Check for updates

Original Article

QJEP

Quarterly Journal of Experimental Psychology
1–20
© Experimental Psychology Society 2018
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1747021818768502
qjep.sagepub.com
SAGE

Computational and empirical simulations of selective memory impairments: Converging evidence for a single-system account of memory dissociations

Evan T Curtis[1] and Randall K Jamieson[2]

Abstract

Current theory has divided memory into multiple systems, resulting in a fractionated account of human behaviour. By an alternative perspective, memory is a single system. However, debate over the details of different single-system theories has overshadowed the converging agreement among them, slowing the reunification of memory. Evidence in favour of dividing memory often takes the form of dissociations observed in amnesia, where amnesic patients are impaired on some memory tasks but not others. The dissociations are taken as evidence for separate explicit and implicit memory systems. We argue against this perspective. We simulate two key dissociations between classification and recognition in a computational model of memory, A Theory of Nonanalytic Association. We assume that amnesia reflects a quantitative difference in the quality of encoding. We also present empirical evidence that replicates the dissociations in healthy participants, simulating amnesic behaviour by reducing study time. In both analyses, we successfully reproduce the dissociations. We integrate our computational and empirical successes with the success of alternative models and manipulations and argue that our demonstrations, taken in concert with similar demonstrations with similar models, provide converging evidence for a more general set of single-system analyses that support the conclusion that a wide variety of memory phenomena can be explained by a unified and coherent set of principles.

Keywords

Memory; amnesia; recognition; classification; dissociation

Received: 22 September 2017; revised: 1 March 2018; accepted: 9 March 2018

Cognitive psychology has a history of dividing and deconstructing behaviours into different functions served by different systems. The benefit is a disciplinary division of labour in which different laboratories examine different problems (e.g., classification and recognition) with the eventual hope of reunification. However, that goal has been neglected, and instead, the strategy of division and deconstruction has fostered a fractionated explanation of human behaviour. The current state of memory theory provides a striking example of the problem.

The dominant view is that memory is composed of multiple systems, each serving unique purposes and operating under unique principles (e.g., Eichenbaum & Cohen, 2001; Schacter & Tulving, 1994; Tulving, 1985). By a typical conceptualisation of the multiple-system perspective, memory is divided into an explicit system and an implicit system which are further divided into systems for episodic memory, semantic memory, habits, non-associative learning, classical conditioning, and priming (Cohen & Squire, 1980). Each of those systems is the subject of a wealth of research and theory, but there is very little integration among them.

A dissenting view is that memory is a unified cognitive system that serves multiple purposes and operates under a common set of principles (e.g., Benjamin, 2010; Berry, Kessels, Wester, & Shanks, 2014; Higham & Vokey, 1994; Higham, Vokey, & Pritchard, 2000; Jamieson & Mewhort, 2009a, 2009b, 2010, 2011; Kinder & Shanks, 2001, 2003; Nosofsky, 1986, 1987, 1988, 1991; Nosofsky & Zaki, 1998;

[1]Department of Psychology, Booth University College, Winnipeg, Manitoba, Canada
[2]Department of Psychology, University of Manitoba, Winnipeg, Manitoba, Canada

Corresponding author:
Evan T Curtis, Department of Psychology, Booth University College, Winnipeg, Manitoba R3B 2P2, Canada.
Email: evan_curtis@boothuc.ca

Shanks & Perruchet, 2002; Shanks & St. John, 1994; Surprenant & Neath, 2013; see also Bussey & Saksida, 2007; Gaffan, 2002 for the same argument applied to non-human animals). However, the single-system account is, itself, divided into a number of competing single-system theories that agree in spirit despite differing in detail. Most of the theories have been examined in isolation, with each researcher working within his or her preferred theory. In more extreme cases, theorists have focused on critiquing one theory in favour of another (Kinder, 2010; Murdock, 2008; Nairne & Neath, 1994; Starns, White, & Ratcliff, 2010). In our view, the focus on differentiating single-system accounts of memory dissociation has overshadowed the strong agreement among the different theories. What forms of evidence are preventing a reunification of memory?

The bulk of evidence for the multiple-system perspective rests on empirical dissociations, where an independent variable (e.g., word frequency) has differential influence on two dependent variables (e.g., recognition and recall). Dissociations can occur under a wide variety of circumstances. However, the strongest cited evidence comes from observations of amnesic patients, who suffer from impaired recognition and recall but perform normally in other laboratory tasks such as categorisation and priming (e.g., Cohen & Squire, 1980; Knowlton, Ramus, & Squire, 1992; Knowlton & Squire, 1993, 1994; Knowlton, Squire, & Gluck, 1996; Squire, 2004). By the multiple-system perspective, the dissociations imply a distinction between the systems supporting the two classes of tasks. Explicit memory supports recognition and recall, whereas implicit memory supports categorisation and priming. Following on that premise, amnesia reflects selective damage to the explicit memory system.

However, the implication of a dissociation is uncertain. First, a dissociation is often consistent with both the single-system and the multiple-system perspective (Dunn & Kirsner, 1988, 2003; Van Orden & Kloos, 2003; Van Orden, Pennington, & Stone, 2001). By rules of reason, evidence that is consistent with two competing theories lacks the necessary logical force to reject one in favour of the other. Therefore, many documented dissociations lack the necessary logical force to reject the single-system perspective in favour of the multiple-system perspective. Second, important differences in methods and materials between the classes of tasks are often glossed over, making it unclear which factors drive the dissociations. For example, recognition requires discrimination between studied and unstudied items, whereas classification requires discrimination among unstudied items, a difference that can be sufficient to produce a dissociation (Jamieson, Holmes, & Mewhort, 2010).

In defence of the single-system perspective, theorists have adopted a strategy of computational analysis and implemented formal models of memory to reproduce the dissociations observed in experiments (e.g., Benjamin, 2010; Jamieson, Holmes, & Mewhort, 2010; Kinder & Shanks, 2001, 2003; Malmberg, Zeelenberg, & Shiffrin, 2004; Nosofsky, Little, & James, 2012; Nosofsky & Zaki, 1998; Shanks & Perruchet, 2002; Zaki & Nosofsky, 2001). By that strategy, each demonstration is added to a database of demonstrations that provides support for the claim that dissociations do not force the multiple-system perspective as traditionally thought.

However, the success of the models does not provide direct evidence against the multiple-system perspective. In fact, it is unclear what kind of evidence would be sufficient to reject the multiple-system perspective. The problem is obvious: a single-system theory can be nested in a multiple-system competitor. Therefore, a multiple-system theory can typically accommodate any evidence that its single-system competitor can (and more) strictly by virtue of its added complexity.[1] Parsimony dictates favouring the account with fewer parameters (i.e., the single-system perspective). However, parsimony can be an ambiguous term in the context of modern psychological theory.

Theorists have also adopted a strategy of experimental analysis, introducing manipulations to simulate amnesic behaviour in healthy controls (e.g., Graf, Mandler, & Haden, 1982; Jamieson et al., 2010; Nosofsky & Zaki, 1998; Woods & Percy, 1974; Zaki & Nosofsky, 2001). The manipulations vary (e.g., Nosofsky and Zaki varied the delay between study and test), but the underlying logic is the same as the computational work; manipulating a single parameter is sufficient to produce dissociations traditionally thought to force a multiple-system perspective. To the extent that a dissociation can be reproduced independent of a memory disorder, the evidence challenges the claim that the selective nature of the dissociation maps on to the selective nature of the physiological damage.

However, the same problem of parsimony plagues the empirical approach. So long as the experimental manipulation can be said to selectively impact the explicit, but not the implicit, memory system, the data are consistent with the multiple-system perspective. Thus, the computational and empirical approaches face the same uncertainties surrounding parsimony. This difficulty illustrates the value of developing converging evidence for a single-system account based on work with different models and experimental manipulations developed in different laboratories.

In the current work, we adopt A Theory of Nonanalytic Association (ATHENA; Chubala & Jamieson, 2013; Chubala, Johns, Jamieson, & Mewhort, 2016; Jamieson & Hauri, 2012; Jamieson & Mewhort, 2010) to examine the claim that empirical dissociations necessitate the multiple-system perspective. ATHENA is based on the classic MINERVA2 model of memory (Hintzman, 1984, 1986, 1988) but adopts holographic reduced representations to represent stimuli in memory (e.g., Jones & Mewhort, 2007; Lewandowsky & Murdock, 1989; Murdock, 1983, 1995; Plate, 1995).

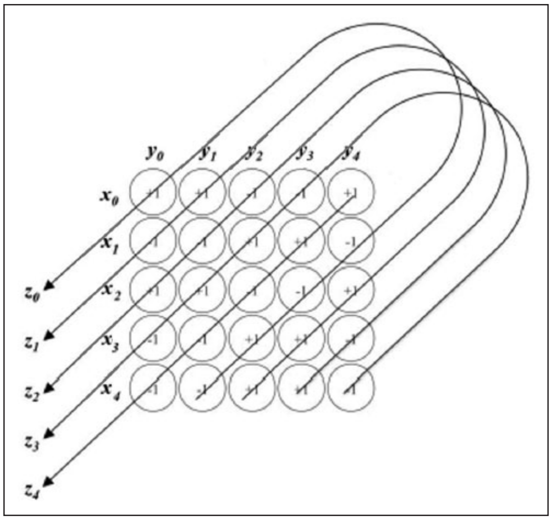

Figure 1. A depiction of circular convolution. Multiplying *x* and *y* produces an *n* by *n* outer product matrix. The arrows demonstrate summation across elements to produce values in the resulting composite vector, z. The method can be thought of as a way to compress the associative outer product matrix z to the same dimensionality as the inputs, *x* and *y*.

Our work builds on Jamieson et al. (2010), who applied MINERVA2 to explain spared classification coupled with impaired recognition in artificial grammar learning (see Knowlton et al., 1992). Whereas Knowlton et al. argued that the dissociation of classification and recognition reflects selective impairment to an explicit memory system, Jamieson et al. argued that the dissociation reflects a general deficit to memory for the studied exemplars. They tested the idea by simulating classification and recognition as a function of how accurately the exemplars were stored in memory. In addition to the computational argument, they also replicated Knowlton et al.'s experiment, simulating amnesic behaviour by manipulating how long each exemplar was studied.

In both cases, impairing memory for studied exemplars affected recognition but not classification, the same dissociation shown by amnesic patients. Jamieson et al. (2010) concluded that the dissociation between classification and recognition is consistent with the perspective that amnesic patients' memory of studied exemplars is poor relative to that of controls and, therefore, the data do not force the multiple-system conclusion.

In the work that follows, we borrow Jamieson et al.'s (2010; see also Graf et al., 1982; Higham et al., 2000; Kinder & Shanks, 2001; Nosofsky & Zaki, 1998; Woods & Percy, 1974; Zaki & Nosofsky, 2001) methods to argue two points. First, we will reproduce additional dissociations using both ATHENA and a manipulation of study time to better examine and evaluate those particular approaches. Second, and more importantly, we will present and discuss the computational and empirical simulations in relation to other accounts to contribute to the growing body of converging evidence in favour of the single-system perspective.

A Theory of Nonanalytic Association

Informally, ATHENA is a framework that describes the representation, storage, and retrieval of experiences. The model assumes that each experience is represented and stored in memory as a unique instance. When a cue is presented to memory, the model retrieves an aggregate of the instances that are similar to the cue. The match between the cue and the resulting aggregate can be used to index both classification and recognition.

