RESEARCH ON SPEECH PERCEPTION

Technical Report No. 8

November 6, 1992

Speech Research Laboratory
Department of Psychology
Indiana University
Bloomington, Indiana 47405
USA

Supported by:

Department of Health and Human Service
U.S. Public Health Service
National Institutes of Health
Research Grant No. DC-00111-16

A COMPUTATIONAL ANALYSIS OF YOUNG CHILDREN'S LEXICONS

John S. Logan

Submitted to the faculty of the Graduate School
In partial fulfillment of the requirements
for the degree
Doctor of Philosophy
in the Department of Psychology
Indiana University
November, 1992

Acknowledgments

A dissertation is never the product of an individual working alone. Instead, it is the product of all those with whom the individual has met or interacted. I am no exception to this truism; I have many people to thank, not only for ideas and advice, but also for encouragement and aid of one kind or another. My dissertation committee — David Pisoni, Linda Smith, Bob Port, and Judy Geirut — gave of their time in an enthusiastic manner. I thank each of them for all their forbearance. David Pisoni has consistently supported me and I owe him a debt of gratitude. Steve Goldinger, Scott Lively, Paul Luce, Jan Charles-Luce, Keith Johnson, Tom Palmeri, Lynn Nygaard, Mitch Sommers, Luis Hernandez, and Mickey Feaster all worked in or near room 186 and gave me invaluable academic and/or technical support mixed with a good dose of friendship. Brian MacWhinney provided the copy of the CHILDES database that I analyzed. Teresa Treat and Dave Horner were essential for moral support. Finally, I thank my colleagues at Carleton University — Adelle Forth, Shelley Parlow, Monique Senechal, Chris Herdman, and Jo-Anne LeFevre. Support for this project was provided by NIH Grant DC-00111-16 to Indiana University, a Grant-in-Aid of Research and a Marie Z. Sinnot Fellowship from Indiana University.

Abstract

A critical issue in understanding the perception and production of spoken language is how information about words is represented and organized in memory. The present investigation examined confusability among words in the lexicons of children between the ages of 1 1/2 and 5 years. Confusability was operationalized by assuming that words are organized in memory in the context of other words known as neighbors. A word qualified as a neighbor if it could be transformed into another word by substituting, adding, or deleting a single phoneme; neighborhood density refers to the number of words related to each other by this relationship. Neighborhood densities were computed for words obtained from a database containing transcripts of child language. Neighborhood density increased significantly in the period between 1 1/2 and 2 years, coinciding with the onset of the naming explosion, a period in which many new words are added to the lexicon. In addition, a consistent relationship was observed between neighborhood density and several other lexical variables in the adult and child lexicons, suggesting that basic principles of lexical organization are similar in children and adults. Finally, because some researchers have proposed that children utilize some form of nonphonemic representation for lexical information, segments were also coded according to broad phonetic categories corresponding to place and manner of articulation. The manner code produced consistently lower neighborhood densities than the place code. Moreover, neighborhood densities were similar for the manner code and the fully-specified phonemic code, suggesting that for the small lexicons of young children, a representational system based on manner class could be a viable alternative to phonemic representation. Overall, the presented investigation demonstrated that a computational approach could be applied productively to several long-standing questions concerning lexical representation and organization in children.

Table of Contents

A Computational Analysis of Young Children's Lexicons

Chapter I. Lexical Representation in Children: An Introduction

Introduction

The goal of the present investigation was to examine how young children represent lexical information and to document the changes that occur in the representations of words during development. The phrase "lexical representation" as used in the present study refers to the representation of a target word's sound pattern rather than other kinds of lexical information, such as a word's meaning or syntactic class. Knowledge of how children represent lexical information has several potential uses. For example, determining the form of lexical representation is critical in understanding how children learn to recognize and produce spoken words. In addition, an understanding of lexical representation in childhood has potential for illuminating characteristics of lexical representation in adulthood. Learning more about lexical representation in children is also essential for understanding how children learn to read, where the child must learn to match the spoken form of a word to its printed form. Finally, knowledge of lexical representation is necessary for the clinical evaluation of children with speech disorders. In short, determining the nature of lexical representation in children has important implications for several areas in language processing and development.

In the present investigation, I examined lexical representation in children from several perspectives. The first goal was to determine the structural characteristics of children's lexicons with respect to the phonological form or sound pattern of target words used by children. The second goal was to characterize the changes in lexical representation that occur during early and middle childhood. A related question was whether the observed structural changes, if present, would be associated with or would be a consequence of adding many new words to the lexicon during development. The final goal was to map out some of the similarities and differences between the lexicons of children and adults.

General Background and Theoretical Overview

Two basic issues associated with understanding lexical representation and organization are considered in the present investigation. The first issue concerns the relationship among words in the lexicon in terms of their sound pattern or phonological structure. The second issue concerns how lexical information is represented in children and adults.

The Neighborhood Similarity Metric

Questions about the organization and representation of information in the lexicon during development require assessment of phonological structure. The measure chosen for use in the present investigation was the neighborhood similarity metric (Greenberg & Jenkins, 1963; Landauer & Streeter, 1973, Sankoff & Kruskal, 1983; Luce, 1986a). The idea behind the neighborhood similarity metric is that many words can be transformed into other words by changing just one phoneme, either through addition, deletion, or substitution. The words that can be obtained via these transformations are termed *neighbors*. For example, neighbors of the word /kæt/ include /hæt/, /ræt/, /mæt/ (one phoneme substitution in initial position), /skæt/ (one phoneme addition in initial position), and /æt/ (one phoneme deletion in initial position). The number of neighbors that a word has determines its *neighborhood* *density*. Using a computerized version of Webster's Pocket Dictionary as a model of the adult lexicon, Luce (1986a) found that some words have many neighbors while other words have relatively few neighbors. Words that have many neighbors are said to be from *high density neighborhoods*, and include words such as "cat," "but," and "milk." Words that have relatively few neighbors are said to be from *low density neighborhoods* , and include words such as "isobar," "emu," and "orca." Luce (1986a) demonstrated the psychological validity of the neighborhood similarity metric in a series of perceptual experiments. Listeners identified words from low density neighborhoods more quickly and accurately than words from high density neighborhoods, all other things being equal. Luce's findings revealed that there are systematic differences in neighborhood structure among the words in the adult lexicon and that these differences in neighborhood structure could account for the performance of subjects tested in perceptual experiments (see also Cluff & Luce, 1990; Goldinger, Luce, & Pisoni, 1989; Luce, Pisoni, & Goldinger, 1990). Thus, Luce established that the neighborhood similarity metric was a useful way to characterize the relationship among words in the adult lexicon.

Luce (1986a) developed the Neighborhood Activation Model (NAM) to account for neighborhood density effects, word frequency and neighborhood frequency effects. NAM utilizes the neighborhood similarity metric, in conjunction with several other mechanisms, to describe how listeners discriminate among different sound patterns. According to NAM, upon presentation of a word, a listener activates a set of words in memory that are defined by the neighborhood similarity metric. Word frequency information then serves to bias the selection of a word from this set of candidates. The success of NAM in accounting for the performance of subjects in word recognition experiments attests to the usefulness of the neighborhood similarity metric in capturing a basic characteristic of how words are organized and represented in the lexicon.

The neighborhood similarity metric has also been used to address questions about the organization of words in children's lexicons. Charles-Luce and Luce (1990) found that words in the lexicons of 5- and 7-year-olds had many fewer lexical neighbors than the same words in the adult lexicon. Based on these results, they argued that the representations of words in children's lexicons could differ substantially from the representations of words in the adult lexicon (cf. Iverson, & Wheeler, 1987; Macken, 1980; and others). Charles-Luce and Luce's findings showed that the neighborhood similarity metric was a useful formalism for elucidating the structural properties of young children's lexicons because it provides a measure of the similarity relations among words in the children's lexicons and how these relations changed between 5 and 7 years. In addition, it offers a convenient way to compare adult and child lexicons.

The present investigation extends these earlier studies in several directions by applying the neighborhood similarity metric to language samples from children between the ages of 18 months and 5 years. The first goal was to see if the trends observed by Charles-Luce and Luce were also present in younger children. A second goal was to determine whether the relationships among lexical variables found in the adult lexicon, such as the positive correlation between word frequency and neighborhood density, were also observed in children's lexicons. A third goal was to use the neighborhood similarity metric to test the plausibility of alternative representations of lexical information to that found in adults.

Nonphonemic Representation

The phonological system in children, especially at early ages, may not correspond to the phonological system in adults (Maxwell, 1984). Whereas adults appear to have a segmental phoneme-based system of lexical organization, it is not clear at the present time what kind of system children utilize. Some researchers have argued that children's lexical representations are more wholistic than adult representations (Peters, 1983), while others have argued that children's lexical representations are essentially the same as adult representations (Smith, 1973). In order to evaluate these possibilities, I examined neighborhood structure using a phoneme-based segmental system and two systems that classified segmental information according to broad phonetic categories based on manner and place of articulation.

Motivation for examining these nonphonemic classification systems came from recent findings in the area of speech recognition by machine. Several researchers have explored the possibility that the classification of utterances using broad phonetic categories could increase the efficiency of word recognition, particularly for words from large lexicons (e.g., Huttenlocher & Zue, 1983). The classification of utterances into broad phonetic categories refers to the categorization of a segment into a relatively small set of categories, such as manner class categories. For example, the phonemes /p/, /t/, and /k/ would each be classified as stops according to a classification system based on manner of articulation. Classifying an utterance on the basis of coarse coding is relatively easy to do and eliminates from further consideration a large number of alternatives in the lexicon. The recognition system is then left with a relatively small number of candidates for further, more detailed analyses (Shipman & Zue, 1982). Efficiency is gained because the system restricts fine-grained acoustic analyses to only those items in the lexicon that fit the broad phonetic categories. The classification of words using broad phonetic classes in speech recognition can be thought to parallel the wholistic lexical representations that several researchers have proposed for young children.

Overview

Chapter II begins with a review what is known about lexical representation in adults. Efforts to understand lexical representation in adults have focused on the size of lexical units (e.g., Cutler, Mehler, Norris, & Segui, 1987), the effects of word frequency (e.g., Forster, 1976; Morton, 1979), the acoustic-phonetic similarity among items in the lexicon (e.g., Luce, 1986a, 1986b; Marslen-Wilson, 1987), and the prosodic characteristics of words (e.g., Cutler & Norris, 1988). Several models of adult spoken word recognition are then reviewed in order to make explicit some of the assumptions researchers have made about lexical representation. In this context, I also briefly review Dell's (1986) recent model of speech production in order to examine how it deals with issues of lexical representation. Next, I examine recent work in linguistics that is relevant to lexical representation. Specifically, the emphasis on economy of representation in linguistic theory, the role of syllable structure, and the relationship between speech production and perception are discussed.

Chapter III describes lexical representation in children. First, I review studies on speech perception and production from the period between infancy and middle childhood (5 years of age) to determine what they reveal about lexical representation (e.g., Jusczyk, in press; Vihman, 1988). I then consider two views of early lexical representation. One position holds that the child's representational system is basically the same as the adult's (Donegan & Stampe, 1978; Smith, 1973). The other position proposes that there are systematic changes in lexical representation during development that lead to the representational system found in the adult (Moskowitz, 1973; Peters, 1983). Next, I review research on beginning reading skills and discuss its implications for lexical representation (Treiman, 1982; Goswami & Bryant, 1990). Finally, I consider the relation between lexical representation and cognitive development (Aslin & Smith, 1988; Kemler-Nelson, 1983).

Chapter IV contains a review of computational analyses of spoken and printed language (e.g., Charles-Luce & Luce, 1990; Cutler & Norris, 1988; Landauer & Streeter, 1973). These studies used a computational approach to gain new insights into how lexical information may be represented and processed. Many of the computational procedures used in the present investigation were derived from the studies reviewed in Chapter IV. Chapter IV also contains a description of the methodology employed in the present investigation. Luce's (1986a) neighborhood similarity metric was used to study the phonological structure of a large corpus of utterances produced by children aged 18 months to 5 years. Using a computerized database, neighborhood densities were computed for each word at different ages for each child using a standard phonemic coding system (i.e., a phonemic transcription as used in Luce, 1986a, and Charles-Luce & Luce, 1990). The results of the present analyses were then compared with the findings of Charles-Luce and Luce, in part to validate the present data, and, in part, to see how the results differ from the findings obtained from 5- and 7-year-olds. Finally, the relationship among several lexical variables, such as neighborhood density and word frequency, were compared in adults and children.

Chapter V describes how the neighborhood similarity metric was used to investigate three nonphonemic representational systems that assigned phonemes to broad categories based on manner of articulation and place of articulation. A third nonphonemic system classified phonemes using a randomly-defined set of categories not based on any known linguistic principles. The nonphonemic representational systems were used to assess the effects on neighborhood density when broad categories that corresponded to gross linguistic features replaced the standard segmental phonemic representational system. Such broad categories reflect the simple representational systems proposed for children yet still allows the unique specification of individual lexical items. Analyses that parallel those described in Chapter IV were carried out to assess the consequences of assuming nonphonemic representation. A comparison of phonemic and nonphonemic systems is also described in Chapter V.

Chapter VI summarizes the findings of the present analyses and discusses the implications for research and theory in the area of spoken word recognition. In addition, I consider some of the limitations of the present investigation and how they might affect future research. I conclude that the computational approach promises to increase our understanding of how lexical information is represented and organized in young children and how the mental lexicon changes during development.

Chapter II. Lexical Representation in Adults

In this chapter, I describe what is known about lexical representation in adults. The form of lexical representation in adults may be thought of as a developmental endpoint that is useful for purposes of comparison with the form of the child's representation of lexical information. Research on lexical representation in adults has concentrated on two major issues, determining the units of representation in the lexicon and understanding the effects of several lexical variables, including word frequency, acoustic-phonetic similarity, and lexical stress. In this chapter, I examine these issues from both psychological and linguistic viewpoints.

Psychological Models of Lexical Representation and Organization in Adults

The organization and representation of lexical information in adults can be viewed from several perspectives. For example, principles of lexical organization include word meanings (Smith, 1978; Smith & Medin, 1983), structural characteristics (Forster, 1976; Luce, 1986), units of representation and processing (Cutler, Mehler, Norris, & Segui, 1987), parts of speech (Bradley, Garrett, & Zurif, 1980), and neurological substrates (Fromkin, 1987; Nadel & Wexler, 1984). In the present investigation, only a subset of these perspectives relevant to the representation of lexical information will be considered, namely, the units associated with lexical representation, and the effects of word frequency, acoustic-phonetic similarity, and lexical stress.

Units of Lexical Representation

In many cases, questions about the units of lexical representation have been phrased in terms of how a listener accesses words from the mental lexicon. Early work by Savin and Bever (1970) suggested that listeners access the lexicon through syllable-sized units and that phonemic information is accessed or derived from the lexicon post-perceptually. Evidence for this viewpoint came from experiments in which subjects monitored nonsense syllables for either an entire syllable or a phoneme. Subjects in these experiments were faster at monitoring for syllables than for phonemes. Savin and Bever concluded from these results that listeners accessed the lexicon through units larger than phonemes and that syllable or word sized units are the primary units of lexical representation. The opposite position was adopted by Foss and Gernsbacher (1983). Based on a series of experiments, Foss and Gernsbacher concluded that sublexical units, such as phonemes, are utilized to develop lexical representations. Similarly, in work with visually presented stimuli, Taft (1979) proposed that affixed words undergo morphological decomposition and are accessed via their stems.

Others investigators have proposed an intermediate position and suggested that access to the lexicon is not carried out exclusively through any one type of perceptual unit (Cutler et al., 1987; McNeill & Lindig, 1973) . Instead, access can occur at several different levels depending on the nature of the task in which the listener is engaged and the type of stimuli presented. Support for the intermediate position was provided by Samuel and Ressler (1986) who demonstrated that subjects' attention is typically focused at the level of word-sized units but, if required, could be focused at the sublexical level as well. Cutler et al. (1987) suggested that in most situations, speech perception proceeds via whole-word units. Only when listeners encounter novel words or low-intelligibility speech does perception proceed via sublexical units to larger units. Consistent with this position, I assume for the present investigation that although words are the basic units of lexical representation, adult listeners do have access to the phonemic structure of both familiar and unfamiliar words.

The Role of Lexical Variables in Spoken Word Recognition and Production

Perhaps the most distinctive characteristic of spoken language processing in adults is how quickly and efficiently lexical information can be accessed (Marslen-Wilson & Welsh, 1978). Such rapid and accurate access during speech perception and production suggests that lexical information is efficiently organized in long-term memory. Hypotheses about lexical representation and organization have been derived from the results of various word recognition experiments and production phenomena. Experiments have shown that three variables are critical in perception and production processes: word frequency, acoustic-phonetic similarity among words in the lexicon, and prosodic features, such as stress, timing, and intonation.

Word frequency. A ubiquitous finding in the area of word recognition is that high frequency words are recognized more quickly and accurately than low frequency words (Howes & Solomon, 1951; Rubenstein, Garfield, Milliken, 1970). Almost all contemporary models of word recognition have incorporated word frequency as a basic principle of how words are organized in long-term memory (e.g., Becker, 1976; Forster, 1976; Morton, 1979). In general, high frequency words are either more easily accessed or have lower thresholds for activation than low frequency words, thus facilitating the recognition process for high frequency words compared to low frequency words. Word frequency effects are also found in speech production experiments. For example, Stemberger and MacWhinney (1986) found that high frequency words were pronounced more accurately than low frequency words. As was the case with models of word recognition, models of speech production must also include mechanisms that deal with word frequency (Dell, 1986). In short, virtually all accounts of lexical processing or representation incorporate mechanisms designed to deal with the effects of word frequency.

Acoustic-phonetic similarity. Words also appear to be organized in long-term memory according to their acoustic-phonetic similarity. As was true for word frequency, the effect of similarity structure is found in both perception and production. Several studies have shown that listeners tested in perceptual identification tasks identify words in the context of other words (e.g., Miller, Heise, & Lichten, 1951; Pollack, Rubenstein & Decker, 1959, 1960). For example, when listeners are presented with a target word and make an identification error, the error tends to have a sound pattern that is similar to the target word (Pollack, Rubenstein & Decker, 1960). Similarly, production errors also tend to create words that are phonologically similar to the intended word (Dell, 1986; Stemberger, 1991).

Several accounts of how the sound patterns of words affect speech perception and production have been proposed (e.g., Klatt, 1979; Marslen-Wilson, 1987; McClelland & Elman, 1986; Treisman, 1978a, b; Luce, 1986; Klatt, 1989; Dell, 1986; Stemberger, 1991) . First, I discuss several models that deal with acoustic-phonetic similarity in speech perception and then I consider several models that deal with acoustic-phonetic similarity in speech production.

<u>Cohort Theory</u>. One of the most influential accounts of spoken word recognition is Marslen-Wilson's cohort theory (Marslen-Wilson & Welsh, 1978; Marslen-Wilson, 1987). Cohort theory was designed to deal with the integration of sensory and contextual/semantic information within the temporal constraints of spoken language processing. According Marslen-Wilson, when the initial portion of a word is heard, the set of all words in the lexicon consistent with the word-initial acoustic-phonetic information is activated. This is referred to as a 'word initial cohort'. Additional sensory information coupled with contextual information serves to reduce the size of the word initial cohort until only one word remains. At that point, the word is said to be recognized. Although the cohort model does not explicitly address lexical representation, it assumes that words are organized in the

lexicon in terms of their acoustic-phonetic properties. Furthermore, it implicitly assumes that these properties are sequentially organized as phoneme-like units. Although initial versions of the cohort model did not address word frequency effects (Marslen-Wilson & Welsh, 1978), a more recent version does so. Specifically, Marslen-Wilson (1987) proposed that frequency increases the activation of high frequency words in the word-initial cohort so that these words are more likely to remain as candidates during the selection process.

Trace. The Trace model is an interactive activation model of word recognition (McClelland & Elman, 1986). It is similar to Cohort Theory in that both were designed to deal with the integration of sensory and contextual information in spoken word recognition. In the interactive architecture of the Trace model, information about words exists at multiple levels corresponding to words, phonemes, and features, each of which influences the other levels. Thus, lexical information can influence phoneme-level decisions and phonemic information can influence lexical decisions. In the Trace model, lexical and sublexical information, including phonemes and features, are represented explicitly as nodes in an interconnected network.

Lexical Access from Spectra (LAFS). Klatt's (1979) LAFS model differs from Cohort Theory and the Trace model in that the lexicon is addressed directly from spectral information contained in the speech signal. In LAFS, there is no intermediate level of symbolic representation between the acoustic information in spectra and the lexical information. Instead, LAFS uses a network of context-sensitive spectral templates to directly access the lexicon. Phonetic information in LAFS is derived post-lexically via a separate mechanism. Thus, lexical representation in LAFS differs from that assumed in Cohort Theory and Trace. Phonemic information is represented neither explicitly or implicitly in LAFS but, instead, exists only as a by-product of information derived from lexical sources.

