HIPPOCAMPUS 22:1952–1964 (2012)

COMMENTARY

Computational Models of Perirhinal Cortex Function

Rosemary A. Cowell*

ABSTRACT: I review seven models of the contribution of perirhinal cortex (PRC) or neighboring neocortical regions to cognition. Five of the models address recognition memory function (Sohal and Hasselmo (2000) Network 11:169-190; Bogacz et al. (2001) J Comput Neurosci 10:5-23; Bogacz and Brown (2003a) Neurocomputing 52:1-6; Norman and O'Reilly (2003) Psychol Rev 110:611-646; Cowell et al. (2006) J Neurosci 26:12186-12197) and two account for the role of PRC in visual discrimination learning (Bussey and Saksida (2002) Eur J Neurosci 15:355-364; Cowell et al. (2010b) J Cogn Neurosci 22:2460-2479). The models span a range of biological scales and target a variety of datasets, such that like for like comparison between them is not always possible. I lay out a novel framework for facilitating comparison by defining some general abstract principles concerning the organization of cognition in the brain about which all of the models make a statement. The controversies that are revealed by scrutinizing the models within this framework highlight the fundamental questions that remain to be answered by future research. Ultimately, it is by combining these disparate accounts to build a unified model that bridges several levels of biological scale and accounts for multiple psychological phenomena that a full account of PRC function will be achieved.

KEY WORDS: computational model; perirhinal cortex; familiarity; visual discrimination; recognition memory

INTRODUCTION

The experimental literature has demonstrated that perirhinal cortex (PRC) is critical for object recognition memory, a canonical test of declarative memory function (Zola-Morgan et al., 1989; Gaffan and Murray, 1992; Meunier et al., 1993; Mumby and Pinel, 1994; Aggleton et al., 1997; Buckley et al., 1997; Baxter and Murray, 2001; Malkova et al., 2001; Winters et al., 2004). There is also much evidence to support the notion that PRC is critical for object perception—as measured by tasks involving visual discrimination of objects—when the task involves sufficiently taxing demands on the perceptual representations of objects (Buckley and Gaffan, 1997; Buckley et al., 2001; Bussey et al., 2002, 2003; Barense et al., 2005, 2007, 2010).

There are a number of published computational models of PRC, some of which address its contribution to recognition memory (Sohal and Hasselmo, 2000; Bogacz et al., 2001; Bogacz and Brown, 2003a; Norman and O'Reilly, 2003; Cowell et al., 2006) and some of its role in visual perception (Bussey and Saksida, 2002; Cowell et al., 2010b). These various models are couched at different levels of biological organization (e.g., synapses vs. anatomically-defined brain regions), have different target data sets (e.g., single unit recordings vs. behavioral performance measures), and use different terminology, suited to their own specific purposes, to describe the mechanisms they propose. This can make it difficult to compare the models directly and to take a comprehensive, synoptic view of the field of models of PRC.

In this article I will review the field using, as far as possible, a set of general, theoretical concepts to describe all of the models. I restrict the discussion to computational models of the function of PRC or neighboring neocortical regions. Models addressing only hippocampal function are excluded, as are abstract cognitive models of recognition memory that are not brain-based (e.g., Shiffrin and Steyvers, 1997; McClelland and Chappell, 1998). In addition, I consider only process models, excluding models that strive for an accurate description of behavioral data without proposing an explicit neural or cognitive mechanism for the function that is described (e.g., Ratcliff et al., 1992; Yonelinas, 1994).

Emphasis will be placed upon the way in which each model specifies how PRC differs from other brain regions in its contribution to cognition—that is, how each model slots PRC function into the global model of cortical processing. The article is intended to provide a guide to the field for nonmodelers, a summary of the state of our understanding of PRC function in terms of formal process models and an overview of the consensus and controversy therein. In reviewing the models, it can be seen that their various accounts address different datasets and therefore sometimes address different aspects, or different levels, of PRC function. When this is true, accepting the validity of several models simultaneously might save us from similitude with the allegorical blind men, who each examined only one part of an elephant and could not agree about the nature of the beast, declaring it most similar to a rope (its tail), a tree (its leg), a wall (its side), a fan (its ear), or a spear (its tusk), according to where they had touched it (Fig. 1). Like the

Department of Psychology, University of California, San Diego, California
*Correspondence to: Rosemary A. Cowell, Department of Psychology, University of California San Diego, 9500 Gilman Drive MC0109, La Jolla, CA 92109, USA. E-mail: rcowell@ucsd.edu
Received 1 June 2012Accepted for publication 26 July 2012
DOI 10.1002/hipo.22064
Published online in Wiley Online Library (wileyonlinelibrary.com).

FIGURE 1. In the fable of the blind men and the elephant, the observers could not agree upon the most accurate description of the creature because they each groped a different part of its anatomy. In cognitive neuroscience, we often face the same dilemma. By examining perirhinal cortex from perspectives that differ in anatomical scale, in the neurobiological properties focused upon, or in the aspect of cognitive function that we desire to explain, we may arrive at quite disparate models. Ultimately, it is through integrating across these models, rather than choosing between them, that we stand to make the most progress in understanding perirhinal function. Courtesy of Jarret Frank at Boston University, http://www.jafradesign.com/. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

elephant, a brain region such as PRC has many observable characteristics, and different models may each explain a different dataset without precluding the explanations of the others. On the other hand, when a valid debate can be established between well-matched yet conflicting accounts of PRC function, the models are even more useful. The controversy highlighted by competing theoretical accounts is vital to guiding the field forward via focused and hypothesis-driven experimental work.

GENERAL THEORETICAL CONCEPTS

In this section, I propose a set of general theoretical concepts or principles that I will subsequently use to characterize the various computational models of PRC. These concepts are intended to provide labels for key characteristics or theoretical claims of the models that are relatively neutral with respect to any particular viewpoint. I eschew labels that are applicable only to a particular model, because they preclude comparison of that aspect of the model with those of other models. I try to avoid labels that invoke a particular psychological process, such as "familiarity," first, because that process might mean different things, computationally, in different models, and second, because the use of psychological labels counteracts a key benefit of computational models—that couching a theory in concrete, mathematical terms avoids the need to use verbal descriptions of cognitive operations that are psychologically loaded and open to misinterpretation. Moreover, recasting some of these psychological terms within a theoretically neutral framework may render their definitions more useful. These neutral theoretical concepts are used to provide a framework within which to compare models and promote better understanding of the proposed mechanisms.

Biological Scale and Problem Space

The biological scale and problem space of a model are central to our ability to define and evaluate the scientific contribution of that model (Churchland and Sejnowski, 1988; Cowell et al., 2011). The biological scale of a model refers to the level of organization at which it is formulated, i.e., the size of the biological components that feature in its explanation. The biological scale of the models of PRC in this review ranges from synapses and individual neurons to anatomically-defined brain regions such as perirhinal cortex and hippocampus. In reviewing a set of models couched at different biological scales it will become clear that, in some cases, direct comparison of those models cannot be made.