Formally, each experience is stored into memory as a unique *trace*. When a *probe* is presented to memory, each trace is activated in parallel as a function of the similarity between the probe and trace. The activation from each trace is summed into a response profile known as the *echo*. The information in the echo is called its *content*. The strength of the echo's activation is called its *intensity*.

Representation and storage

A stimulus is represented as a vector of *n* elements, where each element is randomly sampled from a Gaussian distribution with a mean of zero and a standard deviation of $1/\sqrt{n}$. Association between multiple features of the same stimulus is represented using circular convolution, a vector operation that returns an associative representation between two argument vectors. More concretely, given two vectors, *x* and *y*, circular convolution produces a unique vector of the same dimensionality, *z*

$$z_i = \sum_{j=0}^{n-1} x_{j \bmod n} \times y_{(i-j) \bmod n}$$

where *x* and *y* are item vectors, *z* is the association of *x* and *y*, and *n* is the dimensionality of *x, y*, and *z*. Figure 1 presents an example of the operation on two vectors with five elements. Note that the operation is only used if the stimulus structure involves combining multiple stimulus features into a single representation. We use the operation in Simulation 2 but not Simulation 1.

Encoding is assumed to be imperfect. To implement this idea, each stimulus is stored to a trace in memory with some degree of error as controlled by a parameter *L* that specifies the probability with which each element in a stimulus representation is stored to memory. If an element is not stored to memory, the element is encoded as a zero—a method to represent missing information.[2] Thus, as *L* increases from 0 to 1, the representation of a studied stimulus in memory becomes increasingly complete. In the simulations that follow, this is the critical factor for simulating amnesia, where amnesic patients are assumed to have smaller values of L than controls. Importantly, the model assumes storage of studied items without any redundant secondary storage of implicit information.

Retrieval and decision

When a probe, p, is presented to memory, M, each trace is activated as a function of its similarity to the probe. Similarity, s, between a probe and a trace is given by

$$s_i = \frac{\sum_{j=1}^{n} p_j \times M_{ij}}{\sqrt{\sum_{j=1}^{n} p_j^2}\sqrt{\sum_{j=1}^{n} M_{ij}^2}}$$

where p_j is the jth element of the probe, M_{ij} is the jth element of the ith trace.

The activation of each trace is a non-linear function of its similarity to the probe. Activation, a, is given by

$$a_i = s_i^3$$

The purpose of the non-linearity is to exaggerate the role of similarity. Only very similar traces (i.e., with similarities near +1) make a strong contribution to activation; moderately similar traces make an attenuated contribution, and weakly similar traces make almost no contribution at all.[3]

Once the probe has activated traces in memory, an echo is returned. The echo has two properties: content and intensity.

The echo content is a vector, c, with the same dimensionality n as the traces in memory. It is a sum of the activated traces, where each trace contributes to the sum in proportion to its activation. The echo is computed as

$$c_j = \sum_{i=1}^{m} a_i \times M_{ij}$$

where a_i is the activation of the ith trace and M_{ij} is the jth element of the ith trace. The content of the echo is a vector containing the aggregated information from each of the activated traces.

Echo intensity, f, is the strength of activation elicited by the probe and is equal to the sum of the activation from each trace, thus

$$f = \sum_{i=1}^{m} a_i$$

where a_i is the activation of the ith trace. However, an alternative index of the strength of the echo can be given by computing the cosine similarity between the echo content and the probe given by

$$\cos = \frac{\sum_{j=1}^{n} c_j \times p_j}{\sqrt{\sum_{j=1}^{n} c_j^2}\sqrt{\sum_{j=1}^{n} p_j^2}}$$

where c_j is the jth element of the echo and p_j is the jth element of the probe.

We adopt the cosine similarity metric for two reasons. First, echo intensity sums across all traces. As a result, echo intensity varies as a function of the number of traces in memory. The variation in range makes interpretation of echo intensity difficult, especially when comparing across different study list lengths. In contrast, cosine similarity is bound to a range of −1 to +1, facilitating interpretation. In addition, cosine similarity is an established vector similarity metric in a range of psychological application (e.g., Jones & Mewhort, 2007; Kwantes, 2005; Landauer & Dumais, 1997). In contrast, echo intensity is idiosyncratic to the model.

Finally, the probability of a response (e.g., endorsing an item as "studied") is computed from the cosine similarities according to a standard logistic transformation

$$p(response) = \frac{1}{1 + e^{-\alpha(cos)+\beta}}$$

where cos is the cosine similarity, α is a free scaling parameter, and β is a free decision parameter. The logistic function is a well-established and commonly used method for deriving response probability from a raw signal (e.g., Dienes, 1992; Estes, Campbell, Hatsopoulos, & Hurwitz, 1989; Gluck & Bower, 1988; McClelland & Elman, 1986; McClelland & Rumelhart, 1985).

Overview of simulations and experiments

Jamieson et al. (2010) applied MINERVA2 to explain the dissociation between classification and recognition in artificial grammar learning. We extend the account to two additional dissociations between classification and recognition. First, we simulate and replicate a dissociation between classification and recognition of Posner-Keele dot patterns (Knowlton & Squire, 1993). Second, we simulate and replicate a dissociation between classification and recognition of binary-featured images (Reed, Squire, Patalano, Smith, & Jonides, 1999).

In total, the model has the following five parameters: the number of elements in each vector, L for controls, L for amnesic patients, the scaling parameter (α), and the decision parameter (β). We fixed the number of elements to 100 in all simulations. We conducted simulations at a

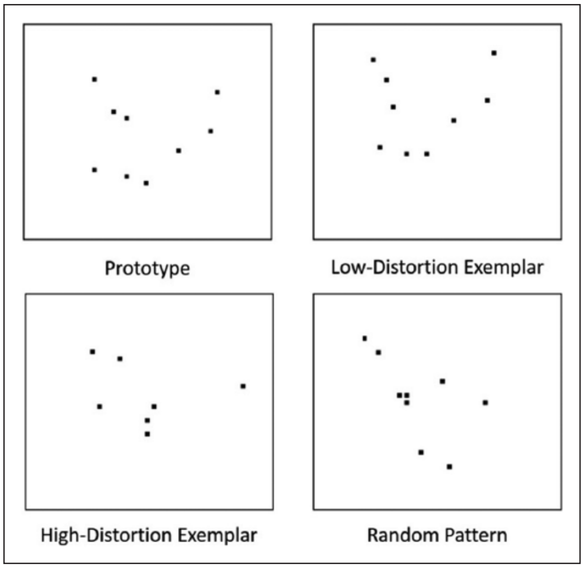

Figure 2. Examples of dot patterns used in the Posner-Keele task, taken from our Experiment 1.

range of values for L, but fixed the values used to simulate controls ($L = .7$) and amnesic patients ($L = .2$) to fit to the original experiments. The particular values were largely arbitrary, and the pattern of results does not change as long as the control and amnesic learning parameters are sufficiently different and the amnesic learning parameter is not greater than the control parameter.

We allowed the scaling and decision parameters to vary freely across simulations. We allowed the decision parameter to vary freely across tasks but not between controls and amnesic patients. We did not allow the scaling parameter to vary across tasks or between controls and amnesic patients. Allowing the decision bias to vary while holding the scaling parameter constant is sensible for two reasons. First, rescaling cosine similarities represents a distortion of the principle basis for decision in the model. Thus, varying the scaling parameter within the same simulation could produce a distorted picture of the model's performance. Second, varying the bias parameter between tasks is consistent with empirical data that show participants adopt a more lax criterion during classification than recognition (e.g., Nosofsky, Little, & James, 2012; Nosofsky & Zaki, 1998; Zaki & Nosofsky, 2001). It is also intuitive; classification requires participants to judge whether a test item is similar to the studied items, whereas recognition requires participants to judge whether a test item is identical to a particular studied item (arguably a more stringent requirement). Because the decision parameter is allowed to vary across two conditions, the model has a total of six parameters, three of which were arbitrarily fixed (the number of vector elements, L for simulated controls, and L for simulated amnesic patients) and three of which were free to vary (the scaling parameter, the decision parameter for recognition, and the decision parameter for classification). We selected the values of the free parameters that maximised the correlation between the simulated and empirical means.

Simulation 1

In a well-cited and important experiment, Knowlton and Squire (1993) had amnesic patients and controls study exemplar dot patterns derived from a prototype. The exemplar patterns were derived by allowing the dots in the prototype to move varying distances with varying probabilities. As the distances and probabilities increased, the exemplars became increasingly distorted (see Figure 2 for examples).

In the classification experiment, participants studied 40 high-distortion exemplars. Then, participants were tested for classification of 84 new patterns: the unstudied prototype (presented four times), 20 unstudied low-distortion exemplars, 20 unstudied high-distortion exemplars, and 40 unstudied random patterns.

In a subsequent recognition experiment, the same participants studied five unique exemplars, each generated from a unique prototype not seen in the classification task, eight times each (a detail that equated the number of study opportunities but not the number of unique exemplars in the preceding classification task). Then, participants were tested for recognition of the five studied exemplars relative to five unstudied exemplars derived from the same prototypes.

Figure 3 presents a reproduction of Knowlton and Squire's (1993) results. As shown, amnesic patients and controls performed similarly in classification. Both groups endorsed the prototype more frequently than the low-distortion exemplars that they endorsed more frequently than the high-distortion exemplars that they endorsed more frequently than the random patterns. However, controls outperformed the amnesic patients at recognition. The dissociation was interpreted as evidence that memory is divided into separate systems—explicit memory (supporting recognition) and implicit memory (supporting classification)—and that amnesia selectively impairs the explicit system.

In rebuttal, Nosofsky and Zaki (1998) examined the dissociation in the framework of the Generalised Context Model (GCM). The GCM is a similarity-driven model of categorisation. Nosofsky and Zaki had participants rate the similarities among Knowlton and Squire's (1993) dot patterns. The model uses those similarity ratings—raised to a sensitivity parameter, c—to evaluate test patterns in the classification and recognition tasks. The model was used to compute each test pattern's similarity to the studied patterns, which were then compared against a criterion to index the probability of endorsing the pattern as either category-belonging (in the case of classification) or

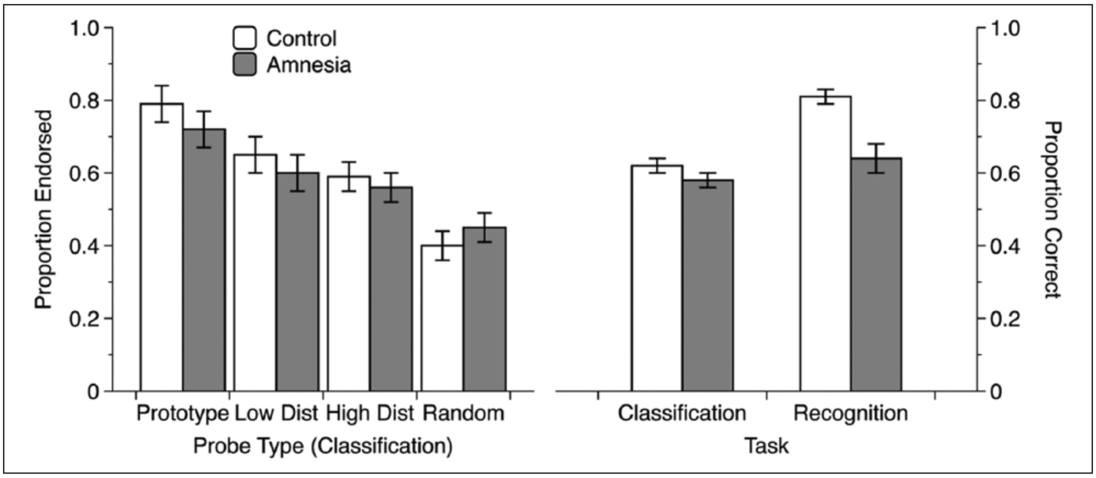

Figure 3. Classification and recognition performance for amnesic patients and control participants, reproduced from Knowlton and Squire (1993). The left panel presents endorsement rates in the classification task. The right panel presents proportion correct in both tasks. Control participants are presented in white bars and amnesic patients are presented in grey bars.

studied (in the case of recognition). Nosofsky and Zaki simulated amnesia by decreasing the sensitivity parameter (i.e., simulated amnesic patients had a smaller value of c relative to controls), successfully producing the dissociation and questioning the necessity of a distinction between explicit and implicit memory.