Partial Identification Theory. In Treisman's (1978a, 1978b) partial identification theory, words are represented in an acoustic space and a stimulus or target word is selected from a location in this perceptual space based on its acoustic characteristics. Frequency effects occur when the discriminability of words is poor and a subset of locations in lexical space is selected. That is, frequency serves to aid identification when acoustic information is incomplete or unreliable. Treisman called words that shared a common acoustic structure 'neighbors'. Discriminability among words, according to Treisman, depends upon whether a word is located in a dense or rarefied subvolume of acoustic space, where "dense" and "rarefied" refer to subvolumes containing many or few neighbors, respectively. The exact characteristics of the dimensions underlying the acoustic space are unspecified. However, a level of representation corresponding to words, phonemes, and features is consistent with partial identification theory.

Neighborhood Activation Model. Luce's (1986) Neighborhood Activation Model (NAM) is similar in several respects to Treisman's partial identification theory. Upon presentation of a spoken word to a listener, a subset of acoustic-phonetic space corresponding to characteristics of the word is activated. As in partial identification theory, the characteristics of the dimensions underlying the acoustic-phonetic space are unspecified (see the next section for further discussion of the proposed characteristics of acoustic-phonetic space). Once a subvolume of acoustic-phonetic space is activated, higher-level knowledge, including word frequency, serves to bias the selection of a particular word in the activated subvolume, particularly when the stimulus is degraded. Consistent with Treisman's terminology, Luce (1986) refers to a subvolume of acoustic-phonetic space as a lexical neighborhood. The number of words that share similar acoustic-phonetic characteristics determines neighborhood density. Words that are phonetically similar to other words in the lexicon are located in high density neighborhoods. Words that are not similar to many other words in the lexicon are located in low density neighborhoods. Luce obtained support for NAM in a series of perceptual experiments that showed word recognition performance depended on both neighborhood density and word frequency. Specifically, Luce determined that three characteristics could be used to predict how easily words would be recognized in a perceptual task: the word's frequency of occurrence, the density of the word's neighborhood, and the frequency of the word's neighbors. Words that were easy to recognize were high frequency words that had a relatively few low-frequency neighbors, whereas words that were hard to recognize were low frequency words that had high frequency neighbors.

Luce's (1986a) findings represent an important extension of Treisman's (1978a, 1978b) partial identification theory. Luce obtained effects similar to those obtained by Treisman despite the fact that Luce used a much larger and more variable set of words. Moreover, Luce computed similarity on the basis of two independent measures, empirically obtained phoneme confusions and neighborhood densities calculated from a computerized database (Luce, 1986). Thus, the major claims of partial identification theory were validated and extended in NAM.

In summary, several models of spoken word recognition have been proposed to deal with the effects of acoustic-phonetic similarity among words. Only one model, the Trace model, explicitly assumes a level of representation that corresponds to phonemes. Cohort Theory, Partial Identification Theory, and NAM each assume that acoustic-phonetic similarity plays an important role in spoken word recognition but do not commit to an explicit phonemic level of representation to account for the effects of acoustic-phonetic similarity among words. LAFS explicitly avoids using an intermediate level of representation corresponding to phonemes, and instead directly matches spectral information to lexical items stored in memory.

The nature of acoustic-phonetic representations. Models of spoken word recognition utilize a variety of representations for dealing with the effects of acoustic-phonetic similarity. However, most do not describe in explicit terms the nature of acoustic-phonetic representation. For example, Treisman (1978a, p. 528) stated that partial identification theory posits a store of "...word images distributed in a multidimensional acoustic space whose dimensions are acoustic parameters. (No attempt will be made here to define these more exactly.)." Similarly, Luce (1986, p. 39) noted that NAM "...is neutral with respect to the dimensions of the [acoustic-phonetic] space." Since Treisman (1978a, 1978b) dealt with error data collected in experiments using closed-response sets, he was able to more easily skirt the issue of representation. However, in order for Luce (1986) to carry out the experiments on which NAM was based, it was necessary to select stimuli using some set of perceptual dimensions that permitted comparison of words based on their physical similarity to each other. In order to operationally define similarity relations among words in the space, Luce (1986) relied on a phonemic coding scheme in which a word was considered a neighbor of another word if the first word could be transformed into the second word by means of substituting, adding, or deleting a single phoneme (cf. Greenberg & Jenkins, 1963). That is, the neighborhood densities computed by Luce were based on phonemic transcriptions of words in a computerized version of Webster's Pocket Dictionary. If one makes the same set of assumptions that Luce did, that an isomorphic relationship exists between a phonemic representation and an acoustic representation (cf. Palmer, 1978), then it is immaterial whether one describes the similarity relations among words in the lexicon in terms of phonemic or acoustic-phonetic similarity. Moreover, the use of the adjective "acoustic-phonetic" when describing the lexical space implies that neither the acoustic or phonetic aspect of the representation is afforded priority. However, one may choose to use a symbolic, phoneme-based representational system for practical reasons when computing similarity among words because it is simply easier to deal with the phonemes that comprise a word than to deal with a highly detailed parametric description of the acoustic characteristics of a word.

Lexical Stress

Lexical stress is primarily a function of the loudness, length, and pitch of the syllables that comprise spoken English (Crystal, 1991; Fry, 1955). Based on several kinds of evidence, some researchers have argued that listeners can use lexical stress to gain access to words in the lexicon. For example, Cutler and Carter (1987) estimated that about 80% of the words a listener is likely to encounter in English begin with stressed syllables. Thus, lexical stress is a reliable indicator of word boundaries. Cutler and Norris (1988) have proposed that whenever listeners encounter a stressed syllable, it signals the onset of a lexical candidate. They obtained support for their hypothesis in several experiments in which listeners were required to detect a real word embedded in a bisyllabic nonword. They found that listeners detected the real word faster when its boundary did not conflict with the syllable boundary established by the stressed syllable.

Grosjean and Gee (1987) cite additional evidence that lexical stress is an important factor in spoken word recognition and lexical organization. For example, mispronunciations are detected faster in stressed syllables than in unstressed syllables, the same phoneme is detected faster in a stressed syllable than in an unstressed syllable, and slips of the ear are more likely to occur in unstressed syllables. Based on this kind of evidence, Grosjean and Gee (1987) argued that the absence of mechanisms to deal with stress in current models of spoken word recognition is a serious oversight because stress is an important source of knowledge that contributes to the comprehension of spoken language, particularly fluent, connected speech.

A Model of Speech Production

Current models of speech production have also made a number of assumptions about lexical representation and organization (Dell, 1986; Stemberger, 1985). For example, Dell (1986) has proposed a detailed model of speech production that is designed to account for a wide range of language production phenomena, especially speech errors. In his model, words are assembled for production as a series of constituent morphemes containing sequences of phonemes organized in a syllabic structure. Dell implemented his model using an interactive activation architecture in which information from different linguistic levels influences other levels. Phonetic similarity among words and the interaction between phonological and morphological levels accounts for a wide variety of sound errors in production. Frequency effects, where more frequent phonemes tend to replace infrequent phonemes, are a direct consequence of positive feedback between information at the syllabic level and information at the phonological level. Thus, Dell's model of speech production incorporates the same units of lexical representation (phonemes, syllables, and words) and attempts to account for similar phenomena (word frequency and phonological similarity) as the models of word recognition reviewed earlier.

Summary

The work that cognitive psychologists have done to specify the nature of the lexicon in adults may be summarized as follows: Words are represented in the lexicon at both sublexical and lexical (whole-word) levels (McNeill & Lindig, 1973). Sublexical levels include phonemic representations, syllabic representations, and possibly affix and stem representations. In addition, the organization of the lexicon reflects the internal structure of words and the similarity relations that exist among words in the lexicon (Cutler & Norris, 1988; Luce, 1986a; Treisman, 1978a, 1978b). Furthermore, words in the lexicon are also organized according to frequency of occurrence. In short, the picture of the adult lexicon that begins to emerge from this work is that of a tightly integrated entity from which information about words or their constituent parts can be quickly and efficiently accessed. Information about a word can be used to inform the listener about the sublexical components that comprise the word. Conversely, sublexical units can be compiled together to inform the listener about the identity of a word. Although adult listeners generally rely on whole-word units for recognition, they can access the lexicon through sublexical units if conditions are unfavorable for whole-word access.

Linguistic Models of Lexical Representation in Adults

An adequate description of lexical representation is a fundamental goal of linguistics. In this section I focus on three topics relevant to this goal: the underspecification of phonemes, the syllable structure of words, and the relationship between language production and perception. Findings in each of these areas has important consequences for understanding lexical representation in adults, as well as raising issues critical to understanding the development of lexical representation in children.

Lexical Representation in Linguistic Models

Linguistic models tends to minimize the contents of the mental lexicon and instead emphasize the use of rule systems designed to generate the surface form of language. Butterworth (1983) refers to this emphasis as 'economy of representation'. In contrast, 'economy of processing' refers to a tendency to minimize the use of processing, which is equated with the use of rules, and instead emphasizes precompiled knowledge. Historically, linguistic models have favored economy of representation over economy of processing. According to Bloomfield (1933), the lexicon consists of a list of basic elements, known as morphemes, that link phonological forms to meaning (p. 138). Bloomfield described the lexicon as containing the basic underlying forms from which the regular form of words are generated, as well as containing a listing of irregular forms (p. 274). Thus, the forms of the verb "turn" – "turns", "turned", "turning" – are not contained in the lexicon but instead are generated by rules, whereas the forms of the verb "eat" – "eat", "ate", "eating" – are listed in full. Subsequent linguistic models of syntax and phonology have retained this emphasis on maintaining a minimal amount of information in the mental lexicon (e.g., Chomsky, 1965; Chomsky & Halle, 1968).

Because linguistics subscribes to a model in which a potentially infinite number of utterances can be generated from a fixed set of lexical items via the application of syntactic and phonological rules, linguistic models of the lexicon tend to focus on how articulatory information for production is represented in memory (e.g., Halle, 1985, 1986). This information is typically in the form of a feature-based representational system in which individual phonemes per se do not exist (Jakobson, Fant, & Halle, 1952). Instead, phonemes are derived from an inventory of features by means of phonological rules (Chomsky & Halle, 1968).

A recent version of this approach to phonology is underspecification theory, proposed by Archangeli (1984, 1988) and others. Archangeli (1988, p. 203) states that the major goal of underspecification theory is to develop a theory of phonology in which "...only idiosyncratic information is included in the most basic representations and all predictable information is encoded in rules." An example of underspecification is that all vowels in English are voiced. Therefore, the feature [+voice] is predictable and need not be specified in the lexicon but can be filled in later during processing. Underspecification theory reflects a trend in modern phonology in which redundancies in feature descriptions are exploited to create the most compact and economical representation system possible.[1] Moreover, several theorists have proposed that principles of underspecification theory can be used to account for the relationship between lexical representation and spoken language in children (e.g., Iverson & Wheeler, 1987; Stemberger & Stoel-Gammon, 1991). Despite its potential use as an explanatory device for understanding lexical representation in young children, the specific claims of underspecification as a linguistic theory fall outside the scope of the present investigation. Further use of the term 'underspecification' will be limited to describing a representational system in which information is not fully specified.

Syllable structure

Another relevant characteristic of current models of lexical representation in linguistics is an emphasis on the critical role of the syllable in organizing sequences of phonemes in production (e.g., Clements and Keyser, 1983; Halle, 1986). Kaye (1989) notes that the notion of syllabic structure is crucial in accounting for a wide variety of phonological phenomena at both segmental and suprasegmental levels. Treiman (Treiman, 1988; Treiman & Danis, 1988) has amassed considerable behavioral evidence to support the claim that syllables can be broken down into an initial structure (onset), a vocalic structure (rime), and a final syllable structure (coda). Thus, recent claims by linguists regarding the hierarchical form of lexical representations have received empirical support from behavioral studies with real listeners and speakers. In the present investigation syllable structure is relevant because sensitivity to syllable structure appears to be a particularly salient characteristic of child language.

Perception and Production in Linguistic Models

As noted in the previous section, models of lexical representation developed by linguists generally deal with production phenomena. Emmorey and Fromkin (1988, p. 127) point out that accounting for perception in a system designed to account for language production is problematic. How is a phonological representation computed from the acoustic signal? Current theories of phonology, such as Archangeli (1988), and Sagey (1986), retain Chomsky and Halle's goal of accounting for production phenomena and have not directly addressed the role of phonology in speech perception.

The relationship between speech perception and production in linguistics is closely connected with the issue of whether there is one generic lexicon that is used for both perception and production, or whether there are two separate lexicons, one used in perception and one used in production (cf. Butterworth, 1983). If one adopts the view that there is only one lexicon for perception and production, phonological models of production are, ipso facto, models of perception. However, if one adopts the opposing view that there are separate lexicons for perception and production, then accounts of lexical representation in terms of production may be inappropriate in accounting for perception phenomena. For pragmatic reasons, I assume in the present investigation that there is one lexicon for both perception and production. This issue is related to the question of whether there are separate lexicons for the visual and auditory forms of words. Recent data indicate no significant difference between familiarity ratings assigned to words presented visually or auditorily

---

[1] Some critics have argued that underspecification theory places an undue emphasis on economy of representation at the expense of economy of processing (Mohanan, in press; cited in Stemberger, 1991). However, preliminary evidence in favor of underspecification has been obtained in studies by Stemberger (1991) and Lahiri and Marslen-Wilson (1991) (for production and perception phenomena, respectively).

(Pisoni & Garber, 1990; Garber & Pisoni, 1991). To the extent that questions regarding the separation of the lexicon into visual/auditory components and perception/production components are related, then these data suggest the existence of a common lexicon. However, these issues remain unresolved at the present time.

Summary

Recent theoretical work in linguistics has emphasized the underspecification of features and the importance of the syllable in organizing phonological structure. While models of phonology provide an account of lexical representation in terms of production phenomena, little research has been done to determine the applicability of such models to perceptual phenomena. Each of these issues has implication for lexical representation that will be discussed in more detail in Chapter III.

Summary

Behavioral evidence from adults has provided several important insights about the representation and organization of information in the mental lexicon. First, information about words exists at featural, phonemic, syllabic, and lexical levels. Second, word frequency information is retained and encoded in some way in the lexicon. Third, the acoustic-phonetic representation of words is important in accounting for several different kinds of perceptual and production phenomena. Finally, phonemes appear to exist in an underspecified form in the lexicon. Several models from psychology and linguistics were described that attempted to incorporate one or more of these findings. No one model deals with all the behavioral evidence cited; each accounts for a specific subset of phenomena. In short, lexical representation in the adult appears to reflect a complex, multi-tiered system in which the sound pattern of a word is a central component.

Chapter III. Lexical Representation in Children

The findings reviewed in Chapter II indicate that adults possess detailed and systematically organized information about the sound structure of spoken words. Moreover, the existence of basic organizing principles based on the sound patterns of words and the role of word frequency is uncontroversial even though there is a lack of unanimous agreement about the exact nature of the representation of words in the adult mental lexicon. In contrast, lexical representation in children is not as well-understood nor is there any consistent agreement among theorists regarding the nature of change in lexical representation during childhood.

There are several reasons for this situation. First, young children are simply not as easy to study as adults in experimental settings. This limits the availability of experimental evidence necessary for the development and testing of detailed theories of lexical representation and organization. A second, related reason for the lack of consensus among child language theorists is that studies of adults tend to focus on recognition abilities whereas studies of children tend to focus on production abilities, making it difficult to draw parallels between the two groups. The major reason for this disparity is because perceptual experiments require the subject to listen to a set of stimuli and make some type of behavioral response, a demanding task for a 2-year-old. In most production studies, the child simply talks and the experimenter records or transcribes what the child says. Although much has been learned about lexical representation in children from speech production studies, the optimal situation would be to have additional perception data available as well. While acknowledging that these problems place limits on developing a theory to explain lexical representation in children , it is nonetheless necessary to consider all the relevant data in this area. This is the goal of Chapter III. First, I consider some basic issues concerning lexical representation and organization in childhood. This is followed by a summary of the developmental changes that occur in perception and production during childhood. Next, a theoretical framework is described that accounts for several phenomena in child language, as well as addressing the issue of lexical representation. Finally, the influence of nonlinguistic cognitive factors on lexical development, such as memory and selective attention, are considered.

Basic Issues Concerning Lexical Representation in Children

The basic question concerning lexical representation in children is whether or not the child's mental lexicon is the same as the adult's mental lexicon. If the child's lexicon is organized according to the same principles as the adult's, then one would expect to find evidence of the same kinds of phenomena observed in adults. For example, one would expect to find evidence of lexical neighborhoods and effects of word frequency in the child's language. Similarly, linguistic principles such as underspecification and syllabic structure would also be expected to account for some of the phenomena associated with the child's language. The opposite position is that the child's lexicon may be organized in a way that is fundamentally different from the adult lexicon. According to this position, the characteristics that typify the adult lexicon are not observed in the child's lexicon. Between these two extremes are a wide spectrum of possibilities in which some characteristics are the same or similar in adult and child lexicons while other characteristics are different.

A number of questions follow from each of these positions. If the child's representational system is similar to the adult system, then it is necessary to account for why children have various difficulties in speech perception and production. If the child's representational system differs from the adult system then it is necessary to account for the ways in which it differs from the adult system and how it eventually develops into the adult system. Individual differences are another important issue in lexical organization: If children have non-adult lexical representations, does each child have a unique representational system or is there a uniform non-adult representational system across children? These issues are addressed below, beginning with a brief summary of what is known about lexical representation in infancy followed by what is known about representation in early childhood.

Language Perception and Production in Infancy and Childhood

Speech Perception in Children

Infant speech perception – Birth to 12 months. In early infancy, defined here as the period from birth to 6 months of age, the ability to perceive speech appears to be much more advanced than the ability to produce speech. Infants in this age range can discriminate a wide variety of phonemes from virtually any language (e.g., Aslin, Pisoni, Hennessy, & Perey, 1981; Best, McRoberts, & Sithole, 1988; Lasky, Syrdal-Lasky, & Klein, 1975; Trehub, 1976; Streeter, 1976; Werker, Gilbert, Humphrey, & Tees, 1981; Werker & Tees, 1984). (For a review of infants' capabilities see Aslin, Pisoni, & Jusczyk, 1983; Eimas, Miller, & Jusczyk, 1987; Werker, 1989) Moreover, there is some evidence suggesting that infants can use perceptual units other than phonemes to segment the incoming speech signal. For example, Bertoncini, Jusczyk, and colleagues (Bertoncini, Bijeljac-Bibic, Jusczyk, Kennedy, & Mehler, 1988; Bertoncini & Mehler, 1981; Jusczyk & Derrah, 1987; Jusczyk, 1982) argue that the syllable may be the primary unit of speech perception in early infancy. In addition, infants also appear to display perceptual constancy for consonants and vowels produced by multiple talkers, and for speech varying in pitch contour from the same talker (Kuhl, 1979, 1983).

Although it is tempting to conclude from these results that infants have at least a level of lexical representation corresponding to phonemes, it is not clear whether infants, especially in early infancy, discriminate phonemes on the basis of phonemic categories or auditory features (Aslin et al., 1983; Jusczyk, 1982; Studdert-Kennedy, 1986). Moreover, infant abilities have only been tested using simple stimuli in relatively undemanding tasks (Locke, 1986). In the context of more complex stimuli, such as words produced in a sentence where the talker reduces or deletes segments, speech perception performance may be substantially different (Locke, 1986). If infant performance was substantially reduced under more demanding situations, the argument for adult-like linguistic representations in early infancy would be weakened. If infants had access to linguistic representations, they presumably would be able to deal with various types of phonological processes, such as segment deletion or reduction, and other situations in which only partial information is available in the signal.

Infants tested in perceptual tasks also appear to be sensitive to the prosodic features of speech, such as intonation and stress (Fernald, 1985; Stern, Spieker, Barnett, & MacKain, 1982). Some researchers have suggested that infants use the pattern of stress and intonation to help segment the speech signal into linguistic categories including phrases, words, and syllables (e.g., Jusczyk, in press; Trehub, Thorpe, & Morrongiello, 1987). Jusczyk (in press) suggests that these suprasegmental characteristics serve to focus the infant's attention to where linguistically relevant 'breaks' in the continuous speech signal are located. In summary, before 6 months of age, infants are capable of discriminating a wide variety of speech sounds from different languages and from different talkers, although the basis for these abilities is not clear. Additional evidence suggests that young infants are also beginning to segment fluent, continuous speech into adult-like categories.

Jusczyk and his colleagues have also demonstrated that English-learning 9-month-olds are sensitive to the acoustic cues associated with the segmentation of English utterances into phrasal units (Jusczyk, Hirsh-Pasek, Kemler Nelson, Kennedy, Woodward, & Piowz, 1992). Specifically, Jusczyk et al. found that 9 month-old infants, but not 6-month-olds, preferred to listen to listen to speech containing breaks that coincided with naturally-occurring boundaries in English more than speech containing breaks that did not coincide with natural boundaries. Locations of phrasal boundaries appeared to be marked by several prosodic characteristics, including changes in fundamental frequency and syllable duration. These findings are important because they support the idea that infants analyze large perceptual objects into smaller, linguistically relevant units as they begin to develop a linguistic system modelled after the adult system.