Problem space is closely related to the issue of biological scale. The problem space of a model is defined by the set of target data that the model attempts to explain. A given dataset may inform a model's proposed mechanism at one or two levels of biological organization. For example, electrophysiological data from single unit recordings may speak to both the level of individual neurons and the level of neural assemblies; behavioral data from patients with brain lesions may reveal properties of individual anatomical systems (such as hippocampus or PRC) as well as the way in which different anatomical systems interact at the level of the whole organism. Moreover, some models may attempt to account for more than one type of data, for example, single unit recordings and behavioral data, simultaneously. A second dimension of problem space is the number of phenomena at a given level of biological organization that the model attempts to address. For example, a model might attempt to account for several psychological phenomena at the level of the whole organism, such as recognition memory, categorization and perceptual learning, or only one of these. All else (parsimony, fit to the target data, concordance with known biology, etc.) being equal, a model that can explain more phenomena is to be favored over a model that explains fewer phenomena. However, constraining the problem space is usually necessary to settle upon a tractable problem. In addition, keeping the explanation limited in scope often renders it simpler, clearer and more elegant.

Localization and Specialization

Localization of function to a region

I will use the concept localization of function to describe one characteristic of the sometimes underspecified notion of an anatomical module. A model shall be described as claiming localization of function if it is consistent with the idea that a given cognitive function (such as judgment of prior occurrence) occurs in only one region (such as PRC). The contrasting posi-

tion that a model may take is that cognitive functions are not localized, but rather that a cognitive goal such as recognition memory can be carried out in many regions, such as PRC, anterior IT, or hippocampus. This definition is neutral as to whether, in the latter case, a cognitive goal is carried out using the identical neural mechanism across different regions.

Specialization of a region for one function

Another theoretical claim implied by the term "anatomical module," and complementary to the idea of localized function, is the notion that a given brain region (such as PRC) performs only one function (such as judgment of prior occurrence). I will use the term specialization for function to describe this concept. The contrasting position is taken by theories in which a given region, such as PRC, may carry out more than one function, such as judgment of prior occurrence, recency detection, and perceptual discrimination.

In both terms, the word "function" refers to a psychologically-defined cognitive goal such as recognition memory or visual discrimination, rather than to a computational operation such as pattern separation that might be used to carry out that goal. These two concepts are intended to replace the somewhat loaded and less specific notion of a module—either anatomical or cognitive—in the discussion of computational models. It should be noted that these concepts can also be applied at a smaller scale: one can make claims about localization and specialization of neural networks within an anatomically-defined brain region, or even at the level of individual neurons.

Cognitive Algorithm Versus Representational Content

I define here two contrasting approaches to explaining functional differences between brain regions. Under the first approach, many models of brain function claim that observable differences in the contributions to cognition made by two different brain regions arise from differences in the type of cognitive operation those regions perform. For example, a model might claim that the primary goal of one region (hippocampus) is to perform episodic recall of stimulus items, while the primary goal of a second region (PRC) is to provide information upon which a judgment of prior occurrence can be based (Brown and Aggleton, 2001; Norman and O'Reilly, 2003; Ranganath et al.; 2004, Eichenbaum et al., 2007). Such a claim entails the assumption that "episodic recall" and "judgment of prior occurrence" involve fundamentally different processes, not just at the level of neurons, but also at the abstract level of cognition. Fundamental differences in the proposed cognitive operations can be specified and examined in terms of their predictions for cognitive level data, such as receiver operating characteristic (ROC) curves (e.g., Elfman et al., 2008).

Other models using the second, contrasting approach seek to explain the different contributions of two or more brain regions by citing differences in the content of those regions' stimulus representations rather than by appealing to distinct cognitive operations (e.g., Bussey and Saksida, 2002, 2005; Cowell et al., 2006, 2010a,b; Graham et al., 2010; Ranganath, 2010). Models adhering to this approach often make claims about the type of stimulus representation housed in each brain region, and appeal to the representational demands of a cognitive task when attempting to explain evidence that a brain region does or does not contribute to performance on that task. Such theories do not necessarily claim that there are no architectural or neural processing differences between different brain regions such as hippocampus and perirhinal cortex. Instead, this approach is consistent with the idea that where architectural and neural processing differences exist, they reflect the differing information processing requirements imposed by performing a common cognitive operation (say, judgment of prior occurrence) on different types of stimulus (say, a rich episodic memory trace involving time, space, context, objects, and people in the hippocampus, vs. memory traces for individual objects in perirhinal cortex). That is, rather than assuming that different brain regions carry out distinct, introspectively-defined psychological goals using distinct cognitive processes, this latter approach prefers to explain functional dissociations in terms of the type of material upon which any cognitive process being carried out is obliged to operate.

This pair of concepts is related to localization and specialization of function. If a model supposes that two different brain regions employ distinct cognitive operations, then it most likely claims specialization of those regions for a function, and often claims localization of a function to each of those regions. Models that instead emphasize the representational content of different brain regions do not endorse these claims. Rather, to the extent that either specialization or localization is endorsed, it would be specialization for processing a type of stimulus representation, or localization of the processing of certain types of stimulus material to particular brain regions.

Differentiation Versus Assimilation

Many models make claims regarding the ability or tendency of a neural system to either discriminate between or categorize as similar a set of input stimuli impinging upon the system. The former operation is often described as *pattern separation* or *discrimination* and the latter as *generalization*. In this article, I restate these two concepts. I will use the word differentiation to describe the process of telling inputs apart, which is less tied to a specific computational means of discriminating between inputs than pattern separation and less strongly associated with a particular cognitive task than discrimination. (This use of the term should not be confused with its use in the animal conditioning literature (Gibson, 1940) or in cognitive models of memory (Shiffrin et al., 1990; McClelland and Chappell, 1998)). I will use the term assimilation to describe the process of grouping similar input stimuli through the use of similar representations, because I wish to discuss primarily the process of forming the group in the first instance, rather than the process of making similar cognitive judgments about all items in a

group (i.e., the step of generalizing a judgment from one item to another, once they have been grouped).

COMPUTATIONAL MODELS OF PERIRHINAL CORTEX FUNCTION

I will review five models that address the role of PRC in recognition memory or familiarity judgments—Sohal and Hasselmo (2000), Bogacz et al. (2001), Bogacz and Brown (2003a), Norman and O'Reilly (2003), and Cowell et al. (2006)—and two models that address the functional contribution of PRC to perception—Bussey and Saksida (2002) and Cowell et al. (2010b). Table 1 details the stance taken by each of the models on the general theoretical concepts introduced above, and provides an at-a-glance comparison of all of the models considered in this article.

Bogacz et al. (2001): ''Hebbian Model''

The model of Bogacz et al. (2001) provides an account of how familiarity discrimination might be performed by the brain and why that process might be separable from the process of recollection. The target data of the model are electrophysiologically observed changes in the responses of PRC neurons to stimulus items as they become familiar. The network is designed to model familiarity discrimination and not other functions of PRC. The level of biological organization of the model is individual neurons and local ensembles of neurons.

