



CONTRIBUTED ARTICLE

The Theoretical and Experimental Status of the n -tuple Classifier

RICHARD ROHWER¹ AND MICHAL MORCINIEC²

¹Prediction Company, and ²Hewlett-Packard Labs,

(Received 5 September 1995; accepted 21 January 1997)

Abstract—A number of theoretical approaches related to the n -tuple classification system are reviewed including Kanerva's SDM, the n -tuple regression network, the Hamming distance framework and likelihood estimation. The limitations of these methods are pointed out and resemblances that exist between them are underlined. Large-scale experiments carried out on StatLog project datasets confirm the n -tuple method as a viable competitor to more popular methods due to its speed, simplicity and accuracy on the majority of a wide variety of classification problems. A further investigation into the failure of the method on certain datasets shows its inner workings and reveals two main problems: difficulties with highly skewed class priors and more importantly, a mismatch between the scales involved in generalization, the amount of training data available, and the volume of the region in which data is likely to exist. This highlights areas where improvements in the method are needed and further theoretical progress would be helpful. © 1998 Elsevier Science Ltd. All rights reserved

Keywords—Benchmarking, CMAC code, Hamming distance, Kanerva model, Regression network, n -Tuple classifier, Review, StatLog.

1. INTRODUCTION

The n -tuple classifier, invented by Bledsoe & Browning (1959), is one of the oldest practical pattern recognition methods based on distributed computation and amenable to description in terms of neural network metaphors. Although eclipsed in popularity by methods such as multilayer perceptrons and radial basis function networks, the n -tuple method continues to offer properties which make it vastly superior for certain common purposes. First among these properties is its speed of operation. The training algorithm is a one-shot memorization task, computationally trivial compared to solving linear systems or minimizing non-linear functions. Another advantage is the sheer simplicity of learning by memorization. If this can form the basis of a sound

pattern recognition principle, then it is arguable that biological systems could make use of it.

It is prudent to suspect that relatively poor performance will accompany the speed and simplicity of the n -tuple algorithm. There are many reports of satisfactory results with the method (Bledsoe & Bisson, 1962; Ullmann, 1969; Ullmann & Kidd, 1969; Aleksander & Stonham, 1979a; Rohwer & Lamb, 1993; Tarling & Rohwer, 1993) but few studies involving comparisons with other methods (Rohwer & Cressy, 1989). Furthermore, most studies use just one or two small data sets. Therefore a large experiment was carried out, in which the n -tuple method was tested on 11 large real-world data sets which had been previously used by the European Community ESPRIT StatLog project (Michie et al., 1994) to test 23 other classification algorithms including the most popular neural network methods. The results, presented in Section 5.4, show no systematic performance gap between the n -tuple method and the others, on seven of the 11 data sets tested. It is easy to recognize the other four, without referring to any competing methods, because the n -tuple method performed no better than random guessing. The experiments therefore suggest that in most cases the n -tuple method gives competitive performance, and the cases when it does not are clearly recognizable.

Acknowledgements: This work was carried out at Aston University in the Neural Computing Research Group. It was supported in part by the British Council and Aston University. The authors are grateful to Louis Wehenkel of the Université de Liege for useful correspondence and permission to report results on the BelgianI and BelgianII data sets, Trevor Booth of the Australian CSIRO Division of Forestry for permission to report results on the Tsetse data set, and Reza Nakhaeizadeh of Daimler-Benz, Ulm, Germany for permission to report on the Technical, Cut20 and Cut50 data sets.

Requests for reprints should be sent to Dr R. Rohwer, Prediction Company, 236 Montezuma Avenue, Santa Fe, NM, USA.

When a fast and simple method proves to be a competitive performer in a large set of experiments, one would like to know why. Unfortunately, the n -tuple method has not yet yielded to theoretical analysis as well as the optimization-based approaches which can be embedded in Bayesian statistical theory (MacKay, 1992), but there has been some progress. The main theoretical results are reviewed, discussed and used as a vehicle for explaining some relationships between the standard n -tuple method, n -tuple methods incorporating linear regression, radial basis function networks and the Kanerva associative memory model. The theoretical tools are then used to develop a semi-quantitative account of why the method failed on four of the 11 datasets. The main theoretical stumbling blocks are also indicated.

2. THE N -TUPLE RECOGNITION METHOD

The n -tuple recognition method is also known as a type of ‘‘RAMnet’’¹ or ‘‘weightless neural network’’. It forms the basis of a commercial product (Aleksander et al., 1984). It is a method for classifying binary patterns, which can be regarded as bit strings of some fixed length L . This is not an important restriction, because there is an efficient preprocessing method, tailored to the RAMnet’s generalization properties, for converting scalar attributes into bit strings. This method is reviewed in Section 4. It is also possible to generalize a special case of the n -tuple method from a classifier to a function interpolator.

2.1. Definition of the Standard n -tuple Method

A concise definition of the n -tuple method is given here. More descriptive definitions can be found elsewhere (Aleksander & Stonham, 1979b; Alexander & Morton, 1995). Several (let us say N) sets of n distinct² bit locations are selected randomly. These are the n -tuples. Collectively, they are called the ‘‘input mapping.’’ The restriction of a pattern to an n -tuple can be regarded as an n -bit number which, together with the identity of the n -tuple, constitutes a ‘‘feature’’ of the pattern. The standard n -tuple recognizer operates simply as follows:

A pattern is classified as belonging to the class for which it has the most features in common (1) with at least one training pattern of that class.

This is the $\theta = 1$ case of a more general rule whereby the class assigned to unclassified pattern u is

$$\operatorname{argmax}_c \left(\sum_{i=1}^N \Theta_\theta \left(\sum_{\mathbf{v} \in \mathcal{D}_c} \delta_{\alpha_i(\mathbf{u}), \alpha_i(\mathbf{v})} \right) \right) \quad (2)$$

¹ RAMnets also include stochastic generalizations, pRAMs, to which the n -tuple recognition algorithm is not applied. These are not considered here.

² Relaxing the requirement that an n -tuple has n different bit locations amounts to introducing a mixture of differently sized n -tuples. Note the restriction does not disallow a single pattern component from being shared by more than one n -tuple.

where \mathbf{v} ranges over \mathcal{D}_c , the set of training patterns in class c , $\Theta_\theta(x) = x$ for $0 \leq x \leq \theta$, $\Theta_\theta(x) = \theta$ for $x > \theta$, $\delta_{i,j}$ is the Kronecker delta³ ($\delta_{i,j} = 1$ if $i = j$ and 0 otherwise) and $\alpha_i(\mathbf{u})$ is the i th feature of pattern \mathbf{u} :

$$\alpha_i(\mathbf{u}) = \sum_{j=0}^{n-1} u_{\eta_i(j)} 2^j. \quad (3)$$

Here u_k is the k th bit of \mathbf{u} and $\eta_i(j)$ is the j th bit location of the i th n -tuple.

Small values of θ , greater than 1, are sometimes found to be most effective (see Section 3.3).

With C classes to distinguish, the system can be implemented as a network of NC nodes, each of which is a 1-bit random access memory (RAM); hence the term *RAMnet*. (Equivalently, it is a network of N RAMs, each containing a C -dimensional bit vector.) The memory content $m_{ci\alpha}$ at address α of the i th node allocated to class c is set to

$$m_{ci\alpha} = \Theta_\theta \left(\sum_{\mathbf{v} \in \mathcal{D}_c} \delta_{\alpha, \alpha_i(\mathbf{v})} \right). \quad (4)$$

Note that symbol α denotes any address value in range $[0, 2^n - 1]$ and $\alpha_i(\mathbf{u})$ refers to one particular address value generated by the input pattern \mathbf{u} .

