

A model of the interaction between mood and memory

E T Rolls¹ and S M Stringer

Department of Experimental Psychology, Oxford University, South Parks Road, Oxford
OX1 3UD, UK

E-mail: Edmund.Rolls@psy.ox.ac.uk

Received 12 June 2000, in final form 5 January 2001

Abstract

This paper investigates a neural network model of the interaction between mood and memory. The model has two attractor networks that represent the inferior temporal cortex (IT), which stores representations of visual stimuli, and the amygdala, the activity of which reflects the mood state. The two attractor networks are coupled by forward and backward projections. The model is however generic, and is relevant to understanding the interaction between different pairs of modules in the brain, particularly, as is the case with moods and memories, when there are fewer states represented in one module than in the other. During learning, a large number of patterns are presented to the IT, each paired with one of two mood states represented in the amygdala. The recurrent connections within each module, the forward connections from the memory module to the amygdala, and the backward connections from the amygdala to the memory module, are associatively modified. It is shown how the mood state in the amygdala can influence which memory patterns are recalled in the memory module. Further, it is shown that if there is an existing mood state in the amygdala, it can be difficult to change it even when a retrieval cue is presented to the memory module that is associated with a different mood state. It is also shown that the backprojections from the amygdala to the memory module must be relatively weak if memory retrieval in the memory module is not to be disrupted. The results are relevant to understanding the interaction between structures important in mood and emotion (such as the amygdala and orbitofrontal cortex) and other brain areas involved in storing objects and faces (such as the inferior temporal visual cortex) and memories (such as the hippocampus).

1. Introduction

This paper investigates a neural network model of the interaction between mood and memory. The model builds on a theoretical study presented in Renart *et al* (1999b), which investigated

¹ Corresponding author.

the pattern retrieval properties of two connected attractor neural networks. However, in this paper we address questions that are specifically related to the dynamical interaction of memory recall and mood. This is achieved through simulations of a network composed of two connected modules, where each module is an attractor network that represents either memory patterns or a mood state. Areas of the brain involved in the storage of memories of objects and faces include the inferior temporal visual cortex, while areas of the brain involved in storing episodic memories include the hippocampal system and medial temporal lobe (Rolls and Treves 1998). In our network simulations, we designate the memory storage module as representing the inferior temporal cortex (IT). Areas of the brain important in representing the mood state produced by emotion provoking stimuli include the amygdala and orbitofrontal cortex (Rolls 1999). In our simulations, the mood module represents the amygdala. Although we have designated the two modules in our simulations as the IT and the amygdala networks, these modules could equally be regarded as representing other areas of the brain involved in either memory storage or mood representation, as discussed above. There are forward projections from the IT module to the amygdala module, and backward projections from the amygdala to the IT module (Rolls and Treves 1998).

It is well known that emotional or mood states can influence the cognitive evaluation of events or memories (Rolls 1999, Blaney 1986). For example, when humans are in a happy mood state, they are more likely to recall memories that were stored when the subject was happy, and vice versa when sad. One of the pioneers of such research was Beck (see, e.g., Beck *et al* 1985), who emphasized in the context of clinical problems the biases in cognitive processes such as memory and attention for material congruent with the individual's emotional state. Research examining biases in cognitive processing associated with emotional disorders has demonstrated selective memory, attention, judgment and interpretation of ambiguity in favour of negative, often self-referent, information in subjects with depression or some form of anxiety disorder (see Williams *et al* 1997 for a review). Indeed, a principal focus of research on the cognitive processing of emotional information has been the investigation of selective memory biases. For example, Bradley and Mathews (1988) found that depressive participants, who had previously been given a list of positive and negative trait adjectives to remember, recalled more self-referent negative adjectives relative to their recall of positive adjectives and to the performance of control participants. Similar memory biases have been found in patients with panic disorder for verbal material related to their emotional concerns (Cloitre and Liebowitz 1991). The specificity of these biases suggests that at least part of the way in which mood affects memory is by a specific effect of each mood state on the type of material recalled from memory, and thus that a general effect of the mood state such as arousal cannot account for the effects found. It is these specific effects of mood state on memory that are investigated by modelling in this paper. In addition to influences of mood on explicit processing, it has been possible to show that mood can influence an example of implicit memory, priming (Watkins *et al* 1996, McNally 1997).

Associative memories can potentially model the effects of mood on memory by storing as part of the memory the emotional state present at the time of storage, or at least associating the mood state with the memory pattern being stored. Retrieval is then likely to be influenced by the mood state. However, a question arises about the relative strengths of the memory and mood parts of what is stored, if the mood is not to dominate the memory retrieval. This issue can be addressed by the two-module simulation described, for it captures the fact that memory patterns are stored in one set of modules, and the mood state in another, with reciprocal connections between them. Part of what must be addressed quantitatively then is the relative strengths of the connections between the modules, relative to the strengths of the associative recurrent connections within each module, in order to allow effective pattern retrieval in the IT module.

Another issue that is addressed is the effect of memory retrieval on mood. A key aspect of the simulations performed here is that we explore the dynamical interaction of memory recall and mood. For example, to what extent is memory retrieval in the IT module able to influence mood state in the amygdala when there is no existing mood state in the amygdala as compared with a situation when there is an existing mood state already present in the amygdala? Another key aspect of the simulations is that we altered independently the strength of the forward connections and the strength of the backprojections between the modules, in order to illuminate the factors that are likely to determine in the brain the relative strengths of forward-projecting and backprojecting connections between modules. This is different from the theoretical studies of Renart *et al* (1999a, b), where the analytical techniques employed were applied to the case of equal strength of the forward and backward projections. The model presented here is relevant to understanding the interaction between structures important in emotion such as the amygdala or orbitofrontal cortex, and various other areas of the brain involved in memory storage such as the IT and hippocampus.

In some of the memory systems the emotional context may be stored as part of the memory. For example, inputs about emotional state reach the entorhinal cortex from the amygdala, and thus the medial temporal lobe memory mechanisms including the hippocampus are capable of linking together arbitrary co-occurrences of inputs with their emotional context. In turn, backprojections from the hippocampus may be able to influence cortical structures earlier in processing such as the inferior temporal visual cortex (see Treves and Rolls 1994), and via this interaction the retrieval of memories in the medial temporal lobe, itself affected by mood context, may be able to influence processing in earlier cortical areas.