The network simulates learning by so-called "novelty neurons" in PRC, which fire less strongly in response to familiar stimuli than to novel stimuli. This pattern of firing is acquired in the model via Hebbian learning: the connections from active inputs to novelty neurons are increased in a process akin to long-term potentiation (LTP) (Bliss and Lomo, 1973), such that a model novelty neuron is more likely to fire in the period immediately following stimulus onset for a stimulus that has been seen before than for a novel stimulus. However, all novelty neurons in the model receive inhibitory input that suppresses their responses. Because inhibition to each novelty neuron is driven by the initial responses of the novelty neurons themselves, the responses of novelty neurons are suppressed for familiar patterns, but not for novel patterns. This produces greater firing in response to novel than to familiar stimuli (i.e., a reversal of the response pattern) in the period following the initial responses. It is proposed that in experimentally observed novelty neurons, activity in the initial, brief poststimulus phase is masked by activity in the subsequent, longer inhibitory phase, so that the observed firing of novelty neurons in PRC resembles the responses of novelty neurons in the model. One advantage of this mechanism is that it allows for a very large recognition memory capacity, which the authors claim is of the same order of magnitude as the capacity of human recognition memory (Standing, 1973).

One key claim made by the model is that there are separate networks within PRC for familiarity judgments and for the learning of stimulus representations. This claim is founded upon the notion that networks that attempt to perform both familiarity discrimination and feature extraction—that is, the learning of stimulus representations through exposure to visual stimuli, during which the features possessed by the stimuli are extracted by exploiting statistical regularities in the input—do not have a sufficiently high capacity for storing familiar items to account for the impressive recognition memory of humans (Bogacz et al., 2001; Bogacz and Brown, 2003b). Such "dual" networks (i.e., networks that perform both functions) include those of Sohal and Hasselmo (2000), Norman and O'Reilly (2003), and Cowell et al. (2006). In contrast, by separating these functions, the model of Bogacz et al. (2001) claims localization of function at the level of neural networks: it is implied that familiarity discrimination occurs only in PRC and only in the "familiarity discrimination" network. In addition, since the authors claim that PRC possesses at least two networks with distinct computational goals—one for familiarity discrimination and another for learning representations—the model also implies specialization for function at the level of neural networks.

In Bogacz et al. (2001), the familiarity discrimination network deliberately emphasizes the unique characteristics of input stimuli (i.e., it differentiates them), relative to a putative feature extraction network in PRC (not simulated in Bogacz et al., 2001), which instead would preserve the similarities of the inputs in its learned representations. This network would presumably support goals such as categorization and semantic knowledge, for which a greater tendency of the network to emphasize similarities would be useful.

Bogacz and Brown (2003a): ''Anti-Hebbian Model''

Bogacz and Brown (2002, 2003a) introduce a new network for simulating familiarity discrimination in PRC. This network is similar to Bogacz et al. (2001), except that it uses an alternative learning rule to produce the novelty neurons' behavior: anti-Hebbian learning instead of Hebbian. In this model, connections from active inputs to novelty neurons are decreased as if by Long-Term Depression (LTD; Dudek and Bear, 1992) so that, without the need for inhibition, novelty neurons come to respond less strongly to a stimulus when it has previously been seen by the network. The use of the anti-Hebbian learning rule has an important consequence: for correlated input patterns this model gives a very large storage capacity, greater not only than in networks that combine familiarity discrimination with feature extraction (e.g., Sohal and Hasselmo, 2000; Norman and O'Reilly, 2003; Cowell et al., 2006), but also than in the Hebbian model of Bogacz et al. (2001) (Bogacz and Brown, 2003b).

Sohal and Hasselmo (2000)

The computational model put forward by Sohal and Hasselmo (2000) is situated in inferior temporal (IT) cortex, which

TABLE 1.

Summary of Key Properties of the Seven Reviewed Computational Models

	Target data	Biological scale	Localization of function	Specialization for function	Cognitive algorithm versus representational content	Differentiation versus assimilation
Bogacz et al. (2001); Bogacz and Brown (2003a)	Response changes in PRC neurons	Individual neurons; neural networks within PRC	Yes, to networks within PRC	Yes, by networks within PRC	Cognitive algorithm	— (Not a key property)
Sohal and Hasselmo (2000)	Response changes in IT neurons	Individual neurons; neural networks spanning IT and basal forebrain	Not stated	No, same IT neurons underlie both recency and familiarity	— (Neither applies)	Differentiation in presence of ACh, assimilation in absence of ACh
Norman and O'Reilly (2003)	Recognition memory performance measures following brain lesions	Neural networks; anatomical systems	Yes: recall to hippocampus; familiarity judgments to MTL neocortex	No	Cognitive Algorithm	Differentiation in hippocampus; (mostly) assimilation in MTL neocortex
Cowell et al. (2006)	Recognition memory performance measures following brain lesions	Anatomical systems	No	No	Representational Content	For *objects*: differentiation in PRC; assimilation in caudal regions
Bussey and Saksida (2002); Cowell et al. (2010b)	Visual discrimination learning measures following brain lesions	Anatomical Systems	No	No	Representational Content	For *objects*: differentiation in PRC; assimilation in caudal regions. For *simple stimuli*: the reverse

All models address the function of anterior/medial temporal lobe neocortex. The position taken by each model on each of six key properties is described under the column heading corresponding to that property. The two models of Bogacz, Brown and colleagues do not differ on these properties and so are described within the same row; similarly the models of Bussey & Saksida (2002) and Cowell et al. (2010b) are described jointly. In the final row, 'simple stimuli' implies the stimuli possess few features (e.g., simple, black and white line drawings of basic geometric shapes), whereas 'complex' refers to stimuli that possess many features varying in size, shape and color, all combined into a conjunction.

lies laterally adjacent to PRC and provides one of its primary sources of input (Suzuki and Amaral, 1994). I discuss this model because it addresses certain response properties of IT neurons that are also possessed by PRC neurons, relating to the ability of cortex to provide information about the recency and familiarity of visual stimuli, a function for which PRC is known to be critical. The biological scale of the model is synaptic and neuronal: the model uses biologically plausible synaptic modification rules to simulate the responses of individual neurons in IT and the basal forebrain. The target data of the model are short- and long-term decreases in the responses of IT neurons and basal forebrain cholinergic neurons to repeatedly presented visual stimuli.

Stimuli are applied to an input region, from where activation spreads via feedforward, excitatory connections to IT cortex. In IT cortex, there are recurrent excitatory connections as well as recurrent inhibitory feedback from an inhibitory interneuron. These two features allow IT activity to be maintained in the absence of a stimulus. IT neurons also provide inhibitory input to cholinergic neurons within the basal forebrain, which in turn provide cholinergic modulation of activity in IT. The feedforward connections from input to IT are learned in a competitive, self-organizing manner according to a Hebbian rule.

Short-term response decreases (recency effects) in the model are simulated by a habituation, or adaptation, mechanism: activation of an IT neuron causes a calcium influx, and subsequent responses in that neuron are suppressed due to the activation of a calcium concentration-dependent potassium current. Long-term response decreases (signaling familiarity) in IT arise because the competitive learning changes the distribution of IT neurons activated by a stimulus. As many patterns compete to be represented in IT cortex, there are two ways in which neural responses may decrease in the model. First, some neurons ("negative neurons") either do not respond strongly enough for Hebbian learning to strengthen their connections to IT and they do not become included in the representation of any pattern, or they respond to too many stimuli and do not develop stable connections from a single input pattern to IT cortex. In these neurons, the average response to all sample and match stimuli decreases as those stimuli become familiar to the network; such neurons have been observed experimentally (Miller et al., 1991; Li et al., 1993). Second, the model predicts the existence of neurons that respond strongly to only a small subset of familiar stimuli and weakly to all others. These arise whenever a neuron is initially activated by several input patterns but is subsequently recruited by the representation of one particular pattern, causing it to 'drop out' of the IT representation for all other patterns.