In the usual $\theta = 1$ case, the 1-bit content of $m_{ci\alpha}$ is set if any pattern of \mathcal{D}_c has feature i and unset otherwise. Recognition is accomplished by summing the contents of the nodes of each class at the addresses given by the features of the unclassified pattern. That is, pattern \mathbf{u} is assigned to class

$$\operatorname{argmax}_c \left(\sum_{i=1}^N m_{ci\alpha_i(\mathbf{u})} \right). \quad (5)$$

To generalize the n -tuple classifier to a function interpolator (Allinson & Kołcz, 1994b), a vector $\mathbf{m}_{i\alpha}$ of the appropriate numeric type for the range of the function is stored at each address α in each node i . Training is accomplished by assigning these vectors to the averages:

$$\mathbf{m}_{i\alpha} = \sum_{\mathbf{v} \in \mathcal{D}} \mathbf{Y}(\mathbf{v}) \delta_{\alpha, \alpha_i(\mathbf{v})} / \sum_{\mathbf{v} \in \mathcal{D}} \delta_{\alpha, \alpha_i(\mathbf{v})} \quad (6)$$

where $\mathbf{Y}(\mathbf{v})$ is the desired output for input bit string \mathbf{v} . The network output, or response to an arbitrary input pattern \mathbf{u} is

$$\frac{1}{N} \sum_{i=1}^N m_{i\alpha_i(\mathbf{u})}. \quad (7)$$

2.2. Practical Experience

Simple theoretical considerations and practical experience provide fairly strong guidance for setting the architectural

³ The comma is unconventional, but is used here optionally for extra clarity.

parameters n , N and θ . To begin with, the fact that the network response to an arbitrary pattern is essentially an average over the n -tuples [eqn (5)] means that the results should become increasingly consistent with increasing N . Because the n -tuple method can process thousands of patterns within seconds, there is little need for any data analysis method other than explicit measurement of the variation of performance as the input mapping is re-randomized a few times for a given N . N is increased if the variation is unacceptably high. Practical experience is that values of 100 to 1000 usually turn out to be adequate. It is commonly observed that the network's output for the winning class exceeds that of the second runner up by an uncomfortably small margin, such as three n -tuples out of 1000, but that the correct class nevertheless wins consistently. Perhaps most n -tuples give a constant response to most patterns, so effectively only a fraction of those in the network are contributing to the decisions. This suggests that many n -tuples could be trimmed from the network, but such variations on the method are vulnerable to overtraining and complicate the theory. The phenomenon deserves further study.

Practical experience tends to favour small values of the threshold θ , particularly $\theta = 1$. A possible rationale for this is given in Section 3.3. Many considerations apply to the choice of n -tuple size n . Experimentally it usually is the case that bigger is better, up to an impractically large size (Rohwer & Lamb, 1993), which requires an unreasonable amount of training data, but $n = 8$ is usually enough, and $n = 3$ is sometimes adequate. This can be explained qualitatively by observing that information about the correlations among up to n bits are available to the classifier. It never hurts to take account of higher-order correlations, but it is plausible that eighth-order correlations contain all that is needed for most binary data sets. Another intuition is that the training process should write to neither too small nor too great a proportion of the 2^n addresses at each node. If n is too large, the sub-patterns occurring in the training data will be unlikely to recur in the test data, whereas if n is too small, the memory can *saturate*, in which case $m_{\alpha_i} = \theta$ for most memory locations, so most discriminative power is lost (Ullmann, 1969; Tarling & Rohwer, 1993). These issues are further complicated if the class priors are highly skewed, so that one class has far more training data than another. Although a precise theory is not available, there are strong enough theoretical tools to gain considerable insight, as is demonstrated in the discussion of the experimental results.

3. THEORETICAL STATUS OF THE N -TUPLE METHOD

The n -tuple classifier is a memory-based method. Such methods differ from optimization-based methods, such as back propagation of error through multi-layer perceptrons,

in two important ways. Firstly, ‘‘hidden’’ representations (or ‘‘features’’) are selected randomly, and secondly, training is a simple one-shot memorization task involving these features. These differences give memory-based methods an awesome advantage in training speed. Radial basis functions obtain part of this speed advantage by selecting features randomly (Broomhead & Lowe, 1988), and multi-layer perceptrons can often be trained faster with little or no loss of performance by using fixed random weights into the hidden layers (Gallant & Smith, 1987; Sutton & Whitehead, 1993). However, this does not give the speed and simplicity that training by mere memorization provides.

In spite of many useful advances (Bledsoe & Bisson, 1962; Ullmann & Kidd, 1969; Aleksander & Stonham, 1979a; Flanagan et al., 1992; Austin, 1994), there is no theory of n -tuple networks of the standard of the sophisticated statistical techniques available with optimization-based methods (MacKay, 1992). It is not particularly difficult to design new training algorithms for the n -tuple architecture in order to make these statistical methods applicable (Tattersall et al., 1991; Luttrell, 1992; Rohwer, 1995), but this approach sidesteps the interesting questions instead of answering them. These modified methods reduce the speed and simplicity advantages as well.

The main tools for understanding the standard n -tuple network are reviewed here, and used to place the n -tuple net into context with similar methods. These tools are used again to explain the experimental results.

3.1. Approaches Based on Hamming Distance

The most productive theoretical concept for understanding the n -tuple method has been ‘‘tuple distance’’ and its non-linear statistical relationship with Hamming distance.

3.1.1. *Tuple Distance and Hamming Distance.* The main theoretical results on the n -tuple method provide a relationship between a ‘‘tuple distance’’ relevant to the network's generalization properties, and the Hamming distance between training and test patterns. The tuple distance $\rho(\mathbf{u}, \mathbf{v})$ between patterns \mathbf{u} and \mathbf{v} is the number of tuples (of a given input mapping) on which the patterns disagree:

$$\rho(\mathbf{u}, \mathbf{v}) = N - \sum_{i=1}^N \delta_{\alpha_i(\mathbf{u}), \alpha_i(\mathbf{v})}. \quad (8)$$

The number on which they agree, $N - \rho(\mathbf{u}, \mathbf{v})$, will be called the ‘‘tuple score.’’ An elementary argument based on the random selection of the n -tuple inputs from the L bits available shows that patterns \mathbf{v} which lie a fixed Hamming distance $H(\mathbf{u}, \mathbf{v})$ from any one pattern \mathbf{u} are distributed binomially in tuple distance:

$$P(\rho|H) = \frac{N!}{\rho!(N-\rho)!} \left(1 - \left(\frac{H}{L}\right)^n\right)^\rho \left(1 - \frac{H}{L}\right)^{n(N-\rho)} \quad (9)$$

More complicated expressions are available for more constrained n -tuple sampling procedures (Tattersall & Johnson, 1984). Distribution eqn (9) gives an expectation value for ρ of

$$\rho(H) = \langle \rho | H \rangle = N \left(1 - \left(1 - \frac{H}{L} \right)^n \right) \quad (10)$$

and indicates that ρ typically strays from this value by the standard deviation

$$\delta \rho(H) = \left[N \left(1 - \frac{H}{L} \right)^n \left(1 - \left(1 - \frac{H}{L} \right)^n \right) \right]^{\frac{1}{2}} \quad (11)$$

This is illustrated in Figure 1. If the patterns are nearby ($H \ll L$), then a convenient approximation is

$$\rho(H) \approx N \left(1 - e^{-n \frac{H}{L}} \right) \quad (12)$$

The n -tuple sampling variations then make $\rho(H)$ uncertain by about $\sqrt{\rho(H)}$.

It is clear from eqn (12) that proximity in Hamming distance plays a role in the generalization behaviour of n -tuple networks. Consider a network trained on just one example \mathbf{v} of class c , and tested on a pattern \mathbf{u} Hamming distance H from \mathbf{v} . Classifications are based on the network response $\sum_{i=1}^N m_{ci\alpha_i(\mathbf{u})}$ to pattern \mathbf{u} , which will be about $N e^{-n(H(\mathbf{u}, \mathbf{v}))/L}$. Hence one could say that the network generalizes from training pattern \mathbf{v} to all patterns within a Hamming distance of about L/n of \mathbf{v} . Figure 1 shows a ‘‘clean’’ case in which the training patterns within L/n of a test pattern are predominately of the

correct class. In our experience it is far more usual for many incorrect patterns to populate this region as well, but the classifier seems to work anyway. It would appear that the training patterns nearest a test pattern tend to match on tuples (‘‘overlap’’) no more often than patterns of any other class.

