

Why Does Brain Damage Impair Memory? A Connectionist Model of Object Recognition Memory in Perirhinal Cortex

Rosemary A. Cowell,¹ Timothy J. Bussey,^{2,3} and Lisa M. Saksida^{2,3}

¹Department of Experimental Psychology, University of Oxford, Oxford OX1 3UD, United Kingdom, ²Department of Experimental Psychology, University of Cambridge, Cambridge CB2 3EB, United Kingdom, and ³The Medical Research Council and Wellcome Trust Behavioural and Clinical Neuroscience Institute, University of Cambridge, Cambridge CB2 3EB, United Kingdom

Object recognition is the canonical test of declarative memory, the type of memory putatively impaired after damage to the temporal lobes. Studies of object recognition memory have helped elucidate the anatomical structures involved in declarative memory, indicating a critical role for perirhinal cortex. We offer a mechanistic account of the effects of perirhinal cortex damage on object recognition memory, based on the assumption that perirhinal cortex stores representations of the conjunctions of visual features possessed by complex objects. Such representations are proposed to play an important role in memory when it is difficult to solve a task using representations of only individual visual features of stimuli, thought to be stored in regions of the ventral visual stream caudal to perirhinal cortex. The account is instantiated in a connectionist model, in which development of object representations with visual experience provides a mechanism for judgment of previous occurrence. We present simulations addressing the following empirical findings: (1) that impairments after damage to perirhinal cortex (modeled by removing the “perirhinal cortex” layer of the network) are exacerbated by lengthening the delay between presentation of to-be-remembered items and test, (2) that such impairments are also exacerbated by lengthening the list of to-be-remembered items, and (3) that impairments are revealed only when stimuli are trial unique rather than repeatedly presented. This study shows that it may be possible to account for object recognition impairments after damage to perirhinal cortex within a hierarchical, representational framework, in which complex conjunctive representations in perirhinal cortex play a critical role.

Key words: inferotemporal cortex; configural; hippocampus; visual perception; medial temporal lobe; ventral visual stream; interference

Introduction

Why does brain damage impair memory? Object recognition is thought to be the canonical test of declarative memory, the type of memory putatively impaired after damage to the temporal lobes. Studies of object recognition memory have helped to elucidate the specific anatomical structures involved in declarative memory implicating, in particular, the perirhinal cortex (Zola-Morgan et al., 1989b; Gaffan and Murray, 1992; Meunier et al., 1993; Mumby and Pinel, 1994; Aggleton et al., 1997; Baxter and Murray, 2001; Málková et al., 2001; Winters et al., 2004). Furthermore, electrophysiological data have identified properties of neurons that seem likely to form part of the mechanism underlying recognition memory (Brown and Aggleton, 2001). However, no full mechanistic account has been provided that explains why impairments after damage to perirhinal cortex should be exacerbated not only by lengthening the delay between presentation of to-be-remembered items and test (Meunier et al., 1993; Mumby and Pinel, 1994) but also by lengthening the list of to-be-remembered items (Meunier et al., 1993), or why such impair-

ments are only revealed when stimuli are trial unique rather than repeatedly presented (Eacott et al., 1994).

The aim of the present study is to offer a mechanistic account of these effects of perirhinal cortex damage on object recognition memory. To do this, we began with the assumption that perirhinal cortex houses representations of the conjunctions of visual features possessed by complex objects. Such representations are proposed to play an important role in memory when it is difficult to solve a task using only the representations of individual visual features of stimuli, thought to be stored in regions of the ventral visual stream (VVS) caudal to perirhinal cortex (see Fig. 1). The perceptual-mnemonic feature-conjunction (PMFC) model (Bussey and Saksida, 2002; Bussey et al., 2002, 2003), which formalizes these assumptions in a connectionist model, has been found to account for the effects of perirhinal cortex lesions on visual discrimination learning (Bussey et al., 2002, 2003; Barense et al., 2005; Bussey and Saksida, 2005; Lee et al., 2005).

In the present study, we ask whether this same representational framework can provide an account of object recognition memory. In the PMFC model, visual object representations are hardwired and static; the only learning that occurs in the model is the formation of associations between visual representations and reward. In the current model, we investigated the development of those object representations with visual experience. It is demonstrated that the notion of complex conjunctive representations

Received July 3, 2006; revised Oct. 2, 2006; accepted Oct. 10, 2006.

R.A.C. was supported by a Newton Abraham Studentship in Biological Sciences at Oxford University and European Commission Sixth Framework NEST Grant 516542.

Correspondence should be addressed to Rosemary A. Cowell, LEAD—CNRS (UMR 5022), Université de Bourgogne, Pôle AAFE, Esplanade Erasme, 21065 Dijon, France. E-mail: rosemary.cowell@u-bourgogne.fr.

DOI:10.1523/JNEUROSCI.2818-06.2006

Copyright © 2006 Society for Neuroscience 0270-6474/06/2612186-12\$15.00/0




Figure 1. The assumed organization of object feature representations in the VVS. Simple features, labeled A–D, are represented in caudal regions of the VVS. More complex representations of the conjunctions of those features are stored in more rostral regions, with complexity reaching a maximum in perirhinal cortex.

resolving feature ambiguity may be able to explain the effects of perirhinal cortex lesions on object recognition memory.

Materials and Methods

General methods

The current model assumes that regions of the ventral visual stream, including perirhinal cortex, contain visual representations that may be used in the service of visual recognition memory. In a similar manner to the PMFC model (Bussey and Saksida, 2002), we assume a hierarchical organization of visual representations, in which simple features are represented in caudal regions of the ventral visual stream, and representations of the conjunctions of those features are stored in more rostral regions (Fig. 1). As in the PMFC model, we do not view perirhinal cortex as the only region that contains conjunctive representations but as a region that contains perhaps the most complex conjunctive representations in the ventral visual stream.

In the current model, as in the PMFC model, the system of representations shown in Figure 1 is reduced to a two-stage scheme, in which the

PRh Layer Caudal Layer




Figure 2. The architecture of the model. The input layer, containing eight nodes, is shown on the far right; the two layers of stimulus representations [perirhinal cortex (PRh) and caudal] are shown to the left of the input layer. Stimulus inputs to the network have eight “stimulus dimensions” (attributes); each dimension is represented in the diagram by an individual input node. Stimulus dimensions are paired into four features. Each feature is shown in a distinct color and is represented individually on the caudal layer. On the perirhinal cortex layer, the four features are combined into a conjunction, shown in gray, which represents the whole stimulus.

first layer corresponds to a caudal region of ventral visual stream and the second layer corresponds to perirhinal cortex (Fig. 2). The caudal layer of the model combines two stimulus dimensions into a single representation; we refer to such two-dimensional combinations as “features” hereafter. (Note that, as in the PMFC model, the choice of the word “features” is arbitrary and is not intended to indicate specific entities, e.g., visual “primitives.” Synonyms include “elements” and “components.”) The perirhinal cortex layer combines eight stimulus dimensions into a single representation, forming a unique and fully specified representation of a visual object possessing four features. Real-world objects may be thought to contain more features than this, but the model is designed to illustrate a principle rather than reproduce the real-world situation strictly veridically. Thus, as in the PMFC model, the perirhinal cortex layer contains conjunctive representations of those visual features that are represented individually in the more caudal layer. The following experiments test whether the effects of damage to perirhinal cortex on object recognition memory can be explained by the following idea: that increasing a delay between sample and choice, or increasing the length of the list of stimuli to be remembered, or using trial-unique stimuli taxes the representational system in such a way that it becomes increasingly difficult to judge whether an object has been seen before using the representations of features alone. According to the current model, however, complex conjunctive representations residing in perirhinal cortex are useful in solving the task under these conditions. This is the same principle that is central to the explanation by the PMFC model of the effects of perirhinal cortex lesions on visual discrimination learning.

