Reinterpreting selective impairments in memory: Computational and empirical simulations of dissociations in amnesia

by

Evan T. Curtis

A Thesis submitted to the Faculty of Graduate Studies of
The University of Manitoba
in partial fulfillment of the requirements of the degree of

Doctor of Philosophy

Department of Psychology
University of Manitoba
Winnipeg

Abstract

By a dominant account, memory is composed of multiple storage systems, each operating according to unique principles. By an alternative account, memory is a single storage system and operates according to a single set of principles. Selective memory impairments in amnesia serve as the primary evidence for the multiple-system perspective. This thesis reports a critical appraisal of the multiple-system perspective using a combination of computational and empirical methods. In the computational analysis, I adopt the Holographic Exemplar Model, a single-system model of memory based on Hintzman's (1986) classic MINERVA2 model. I simulate amnesia by manipulating the quality with which items are encoded in memory. In the empirical analysis, I simulate amnesia by manipulating peoples' quality of encoding by limiting the time given to study stimuli. Simulations 1-2 and Experiments 1-2 simulate a dissociation between classification and recognition. All four analyses are consistent with the original results. Simulation 3 and Experiment 3 simulate single and double dissociations between tachistoscopic identification and recognition. The analyses were consistent with the single but not double dissociation. Simulation 4 and Experiment 4 simulate a dissociation among word-stem completion, cued recall, and recognition. Both analyses were only partially consistent with the original results, representing a failure overall. Simulation 5 and Experiment 5 derived a novel prediction from artificial grammar learning, predicting a non-dissociation between string completion and recognition. The mixed results provide some support for a single-system account of memory and opens opportunities for future work. I argue that the analysis is best considered in convergence with previous work moving toward a more integrated account of memory

Acknowledgements

There are countless people without whom this dissertation would not have been possible to complete. There is not enough space in the entirety of this document to thank everyone who deserves it, nor enough space to properly thank the few people I do acknowledge. Below is just a snippet of the gratitude that I owe:

Randy Jamieson, for hours and hours of time spent shaping me into a thoughtful scientist, for pushing me to take risks and recognize opportunities, for long conversations, for embarrassing amounts of coffee, and occasionally a few too many beers.

Chrissy Chubala for being around to talk through the mutual challenges we both faced along the way.

The University of Manitoba Faculty of Graduate Studies for UMGF funding, as well as NSERC and Randy Jamieson for GETS funding.

Sarah McAmmond, just for being there and being you.

Jeff and Kathy Curtis, for constant support in ways that cannot possibly be counted. There is no way I could have done this without you.

Table of Contents

List of Tables

List of Figures

Chapter 1 – Background and Premises

Early on, memory was conceived as a single store (for exceptions see James, 1890; Maine de Biran, 1804). However, a more recent perspective is that memory is composed of multiple stores, each serving a unique function and operating by different principles. For example, theorists argue for fundamental distinctions between long- and short-term memory (Atkinson & Shiffrin, 1968; Glanzer, 1972), procedural and declarative memory (Cohen, 1984; Scoville & Milner, 1957), semantic and episodic memory (Tulving, 1983), and so on. At present, the multiple-system perspective is dominant and often treated as axiomatic (see Eichenbaum & Cohen, 2001; Fodor, 1983; Schacter & Tulving, 1994; Squire & Zola-Morgan, 1991).

Primary evidence for the multiple-system perspective lies in a critical database of task dissociations, where changing a condition at learning or test affects peoples' performance on one task but not another. The most frequently cited evidence of dissociations, and thus the need for separate memory systems, comes from experimental observations associated with amnesia.

Scoville and Milner's (1957) examination of patient H.M. has played a pivotal role in memory theory. Following a bilateral lobectomy that lesioned parts of the medial temporal lobes, including large segments of the hippocampal formation, H.M. suffered a severe memory deficit. However, H.M.'s memory deficit was selective. For example, when asked to trace a figure viewed in a mirror, H.M. showed gradual improvement despite being unable to remember the training sessions (Milner, 1962). In other experiments, H.M. improved at identifying words after studying them despite being unable to remember those same words (Milner, Corkin, & Teuber, 1968; see also

Warrington & Weiskrantz, 1970, 1974, 1978). Based on the result, Milner and others (e.g., Cohen & Squire, 1980) argued that the selective memory impairment implies independent memory systems: a declarative system that is compromised in amnesia and a procedural system that is not. Although the distinction associated with H.M.'s deficit has been discussed using a variety of terminology, it is most commonly discussed and examined today as a distinction between *explicit* and *implicit* memory.

The story of H.M. and the taxonomy of memory that followed from it are typically taught in all introductory psychology text books. The taxonomy has also become a standard assumption in scientific discussion. However, a closer examination reveals problems.

Despite the popularity of the multiple-system perspective, the concept of a system remains controversial. For example, the division between long- and short-term memory (Atkinson & Shiffrin, 1968) has been challenged by a number of important contributions including Craik and Lockhart's (1972) levels of processing theory, Cowan's (1988) integrative account of working and long-term memory, Nairne's (1990) Feature Model of memory, and Brown, Preece, and Hulme's (2000) Oscillator-Based Associative Recall model of memory. The division between episodic and semantic memory (Tulving, 1985) has been challenged by Hintzman's (1984) MINERVA2 model of memory, Brooks (1987; see also Brooks & Vokey, 1991) theory of decentralized cognition, Whittlesea's (1997; see also Leboe-McGowan & Whittlesea, 2013) SCAPE model, Ratcliff and McKoon's (1986) race models, Jacoby's (1991) Process Dissociation model, and Shiffrin and Steyver's (1997) Retrieving Effectively from Memory model. The dissociation of implicit and explicit memory (Reber & Squire, 1994) has been challenged by McClelland

and Rumelhart's (1985) Parallel Distributed Processing models of memory, Kinder and Shanks' (2001, 2003) Recurrent Network models, Usher and McClelland's (2001) Ballistic Choice model, Benjamin's (2010) DRYAD model, Nosofsky and Zaki's (1998) Generalized Context Model, and Jamieson, Holmes, and Mewhort's (2010) application of MINERVA2. Finally, the very logic of dissociation and the underlying rational model that permits and limits the strength and rationality of conclusions based on dissociations (including double dissociations) has been challenged by Dunn and Kirsner's (1988; 2003) Reversed Association model, Prince, Brown, and Heathcote's (2012) revival of Bamber's (1979) State-Trace Analysis method, and by Van Orden, Pennington, and Stone's (2001; see also Van Orden & Kloos, 2003) analytic conclusions.

Despite a long history of debate and evidence for a single system account of memory by leading scientists published in leading journals, the position that memory is divided has often escaped those criticisms by claiming the authors misconstrue or miscast what a separate system is and how it is defined. Therefore, it is important to define what I will treat as a separate system in my analysis.

In this dissertation, I interpret a memory system as an independent storage space that encodes specific types of representations. In relation to the division that I focus on, the explicit memory system is a storage space for instances or single items studied in a specific time and place. In all of the work that follows, the time and place will be in the study list of a laboratory experiment. In contrast, an implicit memory system is a separate storage space that encodes an aggregate or high-level representation of studied items. I will grant some latitude to the kinds of representations that the system might extract in order to meet the variety and differences as presented in the experimental database.

However, in the context of a standard category-learning experiment, the implicit memory system would learn and store rules that define the category and/or a statistical abstraction of the category's ideal member (i.e., the category prototype).

In the framework of this conceptualization, the explicit memory system that stores particular instances of studied items is compromised in amnesia. In contrast, the implicit memory system that stores high level abstractions of the studied list is intact. It is this selective impairment of the explicit store of studied exemplars that produces dissociations among different memory behaviours; for example, the fact that amnesic patients perform categorization relative to a prototype as well as controls even though they cannot recognize the exemplars that they studied and that the prototype was learned from.

Additionally, it is also important to define what I mean by a dissociation. By the classic definition, a single dissociation is a pattern in which a factor has no impact on one behavior while having a non-zero impact on another behaviour. This is the definition that researchers have adopted, whether consciously or unconsciously. For example, amnesia is claimed to impair recognition but have a completely null effect on categorization.

However, there is reason to question the classic definition, or at least that amnesia produces dissociations by the classic definition. In a critical analysis, Zaki (2004) argued that amnesic patients are impaired at implicit memory tasks such as categorization, and that the failure to observe that impairment is the result of low sample size and, therefore, low statistical power (see also Ostergaard & Jernigan, 1993, for a similar argument applied to priming). To alleviate this problem, Zaki conducted a meta-analysis on the database of amnesic patient's categorization performance. By utilizing a more powerful analysis, Zaki observed an impairment in amnesic patients' categorization performance smaller than the impairment to recognition, but reliable nonetheless. I argue that the pattern of dissociations in amnesia is not based on a null impairment to implicit memory, but instead a pattern in which the impairment to implicit memory is present but of lesser magnitude than the impairment to explicit memory. Independent of whether this pattern can be considered a dissociation, it is the pattern classically used to defend the multiple-system perspective, and thus the pattern that is the target of my analysis.

Finally, I wish to provide the structure and logic of my examination. I wish to argue against the theoretical demands of a multiple-system account. However, I do not think the multiple-system account is inherently negative. Rather, my aim is to show the sufficiency of a single-system account where the multiple-system account has been applied. The analysis is important because it tests the foundational assumptions about the structure and nature of human memory. Additionally, to the extent that a single-system account can accommodate the data, we ought to be appropriately skeptical about the necessity of dividing memory into separate systems. Of course, if the single system theory cannot organize and explain the empirical database, it is in that moment that the multiple-system account gains strength.

Single-System Accounts

The Process Perspective

The process perspective is a popular single-system account of memory (e.g., Jacoby, 1983, 1988; Kolers & Roediger, 1984). From this perspective, dissociations arise in the interaction between how materials are processed at study and test, making the phenomenon of transfer-appropriate processing central to an understanding of memory and memory performance (Morris, Bransford, & Franks, 1977). For example, many implicit tasks are perceptually driven and many explicit tasks are conceptually driven (Jacoby, 1983; Roediger, Srinivas, & Weldon, 1989) and manipulating the match between encoding and retrieval can produce the kinds of dissociations observed in amnesia and associated with a division of memory into separate implicit and explicit stores (Graf, Mandler, & Haden, 1982). Whereas the system denies the need for different stores, it appreciates differences in how memory of studied events are accessed and used to produce the dissociations thought to force a different systems perspective (Jacoby & Dallas, 1981).

The process dissociation procedure (Jacoby, 1991) is a common method to evaluate process distinctions. In that procedure, participants study different items in two study phases and then are tested using two different sets of instructions. In an inclusion test, participants are asked to endorse items from either of the studied lists. In an exclusion test, participants are asked to only endorse words from one of the studied lists.

The underlying logic of the procedure is that automatic processes like familiarity and intentional processes like recollection operate differently under the different test demands. Under inclusion instructions, both familiar and recollected items ought to be endorsed as old. Under exclusion instructions, only recollected items from the targeted list ought to be endorsed as old. Using the data from a subject, one can estimate the rate of recollection by the difference in endorsement rates in the two tests. Like standard dissociation logic, the process dissociation procedure is widely implemented but controversial, with researchers raising questions about the assumptions that underlie the mathematics of the theory, the assumptions that underlie the formulas, and the validity of the logical model itself (see Hirshman, 1998; Curran & Hintzman, 1995). However, the underlying claim from the technique stands: it is possible that a dissociation in performance on two tasks might be observed because of a difference in how participants process the stimuli at study and test rather than assuming that different memory systems store different aspects of experience and knowledge about the studied items.

The Computational Perspective

A third perspective is that memory is unitary and operates by general rather than selective principles. This perspective is most commonly examined in the field of computational psychology. Critically, the perspective proposes that dissociations observed in amnesia can be explained within a framework that assumes a single system. For example, the Generalized Context Model (Nosofsky & Zaki, 1998), Simple Recurrent Networks (Kinder & Shanks, 2001, 2003), and the DRYAD model of memory (Benjamin, 2010) can explain dissociations observed in amnesia and aging even though they assume memory is a single system with a single process. Although the unitary approach to memory is an emerging account and runs contrary to the dominant view, it has two primary benefits. First, formal models provide existence proofs against the multiple-system perspective, providing a more parsimonious account of the empirical database. Second, the computational approach requires theorists to formally explicate their assumptions, providing firm and falsifiable predictions in addition to explaining existing data (see Lewandowsky, 1993).

In this dissertation, I adopt the computational approach and provide a combined computational and experimental simulation analysis of dissociations, specifically several key dissociations observed in amnesia. As outlined earlier, a number of computational analyses have been successful in accommodating various dissociations. However, those successes span a number of different formalisms and frameworks. The situation leaves one with the hope of a model that can explain the database of dissociations. However, the situation also leaves one wondering if there is a common formalism that can accommodate all of the dissociations at once. In pursuit of this goal, my first contribution is to apply a single theory to a range of important task dissociations and demonstrations associated with amnesia. Following the formal analysis, I will discuss the importance and value of approaching memory from a coherent and formal perspective. To foreshadow my results, my model analysis were consistent some of the original experiments but not others. The model accommodated differences between recognition and classification (Simulations 1 and 2). The model also accommodated differences between recognition and tachistoscopic identification (Simulation 3). However, the model failed to accommodate a related double dissociation. The model also failed to convincingly accommodate differences among recognition, cued recall, and word-stem completion (Simulation 4).

MINERVA2 and the Holographic Exemplar Model

MINERVA 2 (Hintzman, 1984, 1986, 1988) is a model of human memory that was originally designed to collapse the distinction between episodic and semantic memory.

Informally, MINERVA2 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 probe item is presented to memory, the model retrieves an aggregate of the instances that are similar to the probe item. When a probe item matches a particular instance in memory, the model primarily retrieves that instance. This process allows the model to perform recognition. Additionally, because the probe retrieves entire instances, and the model stores associations within instances, partial probe items (such as a cue) retrieve events they have co-occurred with. This process allows the model to perform cued recall.

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 known as its *content*. The strength of the echo's activation is known as its *intensity*.

Representation and Storage

MINERVA 2 borrows its representation scheme from Shepard's (1957) classic work in machine learning and the geometry of psychological space. Namely, the stimulus is represented as a vector of *n* elements. As a result, each stimulus can be conceived as having a location in an *n* dimensional space. Consistent with standard rules of algebra, the similarity between two stimuli can be indexed by their distance in the psychological space or by their similarity of their angle as they depart from the origin.

In the original formulation of MINERVA2, each element in a vector representation takes on a value of +1 or -1 with equal probability (see Tversky, 1977). Associations between stimuli are represented by concatenating the component stimulus vectors. Memory is an *m* by *n* matrix, where *m* is the number of rows and *n* is the number of columns. Each trace is stored to a unique row in the memory matrix. The representation scheme has been updated more recently by Jamieson and Mewhort (2011) who adopted the method of holographic reduced representation to represent stimuli in memory (e.g., Jones & Mewhort, 2007; Lewandowsky & Murdock, 1989; Murdock, 1983, 1995; Plate, 1995). The updated model is known as the Holographic Exemplar Model (HEM).

In holographic reduced representations, each vector element in a stimulus representation is randomly sampled from a Normal distribution with a mean of zero and a standard deviation of $1/\sqrt{n}$, where *n* is the number of elements. Association is represented using a vector operation called circular convolution that returns an associative representation given two vectors as arguments. More concretely, given two vectors, *x* and *y*, circular convolution produces a unique vector of the same dimensionality, *z*,

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

where *x* and *y* are item vectors, *z* is the association of *x* and *y*, *n* is the dimensionality of *x*, *y*, and *z*. Figure 1 presents an example of the operation on two vectors with five elements.

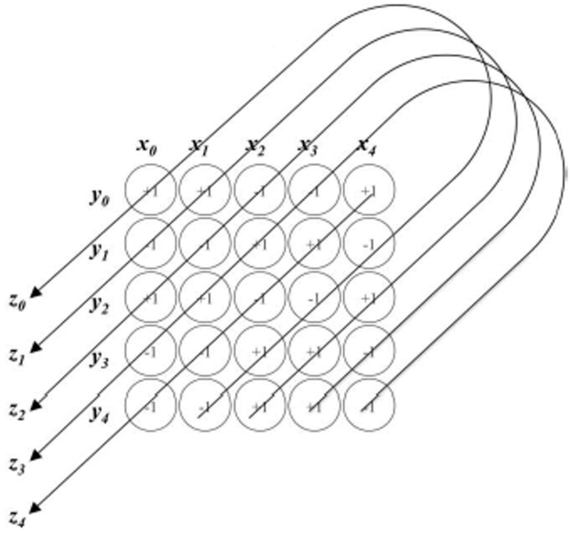

Figure 1. *A depiction of circular convolution. Multiplying x and y produces an n by n outer product matrix. The arrows demonstrate summation across elements to produce values in the resulting composite vector, z. The method can be thought of as a way to compress the associative outer product matrix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

Whether one adopts the original or the HEM representation scheme, encoding is assumed to be imperfect. To implement the 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. 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 secondary storage of rules about or structure in the list of studied items.

Retrieval

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

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

where p_j is the jth element of the probe, M_{ij} is the jth element of the ith trace, and n_R is the number of indices with at least one non-zero value in the trace or probe.

If the HEM representation scheme is adopted, similarity, s, between a probe and a trace is given instead by

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

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

$$a_i = s_i^3$$

Figure 2 illustrates the non-linear relationship between trace similarity and trace activation. When the trace and probe are only moderately similar (i.e., similarity values between -.6 and .6), the trace is weakly activated. In contrast, when the trace and probe are highly similar (i.e., similarity values greater than .6), the trace is strongly activated. Additionally, if the trace is highly dissimilar (i.e., similarity values less than -.6), the trace is strongly negatively activated; however, negative similarities are rare, making the operative range of the activation function 0 to 1. The non-linear relationship between similarity and activation is an important addition to the model because it ensures that a probe retrieves traces most similar to itself without retrieving traces with incidental and passing similarities.

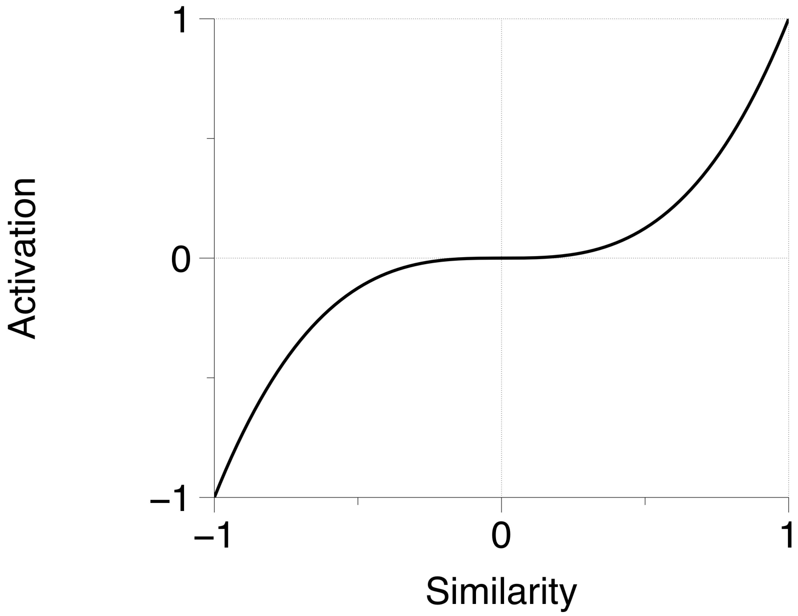

Figure 2. *The non-linear relationship between probe-trace similarity and the activation of the trace.*

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 *i*th trace and M_{ij} is the *jth* element of the *i*th trace. The content of the echo is a vector containing the aggregated information from each of the activated traces.

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

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

where a_i is the activation of the *i*th 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 *j*th element of the echo and p_j is the *j*th element of the probe.

In the work that follows, I adopt the cosine similarity metric. I will do this for two clear reasons. First, echo intensity sums across all traces. As a result, echo intensity varies as a function of the number of traces in memory, ranging from $-m$ to $+m$ (where m is equal to the number of items studied). The variation in range and magnitude 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. Additionally, cosine similarity is an established vector similarity metric in a wide range of psychological application (e.g., Jones & Mewhort, 2007; Kwantes, 2005; Landauer & Dumais, 1997). In contrast, echo intensity is idiosyncratic to the model.

MINERVA2 has been applied to a wide range of memory phenomenon, including, but not limited to, recognition and frequency judgment (Hintzman, 1984, 1988), categorization (Hintzman, 1986), recognition failure of recallable words (Hintzman, 1987), the relationship between confidence and accuracy (Clark, 1997), the Deese-Roediger-McDermott (DRM) procedure (Arndt & Hirshman, 1998), function learning (Kwantes & Neal, 2006), word naming (Kwantes & Mewhort, 1999), semantic access (Kwantes, 2005), decision making (Dougherty, Gettys, & Ogden, 1999), serial reaction time tasks (Jamieson & Mewhort, 2009a), and artificial grammar learning (Jamieson & Mewhort, 2009b). Importantly, MINERVA2 has also been applied to a hallmark dissociation between artificial grammar learning and recognition observed in amnesia.

Artificial Grammar Learning

In a typical artificial grammar task, participants study strings of consonants constructed by a finite state Markov grammar (see Figure 3). Strings are constructed by traversing the paths between nodes and appending the corresponding letter. The string ends at one of the exit points. Any string that can be constructed in this manner is considered grammatical. All other strings are considered ungrammatical. Following study, participants are asked to sort novel grammatical strings from ungrammatical strings. They do so successfully. However, when asked to describe the rules of the grammar, they either claim ignorance or cite rules inconsistent with both the grammar and their own judgments (Reber, 1967).

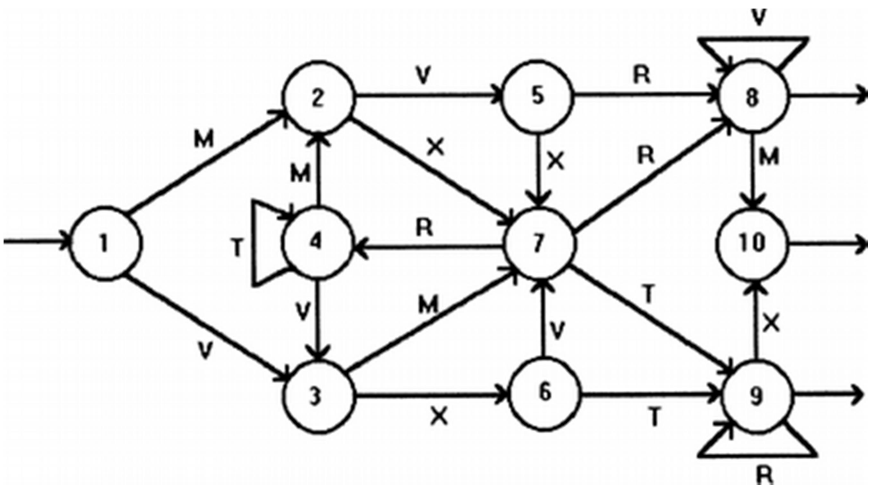

Figure 3. *An example of an artificial Markov grammar. Strings are produced by following the arrows from node to node and appending the corresponding letter. The string ends at one of the three exit points. Reproduced with permission from Vokey and Brooks (1992), Salience of item knowledge in learning artificial grammars, in Journal of Experimental Psychology: Learning, Memory, and Cognition, 8(2), pages 328-344, published by APA*

By one account, an implicit memory system extracts the rules of the grammar. Whereas this information supports classification, it is inaccessible to consciousness and thus the rules cannot be articulated (e.g., Dienes, Broadbent, & Berry, 1991; Knowlton and Squire, 1994; Manza & Reber, 1997, Reber, 1967). This account represents the multiple-system perspective. By another account, classification is accomplished by a comparison to memory for the studied strings (e.g., Brooks & Vokey, 1991; Jamieson & Mewhort, 2005, 2009b; Vokey & Brooks, 1992; Wright & Whittlesea, 1998). Because the explanation does not assume implicit knowledge of the grammar, this account represents the single-system perspective.

To develop additional support for the multiple-system perspective, Knowlton, Ramus, and Squire (1992) tested amnesic patients on the artificial grammar task. Amnesic patients discriminated grammatical strings from ungrammatical similarly to control participants. However, they were much worse at discriminating studied strings from unstudied strings. Based on the dissociation, Knowlton, Ramus, and Squire concluded that classification is supported by an implicit memory system that is spared in amnesia. In contrast, recognition is supported by an explicit memory system that is damaged in amnesia.

However, the dissociation does not force the distinction. Jamieson, Holmes, and Mewhort (2010) provided a computational demonstration that the dissociation is also consistent with the assumption that amnesia is a global encoding deficit rather than a selective system impairment. This assumption maps directly onto the learning parameter, *L*, in MINERVA2. Healthy memory is represented by high values of *L*, resulting in rich and complete traces. In contrast, amnesic memory is represented by low values of *L*, resulting in sparse and incomplete traces. Figure 4 presents the model's performance on classification and recognition in Knowlton, Ramus, and Squire's (1992) task as a function of *L*. Increasing *L* improves both classification and recognition performance. However, recognition performance increases at a faster rate. Therefore, the model predicts that both recognition and classification will improve as a function of L, but that the rate of improvement is faster for recognition than classification. Critically, that assertion is consistent it Knowlton, Ramus and Squire's results and accommodates the pattern of impairment without recourse to a systems distinction. Instead, the model accommodates the pattern by assuming a global encoding deficit.

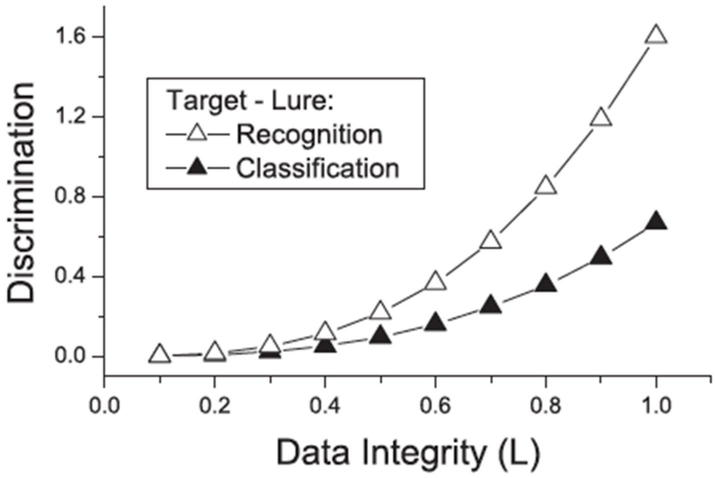

Figure 4. *MINERVA2's discrimination (difference in echo intensities) as a function of L for artificial grammar classification and recognition. Reproduced with permission from Jamieson, Holmes, and Mewhort (2010), Global similarity predicts dissociation of classification and recognition: Evidence questioning the implicit-explicit distinction in amnesia, in Journal of Experimental Psychology: Learning, Memory, and Cognition, 36(6), pages 1529-1535, published by APA.*

Following their computational analysis, Jamieson, Holmes, and Mewhort (2010) also reported an empirical simulation of Knowlton, Ramus, and Squire's (1992) task using healthy participants. Half the participants were given six seconds to study each string, and half were given only 100 ms per string. The short study time was meant to induce poor encoding, thus empirically simulating amnesia; a method that has been used a number of times by other researchers for the same purpose (see Graf, Mandler, & Haden, 1982; Nosofsky & Zaki, 1998; Woods & Percy, 1974; Zaki & Nososfky, 2001). Consistent with Knowlton, Ramus, and Squire's experiments with amnesic patients and controls, there was no difference in the groups' classification performance. However, the 100 ms group showed worse recognition performance. Jamieson, Holmes, and Mewhort concluded that the dissociation between classification and recognition in amnesia is consistent with an impairment pattern based on incomplete encoding of studied exemplars as manipulated by study time.

Premises, Framework, and Outline of Dissertation

The method of blending computational and empirical simulation approaches provides a strong framework for rational analysis. Both approaches test the premise that a global encoding deficit is sufficient to produce differences between putative explicit and implicit memory tasks. However, they do so in different ways. The computational approach produces the deficit by manipulating the learning parameter L. The empirical approach produces the deficit by manipulating study time. The correspondence between the computational parameter and the experimental manipulation provides the kind of converging evidence that is needed to lend strength to the common point.

My dissertation adopts the same strategy of developing converging computational and empirical evidence that differential patterns of impairment observed in amnesia are consistent with a general encoding deficit rather than a selective impairment of separate systems. In Chapter 2, I apply the HEM to hallmark results in the empirical database. In the computational work, I argue that amnesia can be understood as a global encoding deficit and represent this deficit by measuring shifts in task performance as a function of the model's learning parameter L. The computational simulations provide a clear analysis on the necessity of a systems distinction, arguing instead for the sufficiency of the single-system perspective.

In Chapter 3, I replicate hallmark results by manipulating study time in controlled experiments as a proxy for amnesia. Whereas control participants are given study times typical in memory research, simulated amnesic patients are presented with study items rapidly. The empirical simulations provide another window on the necessity of a systems distinction.

In Chapter 4, I leverage the computational and empirical simulations to move from an explanation of existing data to generating a novel prediction in artificial grammar learning. The prediction adds another target into the empirical database to further refine theory development and to solve what I see to be some of the experimental and methodological problems in the database on impairment patterns observed in amnesia.

