University of Alberta

The Referent Model Of Lexical Decision

by

Clarence Darren Piercey ©

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of *Doctor of Philosophy*

Department of Psychology

Edmonton, Alberta
Fall 2002

0-612-81253-7

University of Alberta

Library Release Form

Name of Author: Clarence Darren Piercey

Title of Thesis: The Referent Model Of Lexical Decision

Degree: Doctor of Philosophy

Year this Degree Granted: 2002

Darren Piercey

C. Darren Piercey
18 Cortland Street
New Maryland, N.B.
Canada
E3C 1B5

June 14/02

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled The Referent Model Of Lexical Decision submitted by Clarence Darren Piercey in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Michael R. W. Dawson
Supervisor

Connie K. Varnhagen

Chris F. Westbury

Terrance M. Nearey

Michael E. J. Masson

June 12/02

Abstract

The lexical decision task is used as a tool that provides us with an understanding of how memory is structured and how words are stored in and retrieved from memory. Experiments that have used the lexical decision task produce complicated results. By manipulating either the type of words or nonwords that are being presented, we can increase the size of an effect or remove the effect all together. The purpose of this dissertation is to introduce a new model of lexical decision known as the referent model. This model will provide us with a detailed description of the decision process that is used when performing this task. It is hoped that by having a good understanding of this decision process, we will gain a better understanding of the complicated results that are produced. Support for the referent model is provided using computer simulations as well as human data.

Acknowledgements

I would like to thank my supervisor Mike Dawson for his support during my time in graduate school. I quickly learned to appreciate his style of supervision. I was given freedom to pursue any area of research that interested me. Mike wasn't the kind of person who was always looking over my shoulder to make sure work was getting done. He believed that I was capable enough to complete the task at hand. However, he was always there to provide a gentle push when the time was appropriate. My fondest memory of being in his lab is coffee time. Mike taught me the evils of decaffeinated coffee and the joys of a good espresso buzz. My favorite quote, that I am sure I will often repeat to my students, is "R.T.F.M.", which stands for "read the fuckin' manual"! Mike was my graduate supervisor but has become a good friend. I enjoy chatting with him about life, the universe, and everything. I hope to stay in touch with Mike for years to come.

I would also like to thank the members of the BCP. I really enjoyed getting to know everyone and the special bond that I formed with Leanne Willson and Patricia (PAT) Boechler. We went through the tough times together, when we wondered if we would ever finish school, and now we are going through the good times, when we are all finishing and moving on to academic position. Life is pretty good guys! I wish you the best of luck!

A special thanks goes to Edna and Clarence, my mom and dad. I thank you for the support you have given me, especially during the early years of university. I thank you for the morals and ethics that you have instilled in me. I thank you for making me the man that I am today.

Finally I would like to thank my beautiful wife Karen. I am so lucky to have met someone as wonderful as you are. It's been a stressful time in both our lives. However, you have been the strong one who has been there to support me. You always believed in me, and loved me. I thank you for being my wife and my best friend. I also thank you for our two beautiful children Emalee and Owen. The three of you have taught me that the most important thing in life is family.

TABLE OF CONTENTS

Introduction 1

Chapter 1 6

Chapter2 13

Theoretical Frameworks 14
Multiple Read-Out Model 14
Philosophy, Structure, and Processing 14
The Referent Model of Lexical Decision 17
Description 17
Distributed versus Localist Representations 17
Retrieval of Distributed Representations 19
Measuring the Fit with Items in Memory 20
From Harmony to a Decision 22
Efficiency and Bias 26
Discriminating Between the Models 29

Chapter 3 31

Structure of a Hopfield Network 31

A Hopfield Network For The Referent Model 35
Network Structure 40
Architecture 40
Parameters for lexical decision 40

Network Simulations 41
Simulation 1: The Relationship Between Nonword Context and Word Responses 41
Method 42
Results and Discussion 44
Simulation 2: The Relationship Between Word Frequency and Nonword Responses 46
Method 47
Results and Discussion 47
Simulation 3A and 3B: The Relationship Between Word and Nonword Ratio 48
Method 50
Results and Discussion 51
Simulation 3A 51
Simulation 3B 52
Summary 53

Chapter 4 55

Experiments With Human Subjects 56
Experiment 1: The Effect of Nonword Context on Word Responses ... 56
Method 57
Results and Discussion 60
Experiment 2: The Effect of Word Context on Nonword Responses ... 62
Method 63
Results and Discussion 64
Experiments 3A and 3B: Varying the Word to Nonword Ratio 65
Method 67
Results and Discussion 68
Experiment 3A 68
Experiment 3B 70
Summary 71

Chapter 5 73

Combining Simulations and Experiments 73
Simulation 1 and Experiment 1 73
Simulation 2 and Experiment 2 74
Multiple Read-Out Model 74
The Referent Model 75
Conclusions 76
Simulation 3 and Experiment 3 76

General Discussion 78
Re-Uniting Decision Predictions with Attractor Network Predictions .. 79

Extensions to Other Tasks 82

Future Research 83

Final Note 86

References 87

INTRODUCTION

For hundreds of years philosophers, psychologists, and computer scientists have struggled with the problem of understand how humans think and reason. For example, psychologists perform experiments on humans and animals to gain knowledge of the procedures that occur when information is processed. These experiments vary from single cell recordings in monkey visual cortex to gaining insight into human group dynamics. The range of experiments that are being performed is quite large. So large in fact that it would be impossible for a single psychologist to study all of the different areas of psychology. Therefore psychologists must specialize. They must study simple processes and try to gain a strong understanding of how these processes occur. Then, as a community, we can begin to gain a better understanding of how humans process information.

The area of psychology that this dissertation focuses on is human memory. What is the structure of memory? How are words stored in and retrieved from memory? To gain a better understanding of these processes we need to develop tools that permit us to infer the properties of a system that we cannot observe directly. We need to be able to systematically manipulate the stimuli that human's process and determine how these stimuli affect performance. Then, based on these findings, we will be able to produce theories of how humans process information.

Meyer and Schvaneveldt (1971) provided us with one of these tools when they developed the first lexical decision task. The lexical decision task involves presenting subjects with strings of letters that they must categorize as being either words or nonwords. There are several types of manipulations that can be performed on either the words or nonwords. For example, word frequency, concreteness, orthographic neighborhood size, semantic neighborhood size, etc., can all be manipulated and each of these manipulations can produce changes in both reaction time and accuracy. Also, the type of nonword that is presented can be manipulated. For example, as the nonword items vary from scrambled nonwords (e.g. FRGKL), to pronounceable nonwords (e.g. SHROG), and to pseudohomophones (e.g. BRANE), they become more wordlike. These changes also affect reaction times and accuracy.

The original experiments performed by Meyer and Schvaneveldt (1971) involved presenting subjects two items at the same time, one above the other. Each trial involved presenting either two words, a word and a nonword, a nonword and a word, or two nonwords. The subjects' task was to respond yes if both items were words and no if one or both items were nonwords. The stimuli that were used by Meyer and Schvaneveldt consisted of associated word pairs, unassociated word pairs, and word/nonword pairs. The purpose of their experiment was to determine if word association would affect reaction time and accuracy. They found that subjects could make faster and more accurate lexical decisions when the words were related than when they were not.

Meyer and Schvaneveldt (1971) accounted for their findings using a localist description of memory. They suggested that there is a specific location for words in long term memory. Items that are semantically related are located more closely together than are items that are not related. If the subject is presented with two word items that are related, then a word decision is made to one item at a time. As the first item is processed, activation from the location of this item spreads to items that are related. Items that are more closely related are located more closely together and receive more activation. After the subject makes the first word decision, the second item is processed. The level of activation for this item has already risen slightly because of the spreading activation. Therefore they are able to make the second word decision more quickly. This produces the related word advantage that is found.

Meyer and Schvaneveldt (1971) concluded that subjects are able to make lexical decisions without accessing semantics. Decisions are made based on information that is just sufficient to determine if the item is a word or not. The information that is available at this point in time is not sufficient to include aspects of meaning. It is assumed that subjects are able to make decisions based on processing that occurs before a complete semantic meaning is retrieved from memory.

An interesting thing to point out here is that Meyer and Schvaneveldt (1971) argued that decisions are based on this early processing without actually describing the mechanism that is used to make these decisions. However, their assumption may actually be correct. Models presented in Chapter 2 will address this issue. According to these models, decisions are made based on early processing. However, some semantic processing must occur before a lexical decision is made. There is evidence for this from experiments that manipulate how semantically related the word items are. These manipulations affect both reaction time and accuracy, so some semantic processing must occur before a lexical decision is made. Therefore Meyer and Schvaneveldt are partially correct; decisions are made before a meaning for the item is retrieved but some semantic processing does occur early on as well.

The lexical decision task is used as a tool for understanding how memory is structured. The questions that are addressed in this dissertation are specifically related to the lexical decision task. How is it that subjects are able to distinguish between words and nonwords? How do they perform a lexical decision? What are the mechanisms that are involved when a lexical decision is made? It is obvious that the results from a lexical decision task can be quite complicated. There are many factors that can affect results and possible interactions between these factors. To be able to understand these complicated effects we must first have a complete understanding of the task itself. There is an intimate relationship between words and nonwords used in a lexical decision experiment. Manipulations to either type of stimuli affect reaction times and accuracy of the other. Any theory of how a lexical decision is performed must be able to account for this intimate relationship.

The purpose of this dissertation is to show how dynamic and complicated the lexical decision task actually is and to present a relatively simple model of lexical decision that is able to account for the complicated interactions and the relationship between word and nonword decisions. In Chapter 1 an overview of some recent lexical decision experiments will be presented. The findings from these experiments will provide a background of evidence for assumptions that will be made later on. In Chapter 2, two models of lexical decision will be described; a localist model, the multiple readout model, and a distributed model, the referent model. Obvious differences and similarities will be pointed out and a discussion of how the two models can be compared will be given. In Chapter 3 simulations of the referent model will be performed. These simulations will give us a more concrete understanding of the referent model. They will force us to provide explicit descriptions of the mechanisms involved when making a lexical decision. In Chapter 4, data from experiments using human subjects will be presented. These experiments will replicate the findings from the simulations in Chapter 3 and provide further evidence in favour of the referent model. Then, in Chapter 5, we will compre and contrast both the multiple read-out model and the referent model. We will then be able to determine which of the two models is superior. Finally, at the end of Chapter 5, a summary and a brief conclusion will be provided.

CHAPTER 1

How are we to understand the structure of human memory? The traditional approach used by many cognitive psychologists is to identify tasks that require us to retrieve information from memory as quickly as possible. Then, by comparing response times across different contexts or stimulus classes, researchers determine which conditions cause information to be retrieved more efficiently, and this information could provide clues about the structure of a system that would show such retrieval preferences (e.g. Collins & Qullian, 1969; Coltheart, Davelaar, Jonasson, & Besner, 1977; etc.).

The problem with this logic comes in finding a task that accurately reflects the time it takes to retrieve information from memory. There are likely at least two components to performance on any task, the memory retrieval processes, and some response-decision processes that controls the emission of a response after interpreting the results of the ongoing retrieval processes (Dawson, 1988). Is there any task one could create that minimizes the second of these components such that performance on the task truly reflects the time it takes to completely retrieve some item from memory, with little to no contribution of a response component? Identifying such a task has been more problematic than one might expect.

The lexical-decision task provides a prime example of the problem of finding a task that provides a good index of memory retrieval. In its original instantiation (i.e., Meyer & Schvaneveldt, 1971) the lexical-decision task involved the presentation of two strings of letters. Participants were instructed to decide whether or not both strings represented correctly spelled English words. More recently the task has been altered slightly in that the letter strings are typically presented one at a time, with participants instructed to categorize each stimulus as either a word, or a nonword (e.g., Joordens & Becker, 1997). In either form, this would seem to be a reasonable task for studying memory. That is, in order to decide whether or not a letter string is an English word, the most obvious strategy would be to compare the string to items stored in memory. If a match is found, a subject makes a "word" response. If no match is found a "nonword" response is produced. Thus, the time it takes to correctly categorize an item as a word should provide a good indication of the time it takes to retrieve the item from memory that matches the current letter string.

It is important to emphasize the serial processing nature of this original explanation of the processes underlying performance on the lexical-decision task. Specifically, the assumption once again is that there are two distinct components to the task. The first is a memory search stage in which it is assumed the information is retrieved from memory and matched to the current letter string. Then, based on the results of this first process, a second process outputs a word or nonword response. The response then, is presumed to wait until the search, and retrieval (in the case of a word), is complete. If one further assumes that the second component of performance on this task merely adds some constant to the eventual response time, then differences in response times across various contexts would reflect the affect those contexts have on the efficiency of memory retrieval.

Alternative explanations of performance on the lexical-decision task have been proposed (see Balota & Chumbley, 1984) and, in light of a number of studies, it appears that the serial view proposed above is simply inaccurate. The studies that best illustrate the problems with the serial view are those that have manipulated the character of the nonword foils. James (1975) performed the first such study, just four years after Meyer and Schvaneveldt introduced the lexical-decision task. The words in James' study consisted of words that were either high imageability (e.g., LAMP) or low imageability (e.g., LOVE). This distinction is assumed to be reflected somewhere at the semantic level of the words' representations. James examined the time it took participants to correctly categorize these words in two contexts; one where the nonwords were so-called illegal nonwords (e.g., FMKLR), and one where they were pronounceable nonwords (e.g., FRILK).

Two results emerged from the James (1975) study. First, it took participants longer to correctly categorize the words when the nonwords were the pronounceable nonwords than when they were the illegal nonwords. This result clearly shows that a correct word decision is not based on a simple search of memory because, if it was, then the character of the nonword foils should have no affect on word decisions. Rather, it seems that lexical decisions involve some sort of discrimination process whereby the more wordlike the nonwords are, the longer it takes to discriminate words from them.

The second of James' results suggests that during the time participants are trying to discriminate the words from the nonwords, they continue to actively process the words. Specifically, participants typically make correct word decisions faster to high than to low imageability words, and this result was apparent in both contexts that James examined. However, the magnitude of the imageability effect was significantly larger in the pronounceable nonword condition. Given that imageability is assumed to be primarily a semantic-level variable, this suggests that the words were processed more "deeply" (see Craik & Lockhart, 1972) prior to a response in the pronounceable nonword context.

Thus, on the basis of James' (1975) study, it appears that the extent to which a representation associated with a word is retrieved when participants are performing lexical decisions depends critically on the nature of the nonword foils. More wordlike nonwords force the retrieval of a more complete representation prior to the making of a word response.

There are now a large number of studies that support this general view. As examples, Stone and Van Orden (1993) showed that the time to make correct word decisions, and the magnitude of the frequency effect apparent in those decisions, progressively increased as the nonwords were varied from illegal nonwords (e.g., FRNHT) to pronounceable nonwords (e.g., FRANE) to pseudohomophones (e.g., BRANE)[1]. Piercey & Joordens (2000) showed the word response times and semantic ambiguity effects[2] both increased as nonwords varied from scrambled nonwords (e.g., FRNEA) through pronounceable nonwords through pseudohomophones (see also Borowsky & Masson, 1996). Finally, Joordens and Becker (1997) showed that

[1] Pseudohomophones are considered even more wordlike than pronounceable nonwords because, in addition to having a wordlike spelling pattern, they also possess a phonology that matches a word if they were pronounced aloud (i.e., BRANE -> BRAIN).
[2] Semantic ambiguity refers to the distinction between words that tend to have multiple meanings (e.g., BAT) versus those that tend to have one meaning (e.g., BET). The semantic ambiguity effect refers to the finding that lexical decisions are typically faster and more accurate to ambiguous than to unambiguous words (see Joordens & Besner, 1994).

semantic-priming effects[3] become larger and last over longer lags as the nonword foils are made more wordlike.

Taken together, these findings suggest the following general view of how lexical decisions are made: When the stimulus is presented, the participant begins to retrieve information from memory that is relevant to that stimulus. That is, the participants begin to process the stimulus. Simultaneously, some decision mechanism is monitoring the retrieval process and, at some point, arrives at either a word or nonword decision. If the nonwords in a study are very similar to the words, then the decision mechanism takes longer to arrive at a categorization of the stimulus. During this additional time, the processing of the stimulus continues. As a result, more of that item's representation is retrieved from memory. Thus, the stimulus is processed to a deeper level because of the difficulty in discriminating words from nonwords.

This new view of the processes underlying lexical decisions is clearly more complicated than the serial "retrieve then decide" view. Specifically, it is what has been termed a cascade model, information cascading from the memory retrieval process to the decision process (e.g. Ratcliff, 1978). The difficulty with this is that it opposes the basic assumption that response times on the lexical decision task provide a clear index of the time it takes to retrieve a representation from memory. As this is the first step in the logic underlying the use of the lexical decision task to understand the structure of memory, that whole enterprise is called into question by the dynamic nature of the task.

---

[3] Semantic priming refers to the finding that lexical decisions to some target word (e.g., WOLF) can be made faster and more accurately when that target follows a related prime (e.g., DOG) as opposed to an unrelated prime (e.g., HAT). When we speak of the semantic priming effect occurring over longer lags, lag refers to the number of unrelated items presented between the prime and target stimuli.

This leaves the cognitive psychologist interested in the structure of memory with two options. The first option is to avoid using the lexical-decision task and, instead, to seek out some other task that may provide a better index of memory retrieval. For example, Neely, Keefe, and Ross (1989) came to just this conclusion, stating "...we agree with others that the lexical decision task is not well suited for the study of lexical access". (p. 1017). The problem with simply discarding the task in this manner is that it may be difficult to develop any task that provides a clean index of memory retrieval.

