THE ROYAL SOCIETY

Analysis of the connectional organization of neural systems associated with the hippocampus in rats

Gully A. P. C. Burns† and Malcolm P. Young

Neural Systems Group, Department of Psychology, University of Newcastle upon Tyne, Ridley Building, Newcastle upon Tyne NE1 7RU, UK

The hippocampus of the rat enjoys a central significance for researchers interested in the neural mechanisms of memory and spatial information processing. Many of the theoretical models advanced to explain function in this system, however, do not reflect the wealth of information on the connectivity of these structures, and employ greatly simplified treatments of its complex connectivity. We were interested in whether a more analytical approach, which begins with analysis of the connectivity of the system, might provide insights complementary to those derived by synthetic models. Accordingly, we collated detailed neuroanatomical information about the connectivity of the hippocampal system in the rat, and analysed the resulting data. Analyses of connectivity based on a variety of different analytical techniques have recently been used to elucidate the global organization of other systems in the macaque and cat, and have given rise to successful predictions. We applied non-metric multidimensional scaling and non-parametric cluster analysis to our summary matrix of connection data. The analyses produced organizational schemes that were consistent with known physiological properties and provided the basis for making tentative predictions of the further structures that may contain 'place' and 'head-direction' cells, which structures we identify. The consistency between the analyses of connectivity and the distribution of physiological properties across the system suggests that functional relationships are constrained by the organization of the connectivity of the system, and so that structure and function are linked at the systems level.

Keywords: spatial memory; cluster analysis; place cells; head-direction cells; corticocortical connections; neuroinformatics

1. INTRODUCTION

The hippocampus of the rat has taken on a central significance for researchers interested in the neural mechanisms of memory and spatial information processing. Interest in the hippocampus stems from a number of sources. For example, individual neurons within the CA3 and CAl regions of Ammon's horn fire preferentially when the rat is in a specific region of space (O'Keefe & Dostrovsky 1971), even when the animal navigates there in darkness (Quirk *et al*. 1990), and destruction of the hippocampus impairs the ability of rats to re-navigate to an invisible submerged platform (Morris *et al*. 1982). This interest has led the hippocampal formation to become the focus of a great deal of theoretical work concerning a variety of different ideas about central processing. These theories involve such topics as the mechanism of depression (Gray 1982), memory trace formation (Buzsaki 1989), cognitive mapping and systems of path integration (O'Keefe & Nadel 1978; McNaughton *et al*. 1996), and declarative memory (Eichenbaum *et al*. 1992).

Although different interpretations remain concerning the precise nature of the information processing undertaken in the hippocampus and associated structures (e.g. Cohen & Eichenbaum 1991; Rawlins *et al*. 1991), there seems less controversy about the neuroanatomical connections that form the network in which the hippocampus is embedded. This anatomical circuitry has been comprehensively reviewed (e.g. Amaral & Witter 1989, 1995) and is generally considered to consist of connections between the constituent parts of the limbic cortex, including the hippocampal formation and limbic areas of the periarchicortex, such as the prelimbic, infralimbic, cingulate, retrosplenial, perirhinal, entorhinal and subicular cortices (Lopes da Silva *et al*. 1990). Despite the wealth of information on the connectivity of these structures, however, many of the theoretical models advanced to explain function in this system employ greatly simplified treatments of its complex connectivity. We were interested in whether important insights might be lost by simpler treatments of the organization of this system, particularly if the simplifications were derived by an arbitrary process. Also, attempts to use the rich primary information about connectivity in this system to anticipate the location of interesting neurophysiological features have not been universally successful. For example, even

†Present address: USC Brain Project, University of Southern California, Hedco Neuroscience Building, Los Angeles, CA 90007, USA.

though the prelimbic cortex of the rat (PL) receives a direct connection from CA1, electrophysiological experiments that attempted to find so-called 'head-direction' or 'place' cells in PL report a null result (Poucet 1997; Jung *et al*. 1998). We were interested in whether insights into the organization of the system from actual analysis of the system's connectivity could be derived, and whether predictions from these analyses might fare better. Accordingly, to try to derive information about the organization of the system from available neuroanatomical connection reports from the literature, we collated detailed neuroanatomical information about the connectivity of the hippocampal system and analysed the resulting data. Analyses of connectivity based on a variety of different analytical techniques have recently been used to elucidate the global organization of corticocortical systems in the macaque (Young 1992, 1993; Young *et al*. 1995; Hilgetag *et al*. 1996; Stephan, Hilgetag, Burns, O'Neill, Young & Kötter, this issue) and the corticocortical systems in the cat (Scannell & Young 1993; Scannell *et al*. 1995). Analyses of this kind investigate the large number of anatomical constraints in a system with the aim of placing the functional properties of individual structures within a wider organizational scheme. Such an analysis allowed Scannell *et al*. (1995) to successfully predict that plaid-pattern selective cells could be found in the anterior ectosylvian sulcus of the cat cerebral cortex (Scannell *et al*. 1995, 1996), and we hoped that an analysis of the connectivity of the rat hippocampal system might be similarly revealing.

For the present analyses, we considered a system made up of 24 structures that are widely believed to be implicated in neural processing that underlies spatial navigation (Redish & Touretzky 1997; Neave *et al*. 1996). We included the hippocampal formation and associated limbic cortex, as described above. The anterior nuclei of the thalamus contain cells that fire preferentially when the animal's head is pointing in a specific direction, and are also known as 'head-direction cells' (Blair & Sharp 1995; Taube 1995*a*; Mizumori & Williams 1993). Part of the subcortical system described by Redish & Touretzky (1997) includes the mammillary bodies, which receive input from the subiculum (Shibata 1989). The mediodorsal nucleus of the thalamus was included on the basis of its strong interconnections with parts of the limbic cortex (Groenewegen 1988). No septal nuclei were included, despite their involvement in the organizational schemes described by Gray (1982) and their possible role in the generation of theta rhythms (Buszaki *et al*. 1994). The septum receives a strong input from the hippocampal formation (Meibach & Siegel 1977; Swanson & Cowan 1977), and it has been described as 'a conspicuous, integrated part of the limbic system' (Jakab & Leranth 1995). However, it does not appear in descriptions of systems concerned with spatial navigation (e.g. Redish & Touretzky 1997; Neave *et al*. 1996), and so it was omitted from the present study.

We compiled a comprehensive computational database of the connectivity literature describing the connections between these brain structures. We then extended the non-metric multidimensional scaling (NMDS) method used in previous analyses (e.g. Young *et al*. 1995) to analyse it. Detailed simulation studies (Burns 1998; Young *et al*. 1995) have shown that recovery of the variability in test data by NMDS is generally good, but that it can be compromised most by deriving solutions in numbers of dimensions that are much lower than those implied by the structure of the data (Burns 1998). Accordingly, we employed non-parametric cluster analysis (NPCA) to examine NMDS results in larger numbers of dimensions than can be apprehended unaided. We applied this two-component strategy to our summary matrix of connection data. This process produced organizational schemes that were broadly consistent with known physiological properties and may provide the basis for making tentative predictions, alongside existing neurophysiological data, concerning the functional properties of the constituent structures of this system.

2. METHODS

(a) *A neuroanatomical connection database*

As the first stage of our investigations of central connectional organization in the rat, we designed a relational database using Microsoft Access 7.0 to store and manipulate individual reports of connections. These reports were taken from the abstracts, introductions, results sections and conclusions of neuroanatomical research papers. We entered a total of 14 000 connection reports into this database to provide a fairly inclusive description of the rat connectivity literature. Connection reports meeting the criterion that they involved one or more of the structures identified above numbered more than 900 separate connection reports (see § 3). This extract from the database is available for downloading from (http://www.flash.ncl.ac.uk/ptrs/rathippo.htm). Each and every datum in the database can be substantiated by other researchers by reference to the contents with the actual report in the literature from which the datum derives, since the page and figure number of the connection report are stored in the database.

(b) *Connection matrices, transformations, similarity matrices and the proximity model*

All the methods of analysis we applied extend the concept of applying proximity-based analyses of similarity (such as NMDS) to neural connection data. We discuss the shortcomings and benefits of this paradigm in § 4(a).

Input data are contained in a connection matrix, $\hat{C}$, where each entry in the matrix, c_{ij} represents the 'strength' of a neural connection, an ordinal measure related to the number of neurons participating in the connection, from the ith structure in the system to the jth. Four similarity matrices were calculated from this connection matrix. The first matrix, $\hat{N}$, was obtained by coding the connections with similarity values according to the following scheme. Strong connections were coded with ordinal similarity value of 3, moderate connections with 2, sparse connections with 1 and connections that had been found absent with 0. Connections that had not been identified were assumed to be missing and were assigned a similarity value of 0 (see Young 1992; Young *et al*. 1995). Connections that were reported to exist, but with unspecified strength were assigned similarity values of 1 (i.e. they were considered equivalent to 'weak' connections for the purposes of these analyses). This matrix was symmetrized (that is, it was transposed and added to itself) to give the matrix $\hat{T}$. A third matrix was obtained by acting on $\hat{N}$ with the *pth1* transform (Young *et al*. 1995) to obtain a matrix $\hat{P}$. A fourth matrix was obtained by acting on $\hat{N}$ with

the *wdsm1* transform (Young *et al.* 1995) to give a matrix $\hat{W}$. These transforms generate additional ranks within the data on the basis of certain assumptions: *pth1* interprets chains of similarities to differentiate between dissimilar objects (i.e. if three objects A, B and C have similarity values of 0 in all combinations, then if A and B are likely to be more similar than A and C if A and B are similar to a fourth object, D, while C is not); *wdsm1* is similar to the Czekowski coefficient which is widely used to generate Euclidean representations from binary data (Gower & Legendre 1986; Cox & Cox 1995).

(c) *NMDS, Procrustes analysis and non-parametric cluster analysis*

NMDS is used to generate representations of similarity data as configurations of spatially distributed points where interpoint distances represent the relationships between the objects being studied (see Shepard 1962; Kruskal 1964; Cox & Cox 1995). As with many data analysis methods, a drawback of NMDS is that it is possible to alter the shape of output configurations by changing the internal parameters used in the analysis, without altering any of the input data. Thus, in general, a given connection matrix only produces a unique NMDS solution at a given dimensionality with the aid of assumptions. This consideration emphasizes the importance of using a reasoned basis for selecting the parameters for an analysis. Such a basis was provided at length in Young *et al.* (1995*b*) and followed from close examination of the properties of the data. We further addressed this issue by generating many alternative solutions with, for example, different dimensionalities and cost functions, and then re-examining the output configurations with NPCA to identify the most robust features of data structure.

All published studies of neural connectivity using NMDS have presented only two- or three-dimensional (3D) plots (e.g. Young 1992, 1993; Simmen *et al.* 1994; Young *et al.* 1994, 1995*a*,*b*; Goodhill *et al.* 1995; Scannell *et al.* 1995). As described previously (e.g. Young 1992; Young *et al.* 1995), however, plots in a small number of dimensions may misrepresent data structure. This misrepresentation will be clearly manifest as long lines in the plot, reflecting structures that are 'close' in a higher number of dimensions but which are separated in the low-dimensional plot. This problem could be abolished by deriving configurations in a higher number of dimensions, but these suffer the problem that they cannot be interpreted by visual inspection. A method for interpreting the relationships apparent in configurations with a more realistic number of dimensions would mitigate these problems, and we now describe how non-parametric cluster analysis can be used to interpret configurations of higher dimensionality.

Cluster analysis is a method of classification (Gordon 1981). Given a set of n objects, cluster analysis seeks to partition the set into clusters, so that members of the same cluster have similar properties. The output of the NMDS method naturally lends itself to this approach, since the criteria used to group objects together can be represented by the distances between the objects' points in the configuration. Within this framework, a cluster can be defined as a local maximum of the point density in the space inhabited by the configuration.