As infants become older, they begin to show the influence of the language-learning environment in which they are being raised. By 12 months of age, infants no longer show the same ability to discriminate speech sounds from diverse languages as they did when they were younger (Werker & Tees, 1984). Approximately coincident with the infant's reduced ability to discriminate nonnative phonetic contrasts is evidence for a nascent phonological system. Jusczyk, Friederici, & Wessels (in preparation) found that 9-month-old infants raised in the U.S. preferred to listen to English words rather than Dutch words whereas 2-month-old infants showed no preference for words from either language. This effect was obtained even when prosodic information was removed from the stimuli by means of high-pass filtering. These results suggest that 9-month-olds are sensitive to the phonological structure of the language environment in which they are raised. Jusczyk et al. determined that the effect they obtained was due to the phonotactic differences between English and Dutch. That is, 9 month-old infants are capable of differentiating permissible sequences of phonemes in their language-learning environment from the sequences found in other languages.

Child speech perception – 12 months to 36 months. Several researchers (e.g., Menyuk, Menn, & Silber, 1986; Vihman, 1988) have noted that few experimental studies have been carried out to examine the speech perception abilities of children between 12 and 36 months of age. Children in this period are difficult to test because of their tendency to perform poorly in either the habituation tasks used to test infants or the behavioral procedures used to test older children. Despite these difficulties, several researchers have succeeded in carrying out perceptual experiments designed to answer questions about child capabilities in the period between 12 and 36 months of age. Shvachkin (1948/1973) trained children aged 10–21 months using simple nonsense words (mostly CVC (consonant-vowel-consonant) and CV (consonant-vowel) forms) that contrasted minimal pairs in Russian. After insuring that the children were familiar with the words, he tested the children's ability to discriminate minimal pairs over a period of six months. Shvachkin found that the children followed a developmental sequence and that different contrasts were discriminable at different ages.

Similarly, Garnica (1973) found that 17 to 22-month-old infants from an English-speaking environment tested in a procedure like Shvachkin's also followed a specific developmental sequence for discrimination of different speech contrasts. However, she found substantial variation in the exact form of these abilities among children. Barton (1975) criticized the work of Shvachkin and Garnica on the grounds that their claims were based on statistically unreliable evidence. Barton was unable to replicate the general order of the acquisition of perceptual distinctions reported by Garnica although he did show that performance improved as the children became more familiar with the stimuli used in his task (Barton, 1980). Despite the mixed set of results, the work of Shvachkin, Garnica, and Barton points to the fact that children's perception is not congruent with adult abilities. These results indicate that the abilities displayed in early infancy do not hold up in more demanding contexts, such as when contrasts are embedded in words, thus supporting the argument made by Jusczyk (1982) and others that speech perception studies need to distinguish between auditory and phonological coding. The implications for lexical representation are not as clear-cut as they may seem; although adult and child representational systems may differ, other explanations for differences in performance between adults and children, such as limitations of memory or attentional processes, are also possible.

Child speech perception – 36 months to 60 months. As the child gets older, the range of contrasts within words that he or she can discriminate increases. However, there are some contrasts that the child still has difficulty discriminating. Graham and House (1971) tested children aged 36-54 months in a discrimination task and found that their subjects made more within-manner class confusions than between manner-class confusions. Although Miller & Nicely (1955) obtained a similar pattern of confusions with adults, the children in Graham and House's study tended to make a greater number of these errors than adults. Some phoneme pairs had very high error rates. For example, /f/–/θ/ and /r/–/w/ had error rates of 79% and 60%. respectively. Although most of the errors obtained by Graham and House are also present in adult data, some, such as /l/–/m/ and /p/–/m/ errors, are not typically observed in adult data. As noted above, these results could be due to the form of the lexical representation found in the child, or to cognitive factors such as memory and selective attention.

Speech Production in Children

Infant speech production – birth to 12 months. Infant speech production capabilities generally lag behind infant speech perception capabilities (Benedict, 1979; but for evidence that production can sometimes precede perception, see Strange & Broen, 1981). One example of this lag is the limited range of speech sounds used by infants (Locke, 1983). Early speech production in infants consists of a variety of speech and nonspeech sounds collectively known as babbling. Locke (1983) attributed the development of babbling and its transition to words to the maturation of the vocal tract. As the infant acquires more control over articulatory muscles and the size of the articulators and the shape of the vocal tract becomes more similar to adult dimensions, he or she produces a wider range of speech sounds from the adult inventory. Between 1–5 months of age, speech sounds produced at the back of the vocal tract, such as velars, dominate the infants repertoire. After 5 months of age, sounds that are produced toward the front of the mouth, such as dentals and alveolars, occur much more frequently.

In contrast to the situation at the segmental level, production at the suprasegmental level is more advanced. During the babbling period, the infant begins to superimpose prosodic contours on their speech output that are unique to the language environment in which they are raised. For example, French listeners could distinguish the babbling of 8 month-old French infants from the babbling of infants from other linguistic environments (De Boysson-Bardies, Sagart, & Durand, 1984). This finding was one of the first instances in which the infant's production abilities were found to be influenced by perception of the ambient language and precedes by several months the production of the child's first words. These changes in speech production coincide with the period in which the infant's perceptual capabilities also become closely tuned to the ambient language environment (e.g., Jusczyk et al., in preparation; Werker & Tees, 1984).

In an earlier section I noted that it is not clear whether the current data on infant perceptual abilities are due to linguistic factors or are due instead to nonlinguistic auditory/perceptual abilities (Jusczyk, 1982). A similar point could be made regarding infant speech production abilities. Locke (1983) has argued that much of the infant's early production abilities are primarily due to biological factors, such as the maturation of the vocal tract, that may have little to do with linguistic development per se. Citing linguistic mechanisms as the source of babbling in the productions of early infancy may be premature. On the other hand, Vihman (1988) notes the existence of substantial individual differences among infants in the babbling stage, which vitiates Locke's argument somewhat. If babbling were exclusively a product of biological maturation, one would expect to see much more uniformity across infants. Until there is clear evidence that infant utterances have a communicative intent and are used productively, it is probably premature to impute anything but the most simple linguistic abilities to the infant (cf. Vihman, 1988).

Child speech production – 12 months to 36 months. Between 9-18 months of age the child begins to produce recognizable words to which a communicative intent can be ascribed (Vihman, 1988). Ferguson and Farwell (1975) have suggested that the *word* is the primary unit of acquisition for phonology during this period. A phoneme may first be used in an individual word and only gradually find its way into other words. The form of these early words differs from the form of adult words. The sounds used in early words share much in common with the speech sound repertoire of the babbling period. Indeed, after surveying a large set of studies in which babbling and early word production were examined, Locke (1983) concluded that babbling coexists with the production of the first words by the child. Despite having produced the first few words, there are still phonemes that remain to be correctly produced. In a study of 16–18 month-olds, Irwin and Wong (1983) found that although most English consonants were attempted, only one third of the consonants were considered to be correctly produced. Similarly, Menyuk et al. (1986) describe the vowel space of a child that deviated from the vowel space typically observed in adults.

The sequencing of phonemes in early words also differs in several ways from the adult model. Early words often bear some resemblance to the adult form, as in "wawa" for water, where the initial CV sequence of the adult form is reduplicated in the child's form. Reduplication is one of several phonological processes (cf. Donegan & Stampe, 1979) that have been found in children. Other phonological processes include deletion, substitution, and cluster reduction. The theoretical interpretation of these processes is controversial and will be discussed in the next section. Finally, Ingram (1989) notes that there is substantial variation in the ages when individual phonemes are correctly produced, as well as when and what kinds of phonological processes are used by the child, especially in the early part of this period.

Child speech production – 36 months to 60 months. As the child becomes older, phonemes that were formerly produced inaccurately or avoided altogether become more prevalent and more like the adult model (Irwin & Wong, 1983; Schwartz & Leonard, 1982). These phonemes include the affricates /č/, and /ǰ/, the fricatives /š/, /z/, /v/, /θ/, /∂/, and the liquids /r/ and /l/. These are the same phonemes that are often misperceived by children in this age period (cf. Graham & House, 1971). Locke (1983) speculates that these sounds may become more common during development as the relevant musculature develops, in conjunction with the need to produce the correct adult forms of words like *cheek*, *matches*, and *witch* that the child hears in his or her environment. In addition, the phonological processes common at earlier ages are used less frequently as the child develops and acquires mastery of the adult phonological system (Vihman, 1988).

Inferring lexical representation from perception and production data. Using speech perception and production data to make inferences about the child's representational system is fraught with difficulty. Even when the child begins to produce recognizable sequences of sounds that are regularly associated with some behavior, it is not clear what level of lexical organization is evidenced by these productions. Locke (1986, p. 6) suggests that "to 'motivate' [dædæ] one need only assume enough perception—and internal representation—

for something like [+stop] or [+oral] that is embedded in a sound pattern of roughly disyllabic length." Locke argues that one should not attribute to the child any more representational complexity than is absolutely necessary. Thus, one hypothesis for phonological development is that the child progressively builds up or constructs a linguistic representational system that eventually corresponds to the adult form. This is the position advocated by researchers such as Macken (1980) and Maxwell (1984).

Conversely, the child may have a fully developed lexical representational system but various factors, including an immature neuromuscular system and underdeveloped memory and attention systems, prevent the child from fully utilizing these representations. This is the position taken by researchers such as Donegan and Stampe (1979) and Smith (1973). These researchers argue that the child's underlying representation of the surface form of words is basically the same as the adult's. Donegan and Stampe (1979) propose that from an early age the child has access to accurate perceptual information about the sound structure of the words that he or she hears in the environment. However, the reason that the child's productions do not match the adult form is that the child has an innate phonological system that works to simplify production through the application of phonological processes such as reduplication and cluster reduction that were described earlier. Donegan and Stampe argue that these same processes are also found to some extent in the casual speech of adults and are a result of the tension between ease of articulation and clarity of production. Thus, learning adult phonology is learning not to apply these innate phonological processes. Similarly, Smith (1973) refers to 'realization rules' that the child applies to the underlying adult form in order to derive the productions observed in early childhood.

Various types of evidence have been marshaled to support the positions outlined above. Smith (1973) argued that children must have an adult representational system because they can discriminate among adult productions of words they themselves are not able to correctly produce. Moreover, because children possess an immature production system, their speech does not accurately reflect the state of their representational system. Thus, evidence from children's productions likely underestimates the complexity of their representational system (cf. Benedict, 1979).

Maxwell (1984) disagrees with Smith on this point. She suggests that Smith's use of perceptual data to support the existence of an adult representational system is unwarranted because of several shortcomings associated with the interpretation of data from discrimination tasks of the type cited by Smith (cf. Locke, 1980a, 1980b; Jusczyk, 1982). For example, Locke (1980a) argues that unless perceptual tests include repeated instances of the same stimuli and contrasts from multiple contexts, little can be inferred about the child's linguistic representations. Although in later work Smith (1981) allows for the possibility that there may be some misperception by children, he still argues that, in general, children's representations are faithful reproductions of what they hear from speakers around them.

Maxwell also criticizes Smith's (1973) argument that sound change is "across-the-board" in children. According to Smith (1973), if children have an adult-like representation system for the sound shape of words, changes in the production of a sound that are observed to occur in a particular context within a word should be global, occurring in all instances of the particular context. Maxwell cites several reanalyses of Smith's data (e.g., Macken, 1980; Braine, 1976) that indicate that instead of an across-the-board change, change is idiosyncratic. That is, sound change in children generally includes exceptions to the rule rather than a completely uniform application of the rule. Maxwell interprets this observation as evidence that children do not have the same representational system for the sound structure of words as adults.

Evidence that the child's representational system has a nonadult form comes from several sources. Both Macken (1980) and Braine (1976) have suggested that some of the speech errors made by children are more easily explained by misperception of phonemic categories than by the complex rules for transforming adult representations into a child's form, as in Smith (1973) and Donegan and Stampe's (1979) accounts. In addition to the work of Barton (1980) and Garnica (1973), evidence for misperception in two-year-old children has been obtained by Eilers and Oller (1976) who found that children more frequently misperceived words they misproduced than words they produced correctly. Words misperceived or misproduced varied across the children in their study, suggesting that children not only have non-adult representational systems but that they can also have idiosyncratic representational systems. Maxwell (1984) also argues that children have idiosyncratic nonadult representational systems based on evidence from misarticulating children.

Besides misperception, nonadult-like productions could also be the result of imperfect production hypotheses. Macken & Ferguson (1983) argue that children may engage in hypothesis testing when formulating possible arrangements for phonemic segments and that misproductions may be a result of a tentative hypothesis that will be modified eventually.

Summary. The form of lexical representation in children is far from clear. Work in this area must contend with variability among children, the application of phonological processes, misperceptions, and other assorted problems. Despite these difficulties, several findings have been obtained that bear on the issue of lexical representation in children. Perceptual data indicate that even 4 and 5 year-old children sometimes misperceive segments in words, especially in the context of unfamiliar words. Evidence such as this coincides with the observation that the child's phonological system is tied to individual words early in development (cf. Ferguson & Farwell, 1975; Maxwell, 1984). Together, these findings suggest two important characteristics of lexical representation and organization. First, children do not learn phonemes as abstract idealized units but rather they acquire knowledge of them as the constituent elements comprising individual words. Second, word frequency is an important factor in the organization of lexical information.

The 'Whole/Part Hypothesis' of Lexical Development

In the previous section, I described the development of the perception and production of spoken language in children. In this section, I outline a hypothesis proposed by several child language researchers to account for how the child acquires knowledge about the internal structure of words. The hypothesis, which I will refer to as the 'whole/part hypothesis', posits the successive differentiation of linguistic information during development. After describing the whole/part hypothesis, I present evidence from studies of spoken word recognition, speech production, and metalinguistic awareness that support the hypothesis.

The Whole/Part Hypothesis

The whole/part hypothesis is an account of how the child learns to deal with language as a hierarchically-organized system. Several theorists have suggested that before the child can deal effectively with the phonological structure of language, he or she must first 'break into' or segment the continuous stream of speech found in his or her environment and isolate individual linguistic units (Chiat, 1979; Menyuk et al., 1986; Moskowitz, 1973; Peters, 1977, 1983; Waterson, 1971). The general form of this hypothesis is that infants first deal with global characteristics of the speech signal such as intonation and stress. They then use this prosodic information to organize speech into a sequentially-ordered and hierarchically-organized series of segments that eventually corresponds to the linguistic organization of the adult. Although the units that the child deals with at first could be as large as clauses or other types of word strings, most theorists consider that words or word-sized units are the first utterances that children use linguistically (cf. Jusczyk et al, 1992; Moskowitz, 1973). The only exceptions to this would be idioms, in which the collective meaning of a word sequence is not captured by compiling the meanings of the individual words.

According to the whole/part hypothesis, language development is viewed as a progression from relatively large units, such as words, to successively smaller units, such as syllables or phonemes. As each level of analysis is attained, the infant has a relatively an undifferentiated representation of the units that comprise the levels immediately below the current level of analysis. Thus, a child that has just learned to segment an utterance into a string of words would have access to only relatively coarse and undifferentiated information about the internal structure of individual words, perhaps nothing more than some idiosyncratically salient acoustic features. This is not to say that all words in the child's lexicon are represented at the same level of detail at a particular moment in time. Rather, evidence suggests that at any given time some words are analyzed in more detail than other words (cf. Ferguson & Farwell, 1975). The importance of the whole/part hypothesis lies in its ability to account for the performance of children tested in word recognition experiments, speech production experiments, and in experiments on initial reading skills.

The possibility that children deal with perceptual units as large as whole words and word strings has parallels in how the adult lexicon may be organized. For example, Bolinger (1975) argues that, in addition to a lexical organization based on individual words, the adult lexicon contains many instances where groups of words are produced together as a unit, such as idioms (e.g., "Hold your horses.") and greetings (e.g., "Good morning. How're you?"). Support for Bolinger's position was obtained by Swinney and Cutler (1979). They found that subjects accessed printed idioms faster than printed control strings of words that had the same meaning, suggesting that subjects treat idioms as unitary wholes. Together with the evidence cited earlier regarding lexical constituents such as words, syllables, and phonemes, the use of supralexical units such as idioms attests to the flexibility of adult language users in accessing the lexicon. Although it is unlikely that adults rely exclusively on supralexical units for representation, their existence lends credence to the idea that similar forms of representation may be utilized by children in their initial encounters with language.

What causes the child to develop successively more detailed lexical representations? Jusczyk (1986) and Walley (in press) have proposed that the whole/part developmental sequence is due to the increase in new vocabulary items that begins around 18 to 24 months. Some have referred to the increase in new words in the child's lexicon as the 'vocabulary spurt' or 'naming explosion' because of the rapid rate of acquisition often observed(see Gopnik & Meltzoff, 1987, for a discussion of these terms). Although there are questions as to whether all children acquire words at the same rapid rate (Nelson, 1973; Goldfield & Reznick, 1990), it is a basic fact of language development that the child's lexicon begins to undergo a substantial increase in size around 2 years of age.

The increase in vocabulary size has important consequences for lexical organization. When only a limited number of lexical items are stored in long-term memory, the representational system required to maintain distinctiveness among individual items need not utilize the fine-grained, sequentially organized detail that characterize adult representations. However, as more and more words are added to the child's lexicon, the continued use of relatively undifferentiated representations becomes impractical. Undifferentiated representations increase the potential for confusion among words during perception and production, a condition known as *homonymy*, that would hinder communication between the child and other individuals in his or her environment.

A second factor that may act as an impetus for reorganization of the lexicon is the child's initial attempts at producing multiword utterances. In order to maintain more than one word at a time in short-term memory, the child must have some means by which to chunk the acoustic-phonetic features comprising words. Otherwise, the limited capacity of working memory would be exceeded in trying to keep track of the words and the temporal order in which they are to be produced. Phonological organization supplies the means to chunk the sound structure of the words in a way that enables resources to be freed for the task of managing syntax. Thus, vocabulary growth and increased demands due to initial attempts at multiword utterances require an efficient organization of words in the lexicon. It is reasonable to assume that because the adult lexicon is organized, in part, on the basis of the acoustic-phonetic similarities among words (Luce, 1986a; Treisman, 1978a, 1978b), a similar system is probably used by children. The extent to which the child's lexical organization resembles adult lexical organization most likely depends on the development of an adult-like phonological system in which the sound patterns of words are systematically related to each other.

Evidence for the Whole/Part Hypothesis

The whole/part hypothesis of lexical development is consistent with a substantial body of evidence that has appeared in the published literature. Macken (1979) examined the phonological development of a young child learning Spanish in the period between 1;7 and 2;1 and noted a "...global, only partially differentiated auditory processing..." that was "... paralleled by [the child's] loose, prosodic treatment of words." (p. 47). Similarly, Peters (1977, 1983) suggested that in the initial stages of language development, children often possess only a gross, undifferentiated representation of words and strings of words. Furthermore, Peters (1983) argued that although linguistic units such as words and phonemes are useful for describing a language, they may not be what the child uses. Peters based her claims, in part, on her longitudinal study of a boy who appeared to deal with adult forms, such a questions and descriptions, as non-adult, Gestalt-like units. That is, the child would often produce words in the context of stereotypical utterances that he had heard in his environment, but could not use the same words productively in other contexts (Peters, 1983).

Even older children retain a tendency to produce words that bear an overall correspondence to the sound shape of the intended utterance. Aitchison and Straf (1982) obtained results consistent with the whole/part principle in an analysis of malapropisms produced by children aged 12 years and younger and adults. Malapropisms result from the retrieval of words from long-term memory that are similar in overall sound shape to an intended word (e.g., 'monuments' for 'condiments,' 'strangled' for 'scrambled'). Aitchison and Straf found both similarities and differences in the malapropisms produced by adults and children. The malapropisms produced by both groups were similar in that they resembled the gross phonological features of the intended word. However, the malapropisms produced by adults generally preserved the initial sound sequence of the intended word (e.g., 'Boeing' for 'boa') whereas the malapropisms produced by children tended to preserve the syllabic count and the stressed vowel (e.g., 'money' for 'honey,' 'Christmases' for 'trespasses,' and 'snake' for 'steak'). These findings suggest that the child's representation of the sound pattern of words differ from that of adults. The authors also provided an alternative explanation for their results where the process by which children retrieve lexical information differs from how adults retrieve lexical information. Either account is consistent with a version of the whole/part hypothesis that posits a developmental change in the ability to access the internal parts of words.

Aspects of malapropisms are similar to what has been observed in the tip-of-the-tongue phenomenon, where a talker searches for a word but retrieves only words that are semantically and/or phonologically similar to the target word (Brown & McNeill, 1966). Elbers (1985) observed the tip-of-the-tongue phenomenon in her 30-month-old child and noted that the utterances produced by the child resembled the global form of the intended utterances. The malapropisms produced by Elbers' child preserved the number of syllables, the stress pattern, and some of the phonemes found in the intended word (but not the initial phonemes; cf. Aitchison & Straf, 1982).