As one example, another task assumed to measure lexical access is the naming task. Participants performing the naming task simply read words aloud as quickly as possible when they are presented. One would think that such a task should provide a clean measure of the time it takes to retrieve a phonological representation given some word. However, Lupker, Brown, and Columbo (1997) have shown that naming performance is affected by the other stimuli that make up the list. The very same words are named faster when the other words in the list are "easy" (e.g., mono-syllabic) words than when they are "difficult" (e.g., tri-syllabic) words. When a task as straightforward as the naming task shows these kinds of list effects, it calls into question whether any task provides a clear measure of memory access.

If no task provides a clean measure of memory access, how is an investigator interested in memory to proceed? The only option left is the try our best to gain a good understanding of the exact manner in which the retrieval and decision processes interact in some given task context. That is, through a detailed examination of the task, we may come to understand the manner in which it reflects memory retrieval.

This may lead us to identify conditions in which lexical decision does provide a good measure of memory retrieval. In addition, as we strive to understand the task, we may come to a better understanding of the dynamic way that memory retrieval is linked to performance. This understanding will likely have implications that go beyond examinations of the structure of human memory. Said another way, the problem of how to measure memory may turn out to be an opportunity to better understand the complicated ways in which we are able to use memory to guide our performance.

The remainder of this dissertation will focus on the lexical-decision task, and will examine one issue relevant to lexical decisions in detail; what is the nature of the decision mechanism that underlies lexical decisions, and how might it be linked to the more general memory-retrieval process? In Chapter 2 we will examine a current model of lexical-decision, Grainger & Jacob's (1996) multiple read-out model, and contrast it with a new referent model of lexical decisions.

CHAPTER 2

As was mentioned in Chapter 1, there are two general accounts of how memory is structured and how remembered items are retrieved. According to the localist approach, each memorized item has a specific location in memory and associations between items are made with respect to their relative positions. According to the distributed view, items are not stored in one specific place. Rather, information about each item is distributed across many areas of memory. When you compare retrieval of two items that are related, similar areas of memory will be active. According to the localist model, items are retrieved from memory when the activation in the location of a specific item increases above a threshold. According to the distributed model, increases in activation are also used to retrieve items from memory, but this activation is spread across different areas of memory. The combined information from all of these areas provides us with a concept or a memory.

The purpose of this chapter is to provide descriptions of two models of lexical decision that are based on these two theories of how memory is structured. The two models will then be compared and an argument in support of the distributed model will be made. The first is known as the multiple read-out model. It is based on a localist description of memory. Information about a single word is located in a specific place in memory. As we retrieve a word from memory, activation in this specific location increases. The second is known as the referent model. This is a new model of lexical decision that is based on a distributed description of memory. Information about a single word is distributed across many different areas of memory.

Therefore, when a word is retrieved from memory, activation increases in many areas.

Theoretical Frameworks

Two theoretical views of lexical decision will be discussed. We will first discuss the multiple read-out model (Grainger & Jacobs, 1996), which uses separate indices to make word versus nonword decisions. We will then discuss an alternate view of lexical decision, termed the referent model. This model is based on the synthesis of an attractor model of memory retrieval with a random-walk decision process. The random-walk decision mechanism differs from that of the multiple read-out model in that word and nonword responses are based on the same decision process.

The Multiple Read-Out Model.

Philosophy, Structure, and Processing

The multiple read-out model (Grainger & Jacobs, 1996) is the most complete and explicit of current theories of lexical decision. It embodies a global philosophy that is followed in the present manuscript as well. Specifically, a wide variety of tasks are used to gain a better understanding of the processes and representations underlying word-recognition. These include such tasks as naming, perceptual identification, and, the focus of the current manuscript, lexical decision. If any specific pair of these tasks is considered, they likely draw on some common processes, as well as on some processes or mechanisms that are unique to each task. That is, there is some overlap, but the overlap is not total. Thus, the best way to understand word-recognition performance in general is to first develop an explicit theoretical description of performance on one task, and then to describe how performance on other tasks can be mapped onto that framework. With this philosophy in mind, Grainger and Jacobs (1996) present the multiple read-out model and describe how it can be related to lexical decision and perceptual identification tasks, stressing the processes that are common to both tasks, as well as those that are unique.

With respect to lexical decisions, the multiple read-out model begins by assuming that information is represented locally in memory. That is, it is assumed that concepts are represented by single nodes organized into a semantic network (e.g., Collins & Qullian, 1969). When an item (e.g., TRUCK), is presented in the context of a lexical decision experiment, activation is assumed to grow where the lexical representation for that stimulus is located. Also, activation is assumed to grow where related lexical representations are located (e.g., orthographically related, like TRUCE). This growth of activation would be faster, and reach a higher level, for words than it would for nonwords. Words have an actual node associated with them in memory, whereas nonwords would only have associated nodes (i.e. lexical representations that are related to the nonwords at the level of orthography or phonology).

Correct word decisions are mapped onto this process via two thresholds, one sensitive to local levels of activation, and one sensitive to global levels of activation. If either of these thresholds are surpassed prior to a time deadline, a "word" response is emitted. Thus, in a general sense, word decisions are based on monitoring the levels of activation within lexical representations, and emitting a word decision when the activation level is sufficiently high.

In contrast, nonword decisions are based on a separate criterion that involves a time deadline. Essentially, if the activation is not high enough to warrant a word decision by some time deadline, a "nonword" response is emitted. This time deadline is assumed to be noisy (i.e., the exact deadline is assumed to vary from trial to trial), and is also assumed to be sensitive to the rate at which activation is growing. If an item's rate of activation growth is high, then this suggests that the item may turn out to be a word. In recognition of that possibility, the time deadline is moved further back in time (i.e., producing a longer deadline) in order to give the item every chance to hit one of the word activation thresholds.

It is important to highlight why the model needs to incorporate these complexities rather than simply assuming a constant time deadline. One reason for this is the existence of a phenomenon termed the pseudohomophone effect (Besner & Davelaar, 1983; Coltheart, Davellar, Jonasson & Besner, 1977; McCann, Besner & Davelaar, 1988; see Dennis, Besner & Davelaar, 1985 for a review). The pseudohomophone effect refers to the finding that correct nonword decisions in a lexical decision experiment are made faster to pronounceable nonwords (e.g., FALID) than to pseudohomophones (e.g., SALID). The more wordlike a nonword is, the slower a participant is to call it a nonword. If a constant deadline was used such that nonword decisions were emitted if the activation criterion had not been reached after some criterion amount of time, then there would be no reason to expect differences in nonword response times across different kinds of nonwords. However, it seems reasonable to assume that the more wordlike a nonword is, the faster its activation would grow. By pushing the deadline back for items with high rates of activation growth, the model is able to account for the slower responses to more wordlike nonwords.

The multiple read-out model uses more than one process to be able to make word and nonword decisions. The distributed model, known as the referent model, that is presented next uses one process to make word and nonword decisions. This description of how lexical decisions are made is simpler but is in fact more accurate.

The Referent Model of Lexical Decision

Description

The referent model provides a different way of thinking about lexical decision performance. Its basic assumptions are the following: First, concepts are represented in memory as distributed patterns of activation. Second, memory retrieval reflects a pattern completion process like that embodied in attractor models of memory. Third, during the pattern completion process, a measure of the global match of the current pattern with patterns in memory (i.e., harmony) is passed to a decision mechanism. Fourth, the decision mechanism uses the harmony to drive a random-walk process that eventually leads to either a word or nonword decision. Finally, a bias often enters into the random-walk process in order to increase efficiency of responding, and it is this bias that is responsible for the dynamic mapping of lexical decision responses to memory retrieval. The specifics of each assumption are now described in turn.

Distributed versus Localist Representations

It is important to point out a major difference between localist and distributed models. In a localist model, concepts are represented by one network node. However, in an attractor network, concepts are represented simultaneously in many nodes. The multiple read-out model assumes that individual nodes within a memory network represent concepts, what is sometimes termed localist representations. A popular alternate view is that concepts are represented as patterns of activation across a large set or nodes (Rumelhart, McClelland, and the PDP research group, 1986; Dawson & Piercey, 2001). The basic notion here is akin to that of a picture on a television or computer screen. If a picture of a tree is displayed on a computer screen, no single pixel of that picture contains the information that tells the viewer they are looking at a tree. Rather, the "treeness" of the image is depicted by the collection of pixels, and the RGB values each pixel is displaying. Thus, the representation of the tree is distributed across the screen as a collection of pixels "firing" in a certain manner.

There are a number of advantages associated with distributed versus localist representations, including such things as graceful degradation. That is, if the tree is again displayed but 25% of the pixels are turned off, you would likely still see a tree. A speckled tree perhaps, but a tree nonetheless. However, if a single pixel were used to represent a tree, then if that pixel suddenly stopped working, the concept of tree would be lost.

There is evidence that our brains also experience graceful degradation. If we are presented with stimuli that are somewhat degraded, we are still able to perceive them and form a complete representation. Also, if a specific area of the brain is damaged, we are still able to function and recognize objects. This leads us to believe that human memory represents information in a distributed manner.

However, the characteristic of distributed representations that will be emphasized is the fact that when memory is envisioned in this manner, it allows for different ways to think of how memory retrieval might operate. In fact, one way to retrieve a distributed representation is through a process called pattern completion. In turns out this process provides a very interesting way of thinking about memory retrieval, as described in the next section.

Retrieval of Distributed Representations

A number of authors have recently described distributed models of memory that can be generally characterized as attractor networks (e.g., Becker, Moscovitch, Behrmann, & Joordens, 1997; Hinton & Shallice, 1991; Masson, 1991; 1995; Plaut & Shallice, 1991; Sharkey & Sharkey, 1992). These models assume that memory retrieval involves a pattern completion process whereby, after learning some patterns, a model is able to re-instantiate a given learned pattern when given only part of it. Thus, at the beginning of retrieval, some of the pattern represented across the nodes of the network corresponds to a learned pattern, and the rest of the pattern is set randomly. The model is then able to 'clean-up' the random part of the pattern, changing it to a pattern consistent with the desired learned pattern.

A detailed description of one type of attractor network, known as a Hopfield Network, will be provided in Chapter 3. Simply described, such models are termed attractor networks because the learned patterns act as attractors during retrieval. Specifically, when the patterns are presented during learning, they form a weight matrix that encodes the correlation between each pair of nodes across all of the learned concepts. Thus, if nodes 1 and 2 tend to be in different states across the leaned patterns, the weight between those nodes will be negative with the strength of the weight reflecting the proportion of patterns in which the nodes are in different states. In contrast, nodes 2 and 3 might have a strong positive weight if they tended to be in the same state across the learned concepts.

The retrieval process involves visiting each node and deciding whether to change its state or to leave it as is. Critically, the decision is based on whether changing the state of the node will make the new pattern 'fit' better with memory. This fit can be explicitly quantified in terms of the learned correlations encoded in the weight matrix. For example, suppose that we are deciding whether to change the state of node 2. We know that node 2 is negatively correlated with node 1. If node 1 is currently on, then node 2 would fit better with the current state of node 1 if it was off. However, when considering whether to turn node 2 on or off, the impact of the change is considered with respect to the current state of all the nodes in the network, not just node 1. Thus, the general rule is that a node state will change if and only if the change results in a better overall fit with the learned weight matrix and the current state of all other nodes in the network. As a result, the fit of a pattern with memory is constantly increasing as a pattern is being retrieved.

Thus, as a pattern is being retrieved, it keeps changing in a way that increases the fit with the weight matrix. Given that the learned patterns are what formed the weight matrix originally, it should not be surprising that those patterns fit very well with it. This is what makes an attractor network work. As long as some of the nodes are in a state associated with a learned concept, the pattern across the nodes will move towards that concept because it is consistent with the learned weight matrix and with some of the current node states. Thus, the pattern being retrieved tends to move towards the learned patterns, as if it were being attracted to them. Hence, the learned patterns are described as attractors, and the networks that support them as attractor networks.

Measuring the Fit with Items in Memory

Thus far, the notion of fit has been described very generally. In fact, there is an explicit way to measure the fit of any given pattern as the sum, across all pairs of nodes, of the product of the two node states and the weight between them. This sum is termed harmony (Smolensky, 1986). Mathematically:

$$Harmony = \sum_{i=1}^{n} \sum_{j=1}^{i} n_i n_j w_{ij}$$

where W_{ij} is the weight between Unit i and Unit j, n_i is the state (i.e., 1 or –1) of Unit i, and n_j is the state (i.e., 1 or –1) of Unit j. Thus, harmony will increase if either (a) two nodes are in the same state and the weight between them is positive, or (b) they are in different states and the weight between them is negative. Harmony will decrease if either (a) the two nodes are in different states and the weight between them is positive, or (b) the two nodes are in the same state and the weight between them is negative.

Given all of the above, it is possible to view word recognition as a process of memory retrieval that follows the processing dynamics of an attractor network. Specifically, it is assumed that words are represented in memory across a number of levels including an orthographic level, a phonological level, and a semantic level. Concepts are represented within each level as a distributed pattern of activation across the nodes at that level. During learning, the model is presented complete patterns associated with some concept and encodes the learned information in a weight matrix that stores the correlation between node states both within and between the various levels of representation. During retrieval, the model is then presented with an intact (or forming) orthographic pattern and random patterns across the other levels. Based on the orthographic pattern, and the weight matrix, it is then able to complete the pattern across the remaining levels of representation.

This pattern completion process is seen as the basic process whereby phonological and semantic information is retrieved following the presentation of some orthographic stimulus. In the spirit of Grainger and Jacob (1996), it would be this basic process that can be seen as common to performance on all word recognition tasks. However, as illustrated by Joordens and Besner (1994), and the subsequent commentaries and counter-commentaries on that paper (Besner & Joordens, 1995; Masson & Borowsky, 1995; Rueckl, 1995), it is highly unlikely that this process alone can provide a good account of word recognition performance in general. Rather, additional assumptions must be overlaid onto the basic attractor-network dynamics to provide more precise accounts of performance on specific word-recognition tasks (e.g., Masson, 1995).

From Harmony to a Decision

For current purposes then, the important issue is how this basic retrieval process can be mapped onto lexical-decision performance. The basic assumption is this: As some pattern is being retrieved from memory, the harmony value reflecting the fit of the current pattern with patterns in memory is constantly changing. This value is continually passed to a decision module. Harmony thus serves as the engine that drives a random-walk process, and it is this random-walk process that eventually arrives at a word or nonword decision.

Before describing the specifics of this approach, it is important to first point out that even when an unlearned orthographic pattern is provided to an attractor network (i.e., a nonword), it will still change the semantic and phonological patterns in a manner that increases the overall harmony of the network. Thus, as depicted in Figure 1, harmony increases when either words or nonwords are presented. The difference is that words gain harmony at a faster rate and eventually reach a higher level of harmony prior to reaching asymptote.

Insert Figure 1 about here

How then is harmony mapped onto a decision? The first assumption of our approach is that participants compute a referent function of harmony that reflects the average harmony value of experienced stimuli at various points during processing. The second assumption is that subjects make their word and nonword decisions on discreet trials by comparing the harmony value of the current stimulus at the current point in time with the referent function. If the current harmony value is some criterion amount above the referent, a word decision is reached. If it is some criterion amount below the referent level, a nonword decision is reached. Each of these assumptions will now be described in detail.

The referent function is assumed to be a function that represents the increase in harmony expected for the "average" stimulus experienced by the participant. There are a number of ways this referent could be computed, including such possibilities as a simple running average of the harmony functions for all experienced stimuli. If it does reflect a running average, it may be a running average of all trials, or perhaps of only a certain recent window of trials.

It is assumed that participants have some referent they can compare the current item to, to gauge the extent to which the harmony of the current item is more or less than the harmony of the average stimulus at that point in time. This referent function would lie somewhere between the word and nonword harmony functions as depicted in Figure 1.

Given the existence of a referent function, the word versus nonword decision process can now be thought of in terms of a random-walk process similar to that originally described in Link & Heath's (1975) wave theory (see also Ratcliff, 1978). The basic idea of a random-walk process is that any binary-decision task can be thought of as a random walk towards one of two boundaries, each representing one of the two possible decisions. The walk is driven by the accumulation of evidence in favour of one or the other response such that any evidence in favour of response 1 moves the walk towards boundary 1, and any evidence in favour of response 2 moves the walk towards boundary 2. Thus, this accumulator of evidence slowly drifts towards one or the other boundary but does so in a somewhat jagged manner.

Random-walk models provide a very good account of the reaction time and error patterns for a number of binary decision tasks (see Ratcliff, Van Zandt, & McKoon, 1999), and therefore seem like an ideal candidate for the decision process underlying lexical decision. One can imagine a word and nonword boundary, and some evidence accumulator that moves towards one boundary or another, eventually leading to a word or nonword decision.

However, what is often missing in random-walk accounts, and what is provided in the referent model, is an explicit description of the evidence accumulator itself. In the referent model, processing within the attractor network provides the engine that drives the evidence accumulator. The evidence of interest is the discrepancy between the level of harmony of the current stimulus and the referent level. For words, this discrepancy will become progressively larger in a positive direction as the word is processed. For nonwords, the discrepancy will become progressively larger in a negative direction as the nonword is processed (see Figure 1).

Thus, the lexical-decision process can be depicted in a typical random-walk graph as in Figure 2. The ordinate of the graph represents the extent to which the harmony of an item deviates from the referent. The abscissa reflects time. Any point on the graph reflects the extent to which an item's harmony differs from the referent level at some point in time. Given this, the horizontal line that runs through the middle of the graph represents the referent level because its deviation from the referent is zero at all points in time. The horizontal line at the top of the graph reflects the amount of harmony greater than the referent level that is required to support a word decision. Thus, it is called the word boundary. A similar nonword boundary lies at the bottom of the graph reflecting the amount of harmony less that the harmony level that is required to support a nonword decision.

