simple binary coding

A pattern may also be presented to a SLLUP as image pixels whose values correspond to simple binary encoding of the pattern elements. Random n -tuples of connections are made onto the retina in the normal way. However, the effective basis function created when the input vector is binary coded is poorly defined and irregular for the following reasons.

The output from the SLLUP is determined by the number of RAMs in each neuron which are addressed

with the n -tuple value on which they were trained. The probability of an n -tuple changing value is proportional to the Hamming distance D_h between the current input pattern and the training input pattern. Consequently the output of the SLLUP is related to the Hamming distance between the training pattern and current input pattern. The relationship between Hamming distance and signal space distance for bar chart code is linear whereas it is nonlinear for binary code, as shown in Fig. 11.

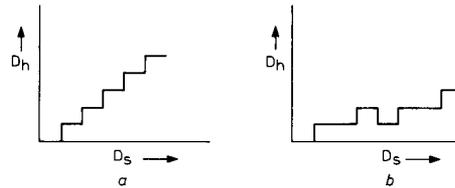


Fig. 11 Hamming distance against signal space distance
a For bar chart
b For simple binary code

In the latter case, the nonlinearity causes the SLLUP output to follow an irregular function as the input pattern is moved away from the training point in pattern space as shown by the computer simulation example of Fig. 10b. In this experiment, simple 5-bit binary coding with 32 RAMs addressed by 4-tuples of retina connec-

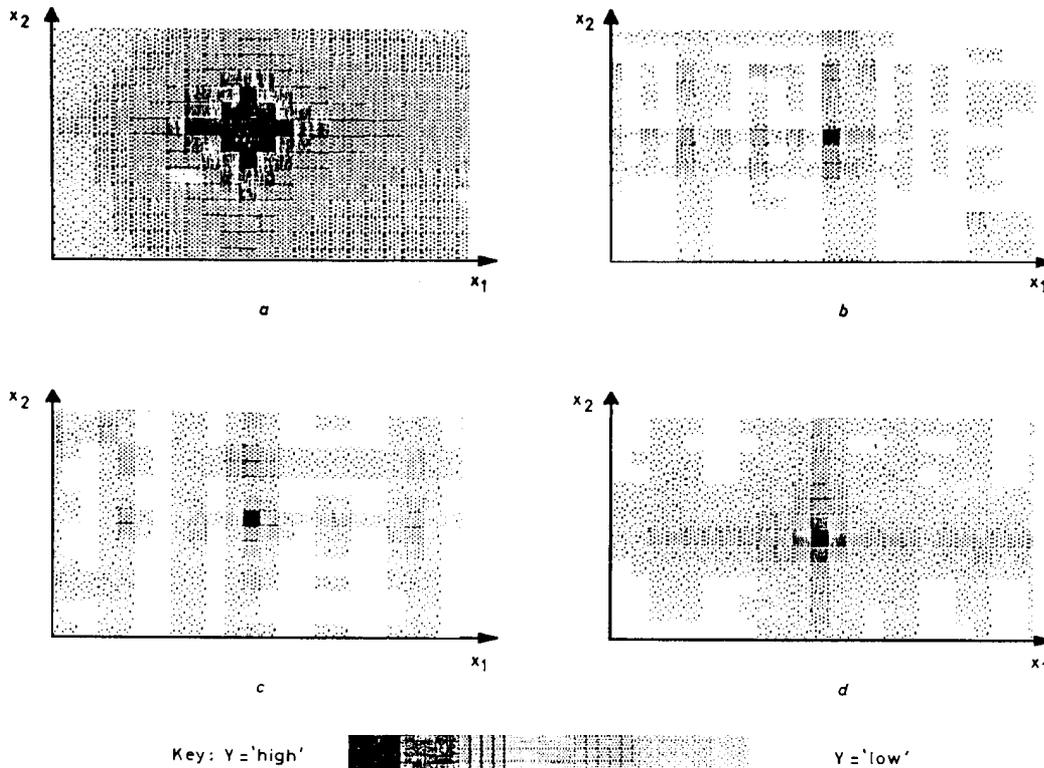


Fig. 10 Interpolation kernels of a function Y in a two-dimensional pattern space (x_1, x_2) using a variety of input codings

- a* Bar chart coding
- b* Binary coding
- c* Gray coding
- d* Redundant Gray coding

tions was used. The irregularity of the basis function is evident, but it should also be noted that it is much narrower than in the bar chart system. This is a possible advantage since it means that complex mapping functions can be synthesised if sufficient training data are available.