Further evidence consistent with the whole/part hypothesis was obtained by Aitchison and Chiat (1981). They presented novel spoken words to children aged 5, 6, and 9 years and tested their ability to repeat the words both immediately and after a delay of several minutes. In the first condition, even the youngest children were able to accurately reproduce the words. In the delayed repetition condition, however, both the younger and older children often produced errors that were similar to those produced in the spontaneous speech of much younger children (i.e., 2- and 3-year-olds). The youngest group of children that Aitchison and Chiat tested showed the most similarity to 2- and 3-year-olds. This similarity was particularly evident in cases of consonant harmony between the first and second consonant (e.g., /dudu/ or /gugu/ for 'kudu'). Other characteristics that accounted for errors in recall included omission of unstressed syllables, imposition of a CV syllable structure on the target words, confusion with other phonemes from the same phonetic class, and better preservation of consonants surrounding a stressed vowel than in other contexts. Overall, these findings suggest that if the capabilities of the child's phonological system are taxed, there is a regression to mechanisms that were used earlier in development, in this case the phonological processes of early childhood that preserve the global characteristics of target words.

Findings from several studies by Walley and her colleagues are also consistent with the whole/part hypothesis. Walley, Smith, and Jusczyk (1986) found that both 5- and 8-year-olds were sensitive to segmental information when they were required to classify a novel nonword stimulus on the basis of a match with two previously presented nonword stimuli (i.e., standards with which to compare the novel stimulus). However, the younger children appeared to judge similarity based more upon the overall similarity among phonemes in the standard stimuli and the novel stimulus than the older children. The older children classified nonwords as similar to each other on the basis of a single initial phoneme match. The results of this experiment are consistent with a view of word recognition in which the young child utilizes a more wholistic representation than older children.

Walley has also used a mispronunciation detection task to investigate word recognition in children (Walley, 1988; Walley & Metsala, 1988). In this task the ability to detect mispronunciations in different positions within a word is assessed. In a study with adults, Cole (1973) found that listeners detected mispronunciations in the initial part of words more accurately than in other parts of words. Walley found that young children (aged 4 and 5 years) did not attend to the initial part of a target word more than any other part of the target word if the word was presented with no supporting context (e.g., no picture referent). However, when the target word was presented in context, the 5-year-olds attended more to the initial part of the word than to other parts of the word. Walley (in press) suggests that this result indicates that children as old as 5 years do not have a lexical representational system that permits easy access to the internal structure of lexical items solely on the basis of the segmental, left-to-right acoustic-phonetic input. Finally, Walley (1988) found that children aged 5–6 years tested in a gating task needed much more of the signal than adults in order to correctly identify the target word (for similar results, see Elliott, Hammer, & Evan, 1987). This result parallels the findings obtained in the mispronunciation detection task;

children have difficulty taking advantage of the temporally-ordered segmental information in spoken words to facilitate recognition. Whereas adults can utilize word-initial information to aid recognition, children are not able to take advantage of this information unless additional contextual information is present.

Metalinguistic Abilities

Evidence consistent with the whole/part hypothesis has also been reported in research examining metalinguistic awareness in children. Metalinguistic awareness describes the child's conscious awareness of the structural features of language and an ability to explicitly manipulate and talk about these features. Metalinguistic awareness contrasts with the implicit knowledge of language, the knowledge of how to *use* language that language users generally develop first. Striking parallels have been observed in children between the development of explicit and implicit forms of linguistic knowledge.

One aspect of metalinguistic awareness is an ability to segment larger linguistic units into smaller units. There is a substantial body of evidence that shows young children have difficulty with tasks requiring the explicit segmentation of words into phonemes, and that this difficulty is especially evident before children learn how to read (Goswami & Bryant, 1990). Even children who have already started school experience difficulty with tasks that seem easy for an adult. Chaney (1989) examined the ability of kindergartners aged 4–6 years to segment into words a passage that they had recited once a day for one year, the Pledge of Allegiance. She found that despite considerable practice, the children all showed various kinds of segmentation errors, averaging only 81% correct on a word-by-word basis. Some of the errors Chaney found were quite remarkable, at least from an adult perspective: 'United States ' → "nine a states", "the night of states", "night in stage"; 'for which it stands ' → "for witches stand"; liberty '→ "liver T". In the same study, Chaney found that errors were more likely to occur in function words than content words, and in abstract words more than concrete words. Overall, the explicit segmentation of sentences into words by kindergartners was found to be far from perfect.

In addition to knowing how to segment sentences into words, a crucial part of reading an alphabetic writing system such as English is an ability to segment words into phonemes (Venezky, 1977). A fundamental characteristic of skilled readers is an awareness of spelling-to-sound rules, the rules that relate the sound of a phoneme to its spelling (Venezky, 1977). However, in order to learn how phonemes map on to letters, it is first necessary to learn how spoken words can be segmented into their constituent phonemes (Rozin & Gleitman, 1977). Children learning how to read often have difficulty learning spelling-to-sound rules, in part, because of the irregular nature of English spelling-to-sound rules (e.g., 'g' can refer to /g/ or to /ʒ/, even within the same word—'garage'). However, an even more fundamental problem for children learning to read is segmenting a word into its constituent phonemes. For example, children have great difficulty tapping out the number of phonemes in a word whereas they can generally learn to tap out the number of syllables in a word quite easily (Liberman, Shankweiler, Fischer, & Carter, 1974). Similarly, Treiman and her colleagues (Treiman & Baron, 1981; Treiman and Breaux, 1982) have shown that young children can analyze spoken words into syllables before they can analyze them into phonemes. Treiman (1985) further showed that children appeared to be also sensitive to a level of perceptual analysis between syllable and phoneme. As noted above, the internal structure of syllables consists of the onset (initial consonant(s)) and the rime (the vowel plus any consonants that follow the vowel). Treiman (1985) found that children aged 4-5 years could more easily identify the initial phoneme of a word if it coincided on a one-to-one basis with the onset (e.g., 'sat') than if it was part of a onset structure consisting of more than one phoneme (e.g., 'stat'), implying that onset and rime are indivisible perceptual units for children at this age.

The difficulty experienced by children attempting to explicitly access knowledge about the internal structure of words appears to parallel the difficulty they experience in constructing a phonological system based on their implicit knowledge of the sound structure of spoken words. From the research reviewed in this section it appears that children can access the syllables and intrasyllable structures within words more easily than individual phonemes. Goswami and Bryant (1990) have suggested that this is because syllabic structure coincides with the rhythmical structure of spoken language. Phonemes, on the other hand, do not have redundant markers such as intonation and stress to indicate their position in an utterance.

Summary

The whole/part hypothesis is a useful framework for integrating a wide variety of findings about children in the areas of speech perception, speech production, spoken word recognition, and metalinguistic awareness. The research described in this section is consistent with the proposal that lexical representation in children begins with the preservation of the global characteristics of utterances and ends with a hierarchically-organized lexicon containing information about words at several different levels of representation.

The Relationship Between Lexical Representation and Cognitive Development

In the previous section I described the whole/part hypothesis as a useful way to characterize some of the changes in the form of the child's lexical representation during development. In this section, I focus on some additional constraints imposed on lexical representation by the child's cognitive and perceptual mechanisms. These mechanisms are important because the way they function may underlie the developmental trend described by the whole/part hypothesis. Therefore, for purposes of understanding the development of lexical representation, it is important to consider the constraints imposed on processing spoken language by categorization, memory, and attentional systems in children.

The Whole/Part Principle in Cognitive/Perceptual Development

In the previous section, I presented an account of lexical development that was based on the successive differentiation of linguistic representations. This account coincides with a more general framework proposed to account for cognitive and perceptual development. According to Aslin and Smith (1988) and others (e.g., Gibson, 1969; Kemler-Nelson, 1983), a general characteristic of perceptual development is a progression from analysis of objects at the level of wholes to analysis at the level of parts. Aslin and Smith (1988) have revised this principle further, suggesting that the development of perceptual abilities follows a curvilinear function. Early infant abilities permit an analysis of perceptual objects into sensory primitives that form the basis for development of higher level perceptual representations. However, the form of these early perceptual representations does not afford the efficient extraction or analysis of elements of the perceptual object. Once representations achieve some stability and contain sufficient detail, the child can again begin to extract from the representations information about the constituent parts of the perceptual object. This sequence of development is identical to the whole/part hypothesis that was invoked in the previous section to account for lexical development. Generally, then, some parallels can be identified between the development of different levels of lexical representation and different types of cognitive/perceptual abilities.

Recent work in the area of infant categorization skills illustrates the similarities between cognitive/perceptual and linguistic development. The ability to group perceptual objects into different categories on the basis of some physical dimension(s) is a fundamental cognitive skill. Mandler and her colleagues (Mandler & Bauer, 1988; Mandler, Bauer, & McDonough, 1991) obtained evidence indicating that by 18 months, infants have attained global conceptual categories but have not yet differentiated the basic-level categories within these superordinate categories. Mandler et al. (1991) suggest that "...one of the characteristics of developing a knowledge base is to make finer and finer distinctions among initially more global conceptions (p. 290)." The results obtained by Mandler et al. contrast with the claims of Mervis and Rosch (1981), who argued that basic-level categories were the level at which infants began to categorize perceptual objects. Mandler et al.'s results are, instead, consistent with the whole/part hypothesis of development.

One possible implication of the similarity between linguistic and cognitive/perceptual development is that there is a causal link between the two. Although no research has demonstrated that developments in cognition or perception *cause* linguistic structures to develop, several studies have shown correlations between the development of cognitive performance and linguistic abilities. For example, Corrigan (1978) and Gopnik and Meltzoff (1987) found that the naming explosion and the attainment of object permanence both appeared at approximately 18 months. Thus, several researchers have obtained results that suggest that the whole/part principle is a central component of cognitive and perceptual development. Perhaps more importantly, from the perspective of the present investigation, is that many of the changes in cognitive/perceptual abilities also coincide with changes in lexical representation and organization of the mental lexicon.

The Role of Memory and Attention in the Development of Lexical Representation

Memory and attention play a central role in virtually every cognitive ability, including spoken language processing. Thus, it is not unreasonable to assume that mechanisms proposed to account for memory development could also account for aspects of language development. Similarly, attentional processes determine in large part the form of the child's representation of an utterance in memory. In this section, I review some findings about the relevance of memory and attention to the development of children's lexical representations.

Memory in Children. A basic finding in infant memory research is that performance improves as the infant becomes older. In a series of experiments testing recognition memory, Rovee-Collier, Early, and Stafford (1989) found that 3-month-olds retained characteristics of a visually presented display over a 24-hour period that 2- month-olds did not retain. If the 2-month-olds were provided substantial retrieval cues, their performance was similar to the 3-month olds. Rovee-Collier et al. attributed this result to a "leaner memory representation" in the 2-month olds. Rovee-Collier et al.'s interpretation coincides with the whole/part hypothesis of lexical development, where the form of early lexical representations is assumed to be less detailed than later representations. An inability to retrieve stored information about past events may be an additional cause of poor retention in early infancy (Sullivan, 1982). However, inefficient retrieval may also be a function of inadequately differentiated representations in memory.

After 6-months of age, infant recognition memory improves substantially and continues to improve into early childhood (Kail, 1990). Recall appears to develop later than recognition memory. Seven month-old infants were able to locate objects that were removed from sight that younger infants (5- and 6-month-olds) were not able to locate (Kail, 1990). Incorporating a delay between the last time the infant saw the object and when the infant was required to search for the object substantially reduced performance in 10-month-olds (Diamond, 1985). Although the work described here dealt with memory for visual events, similar phenomena characterize the development of language abilities. In both cases, the mental representations available to infants before 12 months of age seem to be less stable and more prone to interference than adult representations (cf. Jusczyk, in press).

Changes in auditory/phonological short-term memory are also an important aspect of memory development. Most research has focused on the performance of children between middle and late childhood who have been tested in recall tasks. A general finding in both children and adults is that acoustic similarity among words has an adverse effect on the recall of words lists (Conrad, 1964). However, changes in the effect of acoustic similarity occur during childhood. Hulme (1984) found that the effect of acoustic similarity on recall increased in children between the ages of 4 and 10 years. He tested children in a serial-ordered recall task with acoustically similar word lists and acoustically dissimilar word lists and found that although recall improved for both lists as age increased, recall of the acoustically dissimilar lists improved more than recall of the acoustically similar lists. That is, the older children appeared to be able to take advantage of the acoustic-phonetic distinctiveness of the words in the acoustically dissimilar lists to help keep the words from being confused with each other (cf. Goldinger, Pisoni, & Logan, 1991). Hulme and Tordoff (1989) and Cowan, Saults, Winterowd, and Sherk (1991) have proposed that an increase in the efficiency of covert rehearsal processes accounts for these results. As children get older, the efficiency of rehearsal improves. This means that older subjects can rehearse more items over any given time period than younger subjects. However, the probability of confusing acoustically similar items as a function of increased rehearsal is greater than the probability of confusing acoustically dissimilar items. Thus, acoustic confusions will occur more often in older children and adults.

These findings have implications for development in terms of how words are represented and organized in lexical memory. The more limited capacity of short-term memory in young children may impair rehearsal enough to substantially reduce their access to acoustic/phonetic information. The effects of these limitations can be observed in both speech perception and production. In speech production, phonological processes that simplify the structure of utterances may be a result of the inability of the short-term memory system to keep track of the information associated with individual phonemes. Crowder (1978) has noted that the effect of acoustic similarity in recall tasks is larger when subjects must keep track of the order in which stimuli are presented. Since a central feature of phonological development is learning the proper ordering of phonemes in spoken words, the need to simultaneously maintain information about individual phonemes and their order will cause an increase in the confusability of segments and/or their order. As rehearsal processes become more efficient, the child's ability to deal with segmental and order information improves substantially. Similarly, the child's ability to compare utterances and to determine the identity of individual segments and their ordering will also become more reliable as the capacity of short-term memory increases. Thus, the development of lexical representation may be due, in part, to the limitations of short-term memory.

Selective Attention in Children. The same trend observed for the development of memory in children is also observed for the development of selective attention: the ability to selectively attend to aspects of stimuli improves with age (for reviews of selective attention in children, see Enns & Girgus, 1985; Lane & Pearson, 1982). That is, as children get older, they are distracted less by irrelevant stimulus variation in their immediate environment. It is this aspect of attention that I will concentrate on here, in part because of its potential application to understanding aspects of lexical development. Theoretical accounts of attention in adults and children fall into two broad categories according to Enns and Akhtar (1989). The first posits a limited capacity attentional system that increases in capacity during childhood and is similar to proposals of how short-term memory develops (e.g., Hulme & Tordoff, 1989). The second account posits that attention is a strategic process in which the subject allocates cognitive processes to where they are needed (e.g., Chi, 1976). According to the latter account, the difference between children and adults is similar to the expert-novice distinction in adults: experts differ from novices in that experts have developed strategies for efficiently allocating their attention whereas novices have not.

Enns and Akhtar (1989) provide a representative sample of work on the development of attention. Children (4-, 5-, and 7-year-olds) and adults were required to map a visually presented target to one of two response categories. Two conditions were tested, one in which a target was flanked by distracters and the other in which the target was presented alone. Enns and Akhtar found that the presence of flanking stimuli slowed response time for all subjects. However, the flanking stimuli interfered more with the performance of younger subjects than it did with the performance of older subjects. Enns and Akhtar interpreted their results as consistent with the strategic account of attention, in which children do not have the well-developed attentional strategies of adults. Similarly, Geffen and Sexton (1978) found that selective attention to auditory stimuli improved in children between 7 and 10 years of age. They also attributed their results to the development of attentional strategies.

The development of attention in children as observed in studies such as Enns and Akhtar (1989) or Geffen and Sexton (1978) resembles aspects of lexical development. During lexical development, children are learning a phonological system that may be thought of as an attentional strategy for efficiently dealing with linguistic information. A phonological system facilitates the allocation of attention to the relevant parts of spoken language and inhibits attention to the irrelevant parts of language (cf. Jusczyk, in press). Furthermore, phonological systems undergo developmental changes that parallel the expert-novice shift observed in other cognitive domains. Thus, there are several similarities between the development of attention and the development of a phonological system.

Shepp and Barrett (1991) examined the development of attention using integral and separable stimuli (see Garner, 1974). They tested 6- and 8-year olds and adults in several speeded classification tasks in which geometric figures varying in size and shape served as stimuli. Shepp and Barrett found that separable stimuli were always perceived in terms of the independent features comprising the stimulus object, regardless of age. The integral stimuli were initially perceived as wholes but were analyzable into their component dimensions by the adult subjects, although at some cost in response time. This pattern of results is consistent with the whole/part hypothesis (Smith & Aslin, 1988): although some kinds of stimuli may be perceived in terms of their separate dimensions, other kinds of stimuli are perceived as wholes and only later in development can they be analyzed into their constituent parts. Moreover, the findings obtained by Shepp and Barrett are consistent with results showing that young children may not be able to efficiently access the internal structure of spoken words, suggesting that the changes in the ability of children to allocate attention to relevant stimulus characteristics is a general developmental trend across modalities.

Overall, research on the development of selective attention has revealed a pattern of results that coincides with aspects of lexical development. If one takes the view that attention is essentially the development of various strategies for the efficient processing of information, while not discounting possible changes in capacity, the parallels with the development of a phonological system are striking. In each case, the child is developing expertise for dealing with information in his or her's environment. To learn a first language, the child must learn a phonological system that permits the efficient extraction of phonemes, word boundaries, and other types of linguistic information. Viewed this way, language development is a function of the same principles that underlie development of other kinds of cognitive skills.

Summary

Lexical representation in children is the result of a process that begins in infancy and that culminates with the attainment of an adult-like representational system. Infants can discriminate a remarkable range of phonemic contrasts, but it is not clear whether this is due to a linguistic system or to basic auditory abilities. That the perceptual abilities of infants may be a product of nonlinguistic factors is suggested by the misperceptions observed in children aged 1–2 years tested in tasks using words as stimuli. Even children aged 4–5 years misperceive phonemes more than adults. In addition, children in this age range are not able to take advantage of partial acoustic information to identify words as effectively as adults.

Production abilities also undergo substantial changes during development. The first words are global approximations of adult forms, due, in part, to errors in perception and, in part, to the operation of phonological processes that simplify the production of utterances. Phonological processes continue to be used until around three years of age. However, evidence from child malapropisms suggests that differences in lexical representation and/or processing between and adults and children continue to exist well into late childhood.

A general characteristic of lexical representations in early childhood is that they appear to incorporate the global characteristics of utterances together with a loosely organized set of acoustic-phonetic features. The global representations of early childhood are then analyzed into progressively more fine-grained linguistic units so that the adult hierarchy of word, syllable, intrasyllabic units, phonemes, and features can be achieved.

The whole/part hypothesis was proposed to account for the development of lexical representation and organization. This hypothesis posits that a general principle of cognitive and perceptual development is the successive differentiation of perceptual objects. According to the whole/part hypothesis, infants have efficient mechanisms for extracting sensory primitives from their environment. However, the crude representations formed from the sensory primitives prove inadequate for the extraction and re-analysis of perceptual information and for the construction of robust representations in long-term memory. The form of these early representations undergo successive reorganizations as the child incorporates more detailed lexical information. Eventually, these representations assume the form found in adults.

The whole/part hypothesis of perceptual development is consistent with several critical aspects of lexical development. Infants display good perceptual abilities when required to make simple discriminations between phonemes but fail in more demanding situations. Nonetheless, these basic perceptual skills can be used to extract sensory primitives from the speech signal that correspond to linguistic units such as clauses (cf. Aslin & Smith, 1988). As a consequence of continued exposure to a particular language, these large units then undergo a successive differentiation into more elementary linguistic units, such as words. Two factors may be responsible for the reorganization of lexical representations. First, the addition of new vocabulary items as the child gets older forces differentiation to prevent confusions among words. Second, the demands of producing multiword utterances requires an efficiently organized lexicon in order to facilitate the ordering of words into syntactically correct strings. Constraints on the development of representational systems are imposed by cognitive processes such as categorization, memory, and attention. In short, the development of the lexical representation system and its organization shares many similarities with the development of representational systems in other cognitive domains. In each case, development proceeds from relatively undifferentiated representations to more analyzed and detailed representations.

Chapter IV. Computational Analyses of Child Language: Phonemic Representation

Hypotheses about lexical representations can be tested by collecting a sample of text or spoken language and performing computational analyses to determine the structural characteristics of the words in the sample. 'Structural characteristics' here refers to variables such as word frequency, similarity, and lexical stress. This chapter begins with an overview of computational analyses of lexical databases, including a description of the database used in the present investigation. Then, I describe the computational analysis of the database and how I analyzed neighborhood density and word frequency in samples of spoken language from several children between the ages of approximately 1 1/2 to 5 years. The analyses described in this chapter assumed that the children utilize an adult-like, phonemic segmental representation system. I compare the results of the present analyses to results obtained in earlier work that utilized a similar procedure. Finally, I describe the results of analyses examining the relationship between word frequency and other lexical variables in children between 1 1/2 and 5 years.

