Extremely Dilute Modular Neuronal Networks: Neocortical Memory Retrieval Dynamics

Carlo Fulvi Mari

Nonlinear and Complex Systems Group,

Department of Mathematical Sciences, Loughborough University, U.K.*

and

Quantitative Psychology and Cognitive Science,

Department of Psychology, University of Liège, Bât B32, Sart Tilman, 4000 – Liège, Belgium†

E-mail: C.Fulvi-Mari@ulg.ac.be

5 June 2001

Abstract

A model of the columnar functional organization of neocortical association areas is studied. The neuronal network is composed of many Hebbian autoassociators, or *modules*, each of which interacts with a small subset of the others. Every module stores a number of elementary *features*. Memory *patterns* are peculiar combinations of features sparsely distributed over the multi-modular network. Any feature stored in any module can be involved in several of the stored patterns; feature-sharing is in fact source of local ambiguities and, consequently, a potential cause of erroneous retrieval activity spreading through the model network.

The retrieval dynamics of the large multi-modular autoassociator is investigated by means of quantitative analysis and numerical simulations. Correlated activation of adjacent modules and extramodular contacts more effective than the intramodular ones seem to be general requirements in order to achieve satisfactory quality of memory retrieval. An *oscillatory* retrieval process is found to be very efficient in overcoming feature-sharing drawbacks; it requires a mechanism that modulates the robustness of local attractors to noise, and neuronal activity sparseness such that quiescent and active modules are about equally noisy. Fluctuation of the number of modules pre-synaptic to any module across the network has to be less pronounced than that of the degree of the nodes in Bernoulli random graphs.

Even in ideal conditions, some spots of the network cannot be reached by retrieval activity spread, their locations depending on the pattern to retrieve and on the cue; the oscillatory retrieval process nearly saturates the consequent upper-bound to retrieval quality.

1 Introduction

The neocortex presents several levels of architectural and functional modularity (Kaas, 1987; Braitenberg and Shutz, 1991; Fuster, 1997; Mountcastle, 1997). What seem to be the elementary processing units of cognitive functions are the *columns*, compact assemblies of densely interconnected excitatory and inhibitory neurons that extend vertically through the layers of the cortical sheet and whose widths range between 0.3 and 1.0 mm (Mountcastle, 1977; Braitenberg and Shutz, 1991; Mountcastle, 1997). Due to its small dimensions, the column cannot be revealed by lesion studies, and not yet by imaging (e.g., PET and fMRI), but neurophysiology has provided plentiful evidence of its functional individuality. In somatosensory cortex SI, neurons belonging to the same column have very similar and most overlapping receptive fields on the skin; on the contrary, the receptive fields of neurons that belong to different columns are markedly distinct (Mountcastle, 1957). In visual cortex V1, columns respond preferentially to visual stimuli shown in a well-defined region of the visual field of one of the eyes; when tested with short straight line stimuli, they are selectively sensitive to lines at particular angles of orientation (Hubel and Wiesel, 1977). In medio-temporal area MT, columns are selectively activated by motion directions (Albright et al., 1984), while columns of the nearby area MSTd are responsive to characteristic combinations of the visual flow stimuli contraction, expansion, rotation and translation, *i.e.* spirals, according to a continuous tuning curve (Graziano et al., 1994). Columns in inferotemporal cortex (IT) respond selectively to non-elementary visual stimuli that range from moderately complex shapes to figures very rich in details like individual faces (Miyashita and Chang, 1988; Fujita et al., 1992; Tanaka, 1996)[1]. The posterior parietal cortex seems to be responsible for higher perceptual memory and processing (Mishkin et al., 1983; Mountcastle et al., 1984), and no simple sensorial input seems to be able to select columnar activity there. What all the columns in that region have in common is their activability in tasks involving the actions in, or the perception of, or the attention to the environment that surrounds the subject. Still, neurons with similar properties are arranged in vertical columns (Mountcastle et al., 1975; Mountcastle, 1995; Mountcastle, 1997): a column could be active during fixation of gaze, or slow pursuit tracking, or reaching by an arm, or manipulation, or complex visual stimulation. Rich of connections with subcortical, limbic and neocortical areas, like the reciprocal connections with the posterior parietal and inferotemporal cortices, the prefrontal cortex is at top of the hierarchy of motor memory[2] and also supports high-level cognitive/executive functions other than pure motor planning (Martin and Chao, 2001; Fuster, 1995; Goldman-Rakic, 1988; Grafton, 1995; Buckner and Petersen, 1996; Savage et al., 2001; Fuster, 1997). During *working memory* performance, when subjects are required to keep prolonged memory of briefly presented stimuli in order to execute a *delay task* correctly, marked self-sustained neuronal activity is found in columns of the prefrontal cortex (Fuster (1999) and references therein, e.g., Funahashi et al. (1989)). In fact, working memory is not peculiar of prefrontal cortex (Fuster, 1998; Fuster, 1997): parietal areas simultaneously also maintain high activity level since a perceptual stimulus is processed within its environmental context (Friedman and Goldman-Rakic, 1994)[3],

---

[1] According to Tanaka (1996), neurons in area TE may be also sensitive to orientation, size, and contrast polarity of their critical visual features, though being neutral to the object position in their large visual fields.

[2] In general, the activation of a motor representation in prefrontal cortex is more evident when the subject is learning a new motor task. After practice, it seems that the 'schemas' are relocated in motor areas that lie in lower levels of the motor hierarchy (Jenkins et al., 1994).

[3] In support to the theory of parallel processing in prefrontal and parietal associative areas, there are also some anatomical findings according to which posterior parietal and dorsolateral prefrontal cortices project in common to virtually the same

and, if the stimulus involves visual recognition, inferotemporal regions keep high activity too (Fuster and Jervey, 1982; Miyashita and Chang, 1988). It is often hypothesized that such self-sustainment of enhanced activity is provided by dynamical attractor properties of the columnar neuronal networks and is possibly related to the retrieval of learned features.

Prefrontal, posterior parietal and inferotemporal cortices are usually referred to as *association* areas of the neocortex and are thought to subserve semantic memory (Martin and Chao, 2001; Grafton, 1995; Ricci et al., 1999; Goel and Dolan, 2001; Fuster, 1997), that is, general knowledge of objects and events that is not strictly related to, and does not necessarily depend on specific episodic contexts. The present paper is focused onto modelling memory processes that supposedly take place in the neocortical association areas.

For intracolumnar *recurrent* neuronal contacts are Hebbian and relatively dense (Braitenberg and Shutz, 1991; Mountcastle, 1997), any column may hold autoassociative abilities, like pattern completion and self-sustainment of structured activity, mainly provided by the existence of stable dynamical attractors that formed during learning by means of (Hebbian) synaptic modifications and permit recovering specific distributions of neuronal activity (Little, 1974; Hopfield, 1982). This hypothesis is also supported by findings of persistent delay activity (Fuster and Jervey, 1982; Miyashita and Chang, 1988; Funahashi et al., 1989) and of neuronal responses that reveal ongoing memory retrieval and recognition processes, especially in associative areas (Naya et al., 2001; Tomita et al., 1999; Fujita et al., 1992; Tanaka, 1996; Van Hoesen, 1993). Together with the evidence of feature selectivity of columns, briefly reviewed above, these further observations suggest that memories may consist of distributed representations[4] of which columns locally store the finest elements. It is sometimes hypothesized that one of the advantages of a modular structure of this kind might be the possibility of representing new patterns making use of elements that already appeared in previously stored patterns. The question naturally arises of how the multitude of modules is organized to perform cognitive tasks, first among these the retrieval of neuronal representations from memory, like, for example, in pattern completion tasks. Some authors have addressed problems concerning memory retrieval in modular network models of neocortical areas (O'Kane and Treves, 1992; Lauro-Grotto et al., 1997; Fulvi Mari and Treves, 1998; Renart et al., 1999; Fulvi Mari, 2000), but further investigation seems very necessary.

In the present paper, dynamical properties of memory retrieval processes in a network of many modules are studied, every module being a neuronal (Hebbian) autoassociator. The appearance of any locally stored *feature* in possibly several global *patterns* is what creates most of the problems in building a theoretical model able to reproduce proper memory retrieval and that would also comply with biological constraints. For example, be A and B a pair of modules connected with each other, and be assumed that A is elicited by the cue stimulus to retrieve local feature a, while B is quiescent, in a spontaneous, uniform activity state. As a appeared in several of the learned patterns, B is likely to have stored more than one feature (e.g., $b_1^a, b_2^a, \ldots, b_n^a$) that were respectively active in n patterns that simultaneously also produced feature a in module A. Thus, due to Hebbianly shaped synapses between neurons of the two modules, so to speak, module A would equally tend to drive module B toward any one of the local attractors corresponding to the n mentioned features $(b_1^a, b_2^a, \ldots, b_n^a)$ stored in module B, not necessarily the correct one for what concerns the global pattern the cued

---

targets in over a dozen distinct cytoarchitectonic areas (Selemon and Goldman-Rakic, 1988).

[4]The hypothesis that any mnemonic representation is distributed across areas rather than being localized in a specific locus is widely supported by neurophysiological and neuropsychological/imaging evidence (cf., respectively, Fuster (1997) and Martin and Chao (2001) for short reviews).

fragment was taken from. The network must then be able to favour the spreading of correct retrieval across modules while suppressing spuriously activated features, one of the main obstacles being that every module is 'aware' of the states of only a very small fraction of the others. For experimental data constrain the number of extracolumnar inputs to any neuron to be about equal to the number of the intracolumnar inputs, it is unfeasible designing long-range (*white matter*) afferents from a large number of columns, as this would make the signals from any pre-synaptic module largely overcome by local signals.

In Section 2, the multi-modular network model is defined: neuronal model, architecture of axonal projections, and correlational scheme are described. In Section 3, some statistical quantities are identified concerning neuronal input currents and their correlation with stored features. Neuronal signal-to-noise ratio is calculated in Section 4 and used as an indicator of attractor stability. In Section 5, dynamical transitions between modular states are assigned probabilities that are functions of the signal-to-noise ratio and of the states of neighbouring modules; results from simulations of the dynamical model are presented and discussed. The simpler case in which local features are not shared among different patterns is illustrated in Section 6. Conclusions are drafted in Section 7.