3.1.2. Training Data Overlap on Tuples. A network trained on a set of patterns $\{\mathbf{v}_1, \dots, \mathbf{v}_T\}$ could respond to pattern \mathbf{u} by any amount between $N \min(1, \max_a e^{-nH(\mathbf{u}, \mathbf{v}_a)/L})$ and $N \min(1, \sum_a e^{-nH(\mathbf{u}, \mathbf{v}_a)/L})$, depending on the correlations between the training patterns, as manifest in ‘‘overlap effects’’. Unfortunately, this circumstance limits the usefulness of tuple distance for explaining the standard n -tuple method. Because of a combinatorial explosion, there is no feasible method of measuring tuple correlations. Similar problems crippling an attempt of full formal analysis of the method (Stonham, 1977) for datasets of arbitrary size have been reported. However, some insight into the mechanism of the RAMnet can be gained by analysis of the experimental data.

Overlap effects are displayed in Figure 2. Figures 2(a) and (b) show the network output for a test pattern as training patterns are added in Hamming distance order from the test pattern. For Figure 2(a), the threshold is $\theta = 1$, and for Figure 2(b) it is $\theta = \infty$, effectively ignoring overlap in that the tuple scores are added for all the training patterns. Figure 2(b) shows that distant patterns of the incorrect class match the test pattern on many n -tuples,

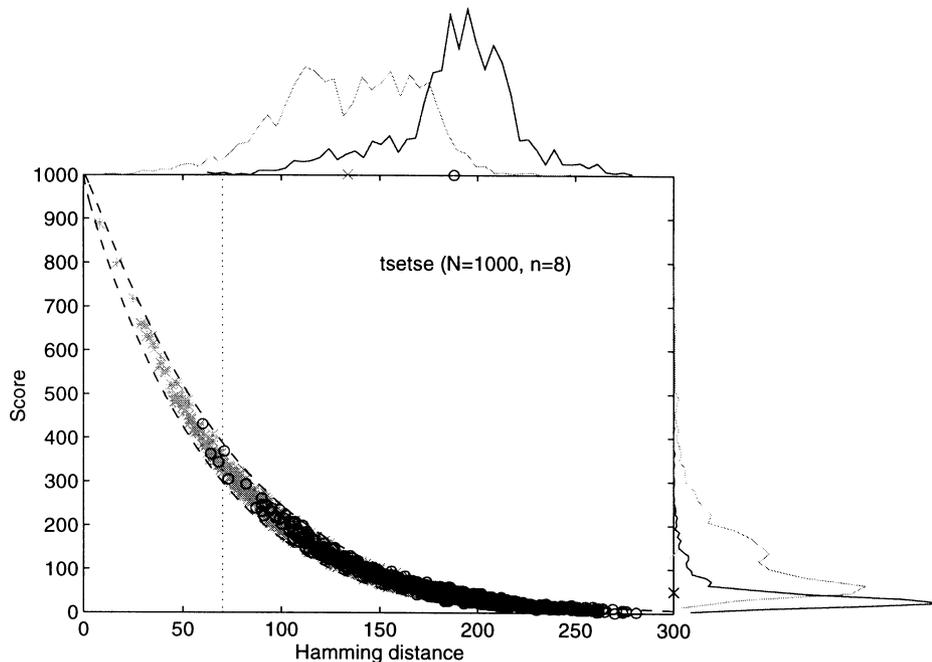


FIGURE 1. Tuple score vs Hamming distance for a fixed test pattern \mathbf{u} and training patterns $\{\mathbf{v}_1, \dots, \mathbf{v}_7\}$. The score of \mathbf{u} was computed for the system trained on just one pattern \mathbf{v} . Distances between patterns of the same class are marked $*$, and \circ is used for different classes. Eqns (10) and (11) are plotted as three s.d. error curves. The total number of patterns of the correct class (dotted line) and incorrect classes (solid line) are plotted in the margins as functions of Hamming distance (top) and tuple score (right). The means of these distributions are indicated by $*$ and \circ marks. The ‘‘generalization distance’’ L/n is indicated by a vertical dotted line.

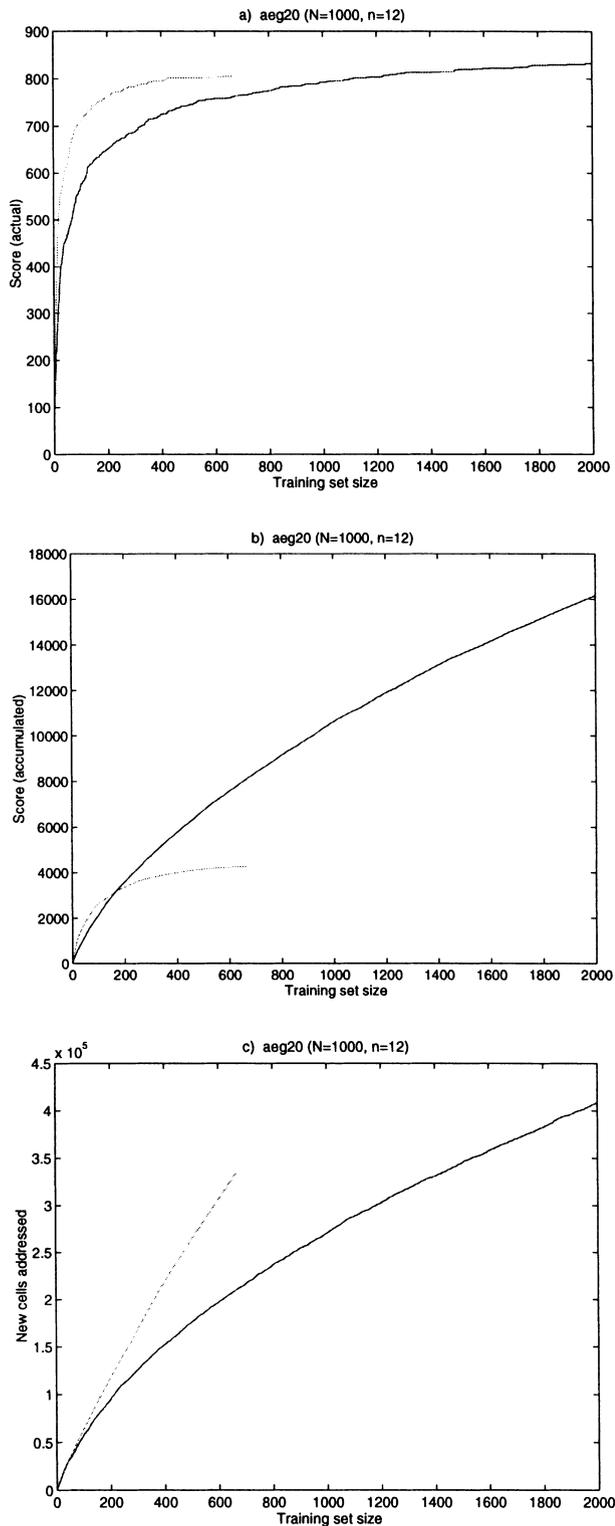


FIGURE 2. (a) Actual network response to a test pattern u ; (b) accumulated network response (as though patterns never overlapped on any tuples); and (c) the number of new RAM cells addressed as a function of training patterns $\{v_1, \dots, v_T\}$. Light lines are used to indicate the discriminator associated with the class that generated u , dark lines denote the other class. The training data is sorted by the Hamming distance to the test pattern.

but Figure 2(a) shows that most of these subpatterns had already turned up in closer training patterns. Figure 2(c) shows the number of new RAM locations accessed as training patterns are accumulated, regardless of whether these are accessed by the test pattern. One class is disadvantaged by a smaller prior probability, and correspondingly fewer training samples, but it has the advantage that it populates new RAM locations more rapidly. Whereas the more probable class is showing signs of levelling off in this respect, the less probable class is not, so it appears likely that if more data were available (in the same proportions), then the test pattern would be more likely to be classified correctly with the less probable class. This scenario was frequently observed in datasets with skewed priors, and motivates the hypothesis that errors would have been reduced if more data were available.

3.2. Related Methods

Tuple distance can be used to compare the n -tuple network to Kanerva's sparse distributed memory, the n -tuple regression network, single-layer lookup perceptron, and the familiar radial basis functions network. The n -tuple network can also be interpreted as a crude likelihood estimator.

3.2.1. *The Kanerva Model.* The n -tuple network is related to Kanerva's sparse distributed memory model (Kanerva, 1988; Songcan & Jun, 1992), which has also been developed theoretically using an "overlap" measure similar to eqn (12). Although intended mainly as an associative memory, it is easily generalized for classification problems or function interpolation problems. The interpolation version is presented here.