6.5 Gray coding

Some improvement in the regularity of the basis function can be obtained by using a Gray code instead of simple binary code. Although the Gray code Hamming distance against signal space distance function is not monotonic, it is better than the binary code because the Hamming distance only changes by one for every 1 unit change in signal. Again, this type of encoding has been tested by simulation and the kernel function obtained using a Gray code is shown in Fig. 10c.

6.6 Other codes

Further improvements in kernel regularity can be obtained by introducing redundancy into the code. There are a large number of ways of doing this, but a particularly effective technique is *redundant Gray coding* [11] in which each signal value is represented by a concatenation of several shifted versions of an R -bit Gray code for the value. Connecting n -tuples onto all the bits in the new code causes the bumps in the relationship between D_n and D_s to be averaged, resulting in a more regular kernel as shown by the simulation results in Fig. 10d.

7 Speech recognition and synthesis problems using the SLLUP

Earlier work has shown [12] that the SLLUP is able to perform simple nonlinear mappings such as the fuzzy EXOR problem. This is achieved using only small amounts of training and with little computation compared to the MLP. In this paper we examine the SLLUP performance on two speech mapping problems of very great complexity on which MLPs and some other neural nets have already been tested.

The first mapping problem is speaker independent recognition of utterances of the letters of the alphabet. A defined cepstral coefficient representation of many utterances of the letters of the alphabet from one large set of talkers must be classified by the SLLUP after it has been trained on examples from a separate set of talkers. The database used in this experiment was compiled by British Telecom Research Labs and is known as the Connex S1 data [13]. The SLLUP is trained on approximately 4000 utterances from a balanced mix of 52 talkers and then tested on approximately 4000 utterances from another 52 talkers. The utterance length is normalised by linear time warping and is presented to the SLLUP as a set of 15 frames of 8 mel cepstral coefficients.

The second complex problem to which the SLLUP has been applied is text-to-speech synthesis. In this case orthographic text has to be mapped to a sequence of phoneme codes which are then used to drive a hardware synthesiser. The experiment uses the same database as Netspeak [14] and is identical in all respects except that the MLP is replaced by a SLLUP. The SLLUP is presented with a character taken from English orthographic text and has to produce an appropriate phoneme code as output. Clearly the pronunciation of a particular character often depends on the word in which it is embedded and so three characters on either side of the target character are simultaneously presented to the SLLUP. Thus,

the complete input pattern consists of a context window of seven characters, each encoded using 11 bits. The output phoneme is represented using a 19-bit code to represent each of 55 phonemes.

It is interesting to consider the types of mapping which the SLLUP has to develop to deal with each of these two problems. In the speech recognition case, input patterns belonging to the same utterance class are likely to cluster together in their N -space and the SLLUP has to map the region of N -space enclosing the cluster to a single specified point in the output space. There may be several clusters belonging to one class but overall the mapping between input and output is smooth, without abrupt transitions. This proposition is supported by the fact that the moderately good speech recognition systems can be made using nearest neighbour classification of the input pattern. The task of the SLLUP in this case is to interpolate so that previously unseen input patterns which lie between training examples of the same class are mapped to the same output code.

The text-to-speech mapping is very different. The distances between the codes representing different characters does not have a simple relationship to the distances between the codes for the corresponding output phoneme codes. In other words, the patterns are really symbolic and just happen to be represented in a Euclidean space for processing by the neural net. Thus, the task of the SLLUP is to detect any logical structure in the data and, failing this, to act as a lookup table.

7.1 Experimental results on the speech recognition problem

Tables 1 to 3 summarise the results obtained using the SLLUP as a speech recogniser. Table 1 shows that a

Table 1: Speech recognition results using SLLUP with binary coded input

n -tuple order	Training set	Test set	RAMS per O/P
2	95%	65%	480
3	98%	62%	320
4	95%	52%	240

8-bit natural binary coding; retina size 960 bits

Table 2: Speech recognition results using SLLUP with bar chart coded input

Tuple order	Training set	Test set	RAMS per O/P
2	76%	71%	960
3	81%	75%	640
4	83%	77%	480
6	85%	78%	320

16-level thermometer code; retina size = 1920 bits

Table 3: Comparison between best results of SLLUP and MLP

Device	Training set	Test set
MLP, 75 hidden units (see Reference 18)	97.4%	88.3%
SLLUP, tuple order 8; 32-level bar chart code; 48 RAMS per O/P	97.3%	82.8%

SLLUP using natural binary coding in the retina is able to learn the training set very well, but performs poorly on the test set. Moreover, the performance tends to improve as the order of n -tuple decreases. Taken together, these two factors suggest that the SLLUP is unable to interpolate sufficiently between the training examples because

the effective width of the interpolation kernel is too small. Reduction of the n -tuple order causes the kernel to become wider, with a consequent improvement in performance on the test set. Increasing the n -tuple order makes the system behave more like a lookup table, giving better recognition of the training set but an inability to generalise.