2   Neuronal network model

Adopting firing-rate coding, it is assumed that in any active module any neuron can be either in an excited or in a suppressed firing level; this approximates the bimodal distribution of firing-rates found in neuronal networks during persistent (*delay*) selective activity, when it is usually possible to discriminate between *foreground* neurons, that are highly active, and *background* neurons, that fire at very low rate. In the quiescent modules, all neurons fire spontaneously at the same firing rate, that is higher than the background, and lower than the foreground neurons firing-rate.

The firing rate of neuron i in module m is represented by

$$V_{i_m} = B\tau'_m \eta^d_{i_m} + A(1 - \tau'_m) + C, \tag{1}$$

where: τ'_m is equal to 1 if module m is actively retrieving a local feature (d), and is equal to 0 if the module is quiescent; $\eta^d_{i_m}$ is equal to 1 if neuron i of module m is in the foreground of local feature d, and is equal to 0 if the neuron is in the background; $B + C$ is the firing-rate of foreground neurons; $A + C$ is the spontaneous activity firing-rate; C is the background neurons firing-rate. If the multi-modular network is reproducing the whole pattern p correctly, then the expression of the firing-rate of neuron i_m becomes

$$V_{i_m} = B\tau^p_m \eta^{d_m(p)}_{i_m} + A(1 - \tau^p_m) + C, \tag{2}$$

where τ^p_m is equal to 1 if module m is actively involved by pattern p, and is equal to 0 otherwise, and use is made of the mapping

$$(m, p) \longrightarrow d_m(p), \tag{3}$$

that produces the local feature d of module m associated with global pattern p. The fraction a of neurons in foreground of any feature is called *neuronal (activity) sparseness*, and is of the order of 10^{-2}, consistently with estimates from experiments on inferotemporal cortex by Miyashita (1988). The neurons forming the foreground of any feature in the model are randomly chosen with probability a, independently. The statistics of the binary variables $\{\tau^p_m\}$ in the set of global patterns is detailed in Section 2.2. For convenience, the active and the quiescent modules of any pattern p will be said to

constitute, respectively, the *foreground* and the *background* of pattern p, analogously to the definition used for neurons.

The parameters A, B and C enter the mathematics of the model only through the ratios A/B and C/B. The local recurrent circuits are assumed to keep A/B and C/B nearly constant across the modules in spite of the modular input fluctuations. From suitable literature involving electrode recording in prefrontal and inferotemporal cortices, it can be reasonably extrapolated that $A/B \sim 0.1$, with C/B slightly ($\sim 20\%$) larger than A/B. In the following, it will be assumed that $A/B = 0.1$ and $C/B = 0.12$ (values inferred from the plots of Funahashi et al. (1989)).

The total input current to neuron i of module m is written as

$$h_{i_m} = \sum_{j_m \neq i_m} J^S_{i_m j_m} V_{j_m} + \sum_{n \neq m} \sum_{j_n} J^L_{i_m j_n} V_{j_n}, \tag{4}$$

where: index n runs over the total number (M) of modules except the module (m) that contains neuron i_m; j_m runs over all the neurons of module m with the exception of neuron i_m (no self-interaction); j_n runs over all the neurons of module n. Non-specific contributions, like fast inhibition, are not reported explicitly. The synaptic weights are determined by Hebbian covariance rule according to the following formulae:

$$J^S_{i_m j_m} = \frac{b_{i_m j_m}}{T} \sum_p \tau^p_m (\frac{\eta^p_{i_m}}{a} - 1)(\frac{\eta^p_{j_m}}{a} - 1) \tag{5}$$

for the intramodular (*short-range*) contacts, and

$$J^L_{i_m j_n} = \frac{c_{i_m j_n}}{\lambda T} \sum_p \tau^p_m \tau^p_n (\frac{\eta^p_{i_m}}{a} - 1)(\frac{\eta^p_{j_n}}{a} - 1) \tag{6}$$

for the extramodular (*long-range*) contacts. Normalization constant T is equal to $L + N - 1$, by definition. Index p runs over the P stored patterns. The quenched random variable $b_{i_m j_m}$ determines the dilution of the intramodular connectivity: its value is 1 with probability b, and 0 with probability $1 - b$ for any choice of the ordered pair (i_m, j_m) (that is, the probability for a neuron to receive axonal projection from any other neuron of the same module is equal to b, and contacts need not be reciprocal). Variable $c_{i_m j_n}$ takes value 1 if neuron j_n in module n projects onto the dendrites of neuron i_m of module m; otherwise, it is equal to zero. The architecture of the inter-modular connections is included in the set $\{c_{i_m j_n}\}$ (cf. Section 2.1 for details). Note that synaptic modification happens between neurons belonging to different modules only when the pre-synaptic and the post-synaptic modules are simultaneously active, which seems biologically plausible.

As can be easily recognized, the expressions for the synapses (Eqs. 5 and 6) are a generalization, to modular neuronal networks, of the well-known formula for uniform networks with sparse neuronal activity (Tsodyks and Feigel'man, 1988). The present model also includes selective inter-modular connectivity (cf. Section 2.1) and sparseness in the pattern of activation of the modules during global processing (cf. Section 2.2).

The prefactor $1/\lambda$ in Eq. 6 is necessarily given a value greater than one. Indeed, while half of the inputs to any neuron come from neurons of the same module (Braitenberg and Shutz, 1991), the remaining half of inputs come from a number of other modules, this being mathematically reflected into the denominator s' appearing in the *signal* of Section 4; this repartition of extramodular

5

afferents makes the signals from a pre-synaptic module significantly weaker than the local signals. Besides, the synapses between neurons that belong to different modules are modified with a slower rate than the local synapses, as the probability for both pre-synaptic and post-synaptic modules to be simultaneously recruited in a global pattern is significantly smaller than the probability for any module to be recruited, of course. Finally, it should be noted that every local feature appears in a module several ($\sim \nu$) times during learning, the Hebbian local synapses thus being incremented correspondingly; on the contrary, very rarely adjacent modules learn the same pair of features for more than one pattern (cf. Section 2.2). Thus, λ should be given a value appropriately to alleviate the disadvantages of extramodular synapses in respect to the local ones (to compensate for the multiplicity ν, one could postulate the existence of a mechanism that keeps short-range contacts from being appreciably modified by further re-presentation of overlearned features, thus permitting larger values of λ; this alternative is not pursued here, though). Possible interpretations of $\lambda < 1$ from a physiological point of view may be that inter-modular synapses are 'faster' in learning (Fulvi Mari and Treves, 1998), that the long-range contacts are intrinsically stronger than the short-range ones, and that apical dendrites convey the post-synaptic potentials to the soma more effectively than the basal ones, as white matter axons make contacts preferably onto the apical dendrites of the pyramidal neurons while the intramodular and horizontal contacts happen mainly on the basal dendrites (Braitenberg and Shutz, 1991).

2.1   Architecture

The number of modules (M) and that of neurons inside any module (N) are supposed to be large; this allows for neglecting statistical fluctuations of some quantities that are used in the mathematical model, and to approximate some binomial distributions with Poisson ones. These assumptions do not seem in discordance with reality: the number of neurons constituting any column is estimated to be of the order of 10^4; being the diameter of the base of any column typically between 0.5 and 1.0 mm (Braitenberg and Shutz, 1991; Goldman-Rakic, 1988), and assuming that the 'flat' extension of an area is about 0.1 m^2, the number of columns per area is of the order of 10^5.

Every neuron receives axonal contacts from a number of other neurons (Fig. 1). The probability (b) for any neuron to receive a projection from another neuron of the same module is assumed to be finite, of the order of a few tenths (Braitenberg and Shutz, 1991). The extramodular afferents to any module are assumed to originate from a small number of the other modules (called *adjacent*, or *neighbouring*, modules); this seems in accordance with some results from dying experiments (cf. Braitenberg and Shutz (1991), especially pages 144-145, and Goldman-Rakic (1988)). One may think of this net of long-range connections imagining the existence of a few *channels* that connect any module to other modules; inside these channels, and only through these, the white matter axons are allowed to pass. The net of channels is modelled as a *random graph*, being $\frac{s'}{M-1}$ the small probability for any pair of modules to be connected by a channel. Since the number of neighbours per module is distributed around a relatively small value (s') even when the ideal 'thermodynamic' limit of infinite number of modules is performed, the graph is said to be *extremely dilute*. The probability for any neuron to receive a contact from any neuron of an afferent module is $\frac{L}{s'N}$, so that on average any neuron receives L synapses from extramodular neurons. The channels, in principle, need not be symmetric, that is, channel (m, n) may be generated randomly and independently from channel (n, m). Nevertheless, for simplicity, the channels are here assumed to be symmetrical. Reality is probably in between full asymmetry and symmetry, or at least this is the impression one has when looking at the results in
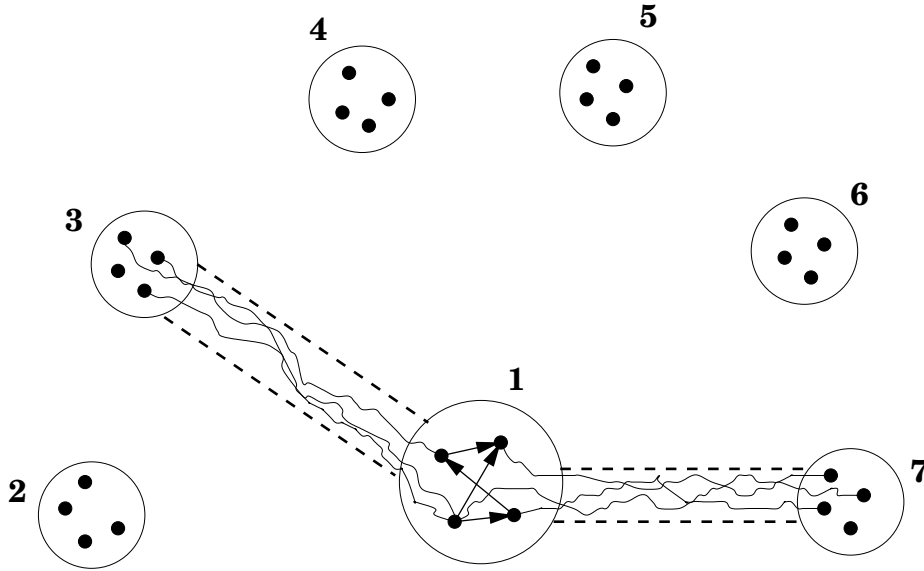
Figure 1: Sketch of the structure of connections in the model. Black dots represent neurons; circles outline modules. The neurons of module 1 connect with each other and receive extramodular axonal projections from modules 3 and 7 only. The long-range axons travel through imaginary channels. For clarity, only the intramodular and the extramodular projections relative to the neurons of module 1 have been drawn. Note that, although the figure includes small numbers of modules and neurons, the model actually assumes large numbers.