Insert Figure 2 about here

The two lines depicted on this graph are meant to reflect the processing of a word and nonword respectively. Note that as a word is processed, its harmony increases at a rate greater than the referent, eventually resulting in it hitting the word boundary. The opposite occurs for a nonword. The rate at which any item approaches either boundary is termed the drift rate. This drift rate is determined by the extent to which the stimulus differs from the "average" stimulus reflected by the referent. In the case of words, the more wordlike a stimulus is, the faster it would deviate from the referent level of harmony, and therefore the greater would be its drift rate. A similar though opposite situation would be the case for nonwords. The less wordlike a nonword is, the faster it would approach the nonword criterion, and hence the higher its drift rate in a negative direction.

Efficiency and Bias

There is one additional concept necessary to complete the description of the referent model; a concept termed the efficiency principle. According to the efficiency principle, people will attempt to be as fast and accurate as possible to both words and nonwords when they are performing a lexical decision. This attempt to be as fast and accurate as possible is made using a response bias. The response bias operates on the decision processes, not the memory retrieval processes.

It is important to specify the way in which it is assumed to affect performance. The starting point with respect to this issue is the assumption there is a constant distance between the word and nonword boundaries. However, the position of the boundaries relative to the referent harmony level can be changed. Specifically, in a situation where participants are biased towards word responses, the assumption is that the boundaries shift down such that the word boundary is now closer to the starting position than is the nonword boundary. Similarly, a bias towards nonwords responses would be realized by shifting the boundaries up.

When a participant begins a lexical-decision task, they are typically instructed to make their decisions as quickly and accurately as possible. How are these instructions applied to the decision process? According to the referent model, the participant incorporates these instructions throughout the course of the experiment by moving the boundaries up or down in order to find a position that produces the fastest responses and lowest error rates.

For example, consider the situation depicted across the panels of Figure 2 where the both the words and nonwords are not very wordlike. The words might be low frequency words like SLALOM, and the nonwords might be scrambled letter strings like EKLTTE. In this situation, the drift rate of the words toward the word boundary would be slow, whereas the drift rate of the nonwords towards the nonword boundary would be fast. If, as depicted in Panel A, the boundaries were positioned equidistant from the referent position, the result would be fast nonword decisions but slow word decisions. If, however, the boundaries were shifted down such that the word boundary was closer to the referent, as depicted in Panel B, word decisions would speed up substantially, whereas the nonword decisions would slow only slightly. Thus, moving the word boundary closer to the referent would lead to more efficient overall responding.

It is the variation of the boundary setting that we see as responsible for the dynamic nature of the lexical-decision process. The boundaries for any given trial would be set on the basis of the previous stimuli the participant had experienced. For example, if several low frequency words had been experienced, the boundaries would shift down, which would cause slower and more error prone responses to subsequently presented nonwords. In contrast, if several wordlike nonwords were experienced (i.e., nonwords with a slow drift rate towards the nonword boundary), the boundaries would shift up, leading to slower and more error prone word decisions. The flexibility of this mechanism is critical to most of the referent model's predictions and accounts.

For example, let's consider how the referent model accounts for the two components of the James (1975) effect; that (1) word reaction times slow in the presence of more wordlike nonwords, and (2) semantic influences increase in magnitude in the context of more wordlike nonwords. In terms of the referent model, the important difference between less and more wordlike nonwords is the drift rate of the random walk on nonword trials. The less wordlike a nonword, the higher its drift rate towards the nonword boundary. Thus, using referent model terminology, the James Effect can be restated as follows: As the drift rate of the nonwords towards the nonword boundary decreases, words are processed longer and to a greater extent before the word boundary is reached.

In fact, this makes perfect sense given the efficiency principle. The more wordlike a nonword is, the slower its drift rate towards the nonword boundary. To keep efficiency high, the decision process should move the boundaries up such that the nonword boundary is closer to the referent. However, the result is that the word boundary is now further from the referent, meaning the words will be processed longer prior to reaching the word criterion, even though the drift rate of the words has not been altered. The increased processing of words should be primarily beneficial to the formation of a more complete semantic representation given the assumption that the semantic level of representation is the last to be completed (see Joordens & Becker, 1997, or Masson, 1995 for a defense of this assumption). Hence, semantic effects should increase in magnitude. Thus, both components of the James Effect follow naturally from the referent model when the efficiency principle is assumed.

The referent model has merits that go beyond it providing a novel account of the James Effect, or the dynamic nature of lexical decision in general. These merits are evident at a number of levels. On a theoretical level, the model makes a useful distinction between the basic psychological process of word recognition, and the random-walk decision process that can be overlaid on the word recognition process to support lexical-decision performance. The word-recognition process is not assumed to vary with respect to how it works from trial to trial, which seems reasonable given it should be a highly over-learned procedure. The component that changes in response to contextual factors is the decision process. Thus, strategies show themselves in the decision process, not in the basic psychological process (cf., Stone & Van Orden, 1993).

It also provides a psychological description for factors critical to a random-walk decision process. For example, the randomness of the random-walk process is critical for allowing random-walk models to account for errors. Sometimes the walk is random enough that the evidence accumulator hits the wrong boundary. But what causes the randomness? The referent model provides an explicit answer to this. Most attractor models of word recognition contain some element of stochasticity. For example, in Masson's (1991; 1995) model, nodes are randomly selected for updating. This randomness results in the patterns being retrieved in "fits and starts" with some of the updates resulting in little to no increase in harmony (e.g., if the same node is sampled twice in a row) and others leading to large increases. Thus, relative to a constantly increasing referent level, the harmony discrepancy can increase and decrease producing the characteristic "jagged" pattern.

Discriminating Between the Models

Although the multiple read-out and referent models share much in the way of global philosophy, they clearly differ in terms of details concerning both the memory processes (and representations) and the decision processes. We will focus on the latter of these, and attempt to use the differences in decision processes to generate differential predictions in a novel experimental context.

Prior to providing this experimental evidence for the referent model, computer simulations of the referent model will be presented in Chapter 3. This will provide us with a concrete example of the referent model and will allow us to make specific predictions about the performance of the subjects that will be tested in Chapter 4. Then, in Chapter 5, we will compare how well the referent model and the multiple read-out model can account for the findings of both the simulations and the experiments.

CHAPTER 3

The purpose of this chapter is to provide a concrete description of the referent model of lexical decision. In Chapter 2, a general class of models that couples an attractor network of word processing with a decision module that utilizes a random-walk process was briefly described. In this chapter, the structure of a specific network, and the decision mechanism that is built on top of it, is discussed in more detail.

Before we can discuss the structure of the model, we need to decide what type of network we will use to perform these simulations. The network must be able to provide us with the following: First, the network that we choose must use distributed representations. The distributed nature of the network allows us to form representations over time. We are then able to compare the representation that is forming to memory and categorize items based on this comparison. Second, we must also be able to obtain some sort of reaction time measurement from the network. We need to compare the response latency of the network in a variety of conditions to the response latencies of human subjects. Finally, the network must be able to process novel items (i.e. nonwords). Lexical decisions are made by categorizing both words and nonwords. One type of network that provides us with all of these requirements is a Hopfield network.

Structure of a Hopfield Network

A Hopfield network (Hopfield, 1982) consists of a set of interconnected nodes. Unlike many other parallel distributed processing (PDP) networks that are arranged in layers of processing units (e.g., Dawson, 1998, Chapter 3), a Hopfield network consists of only one layer of processors. These processors simultaneously represent input to the network, and the network's (eventual) response to this input. In a standard Hopfield network, the processing units are in one of two possible states of activation: off (the processor has an activation value of −1) or on (the processor has an activation value of +1).

A Hopfield network is an autoassociative system in which every processor is connected to every other processor in the network by a weighted connection. Typically, processors in a Hopfield network are not connected to themselves, although such referent connections have been used in autoassociative systems that are related to Hopfield nets (e.g., Dawson's (1991) brainstate-in-a-sphere model). Figure 3 illustrates the typical structure of a Hopfield network.

Insert Figure 3 about here

A Hopfield network is typically used as a model of item recognition or item completion, where an item is represented as a pattern of activation across all of the network's processing units. Item recognition by a Hopfield network requires two processing stages. In the first stage known as learning, the connection weights are modified in order to store memories of different items in the network. In the second stage known as retrieval, a stimulus is presented to the network, and information is retrieved from its memories. The stimulus might be a novel item, an item that has been presented before, or even a distorted version of a previously remembered item. The stimulus serves to disturb the "energy" of the Hopfield network. During retrieval, the network modifies the activities of its input units in such a way as to minimize this energy disturbance. As this is done over time, the network gradually settles into a minimum energy state in which all processing units will have stable activations. This stable pattern represents the memory retrieved from the Hopfield network by the original stimulus.

One example of the type of representations that a Hopfield network can process is presented in Figure 4. In this example the network learns patterns that represent letters of the alphabet. The alphabet letters are formed by giving certain units a positive weight and other units a negative weight.

Insert Figure 4 about here

During learning, the network is presented with each of the patterns one at a time and the connection weights between the units are modified. These connections weights contain the information that represents each of the items that are presented. Each time an item is presented the weights between the units are modified using the following simple Hebbian learning algorithm (Hebb, 1949):

$$\Delta w_{ij} = n_i n_j$$

This learning algorithm strengthens the connection between two units that are the same and weakens the connection between units that are different. After learning has occurred, the network has developed a single matrix of connection weights that contain information about each of the patterns that it has learned. Therefore, memory for previously learned items is distributed across these connection weights.

The network is now ready to retrieve information from its stored memories. At the start of retrieval, the network is presented with some stimulus pattern. Then, in order to simulate the fact that even parallel, brain-like systems require time to process information, one processing unit is chosen randomly for updating. The total input to this unit is calculated by summing the input from all other units to this unit. If the sum is greater than zero, the chosen unit is turned on (i.e., given a value of 1). If the sum is less than zero, the chosen unit is turned off (i.e., given a value of –1). The formula for updating the activity of a randomly selected unit during retrieval is as follows:

$$a_i = \sum w_{ij} n_j, \quad i \neq j,$$

where a_i is the summed input to Unit i, w_{ij} is the weight between Unit i and j, and n_j is the state (i.e., 1 or –1) of Unit j. The retrieval process is then continued by randomly choosing and updating another processing unit.

When a unit is updated, its activity will either be changed or will remain the same. After a Hopfield network has reached a stable configuration, none of the units will change activity when they are updated. In a simulation, one tallies the number of updates that have occurred without a change occurring. This tally can be used to stop the simulation, by operationalizing a stable network as one in which there have been no processing unit changes in the past x number of cycles, where x is suitably large (e.g. 500 cycles). When the network units have remained unchanged for this number of cycles, the network is said to be stable.

To illustrate the process of retrieval, consider the network in Figure 5 that has learned alphabetic patterns. During retrieval the network is presented with a pattern that is similar to one of the previously learned alphabet letters. In Figure 5, Panel A shows the initial pattern that is presented to the network. This pattern is created by taking the pattern that represents the letter "A" and then adding noise to it by randomly setting 25% of the units. Panel B shows the pattern that has formed after 30 cycles have occurred. You can see that the pattern is becoming similar to letter "A". Panel C shows that the pattern that has formed, after the network has settled, represents the letter "A". This particular network stabilized after 86 updates.

Insert Figure 5 about here

The Hopfield network provides us with the general structure that is needed to perform lexical decision simulations. Items are represented in these networks across a set of units. During retrieval, representations form over time. Decisions can be made by determining how similar the representation that is forming is to previously learned items. As was mentioned earlier, these characteristics are needed to model lexical decision. In the next section we will get a better idea of how the Hopfield network can be used to model lexical decision performance.

A Hopfield Network For The Referent Model

The previous section nicely illustrated the basic structure of a Hopfield network. The network structure that is need for lexical decision simulations is very similar. Figure 6 illustrates the type of representations that are presented to the network during learning. There are a total of 125 nodes for each item that is presented to the network. The first 25 nodes represent the orthography of the item. The remaining 100 nodes represent phonology and semantics. During learning, each of the items are presented to the network one at a time. As in the previous example, a weight matrix is calculated that contains a distributed representation of all the items that have been learned. The weight matrix is calculated by comparing the similarities and differences between each of the patterns that are presented to the network during learning. These patterns represent words that are stored in memory. Nonwords are simply patterns that were not presented to the network during training.

Insert Figure 6 about here

During retrieval we can present the network with either previously learned items, which will be referred to as words from this point forward, or novel items, which will be referred to as nonwords. Word presentations are made by setting the orthographic nodes in a state that is consistent with a previously learned item, and randomly setting the remaining nodes. Nonword presentations are made by presenting the network with a pattern that is not consistent with any previously learned item. This would be similar to presenting a subject with a string of letters on a computer screen during a lexical decision experiment.

When word items are presented, the orthographic pattern is clamped and only the remaining 100 nodes can be chosen for updating. To say that the pattern is clamped is to say that the units representing the orthographic information have activity values that cannot be changed. The logic for clamping these units is that the orthographic information is presented, unchanging, to the subject as a stimulus, and is therefore "clamped" by the environment. This also means that the network can only change by updating its other processing units.

Updating continues until the network becomes stable. At this point, the current pattern is the same as a previously learned item. When nonwords are presented, the orthographic nodes are clamped in a state that is not consistent with a previously learned item. The network randomly chooses the remaining 100 nodes for updating. Again, updating occurs until the network becomes stable. However, at this point the current pattern is not the same as a previously learned item but it is similar in some way to all previously learned items. The network takes the random pattern that is initially presented, and tries to match it to something in memory. However, since the orthographic nodes are clamped, it will never achieve an exact match to any item in memory.

Both words and nonwords can be presented to the network for retrieval. As retrieval occurs over time, the representation that forms becomes more and more like items that are stored in memory. For word items, they will eventually become exactly the same as an item that is stored in memory. For nonword items, they will become similar to previously learned items but will never be exactly the same as an item stored in memory. To be able to differentiate between words and nonwords we need to be able to measure how similar the representations that are forming, are to items in memory. To do this we need to go beyond the equations that define a standard Hopfield network, and include an additional equation for a measure called harmony.

As each of the items is retrieved from memory a harmony value is calculated. This value represents the similarity between the current state of the network and previously learned items. Harmony can be calculated for both words and nonwords. Although nonwords have not been learned by the network, it will still attempt to retrieve them from memory. In the case of nonwords, the network will end up in a state that is in some way similar to all previously learned items, but not exactly the same as any single previously learned item.

Think of this in terms of words and nonwords being presented to a human during a lexical decision experiment. The nonwords that are presented are similar to the word items that are stored in memory; they both consist of letters. Therefore the nonword begins to be processed and a representation begins to form. However, the representation that forms never reaches the same level of similarity as a previously learned word.

The formula for harmony is as follows (Smolensky, 1986):

$$Harmony = \sum_{i=1}^{n} \sum_{j=1}^{i} n_i n_j w_{ij}$$

where n_i is the state of Unit i, n_j is the state of Unit j, and w_{ij} is the weight between Unit i and j which was calculated during the learning phase. As the network begins to retrieve an item from memory, the harmony value is calculated after each update. The level of harmony will continue to increase until the network has settled and there is no further change in the pattern of nodes.

At this point we have a Hopfield network that processes both words and nonwords. We also have a way to measure how similar the representation that is forming is to previously learned items. It has been made apparent that words will achieve a level of harmony that is greater than that of nonwords. However, what has not been made apparent is how the network is able to distinguish between words and nonwords. To perform this task we need to add to the model the referent decision mechanism that was referred to in Chapter 2.

The referent is used during recall to allow the network to categorize items as being either words or nonwords. The referent is calculated as the average increase in harmony for all items that have been presented. To categorize the items, the networks current level of harmony is compared to the referent. When the current level of harmony is greater than the referent by a critical amount, that item is categorized as being a word. When it is lower than the referent by a critical amount, the item is categorized as being a nonword. The referent is calculated as the running average of harmony for all experienced stimuli. The formula for the referent is as follows:

$$r_t = 1/q\sum\sum h_{tj}$$

where q is the number of items that have been presented to the network during retrieval and h_{tj} is the level of harmony for all j items over t time. For example, if twelve items have been presented to the network for retrieval, r_3 would be equal to the sum, across all 12 items, of the level of harmony that was achieved after three updating cycles were preformed. This sum would then be divided by the number of items that were presented to the network. In this case that would be 12 items.

The referent decision mechanism allows us to categorize both words and nonwords based on the measurement of harmony. The measurement of harmony is made possible by the basic structure of the Hopfield network. The Hopfield network provides us with a distributed representation of memory. When an items representation is retrieved from memory, the process occurs over time. Therefore we are able to calculate both reaction time and accuracy as the network performs the lexical decision task. Based on the preceding description, the Hopfield network seems well suited to perform simulations of the referent model of lexical decision.

The next step is to perform actual network simulations that are based on the previous description. For the remainder of this chapter, a series of simulations will be presented that utilize the Hopfield network structure. These simulations will allow us to make specific predictions about human performance on similar tasks in Chapter 4. Then, later on, in Chapter 5, we will be better able to make a comparison between the referent model and the multiple read-out model.

Network Structure

Architecture

All simulations were programmed using the C programming language. The simulations were performed using a single-layered Hopfield network (Hopfield, 1982; Hopfield & Tank, 1986) that was augmented with equations for calculating harmony and for calculating the harmony of the referent. The network consists of a total of 125 units. These units are divided into two different groups of 25 perceptual units and 100 conceptual units. The perceptual units represent the visual stimuli that are presented to a person during a lexical decision task and the conceptual units represent the orthography, semantics, etc. of the stimuli. All of the units in the network are set to either 1 or –1 and are completely interconnected.

Parameters for Lexical Decision

A number of parameters need to be set prior to beginning the simulation. These parameters consist of the critical word criterion, critical nonword criterion, correct decision bias, and incorrect decision bias. When the current level of harmony is greater than the current referent value by an amount equal to or greater than the critical word criterion, a word decision was made. When harmony was less than the referent by an amount equal to or greater than the critical nonword criterion, a nonword decision was made. The word and nonword criterion can be shifted closer to or further away from the referent. This criterion shift is controlled by the response bias. After each response is made, the criterion values are adjusted according to the correct and incorrect decision biases.