Non-parametric density estimation provides a means of modelling the density function without making assumptions concerning the form of the function. Methods of cluster analysis based on non-parametric density estimation can detect clusters of unequal size and dispersion, or which have irregular shapes. The most widely used method of density estimation for multivariate data is the kernel method, in which estimates are based on density that has been sampled for a small region of multidimensional space (usually a hypersphere of a specified radius, called a 'kernel'; Silvermann 1986; Scott 1992). Two types of kernel were used in these analyses: one was fixed, so that all kernels had identical radii. The other was based on the 'kth nearest-neighbour' approach where the radius of the kernel centred at each point in the configuration was defined as the minimum distance required to enclose the closest $k-1$ points.

We used the MDS and MODECLUS functions from the SAS 6.09 statistics software to perform NMDS calculations and NPCA with significance testing. We used the ROTATE function in the GENSTAT statistical software package to perform Procrustes rotations. We used Perl 5 scripts to automate the execution of these functions and the generation of output graphics. The MODECLUS procedure was used with the JOIN option to test the significance of clusters by comparing the maximum estimated density of points within a cluster to the maximum around the cluster's border (the 'saddle point') in order to estimate the cluster's significance. Under the JOIN option, MODECLUS produced a hierarchical scheme (called a 'cluster tree') where clusters were sequentially dissolved in ascending order of significance. The points of a dissolved cluster were left unassigned if there were no points from neighbouring clusters within a single kernel radius. If the clusters were sequentially joined to provide a tree, this gave an indication of the relative proximity of separate clusters.

Beginning with the connection matrix in figure 1, we generated two- and five-dimensional (2D and 5D) NMDS configurations with the FIT variable of the MDS routine set to 1, 2 and 0.5. This was performed with both the primary and secondary approach to ties (Young *et al.* 1995). This produced 12 output configurations for each input matrix. We used five dimensions because the accuracy of density estimation falls with increased dimensionality, requiring very large numbers of observations at high dimensions (Epanechnikov 1969; Silvermann 1986). Configurations with five dimensions were selected because of good performance in trials of this method with test data.

We did not wish to prejudge the nature of the clusters in the data. To examine the cluster structure in as unbiased a way as possible, we ran many analyses with a wide variety of different clustering parameters, and in this way sought the most consistent features that occurred in the output clustering schemes. Hence, ten separate MODECLUS analyses were run on each configuration according to four different paradigms. The first paradigm simply ran 30 separate cluster analyses with increasing fixed-radius density estimation and clustering kernels. The radius of this kernel ranged from the minimum inter-point distance to the mean inter-point distance in the configuration.

The next three analyses used a fixed-radius kernel that was calculated from the cluster analysis under the first paradigm. The kernel radius was chosen to produce an initial number of clusters which was equal to the total number of points in the analysis divided by two, four and eight, respectively. A hierarchical cluster tree was obtained for each scheme by testing the significance of these clusters with the JOIN option. The next three analyses were based on nearest-neighbour kernels with the density estimation and clustering kernels set so that the number of neighbours of each point was a minimum of one, two and three, respectively. The last three analyses used nearest-neighbour clustering methods and fixed kernel density estimation with the JOIN option to give hierarchically organized schemes.

Table 1. *Connection matrix for central systems involved in spatial memory*

(See appendices for each individual connection report. The matrix entries have the following meanings: 3, strong connection; 2, moderate connection; 1, weak connection; 0, connection reported as absent; c, connection of unspecified strength; X, connection cannot exist; x, connection reported in abstract of paper.)

	CA1	CA3	DG	ENT	PAR	POST	PRE	SUB	LM	MM	SUM	TM	ACA	ILA	PL	PRh	RSP	AD	AM	AV	IAM	LD	MD
CA1	X		c	3	2	2	c	3					0	c	3	1	1						
CA3	3	X	c	1	1		c	c					0		0	0	x						
DG	c	c	X			c	c	c					0		0								
ENT	3		3	X	2	1	c	c	0	2			1		1	2	1						1
PAR	1	0	1	3	X	2	3	1	3	0	0		0				3	3		1		0	
POST				2	2	X	2		2				0			3	3	3		1		3	
PRE	0	0	1	3	3	2	X	2	3	2	0	0				0	2	2		2		3	
SUB	2	2	2	3	2	c	3	X	2	3	1	2	1	1	3	1	3		2	c	1		
LM									X									3	0				
MM										X						c		2	3	3	3		0
SUM	2	3	3	3	1	c	1	2	2	1	X				1	1		1					
TM									1	1		X				2							
ACA				1	0	2	2		2	2			X	c	c	2	3	0	3	1		2	3
ILA				0				0	2	3	3	3	1	X	c	2	2	0	2	2	2		2
PL				2			1		2	2	1	1	3	2	X	2	2		2	2	2	1	3
PRh	c	0	0		c	2		c							2	X	1						
RSP				2	2	3	2	0		2			3	0	0	x	X	2	3	3		3	
AD	c	c		2	2	3	3				c		3	0	0	c	3	X					
AM		c		3	2		1	2					3	2	2	2	3		X				
AV	c	c		1	2	2	3	2			c		3	0	0	1	3			X			
IAM				2									2			3	2				X		
LD	c	c		1	3	3	3						3		2		3					X	
MD				3				0					3	3	3	1	1						X

In this way, 60 cluster trees were generated for each connection matrix at each dimensionality. Seven out of every ten of these trees were hierarchically organized, and were made up of several individual cluster schemes. In a given cluster scheme, each area would either be assigned to only one cluster, or would be unassigned. We defined an $n \times n$ 'cluster-count' matrix, K for each cluster tree. Each matrix element, k_{ij}, denoted the number of times the ith and jth areas were assigned to the same cluster in the cluster tree, divided by the total number of different cluster schemes in the cluster tree. All cluster counts of a given configuration were averaged to yield an overall cluster-count value that indicated the most consistent features of the cluster analyses. Cluster counts were averaged over configurations that had been derived from specified transforms, or from a specified dimensionality, to give transform-specific, and dimension-specific cluster counts. Within these averaging processes, the cluster counts of each cluster tree were weighted equally, so that each paradigm contributed equally to the overall cluster-count score. The cluster count took the significance of clusters into account. If a brain area was in a cluster that was subsequently dissolved, the brain area would not contribute further to the cluster count (i.e. it would not be counted as being in the same cluster as any other structures, including itself). Hence, some analyses describe cluster counts between an area and itself of less than 100%.

The order of structures within each cluster-count matrix was selected so that a given brain area in the matrix would be followed by the brain area which, when paired with the first area, had the largest cluster count of those remaining. This procedure allowed easier interpretation of the cluster-count matrices by grouping areas with high cluster counts together. We used 20 grey levels to shade individual cells in the matrix, where the lightest shade was set to the minimum cluster-count value and the darkest shade was set to the maximum cluster-count value. To facilitate the interpretation of these cluster-count figures we selected three thresholds to classify groupings of brain areas into 'strongly clustered', 'moderately clustered' and 'weakly clustered' sets. The criteria for inclusion into each such set were that cluster-count values between two areas were in the top seven categories (sharing clusters in around 70–100% of cluster trees) for strongly clustered sets. The next seven categories (sharing clusters in around 30–70% of cluster trees) formed moderately clustered sets and the next three categories (sharing clusters in around 20–30% of cluster trees) were classified as being weakly clustered.

As a final step, we superimposed the set structure of each summary cluster-count matrix on to a 2D NMDS configuration as a Venn diagram consisting of the clusters determined by cluster analysis. The Venn diagrams provided a way of combining the structure derived from the cluster analysis of higher-dimensional configurations with configurations produced in an interpretable number of dimensions.

3. RESULTS

The collated connection data we analysed were derived from 89 papers and 933 separate connection reports. Table 1 summarizes these data in a connection matrix. We derived four similarity matrices from this connection matrix, according to the methods described above ($\hat{N}$, $\hat{T}$, $\hat{P}$ and $\hat{W}$). We then analysed these matrices with NMDS, Procrustes rotations and NPCA.

Table 2. *Mean Procrustes R^2 statistics for the NMDS configurations*

(Each mean value is calculated from 36 individual Procrustes rotations between configurations generated with different combinations of cost function and tied or untied approaches. The mean values along the leading diagonal were calculated from 30 separate R^2-values since configurations were not rotated against themselves.)

		2D configurations					5D configurations				
		N_3	T_3	P_3	W_3	total	N_3	T_3	P_3	W_3	total
2D configurations	N_3	0.79	0.80	0.53	0.81	0.73	0.40	0.42	0.37	0.60	0.45
	T_3	0.80	0.87	0.57	0.81	0.76	0.39	0.41	0.39	0.60	0.45
	P_3	0.53	0.57	0.82	0.60	0.62	0.32	0.30	0.37	0.42	0.35
	W_3	0.81	0.81	0.60	0.99	0.79	0.39	0.40	0.36	0.66	0.45
	total	0.73	0.76	0.62	0.79	0.72	0.38	0.38	0.37	0.57	0.43
5D configurations	N_3	0.40	0.39	0.32	0.39	0.38	0.86	0.86	0.68	0.63	0.75
	T_3	0.42	0.41	0.30	0.40	0.38	0.86	0.94	0.67	0.67	0.78
	P_3	0.37	0.39	0.37	0.36	0.37	0.68	0.67	0.87	0.55	0.68
	W_3	0.60	0.60	0.42	0.66	0.57	0.63	0.67	0.55	0.98	0.69
	total	0.45	0.45	0.35	0.45	0.43	0.75	0.78	0.68	0.60	0.73

(a) *NMDS and Procrustes analysis results*

The Procrustes R^2-values describing the similarities between configurations derived with the different parameters are shown in table 2. The modest overall mean Procrustes R^2-score of 0.43 between the 2D and 5D configurations suggests that there are substantial differences between them. The most likely explanation is that there are aspects of data structure that are poorly reflected in the 2D configurations, and that the methods for interpreting higher-dimensional configurations set out above will be valuable for this system. 2D configurations derived from the FIT = 1 tied NMDS analyses of $\hat{N}$, $\hat{T}$, $\hat{P}$ and $\hat{W}$ (with connections represented as lines) are shown in figures 1 and 2. Comparisons between configurations derived under each combination of cost function and tied or untied conditions are shown in figures 3 and 4.

The 2D configurations representing this system have a consistent structure between different similarity matrices, and under different stress minimization paradigms. While the configurations produced by analysis of the transformed matrices $\hat{P}$ and $\hat{W}$ appear less consistent than the configurations produced by analysis of $\hat{N}$ and $\hat{T}$, some organizational features are immediately apparent from an inspection of figures 1–4.

We exemplarize the features by reference to figure 1. The elements of the hippocampus (CA1, CA3 and DG) lie at the edge of the figure, with CA1 lying closer to the main body of the configuration than the other two points. Their immediate neighbours below and to the right are the elements of the retrohippocampal region (ENT, SUB, PAR, POST, PRE), and their neighbours above and to the left are the supramammillary and lateral mammillary nuclei (SUM, LM). The retrohippocampal areas form a tight clump that includes the anterodorsal nucleus of the thalamus. The tuberomammillary nucleus (TM) lies on the edge of the configuration to the top left-hand corner of the figure. The connection matrix in figure 1 shows TM to have three efferent and three afferent connections, suggesting that the separation in the figure is due to the differences in its pattern of connections to other areas in the structure, rather than 'pop out' due to sparseness (Young *et al.* 1994). This separation of TM is strongly emphasized in the configuration derived from $\hat{W}$.

The mediodorsal nucleus of the thalamus is situated on the edge of the configuration to the mid-to-bottom left. Just to the right of this lie the infralimbic, prelimbic and perirhinal cortices (ILA, PL, PRh), grouped quite closely together, about halfway up the figure to the left-hand side. Somewhat below them, the anterior cingulate (ACA) and retrosplenial cortex (RSP) appear, and between them the medial mammillary (MM) nucleus lies separately from the other mammillary nuclei.