Another brain region where there may be interaction between emotional states and memory retrieval is the IT, which is involved in the visual representation of objects. There is now much experimental evidence to suggest that over a series of successive cortical stages the visual system develops neurons that are able to respond with considerable view, size and position invariance to objects or faces (Perrett *et al* 1982, Desimone 1991, Rolls 1992, 2000a, Rolls and Tovee 1995, Tanaka *et al* 1991, Rolls and Treves 1998, Booth and Rolls 1998). These face or object neurons are found in the IT, and it is thought that such neurons underpin object representation and recognition in the visual system. However, while the IT receives forward inputs from the earlier stages of the ventral visual processing stream, it also receives extensive backprojections from the amygdala and orbitofrontal cortex. This backprojection system potentially allows the mood represented by the amygdala and orbitofrontal cortex to affect the storage and retrieval of visual memories in the IT. It has been suggested (Rolls and Treves 1998) that co-activity between forward inputs and backprojecting inputs to strongly activated cortical pyramidal cells would lead to both the forward and the backprojecting sets of synapses being modified (see figure 1). This could result in the preferential retrieval of cortical representations (for example of particular faces) in the IT that had been originally associated with the same emotional states as currently represented by activity in the amygdala. That is, with relatively strong emotional states present, cortical representations that had been previously associated with the same emotional states would be more easily recalled than other memories. This type of architecture with forward and backward connections between different networks provided the starting point for the hypotheses investigated here. In this paper, each of the connected networks or modules has internal recurrent connections, which are not only a feature of cortical and probably amygdala connectivity, but which also enable useful memory properties to arise such as persisting firing (which would be useful in the mood module to maintain the mood state), and the completion of memories (which would be useful in the memory module) (Rolls and Treves 1998).

In the next section we describe the general form of the neural network model developed here. The model consists of two interacting modules with recurrent associative connections




Figure 1. Pyramidal cells in, for example, layers 2 and 3 of the temporal lobe association cortex receive forward inputs from preceding cortical stages of processing, and also backprojections from the amygdala. Modifiable synapses from the amygdala to the cortical cells would allow mood state to facilitate memory retrieval when current amygdala activity represented the mood state associated with, for example, a particular visual stimulus.

internally and with associative connections between the modules. One module is identified as a memory module, and is named the IT module. The other module is identified as a mood module, and is named the amygdala module. The model could equally represent the interaction between the amygdala and other linked brain regions operating as auto-associative memory systems, for example, the CA3 region of the hippocampus discussed above. We emphasize though that the model is generic and is relevant to understanding the operation of any connected networks with recurrent internal connections. The networks are studied under conditions where the magnitudes of the associative connections between the modules in the two directions can be varied independently, and when there is a small number of patterns stored in one module each of which is associated with a large number of different randomly chosen patterns in the other module. The dynamics of the system are investigated under unclamped conditions (as is common in physics-based approaches to attractor networks, see Hopfield 1982, Amit 1989, Hertz *et al* 1991), that is when a retrieval cue is presented to one of the modules for a short time (the first timestep in the discrete update simulations), and is then removed. This test situation may be compared with presenting a perceptual stimulus to the IT for a short period only, but then determining how the state it produces in the second module, which might continue for a long time, affects the ongoing representation of that stimulus in the perceptual module, as well as how the perceptual module responds to other retrieval cues.

2. The neural network model

The architecture of the neural network investigated here is shown in figure 2. The network has two classes of neuron, with connections between the classes and recurrent connections within the classes. The two classes represent the IT and the amygdala; the firing pattern of the IT neurons represents visual stimuli, while the pattern of activity within the class of amygdala neurons reflects the emotional state. Each class contains 1000 neurons, and each neuron in each class receives inputs from every other neuron in the network (the fully connected condition). In the simulations presented here, the amygdala can represent two types of emotional state (e.g. happy versus sad), which we call positive and negative in the following, corresponding to two different firing patterns. Similarly, the IT neurons can represent 100 visual stimuli corresponding to 100 different firing patterns. However, during learning, half of the IT patterns




Figure 2. Model network architecture showing two neuron classes representing the inferior temporal cortex and amygdala (RC denotes recurrent connections).

co-occur with the positive amygdala pattern, and the other half co-occur with the negative amygdala pattern. That is, 50 visual stimuli become associated with the positive emotional state, while the other 50 stimuli become associated with the negative emotional state.

At the beginning of the learning phase all synaptic weights are initialized to zero. Then, during learning, there are 100 pairs of firing rate patterns, which consist of a unique IT stimulus pattern and its corresponding amygdala pattern (positive or negative). Each pair of firing rate patterns is clamped to produce the firing rates r_i of the IT and amygdala neurons, and the synaptic weights are updated according to the following rules. For the intra-class connections from IT to IT neurons, and from amygdala to amygdala neurons, the synaptic weights are updated according to the covariance rule. That is, for each pair of pre- and post-synaptic neurons, j and i , the synaptic weight w_{ij} is updated by the covariance rule

$$\Delta w_{ij} = k(r_i - \langle r_i \rangle)(r_j - \langle r_j \rangle) \quad (1)$$

where Δw_{ij} is the change in synaptic strength, k is the learning rate and $\langle \cdot \rangle$ indicates an average value. Unless otherwise stated, the learning rate k is set to 1.0. However, in some of the experiments below, the learning rate k is used to vary the relative strength g of the synaptic connections for the four different types of intra- and inter-class connection, where we have simply $g = k$. We denote the four separate intra- and inter-class connection strengths g as follows: g_{II} for IT to IT connections, g_{aa} for amygdala to amygdala connections, g_{Ia} for IT to amygdala connections and g_{aI} for amygdala to IT connections. Given that the default values for each g were 1, if g_{Ia} was set to 0.1, then the ratio of the external to the internal input currents to the IT neurons was 0.1. This definition in terms of input currents induced by external versus internal (recurrent) inputs to a module is consistent with that captured by g as used by Renart *et al* (1999a, b). The average firing rate $\langle r_i \rangle$ for each of the two classes is set equal to the firing rate sparseness (see equation (4)) of the input patterns applied to the classes. The covariance rule is highly effective for the recurrent connections within auto-associative memory networks, and such auto-associative memory networks are thought to be instantiated in regions like IT and the CA3 region of the hippocampus. For the inter-class connections from IT to amygdala neurons, and from amygdala to IT neurons, the synaptic weights are updated according to the modified Hebb rule with hetero-synaptic long-term depression (LTD)

$$\Delta w_{ij} = kr_i(r_j - \langle r_j \rangle). \quad (2)$$

In this rule LTP is incorporated with heterosynaptic LTD. Use of this rule can help to remove the correlations between the input patterns produced by positive-only firing rates (Rolls and Treves 1998). We note that the effect of normal training with all g parameters set to 1.0 was that after 100 training trials, each of the two amygdala patterns would have been presented 50 times, and the depth of the attractor basin implemented by the recurrent collateral amygdala connections would be much deeper than those in the IT. For this reason we explored in some of the simulations the effects of reducing g_{aa} below 1.0, for those experiments in which the amygdala states could evolve.