Goldman-Rakic (1988) and Romanski et al. (1999) obtained by means of the *double tracing* technique. Indeed, in the dye pictures of Romanski et al. (1999) it seems that reciprocal connections between columns are remarkably more frequent than what one would expect in an extremely dilute directed graph, but also several projections seem to be asymmetrical. To the present knowledge of the author, suitable statistics of inter-columnar connectivity are not yet available.

The structure of the net of inter-modular connections is included in the set of coefficients $\{c_{i_m j_n}\}$ through the factorization

$$c_{i_m j_n} = s_{mn} \cdot k_{i_m j_n}, \tag{7}$$

where: s_{mn} is the *structure* binary variable, whose value is 1 if a channel exists between modules m and n (it does with probability $\frac{s'}{M-1}$), and 0 otherwise; $k_{i_m j_n}$ is equal to 1 if neuron j of module n is pre-synaptic to neuron i of module m, given that module n projects to module m (that is, given that $s_{mn} = 1$; thus, $k_{i_m j_n} = 1$ with probability $\frac{L}{s'N}$ if $s_{mn} = 1$), and is equal to 0 otherwise. In the following, the numbers of the extramodular inputs and of the intramodular ones are equal to each other (Braitenberg and Shutz, 1991); for this assumption, the ratio $\gamma \equiv L/T$ must be equal to $\frac{b}{1+b}$.

Possible topographies in the distribution of the modules through the neocortical sheet and of their connections are not accounted for. In fact, as the association cortices are thought to be at the top level of the information processing hierarchy and to cooperate in a parallel manner (Fuster, 1998; Fuster, 1997; Goldman-Rakic, 1988; Friedman and Goldman-Rakic, 1994; Martin and Chao, 2001), it seems a reasonable approximation to assume non-preferential directions or concentrations of long-range projections between their columns.

2.2   Pattern statistics. Correlation

Every *pattern* stored in the multi-modular network corresponds to a specific distribution of neuronal activity. Accordingly with evidence and hypotheses portrayed in Section 1, any pattern is a peculiar combination of local *features* sparsely distributed over the network, each feature being stored in a different column.

It is assumed that every module can nominally store D features, which the P global patterns are randomly and independently associated to. For any module, as it were, P objects are randomly distributed into D boxes; this distribution is uniform, so the probability for any pattern to fall into any box is $1/D$. P and D are assumed to be large and related to each other by the formula $P = \nu D/\tau$, with ν finite. Only some of the patterns associated to any box actually activate the corresponding feature; as any pattern involves any module with finite probability τ, the probability distribution of the number of patterns eliciting a feature in any module is a Poissonian with mean equal to ν. Hence, any local feature may be involved by several patterns, implying non-invertibility of the mapping in Eq. 3. In any module, some boxes may be not associated with any pattern at all, and some may be associated with patterns that actually do not activate the corresponding features; in fact, the number of features that are actually stored in any module is $D_a \simeq \tau P(1 - e^{-\nu})/\nu$.

One of the basic components of the model is the presence of correlation in the activity of connected modules: *Any two connected modules are simultaneously active or quiescent with probability higher than chance, while the activities of any pair of non-adjacent modules are nearly independent.* Indeed, since modular activity is sparse and every module interacts with only a few other modules ($\sim s'$), it seems convenient that adjacent modules analyse correlated kinds of features and, consequently, can more often transmit useful information to each other, for instance, for the completion of a retrieval task (cf. also Sections 5 and 6). The correlation of the activities of connected modules is represented by the following table of conditional probabilities (Fulvi Mari and Treves, 1998; Fulvi Mari, 2000):

$$\mathcal{P}(\tau_m^p = 1 | \tau_n^p = 1, s_{mn} = 1) = t_1$$

$$\mathcal{P}(\tau_m^p = 1 | \tau_n^p = 0, s_{mn} = 1) = t_0$$

$$\mathcal{P}(\tau_m^p = 1 | \tau_n^p = 1, s_{mn} = 0) = \tau \qquad (8)$$

$$\mathcal{P}(\tau_m^p = 1 | \tau_n^p = 0, s_{mn} = 0) = \tau$$

where $t_1 > \tau$ and $t_0 < \tau$, τ being the probability for any module to be active, introduced earlier. So, the activation state of module m does not depend on that of module n if they are not connected by a channel. Since the structure variables of the graph are symmetric ($s_{mn} = s_{nm}$), Eqs. 8 require that $(1-t_1) \cdot \tau = t_0 \cdot (1-\tau)$. In fact, the table of Eqs. 8 has to be seen as the averaged version of the actual distribution, that includes the effects of connectivity fluctuations across the graph. (Cf. Fulvi Mari (2000) for more details about the probabilistic scheme.)

Proving the existence of at least one joint probability distribution whose averaged marginals are given by Eqs. 8 is crucial in order to consider the present statistical model meaningful. Indeed, not any arbitrary set of marginal probabilities over a family of random variables is *consistent*: there is not complete freedom in choosing *a priori* marginal probabilities on a correlated system. In Fulvi Mari (2000) it is shown that the statistics introduced into the present neural model is consistent if proba-

bility t_1 takes values below an upper-bound that is function of the connectivity s' and of the modular sparseness τ.

Note that the correlation between the activation states of adjacent modules does not arise from dynamics of the network; it comes from the statistics of the 'natural' input patterns, that are composed of features that are not independent from each other. The statistical dependence is reflected by the inter-modular connections, as connected modules are assumed to 'analyse' features that are someway related to each other. It may be conjectured that, during learning, correlated features tend to be stored in adjacent modules.

3   The field and its overlap with locally stored features

Due to analogies with physics models, the total input current h_{i_m} to neuron i_m (Eq. 7) is often called *field*. It is possible to write h_{i_m} as a sum of few simple terms (Appendix A). Noise due to memory load appears in the final expression of h_{i_m} as a set of Gaussian random variables; the other contributions to h_{i_m} are combinations of discrete dynamical variables and quenched random variables.

To understand what state a module is being 'pushed' toward by its neighbours, it is useful to know how much the extramodular input signals correlate with any of the local features stored in the module (cf. Section 5). As a useful indicator of the similarity between extramodular input currents $h_{i_m}^{ext}$ and local feature ξ, the *field-overlap* order-parameter is defined as

$$Q_m^\xi \equiv \frac{\lambda s'}{\gamma(1-a)^2 B}\frac{1}{N}\sum_{i_m}(\eta_{i_m}^\xi - a)h_{i_m}^{ext}. \tag{9}$$

It is convenient to define the binary variables $\{\varphi_m\}$: φ_m is equal to 1 if module m is correctly quiescent or retrieving the feature that corresponds to the pattern that should be retrieved by the multi-modular network (conventionally assumed to be pattern #1), and is equal to 0 otherwise. The overlap between the inputs to the neurons of module m and local feature ξ stored in the same module then is

$$
\begin{aligned}
Q_m^\xi &= \tau_m^1\delta[1-\xi]\sum_{n\neq m}\varphi_n s_{mn}\tau_n^1 + \sum_{n\neq m}\varphi_n s_{mn}\tau_n^1\sum_{p\in d_n(1),p\neq 1}\tau_m^p\tau_n^p\delta[d_m(p)-\xi]+\\
&+ \sum_{n\neq m}(1-\varphi_n)s_{mn}\tau_n^{p_n}\sum_{p\in d_n(p_n)}\tau_m^p\tau_n^p\delta[d_m(p)-\xi].
\end{aligned}
\tag{10}
$$

where $\delta[i-j]$ is the Kronecker symbol for integer variables i and j. The Gaussian terms in $h_{i_m}^{ext}$ (cf. Appendix A) do not contribute to Q_m^ξ due to self-averaging (for $N\to\infty$). It is useful to the rest of the paper to introduce the following definitions:

$$\hat{u} \equiv \sum_{p\in d_m(1),p\neq 1}\tau_m^p, \quad \hat{X} \equiv \sum_{n\neq m}\varphi_n s_{mn}\tau_n^1, \quad \hat{v} \equiv \sum_{n\neq m}(1-\varphi_n)s_{mn}\tau_n^{p_n},$$

$$\hat{f} \equiv \sum_{n\neq m}(1-\varphi_n)s_{mn}\tau_n^{p_n}\sum_{p\in d_n(p_n)}\tau_m^p\tau_n^p, \quad \hat{r} \equiv \sum_{n\neq m}\varphi_n s_{mn}\tau_n^1\sum_{p\in d_n(1),p\neq 1}\tau_m^p\tau_n^p,$$

$$\tag{11}$$

where $p\in d_n(p_n)$ indicates any pattern p that elicits in module n the same feature as pattern p_n, that is being locally reproduced, does. The random variables $\hat{r}$ and $\hat{f}$ come respectively from the last two terms of Eq. 10 being summed over the features $\xi\neq d_m(1)$ stored in module m. The usefulness of

9

the definitions above will appear in the next Sections, and their interpretation is straightforward: for example, $\hat{X}$ is the number of modules pre-synaptic to module m that are in the correct local states (correspondingly to the global pattern #1 to be retrieved) and are actively retrieving local features, while $\hat{v}$ is the number of neighbours that are actively retrieving wrong features. The variables $\hat{r}$ and $\hat{f}$ may have relevance in semantic association errors; for example, $\hat{r}$ relates to features that appear in the neighbours of module m in stored patterns other than pattern #1 simultaneously with feature $d_m(1)$ of module m.