In Simulation 1, we analysed the same dissociation in ATHENA. Performance in ATHENA and the GCM are both based on inter-item similarity. However, the models calculate similarity differently. The GCM calculates similarity either as distance in a multidimensional space or using people's intuitive judgements of similarity between items. In either case, the similarity structure is static, remaining consistent across all instances of retrieval. ATHENA assumes representational qualities in the items and computes similarity based on those representational qualities at retrieval, individually for each probe. The similarity structure is dynamic, changing at each instance of retrieval. We emphasise, however, that our intention is not to argue in favour of ATHENA over the GCM. Rather, our intention is to replicate the GCM's success using alternative algorithms to highlight the shared principles rather than the particular differences within a more general and converging single-system argument.

Methods

To simulate the classification test, we generated a vector to stand for the prototype. The vector contained 100 elements with each element taking a random number sampled from a Gaussian distribution with mean of zero and standard deviation of $1/\sqrt{n}$. Once generated, we used the prototype to generate 40 high-level distortions of the prototype for the study list and 80 items for the test list: 20 low-level distortions of the prototype, 20 unstudied high-level distortions of the prototype, and 40 random vectors corresponding to the randomly constructed unstudied test patterns. A prototype distortion was generated by copying the prototype and then flipping the sign of each element in the copy with probability d, where d = .15 for a low-distortion item and d = .25 for a high-distortion item.[4] The strategy is consistent with previous work that applied the model to category learning (e.g., Arndt & Hirshman, 1998; Hintzman, 1986). Next, we simulated study by storing all 40 of the study items to the memory matrix at learning rate L.

Next, we simulated classification by computing the echo for each of the 84 test items (i.e., four prototypes, 20 low-level distortions, 20 high-level distortions, and 40 novel patterns). Then, we computed the cosine similarity between the echo and the test item. Finally, we converted the cosine similarities into endorsement probabilities using the logistic function described above. Consistent with standard practice, we conducted 1,000 independent simulations at each of 10 levels of L = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, and 1.0.

To simulate the recognition test, we generated a set of five random prototypes not used in the simulated classification task. We generated two exemplars from each prototype: one for the study list and one for the test list. We simulated study by storing the five items in the study list to a memory matrix eight times, each at learning rate L. Next, we simulated recognition by computing the echo for each of the 10 test patterns (i.e., all of the five studied and five unstudied patterns). Then, we computed the cosine similarity between the echo and test item. Finally, we computed endorsement probabilities by converting the cosine similarity into response probabilities. As with the simulations of classification, we conducted 1,000 independent simulations at each of 10 levels of L = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, and 1.0.

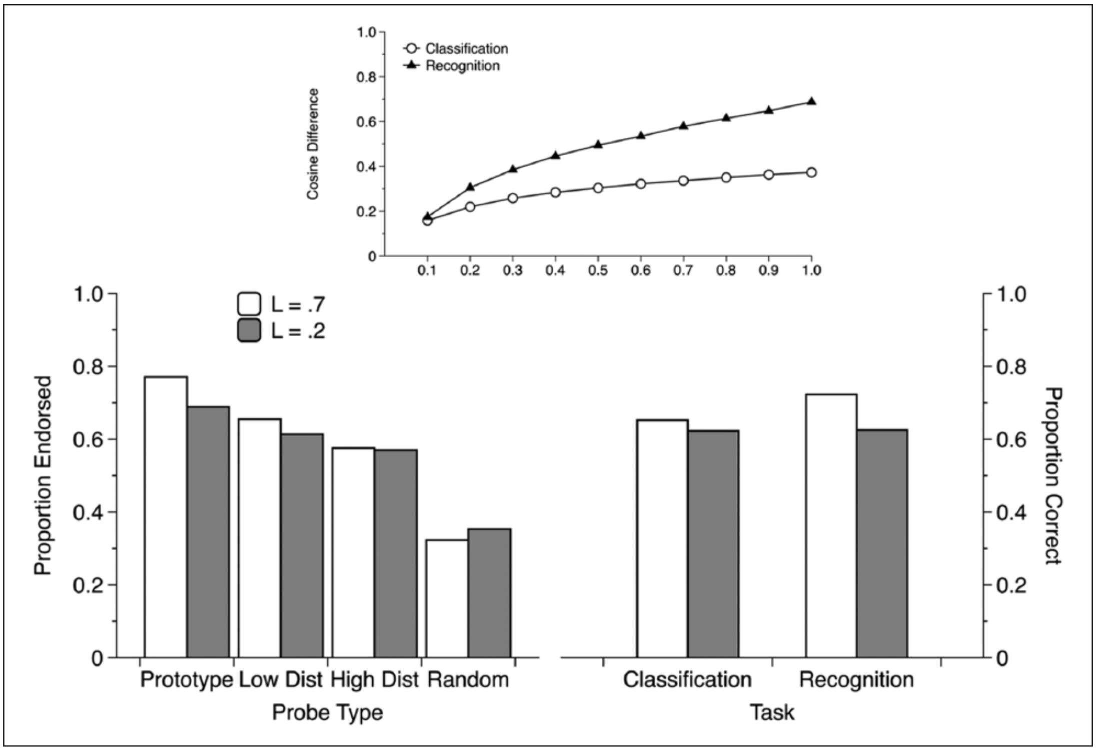

Figure 4. Simulation 1 endorsement rates and proportion correct on Knowlton and Squire's (1993) classification and recognition tasks. The bottom left panel presents the model's proportion of items endorsed in the classification task as a function of *L* and probe type. The bottom right panel presents proportion correct on both tasks as a function of *L*. The top panel presents the difference between cosine similarities for to-be-endorsed and to-be rejected items in both tasks as a function of *L*.

Results

The bottom panel in Figure 4 presents the results from the model where $\alpha = 3.4$, $\beta = 2.0$,[5] and where amnesic patients were simulated with $L = 0.2$, whereas controls were simulated with $L = 0.7$. The left panel presents the model's probability of endorsement on the classification task as a function of L (i.e., amnesic patients and controls) and probe type. The right panel shows overall performance in classification and recognition.

The model fits the empirical data very well. Classification performance shows the standard typicality gradient with little influence of a difference in learning rate. At both $L = .2$ and $L = .7$, the model endorsed the prototype most strongly followed by the low-distortion, high-distortion, then random test patterns. A comparison of simulated and empirical results (see Figure 3) shows a clear correspondence: the Pearson correlation between the two is equal to .94. More importantly, the model predicts the critical pattern of impairment. The change in L from .2 to .7 had a smaller impact on classification performance than on recognition performance.

The top panel in Figure 4 shows how the model accommodates the pattern of results and illustrates the general argument across the full range of L. Increasing L results in improved discrimination on both the classification and recognition tasks. However, discrimination improves faster on recognition than on classification as a function of L. Thus, at any two values of L, the difference in discrimination will be smaller in classification than in recognition.

Discussion of Simulation 1

The simulation shows that ATHENA—an adaptation of a standard theory of memory that assumes memory for studied exemplars without implicit knowledge of an underlying prototype—accommodates the difference between classification and recognition. The model's success questions the necessity of a systems distinction and provides converging evidence on which to argue for the single-system perspective on memory. The analysis also shows that the differences in experimental procedure for measuring classification and recognition (e.g., discrimination of studied from unstudied exemplars versus unstudied from unstudied exemplars) are far from inconsequential and, in fact, a consideration of those differences is important for explaining the nature and form of selective impairments. To further build on that point, we simulated the dissociation empirically. We replicated Knowlton and Squire (1993), simulating amnesic behaviour in healthy participants by presenting the studied exemplars very briefly.

Experiment 1

Experiment 1 provided a simulated replication of Knowlton and Squire (1993). In their experiment, amnesic patients and controls studied dot patterns derived from a prototype. Following study, they sorted new patterns based on category membership. Subsequently, they sorted studied patterns from unstudied patterns. We replicated their procedure using healthy participants. However, half the participants studied each item for 5,000 ms and half the participants studied each item for only 100 ms (i.e., the same technique used by Jamieson, Holmes, & Mewhort, 2010). If a difference in encoding quality, as manipulated by study time, is sufficient to produce the results, classification ought to be similar between groups but the control (5,000 ms) group ought to perform better on recognition.

Methods

Participants. We recruited 80 participants from the University of Manitoba Introduction to Psychology participant pool. Participants were randomly assigned to four groups formed by crossing two factors: encoding time (5,000 ms, 100 ms) and task (recognition, classification). All participants received credit towards a course requirement for participation.

Materials and apparatus. The experiment was administered on eight personal computers. Each computer was equipped with a 22-inch monitor, a standard mouse, and a standard keyboard.

We constructed prototypes by randomly placing nine squares on a white, 420 × 420 pixel background. Each square was black, filled, and 6 pixels by 6 pixels. In the classification task, we generated one random prototype for each participant by placing the nine squares in random locations. Then, we derived exemplars using the statistical techniques described in Posner, Goldsmith, and Welton (1967). We generated 20 high-distortion study patterns. We also generated 20 low-distortion test patterns, 20 new high-distortion test patterns, and 40 random test patterns.

In the recognition task, we generated five prototypes for each participant. Next, we generated two high-distortion exemplars from each prototype. One exemplar from each prototype served as a study pattern. The remaining patterns served as unstudied test patterns.

Procedure. Participants sat at individual computers and were instructed that they would study patterns of dots for a memory test. Participants viewed the dot patterns one at a time in the centre of the screen. Participants in the control group viewed each pattern for 5,000 ms, consistent with Knowlton and Squire's (1993) procedure. Participants in the simulated amnesia group viewed each pattern for 100 ms. The screen was cleared for 250 ms after each presentation. The classification group viewed 40 high-distortion exemplars presented in random order. The recognition group viewed five high-distortion exemplars eight times each in pseudo-random order (i.e., in eight randomised blocks; the last pattern in a block could not be the first pattern in the next block).

Following study, participants in the classification group were informed that the studied patterns belong to the same category and that they were to sort new patterns based on category membership. Participants evaluated 84 test exemplars—the prototype (four times), 20 low-distortion test exemplars, 20 high-distortion test exemplars, and 40 random test patterns—one at a time in pseudo-random order (the instances of the prototype could not appear on consecutive trials). The dot patterns were presented in the centre of the screen. Participants responded by clicking one of two alternatives labelled "Same category" and "Not same category" with the computer mouse. Once the participants had responded, the screen cleared for 250 ms and the next pattern appeared. This process repeated until the participants had provided a response to all 84 patterns.

The recognition group followed a similar procedure. The participants evaluated 10 patterns—five studied and five unstudied—in random order. Participants responded in the same way as the classification group, but the response alternatives were labelled "Old" and "New." The screen cleared for 250 ms after the participants responded. This process repeated until the participants had provided a response to all 10 patterns.

