

A Computational Model of Perceptual Deficits in Medial Temporal Lobe Amnesia

Patrick Sadil (psadil@psych.umass.edu)

Department of Psychological and Brain Sciences, University of Massachusetts, 135 Hicks Way
Amherst, MA, 01003 USA

Rosemary A. Cowell. (rcowell@psych.umass.edu)

Department of Psychological and Brain Sciences, University of Massachusetts, 135 Hicks Way
Amherst, MA, 01003 USA

Abstract

Damage to the Medial Temporal Lobe (MTL) impairs declarative memory and perception. The Representational-Hierarchical (RH) Account explains such impairments by assuming that MTL stores conjunctive representations of items and events, and that individuals with MTL damage must rely upon representations of simple visual features in posterior visual cortex. A recent study revealed a surprising anti-perceptual learning effect in MTL-damaged individuals: with exposure to a set of visual stimuli, discrimination performance worsened rather than improved. We expand the RH account to explain this paradox by assuming that visual discrimination is performed using a familiarity heuristic. Exposure to a set of highly similar stimuli entails repeated presentation of simple visual features, eventually rendering all feature representations equally (maximally) familiar and hence inutile for solving the task. Since the unique conjunctions represented in MTL do not occur repeatedly, healthy individuals are shielded from this perceptual interference. We simulate this mechanism with a neural network previously used to simulate recognition memory, thereby providing a model that accounts for both mnemonic *and* perceptual deficits caused by MTL damage using a unified architecture and mechanism.

Keywords: Neural Network; memory; visual perception; Medial Temporal Lobe; hierarchical object representations

Introduction

Revisiting the Classical, Modular Account of Memory and Perception Long-term, declarative memory has long been known to depend on medial temporal lobe (MTL) structures (Scoville & Milner, 1957; Squire & Zola-Morgan, 1991). However, evidence now suggests an additional role for MTL in other cognitive functions such as visual perception of objects and scenes (Barens et al., 2012; Bartko, et al., 2007; Lee et al., 2005).

A theory termed the Representational Hierarchical (RH) account explains both mnemonic and perceptual deficits caused by MTL damage (Bussey & Saksida, 2002; Cowell, Bussey, & Saksida, 2006, 2010). The RH account assumes that the ventral visual stream contains a hierarchical organization of representations that continues into MTL. Early stages of the pathway (e.g., V1, V2, V4) are assumed to represent simple visual features (e.g., color, orientation), and these simple features are brought together into conjunctions of increasing complexity in anterior brain regions. The hierarchy culminates in MTL where the conjunctions correspond to whole objects, scenes, or complex episodic events. The RH

account claims that conjunctive representations in MTL are important whenever a cognitive task – perceptual or mnemonic – uses stimuli with overlapping features, such that individual feature representations in posterior regions provide ambiguous information (Bussey & Saksida, 2002).

Paradoxical Finding of an Exposure-induced Deficit in Visual Discrimination Performance Barens et al. (2012) documented a new and puzzling way in which MTL lesions impair visual discrimination. MTL amnesics and healthy controls were asked to judge whether pairs of simultaneously presented abstract stimuli were the same or different (Fig. 1). In the High Ambiguity condition, each pair of to-be-discriminated stimuli shared 2 out of 3 explicitly defined features, whereas in the Low Ambiguity condition, the items in a pair did not share any of the defined features. Participants were to declare ‘mismatch’ if any of the 3 features differed in a pair. Amnesic participants were unimpaired at discriminating Low Ambiguity objects, but discrimination of High Ambiguity objects deteriorated dramatically in the second half of trials in this condition (Fig. 2).




Figure 1: Experiment 3 of Barens et al. (2012). Stimuli were defined by 3 features: inner shape, outer shape, and fill pattern. High Ambiguity, mis-matching pairs shared 2 of these features, but Low Ambiguity pairs share 0.