Laboratory tests of recognition memory require a judgment of previous occurrence for their solution; such a judgment is commonly tested by presenting a sample object and then later asking the subject whether that object has been encountered previously (Sidman et al., 1968; Gaffan, 1974; Mishkin and Delacour, 1975; Reed et al., 1997; Buffalo et al., 1998). The current model thus requires a mechanism for judging previous occurrence, because this phenomenon cannot be achieved with the static object representations assumed by the PMFC model. Representations in the current model are therefore allowed to develop with visual experience;

this is achieved by implementing a self-organizing Kohonen grid at each layer (Kohonen, 1984). Kohonen grids are designed to model cortex, including computational abstractions of cortical mechanisms such as lateral inhibition; this type of network is therefore appropriate for the current investigation. Each Kohonen grid comprises a two-dimensional array of processing units that receives stimulus inputs and is characterized by lateral inhibitory feedback between neighboring units. The grids are trained by the successive presentation of a number of stimulus inputs; weights of the units are incrementally adapted on each presentation. This results in an automatic mapping of stimulus inputs onto a set of representations that possess the same topological order as the stimuli, that is, similar stimuli are represented in neighboring locations on the grid. The self-organization process involves the sharpening of representations of stimuli on which the network is trained. A novel stimulus will elicit a moderate level of activity, broadly distributed across a large number of units in the grid (Fig. 3, top); as that stimulus is presented repeatedly, the activation pattern it elicits becomes more selective until only a small area of the grid contains highly active units, producing a peak of activation (Fig. 3, bottom). The development of sharply tuned representations thus can be used as the basis for familiarity judgments: as a stimulus representation becomes sharper, so it is judged to be more familiar (Norman and O’Reilly, 2003).

Simulations. In each of the simulations in this article, we lesion the component of the network corresponding to perirhinal cortex by removing the layer completely. This manipulation corresponds to the procedure followed in monkey and rat experiments, which involve total lesion of perirhinal cortex. The effects of this lesion on memory in the model are then compared with the observed effects of lesions of perirhinal cortex in rats and monkeys.

Control networks comprise both perirhinal cortex and caudal layers. In a control network, training proceeds in parallel on the two layers, and recognition performance is determined by averaging the recognition scores elicited by sample and novel stimuli on the two layers. In lesioned networks, only the caudal layer is trained, and performance is based on the caudal layer alone. Because the layers operate in parallel, lesioning the perirhinal cortex layer does not affect the function of the caudal layer. The computational details of the connectionist network are provided in Appendix 1.

We simulated three well established effects of perirhinal cortex lesions: a delay-dependent impairment in recognition memory; an increase in memory impairment with longer lists of sample stimuli; and a selective impairment in memory for trial-unique stimuli but not for repeatedly presented items.

Object stimuli. All object stimuli in these experiments were created by constructing four-featured objects from a pool of 16 possible visual features. Each feature comprises two stimulus dimensions or attributes (one might think of these as a color and a line orientation, although we are not making specific claims about the exact nature of these features; see above), and each four-featured object comprises eight stimulus attributes. The caudal layer receives two-dimensional inputs, and the perirhinal cortex layer receives eight-dimensional inputs. Thus, on the caudal layer, each two-dimensional feature is represented as a simple conjunction, and a four-featured object is represented as four separate simple conjunctions. On the perirhinal cortex layer, a four-featured object is represented as a single complex conjunction.

Experiment 1: delay-dependent impairments in object recognition memory

A delay-dependent impairment in recognition memory after brain damage is taken as a necessary and sufficient condition for demonstrating the involvement of that brain region in recognition memory. The assumption is that increasing the delay between sample presentation and judgment of recognition increases the load on memory (Gaffan, 1974) and thus increases the extent to which a putative memory system is taxed. Moreover, a lack of memory impairment at short delays is taken to indicate the absence of gross perceptual impairments. Numerous studies have found delay-dependent memory impairments after perirhinal cortex lesions in humans (Buffalo et al., 1998), in monkeys (Meunier et al., 1993; Eacott et al., 1994; Buffalo et al., 2000; Málková et al., 2001), and in rats (Mumby and Pinel, 1994; Wiig and Bilkey, 1994; Ennaceur et al., 1996).

In the present experiment, we investigate whether these data can be accounted for by a hierarchical representational account. In these simulations, we assume that, in the delay after exposure to a sample object, the regions of cortex in which the sample representation has been encoded are activated by a series of other visual stimuli. This activation corresponds to that which would be expected during a delay period when the subject is free to observe or imagine a variety of visual images. The Kohonen grids in these regions are tuned a minimal amount by each of the visual representations that is “played” in the cortex, and this sequential tuning to the series of images affects the object representations in the model, thereafter influencing the assessment by the network of the familiarity of the sample object.

The assumption that forgetting over a delay is caused by interference from events that interpose between encoding and retrieval is in line with interference theory accounts, which have been invoked in both explanations of normal human cognition (Jenkins and Dallenbach, 1924; McGeogh and McDonald, 1931; Loftus, 1977) and accounts of amnesia in brain-damaged subjects (Warrington and Weiskrantz, 1974).

Simulation. Four pairs of stimuli were created for the present experiment. Each pair comprised a four-featured sample object and a four-

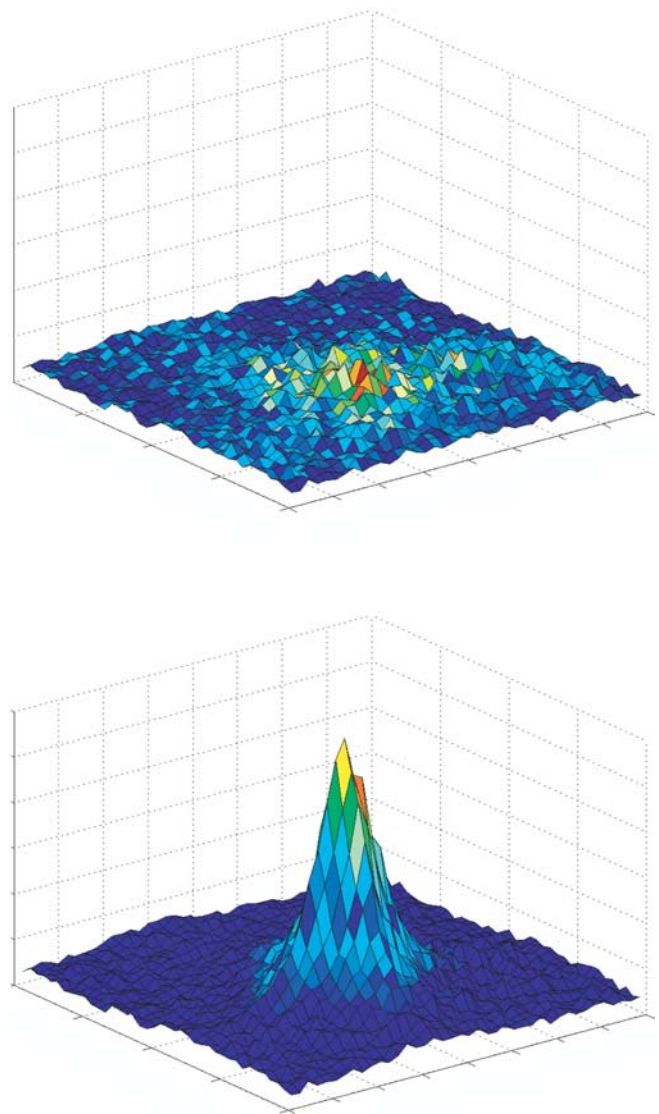


Figure 3. Stimulus representations in the model shown as activation of units in a Kohonen grid. Top, Activation pattern elicited by a sample stimulus before encoding; activity of units is broadly similar across all units with a slight peak around the winner. Bottom, The same stimulus after some encoding; the representation of the sample stimulus is sharply tuned around the winning unit.

featured novel object that shared no features with the sample object (corresponding to easily discriminable objects). Neither the sample nor the novel stimulus in any pair was replicated in any other pair, but individual features were allowed to appear in more than one object pair (for a table schematically illustrating this system of stimuli, see Appendix 2). These rules were based on a consideration of the distribution of visual features likely to be found in a set of unique junk objects: certain visual features are likely to appear more than once within the set, but it is possible to select pairs of items that share few features.