Results

Figure 5 presents participants' performance. The left panel presents the proportion of patterns endorsed as "Same category" in the classification task. The right panel presents the proportion correct for both the classification and recognition tasks. In classification, endorsing the prototype and the distortions was scored as correct and endorsing the random patterns was scored as incorrect. In recognition, endorsing the studied patterns was scored as correct and endorsing the unstudied patterns was scored as incorrect.

In the classification task, participants showed a standard typicality gradient. They were most likely to endorse the prototype, followed by the low-distortion test exemplars, followed by the high-distortion test exemplars, followed finally by the random test patterns. There was very little difference between the patterns of endorsement as a function of study time. The endorsement rates in our experiment are a close match to the rates in Knowlton and Squire's (1993) experiment. For brevity, we do not report any formal analysis of performance on either the classification or recognition task alone. Instead, we turn to the more crucial comparison between tasks.

The right panel of Figure 5 illustrates the key result. The left bars show that there is very little difference in the two

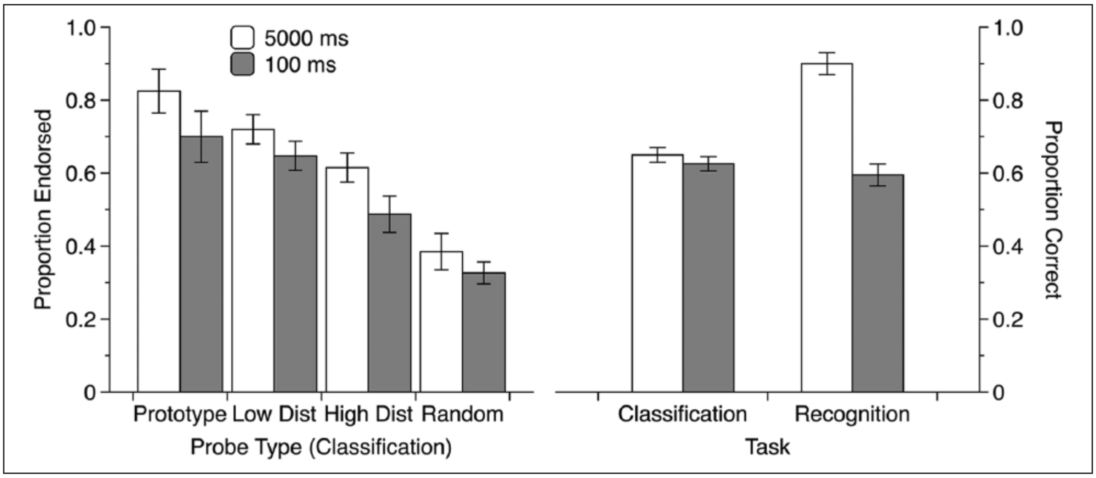

Figure 5. Endorsement probabilities as a function of encoding condition for both tasks in Experiment 1. The left panel presents the proportion of patterns endorsed as "Same category" in classification. The right panel presents the proportion of correct trials for both classification and recognition. The error bars represent standard error.

study-time groups' classification performance. In contrast, the right bars show that the 5,000 ms group performed recognition more accurately than the 100 ms group. This pattern is the same critical pattern reported by Knowlton and Squire (1993). However, in contrast to Knowlton and Squire who observed the pattern as a function of amnesia, we observed the pattern as a function of study time.

To evaluate the data, we analysed the proportion of correct trials in a 2 × 2 between-subjects analysis of variance with study time (5,000 ms, 500 ms) and task (recognition, classification) as factors. Both main effects were significant. Accuracy was superior following 5,000 ms of study time, $F(1, 76) = 40.90$, $p < .001$, and in the recognition task, $F(1, 76) = 18.15$, $p < .001$. More importantly, there was an interaction between encoding condition and task, $F(1, 76) = 29.68$, $p < .001$, such that the difference in accuracy was much greater in recognition than in classification.

There appears to be a bias to endorse all types of patterns in the 5,000 ms encoding condition. This pattern is inconsistent with the patterns observed in previous experiments as well as our Simulation 1. Statistically, a bias to endorse would be reflected by a main effect of encoding condition, which did not reach significance, $F(1, 38) = 3.985$, $p = .0531$. However, the effect is very nearly significant, and does suggest a potential response bias not observed in previous work. Nonetheless, the interaction pattern is the key observation. This pattern is a clear replication of the critical results reported by Knowlton and Squire (1993) and that they interpreted as evidence for a selective rather than a general memory impairment.

Discussion of Experiment 1

Experiment 1 provides a close match to the preceding simulation, suggesting that the pattern of impairments in classification and recognition in amnesia is consistent with what would be expected from a general rather than selective encoding deficit. The computational and empirical analyses approach the results from a common set of fundamental principles.

Knowlton and Squire's (1993) results are only one example of a dissociation between recognition and classification observed in amnesia. If a single-system model is to be taken seriously, it ought to be able to account for a variety of dissociations across a variety of tasks. Jamieson et al. (2010) applied the analysis to artificial grammar learning. We have applied the analysis to classification of Posner-Keele dot patterns. In the following, we apply the analysis to binary-featured line drawings. All of the original experiments produce a dissociation between classification and recognition. However, the category structures in the classification tasks differ in important ways. Categories in artificial grammar tasks are based on rules governing sequential dependencies. Categories in Knowlton and Squire's (1993) experiment are based on derivation from a prototype pattern. In the following, categories are based on matching and mismatching features. Accounting for the dissociations in all three structures makes a stronger case than any independent demonstration.

Classification of binary-featured category items

Reed et al. (1999) tested amnesic patients' and control participants' abilities to learn and categorise line drawings of animals. Each of the animals was composed of nine features and each feature could take one of two values (see Figure 6 for examples). The similarity between any two animals can be measured by counting the number of matching versus mismatching features.

Despite a difference in materials, Reed et al. (1999) conducted an experiment that was very similar to the one conducted by Knowlton and Squire (1993). In the experiment, they constructed a prototype animal composed of a randomly selected value for each of the nine features. Then, they used that prototype to construct exemplars of varying difference. At study, participants studied 20 low-distortion exemplars twice each: the exemplars differed from the prototype on either one or two features. Following study, participants were told that the animals belonged to a species called "*Peggle*" and were asked to sort unstudied *Peggles* from unstudied *non-Peggles*. Participants classified unstudied drawings: the prototype, new low-distortion, moderate-distortion exemplars that differed from the prototype by four or five features, high-distortion exemplars that differed from the prototype by seven or eight features, and the *anti-prototype* that differed on all nine features from the prototype.

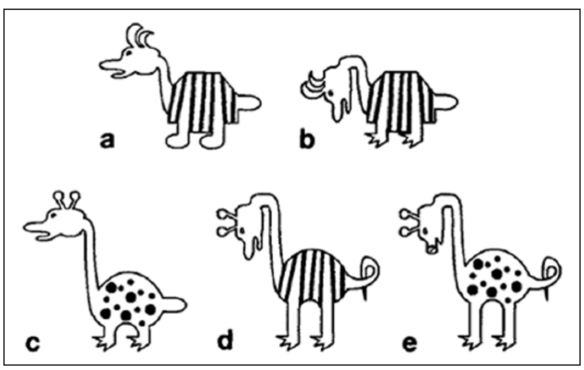

Figure 6. Examples of animal line drawings from Reed, Squire, Patalano, Smith, and Jonides (1999). Animal a is the prototype, b is a low-distortion exemplar, c is a moderate-distortion exemplar, d is a high-distortion exemplar, and e is the anti-prototype.

Figure 7 presents a reproduction of Reed et al.'s (1999) results. The left panel presents endorsement rates on the classification task. The numbers on the x-axis indicate the number of features that matched those in the prototype. As shown, controls showed the standard typicality gradient. Participants were most likely to endorse the prototype, followed by the low-distortion exemplars, followed by the moderate-distortion exemplars, followed by the high-distortion exemplars, followed by the anti-prototype. Amnesic patients showed a similar, albeit shallower, gradient relative to the controls.

Following the classification task, both groups were also given a cued-recall task. Participants were cued with the name of each feature (e.g., head, tail, feet) and asked to describe the two values of each feature. For example, if cued with "body markings," a participant would receive one point for recalling "striped" or "spotted" and two points for recalling both.

The right panel of Figure 7 presents overall accuracy on the classification and the cued recall tasks. As shown, although the groups classified drawings similarly, the amnesic patients were less accurate on the cued recall task. The difference between classification and a cued-recall replicates the result from Knowlton and Squire (1993): amnesia affected participants' performance in an explicit but not an implicit memory task.

Zaki and Nosofsky (2001) applied the GCM to Reed et al.'s (1999) classification task. The general principles of the model remained the same. The studied exemplars were stored to memory. At test, classification decisions were based on the global similarity between the probe and the exemplars.

The GCM fit the classification data well, producing a standard typicality gradient and little difference between simulated amnesic patients and simulated controls. The model's success provides evidence that the single-system

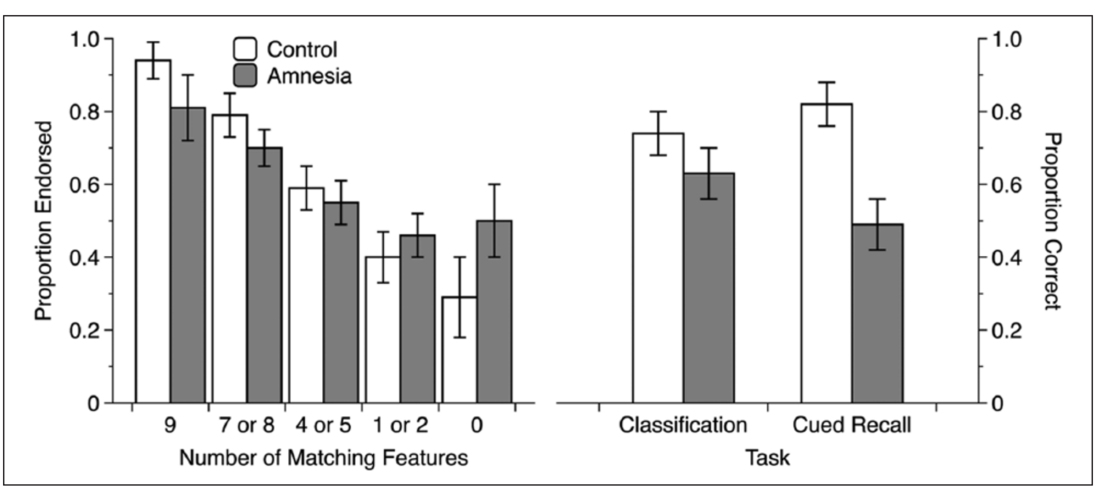

Figure 7. Reproduced results from Reed, Squire, Patalano, Smith, and Jonides's (1999) classification and cued recall tasks. The left panel presents endorsement rates in classification for control participants and amnesic patients as a function of the number of features matching the prototype. The right panel presents percent correct on the classification and cued recall tasks.

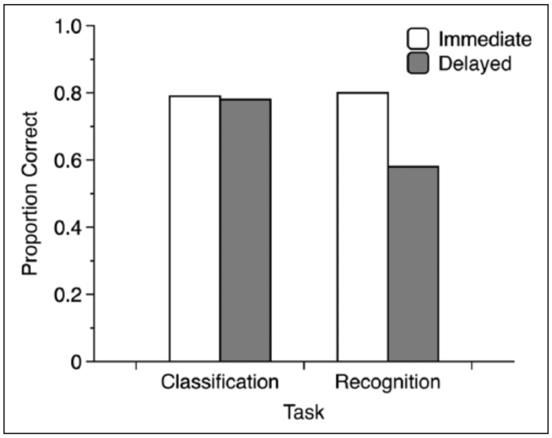

Figure 8. Reproduction of Zaki and Nosofsky (2001), presenting proportion of correct trials for classification and recognition after no delay and a week-long delay. No error bars were presented in the original article.

perspective accommodates patterns of impairment previously thought to force a multiple-system account.