Barens et al. explained their data in terms of the RH account: Individuals with MTL damage lack conjunctive representations of objects usually stored in perirhinal cortex (PRC), a structure in MTL. Objects are instead represented as individual simple features in posterior visual cortex. In the task of Barens et al., each stimulus is a unique conjunction of features, but the features comprising the stimuli repeat often over trials. After viewing many items, feature-level interference renders feature representations in posterior visual

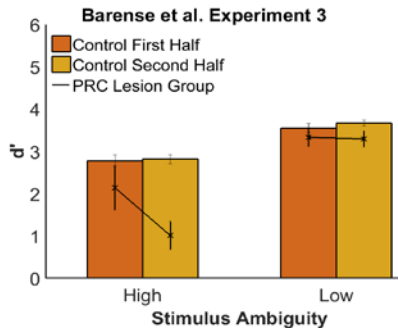


Figure 2: Data from Barensen et al. (2012). Subjects with perirhinal (PRC) lesions were impaired at discriminating High Ambiguity stimuli in the second half of trials. Significance was assessed via Crawford’s t-test for each Lesion participant separately (Control n=8; Lesion n=2).

cortex inadequate for solving difficult (High Ambiguity) discriminations. Control subjects resolve this interference with a unique conjunction in PRC for each stimulus, but if PRC is damaged, performance is impaired.

However, the *decrease* in MTL amnesics’ performance with increasing exposure to the stimuli contrasts with well-established perceptual learning effects. Perceptual learning is usually explained by assuming that experience increases the separability of stimuli, either because stimulus representations become less overlapping (Saksida, 1999; Schoups, Vogels, Qian, & Orban, 2001; Yang & Maunsell, 2004), or because the weights via which representations influence decision-making are optimized (Kumano & Uka, 2013; Liu, Doshier, & Lu, 2015). But if discrimination relies on the separability of stimulus representations, and exposure differentiates representations, then damage to the brain areas containing those representations might be expected to *abolish any improvements* conferred by exposure, but it is not clear why brain damage should cause exposure to *hurt* performance. Put another way, even if individuals with MTL damage possess only feature representations, it is not apparent why exposure should cause feature representations to become *more* overlapping, rather than *less*.

How can this contradiction be resolved? In a previous neural network instantiation of the RH account (Cowell et al., 2006) we accounted for impairments in *recognition memory* induced by MTL damage by considering the *familiarity signal* evoked by stimulus representations in the brain. Exposure to many items sharing visual features entails frequent repetition of the features. Eventually, the representations of all features in posterior visual cortex appear familiar, causing individuals with MTL damage (who possess only feature representations) to perceive all items as equally familiar, impairing recognition memory. This surprising mechanism has been empirically supported (e.g, McTighe, et al., 2010). Here, we invoke the same mechanism to explain the *visual discrimination* impairments reported by Barensen et al. To apply this account to visual discrimination, we make an assumption about the strategy used to solve the task: we propose that partici-

pants did not discriminate stimuli based on their representational overlap, but instead used a *familiarity heuristic* to decide whether two items were identical (Goldstein & Gigerenzer, 2002).

A Familiarity Heuristic Resolves the Paradox of Exposure-induced Discrimination Impairments We assume that, in difficult discrimination tasks like that of Barensen et al. (2012), participants search for a mismatch between two stimuli. To do so, they visually scan back and forth between the stimuli; if switching from one item to the other elicits an impression of novelty relative to the item just examined, this is taken as evidence for a mismatch. That is, the new item appears "unfamiliar" to the extent that it differs from the item inspected immediately before switching. If switching elicits a decrease in familiarity that exceeds some threshold, the two stimuli are judged to mismatch. Hence, the model’s signal for familiarity – the 'tunedness' of the stimulus representations (Cowell et al., 2006) – can be used to perform discrimination.

Just as in the memory simulations of Cowell et al. (2006), representations in the model can ‘saturate’ (i.e., reach maximum) in terms of familiarity. In a discrimination task in which all stimuli are similar, the stimulus features appear repeatedly, resulting in all feature representations becoming highly tuned, i.e., familiar. After sufficient interference, all feature representations possess the same, maximum, level of familiarity and can no longer be used to discriminate two objects via a familiarity heuristic. In contrast, because individual objects do not repeat across trials, familiarity for whole objects does not saturate and so remains useful for discrimination: conjunctive representations in PRC shield a person from perceptual interference. When these conjunctive representations are compromised by MTL damage, visual discrimination is impaired. A key tenet of the RH account is that memory and perception share common neural mechanisms. Our model embodies this by using a mnemonic signal (familiarity) to solve a visual perceptual task.