Two groups of six networks were tested: group “control” consisted of intact networks, and group “lesion” consisted of networks in which the perirhinal cortex layer had been removed to simulate perirhinal cortex lesions. Each network was tested on four object sets at each of five delay conditions, giving 20 trials per network. Each network was initialized and pretrained before testing on the 20 trials (for details, see Appendix 1). On each trial, a network was presented with the sample object and was allowed to “encode” the object for 20 cycles; each cycle sharpened incrementally the peak of activation representing the sample object (see Appendix 1). After encoding, the network was presented with “interfering” stimuli, the number of which was determined by the delay condition,

with more stimuli corresponding to a longer delay (condition 0, 0 interfering stimuli; condition 1, 200; condition 2, 400; condition 3, 600; condition 4, 800). Interfering stimuli were four-featured objects constructed from the 16 available features; they were selected at random, with replacement, from the set of all possible combinations and were presented for one encoding cycle each time they were selected. In practice, the set of possible stimuli was so large that any one object was unlikely to be selected more than once during a given delay. (Note again that the important property of the interference is that the subject is exposed to a number of visual features, and presenting a number of discrete “objects” to the network is a simple way of achieving this. In reality, of course, a subject would not be exposed only to four-featured objects.) After interference, the network was presented with both the sample and the novel object in a “choice” phase. No learning occurred in the choice phase; the representations of the two objects were simply assessed to obtain an index of their relative familiarity, which we shall call the recognition score. For derivation of the recognition score from the activation patterns elicited by the sample and novel stimuli, see Appendix 1. At the beginning of each new trial, each network was reset to the state it had assumed at the end of pretraining.

Experiment 2: effects of list length

Another experimental manipulation frequently used to demonstrate an impairment in recognition memory is increasing the number of stimuli in the list of to-be-remembered items. As with increasing delay, the assumption underlying this manipulation is that increasing the list length increases the load on memory (Gaffan, 1974). Lengthening the list thus increases the extent to which a putative memory system is taxed; therefore, the observation of a lesion-induced memory impairment that worsens with increasing list length is seen as a convincing demonstration of the involvement of a brain region in recognition memory.

Several studies have reported an impairment in recognition memory at long list length after lesions that include perirhinal cortex (Meunier et al., 1993; Eacott et al., 1994; Málková et al., 2001). The present experiment investigates whether this finding can be accounted for in terms of complex conjunctive representations stored in perirhinal cortex.

Simulation. Stimuli for the present experiment were constructed from features and organized into pairs, which were then grouped into sets of pairs of various sizes (list lengths). Each pair comprised a four-featured sample object and a four-featured novel object that shared no features with the sample object. No object stimulus in any set was replicated within that set, but individual features were allowed to appear in more than one object pair in a set (for a table schematically illustrating this system of stimuli, see Appendix 2). Four stimulus sets were constructed at each of four list lengths (1, 6, 12, and 18 pairs of stimuli), yielding 16 stimulus sets in total. Each stimulus set was presented to the networks as a list, first as a list of only the sample stimuli in the set, for encoding, then again as a list of complete pairs of stimuli for testing.

As in experiment 1, two groups of six networks were tested: group control and group lesion. Each network was initialized and pretrained, before being trained and tested on four stimulus sets at each of four list lengths. The test procedure was similar to that for experiment 1, except that networks were presented with several sample stimuli before testing for recognition of any of those stimuli (except, of course, in the case of list length 1). Thus, the procedure took the following format: the first sample stimulus in the list was presented for 20 encoding cycles, followed by presentation of the second sample stimulus for 20 encoding cycles, then the third, and so on until the end of the list of sample stimuli. Networks were not reset to their state at the end of pretraining between successive sample stimulus presentations. After all sample stimuli in the list had been encoded, a series of choice phases occurred, in which networks were presented with successive pairs of stimuli (sample and novel) from the first to the last pair in the list, without encoding. As in experiment 1, the choice phases enabled the relative familiarity of each object pair to be assessed. At the beginning of each new list, each network was reset to the state it had assumed at the end of pretraining.

Experiment 3: effects of using trial-unique versus repeated stimuli

Object recognition, as measured by the delayed nonmatching-to-sample or delayed matching-to-sample (DMS) tasks, is most often tested in

monkeys and in humans using trial-unique stimuli. In this paradigm, the set of objects from which stimuli are drawn is so large that items are repeated either extremely infrequently or not at all. It is therefore assumed that subjects effectively encounter each to-be-remembered stimulus for the first time, on all trials. Many studies have reported that damage to perirhinal cortex impairs object recognition memory with trial-unique stimuli in humans (Reed et al., 1997; Buffalo et al., 1998) and in monkeys (Zola-Morgan et al., 1989a; Meunier et al., 1993; Suzuki et al., 1993; Eacott et al., 1994; Buckley et al., 1997; Buffalo et al., 1999; Málková et al., 2001). However, a study by Eacott et al. (1994) used a small stimulus set; items were drawn repeatedly from the set and viewed many times in total by the subjects. The authors reported no effect of lesions of rhinal cortex (perirhinal and entorhinal cortex) on object recognition in this case.

In this experiment, we investigate whether the model presented in this article can account for the differential effect of perirhinal cortex lesions on recognition memory for trial-unique versus repeated stimuli.

Simulation. Stimulus pairs were constructed from features as in experiments 1 and 2, and, as previously, the sample and novel stimuli in a pair never shared any features. For the present experiment, two sets of 30 pairs of stimuli were composed: a “trial-unique” set and a “repeating” set. In the trial-unique set, no object appeared more than once. In the repeating set, the same pair of objects was presented 30 times, and the designation of novel and sample object within the pair was randomly determined from trial to trial.

Two groups of six networks (control and lesion) were pretrained and initialized, as in experiments 1 and 2. Testing for each network proceeded as follows. The network was presented with the first sample stimulus in the list for 20 encoding cycles. A delay was then simulated by the presentation of 200 interfering items (as in experiment 1) before the sample and novel stimuli from the first pair were presented to the network for assessment of their relative familiarity. After this choice, the network was presented with the second stimulus in the list for 20 encoding cycles, after which a delay of 200 interfering items was simulated and a choice between the second pair of stimuli was made, and so on until the end of the list of stimulus pairs. A delay between sample encoding and choice was simulated because the effects of perirhinal cortex lesions on recognition memory with trial-unique stimuli have typically been revealed after some delay. To allow the effects of encoding trial-unique or repeated stimuli to accumulate, networks were not reset to the state at the end of pretraining between each new stimulus pair, or trial.

Results

Experiment 1

As shown in Figure 4, removal of the perirhinal cortex layer caused impairments in object recognition performance that increased as the delay, simulated by presenting interfering information between encoding and choice, was lengthened.

Experiment 2

The simulation data are shown in Figure 5. Removal of the perirhinal cortex layer of the model caused an impairment in object recognition memory that worsened as list length increased.

Experiment 3

As shown in Figure 6, group lesion was impaired relative to group control on recognition memory for trial-unique stimuli. In contrast, with recognition memory for repeated stimuli, neither control networks nor lesioned networks could perform the task.

Discussion

Experiment 1

This experiment shows that removing the perirhinal cortex layer of the network can reproduce delay-dependent impairments in object recognition memory similar to those observed after lesions of perirhinal cortex in rats and monkeys (Meunier et al., 1993; Eacott et al., 1994; Wiig and Bilkey, 1995; Ennaceur et al., 1996;

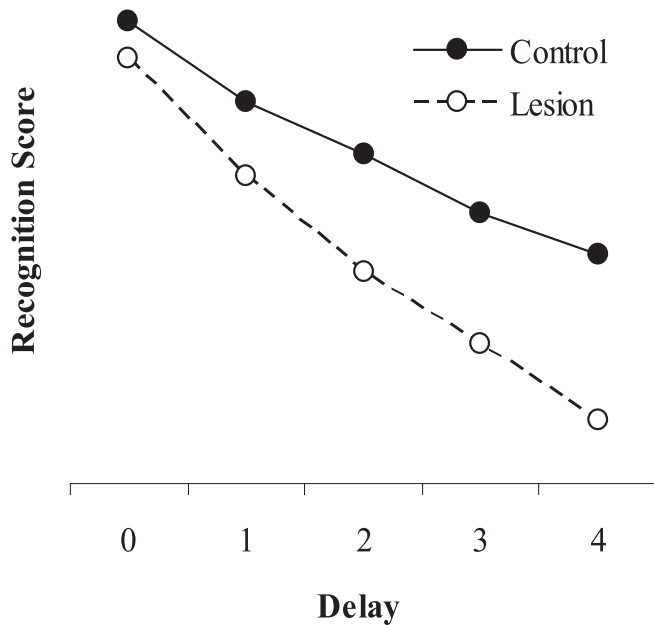


Figure 4. The effect of a simulated delay on object recognition performance by the model. Filled circles show scores from group control; open circles show scores from group lesion. Each successive delay condition includes a greater number of interfering stimuli, corresponding to a longer delay.