Computational Analyses of Lexical Databases

Until recently, the computational analysis of language consisted of word counts in which token frequencies for individual words in a large sample of text were calculated. Examples of word counts include the Thorndike-Lorge (1944) and the Kucera and Francis (1967) counts (see Bontrager, 1991, for a review of early word counts). Originally, counts were used to guide the development of pedagogical materials, such as spelling books and readers. Later, word counts began to be utilized for more theoretical purposes. For example, word counts have proven invaluable in studying the effects of word frequency in word recognition and lexical organization (e.g., Forster, 1976; Howes & Solomon, 1951; Pisoni, Nusbaum, Luce, & Slowiaczek, 1985; Rumelhart & McClelland, 1981; etc.). A less familiar type of count is a phoneme count, in which the frequency of phonemes in a large sample of spoken language is calculated. Examples of phoneme counts include Dewey (1923), Denes (1963), and Mines, Hansen, and Shoup (1978). Reasons for phoneme counts varied. Dewey's 1923 count was carried out to further the goals of the spelling reform movement and to develop efficient methods for taking shorthand. In contrast, the phoneme counts of Denes and Mines et al. were used in telecommunications work and speech recognition.

Overall, word and phoneme counts provided basic information about the characteristics of words used in the English language (Miller, 1951). Even if analyses of lexical databases were limited to word counts and phoneme counts, there are many questions about the structure of language that could be answered by recourse to these databases. For example, if a researcher asserts that a particular phonological feature is marked linguistically, the frequency of the feature in a word count could provide the necessary evidence to support the claim. In summary, the primary benefit provided by the analysis of a large language database is that it can be used to test hypotheses about the structural characteristics of language. If the sample is sufficiently large and representative of the language used by a specific population, a number of generalizations can be formed about the population on the basis of evidence obtained from the sample.

Lexical Statistics in Adult Language

Landauer and Streeter (1973) examined the relationship between word frequency and confusability. One of their goals was to assess the assumption of "acoustical equivalence" between high and low frequency words contained in some accounts of the word frequency effect (e.g., Broadbent, 1967; Morton, 1969). Landauer and Streeter employed a procedure that was similar to one developed by Greenberg and Jenkins (1963) to estimate the number of words that could be confused with a target word. In this procedure, a target word is compared to another word. If by substituting a single phoneme the target word can be transformed into the second word, the second word is considered a neighbor of the target word. Landauer and Streeter uncovered several interesting findings. First, high frequency words tended to have more neighbors than low frequency words. Second, high frequency words tended to have high frequency neighbors. Third, the distribution of constituent phonemes in high and low frequency words also differed. Landauer and Streeter concluded that the assumption of acoustical equivalence between high and low frequency words was unwarranted.

Landauer and Streeter examined the relationship between frequency and orthographic neighborhood density in printed words. Their claims about the relationship between these variables were based on a relatively small subset of the words in English. Luce (1986a) extended the work of Landauer and Streeter by examining the relationship between frequency and neighborhood density in spoken words, as well as sampling a much larger set of words. As described earlier, Luce analyzed each of the 20,000 words in Webster's Pocket Dictionary in terms of three variables: the number of each word's neighbors, word frequency, and the frequency of each word's neighbors. Luce found the same pattern of results that Landauer and Streeter observed. Moreover, in a series of behavioral experiments, Luce found that these structural characteristics also influenced subjects' perception of spoken words. Luce's findings have been replicated in several studies (e.g., Goldinger, Luce, & Pisoni, 1989; Cluff & Luce, 1990), strengthening the claim that neighborhood structure, combined with word frequency, is an important variable in spoken word recognition tasks. In addition, these results suggested that the structural characteristics of lexical databases could be used as a model of how the mental lexicon is organized and how lexical information is utilized in word recognition. The relationship between neighborhood density and word frequency was formalized in NAM, Luce's (1986a) model of spoken word recognition. In NAM, neighborhood density is used to define a set of candidates generated from stimulus input and word frequency serves to bias the selection of a single word from this set of candidates.

Analyses of lexical databases and related experimental findings also indicate a link between structural variables and syntactic processes. Recent analyses by Soreno and Jongman (1990) suggest that phonological structure and frequency interact with syntactic form class. They analyzed the Kucera and Francis (1967) corpus and found that high frequency nouns were more likely to contain back vowels whereas high frequency verbs were more likely to contain front vowels. In contrast, low frequency nouns and verbs did not differ in vowel type. Soreno and Jongman found that these variables also played a role in perception. High frequency nouns with back vowels were recognized faster than a matched set of high frequency nouns with front vowels. Conversely, high frequency verbs with front vowels were recognized faster than a matched set of high frequency verbs with back vowels. Based on their analyses, Soreno and Jongman concluded that syntactic form class and vowel quality appear to be related to each other and may be used to organize words in the lexicon.

In a similar vein, Kelly has carried out a series of experiments to investigate the relationship between various phonological characteristics, such as stress and syllable number, and aspects of syntax (e.g., Cassidy & Kelly, 1991, Kelly & Bock, 1988). As in Soreno and Jongman's research, Kelly has made extensive use of lexical databases to explore relationships between phonology and syntax. For example, Cassidy & Kelly (1991) showed that the number of syllables in a word is highly correlated with the distinction between nouns and verbs. As the number of syllables in a word increases, the probability that the word is a noun also increases. When subjects were tested in an experiment in which they were asked to make a sentence containing a pseudoword, their decision to use the pseudoword as a noun or verb depended upon the number of syllables in the word. As the number of syllables increased, subjects were more likely to use the pseudoword as a noun than a verb. Once again, this result illustrates how analyzing a lexical database can prove useful in developing and testing hypotheses about linguistic performance. Indeed, Kelly (1992) cites a number of experiments by himself and others that utilize this approach and that have provided useful insights into how the sound pattern of language is not only correlated with various syntactic characteristics, but how these relationships also have behavioral consequences for listeners.

Another important factor in the organization of the lexicon is lexical stress. As summarized in Chapter II, Cutler and her colleagues (Cutler & Carter, 1987; Cutler & Norris, 1988) carried out several experiments examining the role of lexical stress in word recognition and lexical organization. Cutler and Carter (1987) analyzed two lexical databases to determine the distribution of lexical stress in English content words. They found that when word frequency was taken into account, about 80% of the words that a listener would likely encounter began with a stressed syllable. The authors hypothesized that word recognition may be optimized to take advantage of this as a segmentation strategy. The form of this strategy is 'Whenever a strong syllable is encountered, assume that it is the onset of a word.' Cutler and Norris (1988) and Butterfield and Cutler (1988; cited in Cutler, 1990) found evidence in perceptual tasks with word and nonword stimuli that listeners did appear to use a segmentation strategy based on stress. Thus, Cutler and Norris's computationally-derived predictions were supported by behavioral evidence. In short, a growing body of literature indicates that the structural analyses of words and their relations provides a fruitful approach to understanding the variables that affect word recognition and, in turn, that affect the organization of words in the mental lexicon.

Lexical Statistics in Developmental Studies

Developmental findings. Computational procedures have been used primarily to answer questions about the adult lexicon. Charles-Luce and Luce (1990) is the only developmental study in which researchers have used computational procedures to analyze child language samples. In an effort to determine if the organization and structure of the child's lexicon is similar to that of the adult's lexicon, Charles-Luce and Luce analyzed a corpus of spoken words produced by 5- and 7-year-olds (Wepman & Hass, 1969) using the same neighborhood similarity metric as Luce (1986a). First, they found that neighborhood size was substantially reduced in the children's lexicons compared to the adult lexicon. Second, a developmental trend was observed. 7-year-olds had denser lexical neighborhoods than 5 year-olds, presumably because of additional words in the 7-year-olds' lexicons. Finally, Charles-Luce and Luce found that the proportion of short words was much larger in the children's lexicons than in the adult lexicon.

In discussing their results, Charles-Luce and Luce (1990) raised the possibility that because children's lexicons contain a relatively small number of words, the representation of these words may not require the same degree of specificity or detail as required to specify words in the adult lexicon. The possibility that children might utilize a nonphonemic representational system is one of the questions examined in the present investigation. If words in the child's lexicon have many fewer neighbors than words in the adult lexicon, then, for purposes of lexical retrieval, a less detailed representational system would suffice to maintain each lexical item distinct from other lexical items in long-term memory. Similarly, as development proceeds and additional words are added to the child's lexicon, the increased size of the lexicon may serve as an impetus for the child to develop more detailed, adult-like representations of the words in the lexicon in order to keep individual items distinct in memory (cf. Jusczyk, 1986; Walley, in press).

Unresolved issues in earlier work on lexical representation. In the only previous study that examined neighborhood structure in children, Charles-Luce and Luce (1990) used language samples from children aged 5 and 7 years and assumed that children utilized a phonemic representational system in their analyses. Their findings raised some interesting questions. First, what is the structure of the lexicon in children at ages younger than 5 years? Does neighborhood density change gradually or in discrete stages? Are changes in neighborhood density associated with any developmental milestones? Second, what effect would nonphonemic representation have on neighborhood density? Would is be possible for a nonphonemic representational system to produce neighborhood densities comparable to those obtained using a phonemic representational system. Finally, would there be any advantage to carrying out these kinds of analyses using a longitudinal design rather than the cross-sectional design employed by Charles-Luce and Luce?

Prior to the present investigation, answers to these questions were based on what could be referred to as 'informed speculation.' Investigators such as Ferguson and Farwell (1975) and Jusczyk (1986) proposed that when the size of the lexicon exceeds approximately 50 words, roughly around 2 years of age, there is 'pressure' to modify or replace the wholistic lexical forms of early childhood with the analyzed forms of adulthood (i.e., representations based on an adult phonological system). By the time children reach 5–7 years of age, the age range studied by Charles-Luce and Luce (1990), their lexicons may already be organized on the basis of adult phonological structure. To learn more about the proposed transition between the use of wholistic representations and segmental phonemic representations, it is necessary to study younger children. The present study addressed this problem by applying the same procedures used by Charles-Luce and Luce to the lexicons of children aged 1 1/2 – 5 years. Because children in this age range are much closer to the age proposed for the transition between wholistic and phonemic representation than the children examined by Charles-Luce and Luce, it becomes possible to evaluate the claims of theorists such as Ferguson and Farwell (1975) and Jusczyk (1986). One way in which these claims were evaluated in the present investigation was by comparing neighborhood densities computed assuming phonemic and nonphonemic representational systems. The basis for comparing the two types of representation was the similarity of neighborhood densities produced by each system. For a nonphonemic system to be viable, items should be uniquely specified in the lexicon. To the extent that a phonemic representational system represents this ideal, then a realistic nonphonemic system should produce neighborhood densities similar to the phonemic system.

The final question posed above concerned the use of cross-sectional data versus longitudinal data when studying children's lexicons. A characteristic of Charles-Luce and Luce's (1990) study was the use of a cross-sectional design in which the words produced by children at a specific age were collapsed into a single sample. Although informative, there are two reasons why it would be useful to perform the same kinds of analyses using longitudinal data. First, longitudinal data permits the assessment of change within an individual over time. When data are collapsed across subjects, individual differences are often obscured. Thus, a longitudinal design allows individual variation in lexical development to be examined. Second, longitudinal data from several children can be combined in order to assess general trends. The reliability of these trends can then be determined using inferential statistics.

In short, the goal of the present investigation was to address the issues raised above. First, the present study analyzed neighborhood structure in language samples from children between the ages of 1 1/2 and 5 years. Second, the present study compared both phonemic and nonphonemic representational systems. Finally, the present study used longitudinal data rather than cross-sectional data. The next section contains a description of the database analyzed in the present investigation, as well as details about how the data were prepared for analyses.

The CHILDES Database

Overview of the CHILDES Database

Analyses of neighborhood structure in children requires a source of child language data. The samples of child language used in the present work were obtained from the CHILDES (CHIld Language Database Exchange System) database (MacWhinney & Snow, 1985). CHILDES consists of data from child language researchers in the form of computer-readable files that contain orthographic (nonphonetic) transcripts of child language. Both cross-sectional and longitudinal data from several languages are included in the database. The language samples consist of verbal exchanges between a target child or group of children and caregivers or other children in the child's environment. In some of the corpora, the samples take the form of naturalistic observations, while in other corpora the samples are elicited in response to questions by an experimenter.

The CHILDES database is useful for investigations of lexical representation in children because it contains a large number of samples from children at different ages, ranging from approximately 1-1/2 to 7 years of age. The availability of data from young children was especially important because of the changes proposed to occur around 2 years of age. Moreover, the language samples are generally large enough (i.e., each contains several thousand words) to reduce the likelihood of sampling problems that could limit the generalizability of any analyses performed on smaller samples.

Selection of Data from the Database for Analyses

In the present investigation, only English language data from longitudinal studies were used. Data from English-speaking children were chosen because research examining neighborhood density in spoken language has generally been restricted to English (e.g., Luce, 1986a; Goldinger, Luce, & Pisoni, 1989). Another reason for choosing data from English-speaking children was to examine the effects of nonphonemic representation of lexical information on neighborhood density. The only previous research in this area assumed an English language lexicon (e.g., Shipman & Zue, 1982).

A subset of the database was selected for detailed analysis on the basis of several criteria in addition to those described above. First, the duration of the study from which the samples were chosen had to be of sufficient length to assess developmental trends over several years, if possible. Second, the samples of child language had to include data from the period around 2 years of age, if possible. Third, the samples had to be available at fairly regular intervals (e.g., samples collected every month, every two months, etc.). Fourth, the samples had be of sufficient size to permit the inclusion of a reasonable range of utterances produced by the child, set at a minimum of 4,000 utterances every 2 months. Application of these criteria resulted in the selection of data from five children. Three children—Adam, Eve, and Sarah—originally took part in a long-term study carried out by Roger Brown and his colleagues to assess syntactic development (Brown, 1973). One child, Peter, was studied by Lois Bloom in research designed to assess the relationship between semantic and syntactic development (Bloom, 1973). The final child, Nina, was studied by Patrick Suppes (CHILDES documentation). The age ranges spanned by the five children are shown in Table 4.1. Data from the earliest age range are from Eve at 1;06/1;07. The next earliest age is Nina and Peter at 1;10/1;11. Thus, the samples selected contain data from three children prior to 2 years of age. Data from two other children, Adam and Sarah, begin after 2 years and continue until 5 years. As shown in Table 4.1, complete coverage is lacking across all ages

Table 4.1

Range of ages spanned by children selected from CHILDES database

Age			Child		
Year;Month	Nina	Adam	Eve	Sarah	Peter
1;06/1;07			•		
1;08/1;09			•		
1;10/1;11	•		•		•
2;00/2;01	•		•		•
2;02/2;03	•	•	•		•
2;04/2;05	•	•			•
2;06/2;07		•			•
2;08/2;09	•	•		•	•
2;10/2;11	•	•		•	•
3;00/3;01	•	•		•	•
3;02/3;03	•	•		•	
3;04/3;05		•		•	
3;06/3;07		•		•	
3;08/3;09		•		•	
3;10/3;11		•		•	
4;00/4;01		•		•	
4;02/4;03		•		•	
4;04/4;05		•		•	
4;06/4;07		•		•	
4;08/4;09		•		•	
4;10/4;11		•		•	
5;00/5;01				•	

for all children. However, the samples overlap; for each sampling period between 1 1/2 and 5 years, data from at least two children can be analyzed.

Insert Table 4.1 here

Preparation of the Data

Once a set of files had been selected for analyses, several additional steps remained before the computational analyses could begin. First, all the utterances produced by each child were extracted from the data files. Each data file consisted of a transcript of the utterances produced by various caregivers or other children in addition to utterances produced by the target child. Thus, it was necessary to separate the target child's utterances from other utterances. In order to maintain a consistent-sized sample of 4,000 words from each child during each sampling period, an equivalent number of words were extracted from each file so that the total summed to 4,000 words, where 'words' refer to tokens rather than types. This meant that if there were five files for a given 2 month sampling period, the first 800 words were extracted from each file (5 × 800 words = 4,000 words). If there were only two files from a sampling period, the first 2,000 words were extracted from each file (2 × 2,000 words = 4,000 words). A frequency count was done at this time so that the relationships between word frequency and lexical representation could also be assessed in the analyses.

The second step in preparing the data for further analyses was to convert the orthographic transcriptions into some phonemic notation. A phonemic transcription was necessary to provide a baseline for purposes of comparison with other representational schemes (see Chapter V). The intent of the phonemic transcriptions was not to provide an exact description of the phonetic form of the utterances produced by a child but rather to provide a representation of what the utterance was intended to be. That is, the goal was to represent a word that the child had heard and was attempting to reproduce. For the samples that were from children approximately 3 years old and above, the utterance was more likely to reflect what the child had heard in his or her environment. However, for samples that were from the earliest age ranges, it is possible that an utterance might not accurately reflect what was originally heard or produced by the child. The problem that exists in the case of children at the earliest age range is that both perceptual and production systems may contribute to the production of utterances that are ill-formed with respect to an adult model of phonology (Maxwell, 1984). An additional constraint may be due to the orthographic transcriptions provided in the CHILDES database. The data selected for inclusion in the present analyses were originally collected to answer questions concerning semantic and syntactic development (e.g., Bloom, 1973; Brown, 1973). As a consequence, the method of transcription took pains to ensure an accurate record of the words the child attempted and the order in which they occurred. However, no effort was made to preserve the exact sound shape or phonetic form of the utterances. Thus, the present analyses are based on two assumptions: first, that the original orthographic transcriptions are an accurate record of what the child *attempted* to produce, and second, that it is possible to generate a phonemic transcription from these orthographic transcriptions that would be an accurate record of what the child had stored in long-term memory if he or she used an adult phonemic representational system. These assumptions may not be warranted for all of the words produced by the children. In general, however, they are not unreasonable assumptions given the nature and scope of the present investigation.

Phonemic transcriptions were derived from two sources. If a word produced by a child was found in Webster's Pocket Dictionary, it was assigned the transcription in the dictionary. If the word was not found in the dictionary, then I supplied a phonemic transcription for the word. Non-dictionary transcriptions were required for the inflected forms of words, proper nouns, as well as for neologisms unique to a particular child.[1] Utterances consisting of more than two reduplicated syllables that appeared to be sound play were not transcribed or included in the final analyses, nor were onomatopoetic utterances such as "bzzzz" transcribed. The number of utterances that were unusable because of these criteria was less than 2.5% of the total number of utterances in a sample and were generally less than 1% of the total number of utterances in a sample. The number of unusable utterances decreased substantially in samples from children older than 3 years.

General Characteristics of Data Selected for Analyses

Tables 4.2 and 4.3 show some of the characteristics of the data selected from the CHILDES database for the present set of analyses. The number of words (types) in each language sample from each child is shown in Table 4.2. The type/token ratio is shown in Table 4.3. The type/token ratio is the proportion of different word types to the total number of words in a sample and can be thought of as one measure of the complexity or diversity of the child's language, specifically, lexical diversity. Ideally, as the number of words in the child's lexicon increases, the number of types in a fixed-size sample of language also grows larger. In turn, the type/token ratio becomes larger. However, an examination of Table 4.3 (and to a lesser extent in Table 4.2) reveals differences in the rate at which words are added to the lexicon in different children. In addition, Table 4.2 shows a non-monotonic increase in the type/token ratio within individual children.

Insert Tables 4.2 and 4.3 here

The fact that children display variability in the rate at which they add new words to their lexicons poses a problem when comparing lexical development across different children. The variability makes it difficult to draw statistically reliable conclusions if the goal is to find systematic changes across time. One way to compensate for this problem is to align the samples from the various children by type/token ratio values rather than by age. This is an imperfect solution to the problem of variability, but it does permit a more valid comparison across children. Table 4.4 shows the samples aligned according to type/token ratio. All of the analyses reported in the present investigation utilized the alignment specified in Table 4.4. The type/token arrangement shown in Table 4.4 was used as the basis for dividing the age range spanned by the children selected from the CHILDES database into three groups that roughly corresponded to three successive age ranges: Age I – 1;06 to 2;11, Age II – 2;10 to 3;05, and Age III – 3;04 to 5;01. The choice of these age ranges was based on the longest continuous sample periods spanned by data.

---

[1] The decision to retain the inflected form of utterances in the transcriptions sometimes resulted in the inclusion of multiple instances of the same word stem (e.g., 'cat'–'cats'). Although children aged 5 years and older can utilize grammatical knowledge of inflections (e.g., Berko, 1958), it does not necessarily mean that children, especially young children, utilize a stem-plus-affix representational system as has been proposed for adults (e.g., Taft, 1979). Neighborhood densities may have been marginally inflated because of this transcription convention (i.e., a target word might have one or two additional neighbors that were derived from the same stem) but the effect was negligible due to the relatively small number of words affected.

