

RESEARCH ON SPEECH PERCEPTION

Technical Report No. 6

December 1, 1986

Speech Research Laboratory

Department of Psychology

Indiana University

Bloomington, Indiana 47405

Supported by:

Department of Health and Human Services

U.S. Public Health Service

National Institutes of Health

Research Grant No. NS-12179-10



NEIGHBORHOODS OF WORDS IN THE MENTAL LEXICON

Paul A. Luce

Submitted to the faculty of the Graduate School
in partial fulfillment of the requirements
of the degree
Doctor of Philosophy
in the Department of Psychology
Indiana University
December, 1986

©

1986

Paul A. Luce

ALL RIGHTS RESERVED

TABLE OF CONTENTS

	Page
Abstract	iv
Acknowledgements	v
Chapter One: Introduction	1
Previous Research on the Role of Neighborhood Structure and Frequency in Word Recognition	3
Neighborhood Structure	3
Word Frequency	5
Description of Present Approach	6
Summary	7
Chapter Two: Evidence from Perceptual Identification	8
Experiment	9
Method	9
Stimuli	9
Screening	9
Subjects	10
Design	10
Procedure	10
Data Analysis	11
Computation of Lexical Statistics	12
Neighborhood Variables	12
Results	13
Analyses over All S/N Ratios	13
Correlation Analyses	13
Regression Analyses	13
Analyses at Each S/N Ratio	13
Correlation Analyses	14
Regression Analyses	14
Discussion	15
Summary and Conclusions	16
Chapter Three: Neighborhood Probability Rules	17
Experiment	18
Method	18
Stimuli	18
Subjects	18
Design	18
Procedure	19
Results	20
Neighborhood Probability Rules	21
Stimulus Word Probabilities	21
Neighbor Word Probabilities	22
Unweighted Neighborhood Probability Rule	22
Frequency-Weighted Neighborhood Probability Rules	24
Predicting Identification Performance Using the Neighborhood Probability Rules	25
Computation of Rules	25
Analyses over All S/N Ratios	27
Correlation Analyses	27
Regression Analyses	28
Analyses at Each S/N Ratio	30
Correlation Analyses	30

Regression Analyses	31
Summary	32
General Discussion	33
Factors Influencing the Performance of the Rules	33
Qualitative Predictions of the Neighborhood Probability Rules	36
Neighborhood Activation Model	38
Summary and Conclusions	43
Footnotes	44
 Chapter Four: Evidence from Auditory	
Lexical Decision	45
Experiment	46
Method	46
Stimuli	46
Words	46
Nonwords	46
Subjects	47
Design and Procedure	47
Results	48
Analysis of Word Responses	48
Correlation Analyses	49
Regression Analyses	50
Discussion of Correlation and Regression Analyses	50
Analysis of Partitioned Word Data	51
Accuracy Data	52
Reaction Time Data	53
Summary of Word Data	53
Analysis of Nonword Responses	54
Accuracy Data	54
Reaction Time Data	54
Summary of Nonword Data	55
Discussion	55
Lexical Decision in the Neighborhood Activation Model	55
Word Responses	57
Nonword Responses	58
Summary and Conclusions	59
 Chapter Five: Evidence from Auditory Word Naming	60
Experiment	61
Method	61
Stimuli	61
Subjects	62
Design and Procedure	62
Results	63
Accuracy Data	63
Reaction Time Data	63
Discussion	64
 Chapter Six: Conclusion	67
Review of Major Results	68
Perceptual Identification	68
Auditory Lexical Decision	68
Auditory Word Naming	70
Summary of the Neighborhood Activation Model	71

Comparison of NAM to Other Models of Word	
Recognition	72
Logogen Theory	72
Cohort Theory	74
Other Theories of Word Recognition	76
Forster's Search Theory	76
Activation-Verification Model	77
Interactive-Activation Models	78
Summary of Word Recognition Models	78
Additional Considerations and Future Directions	79
Contextual Effects	79
Priming Effects	79
Neighborhood Structural Effects in Visual Word	
Recognition	80
Frequency and Familiarity	81
Future Directions in the Development of NAM	81
The Nature of the Similarity of Speech Sounds	81
The Nature of the Units of Representation	82
Effects of Syllable Stress	82
Conclusion	83
References	84



ABSTRACT

A fundamental, but often neglected, problem in the study of human auditory word recognition concerns the structural relations among the sound patterns of words in memory and the effects these relations have on auditory word recognition. In the present investigation, computational and experimental methods were employed to address a number of issues related to the representation and structural organization of spoken words in the mental lexicon. Using a computerized lexicon consisting of phonetic transcriptions of 20,000 words, "similarity neighborhoods" for each of the transcriptions were computed. Among the variables of interest in the computation of the similarity neighborhoods were: (1) the number of words occurring in a neighborhood, (2) the degree of phonetic similarity among the words, and (3) the frequencies of occurrence of the words in the language. The effects of these variables on auditory word recognition were examined in a series of behavioral experiments employing three experimental paradigms: perceptual identification of words in noise, auditory lexical decision, and auditory word naming (i.e., pronunciation). The results of each of these experiments provided strong support for the hypothesis that words are recognized in the context of similar words in the mental lexicon. In particular, it was demonstrated that the number and nature of words in a similarity neighborhood affect the speed and ease with which words are recognized. A neighborhood probability rule was developed that adequately predicted identification performance. This rule, based on Luce's (1959) choice rule, combines stimulus word intelligibility, neighborhood confusability, and frequency into a single expression. Based on this rule, a model of auditory word recognition, the neighborhood activation model, was proposed. This model describes the effects of similarity neighborhood structure on the process of discriminating among the acoustic-phonetic representations of words in memory. The results of these experiments have important implications for current conceptions of both human auditory word recognition and the structural organization of spoken words in the mental lexicon.



ACKNOWLEDGEMENTS

I would first like to thank my wife, Jan Charles-Luce, for her encouragement, support, and assistance at every stage of this project. I am especially appreciative of the many hours she spent discussing the issues with me; her knowledge of phonetics and phonology proved invaluable. She is a first-rate colleague, as well as my best friend. Many thanks, Jan. I would also like to express my deepest appreciation to my advisor, David Pisoni. His contributions to this thesis are inestimable. I have genuinely enjoyed my collaboration with David and it goes without saying that none of this would have been possible without his scientific and moral support. I would also like to thank Howard Nusbaum for many fruitful discussions on a diverse range of issues. Finally, I would like to thank my committee--Dan Dinnsen, Rich Shiffrin, and Linda Smith--for their many helpful comments and suggestions.

This work was supported, in part, by NIH grant NS-12179-10 and a Dissertation Year Fellowship and doctoral grant-in-aid from Indiana University.



CHAPTER ONE

INTRODUCTION

Much attention has previously been devoted to the structural organization of words in the mental lexicon. Typically, however, cognitive psychologists have focused on the structure of higher-level aspects of lexical representations, namely the semantic and conceptual organization of lexical items in memory (e.g., Miller & Johnson-Laird, 1976; Smith, 1978). As a consequence, little attention has been paid to the structural organization of the information used to gain access to these higher-level sources of information. The goal of the present investigation was to explore in detail this structure and its implications for theories of auditory word recognition.

In the present set of studies, "structure" will be defined specifically in terms of similarity relations among the sound patterns of words. Similarity will serve as the means by which the organization of acoustic-phonetic representations in memory will be investigated. Indeed, it is assumed that similarity relations among the sound patterns of words represent one of the earliest stages at which the structural organization of the lexicon comes into play. The identification of structure with similarity may seem to ignore more "interesting" questions relating to the structural organization of words, in particular, phonological and morphological structure. Clearly, the structural constraints imposed by the phonology and morphology of the language are important considerations. However, the precise aim of the present investigation was to gain a further understanding of the lower-level relations between stimulus input, activation of phonetic representations, and, subsequently, recognition of a word. It is assumed that the delineation of these processes must serve as the groundwork on which to base more complete theories of word recognition and lexical access that incorporate notions of phonological and morphological structure.

The identification of structure with similarity relations among sound patterns of words raises the difficult problem of defining similarity. Similarity, although crucial to the present investigation, is an ill-defined concept in research on speech perception and auditory word recognition, and one that deserves considerably more work in these areas of research (see Mermelstein, 1976). However, similarity can be approximated by both computational and behavioral predictors of confusion, the approach taken here. Thus, similarity will be defined in terms of a computational metric for predicting confusions among phonetic patterns as well as a behavioral, or operational, metric.

Having defined structure as the similarity relations among the sound patterns of words, the question arises: Should the structural organization of representations in memory have consequences for auditory word recognition? Consider a content-addressable memory system in which there is no noise either in the signal or the listener (Kohonen, 1980). In such a system, encoding the acoustic-phonetic information in the stimulus word is tantamount to locating the word in memory. In this case, the structural organization of acoustic-phonetic representations in memory would have no consequences for word identification. Instead, the task of auditory word recognition would be identical to phonetic perception, and one need only study phonetic perception in order to understand how words are recognized.

It is undeniable that phonetic perception is important in auditory word recognition. It is not undeniable, however, that the human word recognition system operates as a noiseless content-addressable system or that the acoustic-phonetic signal itself is devoid of noise. To begin, the signal is very often less than ideal for the purposes of the listener. Words are typically perceived against a background of considerable ambient noise, reverberation, and the voices of other talkers. In addition, coarticulatory effects and segmental reductions and deletions substantially restructure the phonetic information in a myriad of ways (see Luce & Pisoni, 1987). Although such effects may indeed be useful to the listener (Church, 1983; Elman & McClelland, 1986), they also undoubtedly produce considerable ambiguities in the acoustic-phonetic signal, making a strictly content-addressable word recognition system based on phonetic encoding infeasible. In short, both the noise inherent in the signal as well as the noise against which the signal is perceived make it unlikely that word recognition is accomplished by direct access, based solely on phonetic encoding, to acoustic-phonetic representations in memory.

Not only is the signal noisy, so too is the recognition system of the listener. Although the human is clearly well adapted for the perception of spoken language, the system by which language is perceived is by no means a perfect one. Encoding, attentional, and memory demands frequently result in the distortion, degradation, or loss of acoustic-phonetic information. The data on misperceptions alone attest to the fact that the auditory word recognition system is less than perfect (Bond & Garnes, 1980; Bond & Robey, 1983). Thus, again, a strictly content-addressable system does not suffice as a model of human auditory word recognition.

The alternative to a noiseless content-addressable system is one in which the stimulus input activates a number of similar acoustic-phonetic representations or candidates in memory, among which the system must choose (Marslen-Wilson & Welsh, 1978). In this system, a considerable amount of work involves discriminating among the lexical items activated in memory. Indeed, many current models of word recognition subscribe to the view that word recognition is to a great degree a process of discriminating among competing lexical items (Forster, 1979; Marslen-Wilson & Welsh, 1978; Morton, 1979; Rumelhart & McClelland, 1981).

Given that one of the primary tasks of the word recognition system involves discrimination among lexical items, the study of the structural organization of words in memory takes on considerable importance, especially if structural relations influence the ease or difficulty of lexical discrimination, and, subsequently, word recognition and lexical access. By the same token, under the assumption that word recognition involves discrimination among competing lexical items, variations in the ease or difficulty of discriminating among items in memory can enlighten us as to the structural organization of the sound patterns of words. In short, lexical discrimination and structure are so inextricably tied together that the study of one leads to a further understanding of the other.

Assuming, then, that structural relations among words should influence auditory word recognition via the process of discrimination, it is important to determine that structural differences among words actually exist. Previous research (Landauer & Streeter, 1973; Luce, 1986c) has indeed demonstrated that words vary substantially not only in the number of words to which they are similar, but also in the frequencies of these words. These findings suggest that both structural and frequency relations among words may mediate lexical discrimination. Investigation of the behavioral effects of these sorts of

relations should help us to understand further not only the process of lexical discrimination, but also the organization of the sound patterns of words in memory.

The issue of word frequency takes on an important role in the investigation of the structural organization of the sound patterns of words. Numerous previous studies (Howes, 1957; Newbigging, 1961; Savin, 1963; Solomon & Postman, 1952) have demonstrated that the ease with which auditorily presented words are recognized is monotonically related to experienced frequency, as measured by some objective count of words in the language. However, little work has been devoted to detailing the interaction of word frequency and structural relations among words (see, however, Triesman, 1978a, 1978b). If word frequency influences the perceptibility of the stimulus word, it should also influence the degree of activation of similar words in memory. Frequency is important, then, in further specifying the relative competition among activated items that are to be discriminated among.

The goal of the present investigation was therefore to examine the effects of the number and nature of words activated in memory on auditory word recognition. Throughout the ensuing discussion, the term similarity neighborhood will be employed. A similarity neighborhood is defined as a collection of words that are phonetically similar to a given stimulus word. (The term stimulus word will be used to refer to the word for which a neighborhood is computed.) Similarity neighborhood structure refers to two factors: (1) the number and degree of confusability of words in the neighborhood, and (2) the frequencies of the neighbors. This first factor will be referred to as neighborhood density; the latter factor will be called neighborhood frequency. In addition to neighborhood structure, the frequency of the stimulus word itself will be of interest.

PREVIOUS RESEARCH ON THE ROLE OF NEIGHBORHOOD STRUCTURE AND FREQUENCY IN WORD RECOGNITION

Neighborhood Structure

For the most part, the present investigation constitutes a fairly novel approach toward the study of auditory word recognition. Little previous research has been devoted to examining the effects of neighborhood structure, primarily because of the lack of computational tools for determining similarity neighborhoods for a large number of words. One early study of visual word recognition by Havens and Foote (1963) examined the effects of the number of competitors, or neighbors, of words on tachistoscopic identification. The results of this study, although based on a very small number of words and a rather imprecise measure of neighborhood membership, are suggestive. Havens and Foote demonstrated that effects of word frequency could be eliminated if the number of competitors for a given word are controlled. That is, low frequency words were identified at levels of accuracy equal to those of high frequency words when the number of competitors was held constant. This result suggests that the effect of frequency is crucially dependent on the neighborhood in which the word resides in the lexicon.

Similar suggestive evidence has been presented in a little known thesis by Anderson (1962). In this study, Anderson examined the effects of the nature and number of alternatives on the intelligibility of spoken words. Although again the means for determining alternatives were crude, Anderson demonstrated that intelligibility of spoken words was affected both by the

number of possible confusors as well as by the frequencies of these confusors. In general, Anderson showed that words with many possible confusors were less intelligible than words with fewer confusors. In addition, he demonstrated that high frequency confusors tended to depress identification performance.

Other evidence for the role of neighborhood structure in auditory word recognition was obtained from a reanalysis of a set of data published by Hood and Poole (1980). Hood and Poole examined the intelligibility of words presented in white noise. They found that word frequency failed to correlate consistently with the word intelligibility scores for their data, in apparent contradiction to many previous findings regarding the effects of frequency on noise-masked words (Howes, 1957; Savin, 1963). This finding indicated that factors other than word frequency were responsible for the wide range of intelligibility scores obtained by Hood and Poole.

To examine the possibility that similarity neighborhood structure was, at least in part, responsible for the differences observed in intelligibility in the Hood and Poole study, I examined their 25 easiest and 25 most difficult words. Similarity neighborhoods were computed (see below) for each of these 50 words. In keeping with Hood and Poole's observation regarding word frequency, no significant difference in frequency was found between the 25 easiest and 25 most difficult words. However, it was found that the relations of easy words to their neighbors differed substantially from the relations of the difficult words to their neighbors. More specifically, approximately 56% of the words in the neighborhoods of the difficult words were equal to or higher in frequency than the frequencies of the difficult words themselves. For the 25 easy words, however, only approximately 23% of the neighbors of the easy words were of equal or higher frequency. Thus, it appears that the observed differences in intelligibility were due, at least in part, to the frequency composition of the neighborhoods of the easy and difficult words, and were not primarily due to the frequencies of the words themselves.

Taken together, these studies strongly suggest that neighborhood structure may play an important role in word recognition. Furthermore, these studies suggest that the effect of the frequency of the stimulus word itself may be mediated by neighborhood structure. Thus, the present set of studies was aimed first at examining the role of neighborhood structure in auditory word recognition, and second at detailing the role of stimulus word frequency in the context of neighbor frequencies. The study of the combined effects of neighborhood structure and frequency should therefore lead to a deeper understanding of the effects of both similarity and frequency on word recognition.

Given that so little research has been devoted to these problems, it is hardly surprising that current models of auditory word recognition have had little to say about the structural organization of acoustic-phonetic patterns in the mental lexicon. Only cohort theory (Marslen-Wilson & Welsh, 1978) has made any precise claims regarding structural effects, and these have primarily been based on the claim that words are recognized at the point at which they diverge from all other words in the mental lexicon, a claim that says little about the structural organization of words in memory. Although a detailed discussion of the various models of auditory word recognition will be deferred until Chapter 6, it should be noted at this point that similarity neighborhood structural effects have been for the most part ignored in both research and theory on auditory word recognition.

As I hope to show, this has been a serious omission in earlier work. Indeed, I will attempt to demonstrate that any adequate theory of word recognition must provide a basic account of the structure of the sound patterns of words in memory as well as how this structure affects perceptual processing.

Word Frequency

Although little work has been devoted to the study of neighborhood structure, a voluminous body of data has been collected on the effect of word frequency in visual and auditory word recognition (e.g., Gordon, 1983; Glanzer & Bowles, 1976; Glanzer & Ehrenreich, 1979; Howes, 1954, 1957; Howes & Solomon, 1951; Landauer & Freedman, 1968; Morton, 1969; Newbigging, 1961; Rubenstein, Garfield, & Millikan, 1970; Rumelhart & Siple, 1974; Savin, 1963; Scarborough, Cortese, & Scarborough, 1977; Soloman & Postman, 1952; Stanners, Jastrzembski, & Westbrook, 1975; Whaley, 1978). In general, these results have demonstrated numerous processing advantages for high frequency words. Many theories have also been proposed to account for the advantages associated with increased word frequency. These theories have cited frequency of exposure (Forster, 1976; Morton, 1969), age of acquisition (Carroll & White, 1973a, 1973b), and the time between the present and last encounter with the word (Scarborough, Cortese, and Scarborough, 1977) as the underlying reasons for the processing advantages observed for high frequency words. Whatever the precise mechanism, it is now assumed by many researchers (Broadbent, 1967; Catlin, 1969; Goldiamond & Hawkins, 1958; Nakatani, 1969; Newbigging, 1961; Pollack, Rubenstein, & Decker, 1960; Savin, 1963; Soloman & Postman, 1952; Triesman, 1971, 1978a,b) that frequency serves to bias, in some manner, the word recognition system toward choosing high frequency words over low frequency words. Although the claim that frequency effects arise from biases is not uncontroversial, many theories of the word frequency effect have espoused such a view (see references cited above). Among these theories are sophisticated guessing theory (Neisser, 1967; Newbigging, 1961; Pollack et al., 1960; Savin, 1963; Soloman & Postman, 1956), criterion-bias theory (Broadbent, 1967), and partial identification theory (Triesman, 1978a,b). Although there has been considerable debate among the proponents of each of these theories, all assume that word frequency, by some as yet poorly specified mechanism, influences the decisions of the word recognition system via some sort of bias (see also Gordon, 1983; Norris, 1982).

Although there is some agreement among researchers as to the means by which processing advantages afforded by high frequency words arise, there has, as previously mentioned, been little research on the relation between frequency and neighborhood structure, a primary issue in the present set of studies. Only Triesman (1978a,b) has addressed the issue of how neighborhood structure may influence the word frequency effect (his work will be discussed in more detail in Chapter 3). Thus, although there has clearly been a great deal of research and theorizing on word frequency, little has been done to examine the effects of word frequency in conjunction with the effects of similarity neighborhood structure. The present set of studies is therefore aimed, in part, at examining word frequency in the context of the similarity neighborhood.

Description of the Present Approach

As previously stated, similarity neighborhood structure was estimated computationally, using a large, on-line lexicon. This lexicon, based on Webster's Pocket Dictionary (Webster's Seventh Collegiate Dictionary, 1967), contains approximately 20,000 entries. In the version of the lexicon used in the present set of studies, each entry contained: (1) an orthographic representation, (2) a phonetic transcription, (3) a frequency count based on the Kucera and Francis (1967) norms, and (4) a subjective familiarity rating (Nusbaum, Pisoni, & Davis, 1984). Examples of entries in the lexicon are shown in Table 1.1.

Insert Tables 1.1 and 1.2 about here

The phonetic transcriptions, coded in a computer-readable phonetic alphabet, were based on a general American dialect and included syllable boundary and stress markings. The computer-readable phonetic symbols and their IPA counterparts are shown in Table 1.2. (All subsequent phonetic transcriptions will employ the computer-readable symbols.) Frequency counts, as noted above, were obtained from an on-line version of the Kucera and Francis (1967) corpus. These counts were based on one million words of printed text. Although the study of frequency effects in auditory word recognition would be best served by a count of spoken words, no such count is available that covers the large number of words in Webster's lexicon. Finally, the subjective familiarity ratings for each of the words were obtained in a large-scale study by Nusbaum, et al. (1984). In this study, groups of college undergraduates were asked to rate the familiarity of each of the words in Webster's lexicon on a seven point scale, ranging from "don't know the word" (1) to "know the word and know its meaning" (7). The familiarity ratings were obtained from visually presented words.

The general procedure for computing similarity neighborhood structure was as follows: A given phonetic transcription (constituting the stimulus word) was compared to all other transcriptions in the lexicon (which constituted potential neighbors). A neighbor was defined as any transcription that could be converted to the transcription of the stimulus word by a one phoneme substitution, deletion, or addition in any position. For example, among the neighbors of the word /k@t/ would be /p@t/, /kIt/, and /k@n/, which are each derived on the basis of one phoneme substitutions in any position. Also included as neighbors would be the words /sk@t/ and /@t/, derived on the basis of one phoneme additions or deletions. (Plurals and inflected forms of the stimulus were not included as neighbors.) The number of such neighbors constitutes the variable of neighborhood density. Neighborhood frequency refers to the average of the frequencies, based on the Kucera and Francis counts, of each of the neighbors. Stimulus word frequency refers to the frequency of the stimulus word for which the neighbors were computed (e.g., /k@t/).

This particular algorithm for computing neighborhood membership was based on previous work by Greenberg and Jenkins (1964) and Landauer and Streeter (1973; see also Sankoff & Kruskal, 1983). Clearly, this method of computing neighborhood membership makes certain strong assumptions regarding similarity. In particular, it assumes that all phonemes are equally similar and that the similarities of phonemes at a given position are equivalent. However, this method of computing similarity neighborhoods provides a computationally simple means of estimating the number and nature of words similar to a given stimulus

Table 1.1.

Examples of entries in Webster's lexicon.

ORTHOGRAPHY	TRANSCRIPTION	FREQUENCY	RATING	CONTENT/FUNCTION
baby	b'e<bi	62	7.00000	c
bachelor	b'@C-lX	6	7.00000	c
bacillus	bx-s'I*lxS	2	3.08333	c
back	b'@k	967	7.00000	c
bacon	b'e<k n	10	7.00000	c
bad	b'@d	142	7.00000	c
bade	b'@d	1	3.25000	c

Table 1.2

Computer-readable phonetic symbols and their IPA counterparts. The computer-readable symbols are given first, followed by the IPA symbols.

p - p	l - l
t - t	r - r
k - k	w - w
b - b	y - y
d - d	
g - g	i - i
C - tʃ	I - I
J - dʒ	E - E
s - s	e - e
S - ʃ	@ - œ
z - z	a - a
Z - ʒ	W - aʊ
f - f	Y - aI
T - θ	^ - ^
v - v	c - ɔ
D - ð	O - oI
h - h	o - o
n - n	U - ʊ
G - ɣ	u - u
m - m	R - ɹ

word. The real test of the algorithm, of course, lies in the demonstration of the behavioral consequences of the similarity neighborhoods computed in this manner. Indeed, the ensuing chapters validate, in part, the use of this simple metric for determining neighborhood membership. As will be discussed in Chapter 3, however, a more sophisticated means of computing neighborhood structure, using confusion probabilities, can be devised that more closely models the actual similarities among phonemes. Nonetheless, the present algorithm serves as one computationally simple means of estimating neighborhood structure, and serves as the basis for deriving the more sophisticated algorithm discussed in Chapter 3.

The present approach thus combined computational and behavioral techniques in the investigation of neighborhood structure and auditory word recognition. The effects of neighborhood structure (i.e., density and neighborhood frequency) were examined in a variety of experimental tasks, ranging from the perceptual identification of words in noise to auditory word naming. The use of various experimental paradigms enabled examination of neighborhood structural effects under a variety of task situations in order to test the generality of the effects in question and the validity of the neighborhood analysis. In addition, unlike much previous work in visual and auditory word recognition, a large number of well-defined stimuli were employed in the present set of experiments in order to estimate the magnitude and generality of the effects of neighborhood structure across a fairly representative sample of monosyllabic words in English. In short, the present investigation brings to bear a number of powerful computational and behavioral techniques to the study of the effects of neighborhood structure on auditory word recognition.

SUMMARY

The present investigation represents an attempt to uncover the precise role of neighborhood structure in auditory word recognition using computational and behavioral techniques. The precise hypothesis put forth is that words are recognized in the context of other words in memory. More precisely, it is predicted that the number of words that must be discriminated among in memory will affect the accuracy and time-course of word recognition. It is furthermore hypothesized that the frequencies of the words activated in memory will affect decision processes responsible for choosing among the activated words. Finally, it is proposed that the classic word frequency effect may be a function of neighborhood frequency and similarity, and not a simple direct function of the number of times the stimulus word has been encountered.

CHAPTER TWO

EVIDENCE FROM PERCEPTUAL IDENTIFICATION

The goal of the present investigation was to specify the effects of similarity neighborhood structure on auditory word recognition and to examine the implications of these effects for the structure of words in the mental lexicon. As a first approximation to achieving this goal, the experiment reported in the present chapter examined the combined effects of neighborhood structure and frequency on the perceptual identification of words masked by noise. As discussed in Chapter 1, the particular approach toward addressing these problems involved the use of a computerized lexicon as a general model of the mental lexicon. This lexicon provided a means for estimating the structural and frequency relationships among words in order to evaluate the role these factors play in auditory word recognition.

In the present study, the now standard paradigm of presenting words for perceptual identification against a background of noise was employed. This paradigm is particularly useful in investigating structural relationships among words because stimulus degradation can be used to magnify the process of discrimination among items in memory. Furthermore, this task is attractive because of the naturalness of the response required of the subject. Many current experimental paradigms require fairly artificial responses (e.g., word-nonword decisions, shadowing). The identification response, on the other hand, is a simple and straightforward extension of the normal activity of recognizing words.

The identification task is not without its disadvantages, however. Some investigators have raised the objection that stimulus degradation slows processing, allowing considerable post-perceptual processes and biases to influence the identification response (e.g., Marslen-Wilson, 1986). This assumes, of course, that the post-perceptual processes are not an intrinsic component of auditory word recognition under normal circumstances. For example, frequency effects have often been attributed to response biases arising post-perceptually. The fact that frequency effects may be "post-perceptual," however, does not diminish their potential importance in normal auditory word recognition. Frequency biases exist throughout the cognitive system (Hasher & Zacks, 1984) and reflect an important and fundamental aspect of the human's response to his/her environment. Thus, there is no reason to assume a priori that frequency effects in the perceptual identification of words are in some sense unimportant because they arise from a propensity on the part of the subject to respond in a specific way to his/her environment based on past experience (see Smith, 1980).

In the present investigation, approximately 900 words were examined. In research on natural language, the stimuli (in this case, words) rarely classify themselves neatly into the cells required for analysis of variance. The approach taken here was thus to examine a large number of words on which few constraints were placed. Correlation and multiple regression analyses, as well as post-hoc classification strategies, were then used to interrogate the data and verify specific hypotheses.

Two approaches to evaluating the structural and frequency relationships among words using the perceptual identification paradigm were taken. The first, reported in the present chapter, attempted to relate similarity neighborhood statistics computed on the basis of a computerized lexicon to word identification. Lexical statistics were therefore computed in order to

evaluate the combined effects of structural and frequency factors. The second approach, discussed in Chapter 3, involved an attempt to develop a specific rule for predicting the identification of stimulus words in the context of their neighbors. This rule, dubbed the neighborhood probability rule, was used to combine the effects of stimulus word intelligibility, confusability, and frequency into a single expression that attempted to capture the phenomena under consideration.

In summary, the goal of the present study was to determine what effects, if any, structural relations among the sound patterns of words have on the process of discriminating among lexical items in the perceptual identification task. In addition, a second goal was to explore the effect of word frequency in the context of the structural organization of lexical items in memory.

EXPERIMENT

Method

Stimuli

Nine-hundred and eighteen words were selected from Webster's lexicon that met the following criteria: (1) All words were three phonemes in length; (2) all were monosyllabic; (3) all were listed in the Brown corpus of frequency counts (Kucera & Francis, 1967); and (4) all words had a rated familiarity of 6.0 or above on a seven point scale. The familiarity ratings were obtained from a previous study by Nusbaum, Pisoni, and Davis (1984). In this study, all words from the Webster's lexicon were presented visually for familiarity ratings. The rating scale ranged from "don't know the word" (1) to "recognize the word but don't know the meaning" (4) to "know the word" (7). The rating criterion was established to ensure that the words would be known to the subjects.

The 918 words were recorded by a male speaker of a Midwestern dialect. The stimuli were recorded in a sound attenuated booth (IAC model 401A) using an Electro-Voice D054 microphone and an Ampex AG-500 tape recorder. The stimuli were then low-pass filtered at 4.8 kHz and digitized via a 12-bit analog-to-digital converter operating at a sampling rate of 10 kHz. Using WAVES, a digital waveform editor (see Luce & Carrell, 1981), each stimulus was spliced from the entire stimulus set and placed in a separate stimulus file. After editing the stimuli into separate files, all stimulus files were equated for overall RMS amplitude using the program WAVMOD (Bernacki, 1981). Equating for RMS amplitude ensured that the stimuli were approximately equal in average intensity.

The 918 stimuli were then randomly partitioned into three stimulus set files consisting of 306 words each. Each of the resulting stimulus set files had approximately equal means for each of the independent variables (see below). From two of the three stimulus sets, two practice lists of 15 words each were selected and placed in separate stimulus set files.

Screening

Prior to conducting the identification experiment proper, each of the 918 words was screened to ensure that no clearly anomalous stimuli were included in the final analysis. Each of the three stimulus set files were presented to separate groups of 10 subjects, resulting in ten observations per word. For

the screening experiment, each word was presented at 75 dB SPL in the absence of masking noise. Except for the manipulation of signal-to-noise (S/N) ratio, the procedure for stimulus presentation and data collection was identical to that for the identification experiment (see Procedure section below). Only those words identified at a level of 90% correct (9 out of 10 subjects) or above were included in the final analysis. Thirty-six of the original 918 words failed to meet this criterion. Although these words were presented in the identification experiment in order to maintain equal numbers of stimuli in each of the three stimulus set files, these words were eliminated from consideration in the analyses of the data.

Subjects

Ninety subjects participated in partial fulfillment of an introductory psychology course. All subjects were native English speakers, reported no history of speech or hearing disorders, and were able to type.

Design

All stimuli were presented at each of three S/N ratios: +15 dB, +5 dB, and -5 dB SPL. S/N ratio was manipulated by varying the amplitude of the stimuli against a constant level of white, band-limited, Gaussian noise. The level of the noise was set at 70 dB SPL and was low-pass filtered at 4.8 kHz to match the gross spectral range of the stimuli. The stimuli were presented at 85 dB SPL for the +15 dB S/N ratio, 75 dB for the +5 dB S/N ratio, and 65 dB for the -5 dB S/N ratio. The choice of S/N ratios was based on a pilot experiment; the three S/N ratios were chosen to produce a level of accuracy, collapsed across S/N ratio, approximately equal to 50 percent correct identification.

Each of the three stimulus set files was presented to three groups of 10 subjects each. Each group of subjects heard one-third of the stimuli at +15 dB, one-third at +5 dB, and one-third at -5 dB. However, presentation at a given S/N ratio varied randomly from trial to trial. For a given stimulus, S/N ratio was a between-subjects factor. Altogether, 10 subjects heard each word at each S/N ratio.

Procedure

Stimulus presentation and data collection were controlled on-line in real-time by a PDP-11/34 minicomputer. The stimuli were presented via a 12-bit digital-to-analog converter over matched and calibrated TDH-39 headphones. The stimuli and noise were first manually calibrated at 85 dB SPL. Programmable attenuators were then adjusted for each trial to achieve the desired S/N ratio.

Subjects were tested in individual booths in a sound-treated room. ADM CRT terminals interfaced to the PDP-11/34 computer were situated in each of the booths. The procedure for an experimental trial was as follows: Subjects were first presented with the message "READY FOR NEXT WORD" on their terminals. One second following the message, 70 dB SPL of white noise was presented over the headphones. One-hundred msec after the onset of the noise, a randomly selected stimulus was presented at one of the three attenuation levels. One-hundred msec following the offset of the stimulus, the noise was terminated until the beginning of the next trial. Following presentation of

the stimulus and noise, a prompt appeared on each subject's terminal. Subjects then typed their responses on the terminals and pressed the RETURN key when finished. Subjects were able to see their responses while typing and were able to correct any typing errors prior to pressing the return key. After each subject had responded, another trial was initiated. In the event that one or more subjects failed to respond, a new trial was initiated within 30 seconds of the offset of the noise. Alphanumeric string responses were collected by the PDP-11/34 and stored in disk files for later analysis.

Subjects were instructed to provide their best guesses for each word they heard. They were also instructed to enter no response (i.e., simply press the RETURN key) only in the event that they were completely unable to identify the word. Following the instructions, 15 practice words, each at one of the three S/N ratios, were presented. None of the 15 practice words were presented in the main experiment. Following the practice list, the instructions were summarized and procedural questions were answered. One of the stimulus set files consisting of 306 words was then presented. Three short breaks were given at equal intervals. An experimental session lasted approximately one hour.

Data Analysis

Data files were transferred from the PDP-11/34 to the SRL VAX 11-750 for analysis. The data files were first combined into a master list consisting of the responses from 10 subjects for each S/N ratio (resulting in 30 total responses per word) for each of the 918 words. The 36 words failing to meet the criterion established in the screening experiment were marked and excluded from further analysis, leaving data for 882 words. In total, 26,460 (882 words X three S/N ratios X 10 observations) subject responses were included in the master data list.

The master list was edited to correct misspellings. Corrections for misspellings were performed by correcting transpositions, deleting single letter insertions, inserting single letter omissions, and correcting single letter substitutions. Single letter substitutions were corrected only when the key of the incorrect letter was within one key of the target letter on the keyboard or when the correct letter would have been produced if the same keystroke had been performed by the opposite hand. Only responses constituting nonwords were corrected in this manner. Approximately 2.5% of the responses were corrected for misspellings.

In addition to correction of misspellings, responses were coded in four ways: (1) inflected forms (i.e., past tense and plural forms) of the target words were marked as "f"; (2) homophones of the target word were marked as "h"; (3) nonwords were marked as "n"; and (4) illegal responses (i.e. responses forming neither words nor phonotactically legal nonwords) were marked as "i." Null responses were entered as blanks.

Following editing of the master list, the following information from Webster's lexicon was included for each of the subjects responses in the master list: phonetic transcriptions, Kucera and Francis (1967) frequencies, and familiarity ratings (Nusbaum et al., 1984). In the case that a word response was not included in Webster's lexicon, a transcription was inserted by the experimenter and the frequency was determined from the original Kucera and Francis count. When no frequency count was available, the response was given a frequency of 1. In the absence of familiarity ratings for these words, a value of 7 was assigned to the responses. For nonwords,

transcriptions based on the orthography were inserted by the experimenter and frequencies and familiarity ratings were set to 0. Upon completion of the editing of the master list, percentages correct (hereafter, "scores") for each word at each S/N ratio were computed. Responses were scored as correct if the phonetic transcription constituted an identical match to the target word or if the response was flagged with an "f," indicating that the response was an inflected form of the target word, or flagged with an "h," indicating a homophone.

