Information Efficiency of the Associative Net at Arbitrary Coding Rates

Bruce Graham and David Willshaw

Centre for Cognitive Science, University of Edinburgh
2 Buccleuch Place, Edinburgh EH8 9LW, Scotland, UK
E-mail: *B.Graham@cns.ed.ac.uk & D.Willshaw@cns.ed.ac.uk*
Phone: +44 131 650 4404; Fax: +44 131 650 6626

Abstract. The associative net is a neural network model of associative memory that is unusual in having binary-valued connections between units. This net can work with high information efficiency, but only if the patterns to be stored are extremely sparse. In this paper we report how the efficiency of the net can be improved for more dense coding rates by using a partially-connected net. The information efficiency can be maintained at a high level over a 2–3 order of magnitude variation in the degree of pattern sparseness.

1 Introduction

The associative net (Willshaw, Buneman, & Longuet-Higgins, 1969; Willshaw, 1971) is a distributed model of heteroassociative memory. It consists of a set of input units connected to a set of output units by modifiable synapses. Pairs of binary patterns are stored in the net using a "clipped Hebbian" learning rule that changes the weight of a synaptic connection from 0 to 1 if both input and output units are active for the same pattern pair. One performance measure of the associative net is how much information can be stored and successfully recalled from it. The information efficiency, η_o, of the net is the number of stored pattern pairs multiplied by the amount of information that can be recalled about each output pattern, divided by the amount of storage, which is the number of synaptic connections. The fully-connected net (every input unit is connected to every output unit) has an optimum information efficiency $\sim \ln 2$ $(= 0.693)$ bits per synapse, when the patterns are suitably sparse. This is remarkably close to the limit of information efficiency, $1/(2\ln 2)$ (~ 0.721), that can be achieved for sparse patterns using continuous-valued synapses (Gardner, 1987). So the lack of precision in the synapses of the associative net is not crucial for the storage of sparse patterns.

Unfortunately, optimum efficiency for reliable recall (at most one bit in error in the recalled output pattern) is only achieved when the number of active units, M_A, in an input pattern, is given by $M_A = \log_2(N_B)$, where N_B is the number of output units (Willshaw et al., 1969). For square nets, where the number of input units, N_A, is the same as the number of output units, the input patterns quickly become very sparse indeed as the net size increases. For example, if

$N_A = N_B = 2^{20}$, only 20 units out of 1,048,576 should be active in each input pattern. Though the exact coding rate is unspecified, the output patterns should also be very sparse if the limit of efficiency is to be approached.

It has been shown that the optimum efficiency of the associative net can be retained for even sparser coding ($M_A < \log_2(N_B)$), if errors are allowed during recall (error-full recall) (Nadal & Toulouse, 1990; Buckingham & Willshaw, 1992). However, as the density of input pattern coding is increased beyond the optimum, information efficiency drops rapidly and entering the error-full regime does not help. This contrasts with continuous-valued synapses, where an efficiency of 2 bits per synapse can be obtained theoretically at dense coding levels (Gardner, 1987). Nadal and Toulouse (1990) have shown that even simple Hebbian learning, in which a synaptic weight is incremented by 1 if the input and output units are active at the same time, saturates at an efficiency limit of $1/(\pi \ln 2)$ (~ 0.459) bits per synapse in the error-full regime, for dense coding.

In this paper, we present an approach that allows the associative net to retain near maximal information efficiency for input coding rates 2–3 orders of magnitude higher than the optimum for the fully-connected net. This is achieved by using a partially-connected net with a suitable thresholding strategy for recall. Optimum efficiency is obtained with reliable recall, with no advantage being available by entering the error-full regime. Though efficiency still drops to low levels as standard coding ($M_A = N_A/2$) is approached, it remains finite for the nets considered here (around 0.05 bits per synapse). For the comparable fully-connected net it is effectively zero.

2 Theoretical Information Efficiency

The information efficiency, η_o, of reliable recall from the fully-connected net is a function of the average probability, $\hat{\rho}$, that a synapse has been modified during pattern storage (Willshaw et al., 1969; Willshaw, 1971; Buckingham & Willshaw, 1992). Optimum efficiency is achieved when $\hat{\rho} = 0.5$ and is approximately $\eta_o = \ln 2$. Under the condition that at most one bit in the output pattern is allowed to be in error, the optimum is achieved when the number of active units in an input pattern is $M_A = \log_2(N_B)$.