Table 4.2

Number of words (types) in files selected from the CHILDES database

Age	Child				
Year;Month	Nina	Adam	Eve	Sarah	Peter
1;06/1;07			433		
1;08/1;09			516		
1;10/1;11	427		483		407
2;00/2;01	542		532		412
2;02/2;03	480	451	532		447
2;04/2;05	567	522			473
2;06/2;07		573			518
2;08/2;09	578	599		573	512
2;10/2;11	577	538		606	472
3;00/3;01	589	515		626	561
3;02/3;03	576	608		711	
3;04/3;05		626		627	
3;06/3;07		574		613	
3;08/3;09		654		666	
3;10/3;11		641		773	
4;00/4;01		554		692	
4;02/4;03		638		691	
4;04/4;05		593		648	
4;06/4;07		589		641	
4;08/4;09		627		710	
4;10/4;11		642		714	
5;00/5;01				613	

Table 4.3

Type/token ratio for files selected from the CHILDES database

Age	Child				
Year;Month	Nina	Adam	Eve	Sarah	Peter
1;06/1;07			.1083		
1;08/1;09			.1290		
1;10/1;11	.1094		.1208		.1022
2;00/2;01	.1355		.1330		.1025
2;02/2;03	.1200	.1122	.1493		.1112
2;04/2;05	.1418	.1299			.1178
2;06/2;07		.1425			.1289
2;08/2;09	.1445	.1490		.1429	.1274
2;10/2;11	.1443	.1337		.1373	.1174
3;00/3;01	.1473	.1281		.1557	.1396
3;02/3;03	.1440	.1512		.1699	
3;04/3;05		.1555		.1560	
3;06/3;07		.1892		.1526	
3;08/3;09		.1627		.1657	
3;10/3;11		.1595		.1690	
4;00/4;01		.1378		.1721	
4;02/4;03		.1587		.1719	
4;04/4;05		.1475		.1612	
4;06/4;07		.1465		.1595	
4;08/4;09		.1560		.1766	
4;10/4;11		.1597		.1776	
5;00/5;01				.1526	

Table 4.4

Alignment of data files according to type/token ratio

		Child			
Nina	Adam	Eve	Sarah	Peter	
				.1022	
				.1025	
.1094	.1122	.1083		.1112	
.1355	.1299	.1290		.1178	
.1200	.1425	.1208		.1289	**Age I**
.1418	.1490	.1330		.1274	
.1445	.1337	.1493	.1429	.1174	
.1443	.1281		.1373	.1396	
.1473	.1512		.1557		**Age II**
.1440	.1555		.1699		
	.1892		.1560		
	.1627		.1526		
	.1595		.1657		
	.1378		.1690		
	.1587		.1721		**Age III**
	.1475		.1719		
	.1465		.1612		
	.1560		.1595		
	.1597		.1766		
			.1776		
			.1526		

Table 4.5

Phonetic symbols and their computer-readable counterparts

p	p
t	t
k	k
b	b
d	d
g	g
C	tʃ
J	dʒ
s	s
S	ʃ
z	z
Z	ʒ
f	f
T	θ
v	v
D	ð
h	h
n	n
G	ŋ
m	m
l	l
r	r
w	w
y	y
L	syllabic ‘l’
M	syllabic ‘m
N	syllabic ‘n’
i	i
I	I
E	ɛ
e	e
@	æ
a	a
W	aU
Y	aI
^	ʌ
O	oI
o	o
U	ʊ
u	u
R	ɚ

Insert Table 4.4 here

Analysis of Neighborhood Structure Assuming Phonemic Representation

Overview of Neighborhood Density Analyses

The nature of the lexical representation in the child language samples was assessed by measuring neighborhood density. The neighborhood density for each word produced by the children was determined using the same neighborhood similarity metric developed by Luce (1986a) and used by Charles-Luce and Luce (1990). The precise form of this metric was as follows. First, the number of segments in the target utterance was determined. This information was used to select other utterances from the sample that were of the same length or that differed in length from the target utterance by only one segment. Next, each of these potential neighbors was compared to the target utterance. If the target and comparison utterance differed by only one segment, such as in /kIn/ versus /ken/, /kIn/ versus /skIn/, or /kIn/ versus /In/, then the comparison utterance was considered to be a neighbor of the target. As each target was compared to the other utterances in the sample, the number of neighbors for that target was tallied. This process was repeated for each utterance in the sample.

Table 4.5 shows the phonetic symbols used in the present investigation and their machine-readable counterparts. Neighborhood density was computed separately for each child within each language sample. For example, for a given child, neighborhood density was determined for each word at age 1;6/1;7, and then again at age 1;8/1;9, and so on. Thus, an independent estimate of neighborhood structure was computed in successive language samples from each child.

Insert Table 4.5 here

A tacit assumption in these analyses is that the child represents the words he or she hears in the mental lexicon in some segmented form and that the internal organization of these words consists of a string of phonemes. However, as noted in Chapter III, substantial evidence exists to suggest that children, especially children younger than about 3 years of age, do not appear to utilize a segmental phonemic representational system for words in their language. The assumption of phoneme-sized units for representation and the fact that there is little evidence for their existence seems at first somewhat incongruous. But, it should be kept in mind that it was the *relationship* among different representational systems, not the absolute number of neighbors that each type of representational system produces that is under examination in the present investigation. Even though children may not use a representational system based on phoneme-sized units, it should be possible to obtain measures of lexical representation and structure regardless of the exact form of the representational system.

This section contains three major divisions corresponding to the three age ranges of children from the CHILDES sample: Age I – 1;06 to 2;11, Age II – 2;10 to 3;05, and Age III – 3;04 to 5;01. The general form of the statistical analysis used for each age range consisted of a two-way repeated-measures analysis of variance in which mean neighborhood density was the dependent variable. Mean neighborhood density was computed by taking the average of the neighborhood densities for a set of words defined by

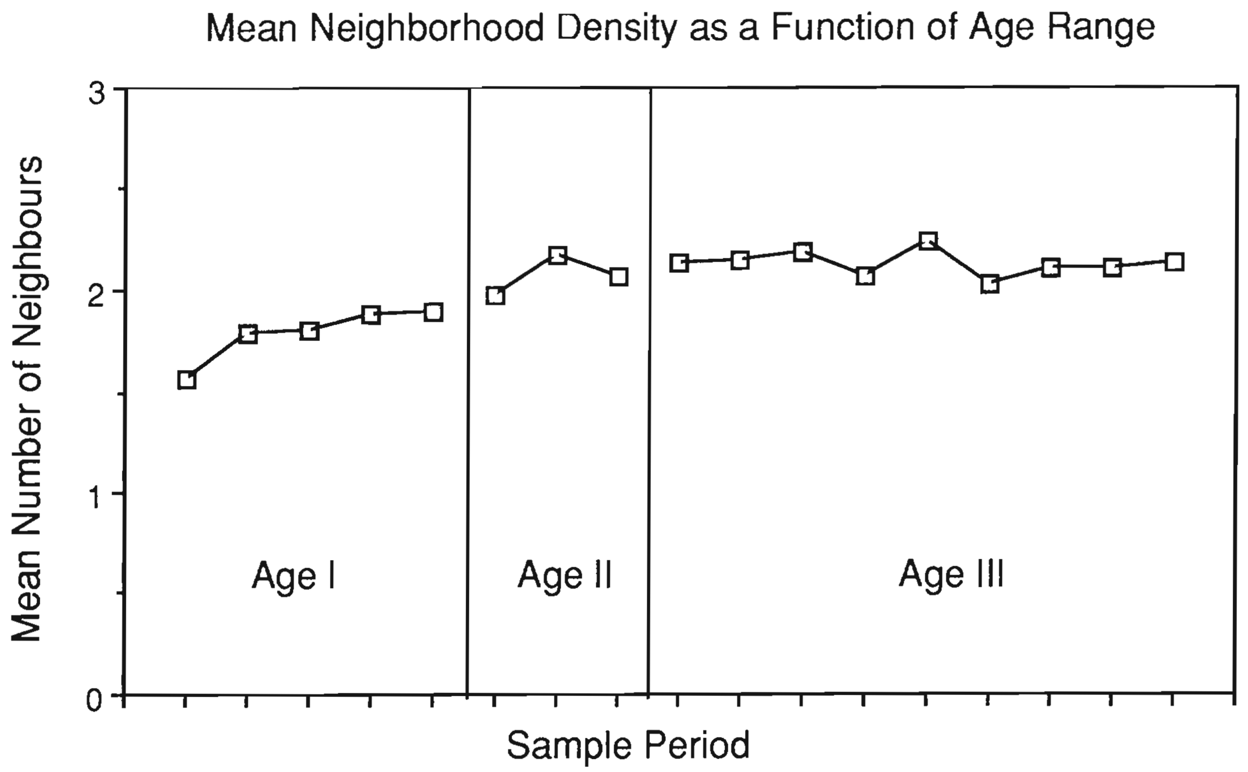

Figure 4.1 Mean neighborhood density as a function of age range. Neighborhood density was computed assuming phonemic representation. Age I data are shown in the left panel, Age II data in the center panel, and Age III data in the right panel. The data in each panel are from different children.

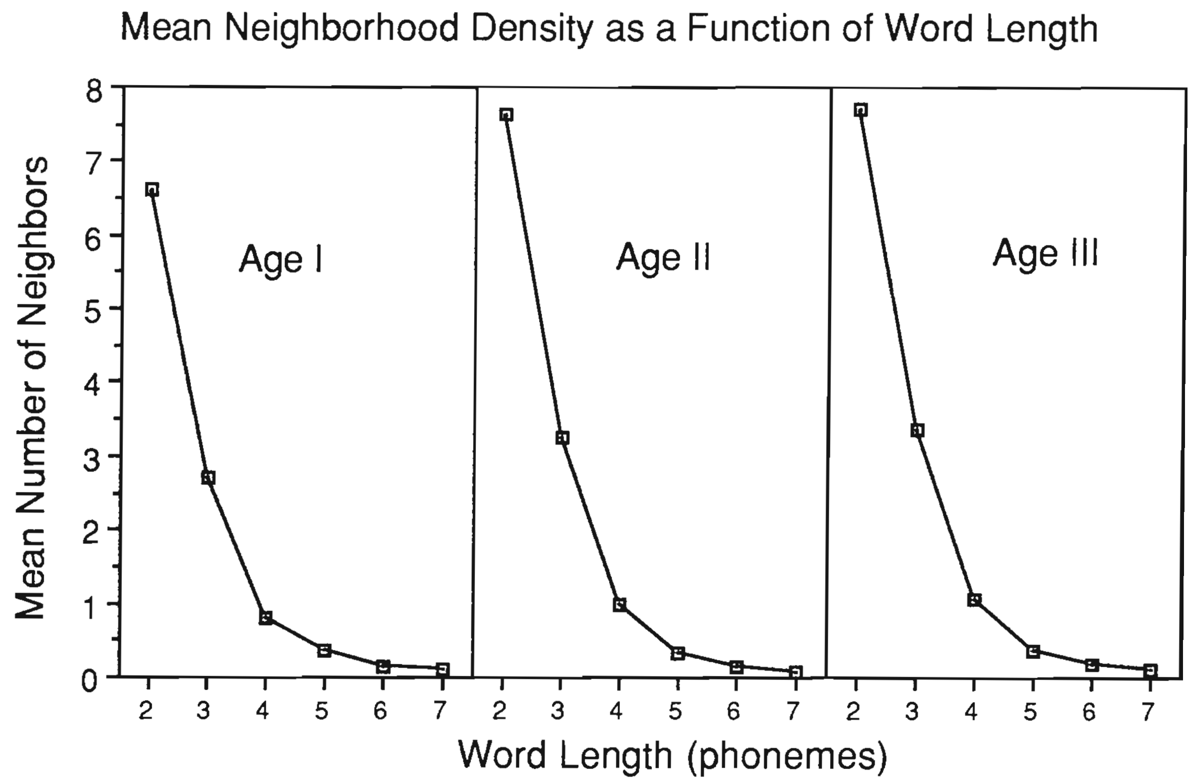

Figure 4.2 Mean neighborhood density as a function of word length measured in phonemes. Age I data are shown in the left panel, Age II data in the center panel, and Age III data in the right panel. Each panel contains data from a different group of children.

the independent variables included in each ANOVA. The two variables in each analysis were sampling period (roughly corresponding to age) and length (i.e., the number of phonemes in a word). The number of samples varied in each of the three age ranges examined. Words included in the analyses were between two and seven phonemes in length. No words longer than seven phonemes were included in the analyses because they comprised only a minuscule proportion of the total number of words in each sample (cf. Charles-Luce and Luce, 1990).

The following questions were addressed in these analyses. First, what kinds of changes, if any, occur in mean neighborhood density between the ages of 1;06 and 5;01 when a segmental phonemic representational system is assumed? That is, what is the developmental sequence for changes in neighborhood density? Second, if changes in neighborhood density over time are observed, do they also vary with word length? Third, how do the results of the present investigation compare with earlier work by Charles-Luce and Luce (1990)?

Age Range I — 1;06-2;11

The sampling period covered in the analysis of data from Age I was 1;06 to 2;11, and included five samples. Data from four children were included in this analysis: Nina, Adam, Eve, and Peter. The left panel of Figure 4.1 shows mean neighborhood density plotted as a function of sample period at Age I. Mean neighborhood density in Figure 4.1 was computed over words two to seven phonemes in length. Because most long words had few neighbors, mean neighborhood density was reduced accordingly. Nevertheless, during the Age I period mean neighborhood density generally increased as a function of successive samples, suggesting that as more words are added to the child's lexicon, neighborhood density increases. Although the range of the increase was small (an increase of 1.54 to 1.85 neighbors between sample period 1 and sample 5), the effect was statistically reliable, $F(4, 12) = 4.53$, $MSE = .097$, $p < .05$.

Insert Figure 4.1 here

The left panel of Figure 4.2 plots mean neighborhood density as a function of word length as measured in phonemes at Age I. As expected, the average number of neighbors decrease substantially as words increased in length. For two-phoneme words the mean number of neighbors was 6.75, whereas for words four phonemes or longer, the mean number of neighbors was less than 1.0. The effect of word length was statistically reliable, $F(5, 12) = 1020.56$, $MSE = .127$, $p < .00$. Possible reasons for the effect of word length include the following. First, it is much easier for a word to qualify as a neighbor of another word if the words are short. As words increase in length, the probability of finding another word that qualifies as a neighbor is reduced, primarily because of increasing phonotactic constraints. Second, the child's lexicon contains a higher proportion of short words than long words. Thus, the pool of potential neighbors for long words is reduced even further. Finally, the analysis revealed no interaction between sample period and word length in the data from Age I. Mean neighborhood density increased at the same rate for both long and short words during this period of development.

Insert Figure 4.2 here

Age Range II — 2;10-3;05

The sampling period covered in the analysis of data from Age II was between 2;10 and 3;05, and included three samples. Data from three children were included in this analysis: Nina, Adam, and Sarah. The center panel of Figure 4.1 shows mean neighborhood density plotted as a function of sample period at Age II. In contrast to the data from Age I, the data from Age II shows an inconsistent effect of successive samples on mean neighborhood density. Although neighborhood density increases from sample 1 (1.95 neighbors) to sample 2 (2.13 neighbors), it decreases from sample 2 (2.13 neighbors) to sample 3 (2.05 neighbors). In keeping with the data shown in Figure 4.1, no reliable effect of sample period on mean neighborhood density was obtained in the ANOVA. Possible reasons for the lack of an effect associated with increasing age include the small sample size (only three children were included in this analysis), the small number of sample periods, and variability among the samples themselves. In terms of sample variability, an examination of Table 4.4 reveals substantial variability in the type/token ratios within individual children from this age range. In turn, variability in type/token measures can be assumed to influence the variability of neighborhood density estimates.

The center panel of Figure 4.2 shows the mean neighborhood density plotted as a function of word length in phonemes. As observed in the Age I data, the average number of neighbors was substantially smaller for longer words compared to short words. For two-phoneme words, the mean number of neighbors was 7.6, whereas for words four phonemes or longer, the mean number of neighbors was less than 1.0. This effect was statistically reliable, $F(5, 10) = 120.09$, $MSE = .662$, $p < .001$. The interaction between sample period and word length was not reliable for the data from Age II. This result is consistent with the observation that mean neighborhood density was similar for both long and short words during this period of development. The lack of interaction between sample period and word length was also observed during Age I.

Age Range III — 3;04-5;01

The sampling period covered in the analysis of data from Age III was between 3;04 and 5;01, and included nine samples. Data from two children were included in this analysis: Adam and Sarah. The right panel of Figure 4.1 shows mean neighborhood density plotted as a function of sample period at Age III. The data in Figure 3 show an inconsistent effect of successive samples on mean neighborhood density similar to that observed at Age II. Mean neighborhood density ranged from 2.14 in the first sample to 2.12 in the final sample. The highest mean neighborhood density value was 2.25 at sample period 5 while the lowest value was 2.03 at sample period 6. In general, neighborhood density appears not to change during Age II. And, similar to what was observed at Age II, no reliable effect of sample period on mean neighborhood density was obtained in the ANOVA. Although it is possible that growth in neighborhood density reaches a temporary plateau between the ages of 3 and 5 years, variability among individual children and the small sample size may also account for this null result.

The right panel of Figure 4.2 shows mean neighborhood density plotted as a function of word length in phonemes at Age III. The average number of neighbors was substantially smaller for longer words compared to shorter words. The same result was also obtained at Ages I and II. Mean neighborhood density ranged from 7.7 for two-phoneme words to 1.0 or less for words four phonemes or less in length. This effect was statistically reliable, $F(5, 5) = 70.90$, $MSE = 2.268$, $p < .001$.

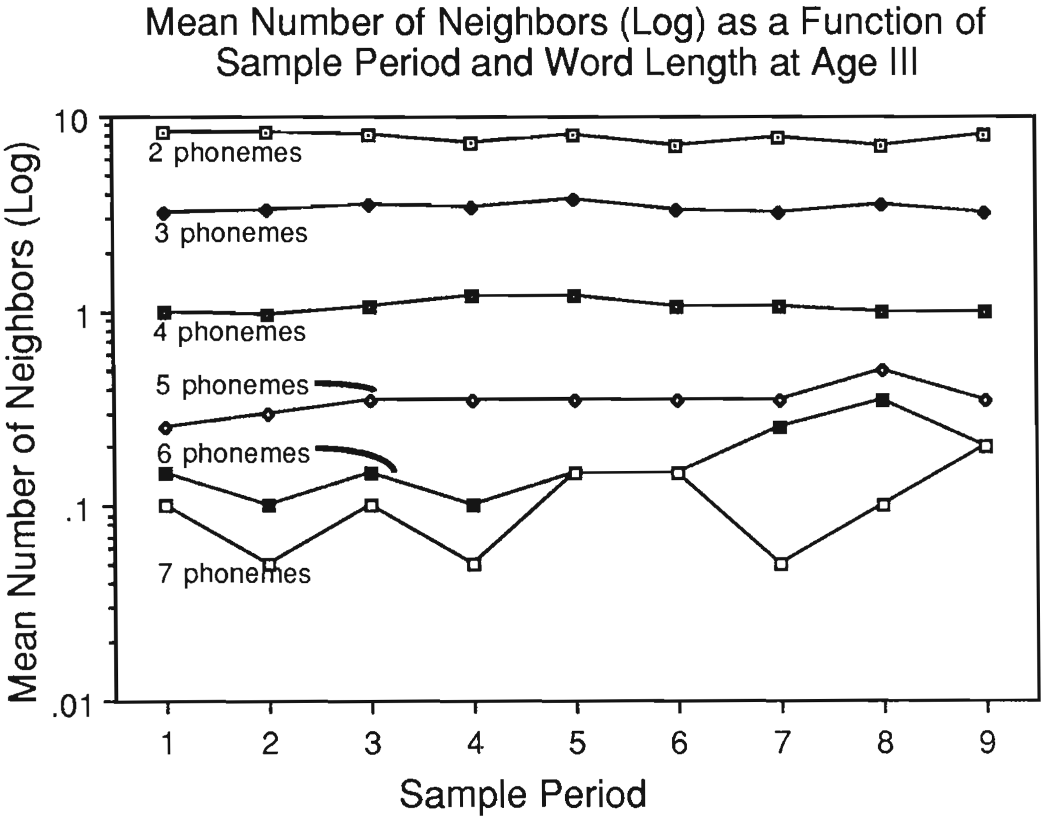

Figure 4.3 Mean neighborhood density as a function of sample period and word length at Age III. Note the use of a log scale for neighborhood density.