The kernel width produced using natural binary code is very narrow and a possible solution to the poor test set performance is to use a different code in the retina, as demonstrated by the results of Table 2. These results were obtained by quantising each of the cepstral coefficients to 16 levels and representing them by a bar chart code. As expected, the performance improves on the test set and gets worse on the training set. This confirms our hypothesis that the natural binary code leads to an over-specific system. The results in Table 2 show an improvement in performance as the n -tuple order increases, indicating that in this system the kernel is actually too wide so that with low values of n overgeneralisation is taking place. This is supported by the fact that the system has been unable to accurately recognise the training set.

The recognition accuracy obtained using this system with $n = 6$ is comparable with results obtained using a 2-layer, 25 hidden-unit MLP on the same data which produced a test set accuracy of 81%. The best results obtained using the SLLUP on this data are compared to those produced by an MLP in Table 3.

7.2 Experimental results on the text-to-speech synthesis problem

In these experiments each of the seven characters in the input window are represented by 11-bit codes containing approximately equal numbers of '1's and '0's. This is important when using a SLLUP because an imbalance in the number of '1's and '0's will cause most n -tuple values to always consist of n '1's or n '0's, respectively, and this renders the n -tuple values insensitive to changes in the input vector X .

The results obtained using a SLLUP in the text-to-speech application are presented in Tables 4 and 5. Table

Table 4: Accuracy of SLLUP as text-to-speech mapper using unstructured input coding

Tuple order	Training set	Test set	RAMS per O/P
4	34.5%	33.4%	20
8	60.8%	55.2%	10

Training: 4 blocks of 10 000 characters; testing: 1 block of 10 000 characters

Codes for each character are approximately equidistant, 11-bits long and consist of five '1's and six '0's; retina size = 77 bits

Table 5: Accuracy of SLLUP as text-to-speech mapper using structured input coding

Tuple order	Training set	Test set	RAMS per O/P
4	52.2%	52.2%	20
8	72.7%	71.3%	10
10	78.4%	75.0%	8
10**	83.9%	79.9%	8

Training: 8 blocks of 10 000 characters; testing: 5 blocks of 10 000 characters

** Frequency weighted training and test data
Each code is 11 bits long and consists of five '1's and six '0's. The distance between each code reflects the letter group; Retina size = 77 bits

4 shows the performance of the SLLUP when the 11-bit codes are placed at approximately equidistant positions in 11-space. This coding is therefore completely unstruc-

ured. As expected, the performance is very poor because the input patterns are really symbolic and the interpolation between arbitrary codes effected by the SLLUP is inappropriate. Using a high tuple order of eight gives improved performance on the training set because the SLLUP starts to operate as a lookup table. However, the test set performance remains poor.

Table 5 shows the performance of the SLLUP working on a modified set of input codes which are chosen so that their mutual distances approximately reflect the perceptual distances between the phonemes which map most frequently to each letter. Using these structured codes, distances in the 11-space have some meaning and so interpolation becomes a more appropriate means of generating an output on unseen input data. Predictably the results in Table 5 are much better, with high accuracies obtained both on training and test data if sufficiently large n -tuples are used. A further improvement can be obtained if the frequency of commonly occurring words is reflected in the content of the training and test sets. This is because the very common words in English often have irregular pronunciation rules which are hard for the SLLUP to learn unless seen very frequently. McCulloch reports [14] that a 2-layer, 77 hidden-unit MLP can give an 86% letter-to-phoneme mapping accuracy which is slightly better than the SLLUP result. However, the SLLUP converges relatively quickly and shows a trend of improving performance as n -tuple order increases.

8 Conclusions

It has been argued that the purpose of any supervised learning network is to perform generalisation by synthesising a continuous nonlinear mapping function from a sparse set of training examples of the function. The continuous function can be generated by interpolation between the discrete examples of the function using a lowpass filter, and radial basis function networks and MLPs are examples of systems which utilise this principle.

Important implications of this argument are that the number of training examples must be sufficient high that the function is sampled at least at the Nyquist rate and that the generalisation which is produced by the filtering is only correct if the spectrum of the function underlying the training data is dominated by low frequencies.

The single-layer lookup perceptron is another example of a lowpass interpolating network which synthesises the required mapping function by effectively convolving the discrete training function samples with a kernel function which is analogous to the impulse response of a lowpass interpolation filter.