The information efficiency of a partially-connected net can be derived in exactly the same manner, if the appropriate threshold strategy is used during recall. For the fully-connected net with a noise-free input cue, an output pattern is recalled by selecting those output units to be active that are connected to all M_A active input units by modified synapses. In a partially-connected net, an output unit that should be active is not necessarily connected to all M_A active input units, so a different threshold for output unit activity must be found. The appropriate threshold is the number of active input units an output unit is connected to, by modified or unmodified synapses. This threshold is different for each output unit. The use of this strategy for partially-connected nets has been investigated in detail elsewhere (Buckingham & Willshaw, 1993; Graham & Willshaw, 1994, 1995).

Given this method of pattern recall, the optimum efficiency is again $\eta_o = \ln 2$ when $\hat{\rho} = 0.5$. This is now achieved under the further condition that $ZM_A = \log_2(N_B)$, where Z is the fraction of input units connected to any given output unit (ZN_AN_B is the number of synapses in the net). So as the input coding rate is raised above $M_A = \log_2(N_B)$, it is simply a matter of lowering the connectivity level in proportion to retain optimum information efficiency.

The approximations used to derive these results are not entirely accurate, even at sparse coding rates (Buckingham & Willshaw, 1992). In order to investigate information efficiency over a wide range of coding rates, accurate numerical calculations of recall performance need to be used. The following information efficiency results for two large nets have been obtained using such calculations. The technical details of these calculations have been given previously (Buckingham & Willshaw, 1992, 1993; Graham & Willshaw, 1994). We only consider input coding rates $M_A \geq \log_2(N_B)$, as the efficiency at sparser rates has been dealt with elsewhere (Nadal & Toulouse, 1990; Buckingham & Willshaw, 1992).

3 Results

Calculations of information efficiency have been made for two different associative nets. The first net has $N_A = 2^{17} = 131072$ and $N_B = 2^{20} = 1048576$. This expansion in size from the input to the output layer represents that seen in the projection from the entorhinal cortex to the dentate gyrus in the mammalian hippocampus. The number of units corresponds roughly to that found in the rat brain (Amaral, Ishizuka, & Claiborne, 1990). For this net, the numbers of active input and output units are varied independently. The second net is square, with $N_A = N_B = 2^{18} = 262144$. This could represent the recurrent collaterals of CA3, or the projection from CA3 to CA1 in rat hippocampus. In this case, the numbers of active inputs and outputs are varied but kept equal.

Figure 1 shows the information efficiency of the first net for different coding rates and connectivities. The theory gives the optimality criterion, $ZM_A = 20$. For each value of the input and output coding levels, M_A and M_B, the information efficiency was calculated over a range of connectivities, Z. In all cases, optimum efficiency was obtained with a connectivity very close to that specifed by the above criterion. For $Z = 20/M_A$ (Fig. 1a), maximum efficiency is obtained at $M_A = 20$, and decreases as M_A is increased, regardless of the output coding rate. While η_o does drop rapidly as dense input coding levels are approached, it is maintained at a high level for a 2–3 order of magnitude increase in M_A beyond $M_A = 20$. It is still above 90% of maximum at 2 orders of magnitude, and above 60% at 3 orders of magnitude. This contrasts markedly with the efficiency of a fully-connected net (Fig. 1b). Now η_o decreases very rapidly with increasing M_A, and is only at about 10% of maximum after a 2 orders of magnitude increase in M_A. It is at this magnitude that the maximum improvement in efficiency is obtained by using partial connectivity. The difference between η_o at optimal partial connectivity and at full connectivity is shown in Fig. 1c. Note that more patterns can be stored in a fully-connected net than a partially-connected net.

For example, at $M_A = 1024$, the fully-connected net can store and reliably recall 7 times more patterns than the optimum partially-connected net, but requires roughly 50 times more connections to do so.