4   Signal-to-noise analysis

When the number of stored patterns is no longer negligible with respect to the number of connections per neuron, back-ground memories can affect the quality of the retrieval, or even make it impossible, by producing noisy input to every neuron, randomly distributed across the network according to a Gaussian density function. The standard deviation (σ) of this noise comes from the calculation of the input field h_{i_m} and is reported in Appendix A. There are in fact non-Gaussian noisy terms besides the memory-load ones, as a rapid inspection of the field in Appendix A reveals (they are given by sums of a relatively small number of binary random variables). However, due to the very small value of the sparseness a and to the values adopted here for the other parameters, only a minor fraction of the neurons in any active module is significantly affected by these non-Gaussian terms. Hence, the latter are neglected in evaluating the stability of local retrieval attractors. It should be noted that those terms, being correlated to local features, are indeed the responsible ones for biasing attractor switching (cf. Section 5). Therefore, in the present approximation, the extramodular neuronal noise may be considered as the sum of two components: an unstructured (memory-load, Gaussian) one, tending to destabilize local retrieval attractors but unable to produce active local states, and a structured one, of little effect in destabilizing local attractors but able to bias modular activation.

The effect of memory-load noise on the stability of local retrieval attractors can be estimated through a standard *signal-to-noise* analysis. For example, be considered the case in which the network is reproducing a pattern, possibly with some errors or wrong activation of some modules, and be considered a neuron belonging to a foreground module. One component of the field tends to keep the neuron in the correct firing state, and can thus assume two values according to whether the neuron is in foreground or in background; the difference between these two modal values is called *signal*. The remaining component of the input field is noisy and tends to disturb the firing level of the neuron. Assuming that the neuron has essentially to distinguish between being in foreground or in background, one can estimate the stability of the state of the module looking at the ratio between the signal and the standard deviation of the noise:

$$\frac{\mathcal{S}}{\mathcal{N}} = \frac{1 + \hat{u} + \frac{1}{\lambda s'}\hat{\Gamma}}{\sqrt{\frac{\alpha a \tau}{\gamma}}\sqrt{\frac{a(B+C)^2 + (1-a)C^2}{aB^2}(1+\nu) + \frac{t_1}{\lambda^2 s'}\left[\frac{a(B+C)^2 + (1-a)C^2}{aB^2}(\hat{\Gamma} + \hat{\Xi}) + \frac{(A+C)^2}{aB^2}(\hat{k} - \hat{\Gamma} - \hat{\Xi})\right]}}, \quad (12)$$

where $\hat{k}$ is the number of neighbours of the module that contains the given neuron, $\alpha \equiv P/T$ is the memory-load parameter, $\hat{u}$ is defined in Eqs. 11, and $\hat{\Gamma}$ and $\hat{\Xi}$ are, respectively, the number of *supporting*[5] neighbours and the number of the other active neighbours. Once the signal-to-noise ratio

---

[5] Any neighbour of module m is here said to be supporting module m if the two modules are retrieving local features that appeared simultaneously in at least one of the stored patterns.

is available, since noise is Gaussian, one can easily estimate the fraction of neurons that receive inputs nearer to the correct modal values than to the wrong ones. If the ratio is too small, this fraction does not differ significantly from the chance level 0.5, while if the ratio is large the fraction approaches 1. In order to evaluate the stability of any active module, it is chosen to fix a threshold θ: if the fraction of stable neurons is not smaller than θ, then the active module is assumed to be in a stable attractor; otherwise, the module decays to quiescence. It is a crude model of destabilization, but not so much for what concerns the sharpness of the threshold; indeed, in many models of Hebbian neuronal network the transition from stability to instability as a function of the noise level is quite abrupt. The probability for any active module to be unstable will be indicated by $H(\hat{k}, \hat{u}, \hat{\Gamma}, \hat{\Xi})$, in the present case just equal to either 0 or 1.

The necessity of $\lambda < 1$ has already been commented in Section 2. Here, it also appears that correlation between adjacent modules is considerably useful (especially for small τ): indeed, if t_1 was equal to τ, then $\hat{X}$ would be different from zero for a significantly smaller fraction of the modules, with the consequence of smaller contribution of extramodular signals to retrieval stability and poorer ability of cue completion (cf. also Sections 5 and 6).

5 Retrieval Dynamics

The dynamical transitions between modular states are defined as follows:

- If a module is quiescent, it can be elicited to retrieve local feature ξ by only an extramodular input field that overlaps with that feature ($Q^\xi > 0$) and only if the new state would be stable (according to the signal-to-noise ratio; cf. Section 4). If the extramodular field overlaps with one only feature whose retrieval would be stable, then the module will retrieve that feature. If $Q^\xi > 0$ for several features ξ whose retrieval would be stable, then the module will retrieve the feature with the largest field-overlap among them (in the case of multiplicity of the maximizers, the module will retrieve a feature randomly chosen among the latter).

- If a module is in an active state, correct or not, it is destabilized and moved to quiescence if the signal-to-noise ratio falls below a certain threshold (cf. Section 4).

Because of the way the features are distributed across the modules, the probability for any pair of features of any pair of modules to be simultaneously present in a pattern is vanishingly small ($\propto 1/D$); consequently, it is very unlikely that the last two terms of the field-overlap (Eq. 10) contribute more than one unit for any feature ξ stored in the post-synaptic module m and different from the correct local feature $d_m(1)$ that should be retrieved. As modules in correct active states always cooperate in pushing a common post-synaptic module toward the correct state, these observations may suggest that it could be useful to have a field-overlap *threshold* with value equal to 2: if a quiescent foreground module receives a total input field whose overlap with any local feature is not larger than 1, then the module should not move from quiescence. A simple consequence of this assumption is that the activation of wrong features is mostly inhibited. However, when the cue is only a small fraction of the pattern to retrieve, a field-overlap threshold larger than 1 could also heavily spoil correct retrieval spread; in fact, simulations show that the performance with the field-overlap threshold equal to 2 is very poor already for cues consisting of $\varrho = 5\%$ of the pattern to retrieve (cf. following, especially Figure 3 and related text).
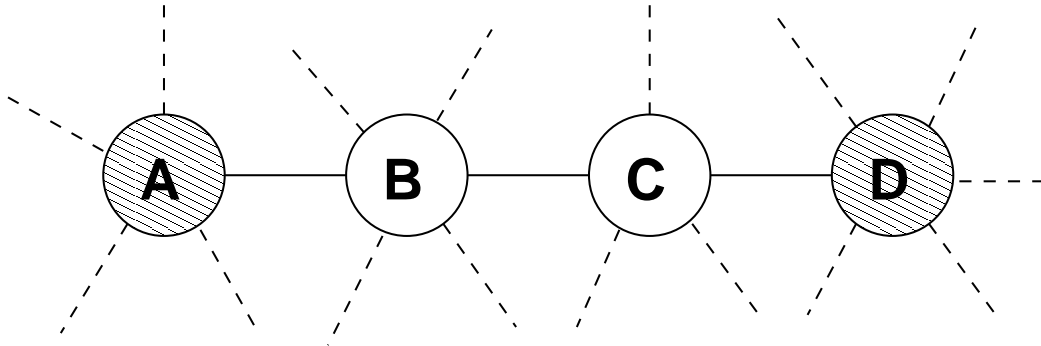
11

Figure 2: Illustration of a simple example of retrieval spreading. Modules A and D are in correct feature retrieval (shaded). Modules B and C are quiescent but should be active in the pattern to retrieve. Cf. Section 5 for details.

If the cue stimulus is small, a low field-overlap threshold would facilitate the spreading of the correct retrieval, but would also allow wrong activation to spread and possibly self-organize in retrieving a pattern different from the one the cue is extracted from. Here, it is studied the possibility of an *oscillatory* process: The field-overlap threshold is equal to 1 and the robustness of local attractors to noise periodically oscillates (by modulation of the neuronal threshold θ; cf. Section 4) in such a way that every time-step in which any active module with one only supporting neighbour is stable is followed by a time-step in which it is unstable. In this way, all the modules that have been wrongly activated during a high-robustness time-step are quickly put to quiescence in the next time-step; this is possible as any module that is in a wrong feature retrieval is most likely to be supported by no more than one of its neighbours, the same one that drove the module from quiescence to wrong activity. The two-steps oscillation is iterated several times. If every low-robustness time-step was followed by two or more high-robustness time-steps, many of the modules in wrong retrieval would elicit further wrong activity and, consequently, would have more supporting neighbours and gain in robustness. During all the retrieval process, any quiescent foreground module with two or more neighbours in correct active state is driven to the correct local retrieval attractor (if it would be stable). A sample case of correct retrieval spread is pictured in Figure 2: Modules A and D are in correct feature retrieval, while modules B and C are quiescent but should be active. During a high-robustness time-step, modules A and D elicit retrieval of features in modules B and C (for clarity, the influence of other neighbours to B and C is neglected). If any of the two features activated, respectively, in B and C is not correct, each of them only has support from one neighbour, that is, respectively, A or D. In this case, during the following time-step both modules B and C are destabilized and decay to quiescence. On the contrary, if the activity elicited in B and C corresponds to correct feature retrieval, modules B and C will support each other and still hold the support of, respectively, module A and module D. Thus, in the following low-robustness time-step both of them will be supported by two neighbours and, consequently, will not be destabilized.