At the beginning of the memory retrieval phase all firing rates are initialized to zero. Then, during memory retrieval, whole or partial cues for the original IT and amygdala firing patterns were clamped onto the IT and amygdala neurons, and the network iterated over a number of timesteps. At each timestep the following steps are performed. First, the activations h_i of each neuron i in the network are calculated according to

$$h_i = \sum_j r_j w_{ij} \quad (3)$$

where r_j is the firing rate of neuron j , and w_{ij} is the strength of the synapse from neuron j to neuron i . Next the neuronal firing r_i of each neuron is either kept clamped to the original pattern, or is re-calculated based on the activations h_i by setting the sparseness a of the firing rates r_i of a class of neuron according to

$$a = \frac{(\sum_i r_i / N)^2}{\sum_i r_i^2 / N} \quad (4)$$

where N is the number of neurons in the layer (Rolls and Treves 1998). The firing rate sparseness a for each class is determined by the firing rate sparseness of the input patterns for that class, and the firing rate sparseness of each class is then kept constant throughout each simulation. This was performed by adjusting the threshold of the binary activation function (Rolls and Treves 1998) of the neurons to a value that resulted in the required sparseness of firing in the relevant class of neurons. For the simulations presented here, all input patterns to both classes of neurons were binarized patterns of zeros and ones, with only 50 neurons firing. This gave a firing rate sparseness of 0.05 for both classes. After 20 timesteps, the firing rates of the IT and amygdala neurons were recorded. In experiments in which we examined the effect of mood on memory retrieval, the correlation was calculated between the recalled IT pattern after 20 timesteps, and the original full pattern for which a partial cue had been presented. The correlation measure ranges from 1.0 for perfect recall to 0.0. The question we then address is how well IT neurons can recall a memory in the IT of a stimulus given an initial partial cue, and given an input from the amygdala containing information about the current emotional state. In experiments examining the effect of the retrieved IT stimulus pattern on mood, the correlation was calculated between the final amygdala pattern after 20 timesteps, and the two original amygdala patterns representing positive and negative moods. Here we investigate under what conditions a retrieved memory of a visual stimulus in the IT can affect the firing pattern represented by the amygdala, and hence alter the emotional state.

3. Simulations

In this section we firstly examine the effect of a constant mood on memory retrieval in section 3.1, and then examine the dynamical interaction of memory retrieval and evolving mood in section 3.2.




Figure 3. Experiment 1. IT recall correlations averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. In this experiment the firing rates of the amygdala neurons are clamped to zero.

3.1. The effect of a constant amygdala mood state on cortical memory retrieval

In this section we examine the effect of a constant mood on memory retrieval in a cortical area such as the IT. This is achieved in the simulation by clamping the firing rates of the amygdala neurons to constant values over all timesteps of the memory retrieval phase. Thus, in these experiments mood is unaffected by the process of memory retrieval in the IT.

3.1.1. Baseline condition with no mood present during memory retrieval. In experiment 1 the retrieval of stimulus representations in the IT without input from the amygdala was measured to provide a baseline for comparison with later experiments. To achieve this, during the memory retrieval phase the amygdala firing rates were clamped to zero over all the timesteps of each simulation. Then, a partial cue was presented to the IT neurons by setting the firing rates of a proportion of the neurons in the IT to one of the original patterns, and then iterating the network over 20 timesteps. This was repeated for all of the 100 IT patterns, and for a range of different sizes of the initial partial cue. The cue was presented for the first timestep of the simulation to trigger retrieval, and then removed for the remaining iterations of the recall. This method of testing recall allows the network to achieve perfect recall of the whole pattern. The results for experiment 1 are presented in figure 3, where each plot shows IT recall correlations averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. Each plot shows the average correlation between the retrieved IT pattern after 20 timesteps and the original pattern, as a function of the size of the partial cue presented to the IT. The retrieval, as measured by the correlation of the retrieved pattern with the stored pattern, is essentially perfect for recall cues greater than about 15% of the original (i.e. completion occurs), and only drops with smaller retrieval cues. As expected, the two plots are almost coincident, showing no retrieval asymmetries between the two groups of IT patterns when the amygdala is not active.




Figure 4. Experiment 2. IT recall correlations averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. In this figure the connection strength parameter g_{1a} from amygdala to IT neurons is set to 0.1.

3.1.2. Simulations with mood present during memory retrieval. In experiment 2 we investigated the effect of different amygdala mood states on the recall of IT patterns associated during learning with the two different mood states. Figure 4 shows retrieval during a positive mood state of both positive-associated and negative-associated patterns in the IT. The connection strength from the amygdala to the IT was set at 0.1 of the other connection strengths, i.e. $g_{1a} = 0.1$. It is shown that the retrieval of positive-associated IT patterns was much better (as shown by higher correlations with the stored patterns) than that of negative-associated IT patterns. The effect of mood was most marked when the IT module was having to perform real completion from an incomplete cue, and the amygdala mood state had less effect on the recalled IT pattern when the IT recall cue was large. The results were of course symmetric with the negative mood state in the amygdala: the negative-associated IT memories were then recalled better.

The results in figure 4 were obtained with $g_{1a} = 0.1$. We show in figure 5 the effect of different values of g_{1a} . The IT retrieval cue was 40% of the pattern stored. At low values of g_{1a} (up to and including 0.05) there is little influence of the amygdala mood state on cortical memory retrieval, and the retrieval is similar for positive-associated and negative-associated IT patterns. As g_{1a} is increased to 0.10 and higher, the average correlation with the correct IT pattern deteriorates considerably, although the retrieval of IT patterns is always better if the corresponding mood state is present in the amygdala. A subsidiary point is that with high values of g_{1a} , the retrieval of the positive-associated IT patterns is worse with a positive mood in the amygdala than with no mood in the amygdala. The reason for this is that 50 IT patterns were associated with a single mood; and when that mood state is present in the amygdala, it effectively makes the IT retrieval cues more similar for these 50 patterns than they would be without the mood state that is common to all 50 patterns.

To reveal more fully how the interaction between the amygdala and the IT operates over a full range of the parameters g_{1a} and the fraction of the original stimulus used to cue retrieval in the IT, we show three-dimensional graphs in figure 6. The graphs show for the condition




Figure 5. Experiment 2. IT recall correlations averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. In this figure the size of the IT retrieval cue is 40% of the pattern stored, while the connection strength parameter g_{1a} from amygdala to IT neurons is varied from 0.0 to 1.0.

when a positive mood is clamped onto the amygdala neurons the correlation between the memory stored and the memory retrieved in the IT when the retrieval cues applied to the IT are those associated with positive moods (left) or negative moods (right). The recall of the 50 positive-associated IT patterns is better than the recall of negative-related IT patterns for a set of parameter values where g_{1a} is in the region of 0.05–0.15, and the fraction of the retrieval cue being used is between 20 and 60%. In addition to this effect, the presence of positive-related firing in the amygdala when g_{1a} is greater than 0.3 prevents good recall of either positive-related or negative-related patterns in the IT, though on average there is a slightly greater correlation between the pattern stored and the pattern retrieved for the positive-associated retrieved patterns than for the negative-associated retrieved patterns (evident in the slightly raised floor in figure 6 left).

3.2. Dynamical interactions between memory retrieval and evolving mood

In the above experiments the firing rates of the amygdala neurons were clamped throughout all the timesteps of the memory retrieval phase. However, in this section we examine the dynamical interactions between memory retrieval and evolving mood by allowing the firing rates of the amygdala neurons to evolve over successive timesteps of the memory retrieval phase. Also, since the firing rates of the amygdala neurons are now free to evolve, the recurrent connections within the amygdala can play an important role; the effect of varying the connection strengths for these recurrent connections is also investigated.