However, Zaki and Nosofsky (2001) did not simulate Reed et al.'s (1999) cued recall task. Their decision was likely due to the fact that the cued recall task is not well suited to a computational analysis. The GCM does not possess representations of heads or bodies and thus cannot answer questions about heads or bodies. Nonetheless, Zaki and Nosofsky's analysis is incomplete.

In lieu of the cued recall task, Zaki and Nosofsky (2001) replicated Reed et al.'s (1999) classification task and a recognition task using healthy participants. They simulated amnesia by introducing a delay between study and test. Figure 8 presents a reproduction of their results. The left bars present accuracy on the classification task and the right bars present accuracy on the recognition task. As shown, there is little difference in classification, but the delayed group was less accurate in recognition. However, Zaki and Nosofsky did not model the data. The goal of Simulation 2 was to provide a full computational analysis of the combined results.

Simulation 2

We evaluated a single-system account of the difference between classification and recognition of binary-featured animal drawings in ATHENA. Similar to Zaki and Nosofsky (2001), we do not model Reed et al.'s (1999) cued-recall task. Instead, we model Zaki and Nosofsky's recognition data. The model extends the GCM's analysis by simulating the influence of L on both tasks, permitting a full analysis of the results.

Methods

To represent drawing features, we generated 18 random 100 element vectors with each element taking a random number from a normal distribution with a mean of zero and a standard deviation of $1/\sqrt{n}$. Each vector represented one of two values for each of the nine binary features. We represented animal-drawings by taking the circular convolution of the vectors representing the relevant features.[6]

For each simulation, we randomly selected values of each feature. The circular convolution of those vectors represented the prototype. We generated 20 low-distortion study exemplars by replacing either one or two of the prototypical features. We stored each exemplar to a memory matrix twice each (i.e., to match the experimental details of Reed et al., 1999). We then generated 20 additional low-distortion exemplars using the same methods to serve as test items. We generated 20 moderate-distortion exemplars by replacing either four or five of the prototypical features. We generated 20 high-distortion exemplars by replacing either seven or eight of the prototypical features. Finally, we generated an anti-prototype by replacing all nine of the prototypical features.

Next, to simulate classification, we computed the echo for each of the 96 test items (i.e., 12 prototypes, 20 low-level distortions, 20 moderate-level distortions, 20 high-level distortions, and 12 anti-prototypes). Then, we computed the cosine similarity between the echo and the test item. We converted the cosine similarities into response probabilities using the same logistic transformation as Simulation 1. We conducted 1,000 simulations at each of 10 levels of $L = 0.1$, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, and 1.0

To simulate the recognition test, we generated five study items and five unstudied items by convolving random selections of features (i.e., there were no prototypes), consistent with Zaki and Nosofsky's (2001) methods. We stored each study item to a memory matrix eight times each. Next, we simulated recognition by computing the echo for each of the 10 test items (i.e., all of the five studied and five unstudied items). Then, we computed the cosine similarity between echo and test item, and computed recognition performance using the same logistic function as the classification task. We conducted 1,000 simulations at each of 10 levels of $L = 0.1$, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, and 1.0

Results

Figure 9 presents the results from the model where $\alpha = 9.8$, $\beta = 5.0$ for classification, and $\beta = 9.3$ for recognition. Consistent with previous work (e.g., Nosofsky & Zaki, 1988; Nosofsky, Little, & James, 2012), the model adopted a lower decision criterion in classification than in recognition. Amnesic patients were simulated with $L = 0.2$, whereas controls were simulated with $L = 0.7$. The left panel presents the model's endorsement probabilities in the classification task. The right panel presents the model's performance on the classification and recognition tasks. The top panel inset presents the cosine differences for both tasks as a function of L.

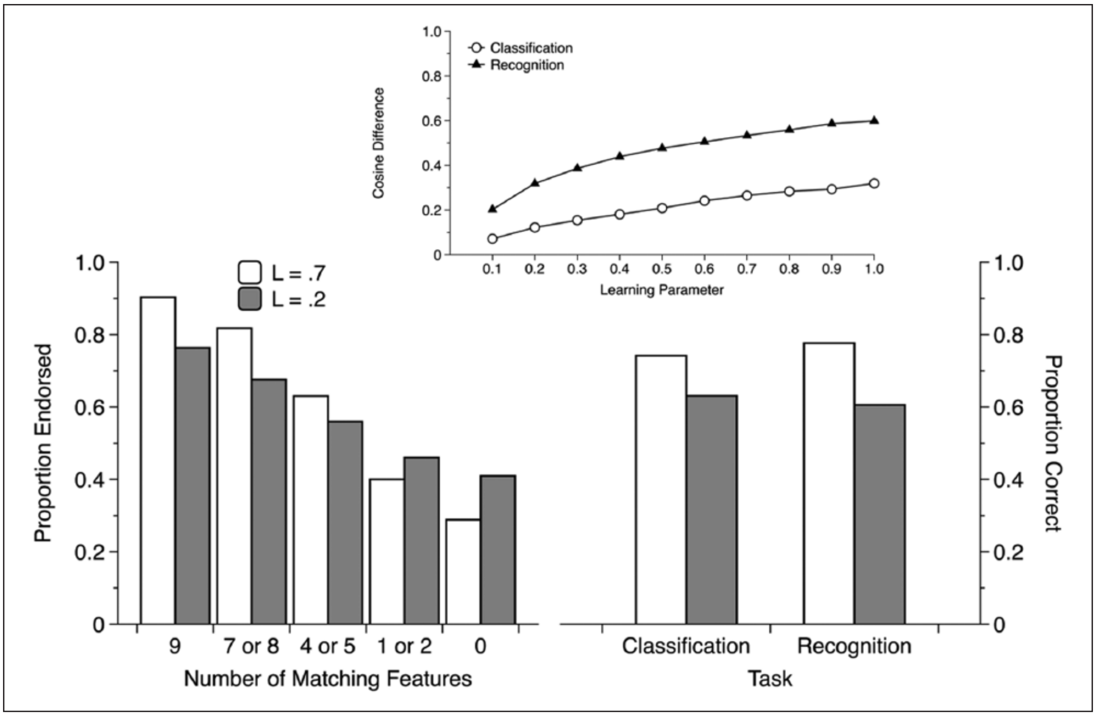

Figure 9. Simulation 2 endorsement rates and proportion correct on Reed, Squire, Patalano, Smith, and Jonides's (1999) classification task and Zaki and Nosofsky's (2001) recognition task. The bottom left panel presents the model's proportion of items endorsed in the classification task as a function of *L* and probe type. The bottom right panel presents proportion correct on both tasks as a function of *L*. The top panel presents the difference between cosine similarities for to-be-endorsed and to-be rejected items in both tasks as a function of *L*.

The model fits the empirical data well. Classification performance shows the standard typicality gradient with little influence of a difference in learning rate. At both levels of L, the model endorsed the prototype most strongly, followed by the low-distortion items, moderate-distortion items, high-level distortion items, and finally the anti-prototype. More importantly, the model predicts the difference in impairment. The change in L had a smaller influence on classification performance (.11) than on recognition performance (.17). The model corresponds closely with the empirical means (r = .98). Although the dissociation is not as dramatic as in the empirical data, the model reliably produces the crucial qualitative pattern thought to support the multiple-system perspective.

Our implementation differs from Zaki and Nosofsky's (2001) implementation in an important way. Zaki and Nosofsky assume that participants selectively attend to only a few of the binary features. Their model analysis supports this assumption, showing that allowing the model to attend to additional features does not result in a significant improvement to performance. However, the assumption has been debated.[7] In addition, the assumption raises questions that are difficult to answer. For example, Zaki and Nosofsky state that attention to a limited number of features is sufficient to support classification but not recognition. However, given that learning in each task is incidental (i.e., participants are simply asked to attend to the stimuli), it is unclear why participants would attend to limited features in classification but all features in recognition. Our analysis demonstrates that the assumption, whether true or false, is not necessary to produce the qualitative pattern of the dissociation, albeit a less dramatic one.

In summary, the model accommodates a difference between classification and recognition that is consistent with the difference that Reed et al. (1999) reported between classification and cued-recall. Importantly, the results can be explained based on memory of just the studied exemplars and without recourse to a systems distinction. In conclusion, the differential pattern of impairment in classification and recognition might be consistent with a multiple-system account of memory, but it is also consistent with a single-system account. To the extent that the data are consistent with both, the data cannot force a conclusion in favour of the multiple-system perspective.

Experiment 2

Experiment 2 provided a simulated replication of Reed et al. (1999). In their experiment, amnesic patients and control participants studied line drawings of cartoon animals. Following study, they sorted new drawings based on category membership. Zaki and Nosofsky (2001) tested

Table 1. The nine features and their possible binary values in Experiment 2. The first set of feature values were the prototypical values for half the participants and the second set of feature values were the prototypical values for half the participants.

Feature	Value 1	Value 2
Outer Shape	Square	Circle
Outer Shape Colour	Red	Blue
Outer Shape Fill	Filled	Border Only
Inner Shape	Triangle	Diamond
Inner Shape Colour	Green	Yellow
Inner Shape Fill	Filled	Border Only
Symbol	Star	Sun
Symbol Location	Right	Left
Line Direction	Bottom Left to Top Right	Bottom Right to Top Left

participants in a recognition task with the same stimuli. We replicate the two sets of data, which form a selective impairment, using healthy participants by manipulating study time. However, because their stimulus set was not available, we instantiated the categories differently.

Methods

Participants. We recruited 80 participants from the University of Manitoba Introduction to Psychology participant pool. Participants were randomly assigned to four groups formed by crossing two factors: encoding time (5,000 ms, 100 ms) and task (recognition, classification). All participants were given credit towards a course requirement for participation.

Materials and apparatus. The apparatus was identical to those in Experiment 1. All stimuli were patterns of shapes and colours consisting of nine features that could each take on one of two values. Table 1 presents the nine features and their possible values.

For the classification task, we constructed a prototype for each participant by combining the feature values for either column 2 or 3 of Table 1 (each for half of the participants in each group). We constructed an anti-prototype using the other set of values. We then constructed 80 more patterns. There were 20 low-distortion study exemplars and 20 low-distortion test items, which differed on either one or two randomly selected features. There were also 20 neutral test items, which differed on either four or five features, and 20 high-distortion test items, which differed on either seven of eight features. Figure 10 presents an example from each category for a single participant.

We also constructed sets of random features for the recognition task. For each participant, we constructed 10 patterns; five served as study patterns and five served as unstudied test patterns. This stimulus structure exactly matches Zaki and Nosofsky's (2001) stimulus structure.

Procedure. The procedure for Experiment 2 was very similar to that of Experiment 1. Participants sat at individual computers for a memory test. They viewed each of the study patterns one at a time in the centre of the screen. Participants in the control group viewed each pattern for 5,000 ms, and participants in the simulated amnesia group viewed each pattern for 100 ms, each followed by a blank screen for 250 ms. The classification group viewed 20 low-distortion study exemplars twice each in pseudo-random order (i.e., two randomised blocks). The recognition group viewed five random study patterns eight times each in pseudo-random order (i.e., eight randomised blocks).