Instead of n -tuples, a set of N bit strings are randomly selected from a uniform distribution. These are used as centres ξ_i of hard-sphere radial basis functions ϕ_i of binary vectors \mathbf{v} :

$$\phi_i(\mathbf{v}; r) = \begin{cases} 1 & H(\xi_i, \mathbf{v}) \leq r \\ 0 & H(\xi_i, \mathbf{v}) > r \end{cases} \quad (13)$$

with a somewhat carefully chosen radius r . Memory space for a vector in the range of the function to be approximated is associated with each centre. The memory at centre i is set to

$$\mathbf{m}_i^{(K)} = \frac{\sum_{\mathbf{v} \in \mathcal{D}} \mathbf{Y}(\mathbf{v}) \phi_i(\mathbf{v}; r)}{\sum_{\mathbf{v} \in \mathcal{D}} \phi_i(\mathbf{v}; r)} \quad (14)$$

during training. Here $\mathbf{Y}(\mathbf{v})$ is the desired output for input pattern \mathbf{v} . The output $\mathbf{Y}(\mathbf{v})$ and memory $\mathbf{m}_i^{(K)}$ can belong to any space in which a weighted average can be defined.

For classification problems, $\mathbf{Y}(\mathbf{v})$ is an indicator function, and for an associative memory, $\mathbf{Y}(\mathbf{v}) = \mathbf{v}$.

The network response to test pattern \mathbf{u} is

$$\mathbf{y}^{(K)}(\mathbf{u}) = \frac{\sum_i \mathbf{m}_i^{(K)} \phi_i(\mathbf{u}; r)}{\sum_i \phi_i(\mathbf{u}; r)} \quad (15)$$

which is

$$\mathbf{y}^{(K)}(\mathbf{u}) = \frac{\sum_{\mathbf{v} \in \mathcal{D}} \mathbf{Y}(\mathbf{v}) \sum_i \phi_i(\mathbf{v}; r) \phi_i(\mathbf{u}; r)}{\sum_{\mathbf{v} \in \mathcal{D}} \sum_i \phi_i(\mathbf{v}; r) \phi_i(\mathbf{u}; r)} \quad (16)$$

in terms of the training data. (A further thresholding operation is required if the output is to be a bit string.) Each training pattern \mathbf{v} contributes its desired output $Y_c(\mathbf{v})$ to an average weighted by

$$N - \rho_r^{(K)}(\mathbf{v}, \mathbf{u}) \stackrel{\text{def}}{=} \sum_i \phi_i(\mathbf{v}; r) \phi_i(\mathbf{u}; r) \quad (17)$$

the number of centres within Hamming distance r of both the training pattern and the test pattern. Evidently, $N - \rho_r^{(K)}(\mathbf{u}, \mathbf{v})$ plays a role similar to the tuple score, with centres within distance r of both patterns being counted instead of n -tuples. Due to the random placement of centres, the expectation of $N - \rho_r^{(K)}(\mathbf{u}, \mathbf{v})$ is also a function $\rho_r^{(K)}(H(\mathbf{u}, \mathbf{v}))$ of the Hamming distance $H(\mathbf{v}, \mathbf{u})$, although it is more complicated than eqn (10) or eqn (12). The exact form is a sum of products of binomial coefficients which can be approximated by

$$\rho_r^{(K)}(H) \approx N \left(1 - \int_{H/L}^1 \frac{dx}{2\pi \sqrt{x(1-x)}} e^{-\frac{(r-L/2)^2}{L/2} \frac{1}{(1-x)}} \right) \quad (18)$$

for $0 \ll H \ll L$.

To carry the comparison further, observe that the interpolative n -tuple network [eqns (6) and (7)] can be regarded as a special case of the Kanerva network [eqns (14) and (15)] if the Hamming distance in eqns (14) and (15) is replaced with a Hamming distance restricted to a tuple. Specifically, a Kanerva centre $\xi_{i\alpha}$ can be associated with each memory location α at each n -tuple i by defining all bits of $\xi_{i\alpha}$ arbitrarily except for those involved in the i th input mapping. These bits must form subpattern α : $\sum_{j=0}^{n-1} \xi_{\eta(j), \alpha} 2^j = \alpha$. The Hamming distance is replaced by $\sum_{j=0}^{n-1} (1 - \delta_{\xi_{\eta(j)}, \mathbf{u}_{\eta(j)}})$ and $r = 0$.

3.2.2. The n -tuple Regression Network. The fact that the n -tuple network response to a test pattern is dominated by the (Hamming) nearest training patterns suggests that it might be formally related to a nearest neighbour method or a local basis function method. A modified method makes the most of this resemblance. Called the single-layer lookup perceptron (Tattersall et al., 1991), it treats

the interpolative n -tuple network response of eqn (7) as a weighted sum of basis functions:

$$y_c = \sum_{i=1}^N m_{ci\alpha_i(\mathbf{u})} = \sum_{i=1}^N \sum_{\alpha=0}^{2^n-1} m_{ci\alpha} \delta_{\alpha, \alpha_i(\mathbf{u})}. \quad (19)$$

From the latter form, the network can be regarded as a linear transformation applied to the outputs of the rather unusual basis functions $\delta_{\alpha, \alpha_i(\mathbf{u})}$, one for each combination of i and α .

Training is easily accomplished by a least mean squares (LMS) method.

Eqn (19) can be related to a basis function expansion in the the exponentials appearing in eqn (12):

$$y_c = \sum_a w_{ca} e^{-n \frac{H(\mathbf{u}, \mathbf{v}_a)}{L}} \quad (20)$$

provided that patterns \mathbf{v}_a and weights w_{ca} can be found such that

$$m_{ci\alpha} = \sum_a w_{ca} \delta_{\alpha, \alpha_i(\mathbf{v}_a)}. \quad (21)$$

[Plugging eqn (21) into eqn (19) and using eqn (8) gives eqn (20) within eqn (12).]

One way to arrange this is to choose all the patterns \mathbf{v}_a so they are separated from each other by tuple distance N ; i.e. none of the patterns \mathbf{v}_a match each other in any n -tuple. (This is possible if and only if the number of these patterns is no more than 2^n .) In this situation there is at most one a for any given n -tuple i and address α_i such that $\alpha_i(\mathbf{v}_a) = \alpha$, which may be called $a(i, \alpha)$ when it exists. Then the choice

$$m_{c\alpha_i} = \begin{cases} w_{ca(i, \alpha)} & a(i, \alpha) \text{ defined} \\ 0 & a(i, \alpha) \text{ not defined} \end{cases} \quad (22)$$

satisfies eqn (21). Such a network acts like the radial basis function network [eqn (20)] with a (Hamming⁴) spherically symmetric local basis function of radius roughly L/n centred on each pattern \mathbf{v}_a , even though the patterns \mathbf{v}_a do not appear in the implementation of eqn (19), and to the benefit of computation speed, no distance calculations $H(\mathbf{u}, \mathbf{v}_a)$ are ever performed.

If the training patterns meet the separation condition, they can be identified with the basis function centres \mathbf{v}_a . To see this, observe that the assumed one-to-one correspondence between pattern a and address $\alpha_i(\mathbf{v}_a)$ in the i th n -tuple memory implies that the set of memory locations corresponding to one pattern is disjoint from the set addressed by any other. Therefore the LMS optimization problem involving eqn (19) will be symmetric with respect to the N parameters $m_{ci\alpha_i(\mathbf{v}_a)}$ for each a and c . Hence these parameters will all be equal in the solution, and one can meaningfully define $m_{ca} = m_{ci\alpha_i(\mathbf{v}_a)}$. The

⁴ The method can be modified so that the effective basis functions have more nearly Euclidean-spherical receptive fields (Kotcz & Allinson, 1994).

network response to a test pattern of eqn (19) can then be rewritten by replacing the sum over addresses with a sum over training patterns:

$$y_c = \sum_{i=1}^N \sum_a m_{ci\alpha_i(\mathbf{v}_a)} \delta_{\alpha_i(\mathbf{v}_a), \alpha_i(u)} \quad (23)$$

Then it is clear from eqns (8) and (20) that one can identify $m_{ca} = w_{ca}$.