3.2.1. Simulations with no mood present in the amygdala at the beginning of memory retrieval.

In experiments 3 and 4 we investigate the effect that the retrieval of a memory in the IT may have on the emotional state represented by the amygdala, and the effect in turn that the retrieved mood in the amygdala may have on the recall in the IT. These experiments are similar to earlier experiments, except that the firing rates of the amygdala neurons start initially at zero,

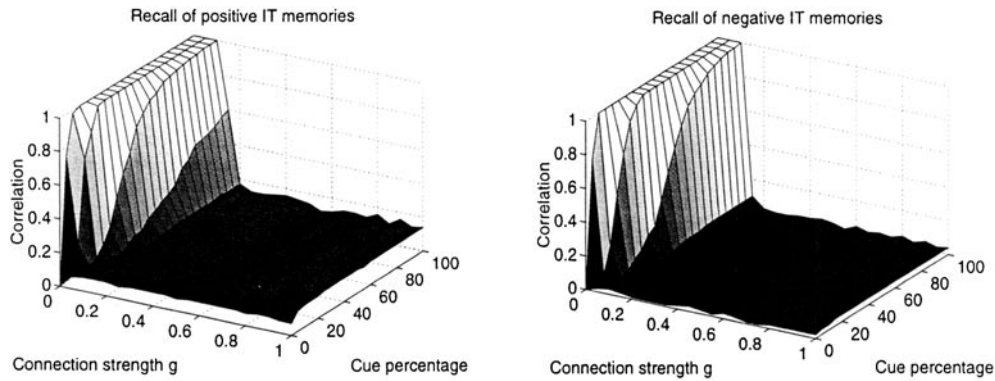


Figure 6. Experiment 2. IT recall correlations averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. The three-dimensional graphs show how the interaction between the amygdala and the IT operates over a full range of the parameters g_{1a} and the fraction of the original stimulus used to cue retrieval in the IT. The graphs show for the condition when a positive mood is clamped onto the amygdala neurons the correlation between the memory stored and the memory retrieved in the IT when the retrieval cues applied to the IT are those associated with positive moods (left) or negative moods (right).

and are free to evolve as a result of the influence of the IT state. In experiment 3 all the connection strength parameters were set to 1.0. In figure 7 we show the effects of applying a retrieval cue to the IT on the evolving emotional state represented by the amygdala. The figure illustrates the correlations of the retrieved amygdala states with the positive amygdala mood state present during training. If a positive-associated cue is applied to the IT, then the retrieved amygdala pattern correlates highly with the positive amygdala training pattern (solid curve). If a negative-associated cue is applied to the IT, then the retrieved amygdala pattern has a low correlation with the positive amygdala training pattern (dashed curve). (A symmetric situation was of course present if the negative amygdala training pattern was considered.) The most interesting effect however was what happened to the evolution of the IT patterns during this retrieval. This is shown in figure 8. Although, for example, presentation of a positive-related IT pattern to the IT was sufficient to produce positive-related mood firing in the amygdala, the amygdala firing then by the backprojections to the IT prevented the IT from developing a continuing good representation of whatever positive-related IT pattern has been presented, and in fact the correlation of the IT firing with the correct pattern was low. The reason for the poor continuing retrieval in the IT was the strength of the amygdala to IT backprojections, with $g_{1a} = 1.0$ for this experiment. The basis for the effect is partly the strength of g_{1a} , and partly the fact that there are 50 different IT patterns all associated with the positive pattern in the amygdala, and the single amygdala firing pattern in the amygdala associated in the backprojecting synapses with the 50 IT patterns prevents the states recalled in the IT being very different from each other, and thus produces poor IT recall. Part of the interest here is that a very short-lasting IT state produced by one of the IT cue patterns is sufficient to produce a continuing and long-lasting amygdala-based mood state, while the IT representation of the pattern is lost almost immediately. This is an interesting parallel to a situation which might occur when a stimulus produces a mood, but the stimulus itself produces only a minor and short-lasting stimulus-related representation in the perceptual module. In this way a stimulus might produce a mood with little perceptual knowledge of the stimulus.




Figure 7. Experiment 3. Correlations between the evolved amygdala firing pattern and amygdala pattern 1 (positive mood), averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. This experiment is similar to earlier experiments, except that the firing rates of the amygdala neurons are clamped to zero for only the first timestep of the memory retrieval stage, after which the firing rates of the amygdala neurons are allowed to evolve. The connection strength parameters g_{1a} , g_{1l} , g_{a1} and g_{aa} were all 1.0.




Figure 8. Experiment 3. IT recall correlations averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. This experiment is similar to earlier experiments, except that the firing rates of the amygdala neurons are clamped to zero for only the first timestep of the memory retrieval stage, after which the firing rates of the amygdala neurons are allowed to evolve. The connection strength parameters g_{1a} , g_{1l} , g_{a1} and g_{aa} were all 1.0.

In experiment 4 the procedure of experiment 3 is explored further, but with different values for the connection strengths g . Experiment 4 is in four parts, where for each part we




Figure 9. Experiment 4, part (i). The IT recall correlations averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. In this experiment the firing rates of the amygdala neurons are clamped to zero for only the first timestep of the memory retrieval stage, after which the firing rates of the amygdala neurons are allowed to evolve. This experiment is different to experiment 3 in that although three of the connection strength parameters are still set to 1.0, the connection strength parameter g_{ia} for connections to IT neurons from the amygdala is set to 0.1.

independently vary the connection strength for the following types of inter- and intra-class connection: (i) to IT from amygdala (g_{ia}), (ii) from IT to IT (g_{ii}), (iii) to amygdala from IT (g_{ai}) and (iv) from amygdala to amygdala (g_{aa}). In each of the cases (i)–(iv), we repeat experiment 3 except with the connection strength for the specified type of connection set to a range of values from 0.0 to 10.0. The connection strengths for the three other types of inter- and intra-class connection are kept at their original value of 1.0.

In a first interesting case, (i), decreasing the connection strength for synaptic connections to the IT from the amygdala eventually led to improved pattern-specific memory recall in the IT, while at the same time producing the correct mood state in the amygdala. This is illustrated for $g_{ia} = 0.1$ in figure 9. The retrieval in the IT is much better than with $g_{ia} = 1.0$, for which the results were illustrated in figure 8. When the backprojections are weak, e.g. with g_{ia} in the region of 0.1, then they can influence recall in the IT in what might be a useful way (by encouraging the retrieval of positive-related memories in the IT for example when there is a positive mood state in the amygdala, an effect illustrated in figure 4). The important effect to be understood here is that the backprojections must not be too strong, or they will dominate the firing in a previous module, and prevent it responding correctly to its inputs from the previous stage. However, it can also be seen that the IT performance demonstrated in figure 9 with amygdala firing rates evolving from zero in response to the IT inputs, is significantly better than that shown in figure 4 where the amygdala firing rates were clamped to the positive mood state from the first timestep. This increased performance is not due to a failure of the amygdala to settle into the firing pattern that was associated with the relevant IT pattern during learning. This is demonstrated in figure 10, where we show mood evolution results in the amygdala with the connection strength for amygdala to IT connections (g_{ia}) set to 0.1. Here it can be seen that when the initial IT cue is one of the IT patterns that was associated with amygdala pattern 1 (positive) during learning, then during the memory recall phase the firing rates in the amygdala still evolve to the positive