Finally, in Chapter 5, I provide a general discussion identifying the conclusions I have drawn based on the results of the computational and empirical simulations. I then discuss the general limitations to my approach and identify directions for future research. Finally, I discuss the importance of adopting a more integrated approach to understanding memory (Newell, 1973, 1990).

Chapter 2 – Computational Simulations: A Single-System Approach to Amnesia

The goal of this chapter is to model hallmark patterns of impairment observed in amnesia in the HEM framework. As explained in the preceding section, I will simulate performance in tasks such as recognition, cued recall, categorization, tachistoscopic identification, and word-stem completion as a function of the learning parameter, L. The general aim is to show that performance in different tasks is affected differently by L. More specifically, the goal is to show that shifts in L affect performance in explicit memory tests more than they affect performance in implicit memory tests. If so, the simulations would demonstrate that selective impairments might be consistent with a multiple-systems interpretation but cannot force that conclusion. To begin, I will consider a classic experiment reported by Knowlton and Squire (1993).

Simulations of Prototype Abstraction and Classification

The Posner-Keele dot pattern task (Posner, Goldsmith, & Welton, 1967; Posner & Keele, 1968) is a classic method to measure classification based on family resemblance. Participants study exemplar dot patterns that are derived from a prototype pattern. The exemplars are produced by distorting the prototype; each dot moves specified distances with specified probabilities. As the distances and probabilities increase, the exemplars become more distorted approximations of the prototype. Following study, participants judge novel dot patterns on category membership.

Participants reliably discriminate between patterns that were derived from the prototype from patterns that were not. Additionally, participants show a *typicality gradient*. Namely, they endorse the unstudied category prototype most strongly than unstudied distortions and endorse the distortions more strongly than unrelated patterns.

The fact that participants endorse the unstudied prototype more strongly than even the studied exemplars is typically taken as evidence that participants learn and base their category decisions on the prototype (e.g., Homa, Cross, Cornell, Goldman, & Schwartz, 1973; Posner & Keele, 1968, 1970; Smith & Minda, 1998; Strange, Keeney, Kessel, & Jenkins, 1970). By this account, classification is supported by an implicit memory system that stores the prototype representation but does not store the exemplars.

An alternative explanation is that participants store the studied exemplars and judge the category of unstudied test items by their global similarity to the studied list (e.g., Hintzman, 1986, Medin & Schaffer, 1978; Nosofsky, 1986, 1988, 1991; Nosofsky & Zaki, 1998; Shin & Nosofsky, 1992). By this account, classification is supported by the same memory system that serves other functions like recognition. The prototype effect follows from the fact that a prototype (even when unstudied) is globally similar to all of the studied items and thus endorsed strongly.

Knowlton and Squire (1993) adapted this procedure to show that amnesic patients perform as well as controls in classification. However, amnesic patients underperform controls in recognition of the studied exemplars. Based on the result, they suggested that amnesic patients can learn the prototype by implicit abstraction, even though they cannot remember the studied items because of their impaired explicit memory. But, a closer look at their procedure reveals critical methodological oversights that allow for alternate explanations.

In the experiment, participants studied 40 exemplars, each representing a high-level distortion of a single prototype. Following study, the participants classified 84 unstudied patterns. Four of the patterns were the prototype (i.e., the single prototype presented four times over the test list), 20 were unstudied high-distortion exemplars of the prototype, 20 were unstudied low-distortion exemplars of the prototype, and 40 were unstudied random dot patterns unrelated to the prototype (see Figure 5 for examples).

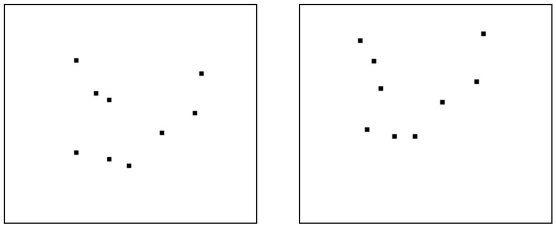

Prototype

Low-Distortion Exemplar

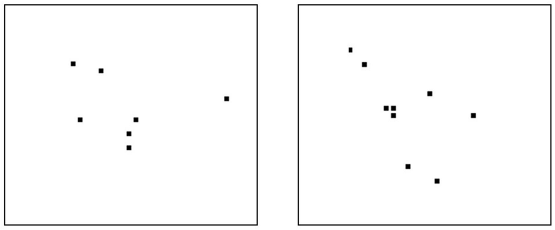

High-Distortion Exemplar

Random Pattern

Figure 5. *Examples of dot patterns used in the Posner-Keele task, taken from my Experiment 1 (see Chapter 3).*

The left panel of Figure 6 shows the results. Both amnesic patients and controls endorsed the prototype more strongly than the distortions, endorsed low-level distortions more strongly than high-level distortions, and endorsed unstudied random patterns least strongly. The presentation of overall categorization performance in the right panel of Figure 6 shows that amnesic patients and controls performed equally at discriminating positive from negative members of the category.

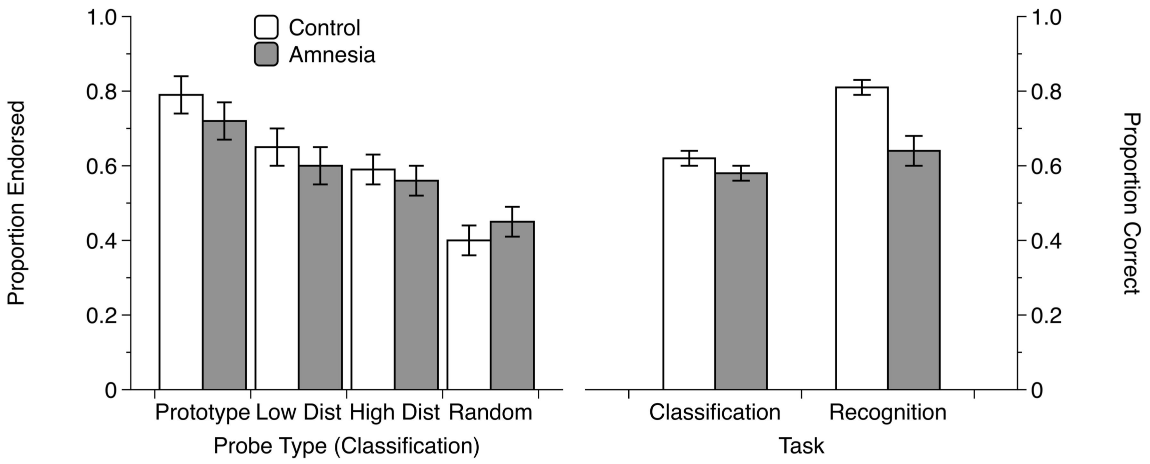

Figure 6. *Classification and recognition performance for amnesic patients and control participants, reproduced from Knowlton and Squire (1993). The left panel presents endorsement rates in the classification task. The right panel presents proportion correct in both tasks. Control participants are presented in white bars and amnesic patients are presented in black bars. From Knowlton, B. J., & Squire, L. R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge, Science, 262(5140), 1747-1749. Reproduced with permission from AAAS and Knowlton, B. J..*

Knowlton and Squire (1993) also tested the same amnesic patients' and controls' ability to recognize studied from unstudied dot patterns. But, contrary to the tale that is usually told, participants were not asked to recognize the patterns they studied. Rather, they participated in a second experiment with a different procedure and that used different materials. In the recognition test, participants studied 40 patterns. However, the study list was composed of five high-distortion exemplars of five different prototypes and participants studied each exemplar eight times. Following study, participants tried to recognize the five studied patterns from five unstudied patterns from the same five prototypes.

The right panel of Figure 6 presents the results. As shown, whereas amnesic patients performed as well as controls at classification, they performed much worse than controls at recognition. Based on that result, Knowlton and Squire (1993) concluded that classification is supported by an implicit memory system that learns the prototype and is intact in amnesia. In contrast, recognition is supported by an explicit memory system that is damaged in amnesia.

Although the interpretation of the results has been challenged, the empirical result is unquestionable. The same empirical results have been obtained in other laboratories (Lories, Van der Linden, & Cornille, 1995; Reber, Stark, & Squire, 1998; Knowlton & Squire, 1995), with different patient populations (Kéri et al., 1999; Zannino et al., 2012), and materials (Graham et al. 2006; Zannino et al.; see Graham et al. and Kolodny, 1994, for exceptions).

The Generalized Context Model

Nosofsky and Zaki (1998) argued against Knowlton and Squire's (1993) conclusions. They proposed an alternative account based on the Generalized Context Model (GCM), an exemplar-based model of categorization. In the GCM, items are represented as points in a psychological space. Similarity between items is indexed as their distance within that space. When a test item is presented to memory, the evidence that the item has been studied is indexed by the summed similarity (i.e., the inverse of distance) between the test item and every item in memory. The summed similarity can also be used to index whether an item belongs to the same category as the items in memory. At an abstract level, the GCM is based on similar principles as MINERVA2. However, the computational details differ.

For example, in Knowlton and Squire's (1993) classification task, the probability of endorsing a probe as a category member is given by,

$$P(C) = \frac{40\ \times\ s(i,h)}{40\ \times\ s(i,h) + k_C}$$

Where $s(i, h)$ is the average similarity between the probe and each of the studied items and k_C is a criterion parameter. In the formula, similarity is multiplied by 40 to reflect the number of studied exemplars. The probe is more likely to be endorsed as the global similarity between the probe and the studied exemplars increases.

In contrast, the probability of endorsing a studied or “old” pattern in the recognition test is given by

$$P_{old}(R) = \frac{\delta + 4 \times s(i,r)}{\delta + 4 \times s(i,r) + k_R}$$

where $s(i,r)$ is the average similarity between the probe and each of the studied items, δ is the similarity of the probe to its copy in memory, k_R is a separate criterion parameter, and similarity is multiplied by four to reflect the number of studied exemplars (excluding the current test exemplar’s similarity that is represented by δ). Similarly, the probability of endorsing a new pattern is given by

$$P_{new}(R) = \frac{5 \times s(i,r)}{5 \times s(i,r) + k_R}$$

where $s(i,r)$ is the average similarity between the probe and each of the studied items, k_R is a separate criterion parameter, and similarity is multiplied by five to reflect the number of studied exemplars.

To compute similarity, Nosofsky and Zaki (1998) had participants provide similarity ratings between pairs of patterns following category learning. The similarity values provided to the model are given by

$$s(i,j) = [rating(i,j)]^c$$

Where c is a parameter reflecting sensitivity to similarity. High values of c improve the model’s ability to distinguish between exemplars; low values of c decrease this ability. Nosofsky and Zaki's (1998) main premise is that amnesia can be expressed as a decrease to the sensitivity parameter.

The model produces a close match to the empirical data, predicting little difference in classification performance but a large difference in recognition performance as a function of c. The model's best-fitting parameter values – derived by least-squares – support the assumption that amnesic patients are less sensitive to similarity; the best-fitting value of c was 5.18 for controls compared to 2.99 for amnesic patients (the absolute values are not readily interpretable; the relative difference is the key point).

The GCM's ability to accommodate the pattern by a shift in memory sensitivity challenges the necessity of a systems distinction. However, there are limitations to the demonstration.

In a rebuttal, Smith and Minda (2001) objected to Nosofsky and Zaki's (1998) analysis based on the fact that they used participants' own similarity ratings to estimate $s(i, r)$. The problem was that participants made those ratings after participating in a categorization experiment using the same materials. Smith and Minda argued that participants could have abstracted the prototype in that process. As a result, the model's success might reflect an extraction of the prototype into implicit memory and an implicit deployment of that knowledge when judging similarity of exemplars, precisely the conclusion that Nosofsky and Zaki's analysis was implemented to reject.

To support this objection, Smith and Minda (2001) equipped the GCM with objective similarity ratings based on the logarithmic similarity functions described by Posner, Goldsmith, and Welton (1967). As predicted, the GCM equipped with objective measurements of stimulus similarity provided a poor fit to Knowlton and Squire's (1993) data. Based on the failure, Smith and Minda rejected the GCM as a viable model of the difference between classification and recognition in amnesia, leading them to conclude in favour of the multiple-system explanation instead.

In the simulation that follows, I apply the HEM to the pattern of impairment reported by Knowlton and Squire (1993). The representation scheme is objective and, therefore, escapes Smith and Minda's (2001) criticism.

Simulation 1

I conducted a simulation to evaluate a single-system account of Knowlton and Squire's (1993) demonstration that amnesia affects recognition more than classification. In their experiment, both amnesic patients and controls were tested for recognition and classification of dot patterns.

In the classification test, participants studied 40 dot patterns that were high-level distortions of a common prototype. Following study, participants were tested for classification of 84 unstudied patterns that included four instances of the prototype, 20 low-level distortions, 20 high-level distortions, and 40 random patterns.

In the corresponding recognition test, participants studied five dot patterns from five unique prototypes eight times each. Following study, they were tested for recognition of those five patterns relative to five unstudied patterns from the same prototype.

Knowlton and Squire (1993) reported the expected pattern: Amnesic patients performed as well as controls in the classification test but performed worse than controls in the recognition test.

To simulate the classification test, I generated a random vector to stand for the prototype. The vector was a 100 element vector with each element taking a random number sampled from a normal distribution with mean of zero and standard deviation of $1/\sqrt{n}$. Once generated, I used the prototype to generate 40 high-level distortions of the prototype for the study list and 80 items for the test list: 20 unstudied high-level distortions of the prototype, 20 low-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* = .25 for a high-distortion item and *d* = .15 for a low-distortion item. The strategy is consistent with previous work that applied the model to category learning (e.g., Arndt & Hirshman, 1998; Hintzman, 1986). Next, I simulated study by storing all 40 of the study items to the memory matrix at learning rate L.

Next, I 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, I computed the cosine similarity between the echo and the test item. Finally, I converted the cosine similarities into response probabilities according to a standard logistic transformation,

$$p('Endorsed') = \frac{1}{1 + e^{-aS+b}}$$

where *S* is the cosine similarity, *a* is a free scaling parameter, and *b* 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., Estes, Campbell, Hatsopoulos, & Hurwitz, 1989; Dienes, 1992; Gluck & Bower, 1988; McClelland & Elman, 1986;

McClelland & Rumelhart, 1985). Consistent with standard practice, I conducted 1000 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, I generated a set of five random prototypes. I generated two exemplars from each prototype: one for the study list and one for the test list. I simulated study by storing the five items in the study list to a memory matrix at learning rate *L*. Next, I 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, I computed the cosine similarity between the echo and test item. Finally, I computed recognition performance by converting the cosine similarity into response probabilities using the logistic function described above. I conducted 1000 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 7 presents the results from the model where *a* = 3.4, *b* = 2.0 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 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 weak influence on classification performance but had a strong influence on recognition performance.

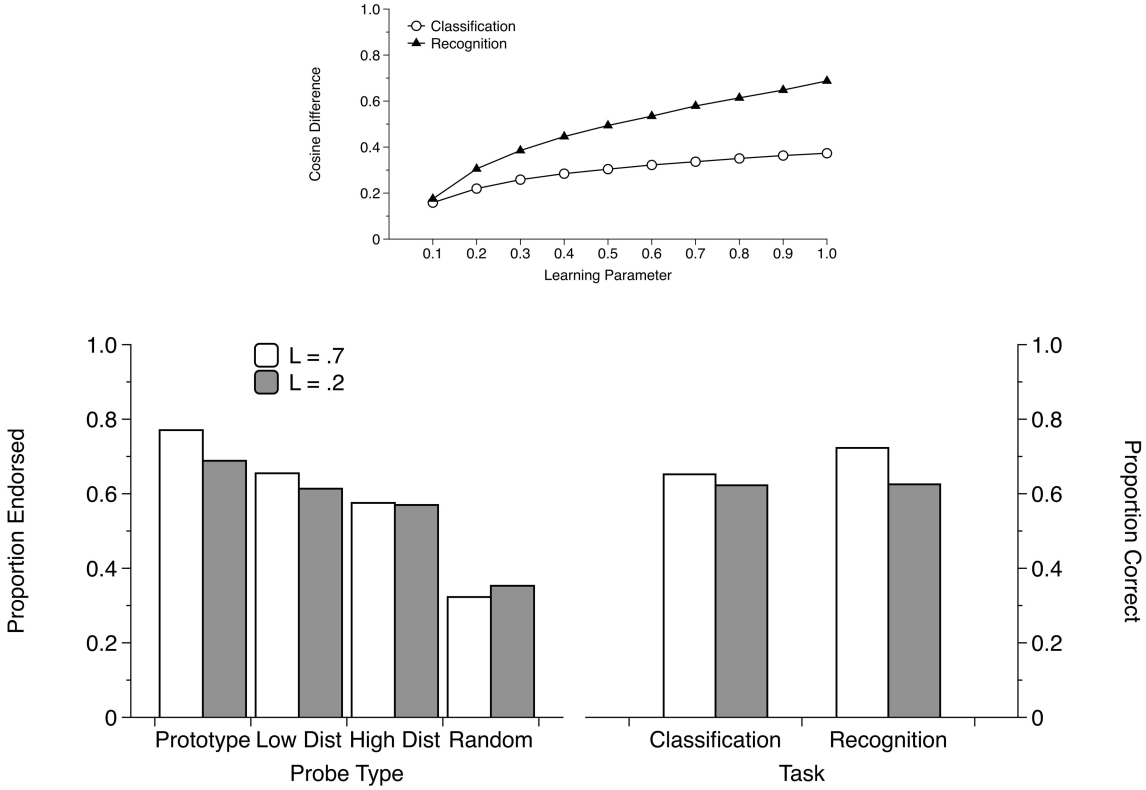

Figure 7. *Simulation 1 endorsement rates and proportion correct on Knowlton and Squire's (1993) classification and recognition tasks. The bottom left panel presents the model's proportion of items endorsed in the classification task as a function of L and probe type. The bottom right panel presents 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.*

Although not shown in Figure 7, the difference in recognition performance as a function of *L* is caused by both an increase in hit rates and a decrease in false alarm rates. Unfortunately, Knowlton and Squire (1993) did not report hits and false alarms so I cannot determine if that prediction is consistent with the data they collected. However, that result is consistent with a mirror effect, a hallmark of recognition data (Glanzer & Adams, 1990).

The inset panel in Figure 7 shows how the model accommodates the pattern of results. Increasing *L* results in an increase in discrimination on both the classification and recognition tasks. However, discrimination increases faster on recognition than on classification as a function of *L*.

Discussion of Simulation 1

In conclusion, the simulation shows that the HEM – an adaptation of a standard theory for memory that assumes memory for studied exemplars without abstraction or 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 additional grounds on which to argue for the single-system perspective on memory. The analysis also shows that the differences in experimental procedure for measuring recognition and classification are far from inconsequential and, in fact, a consideration of those differences is important for explaining the nature and form of selective impairments.

Generalization to Alternative Similarity Structures

Knowlton and Squire's (1993) results generalize to stimuli with different similarity structures. Posner and Keele's (1968) patterns are structured such that the differences among patterns are continuous, indexed by the distance moved per dot. However, the experimental method and my computational approach can be extended to similar observations using stimuli defined by discrete rather than continuous feature values.

Reed, Squire, Patalano, Smith, and Jonides (1999) tested amnesic patients' and control participants' abilities to learn and categorize line drawings of animals. Each of the animals was composed of nine features; each feature could take one of two values (see Figure 8 for examples). For example, an animal's neck could be short or long and its body markings could be striped or spotted. Thus, the similarity between any two animals can be measured by counting the number of matching versus mismatching features.

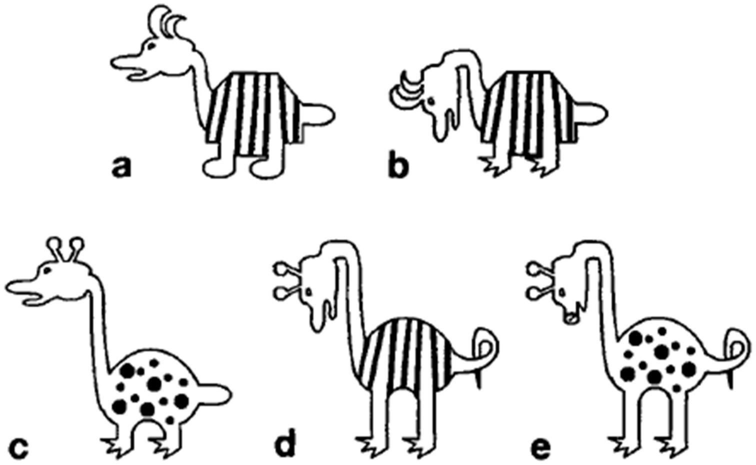

Figure 8. *Examples of animal line drawings. Animal a is the prototype, b is a low-distortion exemplar, c is a neutral exemplar, d is a high-distortion exemplar, and e is the anti-prototype. Reproduced with permission from Reed, Squire, Patalano, Smith, and Jonides (1999), Learning about categories that are defined by object-like stimuli despite impaired declarative memory, in Behavioral Neuroscience, 133(3), pages 411-419, published by APA.*

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, the *anti-prototype* that differed on all nine features from the prototype, low-distortion exemplars that differed by two or three features from the prototype, moderate-distortion exemplars that differed by four or five features from the prototype, and high-distortion exemplars that differed by seven or eight features from the prototype.

Figure 9 presents Reed et al.'s (1999) results. The left panel presents endorsement rates on the classification task. The numbers under the x-axis indicate the number of features that matched those in the prototype. As shown, controls showed the standard typicality gradient observed in other classification tasks. 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.

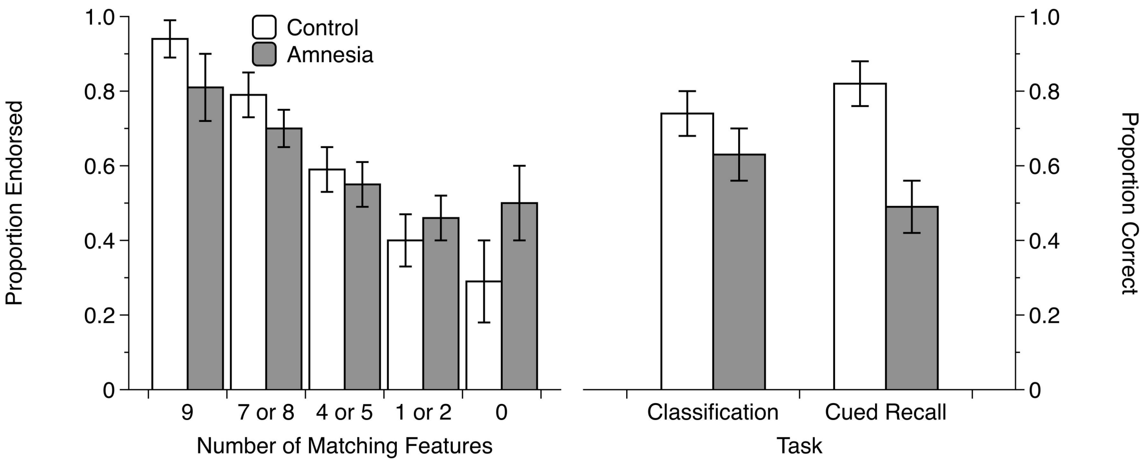

Figure 9. *Reproduced results from Reed et al.'s (1999) classification and cued recall tasks. The left panel presents endorsement rates in classification for control participants and amnesic patients as a function of the number of features matching the prototype. The right panel presents percent correct on the classification and cued recall tasks. Reproduced with permission from Reed, Squire, Patalano, Smith, and Jonides (1999), Learning about categories that are defined by object-like stimuli despite impaired declarative memory, in Behavioral Neuroscience, 133(3), pages 411-419, published by APA.*

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 9 presents overall accuracy on the classification and the cued recall tasks. Controls are presented in black and amnesic patients are presented in white. 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.

Extending the Generalized Context Model

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 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 10 presents their replication. The left bars present accuracy on the classification task and the right bars present accuracy on the recognition task. Participants tested without delay are presented in black and participants tested after a delay are presented in grey. 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.

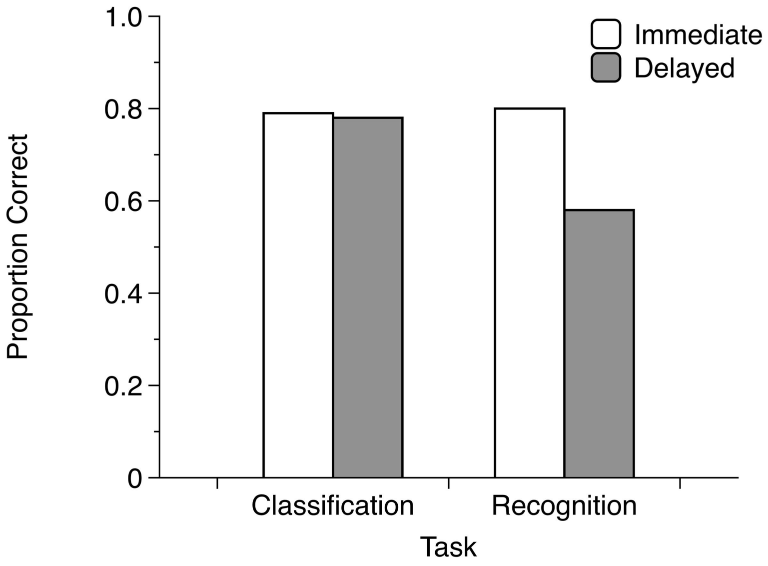

Figure 10. *Reproduction of Zaki and Nosofsky's (2001) replication of Reed et al. (1999), presenting proportion of correct trials for classification and recognition after no delay and a week-long delay. Reproduced with permission from Springer, license number 3884970722026.*

Simulation 2

I conducted a simulation to evaluate a single-system account of the difference between recognition and classification of binary-featured animal drawings in the HEM. The model extends the GCM's analysis by simulating the influence of L on both tasks, permitting a full analysis of the results.

In Reed et al.'s (1999) classification test, participants studied 20 drawings that were low-level distortions of a common prototype twice each. Following study, participants were tested for classification of 96 unstudied drawings that included 12 instances of the prototype, 20 low-level distortions, 20 moderate-level distortions, 20 high-level distortions, and 12 instances of the anti-prototype.

In Zaki and Nosofsky's (2001) corresponding recognition test, participants studied five drawings made from random features. Following study, they were tested for recognition of those five drawings relative to five unstudied drawings.

There was a pattern of selective impairment across experiments. Reed et al. (1999) reported that amnesic patients performed as well as controls in classification. In contrast, Zaki and Nosofsky (2001) reported that simulated amnesic patients (i.e., healthy participants after a delay between study and test) performed worse than controls (i.e., healthy participants tested immediately after study) in the recognition test.

To simulate the classification test, I generated eighteen 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 $1/\sqrt{n} = 0.1$. Each vector represented one of two values for each of the nine binary features.

To generate animal-drawing representations, I convolved the first two feature vectors – the method for holographic reduced representation. Next, I convolved the resulting vector with the next feature vector. I continued this iterative convolution for all nine feature vectors, normalizing each vector to a range of -1 to +1 prior to the next convolution[1]. As an example, consider an animal with four prototypical features, $\boldsymbol{A_P}$, $\boldsymbol{B_P}$, $\boldsymbol{C_P}$, and $\boldsymbol{D_P}$. To represent the combination of those four features, I would first take the convolution of $\boldsymbol{A_P}$ and $\boldsymbol{B_P}$. I would then normalize the resulting vector to a range of -1 to +1. Then, I would take the convolution of the new vector and $\boldsymbol{C_P}$. I would then normalize the resulting vector to a range of -1 to + 1. Then, I would take the convolution of the

---

[1] If the vectors are not normalized, the length of the vector increases with each convolution. The last vectors contribute more information to the convolution as a result. Normalizing the composite vector prior to convolution ensures that each component vector makes an equal contribution.

newest vector with $\boldsymbol{D_P}$. Finally, I would normalize the resulting vector to a range of -1 to +1. That vector would serve as the representation of the prototype.

To generate an exemplar that differs only on the final feature (i.e., the “D” feature), the series of convolutions would be identical, but the final convolution would involve the vector $\boldsymbol{D_A}$, representing the anti-prototypical feature value. By generating exemplars by replacing features in the convolution, the representations map directly onto the structure of the items in the original experiment.

I generated 20 low-distortion study exemplars by replacing either one or two of the prototypical features. I stored each exemplar to a memory matrix twice each. I then generated 20 additional low-distortion exemplars using the same methods to serve as test exemplars. I generated 20 neutral exemplars by replacing either four or five of the prototypical features. I generated 20 high-distortion exemplars by replacing either seven or eight of the prototypical features. Finally, I generated an anti-prototype by replacing all nine of the prototypical features.

Next, to simulate classification, I 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, I computed the cosine similarity between the echo and the test item. I converted the cosine similarities into response probabilities using the same logistic transformation as Simulation 1. I conducted 1000 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, I generated five study items and five unstudied items by convolving random selections of features. I stored each study item to a memory matrix eight times each. Next, I 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, I computed the cosine similarity between echo and test item, and computed recognition performance using the same logistic function as the classification task. I conducted 1000 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 11 presents the results from the model where $a = 9.8$, $b = 5.0$ for classification, and $b = 9.3$ for recognition. Amnesic patients were simulated with $L = 0.2$ whereas controls were simulated with $L = 0.7$. The bottom left panel presents the model's endorsement probabilities in the classification task. The bottom 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.