Cholinergic modulation plays a critical role in the modeling of long-term response decreases. Recurrent excitation in IT cortex is implemented in order to simulate the maintenance of activity in the absence of a stimulus, which is noted to be a property exhibited by IT neurons, but which is not explicitly involved in the simulation of short-term memory effects in this article. However, recurrent excitation also contributes to the mechanism for familiarity discrimination, in a negative way: it interferes with the self-organization process such that, rather than learning representations that differentiate similar input stimuli, the network learns representations that assimilate similar input stimuli. The assimilation can be so extreme as to render the patterns elicited by two different stimuli completely overlapping. Some means of reducing the transmission at recurrent synapses is required to negate this nuisance effect, and cholinergic modulation from basal forebrain input to IT cortex is proposed. When ACh modulation is applied to IT cortex, the recurrent excitatory connections are selectively suppressed, which allows the self-organization of feedforward connections to IT cortex to proceed unhindered. When ACh modulation is included, IT representations of similar inputs are well differentiated and long-term response decreases are successfully simulated.

The model makes no clear statement on the issue of localization of function: while IT cortex is proposed to be a site of recency and familiarity computations, the authors do not explicitly claim that these processes do not occur elsewhere in the brain. However, the notion of specialization for function is not endorsed: the same IT neurons provide both recency and familiarity signals and therefore do not specialize in computing only one function.

This network proposes different neural algorithms for different cognitive goals (recency and familiarity judgments) within a single region (IT cortex), suggesting that more than one cognitive or neural algorithm can operate on the representations in a region. This is consistent with models advocating representational content (rather than cognitive algorithm) as the factor that determines a region's contribution to cognition. However, since the model does not claim that representational content is the best way to describe the role of IT in visual cognition, this particular general principle about the organization of cognition (cognitive algorithm vs. representational content) is not directly addressed.

Norman and O'Reilly (2003)

The model of recognition memory put forward by Norman and O'Reilly (2003) was inspired by the Complementary Learning Systems (CLS) approach (McClelland et al., 1995; O'Reilly and Rudy, 2001). The CLS view proposes that the hippocampus is specialized for rapidly memorizing specific events and the neocortex for slowly learning about statistical regularities of the environment. The biological scale of the model is at the level of anatomical brain systems (such as MTL neocortex and hippocampus) rather than individual neurons. However, the model is also constrained by biological data from lower levels, employing biologically plausible learning algorithms. The target data of the model are behavioral measures of recognition memory, particularly from individuals with brain lesions.

The model has two components: hippocampus and MTL neocortex (MTLC). The neocortical network extracts regularities in the statistics of input patterns by employing Hebbian learning and collateral inhibition to produce competitive learn-

ing of representations in the MTLC hidden layer. Because hidden layer units compete to represent each input pattern, a familiar stimulus is represented by a small number of strongly active MTL neocortex units (those that have won the competition for that pattern, after viewing it multiple times), whereas a novel stimulus is represented by a larger number of weakly active units (none of which have yet been forced to "drop out" through competition). The sharpness of an MTL neocortical representation is used to index the familiarity of a stimulus.

In the hippocampal component of the model, input patterns are encoded sparsely (i.e., using very few active units) on a layer corresponding to the CA3 region of the hippocampus. The encoding process gives rise to distinct CA3 representations, which share very little overlap with each other even for highly overlapping input patterns. This allows the network to differentiate highly similar stimuli. In addition, collateral excitatory connections allow the spreading of activation from one part of the sparse CA3 representation to the remaining parts. This process enables the retrieval (recollection) of the missing parts of a representation, when partial or degraded inputs are provided to the model. Recall is measured by sending activation from CA3 back to an output layer; the recall score is high when there is a high degree of match between the activity patterns on the input and output layers.

In the model, the hippocampal component differentiates between input stimuli as memory traces are laid down. This enables the simultaneous retention of many items and episodes without interference between them, even when those separate instances share many features, as is often the case in everyday life. In contrast, representations in the MTL neocortex component tend to assimilate input stimuli relative to the representations in hippocampus, because they retain some of the similarity information present at the inputs. This allows the MTL model to generalize between similar stimuli and to learn about statistical regularities in the environment, two capabilities that are presumably most useful for tasks involving categorization or requiring semantic knowledge, on which the study of Norman and O'Reilly (2003) does not focus. In this study, MTL neocortex representations simply supplement the recall-based performance of recognition memory in hippocampus by providing information that can be used to make familiarity judgments. The authors note that MTL representations do to some extent differentiate items, in that they possess less overlap than the input patterns; however, this property is deemphasized in this study because MTL representations overlap so much more than in the hippocampus, with which MTL is compared.

This model makes a clear statement regarding localization of function: episodic recall is localized to hippocampus, whereas familiarity discrimination is localized to MTL neocortex. However, the model does not make a strong claim of specialization for function: although the target data addressed in this article concern recognition memory, the model does not claim that hippocampus and MTLC contribute only to recognition memory. This leaves open the possibility that either or both regions may perform other functions. Indeed, the authors describe the suitability of the MTL neocortical learning algorithm for acquiring statistical regularities of the environment, knowledge which would be useful for many cognitive goals beside recognition memory. However, the model certainly claims that distinct cortical regions employ different cognitive (and neural) algorithms for carrying out the cognitive goals they subserve.

The Representational-Hierarchical View

The remaining three models are separate instances of a common theoretical framework that Lisa Saksida, Tim Bussey, and I have termed the "representational-hierarchical view" (Cowell et al., 2010a; Saksida and Bussey, 2010). While these models' target data and some details of their mechanisms differ, they share many common assumptions and invoke the same key properties of PRC in accounting for cognition.

Cowell et al. (2006)

Cowell, Bussey, and Saksida (2006) proposed a model of object recognition memory in PRC. Its target data are behavioral data, including the finding from recognition memory tasks that forgetting over a delay is exacerbated in individuals with PRC lesions, relative to controls. The model resides at the biological scale of anatomical systems; each layer of units in the network corresponds to an anatomically defined region such as PRC or a subregion of the ventral visual stream (VVS).

Similar to Norman and O'Reilly (2003), representations within the network become sharpened with visual experience and the sharpness of an object's representation indicates its level of familiarity. Objects are represented in the model on two layers: the PRC layer and the "caudal" layer, corresponding to a region of ventral visual stream caudal to PRC. The model's mechanism is based on two key assumptions: first, that PRC stores representations of the conjunctions of visual features possessed by complex objects, which play an important role whenever it is difficult to solve a task using representations of only individual visual features; and, second, that all objects in the visual world are composed of simple features (such as lines, colors, and blobs) of which there exist only a limited number that are unique. As a consequence of the second assumption, the same visual features occur repeatedly in different objects, even though many thousands of unique objects may be constructed through different combinations of those features.