Description of the Model

Model Architecture We use the model of Cowell et al. (2006) with only minor modifications. The network contains a PRC layer and a layer corresponding to posterior ventral visual stream (Fig. 3). Visual objects are instantiated as 8-dimensional vectors, given as input to the model. We assume that posterior regions represent simple conjunctions of two individual visual dimensions, so the Posterior layer is divided into 4 grids: nodes in each posterior grid receive 2 input dimensions and combine them into a simple conjunction, termed a ‘feature’. In contrast, because PRC is assumed to represent whole objects, all 8 input dimensions converge into a single conjunction in the PRC layer. Thus, nodes in the PRC layer contain a unique, conjunctive representation of a four-featured object, whereas the Posterior layer represents those four, 2-dimensional features separately. We simulate PRC damage by removing the PRC layer.

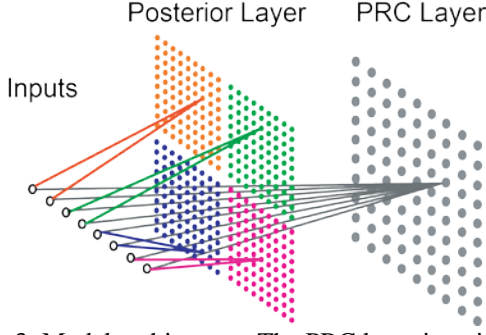


Figure 3: Model architecture. The PRC layer is a single Kohonen grid, representing an object as a unique conjunction. The Posterior layer contains 4 Kohonen grids, each representing a 2-dimensional visual ‘feature’.

All model layers are constructed from Kohonen grids, which mimic properties of cortex, including information processing mechanisms such as Hebbian learning and lateral inhibition (Kohonen, 1984). A Kohonen grid (or self-organizing map) is trained by successively presenting stimulus inputs and incrementally adapting the weights of the grid units on each presentation. As the grid learns, its stimulus representations are sharpened or ‘tuned’: A novel stimulus elicits a broadly distributed pattern of activity, whereas a stimulus that has been repeatedly presented elicits a highly selective activation pattern with a peak over one set of units, relative to the activation levels elsewhere across the grid. In our model, the selectivity of the activation indexes familiarity (Cowell et al., 2006; Norman & O’Reilly, 2003).

Each dimension of an input stimulus can take one of four values in a given stimulus: 0.05, 0.35, 0.65, or 0.95. This scheme yields $4^8 = 65,536$ unique object stimuli, which are represented holistically on the PRC layer, but only $4^2 = 16$ unique, 2-dimensional features on each Posterior grid. This scheme reflects a key assumption of the model: that there is a vast number of possible visual objects in the world, but they are composed from a small number of visual elements.

Initialization and Pretraining All grids contained 200×200 nodes whose weights were initialized with random values between 0 and 1. Networks were pre-trained for 500 cycles according to the standard Kohonen learning rule,

$$w_i(t+1) = w_i(t) + f(r, t) * (stim - w_i(t)) \quad (1)$$

and,

$$f(r, t) = \eta(t) * v(r, t), \quad (2)$$

where w_i refers to the weights of node i , t is the current cycle, $stim$ is stimulus input, $\eta(t)$ is the learning-rate, r is the city-block distance of node i from the most strongly active (winning) node, and $v(r, t)$ is a neighborhood function that scales the learning rate (Eq. 3).

In the pre-training phase, both η and v decrease with each cycle. The neighborhood is defined by a Gaussian function:

$$v(r, t) = \exp\left(-\frac{r}{G(t)}\right)^2 \quad (3)$$

where, $G(t) = 0.5 + 10 t^{-B}$, and B is a constant determining the rate of shrinkage of the neighborhood function. The learning

rate decreases as $\eta(t) = t^{-A}$, where the constant A determines the rate of decrease.

In each pre-training cycle, the network sees a different, unique stimulus. For simulations of the Barense et al. (2012) experiments, η and G are constants fixed at the values of the final training cycle ($\eta = \eta(500)$; $G = G(500)$).

Measuring Familiarity Activation (a) of nodes was determined by the sigmoidal function:

$$a_i = \frac{1}{1 + \exp\left(-k * \ln\left(\frac{1}{dist}\right)\right)} \quad (4)$$

where k is a constant that determines the steepness of the sigmoidal function, $dist$ is the mean squared error between a node’s weights (w_i) and the stimulus input vector.