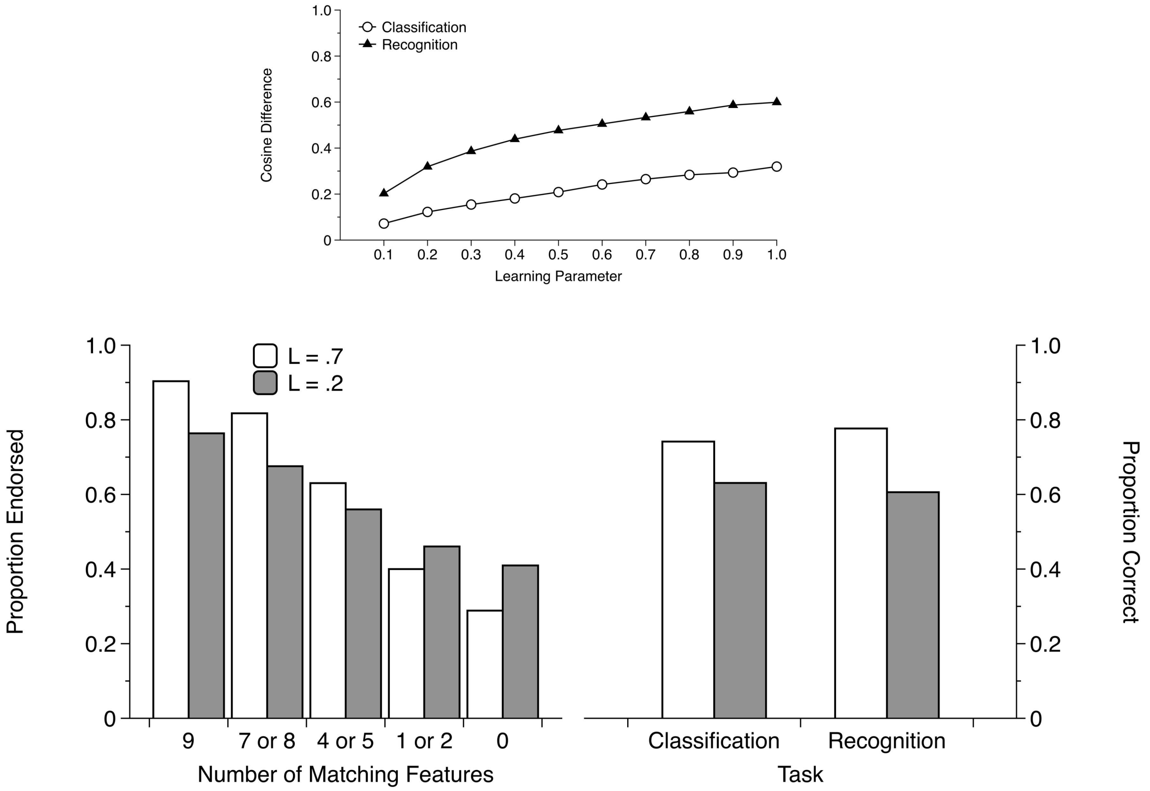

Figure 11. *Simulation 2 endorsement rates and proportion correct on Reed et al.'s (1999) classification task and Zaki and Nosofsky's (2001) recognition task. The bottom left panel presents the model's proportion of items endorsed in the classification task as a function of L and probe type. The bottom 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 very 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 exemplars, moderate-distortion exemplars, high-level distortion exemplars, and finally the anti-prototype. More importantly, the model predicts the difference in impairment. The change in *L* had a stronger influence on recognition performance than on classification performance. The model corresponds closely with the empirical means ($r = .98$).

In summary, the model accommodates a difference between classification and recognition that is consistent with the difference that Reed et al. 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 a multiple-system conclusion.

Discussion of Simulations 1 and 2

Thus far I have presented two simulations examining selective patterns of impairment in classification and recognition in the framework of the HEM, an exemplar-based model of memory. I simulated amnesia as a global deficit to the quality of encoding rather than selective damage to an explicit memory system. The key difference between the two simulations is the category structures of the two sets of stimuli. In Simulation 1, the stimuli differ on a continuous measure. In Simulation 2, the stimuli differ on discrete features. The model accommodates performance under both category structures. The computational analysis questions the necessity of a systems distinction in memory theory by demonstrating that a single-system account accommodates differences between classification and recognition.

Both patterns naturally fall out of the model's retrieval mechanisms. When a probe is presented to memory during recognition, the majority of the activation contributing to the echo comes from the single corresponding trace. However, if that trace is sparse, it does not contribute sufficient activation to support recognition. In contrast, when a probe is presented to memory during classification, the echo is composed of cumulative activation from each of the traces, which are all similar as a result of belonging to the same category. Even when the traces are sparse, the combined activation remains sufficient to support classification (see Brooks & Vokey, 1991).

Although the HEM circumvents problems associated with the GCM, the models share common principles. Both models assume that classification is supported by exemplar retrieval. Both models explain the difference between classification and recognition without recourse to a systems distinction. Additionally, the two models accommodate the differences using related mechanisms. In the GCM, amnesia is represented by low sensitivity to similarity as represented by *c*. In MINERVA2, amnesia is represented by poor encoding quality represented by *L*. The two parameters differ in detail but are conceptually similar. As *L* decreases, traces become sparse and populated with missing elements. As a result, the traces contain fewer informative features and become less distinguishable from one another. By this conceptualization, decreasing *L* decreases the model's sensitivity to the similarity structure among the traces, corresponding to the drop of *c* used to model amnesia in the GCM.

Rather than view the GCM and the HEM as competitors, I argue that it is more beneficial to consider their converging contributions. Both models provide demonstrations that differential patterns of impairment between classification and recognition do not force a systems distinction but they do so at different levels of abstraction. The HEM provides a formal description of the cognitive processes underlying memory performance in greater detail than the GCM. However, the models stem from the same fundamental principle that both explicit and implicit memory are the result of retrieval across a set of stored exemplars. That is, both theories assume that the impairment patterns recognition and classification reflect differences in the demands and structure of the tasks, not differences in the systems and mechanisms that are applied to negotiate those demands (see Nosofsky, 1985, 1987, for the classic argument of this position in relation to recognition and identification).

In the next section, I turn to selective impairments among priming, recognition, and cued recall. First, I discuss priming of tachistoscopic identification. Second, I discuss priming of word-stem completion.

Priming of Tachistoscopic Identification

Selective patterns of impairment in recognition and classification are one primary source of evidence taken to support the multiple-system perspective. In the following sections, I turn to another primary source of evidence.

In early investigations of amnesia, Warrington and Weiskrantz (1970, 1974, 1978) demonstrated that studying words facilitated patients' subsequent identification when the words were only partially presented. This priming is taken as additional evidence for spared implicit memory in amnesic patients.

Priming in amnesia is typically examined with two different experimental procedures. In one procedure, participants identify tachistoscopically presented words. Priming is observed when previously studied words are more likely than unstudied words to be identified correctly. In another procedure, participants are presented with the first few letters of a word (i.e., a word stem) and provide the first complete word that comes to mind. Priming is observed when participants are more likely to offer completions that were words in the studied list.

A classic interpretation of intact priming in amnesia is that studied words activate pre-existing lexical representations in memory. That activation persists into the test phase, facilitating identification or completion of those activated words (e.g., Bowers, 2000; Diamond & Rozin, 1984; Scarborough, Cortese, & Scarborough, 1977). However, this view has been challenged based on observations of priming for non-word stimuli that are unlikely to have pre-existing representations in memory (e.g., Cermak, Talbot, Chandler, & Wolbarst, 1985; Cermak, Verfaellie, Milberg, Letourneau, & Blackford, 1991; Feustel, Shiffrin, & Salasoo, 1983; Jacoby, 1983; Jacoby & Dallas, 1981; Musen & Triesman, 1990; Schacter, Cooper, & Delany, 1990). A resulting alternative view is that priming is supported by information learned in the experiment and stored in an implicit memory system.

In support of this view, Keane, Gabrieli, Noland, and McNealy (1995; see also Hamann & Squire, 1997) had amnesic patients and control participants study 24 three-letter consonant strings (e.g., XVL, JWN, LPS). Participants then completed an identification task in which the studied strings and 24 novel strings were presented tachistoscopically (for either 33 or 50 ms). Participants were asked to identify the strings verbally. In a subsequent recognition task, participants were tested for recognition of studied versus unstudied strings, where the strings were presented clearly and without a time limit until a recognition decision was registered.

Figure 12 presents Keane et al.'s (1995) results. The left panel presents the percentage of strings correctly identified in the tachistoscopic identification task. The right panel presents the probability of endorsement in the recognition task. Performance on the studied strings is presented in light single-hatched bars; performance on the unstudied strings is presented in dark double-hatched bars.

As shown in the left panel of Figure 12, amnesic patients showed similar priming as control participants (i.e., the difference between bars is similar). In contrast, the right panel shows that amnesic patients were much worse at recognizing the letter strings (i.e., the difference between bars is larger for the control participants than the amnesiac patients). The pattern illustrates another impairment difference between tasks and provides another challenge for the single-system perspective.

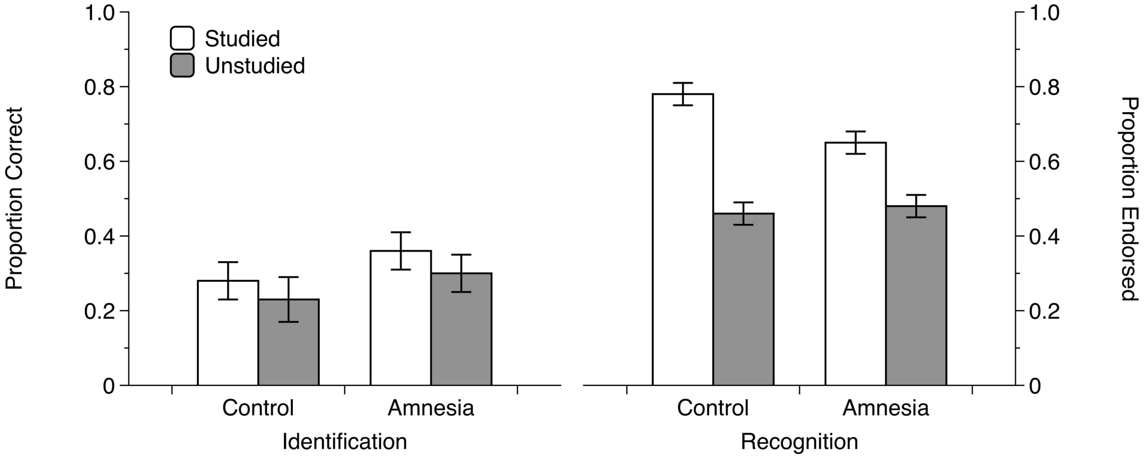

Figure 12. *Results from Keane et al. (1995). The left panel presents the percent of correctly identified strings as a function of probe type for control participants and amnesic patients. The right panel presents the proportion of strings endorsed as a function of the same variables. Reproduced with permission from Keane, M. M., Gabrieli, J. D., Noland, J. S., & McNealy, S. I. (1995). Normal perceptual priming of orthographically illegal nonwords in amnesia. Journal of the International Neuropsychological Society, 1(5), 425-433, published by Cambridge University Press.*

Simple Recurrent Networks

Kinder and Shanks (2001, 2003) provided a single-system account of identification priming using a Single Recurrent Network (SRN), a neural network model designed to learn sequential information (Cleeremans, 1993; Cleeremans & McClelland, 1991; Elman, 1990). The SRN is a multilayered neural network containing an input layer, an output layer, a hidden layer, and a context layer (see Figure 13). The input layer represents the current stimulus, the hidden layer represents the internal representation of that stimulus, and the output layer represents the network's response. The context layer copies the hidden layer to the input layer from trial to trial.

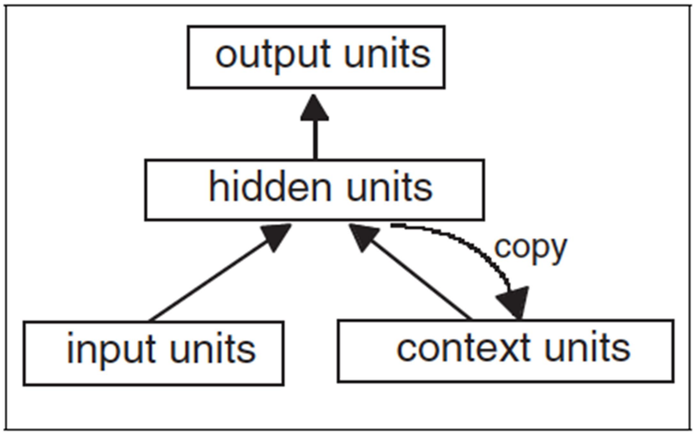

Figure 13. *The basic architecture of the SRN. From Kinder, A., & Shanks, D. R. (2001). Amnesia and the declarative/nondeclarative distinction: A recurrent network model of classification, recognition, and repetition priming. Journal of Cognitive Neuroscience, 13(5), 648-669. Reproduced with permission from The MIT Press.*

The SRN learns by continuously predicting the next item in a sequence; in this case, it predicts the next letter in a string. Each possible letter is represented by a unit in the input layer. When a letter is presented, that unit is set to an activation of .9 and all other units are set to activations of .1. The network predicts the next letter and, if incorrect, learns via a standard backpropagation algorithm. The algorithm includes a learning parameter controlling the rate of learning, which Kinder and Shanks (2001,

2003) vary to simulate the difference between amnesic patients and controls. By their analysis, amnesic patients learn in the same way as controls but more slowly.

Recognition and identification are both indexed by the network's ability to predict test sequences. For each string, the output vectors (i.e., the activations of the units in the output layer) for each predicted letter are concatenated together. This new vector is compared to the vector for the actual test string, and an endorsement / identification probability is computed based on their cosine similarity. The probability of endorsement / identification increases with the network's predictive accuracy. In the spirit of a single-system account, recognition and identification are simulated identically, but with one exception. To simulate tachistoscopic presentation in the identification test, the active input nodes are set to .7 instead of .9, decreasing the strength of the signal produced by the input string. Similarly, the inactive nodes are set to .3 instead of .1, increasing the noise in the signal produced by the input string. As a result, the input strings are less discriminable from one another.

The SRN accommodates the results by predicting a minimal influence of amnesia on priming but a large influence of amnesia on recognition, where amnesia is modelled by a change in the model's learning rate. The model's ability to produce the pattern of impairment without recourse to a systems distinction, like similar demonstrations with the GCM, MINERVA2, and the HEM, provides yet another demonstration that a result thought to force a multiple-system account is equally consistent with a single-system account. However, the SRN has two major limitations.

First, the SRN is limited in scope, designed specifically to learn sequential structure, and it is difficult to generalize the model to stimuli that are not sequential in nature (e.g., dot patterns and animal drawings). However, despite limited generalizability, the model's success provides a useful albeit restricted explanation of selective patterns of impairment that involve priming.

Reber (2002) identified a more serious limitation to the SRN. Training in the model involves gradual learning in multiple epochs (i.e., passes through the stimulus list). In Kinder and Shanks' (2001, 2003) simulations, each item is studied 100 times with different stimuli interleaved together. This learning schedule is essential to the model, because rapid non-interleaved learning results in a phenomenon known as catastrophic interference, where new information overwrites old information (McClelland, McNaughton, & O'Reilly, 1995; McCloskey & Cohen, 1989; Ratcliff, 1990; however see French, 1991). However, participants in an experiment typically study each item once. If the learning schedule of the model mismatches the learning schedule in the experiment, one has grounds to question the conclusions drawn from the analysis.

The goal of Simulation 3 is to model the difference between recognition and identification priming in the HEM. The simulation extends the analysis grounded in the SRN to the HEM, a more flexible model based on MINERVA2, which has been applied to a suite of memory phenomena. The simulation also matches the learning schedule of the experiment by storing a single trace for each stimulus

Simulation 3A

Simulation 3A models the difference between recognition and identification priming reported by Keane et al. (1995). In their experiment, both amnesic patients and controls were tested for recognition and tachistoscopic identification of letter strings.

In both tests, participants studied 24 letter strings. Following study, participants were presented with the studied strings and 24 unstudied strings. In the tachistoscopic identification test, the strings were presented very quickly and the participants reported the string back. In the recognition test, the strings were presented until a response was registered and participants reported whether the string had been studied. Keane et al. (1995) reported a result where amnesic patients showed the same priming as controls in the identification test but performed worse than controls in the recognition test.

To simulate the identification test, I generated a random vector for each of the 26 letters in the alphabet. Each vector was a 100 element vector with each element taking a random number sampled from a normal distribution with a mean of zero and a standard deviation of $1/\sqrt{n}$ = 0.1.

I generated strings by selecting three random consonants and convolving their vector representations using the same iterative circular convolution as Simulation 2. For example, to represent the string PVX, I took the circular convolution of the vectors corresponding to P and V. Next, I normalized the composite vector and took the circular convolution of the composite PV vector and the vector corresponding to X. Finally, I normalized the new vector which served as the representation of the string PVX.

However, using circular convolution to represent strings of letters is problematic for materials like letter-strings that participants compulsively read from left to right. Circular convolution is commutative. As a result, the representations of the strings, PVX, PXV, XPV, XVP, VPX, and VXP are identical. To solve the problem, I used a non-commutative version of circular convolution in which the indices of the letter vectors are scrambled depending on the relative position of the letters (e.g., Jones & Mewhort, 2007).

As a result, the string PV is not identical to the string VP, and the sequential information in each string is maintained in its representation. Of course, commutativity is less of a problem, if one at all, for the encoding of the animal features in the preceding simulation – people do not read legs, neck length, body markings, and so on in some prescribed or overlearned order like they read strings of letters from left to right.

I generated 24 study strings and stored them to a memory matrix. Those vectors served as the studied strings in the test phase. I also generated 24 additional strings to serve as the unstudied strings in the test phase.

The HEM has not yet been applied to a tachistoscopic identification task. Therefore, there is currently no mechanism to simulate rapid presentation. However, I adopted a mechanism from Mewhort and Johns (2005). They adapted the retrieval assumptions in MINERVA2 by assuming that retrieval sharpens over time, reducing the influence of noise from interfering items until the signal is sufficient to produce a response. They simulated this process by manipulating the exponent, τ, in the formula for trace activation. The parameter begins at 1, where activation is equal to similarity. The parameter increases over time, attenuating the influence of moderately similar traces and thus reducing retrieval interference.

I adopted this mechanism to simulate tachistoscopic presentation of test strings. When strings are presented very quickly, the retrieval process does not have sufficient time to sharpen. Thus I set $\tau = 1$ to simulate the identification test phase. I computed the echo for each of the 48 test items (i.e., 24 studied items and 24 unstudied items). Then, I computed the cosine similarity between the echo and the test item. Finally, I converted the cosine similarities into a correct-identification probability using the same logistic transformation as the preceding simulations. I conducted 1000 simulations for each of 10 levels of $L = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0$.

I simulated the recognition test nearly identically to the identification test, except that I set $\tau = 3$ at retrieval. I converted cosine similarities into response probabilities using the same logistic transformation as the identification simulation. I conducted 1000 simulations for each of 10 levels of $L = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0$.

Results. Figure 14 presents the results from the model where $a = 2.3$, $b = 2.3$ for the identification test, and $b = 0.9$ for the recognition test. Amnesic patients were simulated with $L = 0.3$ whereas controls were simulated with $L = 0.6$. The left panel presents the proportion of strings accurately identified. The right panel presents the proportion of strings endorsed as “Old”. The inset panel presents the differences in cosine similarities on the two tasks as a function of L.

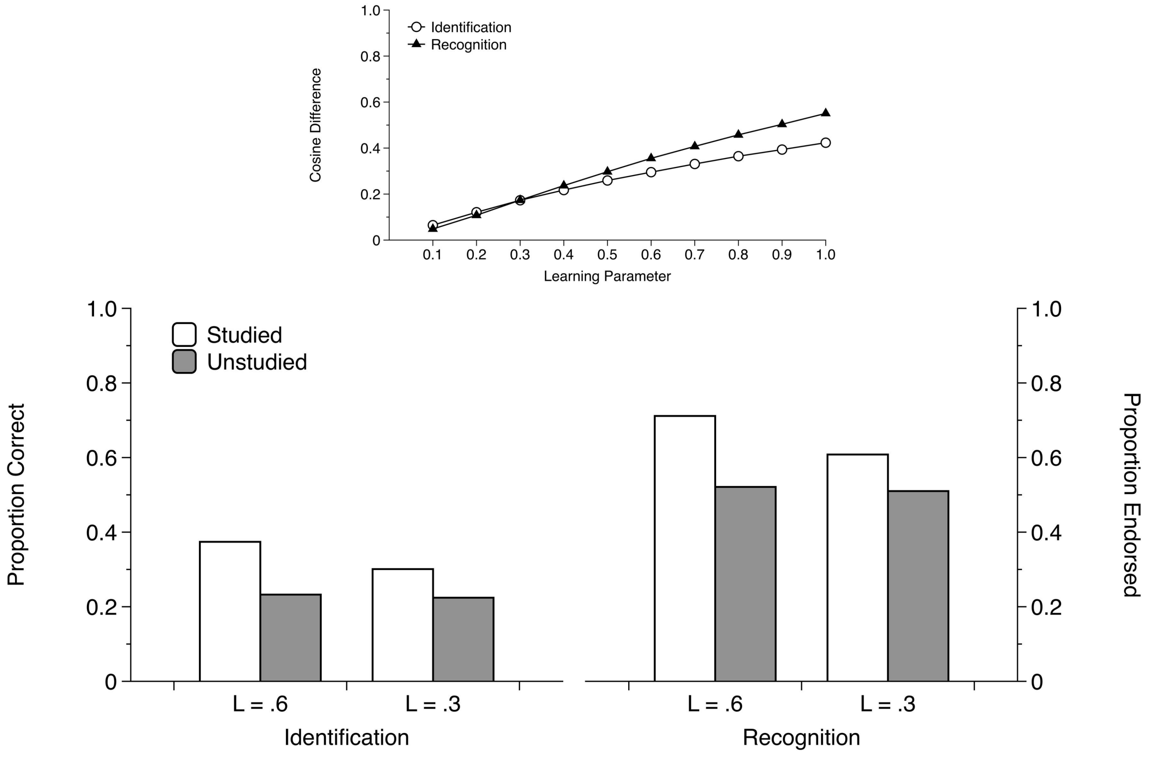

Figure 14. *Simulation 3A identification accuracy and recognition endorsement probabilities on Keane et al.'s (1995) tasks. The bottom left panel presents the model's proportion of items correctly identified in the identification task as a function of L and probe type. The bottom right panel presents proportion endorsed in the recognition task as a function of L and probe type. 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 data quite well (*r* = .94). However, whereas the model accurately predicts the core aspects of the data, it predicts too much priming in the control group, and lower overall identification accuracy in the amnesic group. The model also predicts better recognition performance for both groups than in the empirical data. These patterns represent points of weakness in the quantitative fit, but the inaccuracies in absolute fit are relatively minor when one considers that the model captures the overall pattern of performance quite well. It is also minor when one considers that the model produces the critical pattern of impairment: a larger influence of encoding quality on recognition performance than on tachistoscopic identification performance.

Despite minor deviations from the exact means in the empirical data, the model produces the pattern of impairment originally taken to force a multiple-system account. The model analysis generalizes from dissociations involving classification (see Simulations 1 and 2) to a result involving priming. The framework converges with Kinder and Shanks' (2001, 2003) analysis, extending to a more general and flexible computational approach. However, there is an issue that grows from the data and analysis that I have not yet raised or considered.

Keane and colleagues presented additional data that show two patients with occipital lobe damage show the opposite pattern of impairment as the amnesic patients they generally study. Namely, patients M.S. (Gabrieli, Fleischman, Keane, Reminger, & Morrell, 1995) and L.H. (Keane, Gabrieli, Mapstone, Johnson, & Corkin, 1995) both show intact recognition but reduced priming. In conjunction with amnesic patients, the data from the occipital lobe patients present a double dissociation between recognition and priming.

Kinder and Shanks (2001, 2003) argued that the occipital lobe patients' performance can be explained by a visual deficit that results in a globally decreased ability to identify tachistoscopically presented stimuli. They tested this proposition in the SRN by decreasing the value of active units (signal strength) at study and increasing the value of inactive units (signal noise) at test in order to additionally degrade the input stimuli.

Under this implementation, the SRN accommodates the double dissociation. However, these simulations are subject to the same objections as preceding simulations using the SRN (i.e., the mismatch between study procedures in the experiments and in the simulations of the experiments). Therefore, the goal of Simulation 3B is to extend the analysis to the HEM.

Simulation 3B

Gabrieli et al. (1995) and Keane et al.'s (1995) procedures are difficult to simulate in the MINERVA2 framework. In both experiments, items were presented at increasing durations until the participant was able to accurately identify the item. The model could plausibly be adapted to accommodate the procedure, but not without substantial augmentation and additional assumptions. Thus, following Kinder and Shanks (2001, 2003), Simulation 3B approaches the double dissociation at a conceptual level, simulating a procedure consistent with Simulation 3A (i.e., strings presented a single time, producing a proportion of strings identified).

Simulation 3B was nearly identical to Simulation 3A. However, to simulate occipital lobe damage and, presumably, visual impairment, I set τ to a value of less than one in the identification task. To simulate patient M.S., I set $\tau = 0.1$. To simulate patient L.H., I set $\tau = 0.001$, consistent with the observation that L.H.'s impairment was more severe than M.S.' impairment.

Using non-integer values of τ introduces a computational complication; when the cosine similarity is negative, solving the root involves an imaginary number. To avoid this issue, I computed activation based on the absolute value of the cosine similarity but maintained the appropriate sign.

I did not simulate the recognition task; given that there are no rapid visual demands in the recognition task, the parameter values – and therefore simulation results – would be identical for simulated control participants and simulations based on the two patients.

Results. Figure 15 presents the simulation results. Each set of bars presents the proportion of strings accurately identified for the simulated control participants, amnesic patients, M.S., and L.H., from left to right respectively.

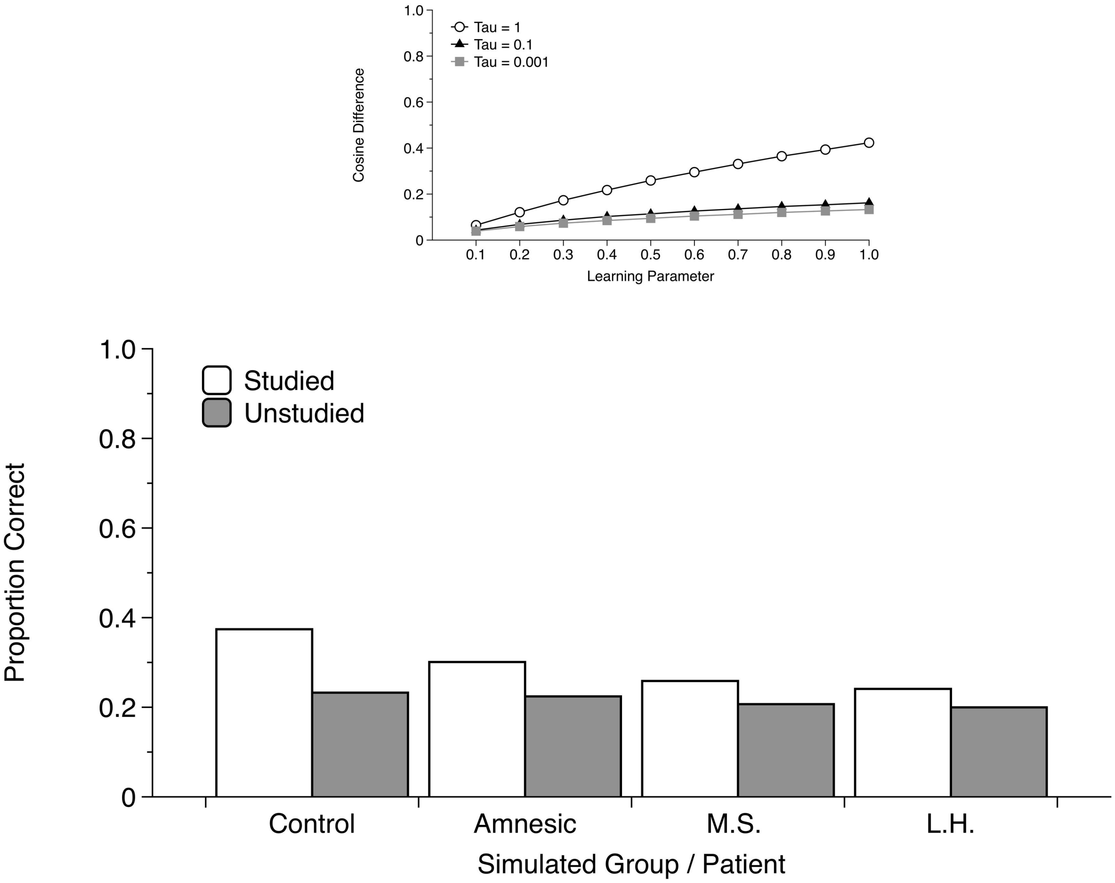

Figure 15. *Simulation 3B identification accuracy in the conceptual simulation of Gabrieli et al. (1995) and Keane et al. (1995). The bottom panel presents the proportion of strings accurately identified in the four simulated groups as a function of probe type. The top panel presents the difference between cosine similarities for to-be-endorsed and to-be rejected items in both tasks as a function of L.*

As shown, the model appears to show some evidence of being able to accommodate the double dissociation. When τ is permitted to be less than 1, overall identification accuracy and, consequently, the difference between studied and unstudied items, decrease as well. However, the difference in priming between the simulated occipital lobe patients and the simulated amnesic patients is not very convincing. Although there is a difference, it is small, only decreasing from .08 to .04. This difference is, in fact, smaller than the difference between simulated control participants and amnesic patients (.14 to .08). Thus the model fails to produce the double dissociation convincingly.