On the PRC layer, all eight input dimensions possessed by an individual object are combined into a unique conjunction. On the caudal layer, there are four separate representational grids, each of which combines two input dimensions into a simple (two-dimensional) visual feature; together, the four grids represent the entire object. This architecture produces a representational space that is vast in anterior regions (the PRC layer) and small in posterior regions (the caudal layer): the number of unique stimulus representations that can be specified in posterior cortical regions is far fewer than the number in PRC. According to representational-hierarchical view, this is the key property of the representational content that confers a brain region's function. The change in the size of the representational space with progression from posterior to anterior regions

(caused by the shift from low-dimensional to high-dimensional representations) is central to the account of the effects of PRC lesions on recognition memory provided by Cowell et al. (2006). Forgetting over a delay in an object recognition memory task is simulated via the presentation of interfering objects between study and test. When interfering objects are presented, the same features appear repeatedly as part of many different interfering objects and become sharpened. Because the representational space spanned by each caudal grid is small—that is, there are only a few possible unique features that may be represented on each—these features "interfere" with the caudal layer representations. Specifically, because the novel object that is later presented on test is also constructed from the same small set of features as the interfering objects, its feature representations are sharpened by interference and it is falsely recognized as familiar. In contrast, input stimuli are represented in PRC as a unique conjunction of all input dimensions. For an interfering object to share a representation with the novel object and interfere with its representation on the PRC layer, it would need to be identical to the novel object. The chance of such an interfering object appearing repeatedly when interfering objects are chosen randomly from eight-dimensional representational space is vanishingly small. Therefore interference does not disrupt the PRC representation of the novel object and it is never falsely recognized.

The model claims that cognitive functions such as recognition memory are not localized to any brain region, because the same algorithms operate at both the PRC and caudal layers. PRC is critical for recognition memory, but only in the case of object-level stimuli (not for simple visual features, nor complex episodes), and only because it houses the most useful representations of objects for discriminating them on the basis of familiarity. Furthermore, the model predicts that recognition memory for simple features would be well executed by caudal regions of VVS, and the general theoretical framework implies that recognition of scenes and complex episodes (in which objects are associated with time, space, context, and events) should be underpinned by the hippocampus, one level higher up the anatomical hierarchy. That is, other brain regions can perform recognition memory when the familiarity discrimination required by the recognition memory task is best performed using the representations in those regions.

Similarly, the model denies any specialization for function by suggesting that any brain region may contribute to any cognitive function when the representations housed in the region are useful for that function. The claim is that because PRC representations are useful both for familiarity judgments and for discriminating between objects that share features, PRC contributes to both recognition memory and higher-order perception.

Although Norman and O'Reilly argued for differentiation in the hippocampus and assimilation in MTL neocortex, Cowell et al. (2006) highlight the role of PRC (part of MTL neocortex) in differentiation: in the model, PRC is critical for differentiating in terms of familiarity between objects that share features, because earlier regions in VVS cannot. However, the model does not claim that PRC is specialized for differentiation and the caudal layer specialized for assimilation. Rather, the tendency of each region to differentiate or assimilate depends on the complexity of the stimulus material used in a given task. For object-level stimuli with shared features, PRC differentiates and caudal regions assimilate. (This is because PRC representations combine simple visual features into unique conjunctions of the optimal level of complexity—i.e. the optimal number of "parts"—for differentiating individual objects, whereas caudal regions represent the features individually without representing the conjunction of those features that defines an object.) For complex episodic stimuli or scenes that share object-level items, PRC would be predicted to assimilate, and hippocampus would be required for differentiation. (This is because the highest level of conjunction represented by PRC is the conjoining of the features within an object into a whole; in order to conjoin multiple objects with other complex episodic attributes such as time, place and context, a higher level in the representational hierarchy—namely, hippocampus—is required).

Bussey and Saksida (2002)

Bussey and Saksida (2002) proposed a model of visual discrimination learning in PRC. Its target data are behavioral performance measures from monkeys and rats with lesions in PRC across a range of visual discrimination learning tasks. The model is again formulated at an anatomical systems level.

Visual discrimination learning tasks involve repeatedly presenting an animal with one or more pairs of visual stimuli, where, in each pair, one stimulus is consistently followed by food reward and the other is not. Over successive presentations, animals learn to choose the rewarded stimulus of each pair reliably. The literature indicated that PRC was important for visual discrimination learning with object stimuli under certain conditions, such as using a large set of stimulus pairs in the task (Buckley and Gaffan, 1997), or presenting the stimuli from different views from one trial to the next (Buckley and Gaffan, 1998a; Buckley et al., 2001). This pointed to a role for PRC in object identification (Buckley and Gaffan, 1998b), but it was unclear through what mechanism the PRC fulfilled this role (but see Gaffan et al., 1986 for a model of visual discrimination learning).

The model of Bussey and Saksida (2002) comprises a connectionist network with two representational layers: caudal VVS and PRC. Representations in this network are static—a node in the caudal layer is designated at the outset as corresponding to a simple visual feature and a node in the PRC layer to a particular conjunction of features (an object). The only learning that occurs is associative learning between these static stimulus representations and an outcome node which signals reward. The network simulates visual discrimination learning by associating the correct stimulus of each pair with reward. Intact networks do so using both caudal layer and PRC layer representations. Lesioned networks are subject to removal of the PRC layer and must rely only on caudal representations.

As in the account of recognition memory offered by Cowell et al. (2006), the effect of PRC lesions on visual discrimination

learning is explained in terms of the difference in representational dimensionality between feature-level caudal representations and object-level PRC representations. In all experimental tasks in which PRC lesions caused impairments, the set of to-be-discriminated stimuli were chosen in such a way that feature-level representations were inadequate for solving the task. For example, using a large set of stimuli ensured the repeated appearance of many simple visual features, which were sometimes rewarded and sometimes not, such that reliable feature-reward associations could not be formed; similarly, presenting stimuli from different views from one trial to the next meant that different features of the objects appeared each time, rendering the acquisition of feature-reward associations very difficult. Therefore, according to the model, in every case where impairments were revealed, the higher-dimensional, conjunctive representations in PRC were required in order to form reliable associations of the stimuli with reward.

The model of Bussey and Saksida (2002) accounts for visual discrimination learning data by exploiting the change in the dimensionality of stimulus representations with progression from posterior to anterior regions. Because the high-dimensional PRC layer represents the conjunctions of objects uniquely, it differentiates objects with shared features. In contrast, because the caudal layer represents all features individually, it assimilates any objects that share features, because they possess overlapping activation patterns. Like the model of Cowell et al. (2006), the network of Bussey and Saksida (2002) argues against localized function and specialization for function: the only difference between the layers is the level of stimulus complexity optimally represented, and each layer may contribute to any cognitive task for which it provides useful stimulus representations.

Cowell et al. (2010b)

Cowell et al. (2010b) present a network that extends the visual discrimination learning model of Bussey and Saksida (2002). It possesses an additional layer of stimulus representations so that there are three layers spanning regions from occipitotemporal cortex to anterior temporal lobe. Lesions of both anterior and posterior layers are therefore possible, which is required for the simulation of an important body of literature from the 1960s and 1970s, which repeatedly found a double dissociation between visual discrimination impairments caused by lesions in anterior versus posterior regions of the VVS (Iwai and Mishkin, 1968; Cowey and Gross, 1970; Gross et al., 1971; Wilson et al., 1972; Blake et al., 1977). Typically, anterior lesions caused impairments when discriminations between many complex stimulus pairs had to be learned concurrently (implying a role for these regions in memory), whereas posterior lesions caused impairments when the discrimination of only a few very simple stimulus pairs had to be acquired (implying these regions were critical for visual perception). This dissociation was interpreted as evidence for anatomical modules, one anterior and one posterior, for the functions of associative memory and visual perception, respectively.