The anterior nuclei of the thalamus (AD, AM, AV, IAM, LD) are scattered over the bottom half of the configuration. AM and AV are close to ACA and RSP; AD is strongly associated with the retrohippocampal regions; LD lies slightly separated to the bottom of the figure; and the interanteromedial nucleus of the thalamus (IAM) lies close to the medial dorsal nucleus (MD) on the left-hand side of the configuration.

The configurations produced by NMDS analysis of the $\hat{P}$ similarity matrix appear to be qualitatively different from configurations derived from the other matrices. The main differences are that CA1 appears separate from CA3 and DG; and that the position of ENT shifts from halfway up on the right-hand side (close to SUB and PAR) to midway across the bottom half of the configuration (close to ACA and RSP). Other features consistent in other configurations also changed. For example, the tight grouping of PAR, PRE, POST and AD became a line of points separating CA1 from CA3 and DG. We return later to the possible reasons for this, but the most likely reasons are that the similarity transform is based on different numbers of steps between structures (Burns 1998; Young *et al.* 1995) and that in this quite highly connected system there is insufficient variability in journey length for the similarity measure to produce reliable information.

Figures 4 and 5 provide an impression of the robustness of the analyses to changes of cost function and approach to ties. The variability of the position of LM between cost functions is the most pronounced feature of these comparisons, suggesting that its position is less constrained in

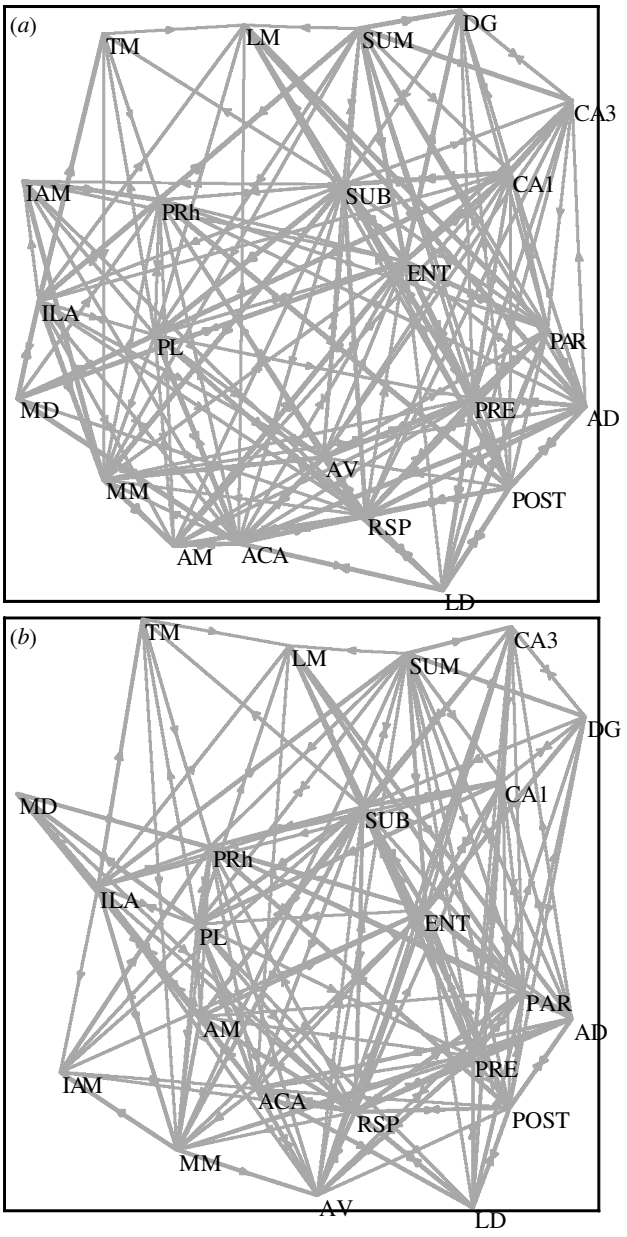

Figure 1. NMDS output configuration produced by NMDS analysis of $\hat{N}$ (*a*) and $\hat{T}$ (*b*), under the FIT = 1 tied cost function.

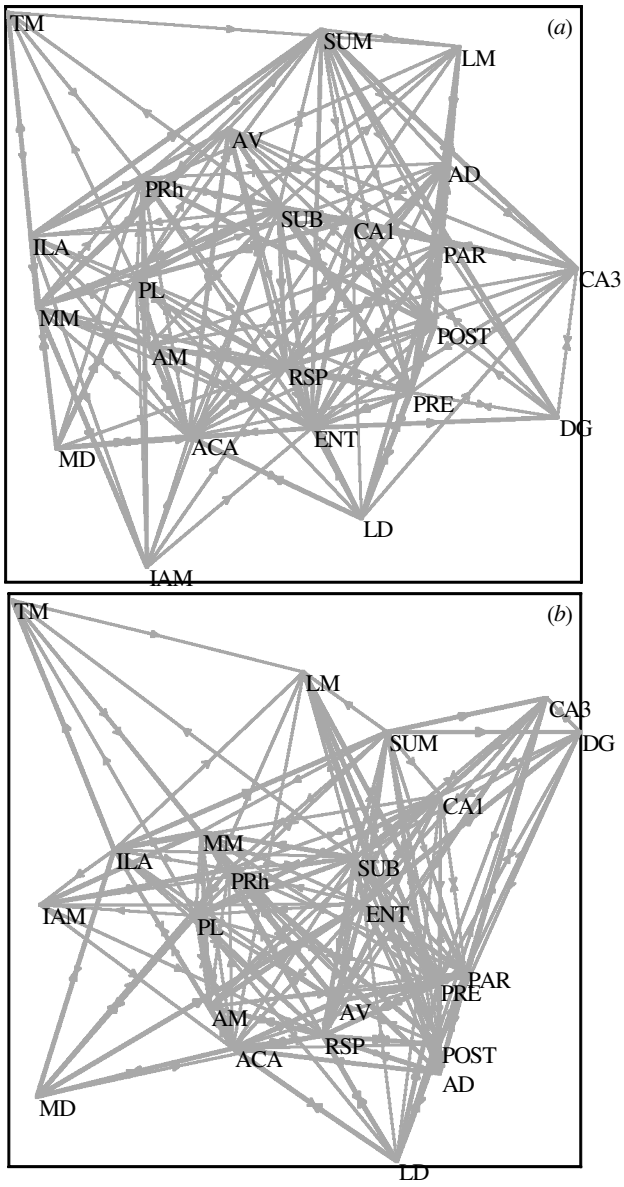

Figure 2. NMDS output configuration produced by NMDS analysis of $\hat{P}$ (*a*) and $\hat{W}$ (*b*) under the FIT = 1 tied cost function.

these 2D configurations than those of its neighbours. In fact, the position of LM is the most variable in every case, except for configurations calculated from $\hat{P}$, where it is the second most variable. Other areas that shift under different combinations of cost function and tied condition are CA3, IAM, MM and MD. In the configurations generated by $\hat{N}$, $\hat{T}$ and $\hat{P}$, the positions of individual areas overlap somewhat with their nearest neighbours, but generally do not exceed this limit. The configurations produced by NMDS analysis of $\hat{W}$ have no overlap at all, suggesting that the solutions derived from the data transform are quite robust to differences of cost function and tying–untying.

(b) *Non-parametric cluster analysis*

We submitted the output configurations of the NMDS analyses to the cluster analyses described in § 2. Forty-eight different configurations were produced in the previous section (four matrices, three cost functions, two approaches to tied data, in 2D and 5D space), which, when analysed, produced 480 separate cluster trees. The number of separate schemes in a given cluster tree depends on the parameters of the analysis, but was typically less than ten. The results are presented as cluster-count matrices.

We first describe the most general results and then describe the way in which individual analyses differ from them. Figure 5 shows the cluster-count matrix taken for all cluster analyses using $\hat{N}$, $\hat{T}$, $\hat{P}$ and $\hat{W}$.

We superimposed the cluster structure from figure 6 on to the 2D NMDS configuration of $\hat{N}$ under the FIT = 1 tied condition as a Venn diagram cartoon in figure 7. The Venn diagrams delineate sets of nuclei that share the same cluster in a proportion of cluster trees corresponding to groups of cluster-count cells, which are darker in figure 6. In this way, the diagram shows information derived from higher-dimensional, better-fitting solutions in a low-

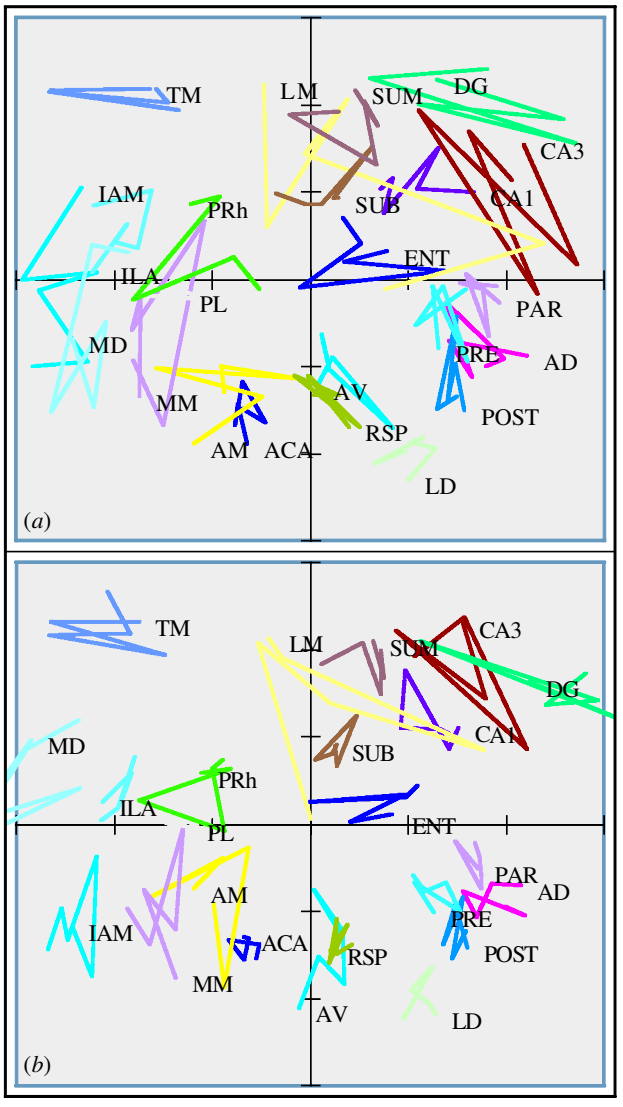

Figure 3. Six output configurations produced by NMDS analysis of $\hat{N}$ (*a*) and $\hat{T}$ (*b*), under all cost functions.

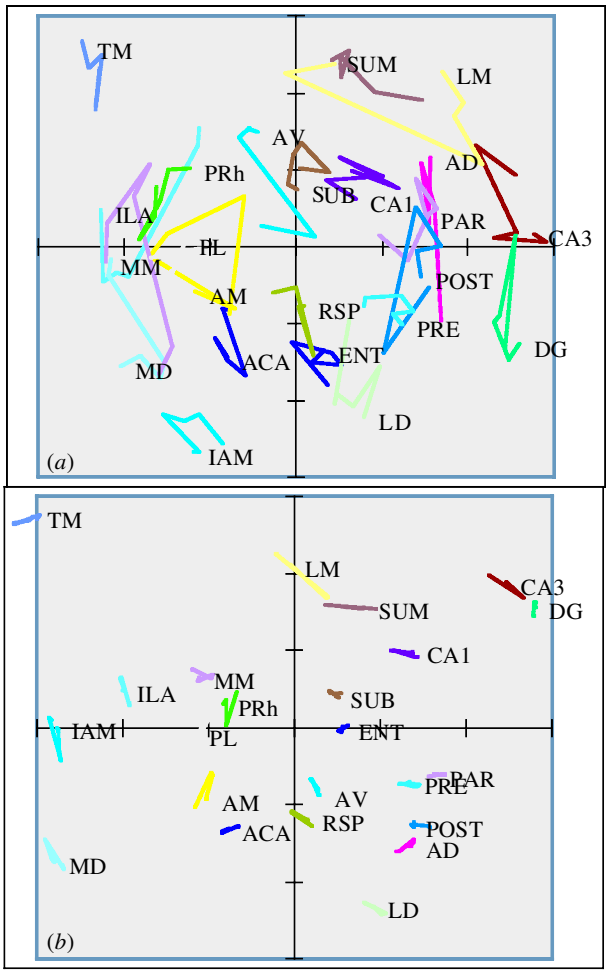

Figure 4. Six output configurations produced by NMDS analysis of $\hat{P}$ (*a*) and $\hat{W}$ (*b*) under all cost functions.

dimensional form. The sets delineated by black lines represent very consistent clusters (>70%), whereas those delineated by dark grey lines contain areas that shared the same cluster in 39–67% of cluster trees. Those surrounded by light grey lines contain structures that share the same cluster in 27–39% of cluster trees. These sets are the strongly, moderately and weakly clustered sets, respectively. The dotted lines help to differentiate between sets when they overlap. These sets correspond to groupings derived from the cluster analysis, but the thresholds that determined the inclusion of areas into these summary sets were arbitrary.