The network structure that has been described will be used for all of the simulations that are presented in this chapter. The parameters that are set for each simulation will be described in each of the methods sections. Predictions will be made based on the referent model assumptions that were provided in Chapter 2.

Network Simulations

Simulation 1: The Relationship Between Nonword Context and Word Responses

In Simulation 1, a simple nonword manipulation was performed. The network was presented with scrambled nonwords for the first half of a lexical decision task followed by pseudohomophones for the second half. The purpose of Simulation 1 was to determine if changes to nonword type would affect reaction times and accuracy for word decisions. A major assumption of the referent model is that word and nonword decisions are intimately related. Therefore, changes to one item type should have an affect on the other. Word and nonword decisions should take longer when the nonword type is switched from scrambled to pseudohomophones. Pseudohomophones are more wordlike; therefore it should become more difficult for the network to distinguish them from words than when scrambled nonwords are presented. Also, accuracy for both words and nonwords should decrease.

Method

Items that were presented to the network during training will be referred to as words. Novel items that are presented to the network during retrieval will be referred to as nonwords. During each iteration of the simulation the following occurred; a) a learning stage where word items are presented to the network; b) a retrieval stage where harmony and the referent are both calculated and a word/nonword decision is made.

During the learning stage, the network was presented with each of the items once. Words and nonwords were created by randomly setting all 125 units in the network. A total of 3 words and 3 nonwords were used for each cycle of the simulation. Each cycle consisted of a learning stage where each of the 3 words was learned, and a retrieval stage where each of the 3 words and nonwords were presented to the network for retrieval. During the retrieval stage, lexical decisions were made as each of the word and nonword items was processed. For each cycle a new set of words and nonwords were created. Therefore the number of unique words and nonwords that were used is equal to three times the number of cycles. For this simulation, 500 cycles were used for a total of 1500 unique words and 1500 unique nonwords.

During the retrieval stage, the network retrieved each of the items twice. The conceptual units were randomly set between the first and second presentations of each item to avoid obvious repetition problems. This made a total of 6 words and 6 nonwords that were retrieved for each of the 500 cycles, for a total of 3000 words and 3000 nonwords. During the first 250 cycles scrambled nonwords were presented to the network. Scrambled nonwords were created by taking one of the previously learned word items and randomly setting 20 of the 25 orthographic units and then randomly setting the remaining 100 units. During the second 250 iterations pseudohomophones were presented to the network. Pseudohomophones were created by taking one of the previously learned word items and randomly setting 10 of the 25 orthographic units and then randomly setting the remaining 100 units. By varying the number of orthographic units that were changed, the wordlikeness of the nonwords was manipulated.

During the retrieval stage the harmony level was calculated for each item as it was being retrieved. After a single unit was chosen and updated, the current level of harmony was calculated and compared to the referent level for that cycle. As soon as either the word or nonword criterion was reached, a lexical decision was made.

The referent was calculated using a moving window over the previous 108 items that were presented. An average was calculated using both words and nonwords. Decisions were made based on this running average.

There are a number of parameters that are set before the simulation began. The critical word and nonword values were both set at 2000. This created a distance between the word and nonword boundaries of 4000. This distance was determined by running a pilot simulation where the distance between the average harmony for words and the average harmony for nonwords was calculated. The correct decision biases for both words and nonwords were set to 25, and the incorrect bias for both words and nonwords were set to 50. The incorrect and correct decision biases shift the boundaries in relation to the referent. The boundary shifts that occur when correct or incorrect nonword decisions are made are as follows; (a) if a nonword item is presented and the network makes an incorrect word decision, the nonword boundary moves towards the referent. This in turn shifts the word boundary away from the referent because the distance between the two boundaries remains constant, (b) if a nonword is presented to the network and a correct nonword decision is made, the nonword boundary shifts away from the referent. This in turn shifts the word boundary closer to the referent. According to the efficiency bias, subjects are trying to be as fast and accurate to both words and nonwords. When the correct nonword decision is made, the nonword boundary moves away from the referent which makes it easier to make a word decision. The same type of boundary shifts occur when correct or incorrect word decisions are made.

Results and Discussion

Phase 2 and Phase 4 are used for the analysis of this simulation. Phase 2 consists of the last 1008 items before the transition. Phase 4 consists of the last 1008 items that are presented to the network. The network is learning to perform the task during Phase 1 so it is not included in this analysis. During Phase 3 the type of nonwords that are being presented are switched from pseudohomophones to scrambled nonwords. This Phase is not included in the analysis so that a clean comparison of performance can be made between each condition. The means and standard deviations for both reaction time and accuracy are presented in Table 1.

Insert Table 1 about here

Phase 2 vs. Phase 4. Based on the predictions of the referent model, we expected two things to occur when the type of nonwords that were being presented were switched from scrambled nonwords to pseudohomophones. First, decisions for nonwords should have taken longer when pseudohomophones were presented than when scrambled nonwords were presented. Second, this manipulation should have affected word decision by increasing reaction times when pseudohomophones were presented compared to when scrambled nonwords were presented.

As expected, correct word decisions take longer when presented with pseudohomophones (M=353) than with scrambled nonwords (M=173), $F(1,5)=172.451$, $p<.0001$, $MSe = 564.68$. Also, word errors are higher when presented with pseudohomophones (M=27%) than when presented with scrambled nonwords (M=18%), $F(1,5)=125.93$, $p<.0001$, $MSe = 2.15$. This is consistent with the notion that the lexical decision becomes more difficult as the nonwords become more wordlike.

The same pattern of results is found in the nonword error data but not in the reaction time data. Errors for pseudohomophones are higher (M=26%) than errors for scrambled nonwords (M=16%), $F(1,5)=195.562$, $p<.0001$, $MSe = 1.483$. This is what is expected if pseudohomophones are more similar to words than scrambled nonwords. However, correct pseudohomophone decisions are made more quickly (M=83) than correct scrambled nonword decisions (M=103), $F(1,5)=104.472$, $p<.001$, $MSe = 10.733$. At first glance the nonword reaction time results seem to pose a problem for the theory. Pseudohomophone decisions are more difficult to make than scrambled nonword decisions; therefore they should take longer. However, what is happening is that some of the pseudohomophones are very similar to words and are incorrectly being called words. These pseudohomophones would eventually recross the word boundary and continue to be processed until they eventually crossed the nonword boundary. But since they are being incorrectly called words, they are not included in the reaction time analysis. If we look at the reaction time for all nonwords, both correct and incorrect, we find that decisions to pseudohomophones do take longer ($\underline{M}$=267) than to scrambled nonwords ($\underline{M}$=219), $\underline{F}(1,5)=93.298$, $\underline{p}<.001$, $\underline{MSe} = 72.55$.

By manipulating the wordlikeness of the nonwords, we can affect both reaction time and accuracy for words. To maximize both reaction time and accuracy when the nonwords are manipulated, there must be a shift in the bias. This shift in turn affects word reaction time and accuracy. This provides evidence for the intimate relationship between word and nonword decisions.

Simulation 2: The Relationship Between Word Frequency and Nonword Responses

In Simulation 1 it was evident that nonword manipulations affect word decisions. This is a typical effect that is found in the lexical decision literature (Craik & Lockhart, 1972; James, 1975; Borowsky & Masson, 1996; Joordens & Becker, 1997; Piercey & Joordens, 2000). These studies make it apparent that manipulating the type of nonword that is presented affects word decision. However, to show that there is an intimate relationship between words and nonwords, we also need to manipulate the type of words that are presented and affect decision to nonwords. The purpose of Simulation 2 was to determine if we could produce a similar interaction between words and nonwords when word frequency was manipulated.

The type of words that were presented was switched from low frequency items for the first half of the simulation, to high frequency items for the second half. According to the referent model, changing from low to high frequency words should affect both word and nonword reaction time and accuracy. Again, the assumption that word and nonword responses are intimately related leads to specific predictions.

When the word frequency is switched from low to high, reaction times to both words and nonwords become faster. Also, accuracy for both words and nonwords become faster.

Method

This simulation was structured the same as in Simulation 1 except for two components. First, all nonwords were created by randomly setting all 125 units. Therefore the nonwords were not similar to the words. Second, for the first 250 cycles the network was trained with each of the words one time. For the second 250 cycles the network was trained with each of the words two times. Therefore, the items that are learned once during the first 250 cycles are referred to as low frequency items and the items that are learned two times during the second 250 cycles are referred to as high frequency items.

Results and Discussion

Phase 2 and Phase 4 were again used during this analysis. A total of 1008 items were analyzed for each of these Phases. The means and standard deviations for both reaction time and accuracy are presented in Table 2.

Insert Table 2 about here

As expected correct decisions to high frequency words were made faster (M=42) than correct decisions to low frequency words (M=167), $F(1,11)=1789.034$, $p<.0001$, Mse = 52.19. Also, the error rate was lower for high frequency words (M=14%) than for low frequency words (M=16%), $F(1,11)=13.146$, $p<.004$, Mse = 2.485. The same pattern of results was found in the reaction time data for nonwords where correct nonword decisions during the high frequency phase were faster (M=65) than during the low frequency phase (M=107), $F(1,11)=638.35$, $p<.0001$, Mse = 16.71. There was no difference found in error rates for nonwords during the high frequency phase (M=13%) and during the low frequency phase (M=13), $F(1,11)=0.673$, $p=.429$, Mse = 2.227.

As predicted by the referent model, manipulation the type of words that were presented affected decisions to nonwords. Again we have evidence for the intimate relationship between words and nonwords. As it became easier to make word decisions, the bias shifts which in turn makes nonword decisions easier.

Simulation 3A and 3B: The Relationship Between Word and Nonword Ratio

The purpose of Simulation 3 is to further test the referent model by examining a novel prediction that can be derived from it. Specifically, as previously highlighted, the referent accounts for the dynamic lexical-decision findings via changes in the positioning of the word and nonword boundaries. Said another way, the "strategic" changes in process are seen as merely reflecting shifts in bias to respond either word or nonword, and this bias shifts are the result of the specific task context. This notion could be tested directly by examining other variables that have a more obvious connection to bias, and then seeing if they produce effects similar to those associated with changes in word or nonword context.

The specific variable we will consider here is the word-to-nonword ratio. Given that the bias is assumed to reflect a predisposition towards either a word or nonword response, the best way to manipulate it would be to make one lexical class more prevalent than the other. If words are more prevalent, the network should be biased to respond “word”, if nonwords are more prevalent, it should be biased to respond “nonword”.

In order to describe the referent model's predictions with respect to word-to-nonword ratio, it is necessary to first revisit the efficiency principle. The central tenet of that principle is that participants are trying to maximize the efficiency of their overall responding when they are in a lexical-decision context. They are assumed to do so by shifting the word and nonword boundaries either up or down relevant to the referent position. When there is an equal number of words and nonwords in the experiment, they are assumed to base the position of the criteria primarily on the drift rates associated with the classes of stimuli they encounter.

However, what should happen if there were more words than nonwords, or vice-versa? In order to truly maximize their overall efficiency of responding, participants should also take this variable into account when setting their word and nonword criteria. Specifically, they should shift the criteria such that the criteria associated with the more prevalent stimulus class ends up closer to the referent function than it otherwise would be. Such a shift would make them more efficient in responding to the prevalent stimulus class, with only a slight cost to the non-prevalent.

This leads to the following predictions. Consider the situation where there is an even number of words and nonwords to be a baseline condition. If we compared a context containing more words than nonwords to this baseline, we would expect the starting position to shift slightly towards the word boundary which should lead to; (a) faster word responses, (b) slower nonword responses, (c) less word errors, and (d) more nonword errors. In contrast, if we compared a situation with less words than nonwords, we would expect a shift of the starting position towards the nonword boundary which should lead to; (a) slower word responses, (b) faster nonword responses, (c) more word errors, and (d) less nonword errors. Simulation 3A embodies the first of these contrasts, Simulation 3B embodies the other.

Method

These simulations again had the same basic structure as Simulation 1 except all nonwords were created by randomly setting all 125 units. Therefore the nonwords were not similar to the words. For simulation 3A, the network was presented with a ratio of 1:1 words to nonwords for the first 250 cycles followed by a ratio of 2:1 words to nonwords for the second 250 cycles. During the 1:1 ratio phase both the words and nonwords were presented for recall one time. This makes a total of 250 word and nonword presentations during this phase. During the 2:1 phase the words were each presented twice and the nonwords were presented once. During this phase 250 unique words and nonwords are used for a total of 500 word presentations and 250 nonword presentations.

The opposite was true for Simulation 3B where the network was presented with a 1:1 ratio of words to nonwords for the first 250 cycles followed by a 1:2 ratio of words to nonwords for the second 250 cycles. During the 1:1 ratio phase both the words and nonwords were presented for recall one time. This makes a total of 250 word and nonword presentations during this phase. During the 1:2 phase the words were each presented once and the nonwords were presented twice. During this phase 250 unique words and nonwords are used for a total of 250 word presentations and 500 nonword presentations.

For both Simulation 3A and 3B the location of the word and nonword were allowed to shift during the 1:1 phase but were held constant when the word to nonword ratios were not equal. In Simulation 3A during the 2:1 phase the word critical value was held constant at 100 and the nonword critical value was held constant at 3900. In Simulation 3B during the 1:2 phase the word critical value was held constant at 3900 and the nonword critical value was held constant at 100. In both cases this maintained a distance of 4000 between the two critical values. By holding the critical values constant we were able to be sure that the bias towards either the word or nonword boundary would be maintained.

Results and Discussion

Phase 2 and Phase 4 were again used for this analysis. During the 1:1 Phase there were 204 words and nonwords. During the 2:1 phase, 276 words and 138 nonwords were used for the analysis. During the 1:2 phase 138 words and 276 nonwords were used during the analysis. The results from Simulation 3A and 3B are depicted in Table 3. The statistical analysis will involve planned comparisons examining the affect of word:nonword ratio separately for words and nonwords.

Insert Table 3 about here

Simulation 3A. Simulation 3A involves a contrast of word and nonword responses across conditions where there were either an equal number of words and nonwords presented or twice as many nonwords than words presented. The referent model predicts that responses to nonwords will be faster and less error prone, and responses to words will be slower and more error prone, in the condition where there are more nonwords than words relative to the condition where they are equally represented.

Planned comparisons conducted revealed that, as predicted by the referent model, nonwords were responded to significantly faster (i.e., 47 cycles faster) in the condition where they were the dominant stimulus class, $\underline{F}(1,5) = 232.653$, $\underline{p} < .0001$, $\underline{MSe} = 27.483$. However, the model's prediction concerning response times to words was not borne out. Also, words were responded to slower in the condition where nonwords were dominant (i.e., 102 cycles slower), $\underline{F}(1,5) = 84.089$, $\underline{p} < .001$, $\underline{MSe} = 367.55$.

The analogous analysis performed on the error rate data revealed a significant reduction in errors to nonwords (i.e., 12% more accurate) in the condition where they were the dominant stimulus, $\underline{F}(1,5) = 529.0$, $\underline{p} < .0001$, $\underline{MSe} = .75$. Also, there was a significant increase in errors to words (i.e., 12% higher) in the condition where nonwords were dominant was reliable at conventional levels of significance, $\underline{F}(1,5) = 94.697$, $p < .001$, $\underline{MSe} = 4.95$. Thus, as predicted by the referent model, responses to nonwords became less error prone, and responses to words more error prone, in the condition where nonwords were the dominant stimulus class. The means and standard deviations for these variables are presented in Table 3.

<u>Simulation 3B.</u> Simulation 3B involves a contrast of word and nonword responses across conditions where there were either an equal number of words and nonwords presented or twice as many words than nonwords were presented. The referent model predicts that responses to nonwords will be slower and more error prone, and responses to words will be faster and less error prone, in the condition

where there are more words than nonwords relative to the condition where they are equally represented.

Planned comparisons focused on this interaction revealed that the 100 cycle slowdown in nonword responding in the condition where the words were dominant was statistically significant, $\underline{F}(1,5) = 133.215$, $\underline{p} < .0001$, $\underline{MSe} = 225.2$. Also, the 52 cycle speed-up of word responses in the condition where they were the dominant stimulus class was significant, $\underline{F}(1,5) = 120.266$, $\underline{p}$ < ..001, $\underline{MSe} = 67.883$.

The analogous analysis performed on the error rate data revealed that the 53% increase in errors to nonwords in the condition where words were the dominant stimulus-class was significance, $\underline{F}(1,5) = 600.172$, $\underline{p} < .0001$, $\underline{MSe} = 12.083$. In addition, the 13% reduction in errors to words in the condition where they were dominant was also significant, $\underline{F}(1,5) = 1743.823$, $p < .0001$, $\underline{MSe} = 0.2833$. Thus, as predicted by the referent model, responses to nonwords became more error prone, and responses to words less error prone, in the condition where words were the dominant stimulus-class.

The data from Simulation #3 is again consistent with the notion of an intimate relationship between word and nonword decision. By manipulating the ratio of words to nonwords, we would expect a bias shift. When this bias shift occurs, we are able to make specific predictions about performance of both words and nonwords. Again this provides us with evidence of an intimate relationship between words and nonwords.

<u>Summary</u>

The purpose of this chapter was to provide a concrete description of the referent model of lexical decision. The network uses a distributed representation of memory. When the network is presented with an item during a lexical decision, it forms a representation of that item over time. As this representation forms, it becomes more and more like something that is already in memory. Then, a lexical decision is made based on the similarity of the representation and memory. It is this decision process and the intimate relationship between words and nonwords that will allow us, in Chapter 5, to justify our claim that the referent model provides a better description, than the multiple read-out model, of how lexical decisions are being performed.