Computation of Lexical Statistics

Similarity neighborhood statistics for each of the phonetic transcriptions of each of the 882 stimulus words were recomputed based on Webster's lexicon. The method of computing the similarity neighborhood statistics was identical to that discussed in Chapter 1, with two exceptions. These two exceptions were designed to ensure that the lexical items included in the similarity neighborhoods of the stimulus words were at least those produced as responses in the identification experiment. First, subject responses not originally present in the lexicon were added to the lexicon. Second, the familiarity rating criterion for inclusion in the neighborhood analysis was set at 5.5. That is, only words with a familiarity rating of 5.5 or above were considered as possible neighbors. Inspection of the familiarity ratings of subject responses revealed that the majority of responses had familiarity ratings of 5.5 or above. As previously discussed, adoption of a rating criterion for inclusion in the neighborhood analysis helped to ensure that words unknown to a majority of subjects were not included as neighbors. All subsequent analyses were based on this augmented lexicon with a rating criterion of 5.5.

Aside from these two modifications, similarity neighborhoods were computed as described in Chapter 1. The transcription of a given stimulus word was compared with each transcription in Webster's lexicon. A neighbor was defined as a word that could be converted to the target word itself by a one phoneme insertion, deletion, or substitution in any position.

Neighborhood Variables

Three variables from the neighborhood analysis were of interest: (1) the frequency of the stimulus word itself, or stimulus word frequency; (2) the number of words in a given neighborhood, or the density of the neighborhood; and (3) the mean frequency of the neighbors, or mean neighborhood frequency. Two variants of the frequency statistics were employed: raw frequencies and Standard Frequency Indices (SFI's). (SFI's are log transforms of the raw frequencies computed according to the formula, $SFI = 40 + 10 \times (\log(\text{no. occurrences}/\text{corpus size}) + 10)$; Carroll, 1970; see also Whaley, 1978). Mean neighborhood frequency was computed both on the raw frequencies of the neighbors and on the SFI's of the neighbors. Word frequency was also examined in terms of both raw frequencies and SFI's.

In addition to the neighborhood variables, duration of the target word in msec was also included in the following analyses. This variable was included to determine the extent to which the absolute physical duration of the target word contributed to identification performance. The six variables and their mnemonics are listed in Table 2.1.

Insert Table 2.1 about here

Results

Analyses over All S/N Ratios

The scores for the 882 words at each S/N ratio were averaged across all S/N ratios, combined with the six variables of interest, and subjected to correlation and regression analyses. Summary statistics for each variable as well as the identification scores are shown in Table 2.2. Means, standard deviations, and minimum and maximum values are shown for each variable.

Insert Tables 2.2 and 2.3 about here

Correlation Analyses

The correlations of the word scores with each of the six variables are shown in Table 2.3. Significant correlations are indicated by an asterisk. All but one correlation was significant at the 0.05 level, namely NHF-SFI. Stimulus word frequency based on SFI's produced the highest correlation, $r = 0.2465$, followed by stimulus duration, $r = 0.1791$. Small but significant correlations were also obtained for stimulus word frequency based on raw frequencies (SWF-RAW), $r = 0.0850$, neighborhood density (DEN), $r = -0.0840$, and mean neighborhood frequency based on raw frequencies (NHF-RAW), $r = -0.0979$.

Aside from the correlations of the variables with the word scores, two additional correlations among the predictor variables themselves are of interest: Both density and mean neighborhood frequency based on SFI's showed moderate negative correlations with the duration variable ($r = -0.3185$, $p < 0.05$, and $r = -0.3081$, $p < 0.05$, respectively). Essentially, given the high intercorrelation of neighborhood frequency and density ($r = 0.9908$), this correlation suggests that as neighborhood density decreases, the duration of the stimulus word increases. Because neighborhood density indicates the degree of phonemic overlap among words, and thus the degree to which the constituent phonemes of a given word are shared by other words, the correlation of density and duration implies that phonemes shared by a large number of words tend to be shorter in duration. That is, increases in density, and the concomitant increases in phoneme frequency, are accompanied by decreases in stimulus duration. This correlation indicates that more frequently used phonemes tend to be shorter in duration than less frequent phonemes (see Dewey, 1970; Miller, 1951; Zipf, 1935).

Regression Analyses

To determine the degree to which the combined effects of these variables contributed to the overall proportion of variance accounted for, a hierarchical multiple regression analysis was performed. Because of the superior correlation of stimulus word frequency based on SFI's, the dependent variables based on raw frequencies were excluded from the regression analysis. The four remaining variables were entered in the following order: (1)

Table 2.1

Neighborhood Analysis: Lexical Statistics. Variables and mnemonics.

VARIABLE	MNEMONIC
1. Raw stimulus word frequency	SWF-RAW
2. SFI stimulus word frequency	SWF-SFI
3. Neighborhood Density	DEN
4. Mean neighborhood frequency based on raw frequencies	NHF-RAW
5. Mean neighborhood frequency based on SFI's	NHF-SFI
6. Duration of stimulus word in msec	DUR

Table 2.2

Neighborhood Analysis: Lexical Statistics. Summary statistics for variables for all signal-to-noise ratios.

VARIABLE	MEAN	STANDARD DEVIATION	MINIMUM VALUE	MAXIMUM VALUE
-----	-----	-----	-----	-----
1. SWF-RAW	129.1564	556.2139	1.0000	10595.0000
2. SWF-SFI	52.1884	8.4655	40.0000	80.2500
3. DEN	16.5011	6.7651	1.0000	35.0000
4. NHF-RAW	205.9523	349.8718	2.6000	4389.0000
5. NHF-SFI	52.0241	3.1903	42.5508	66.4250
6. DUR	0.4968	0.0691	0.3012	0.7358
7. SCORE	0.4969	0.2218	0.0000	1.0000

Table 2.3

Neighborhood Analysis: Lexical Statistics. Correlations between
identification scores and predictor variables for all signal-to-noise ratios.

VARIABLE	r
1. SWF-RAW	0.0850*
2. SWF-SFI	0.2465*
3. DEN	-0.0840*
4. NHF-RAW	-0.0979*
5. NHF-SFI	-0.0258
6. DUR	0.1791*

* $p < 0.05$

SWF-SFI, (2) DEN, (3) NHF-SFI, and (4) DUR.

The results of the stepwise multiple regression analysis are shown in Table 2.4. The variables are listed in the order they were entered into the regression equation. Multiple R's and multiple R²'s are shown at each step. Also shown is the change in R² at each step, indicating the unique contribution to the total proportion of variance accounted for by that variable. Finally, the change in R² at each step as a function of the explained variance is shown in the rightmost column. These values indicate the proportion of the explained variance (i.e., the multiple R² at the final step) contributed by each variable. For example, the change in R² as a function of the explained variance for the variable SWF-SFI was computed by dividing the change in R² as a function of the total variance (0.0607) by the total explained variance (0.1105), which indicates that the SWF-SFI variable accounted for 54.93% of the explained variance.

Insert Table 2.4 about here

As shown in Table 2.4, the four independent variables combined produced a multiple R = 0.3324 and a multiple R² = 0.1105. Thus, all four variables combined accounted for approximately 11% of the total variance. Stimulus word frequency based on SFI's accounted for 6.07% of the total variance and 54.93% of the explained variance. Neighborhood density accounted for approximately 1% of the total variance and 8.42% of the explained variance. Mean neighborhood frequency based on SFI's accounted for 1.13% of the total variance and 10.23% of the explained variance. And, stimulus duration accounted for 2.91% of the total variance and 26.33% of the explained variance. Thus, although the total proportion of variance explained was small, the proportion of the explained variance accounted for by the two neighborhood variables combined (DEN and NHF-SFI) was approximately 18%.

Analyses at Each S/N Ratio

Correlation Analyses

To evaluate the role of the frequency and neighborhood structure variables independently of S/N ratio, the predictor variables were submitted to separate correlation and regression analyses at each S/N ratio. The correlations of each of the six variables with identification scores are shown in Table 2.5.

Insert Table 2.5 about here

Significant correlations of SWF-RAW with identification performance were obtained at the +15 and +5 S/N ratios, and significant correlations for SWF-SFI were obtained at all S/N ratios. Note that for SWF-SFI, the correlations increased with decreasing S/N ratio, suggesting an increase in frequency-biased responding with increased stimulus degradation. Only two correlations were significant for the neighborhood variables. Neighborhood density and mean neighborhood frequency based on raw frequencies resulted in small negative correlations at the +5 and -5 S/N ratios, respectively.

Table 2.4

Neighborhood Analysis: Lexical Statistics. Results of hierarchical multiple regression analysis for all signal-to-noise ratios. Shown are the multiple R's, multiple R²'s, changes in R² as a function of the total variance, and changes in R² as a function of the explained variance at each step.

VARIABLE	R	R ²	ΔR^2 (TOTAL)	ΔR^2 (EXPLAINED)
1. SWF-SFI	0.2465	0.0607	0.0607	0.5493
2. DEN	0.2647	0.0701	0.0093	0.0842
3. NHF-SFI	0.2853	0.0814	0.0113	0.1023
4. DUR	0.3324	0.1105	0.0291	0.2633

Table 2.5

Neighborhood Analysis: Lexical Statistics. Correlations between
identification scores and predictor variables for each signal-to-noise ratio.

VARIABLE	r		
	+15 S/N	+5 S/N	-5 S/N
1. SWF-RAW	0.0750*	0.0790*	0.0475
2. SWF-SFI	0.1711*	0.2145*	0.2253*
3. DEN	-0.0228	-0.1097*	-0.0653
4. NHF-RAW	-0.0281	-0.0180	-0.0936*
5. NHF-SFI	-0.0466	-0.0131	-0.0002
6. DUR	0.1429*	0.1961*	0.0727*

* $p < 0.05$

Finally, duration correlated positively with performance at each S/N ratio.

Regression Analyses

Hierarchical multiple regression analyses were performed for each S/N ratio separately. Again those variables based on raw frequencies were excluded from the regression analyses. The ordering of the variables was identical to that for the overall analysis. The results for each S/N ratio are shown in Table 2.6.

Insert Table 2.6 about here

For the +15 S/N ratio, a multiple $R = 0.2525$ and a multiple $R^2 = 0.0637$ were obtained. All variables except neighborhood density (DEN) contributed significantly to the multiple R . SWF-SFI uniquely accounted for 2.93% of the total variance and 46.00% of the explained variance. NHF-SFI uniquely accounted for 1.11% of the total variance and 17.43% of the explained variance. And, DUR account for 2.43% of the total variance and 36.73% of the explained variance.

For the +5 S/N ratio, all variables contributed significantly to the proportion of variance accounted for, rendering a multiple $R = 0.3147$ and a multiple $R^2 = 0.0991$. In terms of total variance accounted for, SWF-SFI accounted for 4.60% of the variance, DEN 1.46% of the variance, NHF-SFI 0.68% of the variance, and DUR 3.17% of the variance. In terms of the proportions of explained variance, SWF-SFI accounted for 46.42%, DEN 14.73%, NHF-SFI 6.86%, and DUR 31.99%.

Finally, for the -5 S/N ratio, all variables except DUR contributed significantly to the proportion of variance accounted for, rendering a multiple $R = 0.2488$ and a multiple $R^2 = 0.0619$. SWF-SFI contributed 5.07% to the total variance and 81.91% to the explained variance. DEN contributed 0.59% and 9.53% to the total and explained variance, respectively. NHF-SFI contributed 0.53% to the total variance and 8.56% to the explained variance.

DISCUSSION

The multiple regression analysis performed on the data collapsed across S/N ratios revealed significant effects of stimulus word frequency, neighborhood density, mean neighborhood frequency, and stimulus duration. Although the effects of these variables in terms of the total proportion of variance accounted for were small, they nevertheless revealed that stimulus word frequency and neighborhood structure have predictable and measureable effects on identification performance. As stimulus word frequency increased, so too did identification performance. Furthermore, as neighborhood density and mean neighborhood frequency increased, identification performance dropped, indicating that both the number and nature of competitors affects identification performance. Although the proportions of the total variance accounted for by the neighborhood variables was small, their contribution to the overall proportion of explained variance was substantial in the overall analysis. In particular, neighborhood density and mean neighborhood frequency together accounted for over 18% of the explained variance.

Table 2.6

Neighborhood Analysis: Lexical Statistics. Results of hierarchical multiple regression analysis for each signal-to-noise ratio. Shown are the multiple R's, multiple R²'s, changes in R² as a function of the total variance, and changes in R² as a function of the explained variance at each step.

+15 S/N

VARIABLE	R	R ²	ΔR^2 (TOTAL)	ΔR^2 (EXPLAINED)
1. SWF-SFI	0.1711	0.0293	0.0293	0.4600
2. DEN		(not significant)		
3. NHF-SFI	0.2008	0.0403	0.0111	0.1743
4. DUR	0.2525	0.0637	0.0243	0.3673

+5 S/N

VARIABLE	R	R ²	ΔR^2 (TOTAL)	ΔR^2 (EXPLAINED)
1. SWF-SFI	0.2145	0.0460	0.0460	0.4642
2. DEN	0.2462	0.0606	0.0146	0.1473
3. NHF-SFI	0.2595	0.0674	0.0068	0.0686
4. DUR	0.3147	0.0991	0.0317	0.3199

-5 S/N

VARIABLE	R	R ²	ΔR^2 (TOTAL)	ΔR^2 (EXPLAINED)
1. SWF-SFI	0.2253	0.0507	0.0507	0.8191
2. DEN	0.2380	0.0566	0.0059	0.0953
3. NHF-SFI	0.2488	0.0619	0.0053	0.0856
2. DUR		(not significant)		

In addition to the effects of the neighborhood variables, it is of interest that a significant effect of word frequency was observed in the overall analysis. Again, in terms of the total proportion of variance accounted for, the effect was small. However, given that all stimuli were highly familiar to subjects, the observation of a frequency effect indicates that whereas the effects of word frequency and subjective familiarity are highly correlated, they can be dissociated. That is, there appears to be a word frequency effects above and beyond the familiarity of the items. This finding suggests that despite the judged familiarity of the word, degree of exposure somehow influences identification performance, implicating a mechanism sensitive to the degree of repeated exposure (see also Nusbaum & Dedina, 1985).

The regression analyses at each S/N ratio also revealed effects of stimulus word frequency, neighborhood density, and mean neighborhood frequency. Significant effects of mean neighborhood frequency were obtained at each S/N ratio, and effects of neighborhood density were observed at the +5 and -5 S/N ratios. Although the contributions of these variables to the total proportion of variance accounted for were small, their contributions to the explained variance were substantial, ranging from 17% to 21%. Thus, within the levels of the proportions of variance accounted for by the combination of all of the variables examined, the neighborhood structural variables had measurable effects on identification performance.

Summary and Conclusions

The results from the present study provide reliable, albeit modest evidence, for the role of neighborhood structure (i.e., neighborhood density and mean neighborhood frequency) in the perceptual identification of words. In general, the proportions of total variance accounted for by each of the independent variables were small. However, the neighborhood variables, as well as stimulus word frequency and duration, each contributed substantially to the proportion of variance explained by the combination of each of the variables. In particular, the results demonstrated that words in high density neighborhoods (excluding those presented at +15 S/N ratio) were identified less accurately than those in low density neighborhoods, and words in high frequency neighborhoods were identified less accurately than those in low frequency neighborhoods. The present study thus strongly suggests that neighborhood structure is an important determinant of identification performance and therefore provides strong motivation for a more detailed analysis of the roles of stimulus word frequency and neighborhood structure in the perceptual identification of words. Thus, in order to more closely examine neighborhood structure and frequency, a more sophisticated means of determining neighborhood membership was employed in Chapter 3. In particular, effects of basic segmental intelligibility and confusability were incorporated in the computation of similarity neighborhoods in an attempt to better predict identification performance.

CHAPTER THREE

NEIGHBORHOOD PROBABILITY RULES

As stated in the conclusion to Chapter 2, predicted effects of neighborhood structure (i.e., neighborhood density and frequency) were observed for various S/N ratios, although these effects were small. Upon reconsideration of the algorithm used to compute similarity neighborhoods in Chapter 2, it is perhaps not too surprising that neighborhood structure played a relatively small role in accounting for the observed variance. Recall that a neighbor of a stimulus word was defined as a word that could be transformed into the target word itself by a one phoneme substitution, insertion, or deletion. This is a somewhat imprecise measure of similarity. In particular, no index of basic segmental confusability is included in such a determination of neighbors. For example, the stimulus word /k@t/ would have neighbors such as /p@t/, /r@t/, /k@d/, and /k@n/, among others, and all would be equivalent as neighbors. However, because all segments are not equally confusable (see, for example, Miller and Nicely, 1958), each of these neighbors will vary according to the degree to which they are confusable with the stimulus word. Thus, the simple computation of neighborhood membership based on one phoneme substitutions, deletions, and insertions most likely constitutes too imprecise a measure of similarity to predict adequately the identification of words in noise.

In addition to the failure to include measures of basic segmental confusability, the algorithm employed in Chapter 2 included no direct measure of the relations among the frequencies of the stimulus word and its neighbors. Although mean frequency of the neighbors was included in the analyses as an index of the influence of neighbor frequencies on the identification of the stimulus word, a more precise measure of neighborhood frequency should perhaps include a determination of the relation of the frequency of the stimulus word to the frequencies of the words in its neighborhood. That is, how does the frequency of the target word in relation to the frequencies of its neighbors influence identification performance?

A final shortcoming of the present algorithm for computing neighborhood membership discussed in Chapter 2 was the failure to take into account the combined effects of segmental confusability and frequency relations. On a priori grounds, segmental confusability and frequency relations should act in concert. A high frequency neighbor that is also highly confusable with the stimulus word should prove to be a more formidable competitor than a high frequency word that is less confusable with the stimulus word. Likewise, low frequency neighbors may be quite strong competitors if highly confusable with the stimulus word, but much less so than when less confusable with the target word. If one adds to these variables a consideration of the frequency and basic intelligibility of the stimulus word itself in relation to its neighbors, a more complex and presumably more powerful measure of neighborhood structure should be obtained. Thus, as a further step toward specifying more precisely the role of neighborhood structure in auditory word identification, an attempt was undertaken to determine a neighborhood probability rule that would predict word identification based on neighborhood structure.

In order to devise a rule to predict word identification performance based on stimulus word frequency, word intelligibility, neighbor frequency, and neighbor confusability, an independent means of determining segmental intelligibility and confusability was required. Therefore, a second, independent experiment was performed in order to obtain confusion matrices for

all possible initial consonants, vowels, and final consonants. The confusion matrices were then used to investigate the combined effects of stimulus intelligibility and neighborhood confusability.

EXPERIMENT

Method

Stimuli

All possible consonant-vowel (CV) sequences and all possible vowel-consonant (VC) sequences were generated as follows: The 23 possible initial consonants (including a null consonant) were combined with the 15 possible vowels to produce 345 CV sequences. Likewise, the 22 possible final consonants (again including a null consonant) were combined with the 15 possible vowels to produce 330 VC sequences. Initial consonants, vowels, and final consonants are shown in Table 3.1. (The IPA counterparts of the phonetic symbols are shown in Table 1.2.)

Insert Table 3.1 about here

The 345 CV sequences and 330 VC sequences were randomized and recorded by the same male speaker who read the words for the perceptual identification study reported in Chapter 2. The procedure for recording, digitizing, and editing the CV and VC syllables was identical to the procedure reported in Chapter 2. All CV and VC syllables were also matched in overall amplitude to the average amplitude of the words used in the word identification experiment.

Subjects

One-hundred and twenty subjects participated in partial fulfillment of an introductory psychology course. All subjects were native English speakers and reported no history of speech or hearing disorders.

Design

The CV and VC syllables were presented for identification of either the consonant alone or the vowel alone at the three S/N ratios used in the word identification experiment: +15 dB, +5 dB, and -5 dB SPL. Syllable type (CV or VC) and response type (consonant response or vowel response) were between-subject factors. S/N ratio was a within-subject factor. This resulted in four conditions in which (1) subjects heard CV syllables and attempted to identify the consonant, (2) subjects heard CV syllables and attempted to identify the vowel, (3) subjects heard VC syllables and attempted to identify the consonant, or (4) subjects heard VC syllables and attempted to identify the vowel. A given subject participated in only one of the four conditions. The experimental design is shown in Figure 3.1.

Table 3.1

Confusion Matrix Experiment. Initial consonants, vowels, and final consonants composing CV and VC syllables.

 INITIAL CONSONANTS

VOWELS

FINAL CONSONANTS

p
 t
 k
 b
 d
 g
 S
 C
 f
 v
 s
 T
 D
 J
 h
 m
 n
 w
 l
 r
 j
 Ø*

i
 I
 e
 E
 @
 ^
 a
 c
 o
 U
 u
 Y
 W
 O

p
 t
 k
 b
 d
 g
 S
 C
 f
 v
 s
 T
 D
 J
 Z
 m
 n
 G
 l
 r
 Ø*

 *null phoneme

Insert Figure 3.1 about here

S/N ratio varied randomly throughout a given session and was counterbalanced across conditions. No subject heard a given stimulus more than once. Ten observations per syllable type, response type, and S/N ratio were obtained. Manipulation of signal and noise levels was identical to that reported in Chapter 2.

Procedure

Method of stimulus presentation was identical to that employed in the word identification experiment. However, the procedure for data collection was modified. At the beginning of each session, subjects were given answer booklets containing numbered blanks. At the top of each page of the answer booklet, a key was provided. The key indicated the particular letter or letters the subject was to use as responses. Examples of words for each sound were also provided in the key in ambiguous cases. In the conditions in which a vowel response was required, all vowels were accompanied by example words. Examples of the keys for each condition are shown in Table 3.2. (Note that in order to distinguish /T/ and /D/, subjects were instructed to circle the "th" corresponding to /D/ but not the "th" corresponding to /T/.)

Insert Table 3.2 about here

Subjects were instructed that they would hear a CV syllable or a VC syllable (depending on the condition) embedded in noise. The subjects were instructed to identify either the consonant or vowel (again depending on the condition) and to indicate the sound they thought they heard in the appropriate blank on their answer booklets. Subjects were instructed to refer to the key at the top of each page of the answer booklets in order to determine the letter or letters that should be associated with each sound.

The experiment consisted of three phases: (1) familiarization, (2) practice, and (3) testing. In the familiarization phase, subjects heard examples of the stimuli at 75 dB SPL in the absence of masking noise. The stimuli were presented in the order in which they appeared in the key at the top of their answer booklets. Subjects were instructed simply to listen to each syllable and to familiarize themselves with the particular letter or letters that were to be used for each sound. For the CV-syllable consonant-response condition, subjects heard each of the consonants preceding the vowels /u/, /i/, and /a/ in the order in which they were listed in the answer key. For the VC-syllable consonant-response condition, subjects heard each of the consonants following the same vowels. For the two vowel-response conditions, subjects heard each of the vowels in isolation in the order in which they appeared in the answer key. The series of vowels was repeated three times.

In the practice phase of the experiment, subjects heard one token of each initial consonant, vowel, or final consonant, depending on the condition. In this portion of the experiment, the syllables were presented in noise, with S/N ratio varying from trial to trial. In the practice phase, subjects were

SYLLABLE TYPE:

RESPONSE TYPE:




Figure 3.1. Experimental design for the confusion matrix experiment.

Table 3.2

Confusion Matrix Experiment. Response keys for initial consonant, vowel, and final consonants conditions.

INITIAL CONSONANTS:

b ch d f g h j k l m n p r s sh t (th) th v
get jet then thin
w y z Ø

VOWELS:

EE I E AI A AA OW IY U AW OI O OO UU ER
feet hit peck maid sad sock shout hide bud pawed void toad good food bird

FINAL CONSONANTS:

b ch d f g ge j k l m n p r s sh t th (th)
dog beige judge math bathe
v z Ø

instructed to respond with their best guess on each trial. In the event that they were completely unable to identify a consonant or vowel, subjects were instructed to respond with a check mark in the appropriately numbered blank.

Following practice, the instructions were summarized and any procedural questions were answered. The subjects were then instructed that they would hear a much longer list of syllables. They were again encouraged to provide their best guesses on each trial. Subjects were also instructed that although some of the syllables would form real words, they were to treat all syllables on an equal basis and to focus on either the consonant or the vowel, ignoring the irrelevant portion of the syllable. During both the practice and testing phases of the experiment, subjects were provided with four seconds in which to respond. Short breaks were provided in order to allow time to turn the pages in the answer booklets. An experimental session lasted approximately one hour.

Results

Confusion matrices were constructed from the data for initial consonants, vowels, and final consonants at each S/N ratio. The confusion matrices were collapsed across the irrelevant consonant or vowel for a given condition. For example, the phoneme /p/ occurred in initial position in 15 syllables (i.e., before each of the 15 vowels). The data for syllable-initial /p/ were therefore collapsed across the 15 different syllables in which /p/ occurred. Given 10 observations per syllable, collapsing across the 15 syllables in which /p/ occurred in initial position rendered 150 possible responses for /p/. Thus, for each initial and final consonant, a total of 150 responses was obtained. Confusion matrices for the vowels were constructed in a somewhat different manner: Two separate vowel confusion matrices were first constructed, one from the CV syllables and one from the VC syllables. For the vowel confusion matrices based on the CV syllables, 230 observations per vowel were obtained; for the VC syllables, 220 observations per vowel were obtained. Recall that for the CV syllables, each vowel was paired with 23 initial consonants. Given 10 observations per syllable, collapsing across initial consonants renders 230 observations per vowel. Likewise, for the VC syllables, in which each vowel was paired with 22 final consonants, 220 observations per vowel were obtained. The confusion matrices for the vowels based on the CV and VC syllables were then combined, resulting in 450 observations for each vowel.

The confusion matrices for initial consonants for each S/N ratio are shown in Tables 3.3, 3.4, and 3.5. The confusion matrices for final consonants for each S/N ratio are shown in Tables 3.6, 3.7, and 3.8. And, the confusion matrices for the vowels for each S/N ratio are shown in Tables 3.9, 3.10, and 3.11.

Insert Tables 3.3 through 3.11 about here

Stimuli are represented by the rows and responses by the columns. The numbers in each cell represent the raw frequencies of identification. Probabilities for initial and final consonants can be obtained by dividing each cell frequency by 150. Probabilities for vowels can be obtained by dividing each cell frequency by 450.

Table 3.3

Confusion matrix for initial consonants for the +15 signal-to-noise ratio. Stimuli are represented by the rows, responses by the columns. Also shown are the column sums.

		+15 S/N																						
		p	t	k	b	d	g	C	J	s	S	z	f	T	v	D	h	n	m	l	r	w	y	Ø
p	85	27	27	3	0	1	0	0	0	0	0	0	0	0	0	1	6	0	0	0	0	0	0	0
t	0	145	0	0	0	2	1	0	0	0	0	0	0	1	0	0	1	0	0	0	0	0	0	0
k	1	3	139	0	0	0	3	0	1	0	0	1	0	1	0	0	0	0	0	0	0	0	0	1
b	1	3	0	82	13	2	0	0	1	0	0	3	9	7	21	0	0	0	0	0	1	1	0	6
d	0	0	0	0	147	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0
g	1	0	0	3	2	123	0	1	0	0	0	0	0	1	0	0	0	1	0	4	2	11	1	
C	0	3	0	0	0	0	145	1	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	
J	0	0	0	0	0	12	4	133	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	
s	0	4	2	0	0	0	1	0	127	2	12	1	0	0	0	0	0	0	1	0	0	0	0	
S	0	0	0	0	0	0	24	0	2	122	1	0	0	1	0	0	0	0	0	0	0	0	0	
z	0	1	0	0	14	1	2	1	4	1	114	0	1	1	7	0	0	0	1	0	0	0	2	
f	26	5	5	13	3	1	0	0	5	0	0	66	6	6	9	2	0	0	0	0	2	1	0	
T	14	8	0	23	2	1	0	1	4	1	1	38	28	4	15	4	0	0	0	0	3	1	2	
v	0	2	0	15	5	5	0	0	1	0	0	0	17	51	46	0	0	1	2	3	1	0	1	
D	0	3	0	6	17	4	0	0	1	0	5	0	21	14	76	0	1	0	1	0	0	0	1	
h	49	23	26	3	0	0	1	0	0	0	0	6	0	0	0	30	0	0	0	0	1	0	11	
n	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	135	12	1	1	0	0	0	
m	0	0	0	0	0	0	0	0	1	0	0	0	0	0	1	0	4	142	2	0	0	0	0	
l	2	0	0	10	1	3	0	0	0	0	0	1	0	3	1	0	0	3	121	1	3	0	1	
r	0	1	0	0	0	5	0	0	0	0	0	0	1	0	1	0	1	0	23	105	12	1	0	
w	2	1	0	1	0	4	0	0	0	0	0	1	0	0	0	0	1	0	16	33	86	5	0	
y	2	0	0	1	0	2	0	1	0	0	0	0	1	0	0	1	1	1	2	6	1	128	3	
Ø	5	7	2	6	0	0	0	0	1	0	1	4	2	1	0	7	0	1	3	2	0	0	108	
		188	236	201	166	204	167	181	138	148	126	134	121	87	91	178	51	143	161	174	157	112	148	138

Table 3.4

Confusion matrix for initial consonants for the +5 signal-to-noise ratio. Stimuli are represented by the rows, responses by the columns. Also shown are the column sums.

+5 S/N																							
	p	t	k	b	d	g	C	J	s	S	z	f	T	v	D	h	n	m	l	r	w	y	Ø
p	40	55	30	2	0	0	1	0	1	0	0	5	1	3	0	5	0	0	1	0	0	0	6
t	4	103	24	0	3	0	3	0	1	0	0	0	1	0	2	1	0	1	0	3	0	0	4
k	12	20	102	4	1	0	4	1	0	0	0	0	1	0	0	1	0	0	0	1	0	0	3
b	9	6	7	39	30	10	1	1	2	0	0	11	3	9	12	0	0	0	3	1	1	1	4
d	1	6	0	4	114	5	2	3	3	0	1	1	5	0	4	0	0	0	0	0	0	0	1
g	1	0	1	7	21	74	0	3	3	0	0	1	0	8	1	2	0	0	0	2	5	17	4
C	0	7	1	0	0	0	138	0	0	2	0	0	0	0	0	0	1	0	0	0	0	0	1
J	0	1	0	0	3	10	4	130	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1
s	3	18	1	3	7	0	9	3	55	1	16	11	13	0	7	0	1	1	0	0	0	1	0
S	1	0	0	0	0	0	80	3	2	62	0	1	0	0	0	0	0	0	0	0	0	0	1
z	0	1	1	2	24	15	0	4	9	0	45	0	12	9	21	0	0	0	1	3	1	2	0
f	14	16	9	12	2	1	5	0	11	0	4	54	9	3	6	1	0	1	1	0	0	0	1
T	11	18	7	11	6	5	2	0	12	0	2	36	20	4	10	3	0	1	0	0	0	0	2
v	1	2	1	17	15	7	0	0	0	1	3	4	17	34	31	0	1	1	9	2	2	0	2
D	4	2	2	6	12	7	3	0	4	1	14	2	16	36	45	1	2	2	5	4	1	0	1
h	29	42	33	1	0	0	2	0	0	0	0	5	2	1	0	19	2	0	0	0	0	0	14
n	0	1	1	0	0	0	1	0	0	0	0	0	0	1	0	2	118	26	0	0	0	0	0
m	0	0	0	0	0	0	0	0	2	0	0	0	0	0	0	0	20	126	0	0	0	1	1
l	3	0	1	11	1	4	1	1	0	0	0	5	0	0	0	2	2	3	107	3	1	2	3
r	0	1	1	2	1	7	0	1	0	0	1	1	0	0	0	0	0	1	30	68	14	20	2
w	0	4	1	2	1	2	1	1	1	0	1	1	0	1	2	2	1	0	14	22	75	17	1
y	0	0	0	2	1	5	0	0	1	0	1	0	0	3	0	1	3	1	3	6	5	108	10
Ø	3	5	10	3	1	1	0	1	0	0	0	4	2	1	2	11	1	2	3	3	1	2	94

	136	308	233	128	243	153	257	152	107	67	88	143	102	93	143	51	152	166	177	118	106	171	156

Table 3.5

Confusion matrix for initial consonants for the -5 signal-to-noise ratio. Stimuli are represented by the rows, responses by the columns. Also shown are the column sums.

-5 S/N																							

	p	t	k	b	d	g	C	J	s	S	z	f	T	v	D	h	n	m	l	r	w	y	Ø
p	14	53	26	2	3	4	6	2	5	2	0	7	1	0	0	5	2	3	0	3	0	1	11
t	12	41	28	2	8	3	9	1	4	2	1	8	3	3	4	2	2	0	1	1	0	1	14
k	6	46	36	2	5	3	9	0	12	0	1	5	1	2	0	3	1	2	0	0	0	0	16
b	5	19	13	9	23	5	7	3	12	0	4	3	4	5	4	3	4	2	3	2	3	1	16
d	4	17	6	4	38	19	5	8	9	2	3	2	6	1	6	2	1	1	4	3	0	0	9
g	0	6	7	1	19	24	1	25	2	0	7	2	0	3	1	0	1	0	6	1	7	25	12
C	6	21	13	0	5	6	58	9	6	3	2	2	1	1	1	2	1	0	3	0	1	2	7
J	1	7	3	2	25	10	5	60	6	1	5	0	0	2	2	2	0	0	2	0	4	5	8
s	9	18	16	2	16	6	11	1	34	1	6	5	5	2	3	1	0	0	1	0	1	2	10
S	1	12	9	0	5	8	55	29	9	8	1	2	0	1	3	0	2	0	0	0	0	1	4
z	2	4	0	7	30	24	4	21	6	1	7	0	5	6	4	0	0	0	2	3	7	13	4
f	10	24	13	2	4	2	4	2	33	2	3	15	5	6	6	1	0	1	0	4	0	2	11
T	7	34	17	4	5	4	8	2	16	3	4	19	3	4	9	2	0	0	0	0	1	1	7
v	3	8	7	11	24	8	1	10	7	1	13	6	4	10	12	1	2	1	2	1	4	2	12
D	2	11	4	7	21	17	2	13	6	0	12	6	2	3	11	0	2	1	2	7	6	4	11
h	8	35	37	2	0	1	4	2	6	1	0	4	0	0	16	0	0	3	2	0	1	28	
n	0	2	2	1	1	0	0	0	4	1	1	0	0	1	0	0	61	59	5	0	1	2	9
m	0	9	1	0	0	1	0	0	10	0	0	2	0	1	0	1	40	65	1	1	1	7	10
l	2	6	1	3	2	1	1	3	0	0	1	2	2	2	3	1	4	3	65	15	9	13	11
r	1	2	0	1	0	4	1	5	1	0	2	2	1	1	1	0	1	0	24	13	20	62	8
w	1	2	1	0	4	6	0	9	2	0	3	1	0	3	2	2	2	3	20	13	18	42	16
y	1	3	2	0	1	5	1	15	0	0	2	0	2	1	1	0	1	0	7	8	12	83	5
Ø	5	21	11	3	5	3	0	2	4	0	0	5	0	0	3	10	4	2	3	5	1	4	59

	100	401	253	65	244	164	192	222	194	28	78	98	45	58	76	54	131	143	154	82	96	274	298

Table 3.6

Confusion matrix for vowels for the +15 signal-to-noise ratio. Stimuli are represented by the rows, responses by the columns. Also shown are the column sums.

		+15 S/N															
		i	I	E	e	@	a	W	Y	^	c	o	U	u	R	∅*	
i	408	3	17	3	2	0	1	3	0	3	5	2	0	2	1	0	
I	5	368	29	5	4	3	1	10	2	2	1	1	5	6	7	1	
E	1	20	372	12	15	1	0	2	9	1	2	4	0	3	6	2	
e	9	4	10	371	13	9	2	5	1	5	3	2	6	3	6	1	
@	2	1	32	17	362	12	12	1	0	3	0	0	1	2	4	1	
a	0	6	5	0	35	173	13	1	5	184	5	10	6	2	3	2	
W	1	1	1	2	2	0	385	1	3	19	8	13	8	4	2	0	
Y	0	26	0	38	3	0	1	372	0	0	6	1	1	1	1	0	
^	1	4	39	0	26	7	1	0	297	9	1	23	25	5	10	2	
c	0	3	4	2	12	82	16	1	8	254	6	41	13	1	5	2	
o	1	3	1	3	2	2	17	6	1	5	386	10	9	2	2	0	
U	1	1	0	0	0	0	16	0	1	1	8	392	9	13	8	0	
u	2	3	4	3	5	2	7	3	101	3	8	19	240	43	7	0	
R	2	1	2	1	3	0	12	0	10	2	1	10	81	323	1	1	
R	0	3	10	0	1	0	0	1	11	1	2	5	5	5	405	1	
		433	447	526	457	485	291	484	406	449	492	442	533	409	415	468	13

*∅ indicates "no response"

Table 3.7

Confusion matrix for vowels for the +5 signal-to-noise ratio. Stimuli are represented by the rows, responses by the columns. Also shown are the column sums.