Following study, participants in the classification group were given the same instructions as in Experiment 1. Participants evaluated 96 test items—the prototype (12 times), 20 low-distortion test items, 20 moderate-distortion test items, 20 high-distortion test items, and the anti-prototype (12 times)—one at a time in pseudo-random order (neither the prototype nor anti-prototype could appear on consecutive trials). The images were presented in the centre of the screen. Participants responded by clicking one of two alternatives labelled "Same category" and "Not same category." Once the participants had responded, the screen cleared for 250 ms and the next image appeared. This process repeated until the participants had provided a response to all 96 patterns.

The recognition group followed a similar procedure. The participants evaluated 10 images—five studied and five unstudied—in random order. Participants responded in the same way as the classification group, but the response alternatives were labelled "Old" and "New." The screen cleared for 250 ms after the participants responded. This process repeated until the participants had provided a response to all 10 images.

Results

Figure 11 presents participants' performance. The left panel presents the proportion of images endorsed as "Same category" in the classification task. The right panel presents proportion of trials correct in both tasks. In the classification task, endorsing the prototype and the low-distortion test items was scored as correct and endorsing the

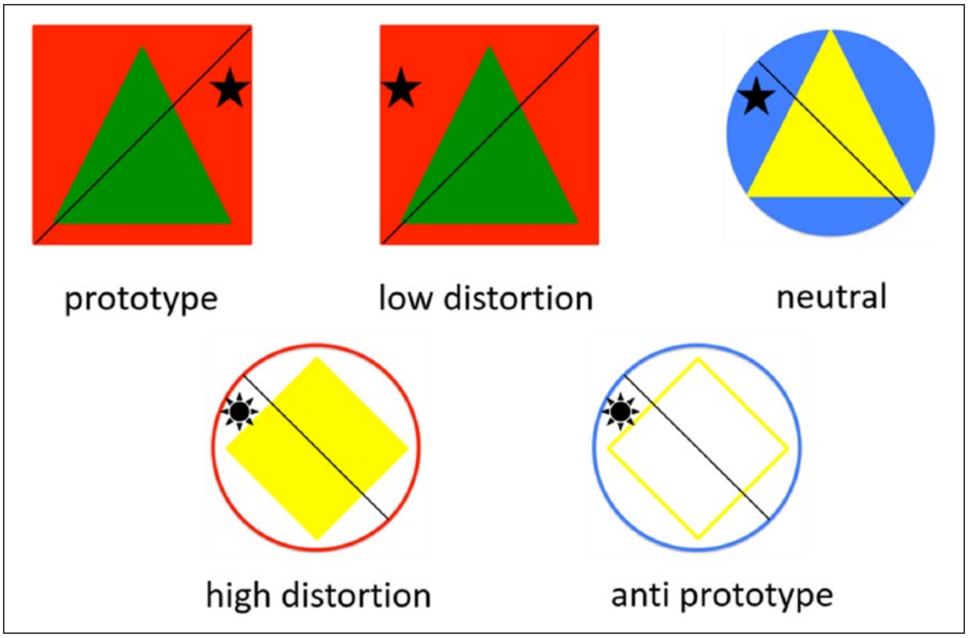

Figure 10. Examples of the binary-featured patterns used in Experiment 2.

anti-prototype and the high-distortion test items was scored as incorrect. The moderate-distortion items were excluded from analysis, consistent with Reed et al.'s (1999) original method. In the recognition task, endorsing the studied items was scored as correct and endorsing the unstudied items was scored as incorrect.

The left panel of Figure 11 shows that participants, as in Experiment 1, endorsed stimuli following a standard typicality gradient. The gradient is dramatic, ranging from nearly uniform endorsement of the prototype (.96) to nearly uniform rejection of the anti-prototype (.02). Between those extremes, participants were more likely to endorse patterns that differed from the prototype on fewer features. The endorsement rates differ very little as a function of study time.

The right panel of Figure 11 presents the crucial dissociation. As in Experiment 1, there is little difference in classification performance as a function of study time. In contrast, the 5,000 ms group was better at the recognition task than the 100 ms group. This is the same pattern reported by Reed et al. (1999). However, we observed the pattern as a function of study time rather than amnesia.

To evaluate the data, we analysed the proportion of correct trials in a 2 × 2 between-subjects analysis of variance with study time (5,000 ms, 100 ms) and task (classification, recognition) as factors. Both main effects were significant. Participants in the 5,000 ms group were more accurate than participants in the 100 ms group, $F(1, 76) = 18.42$, $p < .001$, and participants in the recognition group were more accurate than participants in the classification group, $F(1, 76) = 62.70$, $p < .001$. More importantly, there was an interaction between study time and task, $F(1, 76) = 5.53$, $p = .021$, such that the difference in accuracy was smaller in classification (.05) than in recognition (.17). This pattern, although not of the same magnitude as Experiment 1, nevertheless provides another statistically unambiguous empirical pattern in which study time had a greater impact on recognition than on classification.

Discussion of Experiment 2

Using stimuli with the same category structure as the stimuli used by Reed et al. (1999), Experiment 2 produced another empirical simulation of a pattern of selective impairment typically interpreted as evidence for multiple systems. The consistency between Experiments 1 and 2 is noteworthy. In both classification tasks, participants are required to respond to the structure of the category. In this case, decreasing encoding time only results in a small decrease in performance. However, in both recognition tasks, participants are required to respond to individual items. Because the category structure is orthogonal to the division between studied and unstudied items, the category structure does not provide any useful information. In this case, decreasing encoding time results in a large decrease in performance.

General discussion

By tradition, memory is conceived as a set of interconnected parts with separate systems for implicit and explicit learning. Selective impairment of explicit relative to implicit learning following amnesia has served as principal evidence for the division where amnesia represents a

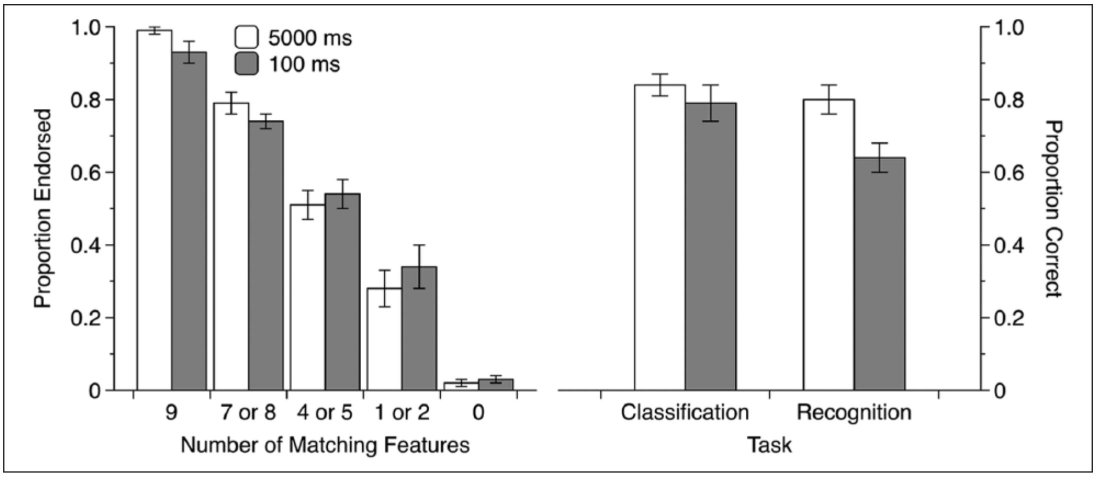

Figure 11. Endorsement probabilities as a function of encoding condition for both tasks in Experiment 2. The left panel presents the proportion of patterns endorsed as "Same category" in classification. The right panel presents the proportion of patterns endorsed as "Old" in recognition.

selective impairment of the explicit system (e.g., Schacter & Tulving, 1994).

In contrast to that tradition, a number of research laboratories have argued that memory is singular and that dissociations of performance in explicit and implicit tasks are consistent with that perspective. In the computational domain, Nosofsky and Zaki (1998; see also Zaki & Nosofsky, 2001) used the GCM to explain selective sparing of classification coupled with recognition impairment in amnesia, Malmberg et al. (2004) made the same argument using the Retrieving Effectively from Memory (REM) model of memory, Higham et al. (2000) made the same argument using a principal component analysis (PCA) autoassociative neural network model, Jamieson et al. (2010) made a similar point using the MINERVA2 model of memory, Benjamin (2010) made the same point in the Density of Representations Yields Age-Related Deficits (DRYAD) model, and Kinder and Shanks (2001, 2003) made a similar point using a Simple Recurrent Network. Nosofsky, Little, and James (2012) reinforced the computational analysis with neuroimaging data (see also Gureckis, James, & Nosofsky, 2011). Cowell, Bussey, and Saksida (2010) made the same argument in the nonhuman animal literature using a hierarchical neural network model (the Extended Perceptual Mnemonic/Feature Conjunction [PMFC] model). In the empirical domain, Zaki (2004) presented a meta-analysis that contradicted conventional wisdom and revealed that classification is less affected but not completely spared in amnesia. Taken together, the database presents a growing body of converging computational, neuropsychological, and empirical evidence that challenges the claim that a selective impairment in explicit memory performance in amnesia has the logical force needed to reject a single-system account.

Our analysis contributes to this growing body of converging evidence. Our computational analysis presents the theoretical argument. Our empirical analysis reinforces the case. Based on our results, it is clear that poor memory of studied exemplars begets poor recognition coupled with spared classification. The substantive argument about a single-system account has been argued before. Our current investigation makes three primary contributions to that argument. First, our computational analysis solves criticisms that have been made against alternative single-system models such as the GCM, which has been criticised on the basis of its representational assumptions (e.g., Smith & Minda, 2001).[8] Second, our empirical analysis begins to clarify the relationship between encoding and interference in amnesia, a point we expand on below. Third and most importantly, we illustrate that a common conclusion can be drawn from multiple implementations of a more general set of principles, a point we discuss throughout the remaining section.

Our argument relies on parsimony and, although parsimony can be a slippery concept in psychological science, it is clear to us that a single-system account is simpler than a multiple-system account: a multiple-system account that includes both an implicit and explicit system (and possibly others) is a superset of a theory that includes only the explicit system. Admittedly, absence of evidence is not evidence of absence. But, if a single-system theory is sufficient to explain the data, a multiple-system theory represents an unnecessary and unmotivated expansion. But, how then does memory theory predict spared classification when memory for studied exemplars is poor?

Modern memory theories assume that retrieval is parallel. In our theory, a probe activates all traces in memory and the information retrieved is a sum of the activated traces, where each trace contributes to the sum in proportion to its activation. The failure to accomplish recognition given poor memory of studied exemplars is straightforward: if an item is stored poorly in memory, it is hard to determine whether the item was studied. The success of classification, however, requires consideration of parallel retrieval of poorly encoded category exemplars. Category exemplars share features with one another by virtue of belonging to the same category. Because the individual traces in memory share features and retrieval is parallel, a category probe retrieves the category members in memory and, by a process of redintegration, as represented in our model by summing over traces in the echo, reconstructs category-level information. The idea is consistent with constructive accounts of memory (see Arndt & Hirshman, 1998, for a more complete discussion of false memory in exemplar models). In conclusion, assuming that amnesic patients remember studied exemplars worse than healthy controls and that retrieval from memory is parallel, one should expect spared classification in the face of compromised recognition.