In Chapter 4, further evidence for the referent model will be provided in the form of data from human experiments that are consistent with the data from the simulations that were just provided. Again we will be able to make specific predictions about reaction time and accuracy when we manipulate both word and nonword stimuli. Changes in one stimulus type produce predictable changes for both the stimulus being manipulated and the stimulus within the context of this manipulation.

CHAPTER 4

In Chapter 3, a working version of the referent model of lexical decision was described. When we perform computer simulations of psychological theories, we are forced to address some design questions. We need to be explicit about the mechanisms that are being used when subjects are performing the specific task that the theory addresses. So, by having an actual working model, we are better able to theorize about how humans are performing the task. Then, taking this knowledge into account, we are able to make specific predictions of how humans will perform when doing the same task. If the subjects' performance is similar to the performance of computer simulations, then we can argue that we have a plausible theory about the mechanisms that humans are using when performing the task

The purpose of this chapter is to provide evidence from human experimentation that tests the referent model and the simulations that were presented in the previous chapter. We will systematically go through each of the simulations and perform similar experiments on human subjects. The predictions of the referent model will be discussed and the data will be analyzed to determine if the results from the experiments support the predictions. In this chapter we will not make a comparison between the referent model and the multiple read-out model. We are simply presenting further data that supports the predictions of the referent model. In Chapter 5 we will compare the predictions that each of the models make with the data from the simulations and the human experiments.

Experiments with Human Subjects

Experiment 1: The Effect of Nonword Context on Word Responses

This experiment will be similar to Simulation 1 that was presented in Chapter 3. For the first half of the experiment, subjects will be presented with scrambled nonwords that will be switched to pseudohomophones during the second half. According to the predictions of the referent model, the task should become more difficult when the pseudohomophone nonwords are presented and reaction times as well as error rates will increase. This will in turn cause reaction times and error rates for words to increase.

A second manipulation will be performed during this experiment. Word decisions to high versus low concreteness words as a function of nonword context will be examined. This manipulation will be performed within-subjects which provides two benefits relative to studies that manipulate nonword context between participants. The first is obvious; it allows us to compare the responses of the same participants across the experimental conditions, thereby ruling out the possibility that differences in performance are due to some confound across the participant groups. The second is perhaps more interesting. By allowing participants to get used to one nonword context, then switching it, we can more closely watch the transition in responding as the participants adapt to the new nonword context. It is predicted that since subjects will take longer to perform a lexical decision during the pseudohomophone phase, more semantic processing will occur. Therefore we should find larger semantic effects during the pseudohomophone condition and the concreteness effect should become larger.

Method

Participants. Twenty-four undergraduates from the University of Toronto at Scarborough participated in Experiment 1 in exchange for either half a bonus credit towards their Introductory Psychology mark, or $5.00. All participants had normal or corrected to normal vision. The data from two participants was not included in the statistical analyses. In the first case, the participant's data was not used because they did not obtain an overall level of accuracy on the task exceeding 80%. The mean overall accuracy level of the remaining subjects was approximately 90% with a standard deviation of 3.5%. In the second case, the participant's data was discarded because their mean reaction times were in excess of 1700 ms, while the average for the remaining participants was 624 ms with a standard deviation of 87. These two participants were clearly not performing the task in a manner consistent with the majority of the participants. Therefore the total number of subjects included in the analysis is twenty-two.

Procedure. The study utilized a running lexical-decision task that appeared to the participant as a single stream of 400 trials. Participants began each trial by depressing two buttons with their two index fingers. One of the buttons was labelled “word”, the second labelled “nonword”. The trial then consisted of (a) a 250 ms blank field, (b) a 250 ms presentation of a fixation cross “+”, (c) a second 250 ms blank field, (d) either a word or nonword which was presented until the subject responded. Participants categorized the lexical status of each item by releasing the button corresponding to their decision. The hand used for the word response was distributed evenly and randomly across the participants. The stimulus disappeared as soon as one of the buttons was released. The next trial did not begin until both buttons were again depressed. Participants were informed of this and told that if they ever needed a rest, all they had to do was to not depress the response key until they were ready to continue.

The 400 trials were actually composed of four phases. The first phase consisted of 80 trials; 40 words and 40 scrambled nonwords. The purpose of this phase was to accustom the participants to the task and to the scrambled nonwords that would be presented in the first half of the experiment. The second phase consisted of 120 trials; 60 words and 60 scrambled nonwords. Half of these 60 words were high concreteness words, and half were low. This contrast allowed us to estimate the strength of semantic influences when participants were accustomed to the scrambled nonwords. The third phase was the Transition Phase, and it consisted of 80 trials; 40 words and 40 pseudohomophones. It is in this phase where we expected lexical-decision performance to markedly change as a result of the increased difficulty discriminating words from these more wordlike nonwords. Finally, the last phase consisted of 120 trials; 60 words and 60 pseudohomophones. Once again, half of the words were high in concreteness and half were low. This allowed us to assess the strength of semantic influences once the participants were accustomed to the pseudohomophone foils.

Stimuli. Experiment 1 required three subsets of stimuli. First, 60 high and 60 low concreteness words were needed for Phases 2 and 4. These stimuli were selected from the MRC Psycholinguistic Database (Coltheart, 1981) by performing and initial search for words five letters in length that have a Kucera and Francis (1967) frequency higher than 1 occurrence per million, and lower than 75 occurrences per million. The resulting words were then split into two groups based on concreteness ratings (Pavio, Yuille & Madigan, 1968). One group had ratings between 200 and 400, the other had ratings between 500 and 700. Finally, these two groups were further trimmed in a manner that equated them in terms of frequency and familiarity.

The result of this process was the two groups of words presented in Appendix A. The low concreteness stimuli have a mean frequency rating of 27 occurrences per million ($\underline{SE} = 3$), a mean familiarity value of 468 ($\underline{SE} = 12$), and a mean concreteness rating of 319 ($\underline{SE} = 5$). The high concreteness stimuli have a mean frequency rating of 26 occurrences per million ($\underline{SE} = 3$), a mean familiarity value of 483 ($\underline{SE} = 7$), and a mean concreteness rating of 581 ($\underline{SE} = 4$). Statistical analyses of the differences between the groups revealed no reliable differences on the frequency and familiarity dimensions, $\underline{t}(118)=0.82$ and 1.09 respectively, but a sizable difference on the concreteness dimension, $\underline{t}(118) = 40.13$, $p < .0001$. Thirty words from each group were randomly assigned to Phases 2 and 4 on a participant by participant basis.

The experiment required a further 80 words, 100 pseudohomophones, and 100 scrambled nonwords. The only constraints used to select the 80 additional words were that they be five letters in length and that they be different from the 120 high and low concreteness words. The pseudohomophones were an expanded set of the five letter pseudohomophones used by Joordens and Becker (1997). None of the pseudohomophones used in this study were homophonic with the word stimuli. Finally, the scrambled nonwords were created by taking the 100 pseudohomophones and re-arranging their letters to form non-pronounceable or, at least, extremely hard to pronounce nonwords. These additional stimuli are presented in Appendix A.

Apparatus. Testing was carried out on an IBM compatible 486 computer equipped with a Magnitronic 15 inch SVGA color monitor. Participants used a MEL response box to input their responses. The experiment was programmed in MEL (Micro-Experimental Laboratory) version 2.0. The screen background color was grey, and the stimuli were presented in white. The participants sat approximately 50 cm from the screen.

Results and Discussion

The data from this and subsequent experiments were analyzed in a similar fashion to the data in Simulations 1 to 3. The analysis focused on a comparison of Phase 2 and Phase 4. As was stated earlier for the simulation data, during Phase 1, subjects are learning the task, and during Phase 3 a transition has just occurred which would greatly effect reaction times and error rates. This will cause the data during Phase 3 to be very noisy. Therefore, to get the cleanest possible comparison of performance for each condition, Phase 2 and 4 are used for the analysis.

Phase 2 vs. Phase 4. Three questions are of relevance in this analysis. First, were the reaction times longer and error rates larger for nonword decisions to pseudohomophones than to scrambled nonwords? Second, were the correct reaction times to words higher when the nonwords were pseudohomophones than when they were scrambled nonwords? Finally, is there any evidence of larger semantic influences in the pseudohomophone context than in the scrambled nonword context?

As expected, correct nonword decisions were longer to pseudohomophones (M=717 ms) than to scrambled nonwords (M=550), $F(1, 21)=101.69$, $p<.0001$, $MSe = 3009.13$. In addition, the decisions to pseudohomophones were also more error prone ($\underline{M}$=18% errors) than were the decisions to scrambled nonwords ($\underline{M}$=4% errors), $\underline{F}(1, 21)=32.55$, $\underline{p}<.0001$, $\underline{MSe} = 65.38$. These results support the notion that pseudohomophones are harder to distinguish from words than are scrambled nonwords.

The remaining two issues were addressed using a 2 x 2 within-subjects analysis of variance performed on the word decisions with Concreteness (high versus low) and Nonword Type (scrambled nonword versus pseudohomophone) as the two factors. Separate analyses were conducted on the reaction time to make correct word decisions and error rate for those decisions. The means and standard deviations of these variables are presented in Table 4.

Insert Table 4 about here

The reaction time analysis revealed a significant main effect of Nonword Type, $\underline{F}(1,21) = 45.60$, $\underline{p} < .0001$, $\underline{MSe} = 4815.26$, with words being responded to slower in the pseudohomophone phase ($\underline{M}$ = 635 ms) than in the scrambled nonword phase ($\underline{M}$ = 535 ms). The main effect of concreteness was also significant, $\underline{F}(1,21) = 12.49$, $\underline{p} < .002$, $\underline{MSe} = 1417.51$, with shorter reaction times to high concreteness words ($\underline{M}$ = 571 ms) than to low concreteness words ($\underline{M}$ = 635 ms). Although the interaction between Nonword Type and Concreteness was in the expected direction with a larger concreteness effect in the pseudohomophone phase ($\underline{M}$ = 34 ms difference in reaction time) than in the scrambled nonword phase ($\underline{M}$ = 23 ms difference in reaction time), it was not statistically reliable, $\underline{F}(1,23) < 1$.

The analogous analysis on error rates produced very similar results. There was a significant main effect of Nonword Type, $F(1,21) = 16.76$, $p < .0006$, $MSe = 53.52$, with words being responded to less accurately in the pseudohomophone phase (M = 15% errors) than in the scrambled nonword phase (M = 8 % errors). The main effect of concreteness was also significant, $F(1,21) = 49.65$, $p < .0001$, $MSe = 24.77$, with a lower error rate for high concreteness words (M = 8 % errors) than for low concreteness words (M = 15 % errors). Finally, the interaction between Nonword Type and Concreteness was significant, $F(1,21) = 25.71$, $p < .0001$, $MSe = 13.54$, such that concreteness effects were larger in the pseudohomophone phase (M = 11 % difference in errors) than in the scrambled nonword phase (M = 3 % difference in errors).

Thus, the current results are in accord with previous studies that have manipulated the wordlikeness of nonwords. As the nonwords were made more wordlike (a) word responses became slower and more error prone, and (b) semantic influences on word responses, as assessed using a concreteness manipulation, generally increased in magnitude. The modulation of the concreteness effect was not as robust as we hoped, resulting in a significant interaction only in the error data. Nonetheless, the overall pattern of data observed in the current experiment provides additional evidence for the dynamic nature of lexical decisions.

These findings are consistent with the general story presented in the introduction. That is, in order to effectively respond in conditions where the nonwords are very wordlike (i.e., to reduce errors to nonwords), the participant must process both the words and nonwords more deeply. This deeper processing is reflected in both word reaction times, and in terms of concreteness effects on correct word responses. Moreover, however this change in responding occurs, and whatever processes underlie it, it appears to occur rapidly, and in response to high error rates.

Experiment 2: The Effect of Word Context on Nonword Responses

Experiment 2 has the same design as Simulation 2. We examined a question that runs parallel to that addressed in Experiment 1, but differs in a crucial manner. The question is; what happens when the character of the word stimuli are altered part way through a lexical-decision experiment? For example, what if the experiment originally consists of a discrimination between nonwords and low-frequency words then, without notifying the participant, the word stimuli change to high-frequency words? Will responses to nonwords be affected by this manipulation? According to the referent model, high frequency words will be easier to process. Therefore, reaction times and error rates will both decrease. Also, since word and nonword decision are intimately related, there will be an effect on nonword decisions where both reaction times and error rates will decrease.

Method

Participants. Twenty-four undergraduate students from the University of Toronto at Scarborough participated in the experiment either in exchange for course credit, or for payment of $5. All participants had normal or corrected to normal vision, and all performed the task with an overall accuracy level greater than 80%.

General. The overall procedure and apparatus used in the present experiment was identical to those used in Experiments 1. The primary modification was to the stimuli presented.

The low and high frequency words used in Phases 2 and 4 respectively were randomly sampled from sets of 102 low-frequency words and 102 high-frequency words. All words were five letters in length. The average Kucera and Francis (1967) frequency of these sets were 4.1 (SD = 2.42) occurrences per million and 75.9 (SD = 27.8) occurrences per million for the low and high-frequency sets respectively. As an additional control, the mean concreteness (Pavio et. al, 1968) of the low frequency words (M = 523.9) was matched to the mean concreteness of the high frequency words (M = 522.39), $t(202) = 0.17$. In addition, 80 filler words were required for Phases 1 and 3. These items were five-letter words selected simply to not overlap with the high and low frequency items described above.

The nonwords used in the current experiment were the scrambled nonwords and pseudohomophones used in Experiment 1. In contrast to Experiment 1, these items were mixed together within each block.

Results and Discussion

The mean reaction times and error rates acquired in Phases 2 and 4 of the current experiment are presented in Table 5. The reaction times and error rates were separately analyzed using 2 X 2 analyses of variance with lexical status (i.e., word versus nonword) and word type (high versus low frequency) as factors. These analyses were then followed up by planned comparisons that separately examined the affect of the word type variable on word and nonword decisions.

Insert Table 5 about here

The reaction time analysis revealed a significant main effect of lexical status, $F(1,21) = 4.62$, $p < .05$, $MSe = 6453.52$, with words being responded to faster overall (M = 655 ms) than nonwords (M = 691 ms). The main effect of word type was also significant, $F(1,21) = 20.66$, $p < .0002$, $MSe = 20888.53$, such that items in general were responded to faster in the high-frequency word condition (M = 606 ms) than in the low-frequency word condition (M = 739 ms). This effect of word type interacted with lexical status, $F(1,23) = 5.24$, $p < .04$, $MSe = 3196.27$, such that the effect was stronger for the words (mean difference of 160 ms), that for the nonwords (mean difference of 107 ms). The subsequent planned comparisons showed that this affect of word type was significant for both the words, $F(1,21) = 22.58$, $p < .0001$, $MSe =$ 309140.70, and for the nonwords as well, $F(1,21) =13.38$, $p < .002$, MSe = 139072.00.

The analogous analysis performed on error rates revealed a significant main effect of lexical status such that error rates were lower for words (M = 7%) than for nonwords (M = 14%), $F(1,21) = 13.32$, $p < .001$, $MSe = .0086$. There was also a main effect of word type such that errors were lower to items in the high-frequency word condition (M = 8%) than in the low-frequency word condition (M = 13%), $F(1,21) = 24.68$, $p < .0001$, $MSe = .0020$. The two variables did not interact, $F < 1$. Planned comparisons to examine these effects more closely revealed that the effect of word type was significant for both the words, $F(1,21) = 10.54$, $p < .004$, $MSe = .0248$, and for the nonwords, $F(1,21) = 17.22$, $p < .0004$, $MSe = .0257$. Thus, both words and nonwords had higher error rates in the low than in the high-frequency condition.

The most critical finding of this experiment is that nonword decisions become faster when the words are changed from low to high frequency. This result is exactly what was predicted on the basis of the decision dynamics of the referent model, and is not at all consistent with those of the multiple read-out model. As such, Experiment 2 provides support for the decision dynamics of the referent model.

Experiments 3A & 3B: Varying the Word to Nonword Ratio

Experiment 1 replicated the finding that variations of the nonword context appear to affect the depth to which words are processed prior to the emission of a response. Experiment 2 provided the novel finding that variations of word characteristics also affect responses to nonwords and, more critically, do so in the manner predicted by the referent model. In combination then, these experiments demonstrate the need for explicit models of the decision dynamics underlying the lexical-decision performance, and provide preliminary support for one such explicit model, the referent model.

The effects found in both Simulation 1 and 2 were consistent with the idea that there is an intimate relationship between word and nonword decision. By manipulating one stimulus type, we found predictable effects for the other stimulus type. As was stated earlier in Chapter 2, the reason for this intimate relationship is the efficiency bias that is used when performing lexical decisions. This bias is set in such a way that decisions to both words and nonwords are made as quickly and as accurately as possible. The purpose of this simulation was to manipulate this bias in a novel fashion and make further predictions about performance.

Simulation 3 was divided into two separate simulations. In Simulation 3A the network was presented with a ratio of words to nonwords of 1:1 for the first half of the simulation followed by a ratio of 1:2. In Simulation 3B the network was presented with a ratio of words to nonwords of 1:1 for the first half followed by a ratio of 2:1. The predictions for these simulations are as follows. If we compare a context containing more words than nonwords to a baseline of an even number of words and nonwords, we would expect the criterion to shift in such a way that the referent would become closer to the word boundary which should lead to; (a) faster word responses, (b) slower nonwords responses, (c) less word errors, and (d) more nonword errors. In contrast, if we compare a situation with fewer words than nonwords, we would expect the criterion to shift so that the referent would become closer to the nonword boundary which should lead to; (a) slower word responses, (b) faster nonword responses, (c) more word errors, and (d) fewer nonword errors.

Method

Participants. A total of 62 undergraduates from the University of Toronto at Scarborough participated in these experiments; 34 in Experiment 3A and 28 in Experiment 3B. All subjects had normal or corrected to normal vision, and all performed the lexical decision task with an accuracy level greater than 80%.