The interpretation of a single-layer lookup perceptron as an effective radial basis function network with a function centre on each training pattern requires the patterns \mathbf{v}_a to be tuple-separated by N . This condition would be valid at least to a good approximation if their Hamming separation were large compared to L/n . The analysis of our experiments typified by Figures 1 and 7 does not suggest that this is likely to happen, and even in principle it may not be easy to arrange this and ensure good coverage of the pattern space by the local effective basis functions. For this, the data needs to just happen to be arranged so that each training pattern of a class is about L/n bits away from its nearest neighbour. To do as well as possible on both conditions, all other neighbours should be much further away, but the triangle inequality requires the next nearest neighbour to be within $2L/n$ bits. If the data just happened to be suitably arranged for this interpretation, then a suitable choice of n -tuple size n could be stated in terms of the nearest neighbour distance d_{NN} as $n = L/d_{NN}$.

Even if the training data is not distributed so that a basis function interpretation is possible with the training points serving as function centres, it may still be possible to obtain a basis function interpretation by finding a

solution, at least approximately, to eqn (21). In any case, there is no reason to suppose that the method will fail when the interpretations do.

3.3. The n -tuple Classifier as a Crude Likelihood Estimator

An alternative to using tuple distance to formulate a theory of n -tuple networks is to view them as probability estimators. With $\theta = \infty$, eqn (4) for the memory content $m_{ci\alpha}$, can be interpreted as an estimate of the probability $P_i(\alpha|c)$ (up to a normalization factor) that a data point from a given class c will have subpattern α in n -tuple i . Assuming these distributions for different n -tuples to be independent of each other, the joint distribution over all the sub-patterns taken together is

$$P(\alpha|c) = \prod_i P_i(\alpha|c). \quad (24)$$

This can be used for maximum likelihood classification, or converted to posterior class probabilities using Bayes' rule with class priors $P(c)$, if they are available. The independence assumption lacks plausibility, because it would be remarkable for the correlations required to make the classes distinguishable not to be reflected in correlations between the n -tuples, but nevertheless, good results have been reported with this formulation, which has been reinvented from time to time (Bledsoe & Bisson, 1962; Sixsmith et al., 1990; Badr, 1993). Aside from its implausibility, its main practical problem is that factors of zero appear if a naive estimate of $P_i(\alpha|c)$ is used for subpatterns which never appear in

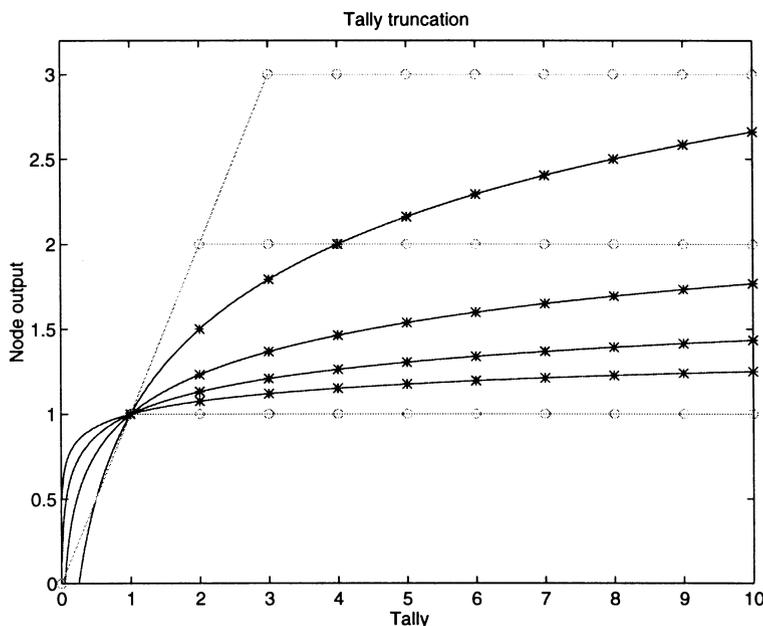


FIGURE 3. $1 + (\log T)/(-\log \epsilon)$ truncated to 0 for $T < \epsilon$ (dark lines) for $\epsilon = 0.25, 0.05, 0.005$ and 0.0001 , and $\Theta_\theta(T)$ (light lines) for $\theta = 1, 2$ and 3 , as functions of T . This shows that eqn (25) can be a reasonable approximation at integral values of T in some circumstances, particularly for $\epsilon \rightarrow 0$ with $\theta = 1$.

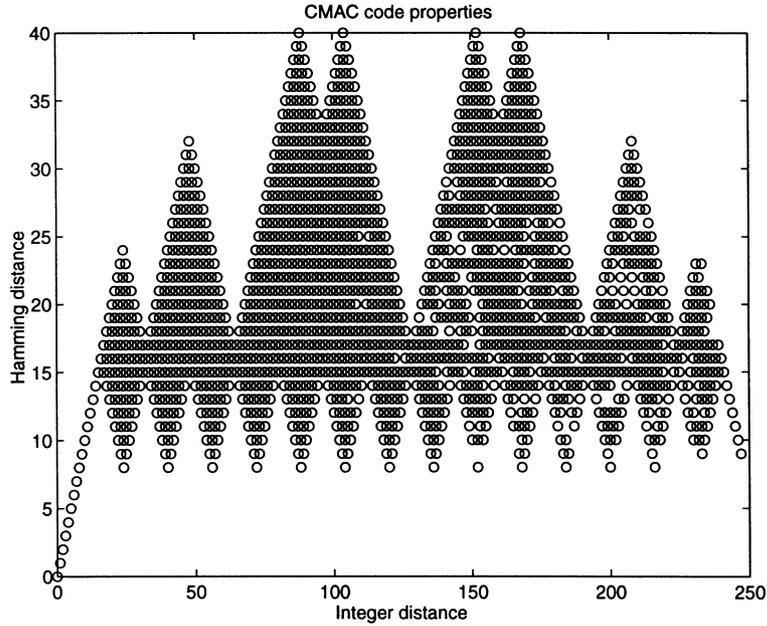


FIGURE 4. Hamming distance between two CMAC/Gray-transformed integers vs their arithmetic difference, for 3×10^4 randomly chosen pairs of integers. $K = 8$ and $a = 5$. The relationship is linear for Hamming distances up to K , and the transformed distance is bounded below by K for greater Hamming distances.

the training data. In practice these are replaced with a small ad hoc positive constant, say ϵ , leaving scope for more principled approaches to estimating these probabilities.

The n -tuple method with finite threshold can be seen as a scaled and translated approximation to the logarithm of eqn (24). Specifically, if test pattern \mathbf{u} has sub-patterns $\alpha(\mathbf{u})$, and the T_c training patterns from class c supply the tallies $T_{c\alpha} = \sum_{\mathbf{v} \in \mathcal{D}_c} \delta_{\alpha, \alpha(\mathbf{v})}$ used to estimate $P_i(\alpha|c)$ by $T_{c\alpha}/T_c$, then for suitable choices of ϵ and θ , the network

responses in eqn (5) will approximately satisfy

$$\begin{aligned}
 N &= \frac{\log(T_c^N P(\alpha(\mathbf{u})|c))}{-\log \epsilon} \\
 &= \sum_{i=1}^N \left(1 + \frac{\log T_c + \log P_i(\alpha(\mathbf{u})|c)}{-\log \epsilon} \right) \\
 &= \sum_{i=1}^N \left(1 + \frac{\log T_{c\alpha_i}}{-\log \epsilon} \right) \approx \sum_{i=1}^N \Theta_{\theta}(T_{c\alpha_i})
 \end{aligned} \tag{25}$$

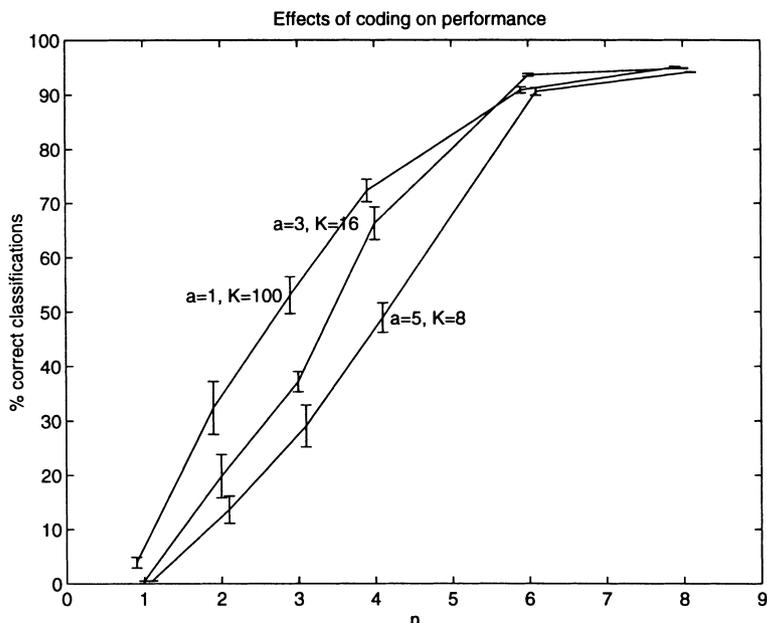


FIGURE 5. RAMnet's performance on Cut20 as a function of tuple size n for different values of a and K . The error bars of size ± 1 standard deviation are centred around the mean performance measured for 10 runs.

comprehensive illustrations are elsewhere (Allinson & Kotcz, 1993).