+5 S/N																
	i	I	E	e	@	a	W	Y	^	c	0	o	U	u	R	∅*
i	415	3	4	4	0	0	1	5	2	0	0	2	0	8	5	1
I	8	265	34	9	5	0	6	5	18	5	7	6	43	28	9	2
E	5	35	299	28	36	7	2	5	7	2	2	5	5	3	7	2
e	9	6	11	372	15	3	0	4	1	0	3	14	4	3	4	1
@	2	3	36	7	354	7	12	2	4	13	1	3	0	1	5	0
a	0	1	4	5	41	170	29	15	16	109	7	22	21	2	3	5
W	0	0	2	2	8	1	399	3	0	17	5	5	3	3	1	1
Y	2	16	1	24	7	5	0	370	6	4	5	0	4	2	4	0
^	0	8	28	2	36	21	13	2	253	6	7	35	14	9	13	3
c	1	1	1	0	13	54	14	0	8	294	11	35	9	1	4	4
0	3	8	6	13	4	1	6	7	13	2	292	27	48	9	11	0
o	2	3	4	15	4	1	26	1	14	4	56	272	13	8	25	2
U	1	26	23	3	2	1	5	1	84	5	4	13	206	43	28	5
u	55	15	13	3	4	1	5	4	10	1	2	4	104	224	5	0
R	1	7	27	2	2	1	2	3	31	4	17	13	12	4	322	2

	504	397	493	489	531	273	520	427	467	466	419	456	486	348	446	28

*∅ indicates "no response"

Table 3.8

Confusion matrix for vowels for the -5 signal-to-noise ratio. Stimuli are represented by the rows, responses by the columns. Also shown are the column sums.

-5 S/N																

	i	I	E	e	@	a	W	Y	^	c	0	o	U	u	R	∅*
i	290	21	17	2	3	2	3	5	11	3	1	4	20	57	7	4
I	43	189	42	23	12	2	2	5	20	0	11	4	62	15	16	4
E	5	27	177	62	37	11	5	5	27	15	4	19	27	8	20	1
e	15	11	44	229	30	8	4	4	5	5	20	43	9	6	15	2
@	1	6	56	9	295	11	31	13	6	14	1	2	0	0	1	4
a	1	10	3	8	41	116	24	54	13	137	6	11	12	2	11	1
W	0	4	15	5	81	22	219	41	14	19	9	7	6	4	1	3
Y	0	4	4	5	31	31	59	233	6	56	5	4	4	1	6	1
^	4	17	71	28	34	15	26	22	119	20	8	39	23	4	17	3
c	2	7	13	8	41	56	27	10	14	131	54	54	15	1	10	7
0	4	32	42	127	20	9	5	2	20	4	34	67	39	11	31	3
o	6	13	23	187	29	6	6	6	18	9	23	88	14	7	12	3
U	13	191	44	18	5	4	4	5	33	4	10	7	57	23	27	5
u	253	28	12	8	5	0	1	7	20	0	4	2	29	74	6	1
R	2	29	41	107	29	10	5	6	39	5	21	43	54	16	42	1

	639	589	604	826	693	303	421	418	365	422	211	394	371	229	222	43

*∅ indicates "no response"

Table 3.9

Confusion matrix for final consonants for the +15 signal-to-noise ratio. Stimuli are represented by the rows, responses by the columns. Also shown are the column sums.

+15 S/N																						
	p	t	k	b	d	g	C	J	s	S	z	f	T	v	D	n	m	l	r	G	Z	Ø
p	94	14	28	1	0	0	1	0	0	0	0	7	2	1	1	0	0	0	0	0	0	1
t	1	146	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2
k	0	1	144	0	0	0	1	0	0	0	0	0	0	0	0	0	1	0	1	0	1	1
b	0	0	0	130	8	10	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	1
d	0	0	0	5	137	4	0	1	0	0	0	1	0	1	0	0	0	1	0	0	0	0
g	0	0	2	3	10	132	0	0	0	0	1	0	0	1	0	0	0	0	0	0	1	0
C	1	0	0	0	0	0	144	1	0	0	0	0	0	0	0	0	0	1	0	0	0	3
J	0	0	0	0	0	3	0	85	0	0	0	0	0	0	0	0	0	0	0	0	62	0
s	0	7	1	0	0	0	0	1	131	2	1	1	0	0	0	0	0	0	0	0	1	5
S	0	0	0	0	0	0	6	0	1	139	0	0	1	0	0	0	0	0	0	0	2	1
z	0	0	0	0	0	0	1	0	19	0	121	0	0	4	0	0	0	2	1	0	2	0
f	16	7	42	0	1	0	4	0	7	2	0	39	5	3	4	0	0	0	5	0	0	15
T	5	38	19	0	0	1	9	0	13	3	0	19	17	1	2	0	0	0	1	0	1	21
v	0	0	0	4	5	25	0	1	4	0	9	2	2	64	17	0	0	4	2	0	0	11
D	0	0	0	5	16	15	0	1	8	0	24	2	7	47	19	0	0	1	0	0	0	5
n	0	1	0	0	3	2	0	0	0	1	1	0	0	1	0	122	2	0	0	12	0	5
m	0	0	0	3	0	2	0	0	1	0	1	0	0	2	0	16	114	0	0	8	0	3
l	0	1	1	3	1	1	0	0	0	0	0	0	0	3	0	0	0	112	3	1	0	24
r	0	2	0	0	0	0	0	1	0	0	2	1	0	3	0	0	0	9	120	0	0	12
G	0	0	0	0	3	9	0	1	1	0	1	0	0	0	1	33	7	0	0	89	0	5
Z	0	0	1	0	0	1	0	13	0	14	8	0	0	1	1	0	0	0	0	0	111	0
Ø	2	4	8	3	6	5	0	0	3	0	2	5	2	11	1	0	0	14	9	0	0	75

	119	221	246	157	191	210	166	105	188	161	171	77	36	144	46	171	124	144	142	110	181	190

Table 3.10

Confusion matrix for final consonants for the +5 signal-to-noise ratio. Stimuli are represented by the rows, responses by the columns. Also shown are the column sums.

		+5 S/N																				
	p	t	k	b	d	g	C	J	s	S	z	f	T	v	D	n	m	l	r	G	Z	Ø
p	52	21	36	1	2	0	13	0	5	0	0	7	2	0	1	0	0	0	1	0	1	8
t	3	81	22	1	0	0	26	0	4	2	0	1	1	1	0	1	0	0	1	0	2	4
k	6	11	114	1	1	4	5	1	0	0	0	0	1	0	0	0	1	0	0	0	0	5
b	0	0	0	94	27	21	0	2	0	0	1	0	0	0	1	0	0	0	0	0	4	0
d	0	1	0	24	90	13	2	8	2	0	0	0	1	0	1	0	0	0	1	0	5	2
g	0	0	4	10	31	85	0	0	3	0	0	2	1	3	1	0	0	1	1	0	6	2
C	0	4	0	0	0	0	132	4	1	2	0	0	0	0	1	0	0	0	0	0	3	3
J	0	0	0	1	10	12	4	64	0	0	0	0	0	0	0	0	0	0	0	0	59	0
s	1	21	11	0	1	2	6	1	86	7	3	5	1	0	0	0	0	0	0	0	0	5
S	0	0	0	0	0	0	16	1	1	129	0	0	0	0	0	0	0	0	0	0	3	0
z	0	1	1	3	10	15	0	2	15	2	52	1	2	10	2	10	0	0	2	2	11	9
f	17	15	35	4	0	0	21	0	11	5	1	17	4	1	0	0	0	0	5	0	1	13
T	5	28	36	0	1	1	17	1	15	3	1	18	3	2	2	0	0	1	0	0	1	15
v	0	0	2	7	2	23	0	2	9	0	6	3	5	42	5	0	0	2	1	3	12	26
D	4	1	0	8	14	20	1	5	6	0	20	1	5	24	6	1	0	3	3	2	15	11
n	0	1	0	0	6	1	0	0	0	0	1	0	0	1	1	100	10	1	1	21	2	4
m	0	0	0	2	1	3	0	0	0	0	2	0	0	6	0	39	78	0	0	17	0	2
l	0	1	0	5	2	8	0	0	0	0	0	2	2	4	0	0	1	92	10	0	0	23
r	1	1	0	1	1	1	0	0	3	0	2	2	1	1	2	1	0	30	89	0	0	14
G	0	2	0	2	5	6	1	1	0	0	0	1	1	6	1	44	14	0	0	56	7	3
Z	0	0	0	1	4	6	1	10	2	16	6	0	0	1	0	0	0	0	0	2	97	4
Ø	3	7	14	5	5	7	1	2	2	0	2	6	0	4	3	5	4	11	7	2	1	59
	92	196	275	170	213	228	246	104	165	166	97	66	30	106	27	201	108	141	122	105	230	212

Table 3.11

Confusion matrix for final consonants for the -5 signal-to-noise ratio. Stimuli are represented by the rows, responses by the columns. Also shown are the column sums.

		-5 S/N																					
		p	t	k	b	d	g	C	J	s	S	z	f	T	v	D	n	m	l	r	G	Z	Ø
p	27	24	31	1	0	1	27	0	11	6	1	7	1	2	0	0	1	1	0	0	1	8	
t	11	36	29	0	1	2	20	4	10	4	0	4	4	2	1	2	2	0	2	0	2	14	
k	11	39	35	1	0	2	22	2	13	3	1	5	1	0	0	2	0	0	1	0	1	11	
b	5	13	7	27	15	33	1	4	1	3	1	2	2	2	1	2	1	0	1	3	9	17	
d	2	7	3	6	20	27	9	7	8	6	1	5	3	8	2	2	2	3	0	1	11	17	
g	3	8	10	8	8	46	5	5	3	3	3	8	1	8	0	0	2	0	1	6	10	12	
C	6	21	12	0	1	0	70	2	6	11	0	3	2	0	2	1	0	0	1	0	5	7	
J	7	8	4	8	11	23	10	11	5	8	4	2	5	2	1	8	0	0	2	1	19	11	
s	9	21	26	2	1	6	18	3	24	5	1	5	3	3	0	2	0	2	1	2	3	13	
S	6	14	16	2	1	0	42	1	5	29	0	3	8	1	1	1	0	0	0	0	6	14	
z	2	2	0	6	11	13	2	5	3	1	5	2	1	11	6	17	2	2	5	6	22	26	
f	9	36	24	1	3	4	16	1	12	10	0	13	3	1	3	1	1	0	0	0	0	12	
T	9	29	29	0	1	7	21	0	12	7	0	6	4	1	3	3	1	2	0	0	1	14	
v	2	2	5	7	11	23	1	5	2	1	13	2	1	17	2	5	3	8	4	4	13	19	
D	4	4	3	6	11	28	2	8	7	0	5	4	3	14	4	4	2	9	3	4	12	13	
n	2	3	2	7	6	6	0	1	1	0	3	1	1	6	2	58	10	4	0	23	5	9	
m	1	1	4	6	6	5	1	0	0	0	1	1	0	5	1	44	29	3	3	22	2	15	
l	1	6	4	5	7	14	1	3	0	0	0	2	1	4	2	1	0	37	24	2	0	36	
r	1	4	2	2	7	7	1	2	6	1	5	3	2	4	2	4	0	30	35	1	2	29	
G	1	1	0	3	6	0	0	0	0	0	0	0	0	2	0	47	23	2	0	48	4	13	
Z	3	5	0	10	12	21	0	5	4	1	17	0	2	7	2	4	1	4	2	7	30	13	
Ø	4	16	12	7	7	6	4	1	8	2	4	7	1	5	1	5	1	8	7	1	3	40	
		126	300	258	115	146	274	273	70	141	101	65	85	49	105	36	213	81	115	92	131	161	363

Neighborhood Probability Rules

Having obtained confusion matrices for the initial and final consonants and vowels, the question becomes: How can the segmental intelligibility of the stimulus word, estimated from the confusion matrices, be combined with the segmental confusability of its neighbors, also estimated from the confusion matrices, to provide an index of the identifiability of the stimulus word? One means of accomplishing this goal is to devise a neighborhood probability rule (NPR) incorporating the probability of identifying the stimulus word and the probabilities of confusing the neighbors with the stimulus word. Stated differently, can a neighborhood probability rule be devised that expresses the probability of identifying the stimulus word given the probabilities of identifying its neighbors? Luce's (1959) choice rule provides a straightforward means of computing such probabilities. Very simply, Luce's choice rule states that the probability of choosing a particular item i is equal to the probability of item i divided by the probability of item i plus the sum of the probabilities of j other items. A general form of the rule is:

$$(3.1) \quad \frac{p(i)}{p(i) + \sum p(j)}$$

Thus, Luce's choice rule provides a means for computing the probability of choosing a given item from a collection of other items.

The applicability of Luce's general choice rule to the problem at hand is transparent. Specifically, it provides a means for attempting to predict the probability of choosing a stimulus word from among its neighbors and thus provides the basis for devising a neighborhood probability rule. Accordingly, a neighborhood probability rule assumes the following general form: The probability of identifying the stimulus word is equal to the probability of the stimulus word divided by the probability of the stimulus word plus the sum of the probabilities of identifying the neighbors. Thus:

$$(3.2) \quad p(\text{ID}) = \frac{p(S)}{p(S) + \sum_{j \neq 1} p(N_j)}$$

where $p(\text{ID})$ is the probability of correctly identifying the stimulus word, $p(S)$ is the probability of the stimulus word, and $p(N_j)$ is the probability of the j th neighbor.

Stimulus Word Probabilities

The data from the confusion matrices can be used to compute the probability of the stimulus word and the conditional probabilities of its neighbors. The probability of the stimulus word can be computed as follows: For each phoneme in the stimulus word, the conditional probability of that phoneme given itself can be obtained from the confusion matrices. Assuming independent probabilities, the obtained conditional phoneme probabilities can be multiplied. This product renders a stimulus word probability (SWP) based on the probabilities of the individual phonemes of the stimulus word. Note that the probability for the stimulus word is based on the product of the conditional probabilities of identifying each phoneme independently, obtained from the confusion matrices. Thus, the stimulus word probability can be

computed as follows:

$$(3.3) \text{ SWP} = \prod_{i=1}^n p(\text{PS}_i | \text{PS}_i),$$

where $p(\text{PS}_i | \text{PS}_i)$ is the conditional probability of identifying the i th phoneme of the stimulus word given that phoneme, and n is the number of phonemes in the word. For example, the stimulus word probability of the word /dɔg/ is:

$$(3.4) \text{ SWP}(\text{dɔg}) = p(\text{d}|\text{d}) * p(\text{ɔ}|\text{ɔ}) * p(\text{g}|\text{g}),$$

where, again, the conditional probabilities of the individual phonemes are determined from the confusion matrices for the initial consonants, vowels, and final consonants. Note that the SWP of /dɔg/ can be construed as the conditional probability of the word /dɔg/ given /dɔg/, or $p(\text{dɔg}|\text{dɔg})$.

Neighbor Word Probabilities

In this manner, conditional probabilities for each neighbor of the stimulus word can also be computed. Thus, the neighbor word probability (NWP) can be computed by finding the conditional probabilities of each of the phonemes of the neighbor given the stimulus word phonemes. Multiplying these probabilities renders an index of the probability of the neighbor, or the NWP. Neighbor word probability can be computed as follows:

$$(3.5) \text{ NWP} = \prod_{i=1}^n p(\text{PN}_i | \text{PS}_i),$$

where PN_i is the i th phoneme of the neighbor, PS_i is the i th phoneme of the stimulus word, and n is the number of phonemes. (This description assumes, of course, an equal number of phonemes in the stimulus word and the neighbor. Moreover, the conditional probabilities can only be obtained from the available confusion matrices if the i th phonemes of the stimulus word and neighbor are both consonants or both vowels. These qualifying assumptions will be discussed in more detail below.)

To return to the example of the stimulus word /dɔg/ given above, the neighbor word probability for the neighbor /t@g/ can be computed as:

$$(3.6) \text{ NWP}(\text{t@g}) = p(\text{t}|\text{d}) * p(\text{@}|\text{ɔ}) * p(\text{g}|\text{g}),$$

which can also be construed to be the probability of identifying /t@g/ given /dɔg/, or $p(\text{t@g}|\text{dɔg})$.

Unweighted Neighborhood Probability Rule

Given these designations of stimulus word and neighbor word probabilities, the appropriate substitutions of terms in equation (3.2) render a neighborhood probability rule based on the general choice rule:

SWP

$$(3.7) \quad p(\text{ID}) = \frac{\text{SWP}}{\text{SWP} + \sum_{j=1}^n \text{NWP}_j},$$

where $p(\text{ID})$ is the probability of identifying the stimulus word and n is the number of neighbors.

The complete expanded form of the neighborhood probability rule is shown in equation (3.8).

$$(3.8) \quad p(\text{ID}) = \frac{\prod_{i=1}^n p(\text{PS}_i | \text{PS}_i)}{\prod_{i=1}^n p(\text{PS}_i | \text{PS}_i) + \sum_{j=1}^{nn} \left[\prod_{i=1}^n p(\text{PN}_{ij} | \text{PS}_i) \right]},$$

where PS_i is the i th phoneme of the stimulus word, PN_{ij} is the i th phoneme of the j th neighbor, n is the number of phonemes, and nn is the number of neighbors. (This rule will be referred to as the unweighted probability rule because it does not take into account stimulus word or neighbor frequency.)

Again, note that this particular form of the rule applies only to stimulus words and neighbors that are equal in length. In addition, given that the probabilities entering into the rule are based on separate confusion matrices for initial consonants, vowels, and final consonants, it is strictly only possible -- given the available confusion data -- to apply this rule to stimulus words and neighbors having the form consonant-vowel-consonant. Indeed, these restrictions (with a few modifications; see below) were adopted for the present analysis. Of course, there need be no inherent limitations in the rule itself; the restrictions on its application arise here only because of the available confusion matrix data. There is no reason to assume that a larger set of confusion data could not be obtained using a variety of structures in order to generalize the present approach beyond CVC words.

The neighborhood probability rule shown in equation (3.8) exemplifies one means of determining the probability of identifying a stimulus word in the context of its neighbors. Note that the rule takes into account both the basic intelligibility of the stimulus word (the SWP) as well as the probabilities of identifying the neighbors of the stimulus word (the NWP_j 's). Note also that the meaning of the term "neighbor" now takes on a somewhat different meaning than employed in Chapter 2. In essence, every word is a neighbor of every other word. However, each "neighbor" has an associated conditional probability that can range, in theory, from zero to one. Thus, neighborhood membership is not categorical, as in the previous analysis of neighbors based on one phoneme substitutions, insertions, or deletions. Instead, each neighbor is represented in the neighborhood with some probability derived from the confusion matrices.

A number of properties of the neighborhood probability rule are worthy of mention. First, intelligibility of the phonemes of the stimulus word itself will determine, in part, the role of the neighbors in determining the predicted probability of identification. Stimulus words with high phoneme probabilities (i.e., words with highly intelligible phonemes) will tend to

have neighbors with low phoneme probabilities, owing to the fact that all probabilities in the confusion matrices are conditional. Likewise, stimulus words with low phoneme probabilities (i.e., those with less intelligible phonemes) will tend to have neighbors with relatively higher phoneme probabilities. However, the output of the neighborhood probability rule is not a direct function of the stimulus word probability. Instead, the output of the rule is dependent on the existence of lexical items that contain phonemes that are confusable with the phonemes of the stimulus word. For example, a stimulus word may contain highly confusable phonemes. However, if there are few actual lexical items (i.e., neighbors) that contain phonemes confusable with those of the stimulus word, the sum of the neighbor word probabilities will be low. The resulting output of the neighborhood probability rule will therefore be relatively high. Likewise, if the phonemes of the stimulus word are highly intelligible, but there are a large number of neighbors that contain phonemes that are confusable with the stimulus word, the probability of identification will be reduced. In short, the output of the neighborhood probability rule is contingent on both the intelligibility of the stimulus word and the number of neighbors that contain phonemes that are confusable with those of the stimulus word. Thus, intelligibility of the stimulus word, confusability of the neighbors, and the nature of lexical items act in concert to determine the predicted probability of identification.

In the context of the neighborhood probability rule, then, the role of "neighborhood density" becomes more complex. The previous definition of neighborhood density in Chapter 2 relied on the categorical inclusion and exclusion of words in the neighborhood of the stimulus word based on one phoneme substitutions, additions, and deletions. However, once neighbors are assigned probability values, the role of neighborhood density depends on the neighbor word probabilities themselves. For example, one stimulus word may have a high number of low probability neighbors. However, another stimulus word may have a low number of high probability neighbors. In these two situations, the predictions of the neighborhood probability rule could be identical. Thus, predicted identification need not be a function of neighborhood density per se. Instead, the number as well as nature of neighbors serves to predict observed identification performance.

Frequency-Weighted Neighborhood Probability Rule

The neighborhood probability rule discussed thus far has incorporated no means for representing the frequency of the stimulus word or the frequencies of the word's neighbors. Although there are undoubtedly numerous ways to incorporate frequency in the neighborhood probability rule, frequency was incorporated in the neighborhood probability rule by weighting (multiplicatively) the probabilities of the stimulus word and its neighbors by their log-transformed frequencies. Thus, frequency was incorporated in equation (3.8) as shown in equation (3.9):

$$(3.9) \quad p(\text{ID}) = \frac{\prod_{i=1}^n p(\text{PS}_i | \text{PS}_i) * \text{FreqS}}{\left\{ \prod_{i=1}^n p(\text{PS}_i | \text{PS}_i) * \text{FreqS} \right\} + \sum_{j=1}^{nn} \left\{ \prod_{i=1}^n p(\text{PN}_{ij} | \text{PS}_i) * \text{FreqN}_j \right\}} ;$$

where PS_i is the probability of the i th phoneme of the stimulus word, PN_{ij} is the probability of the i th phoneme of the j th neighbor, n is the number of phonemes in the stimulus word and the neighbor, $Freq_S$ is the frequency of the stimulus word, $Freq_{Nj}$ is the frequency of the j th neighbor, and nn is the number of neighbors. This rule will be referred to as the frequency-weighted neighborhood probability rule (FWNPR). Given the superior correlation of stimulus word frequency based on Standard Frequency Indices (SFI's, see Carroll, 1970) reported in Chapter 2, only these indices of frequency were used as frequency weights in the computation of the rule.

In FWNPR, the frequencies of the stimulus word and the neighbors will serve to amplify to a greater or lesser degree the word probabilities of the stimulus word and its neighbors. Note that frequency in this rule is expressed in terms of the relation of the frequency of the target word to the frequencies of its neighbors. Thus, the absolute frequency of the stimulus word may have differential effects on predicted identification performance depending on the frequencies of the word's neighbors. For example, given two stimulus words of equal frequency, the stimulus word with neighbors of lower frequencies will produce a higher predicted probability than the stimulus word with neighbors of higher frequencies. The degree to which the frequencies of the neighbors will play a role in determining predicted identification performance will, of course, depend on the neighbor word probabilities. The frequencies of the neighbors with low probabilities of confusion will play less of a role than those with high probabilities of confusion. Simply put, neighborhood structure will play a role in determining predicted identification performance in terms of the combined effects of the number and nature of the neighbors, the frequencies of the neighbors, the intelligibility of the stimulus word, and the frequency of the stimulus word.

Predicting Identification Performance Using the Neighborhood Probability Rules

Computation of rules

In order to evaluate the relative efficacy of the proposed neighborhood probability rules, the data from the perceptual identification study reported in Chapter 2 were reanalyzed in terms of the unweighted and frequency-weighted rules using the confusion matrix data from the experiment reported in the present chapter. Given that confusion matrices were obtained only for consonants occurring in initial and final position, only those words of the form consonant-vowel-consonant from the original data set were analyzed. Altogether, 811 CVC words were analyzed. In addition, in order to simplify the computational analysis, only monosyllabic words contained within the Webster's lexicon were used to compute the neighborhood probability rules. Inspection of the error responses revealed that this was not an unreasonable simplification, given that a significant majority of the error responses were in fact monosyllabic, which indicates that subjects typically perceived monosyllabic words. The restriction of monosyllabic neighbors was necessitated by the particular procedure used to determine neighbor word probabilities (see below).

The general method for computation of the neighborhood probability rules was as follows: The stimulus word probability was first determined for a given stimulus word. This probability was computed from the confusion matrices using equation (3.3). Following computation of the stimulus word probability, the transcription of the stimulus word was compared to the transcriptions of all other monosyllabic words in Webster's lexicon. These

comparisons produced the neighbor word probabilities. The vowel of the stimulus word was first aligned with the vowel of the neighbor being analyzed. The conditional probabilities of the vowel and the consonants flanking the vowel for the neighbor were then determined from the appropriate confusion matrices. In the event that the neighbor was a CVC word, the neighbor word phoneme probability was computed using equation (3.5). That is, the conditional probability of the initial consonant of the neighbor given the initial consonant of the stimulus word was determined, as were the conditional probabilities of the vowel and the final consonant.

A problem arises as to the treatment of neighbors containing either initial consonant clusters, final consonant clusters, or both. In these cases, the transcriptions for the stimulus word and the neighbor were aligned at the vowel. However, when initial and/or final clusters were present in the neighbor word, those consonants not immediately adjacent to the vowel fail to overlap with anything in the stimulus word. For example, if the stimulus word is /k@t/ and the neighbor is /skId/, the /k/ of the stimulus word would align with the /k/ of the neighbor, /@/ would align with /I/, and /t/ would align with /d/. However, the /s/ of the neighbor /skId/ would align with no phoneme in the stimulus word. In this event, the probability of the phoneme /s/ was determined by finding the conditional probability of /s/ given the null phoneme from the confusion matrix for initial consonants, or $p(s|\emptyset)$, where " \emptyset " is the null phoneme. In essence, this is the probability of perceiving /s/ when in fact no phoneme has been presented. The conditional probabilities for the neighbor /skId/ would thus be: $p(s|\emptyset)$, $p(k|k)$, $p(I|@)$, and $p(d|t)$. The procedure for final consonant clusters was identical, except that the probability of the neighbor phoneme given the null phoneme was computed from the final consonant confusion matrix. This method of dealing with initial and final clusters in the neighbor word makes the simplistic assumption that clusters are phonetically and acoustically equivalent to the sum of their constituent phonemes. However, in the absence of confusion matrices for all possible clusters as well as singletons, this simplification appeared reasonable.

A similar problem arises when the neighbor is shorter than the stimulus word. In these cases, however, the solution is much more straightforward. The stimulus word and neighbor are once again aligned at the vowel. The empty slot in the neighbor is then assumed to contain the null phoneme, in which case the conditional probability for that phoneme can be easily determined from the appropriate confusion matrix. For example, if the stimulus word is /k@t/ and the neighbor /@t/, the neighbor /@t/ is assumed to have the transcription / \emptyset @t/. The conditional probabilities for the neighbor phonemes would then be: $p(\emptyset|k)$, $p(@|@)$, and $p(t|t)$.

In this manner, neighbor word probabilities were determined and the neighborhood probability rules were computed for each of the 811 CVC stimulus words. The neighborhood probability rules were computed separately for each S/N ratio for each word using the confusion matrices appropriate to that S/N ratio. When computing the neighborhood probability rules for a given S/N ratio, the confusion matrix obtained at that S/N ratio was used to determine the stimulus and neighbor word probabilities. That is, three predicted identification scores for each neighborhood probability rule were computed for each word, one for each of the three S/N ratio.

Analyses over All S/N Ratios

The output of the unweighted and frequency-weighted neighborhood probability rules were combined with the scores for the 811 CVC words and submitted to correlation and regression analyses. In addition to the predicted identification scores from the neighborhood probability rules, four other variables were examined: (1) duration in msec of the stimulus word (DUR), (2) SFI of the stimulus word (SWF-SFI), (3) the unweighted stimulus word probability (UWSWP) obtained from the confusion matrix data, and (4) the frequency-weighted stimulus word probability (FWSWP). (The frequency-weighted SWP's were computed by multiplying the SWP by its SFI.) The unweighted and weighted stimulus word probabilities were included to provide baselines against which to compare the performance of the unweighted and frequency-weighted neighborhood probability rules. The resulting variables and their mnemonics are shown in Table 3.12.

Insert Table 3.12 about here

The first set of analyses was performed on the entire data set averaged across all three S/N ratios. That is, identification scores, SWP's, and the predictions of the rules were each averaged across S/N ratios and submitted to correlation and regression analyses.

Correlation Analyses

Summary statistics for the predictor variables and word scores are shown in Table 3.13. Means, standard deviations and minimum and maximum values for each variable are shown in this table. The correlations of each of the six predictor variables with the identification scores are shown in Table 3.14. Variable mnemonics are given in Table 3.12.

Insert Tables 3.13 and 3.14 about here

All correlations were significant beyond the 0.05 level of significance. The obtained correlations, ranked from lowest to highest, were: (1) stimulus duration (DUR), $r = .1599$; (2) stimulus word frequency (SWF-SFI), $r = .2760$; (3) unweighted stimulus word probability (UWSWP), $r = .3676$; (4) frequency-weighted stimulus word probability (FWSWP), $r = .4253$; (5) unweighted neighborhood probability rule (UWNPR), $r = .4339$; and (6) frequency-weighted neighborhood probability rule (FWNPR), $r = .4750$. Thus, the frequency-weighted neighborhood probability rule (FWNPR) correlated most highly with performance, followed by the unweighted neighborhood probability rule (UWNPR). In the overall analysis, therefore, inclusion of the frequency-weighting factor improved performance of the neighborhood probability rule. In addition, both neighborhood probability rules proved to be better predictors of identification performance than either the unweighted or frequency-weighted stimulus word probabilities. This result demonstrates that both stimulus intelligibility and neighborhood confusability combined served to better predict identification performance than stimulus intelligibility alone. Finally, and not surprisingly, the frequency-weighted stimulus word probability proved to be a better predictor of identification

Table 3.12

Neighborhood Analysis: Neighborhood Probability Rules. Variables and mnemonics.

VARIABLE	MNEMONIC
1. Duration of stimulus word in msec based on SFI's	DUR
2. SFI stimulus word frequency	SWF-SFI
3. Unweighted stimulus word probability	UWSWP
4. Frequency-weighted stimulus word probability	FWSWP
5. Unweighted neighborhood probability rule	UWNPR
6. Frequency-weighted neighborhood rule	FWNPR

Table 3.13

Neighborhood Analysis: Neighborhood Probability Rules. Summary statistics for variables for all signal-to-noise ratios.

VARIABLE	MEAN	STANDARD DEVIATION	MINIMUM VALUE	MAXIMUM VALUE
1. DUR	0.4961	0.0695	0.3012	0.7358
2. SWF-SFI	52.1626	8.4686	40.0000	80.2510
3. UWSWP	0.2144	0.1051	0.0082	0.5148
4. FWSWP	11.1826	5.8049	0.3293	31.2158
5. UWNPR	0.3559	0.1530	0.0152	0.7199
6. FWNPR	0.3567	0.1536	0.0108	0.7093
7. SCORE	0.4947	0.2241	0.0000	1.0000

Table 3.14

Neighborhood Analysis: Neighborhood Probability Rules. Correlations between identification scores and predictor variables for all signal-to-noise ratios.

VARIABLE	r
-----	-----
1. DUR	0.1599*
2. SWF-SFI	0.2760*
3. UWSWP	0.3676*
4. FWSWP	0.4253*
5. UWNPR	0.4339*
6. FWNPR	0.4750*

*p<0.05

performance than the unweighted stimulus word probability.

Regression Analyses

In order to evaluate the performance of the UWNPR and FWNPR against the other independent variables of interest, hierarchical multiple regression analyses were performed. Three sets of analyses were conducted, one examining the performance of the UWNPR, one examining the performance of the FWNPR, and one comparing the two rules. In the first set of multiple regression analyses, the UWNPR was evaluated in conjunction with stimulus duration (DUR) and unweighted stimulus probability (UWSWP). This analysis was aimed at determining the performance of the neighborhood probability rule in the absence of stimulus word or neighbor frequency. Two different orders of the independent variables (DUR, UWSWP, and UWNPR) were examined in the multiple regression analyses. These two orders were determined a prior according to the predicted contribution of each independent variable to the proportion of variance accounted for.

The first ordering of the variables was designed to proceed from the variable predicted to account for the smallest proportion of variance to the variable predicted to account for the largest proportion of variance, namely, from (1) DUR to (2) UWSWP to (3) UWNPR. By so ordering the independent variables, it is possible to determine the proportion of variance accounted for by the UWNPR after the effects of stimulus duration and stimulus word probability have been removed. Any additional variance contributed by the rule serves as an index of the effects of neighborhood structure above and beyond the effects of stimulus duration and stimulus intelligibility.

The second ordering of the independent variables was designed to proceed from the largest predicted proportion of variance accounted for to the smallest proportion of variance accounted for, that is, from (1) UWNPR to (2) UWSWP to (3) DUR. This ordering enables examination of the degree to which UWSWP and DUR contribute to the proportion of variance accounted for above and beyond that accounted for by UWNPR.

These same orderings of the independent variables were applied to the FWNPR. The ordering from "smallest" to "largest" was: (1) DUR, (2) FWSWP, (3) FWNPR. Again, by so ordering the independent variables in the regression analysis, it is possible to determine the proportion of variance accounted for by FWNPR after the effects of stimulus duration, stimulus word probability, and stimulus word frequency have been removed, thus providing an independent estimate of the effects of the frequency-weighted neighborhood structure. The reverse ordering of variables--proceeding from "largest" to "smallest" predicted proportion of variance accounted for--was also examined for the FWNPR, namely, from (1) FWNPR to (2) FWSWP to (3) DUR.

In addition to the two sets of hierarchical multiple regression analyses on the UWNPR and the FWNPR, a third analysis was performed that compared the two rules. The extent to which the frequency-weightings incorporated in the FWNPR improve prediction was evaluated by entering the UWNPR followed by the FWNPR. Any additional variance accounted for by the FWNPR once the effects of the UWNPR have been removed will serve as an index of what improvement, if any, is provided by the frequency-weights associated with the stimulus word and its neighbors incorporated in the FWNPR.

The results of the two multiple regression analyses ("smallest to largest" and "largest to smallest") for the UWNPR are shown in the top panel of Table 3.15. Multiple R's, multiple R²'s, changes in R²'s as a function of the total variance, and changes in R²'s as a function of the explained variance are shown at each step.

Insert Table 3.15 about here

For the ordering of variables proceeding from smallest to largest for the UWNPR, a multiple R = 0.4472 and a multiple R² = 0.2000 were obtained. Stimulus duration (DUR) accounted for 2.56% of the total variance and 12.80% of the explained variance. Unweighted stimulus probability (UWSWP) accounted for 13.78% of the total variance and 68.90% of the explained variance. Finally, the UWNPR accounted for 3.67% of the total variance and 18.35% of the explained variance. Thus, the UWNPR contributed a substantial proportion of variance above and beyond that contributed by stimulus duration and stimulus intelligibility, as indexed by UWSWP. That is, even when the effects of stimulus intelligibility are removed, neighborhood structure, as indexed by the UWNPR, accounted for approximately 18% of the explained variance, thus clearly demonstrating that neighborhood structure is an important determinant of identification performance above and beyond the intelligibility of the stimulus word itself.

