

This is a repository copy of *Improved storage capacity in correlation matrix memories storing fixed weight codes*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/68013/>

Version: Accepted Version

---

**Proceedings Paper:**

Hobson, Stephen John and Austin, Jim [orcid.org/0000-0001-5762-8614](https://orcid.org/0000-0001-5762-8614) (2009) Improved storage capacity in correlation matrix memories storing fixed weight codes. In: Lecture Notes in Computer Science. Lecture Notes in Computer Science . Springer , ICANN 2009 , pp. 728-736.

---

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# Improved storage capacity in correlation matrix memories storing fixed weight codes

Stephen Hobson and Jim Austin

Advanced Computer Architectures Group  
Department of Computer Science  
University of York  
Heslington  
York, YO10 5DD, UK  
{stephen, austin}@cs.york.ac.uk

**Abstract.** In this paper we introduce an improved binary correlation matrix memory (CMM) with better storage capacity when storing sparse fixed weight codes generated with the algorithm of Baum et al. [3]. We outline associative memory, and describe the binary correlation matrix memory—a specific example of a distributed associative memory. The importance of the representation used in a CMM for input and output codes is discussed, with specific regard to sparse fixed weight codes. We present an algorithm for generating of fixed weight codes, originally given by Baum et al. [3]. The properties of this algorithm are briefly discussed, including possible thresholding functions which could be used when storing these codes in a CMM; L-max and L-wta. Finally, results generated from a series of simulations are used to demonstrate that the use of L-wta as a thresholding function provides an increase in storage capacity over L-max.

**Key words:** associative memory, correlation matrix memory, storage capacity, fixed weight codes, pattern recognition

## 1 Introduction

The fixed weight code generation algorithm of Baum et al. [3] has the benefit of generating unique codes which are well separated in the pattern space, which makes them suited for storage in a CMM. We will subsequently term the codes generated from this algorithm *Baum codes*, for ease of reference. Recall of such codes from a CMM requires the use of a thresholding function, and it is with the nature of this function that this paper is concerned. L-max thresholding [1] has been shown to be an effective thresholding function for fixed weight codes, and has been applied to CMMs storing Baum codes [2]. However, L-max thresholding fails to make use of all the constraints on Baum codes. A thresholding mechanism which takes advantage of these constraints is able to provide an improved storage capacity for a CMM which stores Baum codes.

## 2 Associative Memory

Traditional computer memories store data in a compartmentalised fashion, with each stored item having a unique address. While this leads to perfect recall in all cases where the correct address is known, any amount of error in the address will result in a recall which bears no relation to the stored item. A contrasting model of memory is *distributed associative memory*. In such a model, data items are stored as pairs, with the presentation of the first member of the pair to the memory resulting in the recall of the second. Rather than these associations being stored in a single location in the memory, they are distributed across the memory. This provides robustness in the presence of noise on the input, and enables generalisation.

Such a memory serves a different purpose to a traditional memory. While a standard computer memory is well suited to tasks such as storing a list of tasks or events, it is less capable of “answering questions” [6]. Such a task would require the question to be looked up in a list, which might contain the location of the answer. In a distributed associative memory, the answer is retrieved simply by presenting the question to the input. The recall operation does not require a look-up algorithm, and so is a much more efficient operation.

## 3 Binary Correlation Matrix Memories

A Binary Correlation Matrix Memory (CMM) [8] is one example of a distributed associative memory. It stores the correlations between input and output codes. The memory is a one layer fully connected neural network. This means that the weights can be viewed as an  $m \times n$  matrix  $W$ , where  $m$  is the size of the input vectors and  $n$  is the size of the output vectors. An example of such a memory is shown in Fig. 1. Although it is possible to use a CMM with non-binary weights [5], only the binary case will be considered in this paper.


Learning is achieved using an outer product of the input and output. These matrices are combined using an OR function over all input output pairs to create the weight matrix  $W$ .

$$W = \bigvee_{i=1}^N x_i y_i^T \quad (1)$$

Recall is achieved as shown in Equation. 2.

$$y = f [Wx] \quad (2)$$

Here  $f$  is a thresholding function, which takes the activity output  $Wx$  and converts it to a binary vector. For example, in their original paper Willshaw et al. [8] suggested that the thresholding function could set all output nodes with activity greater than the number of 1s in the input pattern  $x$  to 1. The choice of this threshold function has a profound effect on the storage capability of the network, as we shall see later.



**Fig. 1.** An example of a CMM with 8 input neurons, 8 output neurons and binary weights.

When recalling a pattern from the memory, the resulting vector (before thresholding) can be viewed as a signal (the original stored pattern) and some noise (extra activity from overlaps with other learned codes).