Figure 5. The effect of the length of the list of stimuli on object recognition performance by the model. Filled circles show scores from group control; open circles show scores from group lesion. List length refers to the number of stimulus pairs in the list for a given condition.

Buffalo et al., 2000). Specifically, the longer the delay, the greater the perirhinal cortex lesion impairment.

This effect occurs because, whereas the control networks can represent the conjunction of features of a stimulus as well as the individual features (i.e., they can represent the object ABCD, as well as the individual features A, B, C, and D), the lesioned networks can represent only the individual stimulus features (A, B,

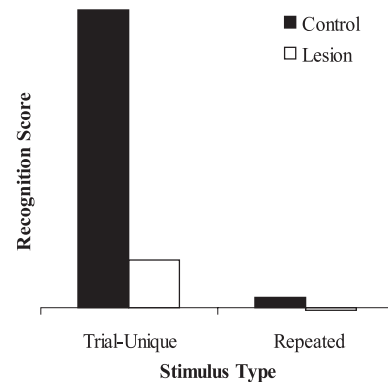


Figure 6. Recognition memory performance by the model for trial-unique stimuli and for repeated stimuli. Filled bars show recognition scores from group control; open bars show scores from group lesion. All simulations were conducted with a “delay” of 200 interfering stimuli.

C, and D). During a delay between encoding a sample stimulus and being required to discriminate that stimulus from a novel item, we assume that the subject encounters numerous visual stimuli (real or imagined) containing simple features such as edges, line orientation, and color. These simple features are in common to many visual objects and will be encountered repeatedly during the delay period. However, the specific conjunction of visual features that comprises a given complex object is unique and is encountered during the delay with a far lower frequency, if at all.

In the model, the caudal layer represents stimulus features individually; in encountering the same features many times over the course of interference during the delay, the representations of those commonly occurring features become sharply tuned on the caudal layer. In other words, to the caudal layer, all of the features appear familiar; indeed, any object encountered will seem familiar because the caudal layer cannot represent the conjunction of features that makes up an object. Conversely, visual features on the perirhinal cortex layer are represented only as part of a larger conjunction, and, because the unique conjunction that defines a complex object occurs during the delay with a very low frequency, no one object is encountered sufficiently often for its representation to become tuned. When comparing the relative familiarity indices of the sample and novel objects after a delay, representations of features on the caudal layer appear familiar in both the novel and sample objects, because all visual features have now been tuned to some extent. The caudal layer can no longer discriminate well between the sample and novel object features. In contrast, the perirhinal cortex layer continues to discriminate well because the conjunctive representation of the novel object remains untuned and hence appears unfamiliar relative to the sample representation. This mechanism is illustrated schematically in Figure 7.

Thus, an intact network, and an animal with an intact perirhinal cortex, has an advantage over a lesioned network or animal because it can represent unique conjunctions of stimulus features. In contrast, the lesioned network, or animal, must rely on the spared representations of individual stimulus features to attempt to discriminate between the novel and familiar stimuli. This advantage of the control networks or animals increases as the delay increases.

Experiment 2

Lesioned networks show an impairment in recognition memory that increases in magnitude as the list of to-be-remembered stim-




Figure 7. Stimulus representations on the Kohonen grids of the model in the choice phase, with and without a delay between the sample and choice phases. The perirhinal cortex (PRh) layer comprises one Kohonen grid with a single conjunctive representation of an object stimulus; the caudal layer comprises four Kohonen grids with separate representations of the individual features comprising the object stimulus. Small circles indicate sharply tuned (“familiar”) representations, and large circles indicate coarsely tuned (“unfamiliar”) representations.

uli lengthens. The mechanism underlying this effect is very similar to that underlying the delay-dependent impairment in experiment 1: it depends on the presence of complex conjunctive representations in perirhinal cortex and the presence of only individual stimulus features in caudal regions. When several stimuli must be encoded in memory and their representations are simultaneously stored for later retrieval, those representations are necessarily overlaid with one another in the network representing them. This gives rise to the possibility of representations interacting with each other as encoding proceeds. In the case of the perirhinal cortex layer, the complex objects presented to the networks are sufficiently unique that their representations do not overlap enough to interact significantly with one another. Thus, all sample stimuli can be easily discriminated from their novel counterparts at the choice stage. On the caudal layer, in contrast, as successive items in a list of stimuli are presented, the same commonly occurring features begin to reappear repeatedly; the representations of all stimulus features are tuned and begin to “look familiar” so that, at choice, the difference between the familiarity indices elicited by the sample and novel stimuli on the caudal layer is much reduced.

In essence, as a subject is presented with successive stimulus items in a list, there occurs a buildup of familiarity over individual features but not over unique objects. When novel versus sample stimulus pairs are presented after encoding the entire list, a discrimination between each pair may only be reliably performed on the basis of the complex conjunctive representations in perirhinal cortex and not on the basis of individual feature representations in caudal regions. As in experiment 1, an intact network, and an animal with an intact perirhinal cortex, has an advantage over a lesioned network or animal because it possesses the necessary complex conjunction of stimulus features. This advantage of intact networks or animals increases as the list of stimuli lengthens.

Experiment 3

Lesioned networks were impaired relative to control networks in the recognition of trial-unique stimuli. Importantly, neither networks in group control nor networks in group lesion could perform recognition of repeated stimuli.

The explanation of the impairment in recognition of trial-unique stimuli after removal of the perirhinal cortex layer is the same as that for experiment 2: successive stimulus presentations

leads to a buildup of familiarity over individual stimulus representations but not over unique objects. Thus, the features of both the sample and novel objects in each stimulus pair appear familiar according to caudal representations, but the conjunctive representations of the sample and novel objects on the perirhinal cortex layer remain discriminable on the basis of familiarity. Therefore, an intact network or animal has an advantage over its lesioned counterpart in the encoding of trial-unique stimuli at a short delay.

In the case of repeating-items object recognition, both group control and group lesion failed to discriminate the novel from the familiar stimulus. This finding, which implies that recognition is not dependent on perirhinal cortex in this case, is consistent with the data from Eacott et al. (1994), in which perirhinal cortex lesions in monkeys did not impair recognition of repeated stimuli.

In fact, the inability of intact networks to perform repeated-items recognition is entirely consistent with the present account; moreover, it highlights an interesting consequence of the hierarchical mechanism of the model. In all previous simulations, repeated presentation of features rendered them ambiguous with respect to object novelty judgments. This ambiguity was resolved in the previous experiments by the existence of complex conjunctive representations in the perirhinal cortex layer. In the case of repeating items, however, not only are the features repeatedly presented, the objects are repeatedly presented, and an additional level of ambiguity (“object ambiguity”) is created. Now, neither the caudal nor the perirhinal cortex layer can solve the task. Just as in the trial-unique case, resolution of ambiguity on the caudal layer requires the rostral layer, so in repeating-items object recognition, the resolution of ambiguity on the rostral (and caudal) layers of the current network would require an additional, more rostral layer. This additional layer would contain conjunctive representations of an even higher degree of complexity than those in our current rostral layer.

The monkeys in the Eacott et al. (1994) study clearly could learn repeating-items DMS, even with a perirhinal cortex lesion. Thus, some other part of the brain is capable of solving the task. What part of the brain might correspond to the layer postulated above? Such a layer would correspond to a structure “downstream” from perirhinal cortex; one obvious candidate is the hippocampus (although there are indeed several alternative candidates, including entorhinal cortex and prefrontal cortex). Many authors have suggested that the hippocampus contains complex, multimodal representations of multiple objects, perhaps including the visuospatial relationships between them (Eichenbaum et al., 1994) or forming a “cognitive map” (O’Keefe and Nadel, 1978). Several researchers have suggested a hierarchical organization of brain structures with the hippocampus at the “top” of the hierarchy (Squire, 1992; Mishkin et al., 1997; McNaughton et al., 2003). A clear prediction arises from this analysis: lesions of the hippocampus should impair repeating-items object recognition. Indeed, Rawlins et al. (1993) report that lesions of the hippocampus or fornix in rats impair performance of repeating-items DMS. Additionally, Charles et al. (2004) found that monkeys with fornix transections were impaired on a recognition memory task in which they were required to judge the relative recency of two stimuli. Such a finding does not, in our view, necessitate the assumption of a brain module for “recency memory.” According to the present account, tasks like repeating-items DMS merely provide an additional degree of ambiguity that must be resolved by even more complex conjunctive representations in a hierarchy that extends throughout the ventral visual stream

through perirhinal cortex and on into other structures such as the hippocampus (Bussey and Saksida, 2005).