For the ordering of the variables proceeding from largest to smallest for the UWNPR, a multiple R = 0.4472 and a multiple R² = 0.2000 were again obtained. The UWNPR accounted for 18.83% of the total variance and 94.15% of the explained variance. The unweighted stimulus word probability and stimulus duration each contributed little to the total proportion of variance accounted for (0.56% and 0.61%, respectively) as well as to the proportion of explained variance (2.80% and 3.05%, respectively). This result demonstrates that stimulus intelligibility and duration contribute little to the proportion of variance accounted for beyond that accounted for by the UWNPR.

The results for the FWNPR are shown in the middle panel of Table 3.15. Both orderings of the independent variables produced a multiple R = 0.4973 and a multiple R² = 0.2473. For the ordering proceeding from smallest to largest, stimulus duration accounted for 2.56% of the total variance and 10.35% of the explained variance. The FWSWP accounted for 18.52% of the total variance and 74.88% of the explained variance. Finally, the FWNPR accounted for 3.65% of the total variance and 14.76% of the explained variance. Again, the neighborhood probability rule contributed a substantial proportion of variance beyond that contributed by the other independent variables. Indeed, even when stimulus duration, stimulus word frequency, and stimulus word probability are removed, the FWNPR accounted for over 14% of the explained variance, once again demonstrating that neighborhood structure has a marked effect on identification performance independent of stimulus duration, frequency, and intelligibility.

For the ordering of variables proceeding from largest to smallest, the neighborhood probability rule again accounted for the majority of the variance. The FWNPR account for 22.56% of the total variance and 91.23% of the explained variance. The FWSWP and stimulus duration contributed small proportions to both the total variance accounted for (1.44% and 0.72%, respectively) and the explained variance (5.82% and 2.91%, respectively). Thus, the FWNPR appears to capture most of the effects of stimulus word frequency, stimulus word intelligibility, and stimulus duration.

Table 3.15

Neighborhood Analysis: Neighborhood Probability Rules. Results of hierarchical multiple regression analyses for all signal-to-noise ratios. Shown are the results for the variables entered from "smallest to largest" and "largest to smallest" for the UWNPR and the FWNPR. Also shown are the regression analyses comparing UWNPR and FWNPR. Multiple R's, multiple R²'s, changes in R² as a function of total variance, and changes in R² as a function of explained variance are shown at each step.

=====

ANALYSIS OF UWNPR

SMALLEST TO LARGEST

VARIABLE	R	R ²	Δ R ² (TOTAL)	Δ R ² (EXPLAINED)
1. DUR	0.1599	0.0256	0.0256	0.1280
2. UWSWP	0.4042	0.1634	0.1378	0.6890
3. UWNPR	0.4472	0.2000	0.0367	0.1835

LARGEST TO SMALLEST

1. UWNPR	0.4339	0.1883	0.1883	0.9415
2. UWSWP	0.4403	0.1939	0.0056	0.0280
3. DUR	0.4472	0.2000	0.0061	0.0305

ANALYSIS OF FWNPR

SMALLEST TO LARGEST

1. DUR	0.1599	0.0256	0.0256	0.1035
2. FWSWP	0.4591	0.2108	0.1852	0.7488
3. FWNPR	0.4973	0.2473	0.0365	0.1476

LARGEST TO SMALLEST

1. FWNPR	0.4750	0.2256	0.2256	0.9123
2. FWSWP	0.4899	0.2400	0.0144	0.0582
3. DUR	0.4973	0.2473	0.0072	0.0291

COMPARISON OF UWNPR AND FWNPR

1. UWNPR	0.4339	0.1882	0.1883	0.6749
2. FWNPR	0.5282	0.2790	0.0908	0.3254

=====

The results of the hierarchical multiple regression analysis comparing the performance of the UWNPR and the FWNPR are shown in the bottom panel of Table 3.15. The UWNPR alone contributed 18.83% to the total variance and 67.49% to the explained variance. Once the effects of the UWNPR were removed, however, the FWNPR contributed an additional 9.08% to the total variance and 32.54% to the explained variance. Clearly, the frequency-weightings incorporated in the FWNPR result in a substantial improvement in the proportion of variance accounted for, even once the effects of stimulus intelligibility and unweighted neighborhood structure are partitioned out, further demonstrating that identification performance is a function of the combined effects of stimulus intelligibility, stimulus word frequency, neighborhood confusability, and neighborhood frequency.

Analyses at Each S/N Ratio

In order to examine the effects of the independent variables at each S/N ratio separately, the same correlation analyses performed on the data average across S/N ratios were performed for each of the three S/N ratios. These analyses were undertaken to determine what differential effects, if any, the independent variables may have as a function of variations in stimulus degradation.

Correlation Analyses

Summary statistics for the predictor variables and identification scores for each of the three S/N ratios are shown in Table 3.16. Means, standard deviations, and minimum and maximum values for each variable are shown in this table. The correlations of each of the six predictor variables with the identification scores are shown in Table 3.17. Variable mnemonics are given in Table 3.12

Insert Tables 3.16 and 3.17 about here

All correlations were significant at the 0.05 level or beyond. Looking across S/N ratios, a number of patterns emerge. In particular, the correlations of stimulus word frequency and identification tend to increase with decreasing S/N ratio. Thus, stimulus word frequency played an increasingly stronger role as stimulus degradation increased. This result suggests either that words of higher frequency tend to be more resistant to degradation or that subjects adopt more strongly frequency-biased strategies of response generation as stimulus degradation increases. The first hypothesis is at least partially ruled out by the intercorrelations between stimulus word probability and SFI. If stimulus word probability is taken as an index of intelligibility, and high frequency words are, in fact, more resistant to degradation, then the correlations of stimulus word probability and stimulus frequency would be expected to increase with decreasing S/N ratio. However, no significant correlations between these two variables were observed at any of the S/N ratios. Thus, the identifiability of the constituent phonemes of high frequency versus low frequency words does not appear to vary differentially with S/N ratio, suggesting that the increase in the correlation of SFI across S/N ratio is due to an increased frequency bias on the part of subjects when presented with more degraded stimuli.

Table 3.16

Neighborhood Analysis: Neighborhood Probability Rules. Summary statistics
 for variables for each signal-to-noise ratio.

+15 S/N

VARIABLE	MEAN	STANDARD DEVIATION	MINIMUM VALUE	MAXIMUM VALUE
1. DUR	0.4961	0.0695	0.3012	0.7358
2. SWF-SFI	52.1626	8.4686	40.0000	80.2510
3. UWSWP	0.4342	0.1958	0.0206	0.8585
4. FWSWP	22.6564	10.9409	0.8240	54.6457
5. FWNPR	0.6292	0.2378	0.0258	0.9911
6. UWNPR	0.6276	0.2377	0.0364	0.9930
7. SCORE	0.7821	0.2771	0.0000	1.0000

+5 S/N

VARIABLE	MEAN	STANDARD DEVIATION	MINIMUM VALUE	MAXIMUM VALUE
1. DUR	0.4961	0.0695	0.3012	0.7358
2. SWF-SFI	52.1626	8.4686	40.0000	80.2510
3. UWSWP	0.1914	0.1214	0.0023	0.6448
4. FWSWP	9.9727	6.5662	0.0920	37.9026
5. UWNPR	0.3897	0.2155	0.0049	0.9881
6. FWNPR	0.3910	0.2166	0.0035	0.9903
7. SCORE	0.5600	0.3326	0.0000	1.0000

-5 S/N

VARIABLE	MEAN	STANDARD DEVIATION	MINIMUM VALUE	MAXIMUM VALUE
1. DUR	0.4961	0.0695	0.3012	0.7358
2. SWF-SFI	52.1626	8.4686	40.0000	80.2510
3. UWSWP	0.0176	0.0182	0.0002	0.1326
4. FWSWP	0.9187	0.9804	0.0096	7.7878
5. UWNPR	0.0505	0.0480	0.0007	0.3029
6. FWNPR	0.0499	0.0486	0.0007	0.3434
7. SCORE	0.1419	0.2041	0.0000	1.0000

Table 3.17

Neighborhood Analysis: Neighborhood Probability Rules. Correlations between identification scores and predictor variables for each signal-to-noise ratio.

VARIABLE -----	r -----		
	+15 S/N -----	+5 S/N -----	-5 S/N -----
1. DUR	0.1134*	0.1819*	0.0764*
2. SWF-SFI	0.2082*	0.2329*	0.2470*
3. UWSWP	0.3832*	0.2941*	0.1279*
4. FWSWP	0.4123*	0.3355*	0.1702*
5. UWNPR	0.3758*	0.4364*	0.1805*
6. FWNPR	0.4043*	0.4687*	0.2276*

* $p < 0.05$

The correlations of stimulus word probability and identification performance also varied systematically as a function of S/N ratio. As S/N ratio decreased, the correlations between stimulus word probability and identification scores also decreased. Thus, the relative identifiability of the constituent phonemes played a decreasing role in determining identification performance as stimulus degradation increased, suggesting a decreased reliance on the information in the stimulus itself as degradation increased.

Finally, the pattern of correlations for the two neighborhood probability rules showed a trend for highest correlations at the +5 S/N ratio, intermediate correlations at the +15 S/N ratio, and lowest correlations at the -5 S/N ratio. In short, the neighborhood probability rules tended to predict performance best at the intermediate level of stimulus degradation. The finding that the neighborhood probability rules tended to correlate more highly with identification performance at the +5 S/N ratio may be related to the simple fact that more variance was available for explanation at the +5 S/N ratio.¹ In light of this fact, direct comparisons among the correlations at each S/N ratio are somewhat tenuous when the correlations vary as a function of the total proportion of variance available for explanation.

Regression Analyses

To further examine the performance of the neighborhood probability rules, hierarchical multiple regression analyses were performed at each S/N ratio. The regression analyses were identical to those performed for data averaged over S/N ratio. The results of these analyses are shown in Tables 3.18, 3.19, and 3.20.

Insert Tables 3.18, 3.19, 3.20 about here

For all analyses, the results revealed largest multiple R's for the +5 S/N ratio and smallest multiple R's for the -5 S/N ratio, again demonstrating that the best prediction was obtained at the intermediate level of stimulus degradation. The results also revealed consistently larger multiple R's for the analyses containing the FWNPR than for those analyses containing the UWNPR.

For the analyses proceeding from smallest to largest at each S/N ratio, it was once again found that the neighborhood probability rule contributed significantly to the proportion of variance accounted for beyond that accounted for by the other independent variables. In particular, the UWNPR uniquely accounted for 1.62% of the total variance and the +15 S/N ratio, 7.61% of the total variance at the +5 S/N ratio, and 2.38% of the total variance at the -5 S/N ratio. Obviously, the UWNPR performed best at the +5 S/N ratio. However, in terms of the proportion of explained variance, the UWNPR contributed significantly larger proportions of variance as stimulus degradation increased, namely, 13.78% at the +15 S/N ratio, 38.61% at the +5 S/N ratio, and 50.00% at the -5 S/N ratio. This result suggests that, in terms of the proportion of variance actually explained by the combination of the independent variables, neighborhood structure plays an increasingly stronger role as stimulus degradation increases.

Table 3.18.

Neighborhood Analysis: Neighborhood Probability Rules. Results of hierarchical multiple regression analyses for the +15 signal-to-noise ratio. Shown are the results for the variables entered from "smallest to largest" and "largest to smallest" for the UWNPR and the FWNPR. Also shown are the regression analyses comparing UWNPR and FWNPR. Multiple R's, multiple R²'s, changes in R² as a function of total variance, and changes in R² as a function of explained variance are shown at each step.

+15 S/N				
ANALYSIS OF UWNPR				
SMALLEST TO LARGEST				
VARIABLE	R	R ²	ΔR^2 (TOTAL)	ΔR^2 (EXPLAINED)
1. DUR	0.1134	0.0129	0.0129	0.0726
2. UWSWP	0.4018	0.1615	0.1486	0.8367
3. UWNPR	0.4215	0.1776	0.0162	0.1378
LARGEST TO SMALLEST				
1. UWNPR	0.3758	0.1412	0.1412	0.8134
2. UWSWP	0.4167	0.1736	0.0324	0.1866
3. DUR		(not significant)		
ANALYSIS OF FWNPR				
SMALLEST TO LARGEST				
1. DUR	0.1134	0.0129	0.0129	0.0625
2. FWSWP	0.4306	0.1854	0.1725	0.8354
3. FWNPR	0.4545	0.2065	0.0211	0.1022
LARGEST TO SMALLEST				
1. FWNPR	0.4043	0.1634	0.1634	0.8057
2. FWSWP	0.4504	0.2028	0.0394	0.1943
3. DUR		(not significant)		
COMPARISON OF UWNPR AND FWNPR				
1. UWNPR	0.3758	0.1412	0.1412	0.7182
2. FWNPR	0.4434	0.1966	0.0554	0.2818

Table 3.19.

Neighborhood Analysis: Neighborhood Probability Rules. Results of hierarchical multiple regression analyses for the +5 signal-to-noise ratio. Shown are the results for the variables entered from "smallest to largest" and "largest to smallest" for the UWNPR and the FWNPR. Also shown are the regression analyses comparing UWNPR and FWNPR. Multiple R's, multiple R2's, changes in R2 as a function of total variance, and changes in R2 as a function of explained variance are shown at each step.

=====

+5 S/N

=====

ANALYSIS OF UWNPR

SMALLEST TO LARGEST

VARIABLE	R	R2	$\Delta R2$ (TOTAL)	$\Delta R2$ (EXPLAINED)
1. DUR	0.1819	0.0331	0.0331	0.1679
2. UWSWP	0.3479	0.1210	0.0879	0.4460
3. UWNPR	0.4440	0.1971	0.0761	0.3861

LARGEST TO SMALLEST

1. UWNPR	0.4364	0.1904	0.1904	0.9690
2. UWSWP		(not significant)		
3. DUR	0.4432	0.1965	0.0060	0.0305

ANALYSIS OF FWNPR

SMALLEST TO LARGEST

1. DUR	0.1819	0.0331	0.0331	0.1471
2. FWSWP	0.3860	0.1490	0.1159	0.5151
3. FWNPR	0.4743	0.2250	0.0760	0.3378

LARGEST TO SMALLEST

1. FWNPR	0.4687	0.2197	0.2197	0.9764
2. FWSWP		(not significant)		
3. DUR	0.4743	0.2250	0.0053	0.0236

COMPARISON OF UWNPR AND FWNPR

1. UWNPR	0.4364	0.1904	0.1904	0.7644
2. FWNPR	0.4991	0.2491	0.0587	0.2356

=====

=====

Table 3.20

Neighborhood Analysis: Neighborhood Probability Rules, Results of hierarchical multiple regression analyses for the -5 signal-to-noise ratio. Shown are the results for the variables entered from "smallest to largest" and "largest to smallest" for the UWNPR and the FWNPR. Also shown are the regression analyses comparing UWNPR and FWNPR. Multiple R's, multiple R²'s, changes in R² as a function of total variance, and changes in R² as a function of explained variance are shown at each step.

=====

-5 S/N

=====

ANALYSIS OF UWNPR

SMALLEST TO LARGEST

VARIABLE	R	R ²	ΔR^2 (TOTAL)	ΔR^2 (EXPLAINED)
1. DUR	0.0764	0.0058	0.0058	0.1218
2. UWSWP	0.1542	0.0238	0.0179	0.3761
3. UWNPR	0.2182	0.0476	0.0238	0.5000

LARGEST TO SMALLEST

1. UWNPR	0.1805	0.0326	0.0326	0.7426
2. UWSWP	0.2096	0.0439	0.0113	0.2574
3. DUR		(not significant)		

ANALYSIS OF FWNPR

SMALLEST TO LARGEST

1. DUR	0.0764	0.0058	0.0058	0.0838
2. FWSWP	0.1925	0.0371	0.0312	0.4509
3. FWNPR	0.2630	0.0692	0.0321	0.4639

LARGEST TO SMALLEST

1. FWNPR	0.2276	0.0518	0.0518	0.7908
2. FWSWP	0.2559	0.0655	0.0137	0.2092
3. DUR		(not significant)		

COMPARISON OF UWNPR AND FWNPR

1. UWNPR	0.1805	0.0336	0.0326	0.3357
2. FWNPR	0.3115	0.0971	0.0645	0.6643

=====

A similar pattern of results was obtained for the analyses containing the FWNPR. The FWNPR uniquely accounted for 2.11% of the total variance and 10.22% of the explained variance at the +15 S/N ratio, 7.60% of the total variance and 38.61% of the explained variance at the +5 S/N ratio, and 3.21% of the total variance and 46.39% of the explained variance at the -5 S/N ratio. What is of primary importance, however, is the finding that the FWNPR significantly contributed to the variance accounted for at each S/N ratio even when stimulus duration, stimulus intelligibility, and stimulus frequency had been removed, again demonstrating that neighborhood structure has a pronounced effect on identification performance beyond the effects of stimulus intelligibility and frequency.

The analyses proceeding from largest to smallest for both the UWNPR and FWNPR revealed that the rules captured most of the explained variance, leaving little or no additional explained variance to be accounted for by the other independent variables. For the UWNPR, the rule accounted for 81.34% of the explained variance at the +15 S/N ratio, 96.90% at the +5 S/N ratio, and 74.26% at the -5 S/N ratio. Similar results were obtained for the FWNPR: The frequency-weighted rule accounted for 80.57% of the explained variance at the +15 S/N ratio, 97.64% at the +5 S/N ratio, and 79.08% at the -5 S/N ratio. Clearly, both rules performed optimally at the +5 S/N ratio. Indeed, at this S/N ratio, neither the UWSWP nor the FWSWP contributed significantly to the variance. It is of special note that the entire effect of stimulus word frequency was accounted for by the FWNPR at the +5 S/N ratio, indicating that the relations of the frequency of the stimulus word to its neighbors entirely captured the effect of stimulus word frequency. However, at both the +15 and -5 S/N ratios, the UWSWP and FWSWP did contribute additional variance, indicating that the neighborhood probability failed to capture all of the effects of stimulus intelligibility and frequency at these S/N ratios.

The multiple regression analyses comparing the performance of the UWNPR and FWNPR are also shown in Tables 3.18, 3.19, and 3.20. For each S/N ratio, the FWNPR uniquely contributed over 5.00% to the total variance accounted for. As in the overall analysis, therefore, these results demonstrate the the frequency-weighting factors incorporated in the FWNPR result in substantial increases in the proportion of variance accounted for, again demonstrating that identification performance is a function of stimulus word frequency, stimulus word intelligibility, neighborhood frequency, and neighborhood confusability.

Summary

The results from the analyses of each S/N ratio separately support the general conclusions obtained from the analyses of the data collapsed across S/N ratios. For each S/N ratio, it was shown that the neighborhood probability rule contributes significantly to the proportion of variance accounted even when the effects of stimulus duration, stimulus intelligibility, and--in the case of the FWNPR--stimulus frequency are removed. Thus, the present results further support the claim that neighborhood structure is an important determinant of identification performance. In addition, these results also demonstrate that the FWNPR is the single best predictor of identification performance of the variables examined.

The analyses at each S/N ratio further revealed that the effects of neighborhood structure become more pronounced as stimulus degradation increases. Not surprisingly, this result demonstrates that increasing stimulus degradation increases the confusability of the stimulus word with its

neighbors. Also of note is the relatively poor performance of the neighborhood probability rules at the -5 S/N ratio. Clearly, the rules fare more poorly under the condition of extreme stimulus degradation. However, nearly 70% of the words in this condition were responded to correctly by at most one subject. Thus, there is far less variance available for explanation at this S/N ratio. In general, given the extreme level of degradation of both the words and the CVs and VCs used to compute the stimulus and neighbor word probabilities at the -5 S/N ratio, systematic prediction of identification scores may fare more poorly given increased noise in the word and confusion matrix data.

In summary, the present study clearly demonstrates that neighborhood structure is an important determinant of the perceptual identification of words in noise. Stimulus intelligibility, neighborhood confusability, and the frequency relations among the stimulus word and its neighbors act in concert to determine identification. In particular, it was demonstrated that a choice rule could be formulated that captures the complex interrelationships of the stimulus, its neighbors, and the frequencies of the stimulus and its neighbors.

GENERAL DISCUSSION

Factors Influencing Performance of the Rules

At first glance, the correlations between the FWNPR and identification performance at each S/N ratio may appear somewhat moderate. However, a number of factors must be considered in evaluating the success of the rule in terms of its correlation with performance. First, the obtained correlations were computed for a large number of stimuli ($N = 811$). The number of stimuli in fact virtually exhaust the population of highly familiar consonant-vowel-consonant words in English. The success of FWNPR, as indexed by the proportion of variance accounted for, is thus rather remarkable given the large number of stimuli examined.

The performance of the rule is even more striking given that no specific information about the idiosyncratic acoustic structures of the individual stimuli, aside from stimulus duration, was included in the rule. Only information concerning the relative intelligibility and confusability of the individual segments was included, and this information was obtained from an independent source of data, namely, the confusion matrices obtained in a quite different experiment. Thus, the rule was able to achieve the obtained level of performance in the absence of specific measurements of the spectral, durational, and amplitude characteristics of the specific segments of the individual word stimuli. Thus, the information entered into the rule was to a large extent independent of the idiosyncratic acoustic structures of the individual words.

The frequency-weighted neighborhood probability rule also incorporates three sources of information that may introduce considerable noise in prediction of identification. The first comes from the confusion matrices. Recall that the confusion matrices were obtained from a separate pool of subjects and were based on consonant-vowel and vowel-consonant syllables. The pattern of confusions obtained from CV and VC syllables may differ in fundamental ways from the pattern of confusions evoked by real word stimuli. In particular, response biases frequently observed in confusion matrices of this type may introduce significant sources of noise in predicting confusions among real word stimuli (see also Klatt, 1968; Miller & Nicely, 1955; Wang &

Bilger, 1973). Inspection of the column sums for the confusion matrices shown in Tables 3.3 to 3.11 reveals strong response biases for particular phonemes. Response bias is indexed by the deviation of the column sums from the row sums. Thus, for the consonant confusion matrices, deviations of the column sums from 150 indicate biases toward or against responding with a particular phoneme. Obtaining confusion matrices for individual segments in a task requiring absolute identification of these segments in nonsense syllables may therefore reflect biases that may be inappropriate for predicting confusions among segments in real word stimuli.

In addition, given that confusions for a particular segment were collapsed across the various phonetic contexts in which the segment was uttered, specific effects of coarticulation within the word stimuli themselves may have been attenuated or ignored. In short, a considerable amount of noise in predicting word identification using the FWNPR may have arisen from the confusion matrices themselves. This does not mean that the use of confusion matrices to determine stimulus and neighbor word probabilities was misguided (see Moore, 1977). Indeed, the use of confusion matrices provides the only independent means of assessing stimulus intelligibility and confusability. However, in assessing the performance of the rule, it must be kept in mind that the confusion matrices provide less than perfect estimations of intelligibility and confusability of real word stimuli. In light of these observations, then, the performance of the neighborhood probability rule is even more impressive.

A second source of possible noise introduced in FWNPR arises from the lexicon used to compute neighborhood structure. Despite the controls placed on the inclusion of words in the neighborhoods of the stimulus words, the lexicon used may tend either to under- or overestimate the mental lexicons of the subjects themselves. The lexicon serves as only a very general model of the mental lexicon of the subject, thus introducing a potentially large source of noise in the estimation of neighborhood structure. However, in the absence of well-controlled techniques for estimating the nature and number of lexical items in the mental lexicon of a particular subject, the lexicon used in the present study provides an invaluable tool for determining neighborhood structure. Indeed, prior to the advent of computerized lexicons containing phonetic transcriptions, such estimations of neighborhood structure would have been nearly impossible.

A third source of noise in predicting identification may have arisen from the use of the Kucera and Francis frequency counts. These counts are not only somewhat dated, having been obtained in the 1960's, but are also based on printed text. However, given that frequency counts were required for a large number of words, the use of available counts of spoken words was not feasible. Thus, the Kucera and Francis counts, although problematic, provided one of the single best estimates of word frequency available for a large number of stimuli.²

Once these factors are taken into consideration, the performance of the FWNPR proved to be more than adequate, and the results clearly demonstrate the role of neighborhood structure in word identification. Indeed, more important than the absolute correlation of the FWNPR with identification performance is the relative performance of the rule in comparison to the other variables examined. In every instance, the FWNPR proved to be a much better single predictor of performance than stimulus intelligibility or frequency, thus demonstrating that identification performance is a function of the complex interaction of stimulus intelligibility, neighborhood confusability, and the relations of stimulus and neighbor frequency.

One additional factor that may have contributed to the moderate correlations of the neighborhood probability rules and identification performance is the relatively low number of observations per stimulus word. Because of the large number of stimuli words employed, it was necessary to restrict the number of observations to 10 responses per stimulus word per S/N ratio. (Even with 10 observations per stimulus word, it was necessary to collect 27,540 observations: 918 stimulus words X 3 S/N ratios X 10 observations.) Thus, a considerable amount of variance may have been introduced by the relatively low number of observations per stimulus word.

One may ask, then, how much of the observed variance is due to the relatively low number of observations alone. To answer this question, the proportion of variance accounted for by the number of observations alone was simulated for each stimulus word. This was done by first assuming that the obtained identification score of a given stimulus word represented its true probability. A random number generator was then consulted and the resulting random number was compared to the probability of the stimulus word. If the random number was less than or equal to the probability of the stimulus word, a "correct" response was tallied. This process was repeated 10 times for each word and a pseudo-identification score was computed for each word based on the 10 observations. This simulation was performed 20 times for each of 811 stimulus words at each of the S/N ratios separately.

The output of the simulation produced 20 simulated identification scores per word per S/N ratio. These 20 identification scores were then correlated with the actual obtained identification scores and the proportions of variance accounted for were averaged for the 20 simulations. Ideally, if the number of observations were infinite, then the pseudo-identification scores should converge almost perfectly on the obtained identification scores (which were taken to be the true probabilities of the words). However, given fewer observations per word, the simulated data will tend to correlate less well with the obtained data, and the unexplained variance provides an estimation of the variance introduced by the number of observations alone. The results of this simulation are shown in Table 3.21.

Insert Table 3.21 about here

The proportions of variance accounted for in the obtained data by the simulated data are shown for each S/N ratio. The degree to which these proportions differ from 1 represents an estimation of the variance introduced by the number of observations alone. On the average, 11% of the overall variance in the obtained data for the +15 and +5 S/N ratios and 16% of the variance for the -5 S/N ratio was due to the low number of observations alone. That is, over 10% of the overall variance at each S/N ratio can be attributed to the number of observations per word. Thus, the predictor variables examined previously need explain approximately only 90% of the overall variance, given that 10 to 16% of the variance observed would be expected to arise from the number of observations per word alone. In short, although the relatively low number of observations per word did not introduce an overly large degree of variability in the overall variance, at least as indexed by the simulation, consideration of this variance demonstrates that the neighborhood probability rule performed better than actually indicated by the original analyses.

Table 3.21

Simulated proportion of variance contributed by number of observations.

	S/N Ratio		
	+15	+5	-5
	-----	-----	-----
Mean proportion of variance accounted for over 20 simulations	0.8937	0.8914	0.8419
Estimated variance contributed by number of observations	0.1063	0.1086	0.1581

Qualitative Predictions of the
Neighborhood Probability Rules

Up to this point, the analyses have been primarily concerned with demonstrating the degree to which the neighborhood probability rule correlates with or classifies identification performance. The results of these analyses support the hypothesis that a word is identified in the context of similar words and that the frequency relations among the stimulus word and its neighbors are important determinants of identification performance. Heretofore, there has been little attention paid to the precise predictions of the neighborhood probability rule. In this section, I turn to a discussion of some of the qualitative predictions of the neighborhood probability rules adopted from the previous analyses and the implications of these predictions for auditory word recognition. A closer examination of these predictions should delineate more specifically the type of model necessary for describing auditory word recognition.

Recall that the frequency-weighted neighborhood probability rule (which proved most predictive of identification) has the form:

$$(3.10) \quad p(\text{ID}) = \frac{\text{SWP} * \text{FreqS}}{\text{SWP} * \text{FreqS} + \sum_{j=1}^n \text{NWP}_j * \text{FreqN}_j} ;$$

where $p(\text{ID})$ is the probability of identifying the stimulus word, SWP is the stimulus word probability obtained from the confusion matrices, NWP_j is the neighbor word probability of neighbor j obtained from the confusion matrices, FreqN_j is the frequency of the j th neighbor, and FreqS is the frequency of the stimulus word. Basically, this rule states the probability of choosing the stimulus from among its neighbors. Note that frequency is built into the rule in terms of a frequency weight applied to the stimulus word and each of its neighbors.

The rule thus predicts that identification performance is a function of the intelligibility of the stimulus word, the confusability of its neighbors, and the frequencies of the stimulus word and its neighbors. Frequency thus serves to bias, positively or negatively, the choice of a word from its neighborhood. Note that the effects of frequency are contingent on the nature of the words residing in the similarity neighborhood. As in Triesman's (1978a,b) partial identification theory (see below), frequency effects are assumed in the rule to be relative. For example, high frequency stimulus words residing in neighborhoods containing high frequency neighbors are predicted by the rule to be identified at approximately equal levels of performance to low frequency words residing in low frequency neighborhoods, assuming that stimulus intelligibility and neighborhood confusability are held constant. That is, the frequency of the stimulus word alone will not determine identification performance. Instead, stimulus word frequency must be evaluated in terms of the frequencies of the neighbors of the stimulus word, as well as the confusability of the neighbors. Thus, the rule implies a complex relation between the stimulus word and its neighbors, such that stimulus frequency, neighbor frequency, stimulus intelligibility, and neighborhood confusability will act in combination to determine identification performance.

The rule therefore makes a number of important predictions depending on the stimulus word probability and the sum of the neighbor probabilities. For simplicity, I will define the sum of the frequency-weighted neighbor word probabilities as the overall frequency-weighted neighborhood probability, or FWNP. Inspection of equation (3.10) reveals that if the FWSWP (i.e., $SWP * FreqS$) is held constant, as FWNP increases, predicted identification will decrease. Likewise, if FWNP is held constant, then increases in FWSWP will result in corresponding increases in predicted identification. The interesting cases arise, however, when both the FWSWP and FWNP are allowed to vary. Consider the four cases in which the FWSWP and FWNP can take on either high or low values: (1) FWSWP high-FWNP high, (2) FWSWP high-FWNP low, (3) FWSWP low-FWNP high, and (4) FWSWP low-FWNP low. The predictions of the neighborhood probability rule for these four cases is shown in Table 3.22.

Insert Table 3.22 about here

As shown in Table 3.22, the rule predicts best performance of those words with high FWSWP's and low FWNP's. These are words that, in a sense, "stand out" in their neighborhoods. The lowest performance is predicted for words with low FWSWP's and high FWNP's. These are words that are least distinguishable in their neighborhoods. Interestingly, however, the rule predicts intermediate levels of performance for the remaining two cases. That is, words with high FWSWP's and high FWNP's are predicted to show approximately equal levels of identification performance to words with low FWSWP's and low FWNP's. Thus, the rule does not always predict an advantage for high frequency words over low frequency words. In addition, the rule predicts that words matched on FWSWP may show differential levels of performance depending on the FWNP, or frequency-weighted neighborhood structure.

To determine if the general pattern of predictions stated in Table 3.22 hold for the present set of identification data, the following analyses were performed. For each of the 811 words, median values for the FWSWP's and FWNP's collapsed across S/N ratio were determined. These median values were then used to divide the stimulus words into classes having high and low FWSWP's and high and low FWNP's. Altogether, four cells were analyzed (two levels of FWSWP X two levels of FWNP). Mean identification scores, collapsed across S/N ratio, were then computed for words falling into each of the four cells.

The results for the classification of word scores by FWSWP and FWNP are shown in Table 3.23. Predicted levels of performance are shown in parentheses.

Insert Table 3.23 about here

As shown in this table, the pattern of results predicted by the neighborhood probability were clearly present in the identification data. As predicted, words with high FWSWP's and low FWNP's were responded to with the highest levels of accuracy; words with low FWSWP's and high FWNP's were responded to with the lowest levels of accuracy. The remaining two cases, as predicted, showed intermediate and nearly identical levels of identification performance. Note that words matched on FWSWP were responded to quite

Table 3.22

Predicted identification performance as a function of frequency-weighted stimulus word probability and frequency-weighted neighborhood probability.

		FREQUENCY-WEIGHTED STIMULUS WORD PROBABILITY	
		LOW	HIGH
FREQUENCY-WEIGHTED NEIGHBORHOOD PROBABILITY	LOW	Intermediate	High
	HIGH	Low	Intermediate

Table 3.23

Obtained identification performance (percent correct) as a function of frequency-weighted stimulus word probability and frequency-weighted neighborhood probability. Qualitative predictions are in parentheses.

		FREQUENCY-WEIGHTED STIMULUS WORD PROBABILITY	
		LOW	HIGH
FREQUENCY-WEIGHTED NEIGHBORHOOD PROBABILITY	LOW	50.56 (Intermediate)	64.03 (High)
	HIGH	37.76 (Low)	54.73 (Intermediate)

differently depending on the FWNP, demonstrating that stimulus word frequency is a direct function of the neighborhood in which the stimulus word occurs. This is also demonstrated by the cases showing intermediate levels of performance. Although the words in these cells differ substantially in their FWSWP's, they show nearly identical levels of identification performance, owing to the composition of their similarity neighborhoods. In short, the present analysis provides further support for the claim that auditory word recognition is the result of a complex interaction of stimulus word intelligibility, stimulus word frequency, and neighborhood confusability and frequency.

Neighborhood Activation Model

The neighborhood probability rules developed above provide the groundwork on which to base a model of auditory word identification, which will be called the neighborhood activation model. The basic postulate of the model is that the process of word identification involves discrimination among lexical items in memory that are activated on the basis of stimulus input. This is a fundamental principle in almost every current model of word recognition (e.g., Forster, 1976, 1979; Marslen-Wilson & Welsh, 1978; Paap, Newsome, McDonald, & Schvaneveldt, 1982). Indeed, one of the most important issues in auditory word recognition concerns the processes by which discrimination among lexical items in memory is achieved. The present model attempts to specify those factors responsible for the relative ease or difficulty of recognizing words arising from the processes involved in discrimination among sound patterns of words. Thus, a second fundamental principle of the model is that discrimination is a function of the number and nature of lexical items activated by the stimulus input. The "nature" of lexical items refers specifically to the acoustic-phonetic similarity among the activated lexical items as well as their frequencies of occurrence. The model will thus be concerned with the long-standing issue of word frequency. However, characterizing the effects of word frequency is only a part of the present model. Instead, the model focuses primarily on structural issues concerning the process of lexical discrimination. Word frequency is important in the model only as a factor affecting the structural relationships among lexical items.

The hypothesis that lexical discrimination is a function of the nature and number of activated lexical items implies that word recognition cannot be represented by a "no cost" system, as at least one recent theory has proposed (see Marslen-Wilson, 1986). Instead, the ease of word recognition is predicted to vary as a function of the degree of difficulty in discriminating among lexical items. A corollary of this claim is that words vary in terms of their structural relationships to other words. That is, words vary in the number and nature of lexical items that they will activate in memory. Given this set of hypotheses, it becomes incumbent on the model to provide an adequate account of the effects of neighborhood structure on word recognition.

The precise form of the model owes much to Triesman's (1978a,b) partial identification theory. Triesman's earlier work presages many of the concepts of the neighborhood activation model. In partial identification theory, a given stimulus word serves to define an acoustic subvolume of words that are similar to the stimulus word. Although other theories of auditory word identification have made similar propositions (e.g., sophisticated guessing theory, Broadbent, 1967; see also Catlin, 1969, Newbigging, 1961, Savin, 1963, and Soloman & Postman, 1952), Triesman's partial identification theory emphasizes that word frequency effects may vary as a function of the nature of

the acoustic subvolume activated. In particular, Triesman argues that a given acoustic subvolume may be "dense" or "rarified." Moreover, the words within an acoustic subvolume may be high or low in frequency. Triesman predicts that word frequency effects will vary as a function of the number and nature of words activated in a given acoustic subvolume.