## 4 Sparse Fixed Weight Coding

In a neural memory such as a CMM there is an intrinsic link between the data representation used and the storage capability of the memory. Using different encodings on the input and output of the memory will have different effects on the storage capacity. In addition, the choice of threshold function will also relate directly to the representation used.

Perhaps the simplest representation of all would be the use of unary output codes (a single bit set to 1 in  $n$  bits). This provides a storage capacity of exactly  $n$  code pairs. Each input code will be stored in exactly one column of the matrix, and given a correct input code there will be no error on recall. However, the fault tolerance capability of the network is lost, since the storage is no longer distributed. It is necessary to use input and output codes with more than one bit set to 1 to distribute storage over the network.

Furthermore, in order to maximise the storage capability of the network, these codes should be sparsely coded. More specifically, the number of 1s in an  $n$  bit code should be in the order of  $\log n$  [7]. The number of 1s in an  $n$  bit code is termed the *weight* of the code. An important property of the use of sparse codes in a CMM is that the memory is capable of storing  $k > n$  codes (where  $k$  is the number of pairs stored, and  $n$  is the number of input neurons), providing a small amount of recall error is tolerated [7].

If codes with fixed weight are used, an alternative threshold function is available; L-max thresholding [1]. This sets the  $l$  neurons with the highest output activity to 1 and the rest to 0, where  $l$  is the weight of the output code. Casasent and Telfer [4] experimented with various output encodings, including Binary codes, Hamming codes and fixed weight codes, albeit with analog input codes. They found that in the presence of noise, fixed weight codes with L-max thresholding gave the greatest storage capacity for a given code length.

It is important to have the ability to generate fixed weight codes in such a fashion that the codes generated are guaranteed to be well separated in pattern space. Baum et al. proposed an algorithm which generates fixed weight codes which have a small amount of overlap [3]. The code is divided into  $l$  sections which are relatively prime<sup>1</sup> (coprime) in length, with each section  $i$  having length  $p_i$ . For example, a code of length 32 where  $l = 3$  could be divided into sections of length 16, 9 and 7. The size of  $l$  defines the weight of the code. To generate code number  $c$ , we set the bit in position  $j$  as follows (where  $x$  is the code to output):

---

<sup>1</sup> Two integers are relatively prime if they have no common factor other than 1. It should be noted that the problem of generating a set of relatively prime numbers which sum to a total is not trivial. However, a discussion of methods is beyond the scope of this paper.

$$\begin{aligned}
 x_j^c &= 1 \text{ if } j - \sum_{k=1}^{i-1} p_k \equiv c \pmod{p_i} \\
 &= 0 \text{ otherwise}
 \end{aligned}
 \tag{3}$$

Essentially what is happening is that a single bit will be set to 1 in each section of the code. As subsequent codes are generated, the next bit in each section will be set to 1 instead, wrapping around to the beginning of the section when the end is reached. For example, Fig. 2 shows a code with  $n = 10$  and  $l = 3$ , taking  $p_1 = 5, p_2 = 3, p_3 = 2$ .

1	0	0	0	0	1	0	0	1	0
0	1	0	0	0	0	1	0	0	1
0	0	1	0	0	0	0	1	1	0
0	0	0	1	0	1	0	0	0	1
0	0	0	0	1	0	1	0	1	0
1	0	0	0	0	0	0	1	0	1
0	1	0	0	0	1	0	0	1	0
0	0	1	0	0	0	1	0	0	1
0	0	0	1	0	0	0	1	1	0
0	0	0	0	1	1	0	0	0	1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

**Fig. 2.** An example of the generation of Baum codes. Here,  $l = 3$  and  $p_1 = 5, p_2 = 3, p_3 = 2$ , giving a code of length 10

Using this mechanism  $p_1 \times p_2 \times \dots \times p_s$  unique codes can be generated, which is substantially fewer than it is possible to represent with a general fixed weight coding scheme ( $\frac{n!}{(n-l)!l!}$ ). However, the overlap between the codes is guaranteed to be small, which improves recall accuracy if they are used in a CMM. Since the method is deterministic, we can be certain about the amount of overlap between generated codes. With no loss of generality we can consider a Baum code with section lengths  $p_1 < p_2 < \dots < p_3$ . The first  $p_1$  codes generated will have no overlap at all. The first  $p_1 p_2$  overlap by at most 1 bit (a Hamming distance of at least  $2l - 2^2$ ). In their analysis Baum et al. [3] state that if  $\prod_{i=1}^t p_i$  codes are used, the minimum Hamming distance between any two codes will be  $d = 2(l - t + 1)$ . For this reason, it is beneficial for  $p_i \approx n/l$ , since this maximises the product between the section lengths  $p_i$ , and hence the number of codes which can be generated with minimal overlap.