The familiarity of a stimulus is given by the selectivity of its activation pattern, calculated separately in each grid. Selectivity is the activation of the peak (the summed activation of the winning node and its nearest 4 neighbors) divided by the summed total activation of the grid. Thus, via normalization, higher familiarity corresponds to more ‘tuned’ representations. Stimulus familiarity is measured separately in each grid in the network, yielding a single, object-level familiarity score from the PRC layer and four separate feature-level familiarity scores in the posterior layer; when comparing two stimuli, the posterior feature-level familiarity scores are compared separately for each pair of features. For Control networks, this includes 4 Posterior grids and 1 PRC grid; for Lesioned networks, only 4 Posterior grids. In all simulations, $k = 0.08$, $B = 0.3$, and $A = 0.6$.

Simulating Visual Discrimination Behavior

Fixations In Barense et al. (2012), participants decided whether two stimuli were the same or different. Eye-tracking data from control subjects revealed a higher ratio of within-stimulus to between-stimulus fixations in High Ambiguity trials (~1.2) than in Low Ambiguity trials (~0.6). Barense et al. conjectured that this reflected a greater tendency to ‘bind’ features together in the High Ambiguity condition. Accordingly, we hypothesized that differential fixation ratios might make an important contribution to task performance (e.g., sampling stimuli with a higher within: between ratio may enable more reliable PRC representations) and aligned model parameters with these empirical data. In addition, on trials declared as a ‘match’, participants made about 25 and 20 fixations for High and Low Ambiguity, respectively.

In the model, one stimulus fixation entails 20 encoding cycles (Eq. 1). Stimuli are ‘sampled’ via fixations, using a probabilistic rule for switching back and forth between two items in a pair, with switch probability derived from the empirical within:between ratio for each experimental condition (1.2 or 0.6). On any trial, the maximum number of fixations (i.e., the point at which the search for a mismatch terminates and a ‘match’ is declared, see below) is 25 and 20 in the High and Low Ambiguity conditions, respectively (since the stopping point is somewhat arbitrary and not specified by our theory, we used empirical parameters to determine this).

Discrimination Decisions On each trial, two stimuli are presented. One is arbitrarily selected first (Item A) and the model samples it via successive fixations until switching to the other stimulus, Item B. Upon switching, the model assesses evidence for a 'mismatch' by computing a novelty score and comparing it to a threshold. The novelty score is a measure of 'familiarity change' obtained by taking the familiarity of Item A and subtracting the familiarity of the Item B. When the two items are identical, the novelty score is zero; when they are different, Item B should have slightly lower familiarity than Item A, yielding a positive novelty score. If the novelty score in any individual grid (4 Posterior grids or, when available, PRC) exceeds the threshold, the items are declared to 'mismatch'. At each switch between the two stimuli, if the network finds no evidence for a mismatch, the network continues fixating. In the next comparison, Item B serves as the previously inspected stimulus and Item A as the newly fixated stimulus. Comparisons proceed until either a mismatch is declared or the maximum number of fixations is reached, whereupon a 'match' is declared.

Criterion shift We assume that participants adjust their decision rule as their stimulus representations adapt. That is, if participants begin to perceive *all* items as looking more similar, they require less evidence to declare that two items are mismatching. To effect this, the threshold value of novelty required to declare a pair of items as mismatching (i.e., the *decision criterion*) is allowed to shift by setting it equal to the average of the novelty ('familiarity change') score on the previous six trials. In addition, for each decision, noise drawn from the uniform distribution ($\pm 1e^{-6}$) is added to the threshold. As familiarity change scores decrease, this noise swamps the discrimination signal present in that score. The starting threshold value (i.e., threshold on Trial 1) is set to 2x the maximum noise: $2e^{-6}$.

Simulation 1

Target Empirical Data In *Experiment 3* of Barense et al. (2012), participants indicated whether two simultaneously presented visual stimuli were a match or a mismatch. Stimuli were trial-unique items constructed of 3 features (Fig. 1). In Low Ambiguity trials, the pairs shared none of these 3 features. In High Ambiguity trials, the items shared 2 out of 3 features. Individuals with PRC damage performed similarly to controls at discriminating Low Ambiguity pairs, but for High Ambiguity pairs their performance was intact at first but fell sharply in the second half of the task (Fig. 2).