Simulation 3B represents the first failure of the model. However, it is hard to know what to make of that failure. It would have been more surprising if a memory model was able to explain performance in occipital lobe damage in the same way that it was able to explain performance in amnesia. It is also risky to draw conclusions about a model's validity based on its ability to fit the data of single-subject special cases. I will return to the issue when I discuss the meaning and value of model fit in psychology in the general discussion.

Discussion of Simulation 3

I have presented two additional simulations of selective patterns of impairment in the HEM framework. Unlike Simulations 1 and 2, which examined differences between classification and recognition, Simulation 3 examined differences between identification priming and recognition. Simulation 3A successfully accommodated original result involving amnesic patients. However, Simulation 3B failed to accommodate the corresponding double dissociation involving patients with occipital lobe damage.

As with the GCM, the HEM differs in computational detail from the SRN, but stems from the same basic premises. Both models provide a single-system interpretation of results thought to support the multiple-system perspective. Also similar to comparisons with the GCM, the HEM and the SRN explain the results through similar mechanisms. In the SRN, amnesia is simulated by reducing the model's learning rate. As a result, the network does not learn to predict the strings as accurately, and consequently develops incomplete representations of those strings. Thus the network's learning rate, like the sensitivity parameter in the GCM, has a similar impact as *L* in the HEM. I argue that the models be considered in convergence, providing explanations at different levels of abstraction but stemming from the same underlying perspective.

In the next section, I examine selective patterns of impairment involving priming of word-stem completion, another key example of results involving priming. Thus far, word-stem completion has not been the focus of computational analysis.

Priming of Word-Stem Completion

Graf, Squire, and Mandler (1984) reported a complex pattern of results involving free recall, recognition, cued recall, and word-stem completion. Amnesic patients and controls studied lists of words. Then, they were tested on their memory for the words using four different measurements. In free recall, participants named as many of the studied words as possible. In recognition, participants discriminated studied from unstudied words in a three-alternative-force-choice task. In cued recall, participants were given word stems (e.g., ele____) and completed the stem with a word from the studied list (e.g., elephant, electric). Word-stem completion was identical to cued recall except that participants were instructed to provide the first completion that came to mind rather than from the studied list. Priming was indexed by the difference between the probabilities of completion following study compared to a baseline completion rate from another set of participants.

Figure 16 presents a summary of Graf, Squire, and Mandler's (1984) results. The left panel presents the results from their Experiments 1 and 3, presenting accuracy in free recall, accuracy in cued recall, and priming in word-stem completion. Controls are presented in open bars and amnesic patients are presented in hatched bars. The right panel presents the results from their Experiment 2, presenting accuracy in recognition and priming in word-stem completion across three delays between study and test. Controls are presented in white points and amnesic patients are presented in black points. The top “chance” line presents chance performance in the recognition task and the bottom “chance” line present chance performance in the word-stem completion task.

As shown, controls performed better than amnesic patients on free recall, cued recall, and recognition. However, amnesic patients showed the same priming (even numerically greater priming in Experiments 1 and 3) as controls. This pattern is surprising given that amnesic patients were impaired in the cued recall task, which only differed from the completion task in instruction.

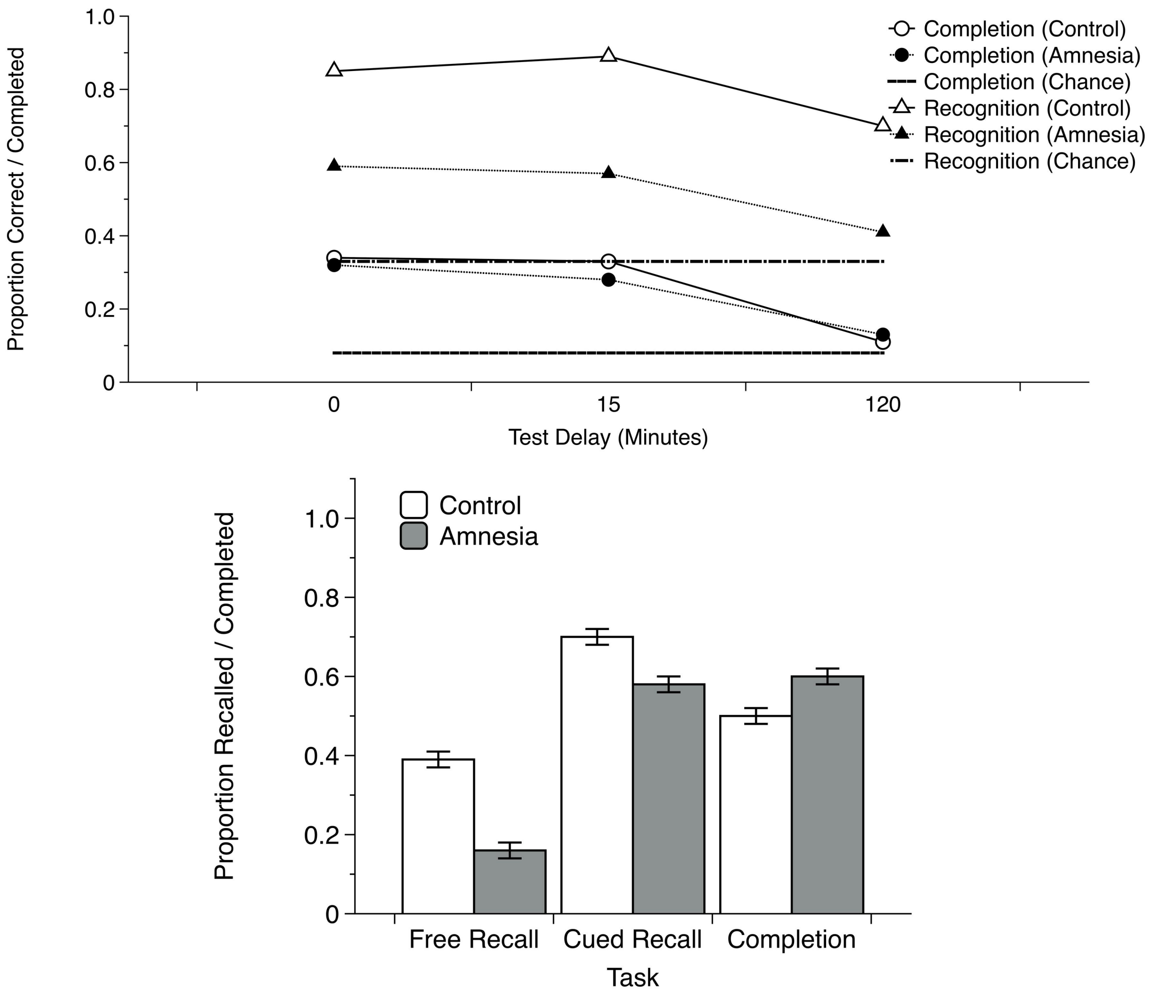

Figure 16. *Results reproduced from Graf, Squire, and Mandler's (1984) Experiments 1, 2, and 3, presenting performance on free recall, cued recall, word-stem completion, and recognition for control participants and amnesic patients. The top panel presents their Experiment 2, presenting accuracy on recognition and priming on completion at three delays between study and test. The bottom panel presents their Experiments 1 and 3, presenting accuracy on free and cued recall, and priming on completion. Reproduced with permission from Graf, Squire, and Mandler (1984), The information amnesic patients do not forget, in Journal of Experimental Psychology: Learning, Memory, and Cognition, 10(1), pages 164-178, published by APA***.**

The pattern has been replicated, generalizing across modalities (Graf, Shimamura, & Squire, 1986), generalizing to alternative stimuli such as pictures (Cave & Squire, 1992) and arithmetic facts (Delazer, Ewen, & Benke, 1997), persisting for novel associations (Graf & Schacter, 1985), and generalizing to other patient populations such as Alzheimer's disease and mild cognitive impairment (LaVoie & Faulkner, 2008).

The results provide yet another challenge to the single-system perspective and has yet to be analyzed in a computational framework. The goal of Simulation 4 was to produce the differential patterns of impairment between word-stem completion and putative explicit memory tasks in the HEM.

Simulations 4A and 4B

I conducted two simulations to evaluate a single-system account of Graf, Squire, and Mandler's (1984) demonstration that amnesia affects recognition and cued recall more than word-stem completion. Simulation 4A evaluates the difference between recognition and completion. Simulation 4B evaluates the difference between cued recall and completion. I did not simulate their free recall data because the HEM has no mechanism for free recall.

Simulation 4A. In the word-stem completion task, participants studied a list of 10 words twice. Following study, participants were presented with word stems from the 10 studied words and 10 unstudied words and were asked to provide the first complete word that came to mind. This procedure repeated four times. Each test took place either immediately, 15 minutes, or 120 minutes after each study session.

In the recognition task, participants studied a list of 20 words twice. Following study, they were tested on forced-choice recognition. On each trial, participants were presented with one studied word and two unstudied words and were asked to select the studied word. The test took place either immediately, 15 minutes, or 120 minutes after each study session.

Graf, Squire, and Mandler (1984) reported a pattern in which amnesic patients showed the same priming as controls in the word-stem completion test but performed worse than controls in the recognition test. The delay between study and test influenced both groups similarly.

To simulate the word-stem completion test, I generated 10 vectors of 100 elements with each element taking a random number sampled from a normal distribution with a mean of zero and a standard deviation of $1/\sqrt{n}$ = 0.1. I stored those 10 vectors to a memory matrix twice each. They were also studied words in the test phase. I generated another 10 vectors that were unstudied words in the test phase.

I simulated word-stem completion by removing information from the test items to create test stems. I removed information by converting each vector element to zero with one of four probabilities. The probabilities were based on the number of letters missing from the stem. The number of missing letters ranged between one and four, thus the four probabilities were .25, .40, .50, and .57. The probability was randomly selected from that set for each test stem.

I computed an echo for each of the 20 test items (i.e., the 10 studied and 10 unstudied items). Then, I computed the cosine similarity between the echo and test item. Finally, I converted the cosine similarity to a probability of completion using the same logistic transformation as the preceding simulations. I repeated this process four times and retained the preceding items in memory at each step. I conducted 1000 simulations at each level of L = 0.05, 0.1, .15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95, and 1.0. I used lower levels of L to simulate delays between study and test, representing the loss of information over time. I increased L in smaller intervals than previous simulations to provide a sufficiently detailed gradient to simulate amnesia and delay simultaneously.

To simulate recognition, I generated 20 vectors and stored them to a memory matrix twice each. To simulate the three-alternative-forced choice procedure, I generated an additional 40 vectors. On each simulated trial, I computed an echo from one of the studied items and two of the unstudied items. Then I computed the cosine similarity between each echo and the corresponding test item.

I computed the probability of selecting the studied item using a Luce choice axiom (Luce, 1957), which computes the probability of selecting an alternative as a ratio of the evidence for that alternative and the evidence for the full set of alternatives. Thus to simulate the three-alternative-forced-choice procedure, the probability of selecting the studied word, *p(“Studied”)*, is given by

$$p(\text{"Studied"}) = \frac{cos_{studied}}{cos_{studied} + cos_{unstudied} \ + cos_{unstudied}}$$

where $cos_{studied}$ is the cosine similarity between the studied word and its resulting echo, $cos_{unstudied1}$ is the cosine similarity between the first unstudied word and its resulting echo, and $cos_{unstudied2}$ is the cosine similarity between the studied word and its resulting echo. I conducted 1000 simulations at each level of L = 0.05, 0.1, .15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75, 0.8, 0.85, 0.9, 0.95, and 1.0.

Results. Figure 17 presents the model results where a = 4.3 and b = 4.5 for word-stem completion. There are no transformation parameters in the recognition test. I simulated amnesic patients with L = 0.3 after no delay, L = 0.25 after a 15 minute delay, and L = 0.1 after a 120 minute delay. I simulated controls with L = 0.7 after no delay, L = 0.65 after a 15 minute delay, and L = 0.5 after a 120 minute delay. The bottom left panel presents the proportion of stems completed in the completion task. The bottom right panel presents the proportion of correct trials in the recognition task. The inset panel presents the difference in cosine similarities as a function of *L*. As indexed by correlation, the model produces a close fit to the empirical data (*r* = .96). However, Figure 17 shows that the model does not produce the results convincingly. Although manipulating *L* had a numerically larger influence on recognition than on completion, the difference in magnitude of that influence is very small. More importantly, the model predicts a much larger difference in completion performance than is present in the empirical data. I conclude that the model fails to produce a pattern consistent with the original results.

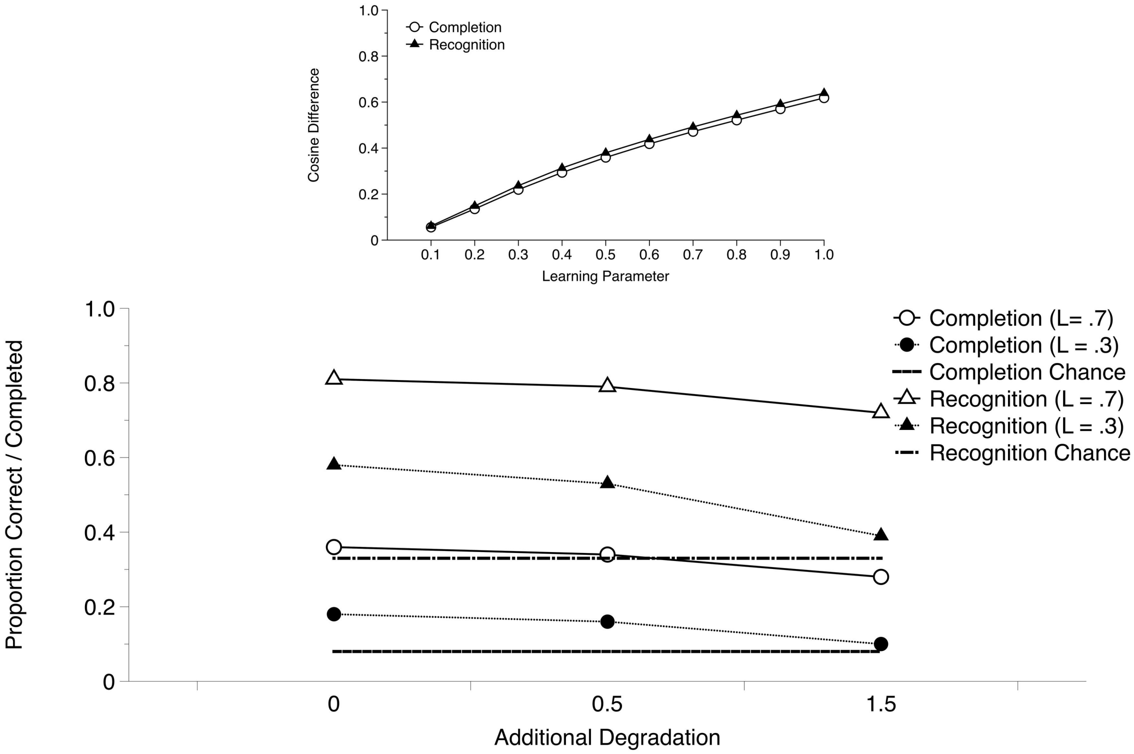

Figure 17. *Simulation 4A completion rates and endorsement probabilities on Graf, Squire, and Mandler's (1984) completion and recognition tasks. Presents the probability of offering a completion from the study list (circles) and the probability of selecting the correct word (triangles). The inset panel presents the difference in cosine similarities on each task.*

Simulation 4B. In Graf, Squire, and Mandler's (1984) cued recall task, participants followed the same procedure as in the word-stem completion test, except that they were asked to provide completed words from the studied list rather than the first word that came to mind. Graf, Squire, and Mandler reported a pattern of impairment such that amnesic patients showed the same priming as controls in the completion task but performed cued recall worse than controls.

Modeling differences between word-stem completion and cued recall presents an interesting challenge. The only procedural difference is in the task instructions; participants provide words from the study list in cued recall but they provide the first word that comes to mind in word stem completion. These instructions cannot be provided to a model in any direct sense.

To simulate cued recall, I incorporated contextual information into the stimulus representations (see Hintzman, 1986, 1988, for precedence). I also provided the model with pre-experimental memory traces of the same words. Whereas all of the items encountered in the experiment shared contextual information, items encountered outside the experiment did not.

My simulation of word-stem completion was the same as in Simulation 4A. I simulated the cued recall test similarly. However, I also generated a random vector to represent the experimental context. Then, I computed the circular convolution of each word vector with the context vector. I also computed the circular convolution of each word with a random vector from the same distribution to represent non-experimental experience with the words

I stored each prior experience to memory at $L = 1.0$ for both control participants and amnesic patients, representing the fact that memory is largely intact for pre-onset events in amnesia (see Postle & Corkin, 1998). I stored each studied item to memory at varying values of L using the same procedures as Simulation 4A.

I simulated cued recall identically to word-stem completion except that the test probes were convolved with the context vector. The context vector was not degraded prior to convolution.

Results. Figure 18 presents the results of the model where $a = 9.9$, $b = 7.8$. I simulated amnesic patients with $L = 0.3$ and controls with $L = 0.7$. The left panel presents the proportion of stems completed in the completion task. The right panel presents the proportion of words recalled in the cued recall task. The inset panel presents the difference in cosine similarities as a function of *L*.

The model produces a poor fit to the data (*r* = .43). The model's rates of completion and cued recall are far too low across all of the means, and the model predicts a large advantage in completion as a function of *L*. I conclude that the model clearly fails to produce the dissociation.

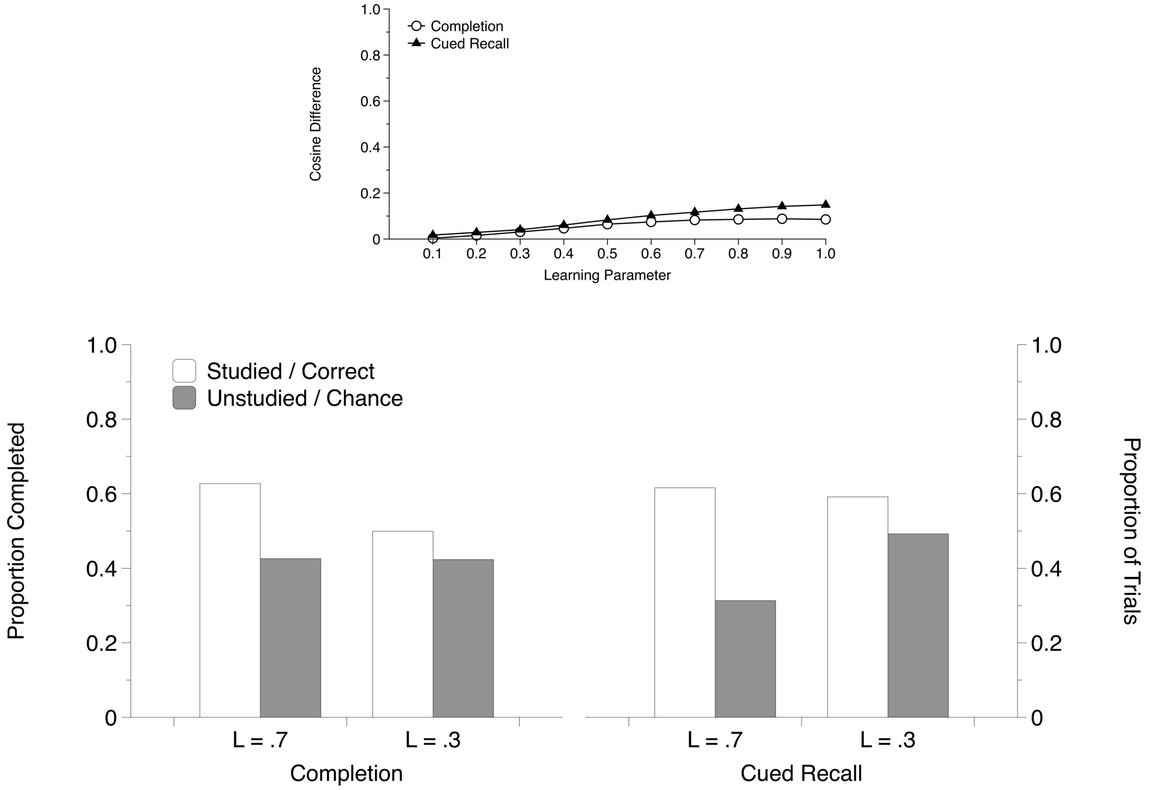

Figure 18. *Simulation 4B completion rates on Graf, Squire, and Mandler's (1984) completion and cued recall tasks. The bottom left panel presents the model's proportion of stems completed in the completion task as a function of L and probe type. The right panel presents the model's proportion of words accurately recalled in the cued recall task as a function of L and probe type. The inset panel presents the difference in cosine similarities on both tasks.*

Discussion of Simulation 4

The pattern of impairments within word-stem completion, cued recall, and recognition continue to pose a challenge to the single-system perspective. At least in the framework of the HEM, assuming a global deficit to memory results in a model that cannot produce the dissociation.

However, there are a number of principled reasons that the model might fail to produce the results. One reason is that word-stem completion and stem-cued-recall likely involve a complex combination of orthographic, phonetic, and semantic information. In contrast to the richly structured representations of words in natural human memory, the model’s representations are random and orthogonal computational conveniences. It is entirely possible that the random representations used in the model to stand for the richly structured materials used in the experiment present an over-simplified and, as a result, insufficient match to support human-like word completion. Unfortunately, there are no current methods that I could import to resolve the discrepancy.

Summary of Computational Simulations

I have reported four simulations in this chapter. The overarching goal of the simulations was to provide a formal account of selective impairments observed in amnesia within a single-system framework. By accommodating a variety of results using different experimental tasks and paradigms, my aim was to provide a computational analysis with a broad scope and general application to problems in memory research. Previous work has successfully accommodated many of these patterns. However, they provide explanations under different computational mechanisms and assumptions. My analysis attempted to explain each of the results under a common set of principles expressed in a single formalism.

The simulation results reflect a mixture of success and failure. In Simulations 1 and 2, I successfully modeled differences between classification and recognition. Simulation 1 modeled a result involving statistically-distorted dot patterns (Knowlton & Squire, 1993). Simulation 2 modeled a result involving binary-featured cartoon animals (Reed et al., 1999). Both simulations provided close fits to the empirical data, questioning the necessity of a multiple-system explanation. Additionally, in Simulation 3A, I successfully modeled a difference between priming of tachistoscopic identification and recognition (Keane et al., 1995). The model's success broadened the scope of the HEM approach, extending my argument from differences in classification and recognition to differences in priming and recognition.

However, the model failed to accommodate other results. In Simulation 3B, the model failed to produce a double dissociation between tachistoscopic identification and recognition involving both patients with amnesia and patients with occipital lobe damage (Gabrieli et al., 1995; Keane et al., 1995). Finally, in Simulation 4, the model failed to produce a pattern of impairments within word-stem completion, cued recall, and recognition (Graf, Squire, & Mandler, 1984). The model's failures represent challenges to the single-system perspective, at least in relation to how the perspective is instantiated in the form of the HEM theory.

When beginning the analysis, I based my preceding simulations on a key premise that the selective impairments observed in amnesia are not the result of a damage to a particular memory system. Instead, they reflect a global encoding deficit within a single system. The evidence in favour of that premise is mixed. To the extent that the model was successful, the results question the necessity of system distinctions in memory theory, instead favouring a more integrated account that explains memory by a common set of principles. To the extent that the model failed, the results pose a challenge to the single-system perspective, at least in the framework of the HEM. For the moment, I put the computational simulations aside and, in the following chapter, I extend my analysis from a computational framework to an empirical examination of the same basic premise.

Chapter 3 – Empirical Simulations: Inducing “Amnesia” in Healthy Memory

Chapter 2 provides an analysis of hallmark patterns of impairment using computational simulation. The main premise of the simulations is that patterns, specifically those observed in amnesia, can be explained as a global encoding deficit. In Chapter 3, I extend this analysis to empirical simulations by producing the same patterns of impairment in healthy participants.

Previous work has simulated amnesic behaviour using various methods. In early demonstrations, Woods and Percy (1974) replicated Warrington and Weiskrantz’s (1970) examinations of word priming. They simulated amnesia by introducing a week-long delay between study and test. In contrast, controls were only given a day-long delay. Participants were tested on three separate tasks: yes-no recognition, word-stem completion, and word-fragment (i.e., words were missing a proportion of pixels) completion. The simulated amnesic participants behaved like Warrington and Weiskrantz’s organic amnesic patients, performing similarly to controls on word-stem and word-fragment completion but performing worse on recognition. Nosofsky and Zaki (1998) used the same method to replicate Knowlton and Squire’s (1993) Results using dot pattern recognition and classification (see also Zaki & Nososfky, 2001). Like Woods and Percy’s (1974) results, introducing a long delay exerted a stronger negative impact on recognition performance than classification performance. Thus introducing a delay between study and test produces the same patterns observed in amnesia.

Other researchers have adopted alternative techniques. Graf, Mandler, and Haden (1982) replicated the difference between free recall and word-stem completion via a manipulation of encoding strategies. Participants studied words either by evaluating the number of vowels in each word or by rating how much they liked the word, a manipulation intended to encourage shallow versus elaborative processing, respectively. The groups performed similarly on word-stem completion, but the elaborative liking-rating group performed better on free recall. Thus, encouraging shallow processing at study can also produce the same patterns observed in amnesia.

Jamieson, Holmes, and Mewhort (2010) adopted yet another method to simulate amnesia. They replicated Knowlton, Ramus, and Squire’s (1992) results using recognition and classification in an artificial grammar task. Participants studied letter strings for either 6000 ms or 100 ms. The groups performed similarly on classification but the 6000 ms group performed better on recognition. Thus manipulating study time can also produce patterns observed in amnesia.

Differential patterns of impairment between memory tasks can be produced using a variety of methods. I adopt Jamieson, Holmes, and Mewhort’s (2010) method of manipulating study time. The manipulation maps cleanly onto the main premise in my computational analysis. By reducing study time I intended to produce a deficit in memory for the study items. In the simulations, reducing the learning parameter produces the deficit, resulting in incomplete traces in the model’s memory. I expected the study time manipulation to produce a parallel deficit, resulting in incomplete traces in participants’ memory.

I report four experiments replicating the same hallmark results I simulated in Chapter 2. In Experiment 1, I replicate Knowlton and Squire’s (1993) difference between recognition and classification of dot patterns. In Experiment 2, I conceptually replicate Reed et al.’s (1999) difference between recognition and classification of animal drawings. In Experiment 3, I replicate Keane et al.'s (1995) difference between recognition and priming of tachistoscopic identification and provide a conceptual analysis of Gabrieli et al.'s (1995) and Keane et al.'s (1995) double dissociation involving occipital lobe damage. In Experiment 4, I replicate Graf, Squire, and Mandler's (1984) differences between recognition, cued recall, and word-stem completion. In each experiment I produce amnesic-like behaviour by manipulating study time. All four experiments were approved by the Psychology / Sociology Research Ethics Board at the University of Manitoba. To foreshadow the outcome of the analysis, decreasing study time produced the expected resultsin Experiments 1, 2, and 3. In Experiment 4, decreasing study time produced the difference between word-stem completion and cued recall but not the difference between completion and recognition.

Experiment 1

Experiment 1 replicates 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. I replicated their procedure using healthy participants. However, half the participants studied each item for 5000 ms and half the participants studied each item for only 100 ms. 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 (5000 ms) group ought to perform better on recognition.

Methods

Participants. I 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 (5000 ms, 100 ms) and task (recognition, classification). All participants were given credit toward 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.

I constructed prototypes by randomly placing nine squares in a 30 x 30 matrix. Each square was black, filled, and 6 pixels by 6 pixels. I derived exemplars using the statistical techniques described in Posner, Goldsmith, and Welton (1967). First, I constructed and numbered a second 20 x 20 matrix. I numbered the central cell zero. I numbered the adjacent ring of cells one through eight. I numbered the cells in the next ring nine through 24, and numbering continued in this pattern up to cell 399. The rings formed five regions containing the cells presented in column 2 of Table 1. I centered the 20 x 20 matrix on a square in the prototype. Then, I moved the square to one of the five regions with specified probabilities depending on the degree of distortion (see columns 3 and 4 of Table 1). Movement within regions occurred to each cell with equal probability. I repeated this process for each square in the prototype, and if two squares occupied the same space the movement procedure restarted. Figure 19 presents examples of the materials.