The idea that the hippocampus might provide an additional level of representational complexity for the resolution of object ambiguity is entirely consistent with the idea that “episodic memory,” known to depend on the hippocampus, can be understood as a conjunction of “what, where, and when” (Tulving, 1972; Clayton et al., 2003). According to the psychological-modular view of episodic memory, episodes are unique yet share many components. For example, we can remember many different episodes that involve ourselves in a particular room in our house, yet despite the ambiguity of the component parts, we can identify these episodes as distinct and unique. Conjunctions resolve the ambiguous components of episodic memories, just as, in our view, they resolve the ambiguous components of visual memories throughout the representational hierarchy. In this way, high-level conjunctive representations in the hippocampus may, under certain circumstances, contribute to episodic recollective processes in single-item recognition tasks (Fortin et al., 2004).

General discussion

The present study tested a neural network model of object recognition in perirhinal cortex. The results indicate that it may be possible to account for object recognition impairments after damage to perirhinal cortex with a hierarchical representational account, in which complex conjunctive representations in perirhinal cortex are useful when it is difficult to solve the problem on the basis of features alone, a property we have referred to previously as feature ambiguity (Bussey and Saksida, 2002, 2005; Bussey et al., 2002, 2003). The same general model has also been shown to account for the effects of perirhinal cortex lesions on visual discrimination tasks (Buckley and Gaffan, 1997, 1998; Bussey et al., 2002, 2003).

In experiment 1, removing the perirhinal cortex layer of the network produced delay-dependent impairments in object recognition similar to those observed after lesions of perirhinal cortex (Meunier et al., 1993; Eacott et al., 1994; Mumby and Pinel, 1994; Wiig and Bilkey, 1995; Málková et al., 2001). Specifically, the longer the delay, the greater the perirhinal cortex lesion impairment. “Forgetting” in the model is caused by interference produced in the network when a subject is free to observe or imagine a variety of visual images during a delay. In this sense, the model endorses an interference account of normal forgetting (Jenkins and Dallenbach, 1924; McGeogh, 1932), although it certainly does not rule out the contribution of other factors under some circumstances. Interfering items are encoded in the caudal layer of the network as features and in the rostral layer as conjunctions of features. In a large set of interfering items, the same features will reoccur with a high frequency, but specific conjunctions of those features (corresponding to unique objects) will occur far less frequently, if at all. When, during the choice phase, the network is confronted with a choice between a novel and a familiar object, many of the features of the novel object appear familiar because they have been encountered during the delay as part of other, interfering items. Therefore, the caudal layer of the network, containing representations of features only, has great difficulty judging object novelty. However, the particular conjunction of features specifying a unique novel object is unlikely to have been encountered during the delay. Thus, the conjunctive representations in the rostral (perirhinal cortex) layer are by far the most useful representations in the network for judging object novelty, and removing the rostral layer can result in impairments in object recognition after a delay. The magnitude of impairment

increases as delay increases because, as the delay lengthens, more interfering items are encountered and the conjunctive representations in perirhinal cortex become increasingly important for resolving the interference.

In experiment 2, removing the perirhinal cortex layer of the network produced impairments in object recognition that increased in magnitude as the list of stimuli to be remembered was lengthened, mirroring the impairment seen in monkeys with perirhinal cortex damage (Meunier et al., 1993; Eacott et al., 1994; Málková et al., 2001). The mechanism underlying this effect is very similar to that underlying the delay-dependent impairment in experiment 1; it again depends on the presence of complex conjunctive representations in perirhinal cortex and only individual stimulus features in caudal regions. Similar to the buildup of interference during a delay in experiment 1, as a subject is presented with successive stimulus items in a list, there occurs a buildup of familiarity of individual features but not of complex objects. In a list of stimuli, objects are unique but share many common features; therefore, individual features are seen many times. Accordingly, the specific conjunctions of novel objects in the stimulus set are never encoded but the component features of novel objects are, because those features are shared with sample objects. For this reason, when novel versus familiar object pairs are presented after encoding the entire list, a judgment of novelty can only be reliably performed on using the complex conjunctive representations in perirhinal cortex and not on the basis of individual feature representations in caudal regions. Intact networks in experiment 2 had an advantage over lesioned networks because they possessed representations of the complex conjunctions of stimulus features on which the task solution became more dependent as the list of stimuli was lengthened.

In experiment 3, lesioned networks were impaired relative to intact networks in the recognition of trial-unique stimuli. Again, this impairment occurred because presenting the network with a long list of stimuli entails the frequent reoccurrence of common object features, causing a buildup of feature ambiguity that, in an intact network, can be resolved by conjunctive representations in perirhinal cortex. In experiment 1 feature ambiguity (interference) built up during the presentation of interfering items during the delay; in experiment 2, feature ambiguity occurred as a result of the presentation of a long list of sample stimuli. In experiment 3, the trial-unique procedure resulted in the presentation of many stimuli to the network resulting, again, in feature ambiguity. Thus, the role of the perirhinal cortex in both visual discrimination and object recognition is the resolution of feature ambiguity, the property of a problem arising when it is difficult to find a solution on the basis of features alone. In the case of visual discrimination, features are ambiguous with respect to judging whether an object predicts reward because they form part of the identity of both rewarded and unrewarded objects. In the case of object recognition memory, features are ambiguous (or less informative) with respect to making judgments of familiarity versus novelty, because all features appear familiar.

Intact and lesioned networks were also tested on repeating-items object recognition, on which monkeys with perirhinal cortex lesions have been shown to be unimpaired (Eacott et al., 1994). Interestingly, even the intact networks were unable to perform this task. This result is exactly what one would expect based on our view. In all of the previous simulations, repeated presentation of features rendered them ambiguous with respect to object novelty judgments. However, this ambiguity could be resolved by complex conjunctive (object) representations. Now, in the case of repeating items not only are the features repeatedly

presented, but the objects are repeatedly presented, and an additional level of ambiguity (object ambiguity) is created, resulting in ambiguity on both layers such that neither layer can solve the task. As discussed above, resolving such ambiguity would require representations of more complex conjunctions. We suggest the hippocampus as one potential site for the storage of such representations.

The simulation results presented here demonstrate that the model can account for three phenomena from the extant data on object recognition memory. The model also makes novel predictions that can be tested experimentally. For example, the model predicts that perirhinal cortex lesions should cause impairments on object recognition memory with zero delay, if the sample and novel objects are sufficiently perceptually similar. Moreover, the size of this impairment should increase as the degree of similarity between the sample and novel objects increases. Such a prediction is at odds with prevailing views of perirhinal cortex function and is therefore a stringent test of the model. Second, because perirhinal cortex is thought to house conjunctive (configural) stimulus representations, object recognition should be particularly sensitive to perirhinal cortex damage when the task is configural; that is, if a judgment of novelty can be made on the basis of the conjunctions of features but not on the basis of the features alone. Third, the model predicts that the amount of forgetting of a sample object caused by presenting interfering visual material that is similar to the to-be-remembered object should be greater, following perirhinal cortex lesions, than the amount of forgetting caused by presenting interfering visual material that is dissimilar to the to-be-remembered object. Each of these three predictions is empirically testable.

To conclude, the simulation experiments reported in the present article demonstrate that the canonical effects of perirhinal cortex lesions on object recognition memory may be accounted for on the basis of a few simple assumptions about the representation of visual information in the brain. Moreover, as is highlighted by the novel predictions outlined above, the model stands in contrast to the prevailing modular view of visual cognition in which the caudal ventral visual stream and perirhinal cortex putatively perform the distinct functions of visual perception and visual memory, respectively.