Stimuli and Task All stimuli were trial unique. Low Ambiguity stimulus pairs shared no 2-dimensional features, whereas High Ambiguity pairs shared 3 out of 4 features (we used 4 total features for consistency with Cowell et al., 2006). Owing to assumptions governing stimulus construction (see *Model Architecture*, above) there are many possible unique stimuli, but the features comprising them appear repeatedly. In addition, we further constrained the feature set to reflect

the high degree of feature-overlap in Barense et al. by constructing stimuli (unique 4-featured objects) using only 6 out the 16 possible 2-dimensional features for each Posterior grid (where a 'feature' is a conjunction of 2 input dimensions). This yielded 6^4 total possible objects, with high feature-overlap among them. As in Barense et al., each condition contained 36 'match' and 36 'mismatch' trials.

Results Networks with no PRC layer ('PRC Lesion Group'), like humans with PRC damage, were impaired relative to controls at High but not Low Ambiguity. Critically, the impairment was worse in the second half of trials (Fig. 4). We do not report statistics on simulated data because the results are highly reliable such that variance scales arbitrarily with the number of networks run. Instead, we focus on qualitative patterns, which match those of the patient data: the interaction between Lesion Group, Stimulus Ambiguity and First/Second Half.

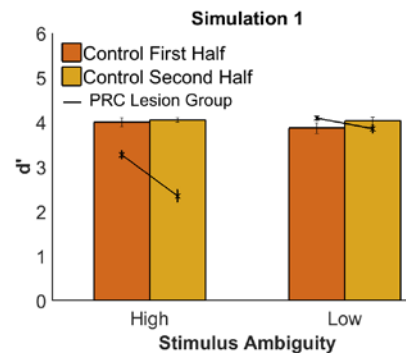


Figure 4: Figure 4: Simulated data for Experiment 3 of Barense et al. (2012). Compare to the empirical data in Figure 2. $n = 50$ networks per group.

Simulation 1 Discussion The model's simulation of a discrimination deficit for High Ambiguity stimuli in the second half of trials, following PRC damage, hinges on three assumptions: (1) participants solve the task using a familiarity heuristic; (2) task stimuli are unique but contain many low-level features that repeat over trials such that all features eventually appear familiar; (3) stimuli are represented in PRC as whole conjunctions but in posterior regions as individual features: when PRC is missing, discrimination performance is impaired once all features are maximally tuned.

Performance in lesioned networks relies upon posterior feature representations. At the start of the task, individual features are not highly tuned. On each new trial, the features of the first stimulus inspected increase in familiarity. When the network switches to inspect the second stimulus, if that stimulus is not identical to the first, its features are lower in familiarity and a mismatch is correctly declared. However, after many trials, all features have been repeatedly encoded by the network. Now, at the start of a trial, there can be no increase in familiarity when the network inspects the first stimulus. When it switches to the second stimulus, even if that item differs from the first, its features have equal familiarity – the maximum value. At this point, a familiarity heuristic can no longer detect mismatching stimulus pairs. The feature-level

interference has more effect at High than Low Ambiguity because High Ambiguity pairs share 3 out of 4 features whereas Low Ambiguity pairs share no features. Networks search for *any* pair of features across the two stimuli that differ. Since there are 4 mismatching features in Low pairs, the network is more likely to discover at least 1 feature that has not yet reached maximum familiarity.

Performance in control networks is maintained throughout the task because of conjunctive representations in the PRC layer. Individual stimuli are trial-unique, i.e., whole conjunctions are never repeated, so representations in PRC never reach maximum familiarity. At the start of each new trial, the PRC representation for the first stimulus inspected increases in familiarity during the inspection. When the model switches to inspect the second stimulus, if the second stimulus differs from the first, the second will exhibit lower familiarity and the pair will be declared to mismatch.

Simulation 2

Target Empirical Data In *Experiment 4* of Barense et al. (2012), all subjects completed 3 blocks: Low Interference, then High Interference, then a *second* Low Interference block. High Interference was the same as High Ambiguity in *Experiment 3*: pairs of abstract stimuli shared 2 out of 3 features. In Low Interference blocks, two-thirds of trials contained a pair of color photographs of real-world objects. Photo images were trial unique and shared few low-level features with the abstract stimuli of critical comparison trials (Fig. 5). *Experiment 4* replicated the result of *Experiment 3*: MTL-damaged patients showed impaired discrimination at High Interference, but not at Low Interference, even in block 3 of the test session (Fig. 6, left). This suggested that the impairment seen in *Experiment 3* was not due to fatigue.