Figure 10. Experiment 4, part (i). Correlations between the evolved amygdala firing pattern and amygdala pattern 1 (positive mood), averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. In this experiment the firing rates of the amygdala neurons are clamped to zero for only the first timestep of the memory retrieval stage, after which the firing rates of the amygdala neurons are allowed to evolve. This experiment is different to experiment 3 in that although three of the connection strength parameters are still set to 1.0, the connection strength parameter g_{1a} for connections to IT neurons from the amygdala is set to 0.1.

mood state, as in experiment 3. Similarly, although not shown here, when the initial IT cue is one of the IT patterns that was associated with amygdala pattern 2 (negative) during learning, then during the memory recall phase the firing rates in the amygdala also continue to evolve to the negative mood state. Hence, it appears that letting the amygdala firing rates evolve freely from zero, rather than clamping the amygdala to a particular mood state, helps both classes of neuron to recall their respective learned patterns, and in particular can lead to enhanced recall performance in the IT. The improved recall in the IT is probably due to the initial few timesteps of the memory recall phase, in which the amygdala has not yet settled into a stable pattern of activity that was associated during learning with a large number (50) of IT patterns. This short time interval, before the amygdala settles into a state that would disrupt recall in the IT, can help significantly to improve recall in the IT from a relatively small partial cue.

For case (ii), increasing the connection strength for recurrent synaptic connections within the IT (g_{11}) above unity had a similar effect to decreasing the connection strength for synaptic connections to the IT from the amygdala (g_{1a}), and eventually led to improved memory recall. Indeed, with the connection strength for the IT to IT connections g_{11} set to 10.0, the recall performance of the network (not illustrated) was similar to that shown in figure 9, while the firing rates in the amygdala still evolved to the associated positive or negative mood state.

For case (iii), with the connection strength for synaptic connections to the amygdala from the IT (g_{a1}) varied, it was found that for g_{a1} greater than zero, the initial cue presented to the IT produced good retrieval of the associated mood state in the amygdala. This then led to the effects illustrated in figure 8, in which, because in these simulations g_{1a} was 1.0, the amygdala firing dominated the IT firing, and the IT retrieval was very poor.

For case (iv), with the strength of the recurrent connections within the amygdala (g_{aa}) varied, the most interesting result is that the correct amygdala mood state could be produced

by a pattern applied to the IT even when g_{aa} was 0.0. The mechanism for this was that the retrieval cue applied to the IT resulted in a continuing IT state which by pattern association produced the correct amygdala mood state. When g_{1a} was 1.0, the actual state in the IT was not the firing that should be produced by the IT retrieval cue, but instead the IT amalgam that was produced by the strong influence of the mood-related amygdala firing on the IT firing. If g_{1a} was then reduced to e.g. 0.1, then the amygdala mood-related firing still continued, but now in the presence of the correct firing in the IT that should be produced by the IT retrieval cue.

3.2.2. Simulations with mood present at beginning of memory retrieval. Experiment 5 is similar to experiment 3, except that the firing rates of the amygdala neurons are now clamped to the pattern representing the positive emotional state for the first timestep of the memory retrieval phase, after which the firing rates of the amygdala neurons are updated according to the neuronal activations as described above. All connection strength parameters were 1.0. It was found that memory retrieval in the IT failed for all sizes of the partial cue presented to IT. Indeed even with 100% of the firing rate patterns presented to the IT as an initial cue, the patterns could not be stably maintained in the IT. This behaviour is linked to the amygdala activity, which did not change from its initial positive state clamped for timestep 1 regardless of which size or type of IT cue was presented. That is, once the amygdala is in a stable emotional state, its firing pattern is unaffected by activity in the IT; the inputs from IT to the amygdala were unable, for example, to change the mood state from positive to negative when an IT cue was presented that had originally been associated during learning with a negative mood. But the constant positive mood state continues to feed into the IT where, for similar reasons to those given above for experiment 3, this leads to a failure of IT to recall, or even stably support, the full IT patterns associated with the partial cues. The results of this experiment thus show that if the backprojections are too strong (in the experiment, of the same strength as the other connections), then the perceptual module does not respond to an external cue applied at timestep 1, nor can it alter the ongoing mood state in the amygdala.

Experiment 6 is a repeat of experiment 5, but where we now investigate the effect of varying the connection strengths g for the different types of synaptic connection. Experiment 6 is in four parts, where for each part we independently vary the connection strength for the following types of inter- and intra-class connection: (i) to IT from amygdala (g_{1a}), (ii) from IT to IT (g_{11}), (iii) to amygdala from IT (g_{a1}) and (iv) from amygdala to amygdala (g_{aa}). In each of the cases (i)–(iv), we repeat experiment 5 except with the connection strength for the specified type of connection set to a range of values from 0.0 to 10.0. The connection strengths for the three other types of inter- and intra-class connection are kept at their original value of 1.0.

For case (i), decreasing the connection strength for synaptic connections from the amygdala to the IT led to improved memory recall in the IT relative to experiment 5. However, pattern retrieval performance in the IT was similar to that obtained in experiment 2, where the amygdala firing rates were clamped to the positive mood state throughout the memory retrieval phase. With the connection strength for the amygdala to IT connections set to 0.0 and 0.1, results were similar to those shown in figures 3 and 4, respectively. The similarity of these results with experiment 2 may be understood from the evolution of the firing rates within the amygdala, where results were obtained similar to those obtained in experiment 5. That is, for all values of the connection strength for connections from the amygdala to the IT, the activity within the IT was unable to affect the firing pattern of the amygdala, which remained in its initial positive mood state regardless of the size or type of cue initially presented to the IT.

For case (ii), increasing the connection strength for the recurrent IT synaptic connections had a similar effect to decreasing the connection strength for synaptic connections from the amygdala to the IT, and eventually led to improved memory recall. Memory recall performance

in the IT was again similar to that obtained in experiment 2. With the connection strength for the IT to IT connections set to 10.0, the recall performance of the IT network was similar to that shown in figure 4. Also, for all values of the connection strength for the recurrent IT connections, the inputs from IT were unable to affect the firing pattern within the amygdala, which remained in a constant positive mood state throughout the memory retrieval phase, with results similar to those found in experiment 5.

For case (iii), in which the strengths of the connections to the amygdala from IT were varied, it was found that even for values of g_{al} as high as 10, the cue presented to the IT was unable to flip the initial mood state represented in the amygdala. In addition, the backprojections from the amygdala to the IT were able to prevent full recall in the IT, with results similar to those described in experiment 5. However, as the strength of g_{al} was increased to the order of 100, then the initial cue presented to the IT could start to overturn the initial mood state such that the final evolved mood state was highly correlated with the mood state originally associated with the IT cue pattern during learning.