Because tuple score decays exponentially with Hamming distance by eqn (12), there should be relatively little ill effect if a training pattern further than L/n bits away from a test pattern is replaced by another training pattern further than L/n bits away (although Figs 1 and 7 indicate that L/n is a *very* approximate estimate). Thus if there are A scalar attributes, one can expect the non-linearity of the CMAC/Gray mapping to do little harm if $K > (LA)/n$.

There are aK bits per attribute, so this condition is

$$a < n. \quad (26)$$

Scalar differences up to $\pm K$ fall within the linear region of the mapping. This represents a fraction $(2K)/(2^a - 1)K$ or about 2^{1-a} of the largest separation allowed. With $a < n$, the ‘‘generalization Hamming distance’’ $(LA)/n = (aK)/n$ corresponds to a scalar separation of $\pm (aK)/n$, which is the fraction $(2a)/(n(2^a - 1)) \approx (a/n)2^{1-a}$ (for $a > 1$) of the largest possible scalar separation.

For $a = 1$, the mapping becomes the ‘‘thermometer code’’, in which integer x is mapped to a bit string with the last x bits set and the remaining $K - x/K$ unset. If K is adjusted to preserve the input interval, then larger a values give shorter codes, which should be similarly

effective as long as $a < n$ and scalar attributes separated by more than fraction 2^{1-a} of their dynamic range can be regarded as dissimilar as far as generalization is concerned.

Figure 5 shows the test set classification accuracy as a function of n for a 100-bit thermometer code, a 48-bit ($a = 3, K = 16$) code and a 40-bit ($a = 5, K = 8$) code for one of the datasets studied. For small values of n , the longer codes perform better, presumably because they are linear over a larger fraction of the dynamic range. However, with increasing n , the ‘‘generalization distance’’ shrinks, so more and more of the linear region of the longer codes is ignored by the RAMnet, eliminating their advantage. Meanwhile, accuracy improves with n , eventually levelling out. This is probably due to making higher-order correlations available, as well as reducing saturation effects.

5. THE EXPERIMENTS

Extensive experimental trials were conducted in order to benchmark and study the n -tuple algorithm.

5.1. Selection and Preprocessing of Statlog Data Sets

The European Community ESPRIT project 5170, the StatLog project, was designed to carry out comparative testing and evaluation of classification algorithms on

TABLE 1
Descriptions of data sets used

Name	Classes	Largest prior	Attributes	Training patterns	Testing patterns	Description
BelgianII	2	0.924	57 real	2000	1000	Classify measurements on simulated large scale power system as leading to stable or unstable behaviour.
Cut50	2	0.941	50 real	11 220	7480	Fifty measurements from a candidate segmentation point in joined handwritten text. Classify as suitable cut point or not. Commercially confidential data.
Cut20	2	0.941	20 real	11 220	7480	Best 20 attributes (by stepwise regression) from Cut50.
Technical	91	0.230	56	4500	2580	Commercially confidential. Appears to be generated by a decision tree. Most attribute values are 0.
DNA	3	0.525	180 Boolean	2000	1186	Sequences of 60 nucleotides (four-valued) classified into three categories.
SatIm	6	0.242	36 integer	4435	2000	3×3 pixel regions of Landsat images. Intensities in four spectral bands. Classified into six land uses at central pixel.
Chromo	24	0.044	16	20 000	20 000	Images of chromosomes, reduced to 16 features.
BelgianI	2	0.5664	28 real	1250	1250	As Belgian II with a smaller simulation. Attributes thought to be least informative omitted from simulation.
Tsetse	2	0.508	14 real	3500	1499	Classify environmental attributes for presence of Tsetse flies.
Letter	26	0.045	16 16-valued	15 000	5000	Images of typed capital letters, described by 16 real numbers discretized into 16 integers.
Shuttle	7	0.784	9 real	43 500	14 500	Classification problem concerning position of radiators on the Space Shuttle. Noise-free data.

large scale applications. About 20 data sets were used to estimate the performance of 23 procedures. These are described in detail by Michie et al. (1994). Each of the larger data sets (with many more than 1000 samples) were randomly split into training and testing partitions. Different methodologies (cross-validation and bootstrap) were applied to the smaller data sets. This study used the large data sets, which are summarized in Table 1. There are 11 of these.

5.2. Experimental Details

The CMAC/Gray parameters used were $K = 8$ and $a = 5$, giving 40-bit representations of the integers in $[0, 248]$. All scalar attributes were linearly rescaled and rounded to obtain integers in this interval. In the Letter data set (see Table 1), where the attributes can take on only 16 values, it would be more reasonable to use a one-out-of- N encoding with strings of 16 bits, but the CMAC/Gray procedure was used anyway for the convenience of uniformity.

The threshold θ was set to 1 in all the experiments reported here, the n -tuple size n was set to 8, and N

was set to 1000 n -tuples. These values for n and N yield the best classification results and have been determined experimentally as discussed in Section 2.2. The results reported are averages over 10 different random input mappings η . Using $n = 6$ gave similar results.

5.3. Time and Memory Requirements

Computation time requirements were insignificant in these experiments, which were carried out with a C++ program on a SUN Sparc workstation. For example, an eight-tuple network can be trained on the 2000 57-attribute training patterns of the BelgianII data set in about 49 s. Sixteen of these seconds are required just to read in the data; another four to do the CMAC/Gray conversion of the floating point attributes; and the final 29 to train the RAMnet itself. Testing the same 2000 patterns takes slightly longer, 37 s instead of 29, because a loop over classes is needed within the loop over n -tuples. Detailed timing statistics are not published for the algorithms used in the StatLog project, but it is clear that popular neural network algorithms such as back propagation and even the relatively fast radial