Figure 6: Experiment 4 of Barense et al. (2012). Low Interference blocks used photo stimuli in 2/3 of trials, which shared few features with stimuli on critical comparison trials.

Stimuli and Task As in Simulation 1, abstract stimuli were modeled by using only 6 of the 16 possible stimulus features on each Posterior grid to construct four-featured stimulus wholes, yielding high feature-overlap among stimuli. To reflect the assumption of Barense et al. that abstract stimuli shared no 2-dimensional features with photo stimuli, we constructed photo stimulus inputs from the remaining 10 abstract features (i.e., a non-overlapping set of features).

Networks performed 3 discrimination blocks. A block contained 88 trials, in which every third trial was a critical comparison pair of abstract stimuli (15 matching, 15 mismatch-

ing). In High Interference blocks the remaining 58 trials contained extra pairs of High Ambiguity objects. In Low Interference blocks the remaining 58 trials contained pairs of 'photo' stimuli. As in Barense et al., performance was judged on critical trials, which occurred at every third position in the sequence and contained a pair of stimuli sharing many (3 out of 4) features. The difference between High and Low Interference was that, for Low, the trials interposed between critical trials contained items sharing no features with critical-trial stimuli whereas, for High, interposed trials contained items similar to critical-trial stimuli.

Results Networks with no PRC layer were able to discriminate at the level of Control networks in Low Interference blocks (the first and third blocks), but their performance was impaired in High Interference (second) block (Fig. 6), mirroring the results of Barense et al. (2012).

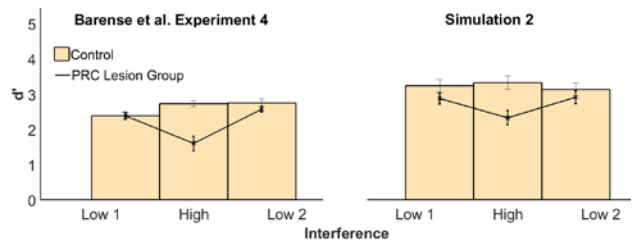


Figure 5: Barense et al., Experiment 4: empirical data (left) and model simulations (right). $n = 50$ networks per group.

Simulation 2 Discussion The same mechanism that produced the effects of PRC lesions in Simulation 1 also drives PRC lesion deficits in Simulation 2. At High Interference, because all trials contain the same class of stimuli, stimulus features appear repeatedly over trials, and posterior feature representations reach maximum familiarity. Once this occurs, there is no increase in familiarity when a network inspects a new stimulus at the start of a trial. Consequently, a network without a PRC layer cannot detect 'novelty' (a drop in familiarity) upon switching to the other stimulus.

At Low Interference, two-thirds of trials contain photos composed of different features than the critical-trial stimuli. The critical-trial stimulus features repeat too infrequently for their representations to reach maximum familiarity, hence lesioned networks remain unimpaired.

Discussion

Barense et al. (2012) reported a striking perceptual deficit in patients with MTL damage: the *accumulation* of perceptual experience impairs visual discrimination. This result is paradoxical because perceptual discrimination typically *improves* with exposure to the stimuli. Barense et al. argued that MTL-lesioned patients suffer from accumulated feature-level interference, which – in the absence of conjunctive MTL representations – cannot be overcome by feature representations in posterior visual cortex. Although we concur with this explanation, we suggest that it is incomplete.

Standard theories of perceptual learning claim that experience improves discrimination performance by reducing the

overlap between stimulus representations (i.e., training increases representational separation). In such theories, the assumed mechanism for visual discrimination is that discriminability is proportional to the overlap between representations (Saksida, 1999; Schoups et al., 2001). But such proportionality could not account for the data of Barense et al.: if exposure separates representations, even feature-based discrimination should improve with exposure because even feature representations should get *less* overlapping with training. To explain why the performance of MTL patients in Barense et al. got *worse* after exposure to the stimuli, a theory based on representational overlap would require the counter-intuitive assumption that while cortical representations of stimuli underlying perceptual learning often become *less* overlapping with exposure, posterior feature representations in this task became *more* overlapping.