---

<sup>2</sup> The Hamming distance between two codes is the number of bits which differ between them.

## 5 Improving the Storage Capacity

In the past the algorithm of Baum et al. [3] has been used to generate fixed weight codes, with L-max used as the thresholding function [2]. This represents an oversight, since L-max thresholding may produce output codes which are not possible under the Baum algorithm. By constraining the threshold so that only the codes generated by the algorithm are output, an increased storage capacity can be achieved.

It has already been mentioned that the algorithm divides the code into a series of sections. Baum et al. point out in the appendix to their paper that a useful property of the algorithm is that there is exactly one 1 in each section of the code. This means that a winner-takes-all (WTA) threshold can be applied to each section of the code, rather than taking the  $l$  highest values from the whole code, as we would with L-max [3]. This thresholding technique incorporates more information about the output encoding into the thresholding function, and therefore provides a more robust thresholding. We shall call this thresholding technique L-wta.

## 6 Results

To demonstrate the improved storage capacity of a CMM when using L-wta compared to L-max a series of simulations were conducted. The storage of a CMM is affected by the size of the input and output codes, and also by the weight of the coding system used. For this reason L-wta technique was compared to L-max for a variety of coding systems. In each experiment an empty CMM was created for the appropriate code sizes. The following steps were then undertaken.

1. Generate an input code according to the algorithm of Baum et al. [3] This code will be unique.
2. Generate a random output code. This code is a random code from the entire space of possible Baum codes for the given set of coprimes, and so may not be unique.
3. Train the CMM using the generated input/output pair.
4. Present every input which the CMM has learnt and compare the correct output to the actual output using L-max and L-wta.
5. If average error (defined below) exceeds 10% for all thresholding techniques then exit, otherwise return to 1.

For each experiment these steps were run with twenty CMMs. A different integer was used to seed the random generator for each CMM, resulting in differing output patterns being trained. After each iteration the average error was calculated for all twenty CMMs. The recall error was defined as the percentage of recalled patterns which contained an error in any bit. In order to measure the performance of the thresholding techniques at a variety of error tolerances we examine the number of codes learnt in each memory before recall error exceeded 0.1%, 1%, 5% and 10%.

**Table 1.** Experimental results when varying the size of the input code. All tables show the number of codes learnt before errors at given levels. Codes are given in the format *length/weight*. Note that in some cases code lengths are approximate. This is due to the complexity of generating large sets of coprime numbers which sum to a given target. Bold numbers show the percentage increase in storage capacity when using L-wta rather than L-max.

Input code	Output code	0.1% error			1% error			5% error			10% error		
		L-max	L-wta	%	L-max	L-wta	%	L-max	L-wta	%	L-max	L-wta	%
64/4	256/4	70	78	<b>11.4</b>	115	129	<b>12.2</b>	153	171	<b>11.8</b>	179	207	<b>15.6</b>
128/4	256/4	139	141	<b>1.4</b>	197	234	<b>18.8</b>	289	334	<b>15.6</b>	345	406	<b>17.7</b>
256/4	256/4	259	286	<b>10.4</b>	424	473	<b>11.6</b>	600	696	<b>16.0</b>	719	831	<b>15.6</b>
512/4	256/4	496	589	<b>18.8</b>	814	927	<b>13.9</b>	1182	1369	<b>15.8</b>	1416	1630	<b>15.1</b>

**Table 2.** Experimental results when varying the weight of the input code.

Input code	Output code	0.1% error			1% error			5% error			10% error		
		L-max	L-wta	%	L-max	L-wta	%	L-max	L-wta	%	L-max	L-wta	%
512/2	256/4	260	260	<b>0.0</b>	282	304	<b>7.8</b>	482	555	<b>15.1</b>	577	707	<b>22.5</b>
512/4	256/4	496	589	<b>18.8</b>	814	927	<b>13.9</b>	1182	1369	<b>15.8</b>	1416	1630	<b>15.1</b>
512/8	256/4	1023	1036	<b>1.3</b>	1453	1603	<b>10.3</b>	1811	1989	<b>9.8</b>	2028	2229	<b>9.9</b>
512/16	256/4	1186	1267	<b>6.8</b>	1408	1512	<b>7.4</b>	1673	1824	<b>9.0</b>	1829	1989	<b>8.7</b>