The present model bears a strong resemblance to partial identification theory with respect to the effects of neighborhood structure (acoustic subvolume) on word identification. However, the present set of data extends partial identification theory in a number of important ways. First, Triesman's arguments stem from an analysis of data obtained in closed response formats using highly restricted sets of words (letters or digits). The present data used both an open response format and a relatively highly unconstrained set of consonant-vowel-consonant words. In addition, Triesman defined acoustic subvolumes on the basis of the error responses. In the present set of experiments, similarity neighborhoods for the words were computed on two independent sources of data: phoneme confusion matrices and Webster's lexicon. Thus, the present data, although subject to more variability than the data examined by Triesman, compose a much more "naturalistic" data base. Nevertheless, to the extent that the present data demonstrate that the classic word frequency effect is a function of the neighborhood in which the stimulus word resides, these data provide support for Triesman's theory.

A flow chart of the neighborhood activation model is shown in Figure 3.2.

Insert Figure 3.2 about here

Upon presentation of stimulus input, a set of acoustic-phonetic patterns are activated in memory. It is assumed that all patterns are activated regardless of whether they correspond to real words in the lexicon or not, an assumption required by the fact that listeners can recognize the acoustic-phonetic form of novel words and nonwords. As in Triesman's partial identification theory, the acoustic-phonetic patterns are activated in a multidimensional acoustic-phonetic space in which the dimensions correspond to phonetically-relevant acoustic differences among the patterns. Specification of the nature of these dimensions poses an important problem for any complete theory of speech perception and auditory word recognition (see Luce & Pisoni, 1987). However, the present model is neutral with respect to the dimensions of the space. The only requirement of the model is that the dimensions of the space produce relative activation levels among the acoustic-phonetic patterns that are isomorphic with the dimensions of similarity to which subjects are sensitive.

The acoustic-phonetic patterns then activate a system of word decision units tuned to the patterns themselves. A diagram of a single decision unit is shown in Figure 3.3.

Insert Figure 3.3 about here

Only those acoustic-phonetic patterns corresponding to words in memory will activate a word decision unit. Neighborhood activation is assumed to be identical to the activation of the word decision units. The claim that decision units are not tuned to every acoustic-phonetic pattern activated is supported by the demonstration of differential effects of neighborhood

Neighborhood Activation Model (NAM)

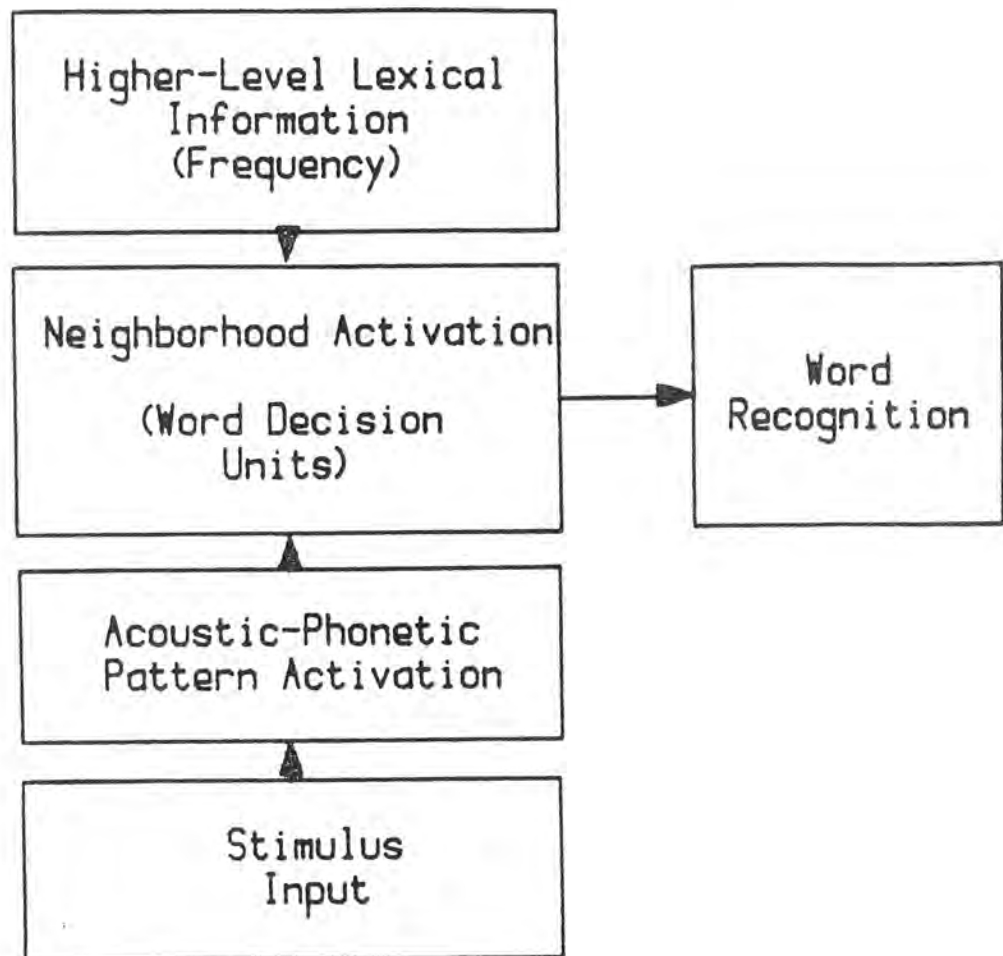


Figure 3.2. Flow chart for the neighborhood activation model.

Word Decision Unit

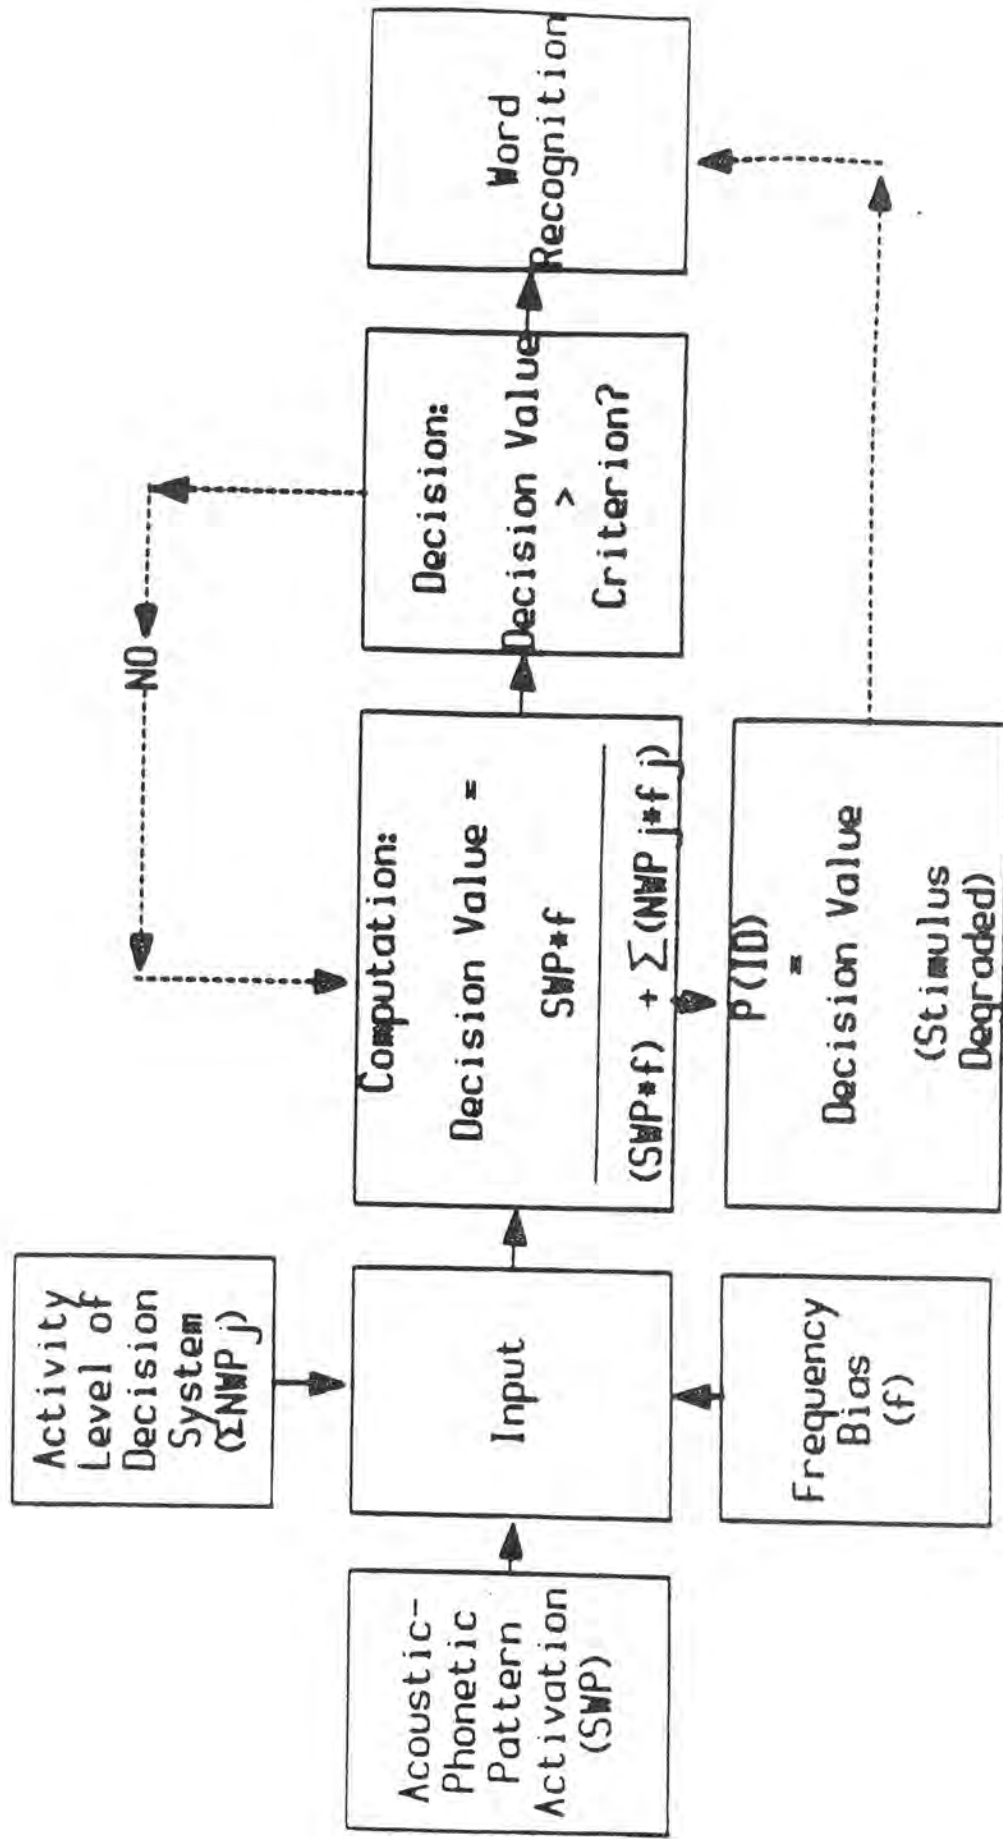


Figure 3.3. Diagram of a single word decision unit.

confusability on perceptual identification. Were the decision units tuned to every pattern, each neighborhood would be maximally dense prior to the biasing effects of frequency. Thus no effects of neighborhood confusability independent of frequency should be observed, a prediction contradicted by the present set of data.

Once activated, these decision units monitor the activation levels of the acoustic-phonetic patterns to which they correspond. Following activation of the word decision units, these units then begin monitoring higher-level lexical information relevant to the words to which they correspond. Word frequency is included in this higher-level lexical information. In addition to monitoring higher-level lexical information in long term memory, the decision units are also assumed to monitor any information in short term memory that is relevant to making a decision on the identity of a word.

The system of word decision units is the crucial aspect of the neighborhood activation model. These units serve as the interface between acoustic-phonetic information and higher-level lexical information. Acoustic-phonetic information drives the system by activating the word decision units, affording priority to bottom-up information, as in cohort theory (Marslen-Wilson & Welsh, 1978). Higher-level lexical information such as frequency is assumed to operate by biasing the decision units. These biases operate by adjusting the activation levels of the acoustic-phonetic patterns represented in the decision units. The biases introduced by higher-level lexical information need not be under volitional control nor need they be conscious (see Smith, 1980). Instead, these biases are assumed to be a fundamental aspect of word perception that enable optimization of the word recognition process via the employment of a priori probabilities and contextual information.

Each word decision unit is therefore responsible for monitoring two sources of information, acoustic-phonetic pattern activation and higher-level lexical information. In addition, the decision units are assumed to be interconnected in such a way that each unit can monitor the overall level of activity in the system of units, as well as the activity level of the acoustic-phonetic patterns to which the units correspond (see Elman & McClelland, 1986). As analysis of the stimulus input proceeds, the decision units continuously compute decision values. These values are assumed to be computed via a rule of the type described by the neighborhood probability rule. In the neighborhood probability rule, the stimulus word probability (SWP) corresponds to the activation level of the acoustic-phonetic pattern. The sum of the neighbor word probabilities (NWP_j's) corresponds to the overall level of activity in the decision system. Frequency information serves as a bias, as in the frequency-weighted neighborhood probability rule, by adjusting the activation levels of the acoustic-phonetic patterns represented in the word decision units.

As processing of the stimulus input proceeds, the acoustic-phonetic pattern corresponding to the stimulus input is resolved or "refined" (see Pisoni, Nusbaum, Luce, & Slowiczek, 1985). As the pattern is refined, the activation levels of similar patterns drop and the decision values computed by the word decision unit monitoring the pattern of the stimulus steadily increase. Once the output of a given decision unit reaches criterion, all information monitored by that decision unit is made available to working memory. "Word recognition" is accomplished once the word decision unit for a given acoustic-phonetic pattern surpasses criterion (i.e., the acoustic-phonetic pattern is recognized). "Lexical access" occurs when higher-level lexical information (i.e., semantic and syntactic information,

pragmatic information) is made available to working memory. The term lexical access is actually somewhat misleading in the context of the neighborhood activation model. Lexical information is monitored by the word decision units once these units are activated. However, this information is used only in the service of choosing among the activated acoustic-phonetic patterns and is therefore not available to working memory. Lexical access in the neighborhood activation model is thus assumed to occur when lexical information is made available for further processing. The word decision units in the model therefore serve as gates on the lexical information available to the system (see Morton, 1979). In so doing, the units prevent the cognitive system from "over-resonating," making information available only once a decision is made as to the identity of the stimulus input.

The system of word decision units is, obviously, quite powerful. However, the system is no more powerful than current interactive-activation models (Elman & McClelland, 1985; McClelland & Rumelhart, 1981). Indeed, further research may reveal that the system of decision units can be adequately described by an interactive-activation system consisting of interconnected nodes. The decision values themselves may be computed via inhibitory and excitatory links among the decision units. At the present time, however, the decision units will be assumed to make decisions on the basis of the overall activation level of the system as well as the activation level of the acoustic-phonetic patterns.

The postulation of a system of word decision units is based on the finding that the frequency-weighted neighborhood probability rule adequately predicted identification performance. Indeed, the system of word decision units is simply a processing instantiation of the neighborhood probability rule. However, the neighborhood activation model, by instantiating the neighborhood probability rule in a system of decision units, makes a number of important claims. First, it is assumed that the word recognition system is, at least initially, completely driven by the stimulus input. Frequency information is thus assumed only to bias the decision units and not to affect the encoding of the acoustic-phonetic patterns. Thus, frequency information is not assumed to be an intrinsic part of the activation levels of the acoustic-phonetic patterns, but is assumed to be a bias that must be interpreted in the context of the frequencies of all other words. If frequency information were assumed to be intrinsic to the activation levels of the acoustic-phonetic patterns and no decisions were made based on the total activity of the system, low frequency words would be responded to less accurately than high frequency words regardless of their neighborhood structures, which is clearly in contradiction to the data reported above.

The neighborhood activation model thus captures an important component of the word frequency effect that is not accounted for by models of word recognition that assume differential thresholds for high and low frequency words. For example, Morton's (1979) logogen theory assumes that high frequency words have lower thresholds than low frequency words. In logogen theory, frequency effects arise from the claim that high frequency logogens require less evidence to surpass threshold than low frequency words. Fixed thresholds cannot account for the results discussed above that stimulus word frequency must be evaluated in the context of the frequencies of the neighbors of the stimulus word. Search (e.g., Forster, 1976, 1979) and verification (e.g., Becker, 1976; Paap et al., 1982) models may fare somewhat better. In these models, a neighborhood of items may be activated and then searched or verified according to frequency. However, the search and verification models, as well as the threshold models, are difficult to countenance in light of the findings that word frequency effects may be attenuated or severely reduced

given certain task requirements (Balota and Chumbley, 1984, 1985), response sets (Pollack, Rubenstein, & Decker, 1959), or contextual information (Grosjean & Itzler, 1984; Luce, 1983; Miller, Heise, & Lichten, 1951). In particular, it is difficult for these models to account for the finding that frequency effects are absent when the response set is known to the subject, at least within limits (Pollack et al., 1959). Although assumptions could be incorporated in each of these models to account for such a finding, the fact that the effect of frequency is a biasing effect not intrinsic to the decision unit system provides a straightforward means of interpreting various results regarding word frequency. In particular, greater or lesser weight can be added to the frequency bias in the decision units depending on the requirements of the task. In the case in which the response set is known and the stimuli have equal probabilities, it may be assumed that the decision units will optimize performance by ignoring a priori probabilities. In short, one of the advantages of the neighborhood activation model is that frequency is not an intrinsic part of the units responsible for producing word frequency effects. The lability of frequency effects is thus accounted for by the relative importance attached to the biasing properties of frequency information on the word decision units.

At present, the neighborhood activation model makes no claims about the temporal course, or left-to-right nature, of word recognition, an issue of primary importance in cohort theory (Marslen-Wilson & Welsh, 1978). Undoubtedly, some of the aspects of cohort theory may have to be incorporated into the present model to account for the fact that words are processed in time, or "left-to-right." Specification of the time-course of lexical activation is, of course, a crucial issue, and cohort theory has provided an elegant account of the possible processes involved in the left-to-right processing of words. However, cohort theory to date has been evaluated almost entirely on the basis of the processing of relatively long words, which prove most amenable to demonstrations of temporal processing. The present data, on the other hand, is based on short, monosyllabic words. It has yet to be shown that the claims of cohort theory hold for such short stimuli (see Grosjean, 1985). For the moment, then, the issue of the time-course of processing will be deferred. I will adopt the assumption that, within the relatively short time windows provided by the stimuli used in the present study, the activation levels across time are more or less equivalent. Thus, no priority will be afforded to the early portions of words in the neighborhood activation model, at least for short words (see Luce, 1986a).

Having laid out a framework for interpreting neighborhood structural and frequency effects, I will now turn to a discussion of how the neighborhood activation model accounts for the results of the perceptual identification study reported in this chapter and Chapter 2. Recall that the neighborhood activation model, under normal circumstances, recognizes a word once the decision value for a given word exceeds criterion. It is assumed that stimulus degradation affects the word recognition system by impeding complete processing of the stimulus input. That is, only so much information can be obtained from the stimulus input when it is masked by noise. Given imperfect information, then, it is assumed that, in the long run, no decision unit will reach criterion, and a decision will thus be forced on the available information. The "available information" is the state of the decision system at the point at which processing of the acoustic-phonetic information is completed. In perceptual identification, therefore, a response is made on the basis of the values of the decision units at the point at which processing is completed. Thus, the neighborhood probability rule developed above expresses the probability of choosing the stimulus word actually presented. If the stimulus input results in a large number of highly confusable, high frequency

neighbors, the probability of actually recognizing the stimulus word will be low. Likewise, if the stimulus input results in only a few confusable, low frequency neighbors, the probability of identification will be high.

Note that because the decision units monitor both the activation levels of the acoustic-phonetic patterns as well as the overall activation of the decision system, probability of identification will not depend solely on the intelligibility of the stimulus word nor on neighborhood confusability. Words of low intelligibility with few confusable neighbors are predicted by the model to be equivalent to words of high intelligibility with many confusable neighbors. Indeed, as shown above, this prediction was borne out. In short, perceptual identification is a function of the values of the decision units computed at the completion of stimulus processing. Furthermore, the role of stimulus degradation is assumed to be one of impeding complete processing of the stimulus input.

Summary and Conclusions

The neighborhood activation model provides a framework for instantiating the neighborhood probability rules developed here. To the extent that the neighborhood probability rules predict identification performance, the model can be deemed an adequate account of the word identification process. Indeed, these rules were shown to make a number of precise predictions about the relative effects of stimulus intelligibility and neighborhood structure that were borne out by the data. In particular, the neighborhood probability rule predicts a complex interrelationship between the stimulus word and its neighbors. In addition, the frequency-weighted neighborhood probability rule was able to account for the stimulus word frequency effect in terms of the frequency relationships between the stimulus word and its neighbors. Indeed, it was shown that stimulus frequency alone was a poor predictor of identification performance. The picture that emerges from the present set of data is one that emphasizes the roles of discrimination and decision in auditory word recognition. Finally, both the data and the neighborhood activation model underscore the degree to which the structure of the mental lexicon influences word identification: Precise accounts of the process of auditory word recognition are crucially tied to detailed accounts of the structural relationships among lexical items in memory.

Footnotes

¹ Some caution must be taken in comparing directly the results from the three different S/N ratios. The results for the +15 dB S/N ratio are highly skewed toward the upper end of the distribution; that is, most words at the +15 dB S/N ratio were identified with a high level of accuracy. Likewise, the results for the -5 dB S/N ratio show a similar skewing toward the lower end of the distribution. The skewness for both of these distributions results in relatively smaller degrees of overall variance, compared to the +5 dB S/N ratio. Simply put, more variance is available for explanation at the +5 S/N ratio than at either of the two other S/N ratios. Thus, direct comparisons among the three S/N ratios in terms of proportion of variance accounted for by each variable is dubious.

² One alternative to the use of the Kucera and Francis frequency counts was to use the familiarity ratings obtained by Nusbaum, et al. (1984). To determine if the familiarity ratings may be in fact superior estimates of frequency for the present data, all the analyses previously reported were conducted using the familiarity ratings. The correlations of the familiarity ratings were consistently smaller than those for the SFI. In addition, the use of familiarity ratings in the neighborhood probability rules resulted in lower correlations than the use of the Kucera and Francis frequencies. The failure of the familiarity ratings to correlate more highly with identification performance than the SFI's may have been due to the fact that the stimulus words used in this study all had familiarity ratings of 6.0 or above on a seven point scale. There was thus far less variability in the familiarity ratings than in the Kucera and Francis frequencies, and this reduced variability may have resulted in lower overall correlations. Nevertheless, the familiarity ratings proved less satisfactory than the Kucera and Francis frequencies and were thus not used as frequency indices in the present study.

CHAPTER FOUR

EVIDENCE FROM AUDITORY LEXICAL DECISION

The results from the perceptual identification experiment reported in Chapters 2 and 3 lend strong support to the notion that subjects' accuracy at identifying words masked by noise is critically dependent on the number and nature of words activated in memory by the stimulus word. In particular, it was shown that the degree of confusability of the neighbors of a stimulus word, as well as the frequencies of these neighbors, influences identification performance in predictable and systematic ways. Furthermore, stimulus word frequency per se proved to be a poor predictor of identification performance. Instead, it was shown that word frequency must be interpreted in the context of the frequencies of the words activated in memory. A neighborhood probability rule was developed that incorporated stimulus word intelligibility, frequency, and neighborhood structure into a single expression that adequately predicted identification performance. Finally, based on the neighborhood probability rule, a working model of auditory word recognition was proposed--the neighborhood activation model--that attempts to describe the effects of neighborhood structure on auditory word recognition.

The purpose of the present chapter is to explore further the effects of neighborhood structure on auditory word recognition. In particular, the lexical decision paradigm was employed to examine these effects. In the lexical decision paradigm, a subject is presented with a real word or a nonsense word, or nonword. The subject's task is to decide as quickly but as accurately as possible whether a given stimulus is a word or nonword. The lexical decision task has proven quite useful in visual word recognition research in examining the effects of such variables as word frequency (Stanners, Jastrzemski, & Westbrook, 1975; Whaley, 1978; see also Forster, 1979). In general, it has been shown that high frequency words tend to be classified as words more quickly than low frequency words. Indeed, this has been a very robust finding in the literature, although there are numerous, and often times conflicting, accounts of frequency effects in lexical decision (Balota & Chumbley, 1984; Glanzer & Ehrenreich, 1976; Gordon, 1983; Paap, McDonald, Schvaneveldt, & Noel, 1986). An auditory analog of the visual lexical decision task thus presents a useful means of examining word frequency effects and the effects of neighborhood structure on auditory word recognition.

The use of the lexical decision task is also attractive for two other reasons. First, investigation of the process of auditory word recognition can be carried out in the absence of stimulus degradation. Although the perceptual identification experiment reported in Chapters 2 and 3 provided useful data regarding the effects of stimulus word frequency and neighborhood structure, a more robust test of these effects hinges on the demonstration that neighborhood structural effects can be demonstrated in the absence of stimulus degradation. In other words, it is important to demonstrate that the effects of neighborhood structure generalize beyond words that are purposefully made difficult to perceive.

The second advantage of the auditory lexical decision task is the ability to collect reaction time data. The reaction time data may aid in uncovering some of the temporal aspects of the effects of neighborhood structure on auditory word recognition. Furthermore, it is of crucial importance to demonstrate that neighborhood structure affects not only the accuracy of word recognition, but also the time course. Thus, the auditory lexical decision

task provides a useful means of corroborating and extending the findings from the previous perceptual identification study.

The approach taken in the present chapter is similar to that in Chapter 2. Similarity neighborhood statistics, computed on the basis of Webster's lexicon, served as independent variables. The statistics of interest were again: (1) the number of words similar to a given stimulus word, or neighborhood density; (2) the mean frequency of the similar words or neighbors; and the (3) frequency of the stimulus word itself. On the basis of the neighborhood activation model, the following general predictions were tested: First, it was predicted that the number of words similar to a given stimulus word or nonword will affect classification time and accuracy. In particular, it was hypothesized that increasing neighborhood density will increase the time it takes to discriminate among lexical items activated in memory. Furthermore, increases in density may be accompanied by lower levels of accuracy, given increased competition among items in memory. It was furthermore predicted that the frequencies of the neighbors will influence classification time and accuracy. Specifically, higher frequency neighborhoods should produce slower reaction times and lower levels of accuracy than low frequency neighborhoods, owing to the higher degree of competition associated with high frequency neighbors. In short, the goal of the present experiment is to examine further the effects of neighborhood density, neighborhood frequency, and stimulus word frequency in order to gain further support for the neighborhood activation model

EXPERIMENT

Method

Stimuli

Words. The same 918 words used in the perceptual identification study reported in Chapters 2 and 3 were used in the auditory lexical decision experiment. As reported in Chapter 2, the 918 words were randomly partitioned into three stimulus set files consisting of 306 words each. Although all 918 words were presented for lexical decision, the 36 words failing to reach criterion in the screening experiment reported in Chapter 2 were excluded from all subsequent analyses. The method of stimulus preparation was described in Chapter 2.

Similarity neighborhood statistics for each of the word stimuli were computed according to the procedure described in Chapters 1 and 2. The phonetic transcription for each stimulus word was compared to the phonetic transcriptions of all other monosyllabic words having familiarity ratings of 5.5 or above. A neighbor of a stimulus word was defined as a word that could be converted to the stimulus word itself by a one phoneme insertion, deletion, or substitution in any position. Neighborhood densities, neighborhood frequencies, and stimulus word frequencies were determined in this manner for each word.

Nonwords. In order to construct a list of phonotactically legal nonwords matched in phoneme length to the word stimuli, a lexicon of nonwords was constructed in the following manner: For all three phoneme words in Webster's lexicon, all initial two-phoneme sequences and all final two-phoneme sequences were determined. That is, for all three phoneme words, P1-P2-P3, all P1-P2 sequences and all P2-P3 sequences were determined. All initial and final sequences sharing P2 were then combined. Three phoneme sequences not

containing a vowel were excluded. Also excluded were any sequences corresponding to a real word in Webster's lexicon. Because both the initial and final biphones of the nonwords actually occurred in real words, the resulting nonword lexicon thus contained three phoneme sequences that strongly followed the phonotactic constraints of real words. Altogether, 3123 nonwords were constructed.

Similarity neighborhood statistics for each of the 3123 nonwords were then computed. The similarity neighborhood statistics were computed by comparing each nonword with each word in Webster's lexicon. Similarity neighborhoods were once again computed on the basis of one phoneme substitutions, additions, and deletions. For the nonwords, two variables were of interest: (1) neighborhood density, or the number of words similar to a given nonword, and (2) mean neighborhood frequency, or the mean frequency of the words similar to a nonword. Note that the neighborhood statistics for the nonwords were computed based on words only.

Three-hundred and four stimuli were selected from the nonword lexicon that fell into one of four cells, resulting in 76 nonwords per cell. The four cells were produced by crossing two levels of density (high and low) with two levels of mean neighborhood frequency (high and low). These cells were: (1) high density-high neighborhood frequency, (2) high density-low neighborhood frequency, (3) low density-high neighborhood frequency, and (4) low density-low neighborhood frequency. Selection of the 76 nonwords for each of the four cells was achieved via an algorithm that first rank-ordered each of the 3123 nonwords on each of the two independent variables. A method of minimizing and maximizing squared deviations of successively ranked nonwords was then employed to ensure that cells that were matched on a given variable (e.g., both high density) were maximally alike and that cells intended to differ on a given variable (e.g., one high and one low density) were maximally different. Nonwords occurring in high density neighborhoods had an average of 17.78 neighbors; nonwords occurring in low density neighborhoods had an average of 8.10 neighbors. The mean frequency of high frequency neighborhoods was 156.96; the mean frequency of low frequency neighborhoods was 11.84.

The nonwords were recorded by the same male speaker who produced the words. The method for recording and digitizing the nonwords was identical to that for the words (see Chapter 2). Overall RMS amplitude for the nonwords was equated to the overall amplitudes of the words.

Subjects

Thirty subjects participated in partial fulfillment of an introductory psychology course. All subjects were native English speakers and reported no history of speech or hearing disorders.

Design and Procedure

Three stimulus set files were constructed by combining each of the three set files of words with the set file containing the nonwords. Each of the three set files thus contained 306 words and 304 nonwords, producing 610 stimuli. (Three-hundred and four nonwords, instead of 306, were selected to enable equal partitioning of the nonwords into the four cells. It was assumed that the slight discrepancy between the number of words and nonwords would have little or no effect on the obtained results given the large number of words and nonwords used.) Each set file was presented to a total of 10

subjects. Thus, 10 observations were obtained for each word, whereas 30 observations were obtained for each nonword, given that the same set of nonwords were presented to each group of subjects.

Stimulus presentation and data collection were controlled by a PDP-11/34 minicomputer. The stimuli were presented via a 12-bit digital-to-analog converter at a 10 kHz sampling rate over matched and calibrated TDH-39 headphones at a comfortable listening level of 75 dB SPL.

Groups of five or fewer subjects were tested in a sound-treated room. Each subject sat in an individual booth equipped with a two-button response box. The button on the left-hand side of the response box was labeled "WORD"; the button on the right-hand side of the response box was labeled "NONWORD." A small light was situated above each button for feedback. In addition, a cue light was situated at the top of the box to warn the subject that a stimulus was about to be presented. Subjects were instructed that they would hear real words in English and nonsense words, or nonwords. They were instructed that after presentation of each stimulus, they were to respond whether they heard a word or a nonword by pressing the appropriately labeled button on the response box. The instructions stressed both speed and accuracy.

A given trial proceeded as follows: The cue light at the top of the response box was illuminated for one second warning the subject that a stimulus was about to be presented. Five-hundred msec after the offset of the cue light, a randomly selected auditory stimulus was presented. Immediately after the subject responded word or nonword, the light above the button that should have been pressed for a correct response was illuminated for one second. Reaction times were recorded from the onset of the auditory stimulus to the response. After each subject had responded, a new trial was initiated. If one or more subjects failed to respond within 4000 msec of the onset of the auditory stimulus, incorrect responses for those subjects were tallied and a new trial was initiated. An inter-trial interval of 500 msec elapsed between the end of one trial and the beginning of the next. The 610 experimental trials were preceded by 30 practice trials consisting of an equal number of randomly presented words and nonwords. None of the words or nonwords presented in the practice phase of the experiment were presented in the experiment proper. An experimental session lasted approximately one hour.

Results

Analysis of Word Responses

To factor out the effect of stimulus word duration on reaction times, the duration of each stimulus word in msec was subtracted from each subject's reaction time to that word. The adjusted reaction times for correct word responses were then entered into a stimulus word-by-subjects array. Means and standard deviations for each stimulus as well as each subject were then computed. Those reaction times falling above or below 2.5 standard deviations of both the subject and stimulus means were deleted and replaced according to the procedure suggested by Winer (1971). Reaction times were then averaged across subjects, producing a mean reaction time for each stimulus word. The number of correct responses for each stimulus word were also tallied and converted to percentages.

The output of the similarity neighborhood statistics were combined with the mean reaction times and percentages correct and submitted to correlation and regression analyses. As stated above, three neighborhood statistics were

of interest for the words: neighborhood density, mean neighborhood frequency, and the frequency of the stimulus word itself. Two variations of each of the frequency-based statistics (stimulus word frequency and mean neighborhood frequency) were computed, one based on raw frequencies and one based on Standard Frequency Indices (SFI's). SFI's are basically log transforms of the raw frequencies (see Chapter 2). In total, five independent variables were examined. These variables, as well as their mnemonics, are shown in Table 4.1. The five variables were: (1) raw stimulus word frequency (SWF-RAW), (2) SFI of the stimulus word (SWF-SFI), (3) neighborhood density (DEN), (4) mean neighborhood frequency based on raw frequencies (NHF-RAW), and (5) mean neighborhood frequency based on SFI's (NHF-SFI). Table 4.2 shows the means, standard deviations, and minimum and maximum values for each of the independent variables. Also shown in Table 4.2 are means, standard deviations, and minimum and maximum values for the adjusted reaction times (RT) and percentages correct (SCORE).

Insert Tables 4.1 and 4.2 about here

Correlation Analyses

Each of the five independent variables was first correlated with the reaction times and percentages correct for each word. The results of these correlation analyses are shown in Table 4.3.

Insert Table 4.3 about here

For the reaction times, only stimulus word frequency based on the SFI's produced a significant correlation, $r = -.2776$. Reaction times were negatively correlated with frequency, demonstrating that high frequency words tended to be responded to more quickly than low frequency words. No significant effects of neighborhood density or neighborhood frequency were observed for the reaction times. For the percentages correct, stimulus word frequency and neighborhood density produced significant correlations, $r = .2999$ and $r = .1173$, respectively. The significant positive correlation of stimulus word frequency and percent correct demonstrates that high frequency words tend to be classified correctly more often as words than low frequency words. The positive correlation of density and percent correct indicates that subjects were better able to correctly classify words that occurred in high density neighborhoods than those occurring in low density neighborhoods. Note that this correlation is contrary to the predicted effect of density on classification accuracy. I will return to a detailed discussion of this finding below.

Table 4.1

Neighborhood Analysis: Auditory Lexical Decision. Variables and mnemonics.

VARIABLE	MNEMONIC
1. Raw stimulus word frequency	SWF-RAW
2. SFI stimulus word frequency	SWF-SFI
3. Neighborhood Density	DEN
4. Mean neighborhood frequency based on raw frequencies	NHF-RAW
5. Mean neighborhood frequency based on SFI's	NHF-SFI

Table 4.2

Neighborhood Analysis: Auditory Lexical Decision. Summary statistics.