TRANSITION	PROBABILITY
$ca \to w\bar{a}$	$H(\hat{k}, \hat{u}, \hat{X} + \hat{v}, 0)$
$wa^1 \to w\bar{a}$	$H(\hat{k}, \hat{u}, 1, \hat{X} + \hat{v} - 1)$
$w\bar{a} \to ca$	$\left[\Theta(\hat{X} - 2) + \eta \delta_{\hat{X},1} \frac{1}{\hat{r}+\hat{f}+1} \right] \left[1 - H(\hat{k}, \hat{u}, \hat{X}, \hat{v}) \right]$
$w\bar{a} \to wa^1$	$\eta \left[\delta_{\hat{X},0} \Theta(\hat{r} + \hat{f} - 1) + \delta_{\hat{X},1} \frac{\hat{r}+\hat{f}}{\hat{r}+\hat{f}+1} \right] \left[1 - H(\hat{k}, \hat{u}, 1, \hat{X} + \hat{v} - 1) \right]$
$c\bar{a} \to wa^0$	$\eta \Theta(\hat{r} + \hat{f} - 1) \left[1 - H(\hat{k}, \hat{u}, 1, \hat{X} + \hat{v} - 1) \right]$
$wa^0 \to c\bar{a}$	$H(\hat{k}, \hat{u}, 1, \hat{X} + \hat{v} - 1)$

Table 1: Transitions of modular state and respective probabilities, conditional to the states of the neighbours through the variables $\hat{X}$, $\hat{v}$, $\hat{r}$, $\hat{f}$. The Heaviside step-function Θ is here defined to satisfy $\Theta(0) = 1$; $\delta_{i,j}$ is the Kronecker symbol for the integer variables i and j.

Although (to the knowledge of the author) there is no specific experimental data about such kind of oscillatory process (though literature on brain waves is wide)[6], there is already some evidence on neuromodulatory mechanisms that could regulate the robustness of local (columnar) attractors to extramodular disturbing signals (Durstewitz et al. (2000) and references therein).

The probability for any foreground module to have two or more neighbours also in foreground in any pattern is

$$\psi \simeq 1 - (1 + s't_1)e^{-s't_1}; \tag{13}$$

therefore, in static conditions with constant low robustness the retrieval quality cannot be larger than ψ. To increase this upper-bound, one can set t_1 and s' so that $s't_1$ is as large as possible, convening with the correlation bound (as already mentioned, correlation cannot be arbitrarily large, and its upper-bound monotonically decreases as s' increases (Fulvi Mari, 2000)). The modular sparseness τ also plays a role in this, as larger τ allows for larger t_1 given s', within certain limits and biological plausibility. In this view, a good choice seems to be $\tau = 0.1$ and $s' = 8$, that allows for t_1 as large as 0.4, at least.

All the possible state transitions of any module are reported in Table 1 together with the respective probabilities. Factor η takes value 0 or 1 if the field-overlap threshold is equal, respectively, to 2 or 1. The following abbreviations have been introduced: w stays for 'wrong', c stays for 'correct', a

---

[6]Merely as an aside, it could be noted that, starting with cue consisting of $\sim 5\%$ of the pattern to retrieve, about 10 cycles are sufficient to reach reasonable levels of retrieval quality (Figure 3: time-steps from #20 to #40). At oscillation frequencies of about 40 Hz (in the range of the brain γ-waves, associated to performance of cognitive tasks), the process would require about 250 ms, a plausible time scale.
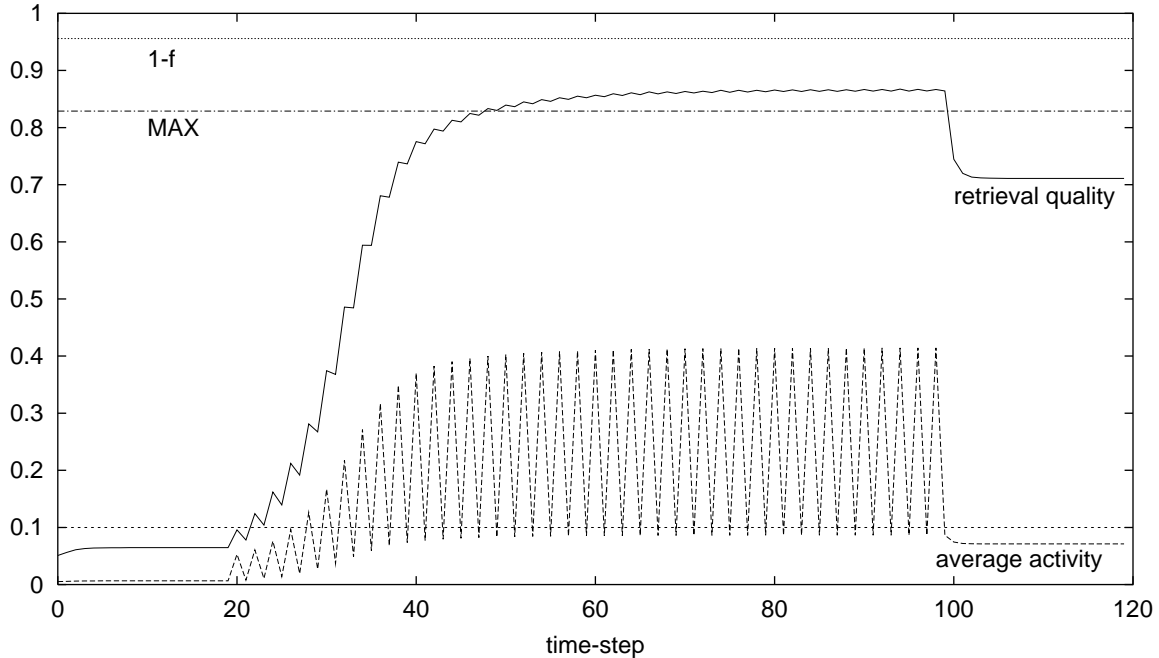
Figure 3: Temporal evolution of retrieval quality and average activity in cued retrieval simulations. The cue consists of $\sim 5\%$ of the pattern to retrieve. The line "MAX" is the upper-bound ψ to retrieval quality, while "1-f" is the upper-bound for the case of no feature-sharing (Section 6). A line corresponding to the average modular activity in the stored pattern (τ) is also shown for reference. Parameters: $\tau = 0.1$; $t_1 = 0.4$; $s' = 8$; $b = 0.2$; $M = 400,000$; $a = 0.027$; λ=0.02; $\nu = 10$.

stays for 'active', $\bar{a}$ for 'non-active', and the upper indices 1 and 0 indicate, when necessary, whether the considered module should be, respectively, active or silent in the global pattern that should be retrieved. For example, $ca \rightarrow w\bar{a}$ indicates the transition of a foreground module from the correct active state to wrong quiescence, while $w\bar{a} \rightarrow wa^1$ indicates a transition from wrong quiescence to a wrong active state. The dynamics is implemented in numerical simulations, with parallel update (some details on the simulations are reported in Appendix B).

In order to appropriately destabilize spurious activation (cf. Section 4), the memory-load parameter α must belong to a certain interval of values. (It should be noted that, although the only noise here considered comes from storage load, it could also be, alternatively or complementarily, of other origins.) A major hindrance is the large variability of the noise level due to the large fluctuation of the number of neighbours across the set of modules (this is a consequence of having adopted a Bernoulli random graph as model of the modular net). Indeed, for example, a noise level that destabilizes the retrieval activity of a module that has one supporting neighbour and a total of 4 neighbours is also likely to destabilize the retrieval activity of a module that has 2 supporting neighbours and a total of 8 or more neighbours. The necessity of a value for λ as small as 0.02 is also partly a consequence of this drawback (limiting the retrieval process to the oscillatory stage would allow for larger values of λ, to some extent). In order to reduce this noxious effect, when testing the oscillatory retrieval process it is chosen to make modules with less than 5 neighbours inactive, that is, they are kept quiescent
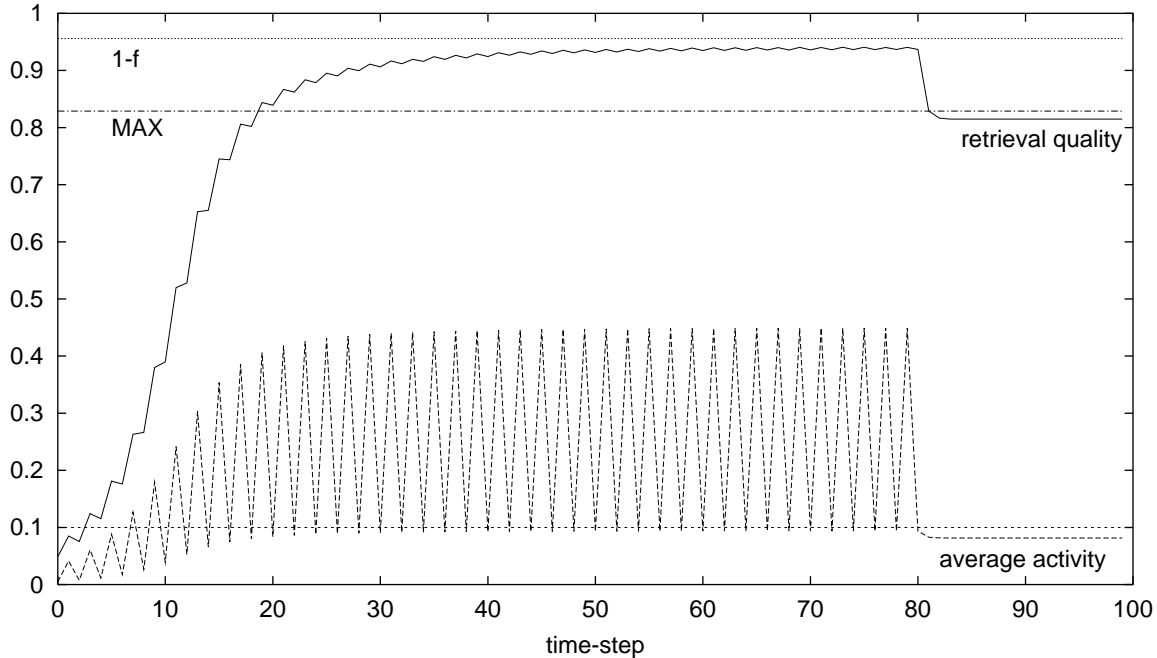
14

Figure 4: Temporal evolution of retrieval quality and average activity in cued retrieval simulations with 'artificial' dynamics (Section 5) and no modular freezing. The cue consists of 5% of the pattern to retrieve. The line "MAX" is the upper-bound ψ, while "1-f" is the upper-bound for the case of no feature-sharing (Section 6). A line corresponding to the average modular activity in the stored pattern (τ) is also shown for reference. Parameters: $\tau = 0.1$; $t_1 = 0.4$; $s' = 8$; $M = 400,000$; $\nu = 10$.

regardless of the cue and of the inputs they receive from the neighbouring modules; as $s' = 8$, the inactive modules account for slightly less than 10% of the network. It is possible that the neuronal model adopted here is too crude to counterbalance the fluctuations of noise level by other means, but it also seems biologically plausible that the number of columns afferent to any other column should not vary with so large a range as in the present model. A more plausible architecture is under study[7].