The model assumes the same hierarchy of stimulus representations, again possessing a very large representational space in anterior temporal lobe; this is critical for good performance on visual discrimination of objects with shared features (the "memory" tasks of the 1960s and 1970s), because, as in Bussey and Saksida (2002), it protects object representations from feature-level interference. However, Cowell et al. (2010b) also propose that the huge representational space in anterior temporal lobe can be disadvantageous for discriminating simple stimuli: for any simple stimulus comprising a conjunction of only a few features (like those used in the "perceptual" tasks of the 1960s and 1970s), many of the highly complex conjunctions represented by units in the anterior layer contain that simple conjunction (as a subpart of the complex conjunction) and so their units are weakly activated. Representations of simple stimuli in the anterior layer are thus smeared across a large area of representational space (many units weakly active but no units highly active). Upon presentation of two simple stimuli which share a feature, the representations in the anterior layer weakly activate many of the same units and are therefore poorly discriminated on this layer. In contrast, these stimuli are well discriminated by the most posterior layer of the network, which represents simple conjunctions of very few features, and thus possesses unique representations for simple stimuli.

The model therefore makes an additional statement regarding the issue of differentiation versus assimilation: not only does anterior temporal lobe differentiate objects while caudal regions assimilate objects, caudal regions tend to differentiate simple stimuli while anterior temporal regions assimilate simple stimuli. That is, each station in VVS performs differentiation for stimuli at its preferred level of complexity and assimilation for stimuli at nonpreferred complexities, and its contribution to a visual discrimination task depends upon the complexity of the stimulus material presented. This model therefore reinforces the claim that there is no localization or specialization for functions such as visual perception and memory in VVS.

COMPARING THE MODELS

One might attempt to compare two models either in terms of their specific, detailed mechanisms, or in terms of their claims regarding the general theoretical principles introduced in General Theoretical Concepts (see also Table 1). Comparing models in terms of their specific, detailed mechanisms is only sensible if those specifics are sufficiently well-matched. We can compare the proposed mechanisms of Sohal and Hasselmo (2000) and Bogacz et al. (2001) and evaluate their relative success in accounting for experimental data, because their target data are similar and the levels of their explanations equivalent. However, comparing the specifics of Bogacz et al.'s mechanism (which simulates decreases in neural firing as stimuli are repeated) with the specifics of Bussey and Saksida's mechanism (which simulates visual discrimination learning behavior) would

be less sensible. In this scenario, we would be guilty of the same mistake as the allegorical blind men who examined the elephant. Although both models address some aspect of PRC function, they each examine a different part—both in terms of the level of explanation and the cognitive function being explained—and to accept the details of one mechanism does not necessarily require us to reject the details of the other. Nor should the explanation of one phenomenon be required to include details that are pertinent only to the explanation of the other, as long as both models have stand-alone explanatory power for the property of PRC function they address.

Happily, comparison of models along the lines of a general theoretical principle is nearly always valid and often useful: evaluation of the entire set of models might help to resolve ongoing debate by suggesting that one view of a particular principle (say, localized vs. distributed function) is overall more tenable than the other. The only case where we must be wary of arguing too forcefully about a general principle of PRC function is where that principle describes a continuous spectrum, rather than a qualitative distinction. This applies to the tendency of PRC to either differentiate or assimilate, and is discussed below. Here, as with comparing mechanisms at different levels, it is possible for both sides of the argument to be correct.

Guided by these considerations, the remainder of this section compares the models of PRC function, first in terms of two specific mechanistic properties and then in terms of the general principles outlined in General Theoretical Concepts. Where there remains controversy over a particular property of PRC function, the models are at their most useful, since they highlight a problem for future research. Moreover, where models with different mechanisms make explicit, conflicting predictions, these predictions can guide focused and fruitful experimental work. In turn, the data generated will enable refinement of current models or the building of more successful theories to replace them.

First, the models of recognition memory differ over whether the same PRC neurons are involved in perceptual discrimination and familiarity discrimination. The models of Sohal and Hasselmo (2000), Norman and O'Reilly (2003), and Cowell et al. (2006) are all consistent with the notion that familiarity discrimination and perceptual discrimination (i.e., item identification) are performed by the same set of neurons within PRC. Although the network of Cowell et al. (2006) is not intended to simulate responses of individual neurons, a single set of neurons for both functions is implied by the model, because a common learning mechanism (which mimics cortical processes) is used for both the acquisition of stimulus representations in general and the acquisition of familiarity information. Bogacz and Brown (2003b) argue that the Hebbian and Anti-Hebbian networks of Bogacz et al. (2001) and Bogacz and Brown (2003a) do not perform both familiarity discrimination and perceptual discrimination; indeed a key feature of these two networks is that they are specialized for familiarity discrimination and not for the learning of stimulus representations. Even though novelty neurons simulated by both networks (like experimentally observed novelty neurons in the brain) are stimulus selective, the authors emphasize that the stimulus selectivity in their proposed networks "is required solely to increase the efficiency of the network, and not because the implied representation of the visual stimuli is used for some further processing." The specialized networks of Bogacz and colleagues perform well at discriminating between two stimuli that are not equally familiar, but since they do not learn complete representations of the stimuli, they would presumably perform poorly at discriminating two different stimuli if both were equally familiar. The implication is, instead, that perceptual discrimination is achieved by a separate network that learns about the locations of stimuli in stimulus space, regardless of their familiarity.

The view of Bogacz et al. on this point is more in line than the other computational models with the standard medial temporal lobe memory system view (Squire and Zola-Morgan, 1991; Squire et al., 2004), in which mnemonic processing in MTL is assumed to be carried out separately from perceptual processing. One distinction, however, is that Bogacz et al. do not require that the network underlying perceptual processing is located outside of MTL cortex, allowing instead that a separate network within PRC may exist to serve this function.

Second, the models of recognition memory point to an open question over whether trace decay is required to account for some aspects of forgetting (Jenkins and Dallenbach, 1924; McGeoch, 1932). Decades of psychological research point to a role for interference in forgetting (Wixted, 2004), and none of the models reviewed contradicts this notion. The networks of Cowell et al. (2006) and Sohal and Hasselmo (2000) make explicit use of an interference mechanism without any contribution of synaptic decay to account for the loss of familiarity discrimination ability over time. In Cowell et al., feature-level representations (but not object-level representations) are modified by interfering stimuli and become less and less useful as interference builds up during a delay; in Sohal and Hasselmo (2000), synaptic modifications due to intervening stimuli cause response decreases to a previously encoded item (i.e., the familiarity signal) to reduce in magnitude (p. 182). However, in the model of Norman and O'Reilly (2003), exponential decay of synaptic weights is proposed to have a role in simulating list length effects (in which adding novel items to a list causes forgetting of other items). Taken together, the models leave the question of a role for synaptic decay in forgetting somewhat open: although synaptic decay was used to simulate one instance of forgetting in one network, the question of whether these effects might have been accounted for by a mechanism solely based on interference has not been explored by other models.