VARIABLE	MEAN	STANDARD DEVIATION	MINIMUM VALUE	MAXIMUM VALUE
-----	-----	-----	-----	-----
1. SWF-RAW	129.1564	556.2139	1.0000	10595.0000
2. SWF-SFI	52.1884	8.4655	40.0000	80.2500
3. DEN	16.5011	6.7651	1.0000	35.0000
4. NHF-RAW	205.9523	349.8718	2.6000	4389.0000
5. NHF-SFI	52.0241	3.1903	42.5508	66.4250
6. RT	424.4785	115.5683	128.0000	931.0000
7. SCORE	0.8961	0.1473	0.1000	1.0000

Table 4.3

Neighborhood Analysis: Auditory Lexical Decision. Correlations of predictor variables with accuracy scores and reaction times.

r

VARIABLE	P(C)	RT
1. SWF-RAW	.0522	-0.0609
2. SWF-SFI	.2999*	-0.2776*
3. DEN	.1173*	0.0324
4. NHF-RAW	.0345	0.0051
5. NHF-SFI	.0657	-0.0551

*p<0.05

Regression Analyses

To further examine the relative roles of the five independent variables on word classification times and accuracy, stepwise multiple regression analyses were performed separately on reaction times and percentages correct. The results of these analyses are shown in Table 4.4. Shown are the multiple R's, multiple R²'s, and change in R² at each step.

Insert Table 4.4 about here

For the reaction times, stimulus word frequency based on SFI's entered first, accounting uniquely for 7.71% of the variance. Word frequency based on raw frequencies also contributed a significant, small additional proportion of the variance, namely .64%. The two variables combined produced a multiple R=.2889 and a multiple R²=.0835. None of the other three variables entered into the equation. For the percentages correct, stimulus word frequency based on SFI's again accounted for the largest proportion of variance: 9.00%. In addition, neighborhood density and raw stimulus word frequency contributed significantly to the total proportion of variance accounted for. Density contributed an additional 1.05% of the variance and raw word frequency an additional .92%. All three variables combined to produce a multiple R=.3311 and a multiple R²=.1096.

Discussion of Correlation and Regression Analyses. For both reaction times and percentages correct, significant correlations of word frequency were obtained. Thus, despite the fact that all words were previously judged to be highly familiar to subjects, frequency based on an objective word count proved to be significantly correlated with reaction times and accuracy levels. This result demonstrates once again that frequency and subjective familiarity are separate, albeit highly correlated variables, underlying subjects' performance in recognizing words.

Neither the correlation nor the regression analyses revealed significant effects of the neighborhood variables on reaction time. Only the SFI of the stimulus word proved to be a significant predictor of reaction times. However, neighborhood density, or the number of words in a similarity neighborhood, significantly predicted accuracy, although the unique proportion of variance accounted for by the density variable was small. In addition, the correlation between density and percentage correct appears somewhat counterintuitive: Words with many neighbors were, in the long run, classified correctly more often than words with few neighbors.

The finding that density correlated positively with word classification accuracy is in contradiction to the general findings from the perceptual identification study reported in Chapters 2 and 3. In that experiment, it was found that words with many confusable neighbors were identified less well than words with fewer confusable neighbors. Why, then, should a reversal of this finding be found for auditory lexical decision, such that words in high density neighborhoods were classified better than words in low density neighborhoods?

The perceptual identification and the auditory lexical decision tasks differ on one important respect. In the perceptual identification task, subjects are required to identify only real words. One may therefore assume that subjects are attempting primarily to discriminate among the words activated in memory by the stimulus word. In the auditory lexical decision

Table 4.4

Neighborhood Analysis: Auditory Lexical Decision. Results of stepwise regression analysis accuracy and reaction time data. Shown are the variables entered into the equation at each step. Also shown are the multiple R's, multiple R²'s, and changes in R² at each step.

ACCURACY DATA			
VARIABLE	R	R ²	ΔR^2
1. SFI	0.2999	0.0900	0.0900
2. DEN	0.3169	0.1004	0.0105
3. RWF	0.3311	0.1096	0.0092
REACTION TIME DATA			
VARIABLE	R	R ²	ΔR^2
1. SFI	0.2776	0.0771	0.0771
3. RWF	0.2889	0.0835	0.0064

task, however, subjects are required to discriminate among both word and nonword patterns in memory. This factor may underlie the differing results for the perceptual identification and lexical decision tasks for the following reasons: First, words occurring in high density neighborhoods have fewer confusable nonwords by virtue of the fact that most of their neighbors are actual words. Lexical decisions may therefore tend to be more accurate for words in high density neighborhoods because of the lower probability of choosing a similar nonword. Likewise, words in low density neighborhoods, which contain more possible nonword patterns, may tend to be classified less accurately due to a higher probability of choosing a similar nonword. Thus, it is possible that lexical decision accuracy is positively correlated with neighborhood density because both word and nonword patterns must be discriminated among in the auditory lexical decision task.

This analysis of the auditory lexical decision task also suggests interactions among the neighborhood variables and stimulus word frequency. In particular, high and low frequency words may be differentially sensitive to neighborhood density and neighborhood frequency when both words and nonwords are presented. Assuming that sound patterns corresponding to low frequency words have lower activation levels in memory, arising from lower levels of frequency bias, discrimination between low frequency words and nonwords may be more difficult than discrimination between high frequency words and nonwords. The neighborhood variables examined may therefore have had quite different effects for high and low frequency words, given that these variables do not explicitly take into account the differential competition of nonword patterns with high and low frequency words. If interactions of this type exist, then the correlations of the neighborhood variables with reaction times and percentages correct may have been deflated. For example, if neighborhood density has differential effects on high and low frequency words because of the effects of nonwords, the correlation of density across the entire set of words may have been considerably reduced. In short, any interactions of the neighborhood variables across words of different frequencies may have served to reduce or eliminate the correlations of the neighborhood variables with reaction times and percentages correct.

Analysis of Partitioned Word Response Data

In order to examine the possibility of interactions among the neighborhood variables and word frequency, the set of words were partitioned into orthogonal cells and submitted to analyses of variance. The partitioning of the stimuli was achieved by performing median splits on the values of each of the three independent variables: stimulus word frequency, neighborhood density, and mean neighborhood frequency. That is, the median frequency of the stimulus words was first determined and words falling above the median were coded as high frequency words; those equal to or less than the median were coded as low frequency words. The same procedure was applied to assign words to high density neighborhoods or low density neighborhoods and to high frequency neighborhoods or low frequency neighborhoods. Based on this coding scheme, each word was assigned to one of eight cells produced by the orthogonal combination of two levels of stimulus word frequency (high and low), two levels of neighborhood density (high and low), and two levels of neighborhood frequency (high and low). The resulting high frequency words had a mean frequency of 254.12; low frequency words had a mean frequency of 5.22. High density neighborhoods contained an average of 21.92 neighbors; low density neighborhoods contained an average of 11.07 neighbors. The mean frequency of high frequency neighborhoods was 370.32; the mean frequency of low frequency neighborhoods was 46.29.

Reaction times and percentages correct for each subject were averaged across words within a cell and submitted to analyses of variance. Because the entire set of words was split into thirds and presented to separate groups of equal numbers of subjects, a grouping factor was included in the analysis, producing a 2 (stimulus word frequency) X 2 (neighborhood density) X 2 (neighborhood frequency) X 3 (groups) repeated measures analysis of variance.

Accuracy Data. Analysis of the percentages correct revealed significant effects of stimulus word frequency, $F(1,27)=135.94$, $p<0.05$, neighborhood density, $F(1,27)=39.39$, $p<0.05$, and mean neighborhood frequency, $F(1,27)=4.93$, $p<0.05$. No effect of groups was obtained, $F(2,27)=3.13$, $p>0.05$. A significant interaction of stimulus word frequency and neighborhood density was also obtained, $F(1,27)=17.07$, $p<0.05$. No other interactions were significant. Mean percentages correct and standard deviations for each cell, collapsed across groups, are shown in Table 4.5. These same results are plotted in Figure 4.1. High frequency words are plotted in the left-hand panel and low frequency words in the right-hand panel. Words occurring in high density neighborhoods are represented by dotted lines with circles; words occurring in low density neighborhoods are represented by solid lines with X's. Neighborhood frequency is plotted on the x axis.

Insert Table 4.5 and Figure 4.1 about here

On the average, high frequency words were responded to 7.39% better than low frequency words. Words in high density neighborhoods were responded to 3.38% better than words in low density neighborhoods. And, words occurring in low frequency neighborhoods were responded to 1.39% better than words occurring in high frequency neighborhoods. The effect of density was the same as that revealed by the correlation analysis: Words in high density neighborhoods were classified correctly more often than words in low density neighborhoods. However, the significant interaction of stimulus word frequency and density indicates differential effects of neighborhood density as a function of word frequency. The interaction of frequency and density is plotted in Figure 4.2. High density words are represented by dotted lines with circles; low density words are represented by solid lines with X's. Stimulus word frequency is plotted on the x axis.

Insert Figure 4.2 about here

Separate analyses based on the interaction of word frequency and density revealed significant effects of word frequency at both levels of density. High frequency words were responded to 4.48% better than low frequency words in high density neighborhoods, $F(1,27)=40.48$, $p<0.05$. High frequency words were also responded to 10.31% better than low frequency words in low density neighborhoods, $F(1,27)=81.51$, $p<0.05$. Thus, significant effects of frequency were observed at each level of neighborhood density.

Separate analyses revealed no significant effect of density for high frequency words for the accuracy data, $F(1,27)<1.0$. However, a significant effect of density was observed for low frequency words, $F(1,27)=6.29$, $p<0.05$. For low frequency words, words in high density neighborhoods were responded to 6.29% better than words in low density neighborhoods. Thus, the effect of density on correct classification of words was observed only for low frequency words.

Table 4.5

Neighborhood Analysis: Auditory Lexical Decision. Means and standard deviations for the accuracy data.

HIGH FREQUENCY WORDS

NEIGHBORHOOD DENSITY

		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	92.59 (4.29)	92.58 (5.20)
	LOW	94.73 (4.83)	93.82 (5.80)

LOW FREQUENCY WORDS

NEIGHBORHOOD DENSITY

		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	88.80 (7.75)	82.19 (9.21)
	LOW	89.57 (7.24)	83.59 (8.05)




Figure 4.1. Percentages correct for the word data for auditory lexical decision. Percentages correct for the high frequency words are shown in the left-hand panel; percentages correct for the low frequency words are shown in the right-hand panel. Words in high density neighborhoods are represented by the dotted lines with circles; words in low density neighborhoods are represented by the solid lines with X's. Neighborhood frequency is plotted on the x axis.

Frequency X Density Interaction




Figure 4.2. Interaction of neighborhood density and stimulus word frequency for the word data for auditory lexical decision. Words in high density neighborhoods are represented by the dotted lines with circles; words in low density neighborhoods are represented by the solid lines with X's. Stimulus word frequency is plotted on the x axis. The results are plotted as a function of percent correct.

Reaction Time Data. Analysis of the reaction time data revealed significant main effects of word frequency, $F(1,27)=70.39$, $p<0.05$, neighborhood density, $F(1,27)=14.32$, $p<0.05$, and neighborhood frequency, $F(1,27)=14.15$, $p<0.05$. No effect of groups was obtained, $F(2,27)=2.32$, $p>0.05$. A significant interaction of word frequency and neighborhood density was also obtained, $F(1,27)=14.15$, $p<0.05$. Means and standard deviations for each cell are shown in Table 4.6. The same data are plotted in Figure 4.3.

Insert Table 4.6 and Figure 4.3 about here

Overall, high frequency words were responded to 55 msec faster than low frequency words. Words occurring in high density neighborhoods were responded to 13.5 msec slower than words in low density neighborhoods. And, words in high frequency neighborhoods were responded to 17.5 msec slower than words in low frequency neighborhoods. Although there was an overall 13.5 msec advantage for words in low density neighborhoods over words in high density neighborhoods, the significant interaction of word frequency and density indicates differential effects of one or both of these variables. This interaction is plotted in Figure 4.4.

Insert Figure 4.4 about here

Separate analyses based on this interaction revealed significant effects of word frequency at each level of density. In high density neighborhoods, high frequency words were responded to 47.5 msec faster than low frequency words, $F(1,27)=60.35$, $p<0.05$. In low density neighborhoods, high frequency words were responded to 62 msec faster than low frequency words, $F(1,27)=54.45$, $p<0.05$. Significant effects of neighborhood density were obtained only for high frequency words for the reaction time data. For high frequency words, words in low density neighborhoods were responded to 21 msec faster than words in high density neighborhoods, $F(1,27)=18.33$, $p<0.05$. No effect of density was observed for low frequency words, $F(1,27)=1.75$, $p>0.05$.

Summary of Word Data. To summarize the data thus far presented, significant effects of word frequency were obtained for both the accuracy and reaction time data. Overall, high frequency words were responded to faster and with a higher level of accuracy than low frequency words. Significant effects of neighborhood frequency were also observed for both the accuracy and reaction time data. Words occurring in high frequency neighborhoods were responded to less accurately and more slowly than words in low frequency neighborhoods. This result demonstrates that high frequency neighbors tend to slow response time and contribute to higher error rates in word classification.

Effects of neighborhood density were also observed for the accuracy and reaction time data, although the overall effects were in opposite directions. For the accuracy data, words in high density neighborhoods were classified more accurately than words in low density neighborhoods. However, the reaction time data revealed the opposite pattern of results: Words in high density neighborhoods were responded to more slowly than words in low density neighborhoods. This pattern of results suggests a speed-accuracy trade-off. However, analyses of the interactions of word frequency and density indicated that density had differential effects on accuracy and reaction time for high and low frequency words. Specifically, density affected accuracy of

Table 4.6

Neighborhood Analysis: Auditory Lexical Decision. Means and standard deviations for the reaction time data.

		HIGH FREQUENCY WORDS	
		NEIGHBORHOOD DENSITY	
		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	409 (74)	382 (75)
	LOW	392 (113)	377 (104)
		LOW FREQUENCY WORDS	
		NEIGHBORHOOD DENSITY	
		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	451 (105)	463 (126)
	LOW	445 (111)	421 (105)




Figure 4.3. Reaction times in msec for the word data for auditory lexical decision. Reaction times for the high frequency words are shown in the left-hand panel; reaction times for the low frequency words are shown in the right-hand panel. Words in high density neighborhoods are represented by the dotted lines with circles; words in low density neighborhoods are represented by the solid lines with X's. Neighborhood frequency is plotted on the x axis.

Frequency X Density Interaction




Figure 4.4. Interaction of neighborhood density and stimulus word frequency for the word data for auditory lexical decision. Words in high density neighborhoods are represented by the dotted lines with circles; words in low density neighborhoods are represented by the solid lines with X's. Stimulus word frequency is plotted on the x axis. The results are plotted as a function of reaction time in msec.

classification only for low frequency words, such that low frequency words with many neighbors were classified correctly more often than low frequency words with fewer neighbors. In terms of reaction times, density affected only high frequency words. In particular, high frequency words with many neighbors were responded to less quickly than high frequency words with fewer neighbors. Before discussing the implications of these findings, the results for the nonword data will be discussed.

Analysis of Nonword Responses

Prior to analysis of the nonword response data, stimulus durations were subtracted from the correct nonword reaction times and outliers were eliminated and replaced according to the procedure described for the word response data. Because the nonwords were originally assigned to four separate cells representing the orthogonal combination of neighborhood density and neighborhood frequency, only analyses of variance were performed. In addition, because each subject heard the same set of nonwords, no grouping factor was included in the analyses of variance.

Accuracy Data. A two-way repeated measures analysis of variance (neighborhood density X neighborhood frequency) on the accuracy data for the nonwords revealed significant main effects of neighborhood density, $F(1,29)=26.54$, $p<0.05$, and neighborhood frequency, $F(1,29)=17.68$, $p<0.05$. In addition, the interaction of neighborhood density and neighborhood frequency was significant, $F(1,29)=24.75$, $p<0.05$. Means and standard deviations for the accuracy data for each cell are shown in Table 4.7. The same data are plotted in Figure 4.5. Nonwords occurring in high density neighborhoods are represented by the dotted line with circles; nonwords occurring in low density neighborhoods are represented by the solid line with X's. Mean neighborhood frequency is plotted on the x axis.

Insert Table 4.7 and Figure 4.5 about here

Nonwords occurring in high density neighborhoods were responded to 3.34% worse than nonwords occurring in low density neighborhoods. Nonwords occurring in high frequency neighborhoods were responded to 3.02% worse than nonwords occurring in low frequency neighborhoods. Separate analyses based on the significant neighborhood density-by-neighborhood frequency interaction revealed a significant effect of density only for nonwords in high frequency neighborhoods, $F(1,29)=51.03$, $p<0.05$. In high frequency neighborhoods, nonwords having many word neighbors were responded to 5.53% worse than nonwords having few word neighbors. In addition, a significant effect of neighborhood frequency was observed only for nonwords occurring in high density neighborhoods, $F(1,29)=30.49$, $p<0.05$. In high density neighborhoods, nonwords with high frequency neighbors were responded to 5.22% worse than nonwords with low frequency neighbors. Both of these effects appear to be due to the lower mean percent correct for nonwords occurring in high density, high frequency neighborhoods.

Reaction Time Data. For the reaction time data for correct nonword classifications, significant main effects were obtained for neighborhood density, $F(1,29)=60.81$, $p<0.05$, and neighborhood frequency, $F(1,29)=5.39$, $p<0.05$. The interaction of neighborhood density and neighborhood frequency was not significant, $F(1,29)<1.0$.

Table 4.7

Neighborhood Analysis: Auditory Lexical Decision. Means and standard deviations for the accuracy data for the nonwords.

		NEIGHBORHOOD DENSITY	
		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	84.08 (6.85)	89.61 (4.96)
	LOW	89.30 (6.74)	90.44 (4.56)

Auditory Lexical Decision Nonwords




Figure 4.5. Percentages correct for the nonword data for auditory lexical decision. Nonwords in high density neighborhoods are represented by the dotted lines with circles; nonwords in low density neighborhoods are represented by the solid lines with X's. Neighborhood frequency is plotted on the x axis.

Means and standard deviations for the reaction times for each cell are shown in Table 4.8. The same data are plotted in Figure 4.6.

Insert Table 4.8 and Figure 4.6 about here

Nonwords occurring high density neighborhoods were classified 39.5 msec slower than nonwords occurring in low density neighborhoods. In addition, nonwords occurring in high frequency neighborhoods were classified 11.5 msec slower than nonwords occurring in low frequency neighborhoods.

Summary of Nonword Data. Significant effects of neighborhood density and neighborhood frequency were observed for both the accuracy and reaction time data for the nonword responses. On the average, nonwords with many neighbors were responded to more slowly and with lower levels of accuracy than nonwords with few neighbors. In addition, nonwords with high frequency neighbors were responded to more slowly and with lower levels of accuracy than nonwords with low frequency neighbors. For the accuracy data, however, it was found that neighborhood density only affected nonwords in high frequency neighborhoods. And, a significant effect of neighborhood frequency was obtained only for nonwords in high density neighborhoods. The implications of these findings will now be considered.

DISCUSSION

The results from both the word and nonword data demonstrate the role of neighborhood structure in word-nonword classification accuracy and reaction time. In addition, a significant effect of stimulus word frequency was observed for the words. The overall pattern of results for both the word and nonword data, while demonstrating effects of neighborhood structure, are not easily accounted for by an independent combination of the variables of word frequency, neighborhood density, and neighborhood frequency. In particular, it is unclear why neighborhood density interacted with word frequency differently for accuracy and reaction time. Thus, in order to account for this pattern of results, the lexical decision paradigm will be considered in light of the neighborhood activation model developed in Chapter 3.

Lexical Decision in the Neighborhood Activation Model

Recall that the neighborhood activation model states that stimulus input first activates a set of acoustic-phonetic patterns in memory. The activation levels of the patterns are assumed to be direct functions of the similarity of each pattern to the stimulus input. Once the acoustic-phonetic patterns have been activated, decision units for words begin monitoring the acoustic-phonetic patterns to which they correspond. Each decision unit also monitors higher-level lexical information corresponding to the word for which the unit is responsible. In particular, once a decision unit is activated, frequency information serves to bias the decision units by multiplying the activation level of the acoustic-phonetic pattern by its frequency of occurrence. Each decision unit continuously computes values for its word based on the frequency-adjusted activation level of its pattern and the overall level of activation in the system of decision units. As stimulus processing proceeds in time, the acoustic-phonetic information within the stimulus is further resolved, or "refined" (see Pisoni, Nusbaum, Luce, & Slowiaczek, 1985). The resolution of the acoustic-phonetic information serves

Table 4.8

Neighborhood Analysis: Auditory Lexical Decision. Means and standard deviations for the reaction time data for the nonwords.

		NEIGHBORHOOD DENSITY	
		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	455 (118)	419 (116)
	LOW	447 (115)	404 (99)

Auditory Lexical Decision Nonwords




Figure 4.6. Reaction times in msec for the nonword data for auditory lexical decision. Nonwords in high density neighborhoods are represented by the dotted lines with circles; nonwords in low density neighborhoods are represented by the solid lines with X's. Neighborhood frequency is plotted on the x axis.

to reduce the activation levels of similar patterns in memory, until the output of one decision unit reaches criterion. (This is true only for stimuli that are neither degraded nor impoverished; see Chapter 3.) The decision units are assumed to compute values via a decision rule of the sort described in Chapter 3. Once a value exceeds criterion, all information monitored by that decision unit is made available to working memory.

Because decision units only correspond to words actually occurring in memory, lexical decision in the context of the neighborhood activation model can only be achieved by accepting or rejecting words. In particular, lexical decisions are assumed to be based on one of two criteria being exceeded (see Coltheart, Develaar, Johansson, & Besner, 1976). A word response is executed if a decision unit determines that the activation level of the pattern it is monitoring exceeds the criterion, the normal procedure for recognizing a word and depositing its lexical information in working memory. However, the procedure for executing a nonword response is contingent on surpassing a lower-level criterion. A nonword response is executed if the total activation level monitored by the decision units falls below a lower-level criterion, indicating that no word is consistent with the stimulus input. According to the neighborhood activation model, therefore, word-nonword classification is based on the activity within the decision system surpassing or falling below one of two criteria, which can be referred to as the "word" and "nonword" criterion levels.

Under circumstances in which subjects are required to classify a clearly presented word or nonword under no time constraint, it is assumed that classification accuracy will be near perfect. Exhaustive analysis of the stimulus input would result in few, if any, errors in classification. However, once time constraints are imposed by instructions to respond as quickly as possible, accuracy levels will vary as a function of the amount of stimulus processing carried out prior to the response. It is assumed that subjects will attempt to classify an item prior to a self-imposed reaction time deadline. The assumption of a reaction time deadline is motivated by the fact that subjects are attempting to respond as quickly as possible and will allow only a given amount of time to pass before executing a response, regardless of the processing achieved at that moment in time. This assumption has its precedent in earlier models of visual lexical decision (Coltheart et al., 1976). The notion of a self-imposed deadline leads to a further assumption regarding subjects' behavior in the lexical decision task, which I will refer to as the accuracy assumption. The accuracy assumption states that differences in classification accuracy will only be observed when the response time deadline has been exceeded. If stimulus processing is completed before the deadline, few errors in classification should be observed. However, if stimulus processing is incomplete at the expiration of the deadline, subjects will be forced to execute a response based on only partial information, thus producing errors in classification. Simply put, according to the accuracy assumption, only those stimuli requiring processing times exceeding the response time deadline should produce errors in classification. In addition, if certain stimuli consistently require processing times that exceed the deadline, reaction times to these stimuli will tend to be equal to the deadline itself. If the deadline is reached prior to exceeding either the word or nonword criteria, classification is based on the overall level of activity in the decision system. If this activity is high, a word response will be executed; otherwise, a nonword response will be made.

To summarize, lexical decisions in the neighborhood activation model are assumed to be made via the decision units for words. Word responses are executed once a decision unit surpasses the upper-level, word criterion.

Nonword responses are executed once the total level of activation for words falls below the lower-level, nonword criterion. Furthermore, if neither of these criteria are exceeded prior to a self-imposed reaction time deadline, a decision is forced based on the overall activity in the decision system, which may lead to erroneous classification responses. Having devised a reasonable framework for understanding how subjects make lexical decisions, I will now turn to a discussion of the specifics of the data obtained in the present experiment, beginning with the word response data.

Word Responses

Three main effects were observed for the word response data. High frequency words were classified more quickly and more accurately than low frequency words. Words in low frequency neighborhoods were classified more quickly and more accurately than words in high frequency neighborhoods. Finally, words occurring in high density neighborhoods were classified more slowly but more accurately than words in low density neighborhoods. Although this last result suggests a speed-accuracy trade-off, the significant interactions of neighborhood density and word frequency for both reaction times and accuracy revealed differential effects of density on accuracy and reaction time as a function of word frequency. High frequency words in high and low density neighborhoods were classified equally as accurately. However, classification time was slower for high frequency words in high density neighborhoods. A different pattern of results was observed for the low frequency words. No reaction time differences were observed as a function of neighborhood density. However, low frequency words in high density neighborhoods were classified better than low frequency words in low density neighborhoods.

Before considering the interesting interactions of density and word frequency for the accuracy and reaction time data, consider first how the neighborhood activation model explains the effects of word frequency and neighborhood frequency on word classification responses. Recall that word frequency increases the activation level of the acoustic-phonetic pattern represented in the decision units. Thus, high frequency words will tend to have higher levels of activation in the decision system than low frequency words. These higher levels of activation thus lead to faster word responses for high frequency words given that the word criterion is surpassed more quickly as stimulus processing proceeds. Thus, high frequency words show faster reaction times than low frequency words. And, given the accuracy assumption stated above, slower processing times associated with low frequency words will result in higher error rates.

Neighborhood frequency affects classification times for words by slowing the time for a decision unit to reach criterion. Recall that the decision units monitor overall activity in the decision system. Because high frequency neighborhoods result in overall higher activity levels, the time for a given decision unit to surpass the criterion will be extended. Thus, high frequency neighbors serve as stronger competitors by virtue of the fact that they raise the overall level of activity within the decision system.

Effects of neighborhood density can be explained by the same basic principle used to explain the effects of neighborhood frequency. In particular, heightened overall activity in the decision system extends the time needed for a given decision unit to surpass the criterion. Words with many neighbors produce high levels of activity in the decision system, thus slowing response time. Such an effect was observed for high frequency words.

No effect of density was observed for the accuracy data for high frequency words presumably because decisions for high frequency words, in the long run, were made prior to the response-time deadline. However, no effect of density on reaction times was observed for low frequency words. The failure to observe reaction time differences for low frequency words may have arisen from the fact that, on the average, decision units failed to surpass the criterion level for the low frequency words by the time the response-time deadline had expired. Thus, reaction times for the low frequency words simply reflect a forced decision at the deadline. According to the accuracy assumption, then, density effects, if they exist for low frequency words, should be observed only for the accuracy data. Indeed, it was found that low frequency words in high density neighborhoods were classified more accurately than low frequency words in low density neighborhoods. Given that decisions that are forced at the response time deadline are based on the overall level of activity in the decision system, low frequency words with many word neighbors would have higher levels of overall activity and would thus be classified more accurately as words. The somewhat counterintuitive finding that low frequency words with many neighbors were classified more accurately than low frequency words with few neighbors can therefore be accounted for by the neighborhood activation model.

Within the context of the neighborhood activation model, then, the present pattern of results for the word responses can at least be described. Clearly, independent tests of the assumptions regarding both the model and subjects' behavior in the lexical decision task are required. However, the neighborhood activation model provides at least a prima facie, coherent account of the results obtained for the word responses. Nevertheless, the results from the word data demonstrate once again the effects of neighborhood structure on auditory word recognition, supporting the general claim that words are recognized in the context of similar words in memory. Further support for this claim was obtained from the nonword data as well, which I will now consider.

Nonword Responses

The results for the reaction time data for the nonwords are easily accounted for by the neighborhood activation model. Recall that it was assumed that a nonword response is executed whenever the overall activity in the decision system falls below a lower-level criterion. Thus, any factor that slows the time for the activity level to drop in the decision system should slow the time to correctly classify a nonword pattern. Indeed, the results for the nonwords showed significant main effects of neighborhood density and neighborhood frequency. Nonwords with many neighbors were responded to more slowly than nonwords with few neighbors. Because the activity level in the decision system takes longer to decay when there are many similar words activated by the nonword stimulus, nonword classification times were longer for nonwords in high density neighborhoods than for nonwords in low density neighborhoods (see also Forster, 1976; Rubenstein, Richter, & Kay, 1975). The same reasoning applies to nonwords in high frequency neighborhoods. Given the overall higher activity level associated with high frequency neighborhoods, nonwords with high frequency neighbors were classified more slowly than words with low frequency neighbors.

For the accuracy data for the nonwords, an interaction of density and neighborhood frequency was observed. This interaction was due to one cell, namely that for nonwords occurring in high density, high frequency neighborhoods. The accuracy levels for all other nonwords were approximately

equal. The mean reaction time for nonwords occurring in high density, high frequency neighborhoods was also the longest, and was approximately equal to the maximum reaction observed in the experiment as a whole. Thus, under the accuracy assumption, it is possible that the reduced accuracy for these nonwords was due to response-time deadline expiration. In the case in which a nonword stimulus activates a set of high frequency similar words, the overall activity level in the decision system may have at times failed to drop below the criterion for a nonword response. In this case, the overall activity level in the system would be examined in order to determine a response. Given that the activity level would tend to be high for nonwords with many high frequency word neighbors, errors in nonword classification would be expected to arise in this case.

Summary and Conclusions

The data from both the word and nonword responses revealed significant effects of neighborhood structure on classification time and accuracy. Although the correlation and regression analyses for the word responses revealed only a small effect of density on response accuracy, the subsequent analyses using median splits of the data revealed significant interactions that would have attenuated the correlations of the variables of interest. In particular, these interactions demonstrated that the effects of neighborhood density varied as a function of word frequency for both the accuracy and reaction time data. In addition, although the correlation and regression analyses failed to show significant effects of neighborhood frequency, the analysis of variance did reveal small but significant effects when the data was partitioned into cells. These latter results thus confirm the effects of neighborhood structure on word classification times and accuracy. Furthermore, the results from the nonword data support the conclusions based on the word data. The effects of neighborhood structure on word recognition are therefore not restricted to degraded stimuli. In addition, the reaction times for the word and nonword responses support the conclusions drawn from the perceptual identification study. Although the effect of density on classification accuracy for low frequency words ran counter to the results observed in the identification study, it was argued that the neighborhood activation model can, in fact, account for this result via the same mechanisms invoked to account for the identification data.

The neighborhood activation model thus provides a coherent framework for interpreting the effects of neighborhood structure on word and nonword classification times. Although much of the interpretation of the effects of neighborhood structure on both word and nonword responses relies on a small number of crucial assumptions regarding the behavior of subjects in the lexical decision task, the neighborhood activation model permits a unified account of the data observed in the lexical decision task. Many of the particulars of the model, and the characterization of subjects' behavior in the lexical decision task, necessarily require further independent testing based on the predictions of the model. The results nonetheless show explicable and consistent effects of neighborhood structure. Furthermore, the data obtained from the lexical decision task reveal that neighborhood structure is an important determinant of the ease and speed with which words are recognized in this experimental paradigm.

CHAPTER FIVE

EVIDENCE FROM AUDITORY WORD NAMING

The present chapter attempts to gather further support for the neighborhood activation model developed in the previous chapters by examining the effects of neighborhood structure in the context of a different methodological paradigm. Specifically, the paradigm used in the present chapter is auditory word naming (see Andrews, 1982; Frederiksen & Kroll, 1976; Forster, 1981; Forster & Chambers, 1973). In the auditory word naming task, a subject is presented with an auditory word stimulus and is required to repeat or pronounce the word as quickly as possible. The dependent variable is the time required to initiate the naming response.

The use of the auditory word naming task is motivated by a number of factors. First, as shown previously, auditory lexical decision proves somewhat problematic in examining neighborhood density and frequency effects by virtue of the fact that the task requires discrimination among both word and nonword patterns. Although the results from the previous lexical decision study provided evidence for the effects of neighborhood structure, these effects may have been attenuated or altered given that subjects were required to consider both words and nonwords in making their responses. In the absence of any systematic control for the nonword patterns that may be activated in memory by subjects, manipulation of neighborhood structure on the basis of words alone makes precise control of neighborhood structure difficult in a task requiring discrimination among words and nonwords both. Thus, the naming task provides a means of collecting reaction times to word stimuli without presenting nonwords.

A second motivation for using the auditory word naming task comes from recent findings in the visual literature on the role of word frequency in the naming task. Balota and Chumbley (1984) have presented evidence that word frequency effects are severely reduced in the visual word naming task, as compared to the lexical decision task. In addition, these researchers have argued that the small word frequency effects obtained in the naming task are due to factors related to the pronunciation of the visually presented item and not to the frequency of the word itself (Balota & Chumbley, 1985). Finally, Paap, McDonald, Schvaneveldt, and Noel (1986) have argued that the visual word naming task circumvents lexical access and thus circumvents access to frequency information. Paap et al. argue that naming a visually presented word simply requires grapheme-to-phoneme conversion with no access to the mental lexicon.

These findings suggest an interesting test of the neighborhood activation model. Recall that the model proposes that acoustic-phonetic patterns similar to the stimulus input are activated in memory. These patterns then activate decision units corresponding to words that monitor the acoustic-phonetic patterns as well as higher-level lexical information, which includes word frequency. The decision units continuously compute probability values based on the activation level of the acoustic-phonetic patterns they monitor and the overall level of activity within the system. Frequency information is assumed to bias the decision units by adjusting the activation levels of the acoustic-phonetic patterns represented in the units.

The system of decision units is therefore driven by the activation of acoustic-phonetic patterns, giving acoustic-phonetic pattern similarity priority in the decision system. Word frequency, on the other hand, serves as

a biasing factor that may or may not come into play in the decision-making process, depending on the requirements of the task situation (see Chapter 3). Thus, in the model, similarity and frequency effects arise from two distinct sources and operate in fundamentally different ways in influencing decisions.

If the visual and auditory word naming tasks are sufficiently similar, auditory word naming should not be sensitive to the biasing properties of word frequency information. The model predicts, however, that pattern similarity is fundamental to the system of decision units and cannot be bypassed, at least in situations involving an open response set. Therefore, if the predictions of the neighborhood activation model are correct, robust effects of neighborhood density should be observed on naming times regardless of whether frequency information acts to bias the decision units. In particular, high density neighborhoods should produce longer naming times than low density neighborhoods. However, if the auditory naming task does not invoke the biasing properties of word frequency information, as predicted by the visual word naming studies, no effects of stimulus word frequency or neighborhood frequency should be observed. The auditory naming task may thus aid in dissociating the effects of similarity and bias (i.e., frequency), providing further support for the predictions of the neighborhood activation model.

EXPERIMENT

Method

Stimuli

Four-hundred words were selected from the 882 words used in the perceptual identification and auditory lexical decision experiments. (Recall that 918 words in total were used in the previous experiments. However, 36 of these words were excluded from the analyses because of failure to reach criterion in a screening experiment.) Only consonant-vowel-consonant words were selected. The words were chosen in order to construct eight cells with 50 words per cell. The eight cells were constructed by orthogonally combining two levels (high and low) of each of the following independent variables: Stimulus word frequency, neighborhood density, and neighborhood frequency. The neighborhood density and neighborhood frequency variables were computed for each word in the same manner as described in Chapters 2 and 4. The phonetic transcriptions of each stimulus word were compared with each monosyllabic word in Webster's lexicon having a familiarity rating of 5.5 or above. A neighbor of the stimulus word was defined as any word that could be converted to the stimulus word via a one phoneme addition, substitution, or deletion in any position. Neighborhood density again refers to the number of neighbors of a given stimulus word and neighborhood frequency refers to the mean frequencies of these neighbors.

Selection of the 50 words for each cell was achieved via an algorithm that first rank-ordered each of the 882 words on each of the three independent variables. A method of minimizing and maximizing squared deviations of successively ranked words was then employed to ensure that cells that were matched on a given variable (e.g., both high density) were maximally alike and that cells intended to differ on a given variable (e.g., one high and one low density) were maximally different. In addition, words were chosen such that the mean stimulus durations for each cell were not significantly different. High frequency words had a mean of 145.95; low frequency words had a mean of 4.33. Words occurring in high density neighborhoods had an average of 22.12 neighbors; words occurring in low density neighborhoods had an average of

11.44 neighbors. The mean frequency of high density neighborhoods was 245.17; the mean frequency of low frequency words was 60.50. Details of the preparation of the auditory stimuli are given in Chapter 3.