For case (iv), memory retrieval in the IT failed, with results similar to those described in experiment 5, for all values of the connection strength for the recurrent synaptic connections within the amygdala, g_{aa} . This was because the mood state applied to the amygdala at timestep 1 prevented the IT module getting started at timestep 1 towards the attractor state that should have been induced by the retrieval cue applied to the IT. The result was that the IT module at later timesteps entered an amalgam of the IT patterns which correspond to the mood applied at timestep 1 to the amygdala, and this in turn is sufficient to reactivate the amygdala mood state.

Finally, in experiment 7 we repeated the general procedure of experiments 5 and 6, except that we now simultaneously adjusted a number of the connection strengths. It was shown that setting the strength of the forward connections g_{al} to a high value (e.g. 200), and then reducing the strength of the backprojections g_{la} to a low value (e.g. 0.1) could allow the IT to retrieve correctly from the cue presented to the IT at timestep 1, and additionally for the correct amygdala state associated with that IT pattern to be retrieved, even when the amygdala had been initialized to the incorrect mood state at timestep 1. These effects are demonstrated in figures 11 and 12. Firstly, figure 11 shows retrieval results in the IT, where it can be seen that there is reasonably good recall of IT patterns for moderately large cues (e.g. when the fraction of the retrieval cue is over 50%). Secondly, figure 12 shows the correlations of the retrieved amygdala states with the negative amygdala mood state present during training. It can be seen that if a negative-associated cue is applied to the IT, then the recalled amygdala pattern correlates highly with the negative amygdala training pattern (dashed curve). Thus the initial positive mood state is flipped to the negative mood state by an IT cue associated with the negative mood during learning. The reason for this behaviour in the two modules is that with weak backprojections, IT retrieval is not blocked by the amygdala backprojections; and with strong forward projections to the amygdala, the existing amygdala state can be overturned by the state produced in the IT by the input retrieval cue. These further investigations thus emphasize that, at least when inputs are not clamped to the IT but are applied briefly, backprojections from the amygdala must be very weak so as not to interfere with IT retrieval; and to change an existing mood state in the amygdala, even when it is unclamped, the strength of the forward projections to the amygdala g_{al} must be relatively high with respect to the strength of the recurrent amygdala connections g_{aa} .

4. Discussion

The results described here are relevant generically to the interactions that occur between coupled attractor networks when more than one memory in one network is coupled associatively




Figure 11. Experiment 7. The IT recall correlations averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. In this experiment the firing rates of the amygdala neurons are clamped to the pattern representing the positive emotional state for only the first timestep of the memory retrieval stage, after which the firing rates of the amygdala neurons are allowed to evolve. For this experiment the connection strengths are set to the following values: $g_{al} = 200$, $g_{la} = 0.1$, $g_{ll} = 1$ and $g_{aa} = 1$.




Figure 12. Experiment 7. Correlations between the evolved amygdala firing pattern and amygdala pattern 2 (negative mood), averaged over the simulations in which the initial IT cue pattern was associated during learning with either of the two mood states, positive and negative. In this experiment the firing rates of the amygdala neurons are clamped to the pattern representing the positive emotional state for only the first timestep of the memory retrieval stage, after which the firing rates of the amygdala neurons are allowed to evolve. For this experiment the connection strengths are set to the following values: $g_{al} = 200$, $g_{la} = 0.1$, $g_{ll} = 1$ and $g_{aa} = 1$.

with one memory in the other network. There may be very many such sets of pairings. The description in this paper was placed in the context of the interaction between different

memories, of which there are very many, and mood states, of which there are relatively few. The particular brain region chosen to be named as a memory network was the inferior temporal visual cortex (IT), in which distributed perceptual representations of objects and faces are present (Booth and Rolls 1998, Rolls 2000a). The region taken as an example of where moods are represented was the amygdala, because this is a region involved in emotion (Rolls 2000b). Part of the reason that these two regions were chosen is that they are *directly* connected, by forward projections from IT to the amygdala, and by backprojections from the amygdala to the IT (Amaral *et al* 1992). However, the results described here are relevant also in architectures when the attractor networks concerned are not directly connected, but are connected by multiple other networks, as applies for example in the multiple-stage backprojection pathway from the hippocampus to preceding cortical areas (Treves and Rolls 1994, Rolls and Treves 1998). Indeed, the results are generically relevant to understanding the effects of cortico-cortical backprojections in the brain (Rolls and Treves 1998). We emphasize that in the experiments described here, it is the many-to-one (in the simulations 100-to-2) relationship between IT memories and moods that is responsible for the asymmetric behaviour of the two classes of network, and is responsible for much of the unexpected behaviour of the network. One effect of the many-to-one relationship is that 50 memory patterns are associated with a single mood pattern, and the influence of a single mood pattern makes the 50 memory patterns less distinct from each other than if there were no mood present. Another effect of the many-to-one relationship is that individual synapses trained 50 times with a particular mood pattern present in the amygdala module will become strong for a given learning rate k , and when that mood pattern is present during retrieval, will strongly influence retrieval. (We note that the g factors reflect the average current that can be injected by a pattern after one training trial, and that if there are 50 training trials for a particular pattern, then 50 times as much current will be injected when that particular pattern is presented.) These points are drawn out in the following discussion.

Experiment 1 provided the baseline comparison condition for retrieval of memories in the IT module when there is no mood state present in the amygdala, and amygdala backprojections cannot influence retrieval in the IT. The IT module under these conditions showed the expected memory retrieval properties of attractor neural networks (Amit 1989, Hopfield 1982, Hertz *et al* 1991, Rolls and Treves 1998). In experiment 2 it was shown that when the backprojections were relatively weak ($g_{1a} = 0.1$), a positive mood state held constant in the amygdala made it much more likely that a memory that was associated during storage with a positive mood state was retrieved in the IT module (see figure 4). It was also shown in experiment 2 that if the backprojections were much stronger than this, then the mood state in the amygdala actually disrupted retrieval in the IT (see figures 5 and 6). The reason for this is that with each amygdala mood state being associated by the backprojections with so many (50 in the simulations) different IT patterns, the mood association drew together the IT patterns in correlation space, and made each IT pattern more difficult to retrieve and maintain individually. This emphasizes that in the brain when mood states influence recall, because there are many fewer mood states than memories that can be recalled, the backprojections between the two modules must be kept relatively weak. Thus a reason for backprojection connections to be relatively weak from the mood system is that otherwise all the retrieved memories that are associated with a particular mood will become too similar to each other, and will not be separably retrievable.

In the second type of simulation, the dynamical behaviour of the interaction between mood and memory was investigated (section 3.2). All the investigations described in this paper were performed by presenting the retrieval cue (to the IT) at the first timestep, and then removing the retrieval cue to investigate how recall proceeded in the absence of the input stimulus to the IT. (This is described as retrieval under conditions when the retrieval cue is unclamped.)