At the broadest level, there are three overlapping, weakly clustered sets. Two of these sets are distinct from each other, and the third overlaps both of them. Between them, the two non-overlapping sets involved most of the nuclei and areas studied. One of these sets includes most of the nuclei and areas in the top half of the cluster-count matrix in figure 6, and this set is situated on the right-hand side of figure 7. This weakly clustered set contains one moderately clustered set and a strongly clustered set. The moderately clustered set contains the parts of the hippocampus (CA1, CA3 and DG), the subiculum (SUB), the entorhinal cortex (ENT) and the SUM. The strongly clustered set contains the presubiculum, the parasubiculum, the postsubiculum and the anterodorsal nucleus of the thalamus (PRE, PAR, POST and AD). This strongly clustered set is one of the most consistent features of all the analyses. In the analysis of $\hat{N}$, the LM is also part of the moderately clustered set.

The second weakly clustered set contains all the members of the first weakly clustered set, except the parts of the hippocampus (CA1, CA3 and DG) and the SUM. The nuclei involved in the strongly clustered set from the first cluster (PRE, PAR, POST and AD), are part of a moderately clustered set that also includes the LM and the anteroventral nucleus of the thalamus (AV).

The 2D structure of the NMDS configuration of $\hat{N}$ in figure 2 appears to be quite consistent with the cluster-count sets that are superimposed on to it. There are relatively few sets with components that are widely separated. The one notable exception is the inclusion of LM into the moderately clustered set containing PAR, PRE, AD, POST and AV described above. This may indicate that some 2D configurations placed LM in a less peripheral position (see figures 4 and 5) and that this feature is preserved in the 5D configurations.

There are a series of partially overlapping, moderately clustered sets at the bottom of the figure. The second

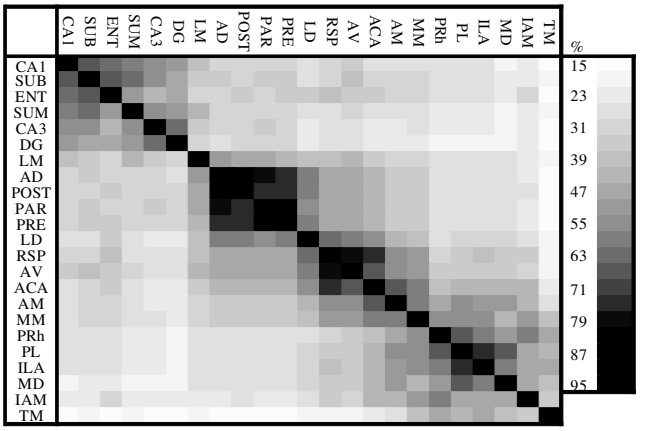

Figure 5. Cluster-count matrix derived from all 480 cluster trees from $\hat{N}$, $\hat{T}$, $\hat{P}$ and $\hat{W}$ involving 2607 cluster schemes.

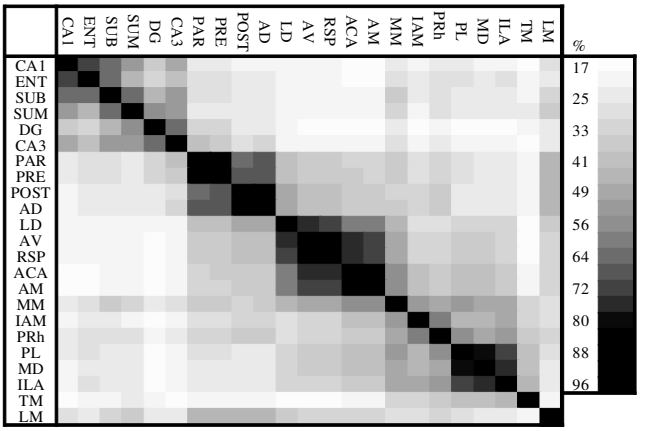

Figure 7. Cluster-count matrix calculated from non-parametric cluster analysis of 2D and 5D NMDS configurations of $\hat{N}$.

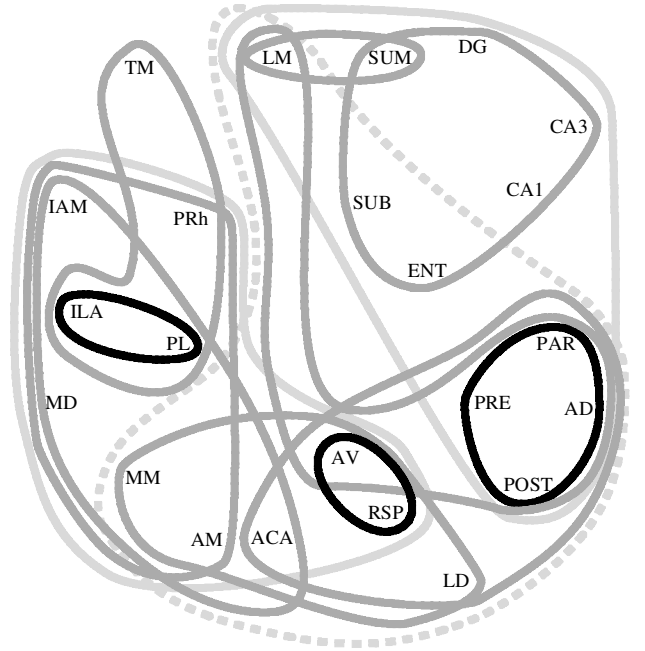

Figure 6. Venn diagram illustrating cluster-count data from figure 5 superimposed onto the NMDS configuration derived from FIT = 1 tied analysis of $\hat{N}$.

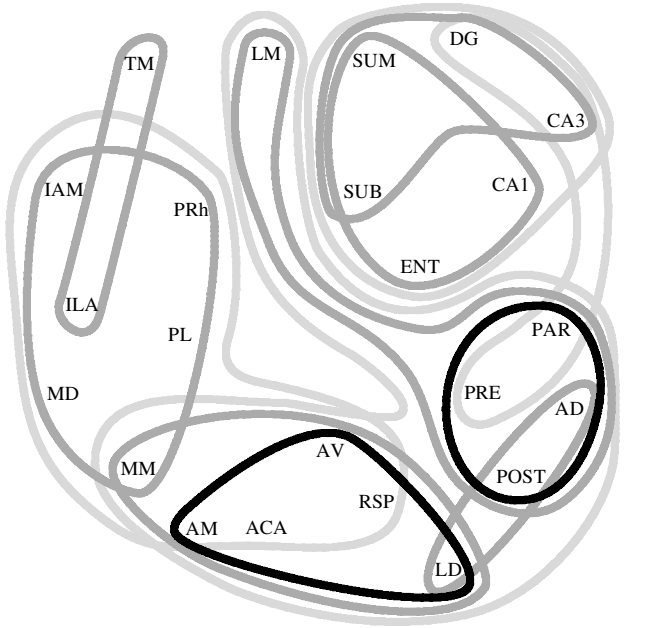

Figure 8. Venn diagram of cluster counts from figure 7 superimposed onto FIT = 1 tied NMDS configuration produced by analysis of $\hat{N}$.

moderately clustered set contains the strongly clustered set from the first weakly clustered set (PRE, PAR, POST and AD), the laterodorsal thalamic nucleus (LD), the AV, the RSP and the ACA.

The third moderately clustered set contains the AV, the RSP, the ACA, the LD, the anteromedial nucleus of the thalamus (AM), and the MM. To make the Venn diagrams more interpretable we reduced the number of overlapping lines. In this case, we included both LD and MM in this set, despite the fact that they were included in the same cluster in less than 39% of the analyses. A strongly clustered set contained AV and RSP.

The third weakly clustered set contains all of the areas and nuclei in the last set described except for LD (RSP, AV, ACA, AM, MM), as well as the MD, the ILA and PL, the PRh and the IAM. There are two overlapping moderately clustered sets which are wholly contained within this weakly clustered set. They both contain AM, MM, MD, ILA, PL and IAM, and only differ in that one contains PRh and not ACA, and the other contains ACA and not PRh. The TM is the only structure that does not lie within one of the weakly clustered sets, but is contained in a moderately clustered set with the PRh, ILA and PL.

There are several interesting features of this analysis that merit further description. The PRh and the ENT are associated with wholly different structures, even though they neighbour each other and appear relatively close on many of the NMDS configurations. Lesion studies that affect this area often damage both the ENT and PRh simultaneously (e.g. Rothblat *et al.* 1993). If these two areas are involved in different connectional systems, then these lesions may affect a wider range of information processing. Other, similarly surprising features include the separation between the MM, the LM and the SUM.

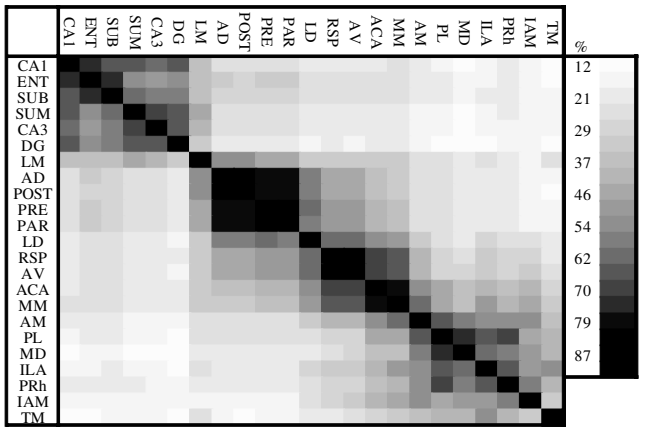

Figure 9. Cluster-count matrix calculated from non-parametric cluster analysis of 2D and 5D NMDS configurations of $\hat{T}$.

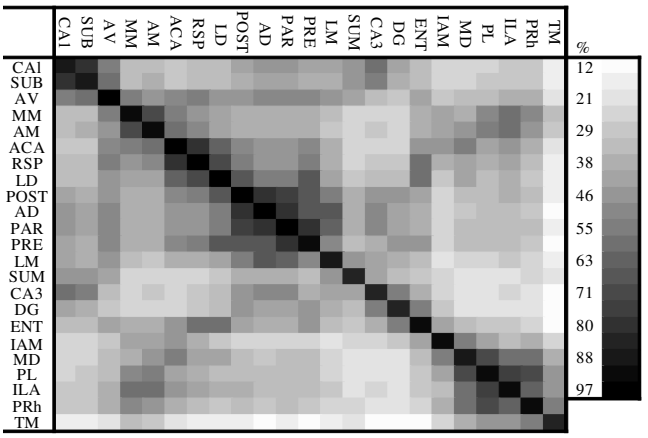

Figure 11. Cluster-count matrix calculated from non-parametric cluster analysis of 2D and 5D NMDS configurations of $\hat{P}$.