**Table 3.** Experimental results when varying the size of the output code.

Input code	Output code	0.1% error			1% error			5% error			10% error		
		L-max	L-wta	%	L-max	L-wta	%	L-max	L-wta	%	L-max	L-wta	%
256/4	64/4	91	120	<b>31.9</b>	178	203	<b>14.0</b>	245	286	<b>16.7</b>	287	337	<b>17.4</b>
256/4	128/4	148	179	<b>20.9</b>	265	289	<b>9.1</b>	362	429	<b>18.5</b>	436	508	<b>16.5</b>
256/4	256/4	259	286	<b>10.4</b>	424	473	<b>11.6</b>	600	696	<b>16.0</b>	719	831	<b>15.6</b>
256/4	512/4	373	415	<b>11.3</b>	608	712	<b>17.1</b>	950	1099	<b>15.7</b>	1137	1327	<b>16.7</b>

**Table 4.** Experimental results when varying the weight of the output code.

Input code	Output code	0.1% error			1% error			5% error			10% error		
		L-max	L-wta	%	L-max	L-wta	%	L-max	L-wta	%	L-max	L-wta	%
256/4	512/2	564	709	<b>25.7</b>	1259	1435	<b>14.0</b>	1932	2138	<b>10.7</b>	2310	2599	<b>12.5</b>
256/4	512/4	373	415	<b>11.3</b>	608	712	<b>17.1</b>	950	1099	<b>15.7</b>	1137	1327	<b>16.7</b>
256/4	512/8	257	267	<b>3.9</b>	353	400	<b>13.3</b>	495	569	<b>14.9</b>	580	677	<b>16.7</b>
256/4	512/16	71	89	<b>25.4</b>	138	157	<b>13.8</b>	197	222	<b>12.7</b>	219	266	<b>21.5</b>



Table 1 shows the results when the size of the input was varied, whilst size and weight of the output code remained constant. Similarly, Table 2 shows the results when the weight of the input code was varied. In both cases it can be seen that the use of L-wta results in an increase of approximately 15% in storage capacity. L-wta appears to provide the largest increase in storage over L-max when the input code is sparse; that is, when the code size is increased or the weight is decreased. This advantage appears less pronounced as the amount of output error increases.

The case is similar when examining Tables 3 and 4. The effect of output sparsity on the effectiveness of the technique is less clear. However, the storage capacity when using L-wta is consistently a considerable improvement over that achieved using L-max.

Fig. 3 shows two examples of the performance of the two thresholding techniques as codes are trained into the memories. It can clearly be seen that as the memory becomes increasingly saturated, the use of L-wta provides an increasing benefit over L-max thresholding.


## 7 Summary

In summary, it has been demonstrated in this paper that when using codes generated by the algorithm of Baum et al. [3] L-wta provides an increase in storage capacity over thresholding using L-max, provided some error is tolerated. This increase in storage capacity is generally in the order of 15%, but has been observed to be as high as 30%.

While this paper has demonstrated the benefit of this thresholding technique across a variety of conditions, a mathematical treatment of the technique is now required.

## References

1. J. Austin, T. Stonham. Distributed associative memory for use in scene analysis. *Image and Vision Computing* 5, 251–260, 1987.
2. J. Austin, J. V. Kennedy and K. Lees. A neural architecture for fast rule matching. In *Artificial Neural Networks and Expert Systems Conference*, 1995.
3. E. B. Baum, J. Moody, and F. Wilczek. Internal representations for associative memory. *Biological Cybernetics*, 59(4):217–228, 1988.
4. D. Casasent, B. Telfer. High capacity pattern recognition associative processors. *Neural Networks*, 5, 687–698, 1992.
5. J. Nadal, G. Toulouse. Information storage in sparsely coded memory nets. *Network*, 1, 61–74, 1990.
6. G. Palm. *Neural Assemblies*. Springer-Verlag, 1982.
7. G. Palm, F. Schwenker, F.T. Sommer and A. Strey. Neural associative memories *Associative processing and processors*, 307–326, 1997.
8. D. J. Willshaw, O. P. Buneman, and H. C. Longuet-Higgins. Non-holographic associative memory. *Nature*, 222(7):960–962, 1969.



**Fig. 3.** Two comparisons of the storage capabilities of a CMM when using L-max and L-wta. Dotted lines show the standard deviation of average recall error between runs of the experiment. (top) Input codes had size 256 and weight 4. Output codes had size 512 and weight 4. (bottom) Input codes had size 512 and weight 16. Output codes had size 256 and weight 4.