The rationale for this is that a retrieval stimulus may be present in the real world for only a short time in relation to the long periods for which a mood state may persist. In experiment 3 the retrieval cue was applied to the IT at timestep 1, and the amygdala was in a quiescent state. The effects of the retrieval cue in the IT were transmitted by the forward projections to the amygdala module, and the correct mood state (i.e. the mood state present with that pattern during learning) was produced in the amygdala (see figure 7). However, because all the connection strength parameters g including g_{1a} were 1.0, the mood state produced in the amygdala prevented good pattern-specific retrieval in the IT module (see figure 8), for the reasons described above. In experiment 4, it was shown that if g_{1a} was reduced to 0.1, then good retrieval could be produced in IT. However, interestingly, because the mood state in the amygdala was not present at timestep 1, the IT retrieval occurred with little influence of the mood state being gradually produced in the amygdala, with IT retrieval results shown in figure 9 similar to the good retrieval shown in figure 3. This demonstrates an interesting aspect of the dynamics of the evolution of recall in such networks.

In section 3.2.2 the dynamical behaviour of the coupled attractor networks was investigated when a mood was present in the amygdala at the start of retrieval. The mood *might* be retained after timestep 1 in the amygdala as a result of the attractor properties of the amygdala network (or because of an appropriate input from the IT network), but there was no persistent (clamping) external input to the amygdala network. It was found in experiment 5 that with all connection strength parameters set to unity the amygdala firing persisted in its attractor independently of which retrieval cue pattern was presented to the IT, and that the continuing mood state disrupted IT retrieval. When g_{1a} was decreased to 0.1 in experiment 6(i), retrieval was then much better in the IT (and as good as that in experiment 2), and at the same time the mood state initiated in the amygdala at timestep 1 persisted. In experiment 6(ii) it was shown that a similar effect could be achieved by increasing g_{11} . In experiment 6(iii) it was demonstrated that relatively large values of g_{a1} were required to affect the mood state represented in the amygdala. In particular, it was found that with g_{a1} as high as 10, the cue presented to the IT was unable to affect the mood state represented in the amygdala. However, as the strength of g_{a1} was increased to the order of 100, the initial cue presented to the IT could flip the initial mood state to that associated with the IT cue pattern during learning. However, for all values of g_{a1} , the backprojections from the amygdala to the IT were able to prevent full recall in the IT, which was left in an amalgam state. In experiment 6(iv) it was shown that a mood state presented at timestep 1 to the amygdala was sufficient to disrupt IT retrieval, because with the strong backprojections from the amygdala to the IT, the IT entered an amalgam state that reflected all the IT patterns that had been associated with the mood presented at timestep 1 in the amygdala.

The generic conclusion from experiment 6, from the further results described in experiment 7 and from the other results in this paper is that for useful interactions between attractor networks, the inter- (between-) module connections must be much weaker than the intra- (within-) module connections, otherwise the effects of the attractors will be dominated by the between-module connections. This may be particularly important for backprojections in general in the brain, which this investigation predicts must be weak relative to the internal recurrent collateral connections (and to the forward inputs if they are still present), if the backprojections are to initiate retrieval or in other ways to beneficially influence and not disrupt memory retrieval and other effects produced by within-module recurrent collateral synapses. Forward inputs to a module may operate appropriately if they are stronger, for then they will push the attractor network to at least the edge of a basin of attraction, away from the current state. Then when the forward input is removed, for example during short-term memory operations or by synaptic adaptation as discussed below, the attractor effects implemented by

the recurrent collaterals within a module can operate to maintain the memory, and let it fall towards its basin of attraction. Strong forward inputs may also be needed during the original learning, so that these inputs, and not the internal recurrent collateral connections, dominate the firing during learning (Treves and Rolls 1992, 1994, Rolls and Treves 1998). It may be consistent with these requirements that the forward projections onto pyramidal cells in cortical networks terminate closer to the cell body than do the backward connections, which typically terminate in layer 1 on the distal extremity of the apical dendrite. Among other effects, this may mean that the backward connections are relatively weak, especially when there are forward inputs to a cell, and recurrent collateral inputs too, both of which would shunt the synaptic currents produced in the distal part of the apical dendrites by the backward connections.

The memory attractor network (IT) operated in the simulations described here under unclamped conditions, that is the retrieval cue was presented at timestep 1, and then removed. Part of the rationale for this was that mood states can persist for much longer than does the retrieval of each particular memory. Operation under unclamped conditions can be useful because it enables the retrieved memory to be recalled perfectly, whereas if an incomplete or noisy retrieval cue is left clamped to the retrieving neurons, the resulting recall state includes both the probably correct memory being retrieved from the recurrent collaterals, and the somewhat incorrect retrieval cue (see Rolls and Treves 1998, section 3.2 and equation 3.3). When networks operate under clamped conditions (not studied in this paper), it is normal to think of the external clamping input as producing activation that adds to the activation produced by the recurrent collateral effect, so that the final recall state is influenced by both the external ('clamping') and internal (recurrent collateral) inputs (Rolls and Treves 1998). We suggest in the context of coupled attractor networks that one way that the brain might operate in part in this way is by making it a property of the forward and backward, i.e. external, inputs to attractor networks that firing rate (or synaptic) adaptation (see Markram *et al* 1998, Abbott *et al* 1997) is present, but that it is relatively absent from the recurrent collaterals that implement the attractor state. In addition, the degree of adaptation being shown by the forward and backprojection inputs to a cortical recurrent net might be influenced by cholinergic inputs to the system. This could operate usefully as follows (cf Barkai and Hasselmo 1997). During learning, when acetylcholine release in the cortex might be expected to be high, there would be little adaptation of the forward and backward projections, which would then have a strong influence in setting the firing rates of the neurons in the recurrent network. The recurrent collaterals would then learn the firing rates of the neurons being produced by the external inputs. In contrast, during later retrieval, when the acetylcholine is low, the forward and backward external inputs to the network would show rapid adaptation, allowing the recurrent collaterals then to have a large influence on the firing. In this way acetylcholine might make the system sensitive to external inputs (introduced by forward and backward projections) during learning, and make the system more sensitive to intramodular recurrent collateral effects during memory retrieval.

The networks described in this paper have been considered in the simulations as discrete timestep networks, where the firing rates are calculated at discrete timesteps based on the states of the other neurons in the network. This approach enables what happens at each timestep to be followed. Given that real neurons in the brain have continuous dynamics in that the synaptic conductances operate in continuous time (with time constants in the region of 7–10 ms) and also are kept close to threshold in that they typically have low spontaneous firing rates, the approach to the global attractor is very fast, taking approximately two time constants of the synapses (Treves 1993, Treves *et al* 1997, Battaglia and Treves 1998, Rolls and Treves 1998). It will be of interest in future studies to investigate whether some of the effects described here that are related to the unclamping of the IT retrieval cue after timestep 1 would all operate similarly. Some insight might be gained by a discrete timestep simulation

with the IT retrieval cue left clamped, but it would be very interesting now that some of the fundamental phenomena have been described to repeat the simulations using similar networks implemented with integrate-and-fire (Battaglia and Treves 1998) neurons.