Table 1. *The cells contained in each region of the 20 x 20 matrix and the probability of movement to each region for low-distortion and high-distortion exemplars in Experiment 1. Note that moving to Region 1 is equivalent to not moving.*

Region	Cells	p(low distortion)	p(high distortion)
1.	0	.20	.00
2.	1 – 8	.30	.24
3.	9 – 24	.40	.16
4.	25 – 99	.05	.30
5.	100 – 399	.05	.30

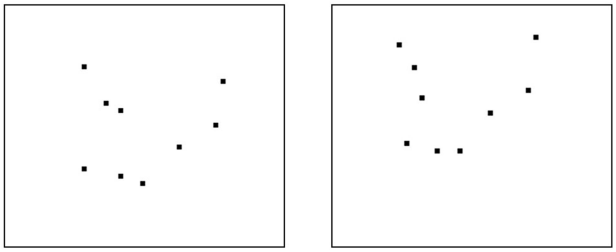

Prototype

Low-Distortion Exemplar

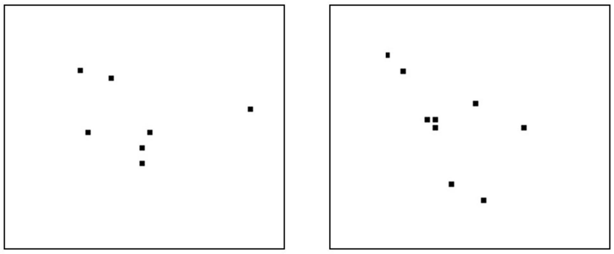

High-Distortion Exemplar

Random Pattern

Figure 19. *Examples of the stimuli used in Experiment 1.*

In the classification task, I generated one prototype for each participant. Then, I generated 20 high-distortion study patterns. I also generated 20 low-distortion test patterns, 20 high-distortion test patterns, and 40 random test patterns.

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

Procedure. Participants sat at individual computers and were instructed that they were going to study patterns of dots for a memory test. Participants viewed the dot patterns one at a time in the center of the screen. Participants in the control group viewed each pattern for 5000 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 exemplar eight times each in pseudo-random order (i.e., in eight randomized 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, in the sense that all dogs belong to the category "dog" despite individual differences, and that they were to sort new patterns based on category membership. These instructions are modeled after Knowlton and Squire's (1993) instructions. 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 center of the screen. Participants responded by clicking one of two alternatives labeled "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 ten patterns – five studied and five unstudied – in random order. Participants responded in the same way as the classification group, but the response alternatives were labeled "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 ten patterns.

Results

Figure 20 presents participants' endorsement rates in both tasks. The left panel presents the proportion of patterns endorsed as "Same category" in the classification task. The right panel presents the proportion of patterns endorsed as "Old" in the recognition task.

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.

To evaluate the data, I analyzed the proportion of patterns endorsed as "Same category" in a 2 x 4 mixed-factor analysis of variance with study time (5000 ms, 100 ms) as a between-subjects factor and probe type (prototype, low-distortion exemplar, high-distortion exemplar, random pattern) as a within-subjects factor. The main effect of probe type was significant, $F(3, 114) = 38.34$, $p < .001$, supporting the claim that participants reliably discriminated among types of patterns. There was no main effect of study time, $F(1, 38) = 3.99$, $p = .053$, although the effect approached significance. There was no interaction between study time and probe type, $F(3, 114) = 0.383$, $p = .765$. Thus there was little evidence for a difference in performance as a function of study time.

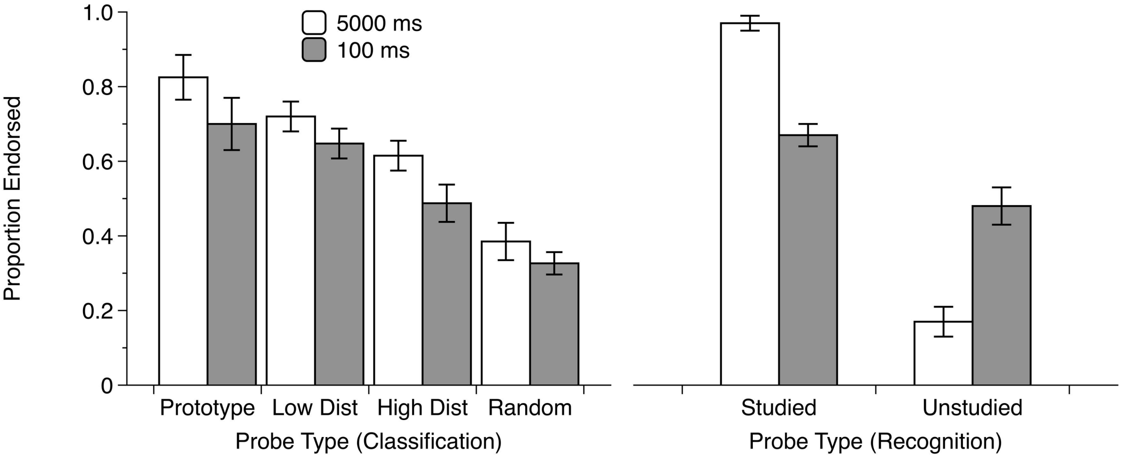

Figure 20. *Endorsement probabilities as a function of encoding condition for both tasks in Experiment 1. The left panel presents the proportion of patterns endorsed as "Same category" in classification. The right panel presents the proportion of patterns endorsed as "Old" in recognition. The error bars represent standard error.*

To analyze the patterns of endorsement in greater detail, I conducted three planned orthogonal contrasts to decompose the main effect of probe type. The first contrast tested the prototype against the other three probe types and was significant, $t(38) = 7.00$, $p < .001$, and shows that the prototype was the most likely pattern type to be endorsed. The second contrast tested the exemplars against the random patterns and was also significant, $t(38) = 7.44$, $p < .001$, and shows that participants discriminated between category and non-category patterns. The third contrast tested the low-distortion exemplars against the high-distortion exemplars and was also significant, $t(38) = 3.26$, $p = .002$, and shows that participants were more likely to endorse patterns similar to the prototype. The pattern of endorsement follows the usual typicality gradient and, more importantly, is quite consistent with Knowlton and Squire's (1993) data.

The right panel of Figure 20 presents the endorsement rates in the recognition task. Participants were more likely to endorse studied patterns than unstudied patterns. However, this difference was much larger following 5000 ms of study time than 100 ms of study time. The conclusion, then, is that reducing study time impaired recognition performance.

To evaluate the data, I analyzed the proportion of patterns endorsed as “Old” in a 2 x 2 mixed-factor analysis of variance with study time (5000 ms, 100 ms) as a between-subjects factor and probe type (target, foil) as a within-subjects factor. There was no main effect of study time, $F(1, 38) = 0.02$, $p = .895$. However, there was a main effect of probe type, $F(1, 38) = 169.44$, $p < .001$, and an interaction between encoding condition and probe type, $F(1, 38) = 64.33$, $p < .001$. The 5000 ms group showed both a higher hit rate and a lower false alarm rate than the 100 ms group, demonstrating much better recognition performance. Although Knowlton and Squire (1993) did not present endorsement rates to targets and foils separately, my results are consistent with their difference in percent correct, replicating their results in addition to providing additional detail. The mirror effect pattern is also consistent with the results of my model-based simulations from Chapter 2.

Figure 21 presents the proportion of correct trials on both tasks as a function of study time. In classification, endorsing the prototype and the exemplars was considered correct and endorsing the random patterns was considered incorrect. In recognition, endorsing the studied patterns was considered correct and endorsing the unstudied patterns was considered incorrect. The left bars show that there is very little difference in the two study-time groups’ classification performance. In contrast, the right bars show that the 5000 ms group was 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, I observed the pattern as a function of study time.

To evaluate the data, I analyzed the proportion of correct trials in a 2 x 2 between-subjects analysis of variance with study time (5000 ms, 500 ms) and task (recognition, classification) as factors. Both main effects were significant. Accuracy was higher following 5000 ms of study time, $F(1, 76) = 18.15$, $p < .001$, and in the recognition task, $F(1, 76) = 40.90$, $p < .001$. 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 (.31) than in classification (.02). This pattern is a clear replication of the results reported by Knowlton and Squire (1993).

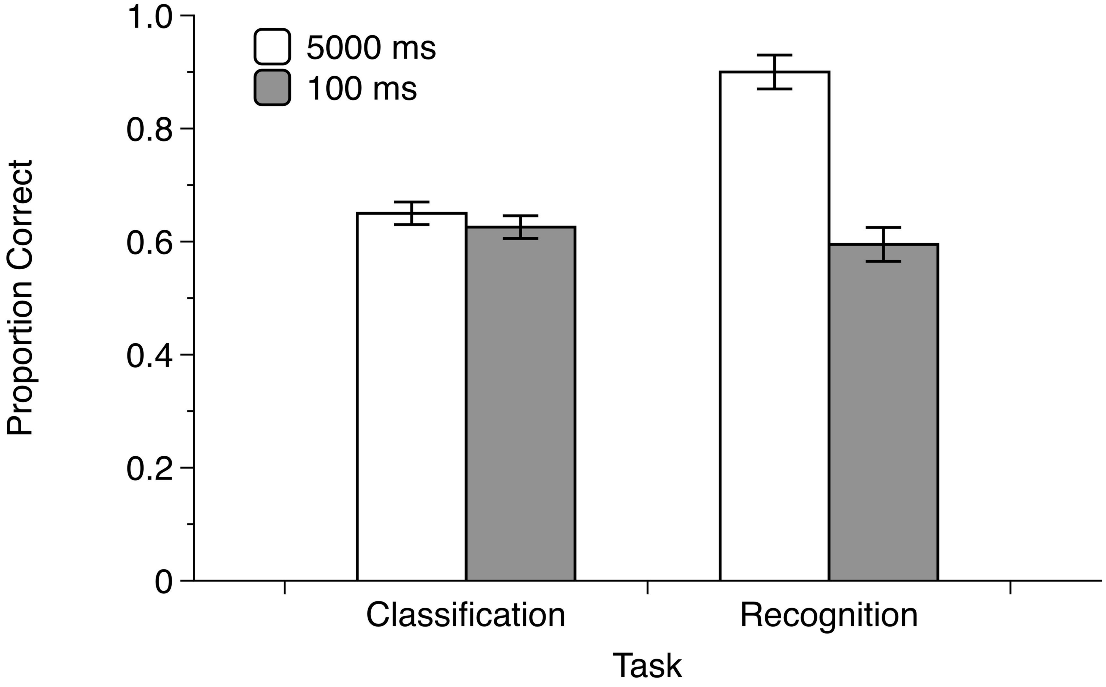

Figure 21. *The proportion of correct trials as a function of task and encoding condition in Experiment 1. The error bars represent standard error.*

Discussion of Experiment 1

Experiment 1 provides converging evidence with Simulation 1 that the pattern of impairments in recognition and classification in amnesia is consistent with what would be expected from a general rather than selective encoding deficit. Taken together with the computational simulations, the data converge to strengthen the argument that pattern of impairment observed in amnesia can be explained as a global encoding deficit. The computational and empirical analyses approach the results from a common set of fundamental principles.

The empirical analysis also converges with work by Nosofsky and Zaki (1998). They replicated the same procedure and the same pattern by introducing a delay between study and test, presumably resulting in some amount of forgetting. By their analysis, the pattern occurs due to rapid forgetting. By my analysis, the pattern occurs due to poor encoding. The two analyses differ in the details of the primary mechanism driving the results. However, they converge on the principles that the pattern is the result of incomplete traces in memory consistent with either rapid forgetting or poor encoding.

In Experiment 2, I extend my empirical analysis to the same pattern using binary-featured patterns from Reed et al.'s (1999) study with cartoon animals.

Experiment 2

Experiment 2 replicates 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 participants in a recognition task with the same stimuli. I replicate the two sets of data, which form a selective impairment, using healthy participants by manipulating study time. However, to avoid complications inherent to these types of stimuli, and because their stimulus set was not available, I developed a set of stimuli that maintain the exact same category structure but are more abstract and precise.

Methods

Participants. I 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 (5000 ms, 100 ms) and task (recognition, classification). All participants were given credit toward a course requirement for participation.

Materials and Apparatus. The apparatus were identical to Experiment 1. All stimuli were patterns of shapes and colours consisting of nine features that could each take on one of two values. Table 2 presents the nine features and their possible values.

Table 2. *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

For the classification task, I constructed a prototype for each participant by combining the feature values for either column 2 or 3 of Table 2 (each for half of the participants in each group). I constructed an anti-prototype using the other set of values. I then constructed 80 more patterns. There were 20 low-distortion study exemplars and 20 low-distortion test exemplars. There were also 20 neutral test exemplars, which differed on either four or five randomly-selected features, and 20 high-distortion test exemplars, which differed on either seven of eight features. Figure 22 presents an example from each category for a single participant.

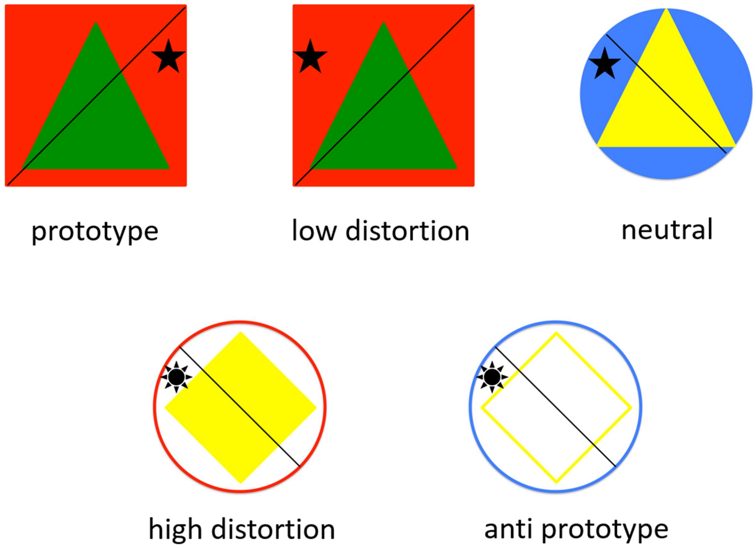

Figure 22. *Examples of the binary-featured patterns used in Experiment 2.*

I also constructed sets of random features for the recognition task. For each participant, I constructed ten 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 center of the screen. Participants in the control group viewed each pattern for 5000 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 randomized blocks). The recognition group viewed five random study patterns eight times each in pseudo-random order (i.e., eight randomized blocks).

Following study, participants in the classification group were given the same instructions as in Experiment 1. Participants evaluated 96 test exemplars –the prototype (12 times), 20 low-distortion test exemplars, 20 moderate-distortion test exemplars, 20 high-distortion test exemplars, 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 stimuli were presented in the center of the screen. Participants responded by clicking one of two alternatives labeled “Same category” and “Not same category”. 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 96 patterns.

The recognition group followed a similar procedure. The participants evaluated ten patterns – five studied and five unstudied – in random order. Participants responded in the same way as the classification group, but the response alternatives were labeled “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 ten patterns.

Results

Figure 23 presents participants’ endorsement rates in the classification and recognition tasks. The left panel presents the proportion of patterns endorsed as “Same category” in the classification task. The right panel presents the proportion of patterns endorsed as “Old” in the recognition task.

The left panel of Figure 24 shows that participants, as in Experiment 1, endorsed patterns 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.

To evaluate the data, I analyzed the proportion of patterns endorsed as “Same category” in a 2 x 5 mixed-factor analysis of variance with study time (5000 ms, 100 ms) as a between-subjects factor and probe type (prototype, low-distortion exemplar, neutral exemplar, high-distortion exemplar, anti-prototype) as a within-subjects factor. As in Experiment 1, there was a main effect of probe type, $F(3, 152) = 253.65$, $p < .001$, but no main effect of study time, $F(1, 38) = 0.03$, $p = .870$, and no interaction between probe type and study time, $F(3, 152) = 1.27$, $p = .283$. The analysis supports the claim that participants reliably discriminated among category members and nonmembers and that categorization success was not influenced by study time.

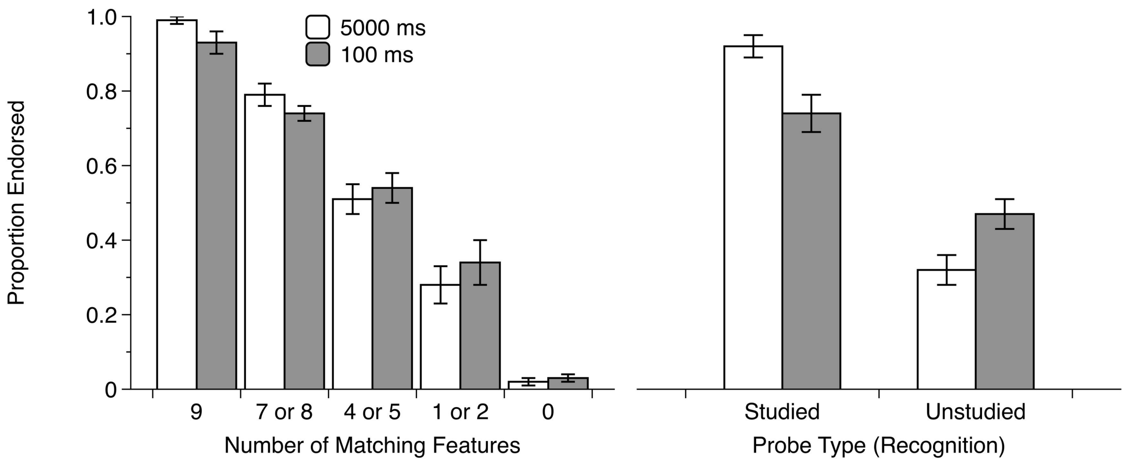

Figure 23. *Endorsement probabilities as a function of encoding condition for both tasks in Experiment 2. The left panel presents the proportion of patterns endorsed as "Same category" in classification. The right panel presents the proportion of patterns endorsed as "Old" in recognition. The error bars represent standard error.*

As in Experiment 1, I decomposed the main effect of probe type using three planned orthogonal contrasts to better evaluate the endorsement pattern. The first contrast compared the prototype and low-distortion exemplars against the anti-prototype and high-distortion exemplars, $t(38) = 22.26, p < .001$, and shows that participants discriminated between category and non-category patterns (by the same division as Reed et al., 1999). The second compared the prototype against the low-distortion exemplars, $t(38) = 4.32, p < .001$, and shows that participants had an additional bias toward the prototype. The third compared the anti-prototype against the high-distortion exemplars $t(38) = 6.50, p < .001$, and shows the corresponding bias against the anti-prototype.

The right panel of Figure 23 shows that, after studying each patterns for 5000 ms, participants were more likely to endorse the studied patterns and less likely to endorse the unstudied patterns compared to the 100 ms group. As in Experiment 1, although both groups discriminated studied from unstudied patterns, the difference was much larger following 5000 ms of study time per item.

To evaluate the data, I analyzed the proportion of patterns endorsed as "Old" in a 2 x 2 mixed-factor analysis of variance with study time (5000 ms, 100 ms) as a between-subjects factor and probe type (target, foil) as a within-subjects factor. There was no main effect of study time, $F(1, 38) = 0.10$, $p = .754$, but there was a main effect of probe type, $F(1, 38) = 142.25$, $p < .001$, and an interaction between study time and probe type, $F(1, 38) = 20.47$, $p < .001$. The pattern is nearly identical to Experiment 1.

Figure 24 presents the proportion of correct trials on both tasks as a function of study time. In the classification task, endorsing the prototype and the low-distortion test exemplars was considered correct and endorsing the anti-prototype and the high-distortion test exemplars was considered incorrect. The moderate-distortion exemplars were excluded from analysis, consistent with Reed et al.'s original method. In the recognition task, endorsing the studied patterns was considered correct and endorsing the unstudied patterns was considered incorrect.

As in Experiment 1, there is little difference in classification performance as a function of study time. In contrast, the 5000 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, I observed the pattern as a function of study time rather than amnesia.

To evaluate the data, I analyzed the proportion of correct trials in a 2 x 2 between-subjects analysis of variance with study time (5000 ms, 500 ms) and task (classification, recognition) as factors. Both main effects were significant. Participants in the 5000 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 greater in recognition (.17) than in classification (.05). This pattern, although not of the same magnitude as Experiment 1, provides another unambiguous empirical pattern in which study time had a greater impact on recognition than on classification.

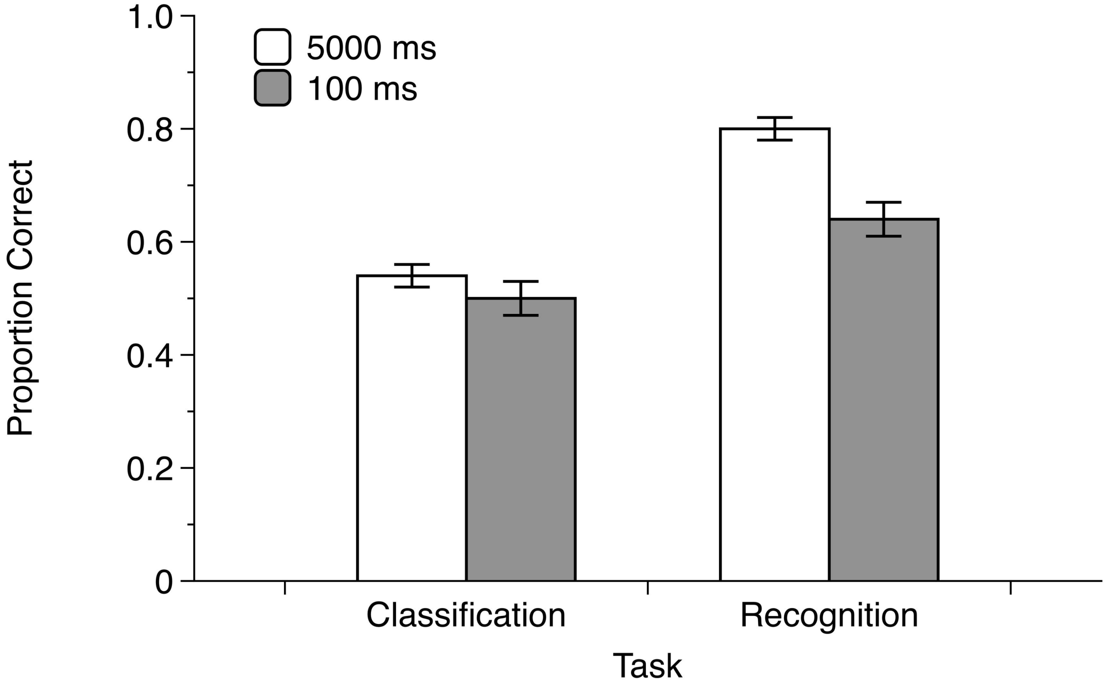

Figure 24. *The proportion of correct trials as a function of task and encoding condition in Experiment 2. The error bars represent one standard error above and below each mean.*

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 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.

In the next experiment, I empirically simulate differences between tachistoscopic identification and recognition (Keane et al. 1995).

Experiment 3

Experiment 3 replicates Keane et al.'s (1995) results using recognition and priming of tachistoscopic identification. In their experiment, amnesic patients and controls studied consonant strings. Following study, they identified strings presented tachistoscopically. Subsequently they discriminated between studied and unstudied strings. I conducted an experiment to replicate their procedure using healthy participants and simulating poor encoding in amnesia by manipulating study time.

Methods

Participants. I recruited 80 participants from the University of Manitoba psychology participant pool. The participants were randomly assigned to four groups formed by crossing study time (5000 ms, 100 ms) and task (identification, recognition). I recruited an additional 40 participants to perform an easier version of the identification task as a result of poor performance in the first groups. All participants were given credit toward a course requirement for participation.

Materials and Apparatus. The experiment was conducted on the same apparatus as Experiments 1 and 2.

I randomly selected consonant strings from a list of all possible three-consonant combinations under the restrictions that there were no letter repetitions, that there were no X's (X's were used to construct a mask), and that the strings did not form well-known acronyms (e.g., BTW, NHL, NDP). I randomly selected 48 strings for each participant and randomly assigned 24 to serve as study strings and 24 to serve as unstudied test strings.

Procedure. Participants sat at individual computers and viewed the 24 strings from the study list one at a time in random order. The control group viewed each string for 5000 ms and the simulated amnesia group viewed each string for 100 ms. The screen cleared for 250 ms between each string. This process repeated until all 24 strings had been presented.

Testing for the recognition group was nearly identical to the procedure in Experiments 1 and 2. Participants evaluated the 24 studied strings and 24 unstudied strings one at a time in random order. The strings were presented centrally above two response alternatives. Participants responded by clicking one of two response alternatives labeled “Old” and “New”. This process repeated until participants had provided a response to all 48 strings.

In the identification test, participants viewed the 48 strings for 50 ms: each string was replaced after 50 ms by a mask of five X's for 250 ms, followed by a clear screen for 250 ms. After each presentation, participants attempted to identify the string by typing it into a response box using the keyboard and pressing “enter” on the keyboard. After the response, the screen cleared for 250 ms. This process repeated until participants had provided a response to all 48 strings.

Due to poor overall performance in both identification groups, I also conducted an easier version of the task where the strings were presented for approximately 67 ms. The seemingly arbitrary increase in presentation time is due to restrictions caused by the refresh rate on the computer monitors. Given that the screen refreshed at 60 Hz, or once every 16.67 ms, display times must be a multiple of 16.67 ms.

Results

I kept the 67 ms identification group for initial analysis, and measured their performance by the proportion of correct identifications. I excluded two participants, one in each study time condition, for failing to type a response to any trial. I measured the recognition group's performance by the proportion of endorsements.

Figure 25 presents the results of Experiment 3. The left panel presents identification accuracy as a function of study time and probe type. Priming is measured by the difference in accuracy between studied and unstudied strings. As shown, participants were slightly more accurate at identifying studied than unstudied strings. However, the difference was larger in the 5000 ms group than the 100 ms group.

To evaluate the data, I analyzed identification accuracy in a 2 x 2 mixed-factors analysis of variance with study time (5000 ms, 100 ms) as a between-subjects factor and probe type (studied, unstudied) as a within-subjects factor. There was no main effect of study time, *F*(1, 36), = 0.88, *p* = .354. However, there was a main effect of probe type, *F*(1, 36), = 12.38, *p* = .001, and an interaction between probe type and study time, *F*(1, 36), = 4.56, *p* = .040. Participants were more likely to correctly identify studied strings compared to unstudied strings, but the difference was larger in the 5000 ms group. This pattern contradicts the pattern observed by Keane et al., (1995) who reported no statistical difference between controls and amnesic patients.

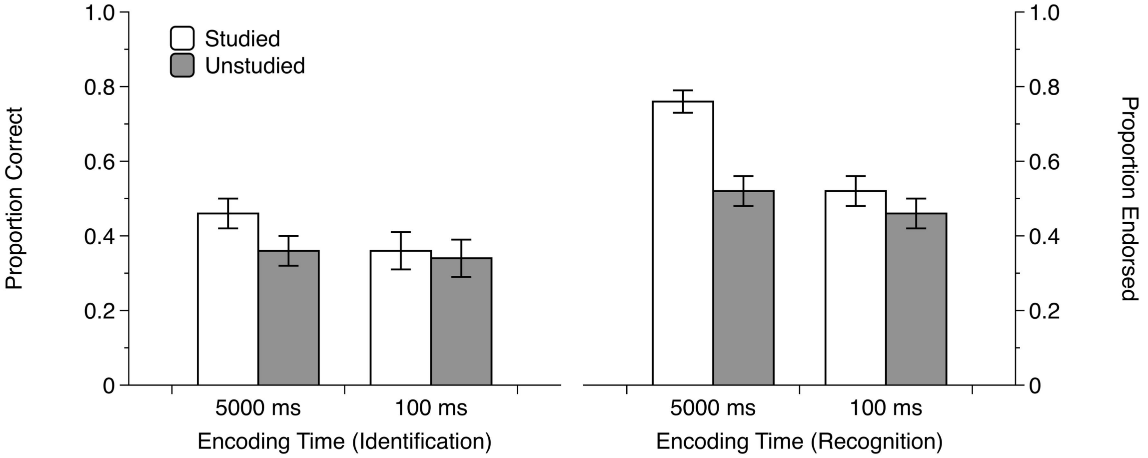

Figure 25. *Identification and endorsement probabilities as a function of encoding condition for both tasks in Experiment 3. The left panel presents the proportion of strings correctly identified in the identification task. The right panel presents the proportion of strings endorsed as "Old" in recognition. The error bars represent standard error.*

The right panel of Figure 25 presents the proportion of strings endorsed as "Old" as a function of study time and probe type. As in the preceding experiments, both groups were more likely to endorse the studied strings, but the difference was much greater in the 5000 ms group than the 100 ms group. Again, reducing study time produced an impairment to recognition performance. The difference between discrimination in the recognition task is much greater than the difference between priming in the identification task, suggesting the same pattern observed by Keane et al. (1995).

To evaluate performance in the recognition task alone, I analyzed the proportion of "Old" endorsements in a 2 x 2 mixed-factors analysis of variance with study time (5000 ms, 100 ms) as a between-subjects factor and probe type (studied, unstudied) as a within-subjects factor. There was a main effect of encoding, $F(1, 38) = 10.27$, $p = .003$, a main effect of probe type, $F(1, 38) = 32.35$, $p < .001$, and an interaction between the two factors, $F(1, 38) = 11.02$, $p = .002$. The 5000 ms group had a higher endorsement rate overall, although this increase was primarily to the studied strings. Both groups were more likely to endorse the studied strings, but the difference was much larger in the 5000 ms group.