Procedure. The procedure involved 3 blocks of a running lexical-decision task. The first block consisted of 24 practice trials. The second and third blocks contained the experimental trials. One of the blocks was 72 trials, the other was 108 trials. In all cases, a trial consisted of the follow sequence of events; (a) a fixation cross presented for 250 ms, (b) a 250 ms blank field, (c) a letter string presented until the participant categorized it as a word or nonword, (d) another 250 ms blank field.

The practice block consisted of 12 trials on which a word was presented, and 12 on which a nonword was presented. Similarly, the experimental block containing 72 trials consisted of half word and half nonword trials. In contrast, the experimental block containing 108 trials consisted of 72 trials of one stimulus type, and 36 of the other. In Experiment 3A, it was nonwords that dominated, whereas in Experiment 3B the words dominated. The two experimental blocks were counterbalanced within each experiment such that half of the subjects received the 72 trial experimental block first, and the 108 trial experimental block second. The other half of the subjects received the reversed ordering.

Stimuli and Apparatus. The apparatus, stimulus-size, and response characteristics were identical to those described in Experiment 1. The stimuli consisted of 120 five-letter words, and 120 five-letter pronounceable nonwords. Only 84 of the words were used in Experiment 3A and 84 of the nonwords in Experiment 3B. In both cases, the particular items used were selected randomly on a participant by participant basis from the respective 120 item pool.

Results and Discussion

The results from Experiments 3A and 3B are depicted in Table 6. This section will present the statistical analysis separately for each experiment and, within each experiment, separately for the reaction times and error rates. In each case, the analyses will consist of a 2 X 2 analysis of variance with lexicality (word vs.

nonword) and word:nonword ratio as the factors, followed by planned comparisons examining the effect of word:nonword ratio separately for words and nonwords.

Insert Table 6 about here

Experiment 3A. Experiment 3A involves a contrast of word and nonword responses across conditions where there were either an equal number of words and nonwords presented or twice as many nonwords than words presented. The referent model predicts that responses to nonwords will be faster and less error prone, and responses to words will be slower and more error prone, in the condition where there are more nonwords than words relative to the condition where they are equally represented.

The analysis of reaction times revealed an overall main effect of lexicality, $F(1,33) = 8.71$, $p < .006$, $MSe = 5912.64$, such that words were responded to faster overall ($M = 680$ ms) than were nonwords ($M = 719$ ms). There was also a main effect of word:nonword ratio, $F(1,33) = 7.43$, $p < .011$, $MSe = 4133.30$, such that responses were faster overall in the 1:2 condition ($M = 684$ ms) than in the 1:1 condition ($M = 714$ ms). Finally, the interaction between these factors was also significant, $F(1,33) = 28.76$, $p < .001$, $MSe = 1371.16$, suggesting that the word-to-nonword ratio manipulation had a different effect on nonwords than it did on words.

Planned comparisons conducted revealed that, as predicted by the referent model, nonwords were responded to significantly faster (i.e., 64 ms faster) in the condition where they were the dominant stimulus class, $F(1,33) = 22.95$, $p < .001$, $MSe = 3043.65$. However, the model's prediction concerning response times to words was not borne out. Although words were responded to slower in the condition where nonwords were dominant (i.e., 4 ms slower), this difference was not near being reliable, $\underline{F}(1,33) < 1$.

The analogous analysis performed on the error rate data revealed an overall main effect of lexicality, $\underline{F}(1,33) = 43.05$, $\underline{p} < .003$, $\underline{MSe} = 0.2045$, indicating higher overall errors to words ($\underline{M} = 10.5\%$) than to nonwords ($\underline{M} = 5.4\%$). Although the main effect of nonword ratio was not significant, $\underline{F}(1,33) < 1$, the lexicality effect did interact with nonword ratio, $\underline{F}(1,33) = 8.96$, $\underline{p} < .006$, $\underline{MSe} = 0.1847$, once again indicating that the word-to-nonword ratio manipulation had differential effects on word and nonword responses.

Further examination of this interaction via planned comparisons revealed that the 1.6% reduction in errors to nonwords in the condition where they were the dominant stimulus class approached conventional levels of significance, $\underline{F}(1,33) = 3.82$, $\underline{p} < .059$, $\underline{MSe} = 0.1040$. Moreover, the 2.8% increase in errors to words in the condition where nonwords were dominant was reliable at conventional levels of significance, $\underline{F}(1,33) = 5.21$, $p < .029$, $\underline{MSe} = 0.2708$. Thus, as predicted by the referent model, responses to nonwords became less error prone, and responses to words more error prone, in the condition where nonwords were the dominant stimulus class.

Experiment 3B. Experiment 3B involves a contrast of word and nonword responses across conditions where there were either an equal number of words and nonwords presented or twice as many words than nonwords were presented. The referent model predicts that responses to nonwords will be slower and more error prone, and responses to words will be faster and less error prone, in the condition where there are more words than nonwords relative to the condition where they are equally represented.

The analysis of reaction times revealed an overall main effect of lexicality, $F(1,27) = 67.63$, $p < .001$, $MSe = 6713.42$, such that words were responded to faster overall ($M = 679$ ms) than were nonwords ($M = 806$ ms). Although there was no main effect of word-to-nonword ratio, $F(1,27) < 1$, the interaction between lexicality and word-to-nonword ratio was significant, $F(1,27) = 18.23$, $p < .001$, $MSe = 1516.16$.

Planned comparisons focused on this interaction revealed that the 25 ms slowdown in nonword responding in the condition where the words were dominant was not statistically significant, $F(1,27) = 1.44$, $p < .240$. However, the 38 ms speed-up of word responses in the condition where they were the dominant stimulus class was significant, $F(1,27) = 4.11$, $p < .053$, $MSe = 4872.57$. Once again, the predictions of the referent model with respect to the reaction times are borne out for the stimulus class that becomes dominant, but not, quantitatively at least, for the one that does not.

The analogous analysis performed on the error rate data revealed an overall main effect of lexicality, $F(1,27) = 10.71$, $p < .003$, $MSe = 0.2948$, indicating higher overall errors to nonwords ($M = 9.3\%$) than to words ($M = 5.9\%$). Although the main effect of nonword ratio was not significant, $F(1,27) < 1$, the lexicality effect did interact with the word-to-nonword ratio, $F(1,27) = 5.19$, $p < .031$, $MSe = 0.1791$, again indicating that the word-to-nonword ratio manipulation had differential effects on word and nonword responses.

Further examination of this interaction revealed that the 1.8% increase in errors to nonwords in the condition where words were the dominant stimulus-class approached conventional levels of significance, $\underline{F}(1,27) = 3.82$, $\underline{p} < .081$, $\underline{MSe} = 0.1361$. In addition, the 1.8% reduction in errors to words in the condition where they were dominant also approached conventional levels of significance, $\underline{F}(1,27) = 2.76$, $p < .108$, $\underline{MSe} = 0.1757$. Thus, as predicted by the referent model, responses to nonwords became more error prone, and responses to words less error prone, in the condition where words were the dominant stimulus-class.

Summary

The data from Experiments 3A and 3B can be summarized as follows. Across two experiments, the referent model made eight predictions concerning the effect of the word-to-nonword ratio manipulation on word and nonword responses. All eight predictions went in the direction predicted by the theory, although not always to a statistically significant extent. The probability of all 8 effects being in the predicted direction if word-to-nonword ratio were having no effect is .0039. Thus, the overall pattern strongly supports the contention that manipulations of word-to-nonword have a systematic effect on word and nonword decisions.

The data from all three experiments is consistent with the data from the simulations that were presented earlier, and more importantly, with the predictions of the referent model. So we now have strong evidence from these simulations and experiments that supports the referent model. In other words, the referent model is able to account for these findings. The question now is, does the referent model account for these findings better than the multiple read-out model? This question will be addressed in Chapter 5. We will go over the results from each of the simulation/experiment pairs and determine which of the two models is best able to account for the findings. It will become apparent that the referent model is the better of the two descriptions of how humans are performing a lexical decision. Chapter 5 will end with some a general discussion and future directions of research using the referent model.

CHAPTER 5

In Chapter 1, some recent findings from experiments that used the lexical decision task were discussed. These experiments provided a background that allowed us to discuss, in Chapter 2, two competing models of how a lexical decision is performed. These two models are the multiple read-out model and the referent model. Then, in Chapter 3, computer simulations were presented that provided us with a more concrete understanding of the mechanisms involved when performing a lexical decision in the context of the referent model. The findings were replicated with human experiments in Chapter 4. The purpose of this chapter is to make a comparison between the referent model and the multiple read-out model. We will look at the predictions that each of the models make for each of the experiments that were presented earlier. It will become evident that the referent model predictions provide the better fit to the data that has been provided. The chapter will conclude with a general discussion and, finally, future research involving the referent model theory and simulations will be discussed.

Combining Simulations and Experiments

Simulation 1 and Experiment 1

[4]The purpose of Study 1 was to show that, by manipulating the type of nonwords that are presented during a lexical decision task, we are able to make specific predictions about performance for word decisions. This type of manipulation is most common in the literature. Both the multiple read-out model and the referent model predict findings that are produced by this study. As nonwords were switched from

scrambled to pseudohomophones, reaction times and error rates for nonwords increased. This also caused the reaction times and error rates for words to increase. Study 1 sets the stage for the theories while also giving us a starting point to highlight how Study 2 differs from previous ones, and why this difference is theoretically relevant.

Simulation 2 Experiment 2

The purpose of Study 2 was to determine if manipulating the type of words that were being presented would produce a predictable effect on nonword performance. A more common manipulation, which was presented in Study 1, involves manipulating the type of nonwords that are presented. The question here is, how will switching the type of words that are being presented from high frequency to low frequency effect word decisions? The answer to this question is very relevant to discriminating between the models as highlighted by the contrasting predictions outlined below.

Multiple Read-Out Model

According to the multiple read-out model, responses to words are a function only of the setting of the word thresholds. Nonword responses are emitted when an item's gain of activation asymptotes prior to it reaching one of the word thresholds. If such a model were attempting to distinguish low-frequency words from nonwords, the most efficient setting of the word thresholds would be at a level just above the activation value that most nonwords reach. Such a setting would allow for very fast word decisions, and very few errors.

---

[4] From this point forward, when a reference to both a particular simulation and its corresponding experiment is made, they will be referred to as a Study (i.e. Study 1 referes to both Simulation 1 and Experiment 1).

What should happen when the words are changed to high-frequency words? Obviously high-frequency words should hit the word threshold faster and with greater accuracy than their low-frequency counterparts. Thus, we would expect to see a standard frequency effect emerge across the first and second half of the experiment.

More critical, however, are the predictions concerning the effects of this manipulation on nonword responses. The most obvious prediction is, in fact, no effect at all. The criterion set up for low frequency words should work well for high frequency words as well. Thus, nonwords should continue to be categorized with the same speed and accuracy as they were when the words were low frequency.

It may be possible to modify the multiple readout model in the following way. Perhaps the word criterion might be raised slightly given that high frequency words could easily and accurately reach a higher criterion, and a higher criterion would result in fewer errors to nonwords. This raising of the word criterion could indirectly affect the reaction time to nonwords if it were assumed that the nonwords that were previously errors (i.e., the ones whose activation level was high) would result in slower correct reaction times than average, if the criterion were raised. Specifically then, one could predict slower nonword reaction times and lower nonword error rates if the word criterion were raised.

The Referent Model

The predictions of the referent model are straightforward as were observed from the simulations in Chapter 3. High frequency words are assumed to have a higher drift rate than low frequency words. Thus, when the high frequency words are introduced, they should initially result in fast and highly accurate word responses.

However, as a function of the efficiency principle, the boundaries should migrate upwards such that the nonword boundary were moved closer to the referent, and the word boundary further away. Such a migration would show through as faster and less error prone nonword responses. Thus, in contrast to the prediction(s) of the multiple read-out model, the referent model predicts faster nonword responses after the high-frequency words are introduced.

Conclusions

As was stated earlier, the most critical finding from Study 2 is that nonword decisions become faster when the words are changed from low to high frequency. The multiple read-out model predicts no effect while the referent model predicts the effect that was found. This result is exactly what was predicted on the basis of the decision dynamics of the referent model, and is not at all consistent with those of the multiple read-out model. As such, Study 2 provides support for the decision dynamics of the referent model.

Simulation 3 and Experiment 3

The purpose of Study 3 was to determine if, by either doubling the number of words or nonwords that were being presented, we could effect performance for the opposite stimulus type. For example, when switching from a condition that contained an equal number of words and nonwords to a condition of twice as many words as nonwords, we found that word decisions became very fast and accurate while nonword decisions became slower and less accurate. The opposite was true when the number of nonwords was doubled; nonword decisions became faster and more accurate while word decisions became slower and less accurate.

Theories that assume separate decision-process underlying word versus nonword decisions, such as the multiple read-out model, have no reason to predict any effect of word-to-nonword ratio. To do so they would need to include a mechanism for bias that is not currently part of such models. Such a mechanism is an integral part of the referent model, not only in accounting for the current data, but also in accounting for the dynamic nature of lexical-decision performance in general. Given this, each case of a significant result in the planned comparisons of Study 3A and 3B, in addition to the overall pattern described above, are problematic for models assuming separate decision processes for word versus nonword decisions.

But, do the data strongly support the outlined predictions of the referent model? They do fit the general prediction that a manipulation of word-to-nonword ratio will effect both word and nonword decisions. However, there is an aspect to the current data that stands out and was not predicted in the introduction of this section. Specifically, as we move from a condition where words and nonwords are equally presented to one where one stimulus class is dominant, responses to the non-dominant stimulus class are not affected to the same extent as responses to the dominant stimulus class. For example, when comparing the 1:1 condition to the condition where there we twice as many nonwords as words (i.e., Experiment 3A), responses to nonwords sped up dramatically, but responses to words slowed only slightly. This asymmetry is also present in the Experiment 3B data, although not to the same extent.

In fact, it is not surprising that the predictions of the referent model as they are presented do not fit perfectly. The predictions were based on a simplified version of the model in which the referent function is assumed to remain static and only the decision dynamics (i.e., the word and nonword criteria) were assumed to vary. However, if the referent function does reflect some sort of "average stimulus", then it would also be affected by the manipulation of word-to-nonword ratio.

Consider the situation in which nonwords are dominant (e.g., the 1:2 condition of Experiment 3A). In this situation, the referent function would be lowered by the inclusion of a high proportion of nonwords. Thus, it would lie closer to the nonword function. This should make it take more time for nonwords to deviate from it, and less time for words to deviate from it. However, that is not the response pattern that was observed presumably because the shift in the word and nonword criteria that we emphasized counteracted it.

Thus, when the referent model is considered in detail, its predictions are actually more complex than was outlined in the introduction to these experiments. The decision criteria are the main things that are assumed to vary, but they do so in the foreground of other changes. Given this, it is not surprising that there are complexities to the data that go beyond the general predictions following from variations of word and nonword criteria. The fit of the data to the predictions of the referent model is actually quite good. The fit is not perfect but the model does account for most of the complexities of the data when the predictions of the model are considered carefully.

General Discussion

The goal in writing this dissertation was to introduce a new model of lexical decision; the referent model. The referent model combines the processing characteristics of an attractor model of word recognition, with the decision dynamics of a random-walk process. It is argued that this type of model could account for many of the more challenging lexical-decision results, although we also pointed out that an existing model, the multiple read-out model, could also account for much of the same data.

The experiments we reported were aimed at again demonstrating some of these challenging lexical-decision results, and then attempting to determine whether the decision dynamics of the referent model or those of the multiple read-out model appear to provide the best account of the data. The results of these experiments favor the account provided by the referent model.

Re-Uniting Decision Predictions with Attractor Network Predictions

One of the merits of the referent model is that it provides an explicit description of two components of lexical decision; an overlearned word-recognition process and an overlaid decision process. In so doing, it permits one to discuss how variables affecting either of these simultaneously operating processes can impact the observed lexical-decision performance.

However, the studies that provide the strongest support for the referent model (i.e., Studies 2, 3A, & 3B) focused primarily on the variables assumed to affect the decision component. If the referent model is truly correct, then variations of the decision process should also affect the "depth" to which a stimulus is processed, and this should show through in the lexical-decision performance. For example, consider Experiment 3B in which word responses were sped up when words were the dominant stimulus. In this situation, those faster responses should be associated with less processing in the attractor network, and hence less semantic influences given that semantic influences are assumed to come about latest in stimulus processing.

More generally stated, one could predict that as the proportion of words relative to nonwords increased, word responses should speed up, and semantic influences should decrease. It turns out that this issue has already been examined by Neely, et. al (1989). They found that the lower the nonword ratio, the less the semantic priming. Thus, it does indeed appear that variations of word-to-nonword ratio do affect the depth to which words are processed prior to a response.

In fact, there is a further extent to which the Neely et. al (1989) results support the referent model, indirectly at least. Recall that Joordens and Becker (1997) showed smaller semantic priming effects when the nonwords are made less wordlike, a result that runs parallel to Neely et. al's finding of less semantic priming when words were dominant. Given that the referent model attributes both of these effects to the same change in processing, the parallel findings are not surprising at all. Any variable the biases participants towards making a word response should result in smaller semantic influences.

This point may seem obvious but it is as critical to supporting the referent model as a whole, and not just supporting its decision dynamics. Recall the implication that the multiple read-out model might be able to account for our Experiment 3A and 3B datasets if some sort of bias mechanism were added to it. However, now the model would be accounting for the effects of nonword context via one mechanism (i.e., changes in the criteria level for word responses) and it would be accounting for the effects of word-to-nonword context via a different mechanism (i.e., changes due to the bias mechanism). This sort of view does not naturally predict that the two manipulations would produce parallel results with respect to semantic effects, whereas the referent model does.