We are impressed by the growing body of converging evidence that data once thought to force a multiple-system conclusion is, in fact, consistent with a single-system explanation. However, a sceptic of the single-system perspective can point to the differences in the computational instantiations of the single-system approach: "Get the single-system theory's house in order and then come back to the single- versus multiple-system debate." There is wisdom in that rebuttal. But, it misses a larger point. The fact that so many different expressions of memory converge on a common conclusion suggests that the single-system perspective is not conditional on particular details of how the single-system theory is instantiated. For example, our account assumes that amnesia can be understood as poor memory of studied exemplars (see also Benjamin, 2010; Jamieson et al., 2010), Nosofsky and Zaki (1998; see also Zaki & Nosofsky, 2001) assume that amnesia can be understood as a relaxed retrieval gradient, Kinder and Shanks (2001, 2002) assume that amnesia can be understand as slowed learning, Malmberg et al. (2004) assume that amnesia reflects inaccurate encoding of exemplars, and Cowell et al. (2010) assume amnesia reflects an inability to form complex feature conjunctions. Looking across the different accounts, the expression for memory differs. But, the underlying nature of the deficit is consistent. Poor memory for a studied list (i.e., amnesia), independent of how it is produced, impairs explicit performance more than implicit performance. It would be an error to focus on the differences between memory theories without also focusing on the overwhelming similarities among the theories and the fact that they all arrive at the same conclusion.

Our experiments follow from a general strategy in previous work to examine the nature of deficits in amnesia. For example, Graf et al. (1982) conducted experiments with healthy participants who studied training words deeply or shallowly. Their analysis confirmed the standard dissociation between explicit and implicit memory performance. Participants who studied training words shallowly showed impaired recall but equivalent stem completion performance compared to participants who studied training words deeply. Based on the result, Graf et al. concluded that amnesia is a processing impairment—as if amnesic patients encode shallowly whereas healthy individuals encode deeply. Zaki and Nosofsky (2001) adopted the strategy to show that classification is spared and recognition is not after a long study-test delay (see also Nosofsky & Zaki, 1998; Woods & Percy, 1974). Our experiments use the strategy as well: participants studied items quickly or slowly (see Benjamin, Diaz, Matzen, & Johnson, 2012; Jamieson et al., 2010). As in other work, our manipulation produced the dissociation between performance on an explicit versus implicit memory task. Based on those results, we concluded that poor memory of studied exemplars is sufficient to produce the dissociation.

One might argue that poor encoding, as manipulated by short presentation time, is an inaccurate description of organic amnesia. For example, amnesic patients perform well if memory is tested immediately after the presentation of a single item (e.g., Mayes, Downes, Shoqeirat, Hall, & Sagar, 1993). Thus it appears that amnesic patients encode well but forget rapidly. Clearly, neither short presentation time nor delay between study and test are perfect descriptions of amnesia. Nonetheless, Nosofsky and Zaki's (1998) delay manipulation might at least be a better proxy for amnesia.

However, the delay interpretation of amnesia also has weaknesses. It seems unlikely that Nosofsky and Zaki's (1998) delay manipulation caused forgetting by a passive decay mechanism—the concept has been famously critiqued in relation to long-term memory (McGeoch, 1932) and more recently in the context of sensory and working memory (see Nairne, 2003 for a review). It is more plausible that participants in their delay condition suffered from additional interference from the intervening events between study and test. To our knowledge, there are no empirical demonstrations that amnesic patients suffer from more interference per se. However, our manipulation and model analysis provide a potential explanation for why amnesic patients might suffer more severely from interference.

On average, when ATHENA's vector representations are sparse (i.e., contain many zeroes), they are drawn closer to one another in geometric space (i.e., they are made less discriminable from one another). As a result, items in memory are made more similar and therefore they interfere more strongly with the retrieval of any particular item. Poor encoding results in greater interference during retrieval. Why, then, can amnesic patients remember very recent items? Our model analysis cannot provide an answer. However, distinctiveness-based models of memory (e.g., SIMPLE; Brown, Neath, & Chater, 2007) suggest that recently encountered stimuli suffer far less interference. To be clear, we are generally agnostic about whether amnesia is best described as poor encoding. However, it seems at least to be a plausible candidate explanation, and our analysis, in conjunction with principles from other models, aids in clarifying the relationship between encoding and interference. More importantly, and mirroring the critical conclusion we draw from our model analysis, the fact that many different manipulations produce the same pattern of results is the primary lesson.

Despite the success of our computational and empirical simulations, other data poses a unique challenge to our analysis and the analyses of others. For example, patient E.P. shows no ability to recognise above chance levels. Nonetheless, E.P. shows near-normal classification in the paradigms we have covered (Squire & Knowlton, 1995). The pattern is problematic for the exemplar framework. For example, in ATHENA, as L approaches zero, *both* recognition and classification performance must also approach zero.

However, previous work has provided explanations of E.P.'s highly selective impairment within an exemplar-based framework. Palmeri and Flanery (1999) suggest that E.P.'s intact classification performance is the result of short-term memory acquired during the test phase. During a classification test phase, many of the patterns (i.e., all but the random lures) are similar. As a result, new test patterns can be compared against recently viewed test patterns, potentially permitting classification without any reference to exemplars from the study phase. In contrast, recognition test patterns are orthogonal and, as a result, recently viewed test patterns are not informative for subsequent recognition judgements.

Palmeri and Flanery (1999) tested this hypothesis by informing participants that study patterns had been presented subliminally—in reality, the participants had not been presented with any study patterns (i.e., a 0 ms encoding condition). Participants, unsurprisingly, were unable to discriminate between "old" patterns and new patterns. However, participants were able to classify patterns at levels consistent with previous experiments. Despite having no experience with the category exemplars, participants were able to use memory of the test phase to aid classification. The result has been replicated and extended, with one participant showing uniform endorsement of the prototype despite no traditional study phase (Palmeri & Flanery, 2002). In addition, Zaki and Nosofsky (2007) have demonstrated that the structure of the patterns in the test phase is crucial, even causing participants to prefer high-distortion patterns over the prototype.

Our current analysis, like many exemplar-based models, makes a highly artificial distinction between the study and test phases. The model only stores new information during the study phase and only evaluates patterns during the test phase. However, previous exemplar models have been extended to allow for continued encoding and retrieval through the test phase of an experiment, including the GCM (see Zaki & Nosofsky, 2007). A similar extension of ATHENA in which both encoding and retrieval occur in tandem seems straightforward. In fact, the MINERVA2 model has been extended in this way in order to account for phenomena in associative learning (MINERVA-AL; Jamieson, Crump, & Hannah, 2012; Jamieson, Hannah, & Crump, 2010). Extending ATHENA to model interactive and ongoing retrieval/encoding effects is a meaningful point of future work.

The work presented here converges on a growing consensus and counterargument about the nature of memory deficits in amnesia. Thus, our data make a meaningful contribution on that point. However, we think the analysis points a more general issue that needs resolution.

The argument for a general account of memory is fundamental. Admittedly, the multiple-system perspective—to the extent that it is true—has advantages. If memory is composed of independent subsystems, the field is justified in developing different psychological theories for different behaviours. However, if memory is not composed of independent subsystems, the strategy is misguided: mistaking the number of tasks for the number of systems (Roediger, 1990). The problem comes down to a consideration of the relationship between laboratory tests of memory and the systems that those tests are meant to measure: a logical error known as the process-pure assumption.

According to the multiple-system perspective, each task (e.g., recognition and classification) measures a system of memory. According to the single-system perspective, each task measures the behaviour of memory under different conditions. Consequently, the problem that is faced by the single-system perspective is to explain memorial behaviour under different conditions holding principles and mechanisms of memory constant. Although the problem makes theory building difficult, the work presented here, as well as the work presented by others, points to the fact that the task of building a general approach to understanding memory is possible and that the field is making progress on that goal.

Of course, we are not the first to make this argument. Shanks and St. Johns (1994) and, more recently, Surprenant and Neath (2013) have called for a reunification of psychological theory. However, Newell (1973) provided perhaps the most famous argument for reunification. In his famous 20-questions paper, he argued that psychology has historically adopted a strategy of dividing behaviour into categories, labouring under the illusion that each division resolves one bit of uncertainty about psychology. He pointed out, however, that the strategy, albeit seductive, would lead the discipline into a study of phenomena and tasks rather than an explanation of behaviour. The lesson seems to have been largely ignored (e.g., see Jamieson, Mewhort, & Hockley, 2016).

Our work, along with the converging evidence developed by Benjamin (2010), Higham et al. (2000), Jamieson, Holmes, and Mewhort (2010), Nosofsky and Zaki (1998), Zaki and Nosofsky (2001), Malmberg et al. (2004), and Kinder and Shanks (2001, 2003) takes up Newell's (1973) argument by challenging the popular distinction between explicit and implicit memory in favour of a principled and unified explanation of behaviour.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Notes

1. Some multiple-system theories might be more constrained. For example, a multiple-system theory that assumes that classification is supported *only* by an implicit memory system would be falsified by better classification performance by controls compared to amnesic patients.
2. Because the vector elements are sampled from a distribution with a mean of zero, it is possible that the element might have already been zero (or very close to zero). As a result, it is not always possible to distinguish between intact information sampled at zero and missing information. Nonetheless, on average, low values of L result in traces missing larger amounts of information.
3. Mathematically, the activation of a trace is constrained to fall between −1 and +1. However, in practice, the lower limit of activation is typically approximately 0. This is because, in most implementations, lures (i.e., unstudied items, non-category items) are orthogonal-by-expectation to each of the studied items.
4. This is an additional parameter in the simulation. Similar to L we chose the values relatively arbitrarily. The same pattern of results occurs so long as d is lower for the low-distortion items.
5. Allowing β to vary across tasks did not result in a large improvement in fit.
6. In order to keep the length of the composite vector constant, we normalised it before adding each feature vector. If the length is allowed to increase with each new feature, the last feature vectors contribute more information to the composite.
7. Reed, Squire, Patalano, Smith, and Jonides (1999) demonstrated that even the most frequently endorsed features were endorsed on fewer than 80% of trials. However, Zaki and Nosofsky responded convincingly by demonstrating that their model produces similar endorsement rates, even at the level of individual features.
8. To be clear, we disagree with Smith and Minda's (2001) criticisms and refer the reader to Palmeri and Flanery (2002) for a rebuttal. Nonetheless, our analysis shows that the same dissociations can be produced under different representational assumptions.

References

Arndt, J., & Hirshman, E. (1998). True and false recognition in MINERVA2: Explanations from a global matching perspective. *Journal of Memory and Language*, *39*, 371–391.

Benjamin, A. S. (2010). Representational explanations of "process" dissociations in recognition: The DRYAD theory of aging and memory judgments. *Psychological Review*, *117*, 1055–1079. doi:10.1037/a0020810

Benjamin, A. S., Diaz, M., Matzen, L. E., & Johnson, B. (2012). Tests of the DRYAD theory of the age-related deficit in memory for context: Not about context, and not about aging. *Psychology and Aging*, *27*, 418–428.

Berry, C. J., Kessels, R. P., Wester, A. J., & Shanks, D. R. (2014). A single-system model predicts recognition memory and repetition priming in amnesia. *Journal of Neuroscience*, *34*, 10963–10974.

Brown, G. D., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, *114*, 539–576.

Bussey, T. J., & Saksida, L. M. (2007). Memory, perception, and the ventral visual-perirhinal-hippocampal stream: Thinking outside of the boxes. *Hippocampus*, *17*, 898–908. doi:10.1002/hipo.20320

Chubala, C. M., & Jamieson, R. K. (2013). Recoding and representation in artificial grammar learning. *Behavior Research Methods*, *45*, 470–479.