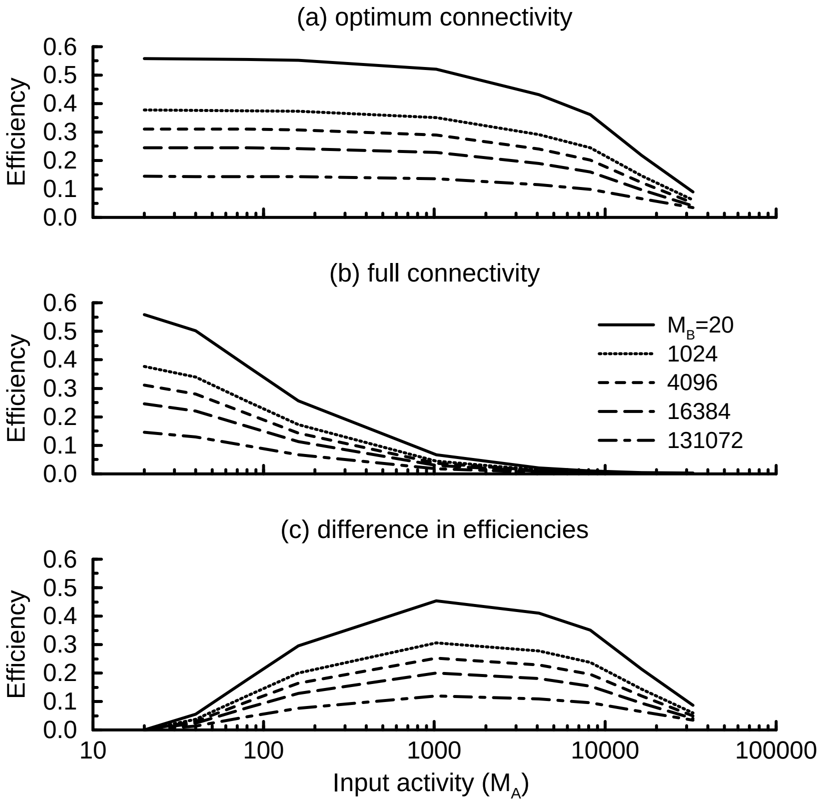

Fig. 1. Information efficiency of the first net ($N_A = 131072$, $N_B = 1048576$)

The sensitivity of η_o to variations in connectivity is demonstrated in Figure 2. It is only mildly sensitive to Z around the optimum connectivity, which is essentially $Z_o = \log_2(N_B)/M_A$. When M_A is held constant, at $Z = 2Z_o$ efficiency has dropped to around 90% of maximum, and at $Z = Z_o/2$ it is down to 80% of maximum.

The maximum efficiency is also dependent on the output coding rate, M_B, and is somewhat less than $\ln 2$. It decreases linearly with $\log(M_B)$ and, at least for $M_A = 20$, a reasonable approximation for η_o is

$$\eta_o = \left(1 - \frac{\log_2(M_B)}{\log_2(N_B)}\right) \ln 2 \tag{1}$$

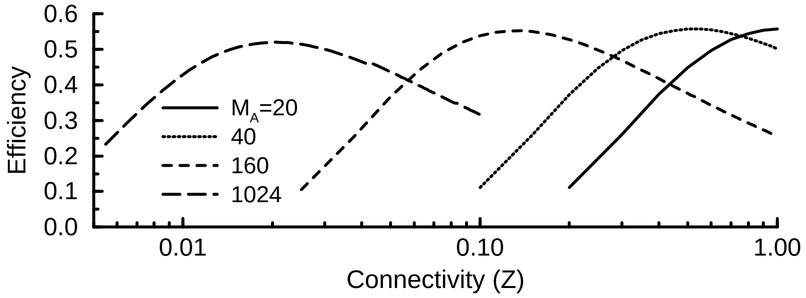

Fig. 2. Sensitivity of information efficiency to variations in connectivity around the optimum

The results for the second net are similar to those for the first net. The information efficiency obtained using either optimum or full connectivity, while coding levels M_A and M_B are varied equally, is plotted in Figure 3a. At optimum connectivity, η_o decreases approximately linearly as the log of the coding levels is increased. The maximum improvement by using partial connectivity is again obtained when M_A is increased by about 2 orders of magnitude from the initial value of $M_A = \log_2(N_B) = 18$ (Fig. 3b).