Appendix 1: computational details of the model

Network overview

The model is composed of two layers: “perirhinal cortex” and “caudal.” The layers operate in parallel and are constructed from two-dimensional Kohonen maps or grids (Kohonen, 1984). The perirhinal cortex layer comprises one Kohonen grid, and the caudal layer comprises four grids. The caudal layer grids each receive two-dimensional inputs, and the perirhinal cortex layer grid receives eight-dimensional inputs (Fig. 2). The layers are constructed in this way because it is assumed that both caudal regions of VVS and perirhinal cortex are capable of representing all of the visual attributes of an object in some form. Therefore, an object with four features that is represented in a single conjunctive representation on the perirhinal cortex layer must be represented as four separate features on four independent feature maps in more caudal regions of VVS.

Kohonen grids

A Kohonen grid, or map, is usually a two-dimensional representation in which multidimensional stimulus inputs are classified such that those that share similar characteristics are located in the same area of the map (Kohonen, 1984). Kohonen grids self-

organize, that is, they achieve their topological mapping of stimulus inputs via an unsupervised learning process. The two-dimensional scene is formed from a grid of units, or nodes, to which stimuli are presented. A stimulus is described by a vector of length n , which, in the present article, represents the visual characteristics of the object. Each unit in the grid is linked to each element of the stimulus vector via a weight w , so that every unit is associated with a weight vector of length n , containing the weights.

Initialization

Weights associated with all units in each Kohonen grid are initialized to random values between 0 and 1.

Pretraining, or self-organization

Training stimuli

Each of the four caudal grids is trained with two-dimensional input vectors; the perirhinal cortex grid is trained with eight-dimensional input vectors. There are 16 possible two-dimensional “visual feature” input vectors and a large number of possible eight-dimensional “visual object” input vectors (any combination of 4 of the 16 two-dimensional visual features).

The training cycle

(1) Present a training stimulus input vector, selected at random from all possible training stimuli, to the grid.

(2) The “winning unit” for that training stimulus is determined. This is the unit possessing a weight vector that most closely matches the training stimulus input vector:

$$\|input - w_{win}\| = \min_i \|input - w_i\|, \quad (1)$$

where $input$ is the stimulus vector, w_{win} is the weight vector of the winning unit, w_i is the weight vector of unit i , and i is the set of all units in the grid.

(3) The weights of the winning unit, and those of its near neighbors, are modified such that the weight vectors become more similar to the input vector of the training stimulus presented to the grid:

$$w_i(t) = w_i(t - 1) + f(r, t)(input - w_i(t - 1)), \quad (2)$$

where

$$f(r, t) = \alpha(t)v(r, t), \quad (3)$$

in which t is the number of time steps (or number of stimuli presented) since training began, $w_i(t)$ is the weight vector of unit i at time step t , $w_i(t - 1)$ is the same weight vector on the previous time step, r is the city-block distance from the winning unit, $\alpha(t)$ is the learning rate, and $v(r, t)$ is the neighborhood function. If a given unit is outside the neighborhood of the winner, $v(r, t)$ for that unit is 0 and no weight update occurs.

(4) Reduce the size of the neighborhood around the winning unit, that is, the area inside which units on the grid undergo weight modification (see below, The neighborhood function).

(5) Reduce the learning rate, $\alpha(t)$, which controls the size of the modifications applied to weight vectors (see below, The learning rate).

(6) Choose a new training stimulus and repeat steps 1–5 until 100 randomly chosen training stimuli have been presented (i.e., $t = 100$).

(7) Modify every element of each weight vector by a value

randomly chosen from the range of -1 to $+1$, to introduce some noise into the distribution of the weights.

(8) After step 7, the distribution of the weights lies in the range of -1 to $+2$. Set the extremes of the distribution of the weights back to the values 0 and 1, by adding 1 to each weight and dividing all weights by 3.

The neighborhood function, $v(r,t)$

The size of the modification of the weight vector of a unit depends on the position of that unit with respect to the winning unit: the weight vector of the winning unit is subject to the greatest modification, whereas more distant units receive little or no modification of their weight vectors. The function $v(r,t)$ is at a maximum when $r = 0$ and decreases as r increases. This weight update profile resembles a Gaussian function. In the current model, lateral inhibition is not implemented between the units directly, but the use of this neighborhood function during pretraining and encoding yields stimulus activation patterns characteristic of those that would typically emerge from a network of units with lateral inhibitory connections.

This neighborhood size of the winner is described by Equation 4, in which v , governing the size of the weight change for any unit in the grid, reduces with distance from the winner. In addition, the neighborhood size shrinks as training progresses (as t increases):

$$v(r,t) = \exp\left(-\left(\frac{r}{G(t)}\right)^2\right), \quad (4)$$

where r denotes city-block distance from the winning unit, and $G(t)$ is a parameter that reduces across training cycles according to Equation 5:

$$G(t) = 0.5 + 10t^{-B}, \quad (5)$$

where B is a constant determining the rate of reduction of G across time steps.

The learning rate, $\alpha(t)$

The learning rate decreases during the pretraining phase according to Equation 6, in which A is a constant determining the rate of decrease of α :

$$\alpha(t) = t^{-A}. \quad (6)$$

This ensures that the weights converge onto an appropriate mapping of stimulus representations without excessive oscillation about the solution, by decreasing the size of the weight adjustments as the network begins to reach a topographical organization.

Treatment of the grid edges

Units at the edges of the grid are unusual in that they do not have as many neighboring units as those in the center of the grid. This creates a discontinuity in the representation of stimulus space on the grid and can cause instabilities in the configuration of the stimulus map from one training cycle to the next. It also presents a problem when measuring the selectivity of a stimulus representation (see below, Choice phase) if the peak in activation elicited by that stimulus falls at the edge of the grid, because the summed activity of the peak will be reduced relative to centrally located representations. Therefore, we consider that a unit at the edge of the grid is neighbor to units at the opposite edge: the grid “wraps around” into a toroid formation.

Encoding and testing

Once the self-organization process described above is completed, the pretrained network is trained and tested in simulation of an object recognition task. In all simulations presented in this article, training is divided into discrete trials. Each trial involves the encoding of a sample stimulus, followed by a choice phase involving the presentation of that sample stimulus along with another novel stimulus for comparison of their relative familiarity (but see below, Simulation of list length, for a modification of this procedure).

The representation of stimuli in the model

Each sample or novel stimulus in the present article is composed of four two-dimensional visual features. It is presented to the perirhinal cortex layer as a whole, i.e., as an eight-dimensional input vector, but to the caudal layer as four separate two-dimensional input vectors, one input for each of the four Kohonen grids in the caudal layer.

The encoding cycle

A sample stimulus is encoded in a similar manner to a training stimulus in pretraining, except that the learning rate and neighborhood size do not decrease. For the caudal layer, the sample stimulus is divided into four two-dimensional input vectors corresponding to the four stimulus features; the first feature is encoded by the first caudal grid, the second feature by the second caudal grid, and so on. The encoding process occurs once on the perirhinal cortex layer but four times on the caudal layer, once for each of the four feature-grid designations.

(1) Present an input vector describing the sample stimulus (or, in the case of the caudal layer, the stimulus feature) to the grid.

(2) The winning unit for the sample stimulus (or sample stimulus feature) is determined, as in pretraining, according to Equation 1.

(3) The weights of the winner and its neighbors are modified, as in pretraining, except that learning rate and neighborhood size are held constant across encoding cycles:

$$w_i(t) = w_i(t-1) + f(r)(\text{input} - w_i(t-1)), \quad (7)$$

where

$$f(r) = \alpha v(r), \quad (8)$$

where α is a constant learning rate, t is the number of time steps since training began, and $v(r)$ is the neighborhood function. During encoding, the weight updates given by $v(r)$ are smaller for units that are farther from the winner, as in Equation 4, but G is held constant at the value that it took at the end of pretraining, i.e., the neighborhood size does not reduce over time.

(4) Repeat steps 1–3 until the same sample stimulus (or sample stimulus feature) has been presented to the grid for 20 encoding cycles.

Choice phase

Perirhinal cortex layer

(1) Present an input vector describing the sample stimulus to the grid.

(2) Calculate the activation of all units in the grid, according to Equation 9:

$$a_i = \ln\left(\frac{1}{\text{dist}}\right), \quad (9)$$

where a_i is activation of unit i , and dist is calculated according to Equation 10:

$$\text{dist} = \frac{\sum \text{samp} - w_i}{n}, \quad (10)$$

where samp is the sample stimulus input vector, w_i is the weight vector of unit i , and n is the number of elements in the input stimulus vector.