To evaluate the pattern of impairment, I analyzed the overall performance on both tasks, where performance on both tasks was measured by the difference between responses to studied strings and unstudied strings.[2] I analyzed performance in a 2 x 2 between-subjects analysis of variance with encoding time (5000 ms, 100 ms) and task (identification, recognition) as factors. The main effects of encoding, $F(1, 74) = 34.87$, $p < .001$, and task, $F(1, 74) = 51.22$, $p < .001$, were both significant. More importantly, the interaction was significant, $F(1, 74) = 22.14$, $p < .001$, such that the difference in performance as a function of encoding was greater in recognition (.37) than in identification (.07). This pattern reflects a difference as a function of study time, replicating Keane et al.'s (1995) difference as a function of amnesia.

Additionally, the 50 ms identification group sheds some light on the double dissociation observed by Gabrieli et al. (1995) and Keane et al. (1995). They tested two patients with occipital lobe damage. The patients showed intact recognition performance but impaired identification performance, the opposite pattern of amnesic patients.

---

[2] Keane et al. (1995) analyzed d' instead of a simple difference in their recognition task. I re-analyzed my data using a d' measure with identical results. I opted to present the simple difference measure to maintain consistency with the data from the identification task (i.e., same scale). However, the analysis showed no differences in the conclusions I draw.

Although the procedure in my Experiment 3 differed from their procedure, the poor performance of my 50 ms identification group might be sufficient to produce the second half of the double dissociation.

Figure 26 presents proportion of correct identifications at 50 ms and 67 ms presentation times and for studied and unstudied strings. The two leftmost sets of bars represent the same data as Figure 29, presenting the proportion of correct identifications at 67 ms presentation times. The two rightmost sets of bars represent the proportion of correct identifications at 50 ms presentation times. Overall identification accuracy was lower at 50 ms presentation times, decreasing from 38% to 23%, collapsed across study-time groups, and this difference was statistically reliable, $t(74) = 5.05$, $p < .001$. Additionally, decreasing presentation time also decreased the differences between studied and unstudied strings. After studying each string for 5000 ms, there was a dramatic decrease in the priming difference. After studying each string for 100 ms, unstudied strings were actually slightly more likely to be identified than studied strings, although the difference is small.

To evaluate the data, I analyzed the magnitude of identification priming (i.e., the difference between studied and unstudied strings) in a 2 x 2 between-subjects analysis of variance with study time (5000 ms, 100 ms) and identification time (67 ms, 50 ms) as factors. The main effect of encoding was significant, $F(1, 74) = 4.60$, $p = .035$, such that priming was greater after studying each string for 5000 ms. Importantly, the main effect of presentation time at test was also significant, $F(1, 74) = 4.00$, $p = .049$, such that the 67 ms group showed greater priming than the 50 ms group. There was no significant interaction, $F(1, 74) = 0.64$, $p = .425$. The 50 ms group behaved similarly to Gabrieli et al.'s (1995) and Keane et al.'s (1995) occipital lobe damage patients, showing both reduced accuracy and priming. That is, participants placed under severe presentation times that make it difficult to perceive the test stimuli perform similarly to Gabrieli et al.'s and Keane et al.'s patients with visual impairment.

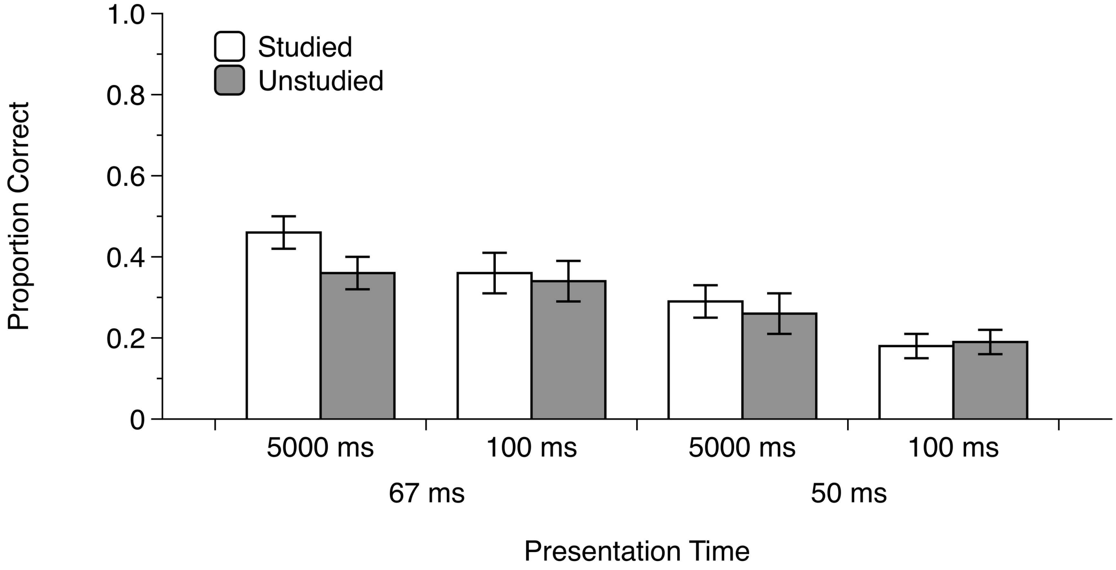

Figure 26. *Identification probabilities as a function of encoding condition and presentation time at test in Experiment 3. The error bars represent one standard error above and below each mean.*

Discussion of Experiment 3

Experiment 3 extends my empirical simulations from selective impairments involving classification to selective impairments involving priming. The results also extend the computational analysis in Simulation 3, converging on the argument that patterns of impairment observed in amnesia can be produced by a global encoding deficit.

Experiment 3 is not a perfect replication of Keane et al.'s (1995) data. Their amnesic patients showed statistically equivalent priming compared to the control participants. In contrast, the participants in my 5000 ms group showed statistically greater priming than the participants in my 100 ms group. However, any conclusion drawn from a null result must be scrutinized carefully. This concern is especially true in Keane et al's experiment, which involved fewer participants and much greater error variance than my Experiment 3. To avoid making the deductive error of accepting a null hypothesis, the valid comparison is in the interaction between study time and task. The interaction in Experiment 3 captures the nature of the selective impairment despite minor differences with its original report.

The data also provide some insight into a corresponding double dissociation (Gabrieli et al., 1995; Keane et al., 1995). They reported two cases of occipital lobe damage that resulted in impaired identification priming but intact recognition performance. In addition to impaired priming, the patients' identification itself was impaired. My 50 ms identification participants showed the same reduced priming and accuracy. However, the comparison is far from perfect. First, the 50 ms group did not perform a recognition task, resulting in an incomplete analysis. Second, there were many differences between the procedures in my Experiment 3 and the original reports of the double dissociation. Nonetheless, the data do suggest that the impairment in priming is caused by a more general and compound impairment in accuracy. Consequently, one can reasonably suggest that the impairment is one of vision rather than memory.

Experiment 3 provides an analysis of one type of priming. However, priming is also observed in other experimental tasks. In Experiment 4, I extend my empirical analysis to priming of word-stem completion.

Experiment 4

Experiment 4 replicates Graf, Squire, and Mandler's (1984) differences in impairment between recognition, cued recall, and word-stem completion. In their experiments, amnesic patients and controls studied words. Following study, they performed three tasks. In one task, they discriminated studied words from unstudied words. In another task, they were presented with word stems and provided the first complete word that came to mind. In another task, they completed word stems with words from the study list. I replicated their procedure using healthy participants, manipulating study time as a proxy for general memory deficits consistent with the hypothesis of a global encoding impairment in amnesia.

Methods

Participants. I recruited 120 participants from the University of Manitoba Introduction to Psychology participant pool. I randomly assigned participants to six groups formed by crossing the factors study time (5000 ms, 100 ms) and task (completion, cued recall, recognition). All participants were given credit toward a course requirement for participation.

Materials and Apparatus. The experiment was conducted using the same apparatus as the previous experiments.

I randomly selected words for each participant from the SUBTLEX-US database (Brysbaert & New, 2009) under the following constraints. Each word was between four and seven letters long. Word stems, consisting of the first three letters of the word, shared a common stem with at least 10 other English words. Each word had a frequency between

10 and 100 per million. Finally, I removed any words that were proper nouns, curse words, words with sexual meaning or connotation, or derogatory slurs.

For each participant in the word-stem completion and cued recall groups, I randomly selected 80 words, none of which contained the same stem, and randomly assigned them to four lists of 20 words. Half of the words in each list served as study words and half served as unstudied test words. For each participant in the recognition groups, I randomly selected 20 word stems and three words from each stem. One word from each trio served as a study word and the remaining two words served as unstudied test words.

Procedure. Participants in the word-stem completion groups began by viewing the first list of 10 words one at a time, twice each, and in blocked random order. Following study, participants were presented with the first three letters of the studied words and the unstudied words from the same list. All word stems were presented in random order. The participants were asked to complete the stem with the first word that came to mind by typing the rest of the word and pressing “enter” on the keyboard. The screen then cleared for 250 ms. This process repeated until participants had provided a response to all 20 word stems. The full process repeated until all four lists had been studied and tested.

Participants in the cued recall groups completed the same procedure. However, they were explicitly asked to complete word stems with words from the study list rather than the first word that came to mind.

Participants in the recognition group viewed the 20 study words twice in blocked random order. Following study, participants were asked to discriminate studied from unstudied words in a three-alternative-forced-choice test. On each test trial, participants were presented with three words with the same stem, one of which was from the study list. The words were arranged in a triangle and the location of the studied word was randomized on each trial. A message at the bottom of the screen prompted the participant to “Click on the word that was from the list you just studied”. Once the participant had clicked on a word, the screen cleared for 250 ms before the next three words appeared. This process repeated until the participant had made a response on all 20 trials.

Results

Figure 27 presents performance on all three tasks as a function of study time. Performance on the word-stem completion and cued recall tasks was measured by the difference between the valid completion of studied stems and the baseline completion of unstudied stems. Performance in the recognition task was measured by the proportion of trials on which the participants selected the studied word.

The left set of bars presents performance on the word-stem completion tasks as a function of study time. Both groups were more likely to complete stems with studied words compared to baseline. However, the difference was larger in the 5000 ms group than the 100 ms group. This pattern contrasts with Graf, Squire, and Mandler’s (1984) data showing no statistical difference in completion performance.

To evaluate the data, I analyzed the proportion of stems completed as the target word in a 2 x 2 mixed-factors analysis of variance with study time (5000 ms, 100 ms) as a between-subjects factor and probe type (studied, unstudied) as a within-subjects factor. There were main effects of study time, $F(1, 38) = 25.98$, $p < .001$, and probe type, $F(1, 38) = 149.40$, $p < .001$. There was also an interaction between the two factors, $F(1, 38) =$

26.37, $p < .001$. Both groups completed stems with words from the study lists above baseline. However, participants in the 5000 ms group were more likely to do so than participants in the 100 ms group, resulting in both the main effect of study time and the interaction.

The middle set of bars presents performance on the cued recall task as a function of study time. The results are similar to the word-stem completion task. Both groups were more likely to complete stems with the studied words, and the difference was larger following 5000 ms study than 100 ms study. However, the difference was even larger than in the completion task.

To evaluate the data, I analyzed the proportion of stems completed / words recalled in a 2 x 2 mixed-factors analysis of variance with study time (5000 ms, 100 ms) as a between-subjects factor and probe type (studied, unstudied) as a within-subjects factor. As in word-stem completion, all three terms were significant: there was a main effect of study time, $F(1, 38) = 94.99$, $p < .001$, a main effect of probe type, $F(1, 38) = 350.10$, $p < .001$, and an interaction, $F(1, 38) = 115.90$, $p < .001$. The pattern of results is consistent with Graf, Squire, and Mandler's (1984) data.

The right set of bars presents the proportion of trials on which participants selected the studied word in the recognition task as a function of study time. Both groups performance was well above chance (33%). However, as in the other two tasks, the 5000 ms group were more accurate than the 100 ms group. To evaluate the data, I conducted an independent Welch's *t* test comparing the proportion of correct trials for the 5000 ms and 100 ms groups. There was a significant difference such that the 5000 ms group (.95) was more accurate than the 100 ms group (.64), $t(38) = 7.01$, $p < .001$.

To evaluate the complex pattern, I analyzed performance across tasks in a 2 x 3 between-subjects analysis of variance with study time (5000 ms, 100 ms) and task (word-stem completion, cued recall, recognition) as factors. The main effects of study time, *F*(1, 114) = 162.78, *p* < .001, and task, *F*(2, 114) = 131.27, *p* < .001, were both significant. More importantly, the interaction was significant, *F*(2, 114), *p* = .029.

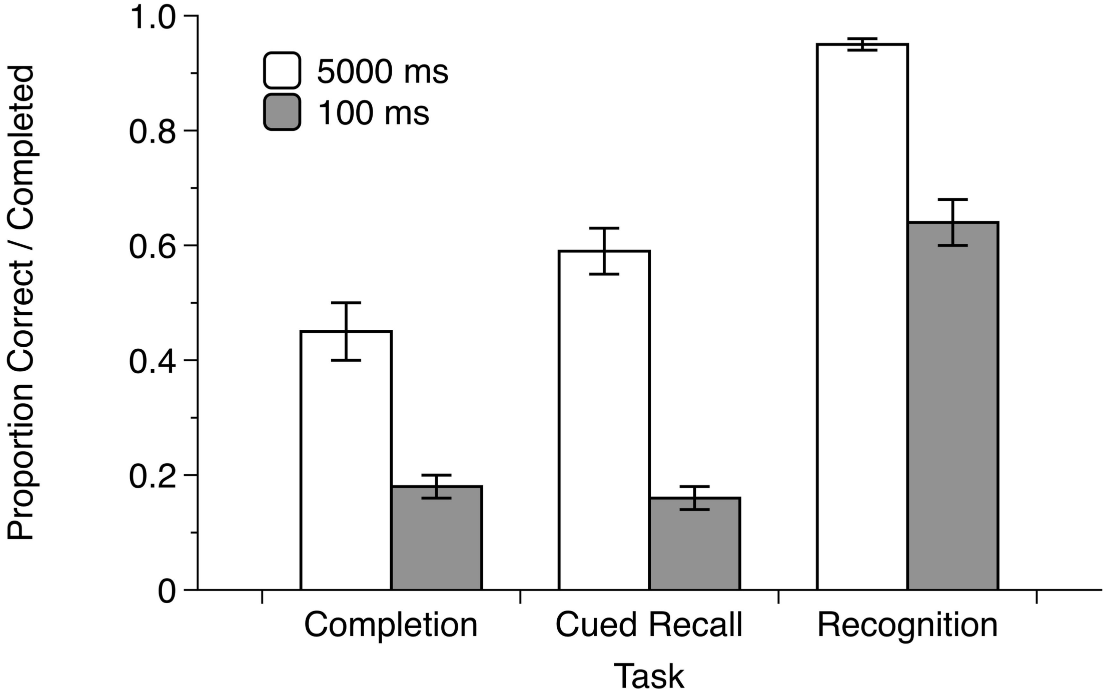

Figure 27. *Performance as a function of encoding time and task in Experiment 4. Performance in word-stem completion and cued recall is indexed as the difference between completion rates to studied versus unstudied items. Performance in recognition is indexed as the proportion of correct trials.*

I decomposed the interaction term in two planned orthogonal contrasts. First, I compared performance on word-stem completion (an implicit memory task) against performance on cued recall and recognition (explicit memory tasks) in a standard interaction contrast. The contrast was significant, *t*(114) = 4.55, *p* < .001, such that the difference between encoding conditions was larger in cued recall and recognition than in word-stem completion. This pattern suggests that decreasing encoding time influences tasks thought to rely on explicit memory more strongly than tasks thought to rely on implicit memory.

The second contrast compared performance on cued recall against performance on recognition in another interaction contrast. This contrast was also significant, $t(114) = 3.90$, $p < .001$, such that the difference between encoding conditions was larger in cued recall than in recognition.

However, another important comparison is the one between word-stem completion and recognition. An inspection of the means demonstrates that the differences between encoding conditions are more similar between those two tasks. Thus although not orthogonal with the first planned contrast, I conducted an interaction contrast comparing performance on word-stem completion and recognition. The contrast was not significant, $t(114) = 1.27$, $p = .206$. Despite limitations to conclusions drawn from null results, the lack of interaction between completion and recognition when isolated from the difference in cued recall, stands as a point of failure to replicate Graf, Squire, and Mandler's (1984) results.

Discussion of Experiment 4

Experiment 4 replicates some aspects of the differences among word-stem completion, cued recall, and recognition. However, it also fails to convincingly replicate others. Interestingly, the experiment fails to replicate the pattern involving completion and recognition, but succeeds in replicating the intuitively more intriguing difference between completion and cued recall. The completion and recognition tasks differ in procedure and, as a result, any differences might be expected for any number of reasons. However, the completion and cued recall tasks differ only in a single instruction. Nonetheless, the result can be produced by a manipulation of study time.

It should be noted the lack of difference between word-stem completion and recognition is particularly unconvincing given that the 5000 ms group performed near ceiling levels (.95). It is plausible that, if the recognition task were made more difficult, performance would decrease in both groups but more rapidly in the 100 ms group. If so, one could produce a difference between word-stem completion and recognition by manipulating study time. However, that must remain an empirical question until additional experimental work is completed.

The mixed results of the empirical simulation also contrast with the convincing failure of the corresponding computational simulation (see Chapter 2, Simulation 4). The complex differences among word-stem completion, cued recall, and recognition certainly deserves more careful examination in the future. The partial success of the empirical simulations suggests that the results can, at least in part, be produced by manipulating the quality of encoding. Combined with the failure of the model, the results suggest that HEM’s failure to produce the pattern reflects a deficit in specific mechanisms rather than principle assumptions. In particular, the result points to the need for research using computational theories of memory to begin importing the sophisticated and modern linguistic vector spaces that have been developed in psychology over the past 15 years (e.g., Chubala, Johns, Jamieson, & Mewhort, 2016; Jones & Mewhort, 2007; Landauer & Dumais, 1997).

Summary of Empirical Simulations

I have reported four experiments in this chapter. The primary aim of the experiments was to extend my computational simulations using empirical simulations based on the same underlying principle. The main premise in both the computational and empirical analyses is that patterns of behaviour observed in amnesia can be explained as a global encoding deficit rather than selective impairment to an explicit memory system. The empirical simulations were largely successful, with the exception of mixed results in Experiment 4. Experiments 1 and 2, examining differences between classification and recognition, converge with simulations of the same tasks (Simulations 1 and 2). Additionally, Experiment 3, examining a difference between identification priming and recognition, converges with another simulation of the same task (Simulation 3). The mixed success of Experiment 4, examining differences among word-stem completion, cued recall, and recognition stands in interesting contrast with the failure of the corresponding computational simulation (Simulation 4).

I have elected to produce amnesic-like behaviour in healthy participants by manipulating study time. Although previous work has successfully used different methods (Graf, Mandler, & Haden, 1982; Nosofsky & Zaki, 1998; Woods & Percy, 1974; Zaki & Nosofsky, 2001), the manipulation of study time maps directly onto the parameter manipulation in my computational simulations. In turn, that convergence between approaches permits a coherent and rational analysis under a common theoretical framework.

Together, the set of empirical simulations provide converging evidence in favour of the single-system perspective. However, the true value of a theoretical framework is to move beyond explaining existing data to making novel testable predictions. To that end, Chapter 4 presents a concrete example of how I can extend both my computational and empirical analyses to develop a novel prediction for a high resolution experiment and analysis to be carried out with amnesic patients.

Chapter 4 – Deriving and Testing Novel Predictions

My computational and empirical simulations were largely successful in accommodating a set of hallmark patterns observed in amnesia. The simulations are based on the premise that the patterns reflect a global encoding deficit rather than a selective impairment of memory. To the extent that they were successful in accommodating the results, they support the conclusion that systems distinctions are an unnecessary addition to memory theory.

However, a stronger test of a theory is the ability to move beyond existing data and derive novel predictions. In this chapter, I apply the HEM to derive a novel prediction from a variant of the artificial grammar task. I also apply my empirical simulation technique by inducing amnesic-like behaviour in healthy individuals by manipulating study time.

String Completion in Artificial Grammar Learning

Jamieson and Mewhort (2010) tested participants on a variant on the standard artificial grammar task. Participants studied strings constructed according to an artificial grammar (e.g., SPSVP, TXSVP, SPT). At test, participants were presented with new strings that were missing one letter (e.g., S_S, TXS_PV, TXX_S) and five response alternatives (S, T, X, P, and V). Only one of the response alternatives would make the incomplete string grammatical – all others were illegal according to the rules of the artificial grammar.

Participants selected the grammatical alternative well above chance levels (70% compared to 20% chance). However, all of the participants were healthy undergraduates. It is unknown whether amnesic patients perform artificial-grammar string completion similarly. In Simulation 5, I test the HEM's performance on the string completion task and a corresponding recognition task to derive predictions about the influence of encoding quality.

Simulation 5

Simulation 5 examines the influence of L on string completion in artificial grammar learning. The simulation uses the same materials as Jamieson and Mewhort's (2010) Experiment 1. Additionally, I model a corresponding recognition task in order to capture a potential interaction.

To simulate string completion, I generated 26 vectors, one for each letter of the alphabet. Each vector was a vector of 100 elements with each element taking a random number sampled from a normal distribution with a mean of zero and a standard deviation of $1/\sqrt{n} = 0.1$. I generated letters strings by taking the individual letters in each string in Table 4 and computing their iterative circular convolution. As in Simulation 3, where the stimuli were also letter strings, I used non-commutative circular convolution in which the indices of the vectors are scrambled prior to convolution in order to preserve the ordinal information of the strings.

Additionally, participants do not encode all of the information in a string in an artificial grammar task, and the information encoded differs among participants (Wright & Whittlesea, 1998). To solve this problem, I generated all possible n-grams from each string. For example, the string SPT contains the n-grams: S, P, T, SP, PT, and SPT. To represent a letter string, I randomly selected a subset of the n-grams. Then, I iteratively convolved the letter representations in each n-gram, normalizing the vector before each convolution. I summed the representations of each *n*-gram, and finally normalized the final vector.

For example, in one simulation the string SPT might be encoded using the *n*-grams SP, P, and T. To represent the bigram SP, I would take the circular convolution of the vectors corresponding to S and P, and normalize the resulting composite vector. I would then sum the elements of the composite SP vector, the P vector, and the T vector. In another simulation, the string SPT might be encoded using the *n*-grams S, T, and SPT. To represent the trigram SPT, I would first take the circular convolution of the vectors corresponding to S and P, and normalize the resulting composite vector. I would then take the circular convolution of the SP composite vector and the vector corresponding to T, and normalize the resulting vector. I would then sum the elements of the SPT composite vector, the S vector, and the P vector.

Table 3. *The letter strings used in Simulation 5 and Experiment 5, taken from Jamieson and Mewhort (2010)*

Training Strings	Partial Test Strings	Legal Completion
SPSVP	S_S	P
SPSVPS	SPS_PV	V
SPSX	S_V	P
SPT	SP_T	V
SPVV	SPV_V	V
SPVVT	S_VVVT	P
SPVVVV	SP_X	S
SPXSVP	SPXX_	S
SPXSX	_PXXXS	S
SPXXSX	TSV_T	P
TSVP	_SVPV	T
TSVPS	TSVP_T	V
TSVPSX	TS_	X
TSVPVV	T_SVPS	X
TSVPXS	TXS_PV	V
TXS	TX_X	S
TXSVP	TXX_	S
TXSVPT	T_XSVP	X
TXXSX	TXX_S	X
TXXXXS	TX_XXSX	X

I generated the 20 study strings in column 1 of Table 3 and stored them to a memory matrix. To simulate string completion, I generated five representations for each of the test strings in column 2 of Table 4, replacing the missing letter (marked by the underscore) with each of the five letters in the study materials (S, P, T, V, and X), using the same sampling method as the study strings. Then, I computed an echo from each completed string. Then, I computed the cosine similarity between each echo and string. I conducted 1000 simulations at each level of L = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0.

I simulated recognition performance using nearly the identical method as string completion. However, to simulate the test phase, I probed memory with both the studied strings and unstudied strings. The studied strings were the same strings stored to memory. I generated the unstudied strings by replacing the missing letter in each string in column 2 of Table 3 with the legal completion presented in column 3 of Table 3. I computed an echo and a cosine similarity for each test item. I conducted 1000 simulations at each level of L = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0.

Results

Figure 28 presents the model's results. Given that there are no existing data to fit, I present the raw cosine similarities; I fit the cosines to new data in a subsequent section. The model results have a few interesting features. First, the model predicts a small but consistent increase in string completion performance as L increases. Second and surprisingly, the model predicts a much smaller increase in recognition performance. This pattern is the exact opposite of the patterns typically observed in amnesia; the implicit task is influenced more strongly than the explicit task. Third, a partial explanation for this pattern might be that recognition performance was extremely poor. In fact, the model's discrimination was near zero at all values of L.

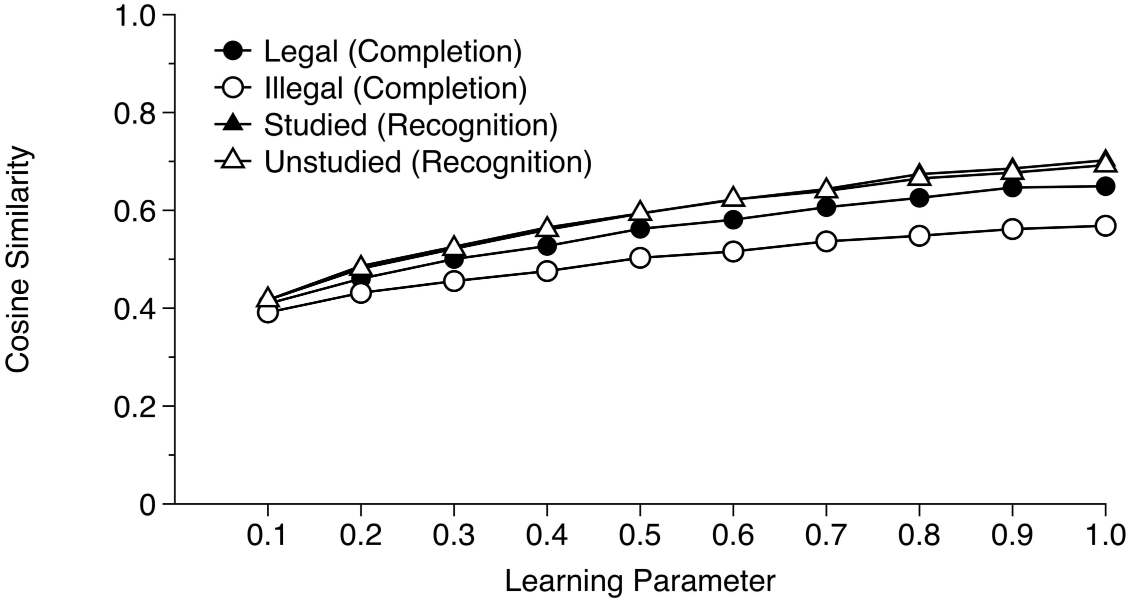

Figure 28. *Simulation 5 cosine similarities in string completion and recognition, in response to legal and illegal strings in completion, and studied and unstudied strings in recognition. Studied and legally-completed strings are presented in black. Unstudied and illegally-completed strings are presented in white. The string completion task is presented in circles. The recognition task is presented in triangles.*

The simulation suggests the possibility of an interesting selective impairment atypical of those generally observed in amnesic patients. However, the differences between cosines for correct compared to incorrect answers are very low, also suggesting poor performance. This prediction is contrary to Jamieson and Mewhort's (2010b) reports of successful string completion. Relatedly, raw cosine similarities are difficult to interpret. As such, I conducted an empirical simulation of the same tasks to provide a modeling target to facilitate interpretation of the model results.

Experiment 5

Experiment 5 was designed to provide an empirical target for Simulation 5 as a first test of the model's predictions. Although I am not testing organic amnesic patients, the empirical framework in Experiments 1 through 4 provides a sensible proxy as a first step. Experiment 5 replicates the string completion and recognition procedures adopted from Jamieson and Mewhort (2010b) using healthy participants by manipulating study time.

Methods

Participants. I recruited 80 participants from the University of Manitoba Introduction to Psychology participant pool. I randomly assigned them to four groups created by crossing the factors of study time (5000 ms, 100 ms) and task (string completion, recognition). All participants received course credit for participating.

Materials and Apparatus. The experiment was conducted using the same apparatus as Experiments 1 through 4. The letter strings in Table 4 were presented in black 36 point Arial font, in capital letters, and in the center of the screen.

Procedure. Participants viewed the 20 strings in the study list one at a time in random order. Participants were not informed about the grammar or its rules before study. Participants in the control group viewed each string for 5000 ms and participants in the simulated amnesia group viewed each string for 100 ms. The screen cleared for 250 ms between presentations. This process repeated until all 20 strings had been presented.

In the recognition task, participants were presented with studied strings and legally-completed unstudied strings. As in previous experiments, participants evaluated the strings one at a time in random order by selecting one of two response alternatives labeled “Old” and “New”. The screen then cleared for 250 ms. This process repeated until participants had made a response to all 40 strings.

In the string completion task, participants were informed that the studied strings were constructed according to a complex set of rules. They were then presented with the incomplete test strings and asked to complete the string to make it grammatical. The incomplete strings were presented one at a time in random order. Five response alternatives labeled “S”, “P”, “V”, “T”, and “X” were presented below the string. Participants responded by clicking one of the alternatives with the mouse. The screen then cleared for 250 ms. This process repeated until participants had provided a response to all 20 strings.

As with Experiments 1 through 4, Experiment 5 was approved by the Psychology / Sociology Research Ethics Board at the University of Manitoba.