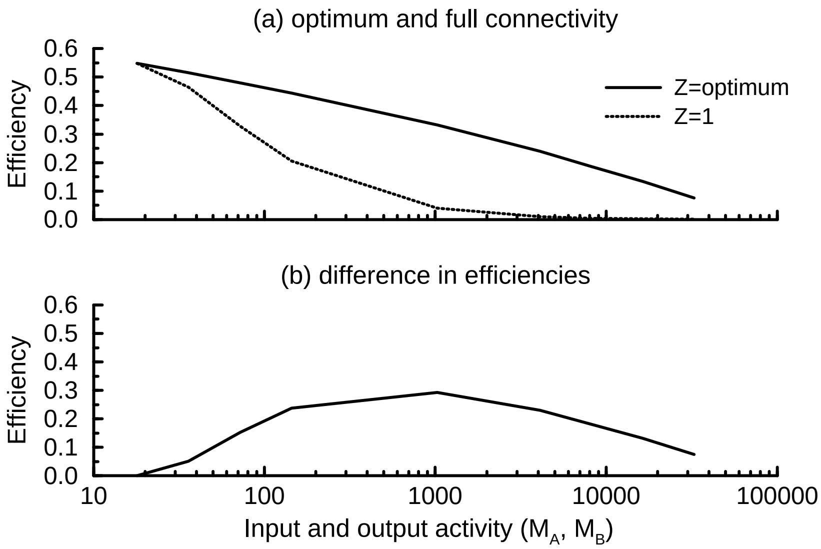

Fig. 3. Information efficiency of the second net ($N_A = N_B = 262144$; $M_A = M_B$)

4 Discussion

We have demonstrated, with numerical calculations, an associative net model with high information efficiency for input coding levels above the optimum for a fully-connected net of $M_A = \log_2(N_B)$. By lowering the connectivity level as M_A is increased, so that $ZM_A = \log_2(N_B)$, and using a suitable thresholding strategy for recall, information efficiency is held at high levels for a 2–3 order of magnitude increase in M_A. Optimum efficiency is achieved with reliable recall, and no advantage is gained by entering the error-full regime. This contrasts with the situation for input coding rates less than $\log_2(N_B)$, and other learning algorithms, such as Hebbian learning, where optimum efficiency may be found in the error-full regime (Nadal & Toulouse, 1990; Buckingham & Willshaw, 1992).

This work goes some way towards extending the associative net as a useful model for arbitrary coding rates. However, the performance of the associative net at standard coding ($M_A = N_A/2$, $M_B = N_B/2$) remains rather poor, with a maximum information efficiency of around 0.05 bits per synapse for the large nets studied here.

Acknowledgement: to the MRC for financial support under PG 9119632.

References

Amaral, D., Ishizuka, N., & Claiborne, B. (1990). Neurons, numbers and the hippocampal network. In Storm-Mathisen, J., Zimmer, J., & Ottersen, O. (Eds.), *Progress in Brain Research, Vol.83*, pp. 1–11. Elsevier Science Publishers B.V.

Buckingham, J., & Willshaw, D. (1992). Performance characteristics of the associative net. *Network*, *3*, 407–414.

Buckingham, J., & Willshaw, D. (1993). On setting unit thresholds in an incompletely connected associative net. *Network*, *4*, 441–459.

Gardner, E. (1987). Maximum storage capacity in neural networks. *Europhysics Letters*, *4*, 481.

Graham, B., & Willshaw, D. (1995). Improving recall from an associative memory. *Biol. Cybern.*, *72*, 337–346.

Graham, B., & Willshaw, D. (1994). Capacity and information efficiency of a brain-like associative net. In Tesauro, G., Touretzky, D., & Leen, T. (Eds.), *Neural Information Processing Systems 7*, pp. 513–520. MIT Press.

Nadal, J.-P., & Toulouse, G. (1990). Information storage in sparsely coded memory nets. *Network*, *1*, 61–74.

Willshaw, D. (1971). *Models of distributed associative memory.* Ph.D. thesis, University of Edinburgh.

Willshaw, D., Buneman, O., & Longuet-Higgins, H. (1969). Non-holographic associative memory. *Nature*, *222*, 960–962.

This article was processed using the LaTeX macro package with LLNCS style