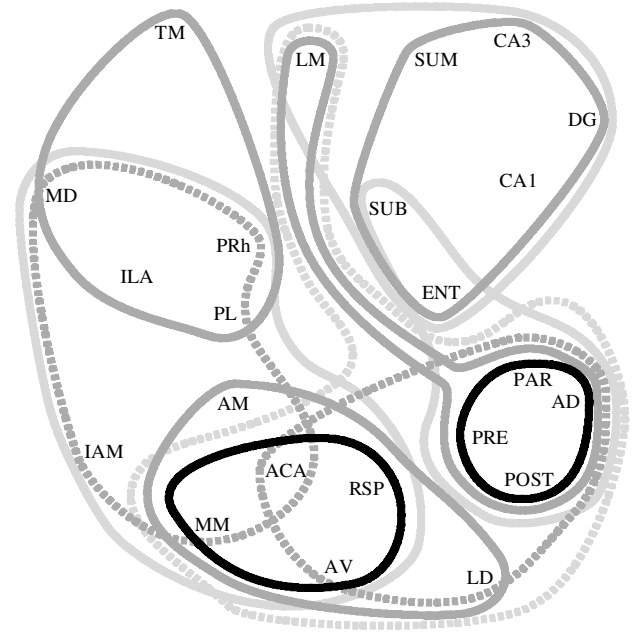

Figure 10. Venn diagrams of cluster counts from figure 9 superimposed onto FIT = 1 tied NMDS configuration produced by analysis of $\hat{T}$.

(c) *A comparison between cluster analyses of different similarity matrices*

The results in the previous section describe cluster-counts that were derived from averages over all analyses of all the different similarity matrices. This section considers cluster-count matrices that represent all the cluster analyses derived for each similarity matrix in turn. Noticing that the Procrustes R^2-value falls to low values for comparisons between certain configurations, it is important to establish whether NMDS configurations from different matrices have consistent cluster structures.

Figures 7 and 8 illustrate the cluster-count matrices and Venn diagrams of the cluster counts, averaged over all solutions of $\hat{N}$. Figures 9 and 10 describe the cluster counts of $\hat{N}$. The overall appearance of these figures is strikingly similar to those of the global analyses, in terms of the number and general shape of the sets.

The most consistent features of these analyses are that the parts of the hippocampus (CA1, CA3 and DG), the SUB, the SUM and the ENT are separate from other structures; being placed either in a single, moderately clustered set or in overlapping, moderately clustered sets. The AD, the postsubiculum (POST), the presubiculum (PRE) and the parasubiculum (PAR) consistently inhabit a strongly clustered set and share a moderately clustered set with the LM. The remaining nuclei and areas are involved in a series of overlapping, moderately clustered sets, although the sets are more discrete in the $\hat{N}$ cluster counts than those from the $\hat{T}$ analysis. One possible interpretation of these overlapping clusters is that their sequential sets form an anatomically definable pathway, and we examine this hypothesis later. Cluster analysis of the configurations generated from $\hat{P}$ resulted in a cluster-count matrix with many off-diagonal patches, which consequently we could not readily transfer to a Venn diagram (figure 11). This was caused by marked differences between the 2D and 5D configurations, and these are illustrated by superimposing the cluster structure of the 5D cluster-count matrix (figure 12) on to the 2D NMDS solution as a Venn diagram (figure 13). The set structure shown in figure 13 illustrates how the five-dimensional cluster sets derived from $\hat{P}$ differ markedly from the 2D configurations on to which they have been drawn. This can be seen in the way in which the sets are elongated and have to cross over one another without including the same nuclei. Despite this, some of the features that were mentioned before were clearly conserved. AD, PAR, PRE and POST formed a strongly clustered set. The SUB and the parts of the hippocampus (CA1, CA3 and DG) form a moderately clustered set. The exception to this is the ENT, which forms a strongly clustered set with the dentate gyrus (DG), and is only part of a weakly clustered set with the remaining areas of the hippocampus. The PRh, ILA and PL form a strongly clustered set with the MD. The association of AD, PAR, POST, PRE, RSP, ACA, AM and LD in a moderately clustered set that traverses the entire configuration was consistent with the series of overlapping, moderately clustered sets at the bottom right-hand corner of figures 8 and 10. But the inclusion of LM and MM in the same set,

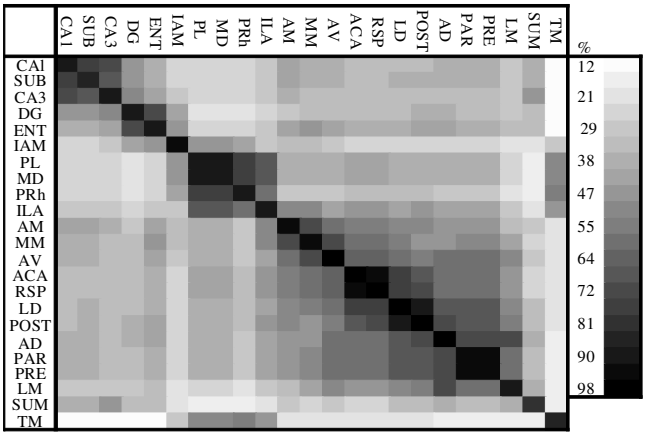

Figure 12. Cluster-count matrix calculated from non-parametric cluster analysis of 5D NMDS configurations of $\hat{P}$.

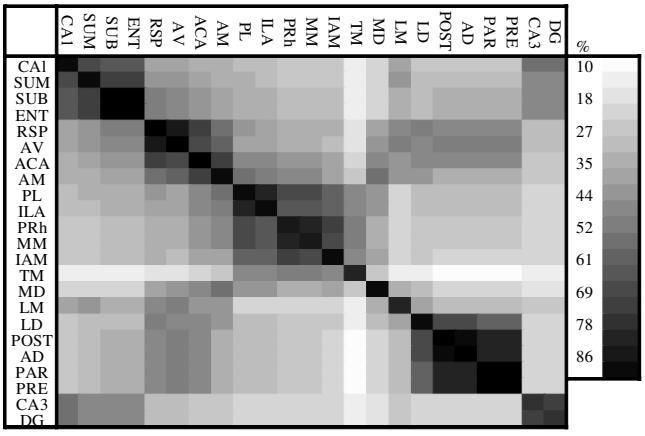

Figure 14. Cluster-count matrix calculated from non-parametric cluster analysis of 2D and 5D NMDS configurations of $\hat{W}$.

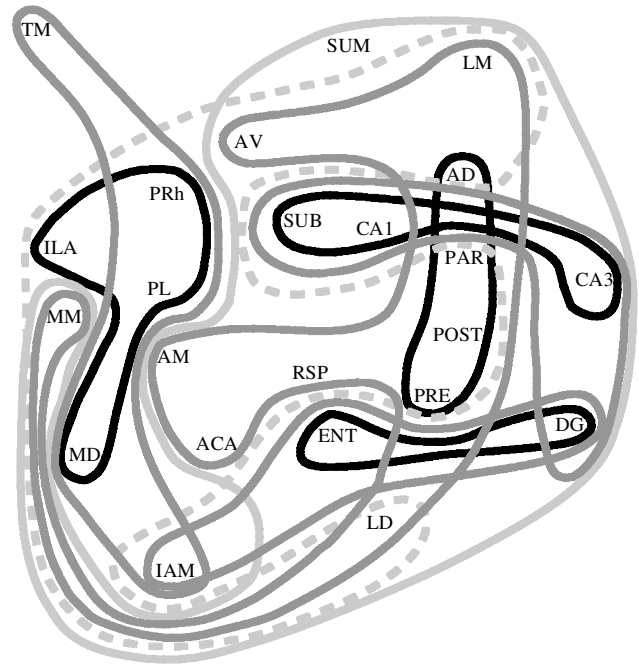

Figure 13. Venn diagrams of cluster counts from figure 12 superimposed onto FIT = 1 tied NMDS configuration produced by analysis of $\hat{P}$.

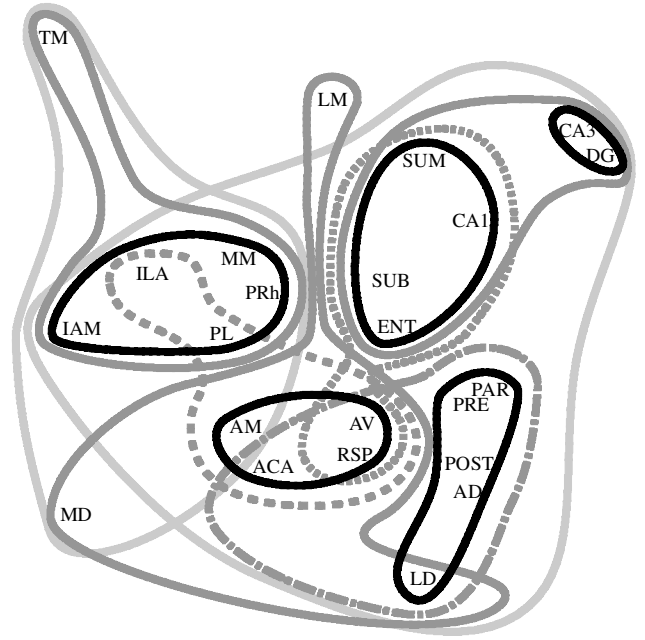

Figure 15. Venn diagrams of cluster counts from figure 14 superimposed onto FIT = 1 tied NMDS configuration produced by analysis of $\hat{W}$.

despite being on different sides of the NMDS configuration, differs from analyses of $\hat{N}$ and $\hat{T}$.

Cluster analyses of the configuration produced by the NMDS analysis of $\hat{W}$ are very consistent, producing clearly defined sets that rarely overlap (see figures 14 and 15). In these analyses, there are five distinct strongly clustered sets which appeared separate from each other, but are grouped together in overlapping, moderately clustered sets. The CA3 area of the hippocampus and DG occupy a strongly clustered set, as did the SUM, the SUB, ENT and the CA1 area of the hippocampus. These two strongly clustered sets are also associated in a moderately clustered set, which is strongly similar to the organization that emerged from analysis of $\hat{N}$ and $\hat{T}$.

The third strongly clustered set occupies a central position in the 2D configuration and contains the RSP, ACA, AM and AV. This set is either wholly or partially involved in four moderately clustered sets, indicating that the *wdsm1* transform designates this set of nuclei and areas to be a central set in this particular system. AV and RSP are involved in a moderately clustered set together with the structures of the second strongly clustered set (SUM, CA1, SUB, ENT). All four structures in the third strongly clustered set are involved with the ILA and PL, and are involved in a moderately clustered set that includes LM, MD and LD.

As before, AD, POST, PRE and PAR are involved in a strongly clustered set, but in this case LD is also included. A moderately clustered set associates these structures with ACA, AV and RSP. The last strongly clustered set involves IAM, ILA, PL, PRh and MM. These structures form a moderately clustered set together with the TM.

Although some features of the connectional organization shown in previous analyses are preserved in this figure, some differences are also apparent. There are no series of partially overlapping clusters, and the strongly clustered set

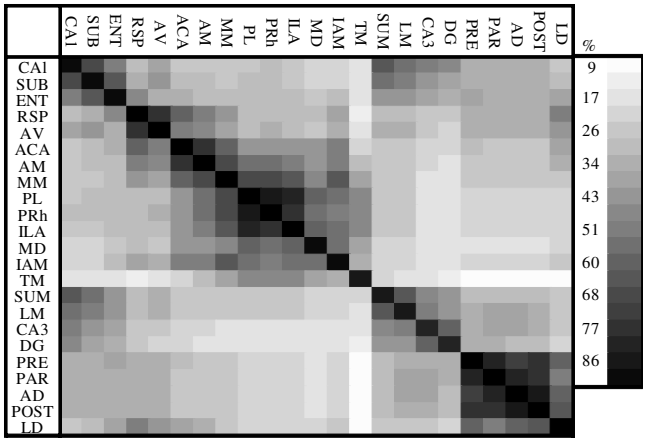

Figure 16. Cluster-count matrix calculated from non-parametric cluster analysis of 2D NMDS configurations of $\hat{N}$, $\hat{T}$, $\hat{P}$ and $\hat{W}$.