Results

Figure 29 presents the results of Experiment 5. The left bars present the proportion of strings legally completed in the string completion task. The right bars present the participants’ discrimination (i.e., the difference between hits and false alarms) in the recognition task.

In the completion task, both study-time groups performed above chance (20%), being more likely to complete each string with the single letter that makes the string grammatical. This pattern was supported by two separate single-sample t tests, demonstrating that both groups performed significantly above chance, whether they had studied each string for 5000 ms, $t(19) = 7.48$, $p < .001$, or 100 ms, $t(19) = 4.37$, $p < .001$. There was also little change in accuracy as a function of study time, and an independent Welch’s t test found no evidence for a difference between the groups, $t(38) = 1.77$, $p = .084$.

In contrast, performance was near chance on the recognition task (overall accuracy was higher than in completion, but chance performance is 50% as opposed to

20%). As in the completion task, there was little difference in performance as a function of study time.

To evaluate the data, I analyzed the proportion of strings endorsed as "Old" in a 2 x 2 mixed-factors analysis of variance with study time (5000 ms, 100 ms) as a between-subjects factor and probe type (studied, unstudied) as a within-subjects factor. There was a main effect of probe type, $F(1, 38) = 7.63$, $p = .009$, such that studied strings were more likely to be endorsed than unstudied strings. It should be noted that performance, although significantly above chance, was quite low compared to recognition performance in the preceding experiments. There was no main effect of encoding time, $F(1, 38) = 0.01$, $p = .933$, and no interaction, $F(1, 38) = 1.64$, $p = .210$, finding no evidence for any differences between the study-time groups.

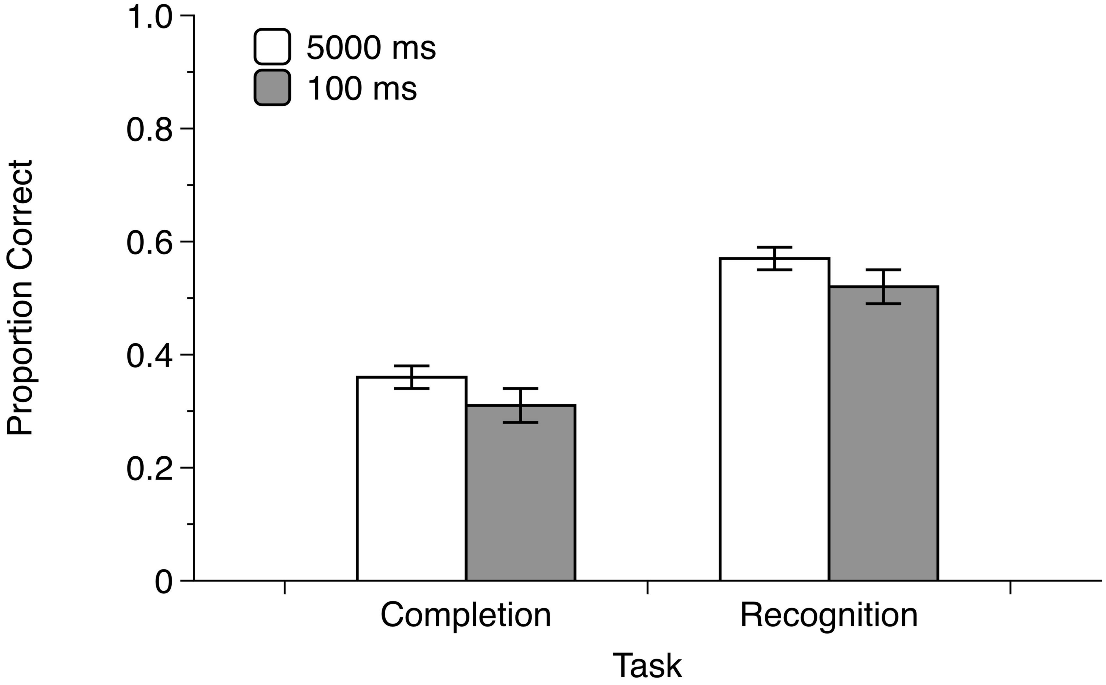

Figure 29. *The proportion of correct trials as a function of task and encoding time in in Experiment 5. The error bars represents standard error.*

Unfortunately, there appears to be no evidence of a dissociation-like pattern. Although there was no evidence for a difference in string completion, there was also no evidence for a difference in recognition. However, to evaluate the data more carefully, I analyzed the proportion of correct trials in a 2 x 2 between-subjects analysis of variance with study time (5000 ms, 100 ms) and task (completion, recognition) as factors. There was a main effect of study time, $F(1, 76) = 4.65$, $p = .034$, such that participants in the 5000 ms group were more accurate overall. There was also a main effect of task, $F(1, 76) = 83.37$, $p < .001$, such that accuracy was higher on the recognition task. However, that analysis does not correct for chance performance, which is much higher in the recognition task. As a result, the difference is not very meaningful and, in fact, misleading given that the difference from chance was greater in completion than in recognition. Finally, there was no interaction between the two factors, $F(1, 76) = 0.10$, $p = .753$, again showing no evidence for a selective impairment.

Model Fit to Experiment 5

Given that the completion task was essentially a five-alternative-forced-choice procedure, I did not transform the raw cosines similarities. Instead, I computed the probability of selecting the grammatical letter using the Luce choice axiom (Luce, 1959). For example, if the grammatical letter is S, the probability of selecting S is given by

$$\frac{cos_S}{cos_S + cos_P + cos_T + cos_V + cos_X}$$

where each *cos* is the cosine similarity between the probe containing the candidate letter and its resulting echo. To compute endorsement probabilities in recognition, I transformed the cosine similarities using the parameter values $a = 7.5$ and $b = 0.1$. To simulate the 100 ms group, $L = 0.2$, and to simulate the 5000 ms group, $L = 0.7$.

Figure 30 presents the simulation results from the full model. The model produced an extremely close fit to the data (*r* = .98), capturing the overall pattern of means. Critically, the computational and empirical simulations largely converge on the same general predictions, the most important of which is a consistency between string completion and recognition.

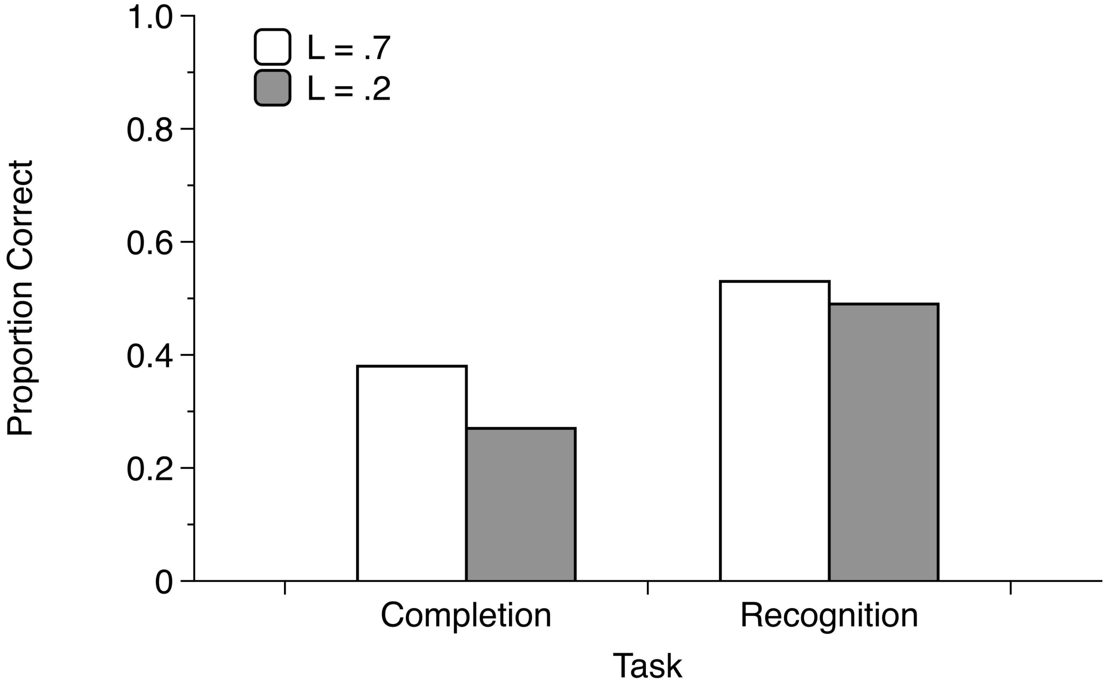

Figure 30. *Simulation 5 proportion of correct trials in string completion and recognition as a function of L.*

Discussion of Simulation 5 and Experiment 5

Simulation 5 and Experiment 5 were intended to derive a novel prediction within my analytic framework. The two analyses produced the same prediction. But, unfortunately, they do not seem to predict a difference between string completion and recognition in artificial grammar learning as I had hoped. The computational analysis, to an extent, seems to predict a surprising pattern; amnesia will cause a greater impairment to recognition than to completion. The pattern does not bear out statistically in the empirical analysis, but it does bear out numerically.

However, the prediction might be an artifact of poor overall recognition performance. Although the studied strings were endorsed more frequently than the unstudied strings, discrimination was very low. Additionally, although there was no statistical difference between the groups' performance, the 100 ms participants performed at levels that were extremely close to chance. It might be the case that the recognition task was simply too difficult, and floor levels of performance obscured any potential interaction dissociation. The task was certainly difficult; the studied and unstudied strings were very similar, often only differing by one letter. Future work ought to make the two probe types differ to a greater extent to detect an interaction, if one exists.

Despite the lack of evidence for a task difference, the data still provide a suboptimal but valid target for empirical work from other research groups working with memory-impaired populations. The technique of deriving, testing, and refining models of selective impairment can be – and ought to be – used to work toward a rational and coherent account of the cognitive processes underlying memory. Given more time, I hope to find a firm testable prediction from the model that can be carried out in the lab with amnesic patients.

Chapter 5 – General Discussion

By one account, memory is composed of multiple systems, each with a unique neural basis and principles of operation. Dissociation-like differences between memory tasks are a primary source of evidence for the multiple-system perspective, and amnesia is a key example of these patterns. Amnesic patients perform similarly to controls on some memory tasks but show severe impairments in others. By the multiple-system perspective, an explicit memory system that is damaged in amnesia supports performance on the impaired tasks, but an implicit memory system that is intact in amnesia supports performance on the unimpaired tasks.

In this dissertation, I set out to argue against the multiple-system perspective. I argued instead for a single-system perspective in which dissociations reflect a global encoding deficit that has quantitatively different influences on various memory tasks. I tested this premise using a combination of computational and empirical simulation.

I have reported four computational simulations of hallmark results. I have also reported four experiments in conjunction with those simulations. The computational and empirical analyses generally converge on the same points in support of the single-system perspective. But, there are inconsistencies and weaknesses in the larger and complete argument.

Simulation 1 modeled a difference between classification and recognition of dot patterns (Knowlton & Squire, 1993). Experiment 1 examined the difference by simulating amnesia in healthy participants. Both the computational and empirical simulations readily accommodated the results.

Simulation 2 modeled a difference between classification and recognition of binary-featured animal drawings (Reed et al., 1999). Experiment 2 examined the difference in a conceptual replication of Reed et al.'s experiment, using stimuli that differed in appearance but maintained the exact same category structure. Again, both the computational and empirical analyses readily accommodated the results.

Simulation 3A modeled a difference between tachistoscopic identification priming and recognition (Keane et al., 1995), and Experiment 3 examined the difference empirically. Both the computational and empirical analyses accommodated the results with only minor discrepancies.

Simulation 3B modeled a related double dissociation involving patients with occipital lobe damage (Gabrieli et al., 1995, Keane et al., 1995). The computational analysis failed to accommodate the double dissociation. However, although there were many procedural differences between experiments; a group of participants from Experiment 3 provided some insight into the double dissociation.

Simulation 4 modeled a pattern of differences among word-stem completion, cued recall, and recognition (Graf, Squire, & Mandler, 1984). Experiment 4 examined the pattern empirically. The computational and empirical analyses accommodated pieces of the results but were generally inconsistent with the original results.

Additionally, Simulation 5 and Experiment 5 leveraged the analysis to derive a novel prediction. I tested a potential difference between string completion and recognition in an artificial grammar task. However, the approach serves as an example of how I might proceed to derive a clearer empirical test of the model in additional work.

Taken together, the dissertation provides some positive evidence for a single-system approach to understanding selective impairments associated with amnesia. However, it also tempers the strength of that endeavor by providing a patchwork of mostly positive but some negative outcomes. Still, the work on the whole makes several meaningful contributions.

Computational Contribution

In my computational work, I have adopted a framework in the context of the HEM. The main premise of the computational simulations is that impairment patterns in amnesia can be understood as a global encoding deficit, represented by the learning parameter, L, in the model. As L decreases, less complete information is encoded to memory. I predicted that the resulting incomplete traces would support accurate classification and normal priming but would be insufficient to support accurate recognition and recall. My predictions were supported in most cases but not in others.

Previous computational analyses have also examined impairment patterns in amnesia. The GCM (Nosofsky & Zaki, 1998; Zaki & Nosofsky, 2001) has been applied to results involving classification. The SRN (Kinder & Shanks, 2001, 2003) has been applied to results involving priming. I applied the HEM to the broader set of patterns. Whereas the models approach selective impairments using computational mechanisms that differ in detail, they are related and converge on common conclusions.

In the HEM, I explained selective memory impairment as a deficit to encoding, as represented by L. In the GCM, the differences are a function of sensitivity to similarity, as represented by their parameter c. Although the details of the mechanisms differ, they share mathematical and conceptual similarities. The primary effect of decreasing L in the

HEM is to produce incomplete traces in memory. The resulting representations have little information on which to base discrimination, and the model's similarity-based retrieval mechanisms are less sensitive as a result. Decreasing sensitivity is the primary mechanism utilized in the GCM. Similarly, in the SRN, selective impairments are a function of the learning rate of the neural network. This parameter is, again, different in detail, but related in spirit. A lower learning rate results in memory representations that are less complete approximations of the studied stimulus.

Some of the motivation for my simulations was in response to problems with and objections to the GCM and SRN. Although the HEM alleviates some of those problems, my intent is not to argue against the accounts provided by the GCM and SRN. Rather, the evidence from the different frameworks converges onto a common point, namely that the dissociations observed in amnesia can be, and ought to be, understood from a common and unified set of principles. The convergence of the models in tandem provides a much richer analysis and a more complete story than the successes and failures of any of the models in isolation.

Empirical Contribution

Based on the premises from my computational analysis, I have adopted an empirical framework that manipulates study time in lieu of organic amnesia. Whereas my results suggest that many of the impairment differences between memory tasks can be produced by a manipulation of study time, other researchers have adopted different techniques such as delay between study and test (Nosofsky & Zaki, 1998; Zaki & Nosofsky, 2001, Woods & Percy, 1974) and differences in elaborative processing (Graf, Mandler, & Haden, 1982).

In contrast to my expectations, it was the results of the empirical and not of the computational simulations that were generally very successful in accommodating the patterns observed in amnesia. The results converge with my computational simulations, providing two related analyses that support the view that selective impairments do not force distinctions between memory systems. The success of the method and approach also suggests that there may be profit in pursuing a broader and more vigorous program of empirical simulation of memory disorder to identify which mechanisms and processes might be compromised in the disorder.

Importantly, in my empirical simulations, I chose to manipulate study time due to its close correspondence with the parameter manipulation in my computational simulations. However, like my computational contribution, I do not intend to argue that amnesia is best understood as equivalent to short study time. Instead, the fact that reducing study time, increasing delay, encouraging shallow processing, and damaging the medial temporal lobes all produce similar dissociations stand as a convergent empirical target for rational analysis and theory development.

Objections, Limitations, and Future Directions

There are, of course, a number of objections that could be made against my arguments. I consider one objections preemptively, namely the argument that my analysis is inconsistent with neurological data.

Although I have reported evidence that single-system models can accommodate behavioural differences in amnesic performance, they are arguably unable to account for neurological differences. A wealth of research suggests that putative explicit and implicit memory tasks result in neural activation in distinct brain regions. For example, Reber, Stark, and Squire (1998; see also Reber, Gitelman, Parrish, & Mesulam, 2003) demonstrated that dot pattern classification elicits decreased activity in the posterior occipital cortex but recognition elicits increased activity in the posterior occipital cortex as well as other regions, including the hippocampus. Thus it might be argued that, although the behavioural evidence does not force a systems distinction, the neurological evidence does.

However, Nosofsky, Little, and James (2012) have argued that the neurological evidence is also consistent with a single-system perspective. They replicated Knowlton and Squire's (1993) dot pattern experiment with healthy participants while measuring BOLD signal in an fMRI. In addition to the classification and standard recognition tasks, participants completed a "lax" recognition task in which they were instructed to emphasize hits at the expense of false alarms; a manipulation designed to invite participants to adopt a lax decision criterion. The multiple-system perspective predicts similar brain activity in the two recognition tasks compared to the classification task based on the premise that both tasks rely on explicit memory. In contrast, the single-system perspective predicts similar brain activity in the classification task and lax recognition task compared to the standard recognition task based on the premise that the differences in brain activity reflect a shift in criterion. The results favoured the single-system perspective. Whereas the standard recognition task elicited increased activity in the frontal eye fields and anterior insular cortex, both the lax recognition and classification tasks elicited a decreased activity in the same regions. By assuming a parameter change rather than a systems distinction, a single-system perspective is consistent with the behavioural *and* neurological differences.

Admittedly, there is additional neurological evidence that has not been accommodated within the single-system perspective. A wealth of research suggests that, whereas the hippocampal formation plays a central role in explicit memory, implicit memory is supported by alternative brain regions (e.g., Foerde & Shohamy, 2011; Opitz, 2010; Squire, 1992, 2004; Squire, Knowlton, & Musen, 1993; Seger, 1994; although see Graham, Barense, & Lee, 2010; Lech & Suchan, 2013). The exact locus of implicit memory is still largely unknown, likely differs for different implicit memory tasks, and might be supported by highly distributed neural patterns rather than particular regions (Reber, 2013). Despite the uncertainty regarding the neurological basis of implicit memory, a reliable finding is that it does not rely on the hippocampal formation.

The neurological data are taken by many as evidence that the single-system perspective is bound to be incorrect. In humans, the patterns observed in amnesia are the result of selective damage to particular neural systems. Thus, even if a single-system account can produce behavioural patterns, it does so by ignoring anatomical fact. However, Nosofsky, Little, and James' (2012) results demonstrate that the inference from the neurological data is not that simple. Rather, the proper interpretation of neurological data is complex and still unclear. The uncertainty does not, of course, relieve a computational analysis from explaining neurological data in addition to behavioural data. On that point, my analysis is admittedly deficient. However, it does undermine the neurological data's ability to reject the single-system perspective outright.

A primary reason that my proposed analysis does not speak to neurological data is that the analysis approaches the theoretical problem at a relatively high level of abstraction (Marr, 1982). Whereas my analysis roughly corresponds to an algorithmic level of explanation, an analysis of neural function operates at the level of implementation. But, there are ways that one might go about bridging those levels of description.

Artificial neural networks are one possible analytic framework to examine selective impairments at a neural level. For example, the SRN is a neural model applied to differences between priming and recognition (Kinder & Shanks, 2001, 2003), and neural network models have been applied to differences between anterograde and retrograde amnesia (McClelland, McNaughton, & O'Reilly, 1995). I have focused my analysis on an algorithmic interpretation. However, future work ought to incorporate analyses across level of explanation and abstraction by adopting more directly and explicitly biologically-based models in addition to system models. The shift from system models such as MINERVA2 or the HEM to biological models is not necessarily a large one. For example, Hintzman (1990) re-described MINERVA2 as a multi-layered neural network where each trace is represented as a unit in the hidden layer. Despite mathematical and computational differences, the analyses can converge on the same conclusions, organizing the empirical database under common principles.

Continuous Learning Paradigms

In all of the results I have reported, there are discrete study and test phases. However, amnesia also produces patterns of selective impairment in continuous learning experiments in which study and test occur on each trial. For example, in a probabilistic classification task, participants learn about combinations of cues that are associated with particular outcomes. Participants might, for example, be presented with combinations of symptoms and predict one of two diseases. However, the cues are associated with the outcomes probabilistically rather than deterministically. As learning proceeds, participants begin to match the probability relations between the cues and outcomes. For example, if an outcome co-occurs with a cue combination 80% of the time, participants will select that outcome on 80% of the trials (Shimp, 1966).

Probabilistic classification tasks have been used in two primary lines of research. One line of research examines the variety of strategies participants employ for probabilistic classification (e.g., Ashby & Ells, 2001; Gluck, Shohamy, & Myers, 2002; Lagnado, Newell, Kahan, & Shanks, 2006; Meeter, Myers, Shohamy, Hopkins, & Gluck, 2006; Meeter, Radics, Myers, Gluck, & Hopkins, 2008). Another line of research examines the interacting roles of the basal ganglia and striatum in supporting probabilistic classification. (e.g., Aron, Shohamy, Clark, Myers, Gluck, & Poldrack, 2004; Foerde, Knowlton, & Poldrack, 2006; Shohamy, Myers, Kalanithi, & Gluck, 2008). Debates between the multiple-system and single-system perspectives are present but not at the forefront of current empirical examination.

However, Knowlton, Squire, and Gluck (1994) tested amnesic patients and controls in a set of probabilistic classification tasks. In each task, four cues were probabilistically associated with two different outcomes. In the first task, symptoms were associated with two diseases. In the second task, patterned cards were associated with rain or sun. During each of 350 trials, one, two, three, or four of the cues were presented and participants predicted the outcome that was determined by a fixed probability of occurrence for each cue combination. A response was measured correct if the participant selected the most probable outcome independent of which outcome actually occurred.

Early in learning (the first 50 trials), amnesic patients and controls showed similar learning curves. However, later in learning (the remaining 300 trials), the controls' predictions became closer to the outcome probabilities while the amnesic patients' probability matching plateaued. Knowlton, Squire, and Gluck (1994) took the difference as additional evidence for a distinction between memory systems. Implicit memory supports probabilistic classification, but only early in learning. Late learning is additionally supported by explicit memory. As a result, amnesic patients show intact early learning but impaired late learning.

The HEM can be adapted to continuous learning situations. Rather than supplying memory with the entire set of study items and then probing memory with the entire set of test items, memory can be probed prior to each encoding event. As a result, study and test are combined into each trial.

The extension is straightforward. However, an initial simulation showed that the model fails to produce sensible predictions. Under some circumstances, the model does show gradual learning, but that learning is largely uninfluenced by the learning parameter, both early and late in learning. Under other circumstances, the model fails to learn at all. A computational analysis of the difference between early and late learning requires the exploration of either adaptations to the HEM or the exploration of other models.

An empirical analysis of the pattern is also problematic, at least in the framework of this dissertation. If each trial was only presented for 100 ms, any change in performance is confounded due to the combination of study and test on each trial. The change might be due to poor encoding of the cue-outcome association, inability to use the cues to probe memory, or some combination of both factors. Any result, whether a successful or failed replication of the original pattern, would be ambiguous at best and uninterpretable at worst. Consequently, I have excluded the analysis I conducted from this dissertation. Nevertheless, I consider the selective impairments in continuous learning paradigms to remain an important target for future work.

Extension to Aging

I have focused on selective impairments observed in amnesia. However, same patterns occur across a wide variety of memory impairments, both severe and mild. The patterns of results discussed in preceding chapters generalize – with some caveats – to other populations such as Alzheimer's disease patients and those with mild cognitive impairment (e.g., Eldridge, Masterman, & Knowlton, 2002; Kéri et al., 1999; LaVoie & Faulkner, 2008; Zannino et al., 2012). Further, the same patterns are observed in the general decline of memory performance even in healthy aging (e.g., Burke & Light, 1981; Light & Singh, 1987).

Benjamin (2010) provided a computational account of differences between recollection and familiarity in healthy aging. Older adults are impaired in source memory but not in item memory (Benjamin & Craik, 2001; Spencer & Raz, 1995). The result is taken as evidence for a distinction between recollection and familiarity. However, Benjamin argued against this conclusion by accommodating the same pattern in an exemplar model of memory, DRYAD.

Benjamin (2010) also speculated about an extension of DRYAD to impairment differences between explicit and implicit memory in healthy aging (Light & Singh, 1987). With some minor exceptions, DRYAD is formally equivalent to MINERVA2, thus an extension from DRYAD to MINERVA2 – and, consequently, the HEM – is straightforward. Variation in L permits a natural extension to other forms of memory impairment. For example, older adults are impaired on measures of explicit memory, but not nearly to the same extent as amnesic patients. Therefore it might be possible to explain the patterns observed in healthy aging as a moderate encoding deficit, represented by an intermediate value of L. In fact, this is precisely what Benjamin argues. Future work ought to extend analysis in the HEM to healthy aging as well as other patient populations. To the extent that the model accounts for the data, the analysis will provide an explanation for a suite of phenomena under the same theoretical principles. To the extent that the model does not account for the data, the analysis will provide directions for future research into the etiology and possible treatment of memory disorders, having both theoretical and practical implications.

Comments on Model Fitting

One might wonder why I have not reported any formal model fitting procedures in my computational simulations. The decision was deliberate and based on two key considerations: (1) the nature of model fits in psychology and (2) the value of models in psychology.

Formal models are generally evaluated based on their goodness of fit to empirical data. However, it is unclear what constitutes a sufficiently close fit (e.g., Estes, 1975, 2002). Consequently, models are evaluated in relation to alternative models, selecting one over the other on the basis of goodness of fit and, more recently, model complexity (e.g., Busemeyer & Wang, 2000; Myung, 2000; Pitt, Myung, & Zhang, 2002; Zucchini, 2000).

In the absence of competing multiple-system models, a principled evaluation of fit is impossible by these methods.

More critically, focusing exclusively on model fits to experimental cases understates the value of computational models in psychology. The insights provided by a model are far more important that the precision with which it accounts for data. Of course, a model must provide at least a reasonable fit to the data. However, arguing over whether the model is correct seems to be the wrong tact.

I am certainly not the first to make this argument. In a famous aphorism, Box (1979) stated that:

> …it would be very remarkable if any system existing in the real world could be *exactly* represented by any simple model. However, cunningly chosen parsimonious models often do provide remarkably useful approximations. For example, the law $PV = RT$ relating pressure P, volume V, and temperature T of an "ideal" gas via a constant R is not exactly true for any real gas, but it frequently provides a useful approximation and further its structure is informative since it springs from a physical view of the behavior of gas molecules. For such a model there is no need to ask the question "Is the model true?" If "truth" is to be the "whole truth" the answer must be "No". The only question of interest is "Is the model illuminating and useful?". (pp. 2-3)

The aim of my simulations follows this argument. Independent of the correctness of the HEM, the simulations demonstrate that a range of seemingly disparate phenomena – and complex relationships among those phenomena – can be produced by a common set of operations.

Unified Theories of Psychology

A primary reason I have argued against the multiple-system perspective is that it, in its current trajectory, poses some treats to developing a unified theory of psychology. Given that unification is a core goal of scientific study, this situation is problematic.

As in the preceding section, I am not the first to make this argument. For example, Shanks and St. Johns (1994) and 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 the entire history of psychological research has adopted a strategy of designing experiments to solve forced binary oppositions such as: serial versus parallel search, analog versus digital representations, decay versus interference, and so on. He also argued that, no matter how rigorously tested, these oppositions will not culminate as an integrated or principled theory of psychology.

The multiple-system perspective is not inherently in opposition to a unified theory of psychology. For example, Newell (1990) is upfront in stating that there will necessarily be distinctions between cognitive systems. However, our current taxonomy of those systems has been built on insufficient evidence. Rather than proposing new systems of memory for the sake of explanatory convenience, we ought to adopt an approach in which we propose new systems when they are necessitated by the data. The approach I have adopted challenged the ability of our existing data to necessitate our existing distinctions. However, there were cases where the analysis failed.

The failures of my analysis might imply the need for a separate system of memory. The exact nature of that system is currently unclear and should be a target for future theoretical work. However, a key difference in the successes and failure of my analysis suggest a potential distinction. The primary failure of the analysis (Simulation and Experiment 4) involved verbal information. In contrast, the successes all involved non-verbal information. It is possible that a more appropriate division of memory systems is a verbal and non-verbal store. Of course, this division must be tested through careful theoretical and empirical analysis.

Additionally, the exact locus of a verbal / non-verbal distinction is unclear. It might be the case that one system stores verbal information while another stores non-verbal information. However, it might be the case that two distinct processes work within a single system, only one of which operates on verbal information. Finally, it might be the case that the same processes operate on verbal and non-verbal information, but the two types of information rely on different forms of representation. In fact, Hintzman (1986) and Jamieson, Mewhort, and Hockley (2016) have presented implementations of MINERVA2 in which different types of information correspond to different features in a representation. The complex relationship between systems, processes, and representation itself suggests the need for a carefully integrated theoretical perspective on memory and cognition.

Summary and Conclusions

The analysis in this dissertation stemmed from three primary goals. The first goal was to work toward a formal and principled account of the structure and function of memory by combining computational and empirical approaches. The correspondence between the parameter manipulation in the computational work and the experimental manipulation in the empirical work is crucial to this goal. The two analyses approach the problem from different angles and using different techniques. However, the correspondence reflects that they are based on the same fundamental principles.