Turning to the general principles of cognitive function, it is clear from Table 1 that these concepts are the subject of debate. The models are split regarding the issues of localization of function to PRC, specialization for function within PRC, and whether differences in the contributions of different temporal lobe structures to cognition are better characterized by distinct cognitive algorithms or distinct representational content. In

addition, the models also appear to make discrepant statements regarding whether PRC primarily differentiates between the stimulus inputs it receives, or assimilates them. However, in this case, there is no real conflict between the models' mechanisms. The tendency of a neural network to assimilate versus differentiate its inputs is not a qualitative distinction, unlike the notion of localized versus distributed function, or a claim of "cognitive algorithm" versus "representational content" as the organizing principle of brain function. Instead, the tendency of a network to differentiate or assimilate is a relative, quantitative value that lies on a continuum from "assimilates all stimuli" to "differentiates all stimuli." Moreover, the continuum can be defined separately for stimuli at different levels of complexity (ranging from simple features to whole objects). Thus, while Norman and O'Reilly (2003) appear to claim that PRC does primarily assimilation, whereas Cowell et al. (2006) state that the chief contribution of PRC is to differentiate objects, this debate may be moot: the former arrive at that description by comparing PRC with hippocampus, the master differentiator of the brain, whereas the latter arrive at their description by comparing PRC with upstream regions of ventral visual cortex. Indeed, the model of Cowell et al. (2010b) embodies the idea that differentiation is a mutable concept, claiming that PRC differentiates representations at one level of stimulus complexity and assimilates them at another.

ADVANCES IN UNDERSTANDING PERIRHINAL CORTEX FUNCTION

Computational models have greatly advanced our understanding of perirhinal cortex function, by several different means. First, these models have helped us to conceptualize cognitive operations in terms of neural processes. Instantiating processes such as encoding, storage, and discrimination in an explicit computational model forces the authors to give those processes concrete characteristics. For example, certain properties of the network of Bogacz et al. (2001) necessitated the proposal that encoding and familiarity discrimination are carried out in PRC neurons in two separate phases; a desire to capture particular properties of recollection behavior led Norman and O'Reilly (2003) to localize aspects of encoding and stimulus discrimination to different hippocampal regions with neural architecture suited to the purpose. Concrete characterization of these fundamental psychological processes defines the terms of the theories explicitly, which facilitates dialogue between opposing accounts, in turn forcing development of the theoretical ideas, particularly where conflicting theories generate differing predictions that can be tested.

Second, simulations with the models have enabled the exploration of low-level neural mechanisms and their consequences for the behavior of neural networks. For example, in the network of Sohal and Hasselmo (2000), simulations suggested that a gating mechanism is required in order to prevent recurrent transmission (which is necessary to simulate empirically observed maintenance of activation presumed to underlie working memory) from interfering with the acquisition of familiarity information by IT neurons. This demonstration suggested a possible role for the cholinergic modulation of IT responses by basal forebrain ACh neurons, and supplied a candidate mechanism for the action of visual attention in IT.

Third, the models have forged links between neural mechanisms and the behavior of whole organisms by using neurobiologically plausible algorithms to predict measures of recognition memory performance or visual discrimination learning. For example, Norman and O'Reilly (2003) were able to account for the counter-intuitive behavioral finding that partial hippocampal lesions cause worse impairments in recognition memory than full lesions, by demonstrating that a noisy mnemonic signal from a damaged hippocampus can do more harm than good by drowning out useful familiarity information supplied by MTL neocortex. Similarly, Bussey and Saksida (2002) were able to explain a puzzling and contradictory series of behavioral findings from the visual discrimination learning literature by appeal to a simple assumption about the organization of representations in visual cortex, and the use of a simple associative learning rule that is compatible with known neurobiology.

Fourth, an important consequence of the modeling enterprise has been to drive experiments that allow refinement of existing models and differentiation between them. Bogacz et al. (2001) suggest a list of five explicit predictions arising from their model, and in the review of Bogacz and Brown (2003b), the authors provide a summary of seven empirically testable claims that differentiate between the accounts of Bogacz et al. (2001), Bogacz and Brown (2003a), Sohal and Hasslemo (2000), and Norman and O'Reilly (2003). The network of Cowell et al. (2006) has similarly been used to simulate at least four novel predictions, each of which has been tested in a study with PRC-lesioned rats (Bartko et al., 2007a,b, 2010; McTighe et al., 2010). The last of these is a strong and highly counter-intuitive prediction, suggesting that perirhinal lesions induce forgetting by causing novel objects to appear familiar, rather than by causing familiar objects to appear novel. This prediction, which was confirmed by two experiments in rats with PRC lesions (McTighe et al., 2010), suggests reconsideration of a fundamental assumption underlying most current theories of amnesia: that forgetting arises when mnemonic information regarding familiar objects is either lost or made inaccessible. Instead, the theory suggests that forgetting is induced by a change in the representations of novel objects.

In sum, extant models of PRC function have helped to advance our understanding of cognition in PRC in many ways. They have generated fruitful debate in areas where their accounts span the same problem space, and on general principles to which many models may speak regardless of differences in their biological scale. However, keeping in mind the story of the blind men and the elephant, it is clear that there are situations where extensive debate favoring one account over another—when those accounts are couched at different levels and therefore not mutually exclusive—is not fruitful. Rather,

our ultimate goal in using computational modeling to advance our understanding of PRC should be to integrate across these disparate models to account for experimental data from multiple levels of biological scale and multiple cognitive phenomena simultaneously; only then will we see the "whole elephant."

Acknowledgments

The author gratefully acknowledges the intellectual contributions of David Huber and Lisa Saksida, provided via helpful discussions and comments on an earlier version of the manuscript.

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.

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

Barense MD, Gaffan D, Graham KS. 2007. The human medial temporal lobe processes online representations of complex objects. Neuropsychologia 45:2963–2974.

Barense MD, Henson RNA, Lee ACH, Graham KS. 2010. Medial temporal lobe activity during complex discrimination of faces, objects, and scenes: Effects of viewpoint. Hippocampus 20:389–401.

Bartko SJ, Cowell RA, Winters BD, Bussey TJ, Saksida LM. 2010. Heightened susceptibility to interference in an animal model of amnesia: Impairment in encoding, storage, retrieval—Or all three? Neuropsychologia 48:2987–2997.

Bartko SJ, Winters BD, Cowell RA, Saksida LM, Bussey TJ. 2007a. Perceptual functions of perirhinal cortex in rats: Zero-delay object recognition and simultaneous oddity discriminations. J Neurosci 27:2548–2559.

Bartko SJ, Winters BD, Cowell RA, Saksida LM, Bussey TJ. 2007b. Perirhinal cortex resolves feature ambiguity in configural object recognition and perceptual oddity tasks. Learn Memory 14:821–832.

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

Blake L, Jarvis CD, Mishkin M. 1977. Pattern-discrimination thresholds after partial inferior temporal or lateral striate lesions in monkeys. Brain Res 120:209–220.

Bliss TV, Lomo T. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. J Physiol 232:331–356.

Bogacz R, Brown MW. 2002. The restricted influence of sparseness of coding on the capacity of familiarity discrimination networks. Network 13:457–485.

Bogacz R, Brown MW. 2003a. An anti-Hebbian model of familiarity discrimination in the perirhinal cortex. Neurocomputing 52:1–6.