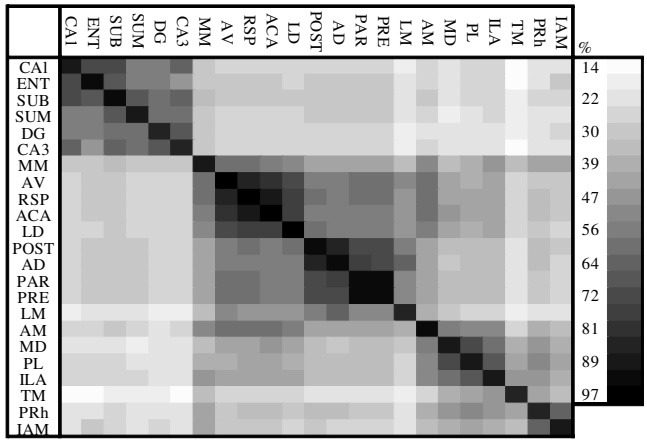

Figure 18. Cluster-count matrix calculated from non-parametric cluster analysis of 5D NMDS configurations of $\hat{N}$, $\hat{T}$, $\hat{P}$ and $\hat{W}$.

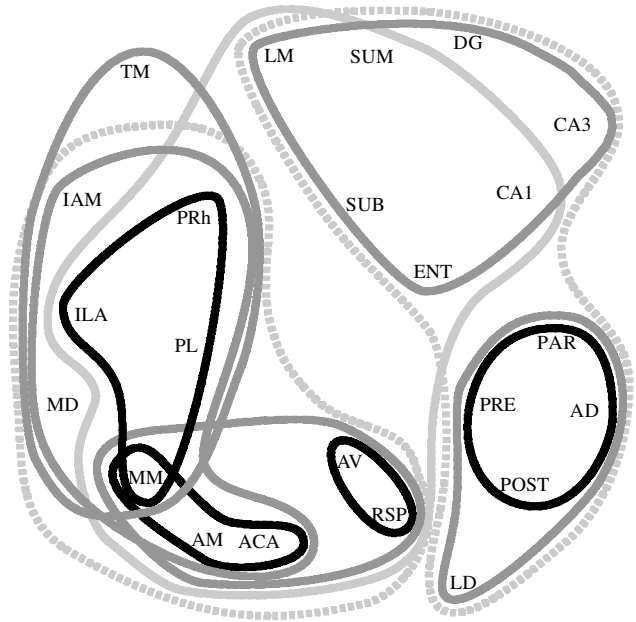

Figure 17. Venn diagrams of cluster counts from figure 16 superimposed onto FIT = 1 tied NMDS configuration produced by analysis of $\hat{N}$.

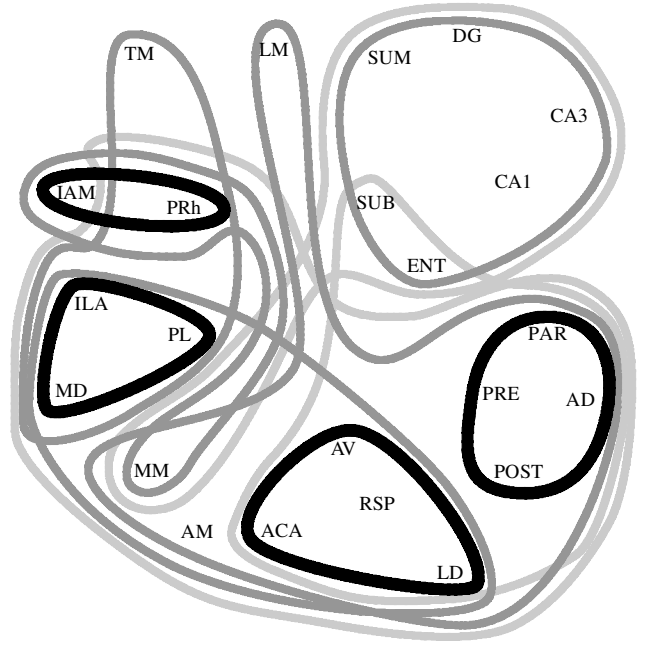

Figure 19. Venn diagrams of cluster counts from figure 18 superimposed onto FIT = 1 tied NMDS configuration produced by analysis of $\hat{N}$.

containing PAR, POST, PRE, AD and LD is not strongly affiliated with the hippocampus proper in these analyses.

(d) *Differences between cluster schemes derived from NMDS configurations with different numbers of dimensions*

We now consider cluster analyses derived from configurations in either two or five dimensions. Figures 16 and 17 show the cluster-count matrix and associated Venn diagram derived from all 2D NMDS configurations for this set of areas and nuclei.

The structure of the configuration has been preserved in the set structure of this figure; that is, there are no long, sinuous pockets where one set includes structures that appear in completely different regions of the configuration. At the largest scale, the set structure in this case is similar to that of the global cluster-count matrix described earlier (figure 5). There are four weakly clustered sets where the constituent nuclei all share the same cluster in at least 27% of the separate cluster trees. The first of these weakly clustered sets contains the parts of the hippocampal formation that were closely associated in previous analyses (CA3, DG, CA1, ENT, SUB and SUM), but unlike other analyses, LM is also included in this set. The second group of areas in this set includes a strongly clustered set that is consistently found across all analyses and includes PAR, PRE, POST and AD. This strongly clustered set also formed a moderately clustered set with LD.

The second weakly clustered set is distinct from the first. It contains several overlapping, moderately clustered sets and three distinct, strongly clustered sets. The first strongly clustered set contains PRh, PL, ILA and MM, and the second, MM, AM and ACA. Both of these sets are included in a moderately clustered set that also contains MD and IAM. The third strongly clustered set contains AV and RSP and is also included in a moderately

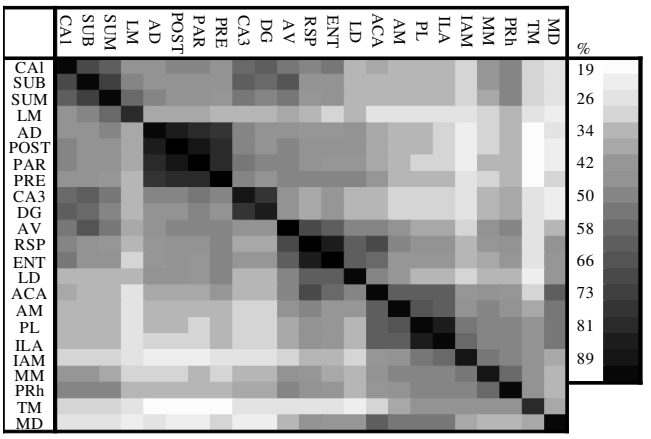

Figure 20. Cluster-count matrix calculated from non-parametric cluster analysis of 2D and 5D NMDS configurations of the binarized forms of $\hat{N}$, $\hat{T}$, $\hat{P}$ and $\hat{W}$.

clustered set that also includes the second strongly clustered set just described. The third and fourth weakly clustered sets overlap the first and second, with a degree of overlap between themselves. The set structure of the 2D solutions differs from the global and transform-specific cluster-count matrices: the moderately clustered sets overlap in a way similar to that described previously, but not to an extent that would strongly reinforce the interpretation of a pathway of areas in this system.

The set structure of the cluster-count summary matrix of all the cluster analyses derived from 5D configurations differs from that derived from 2D structures (compare figures 18 and 20). The principal difference is the large-scale (weakly clustered set) structure.

The parts of the hippocampus (CA3, CA1 and DG), the SUM, the ENT and the SUB once again form a moderately clustered set that appears separately from other structures, except for the inclusion of the MM in a weakly clustered set that involves all six members of this moderately clustered set (see figure 19). A second weakly clustered set involves SUB and ENT, and is associated with all of the other parts of the retrohippocampal region (PAR, PRE, POST), and the AD. As was the case in almost all the other analyses, AD, PRE, POST and PAR formed a strongly clustered set.

AV, LD, RSP and ACA also form a nearby, strongly clustered set that is included in two moderately clustered sets, which both include PAR, PRE, POST, AD and AM. One of these sets includes LM and the other MM. This reinforces the connectional difference between MM and LM that has been a feature throughout these analyses. Most nuclei of the anterior thalamus appear in the same moderately clustered set, rather than being spread out over many structures. IAM and PRh form a strongly clustered set that is also contained within a moderately clustered set with MM and PL. ILA, PL and MD form a strongly clustered set that is also contained within a moderately clustered set with TM and PRh.

(e) *Summary of results*

The cluster-count matrices across all data analyses for a given similarity matrix yielded broadly similar results with respect to the composition of the main sets of clusters. Some features of the summary cluster-count matrix were specific to a subset of cluster analyses: for example, the association of LM to the strongly clustered set of POST, PRE, PAR and AD consistently occurred in 5D and global cluster-count analyses but not in 2D summaries.

The main differences between analyses derived from different similarity matrices were the medium- and large-scale groupings of areas. The strongly clustered structure was consistent throughout, but the different methods of analysis emphasized different aspects of the data at medium and large scales. The interpretation of the set structure in terms of organizational schemes, which might correspond to physiological properties of individual cells in the different areas, is therefore not straightforward. However, the analyses indicated the existence of four main 'connectional groups'. The first was made up of SUM, SUB, ENT, CA1, CA3 and DG. These areas are the parts of the hippocampus proper, together with the retrohippocampal areas often associated with it (see Redish & Touretzky 1997). The inclusion of the SUM in this group is unlike most other organizational schemes (e.g. Lopes da Silva *et al.* 1990), and its inclusion was a product of its efferent connections to the hippocampus, especially to the dentate gyrus (Haglund *et al.* 1984). This nucleus contains cells that fire in phase with the theta rhythm of the hippocampus (Kirk 1997).

The second group was made up of PRE, POST, PAR and AD. These structures are probably the most tightly associated structures. This reflects particularly the difference in connection patterns between the various anterior thalamic nuclei (Shibata 1993).

The third and fourth groups were less clearly defined. The third group consisted of LM, RSP, AV, ACA, AM, LD and MM, and this group had quite a heterogeneous structure. It was rare for MM and LM to be categorized in the same group. The LM tended to be more frequently affiliated to the second group; whereas the MM tended be more frequently associated with the fourth group. The fourth group consisted of IAM, ILA, PL and MD. TM, which was frequently unaffiliated, only affiliated to this group. The 2D and 5D configurations presented different cluster-count sets. Extensive simulation with test data (Burns 1998) suggests that the better-fitting, 5D sets are to be preferred and, therefore, that the methods we have used for rendering higher-dimensional relationships interpretable have been valuable in this case.

The parcellation scheme used in the analysis had a strong effect on the interpretation of the connectional groups produced. For example, these analyses have described the various parts of the mammillary bodies (SUM, TM, LM and MM) as belonging to different connectional groups, which would clearly be impossible if the mammillary bodies were considered to be a single entity (Redish & Touretzky 1997; Gray 1982).

4. DISCUSSION

(a) *Limitations of the proximity model of neural connectivity*

The basic premise of this paper is that analysis is required to interpret neuroanatomical connection data.

This point of view was stated explicitly by Young *et al.* (1995): 'In all other disciplines in which complex numerous data are derived from experiments, conclusions drawn from data without any supporting analysis are not considered reliable.' An important method that has been employed to provide supporting analyses of neuroanatomical data has been the NMDS technique (Shepard 1962; Kruskal 1964; Takane *et al.* 1977; Cox & Cox 1995; Young 1992, 1993; Scannell & Young 1993; Scannell *et al.* 1995; Burns & Young 1996). This methodology is based on representing a system of brain structures as a configuration of points embedded in a multidimensional Euclidean space. If the distances between points are monotonically related to the strength of the connection between the areas that the points represent, the configuration can be treated as an approximate representation of the connectional organization of the system. We call this the 'proximity model' of neural connectivity, and now address some of its theoretical drawbacks.