The second goal was to provide an analysis that converges with previous work. Although the details of my analysis differ from other accounts, they share the same core concepts. Interpreting the results of previous work and my proposed work under the same framework will permit movement toward a more coherent and unified theory of memory – providing progress in theoretical clarity rather than debate over the particular details in our fictional models of nature.

My third goal is to argue in favour of the single-system over the multiple-system perspective. The goal is peripheral to the first two goals, and if firm data necessitates taking the multiple-system perspective, then that perspective ought to be adopted. However, the argument is important to ensure that theories develop in the correct direction. If a theory develops from false premises, the theory will be misguided at best and straightforwardly false at worst. More importantly, no scientific conclusions should be allowed to develop without productive skepticism and challenge. At least, my contribution will serve as a firm test to advance and refine of the multiple-system theory of memory. To the extent that the analysis succeeded, it questions the necessity of systems distinctions in memory theory and suggests avenues of exploration to further refine theoretical divisions in the cognitive system.

In summary, this dissertation was aimed at offering a rational analysis of our current empirical database, with a particular focus on the patterns of selective impairment observed in memory impairment associated with amnesia. The work provides additional means to develop foundational principles that can be leveraged to guide theory construction and refinement, ultimately working towards a more integrated and unified approach to memory.

References

Arndt, J., & Hirshman, E. (1998). True and false recognition in MINERVA2: Explanations from a global matching perspective. *Journal of Memory and Language, 39*(3), 371-391.

Aron, A. R., Shohamy, D., Clark, J., Myers, C., Gluck, M. A., & Poldrack, R. A. (2004). Human midbrain sensitivity to cognitive feedback and uncertainty during classification learning. *Journal of Neurophysiology, 92*(2), 1144-1152. doi:http://dx.doi.org/10.1152/jn.01209.2003

Ashby, F. G., & Ells W. (2001). The neurobiology of human category learning. *Trends in Cognitive Sciences, 5*(5), 204-210. doi:http://dx.doi.org/10.1016/S1364-6613(00)01624-7

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. *Psychology of learning and motivation*, *2*, 89-195.

Bamber, D. (1979). State-trace analysis: A method of testing simple theories of causation. *Journal of Mathematical Psychology*, *19*(2), 137-181.

Benjamin, A. S. (2010). Representational explanations of “process” dissociations in recognition: The DRYAD theory of aging and memory judgments. *Psychological Review, 117*(4), 1055-1079. doi:http://dx.doi.org/10.1037/a0020810

Benjamin, A. S., & Craik, F. I. M. (2001). Parallel effects of aging and time pressure on memory for source: Evidence from the spacing effect. *Memory & Cognition, 29*(5), 691-697. doi:http://dx.doi.org/10.3758/BF03200471

Bowers, J. S. (2000). In defense of abstractionist theories of repetition priming and word identification. *Psychonomic Bulletin & Review*, *7*(1), 83-99.

Box, G. E. P. (1979). Robustness in the strategy of scientific model building. *Robustness in Statistics, 1,* 201-236.

Brooks, L. R. (1987). Decentralized control of categorization: The role of prior processing episodes. In U. Neisser (Ed.), Concepts and conceptual development: Ecological and intellectual factors in categorization, pp. 141-174. New York, NY: Cambridge University Press.

Brooks, L. R., & Vokey, J. R. (1991). Abstract analogies and abstracted grammars: Comments on Reber (1989) and Mathews et al. (1989). *Journal of Experimental Psychology: General, 120*(3), 316-323. doi:http://dx.doi.org/10.1037/0096-3445.120.3.316

Brown, G. D., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological review*, *107*(1), 127-181.

Brysbaert, M., & New, B. (2009). Moving beyond Kučera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods, 41*(4), 977-990.

Burke, D. M., & Light, L. L. (1981). Memory and aging: The role of retrieval processes. *Psychological Bulletin, 90*(3), 513-546. doi:http://dx.doi.org/10.1037/0033-2909.90.3.513

Busemeyer, J. R., & Wang, Y. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology, 44*, 171-189. doi:10.1006/jmps.1999.1282

Cave, C. B., & Squire, L. R. (1992). Intact verbal and nonverbal short-term memory following damage to the human hippocampus. *Hippocampus*, *2*(2), 151-163.

Cermak, L. S., Talbot, N., Chandler, K., & Wolbarst, L. R. (1985). The perceptual priming phenomenon in amnesia. *Neuropsychologia, 23*(5), 615-622.

Cermak, L. S., Verfaellie, M., Milberg, W., Letourneau, L., & Blackford, S. (1991). A further analysis of perceptual identification priming in alcoholic Korsakoff patients. *Neuropsychologia, 29*(8), 725-736.

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*(6), 1049-1055.

Clark, S. E. (1997). A familiarity-based account of confidence–accuracy inversions in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(1), 232.

Cleeremans, A. (1993). *Mechanisms of implicit learning: Connectionist models of sequence processing*. MIT press.

Cleeremans, A., & McClelland, J. L. (1991). *Learning the structure of event sequences.* Lawrence Erlbaum Associates, Inc, Hillsdale, NJ.

Cohen, N. J. (1984). Preserved learning capacity in amnesia: Evidence for multiple memory systems. *Neuropsychology of Memory*, *1*, 1984.

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*(4466), 207-210.

Cowan, N. (1988). Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological Bulletin, 104,* 163–191. doi:10.1037/0033-2909.104.2.163

Craik, F. I., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of verbal learning and verbal behavior*, *11*(6), 671-684.

Curran, T., & Hintzman, D. L. (1995). Violations of the independence assumption in process dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*(3), 531-547.

Delazer, M., Ewen, P., & Benke, T. (1997). Priming arithmetic facts in amnesic patients. *Neuropsychologia, 35*(5), 623-634.

Diamond, R. J., & Rozin, P. (1984). Activation of existing memories in anterograde amnesia. *Journal of Abnormal Psychology,93*(1), 98-105. doi:http://dx.doi.org/10.1037/0021-843X.93.1.98

Dienes, Z. (1992). Connectionist and memory-array models of artificial grammar learning. *Cognitive Science, 16*(1), 41-79.

Dienes, Z., Broadbent, D., & Berry, D. C. (1991). Implicit and explicit knowledge bases in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17*(5), 875-887. doi:http://dx.doi.org/10.1037/0278-7393.17.5.875

Dougherty, M. R., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory processes model for judgments of likelihood. *Psychological Review*, *106*(1), 180.

Dunn, J. C., & Kirsner, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review, 95*(1), 91-101. doi:http://dx.doi.org/10.1037/0033-295X.95.1.91

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

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

Eldridge, L. L., Masterman, D., & Knowlton, B. J. (2002). Intact implicit habit learning in Alzheimer's disease. *Behavioral Neuroscience, 116*(4), 722-726. doi:http://dx.doi.org/10.1037/0735-7044.116.4.722

Elman, J. L. (1990). Finding structure in time. *Cognitive science*, *14*(2), 179-211.

Estes, W. K. (1975). Some targets for mathematical psychology. *Journal of Mathematical Psychology, 12*, 263-282.

Estes. W. K. (2002). Traps in the route to models of memory and decision. *Psychonomic Bulleting & Review, 9*(1), 3-25.

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*(4), 556-571. doi:http://dx.doi.org/10.1037/0278-7393.15.4.556

Feustel, T. C., Shiffrin, R. M., & Salasoo, A. (1983). Episodic and lexical contributions to the repetition effect in word identification. *Journal of Experimental Psychology: General*, *112*(3), 309.

Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. MIT press.

Foerde, K., Knowlton, B. J., & Poldrack, R. A. (2006). Modulation of competing memory systems by distraction. *PNAS Proceedings of the National Academy of Sciences of* *the United States of America, 103*(31), 11778-11783. doi:http://dx.doi.org/10.1073/pnas.0602659103

Foerde, K., & Shohamy, D. (2011). The role of the basal ganglia in learning and memory: Insight from Parkinson’s disease. *Neurobiology of Learning and Memory, 96*(4), 624-636. doi:http://dx.doi.org/10.1016/j.nlm.2011.08.006

French, R. M. (1991) "Using semi-distributed representations to overcome catastrophic forgetting in connectionist networks" in Proceedings of the 13th Annual Cognitive Science Society Conference. Hillsdale, NJ: Lawrence Erlbaum, 173-178.

Gabrieli, J. D. E., Fleischman, D. A., Keane, M. M., Reminger, S. L., & Morrell, F. (1995). Double dissociation between memory systems underlying explicit and implicit memory in the human brain. *Psychological Science, 6*(2), 76-82. Retrieved from http://search.proquest.com/docview/618698060?accountid=14569

Glanzer, M. (1972). Storage mechanisms in recall. In G. H. Bower (Ed.), *Psychology of learning and motivation*, Vol. 5. New York: Academic Press pp. 129-193.

Glanzer, M., & Adams, J. K. (1990). The mirror effect in recognition memory: Data and theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*(1), 8-20.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General, 117*(3), 227-247. doi:http://dx.doi.org/10.1037/0096-3445.117.3.227

Gluck, M. A., Shohamy, D., & Myers, C. (2002). How do people solve the “weather prediction” task?: Individual variability in strategies for probabilistic category learning. *Learning & Memory*, *9*(6), 408-418.

Graf, P., Mandler, G., & Haden, P. E. (1982). Simulating amnesic symptoms in normal subjects. *Science*, *218*(4578), 1243-1244.

Graf, P., & Schacter, D. L. (1985). Implicit and explicit memory for new associations in normal and amnesic subjects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*(3), 501-518. doi:http://dx.doi.org/10.1037/0278-7393.11.3.501

Graf, P., Shimamura, A. P., & Squire, L. R. (1985). Priming across modalities and priming across category levels: Extending the domain of preserved function in amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*(2), 386-396. doi:http://dx.doi.org/10.1037/0278-7393.11.2.386

Graf, P., Squire, L. R., & Mandler, G. (1984). The information that amnesic patients do not forget. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*(1), 164-178. doi:http://dx.doi.org/10.1037/0278-7393.10.1.164

Graham, K. S., Barense, M. D., & Lee, A. C. H. (2010). Going beyond LTM in the MTL: A synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. *Neuropsychologia, 48*(4), 831-853. doi:http://dx.doi.org/10.1016/j.neuropsychologia.2010.01.001

Graham, K. S., Scahill, V. L., Hornberger, M., Barense, M. D., Lee, A. C. H., Bussey, T. J., & Saksida, L. M. (2006). Abnormal categorization and perceptual learning in patients with hippocampal damage. *The Journal of Neuroscience, 26*(29), 7547-7554. doi:http://dx.doi.org/10.1523/JNEUROSCI.1535-06.2006

Hamann, S. B., & Squire, L. R. (1997). Intact perceptual memory in the absence of conscious memory. *Behavioral Neuroscience,111*(4), 850-854.

doi:http://dx.doi.org/10.1037/0735-7044.111.4.850

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

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review, 93*(4), 411-428. doi:http://dx.doi.org/10.1037/0033-295X.93.4.411

Hintzman, D. L. (1987). Recognition and recall in MINERVA 2: Analysis of the 'recognition-failure' paradigm. *Modeling cognition.* (pp. 215-229) John Wiley & Sons, Oxford.

Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, *95*(4), 528.

Hintzman, D. L. (1990). Human learning and memory: Connections and dissociations. *Annual Review of Psychology*, *41*(1), 109-139.

Hirshman, E. (1998). On the logic of testing the independence assumption in the process-dissociation procedure. *Memory & Cognition, 26*(5), 857-859.

Homa, D., Cross, J., Cornell, D., Goldman, D., & Schwartz, S. (1973). Prototype abstraction and classification of new instances as a function of number of instances defining the prototype. *Journal of Experimental Psychology*, *101*, 116-122.

Jacoby, L. L. (1983). Remembering the data: Analyzing interactive processes in reading. *Journal of Verbal Learning and Verbal Behavior*, *22*(5), 485-508.

Jacoby, L. L. (1988). Memory observed and memory unobserved. *Remembering reconsidered: Ecological and traditional approaches to the study of memory.* (pp. 145-177) Cambridge University Press, New York, NY.

doi:http://dx.doi.org/10.1017/CBO9780511664014.007

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*(5), 513-541.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General, 110*(3), 306-340. doi:http://dx.doi.org/10.1037/0096-3445.110.3.306

James, W. (1918). The Principles of Psychology. New York: Holt. (Original work published 1890)

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*(6), 1529-1535. doi:http://dx.doi.org/10.1037/a0020598

Jamieson, R. K., & Mewhort, D. J. K. (2005). The influence of grammatical, local, and organizational redundancy on implicit learning: An analysis using information theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition,31*(1), 9-23. doi:http://dx.doi.org/10.1037/0278-7393.31.1.9

Jamieson, R. K., & Mewhort, D. J. K. (2009). Applying an exemplar model to the serial reaction-time task: Anticipating from experience. *The Quarterly Journal of Experimental Psychology, 62*(9), 1757-1783. doi:http://dx.doi.org/10.1080/17470210802557637

Jamieson, R. K., & Mewhort, D. J. K. (2009b). Applying an exemplar model to the artificial-grammar task: Inferring grammaticality from similarity. *The Quarterly Journal of Experimental Psychology*, *62*(3), 550-575.

Jamieson, R. K., & Mewhort, D. J. K. (2010b). Applying an exemplar model to the artificial-grammar task: String completion and performance on individual items. *The Quarterly Journal of Experimental Psychology, 63*(5), 1014-1039. doi:http://dx.doi.org/10.1080/17470210903267417

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*(2), 209-216.

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

Keane, M. M., Gabrieli, J. D. E., Mapstone, H. C., Johnson, K. A., & Corkin, S. (1995). Double dissociation of memory capacities after bilateral occipital-lobe or medial temporal lobe lesions. *Brain: A Journal of Neurology, 118*(5), 1129-1148.

Keane, M. M., Gabrieli, J. D. E., Noland, J. S., & McNealy, S. I. (1995). Normal perceptual priming of orthographically illegal nonwords in amnesia. *Journal of the International Neuropsychological Society, 1*(5), 425-433.

Kéri, S., Kálmán, J., Rapcsak, S. Z., Antal, A., Benedek, G., & Janka, Z. (1999). Classification learning in Alzheimer's disease. *Brain: A Journal of Neurology, 122*(6), 1063-1068.

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*(5), 648-669. doi:http://dx.doi.org/10.1162/089892901750363217

Kinder, A., & Shanks, D. R. (2003). Neuropsychological dissociations between priming and recognition: A single-system connectionist account. *Psychological Review, 110*(4), 728-744. doi:http://dx.doi.org/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*(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*(5140), 1747-1749. Retrieved from http://search.proquest.com/docview/618452597?accountid=14569

Knowlton, B. J., & Squire, L. R. (1994). The information acquired during artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(1), 79-91. doi:http://dx.doi.org/10.1037/0278-7393.20.1.79

Knowlton, B. J., & Squire, L. R. (1995). Remembering and knowing: Two different expressions of declarative memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*(3), 699-710. doi:http://dx.doi.org/10.1037/0278-7393.21.3.699

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory, 1*(2), 106-120.

Kolers, P. A., & Roediger, H. L. (1984). Procedures of mind. *Journal of Verbal Learning & Verbal Behavior, 23*(4), 425-449.

Kolodny, J. A. (1994). Memory processes in classification learning: An investigation of amnesic performance in categorization of dot patterns and artistic styles. *Psychological Science, 5*(3), 164-169.

Kwantes, P. J. (2005). Using context to build semantics. *Psychonomic Bulletin & Review, 12*(4), 703-710. doi:http://dx.doi.org/10.3758/BF03196761

Kwantes, P. J., & Mewhort, D. J. K. (1999). Modeling lexical decision and word naming as a retrieval process. *Canadian Journal of Experimental Psychology/Revue Canadienne De Psychologie Expérimentale, 53*(4), 306-315. doi:http://dx.doi.org/10.1037/h0087318

Kwantes, P. J., & Neal, A. (2006). Why people underestimate y when extrapolating in linear functions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*(5), 1019-1030. doi:http://dx.doi.org/10.1037/0278-7393.32.5.1019

Lagnado, D. A., Newell, B. R., Kahan, S., & Shanks, D. R. (2006). Insight and strategy in multiple-cue learning. *Journal of Experimental Psychology: General*,*135*(2), 162.

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*(2), 211-240.

LaVoie, D. J., & Faulkner, K. M. (2008). Production and identification repetition priming in amnestic mild cognitive impairment.*Aging, Neuropsychology, and Cognition, 15*(4), 523-544. doi:http://dx.doi.org/10.1080/13825580802051497

Leboe-McGowan, J. P., & Whittlesea, B. W. (2013). 17 Through the SCAPE Looking Glass--Sources of Performance and Sources of Attribution. *The Oxford Handbook of Cognitive Psychology*, 243-266.

Lech, R. K., & Suchan, B. (2013). The medial temporal lobe: Memory and beyond. *Behavioural Brain Research, 254*, 45-49. doi:http://dx.doi.org/10.1016/j.bbr.2013.06.009

Lewandowsky, S. (1993). The rewards and hazards of computer simulations. *Psychological Science, 4*(4), 236-243.

Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review, 96*(1), 25-57.

Light, L. L., & Singh, A. (1987). Implicit and explicit memory in young and older adults. *Journal of Experimental Psychology: Learning, Memory, and Cognition,13*(4), 531.

Lories, G., Van, d. L., & Cornille, M. (1994). Categorization and amnesia: Maintenance and decline of the typicality effect. *Cahiers De Psychologie Cognitive/Current Psychology of Cognition, 13*(2), 189-206.

Luce, R. D. (1959). On the possible psychophysical laws. *Psychological Review, 66*(2), 81-95.

Maine de Biran, F.-P.-G. 1929. The Influence of Habit on the Faculty of Thinking. Baltimore: Williams & Wilkins. (Original work published 1804)

Manza, L., & Reber, A. S. (1997). Representing artificial grammars: Transfer across stimulus forms and modalities. *How implicit is implicit learning?* (pp. 73-106) Oxford University Press, New York, NY.

Marr, D. (1982). *Vision: A Computational Approach,* San Fransisco, Freeman & Co.

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

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102*(3), 419-457. doi:http://dx.doi.org/10.1037/0033-295X.102.3.419

McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General, 114*(2), 159-188. doi:http://dx.doi.org/10.1037/0096-3445.114.2.159

McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of learning and motivation*, *24*, 109-165.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*(3), 207-238. doi:http://dx.doi.org/10.1037/0033-295X.85.3.207

Meeter, M., Myers, C. E., Shohamy, D., Hopkins, R. O., & Gluck, M. A. (2006). Strategies in probabilistic categorization: Results from a new way of analyzing performance. *Learning & Memory*, *13*(2), 230-239.

Meeter, M., Radics, G., Myers, C. E., Gluck, M. A., & Hopkins, R. O. (2008). Probabilistic categorization: How do normal participants and amnesic patients do it? *Neuroscience and Biobehavioral Reviews, 32*(2), 237-248. doi:http://dx.doi.org/10.1016/j.neubiorev.2007.11.001

Mewhort, D. J. K., & Johns, E. E. (2005). Sharpening the echo: An interative-resonance model for short-term recognition memory. *Memory, 13*(3-4), 300-307.

Milner, B. 1962. Les troubles de la memoire accompagnant des lesions hippocampiques bilaterales. In P. Passant (Ed.), Physiologie de Hippocampe. Paris: C.N.R.S. Paris, pp. 257-272.

Milner, B., Corkin, S., & Teuber, H. L. (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of H. M. *Neuropsychologia, 6*(3), 215-234.

Morris, C. D., Bransford, J. D., & Franks, J. J. (1977). Levels of processing versus transfer appropriate processing. *Journal of Verbal Learning & Verbal Behavior, 16*(5), 519-533.

Murdock, B. B. (1983). A distributed memory model for serial-order information. *Psychological Review, 90*(4), 609-626.

Murdock, B. B. (1995). Developing TODAM: Three models for serial-order information. *Memory & Cognition, 23*(5), 631-645.

Musen, G., & Treisman, A. (1990). Implicit and explicit memory for visual patterns. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*(1), 127-137. doi:http://dx.doi.org/10.1037/0278-7393.16.1.127

Myung, I. J. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology, 44*, 190-204. doi:10.1006/jmps.1999.1283.

Nairne, J. S. (1990). A feature model of immediate memory. *Memory & Cognition*, *18*(3), 251-269.

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*. New York: Academic Press.

Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: MIT Press.

Nosofsky, R. (1985). Overall similarity and the identification of separable-dimension stimuli: A choice model analysis. *Perception & Psychophysics*,*38*(5), 415-432.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115*(1), 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*(1), 87-108. doi:http://dx.doi.org/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*(4), 700-708. doi:http://dx.doi.org/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*(1), 3-27. doi:http://dx.doi.org/10.1037/0096-1523.17.1.3

Nosofsky, R. M., Little, D. R., & James, T. W. (2012). Activation in the neural network responsible for categorization and recognition reflects parameter changes. *PNAS Proceedings of the National Academy of Sciences of the United States of America, 109*(1), 333-338. doi:http://dx.doi.org/10.1073/pnas.1111304109

Nosofsky, R. M., & Zaki, S. R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. *Psychological Science, 9*(4), 247-255.

Opitz, B. (2010). Neural binding mechanisms in learning and memory. *Neuroscience and Biobehavioral Reviews, 34*(7), 1036-1046. doi:http://dx.doi.org/10.1016/j.neubiorev.2009.11.001

Ostergaard, A. L., & Jernigan, T. L. (1993). Are word priming and explicit memory mediated by different brain structures?. In P. Graf and M. E. J. Masson (Eds.), *Implicit memory: New directions in cognition, development, and neuropsychology* (pp. 327-349). Hillsdale, NJ, England: Lawrence Erlbaum Associations, Inc.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109*(3), 472-491. doi:10.1037//0033-295X.109.3.472

Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks, 6*(3), 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*(1), 28-38. doi:http://dx.doi.org/10.1037/h0024135

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77*(3), 353-363. doi:http://dx.doi.org/10.1037/h0025953

Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology, 83*(2), 304-308. doi:http://dx.doi.org/10.1037/h0028558

Prince, M., Brown, S., & Heathcote, A. (2012). The design and analysis of state-trace experiments. *Psychological Methods*, *17*(1), 78-99.

Ratcliff, R. (1990). Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review, 97*(2), 285-308. doi:http://dx.doi.org/10.1037/0033-295X.97.2.285

Ratcliff, R., & McKoon, G. (1986). More on the distinction between episodic and semantic memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12*(2), 312-313.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning & Verbal Behavior, 6*(6), 855-863.

Reber, P. J. (2002). Attempting to model dissociations of memory. *Trends in Cognitive Sciences, 6*(5), 192-194. doi:http://dx.doi.org/10.1016/S1364-6613(02)01893-4

Reber, P. J. (2013). The neural basis of implicit learning and memory: A review of neuropsychological and neuroimaging research. *Neuropsychologia, 51*(10), 2026-2042. doi:http://dx.doi.org/10.1016/j.neuropsychologia.2013.06.019

Reber, P. J., Gitelman, D. R., Parrish, T. B., & Mesulam, M. M. (2003). Dissociating explicit and implicit category knowledge with fMRI. *Journal of Cognitive Neuroscience, 15*(4), 574-583. doi:http://dx.doi.org/10.1162/089892903321662958

Reber, P. J., Stark, C. E. L., & Squire, L. R. (1998). Contrasting cortical activity associated with category memory and recognition memory. *Learning & Memory, 5*(6), 420-428.

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*(3), 411-419. doi:http://dx.doi.org/10.1037/0735-7044.113.3.411

Roediger, H. L., Srinivas, K., & Weldon, M. S. (1989). Dissociations between implicit measures of retention. *Implicit memory: Theoretical issues.* (pp. 67-84) Lawrence Erlbaum Associates, Inc, Hillsdale, NJ.

Scarborough, D. L., Cortese, C., & Scarborough, H. S. (1977). Frequency and repetition effects in lexical memory. *Journal of Experimental Psychology: Human Perception and Performance, 3*(1), 1-17. doi:http://dx.doi.org/10.1037/0096-1523.3.1.1

Schacter, D. L., Cooper, L. A., & Delaney, S. M. (1990). Implicit memory for unfamiliar objects depends on access to structural descriptions. *Journal of Experimental Psychology: General*, *119*(1), 5.

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

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery & Psychiatry, 20*, 11-21. doi:http://dx.doi.org/10.1136/jnnp.20.1.11

Seger, C. A. (1994). Implicit learning. *Psychological Bulletin, 115*(2), 163-196. doi:http://dx.doi.org/10.1037/0033-2909.115.2.163

Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences, 17*(3), 367-395.

Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relting generalization to distance in psychological space. *Psychometrika, 22*(4), 325-345.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*,*4*(2), 145-166.

Shimp, C. P. (1966). Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior, 9*(4), 443-455.

Shin, H. J., & Nosofsky, R. M. (1992). Similarity-scaling studies of dot-pattern classification and recognition. *Journal of Experimental Psychology: General, 121*(3), 278-304. doi:http://dx.doi.org/10.1037/0096-3445.121.3.278

Shohamy, D., Myers, C. E., Kalanithi, J., & Gluck, M. A. (2008). Basal ganglia and dopamine contributions to probabilistic category learning. *Neuroscience and Biobehavioral Reviews, 32*(2), 219-236. doi:http://dx.doi.org/10.1016/j.neubiorev.2007.07.008

Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*(6), 1411-1436. doi:http://dx.doi.org/10.1037/0278-7393.24.6.1411

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.

Spencer, W. D., & Raz, N. (1995). Differential effects of aging on memory for content and context: A meta-analysis. *Psychology and Aging, 10*(4), 527-539. doi:http://dx.doi.org/10.1037/0882-7974.10.4.527

Squire, L. R. (1992). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. *Journal of Cognitive Neuroscience, 4*(3), 232-243.

Reber, P. J., & Squire, L. R. (1994). Parallel brain systems for learning with and without awareness. *Learning & Memory*, *1*(4), 217-229.

Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory, 82*(3), 171-177. doi:http://dx.doi.org/10.1016/j.nlm.2004.06.005

Squire, L. R., Knowlton, B., & Musen, G. (1993). The structure and organization of memory. *Annual Review of Psychology,44*, 453-495.

Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science, 253*(5026), 1380-1386.

Strange, W., Keeney, T., Kessel, F. S., & Jenkins, J. J. (1970). Abstraction over time of prototypes from distortions of random dot patterns: A replication. *Journal of Experimental Psychology, 83*(3), 508-510. doi:http://dx.doi.org/10.1037/h0028846

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

Tulving, E. (1983). Elements of episodic memory. Oxford: Clarendon Press.

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

Tversky, A. (1977). Features of similarity. *Psychological Review, 84*(4), 7-45.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological review*, *108*(3), 550-592.

Van Orden, G. C., & Kloos, H. (2003). The module mistake. *Cortex*, *39*(1), 164-166.

Van Orden, G. C., Pennington, B. F., & Stone, G. O. (2001). What do double dissociations prove?. *Cognitive Science*, *25*(1), 111-172.

Vokey, J. R., & Brooks, L. R. (1992). Salience of item knowledge in learning artificial grammars. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*(2), 328-344. doi:http://dx.doi.org/10.1037/0278-7393.18.2.328

Warrington, E. K., & Weiskrantz, L. (1974). The effect of prior learning on subsequent retention in amnestic patients. *Neuropsychologia, 12*(4), 419-428.

Warrington, E. K., & Weiskrantz, L. (1978). Further analysis of the prior learning effect in amnesic patients. *Neuropsychologia, 16*(2), 169-177.

Weiskrantz, L., & Warrington, E. K. (1970). Verbal learning and retention by amnesic patients using partial information. *Psychonomic Science, 20*(4), 210-211.

Whittlesea, B. W. (1997). Production, evaluation, and preservation of experiences: Constructive processing in remembering and performance tasks. In D. L. Medin (Ed.), *The psychology of learning and motivation: Advances in research in* theory, pp. 211-264, San Diego, CA: Academic Press.

Woods, R. T., & Percy, M. (1974). A similiarity between amnesic memory and normal forgetting. *Neuropsychologia, 12*(4), 437-445.

Wright, R. L., & Whittlesea, B. W. A. (1998). Implicit learning of complex structures: Active adaptation and selective processing in acquisition and application. *Memory & Cognition, 26*(2), 402-420.

Zaki, S. R. (2004). Is categorization performance really intact in amnesia? A meta-analysis. *Psychonomic Bulletin & Review,11*(6), 1048-1054. doi:http://dx.doi.org/10.3758/BF03196735

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*(4), 344-359.

Zaki, S. R., Nosofsky, R. M., Jessup, N. M., & Unversagt, F. W. (2003). Categorization and recognition performance of a memory-impaired group: Evidence for single-system models. *Journal of the International Neuropsychological Society, 9*(3), 394-406. doi:http://dx.doi.org/10.1017/S1355617703930050

Zannino, G. D., Perri, R., Zabberoni, S., Caltagirone, C., Marra, C., & Carlesimo, G. A. (2012). Performing prototype distortion tasks requires no contribution from the explicit memory systems: Evidence from amnesic MCI patients in a new experimental paradigm. *Neuropsychologia, 50*(12), 2907-2915. doi:http://dx.doi.org/10.1016/j.neuropsychologia.2012.08.024

Zucchini, W. (2000). An introduction to model selection. *Journal of Mathematical Psychology, 44*, 41-61. doi:10.1106/jmps.1999.1276