Subjects

Eighteen subjects participated as paid volunteers. Subjects received \$3.50 for their participation. All subjects were native English speakers and reported no history of speech or hearing disorders.

Design and Procedure

The 400 stimulus words were combined in a single stimulus set file. Each subject heard each of the 400 stimulus words in a different random order. In addition, subjects were given 30 practice trials prior to the experiment proper on a separate set of words. Stimulus presentation and data collection were controlled by a PDP 11/34 minicomputer. The stimuli were presented on matched and calibrated TDH-39 headphones at a comfortable listening level of 75 dB SPL.

Each subject was run individually in a sound treated room. An Electro-Voice D054 microphone was situated immediately in front of the subject. The microphone was interfaced to a voice key that was interfaced to the PDP 11/34. The voice key registered a response at the onset of the subject's naming response. The subject was positioned such that his/her lips were approximately 12 inches from the microphone. The subject was instructed to maintain the 12 inch distance at all times. In addition, each subject was instructed to avoid unnecessary noise.

The subjects were instructed that they would hear words over their headphones which they were to repeat back or name as quickly but as accurately as possible. The subjects were told that the microphone would register when they began speaking and that the time it took them to name the stimulus would be recorded by the computer. The experimenter was seated in a booth next to the subject and monitored an ADM CRT terminal. On each trial, the stimulus word appeared on the terminal and the experimenter listened to the subject's naming response. After the subject's response, the experimenter would indicate on the computer terminal whether the subject had responded correctly or incorrectly. If an incorrect response was made, the experimenter typed the mistake on the terminal.

A given trial proceeded as follows: Subjects first saw the prompt "GET READY FOR NEXT TRIAL" on a CRT screen situated above the microphone. One second following the prompt, a word was presented over the headphone. The experimenter then scored the subject's response and initiated a new trial. Reaction times were measured from the onset of the auditory stimulus to the onset of the subjects response. A given experimental session lasted approximately one hour.

Results

Reaction times were first entered into a stimulus word-by-subject array and means and standard deviations were computed for each word and each subject. Any reaction time falling 2.5 standard deviations above and below both the subject and stimulus means was eliminated and replaced according to the procedure suggested by Winer (1971). Mean reaction times were then computed for each subject for each cell. Preliminary inspection of the data revealed correlations of reaction time with the identity of the initial segment of the stimulus word. In particular, fricatives were associated with longer reaction times, presumably due to the differential sensitivity of the microphone and voice key to the identity of the initial segment of the naming response. Because the identity of initial segments was not evenly distributed across cells, it was deemed necessary to factor out the effects of the initial segment. In order to do this, the number of initial segments falling in each of six manner classes for each cell was tallied. The five manner classes were: stops, strong fricatives, weak fricatives, nasals, liquids and glides, and affricates. These numbers were then entered as covariates for the reaction times in a repeated measures analysis of variance. In addition to the reaction time data, percentages of correct responses were computed for each subject for each cell and submitted to analysis of variance.

Accuracy Data. A 2 (stimulus word frequency) X 2 (neighborhood density) X 2 (neighborhood frequency) repeated measures analysis of variance was computed on the percentages correct. The main effects of stimulus word frequency, $F(1,17)=24.73$, $p<0.05$, and neighborhood density, $F(1,17)=4.89$, $p<0.05$, were obtained. No effect of neighborhood frequency was observed, $F<1.0$. In addition, none of the interactions were significant. Means and standard deviations are shown in Table 5.1 for each cell.

Insert Table 5.1 about here

Inspection of Table 5.1 reveals that the significant effects of stimulus word frequency and neighborhood density on accuracy were extremely small. High frequency words were responded to .89% better than low frequency words. Words occurring in high density neighborhoods were responded to .28% worse than words occurring in low density neighborhoods. In addition, it should be noted that the accuracy levels overall were very high, the lowest cell percentage being 96.78%. Thus, naming an auditory stimulus appears to be quite easy for subjects to perform

Reaction Time Data. A 2 X 2 X 2 repeated measure analysis of covariance was performed on the reaction times. Recall that the covariates were the number of initial segments in each cell falling into one of the six manner classes. Only the main effect of neighborhood density was observed for the reaction times, $F(1,11)=5.10$, $p<0.05$. Neither stimulus word frequency, $F(1,11)=1.71$, $p<0.05$, nor neighborhood frequency, $F(1,11)<1.0$, reached significance. In addition, no significant interactions were observed. Means and standard deviations are shown in Table 5.2.

Table 5.1

Neighborhood Analysis: Naming. Means and standard deviations for the accuracy data.

HIGH FREQUENCY WORDS			
NEIGHBORHOOD DENSITY			
		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	97.67 (2.30)	98.56 (1.79)
	LOW	98.78 (1.70)	98.56 (1.92)
LOW FREQUENCY WORDS			
NEIGHBORHOOD DENSITY			
		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	96.78 (2.76)	97.11 (3.01)
	LOW	98.00 (2.38)	98.11 (2.11)

Insert Table 5.2 about here

Overall, words occurring in high density neighborhoods were responded to approximately 102 msec slower than words in low density neighborhoods. The effect of density was consistent across word frequency and neighborhood frequency in all cases but one. Virtually no reaction time differences between words in high and low density neighborhoods were observed for low frequency words occurring in high frequency neighborhoods. However, the three-way interaction suggesting a lack of statistical significance for this cell was far from significant, $F(1,11)=0.64$. Thus, although there was clearly a reduction of the effect of density for low frequency words in high frequency neighborhoods, there was no statistical support for any differential effects of density across word frequency or neighborhood frequency.

In summary, small effects of word frequency and neighborhood density were observed for the accuracy data. Although both effects were in the predicted direction, the magnitude of the differences were so small as to be almost negligible. However, a large effect of density was observed for the reaction time data. Words occurring in low density neighborhoods were named an average of 102 msec faster than words in high density neighborhoods.

DISCUSSION

The results of the naming study lend further support to the notion that the neighborhood structure of words in the mental lexicon strongly affects auditory word recognition. The reaction time data demonstrate that words with many neighbors are named more slowly than words with few neighbors. Perhaps the more interesting and crucial finding, however, is that no frequency effects were observed either in terms of stimulus word frequency or neighborhood frequency. This is in contrast to the findings from the perceptual identification and auditory lexical decision studies, in which consistent effects of frequency were observed. Although the present study examined only a subset of the words used in the previous studies, the number of stimuli was still quite large, and the difference in frequency between the high and low frequency words was substantial (mean for high frequency words = 145.95; mean for low frequency words = 4.33), as was the difference in mean frequency between high and low frequency neighborhoods (mean for high frequency neighborhoods = 245.17; mean for low frequency words = 60.15). Thus, there is therefore no reason to expect that the use of a subset of stimuli was responsible for the lack of frequency effects. Instead, the failure to observe word frequency and neighborhood frequency effects may lie in the circumvention of bias effects in the naming task.

As previously discussed, a precedent for the finding that frequency does not affect naming times can be found in the literature on visual word recognition. Balota and Chumbley (1984) compared lexical decision and naming times for high and low frequency printed words and found a marked reduction in the frequency effect for naming as compared to lexical decision. In a later study, Balota and Chumbley (1985) found that frequency effects in the naming of visually presented words may be due to the structure of the words themselves. They argued that any frequency effects observed in the naming task are due to factors correlated with frequency that affect articulation. The main conclusion to be drawn from these studies is that the naming task fails to reveal frequency effects that are unrelated to differences in the articulation of high and low frequency words.

Table 5.2

Neighborhood Analysis: Naming. Means and standard deviations for the reaction time data.

HIGH FREQUENCY WORDS			
NEIGHBORHOOD DENSITY			
		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	840 (183)	744 (175)
	LOW	852 (168)	716 (171)
LOW FREQUENCY WORDS			
NEIGHBORHOOD DENSITY			
		HIGH	LOW
MEAN NEIGHBORHOOD FREQUENCY	HIGH	731 (179)	736 (192)
	LOW	867 (178)	685 (174)

Paap, et al. (1986) have argued that the naming task for visually presented words does not require lexical access. Instead, they argue that a visually presented word may be named via a grapheme-to-phoneme route that bypasses the lexicon. The authors assume that frequency effects are only apparent once lexical access has been achieved, at which time frequency information is made available to the processing system. Paap et al. in fact showed that when the naming task is modified to require multiple lexical decisions prior to the naming response, large and consistent effects of frequency are observed. The effect of the lexical decisions is presumably to force subjects to access the lexicon, thus gaining access to frequency information.

Although the present results corroborate the finding that frequency effects are not obtained in the naming task (when highly controlled stimuli are used), the explanation put forth by Paap et al. does not account for finding that neighborhood density affected naming time. As previously mentioned, Paap et al. argue that visually presented words can be named by a direct mapping of graphemes onto phonemes; units corresponding to words need not be activated in order to name the visual stimulus. If one assumes that an auditorily presented word can be named by some sort of process that maps phoneme-to-phoneme, again circumventing activation of units corresponding to words, no effect of neighborhood density should be observed. This prediction was not supported by the results obtained in the present study, in which a large effect of neighborhood density was found.

It is therefore more reasonable to assume that the naming task requires activation of units corresponding to words, namely the decision units in the neighborhood activation model, but that word frequency information does not bias these decision units. The crucial question arises, then, as to why the naming task does not invoke the biasing properties of word frequency. One explanation may lie in the nature of the response required by the naming task. In this task, subjects are simply required to repeat back the stimulus word. No explicit decision is required regarding the lexicality of the item, as in lexical decision. Frequency biases in lexical decision may help to optimize decision times by allowing certain items to surpass the word or nonword criterion faster than would be expected with no bias. Nor is an explicit decision required as to the identity of a degraded item, as in perceptual identification. In perceptual identification, subjects may optimize performance by choosing words of higher probabilities of occurrence in the face of incomplete stimulus information. In contrast, no higher-level lexical information is required to make a naming response. Behavior in this task is optimized by simply deciding on the acoustic-phonetic identity of the stimulus word. Because the naming response requires a precise analysis of the acoustic-phonetics of the stimulus word in order to build an articulatory plan for executing a response, biases not based on the acoustic-phonetics themselves (e.g., frequency biases) may indeed hinder response generation. Given the response required by the naming task, therefore, subjects may optimize performance by focusing on discriminating among the acoustic-phonetic patterns and ignoring higher-level lexical information. Thus, frequency effects would not be expected to affect naming times. However, because an acoustic-phonetic pattern must be isolated in order to make the naming response, neighborhood density, or the number of similar acoustic-phonetic patterns corresponding to words in memory, would be expected to influence the time needed to generate a naming response. Indeed, this pattern of results was obtained in the present study.

Whatever the precise explanation for the failure to observe word frequency or neighborhood frequency effects in the auditory naming task, the results of the present study demonstrate that neighborhood density effects are clearly separate from frequency effects. In other words, the present study demonstrates that stimulus similarity and decision biases are separate effects that have differential effects on the levels of processing within the word recognition system. The present study also confirms the prediction of the neighborhood activation model that stimulus word frequency and neighborhood frequency effects must occur, or not occur, in tandem. That is, the model does not predict stimulus word frequency effects when no neighborhood frequency effects are present, and vice versa. Thus, the absence of one effect requires the absence of the other. Although unequivocal support for this prediction cannot be offered on the basis of two null results, the present study at least does not disconfirm the prediction.

In summary, the results of the auditory naming task support the basic tenets of the neighborhood activation model, in particular the claim that density effects are distinct from the effects of frequency. The results also provide further strong support for the notion that neighborhood structure affects auditory word recognition, such that increasing the size of the neighborhood increases the time needed to discriminate among items in memory. Finally, once again, the results from the present study demonstrate that effects of neighborhood structure can be obtained even when stimulus words are not degraded. Thus, the results from the present auditory naming task mesh well with the predictions of the neighborhood activation model and provide further support for the notion that a word is recognized in the context of similar words in memory.

CHAPTER SIX

CONCLUSION

The goal of the present research was to examine how the structural organization of the sound patterns of words in memory influences auditory word recognition. The term "structural organization" was defined specifically in terms of similarity relations among the sound patterns corresponding to words in the mental lexicon. In particular, the effects of the structural organization of "similarity neighborhoods" were examined. A "neighborhood" was defined as a collection of words that are phonetically similar to the stimulus word. In addition to structural relations among words, the relative effects of word frequency were examined in the context of similarity neighborhoods. It was hypothesized, and subsequently confirmed, that stimulus word frequency effects are a function of the neighbors of the stimulus word as well as the frequencies of these neighbors.

The present investigation adopted the approach of many current models of stimulus identification (see Broadbent, 1967; Nakatani, 1970; Smith, 1980), in which word identification was assumed to be a function of similarity and bias. Similarity defines the set of alternatives that are to be discriminated and chosen among in deriving a response. Bias refers to those higher-level factors that adjust the probability of choosing among the alternatives activated in memory. In the present investigation, the biasing properties of word frequency were examined.

Unfortunately, "bias" has taken on a somewhat pejorative sense in the literature, referring to some ill-specified process by which the subject explicitly chooses among possible alternatives based on a priori probabilities of occurrence or previous contextual information. It was assumed in the present approach, however, that biases arising from information regarding word frequency constitute an important aspect of the process of word recognition and serve to optimize recognition in situations requiring identification of a degraded stimulus or a speeded response to a stimulus presented in the clear. The term "response bias" has been cautiously avoided because this term implies that frequency biases are only important in influencing subjects' volitional generation of responses in laboratory task situations. There is no evidence that frequency bias effects are restricted to the laboratory, nor that such effects are under the conscious control of the subject (Hasher & Zacks, 1984; Smith, 1980).

Thus, bias is assumed to be a crucial component in word recognition and not an artifactual effect that should be treated as a nuisance in the study of the "pure" perceptual processing of words. Indeed, current interactive-activation models of word recognition incorporate bias effects in the recognition system in such a way that the effects of bias and perception are virtually indistinguishable (McClelland & Rumelhart, 1981; see also Slowiaczek, Nusbaum, & Pisoni, 1987). In short, in the present research, frequency effects were assumed to arise from biases in the word recognition system that constitute important and fundamental aspects of the process of word recognition. In addition, frequency effects were assumed to operate in an extremely sophisticated manner, influencing the entire neighborhood of words activated in memory by stimulus input.

The effects of similarity neighborhood structure (i.e., density and neighborhood frequency) and stimulus word frequency were examined in a number of experiments which included perceptual identification of auditorily

presented words, auditory lexical decision, and auditory word naming. Each of these studies provided robust evidence for the role of similarity neighborhood structure in auditory word recognition, although each of the tasks produced somewhat different patterns of results. It was shown that the neighborhood activation model (NAM) provided a reasonable account of the results for each of the experimental tasks. In addition, NAM was shown to describe adequately the combined effects of similarity neighborhood structure and word frequency. Basically, NAM proposes that stimulus input activates a set of acoustic-phonetic patterns in memory that are monitored by a system of word decision units sensitive to higher-level lexical information. Choices within the decision system are based on the overall level of activity in the system, the activity level of the acoustic-phonetic patterns, and biases stemming from higher-level lexical information, which includes word frequency. Before discussing NAM in detail and its comparison to other current models of auditory word recognition, I turn to a brief review of the empirical results reported in the previous chapters.

REVIEW OF MAJOR RESULTS

Perceptual Identification

In Chapters 2 and 3, the effects of neighborhood structure and word frequency were examined in a perceptual identification experiment. Approximately 900 three-phoneme monosyllabic words were presented for identification at three S/N ratios. Accuracy of identification served as the dependent variable. Although effects of neighborhood structure were observed in Chapter 2, in which neighborhoods were computed on the basis of one phoneme substitutions, additions, and deletions, the effects were small and provided at best moderate evidence for effects of similarity neighborhood structure. A further attempt at computing similarity neighborhoods based on basic segmental intelligibility and confusability proved much more successful. It was shown that a choice rule--dubbed the neighborhood probability rule--could be developed that adequately predicted identification performance. This rule incorporates estimates of stimulus word intelligibility and neighbor confusability in predicting the probability of choosing a stimulus word from among its phonetically similar neighbors. In addition, it was shown that frequency-weighting factors incorporated in the rule improved prediction. These frequency-weighting factors were bias parameters that adjusted the initial probabilities of the acoustic-phonetic patterns of the neighbors relative to the pattern of the stimulus word. One important finding regarding the frequency-weighted neighborhood probability rule was that frequency of the stimulus word per se was a poor predictor of identification performance. Instead, it was shown that the relations among the frequencies of the stimulus word and its neighbors served to better capture the effects of word frequency. The results demonstrated that the frequency effect is a relative one crucially dependent on the frequencies of the neighbors of the stimulus word.

Auditory Lexical Decision

As a further test of the role of neighborhood structure in auditory word recognition, an auditory lexical decision task was employed in Chapter 4. The goal of this study was to examine the effects of similarity neighborhood structure in the absence of stimulus degradation. In addition, the auditory lexical decision task enabled the collection of response times in order to determine if similarity neighborhood structure affects not only the accuracy of word recognition but also the time course.

Regression analyses on the word data showed no effects of similarity neighborhood structure on reaction time, although an effect of neighborhood density was observed for the accuracy data. However, this latter effect was in the opposite direction of that predicted. It was argued that the effect of density on accuracy suggested interactions in the data that may have served to eliminate or attenuate the results of the regression analyses. In particular, it was argued that the discrimination among words and nonwords required by the lexical decision task may have had differential effects for high and low frequency words, thus producing interactions that may have attenuated the correlations of the neighborhood variables with the reaction time and accuracy data.

In a subsequent analysis in which the original word data were partitioned into orthogonal cells and submitted to analyses of variance, neighborhood density was in fact found to interact with word frequency. In addition, the interactions of neighborhood density and word frequency were opposite for the reaction time and accuracy data. Effects of neighborhood density on classification time were observed for high frequency words but not for low frequency words. High frequency words in dense neighborhoods were responded to more slowly than high frequency words in sparse neighborhoods. In terms of the accuracy data, density effects were only observed for low frequency words, and these effects were opposite to those predicted: Low frequency words in dense neighborhoods were responded to at higher levels of accuracy than low frequency words in sparse neighborhoods. In addition to the effects of density, consistent effects of word frequency and neighborhood frequency were obtained.

An analysis of the auditory lexical decision task in the context of NAM demonstrated that the relative effects of neighborhood density on reaction time and accuracy could be accounted for under the assumption of reaction-time deadline expiration. Very simply, it was assumed that neighborhood density would affect reaction times if a response could be initiated prior to the expiration of the deadline. However, if a response could not be initiated prior to the deadline, it was assumed that no effect of density on accuracy should be observed. Conversely, if decisions were forced at the reaction-time deadline, only effects of density on accuracy should be observed. Given that low frequency words require longer processing times than high frequency words, the differential effects of density became explicable in the context of NAM.

Analyses of the nonword responses provided further strong support for the role of neighborhood structure. Nonwords in high density neighborhoods were responded to more slowly than nonwords in low density neighborhoods. And, nonwords in high frequency neighborhoods were responded to more slowly than nonwords in low frequency neighborhoods. It was again shown that NAM could easily account for these effects.

Altogether, the results from the auditory lexical decision experiment provided further support for the notion that similarity neighborhood structure affects auditory word recognition in predictable and systematic ways. In addition, these results demonstrated that similarity neighborhood structure affects processing time as well as accuracy and that the effects of neighborhood structure are demonstrable in the absence of stimulus degradation. Finally, the overall pattern of results for the auditory lexical decision experiment proved consistent with the claim that decision units only monitor words in memory.

Auditory Word Naming

In a final experiment, the effects of neighborhood structure were examined using the auditory word naming task. The motivation for using this task was two-fold: First, it avoids the requirement of word-nonword discriminations which may have attenuated the effects of neighborhood structure in the auditory lexical decision task. Second, recent evidence from visual word naming studies (Balota & Chumbley, 1984, 1985) suggested that the naming task is insensitive to word frequency effects. It was hypothesized that whereas the effects of word frequency may be circumvented by the naming task, effects of neighborhood density should nonetheless remain, given that neighborhood density effects arise at the earliest level of analysis in the system and are inherent in the activation levels of the word decision units.

The results of this study confirmed that the naming task is insensitive to stimulus word frequency and neighborhood frequency effects when the stimuli are carefully constructed. However, a large effect of neighborhood density was still observed: Words in high density neighborhoods were named approximately 100 msec slower than words in low density neighborhoods, a result consistent with the predictions of NAM. The failure to observe stimulus word frequency and neighborhood frequency results was attributed to the level of analysis required to produce a response in the naming task. It was furthermore argued that the results of this study provided support for the notion of frequency as a biasing factor in the system of word decision units postulated in NAM. In particular, it was argued that frequency must be a higher-level biasing factor in order to countenance the result that the effect of frequency can be circumvented by task requirements. If frequency were an inherent aspect of word thresholds or activation levels, effects of frequency would not be expected to disappear in the naming task. It was furthermore argued that the word decision units had to be activated in the naming task in order to produce the neighborhood density effect. If only phonemes or segments were synthesized in order to produce the naming response, the number of neighbors should have made no difference in naming response times. Thus, it was argued that information regarding words must have been activated at some level, despite the fact that frequency information was not used in producing a naming response.

The results of each of the experiments summarized above lend strong support to the proposal that the number and nature of items activated in memory by the stimulus input influence the speed and accuracy of auditory word recognition. These results strongly suggest a model of auditory word recognition in which multiple acoustic-phonetic patterns are activated, discriminated among, and decided upon according to the inherent acoustic-phonetic similarity of the patterns and higher-level lexical information biasing word decisions. In short, these studies demonstrate that words are recognized in the context of similar words in memory and that recognition is a function of the number of items that must be discriminated among as well as the biases influencing the decisions based on these discriminations.

SUMMARY OF THE NEIGHBORHOOD ACTIVATION MODEL

The neighborhood activation model (NAM) was developed to explain the results of the perceptual identification experiment and was extended to the findings from the auditory lexical decision and auditory word naming studies. The model is primarily a processing instantiation of the frequency-weighted neighborhood probability rule developed in Chapter 3. Basically, the model assumes that a set of similar acoustic-phonetic patterns are activated in memory on the basis of stimulus input. The activation levels of these patterns are assumed to be a direct function of their similarity to the stimulus input. These patterns are further assumed to represent both words and nonwords. Over the course of processing, stimulus input serves to resolve or "refine" a pattern in a manner suggested by Nusbaum's Phonetic Refinement Theory (see Pisoni et al., 1985). That is, as processing proceeds, the pattern corresponding to the stimulus input is refined, receiving successively higher levels of activation, while the activation levels of similar patterns are reduced.

Words emerge in NAM when a system of word decision units tuned to the acoustic-phonetic patterns are activated. The activation of the decision units is assumed to be direct, in the sense of logogen theory (Morton, 1979) and cohort theory (Marslen-Wilson & Welsh, 1978). In addition, as in cohort theory, the system of word units is assumed to be based only on the activation of the acoustic-phonetic patterns. That is, word recognition is assumed to be, at least initially, completely bottom driven. Once the word decision units are activated, they monitor a number of sources of information. The first source of information is the activation of the acoustic-phonetic patterns, which have previously served to activate the decision units themselves. The word decision units also monitor the overall level of activity in the decision system itself, much like processing units monitor the net activity level of the system in the TRACE model (Elman & McClelland, 1986). Finally, the decision units are tuned to higher-level lexical information, which includes word frequency. This information serves to bias the decisions of the units by multiplying the activity levels of the acoustic-phonetic patterns by the frequencies of the words to which they respond. The values that serve as the output of the decision units are assumed to be computed via a rule similar to the frequency-weighted neighborhood probability rule discussed in Chapter 3.

Word recognition in NAM may be accomplished in a number of ways, depending on the requirements of the task. In situations in which the stimulus input is degraded, word recognition is accomplished by evaluating the values computed by the decision units and picking a response based on these values. When speeded responses are required, it is assumed that the subject sets a criterion for responding that, once exceeded by the output of a decision unit, results in the recognition of a word. Word recognition is defined explicitly as the choice of a particular pattern by the system of decision units. Lexical access is assumed to occur once a decision unit makes all of the information it was monitoring available to working memory. Thus, the decision units act as gates on the acoustic-phonetic and lexical information available to the processing system. If insufficient evidence for a word is provided by the decision system, the activation levels of the acoustic-phonetic patterns themselves may be consulted, resulting in the recognition of a nonword pattern.

NAM places much of the burden of auditory word recognition on the discrimination among similar acoustic-phonetic patterns corresponding to words and the decisions necessary for choosing among these patterns. In addition,

NAM accounts for word frequency effects by allowing frequency information to bias the decisions of the word decision units. However, because each word decision unit computes values based on its acoustic-phonetic pattern as well as the overall level of activity in the decision system, decisions are assumed to be context sensitive. Thus, frequency is assumed to be a relative factor. That is, if many decision units are receiving strong frequency biases, a decision unit for a given high frequency word may compute relatively low values. Likewise, a decision unit for a low frequency word may quickly begin to output high values if there is little other activity in the system. Thus, effects of frequency are not assumed to be absolute, but dependent on the activity level of the decision system as a whole.

Comparison of NAM to Other Models of Word Recognition

As previously mentioned, NAM bears a strong resemblance to other models of auditory word recognition, and many of the concepts incorporated in the model have precedents in previous theories of auditory word recognition. However, as will be argued below, NAM makes certain predictions that are inconsistent with current models of auditory word recognition, in particular with regard to the roles of frequency and similarity. I will now turn to a discussion of some of the more influential models of word recognition in order to highlight the fundamental differences and similarities between NAM and these models.

Logogen Theory

Morton (1969, 1979) has proposed a model of word recognition based on a system of "logogens" that monitor bottom-up sensory information and top-down contextual and lexical information. Information from either of these sources serves to drive the logogens toward threshold. Once a threshold is reached, the information to which the logogen corresponds is made available to the processing system and a word is said to be recognized and accessed. Morton accounts for word frequency effects in the logogen model by assuming that high frequency words require less evidence than low frequency words for crossing threshold. Morton thus refers to logogen theory as an evidence-bias model.

The resemblance between Morton's system of logogens and the system of word decision units in NAM is quite strong. Both logogens and word decision units monitor top-down and bottom-up information. In addition, both logogens and word decision units are assumed to prohibit information from becoming available to the general processing system until a decision regarding the identity of the word has been made. However, word decision units differ from logogens in a number of crucial ways.

First, logogens can be activated on the basis of either top-down or bottom-up information. A logogen is insensitive to the source of the information it monitors. In contrast, word decision units are activated only on the basis of acoustic-phonetic input. This aspect of NAM ensures that stimulus input will not activate a clearly inconsistent word decision unit. Morton (1979) argues against such a system by citing evidence of responses to degraded word stimuli that appear to be clearly dissimilar to the stimulus word. However, this argument fails to acknowledge that under extreme levels of degradation, the stimulus input may be so degraded as to activate what may appear to be phonetically dissimilar words. That is, under conditions of extreme stimulus degradation, the available acoustic-phonetic information may overlap with a large number of words, given that this information is so

impoverished. Indeed, Broadbent (1967) and Triesman (1978b) both argue that similarity among items in memory increases as stimulus degradation increases. Thus, under high levels of degradation, words quite dissimilar to the stimulus input may be produced as responses. This does not vitiate the claim, however, that the initial activation of the word decision units is based on the acoustic-phonetic input. Morton's results simply demonstrate that many disparate items may be considered as similar to the stimulus input when stimulus degradation is high.

Perhaps the most crucial difference between logogens and the word decision units hinges on the problem of accounting for neighborhood structural effects. Logogens are assumed to be independent processing units with no interconnections. The lack of crosstalk among logogens makes it difficult to account for the findings that words in highly dense or confusable neighborhoods take longer to respond than words in less dense or less confusable neighborhoods. Because logogens are independent processing units, stimulus input should push a given logogen over threshold at the same point in time, regardless of whether the stimulus input activates many or few logogens. Granted, accuracy differences between dense and sparse neighborhoods may arise because there is a higher probability that logogens corresponding to similar words may surpass threshold prior to the logogen corresponding to the stimulus input. It is not so clear, however, how logogen theory would account for neighborhood density effects on reaction times. When presented with clearly specified acoustic-phonetic information, as in auditory lexical decision or auditory word naming, the logogen corresponding to the stimulus input should always cross threshold at the same point in time, regardless of the activity levels of other logogens, assuming that word frequency is held constant. The results for the high frequency words for the auditory lexical decision task reported in Chapter 4 contradict this prediction, as do the results for the auditory word naming task.

The most fundamental problem that the present set of results poses for logogen theory concerns the robust findings that frequency effects are dependent on the neighborhood structure of the stimulus word. In the perceptual identification study, it was shown that certain classes of high and low frequency words are responded to at equal levels of accuracy if the neighborhood structures of the words are equated. Because logogens corresponding to high and low frequency have differing thresholds, low frequency words should always require more evidence than high frequency words in order to cross threshold. Because a single logogen has no knowledge of the activation levels of other logogens, it is difficult to explain within logogen theory how the frequencies of neighbors could influence recognition of the stimulus word. One could again assume that the effects of neighborhood frequency and density on accuracy reflect incorrect logogens surpassing threshold. That is, it is possible that both neighborhood density and neighborhood frequency increase the probability of incorrect logogens reaching threshold, thus depressing accuracy of identification for words occurring in high density and/or high frequency neighborhoods. However, such an account does not explain the effects of neighborhood density and neighborhood frequency on reaction times observed in the auditory lexical decision experiment. The time for a given logogen to reach threshold cannot be influenced by the activations of other logogens, and thus logogen theory fails to account adequately for the present set of results.

Finally, logogen theory has no mechanism for explaining the results of the naming study. Recall that in the naming study it was argued that word units must have been accessed by subjects in order to produce the effect of neighborhood density. However, no effects of word frequency or neighborhood

frequency were observed. It is perhaps possible that the thresholds for logogens corresponding to high and low frequency words were temporarily equated due to some unspecified property of the naming task. However, not only is this solution extremely inelegant and unparsimonious, it seriously calls into question logogen theory's claim that thresholds are intrinsic to the logogens themselves and arise over time as a function of degree of exposure to words.

A final problem for logogen theory concerns the nonword data from the auditory lexical decision experiment. It has been proposed that nonword decisions in the logogen model are made in a similar manner to nonword decisions in NAM (Coltheart, et al., 1976). Specifically, a nonword decision is executed when no logogen fires. However, because the activation levels within the logogens are not available for inspection (i.e., logogens are either above or below threshold), it is difficult to account for the finding that the number and nature of words activated by the nonword stimulus influence reaction time. As logogen theory stands, there is no means for evaluating the overall level of activity in the logogen system, and there is therefore no mechanism for making faster decisions to nonwords with fewer neighbors or lower frequency neighbors. The nonword data from the auditory lexical decision experiment thus prove problematic for a system of independent processing units that respond only upon surpassing an intrinsic thresholds.

NAM, on the other hand, provides a coherent description of the present set of results by assuming that the decision units are interconnected and that frequency effects arise from biases stemming from higher-level sources of information. Modifications of logogen theory may be possible to account for the present results, but it is very likely that the resulting model would bear a strong resemblance to NAM. Nonetheless, there are important similarities between NAM and logogen theory, owing to the fact that NAM incorporates many ideas from logogen theory. In particular, NAM assumes a system of word decision units that serve as the interface between the acoustic-phonetic input and higher-level information, as proposed by logogen theory. However, due to the interconnectedness of the system of word decision units, NAM is able to account for the effects of neighborhood structure, whereas logogen theory apparently is not.

Cohort Theory

Perhaps the most influential of current models of auditory word recognition is cohort theory, proposed by Marslen-Wilson (Marslen-Wilson & Welsh, 1978; Marslen-Wilson & Tyler, 1980; Marslen-Wilson, 1984, 1986). According to this theory, a "cohort" of words is activated in memory on the basis of the initial acoustic-phonetic input of the stimulus word. Words in the cohort are then eliminated by two sources of information: continued acoustic-phonetic input and top-down contextual information. That is, words in the cohort are ruled out or deactivated by continued processing of the stimulus information as well as by inconsistent contextual information. A given word is recognized when it is the only word remaining in the cohort.

Cohort theory has provided a number of valuable insights into the temporal processing of spoken words. In previous versions of the theory, however, no attempt was made to account for word frequency effects. In a recent version of the theory, though, Marslen-Wilson (1986) has incorporated a mechanism for accounting for word frequency effects by assuming words in a cohort have differing levels of activation depending on their frequencies of occurrence. Words with higher levels of activation take longer to eliminate

from the cohort than words with lower levels of activation, thus affording at least an initial advantage to high frequency words. Because the latter version of cohort theory represents a significant improvement over the initial formulation of the theory, only this version will be considered in the present discussion.

Cohort theory and NAM are similar in the respect that both models assume bottom-up priority in the activation of items in memory. Furthermore, both models assume that a set of items are activated and processed in parallel. In addition, both models state that items receive reduced levels of activity as disconfirming acoustic-phonetic information is presented. Unlike cohort theory, however, NAM at this stage of formulation has little to say about the time course of effects in the word recognition system, primarily due to the fact that the model was developed on the basis of data from very short words. Indeed, as stated earlier, some of the aspects of cohort theory may have to be incorporated into NAM in order to account for the recognition of longer words. Nonetheless, cohort theory and NAM do make fundamentally different predictions, at least for short stimuli.

Marslen-Wilson (1986) argues that because cohort theory is realized as a parallel system, no effects of set size should be observed on recognition. Words in a cohort are assumed to be activated at no cost. NAM is also realized as a system of parallel processing units, but the fundamental claim of NAM is that the nature and number of items activated in memory influence the accuracy as well as the speed of recognition. This prediction of NAM stems from the claim that the word decision units are sensitive to the overall level of activity in the decision system and are therefore influenced by the number and nature of competing items. Evidence for this claim was provided by each of the experiments previously reported.

Marslen-Wilson (1986) argues for the claim that set size has no effect on recognition performance on the basis of a set of experiments examining lexical decisions for nonwords. He claims that if nonwords are matched according to the point at which they diverge from words, no effect of set size is observed on reaction times. This is in contradiction to the findings in the lexical decision experiment reported earlier in which large effects of neighborhood density (i.e., set size) and neighborhood frequency were observed for nonwords. Note that because of the manner in which these nonwords were constructed, each of the nonwords diverged from words at the third phoneme (see Chapter 4). Thus, set size effects were demonstrated even when divergence points were equated. Given that Marslen-Wilson's claim of no effects of set size are based on null results, the positive findings reported in Chapter 4 for the nonwords seriously call this claim into question.

Indeed, each of the experiments reported previously fail to support the notion that the number of items activated in memory has no influence on recognition performance. Although Marslen-Wilson may object to the results from the perceptual identification study, claiming that the use of "noisy" stimuli induce post-perceptual processes, the results from the lexical decision study as well as the auditory naming study clearly contradict a fundamental claim of cohort theory. Indeed, it is not even clear that the postulation of some vague "post-perceptual" processes indicts the results from the perceptual identification study, which showed significant effects of neighborhood structure on identification performance. In short, the results of the present set of studies taken together refute certain crucial claims of cohort theory.

The results of the naming study also provide counter evidence to cohort theory's treatment of word frequency. All words used in the naming study can be assumed to have had approximately equal divergence points or isolation points by virtue of their short length (see Luce, 1986a). Indeed, it has yet to be shown for short word stimuli that divergence points influence on-line recognition. Thus, one can safely assume that these stimuli did not differ in important ways in terms of their isolation points. However, despite equivalent isolation points, high frequency words were named no faster than low frequency words, in contradiction to the predictions made by the most recent version of cohort theory. In addition, because there was a strong effect of density, it cannot be assumed that lexical items were bypassed in the generation of the naming response. Thus, the current version of cohort theory further fails to account for the results obtained in the present investigation.

As previously argued, an adequate model of auditory word recognition cannot assume differing inherent activation levels or thresholds for the units monitoring high and low frequency words. Instead, the effects of frequency are best described as biases on the decision units responsible for choosing among activated lexical items. By treating the effects of frequency as biases on the decision process, one can account for results demonstrating the lability of the frequency effect depending on task requirements (e.g., Pollack, et al., 1959) and higher-level sources of information (Grosjean & Itzler, 1984). Thus, the instantiation of frequency in the latest version of cohort theory is difficult to countenance. NAM, however, provides a more principled explanation of the effects of word frequency on both the stimulus word and its neighbors.