In our account of the data, we eschew representational overlap as the mechanism for visual discrimination. Instead, we take the account offered by Barense et al. (i.e., amnesics suffer from the loss of conjunctive MTL representations) and combine it with a less commonly invoked discrimination mechanism: a *familiarity heuristic*. Under this account – as in prior instances of the RH account applied to *memory* (Cowell et al., 2006; McTighe et al., 2010) – representations of features, but not conjunctions, reach maximum familiarity as interference accrues. Thus, after MTL damage, perceptual experience impairs *visual discrimination performance*.

Before concluding, we clarify two important points. First, although we simulate only two layers in the ventral pathway, this is a subset of the full hierarchy of representations, which includes simpler layers prior to our model's posterior layer and more complex layers such as hippocampus after PRC. Other tasks may require other layers. For example, a discrimination task involving whole objects that repeat would require hippocampal representations capable of combining objects with (e.g.) context, to shield participants from object-level interference that would detrimentally affect PRC representations. Second, we do not suggest that a familiarity heuristic must be used in all discriminations, including easy tasks in which stimuli differ on the basis of simple, salient features such as color: such tasks could be solved by a more standard discrimination mechanism of assessing representational overlap. Our model suggests only that a familiarity heuristic is used for difficult discriminations between similar stimuli such as used in Barense et al.

In sum, we simulated the visual discrimination performance of MTL amnesics with a neural network model that differs only in trivial details from the recognition memory model of Cowell et al. (2006). To our knowledge, this is the first computational model to simulate both mnemonic and perceptual deficits caused by MTL damage using a unified architecture and mechanism.

Acknowledgements

We thank Morgan Barense for sharing her data and stimuli, and for helpful feedback. This work was supported by NSF CAREER Award 1554871 to R.A.C.

References

- Barense, M. D., Groen, I. I. A., Lee, A. C. H., Yeung, L. K., Brady, S. M., Gregori, M., Kapur, N., Bussey, T. J., Saksida, L. M., Henson, R. N. A. (2012). Intact memory for irrelevant information impairs perception in amnesia. *Neuron*, 75(1), 157–167.
- Bartko, S. J., Winters, B. D., Cowell, R. A., Saksida, L. M., & Bussey, T. J. (2007). Perirhinal cortex resolves feature ambiguity in configural object recognition and perceptual oddity tasks. *Learning & Memory*, 14(12), 821–32.
- Bussey, T. J., & Saksida, L. M. (2002). The organization of visual object representations: a connectionist model of effects of lesions in perirhinal cortex. *The European Journal of Neuroscience*, 15(2), 355–364.
- Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2006). Why does brain damage impair memory? A connectionist model of object recognition memory in perirhinal cortex. *The Journal of Neuroscience*, 26(47), 12186–12197.
- Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2010). Functional dissociations within the ventral object processing pathway: cognitive modules or a hierarchical continuum? *Journal of Cognitive Neuroscience*, 22(11), 2460–2479.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1), 75–90.
- Kohonen, T. (1984). *Self-Organizing Maps* (Third Edit). Berlin-Verlag: SPRINGER.
- Kumano, H., & Uka, T. (2013). Neuronal mechanisms of visual perceptual learning. *Behavioural Brain Research*, 249, 75–80.
- Lee, A. C. H., Bussey, T. J., Murray, E. A., Saksida, L. M., Epstein, R. A., Kapur, N., Hodges, J. R., Graham, K. S. (2005). Perceptual deficits in amnesia: challenging the medial temporal lobe “mnemonic” view. *Neuropsychologia*, 43(1), 1–11.
- Liu, J., Doshier, B. A., & Lu, Z. (2015). Augmented Hebbian reweighting accounts for accuracy and induced bias in perceptual learning with reverse feedback. *Journal of Vision*. 15(2015), 1–21.
- McTighe, S. M., Cowell, R. A., Winters, B. D., Bussey, T. J., & Saksida, L. M. (2010). Paradoxical False Memory for Objects After Brain Damage. *Science*, 330(6009), 1408–1410.
- Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychological Review*, 110(4), 611–646.
- Saksida, L. M. (1999). Effects of similarity and experience on discrimination learning: a nonassociative connectionist model of perceptual learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 25(3), 308–323.
- Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, 412(6846), 549–553.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery & Psychiatry*, 20(11), 11–21.
- Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science*, 253(5026), 1380–6.
- Yang, T., & Maunsell, J. H. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *Journal of Neuroscience*, 24(7), 1617–1626.