Chubala, C. M., Johns, B. T., Jamieson, R. K., & Mewhort, D. J. K. (2016). Applying an exemplar model to an implicit rule-learning task: Implicit learning of semantic structure. *The Quarterly Journal of Experimental Psychology*, *69*, 1049–1055.

Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of pattern-analyzing skill in amnesia: Dissociation of knowing how and knowing that. *Science*, *210*, 207–210.

Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2010). Functional dissociations within the ventral object processing pathway: Cognitive modules or hierarchical continuum? *Journal of Cognitive Neuroscience*, *22*, 2460–2479.

Dienes, Z. (1992). Connectionist and memory-array models of artificial grammar learning. *Cognitive Science*, *16*, 41–79.

Dunn, J. C., & Kirsner, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review*, *95*, 91–101. doi:10.1037/0033-295X.95.1.91

Dunn, J. C., & Kirsner, K. (2003). What can we infer from double dissociations? *Cortex*, *39*, 1–7.

Eichenbaum, H., & Cohen, N. J. (2001). *From conditioning to conscious recollection: Memory systems of the brain*. New York, NY: Oxford University Press.

Estes, W. K., Campbell, J. A., Hatsopoulos, N., & Hurwitz, J. B. (1989). Base-rate effects in category learning: A comparison of parallel network and memory storage-retrieval models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 556–571. doi:10.1037/0278-7393.15.4.556

Gaffan, D. (2002). Against memory systems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *357*, 1111–1121.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 227–247. doi:10.1037/0096-3445.117.3.227

Graf, P., Mandler, G., & Haden, P. E. (1982). Simulating amnesic symptoms in normal subjects. *Science*, *218*, 1243–1244.

Gureckis, T. M., James, T. W., & Nosofsky, R. M. (2011). Re-evaluating dissociations between implicit and explicit category learning: An event-related fMRL study. *Journal of Cognitive Neuroscience*, *23*, 1697–1709.

Higham, P. A., & Vokey, J. R. (1994). Recourse to stored exemplars is not necessarily explicit: A comment on Knowlton, Ramus, and Squire (1992). *Psychological Science*, *5*, 59–60.

Higham, P. A., Vokey, J. R., & Pritchard, J. L. (2000). Beyond dissociation logic: Evidence for controlled and automatic influences in artificial grammar learning. *Journal of Experimental Psychology: General*, *129*, 457–470.

Hintzman, D. L. (1984). MINERVA 2: A simulation model of human memory. *Behavior Research Methods, Instruments, & Computers*, *16*, 96–101.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*, 411–428. doi:10.1037/0033-295X.93.4.411

Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, *95*, 528–551.

Jamieson, R. K., Crump, M. J., & Hannah, S. D. (2012). An instance theory of associative learning. *Learning & Behavior*, *40*, 61–82.

Jamieson, R. K., Hannah, S. D., & Crump, M. J. (2010). A memory-based account of retrospective revaluation. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *64*, 153.

Jamieson, R. K., & Hauri, B. R. (2012). An exemplar model of performance in the artificial grammar task: Holographic representation. *Canadian Journal of Experimental Psychology / Revue Canadienne de Psychologie Expérimentale*, *66*, 98–105.

Jamieson, R. K., Holmes, S., & Mewhort, D. J. K. (2010). Global similarity predicts dissociation of classification and recognition: Evidence questioning the implicit–explicit learning distinction in amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*, 1529–1535. doi:10.1037/a0020598

Jamieson, R. K., & Mewhort, D. J. K. (2009a). Applying an exemplar model to the artificial-grammar task: Inferring grammaticality from similarity. *The Quarterly Journal of Experimental Psychology*, *62*, 550–575.

Jamieson, R. K., & Mewhort, D. J. K. (2009b). Applying an exemplar model to the serial reaction-time task: Anticipating from experience. *The Quarterly Journal of Experimental Psychology*, *62*, 1757–1783. doi:10.1080/17470210802557637

Jamieson, R. K., & Mewhort, D. J. K. (2010). Applying an exemplar model to the artificial-grammar task: String completion and performance on individual items. *The Quarterly Journal of Experimental Psychology*, *63*, 1014–1039.

Jamieson, R. K., & Mewhort, D. J. K. (2011). Grammaticality is inferred from global similarity: A reply to Kinder (2010). *The Quarterly Journal of Experimental Psychology*, *64*, 209–216.

Jamieson, R. K., Mewhort, D. J. K., & Hockley, W. E. (2016). A computational account of the production effect: Still playing twenty questions with nature. *Canadian Journal of Experimental Psychology*, *70*, 154–164.

Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, *114*, 1–37.

Kinder, A. (2010). Is grammaticality inferred from global similarity? Comment on Jamieson and Mewhort (2009). *The Quarterly Journal of Experimental Psychology*, *63*, 1049–1056.

Kinder, A., & Shanks, D. R. (2001). Amnesia and the declarative/nondeclarative distinction: A recurrent network model of classification, recognition, and repetition priming. *Journal of Cognitive Neuroscience*, *13*, 648–669. doi:10.1162/089892901750363217

Kinder, A., & Shanks, D. R. (2003). Neuropsychological dissociations between priming and recognition: A single-system connectionist account. *Psychological Review*, *110*, 728–744. doi:10.1037/0033-295X.110.4.728

Knowlton, B. J., Ramus, S. J., & Squire, L. R. (1992). Intact artificial grammar learning in amnesia: Dissociation of classification learning and explicit memory for specific instances. *Psychological Science*, *3*, 172–179.

Knowlton, B. J., & Squire, L. R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. *Science*, *262*, 1747–1749. doi:10.1126/science.8259522

Knowlton, B. J., & Squire, L. R. (1994). The information acquired during artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 79–91. doi:10.1037/0278-7393.20.1.79

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1996). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*, 106–120.

Kwantes, P. J. (2005). Using context to build semantics. *Psychonomic Bulletin & Review*, *12*, 703–710. doi:10.3758/BF03196761

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240.

Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review*, *96*, 25–57.

Malmberg, K. J., Zeelenberg, R., & Shiffrin, R. M. (2004). Turning up the noise or turning down the volume? On the nature of the impairment of episodic recognition memory by midazolam. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 540–549.

Mayes, A. R., Downes, J. J., Shoqeirat, M., Hall, C., & Sagar, H. J. (1993). Encoding ability is preserved in amnesia: Evidence from a direct test of encoding. *Neuropsychologia*, *31*, 745–759. doi:10.1016/0028-3932(93)90126-K

McClelland, J. L., & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology*, *18*, 1–86.

McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, *114*, 159–188. doi:10.1037/0096-3445.114.2.159

McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, *39*, 352–370. doi:10.1037/h0069819

Murdock, B. B. (1983). A distributed memory model for serial-order information. *Psychological Review*, *90*, 609–626.

Murdock, B. B. (1995). Developing TODAM: Three models for serial-order information. *Memory & Cognition*, *23*, 631–645.

Murdock, B. B. (2008). Issues with the SIMPLE model: Comment on Brown, Neath, and Chater (2007). *Psychological Review*, *115*, 779–780.

Nairne, J. S. (2003). Sensory and working memory. In A. H. Healy & R. W. Proctor (Eds.), *Comprehensive handbook of psychology, Vol. 4: Experimental psychology* (pp. 423–446). New York, NY: Wiley.

Nairne, J. S., & Neath, I. (1994). Critique of the retrieval/deblurring assumptions of the theory of distributed associative memory. *Psychological Review*, *101*, 528–533.

Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing* (pp. 283–305). New York, NY: Academic Press.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.

Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 87–108. doi:10.1037/0278-7393.13.1.87

Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 700–708. doi:10.1037/0278-7393.14.4.700

Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 3–27. doi:10.1037/0096-1523.17.1.3

Nosofsky, R. M., Denton, S. E., Murphy-Knudsen, A. F., & Unverzagt, F. W. (2012). Studies of implicit prototype extraction in patients with mild cognitive impairment and Alzheimer's disease. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*, 860–880.

Nosofsky, R. M., Little, D. R., & James, T. W. (2012). Activation in the neural network responsible for categorization and recognition reflects parameter changes. *Proceedings of the National Academy of Sciences*, *109*, 333–338.

Nosofsky, R. M., & Zaki, S. R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. *Psychological Science*, *9*, 247–255.

Palmeri, T. J., & Flanery, M. A. (1999). Learning about categories in the absence of training. *Psychological Science*, *10*, 526–530. doi:10.1111/1467-9280.00200

Palmeri, T. J., & Flanery, M. A. (2002). Memory systems and perceptual categorization. In B. H. Ross (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 41, pp. 141–189). San Diego, CA: Academic Press.

Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, *6*, 623–641.

Posner, M. I., Goldsmith, R., & Welton, K. E. J. (1967). Perceived distance and the classification of distorted patterns. *Journal of Experimental Psychology*, *73*, 28–38. doi:10.1037/h0024135

Reed, J. M., Squire, L. R., Patalano, A. L., Smith, E. E., & Jonides, J. (1999). Learning about categories that are defined by object-like stimuli despite impaired declarative memory. *Behavioral Neuroscience*, *113*, 411–419. doi:10.1037/0735-7044.113.3.411

Roediger, H. L., III. (1990). Implicit memory: A commentary. *Bulletin of the Psychonomic Society*, *28*, 373–380.

Schacter, D. L., & Tulving, E. (1994). What are the memory systems of 1994? In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 1–38). Cambridge, MA: The MIT Press.

Shanks, D. R., & Perruchet, P. (2002). Dissociation between priming and recognition in the expression of sequential knowledge. *Psychonomic Bulletin & Review*, *9*, 362–267.

Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, *17*, 367–395.

Smith, J. D., & Minda, J. P. (2001). Journey to the center of the category: The dissociation in amnesia between categorization and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 984–1002.

Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, *82*, 171–177. doi:10.1016/j.nlm.2004.06.005

Squire, L. R., & Knowlton, B. J. (1995). Learning about categories in the absence of memory. *Proceedings of the National Academy of Sciences*, *95*, 747–750.

Starns, J. J., White, C. N., & Ratcliff, R. (2010). A direct test of the differentiation mechanism: REM, BCDMEM, and the strength-based mirror effect in recognition memory. *Journal of Memory and Language*, *63*, 18–34.

Surprenant, A. M., & Neath, I. (2013). *Principles of memory*. New York, NY: Psychology Press.

Tulving, E. (1985). *How many memory systems are there?* New York, NY: Psychology Press.

Van Orden, G. C., & Kloos, H. (2003). The module mistake. *Cortex*, *39*, 164–166.

Van Orden, G. C., Pennington, B. F., & Stone, G. O. (2001). What do double dissociations prove? *Cognitive Science*, *25*, 111–172.

Woods, R. T., & Percy, M. (1974). A similiarity between amnesic memory and normal forgetting. *Neuropsychologia*, *12*, 437–445.

Zaki, S. R. (2004). Is categorization performance really intact in amnesia? A meta-analysis. *Psychonomic Bulletin & Review*, *11*, 1048–1054.

Zaki, S. R., & Nosofsky, R. M. (2001). A single-system interpretation of dissociations between recognition and categorization in a task involving object-like stimuli. *Cognitive, Affective, & Behavioral Neuroscience*, *1*, 344–359.

Zaki, S. R., & Nosofsky, R. M. (2007). A high-distortion enhancement effect in the prototype learning paradigm: Dramatic effects of category learning during test. *Memory & Cognition*, *35*, 2088–2096. doi:10.3758/BF03192940