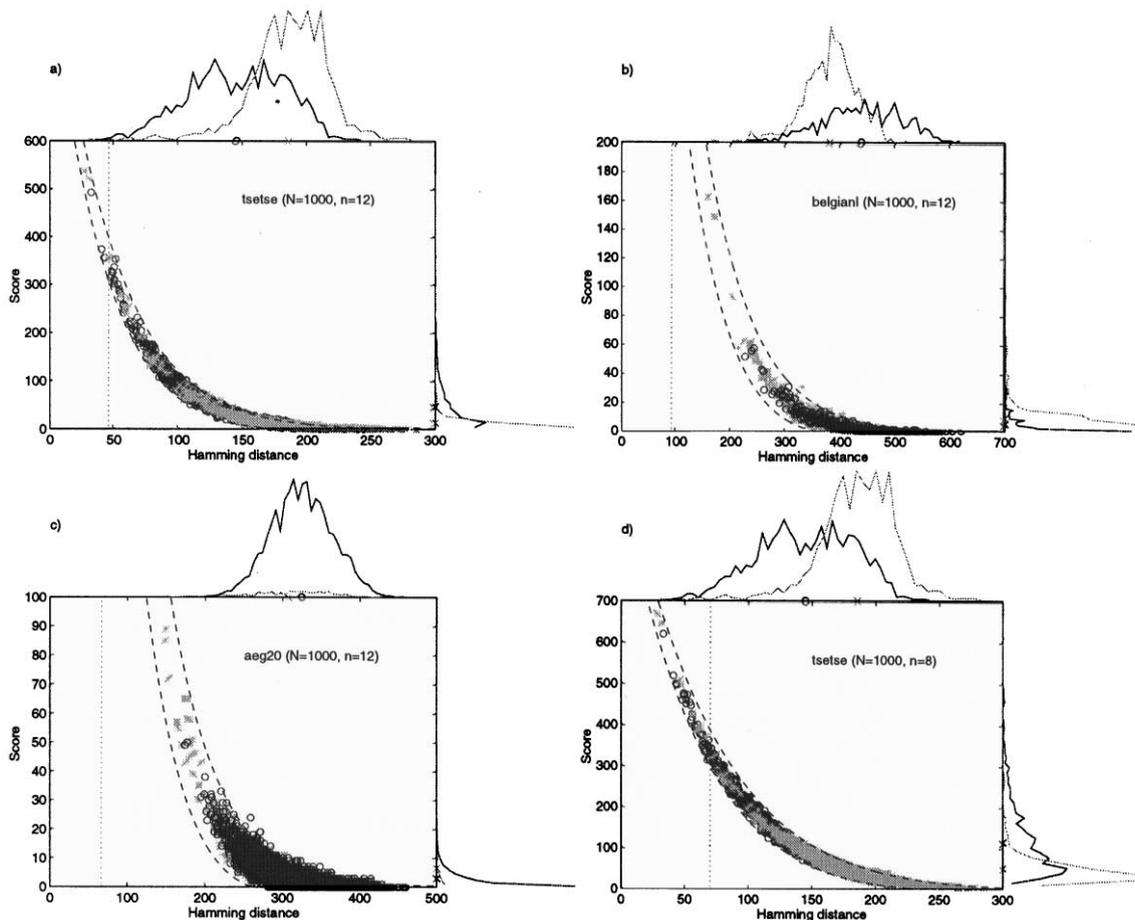


FIGURE 7. Tuple score vs Hamming distance: (a) classification error (* = 952, O = 970) due to high number of NN from the incorrect class; (b) correct classification (* = 803, O = 564) resulting from the domination of NN patterns of the correct class; (c) classification error (* = 805, O = 884) caused by exceptionally large number of patterns from the incorrect class in the tail of the distribution (skewed priors); (d) correct classification (* = 996, O = 994) owing to the smaller tuple overlap on patterns from the correct class.

basis functions are slow by comparison. The algorithm is highly parallelizable, so if it were important for the RAMnet to be even faster, special purpose parallel hardware could be designed or purchased (Aleksander et al., 1984). It would be feasible for a biological system to implement a highly parallel, but otherwise trivial calculation along these lines.

The storage requirements were moderate in most cases. In the most extreme case (Shuttle) 128 kB of RAM per class was required.

5.4. Results

The classification results for each algorithm attempted with each data set are presented in Figure 6. Table 2 gives a brief description of each algorithm with the symbol used to represent it in the figure. The classification error rates increase from left to right, and are scaled separately for each data set, so that they equal 1 at the error rate of the trivial method of always guessing the class with the highest prior probability, ignoring the input pattern. The arrows indicate the few cases in which performance was worse than this.

As remarked in Section 5.2, the results plotted for the n -tuple recognition algorithm are averages over 10

randomly selected input mappings. If the corresponding standard deviations were plotted as error bars in Figure 7, they would be obscured by the dots representing the means.

6. ANALYSIS OF RESULTS

The n -tuple method delivered competitive accuracy on six of the data sets tested (Shuttle, Letter, Tsetse, BelgianI, Chromo, SatIm), performed modestly on one (DNA) and failed entirely on the other four (BelgianII, Cut50, Cut20, Technical). Further experimental and theoretical analysis was carried out to explain the failures.

The available tuple-distance theory is not amenable to treating overlap effects, even though Figures 1–3 (and many like them) show these to be important. Therefore we generated and inspected many detailed Hamming distance vs tuple score plots such as Figure 1 in order to develop a qualitative impression of the validity of this theory when there are many training patterns.

Ignoring overlap effects, a test pattern should be assigned to the class of most of the training patterns that lie nearer to it, in Hamming distance, than about L/n . A glance at Figure 7 shows that the distribution of tuple distances bears no systematic relationship to L/n , but nevertheless it was found that the closest patterns did tend to determine the decision, at least when the class priors were roughly equal. Figures 7(a) and (b) show examples of this, which was by far the most common situation encountered. Figure 7(c) shows an error due to overlap effects, abetted by a highly skewed prior. The four troublesome datasets had highly skewed priors and predominately showed this pattern, although the very easy Shuttle dataset also had highly skewed priors. Figure 7(d) shows a relatively rare situation in which overlap effects rescued a pattern which would have been misclassified, judging by its Hamming-near neighbours. It would appear that although overlap effects are quantitatively important, they tend not to alter the conclusion that the near neighbours determine the decision, at least when the priors are relatively uniform.

Given that Hamming neighbours tend to determine the classification outcome, it seems sensible to suspect that test patterns in the four problematic data sets have a shortage of good neighbours. It turns out that they simply do not have enough neighbours at all, within the distance scales relevant to RAMnet generalization. To generalize properly, a test pattern must have at least one training pattern within a Hamming distance of about L/n . Distributed evenly over A CMAC/Gray-mapped scalar attributes, this is a scalar difference of about $(a/n)2^{1-a}$, with the attributes scaled to lie between 0 and 1, as explained in Section 4. Therefore each training pattern can provide information about any test pattern which falls within a hypercube of volume roughly $(a/n2^{1-a})^A$. The number of such cubes required to cover the region of attribute space where test data is likely to appear can be

TABLE 2
Synopsis of algorithms with symbols used in Figure 6

RAMnets	
(●)	n -tuple recogniser
Discriminators	
(♣)	Back propagation in a one-hidden-layer MLP
(♠)	Radial basis functions
(♥)	Cascade correlation.
(⊕)	SMART (projection pursuit)
(⊗)	Dipol92 (based on pairwise linear discriminators)
(⊖)	Logistic discriminant
(⊙)	Quadratic discriminant
(○)	Linear discriminant
Methods related to density estimation	
(α)	CASTLE (probabilistic decision tree).
(β)	k -NN (k nearest neighbours).
(γ)	LVQ (learning vector quantization)
(δ)	Kohonen topographic map
(ε)	Naive-Bayes (estimate assuming independent attributes)
(ζ)	ALLOC80 (kernel function density estimator)
Decision trees	
(a)	NewID (decision tree)
(b)	AC^2 (decision tree)
(c)	Cal5 (decision tree)
(d)	CN2 (decision tree)
(e)	C4.5 (decision tree)
(f)	C4.5 (decision tree)
(g)	IndCART (CART variation)
(h)	BayesTree (decision tree)
(i)	ITrule (decision tree)

crudely estimated by approximating this region as a hyper-rectangle with edge lengths given by the eigenvalues of the sample covariance matrix of the training data. Any eigenvalues smaller than $(an)2^{1-a}$ should be rounded up to this value, because the covering cubes must be at least this thick. The number of “generalization hypercubes” required to cover the data region is therefore roughly $\prod_{i=1}^A \max(1, \lambda_i/(an)2^{a-1})$ for $1 < a \leq n$, where the λ_i are the eigenvalues. Figure 8 shows this lower bound on the number of training samples required, for each dataset studied, taking $a = 5$ and $n = 8$ as in the experiments.

Aside from Technical and DNA, the problematic datasets stand out as several orders of magnitude more deficient in training data than the others, some of which are mildly deficient according to this crude estimate. DNA is special in that its Boolean attributes were treated as integers, so its data distribution will be highly non-gaussian and therefore poorly described by the covariance matrix. The Technical data set turned out to be coverable by just one hypercube, according to this estimate. Presumably then, each of its patterns looks the same to the RAMnet, and this accounts for its failure. Perhaps a non-linear rescaling of its attributes, such as histogram equalization, would help. This possibility remains to be pursued.