(3) Measure the peak activation and total activation elicited by the sample stimulus and calculate the selectivity of the sample representation according to Equations 11–13:

$$\text{peak} = \sum_k a_k, \quad (11)$$

$$\text{total} = \sum_i a_i, \quad (12)$$

$$S = \frac{\text{peak}}{\text{total}}, \quad (13)$$

where s is the selectivity of the representation, k is the set of units comprising the winner and its eight closest neighbors on the grid according to a city-block distance measure, and i is the set of all units on the grid.

(4) Present an input vector describing the novel stimulus to the grid.

(5) Calculate the activation of all units in the grid, according to Equations 9 and 10.

(6) Measure the peak and total activation elicited by the novel stimulus and calculate the selectivity of the novel stimulus representation according to Equations 11–13.

(7) Calculate the relative familiarity of the two stimuli, to give a recognition score, R , analogous to the difference score frequently calculated in the spontaneous object recognition task performed by rats:

$$R = \frac{S_{\text{samp}} - S_{\text{nov}}}{S_{\text{samp}} + S_{\text{nov}}}, \quad (14)$$

where S_{samp} is the selectivity of the representation of the sample stimulus and S_{nov} is the selectivity of the representation of the novel stimulus.

Caudal layer. For the caudal layer, the sample and novel stimuli are each divided into four two-dimensional input vectors corresponding to the four stimulus features; the first feature of the sample stimulus and the first feature of the novel stimulus are presented to the first caudal grid, the second feature of the sample stimulus, and the second feature of the novel stimulus are presented to the second caudal grid, and so on.

(1) Perform steps 1–6 of the choice phase described above for the perirhinal cortex layer, for each of the four caudal grids.

(2) Calculate the average selectivity of the sample and novel representations according to Equations 15 and 16, respectively:

$$S_{\text{samp}} = \frac{S_{\text{samp}^1} + S_{\text{samp}^2} + S_{\text{samp}^3} + S_{\text{samp}^4}}{4}, \quad (15)$$

$$S_{\text{nov}} = \frac{S_{\text{nov}^1} + S_{\text{nov}^2} + S_{\text{nov}^3} + S_{\text{nov}^4}}{4}, \quad (16)$$

where S_{samp^1} is the selectivity of the representation of feature 1 of the sample stimulus on grid 1 of the caudal layer, S_{samp^2} is the selectivity of sample feature 2 on grid 2, and so on.

(3) Calculate the relative familiarity of the two stimuli, according to Equation 14.

Additional procedures

Simulation of a delay

During the simulation of a delay, interfering stimuli are presented to the network between encoding and choice phase. The interference delay process is similar to pretraining, except that the learning rate and neighborhood size do not decrease, and no additional noise is added to the weights. For the caudal layer, interfering object stimuli are divided into four two-dimensional input vectors corresponding to the four stimulus features; the first feature is encoded by the first caudal grid, the second feature by the second caudal grid, and so on. The interference process occurs once on the perirhinal cortex layer and once on each of the four grids of the caudal layer.

The interference cycle

(1) Select an interfering stimulus at random from all possible four-featured object stimuli.

(2) Present the stimulus (or in the case of the caudal layer, stimulus feature) to the grid.

(3) The winning unit for the interfering stimulus (or interfering stimulus feature) is determined according to Equation 1.

(4) The weights of the winner and its neighbors are modified according to Equations 7 and 8. Note that the learning rate and neighborhood size are held constant.

(5) Repeat steps 1–4 until the desired number of interfering stimuli have been presented to the grid. The number of interfering stimuli is determined by the length of the delay.

Simulation of list length

In simulations of the effects of list length, networks were subjected to pretraining, encoding, and choice phases.

The only alteration to the standard administration of these processes was that the weights were not reset to the values they had taken at the end of pretraining for each new stimulus pair. Instead, weights were preserved from one sample stimulus in the list to the next, so that encoding of a new sample stimulus occurred on top of previous learning.

In addition, encoding of all of the sample stimuli in the list was completed before administering the choice phase for any stimulus pair. Choice phases for all stimulus pairs occurred, in order of the list, after encoding.

Simulation of trial-unique versus repeated stimuli

In simulating the effects of using trial-unique or repeated stimuli, networks underwent pretraining, encoding, delay-interference, and choice phases.

As for list length simulations, weights were not reset to the values they had taken at the end of pretraining for each new stimulus pair. Instead, weights were preserved from one stimulus pair to the next, so that encoding of a new sample stimulus occurred on top of previous learning, thus allowing the effects of encoding trial-unique versus repeated stimulus representations to accumulate.

These simulations differed from list length simulations in that the choice phase for each stimulus pair was administered directly after encoding the sample stimulus of that pair. Also, a delay was simulated between encoding and choice for each stimulus pair by presenting interfering stimuli, as in experiment 1.

Parameters

Network parameters

Each Kohonen grid in the model is a square with sides of length 200 units, giving a total of 40,000 units per grid. The edges of each

grid wrap around into a toroid shape, as explained above (this characteristic is not shown in Figs. 2 and 3 for simplicity). Each Kohonen grid has an input layer with n units. For each of the four caudal grids, $n = 6$: an input constitutes a two-dimensional visual feature with three elements per stimulus dimension (i.e., each dimension is specified by a number in triplicate that spans three elements in the input vector, so that a typical feature might take the form [0.333 0.333 0.333 0.950 0.950 0.950]). For the perirhinal grid, $n = 24$: an input constitutes an object with four two-dimensional features, with three elements per stimulus dimension.

Learning parameters

In all simulations, the parameter A , governing the reduction of the learning rate during pretraining, took the value 0.2. The parameter B , governing the reduction of the neighborhood size during pretraining, was set to 0.4. The learning rate, α , took a constant value of 0.35 in the sample phase of all simulations and a constant value of 0.05 during simulation of interference. The neighborhood size did not reduce across encoding cycles, because G was held constant at 2.085 (which was the value of G reached at the end of pretraining with 100 pretraining cycles). In addition, activation values of units were capped on all Kohonen grids at a maximum value of 9.210. In all simulations, the number of pretraining cycles was 100 and the number of encoding cycles was 20.

Appendix 2: construction of stimulus sets

Table 1. Schematic illustration of the construction of stimuli for experiment 1

Stimulus pair		Stimulus feature			
		1	2	3	4
1	Sample	K	C	H	G
	Novel	M	A	D	P
2	Sample	D	H	F	K
	Novel	B	O	L	J
3	Sample	C	J	E	F
	Novel	M	N	G	I
4	Sample	B	D	I	A
	Novel	N	G	P	L

A capital letter denotes a stimulus feature; there are 16 possible features, A–P. Each stimulus is composed of four features. No stimulus pair (sample and novel) shares any features. No stimulus is replicated within the set of eight stimuli.

Table 2. Construction of stimuli in experiment 2

Stimulus pair		Stimulus feature			
		1	2	3	4
1	Sample	K	C	H	G
	Novel	M	A	D	P
2	Sample	D	H	F	K
	Novel	B	O	L	J
3	Sample	C	J	E	F
	Novel	M	N	G	I
4	Sample	B	D	I	A
	Novel	N	G	P	L
5	Sample	P	E	J	K
	Novel	M	I	C	F
6	Sample	G	H	F	D
	Novel	N	A	C	B

Examples of stimuli in a set of list length 6 are shown. Stimulus sets with list length 1, 12, and 18 are not shown. A capital letter denotes a stimulus feature; there are 16 possible features, A–P. Each stimulus is composed of four features. No stimulus pair (sample and novel) shares any features. No stimulus is replicated within the set of stimuli, regardless of list length.

References

Aggleton JP, Keen S, Warburton EC, Bussey TJ (1997) Extensive cytotoxic lesions involving both the rhinal cortices and area TE impair recognition but spare spatial alternation in the rat. *Brain Res Bull* 43:279–287.

Barese MD, Bussey TJ, Lee AC, Rogers TT, Davies RR, Saksida LM, Murray EA, Graham KS (2005) Functional specialization in the human medial temporal lobe. *J Neurosci* 25:10239–10246.

Baxter M, Murray EA (2001) Opposite relationship of hippocampal and rhinal cortex damage to delayed nonmatching-to-sample deficits in monkeys. *Hippocampus* 11:61–71.