Quantitative connection weights are rarely measured or defined in experimental tracing studies. This is due to inherent methodological problems (Warren 1992; but, see Patton & McNaughton 1995). Consequently, almost all previous connectivity analyses have been based on qualitatively defined data (cf. Young *et al.* 1995). However, if it were possible to define neural connection strength as a quantitative scalar measurement at the ratio level (Coombs 1964), then the proximity model of neural connectivity would not be able to represent the data accurately for several reasons.

First, reciprocal connections are common, but these connections are usually asymmetrical (i.e. $c_{ij} \neq c_{ij}$). The distances in NMDS solutions representing their origins are not (i.e. $d_{ij} = d_{ij}$). This is an intrinsic limitation of the proximity model, and can be addressed only by unwieldy modelling of the asymmetries (Cox & Cox 1995). Multidimensional scaling can hence, in principle, be used to model asymmetrical data, but the methods are somewhat underdeveloped presently, and almost any extension of them will make interpretation more difficult.

Second, the question of how non-connections should be modelled in the proximity model has been debated (Simmen *et al.* 1994; Goodhill *et al.* 1995; Young *et al.* 1994, 1995). Modelling non-connections (i.e. $c_{ij} = 0$) in a proximity model by setting all corresponding distances to a high value would appear to be a natural extension of modelling very sparse connections (i.e. $c_{ij} \approx 0$). This could be considered a realistic reflection of the data, since one absent connection could not be said to be more or less absent than another absent connection (Young *et al.* 1995). It is tempting to conclude that unconnected structures should all lie equally far apart in an NMDS solution. Consider, however, that the retina is connected neither to the visual cortex nor to the hippocampus. It seems counter-intuitive to hold that the retina should lie equally distant from both the hippocampus and the visual cortex, since studies of neural organization should associate the retina and visual cortex more closely than the retina and hippocampus. For example, the activity of neurons in primary visual cortex is directly influenced by the retina even under general anaesthesia (Wiesenfeld & Kornel 1975), whereas physiological activity of hippocampal cells may occur in the absence of visual input (Quirk *et al.* 1990). On the other hand, it could be argued that the closer proximity—in both neurophysiological and NMDS analyses—of retina and visual cortex is a function of the global connectivity of the system, which emerges only by fitting the global anatomical constraints. NMDS analyses hence meet this issue by embedding the data in a low-dimensional Euclidean space that illustrates global organizational features of the data, rather than fitting the minutiae of all data to an overly complex representation. While these considerations suggest a congruence between the requirement for interpretably low-dimensional configurations and the requirement to fit global constraints in an informative way, it is apparent that useful approximation of neural organization is the most that can be claimed for results derived within the proximity model.

Third, considerations of how quantitative connection strength data could be fitted in the proximity model raises questions of how appropriate are coordinate spaces in general to represent connectional organization. For example, consider the pathway from the retina (R) to the primary visual cortex (VISp) in the rat. Roughly 4×10^4 retinal ganglion cells project to the dorsal lateral geniculate nucleus (LGd; Linden & Perry 1983; Martin 1986). When wheat germ agglutinin–horseradish peroxidase is injected into the primary visual cortex, over 90% of labelled cells lie in the dorsal lateral geniculate nucleus (Sanderson *et al.* 1991). Thus, both the connection from R to LGd and the connection from LGd to VISp are dense and would be represented by short distances in the proximity model. The non-connection from R to VISp would be represented by a long distance. The triangle inequality (where the sum of the lengths of two shorter edges of a triangle must be equal to greater than the length of the longest edge) would certainly be violated in this circumstance. Fulfilment of the triangle inequality is a requirement of any coordinate space, so that fitting the same data into a non-Euclidean space would not suffice. However, given a non-Euclidean metric dissimilarity variable (δ_{ij} for $i, j = 1$ to n), the most common practice is to transform it to ϵ_{ij} by adding a constant value to all dissimilarities (Cailliez 1983). In this way, non-Euclidean connection data could be represented in a Euclidean proximity model, but at the cost of the ratio properties of such data.

Fourth, in the general case, the number of dimensions required in an analysis increases as the number of points increases. Since a key goal is to treat the organization of the whole brain, methods that can make higher dimensional configurations representing very large networks interpretable are desirable. High-dimensional representations, however, possess geometrical properties that appear counter-intuitive when compared to 2D and 3D spaces (see Scott 1992). One difference is that the volume occupied by the proportion of the space surrounding the centre of mass is much lower in high-dimensional space. For example, the ratio of the area of a 2D circle of unit radius to the area of a square that encloses it is 0.793. In contrast, the ratio of the volume of a 7D unit hypersphere to that of its enclosing hypercube is 0.037 (Scott 1992). Concomitantly, the tails of high-dimensional multivariate normal distributions contain increasingly more of the probability mass of the distribution as the dimensionality increases. We have already seen that straightforward

visualization of spaces with more than three dimensions is not possible, and that density estimation is often impossible in spaces with more than five dimensions, since accurately estimating the density of points in high-dimensional spaces requires very large sample sizes (Epanechnikov 1969; Silvermann 1986). These difficulties are worthy of further study, but the present analyses mark our first attempts to meet some of the problems posed by systems of intermediate dimensionality, by mapping sets derived from higher dimensional representations into an interpretable number of dimensions.

These considerations suggest some of the difficulties presented by the proximity model: the asymmetries of reciprocal connections are not represented; the definition of dissimilarity depends on the interpretation of non-connections; non-Euclidean properties are difficult to fit; and data that may be of quite high dimensionality are required to be represented in as few dimensions as possible in order to be interpretable. The NMDS approach used here and elsewhere is hence an approximation of a notional ratio-level proximity model, using ordinal data. The objective of this work is not then to obtain a mathematically perfect representation of the connectivity data. Rather, the aim is to better understand the organization of a system, establishing the main organizational features of the macrocircuitry in order to use the information gained as a predictive tool (e.g. Scannell *et al*. 1996). The limitations described above suggest that the proximity model would probably not be capable of providing more than a general description of the organization of neural systems, even if ideally accurate, quantitative data describing connection strengths were available. Methods of analysis that are not based on the proximity model, such as optimal set analysis (Hilgetag, Burns, O'Neill, Scannell & Young, this issue) do not suffer from these drawbacks and may prove more appropriate in the future.

(b) *The organization of neural systems involved in spatial memory in rats*

The network of connections we have analysed could be considered to be part of the so-called 'limbic system' (Kandel *et al*. 1991; Lopes da Silva *et al*. 1990). The structures analysed included many areas that are implicated in functional models concerning spatial memory and navigation (O'Keefe & Nadel 1978; Buzsaki *et al*. 1994; McNaughton *et al*. 1996; Redish & Touretzky 1997), anxiety (Gray 1982) and declarative memory (Eichenbaum *et al*. 1992). Most of the areas display functional properties other than those expected in spatial processing, and these include the hippocampus (Bunsey & Eichenbaum 1996).

Theoreticians naturally take anatomical constraints into account when designing their models, but only rarely is the anatomy taken as a starting point. More often, the functional properties of the system are probed and selected anatomy is used as a means of substantiating models of physiological or computational mechanisms. We have wondered whether in this process there might be a slight danger that only neuroanatomical information that conforms to the theory is used. Even the most accomplished and well-informed models (e.g. Redish & Touretsky 1997), which are inspired by earlier theoretical work and driven carefully by neurophysiological and behavioural results (Touretzky & Redish 1996), appear to omit empirically verified connections from their models. For example, a number of 'anomalous' connections appear to have been omitted from Touretsky & Redish's models: the efferent connections of SUM to the hippocampus (Haglund *et al*. 1984); the projection from PRE to LD (Van Groen & Wyss 1990); the projection from RSP to LD (Seki & Zyo 1984); and the projection from AM to ENT (Shibata 1993). In other respects, however, the Touretsky & Redish models present an organizational scheme similar in many ways to that our analyses of real data have described. For example, the hippocampus, the ENT and the SUB appear closely associated; the POST, the PRE and the PAR are associated with the anterior thalamic nuclei and the retrospenial cortex. There is some disagreement over the role of the mammillary bodies, since only the mammillary bodies' inputs from the SUB are considered. Similarly, SUM, MM and LM possess other connections within this system (see figure l; Shibata 1988, 1989, 1992; Haglund *et al*. 1984).

It may be that the computational approach to purely anatomical data is of limited use as a guide to function. However, successful predictions of physiological phenomena at the systems level have been derived from the analysis of neuroanatomical data, a fact that implies that structure and function are linked at this level, as at every other. The products of our analyses are clues to the trends that might be expected in receptive-field characteristics and other physiological properties in the areas analysed. Patchy knowledge of the physiological properties of areas in a system, allied to analysis of connectivity, can hence be used to make predictions of physiological properties in specific, but less well-studied regions. For example, Scannell *et al*. (1995, 1996) predicted successfully that plaid-pattern selective cells would be found in the anterior ectosylvian sulcus of the cat on the basis of anatomical relationships made apparent by analysis of connectivity and patchy information about the neurophysiology of suprasylvian cortical cells (Scannell *et al*. 1995, 1996). The hippocampus and other parts of the limbic system in the rat have been well studied neurophysiologically, and we hoped that a similar approach would be beneficial here.

Deductions of this kind begin by identifying an interesting neurophysiological property that varies across brain structures. In this system, 'place cells' are particularly interesting. The electrophysiological properties of place cells in the hippocampus have been extensively studied over the past 25 years (e.g. O'Keefe & Dostrovsky 1971; Muller 1996). Cells that exhibit these and similar selective properties have been found in several parts of the rat brain, namely the CA3 and CA1 fields of the hippocampus (O'Keefe & Dostrovsky 1971; O'Keefe & Nadel 1978), the SUB (Sharp & Green 1994; Sharp 1997), the ENT (Quirk *et al*. 1992), and the PAR (Taube 1995*b*). A second neurophysiological property is shown by cells that fire preferentially when the head of the animal is pointing in a specific direction: 'head-direction' cells. These cells have been found in the POST (Taube *et al*. 1990), AD (Blair & Sharp 1995, Taube 1995*a*), LD (Mizumori & Williams 1993), and the posterior parietal (also referred to as the anterior part of the medial extrastriate cortex) and retrosplenial cortices (Chen *et al*. 1994).

This distribution of these electrophysiological properties across structures is consistent with the classification of brain structures derived here. One of the cluster-count sets that was clearly delineated in the analyses comprised CA1, CA3, DG, SUM, SUB and ENT. Four of these are structures that contain place cells. The closest neighbouring sets, which occasionally overlap with this set, contain AD, POST, PRE and LD, which have all been reported to contain head-direction cells. This separation between clusters is not complete, since single units have been recorded in the PAR, which correspond to the spatial position of the animal, rather than the direction in which it faces (Taube 1995*a*). In the plots shown earlier, however, PAR lies closer to the structures in the set containing the hippocampus than the other parts of the second set.

Studies that have attempted to find either head-direction or place cells in the PL of the rat report a null result, even though PL receives a direct connection from CA1 (Poucet 1997; Jung *et al*. 1998). Our results showed PL to be placed consistently in a different cluster from CA1, and so we would not expect PL to be an intimate functional associate of CA1. The presence of the direct connection between these structures does not reliably predict their physiological relationship. We believe that functional relationships are constrained by the organization of the connectivity of the rest of the system, which cannot be determined by reference to only a single projection, but which we have analysed. If physiological properties remain consistent with this assumption and the results of these analyses, we would expect SUM, which appears in the same set as several other brain areas that contain place cells, also to contain place cells. The same logic suggests that ACA and AV may contain head-direction cells.

REFERENCES

Amaral, D. G. & Witter, M. P. 1989 The three-dimensional organization of the hippocampal formation: a review of anatomical data. *Neuroscience* **31**, 571–591.