Bogacz R, Brown MW. 2003b. Comparison of computational models of familiarity discrimination in the perirhinal cortex. Hippocampus 13:494–524.

Bogacz R, Brown MW, Giraud-Carrier C. 2001. Model of familiarity discrimination in the perirhinal cortex. J Comput Neurosci 10: 5–23.

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, Booth MC, Rolls ET, Gaffan D. 2001. Selective perceptual impairments after perirhinal cortex ablation. J Neurosci 21:9824–9836.

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

Buckley MJ, Gaffan D. 1998a. Learning and transfer of object-reward associations and the role of the perirhinal cortex. Behav Neurosci 112:15–23.

Buckley MJ, Gaffan D. 1998b. Perirhinal cortex ablation impairs visual object identification. J Neurosci 18:2268–2275.

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.

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 15:730–737.

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.

Churchland PS, Sejnowski TJ. 1988. Perspectives on cognitive neuroscience. Science 242:741–745.

Cowell RA, Bussey TJ, Saksida LM. 2006. Why does brain damage impair memory? A connectionist model of object recognition memory in perirhinal cortex. J Neurosci 26:12186–12197.

Cowell RA, Bussey TJ, Saksida LM. 2010a. Components of recognition memory: Dissociable cognitive processes or just differences in representational complexity? Hippocampus 20:1245–1262.

Cowell RA, Bussey TJ, Saksida LM. 2010b. Functional dissociations within the ventral object processing pathway: Cognitive modules or a hierarchical continuum? J Cogn Neurosci 22:2460–2479.

Cowell RA, Bussey TJ, Saksida LM. 2011. Using computational modeling to understand cognition in the ventral visual-perirhinal pathway. In: Alonso E, Mondragon E, editors. Computational Neuroscience for Advancing Artifical Intelligence: Models, Methods and Applications. Hershey, PA: IGI Global. pp 15–45.

Cowey A, Gross CG. 1970. Effects of foveal prestriate and inferotemporal lesions on visual discrimination by rhesus monkeys. Exp Brain Res 11:128–144.

Dudek SM, Bear MF. 1992. Homosynaptic long-term depression in area CA1 of hippocampus and effects of *N*-methyl-D-aspartate receptor blockade. Proc Natl Acad Sci USA 89:4363–4367.

Eichenbaum H, Yonelinas AP, Ranganath C. 2007. The medial temporal lobe and recognition memory. Annu Rev Neurosci 30:123–152.

Elfman KW, Parks CM, Yonelinas AP. 2008. Testing a neurocomputational model of recollection, familiarity, and source recognition. J Exp Psychol 34:752–768.

Gaffan D, Harrison S, Gaffan EA. 1986. Visual identification following inferotemporal ablation in the monkey. Q J Exp Psychol B 38:5–30.

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

Gibson E. 1940. A systematic application of the concepts of generalization and differentiation to verbal learning. Psychol Rev 47:196–229.

Graham KS, Barense MD, Lee ACH. 2010. Going beyond LTM in the MTL: A synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. Neuropsychologia 48:831–853.

Gross CG, Cowey A, Manning FJ. 1971. Further analysis of visual discrimination deficits following foveal prestriate and inferotemporal lesions in rhesus monkeys. J Comp Physiol Psychol 76:1–7.

Iwai E, Mishkin M. 1968. Two visual foci in the temporal lobe of monkeys. In: Yoshii N, Buchwald N, editors. Neurophysiological Basis of Learning and Behavior. Japan: Osaka University Press.

Jenkins JG, Dallenbach KM. 1924. Obliviscence during sleep and waking. Am J Psychol 35:605–612.

Li L, Miller EK, Desimone R. 1993. The representation of stimulus-familiarity in anterior inferior temporal cortex. J Neurophysiol 69:1918–1929.

Malkova L, Bachevalier J, Mishkin M, Saunders RC. 2001. Neurotoxic lesions of perirhinal cortex impair visual recognition memory in rhesus monkeys. Neuroreport 12:1913–1917.

McClelland JL, Chappell M. 1998. Familiarity breeds differentiation: A subjective-likelihood approach to the effects of experience in recognition memory. Psychol Rev 105:724–760.

McClelland JL, Mcnaughton BL, O'Reilly RC. 1995. Why there are complementary learning-systems in the hippocampus and neocortex—Insights from the successes and failures of connectionist models of learning and memory. Psychol Rev 102:419–457.

McGeoch JA. 1932. Forgetting and the law of disuse. Psychol Rev 39:352–370.

McTighe SM, Cowell RA, Winters BD, Bussey TJ, Saksida LM. 2010. Paradoxical false memory for objects after brain damage. Science 330:1408–1410.

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.

Miller EK, Li L, Desimone R. 1991. A neural mechanism for working and recognition memory in inferior temporal cortex. Science 254:1377–1379.

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'Reilly RC, Rudy JW. 2001. Conjunctive representations in learning and memory: Principles of cortical and hippocampal function. Psychol Rev 108:311–345.

Ranganath C. 2010. A Unified framework for the functional organization of the medial temporal lobes and the phenomenology of episodic memory. Hippocampus 20:1263–1290.

Ranganath C, Yonelinas AP, Cohen MX, Dy CJ, Tom SM, D'Esposito M. 2004. Dissociable correlates of recollection and familiarity within the medial temporal lobes. Neuropsychologia 42:2–13.

Ratcliff R, Sheu CF, Gronlund SD. 1992. Testing global memory models using ROC curves. Psychol Rev 99:518–535.

Saksida LM, Bussey TJ. 2010. The representational-hierarchical view of amnesia: Translation from animal to human. Neuropsychologia 48:2370–2384.

Shiffrin RM, Ratcliff R, Clark SE. 1990. List-strength effect. 2. Theoretical mechanisms. J Exp Psychol 16:179–195.

Shiffrin RM, Steyvers M. 1997. Model for recognition memory: REM—Retrieving effectively from memory. Psychon Bull Rev 4:145–166.

Sohal VS, Hasselmo ME. 2000. A model for experience-dependent changes in the responses of inferotemporal neurons. Network 11:169–190.

Squire LR, Stark CE, Clark RE. 2004. The medial temporal lobe. Annu Rev Neurosci 27:279–306.

Squire LR, Zola-Morgan S. 1991. The medial temporal-lobe memory system. Science 253:1380–1386.

Standing L. 1973. Learning 10,000 pictures. Q J Exp Psychol 25: 207–222.

Suzuki WA, Amaral DG. 1994. Perirhinal and parahippocampal cortices of the macaque monkey—Cortical afferents. J Comp Neurol 350:497–533.

Wilson M, Kaufman HM, Zieler RE, Lieb JP. 1972. Visual identification and memory in monkeys with circumscribed inferotemporal lesions. J Comp Physiol Psychol 78:173–183.

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.

Wixted JT. 2004. The psychology and neuroscience of forgetting. Annu Rev Psychol 55:235–269.

Yonelinas AP. 1994. Receiver-operating characteristics in recognition memory—Evidence for a dual-process model. J Exp Psychol 20:1341–1354.

Zola-Morgan S, Squire LR, Amaral DG, Suzuki WA. 1989. Lesions of perirhinal and parahippocampal cortex that spare the amygdala and hippocampal formation produce severe memory impairment. J Neurosci 9:4355–4370.