As it stands, cohort theory is clearly inconsistent with a number of findings from the previous studies. Although cohort theory still makes a number of important claims regarding the temporal processing of longer spoken words, which NAM has yet to address, certain fundamental aspects of the cohort theory appear at this point to be mistaken.

Other Theories of Word Recognition

Although Morton's logogen theory and Marslen-Wilson's cohort theory constitute the dominant theories of auditory word recognition at the present time, it is perhaps wise to consider briefly a few additional models based primarily on research on visual word recognition. It is suggested that these models fall prey to many of the same problems encountered by logogen and cohort theory, namely the inability to account adequately for effects of neighborhood structure and word frequency.

Forster's Search Theory. Forster's (1976, 1979) search theory has been very influential in research on visual word recognition, although the theory attempts to account for auditory word recognition as well. In terms of the auditory recognition of words, Forster's theory claims that access to the mental lexicon is gained through a peripheral access file containing phonetic codes for words. A "bin" of these codes is first selected based on acoustic-phonetic similarity. The entries within this bin are assumed to be arranged according to frequency, such that high frequency words are searched prior to low frequency words. When a match is made in the peripheral access file, search is terminated and a pointer is used to locate the word in the master lexicon, where all information regarding the word resides. A post-access check is then made to ensure that the entry in the master lexicon matches the stimulus input.

Can the search model account for the neighborhood structure effects observed in the previous studies? If a bin of codes is activated based on similarity, density effects may be observed. Because a large number of codes may be activated for dense neighborhoods, it should take longer to search the bin in order to locate the code corresponding to the stimulus input. Furthermore, if the bin contains highly similar items, errors in locating the correct entry may arise, thus producing the effects of neighborhood density on accuracy previously observed. Finally, if the set of codes to be searched are high in frequency (i.e., constitute a high frequency neighborhood), it will take longer to locate the code in the bin, owing to the fact that high frequency words are searched first, thus producing effects of neighborhood frequency. In short, Forster's search model may in fact produce the observed effects of neighborhood density and frequency. However, it is difficult to state with precision that the model would predict the precise effects of neighborhood structure observed. In particular, because bins are first chosen on the basis of similarity and subsequently ordered by frequency, the model may predict larger effects of neighborhood frequency than are actually observed. In short, it is not clear how similarity and frequency would interact in the search model. If the peripheral access codes are arranged according to frequency, less similar items may have to be searched prior to more similar items simply because of the frequency ordering. This suggests a less than optimal system that is perhaps too strongly biased by frequency.

Although the search model may account for neighborhood structural effects, it cannot easily account for the findings that frequency effects may be circumvented in the naming task, whereas density effects cannot. It is a fundamental claim of the model that the access codes are arranged according to frequency, and there is no means in the model by which frequency effects may be bypassed, especially once a bin is activated (which is a necessary prerequisite for effects of density). Thus, although the search model may at first appear to be able to produce the effects of neighborhood structure previously demonstrated, it cannot in its present formulation produce effects of density without frequency effects. Moreover, it is possible that the frequency-ordering of the peripheral access file would overestimate neighborhood frequency effects and thereby underestimate the effects of neighborhood similarity.

Activation-Verification Model. The activation-verification model proposed by Paap et al. (1982, see also Becker, 1976) is similar to Forster's search theory. In this model, a set of similar items are activated in memory and verified against the stimulus input. Verification is assumed to be ordered by frequency, such that high frequency words are verified before low frequency words. Again, as in the search theory, neighborhood density and frequency effects may be predicted by the model, although it is unclear precisely how these effects will interact. As with the search model, further tests of the effects of neighborhood structural effects are required to evaluate the performance of the activation-verification model; there are no fundamental theoretical problems in the activation-verification model that prevent it from showing density and neighborhood frequency effects. However, because frequency is inherent in the process of verification, it is difficult to state how density effects may arise in the absence of frequency effects. Furthermore, it is not certain that the locus of the frequency effect in the activation-verification model is actually at the stage of verification. If, as suggested by Dobbs, Friedman, and Lloyd (1985), frequency effects arise prior to verification, it is even more difficult to account for density effects in the absence of frequency effects in the naming task.

Interactive-Activation Model. Interactive-activation models of visual (McClelland & Rumelhart, 1981) and auditory word recognition (Elman & McClelland, 1986) have been increasingly popular in recent years, primarily due to the breadth of phenomena accounted for by these models. Basically, interactive-activation models assume a set of primitive processing units that are densely connected to one another. In models such as TRACE (Elman & McClelland, 1985), processing units or nodes have excitatory connections between levels and inhibitory connections among levels. These connections serve to raise and lower activation levels of the nodes depending on the stimulus input and the activity of the overall system. As noted earlier, the system of decision units proposed in NAM may very well be realized as an interacting set of processing units. NAM may therefore turn out to be virtually indistinguishable from an interactive-activation model. Indeed, the decision values computed by the decision units in NAM may very well arise in a system of interconnected nodes having excitatory and inhibitory connections. McClelland and Rumelhart (1981) in fact discuss the possibility that neighborhood structure (both density and frequency) may be accounted for by their model, although they admit that further work is necessary in order to confirm their speculations.

At present, little can be said regarding the effects of neighborhood structure in an interactive-activation model short of actually conducting the simulations required for testing the model. One potential problem is the treatment of frequency within the interactive-activation framework as an inherent component of the activation levels of words in memory. As previously argued, the present data strongly suggest a labile frequency bias, and not an inherent threshold or activation level. Thus, it is unclear that an interactive-activation model that assumes higher levels of activation for high frequency words can account for the present set of data. Nonetheless, the interactive-activation approach is suggestive of an interesting means of further specifying NAM, although additional research and model development is clearly required.

Summary of Word Recognition Models

The previous discussion of current models of word recognition suggests that none of the models can adequately account for the neighborhood structural effects observed in the previous experiments. In particular, logogen theory and cohort theory appear to be the most difficult to reconcile with the present data, although this may simply be a function of the degree of specificity of description provided by each of these models. Search and verification theories may provide potential accounts of neighborhood structure and frequency, although they are hard-pressed to explain the findings of density effects in the absence of frequency effects. Finally, the interactive-activation approach may ultimately prove most successful in accounting for the present set of findings, although it was again argued that any model not treating frequency as primarily a bias effect on decisions cannot adequately account for the present data. In short, NAM appears at present to provide the most consistent account of the effects observed, primarily due to the interactive nature of the word decision units and the biasing effect of frequency on these units.

ADDITIONAL CONSIDERATIONS AND FUTURE DIRECTIONS

Contextual Effects. Little has been said about the roles of context in NAM, mainly because these issues have yet to be addressed empirically in terms of precise manipulations of contextual and priming effects on neighborhood activation. However, some speculative comments about the incorporation of contextual information within the model are possible at this time.

First, it has been well documented that words are more easily perceived in biasing or constraining contexts (Leventhal, 1973; Lieberman, 1963; Luce, 1983; Miller, Heise, & Lichten, 1951) than in isolation or in neutral or inhibitory contexts. Within NAM, it is assumed that prior constraining contextual information will operate via the same mechanism as word frequency. That is, higher-level lexical information activated on the basis of prior context will serve to bias the decision units. Thus, the activation levels of the acoustic-phonetic patterns will be adjusted according to the bias introduced by contextual information acting on the decision units. Facilitation may thus be observed for contextually consistent words via increased activation levels in the decision units. Inhibition is also predicted, in that higher-level information would be assumed to bias incorrect alternatives, thus raising the activation levels of competitors and lowering the levels of the acoustic-phonetic patterns consistent with the stimulus input. Contextual effects in NAM are therefore assumed to operate via the same fundamental mechanism as word frequency. Thus, like cohort theory and logogen theory, NAM incorporates a means for producing contextual effects in auditory word recognition. However, unlike cohort theory, NAM assumes that context serves to bias decision processes and not to eliminate words from the cohort.

Priming Effects. Slowiaczek, Nusbaum, and Pisoni (1987) have recently conducted an interesting test of cohort theory using a phonological priming technique. They presented word and nonword primes before target words embedded in noise. The word and nonword primes varied in the degree to which they overlapped phonologically with the target words. Slowiaczek et al. showed that increasing phonological overlap increased subjects' ability to identify the target words. The results showed equivalent effects of priming regardless of whether the overlap was from the beginnings or ends of the prime and the target, calling into question cohort theory's claim that cohorts are activated on the basis of word-initial information. In addition, the authors demonstrated effects of priming by nonwords, although these effects were smaller than those observed for the nonword primes.

Phonological priming may arise in the system of decision units via adjustment of the bias parameters in the word decision units. Because of expected contingencies between the prime and the target induced by the task (see Tweedy & Lapinski, 1977), the system of decision units may adjust their bias parameters based on presentation of the priming stimulus. Thus, upon presentation of the target word, the decision units activated by the prime will reach higher-levels of activation, thus producing higher decision values for those words overlapping with the prime. Nonword primes would likewise adjust the decision units because they would increase the activation levels of similar words. However, because no word would be consistent with a nonword prime, the decision units would not have reached levels of activation as high as in the case in which a word prime was presented. Thus, smaller effects of a nonword prime on the decision units for words would be expected. This explanation is similar to one offered by Nusbaum and Slowiaczek (1985), although the locus of the effect in the present account is placed on the decision units and is assumed to act via biases.

An important test of the explanation of NAM of these results would be to vary the probability of overlap between the target and the prime. If the effects of priming are bias effects on the word decision units, situations in which there is a very low probability of prime-target overlap should produce little or no priming; situations in which there is a high probability of prime-target overlap should produce stronger and more consistent priming effects. However, equivalent effects of priming in both cases would rule out the decision bias explanation and suggest that the prime acts to increase the activation levels of similar acoustic-phonetic patterns prior to their activation of word decision units.

Neighborhood Structural Effects in Visual Word Recognition. Although NAM was developed to account for auditory word recognition, there is evidence in the visual word recognition literature for neighborhood structural effects. In particular, Havens and Foote (1963) have demonstrated that frequency effects in the identification of tachistoscopically presented words are mediated by the number of competitors, or neighbors, of the stimulus word. These results bear a strong similarity to those reported in the previous experiments on auditory word recognition.

In a recent study, Luce (1986b) further examined the role of neighborhood structure for visually presented, masked items. As in the present experiments on auditory word recognition, neighborhood density, neighborhood frequency, and stimulus word frequency were orthogonally combined, and printed words falling into each of two classes (high and low) of each of these variables were presented for whole report identification. The results showed consistent effects of each of these variables. High frequency words were responded to more accurately than low frequency words. Words in high frequency neighborhoods were responded to less accurately than words in low frequency neighborhoods. However, an interesting result was obtained for neighborhood density: Visually presented words with many neighbors were actually responded to better than visually presented words with few neighbors.

Although this effect is contrary to the effect observed for neighborhood density on auditorily presented words, it was shown that neighborhood density was highly correlated with another factor previously shown to affect visual word recognition, namely letter positional frequency (Mason, 1975; Massaro, Venezky, & Taylor, 1979; Rumelhart & Siple, 1974). Words in high density neighborhoods have letters of higher positional frequency than words in low density neighborhoods, and it has previously been demonstrated that high letter positional frequency enhances identification of visually presented words (see references cited above). Thus, the reverse effects of density were presumably due to the confounding of density and letter positional frequency. Further analysis of the data revealed that although density enhanced identification of letters, neighborhood frequency had no effect on accuracy of letter identification. Instead, neighborhood frequency affected only word identification, such that the words in high density neighborhoods were identified less well than words in low density neighborhoods. It was argued that letter positional frequency, or density, aids in letter identification, but that neighborhood frequency inhibits word identification by interfering with word response selection based on the correctly perceived letters.

As it stands, NAM cannot account for the finding that visually presented words in high density neighborhoods are identified better than words in low density neighborhoods, although the effect of neighborhood frequency is easily explicable in the model. However, it may be possible to build a revised model for visual and auditory word recognition in which a separate level of frequency is coded at the letter level. Nonetheless, the results from the

visual identification study suggest that visual and auditory word recognition may operate via very different mechanisms. In particular, letter frequency has been shown to be an important determinant of the recognition of printed words, whereas no such effects of frequency on the segment-level have been demonstrated for auditorily presented words.

Frequency and Familiarity. One final result of the present studies concerning the relationship of frequency and familiarity deserves mention. Recall that the word stimuli employed in each of the previous experiments were all rated as highly familiar to subjects in a previous study by Nusbaum, Davis, and Pisoni (1984). However, despite this fact, consistent effects of frequency based on objective word counts were observed in the perceptual identification and auditory lexical decision experiments. Similar findings have been reported for visual lexical decision (Nusbaum & Dedina, 1985). Taken together, therefore, these results suggest that two separate phenomena underlie rated familiarity and experienced frequency as indexed by word counts. Although objective frequency and subject familiarity are highly correlated (Nusbaum et al., 1984), there nonetheless appears to be an effect of repeated exposure of a word that is independent of subjects judged familiarity, contrary to previous claims (Gernsbacher, 1984). What is important in the present set of experiments is the demonstration that word frequency effects are present even when the words are well known to subjects. Further research on the frequency-familiarity distinction should help to uncover precisely how frequency and familiarity information is coded in memory. Nevertheless, the present findings implicate the role of repeated exposure to words above and beyond the judged familiarity of the words.

Future Directions in the Development of NAM

NAM constitutes only a preliminary model of auditory word recognition. Undoubtedly, further tests of the basic claims of the model are in order, in particular the claim that word recognition is best described by processes of activation and decision. In addition, the claim of NAM that higher-order information only acts to bias decisions among acoustic-phonetic patterns is worthy of more intense study. Although NAM provides a basic framework for interpreting neighborhood structural effects, it is neutral with respect to a number of important issues, some of which I will now discuss.

The Nature of the Similarity of Speech Sounds. Throughout much of the preceding exploration of neighborhood structure, a relatively simple means of determining neighborhood membership was employed, namely, one phoneme substitutions, additions, and deletions. Although it was demonstrated that a more sophisticated algorithm for computing similarity relations could be devised based on confusion matrices, this approach has limited applicability. In particular, it is strictly only valid to use the confusion matrix data to predict identification performance for words presented under the same conditions in which the confusions were obtained. This approach proved useful for predicting identification performance but was not generalizable to the prediction of similarity relationships when stimuli were presented in the absence of noise. Thus, future work on similarity neighborhood structure will depend on developing more sophisticated methods for determining the similarity of speech sounds (see Carroll & Wish, 1974; Mermelstein, 1976; Shepard, 1972; Wish & Carroll, 1974).

One possible means of deriving more sophisticated similarity measures is to employ multidimensional scaling techniques (see Carroll & Wish, 1974; Wish & Carroll, 1974, Shepard, 1972). Pairwise comparisons of CV and VC syllables

can be used to obtain similarity judgments. Multidimensional scaling solutions can then be derived in order to determine the underlying variables contributing to these similarity judgments. These variables can then be employed in algorithms for computing similarity neighborhoods of words. Such an approach may provide more accurate estimates of similarity neighborhood structure and may enable the development of neighborhood probability rules for predicting accuracy levels and reaction times obtained in tasks in which the stimulus word is not degraded by noise.

Thus, a crucial question in the present research project concerns the perceptual dimensions relevant to the activation of acoustic-phonetic patterns in memory. It is hoped that continuing the present line of research will aid in uncovering the relevant perceptual dimensions of similarity and relating these to the effects of neighborhood structure on auditory word recognition.

The Nature of the Units of Representation. An important question related to the issue of the perceptual dimensions of the space in which acoustic-phonetic patterns are activated concerns the precise nature of the processing units in NAM (see Pisoni & Luce, 1986). Thus far, the model has been more or less neutral with respect to the units of representation and processing. NAM assumes activation of acoustic-phonetic "patterns," although a precise specification of what constitutes a pattern is left unspecified. Indeed, there may be no need to postulate units of patterns per se. As in Nusbaum's phonetic refinement theory (Pisoni et al., 1985), the acoustic-phonetic patterns may simply be a byproduct of constraint satisfaction in a multidimensional space.

The nature of the word decision units in NAM may not be so easily dismissed. It has been heretofore assumed that these correspond to words or word-like units, although this assumption may or may not be correct. Further research on longer word stimuli should enable more precise specification of the nature of these units. The word decision units may actually correspond to morphemes (Fowler, Napps, & Feldman, 1985; Morton, 1969; Napps, 1985) or to syllable units, although at present the postulation of either of these units in perceptual processing is controversial (see Jakimik & Hunnicutt, 1981; Cutler, Mehler, Norris, & Sequi, 1983). If the word decision units actually correspond to units below the word, it may therefore be necessary to build further connections among the decision units to explain syllable-level effects on the auditory word recognition (see Cole, 1983). In short, further research is necessary to define the precise nature of the word decision units.

Effects of Syllable Stress. On a related issue, it will be important to specify the role of lexical stress in NAM. Specifically, it is important to determine the nature of the similarity neighborhoods for stressed and unstressed syllables. Given that the stressed syllable may be the bootstrap for the recognition system (Grosjean & Gee, 1986), it becomes incumbent on the model to describe in precise terms the differential effects that syllable stress may have on similarity neighborhood structure. Indeed, the overall effects of similarity neighborhood structure may vary as a function of stress, in which case it may be necessary to specify in more detail the effects of the number and nature of words in a similarity neighborhood as a function of syllable stress. Obviously, these questions await further experimentation and modeling.

CONCLUSION

The structure of the mental lexicon has been a long-standing and important issue. To date, however, little attention has been directed to specifying the structural organization of the acoustic-phonetic patterns in memory that are used to gain access to the mental lexicon. The present investigation serves as an initial attempt at characterizing that structure and its effects on auditory word recognition. The picture that begins to emerge from the results reported in previous chapters is one of a perceptual and cognitive system optimized for the recognition of words under a variety of circumstances. This optimization is achieved by a simultaneous activation of alternatives based on the stimulus input, and by a sophisticated system that attempts to maximize decisions among these alternatives. The fact that the auditory word recognition system is capable of considering numerous alternatives in parallel helps to assure the best performance of the system in the face of stimulus input that is often impoverished, degraded, or poorly specified. However, as the previous experiments have demonstrated, this optimization is not without its processing costs. Both the number and nature of words activated by the stimulus input affect not only the accuracy of word recognition, but also the time required to decide among the activated candidates. Nevertheless, such processing costs subserve the ultimate goal of the human auditory system, namely to maximize the speed and accuracy with which words are recognized in real-time. In short, the study of the structural organization of the neighborhoods of words in the mental lexicon has provided deeper insights into one important aspect of the fundamentally and uniquely human capacity to communicate with the spoken word.

REFERENCES

- Anderson, D. C. (1962). The number and nature of alternatives as an index of intelligibility. Unpublished doctoral dissertation, Ohio State University.
- Andrews, S. (1982). Phonological recoding: Is the regularity effect consistent? Memory & Cognition, 10, 565-575.
- 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.
- Balota, D. A., & Chumbley, J. I. (1985). The locus of word-frequency effects in the pronunciation task: Lexical access and/or production frequency? Journal of Verbal Learning and Verbal Behavior, 24, 89-106.
- Becker, C. A. (1976). Allocation of attention during visual word recognition. Journal of Experimental Psychology: Human Perception and Performance, 2, 556-566.
- Bernacki, B. (1981). WAVMOD: A program to modify digital waveforms. (Research on Speech Perception Progress Report No. 7). Bloomington, Ind: Speech Research Laboratory, Psychology Department, Indiana University.
- Bond, Z. S., & Garnes, S. (1980). Misperceptions of fluent speech. In R. A. Cole (Ed.), Perception and production of fluent speech. Hillsdale, N. J.: Erlbaum.
- Bond, Z. S., & Robey, R. R. (1983). The phonetic structure of errors in the perception of fluent speech. In N. J. Lass (Ed.), Speech and language: Advances in basic research and practice, Vol. 9. New York: Academic.
- Broadbent, D. E. (1967). Word-frequency effect and response bias. Psychological Review, 74, 1-15.
- Carroll, J. B. (1970). An alternative to Juillard's Usage Coefficient for lexical frequencies, and a proposal for a Standard Frequency Index (SFI). Computer Studies in the Humanities and Verbal Behavior, 3, 61-65.
- Carroll, J., & White, M. (1973a). Age-of-acquisition norms for 220 picturable nouns. Journal of Verbal Learning and Verbal Behavior, 12, 563-576.
- Carroll, J., & White, M. (1973b). Word frequency and age of acquisition as determiners of picture naming latency. Quarterly Journal of Experimental Psychology, 25, 85-95.
- Carroll, D., & Wish, M. (1974). Multidimensional perceptual models and measurement methods. In E. C. Carterette & M. P. Friedman (Eds.), Handbook of Perception. New York: Academic.
- Catlin, J. (1969). On the word-frequency effect. Psychological Review, 76, 504-506.

- Church, K. W. (1983). Phrase-structure parsing: A method for taking advantage of allophonic constraints. Bloomington, Ind.: Indiana University Linguistics Club.
- Cole, R. A. (1973). Listening for mispronunciations: A measure of what we hear during speech. Perception & Psychophysics, 13, 153-156.
- Coltheart, M., Develaar, E., Jonasson, J. T., & Besner, D. (1976). Access to the internal lexicon. In S. Dornic (Ed.), Attention and performance VI. Hillsdale, N. J.: Erlbaum.
- Cutler, A., Mehler, J., Norris, D., & Segui, J. (1983). A language-specific comprehension strategy. Nature, 304, 159-160.
- Dewey, G. (1970). Relative frequency of English spellings. New York: Teachers College Press.
- Dobbs, A. R., Friedman, A., & Lloyd, J. (1985). Frequency effects in lexical decisions: A test of the verification model. Journal of Experimental Psychology: Human Perception and Performance, 11, 81-92.
- Elman, J. L., & McClelland, J. L. (1986). Exploiting lawful variability in the speech waveform. In J. S. Perkell & D. H. Klatt (Eds.), Invariance and variability in speech processing, (pp. 360-385). Hillsdale, N.J.: Erlbaum.
- Forster, K. I. (1976). Accessing the mental lexicon. In R. J. Wales & E. Walker (Eds.), New approaches to language mechanisms. Amsterdam: North Holland.
- Forster, K. I. (1979). Levels of processing and the structure of the language processor. In W. E. Cooper & E. C. T. Walker (Eds.), Sentence processing: Psycholinguistic studies presented to Merrill Garrett. Hillsdale, N. J.: Erlbaum.
- Forster, K. I. (1981). Frequency blocking and lexical access: One mental lexicon or two? Journal of Verbal Learning and Verbal Behavior, 20, 190-203.
- Forster, K. I., & Chambers, I. M. (1973). Lexical access and naming time. Journal of Verbal Learning and Verbal Behavior, 12, 627-635.
- Fowler, C. A., Napps, S. E., & Feldman, L. (1985). Relations among regular and irregular morphologically related words in the lexicon as revealed by repetition priming. Memory & Cognition, 13, 241-255.
- Frederiksen, J. R., & Kroll, J. F. (1976). Spelling and sound: Approaches to the internal lexicon. Journal of Experimental Psychology: Human Perception and Performance, 2, 361-379.
- Gernsbacher, M. A. (1984). Resolving 20 years of inconsistent interactions between lexical familiarity and orthography, concreteness, and polysemy. Journal of Experimental Psychology: General, 113, 256-281.
- Glanzer, M., & Bowles, N. (1976). Analysis of the word-frequency effect in recognition memory. Journal of Experimental Psychology: Human Learning and Memory, 2, 21-31.

- Glanzer, M., & Ehrenreich, S. L. (1979). Structure and search of the internal lexicon. Journal of Verbal Learning and Verbal Behavior, 18, 381-398.
- Goldiamond, I., & Hawkins, W. F. (1958). Vexiersversuch: The logarithmic relationship between word-frequency and recognition obtained in the absence of stimulus words. Journal of Experimental Psychology, 56, 457-463.
- Gordon, B. (1983). Lexical access and lexical decision: Mechanisms of frequency sensitivity. Journal of Verbal Learning and Verbal Behavior, 18, 24-44.
- Greenberg, J. H., & Jenkins, J. J. (1964). Studies in the psychological correlates of the sound system of American English. Word, 20, 157-177.
- Grosjean, F. (1985). The recognition of words after their acoustic offsets: Evidence and implications. Perception & Psychophysics, 38, 299-310.
- Grosjean, F., & Gee, J. P. (1986). Another view of spoken word recognition. Cognition, in press.
- Grosjean, F., & Itzler, J. (1984). Can semantic constraint reduce the role of word frequency during spoken-word recognition. Perception & Psychophysics, 22, 180-182.
- Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information. American Psychologist, 39, 1372-1388.
- Havens, L. L., & Foote, W. E. (1963). The effect of competition on visual duration threshold and its independence of stimulus frequency. Journal of Experimental Psychology, 65, 6-11.
- Hood, J. D., & Poole, J. P. (1980). Influence of the speaker and other factors affecting speech intelligibility. Audiology, 19, 434-455.
- Howes, D. H. (1954). On the interpretation of word frequency as a variable affecting speech of recognition. Journal of Experimental Psychology, 48, 106-112.
- Howes, D. H. (1957). On the relation between the intelligibility and frequency of occurrence of English words. Journal of the Acoustical Society of America, 29, 296-305.
- Howes, D. H., & Soloman, R. L. (1951). Visual duration threshold as a function of word probability. Journal of Experimental Psychology, 41, 401-410.
- Jakimik, J., & Hunnicutt, S. (1981). Organizing the lexicon for recognition. Paper presented at the 101st meeting of the Acoustical Society of America, Ottawa, Canada.
- Klatt, D. H. (1968). The structure of confusions in short-term memory. Journal of the Acoustical Society of America.
- Kohonan, T. (1980). Content-addressable memories. New York: Springer-Verlag.

- Kucera, F., & Francis, W. (1967). Computational Analysis of Present Day American English. Providence, R.I.: Brown University Press.
- Landauer, T. K., & Freedman, J. (1968). Information retrieval from long-term memory: Category size and recognition time. Journal of Verbal Learning and Verbal Behavior, 7, 291-295.
- Landauer, T. K., & Streeter, L. A. (1973). Structural differences between common and rare words: Failure of equivalence assumptions for theories of word recognition. Journal of Verbal Learning and Verbal Behavior, 12, 119-131.
- Leventhal, G. Effect of sentence context on word perception. (1973). Journal of Experimental Psychology, 101, 318-323.
- Lieberman, P. (1963). Some effects of semantic and grammatical context on the production and perception of speech. Language and Speech, 6, 172-188.
- Luce, P. A. (1983). Context and frequency effects in spoken word recognition. Paper presented at the 1st Conference on Hoosier Mental Life, West Lafayette, Ind., May, 1983.
- Luce, P. A. (1986a). A computational analysis of uniqueness points in auditory word recognition. Perception & Psychophysics, 39, 155-158.
- Luce, P. A. (1986b). Similarity neighborhoods of words in the mental lexicon: A computational analysis. Manuscript in preparation. Bloomington, Ind.: Speech Research Laboratory, Psychology Department, Indiana University.
- Luce, P. A. (1986c). Similarity neighborhoods and word frequency effects in visual word identification: Sources of facilitation and inhibition. Submitted.
- Luce, P. A., & Carrell, T. D. (1981). Creating and editing waveforms using WAVES. (Research on Speech Perception Progress Report No. 7). Bloomington, Ind.: Speech Research Laboratory, Psychology Department, Indiana University.
- Luce, P. A., & Pisoni, D. B. (1987). Speech perception: Recent trends in research, theory, and applications. In H. Winitz (Ed.), Human communication and its disorders. Norwood, N.J.: Ablex, in press.
- Luce, R. D. (1959). Individual choice behavior. New York: Wiley.
- Marslen-Wilson, W. D. (1984). Function and process in spoken word recognition: A tutorial review. In H. Bouma and D. G. Bouwhuis (Eds.), Attention and performance X: Control of language processes. Hillsdale, N.J.: Erlbaum.
- Marslen-Wilson, W. D. (1986). Parallel processing in spoken word recognition. Cognition, in press.
- Marslen-Wilson, W. D., & Tyler, L. K. (1980). The temporal structure of spoken language understanding. Cognition, 8, 1-71.
- Marslen-Wilson, W. D., & Welsh, A. (1978). Processing interactions and

- lexical access during word recognition in continuous speech. Cognitive Psychology, 10, 29-63.
- Mason, M. (1975). Reading ability and letter search time: Effects of orthographic structure defined by single-letter positional frequency. Journal of Experimental Psychology: General, 104, 146-166.
- Massaro, D. W., Venezky, R. L., & Talyor, G. A. (1979). Orthographic regularity, positional frequency, and visual processing of letter strings. Journal of Experimental Psychology: General, 108, 107-124.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. Psychological Review, 88, 375-407.
- Mermelstein, P. (1976). Distance measures for speech recognition--psychological and instrumental. (Status Report on Speech Research SR-47). New Haven: Haskins Laboratories.
- Miller, G. A. (1951). Language and communication. New York: McGraw-Hill.
- Miller, G. A., & Johnson-Laird, P. N. Language and perception. Cambridge: Harvard University Press, 1976.
- Miller, G. A., Heise, G. A., & Lichten, W. (1951). The intelligibility of speech as a function of the context of the test materials. Journal of Experimental Psychology, 41, 329-335.
- Miller, G. A., & Nicely, P. E. (1955). An analysis of perceptual confusions among some English consonants. Journal of the Acoustical Society of America, 27, 338-352.
- Moore, R. K. (1977). The evaluation and optimization of a basic speech recognizer. Unpublished masters thesis, University of Essex.
- Morton, J. (1969). Interaction of information in word recognition. Psychological Review, 76, 165-178.
- Morton, J. Word recognition. (1979). In J. Morton & J. D. Marshall (Eds.), Psycholinguistics 2: Structures and processes, (pp. 107-156). Cambridge, Mass.: M.I.T. Press.
- Nakatani, L. H. (1969). A confusion-choice stimulus recognition model applied to word recognition. Unpublished doctoral dissertation, University of California, Los Angeles.
- Napps, S. E. (1985). Morphological, semantic, and formal relationships in the organization of the "mental lexicon." Unpublished doctoral dissertation, Dartmouth College, Hanover, New Hampshire.
- Neisser, U. (1967). Cognitive Psychology. New York: Appleton-Century-Crofts.
- Newbigging, P. L. (1961). The perceptual redintegration of frequent and infrequent words. Canadian Journal of Psychology, 15, 123-132.
- Norris, D. (1982). Word recognition: Context effects without priming. Unpublished manuscript, University of Sussex, Brighton, England.

- Nusbaum, H. C., & Dedina, M. J. (1985). The effects of word frequency and subjective familiarity on visual lexical decisions. (Research on Speech Perception Progress Report No. 11). Bloomington, Ind: Speech Research Laboratory, Psychology Department, Indiana University.
- Nusbaum, H. C., Pisoni, D. B., & Davis, C. K. (1984). Sizing up the Hoosier mental lexicon: Measuring the familiarity of 20,000 words. (Research on Speech Perception Progress Report No. 10.). Bloomington, Ind.: Speech Research Laboratory, Psychology Department, Indiana University.
- Nusbaum, H. C., & Slowiaczek, L. M. (1983). An activation model of the cohort theory of auditory word recognition. Paper presented at the Sixteenth Annual Meeting of the Society for Mathematical Psychology, Boulder, Col.
- Paap, K. R., McDonald, J. E., Schvaneveldt, R. W., & Noel, R. W. (1986). Frequency and pronounceability in visually presented naming and lexical-decision tasks. To appear in Attention and performance XII: The psychology of reading. Hillsdale, N.J.: Erlbaum.
- Paap, K. R., Newsome, S. L., McDonald, J. E., & Schvaneveldt, R. W. (1982). An activation-verification model for letter and word recognition: The word superiority effect. Psychological Review, 89, 573-594.
- Pisoni, D. B., Nusbaum, H. C., Luce, P. A., & Slowiaczek, L. M. (1985). Speech perception, word recognition, and the structure of the lexicon. Speech Communication, 4, 75-95.
- Pisoni, D. B., & Luce, P. A. (1986). Acoustic-phonetic representations in word recognition. Cognition, in press.
- Pollack, I., Rubenstein, H., & Decker, L. (1959). Intelligibility of known and unknown message sets. Journal of Acoustical Society of America, 31, 273-259.
- Pollack, I., Rubenstein, H., & Decker, L. (1960). Analysis of incorrect responses to an unknown message set. Journal of the Acoustical Society of America, 32, 340-353.
- Rubenstein, D. E., Garfield, L., & Millikan, J. (1970). Homographic entries in the internal lexicon. Journal of Verbal Learning and Verbal Behavior, 9, 487-494.
- Rubenstein, H., Richter, M. L., & Kay, E. J. (1975). Pronounceability and the visual recognition of nonsense words. Journal of Verbal Learning and Verbal Behavior, 14, 651-657.
- Rumelhart, D. E., & Siple, P. (1974). Process of recognizing tachistoscopically presented words. Psychological Review, 81, 99-118.
- Sankoff, D., & Kruskall, J. B. (1983). Time warps, sting edits, and macromolecules: The theory and practice of sequence comparison. London: Addison-Wesley.
- Savin, H. B. (1963). Word-frequency effect and errors in the perception of speech. Journal of the Acoustical Society of America, 35, 200-206.

- Scarborough, D., Cortese, C., & Scarborough, H. (1977). Frequency and repetition effects in lexical memory. Journal of Experimental Psychology: Human Perception and Performance, 3, 1-17.
- Shepard, R. N. (1972). Psychological representation of speech sounds. In E. E. David & P. B. Denes (Eds.), Human communication: A unified view. New York: McGraw-Hill, pp. 67-113.
- Slowiaczek, L. M., Nusbaum, H. C., & Pisoni, D. B. (1987). Phonological priming in auditory word recognition. Journal of Experimental Psychology: Learning, Memory, and Cognition, 13, in press.
- Smith, E. E. (1978). Theories of semantic memory. In W. K. Estes (Ed.), Handbook of learning and cognitive processes: Linguistic functions in cognition theory, (pp. 1-56). Hillsdale, N.J.: Erlbaum.
- Smith, J. E. K. (1980). Models of identification. In R. Nickerson (Ed.), Attention and Performance VIII. Hillsdale, N.J.: Erlbaum.
- Soloman, R. L., & Postman, L. (1952). Frequency of usage as a determinant of recognition thresholds for words. Journal of Experimental Psychology 43, 195-201.
- Stanners, R. F., Jastrzemski, J. E., & Westbrook, A. (1975). Frequency and visual quality in a word-nonword classification task. Journal of Verbal Learning and Verbal Behavior, 90, 45-50.
- Triesman, M. (1971). On the word frequency effect: Comments on the papers by J. Catlin and L. H. Nakatani. Psychological Review, 78, 420-425.
- Triesman, M. (1978a). A theory of the identification of complex stimuli with an application to word recognition. Psychological Review, 85, 525-570.
- Triesman, M. (1978b). Space or lexicon? The word frequency effect and the error response frequency effect. Journal of Verbal Learning and Verbal Behavior, 17, 37-59.
- Tweedy, J. R., Lapinski, R. M., & Schvaneveldt, R. W. (1977). Semantic context effects on word recognition: Influence of varying the proportion of items presented in an appropriate context. Memory & Cognition, 5, 84-89.
- Wang, M. D., & Bilger, R. C. Consonant confusions in noise: A study of perceptual features. Journal of the Acoustical Society of America, 45, 1248-1266.
- Webster's Seventh Collegiate Dictionary. (1967). Los Angeles: Library Reproduction Service.
- Whaley, C. P. (1978). Word-nonword classification time. Journal of Verbal Learning and Verbal Behavior, 17, 143-154.
- Winer, B. J. (1971). Statistical principles in experimental design. New York: McGraw-Hill.

Wish, M., & Carroll, D. (1974). Applications of individual differences scaling to studies of human perception and judgment. In E. C. Carterette & M. P. Friedman (Eds.), Handbook of Perception. New York: Academic, pp. 449-491.

Zipf, G. K. (1935). The psycho-biology of language. New York: Houghton-Mifflin.