Naturally, noise also affects every foreground module that has one only neighbour in foreground; hence, some spots will still not be percolated by correct activity (similarly to the *activity isles* of Section 6), but overall performance is anyway remarkably improved by the oscillatory modulation. This is evident in Figure 3, where "retrieval quality" is the overlap between the features reproduced by the foreground modules and the pattern to retrieve, while "average activity" is the fraction of modules that are active in the network. The cue consists of $\varrho = 5\%$ of the pattern that should be retrieved[8] (smaller cues are equally suitable, to some extent, at the expenses of longer oscillatory stage). During the first 20 time-steps, the field-overlap threshold is equal to 2 and robustness to noise is high; very little retrieval from memory is produced. This first stage does not belong to the proposed retrieval process and has been included in the simulation just to show how poor the performance would be if

---

[7]It also requires proving consistency of marginal distributions and studying how to produce patterns with the correlated statistics in the new kind of graphs.

[8]As almost 10% of the modules are frozen to quiescence, effectively the cue is about 4.5% of the pattern.

the field-overlap threshold was equal to 2; removing this stage does not affect the results. At time-step #20, the threshold is lowered to 1 and the oscillatory phase begins; the retrieval quality quickly increases (when the cue is very small, e.g. $\varrho = 0.01$, the duration of the oscillatory stage necessary to reach the plateau can vary significantly, depending on the random choice of modules that are cued; this variability is due to finite size and becomes soon negligible when larger cues are provided, as for $\varrho = 0.05$). From time-step #100 onward, the robustness to noise is constantly at the lower level: the retrieved information is stable and there is no spurious activity left (e.g., the average activity equals the activity of the modules in correct retrieval), and no further dynamical transition happens. It is simple to prove that, given any foreground module in correct state, if one or more of its neighbours become active, then the given module always gains in robustness, regardless of the correctness of the new active states of the neighbours. As a consequence, the retrieval quality can overcome the upper-bound ψ (Eq. 13) during the oscillatory phase (arriving not too far, in fact, from the optimal upper-bound $1 - f$ calculated in Section 6), falling below ψ in the final stage. The final retrieval quality is lower than ψ possibly also due to small regions that cannot be percolated by the oscillatory process, but mainly due to the freezing of modules with less than 5 neighbours and to the variability of the noise level as a function of the number of neighbours. Simulations have been also carried out in which the noise of neuronal/synaptic origins is neglected and the dynamics 'artificially' follows the rule that every module with one only supporting neighbour is unstable during the low-robustness time-steps, while it is stable during the high-robustness time-steps, regardless of the signal-to-noise ratio: the retrieval quality during the oscillatory stage is not improved by this proviso, while the final value is ($+5\%$, still below the upper-bound ψ). Of course, removal of modular freezing, no longer necessary in the artificial case, provides significant improvement in both the stages of the retrieval process (Figure 4; oscillatory phase since the beginning and until time-step #80).

Figure 5 shows the fraction of modules in correct state at the end of the retrieval process of Figure 3, conditionally to the number of neighbours $\hat{k}$. Variables φ^1 and φ^0 are the fractions of, respectively, foreground and background modules that are in correct state, while τ' is the fraction of active modules among the foreground modules that are in wrong states. The non-monotonicity of the histogram φ^1 is due to the dependence of the noise level on $\hat{k}$. The histograms of φ^0 and τ' confirm that there is not spurious activity remaining.

In order for the oscillatory process to work properly, it is necessary that active modules are nearly as noisy as the quiescent ones. This condition is realized if

$$a \approx \frac{A}{B} \cdot \frac{A/B + 2C/B}{1 + 2C/B}, \tag{14}$$

which, with the values of A/B and C/B here adopted, implies $a \approx 0.027$, well compatible with experimental evidence.

When ϱ is large, for example 50% or more, then the monotonic process with field-overlap threshold equal to 2 and low noise level can achieve larger retrieval quality than the final stage of Figure 3 (though not larger than the value during the oscillatory stage), as cued modules that have each one supporting neighbour only do not decay to quiescence.

To hamper wrong activation spreading, it may be useful to distribute the features in such a way that modules with larger number of neighbours code for local features that are more rarely shared among different patterns; this proviso would decrease the biasing terms with $\hat{r} + \hat{f}$ and hence would make wrong activation more unlikely. This possibility is not investigated further in the present paper, and is left for future, more detailed simulations (cf. Appendix B) that should also investigate on the
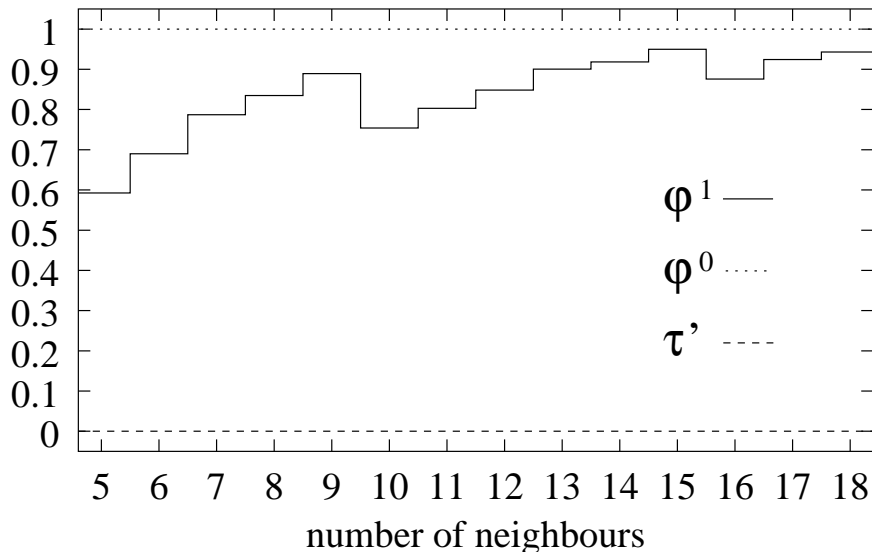
Figure 5: Probability for any foreground or background module to be in the correct state (indicated, respectively, by φ^1 and φ^0) and for any foreground module that is in a wrong state to be active (indicated by τ') at the end of the simulation of Figure 3, as a function of the number of neighbours. (Modules with less than 5 neighbours are frozen to quiescence.)

possibility of taking advantage of the fact that, as ν is considerably larger than 1, wrong activity should take longer to self-organize than correct retrieval activity, that is also advantaged by starting with already a large number of coherent modular retrievals.

6 No feature-sharing. Activity isles

As clear from the previous Sections, the main obstacle to obtain proper memory retrieval in the large modular autoassociator is given by feature-sharing: all the wrong activity generated after a 'clean' cue is due to the fact that several patterns can activate any feature in any module. For comparison, Figure 6 shows the outcomes of a simulation of a network in which every local feature is active in one only global pattern (the dynamics is very simple, as there is no possibility of spurious activations). As the graph shows, the overlap of the activity with the stored pattern rapidly reaches a value slightly above 95%. In random graph theory it is well established that for $s' > 1$ a finite fraction of the number of modules are directly or indirectly connected with each other, forming the so-called *giant component*; the rest form a large number of small, isolated groups. The giant component in a large graph with $s' = 8$ accounts for more than 99.96% of the network; thus, the reason of the imperfect retrieval is not the presence of isolated groups of modules. The reason for the phenomenon is in fact that in any stored pattern a certain number of active modules are surrounded by modules that are not active. Given any pattern, let any connected group of foreground modules that are surrounded by only background modules be called *activity isle*. Many of the isles are not elicited by the small cue, and thus cannot be driven to correct retrieval activity. Be considered any n modules, and be considered the case in which they form a connected group and all of them should be active but none of them
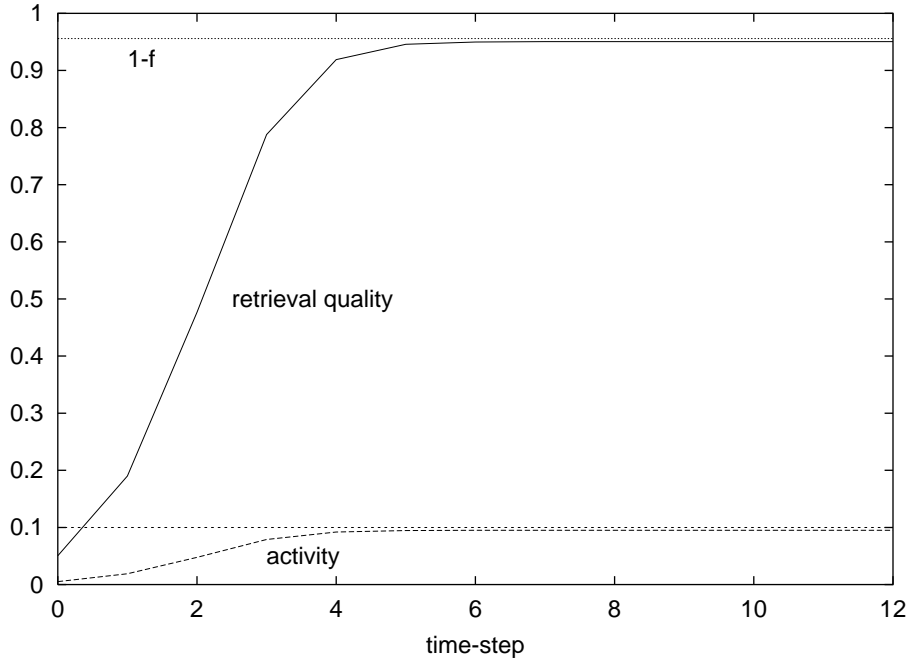
17