It is not possible to address the data deficiencies by supplying more data, especially when several orders of magnitude more samples are needed, but it is possible to tweak the RAMnet parameters to enlarge the “generalization cubes”. However, there is less room to maneuver than one would like. To enlarge the cubes, n must be decreased, but this risks degradation of performance due to loss of high-order correlation information, as indicated in Figure 6. Decreasing n also requires decreasing a , if the constraint $a \leq n$ is to be respected, keeping L/n within the linear region of the CMAC/Gray mapping. Low a values give less memory-efficient representations of scalars, at any given resolution. Systematic experiments varying the parameters did not produce significant improvements on the four problematic data sets (or the others). A more far-reaching improvement in the algorithm is required.

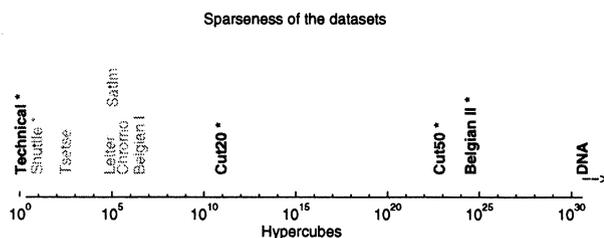


FIGURE 8. The number of hypercubes required to cover the space occupied by data. The datasets on which n -tuple classifier performed poorly are printed in bold face. A star denotes the existence of skewed priors.

7. CONCLUSIONS

Extensive experimental trials, on a scale uncommon for any algorithm, were carried out with the n -tuple classifier. The fact that this was possible at all testifies to the method’s speed, which derives from its simple principle of learning by one-shot memorization of random features. In six of the 11 datasets tested, this speed and simplicity can be enjoyed without sacrificing classification accuracy relative to 23 other slower methods, including the most popular neural network methods. Clearly, the n -tuple algorithm should be considered by every neural network researcher before a more sophisticated and therefore slower method is applied.

The theoretical tools available for the n -tuple method were reviewed and used to place it into context with other algorithms. These tools were also used to explain what went wrong on the four data sets which gave poor results. One problem is that skewed class priors tend to be problematic in a way that is difficult to quantify using current theories of “overlap effects”. A more fundamental problem is that one parameter, the n -tuple size n , controls both the distance-scales associated with generalization behaviour, and the complexity of the random features used to discriminate classes. For some data sets it is not possible to find one setting suitable for both of these considerations.

In spite of its imperfections, the n -tuple method demonstrates that its underlying principle, learning by memorization of random features, is a powerful one. It should be rewarding to develop the theory further, especially by inventing practical approximations to describe overlap effects, and to invent improved methods which incorporate the underlying principle in a more flexible way.

REFERENCES

- Aleksander, I. & Stonham, T.J. (1979a). Guide to pattern recognition using random-access memories. *Computers and Digital Techniques*, 2, 29–40.
- Aleksander, I. & Stonham, T.J. (1979b). Guide to pattern recognition using random-access memories. *IEE Proceedings on Computers and Digital Techniques (part E)*, 2(1), 29–40.
- Aleksander, I., Thomas, W.V., & Bowden, P.A. (1984). WISARD: a radical step forward in image recognition. *Sensor Review*, 4, 120–124.
- Alexander, I., & Morton, H. (1995). *An Introduction to neural computing. The secrets of the WISARD*, Chap. 5. London: Thomson.
- Allinson, N.M., & Kolcz, A. (1993). Enhanced n -tuple approximators. *Weightless Neural Network Workshop* (pp. 38–45).
- Allinson, N.M. & Kolcz, A. (1994a). Application of the CMAC input encoding scheme in the n -tuple approximation network. *IEE Proceedings on Computers and Digital Technology*, 141(3), 177–183.
- Allinson, N.M., & Kolcz, A. (1994b). The theory and practice of n -tuple neural networks. In Taylor, J.G. (ed), *Adaptive computing and information processing: neural networks*. Unicom Publishing.
- Austin, J. (1994). A Review of RAM based Neural Networks. *Proceedings of the Fourth International Conference on Microelectronics for Neural Networks and Fuzzy Systems* (pp. 58–66) Turin: IEEE Computer Society Press.

- Badr, A. (1993). n -tuple classifier for ECG signals. In Allinson, N. (Ed.), *Proceedings of the Weightless Neural Network Workshop '93, Computing with Logical Neurons* (pp. 29–32).
- Bledsoe, W.W., & Bisson, C.L. (1962). Improved memory matrices for the n -tuple pattern recognition method. *IEEE Transactions on Electronic Computers*, 11, 414–415.
- Bledsoe, W.W., & Browning, I., Pattern recognition and reading by machine. *Proceedings of the Eastern Joint Computer Conference* (pp. 232–255).
- Broomhead, D.S., & Lowe, D. (1988). Multi-variable functional interpolation and adaptive networks. *Complex Systems*, 2, 321–355.
- Flanagan, C., Rahman, M.A., & McQuade, E. (1992). A model for the behaviour of n -tuple RAM classifiers in noise. *Journal of Intelligent Systems*, 2, 187–224.
- Gallant, S., & Smith, D., (1987). Random Cells: an idea whose time has come and gone...and come again? In Caudill and Butler (eds), *IEEE International Conference on Neural Networks* (pp. II-671–II-678), San Diego, CA: IEEE.
- Kanerva, P. (1988). *Sparse distributed memory*. Cambridge, MA: MIT Press.
- Kołcz, A., & Allinson, N. (1994). Euclidian input mapping in a n -tuple approximation network. *Proceedings of the Sixth IEEE Digital Signal Processing Workshop* (pp. 285–289).
- Luttrell, S. P. (1992). Gibbs distribution theory of adaptive n -tuple networks. In Aleksander, I., & Taylor, J. (Eds.), *Artificial neural networks*, 2. (pp. 313–316) New York: Elsevier.
- MacKay, D. (1992). Bayesian interpolation. *Neural Computation*, 4, 415–447.
- Michie, D., Spiegelhalter, D.J., & Taylor, C.C., (Eds) (1994). *Machine learning, neural and statistical classification*. Englewood Cliffs, NJ: Prentice-Hall.
- Rohwer, R. (1995). Two Bayesian treatments of the n -tuple recognition method. *Proceedings of the IEE 4th International Conference on Artificial Neural Networks* (publication 409) (pp. 171–176).
- Rohwer, R., & Cressy, D. (1989). Phoneme classification by Boolean networks. In Tubach, J.P., and Mariani, J.J. (Eds.), *Proceedings of the European Conference on Speech Communication and Technology* (pp. 557–560) Paris: CEP, 26–28 Albany St., Edinburgh, UK.
- Rohwer, R., & Lamb, A. (1993). An exploration of the effect of super large n -tuples on single layer RAMnets. In Allinson, N. (Ed.), *Proceedings of the Weightless Neural Network Workshop '93, Computing with Logical Neurons* (pp. 33–37).
- Sixsmith, M.J., Tattersall, G.D., & Rollett, J.M. (1990). Speech recognition using n -tuple Techniques. *British Telecom Technology Journal*, 8, 50–60.
- Songcan, Y., Guoqing, Ch., & Jun, L. (1992). Multilayer parallel distributed pattern recognition system model using sparse RAM nets. *IEE Proceedings-E*, 139(2), 144–146.
- Stonham, T.J. (1977). Improved Hamming-distance analysis for digital learning networks. *Electronic Letters*, 13(6), 155–156.
- Sutton, R.S., & Whitehead, S.D. (1993). Online learning with random representations. *Tenth International Conference on Machine Learning (ML'93)*.
- Tarling, R., & Rohwer, R. (1993). Efficient use of training data in the n -tuple recognition method. *Electronics Letters*, 29(24), 2093–2094.
- Tattersall, G.D., & Johnson, R.D. (1984). Speech recognisers based on n -tuple sampling. *Proceedings of the Institute of Acoustics—Spring Conference* (pp. 405–413).
- Tattersall, G.D., Foster, S., & Johnston, R.D. (1991). Single-layer lookup perceptrons. *IEE Proceedings-F*, 138, 46–54.
- Ullmann, J.R., & Kidd, P.A. (1969). Recognition experiments with typed numerals from envelopes in the mail. *Pattern Recognition*, 1, 273–289.
- Ullmann, J.R. (1969). Experiments with the n -tuple method of pattern recognition. *IEEE Transactions on Computers*, 18(12), 1135–1137.