The SLLUP uses comparable amounts of memory to the MLP for all but the most trivial functions and in general will learn the required mapping function much faster than an MLP because it is a single-layer machine in which error gradients used for its adaption can be calculated directly from the output error. Moreover, because it is a single-layer machine, the error surface for the SLLUP is quadratic and it therefore always converges to a minimum error.

It has been shown that the SLLUP is able to operate as a speaker-independent recogniser with almost as high accuracy as an MLP, which suggests both that speech recognition can be effectively performed by interpolation and, perhaps more important, that the MLP also appears to be doing little more than interpolation. This is sup-

ported by the use of a SLLUP for text-to-speech synthesis, which again gave a performance only slightly inferior to an MLP.

Although multidimensional interpolation is a non-trivial task, our arguments suggest that many problems to which neural nets such as the MLP are currently applied with enthusiasm might be more efficiently solved by using explicit, classical interpolation techniques. More importantly, there are many logical problems which embody functions which have a bandpass spatial spectrum and which cannot be correctly generalised by lowpass interpolation, MLPs, RBFs or SLLUPs. New types of neural network are needed to deal with these classes of problem.

9 Acknowledgment

The director of British Telecom Research Labs supported this work and provided the Connex S1 and Netspeak databases.

10 References

- 1 RUMELHART, D.E., HINTON, G.E., and WILLIAMS, R.J.: 'Learning internal representations by error propagation'. ICS Report 8506, University of California, September 1985
- 2 LOWE, D.: 'Adaptive radial basis function nonlinearities and the problem of generalisation'. Proc. 1st Int. Conf. on ANNs, London 1989, pp. 171-175
- 3 BLEDSOE, W.W., and BROWNING, I.: 'Pattern recognition and reading by machine'. Proc. Eastern Joint Comp. Conf., Boston, Mass., 1959
- 4 ALEKSANDER, I. THOMAS, W.V., and BOWDEN, P.A.: 'WISARD — a radical step forward in image recognition', *Sensor Review*, July 1984, pp. 120-124
- 5 JOHNSTON, R.D.: 'Adaptive recognising device'. UK Patent No. 8427165, 1984; US Patent No. 4782459, 1988
- 6 ALBUS, J.S.: 'A new approach to manipulator control: the cerebular model of articulation controller (CMAC)', *Trans. ASME J. Dyn. Syst. Meas. & Control*, September 1975, pp. 220-227
- 7 LAPEDES, A., and FARBER, R.: 'Nonlinear signal processing using neural networks: prediction and system modelling'. Los Alamos National Lab. Report No. LA-UR-87-2662, Los Alamos FMLP
- 8 TATTERSALL, G.D., and JOHNSTON, R.D.: 'Speech recognisers based on N-tuple sampling'. Proc. Inst. of Acoustics, Spring Conf., 1984, pp. 405-413
- 9 SIXSMITH, M.J., and TATTERSALL, G.D.: 'Speech recognition using n-tuple techniques', *Br. Telecom Technol. J.*, April 1990, 8, (2), pp. 50-60
- 10 CAUCHY, A.: 'Méthode générale pour la resolution des systèmes d'équations simultanées'. C.r. Hebd. Seanc. Acad. Sci., Paris, 1847, 25, p. 536-538
- 11 TATTERSALL, G.D., JOHNSTON, R.D., and FOSTER, S.M.: 'Single layer look up perceptrons'. Patent Application No. 8906558.5, March 1989
- 12 TATTERSALL, G.D., JOHNSTON, R.D., and FOSTER, S.M.: 'Single layer look up perceptrons', in WHEDDON, C., and LINGGARD, R. (Eds.): 'Speech and language processing' (Chapman & Hall, 1990), pp. 307-333
- 13 LINGGARD, R., and WOODLAND, P.: 'Connex S1 Database' (British Telecom Research Labs., Ipswich IP5 7RE, UK)
- 14 MCCULLOCH, N., AINSWORTH, W.A., and LINGGARD, R.: 'Multi-layer perceptrons applied to speech technology', *Br. Telecom Technol. J.*, April 1988, 6, (2), pp. 131-139
- 15 ALEKSANDER, I., and STONHAM, T.J.: 'Guide to pattern recognition using random access memories'. *Comput. & Digital Tech.*, February 1979, 2, (1), pp. 29-40
- 16 ALEKSANDER, I., and MORTON, H.B.: 'An overview of weightless neural nets'. Invited Paper, Int. Joint Conf. on Neural Networks, Washington 1990
- 17 FLETCHER, R. (Ed.): 'Optimisation' (Academic Press, 1969), pp. 307-325
- 18 WOODLAND, P.C.: 'Isolated word speech recognition based on connectionist techniques', *Br. Telecom Technol. J.*, April 1990, 8, (2), pp. 61-66