Figure 6: Temporal evolution of retrieval quality and average activity in cued retrieval simulations in the case in which every stored feature appears in one only pattern. The cue consists in 5% of the pattern to retrieve. The line "1-f" gives the maximum retrieval quality permitted by the model (cf. Section 6). Parameters: $\tau = 0.1$; $t_1 = 0.4$; $s' = 8$; $M = 400,000$.

is elicited by the cue, and all their neighbours but those in the group are in the background of the pattern to be retrieved; the probability that all these occurrences happen in any group of n modules, that is, the probability that any group of n modules is a non-cued activity isle, can be estimated (for $n \ll M$); the fraction of foreground modules that belong to non-cued isles of n modules results to be

$$f_n \simeq (s't_1)^{n-1}(1 - \varrho)^n e^{-ns't_1}. \tag{15}$$

The value of f_n decreases very quickly as n grows; with the values of the parameters adopted in the simulations, one has $f = f_1 + f_2 + f_3 + f_4 \simeq 0.044191$, being $f_4 \simeq 0.000074$. In fact, simulations show (cf. Figure 6) that the retrieval quality tends asymptotically to the value $1 - f$ (but for statistical fluctuations due to finite size), corresponding to correct local retrieval in all modules but those in the activity isles that are not elicited by the cue.

7 Conclusions

Anatomical and physiological data suggest that association areas of the neocortex could be modelled as a large number of autoassociators (the columns), each of which interacts with a small fraction of the others (Sections 1 and 2). If the global memory patterns of neuronal activity are composed of peculiar combinations of local features, each feature being stored in a module and shared by several

18

patterns, then it becomes important to understand how the multitude of modules is structurally and functionally organized in order to be able to perform proper memory retrieval when a cue is briefly presented to the network as an external stimuli. One of the major difficulties to overcome in the modelling is that every feature is involved is several patterns and, thus, is Hebbianly associated with several features stored in any of the adjacent modules. This implies that, when feature d_m is active in module m, the latter could move a neighbouring module n to the local attractor of one of the features stored in n and associated with d_m during learning, but not necessarily the correct one for what concerns the retrieval of the pattern corresponding to the cue.

Investigating retrieval mechanisms that could achieve good retrieval quality without producing spurious activity, an oscillatory retrieval process is found to be particularly efficient (Section 5). It requires: a modulatory mechanism that periodically modifies the robustness of local attractors to noise; neuronal activity sparseness such that active modules are nearly as noisy as the quiescent ones; number of neighbours that fluctuates from one module to another across the network significantly less than the degree of the nodes in Bernoulli random graphs. In order to achieve satisfactory retrieval abilities, the activity of adjacent modules is correlated: across the set of stored patterns, any pair of modules adjacent to each other are simultaneously active or quiescent more often than chance (cf. also Section 2.2). Besides, the extramodular contacts have to be more effective and/or faster in learning than the intramodular ones (cf. also Section 2).

It is found that some spots of the network, depending on pattern to retrieve and choice of the cued fragment, cannot be reached by percolating retrieval activity. The oscillatory mechanism nearly saturates the optimal upper-bound to retrieval quality (Sections 5 and 6).

The simpler case in which local features are in one-to-one correspondence with global patterns is also discussed (Section 6). The retrieval dynamics in this case is simpler, and retrieval performance better, but the putative 'cognitive' advantage of reusing stored features for several patterns is lost, of course. In fact, during the oscillatory stage the retrieval quality of the model with feature-sharing is not too far below (in the 'artificial' case, almost equal to) the best value of the model without feature-sharing.

The assumptions on the architecture of the connections and on the neurophysiological parameters used in the construction of the model seem in accord with available experimental data. The constraints necessary to reproducing proper memory retrieval also seem compatible with biological evidence.

It should be noted that the network of modules defined in this paper holds the *small world* property (that is, the average length of the shortest path joining any two modules on the graph only scales as the logarithm of the number of modules, as for any Bernoulli random graph), but no clustering is present. At the level of the neurons, the network holds both small-world and clustering properties, as the connectivity inside any column is finite. (For definitions and recent results on small-world effect, cf., e.g., Newman (1999) and references therein.) It would be interesting to include clustering in the net of modules too, in order to model topological nearness of columns and then study how relevant it is to what concerns memory retrieval in associative areas of the neocortex.

A further development of the present work could be the inclusion of effects from the remains of a previously retrieved pattern on the current process of retrieval, or the similar case of cues composed of fragments from more than one pattern. If it is allowed that fragments of a previous retrieval are still present in the network during a new cue, then they may cooperate and produce relevant consequences, for example favouring the retrieval of a global pattern that has larger correlation with the one previously retrieved (e.g., *priming* effect).

It seems of great importance to study in deeper details the properties of recurrent neuronal net-

19

works in switching from an attractor to another one, how effective the input must be in order to cause the mentioned transition, and the stability of the quiescence attractor. In the present work the firing-rate neuron model has been used, but it would be important the adoption of a spiking neuron model. In particular, it seems extremely interesting to investigate whether the effects of the oscillatory modulation introduced in this paper could instead be obtained by appropriate timing of spiking activity.

Appendix A

Neuronal field h_{i_m} and memory-load noise

Summing together the input synaptic currents of neuron i_m, one obtains:

$$
\begin{aligned}
h_{i_m} &= \varphi_m B \tau_m^1 \left(\tfrac{\eta_{i_m}}{a} - 1\right)(1-\gamma)b(1-a)\left[1 + \sum_{p \in d_m(1), p \neq 1} \tau_m^p\right] + \\
&+ \varphi_m \hat{x}\left[\left(\tfrac{1-a}{a}\right)^2 \alpha\tau(1-\gamma)b(1+\nu)\left\{\left[a(B+C)^2 + (1-a)C^2\right]\tau_m^1 + (A+C)^2(1-\tau_m^1)\right\}\right]^{1/2} + \\
&+ (1-\varphi_m)B\tau_m^{p_m}\left(\tfrac{\eta_{i_m}^{p_m}}{a} - 1\right)(1-\gamma)b(1-a)\left[1 + \sum_{p \in d_m(p_m), p \neq p_m} \tau_m^p\right] + \\
&+ (1-\varphi_m)\hat{x}'\left[\left(\tfrac{1-a}{a}\right)^2 \alpha\tau(1-\gamma)b(1+\nu)\left\{\left[a(B+C)^2 + (1-a)C^2\right]\tau_m^{p_m} + (A+C)^2(1-\tau_m^{p_m})\right\}\right]^{1/2} + \\
&+ \tfrac{\gamma}{\lambda s'}(1-a)B\tau_m^1\left(\tfrac{\eta_{i_m}}{a} - 1\right)\sum_{n \neq m}\varphi_n s_{mn}\tau_n^1 + \\
&+ \sum_{n \neq m}\varphi_n s_{mn}\hat{y}_n\left[\tfrac{\alpha\gamma}{s'}\tfrac{\tau t_1}{\lambda^2}\left(\tfrac{1-a}{a}\right)^2\left\{\left[a(B+C)^2 + (1-a)C^2\right]\tau_n^1 + (A+C)^2(1-\tau_n^1)\right\}\right]^{1/2} + \\
&+ \tfrac{\gamma}{\lambda s'}(1-a)B\sum_{n \neq m}\varphi_n s_{mn}\tau_n^1\sum_{p \in d_n(1), p \neq 1}\tau_m^p\tau_n^p\left(\tfrac{\eta_{i_m}^p}{a} - 1\right) + \\
&+ \tfrac{\gamma}{\lambda s'}(1-a)B\sum_{n \neq m}(1-\varphi_n)s_{mn}\tau_n^{p_n}\sum_{p \in d_n(p_n)}\tau_m^p\tau_n^p\left(\tfrac{\eta_{i_m}^p}{a} - 1\right) + \\
&+ \sum_{n \neq m}(1-\varphi_n)s_{mn}\hat{z}_n\left[\tfrac{\alpha\gamma}{s'}\tfrac{\tau t_1}{\lambda^2}\left(\tfrac{1-a}{a}\right)^2\left\{\left[a(B+C)^2 + (1-a)C^2\right]\tau_n^{p_n} + (A+C)^2(1-\tau_n^{p_n})\right\}\right]^{1/2}.
\end{aligned}
$$

The memory-load noise terms (containing the normal variables $\hat{x}$, $\hat{x}'$, $\hat{y}_n$, $\hat{z}_n$) can be unified in virtue of independence properties, resulting in a Gaussian term with variance

$$
\begin{aligned}
\sigma^2 &= \tfrac{\alpha\gamma t_1 a\tau}{s'\lambda^2}B^2\left(\tfrac{1-a}{a}\right)^2\left\{b\tfrac{1-\gamma}{\gamma}\tfrac{\lambda^2 s'}{t_1}(1+\nu)\left[\tfrac{a(B+C)^2 + (1-a)C^2}{aB^2}\tau_m^{p_m} + \tfrac{(A+C)^2}{aB^2}(1-\tau_m^{p_m})\right] + \right. \\
&\left. + \tfrac{a(B+C)^2 + (1-a)C^2}{aB^2}(\hat{X} + \hat{v}) + \tfrac{(A+C)^2}{aB^2}(\hat{k} - \hat{X} - \hat{v})\right\},
\end{aligned}
$$

where the abbreviations introduced in Section 3 (Eqs. 11) are used.