There is an additional result that adds further strength to the validity of the referent model. The most critical finding in the Joordens and Becker (1997) paper was the first ever demonstration of long-term semantic priming in the lexical-decision task. Joordens and Becker argued that one reason nobody had found long-term semantic priming prior to this was because people typically set up the lexical-decision task in a manner that did not encourage much semantic processing of the items. When more semantic processing was encouraged, the long-term effect emerged and typical short-term priming effects increased in magnitude.

One of the ways that Joordens and Becker (1997) encouraged semantic processing (i.e., Experiment 3) was to create conditions which, based on the referent model, should lead to the longest word reaction times and, hence, the strongest semantic influences. This condition involved three manipulations designed to shift the word boundary as far as possible from the referent level. These conditions included (a) using a large number of nonwords relative to words, (b) pre-exposing the nonwords to make them more familiar (i.e., to decrease their drift rate towards the nonword boundary), (c) including filler words that were very high frequency. The inclusion of these filler items increases the average drift rate for words, allowing the word criteria to move further away from the referent function without increasing the "overall" response time to words very much. However, the drift rate for the critical words is still low, so they still take a long time to hit the word boundary.

Once again, these manipulations had the exact result that would be predicted on the basis of the referent model. First, word reaction times increased dramatically, up to approximately 900 ms. Second, semantic effects were very large. In fact, a 63 ms semantic priming effect was observed across a lag of eight items between the prime and target. These results would not be predicted on the basis of any other lexical decision model.

Thus, the Neely et. al (1989) and Joordens and Becker (1997) results provide further evidence that bears directly on the full referent model. Variables that affect the decision processes lead to the results expected if the decision module was overlaid on a word-recognition process that worked in a manner similar to the referent model described here.

Extensions to Other Tasks

Once again, one of the attractive features of the referent model is the coupling of a decision process with a more basic word-recognition process. It is not assumed that any human has a "lexical decision" process, pre-existent in their minds prior to entering the lab. Rather, it is assumed that the brain has a word-recognition process, and a basic decision process that can be used to support binary decisions (i.e., yes/no, left/right, like/dislike). When presented with the lexical decision task, these two sets of processes are coupled in an attempt to perform the task efficiently. This is accomplished by taking an ongoing product of the word-recognition process - harmony - and feeding it into their binary-decision module. Performing the task in this manner would allow one to emit a decision prior to complete processing of the stimulus on each trial, while still retaining a reasonable degree of accuracy.

One obvious question then is what about other tasks? In some cases, the referent model as described simply would not apply. For example, consider the naming task in which participants are asked to verbally produce the phonology associated with various orthographies (typically words). In this task, no binary decisions are being made. Instead, a complete phonological code must be generated. In a task like this, the attractor model may simply generate the phonological code without the aid of other cognitive processes (e.g., Masson, 1995).

What about tasks like recognition memory? In fact, one of the first extensions of the random-walk decision process from the realm of psychophysics was to provide an explanation of old/new recognition (i.e., Ratcliff, 1978). Is there any room for the word-recognition component of the referent model in such an account? In fact, when Ratcliff described how the random-walk dynamics could be used to account for recognition memory data, he was not explicit about the process that was driving the random walk. Perhaps the referent model could offer a possibility in this respect.

At the very least, by being explicit about the linkage between basic word-recognition processes, and other processes that may draw on them to support performance on some task, the referent model invites investigators to consider their tasks of preference in more detail. It is quite possible that the philosophy embodied in the referent model, which was obviously borrowed from the multiple read-out model, could be used to gain a better understanding of a variety of tasks other than lexical decision. Once we have a better understanding of the tasks we are using, we are then in a much stronger position to understand the phenomena that present themselves via performance on these tasks.

Future Research

We have a new model of lexical decision and a computer network that can be used to perform simulations. How can we take this knowledge and use it as a tool for gaining a better understanding of how humans process information or how memory is structured? We need to focus on issues that involve the lexical decision task and use this model to help resolve these issues.

One issue that has come up in the literature is the fact that, if human memory is distributed, why can't computer simulations that use a distributed representation of memory simulate the ambiguity effect? Ambiguous words are words that have more than one meaning. For example, the word BAT could mean "a flying mouse", or "a wooden stick that is used to hit a ball". When subjects are presented with ambiguous words during a lexical decision task, they make faster and more accurate decisions to ambiguous words than to unambiguous words. This seems like a pretty straightforward effect to simulate but distributed networks have struggled with it.

For example, Joordens and Besner (1994) tested whether a distributed model of memory (Masson, 1991) could produce the ambiguity advantage that is seen in the lexical decision task. They found that the network produced an ambiguity advantage for items that settled into a correct state. A problem with the network was that less than 50% of the ambiguous items settled into a correct state. This led them to conclude that this type of network that used a hopfield learning algorithm (Hopfield, 1982; Hopfield & Tank, 1986) may not be adequate enough to deal with the ambiguity effect. However, in later commentaries (Masson & Borowsky, 1995; Rueckl, 1995; Besner & Joordens, 1995) it was concluded that it may be possible for lexical decisions to be made prior to the network settling into a state that was consistent with a previously learned item. This decision could be made regardless of whether or not a blend state would eventually occur.

If it is the case that lexical decisions can be made regardless of a blend state, then we need to find evidence that this may be occurring when humans process ambiguous items. Piercey and Joordens (2000) compared the ambiguity advantage in lexical decision to the ambiguity disadvantage found in reading tasks. They performed an experiment in which subjects were presented with items to which they first made a lexical decision and then made a relatedness decision. They found an ambiguity advantage during the lexical decision component of the experiment and an ambiguity disadvantage during the relatedness component. They concluded that a lexical decision could be made based on early processing and that a blend state would produce an advantage for lexical decision but a disadvantage for the relatedness decision. The lexical decision could be made based on a measurement of harmony between the current representation and previously learned items whether or not a blend state was eventually produced. However, when the relatedness decision was made the subject would have to leave the blend state so that further semantic processing could occur. Therefore, an ambiguity advantage would be found early on followed by an ambiguity disadvantage.

The assumption that decisions are made based on early processing is not new. Meyer and Schvaneveldt (1971) concluded that lexical decisions could be made early on before semantic information was accessed. This early processing conclusion is similar to the conclusion that Piercey and Joordens (2000) make. The difference being that Piercey and Joordens believe that some semantic processing does occur when a decision is made. We know that this is probably the case from some previous studies that were presented earlier (see, Neely et. al, 1989; Joordens & Becker, 1997)

If these blend states do in fact occur when humans process ambiguous words, then the results from these simulations may not be as troublesome as we first thought. What we need to do is perform a simulation where a lexical decision is made early on, before a blend state occurs. The referent model provides us with a decision mechanism that utilizes early processing in order to perform a lexical decision. Using this model we may be able to produce the ambiguity effect found during a lexical decision. They network makes lexical decisions based on a comparison of the current level of harmony to the referent. This decision process would occur regardless of whether or not a blend occurred. Therefore, the referent model may provide us with an example of a distributed model of memory that is able to produce the ambiguity effect.

Final Note

The simulations presented in this dissertation are used as a tool that provides us with a better understanding of how a lexical decision is performed. Specifically, they allow us to test our assumptions and develop a more concrete description of the decision mechanisms that are used during this decision process. The purpose of these simulations is not to produce data that provides us with an exact fit to the human data. If this were the intent of the simulations, we could more easily fit data using a mathematical description or equation.

These simulations are powerful tools that allow us to gain some insight into how humans perform a specific task. In this case the task is the lexical decision. In fact, the findings from a simulation are much more interesting when they surprise us. When they produce data that we did not expect, it makes us think about the assumptions that we are basing our theory on. If the network produces surprising results, we may need to modify our theory and test the new theory on human subjects. This is turn may lead to a new finding about how humans process information. When this occurs we have harnessed the true power of computer simulations.

REFERENCES

Balota, D. A., & Chumbley, J. I. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. Journal of Experimental Psychology: Human Perception and Performance, 10, 340-357.

Becker, S., Moscovitch, M., Behrmann, M. & Joordens, S. (1997). Long-term semantic priming: A computational account and empirical evidence. Journal of Experimental Psychology: Learning, Memory, and Cognition, 23, 1059-1082.

Besner, D. & Davelaar, E. (1983). Pseudohomofoan effects in visual word recognition: Evidence for phonological processing. Canadian Journal of Psychology, 37, 300-305.

Besner, D., & Joordens, S. (1995). Wrestling with ambiguity - further reflections: Reply to Masson and Borowsky (1995) and Rueckl (1995). Journal of Experimental Psychology: Learning, Memory, and Cognition, 21, 515-519.

Borowsky, R., & Masson, E. J. (1996). Semantic ambiguity effects in word identification. Journal of Experimental Psychology: Learning, Memory, and Cognition. 22, 63-85.

Collins, A. M., & Qullian, M. R. (1969). Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 8, 240-248.

Coltheart, M. (1981). The MRC Psycholinguistic Database. Quarterly Journal of Experimental Psychology, 33A, 497-505.

Coltheart, M., Davelaar, E., Jonasson, J. T., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), Attention and Performance VI (pp. 535-555). Hillsdale, NJ: Erlbaum.

Craik, F. I. M., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. Journal of Verbal Learning and Verbal Behavior, 11, 671-684.

Dawson, M.R.W. (1988). Fitting the Ex-Gaussian equation to reaction time distributions. Behavior Research Methods, Instruments, & Computers, 20, 54-57.

Dawson, M.R.W. (1998). Understanding cognitive science. Oxford, UK: Blackwell Publishers.

Dawson, M.R.W., & Piercey, C.D. (2001). On the subsymbolic nature of a PDP architecture that uses a nonmonotonic activation function. Minds And Machines, 11, 197-218.

Dennis, I., Besner, D., & Davelaar E. (1985). Phonology in visual word recognition: Their is more two this than meets the I. In D. Besner, T. G. Waller, & E. M. McKinnon (Eds.), Reading Research: Advances in Theory and Practice (Vol 5, pp. 170-195). Toronto, ON: Academic Press

Grainger, J. & Jacobs, A. M. (1996). Orthographic processing in visual word recognition: A multiple read-out model. Psychological Review, 103, 518-565.

Hebb, D. O. (1949). The organization of behavior. New York: Wiley.

Hinton, G. E. & Shallice, T. (1991). Lesioning an attractor network: Invesitgations of acquired dyslexia. Psychological Review, 98, 74-95.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proceeds of the National Academy of Science, 79, 2554-2558.

Hopfield, J. J., & Tank, D. W. (1986). Computing with neural circuits: A model. Science, 233, 625-633.

James, C. T. (1975). The role of semantic information in lexical decisions. Journal of Experimental Psychology: Human Perception and Performance, 1, 130-136.

Joordens, S., & Becker, S. (1997). The long and short of semantic priming effects in lexical decision. Journal of Experimental Psychology: Learning, Memory, and Cognition, 23, 1083-1105.

Joordens, S. & Besner, D. (1994). When banking on meaning is not (yet) money in the bank: Explorations in connectionist modeling. Journal of Experimental Psychology: Learning, Memory, and Cognition, 20, 1051-1062.

Kucera, H., & Francis, W. N. (1967). Computational analysis of present-day American English. Providence: Brown University Press.

Link, S. W. & Heath, R. A. (1975). A sequential theory of psychological discrimination. Psychometrika, 40, 77-105.

Lupker, S. J., Brown, P., & Colombo, L. (1997). Strategic control in a naming task: Changing routes or changing deadlines? Journal of Experimental Psychology: Learning, Memory, & Cognition, 23(3), 570-590.

Masson, M. E. J. (1991). A distributed memory model of context effects in word identification. In D. Besner & G. W. Humphreys (Eds.), Basic processes in reading: Visual word recognition (pp. 233-263). Hillsdale, NJ: Erlbaum.

Masson, M. E. J. (1995). A distributed memory model of semantic priming. Journal of Experimental Psychology: Learning, Memory, and Cognition, 21, 3-23.

Masson, M. E. J., & Borowsky, R. (1995). Unsettling questions about semantic ambiguity in connectionist models: Comment on Joordens and Besner (1994). Journal of Experimental Psychology: Learning, Memory, and Cognition, 21, 509-514.

McCann, R. S., Besner, D., & Davelaar, E. (1988). Word recognition and identification: Do word frequency effects reflect lexical access? Journal of Experimental Psychology: Human Perception and Performance, 14, 693-706

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. Journal of Experimental Psychology, 90, 227-234.

Neely, J. H., Keefe, D. E., & Ross, K. L. (1989). Semantic priming in the lexical decision task: Roles prime-generated expectancies and retrospective semantic matching. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 1003-1019.

Pavio, A., Yuille, J. C., & Madigan, S. A. (1968). Concreteness, imagery and meaningfulness values for 925 words. Journal of Experimental Psychology Monograph Supplement, 76, 3 (part 2).

Piercey, C. D. & Joordens, S. (2000). Turning an advantage into a disadvantage: Ambiguity effects in lexical decision versus reading tasks. Memory & Cognition, 28(4), 657-666.

Plaut, D. & Shallice, T. (1991). Deep dyslexia: A case study of connectionist neuropsychology. Technical Report CRG-TR-91-3, Department of Computer Science, University of Toronto.

Ratcliff, R. (1978). A theory of memory retrieval. Psychological Review, 85, 59-108.

Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. Psychological Review, 106(2), 261-300.

Rueckl, J. G. (1995). Ambiguity and connectionist networks: Still settling into a solution - comment on Joordens and Besner (1994). Journal of Experimental Psychology: Learning, Memory, and Cognition, 21, 509-514.

Rumelhart, D. E., McClelland, J. L. & the PDP Group. (1986). Parallel Distributed Processing, vol. 1. Cambridge, Mass.: MIT Press.

Sharkey, A. J. C., & Sharkey, N. E. (1992). Weak contextual constraints in text and word priming. Journal of Memory and Language, 31, 543-572.

Smolensky, P. (1986). Foundations of harmony theory: Cognitive dynamic systems and the subsymbolic theory of information processing. In D. E. Rumelhart, J. L. McClelland, and the PDP research group (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition, Volume I. Cambridge, MA: Bradford Books

Stone, G. O., & Van Orden, G. C. (1993). Strategic control of processing in word recognition. Journal of Experimental Psychology: Human Perception and Performance, 19, 744-774.

Table 1

Means and standard deviations for the reaction times (ms) and error rates (%) across Phases 2 and 4 of Simulation 1

	Phase 2 (Scrambled Nonwords)		Phase 4 (Pseudohomophones)	
	M	SD	M	SD
Reaction Time				
Words	173	10	353	36
Nonwords	103	3.4	83	3.7
Error Rates				
Words	18	1.3	27	1.4
Nonwords	16	1.5	26	1.6

Table 2

Means and standard deviations for the reaction times (ms) and error rates (%) across Phases 2 and 4 of Simulation 2

	Phase 2 (Low Frequency Words)		Phase 4 (High Frequency Words)	
	M	SD	M	SD
Reaction Time				
Words	167	10.3	42	2.7
Nonwords	107	5.2	65	3.9
Error Rates				
Words	16	1.7	14	1.2
Nonwords	13	1.7	13	1.3

Table 3

Means (with standard deviations in the brackets) for the reaction times (ms) and error rates (%) in Experiments 3A and 3B

	Word:Nonword Ratio 2:1	1:1	1:2
Experiment 3A			
Words			
Reaction Time		170 (16.4)	272 (17.4)
Error Rates		16 (2.0)	29 (2.3)
Nonwords			
Reaction Time		111 (5.4)	64 (2.9)
Error Rates		14 (1.0)	2 (0.4)
Experiment 3B			
Words			
Reaction Time	111 (9.2)	163 (4.5)	
Error Rates	4.0 (1.1)	17 (1.0)	
Nonwords			
Reaction Time	211 (15.8)	111 (9.0)	
Error Rates	63 (4.0)	14 (2.1)	

Table 4

Means and standard deviations for the reaction times (ms) and error rates (%) across Phases 2 and 4 of Experiment 1

	Phase 2 (Scrambled Nonwords)		Phase 4 (Pseudohomophones)	
	M	SD	M	SD
<u>Reaction Time</u>				
Nonwords	550	55	717	104
Words				
Low Concrete	547	66	653	102
High Concrete	524	60	619	81
Difference	23	50	34	60
<u>Error Rates</u>				
Nonwords	3.7	3.3	17.6	11.5
Words				
Low Concrete	10.1	5.6	20.5	9.2
High Concrete	6.6	4.7	9.0	7.2
Difference	3.5	6.6	11.5	5.8

Table 5

Means and standard deviations for the reaction times (ms) and error rates (%) across Phases 2 and 4 of Experiment 2

	Phase 2 (Low Frequency Words)		Phase 4 (High Frequency Words)	
	M	SD	M	SD
Reaction Time				
Words	736	219	575	91
Nonwords	745	187	637	107
Error Rates				
Words	9.2	6.3	4.7	4.1
Nonwords	16.2	10.5	11.5	7.4

Table 6

Means (with standard deviations in the brackets) for the reaction times (ms) and error rates (%) in Experiments 3A and 3B

	Word:Nonword Ratio 2:1		1:1		1:2	
Experiment 3A						
Words						
Reaction Time			678	(125)	682	(129)
Error Rates			9.1	(5.2)	11.9	(6.4)
Nonwords						
Reaction Time			751	(182)	687	(159)
Error Rates			6.2	(4.9)	4.6	(4.2)
Experiment 3B						
Words						
Reaction Time	660	(104)	698	(120)		
Error Rates	5.0	(4.3)	6.8	(5.6)		
Nonwords						
Reaction Time	819	(165)	794	(158)		
Error Rates	10.2	(6.7)	8.4	(6.9)		

Figure Captions

Figure 1. Functions depicting the increase of harmony as a stimulus is processed. The referent function is assumed to reflect the "average" stimulus in a lexical decision experiment.