Amaral, D. G. & Witter, M. P. 1995 Hippocampal formation. In *The rat nervous system* (ed. G. Paxinos), pp. 443–493. Sydney: Academic Press.

Blair, H. T. & Sharp, P. E. 1995 Anticipatory head direction signals in anterior thalamus: evidence for a thalamocortical circuit that integrates angular head motion to compute head direction. *J. Neurosci.* **15**, 6260–6270.

Bunsey, M. & Eichenbaum, H. 1996 Conservation of hippocampal memory function in rats and humans. *Nature* **379**, 255–257.

Burns, G. A. P. C. 1998 Neural connectivity of the rat: theory, methods and applications. DPhil thesis, University of Oxford, UK.

Burns, G. A. P. C. & Young, M. P. 1996 *Neurobase: a neuroanatomical connection database and its use in providing a description of connections in the rat hippocampal system. Abstracts.* Newcastle upon Tyne, UK: Brain Research Association.

Buzsaki, G. 1989 Two-stage model of memory trace formation: a role for 'noisy' brain states. *Neuroscience* **31**, 551–570.

Buzsaki, G., Bragin, A., Chrobak, J. J., Nadasdy, Z., Sik, A., Hsu, M. & Ylinen, A. 1994 Oscillatory and intermittent synchrony in the hippocampus: relevance to memory trace formation. In *Temporal coding in the brain* (ed. G. Buszaki, R. Llinas, W. Singer, A. Berthoz & Y. Christen), pp. 145–155. Berlin: Springer.

Cailliez, F. 1983 The analytical solution to the additive constant problem. *Psychometrika* **48**, 305–308.

Chen, L. L., Lin, L. H., Green, E. J., Barnes, C. A. & McNaughton, B. L. 1994 Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Exp. Brain Res.* **101**, 8–23.

Cohen, N. J. & Eichenbaum, H. 1991 The theory that wouldn't die: a critical look at the spatial mapping theory of hippocampal function. *Hippocampus* **1**, 265–268.

Coombs, C. H. 1964 *A theory of data*. New York: Wiley.

Cox, T. F. & Cox, M. A. A. 1995 *Multidimensional scaling*. Newcastle upon Tyne, UK: Chapman & Hall.

Eichenbaum, H., Otto, T. & Cohen, N. J. 1992 The hippocampus—what does it do? *Behav. Neural Biol.* **57**, 2–36.

Epanechnikov, V. A. 1969 Nonparametric estimation of a multidimensional probability density. *Theor. Prob. Appl.* **14**, 153–158.

Goodhill, G. J., Simmen, M. W. & Willshaw, D. J. 1995 An evaluation of the use of multidimensional scaling for understanding brain connectivity. *Phil. Trans. R. Soc. Lond.* B **348**, 265–280.

Gordon, A. D. 1981 *Classification*. London: Chapman & Hall.

Gower, J. C. & Legendre, P. 1986 Metric and euclidean properties of dissimilarity coefficients. *J. Classif.* **3**, 5–48.

Gray, J. A. 1982 *The neuropsychology of anxiety: an enquiry into the functions of the septo-hippocampal system*. Oxford University Press.

Groenewegen, H. J. 1988 Organization of the afferent connections of the mediodorsal thalamic nucleus in the rat, related to the mediodorsal–prefrontal topography. *Neuroscience* **24**, 379–431.

Haglund, L., Swanson, L. W. & Kohler, C. 1984 The projection of the supramammillary nucleus to the hippocampal formation: an immunohistochemical and anterograde transport study with the lectin PHA-L in the rat. *J. Comp. Neurol.* **229**, 171–185.

Jakab, R. L. & Leranth, C. 1995 Septum. In *The rat nervous system* (ed. G. Paxinos), pp. 405–442. Sydney: Academic Press.

Jung, M., Qin, Y., McNaughton, B. & Barnes, C. 1998 Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cerebr. Cortex* **8**, 437–450.

Kandel, E. R., Schwartz, J. H. & Jessel, T. M. 1991 *Principles of neural science*. London: Appleton and Lange.

Kirk, I. J. 1997 Supramammillary neural discharge patterns and hippocampal EEG. *Brain Res. Bull.* **42**, 23–26.

Kruskal, J. B. 1964 Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika* **29**, 1–27.

Linden, R. & Perry, V. H. 1983 Massive retinotectal projection in rats. *Brain Res.* **272**, 145–149.

Lopes da Silva, F. H., Witter, M. P., Boeijinga, P. H. & Lohman, A. H. 1990 Anatomic organization and physiology of the limbic cortex. *Physiol. Rev.* **70**, 453–511.

McNaughton, B. L. (and 10 others) 1996 Deciphering the hippocampal polyglot—the hippocampus as a path integration system. *J. Exp. Biol.* **199**, 173–185.

Martin, P. R. 1986 The projection of different retinal ganglion cell classes to the dorsal lateral geniculate nucleus in the hooded rat. *Exp. Brain Res.* **62**, 77–88.

Meibach, R. C. & Siegel, A. 1977 Efferent connections of the hippocampal formation in the rat. *Brain Res.* **124**, 197–224.

Mizumori, S. J. Y. & Williams, J. D. 1993 Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *J. Neurosci.* **13**, 4015–4028.

Morris, R. G., Garrud, P., Rawlins, J. N. & O'Keefe, J. 1982 Place navigation impaired in rats with hippocampal lesions. *Nature* **297**, 681–683.

Muller, R. 1996 A quarter of a century of place cells. *Neuron* **17**, 813–822.

Neave, N., Aggleton, J. P., Nagle, S. & Sahgal, A. 1996 *Proposed neuroanatomical circuitry of spatal memory in the rat. Abstracts.* Newcastle upon Tyne, UK: Brain Research Assocation.

O'Keefe, J. & Dostrovsky, J. 1971 The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175.

O'Keefe, J. M. & Nadel, L. 1978 *The hippocampus as a cognitive map.* Oxford, UK: Clarendon Press.

Patton, P. E. & McNaughton, B. 1995 Connection matrix of the hippocampal formation. I. The dentate gyrus. *Hippocampus* **5**, 245–286.

Poucet, B. 1997 Searching for spatial unit firing in the prelimbic area of the rat medial prefrontal cortex. *Behav. Brain Res.* **84**, 151–159.

Quirk, G. J., Muller, R. U. & Kubie, J. L. 1990 The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J. Neurosci.* **10**, 2008–2017.

Quirk, G. J., Muller, R. U., Kubie, J. L. & Ranck Jr, J. B. 1992 The positional firing properties of medial entorhinal neurons: description and comparison with hippocampal place cells. *J. Neurosci.* **12**, 1945–1963.

Rawlins, J. N., Lyford, G. & Seferiades, A. 1991 Does it still make sense to develop nonspatial theories of hippocampal function? *Hippocampus* **1**, 283–286.

Redish, A. D. & Touretzky, D. S. 1997 Cognitive maps beyond the hippocampus. *Hippocampus* **7**, 15–35.

Rothblat, L. A., Vnek, N., Gleason, T. C. & Kromer, L. F. 1993 Role of the parahippocampal region in spatial and non-spatial memory: effects of parahippocampal lesions on rewarded alternation and concurrent object discrimination learning in the rat. *Behav. Brain Res.* **55**, 93–100.

Sanderson, K. J., Dreher, B. & Gayer, N. 1991 Prosencephalic connections of striate and extrastriate areas of rat visual cortex. *Exp. Brain Res.* **85**, 324–334.

Scannell, J. W. & Young, M. P. 1993 The connectional organization of neural systems in the cat cerebral cortex. *Curr. Biol.* **3**, 191–200.

Scannell, J. W., Blakemore, C. & Young, M. P. 1995 Analysis of connectivity in the cat cerebral cortex. *J. Neurosci.* **15**, 1463–1483.

Scannell, J. W., Sengpiel, F., Benson, P. J., Tovée, M. J., Blakemore, C. & Young, M. P. 1996 Visual motion processing in anterior ectosylvian sulcus of the cat. *J. Neurophysiol.* **76**, 895–907.

Scott, D. 1992 *Multivariate density estimation. Theory, practice and visualization.* New York: Wiley.

Seki, M. & Zyo, K. 1984 Anterior thalamic afferents from the mamillary body and the limbic cortex in the rat. *J. Comp. Neurol.* **229**, 242–256.

Sharp, P. E. 1997 Subicular cells generate similar spatial firing patterns in two geometrically and visually distinctive environments. *Behav. Brain Res.* **85**, 71–92.

Sharp, P. E. & Green, C. 1994 Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *J. Neurosci.* **14**, 2339–2356.

Shepard, R. N. 1962 The analysis of proximities: multidimensional scaling with an unknown distance function. I. *Psychometrika* **27**, 219–246.

Shibata, H. 1988 A direct projection from the entorhinal cortex to the mammillary nuclei in the rat. *Neurosci. Lett.* **90**, 6–10.

Shibata, H. 1989 Descending projections to the mammillary nuclei in the rat, as studied by retrograde and anterograde transport of wheat germ agglutinin-horseradish peroxidase. *J. Comp. Neurol.* **285**, 436–452.

Shibata, H. 1992 Topographic organization of subcortical projections to the anterior thalamic nuclei in the rat. *J. Comp. Neurol.* **323**, 117–127.

Shibata, H. 1993 Direct projections from the anterior thalamic nuclei to the retrohippocampal region in the rat. *J. Comp. Neurol.* **337**, 431–445.

Silvermann, B. W. 1986 *Density estimation for statistics and data analysis.* London: Chapman & Hall.

Simmen, M. W., Goodhill, G. J. & Willshaw, D. J. 1994 Scaling and brain connectivity [letter; comment]. *Nature* **369**, 448–450.

Swanson, L. W. & Cowan, W. M. 1977 An autoradiographic study of the organization of the efferent connections of the hippocampal formation in the rat. *J. Comp. Neurol.* **172**, 49–84.

Takane, Y., Young, F. W. & de Leeuw, J. 1977 Nonmetric individual differences multidimensional scaling: an alternating least squares method with optimal scaling features. *Psychometrika* **42**, 7–67.

Taube, J. S. 1995*a* Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *J. Neurosci.* **15**, 70–86.

Taube, J. S. 1995*b* Place cells recorded in the parasubiculum of freely moving rats. *Hippocampus* **5**, 569–583.

Taube, J. S., Muller, R. U. & Ranck Jr, J. B. 1990 Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J. Neurosci.* **10**, 420–435.

Touretzky, D. S. & Redish, A. D. 1996 Theory of rodent navigation based on interacting representations of space. *Hippocampus* **6**, 247–270.

Van Groen, T. & Wyss, J. M. 1990 The connections of presubiculum and parasubiculum in the rat. *Brain Res.* **518**, 227–243.

Warren, M. A. 1992 Simple morphometry of the nervous system. In *Quantitative methods in neuroanatomy* (ed. M. G. Stewart), pp. 211–247. New York: Wiley.

Wiesenfeld, Z. & Kornel, E. E. 1975 Receptive fields of single cells in the visual cortex of the hooded rat. *Brain Res.* **94**, 401–412.

Young, M. P. 1992 Objective analysis of the topological organization of the primate cortical visual system. *Nature* **358**, 152–155.

Young, M. P. 1993 The organization of neural systems in the primate cerebral cortex. *Proc. R. Soc. Lond.* B **252**, 13–18.

Young, M. P., Scannell, J. W., Burns, G. A. & Blakemore, C. 1994 Analysis of connectivity: neural systems in the cerebral cortex. *Rev. Neurosci.* **5**, 227–250.

Young, M. P., Scannell, J. W. & Burns, G. A. P. C. 1995*a* *The analysis of cortical connectivity.* Austin, TX: Springer and R. G. Landes.

Young, M. P., Scannell, J. W., O'Neill, M. A., Hilgetag, C.-C., Burns, G. & Blakemore, C. 1995*b* Non-metric multidimensional scaling in the analysis of neuroanatomical connection data and the organization of the primate cortical visual system. *Phil. Trans. R. Soc. Lond.* B **348**, 281–308.