Appendix B

Simulations

Simulations were carried out on a Linux-operated PC. The first step required for the simulations is the construction of the random graph underlying the extramodular connections; this is simply achieved by randomly drawing a bond between any pair of modules with probability $s = s'/(M-1)$. Then, it is necessary to produce patterns of activity that obey the statistics of Eqs. 8; to do this, use is made of the method described in Fulvi Mari (2000). Once these sets of quenched random variables are available, the analytical framework presented in this paper allows one to simulate the dynamics of very large networks in reasonable CPU time. This is particularly effective when one simplifies the expressions of the random variables $\hat{r}$ and $\hat{f}$ further. Indeed, without further simplification, one should actually distribute the features across the modules for every one of the patterns of a long sequence and keep memory of that, and at every time step evaluate $\hat{r}$ and $\hat{f}$. This is of course a problem in terms of memory space and CPU time. Thus, here it is chosen to average the quantities that enter the formulae of $\hat{r}$ and $\hat{f}$ and concern the statistics of patterns other than the one to retrieve. Namely, $\hat{r}$ and $\hat{f}$ are given the values, respectively, $nint\{\hat{X} \cdot [\nu t_1/(1-e^{-\nu/\tau}) - \tau t_1]\}$ and $nint\{\hat{v} \cdot [t_1 + \nu t_1/(1-e^{-\nu/\tau}) - \tau t_1]\}$, where $nint\{\cdot\}$ is the nearest-integer function. In this way, one only has to discriminate the state of any module among correct local retrieval, wrong local retrieval, and quiescence, without having to specify what particular feature belonging to what particular pattern is being retrieved in the module. For simplicity as well, fluctuations of the number of patterns that in any given module activate the same feature are also neglected. Simulations that take into account more of the details of the distribution and association of features across the network are planned for future work.

References

Albright, T. D., Desimone, R. and Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque, *J. Neurophysiol.* **51**: 16–31.

Braitenberg, V. and Shutz, A. (1991). *Anatomy of the Cortex: statistics and geometry*, Springer-Verlag, Berlin Heidelberg.

Buckner, R. L. and Petersen, S. E. (1996). What does neuroimaging tell us about the role of prefrontal cortex in memory retrieval?, *Semin. Neurosc.* **8**: 47–55.

Durstewitz, D., Seamans, J. K. and Sejnowski, T. J. (2000). Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex, *J. Neurophysiol.* **83**: 1733–1750.

Friedman, H. R. and Goldman-Rakic, P. S. (1994). Coactivation of prefrontal cortex and inferior parietal cortex in working memory tasks revealed by 2DG functional mapping in the rhesus monkey, *J. Neurosci.* **14**: 2775–2788.

Fujita, I., Tanaka, K., Ito, M. and Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex, *Nature* **360**: 343–346.

Fulvi Mari, C. (2000). Random fields and probability distributions with given marginals on randomly correlated systems: a general method and a problem from theoretical neuroscience, *J. Phys. A: Math. Gen.* **33**: 23–38.

Fulvi Mari, C. and Treves, A. (1998). Modeling neocortical areas with a modular neural network, *Biosystems* **48**: 47–55.

Funahashi, S., Bruce, C. J. and Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex, *J. Neurophysiol.* **61**: 331–349.

Fuster, J. M. (1995). *The prefrontal cortex: anatomy, physiology and neuropsychology of the frontal lobe*, Raven Press, New York.

Fuster, J. M. (1997). Network memory, *Trends in Neuroscience* **20**(10): 451–459.

Fuster, J. M. (1998). Distributed memory for both short and long term, *Neurobiol. Learn. Mem.* **70**: 268–274.

Fuster, J. M. (1999). *Memory in the cerebral cortex: an empirical approach to neural networks in the human and nonhuman primate*, M.I.T. Press, Cambridge (MA).

Fuster, J. M. and Jervey, J. P. (1982). Neuronal firing in the inferotemporal cortex of the monkey in a visual memory task, *J. Neurosci.* **2**: 361–375.

Goel, V. and Dolan, R. J. (2001). The functional anatomy of humor: segregating cognitive and affective components, *Nature Neuroscience* **4**: 237–238.

Goldman-Rakic, P. S. (1988). Changing concepts of cortical connectivity: parallel distributed cortical networks, *in* P. Rakic and W. Singer (eds), *Neurobiology of neocortex*, John Wiley, pp. 177–202.

Grafton, S. T. (1995). Mapping memory systems in the human brain, *Semin. Neurosc.* **7**: 157–163.

Graziano, M. S., Andersen, R. A. and Snowden, R. J. (1994). Tuning of MST neurons to spiral motions, *J. Neurosci.* **14**: 54–67.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities, *Proc. Nat. Acad. Sci. USA* **79**: 2554–2558.

Hubel, D. H. and Wiesel, T. N. (1977). Functional architecture of the macaque monkey visual cortex, *Proc. R. Soc. (Lond.) B* **198**: 1–59.

Jenkins, I. H., Brooks, D. J., Nixon, P. D., Frackowiak, R. S. and Passingham, R. E. (1994). Motor sequence learning: a study with positron emission tomography, *J. Neurosci.* **14**: 3775–3790.

Kaas, J. H. (1987). The organization of neocortex in mammals: implications for theories of brain function, *Ann. Rev. Psychol.* **38**: 129–151.

Lauro-Grotto, R., Reich, S. and Virasoro, M. A. (1997). The computational role of conscious processing in a model of semantic memory, *in* M. Ito, Y. Miyashita and E. T. Rolls (eds), *Cognition, Computation and Consciousness*, Oxford Univ. Press, Oxford, U.K.

Little, W. A. (1974). The existence of persistent states in the brain, *Math. Biosci.* **19**: 101–120.

Martin, A. and Chao, L. L. (2001). Semantic memory and the brain: structure and processes, *Curr. Opin. Neurob.* **11**: 194–201.

Mishkin, M., Ungerleider, L. G. and Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways, *Trends in Neuroscience* **6**: 414–417.

Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex, *Nature* **335**: 817–820.

Miyashita, Y. and Chang, H. S. (1988). Neuronal correlate of pictorial short-term memory in the primate temporal cortex, *Nature* **331**: 68–70.

Mountcastle, V. B. (1957). Modality and topographic properties of single neurons of cat somatic sensory cortex, *J. Neurophysiol.* **20**: 408–434.

Mountcastle, V. B. (1977). An organizing principle for cerebral function: the unit module and the distributed system, *in* F. Schmitt and F. Worden (eds), *The Neurosciences: Fourth Study Program*, Addison Wesley, pp. 21–42.

Mountcastle, V. B. (1995). The parietal system and some higher brain functions, *Cerebral Cortex* **5**: 377–390.

Mountcastle, V. B. (1997). The columnar organization of the neocortex, *Brain* **120**: 701–722.

Mountcastle, V. B., Lynch, J. C., Georgopoulos, A., Sakata, H. and Acuna, C. (1975). Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space, *J. Neurophysiol.* **38**: 871–908.

Mountcastle, V. B., Motter, B. C., Steinmetz, M. A. and Duffy, C. J. (1984). Looking and seeing: the visual functions on the parietal lobe, *in* G. M. Edelman, W. E. Gall and W. M. Cowan (eds), *Dynamical Aspects of Neocortical Function*, John Wiley & Sons, New York, pp. 159–193.

Naya, Y., Yoshida, M. and Miyashita, Y. (2001). Backward spreading of memory-retrieval signal in the primate temporal cortex, *Science* **291**: 661–664.

Newman, M. E. J. (1999). Models of the small world: A review, *Santa Fe Institute working paper 99-12-080* (cond-mat/0001118).

O'Kane, D. and Treves, A. (1992). Short- and long-range connections in autoassociative memory, *J. Phys. A: Math. Gen.* **25**: 5055–5069.

Renart, A., Parga, N. and Rolls, E. T. (1999). Associative memory properties of multiple cortical modules, *Network: Comput. Neural Syst.* **10**: 237–255.

Ricci, P. T., Zelkowicz, B. J., Nebes, R. D., Cidis Meltzer, C., Mintun, M. A. and Becker, J. T. (1999). Functional neuroanatomy of semantic memory: Recognition of semantic associations, *NeuroImage* **9**: 88–96.

Romanski, L. M., Tian, B., Fritz, J., Mishkin, M., Goldman-Rakic, P. S. and Rauschecker, J. P. (1999). Dual streams of auditory afferents target multiple domains in the primate prefrontal cortex, *Nature Neuroscience* **2**: 1131–1136.

Savage, C. R., Deckersbach, T., Heckers, S., Wagner, A. D., Schacter, D. L., Alpert, N. M., Fischman, A. J. and Rauch, S. L. (2001). Prefrontal regions supporting spontaneous and directed application of verbal learning strategies: Evidence from PET, *Brain* **124**: 219–231.

Selemon, L. D. and Goldman-Rakic, P. S. (1988). Common cortical and subcortical targets of the dorsolateral prefrontal and posterior parietal cortices in the rhesus monkey: evidence for a distributed neural network subserving spatially guided behavior, *J. Neurosci.* **8**: 4049–4068.

Tanaka, K. (1996). Representation of visual features of objects in the inferotemporal cortex, *Neural Networks* **9**(8): 1459–1475.

Tomita, H., Ohbayashi, M., Nakahara, K., Hasegawa, I. and Miyashita, Y. (1999). Top-down signal from prefrontal cortex in executive control of memory retrieval, *Nature* **401**: 699–703.

Tsodyks, M. V. and Feigel'man, M. V. (1988). The enhanced storage capacity in neural networks with low activity level, *Europhys. Lett.* **6**: 101–105.

Van Hoesen, G. W. (1993). The modern concept of association cortex, *Curr. Opin. Neurob.* **3**: 150–154.