Insert Figure 4.3 here

Figure 4.3 shows mean log neighborhood density as a function of sample period and word length. The ANOVA revealed a significant interaction between sample period and word length, $F(40, 40) = 2.81$, $MSE = .0297$, $p < .001$. Figure 3 indicates that this effect is due to small variations in neighborhood density for five-, six-, and seven-phoneme words whereas neighborhood density for the shorter words (i.e., words less than five phonemes in length) tends to be relatively constant as a function of sample period. Plotting the data in log form enables these small variations in neighborhood density for long words to be observed. Unfortunately, no stable patterns emerge as a function of sample period, making this interaction difficult to interpret in terms of any theoretical perspective. Nevertheless, two interpretations of this result are possible. First, compared to short words, long words may have a greater tendency to drop in and out of use by the child. Second, sampling error may make conclusions based on the small number of five-, six-, and seven-phoneme words problematic. Data from additional children at this age would eliminate the latter possibility.

In summary, analyses of neighborhood density in children between 1 1/2 and 5 years revealed that the largest change in mean neighborhood density occurred during the Age I period. Mean neighborhood density increased only slightly or remained constant during Age II and Age III. The implication of this finding is that the period around two years of age coincides closely with the onset of the naming explosion (Gopnick & Meltzoff, 1987). As the child's lexicon increases in size, neighborhood density also increases. Thus, the child not only acquires more words, but the words become more similar to each other. In turn, increased similarity among words in the mental lexicon necessitates more detailed representations and a more sophisticated organization (Jusczyk, 1986; Walley, in press).

What is intriguing about this finding is that changes in mean neighborhood density were only manifested during Age I. Although the naming explosion begins about 18 months of age, several researchers have asserted that children continue to add five to ten words to their lexicons daily until 6 years of age (e.g., Gleitman & Gleitman, 1992). Thus, one would expect neighborhood densities to continue increasing during Age II and Age III. The reason for this lack of increase in neighborhood densities during later age periods is unknown at the present time. Even when neighborhood density was considered as a function of word length, there was little systematic effect of successive sample periods during Age II and Age III.

The lack of age-related change observed during Ages II and III could be due to several causes. First, data from only two or three children may have provided insufficient statistical power, thus resulting in a null result for the effect of sample period. Obtaining data from additional children would enable this explanation to be evaluated. Alternatively, the well-known disparity between receptive and productive vocabulary may be responsible for the null result of sample period during Age II and Age III (Benedict, 1979). Children may be acquiring words at a rapid rate before 6 years, but they may not necessarily produce them at the same rate. The difference in results between Age I versus Age II and Age III could have occurred because the earliest words that children learn, such as those learned during Age I, form a core production vocabulary and this vocabulary is used productively for most general interactions with individuals in their environment throughout the three age ranges. The reason that an increase in mean neighborhood density was observed during Age I is that children began with no words in their lexicon and added the minimum number necessary to efficiently communicate with those in their environment. Hence, any increase in the number of words present, up to this minimal amount, will be reflected by an increase in mean neighborhood density. Although many additional words are learned after this initial

Table 4.6

Calculation of number of three–, four–, and five–phoneme words in Charles-Luce and Luce (1990) .

Age	Total number of Words	Length 3	 4	 5
5	679	40% (272)	25% (170)	11% (75)
7	943	37% (349)	27% (255)	12% (113)

Table 4.7

Calculation of mean neighborhood densities for three–, four–, and five–phoneme words in Charles–Luce (1990).

		Number of Neighbors						
		0-1	2-3	4-5	6-7	8-9	10-11	12-13
Length 3	5 years	17%	32%	22%	19%	8%	–	2%
		(46)	(87)	(60)	(52)	(22)	–	(5)
	Mean neighborhood density (approximate) = 1098/272 = 4.04							
	7 years	14%	22%	28%	17%	14%	3%	2%
		(49)	(77)	(98)	(59)	(49)	(10)	(7)
	Mean neighborhood density (approximate) = 1650.5/349 = 4.73							

		0-1	2-3	4-5	6-7	8-9	10-11	12-13
Length 4	5 years	81%	17%	2%	–	–	–	–
		(138)	(29)	(3)	–	–	–	–
	Mean neighborhood density (approximate) = 155/170 = .91							
	7 years	68%	28%	3%	1%	–	–	–
		(173)	(71)	(8)	(3)	–	–	–
	Mean neighborhood density (approximate) = 319.5/255 = 1.25							

		0-1	2-3	4-5	6-7	8-9	10-11	12-13
Length 5	5 years	100%	–	–	–	–	–	–
		(75)	–	–	–	–	–	–
	Mean neighborhood density (approximate) 37.5/75 = .50							
	7 years	95%	5%	–	–	–	–	–
		(107)	(6)	–	–	–	–	–
	Mean neighborhood density (approximate) = 68.5/113 = .61							

Insert Tables 4.6 and 4.7 here
----

period, their use in normal conversational settings may be limited. If this explanation is correct, then it is understandable that no effect of neighborhood density would be observed as a function of successive sample periods during Age II and Age III because the data used in the present investigation were collected under conditions that did not seek to maximize the extent of vocabulary produced by the children. If data were available that provided a more accurate picture of children's receptive vocabulary, it would likely show an effect of successive sample size on mean neighborhood density.

Comparison of Present Findings with Earlier Studies

A comparison of the results obtained in the present investigation with the findings obtained by Charles-Luce and Luce (1990) are possible, albeit not straightforward. The difficulty arises because of different methods used in the two studies. Charles-Luce and Luce used composite data from two groups of 30 children, whereas the present investigation used longitudinal data from five children. In addition, Charles-Luce and Luce considered a narrower range of word lengths (3-, 4-, and 5-phoneme words) than the present analysis (2- to 7-phoneme words). Most importantly, Charles-Luce and Luce analyzed their data using frequency distributions in which the percentage of words in the lexicon was plotted as a function of neighborhood density, whereas in the present investigation, mean neighborhood densities were computed. Overall, differences between the two studies require that the data in one study be transformed into an alternate form, one that is comparable to the form of the data from the other study.

In order to facilitate comparison with the results of the present investigation, the data in Charles-Luce and Luce were converted to mean neighborhood densities. First, the proportion of three-, four-, and five-phoneme words in the samples used by Charles-Luce and Luce was used to compute the number of three-, four-, and five-phoneme words in their sample. These calculations are shown in Table 4.6.

After computing the total number of words for each age and length, the mean neighborhood density was computed for each age and length, as shown in Table 4.7. The percentage of three-, four-, and five-phoneme words were obtained from the figures in Charles-Luce and Luce. For each age and for each length, the following procedure was used. From the percentage of words at different neighborhood densities, the number of words at different neighborhood densities was computed. These values were then multiplied by the midpoint of the density ranges used by Charles-Luce and Luce (e.g., 0–1 = .5, 2–3 = 2.5, 4–5 = 4.5, etc.). The products of each of these operations were summed. This sum was then divided by the total number of words of that length (obtained from the estimates shown in Table 4.6). The results of these procedures provided an approximation of the values described in Charles-Luce and Luce (1990) in terms of the mean neighborhood densities used in the present investigation.

The data from the present investigation that is most comparable to the data from Charles-Luce and Luce is from the Age III period (3;04-5;01), shown in Figure 4.2. In the present investigation, the mean neighborhood densities among three-, four-, and five-phoneme words for the age representative of 5-year-olds were 3.37, 1.06, and .35, respectively. In Charles-Luce and Luce, the mean neighborhood densities for three-, four-, and five-phoneme words for 5-year-olds were 4.04, .91, and .50, respectively. The similarity of the estimates of neighborhood density in the two studies is remarkable, especially considering the different sources from which each sample was obtained. In short, the results of the present investigation for 5-year-olds essentially replicate Charles-Luce and Luce's results for 5-year-olds. The similarity of these results validate the extrapolation of these methods of analysis to younger children.

The Relationship Between Neighborhood Density and Other Lexical Variables

As described earlier, Landauer and Streeter (1973) found that high frequency words tended to have more orthographic neighbors than low frequency words. This finding was replicated and extended in a more comprehensive analysis by Luce (1986a) in which he examined the sound patterns of 20,000 words. Perhaps the simplest explanation of this relationship is that high frequency words tend to be shorter in length than low frequency words (cf. Cassidy & Kelly, 1991; Zipf, 1949). Because high frequency words are generally short, the reduction in phonotactic constraints increases the likelihood of a word qualifying as a neighbor of another word (see Luce, 1986b). Conversely, because low frequency words are generally long, the increase in phonotactic constraints reduces the likelihood of a word qualifying as a neighbor of another word. The relationship among neighborhood density, word frequency, and word length appears to be a robust property of the adult lexicon. Because one of the goals of the present investigation was to assess lexical organization in young children, I was interested in determining whether a similar relationship existed in children's lexicons. In this section, I describe the results of analyses designed to address this issue.

At the time that individual words were extracted from the original CHILDES files, word frequency was also calculated. Thus, for every word for which neighborhood density was computed, I also had word frequency information. In addition, during the course of computing neighborhood density, I also computed mean neighborhood frequency, a measure that was shown by Luce and his colleagues to be important in spoken word recognition in adults (e.g., Goldinger et al., 1989; Luce, 1986a). Pearson product moment correlations were used to determine the strength of the relationship among neighborhood density, word frequency, and mean neighborhood frequency, as well as word length. A separate series of correlations were carried out for words from Age I, Age II, and Age III. A fifth variable, sample period, was also included in the analyses to determine if there were any changes in the relationship among the lexical variables within each of the three age ranges. Finally, correlations among the same variables (with the exception of sample period) were also carried out using the 20,000 word *Webster's Pocket Dictionary*. This was done in order to determine if the relationship among these lexical variables was of comparable magnitude in adult and child lexicons.

Insert Table 4.8 here

Table 4.8 contains a series of correlation matrices corresponding to Age I, Age II, Age III, and the adult lexicon. Within the matrices for the children's lexicons, correlations among neighborhood density ("Density"), log word frequency ("Log Freq"), neighborhood frequency ("NH Freq"), word length ("Length"), and sample period ("Sample") are shown. Within the matrix corresponding to the adult lexicon, correlations among neighborhood density ("Density"), word frequency ("Freq"), log word frequency ("Log Freq"), neighborhood frequency ("NH Freq"), and word length ("Length") are shown. Rather than describe each matrix individually, I will outline several trends common to all of the groups.

First, a small but reliable ($p < .05$) positive correlation was observed between log word frequency and neighborhood density. The magnitude of this correlation increased most

Table 4.8

Correlation matrix for lexical variables – Age I

	Density	Log Freq	NH Freq	Length	Sample
Density	1.000				
Log Freq	.266	1.000			
NH Freq	.404	.175	1.000		
Length	-.627	-.273	-.388	1.000	
Sample	.041	-.060	-.009	.014	1.000

Correlation matrix for lexical variables – Age II

	Density	Log Freq	NH Freq	Length	Sample
Density	1.000				
Log Freq	.349	1.000			
NH Freq	.488	.270	1.000		
Length	-.633	-.339	-.471	1.000	
Sample	.016	-.034	-.036	.019	1.000

Correlation matrix for lexical variables – Age III

	Density	Log Freq	NH Freq	Length	Sample
Density	1.000				
Log Freq	.389	1.000			
NH Freq	.482	.300	1.000		
Length	-.637	-.361	-.471	1.000	
Sample	-.019	.008	.000	.016	1.000

Note: All correlations in the Age I, Age II, and Age III matrices, with the exception of those associated with the variable Sample, were significant, $p < .05$ (Bonferroni method).

Table 4.8 (continued)

Correlation matrix for lexical variables – Adults

	Density	Freq	Log Freq	NH Freq	Length
Density	1.000				
Freq	0.090	1.000			
Log Freq	0.327*	0.226*	1.000		
NH Freq	0.234*	0.245*	0.174*	1.000	
Length	-0.598*	-0.073	-0.271*	-0.194*	1.000

* $p < .05$ (Bonferroni method)

between Age I (r = .266) and Age II (r = .349). At Age III, the size of the correlation increased even further (r = .389). In comparison, the correlation between log word frequency and neighborhood density for the adult lexicon is roughly in the same range as the Age II and Age III values (r = .327). (The same result was not obtained for raw word frequency, presumably because a log function more accurately characterizes the distribution of words in the lexicon according to frequency.) In summary, the relationship between word frequency and neighborhood density becomes stronger during development and is in the same direction and of comparable magnitude as the adult data.

Second, a reliable positive correlation was found between neighborhood frequency and neighborhood density. The correlations ranged between .404 (Age I) and .488 (Age II) in the children's lexicons. The correlation was slightly less in the adult lexicon (r = .327) . In addition, there was also a reliable positive correlation between log frequency and neighborhood frequency, ranging between .175 (Age I) and .300 (Age (II) in the children's lexicons, to .174 in the adult lexicon. Thus, not only do high frequency words tend to have more lexical neighbors than low frequency words, the neighbors of high frequency words tend to be other high frequency words (see Luce, 1986b).

Third, word length has a substantial effect on other lexical variables, especially in the children's lexicons. A strong negative correlation between word length and neighborhood density was observed. The values ranged between -.637 (Age III) and -.627 (Age I) in the children's lexicons to -.598 in the adult lexicon. Similarly, length was negatively correlated with log word frequency and neighborhood frequency in both child and adult lexicons, although the magnitude of the effect was greater in the children's lexicons. As described earlier in this section, the negative relationship between word length and other lexical variables is partly a consequence of increasing phonotactic constraints; as word length increases, the more likely it is that the word has a unique phonotactic structure. In turn, this reduces the likelihood of a long word having neighbors, as well as increasing the likelihood that what neighbors it does have will be other long words that are low in frequency.

Finally, the effect of sample period on lexical variables was virtually nonexistent. Across each of the three age ranges, the value of r associated with Sample was near zero. It appears that changes in the relationship among variables such as neighborhood density and word frequency are not correlated with successive sample periods. In other words, changes in lexical variables within an age range are negligible.

In short, all the major relationships among lexical variables observed in the adult lexicon were also observed in the present analyses of children's lexicons. In addition, a developmental trend was observed from Age I to Age III between word frequency and neighborhood density. However, no evidence of systematic changes within an Age range was found. Although these results indicate that the adult and child lexicons show some similarity to each other, the reason for the similarities between the two groups may be due to the basic properties of words in English. Phonotactic constraints and the tendency to use short words more frequently than long words can account for many of the effects observed in these correlations (e.g., Zipf's Law [Zipf, 1949]). What is intriguing is that children's lexicons display these effects despite their relatively small size. The fact that these relationships are observed even in children's lexicons attests to the robustness of the effects.

Summary

The CHILDES database was used to obtain a sample of spoken language from children between the ages of 1;06 to 5;01. Neighborhood densities were calculated for each word produced by each child using the same metric employed by Charles-Luce and Luce (1990). The results of the analyses revealed several novel findings. First, neighborhood density increased most rapidly in the period around two years of age. After this age, neighborhood density appears to change more gradually. This pattern of results observed in later periods may be due to limitations of the database from which the child language samples were obtained; only a small subset of the words actually present in the children's lexicons may have been sampled. Second, neighborhood density in short and long words tended to increase at approximately the same rate until the children reach the period between 3 and 5 years of age. During this age range, long words (five to seven phonemes in length) tend to have more unstable neighborhood densities than short words. Third, the mean neighborhood densities for 5-year-olds observed in the present investigation replicated earlier findings (Charles-Luce & Luce, 1990). Finally, the relationships among neighborhood density, word frequency, and other lexical variables become stronger as a function of age, and are roughly the same as those found in analyses of the adult lexicon.

Overall, the present results demonstrate that reliable effects can be obtained by a computational analyses of children's lexicons. Moreover, the findings described in the present chapter extend our knowledge of lexical organization in children and how it compares to the adult lexicon. Such knowledge is essential for developing a empirically-based account of word recognition, lexical organization, and lexical representation in children. However, the present chapter leaves unanswered questions about alternatives to segmental representational systems and how they compare to the phonemic system. Because young children may not possess a system of lexical representation and organization that corresponds to a phonemic system, I examine several nonphonemic representation systems in Chapter V.

Chapter V. Computational Analyses of Child Language: Nonphonemic Representation

The discussion of lexical development in Chapter III emphasized that although it is possible that children have developed a representational system for lexical information corresponding to the adult phonemic model by the age of 5 years, alternative means of representing lexical information are also possible, especially in young children. The present chapter explores several alternatives to phonemic representations for words and considers the effect of these nonphonemic representational systems on neighborhood density in the same group of children examined in Chapter IV. I begin by describing how solutions to problems in speech recognition by machine may be utilized to help understand the representation of lexical information in children. Under certain conditions it is possible to reduce the size of the search space in automatic speech recognition to a fraction of the size of the original lexicon by utilizing a coarse phonetic code and still effectively isolate an individual word. Thus, only a relatively small amount of information about a word may be necessary to uniquely identify it as distinct from other words. Next, the application of this work to understanding the representation and organization of lexical information in children is described. Finally, the results of several analyses utilizing nonphonemic representational schemes applied to the same data analyzed in Chapter IV are described.

Lexical Statistics in Speech Recognition by Machine

An analogous problem to understanding the mechanisms underlying speech perception in humans is the development of speech recognition by machine. Because the human perceptual system is remarkably adept at dealing with spoken language, several researchers in the past decade have explored how phonological structure, word frequency, and lexical stress could be utilized to construct more efficient algorithms for automatic speech recognition. Central to this approach is the idea of *search space reduction* in which a subset of the lexicon is selected for detailed analysis rather than analyzing the entire lexicon in detail (Altmann, 1990). Essentially, this amounts to eliminating from further consideration obvious mismatches between the input utterance and the contents of the system's lexicon. Search space reduction is a useful goal to pursue because it is computationally expensive to apply fine-grained acoustic analyses to all the words in the system's lexicon in order to match an input utterance. Search space reduction procedures take advantage of particularly salient acoustic-phonetic properties in utterances to form *equivalence classes.* An equivalence class is a subset of the items in the lexicon that are matched on the basis of coarsely coded acoustic-phonetic features, such as the number of constituent segments or lexical stress pattern. The words within an equivalence class are then subjected to more detailed analysis in order to isolate the identity of the input utterance. An effective search space reduction procedure minimizes the size of the equivalence classes containing the input utterance. Search space reduction procedures are discussed here is because their effectiveness has been evaluated using large lexical databases (e.g., Huttenlocher & Zue, 1984; Harrington & Johnstone, 1988). Moreover, these procedures may have some application beyond speech recognition. The use of an equivalence class that groups together words in terms of coarse acoustic features is analogous to the possibility that children initially perceive spoken language in terms of large, relatively unanalyzed units. The extent to which equivalence classes prove useful in automatic speech recognition work suggests that similar kinds of classes could be used to organize words in the lexicons of young children.

Zue and his colleagues (Shipman & Zue, 1982; Huttenlocher & Zue, 1984) examined search space reduction procedures in which words were analyzed according to a coarse coding of phonetic features and lexical stress pattern. As used here, coarse coding refers to a procedure for generating broad equivalence classes in which an input utterance is segmented according to superordinate phonetic classes, such as manner class, place, or voicing. It is relatively easy to segment an utterance according to broad phonetic classes compared to labeling individual phonemes. For example, vowels are distinct from consonants because vowels are characterized by a well-defined formant structure while consonants have a dynamic, rapidly-changing formant structure. In contrast, uniquely identifying individual phonemes requires much more detailed information. Moreover, the information required to identify an individual phoneme tends to be less reliable than the information required to determine if the segment is a consonant or vowel or some other broad phonetic category.

Shipman and Zue (1982) examined two phoneme-based coding schemes. One scheme simply classified words according to whether each segment was a consonant or a vowel. The other scheme classified words according to manner class using the following six broad categories: (1) vowels, (2) nasals, (3) liquids and glides, (4) stops, (5) strong fricatives and affricates, and (6) weak fricatives. Manner class categories are perceptually salient categories that can be reliably extracted from utterances by adult listeners (Miller & Nicely, 1955). Shipman and Zue assessed the effectiveness of coarse coding by calculating the size of the equivalence classes formed by various coarse schemes for words from Webster's Pocket Dictionary.

They found that for the simplest coding scheme, in which segments were simply coded as either consonants or vowels, the average size of an equivalence class was 26 words. The pattern CVC yielded the largest equivalence class, containing 1,500 words out of the 20,000 word lexicon. When the selection of words was weighted by word frequency (i.e., each cohort was weighted according to how often each consonant-vowel pattern occurs in the Kucera & Francis (1967) word count), the average size of an equivalence class was almost an order of magnitude larger. Unfortunately, for certain high frequency patterns such as the pattern CVC, the simplest coding scheme combined with frequency of occurrence information was only marginally better than searching all 20,000 words.

Shipman and Zue (1982) found that the six-way coding scheme was substantially better than the simple consonant-vowel scheme. The average size of an equivalence class in the 20,000 word lexicon was two words (compared to 26 for the consonant-vowel scheme). The pattern [stop-vowel-stop] produced the worst performance, with an equivalence class size of 200 words. Again, incorporation of word frequency resulted in an increase in cohort size by an order of magnitude. Nevertheless, the six-way coding scheme resulted in a substantial reduction in the size of the search space, especially when compared to the simple two-way coding scheme. Shipman and Zue concluded that a coarse coding strategy could be used to reduce the search space for recognition of isolated spoken words.