The simulations described in this paper address how the memories that are recalled depend on the mood state present when the memory was stored. This is described as mood-dependent memory (see pp 292–4 of Power and Dalgleish 1997). Another way in which mood can interact with memory is that semantically (or intrinsically) happy memories are recalled when one is happy, and semantically sad memories are recalled when one is sad. This is described as a mood-congruency effect of mood on memory recall (see pp 292–4 of Power and Dalgleish 1997). We conjecture that mood congruency effects are produced as follows. Semantically happy memories have component features that are generally as a result of previous long-term learning associated with happy mood states. Then if one is sad when a new memory is formed that includes some of the previous happy-associated features, then that particular memory will be recalled better later when one is happy, because of the long-term associations of at least some of the component features with a happy mood state. In this case, the effect of mood on memory would be described as a mood-congruency effect rather than as a mood-dependent effect.

The investigations described in this paper show that one advantage of multimodular attractor architectures which interact is that the different modules can be performing different functions, and yet interactions can nevertheless be potentially useful. In particular, in the simulations described, the IT module was taken as a module that represented memories, and the amygdala as a module that represented moods. Interaction between mood and memory recall is one way in which the continuation of mood states, which might be adaptive, might be facilitated (see chapter 3 of Rolls 1999). The attractor function within the memory module would help with the completion, generalization and graceful degradation of memories (see Rolls and Treves 1998). The attractor function within the mood module would enable the mood to persist. Another use of coupled attractor modules is in short-term memory, in which it may be important to separate a short-term memory module instantiated in the prefrontal cortex which may need to remain in its attractor state to implement the short-term memory of a sample stimulus, while a connected IT module has to respond to not only the sample stimulus but also to intervening stimuli until a match stimulus is presented (Renart *et al* 2000). In other architectures, it may be useful to have a hierarchy of connected modules in order to perform a computation to provide for example invariant visual representations of objects (Wallis and Rolls 1997, Elliffe *et al* 2001). Within such a multilayer architecture, recurrent collaterals within each module may be useful to provide constraint satisfaction between different features and short-term memory traces potentially useful for the invariance learning, and backprojections in the hierarchy may be useful to implement top-down constraint satisfaction and memory recall (see chapter 10 of Rolls and Treves 1998, Renart *et al* 1999a).

Acknowledgments

This research was supported by the Medical Research Council, grant PG9826105, by the MRC Interdisciplinary Research Centre for Cognitive Neuroscience and by the Human Frontier Science Programme.

References

- Abbott L F, Varela J A, Sen K and Nelson S B 1997 Synaptic depression and cortical gain control *Science* **275** 220–4
Amaral D G, Price J L, Pitkanen A and Carmichael S T 1992 Anatomical organisation of the primate amygdaloid complex *The Amygdala* ed J P Aggleton (New York: Wiley) pp 1–66

- Amit D J 1989 *Modeling Brain Function. The World of Attractor Neural Networks* (Cambridge: Cambridge University Press)
- Barkai E and Hasselmo M H 1997 Acetylcholine and associative memory in the pyriform cortex *Mol. Neurobiol.* **15** 17–29
- Battaglia F and Treves A 1998 Rapid stable retrieval in high-capacity realistic associative memories *Neural Comput.* **10** 431–50
- Beck A T, Emery G and Greenberg R L (eds) 1985 *Anxiety Disorders and Phobias: A Cognitive Perspective* (New York: Basic Books)
- Blaney P H 1986 Affect and memory: a review *Psychol. Bull.* **99** 229–46
- Booth M C A and Rolls E T 1998 View invariant representations of familiar objects by neurons in the inferior temporal visual cortex *Cereb. Cortex* **8** 510–23
- Bradley B P and Mathews A 1988 Memory bias in recovered clinical depressives *Cognition Emotion* **2** 235–45
- Cloitre M and Liebowitz M R 1991 Memory bias in panic disorders: an investigation of the cognitive avoidance hypothesis *Cogn. Ther. Res.* **15** 371–86
- Desimone R 1991 Face-selective cells in the temporal cortex of monkeys *J. Cogn. Neurosci.* **3** 1–8
- Elliffe M C M, Rolls E T and Stringer S M 2001 Invariant recognition of feature combinations in the visual system *Biol. Cybern.* at press
- Hertz J, Krogh A and Palmer R G 1991 *Introduction to the Theory of Neural Computation* (Wokingham: Addison-Wesley)
- Hopfield J J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci. USA* **79** 2554–8
- Markram H, Gupta A, Uziel A, Wang Y and Tsodyks M 1998 Information processing with frequency-dependent synaptic connections *Neurobiol. Learn. Mem.* **70** 101–12
- McNally R J 1997 Memory and anxiety disorders *Phil. Trans. R. Soc. B* **352** 1755–9
- Perrett D I, Rolls E T and Caan W 1982 Visual neurones responsive to faces in the monkey temporal cortex *Exp. Brain Res.* **47** 329–42
- Power M and Dalgleish T 1997 *Cognition and Emotion* (Hove: Psychology Press)
- Renart A, Parga N and Rolls E T 1999a Associative memory properties of multiple cortical modules *Network* **10** 237–55
- 1999b Backprojections in the cerebral cortex: implications for memory storage *Neural Comput.* **11** 1349–88
- 2000 A recurrent model of the interaction between the prefrontal cortex and inferior temporal cortex in delay memory tasks *Adv. Neural Inform. Process. Syst.* **12** at press
- Rolls E T 1992 Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical areas *Phil. Trans. R. Soc. B* **335** 11–21
- 1999 *The Brain and Emotion* (Oxford: Oxford University Press)
- 2000a Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition *Neuron* **27** 205–18
- 2000b Neurophysiology and functions of the primate amygdala, and the neural basis of emotion *The Amygdala: A Functional Analysis* ed J P Aggleton (Oxford: Oxford University Press) pp 447–78
- Rolls E T and Tovee M J 1995 Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex *J. Neurophysiol.* **73** 713–26
- Rolls E T and Treves A 1998 *Neural Networks and Brain Function* (Oxford: Oxford University Press)
- Tanaka K, Saito H, Fukada Y and Moriya M 1991 Coding visual images of objects in the inferotemporal cortex of the macaque monkey *J. Neurophysiol.* **66** 170–89
- Treves A 1993 Mean-field analysis of neuronal spike dynamics *Network* **4** 259–84
- Treves A and Rolls E T 1992 Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network *Hippocampus* **2** 189–99
- 1994 A computational analysis of the role of the hippocampus in memory *Hippocampus* **4** 374–91
- Treves A, Rolls E T and Simmen M 1997 Time for retrieval in recurrent associative memories *Physica D* **107** 392–400
- Wallis G and Rolls E T 1997 A model of invariant object recognition in the visual system *Progress Neurobiol.* **51** 167–94
- Watkins P C, Vache K, Verney S P, Muller S and Mathews A 1996 Unconscious mood-congruent memory bias in depression *J. Abnorm. Psychol.* **105** 34–41
- Williams J M G, Watts P N, Macleod C and Mathews A 1997 *Cognitive Psychology and Emotional Disorders* 2nd edn (Chichester: Wiley)