Brown MW, Aggleton JP (2001) Recognition memory: what are the roles of the perirhinal cortex and hippocampus? *Nat Rev Neurosci* 2:51–61.

Buckley MJ, Gaffan D (1997) Impairment of visual object-discrimination learning after perirhinal cortex ablation. *Behav Neurosci* 111:467–475.

Buckley MJ, Gaffan D (1998) Perirhinal cortex ablation impairs configural learning and paired-associate learning equally. *Neuropsychologia* 36:535–546.

Buckley MJ, Gaffan D, Murray EA (1997) Functional double dissociation between two inferior temporal cortical areas: perirhinal cortex versus middle temporal gyrus. *J Neurophysiol* 77:587–598.

Buffalo EA, Reber PJ, Squire LR (1998) The human perirhinal cortex and recognition memory. *Hippocampus* 8:330–339.

Buffalo EA, Ramus SJ, Clark RE, Teng E, Squire LR, Zola SM (1999) Dissociation between the effects of damage to perirhinal cortex and area TE. *Learn Mem* 6:572–599.

Buffalo EA, Ramus SJ, Squire LR, Zola SM (2000) Perception and recognition memory in monkeys following lesions of area TE and perirhinal cortex. *Learn Mem* 7:375–382.

Bussey TJ, Saksida LM (2002) The organization of visual object representations: a connectionist model of effects of lesions in perirhinal cortex. *Eur J Neurosci* 15:355–364.

Bussey TJ, Saksida LM (2005) Object memory and perception in the medial temporal lobe: an alternative approach. *Curr Opin Neurobiol* 15:1–8.

Bussey TJ, Saksida LM, Murray EA (2002) Perirhinal cortex resolves feature ambiguity in complex visual discriminations. *Eur J Neurosci* 15:365–374.

Bussey TJ, Saksida LM, Murray EA (2003) Impairments in visual discrimination after perirhinal cortex lesions: testing “declarative” vs. “perceptual-mnemonic” views of perirhinal cortex function. *Eur J Neurosci* 17:649–660.

Charles DP, Gaffan D, Buckley MJ (2004) Impaired recency judgments and intact novelty judgments after fornix transection in monkeys. *J Neurosci* 24:2037–2044.

Clayton NS, Bussey TJ, Dickinson A (2003) Can animals recall the past and plan for the future? *Nat Rev Neurosci* 4:685–691.

Eacott MJ, Gaffan D, Murray EA (1994) Preserved recognition memory for small sets, and impaired stimulus identification for large sets, following rhinal cortex ablations in monkeys. *Eur J Neurosci* 6:1466–1478.

Eichenbaum H, Otto T, Cohen NJ (1994) Two functional components of the hippocampal memory system. *Behav Brain Sci* 17:449–518.

Ennaceur A, Neave N, Aggleton JP (1996) Neurotoxic lesions of the perirhinal cortex do not mimic the behavioural effects of fornix transection in the rat. *Behav Brain Res* 80:9–25.

Fortin NJ, Wright SP, Eichenbaum H (2004) Recollection-like memory retrieval in rats is dependent on the hippocampus. *Nature* 431:188–191.

Gaffan D (1974) Recognition impaired and association intact in the memory of monkeys after transection of the fornix. *J Comp Physiol Psychol* 86:1100–1109.

Gaffan D, Murray EA (1992) Monkeys (*Macaca fascicularis*) with rhinal cortex ablations succeed in object discrimination learning despite 24-hr inter-trial intervals and fail at matching to sample despite double sample presentations. *Behav Neurosci* 106:30–38.

Jenkins JB, Dallenbach KM (1924) Oblivescence during sleep and waking. *Am J Psychol* 35:605–612.

Kohonen T (1984) Self-organization and associative memory. Berlin: Springer.

Lee AC, Bussey TJ, Murray EA, Saksida LM, Epstein RA, Kapur N, Hodges JR, Graham KS (2005) Perceptual deficits in amnesia: challenging the medial temporal lobe “mnemonic” view. *Neuropsychologia* 43:1–11.

Loftus EF (1977) Shifting human colour memory. *Mem Cognit* 5:696–699.

Málková L, Bachevalier J, Mishkin M, Saunders RC (2001) Neurotoxic lesions of perirhinal cortex impair visual recognition memory in rhesus monkeys. *NeuroReport* 12:1913–1917.

- McGeogh JA (1932) Forgetting and the law of disuse. *Psychol Rev* 39:352–370.
- McGeogh JA, McDonald WT (1931) Meaningful relation and retroactive inhibition. *Am J Psychol* 43:579–588.
- McNaughton BL, Barnes CA, Battaglia FP, Bower MR, Cowen SL, Ekstrom AD, Gerrard JL, Hoffman KL, Houston FP, Karten Y, Lipa P, Pennartz CMA, Sutherland GR (2003) Offline reprocessing of recent memory and its role in memory consolidation: a progress report. In: *Sleep and brain plasticity* (Maquet P, Smith C, Stickgold R, eds), pp 225–246. Oxford: Oxford UP.
- Meunier M, Bachevalier J, Mishkin M, Murray EA (1993) Effects on visual recognition of combined and separate ablations of the entorhinal and perirhinal cortex in rhesus monkeys. *J Neurosci* 13:5418–5432.
- Mishkin M, Delacour J (1975) An analysis of short-term visual memory in the monkey. *J Exp Psychol Anim Behav Process* 1:326–334.
- Mishkin M, Suzuki WA, Gadian DG, Vargha-Khadem F (1997) Hierarchical organization of cognitive memory. *Philos Trans R Soc Lond B Biol Sci* 352:1461–1467.
- Mumby DG, Pinel JPJ (1994) Rhinal cortex lesions and object recognition in rats. *Behav Neurosci* 108:11–18.
- Norman KA, O'Reilly RC (2003) Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychol Rev* 110:611–646.
- O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Oxford: Oxford UP.
- Rawlins JN, Lyford GL, Seferiades A, Deacon RM, Cassaday HJ (1993) Critical determinants of nonspatial working memory deficits in rats with conventional lesions of the hippocampus or fornix. *Behav Neurosci* 107:420–433.
- Reed JM, Hamann SB, Stefanacci L, Squire LR (1997) When amnesic patients perform well on recognition memory tests. *Behav Neurosci* 111:1163–1170.
- Sidman M, Stoddard LT, Mohr JP (1968) Some additional quantitative observations of immediate memory in a patient with bilateral hippocampal lesions. *Neuropsychologia* 6:245–254.
- Squire LR (1992) Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol Rev* 99:195–231.
- Suzuki WA, Zola-Morgan S, Squire LR, Amaral DG (1993) Lesions of the perirhinal and parahippocampal cortices in the monkey produce long-lasting memory impairment in the visual and tactual modalities. *J Neurosci* 13:2430–2451.
- Tulving E 1972 *Organization of memory* (Tulving E, Donaldson W, eds), pp 381–403. San Diego: Academic.
- Warrington EK, Weiskrantz L (1974) The effect of prior learning on subsequent retention in amnesic patients. *Neuropsychologia* 12:419–428.
- Wiig KA, Bilkey DK (1994) The effects of perirhinal cortical lesions on spatial reference memory in the rat. *Behav Brain Res* 63:101–109.
- Wiig KA, Bilkey DK (1995) Lesions of rat perirhinal cortex exacerbate the memory deficit observed following damage to the fimbria-fornix. *Behav Neurosci* 109:620–630.
- Winters BD, Forwood SE, Cowell RA, Saksida LM, Bussey TJ (2004) Double dissociation between the effects of peri-postrhinal cortex and hippocampal lesions on tests of object recognition and spatial memory: heterogeneity of function within the temporal lobe. *J Neurosci* 24:5901–5908.
- Zola-Morgan S, Squire LR, Amaral DG (1989a) Lesions of the amygdala that spare adjacent cortical regions do not impair memory or exacerbate the impairment following lesions of the hippocampal formation. *J Neurosci* 9:1922–1936.
- Zola-Morgan S, Squire LR, Amaral DG, Suzuki WA (1989b) Lesions of perirhinal and parahippocampal cortex that spare the amygdala and hippocampal formation produce severe memory impairment. *J Neurosci* 9:4355–4370.