Figure 2. Depiction of the lexical decision process in a typical random-walk graph. The difference in starting position across Panels A and B reflects the efficiency principle, with Panel B reflecting a more efficient starting position for the random walk. SP stands for starting position. DR-W stands for the drift rate for words and is represented by the slope of the random walk toward the word boundary. DR-NW stands for the drift rate for nonwords and is represented by the slope of the random walk towards the nonword boundary.

Figure 3. Typical structure of a Hopfield network. All processing units are in one of two possible states and are interconnected. The connection weights between units reflect the similarities between the unit activations.

Figure 4. The cells of the table represent the processing units of the Hopfield network. Patterns of positive or negative activation within these cells represent various letters of the English alphabet (i.e. “X”, “A”, and “E”).

Figure 5. Panel A through C represent various stages of recall for the Hopfield network.

Figure 6. The structure of the Hopfield network that was used during the referent Simulations. There were a total of 25 orthographic units and 100 other units that represented phonology and semantics.

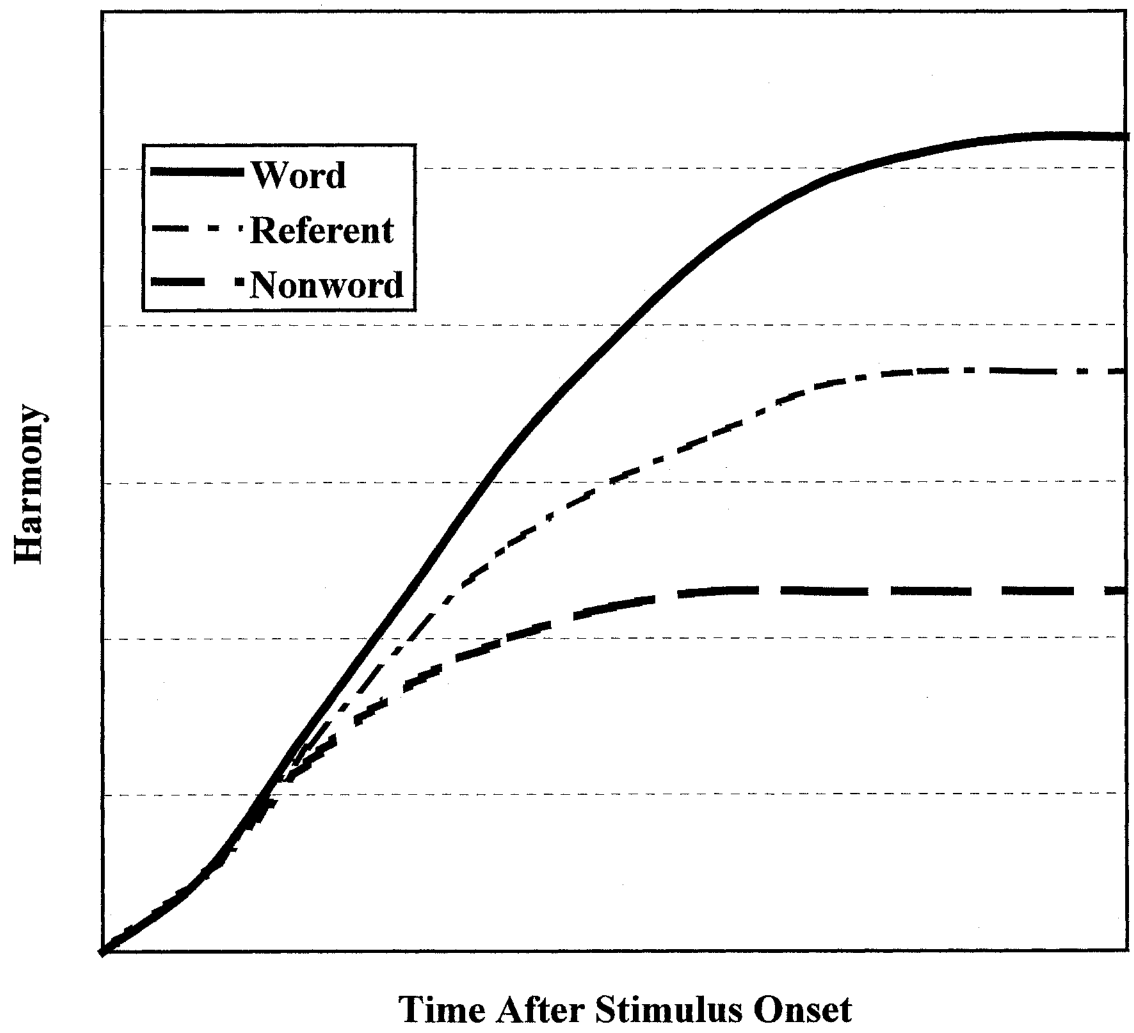
Word
Referent
Nonword
Harmony
Time After Stimulus Onset

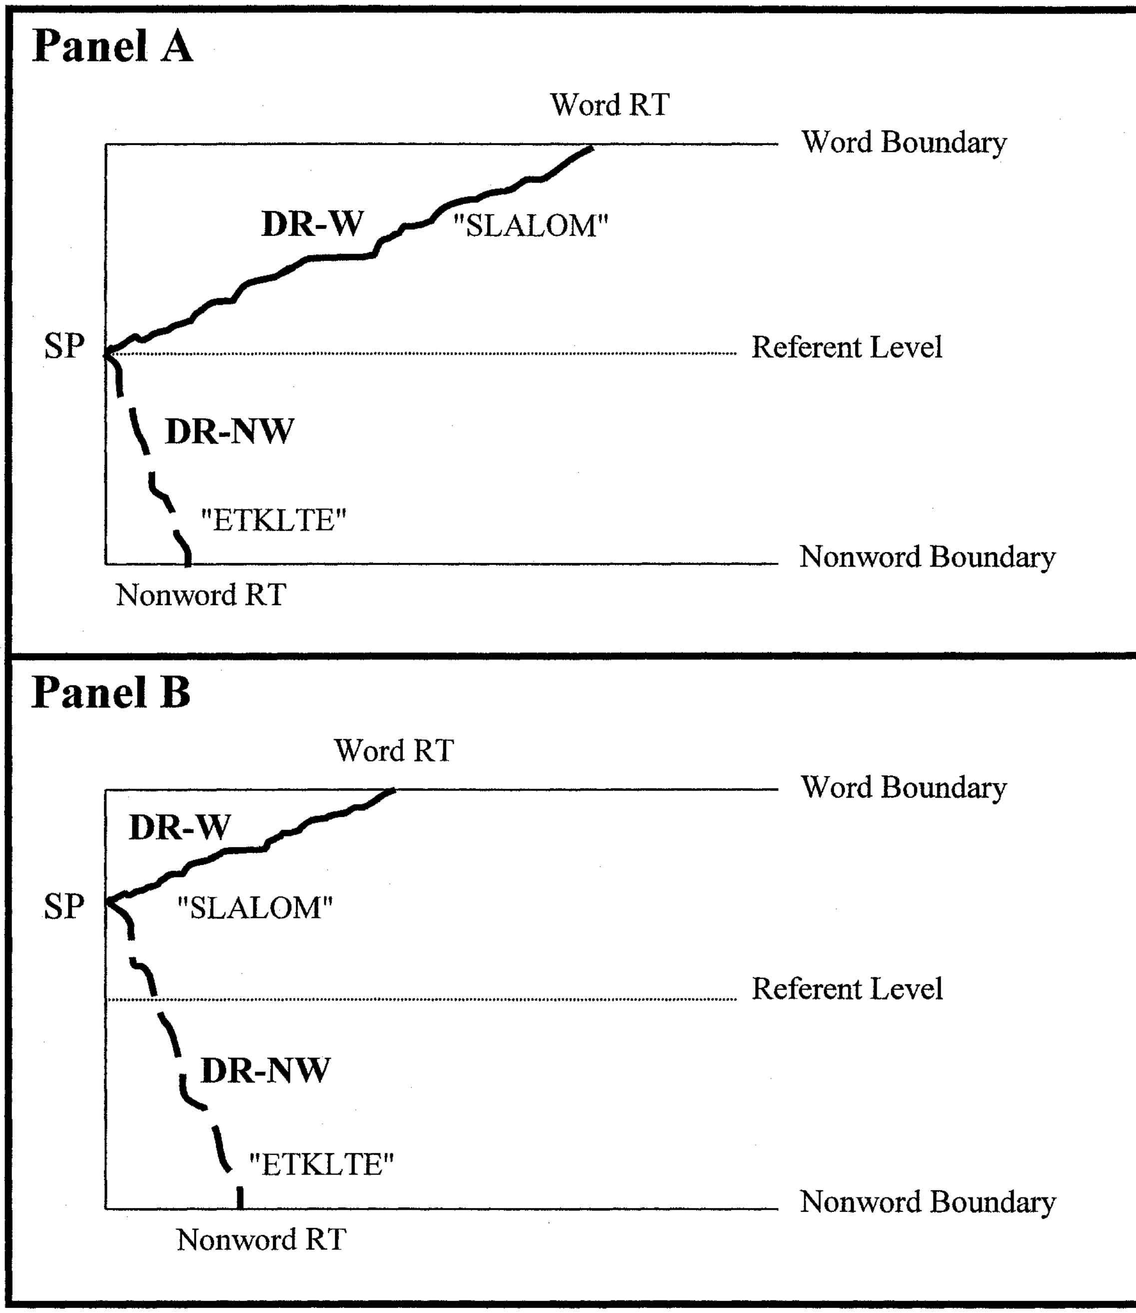
Panel A
Word RT
Word Boundary
DR-W
"SLALOM"
SP
Referent Level
DR-NW
"ETKLTE"
Nonword Boundary
Nonword RT
Panel B
Word RT
Word Boundary
DR-W
SP
"SLALOM"
Referent Level
DR-NW
"ETKLTE"
Nonword Boundary
Nonword RT

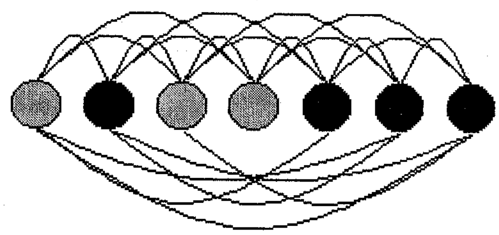

1	-1	-1	-1	1
-1	1	-1	1	-1
-1	-1	1	-1	-1
-1	1	-1	1	-1
1	-1	-1	-1	1

1	1	[illegible]	1	1
1	-1	-1	-1	1
1	1	1	1	1
1	-1	-1	-1	1
1	-1	-1	-1	1

1	1	1	1	1
1	-1	-1	-1	-1
1	1	1	1	[illegible]
1	-1	-1	-1	-1
1	1	1	1	1

Panel A

1	-1	1	-1	1
-1	1	-1	-1	1
1	-1	1	1	1
-1	-1	1	-1	1
-1	-1	1	-1	1

Panel B

1	1	1	-1	1
1	-1	-1	-1	1
1	-1	1	1	1
1	-1	-1	-1	1
1	-1	-1	-1	1

Panel C

1	1	1	1	1
1	-1	-1	-1	1
1	1	1	1	1
1	-1	-1	-1	1
1	-1	-1	-1	1

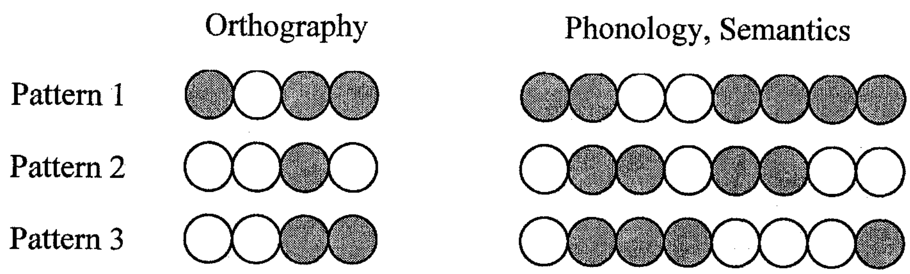
Orthography
Phonology, Semantics
Pattern 1
Pattern 2
Pattern 3

Appendix A

Words used in Phases 1 and 3 of Experiments 1A and 1B

Low Concreteness Stimuli

KF	Fam	Con	Stim	KF	Fam	Con	Stim	KF	Fam	Con	Stim
3	160	373	adage	72	504	268	allow	57	555	304	apart
11	383	371	array	2	274	351	ashen	4	326	371	audit
34	541	293	blame	3	436	338	bland	8	380	324	borne
73	523	361	brief	14	351	303	debut	4	328	317	dogma
9	460	267	dread	50	570	262	extra	22	541	315	fault
2	234	288	feint	10	414	304	folly	71	474	322	forth
8	447	304	fraud	21	493	304	glory	35	533	338	grade
43	552	356	grown	56	585	247	guess	33	559	299	guilt
12	509	359	harsh	5	494	290	hasty	61	521	253	ideal
7	420	360	idiom	12	458	243	irony	17	504	250	logic
20	460	239	mercy	29	475	308	merit	58	536	353	minor
21	517	306	pause	8	326	350	parry	72	516	360	phase
10	411	288	pious	45	492	360	prime	40	546	328	proof
7	350	365	proxy	68	570	343	quick	19	376	303	realm
7	433	305	reign	4	417	290	scorn	42	549	346	skill
27	562	383	slept	21	602	304	smart	23	552	313	spare
3	302	305	stoic	4	374	331	tally	55	524	336	theme
9	539	366	topic	23	531	371	trace	26	534	399	treat
46	503	328	trend	52	548	300	trust	13	480	365	utter
25	522	272	vague	50	529	279	worse	9	466	304	wrath

High Concreteness Stimuli

KF	Fam	Con	Stim	KF	Fam	Con	Stim	KF	Fam	Con	Stim
14	490	595	arrow	61	553	612	beach	7	456	564	beast
23	486	630	belly	35	488	614	bench	9	470	573	berry
59	496	585	bible	8	425	552	bosom	45	580	556	brain
3	384	611	brook	44	579	589	brush	7	441	623	canoe
50	513	595	chain	53	543	580	chest	69	507	578	china
11	479	591	cliff	3	511	627	clown	61	541	562	coast
5	425	572	coral	5	428	606	crane	13	474	627	dough
67	588	595	dress	33	473	516	drill	3	478	551	dummy
5	465	616	eagle	3	457	542	feast	11	458	580	ferry
5	401	595	flask	52	483	597	flesh	23	469	515	giant
13	477	535	globe	53	587	599	grass	33	501	535	grave
3	304	542	noose	5	444	618	olive	5	391	631	otter
54	550	624	phone	43	455	596	porch	41	527	537	queen
63	477	606	rifle	5	470	594	satin	8	452	562	scout
3	455	606	shawl	61	531	574	shore	30	418	539	slave
58	594	514	smile	41	596	541	smoke	44	501	621	snake
2	471	572	spike	58	564	614	stone	5	438	570	swamp
7	444	577	sword	3	454	586	thorn	13	463	585	tower
4	479	617	trout	57	557	580	uncle	2	445	550	vault
5	474	522	witch	28	474	561	wound	4	464	606	yacht

Appendix A (continued)

Words and Nonwords used in Phases 1 and 3 in Experiments 1A and 1B

Words

adopt	alley	annoy	anger	bangs	basic	beers	bless	boots
bound	bread	bride	brown	candy	cheat	child	clerk	clump
cooks	cruel	curse	decay	deals	dream	drugs	earth	ether
locks	facet	faith	fangs	flood	freak	frost	funny	gauge
grunt	heard	herbs	humid	lance	lungs	manic	match	model
money	mouse	packs	paint	pasta	peril	pills	polar	quart
ranch	rapid	rigor	rough	relax	salad	shift	signs	silky
skunk	spice	spoon	smell	snail	straw	stunt	tango	teeth
thaws	train	troll	unite	venom	while	widow	yeast	

Nonwords

Pseudo	Scram	Pseudo	Scram	Pseudo	Scram	Pseudo	Scram
appul	auplp	attak	taatk	armer	mrera	asess	ssaes
asign	inasg	attik	tkiat	batle	tbela	beger	eerbg
beloe	oeelb	bisun	nbuis	brall	albrl	byker	ybkre
caben	aecbn	carot	ctoar	cauze	czeau	cheez	hzece
cleen	nlece	cloke	elkco	cloze	clzeo	cryme	ymecr
dager	aedgr	daizy	zdyia	danse	ndesa	defur	reufd
denem	eednm	devel	evlde	durty	yrtdu	eagul	auelg
embir	emrbi	epick	peikc	errur	urerr	fabul	fualb
fansy	ayfsn	felun	fneul	feeld	edfle	fleks	klfes
flert	etlrf	fraim	miafr	fraze	zreaf	frite	frtie
furst	fsutr	glyde	lgyde	golph	gploh	gruve	ugrve
guzle	zgleu	habet	etbha	hailo	liaoh	hevin	vhien
iglue	iuegl	judje	djeuj	kamel	eakml	kanal	aklna
kanon	noakn	kight	gkiht	klawk	lkakw	komic	kmioc
labul	lbaul	laime	eialm	leson	lseon	majik	imkja
mimik	mkiim	musle	uelms	muzik	mziuk	neece	eecne
nerse	neesr	nibul	nbliu	nikle	knlei	noyse	yneos
nurve	nvure	ordir	roidr	pedel	eplde	penee	eepne
peper	rpepe	phlip	phpil	pikle	iepkl	pijun	unjpi
proze	epzro	pruve	pvuer	pudle	udple	pupit	ipupt
rabit	tiabr	rapht	ahrpt	relik	rliek	riple	lpier
rivul	luivr	sirge	rsieg	skool	oklso	sleat	aelst
sneek	nkese	staph	psaht	surve	seuvr	swerl	rwlse
tenis	tnsie	tikle	tliek	toste	otste	truse	rtesu
twerl	rlwte	undir	nridu	voise	sieov	wheal	elwha