Four smaller-sized lexicons containing 1,250, 2,500, 5,000, and 10,000 words that were subsets of Webster's Pocket Dictionary were also analyzed. Words in the small lexicons were chosen on the basis of frequency; all the words in the small lexicons were the most frequent words in Webster's Pocket Dictionary. The results of the additional analyses are important because these subsets are similar in size and content to the lexicons of children. Children's lexicons are not only smaller than the adult lexicon but they also contain many words that are high frequency words in the adult lexicon (cf. Schwartz & Terrell, 1983). Shipman and Zue found essentially the same pattern of results with the small lexicons as with the full-sized lexicon, especially for the consonant-vowel coding scheme. However, the size of the equivalence classes was substantially reduced in the small lexicons compared to the 20,000 word lexicon. The implication of this finding for small-sized lexicons, such as those found in children, is that a coarse-grained representational scheme could conceivably permit the unique specification of lexical items.

Huttenlocher and Zue (1983) extended Shipman and Zue's (1982) findings in several ways. First, they utilized frequency-weighted lexical statistics that provided a more accurate picture of the size of equivalence classes than the values obtained by Shipman and Zue (1982).

These statistics were calculated by taking each word in an equivalence class and multiplying the word by its frequency of occurrence in the Kucera and Francis (1967) word count. Median equivalence class size increased from four to 25 words. Huttenlocher and Zue also evaluated the effectiveness of lexical stress in reducing the size of the search space. Their goal was to determine whether stress information could be used to augment the coarse code utilized by Shipman and Zue (1982). Huttenlocher and Zue noted that, despite the robust nature of manner class categorization, there was still potential for mislabeling segments. One strategy for reducing the likelihood of mislabeled segments is to ignore unstressed syllables because segments in unstressed syllables are more likely to be deleted or reduced than segments in stressed syllables. Huttenlocher and Zue examined the words in Webster's Pocket Dictionary under two conditions. First, each stressed syllable was specified according to manner class and each unstressed syllable was replaced with a cover symbol. Second, each stressed syllable was replaced with a cover symbol and each unstressed syllable was specified according to manner class. They found that the manner class information in stressed syllables was more constraining than the same information in unstressed syllables. That is, a smaller-sized equivalence class was formed when manner class information in stressed syllables was specified than when manner class information in unstressed syllables was specified. Huttenlocher and Zue's results, together with those of Cutler and Carter (1987), demonstrate the effectiveness of lexical stress for specifying the identity of words in a lexical database.

Pisoni et al. (1985) analyzed a 126,000 word database and found that information about word length could be used to reduce the size of the search space to 6,342 words. That is, without any knowledge of segmental identity, knowledge of word length alone substantially reduced search space size from 126,000 to 6,342 words. When the words were classified according to consonant-vowel identity, the average size of the search space was reduced to 34 words. Finally, Pisoni et al. found that search space size was reduced more when the initial half of a word was fully specified (i.e., each segment identified as a phoneme) than when the final half of a word was fully specified. This result suggests two conclusions. First, words in the lexicon are organized according to phonotactic constraints that can be used in word recognition. Second, the results obtained by Zue and his colleagues with a 20,000 word lexicon could be generalized to a larger lexicon (Shipman & Zue, 1982; Huttenlocher & Zue, 1983).

Altmann (1990) has identified some important limitations of research that utilizes lexical statistics to make inferences about human or machine speech recognition processes. First, he argues that these statistics only apply to isolated word recognition. Harrington and Johnstone (1988) found that the size of the search space increased astronomically when words were strung together. For example, 'no notions' and 'known oceans' both have roughly the same phonetic transcription. If coarse phonetic coding is used, such as the manner class code employed by Zue and his colleagues, 'no notions,' 'known oceans,' and 'many chins' are all in the same equivalence class. Such examples illustrate the problems faced by any system, human or machine, when dealing with fluent connected speech instead of isolated words which are already segmented and have well-defined onsets and offsets. Second, Altmann (1990) pointed out that even if the search space is substantially reduced, much additional work remains to be done in order to isolate and specify the identity of individual phonemes.

Altmann (1990) has correctly noted some of the formidable problems that need to be addressed before the automatic recognition of words in fluent, connected speech is fully realized. Nevertheless, the work discussed in this section indicates that the phonotactic and prosodic constraints that are part of the organization of the lexicon can be exploited to aid recognition processes. Perhaps more importantly, given the perspective of the present investigation, is the demonstration that gross properties of the sound patterns of words can be used to constrain the recognition of individual words in the lexicon, especially when the size of the lexicon is small, a result that is also consistent with several views of lexical development in children (e.g., Jusczyk, in press; Moskowitz, 1973).

Nonphonemic Representations in Children

Charles-Luce and Luce (1990) evaluated neighborhood density in 5- and 7-year-old children using a phonemic coding system. Similarly, in Chapter IV, I described how the phonemic coding system was used in the present investigation to determine neighborhood density for children aged 1;06 to 5;01. However, there is a possibility that if neighborhood density is calculated under the assumption that phonemes are the basic structural units comprising words, the picture of lexical organization that results may not accurately reflect the actual organization of words in the lexicons of young children. If children do not have the same detailed phonemic system as adults, the neighborhood structure of words in their lexicons may be quite different from that proposed for the adult lexicon. In short, a more accurate picture of neighborhood density and lexical organization might be obtained by considering alternatives to a segmental phonemic coding system. The whole/part hypothesis suggests that lexical representation in young children may correspond to a relatively small number of salient features in the utterances they hear. If children's lexical representations are less detailed than adult representations, perhaps recoding phonemes into broad categories would be more useful in describing children's words. In short, the use of a coarse coding system, such as the manner class system employed by Zue and his colleagues, might provide a more accurate picture of lexical organization in children than a segmental phonemic coding system.

The present set of analyses examined lexical organization in children in which the segmental information in words was classified according to manner class categories, place of articulation categories, and a nonlinguistic random classification scheme. The manner classes used by Zue and his colleagues were based on acoustic criteria that presumably would be available to even very young children. For example, the difference between strong and weak fricatives is determined by the amplitude of the noise that characterizes these two classes of segments. The place code was included to provide a different classification system based on primarily on articulatory criteria for purposes of comparison with the manner class code. Overall, the use of three different kinds of classification schemes permitted an assessment of lexical organization that simulated a range of possible representations available to the child. Two versions of the manner class and place of articulation coding systems were used. One was the standard version in which each phoneme was classified as a member of a superordinate manner or place category. The other version classified each phoneme as a member of a superordinate manner or place category, except for vowels and syllabic consonants. In this variant of the coding systems, vowels and syllabic consonants were fully specified (e.g., /a/=/a/, /i/=/i/, etc.). Altogether, two versions of the manner class coding system and two versions of the place of articulation coding system were examined. In addition, two versions of the random coding scheme were used that were comparable to the manner class and place of articulation schemes. One version randomly assigned phonemes to a set of six cover symbols that did not correspond to any known linguistic system. A second version used five cover symbols for 24 phonemes and fully specified the remaining 17 phonemes. The distributions of phonemes according to these nonphonemic classification schemes are shown in detail in Tables 5.1 – 5.6.

Insert Tables 5.1 - 5.6 here

In addition to the different types of representational systems assessed in the present analyses, two versions of the neighborhood similarity metric also were examined. One was

Table 5.1

Manner class coding scheme I

p – c
t – c
k – c
b – c
g – c
C – c
J – c
s – s
S – s
z – s
Z – s
f – w
T – w
v – w
D – w
h – w
n – n
G – n
m – n
l – l
r – l
w – l
y – l

L – v
M – v
N – v
i – v
I – v
E – v
e – v
@ – v
a – v
W – v
Y – v
^ – v
O – v
o – v
U – v
u – v
R – v

Key —
c = consonants
s = strong fricatives
w = weak fricatives plus ‘h’
n = nasals
l = liquids, glides
v = vowels, syllabic consonants

Table 5.2

Manner class coding scheme II (fully -specified vowels & syllabics)

p – c
t – c
k – c
b – c
g – c
C – c
J – c
s – s
S – s
z – s
Z – s
f – w
T – w
v – w
D – w
h – w
n – n
G – n
m – n
l – l
r – l
w – l
y – l

L – L
M – M
N – N
i – i
I – I
E – E
e – e
@ – @
a – a
W – W
Y – Y
^ – ^
O – O
o – o
U – U
u – u
R – R

Key —
c = consonants
s = strong fricatives
w = weak fricatives plus ‘h’
n = nasals
l = liquids, glides

Table 5.3

Place of articulation coding scheme I

p – l
t – a
k – ꝟ
b – l
g – ꝟ
C – p
J – p
s – a
S – p
z – a
Z – p
f – l
T – i
v – l
D – i
h – g
n – a
G – v
m – l
l – a
r – a
w – l
y – p

L – v
M – v
N – v
i – v
I – v
E – v
e – v
@ – v
a – v
W – v
Y – v
^ – v
O – v
o – v
U – v
u – v
R – v

Key —
l = labial
i = interdental
a = alveolar
p = alveopalatal
ꝟ = velar
g = glottal
v = vowels, syllabics

Table 5.4

Place of articulation coding scheme II (fully -specified vowels & syllabics)

p – l
t – a
k – v
b – l
g – v
C – p
J – p
s – a
S – p
z – a
Z – p
f – l
T – i
v – l
D – i
h – g
n – a
G – v
m – l
l – a
r – a
w – l
y – p

L – L
M – M
N – N
i – i
I – I
E – E
e – e
@ – @
a – a
W – W
Y – Y
^ – ^
O – O
o – o
U – U
u – U
R – R

Key —
l = labial
i = interdental
a = alveolar
p = alveopalatal
v = velar
g = glottal

Table 5.5

Random coding scheme I

p – 5
t – 2
k – 1
b – 3
g – 5
C – 6
J – 5
s – 5
S – 4
z – 5
Z – 2
f – 5
T – 5
v – 1
D – 1
h – 5
n – 1
G – 5
m – 1
l – 1
r – 5
w – 3
y – 4

L – 1
M – 1
N – 3
i – 5
I – 6
E – 5
e – 6
@ – 5
a – 4
W – 5
Y – 3
^ – 1
O – 3
o – 2
U – 5
u – 2
R – 5

Distribution* —
1 = 8 items
2 = 4 items
3 = 5 items
4 = 3 items
5 = 17 items
6 = 4 items

*This distribution was designed to approximate the distribution in terms of number) of phonemes in the manner class scheme.

Table 5.6

Random coding scheme II

p – p
t – 2
k – 1
b – 3
g – g
C – 6
J – J
s – s
S – 4
z – z
Z – 2
f – f
T – T
v – 1
D – 1
h – h
n – 1
G – G
m – 1
l – 1
r – r
w – 3
y – 4

L – 1
M – 1
N – 3
i – i
I – 6
E – E
e – 6
@ – @
a – 4
W – W
Y – 3
^ – 1
O – 3
o – 2
U – U
u – 2
R – R

Distribution* —
1 = 8 items
2 = 4 items
3 = 5 items
4 = 3 items
6 = 4 items plus 17 fully-specified items from category 5

*This distribution was designed to approximate the distribution (in terms of number) of phonemes in the manner class scheme.

Table 5.7

Summary of neighborhood density analyses

Representational system	Neighborhood similarity metric: One segment substitutions additions, & deletions substitutions only	Neighborhood similarity metric: One segment additions, deletions, & within-category
Phonemic	•	
41 phonemes		
Manner class		
Six manner class categories + one vowel category	•	•
Six manner class categories + 17 fully specified vowels & syllabics	•	•
Place of articulation		
Six place categories + one vowel category	•	•
Six place categories + 17 fully specified vowels & syllabics	•	•
Random		
Six random categories	•	•
Five random categories + 17 fully specified items	•	•

the standard similarity metric in which a word was considered a neighbor of a target word if by substituting, adding, or deleting a single segment (where 'segment' could be a phoneme, or a manner or place category) the word could be transformed into the target word. The other was a modified version of the standard similarity metric in which a word was considered a neighbor of a target word if by adding or deleting a single segment, or by substituting a segment from the same manner or place class, the word could be transformed into the target word. The only type of substitutions that were allowed in this variant of the standard neighborhood similarity metric were within-category substitutions (e.g., a stop could substitute for another stop, but a stop could not substitute for a nasal). This version of the similarity metric was used was because perceptual confusions generally tend to be within-category confusions in both adults (Miller & Nicely, 1955) and children (Graham & House, 1971). Table 5.7 provides a summary of the twelve different types of neighborhood density analyses carried out to assess nonphonemic representation and their relation to each other.

Insert Table 5.7 here

Analysis of Neighborhood Structure Assuming Nonphonemic Representation

This section describes analyses in which children are assumed to have alternatives to the phonemic representation of lexical information. Paralleling the organization of the previous chapter, it contains three major sections corresponding to the three age ranges of children from the CHILDES sample: Age I — 1;06 to 2;11, Age II — 2;10 to 3;05, and Age III — 3;04 to 5;01. For each age period, a separate ANOVA was carried out. The general form of these analyses is described below.

The analysis at each age period consisted of a four-way repeated-measures ANOVA in which mean neighborhood density was the dependent variable. Two of the factors in each analysis were also examined in the analyses carried out in Chapter 4, sampling period (roughly corresponding to age) and word length (i.e., the number of phonemes in a word). The remaining two variables included in these analyses were code and specificity . Code refers to the type of coding system used to replace a phoneme with a cover symbol. Three coding systems were analyzed: (1) manner class (each phoneme was replaced with a cover symbol corresponding to its manner of articulation), (2) place of articulation (each phoneme was replaced with a cover symbol corresponding to its place of articulation), and (3) random (replace a set of phonemes with a cover symbol corresponding to a randomly-specified system, e.g., /m, ð, k, v/ = 1). The final variable included in each analysis was specificity, which referred to the level of phonemic detail preserved in each coding system. Four levels of this variable were examined: (1) the basic coding system as described above, (2) the basic coding system plus fully-specified vowels and syllabics (or a comparable number of randomly chosen, fully-specified phonemes in the case of the random coding system), (3) the basic coding system combined with the standard similarity metric for determining neighbors (as used in Charles-Luce & Luce, 1990, and in Chapter IV of the present investigation), and (4) the basic coding system plus fully-specified vowels and syllabics combined with a similarity metric that permitted only within-category substitutions for segments to determine whether or not a word qualified as a neighbor. Finally, as in Chapter IV, the number of samples varied across the three age periods examined. In addition, the words included in the analyses ranged between two and seven phonemes in length.

The following questions were addressed in these analyses. First, what kinds of changes, if any, occur in mean neighborhood density between the ages of 1;06 and 5;01 when phonemes are coded according to a manner class, place of articulation, or random code? Second, are there differences in the neighborhood structure when different coding systems are used? Third, what is the effect on neighborhood structure when different degrees of specificity are used? Fourth, what kinds of interactions, if any, occur among these variables?

Insert Figures 5.1 & 5.2 here

Age Range I — 1;06-2;10

The sampling period covered in Age I was 1;06 to 2;10, and included five samples. Data from four children were included in this analysis: Nina, Adam, Eve, and Peter. The left panel of Figure 5.1 shows mean neighborhood density plotted as a function of sample period at Age I. The most noticeable feature of the results shown in Figure 5.1 is that the mean number of neighbors issubstantially greater when a nonphonemic representation system is used than when a phonemic representation system is used; mean neighborhood densities are seven to eight times the densities observed under a phonemic representation system (cf. Figure 4.1). Nevertheless, mean neighborhood density generally increased as a function of successive samples during Age I. Thus, as more words are added to the child's lexicon, neighborhood density increases, an effect also observed when neighborhood densities were computed using the phonemic representational system. Although the range of the increase was not large (14.15 to 17.64 neighbors), the effect was statistically reliable, $F(4, 12) = 5.03$, $MSE = 91.378$, $p < .05$.

The left panel of Figure 5.2 shows the mean neighborhood density plotted as a function of word length in phonemes. As observed in earlier analyses, the average number of neighbors decreased substantially as words increased in length. For two– and three–phoneme words, the mean number of neighbors was roughly 35, whereas for words five–phoneme and longer, the mean number of neighbors was less than two. What is remarkable is how mean neighborhood density for long words remains close to zero despite the substantial increase in overall neighborhood density when the nonphonemic representation systems are used. From another perspective, however, this effect is not surprising. As noted in Chapter IV, the effect of increasing word length imposes powerful constraints on whether a word qualifies as a neighbor for another word; even utilizing coarse coding fails to appreciably increase neighborhood density for long words. The effect of word length was statistically reliable, $F(5, 15) = 1560.91$, $MSE = 38.140$, $p < .0001$.

The left panel of Figure 5.3 shows the effect of the three nonphonemic representational systems (hereafter referred to as "code") on mean neighborhood density for Age I. The two linguistically-based codes both produced a greater number of neighbors than the random code. Of the two linguistically-based codes, place of articulation produced more neighbors than manner class (M(place of articulation) = 21.7, M(manner class) = 17.2, M(random) = 9.6). The effect of code was statistically reliable, $F(2, 6) = 202.61$, $MSE = 88.273$, $p < .0001$. The contrast between the two linguistically-based codes and the random code attests to the effect of phonological structure in computing similarity relationships among words in the lexicon, an effect also observed when the same codes are applied to the adult lexicon (cf. Shipman & Zue, 1982). In addition, the present results also show that a representational system based on manner class is more effective in uniquely specifying a lexical item than a code based on place of articulation, a finding that coincides with earlier perceptual findings from adults (e.g., Miller & Nicely, 1955) and computational analyses carried out in speech recognition (e.g., Huttenlocher & Zue, 1983).

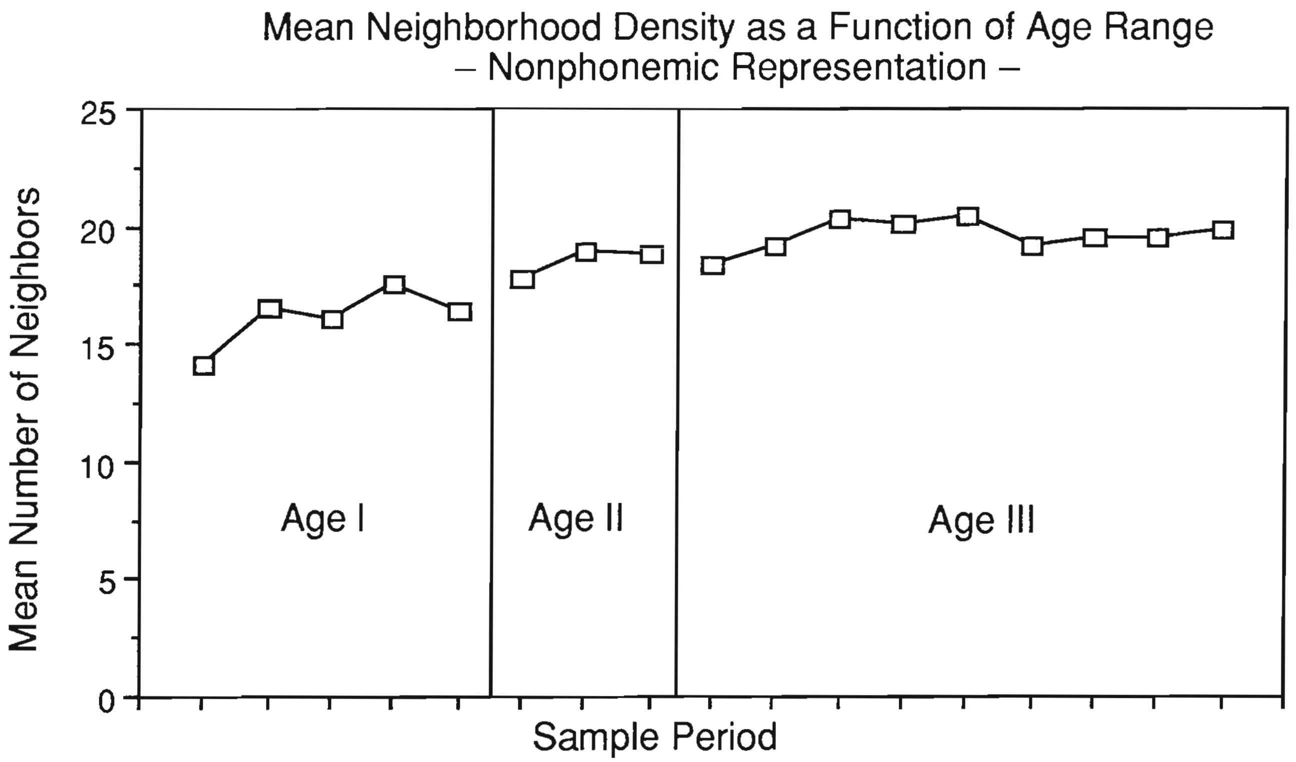

Figure 5.1 Mean neighborhood density as a function of age range. Neighborhood density was computed assuming nonphonemic representation (manner class, place of articulation, and random coding). Age I data are shown in left panel, Age II data in the center panel, and Age III data in the right panel.

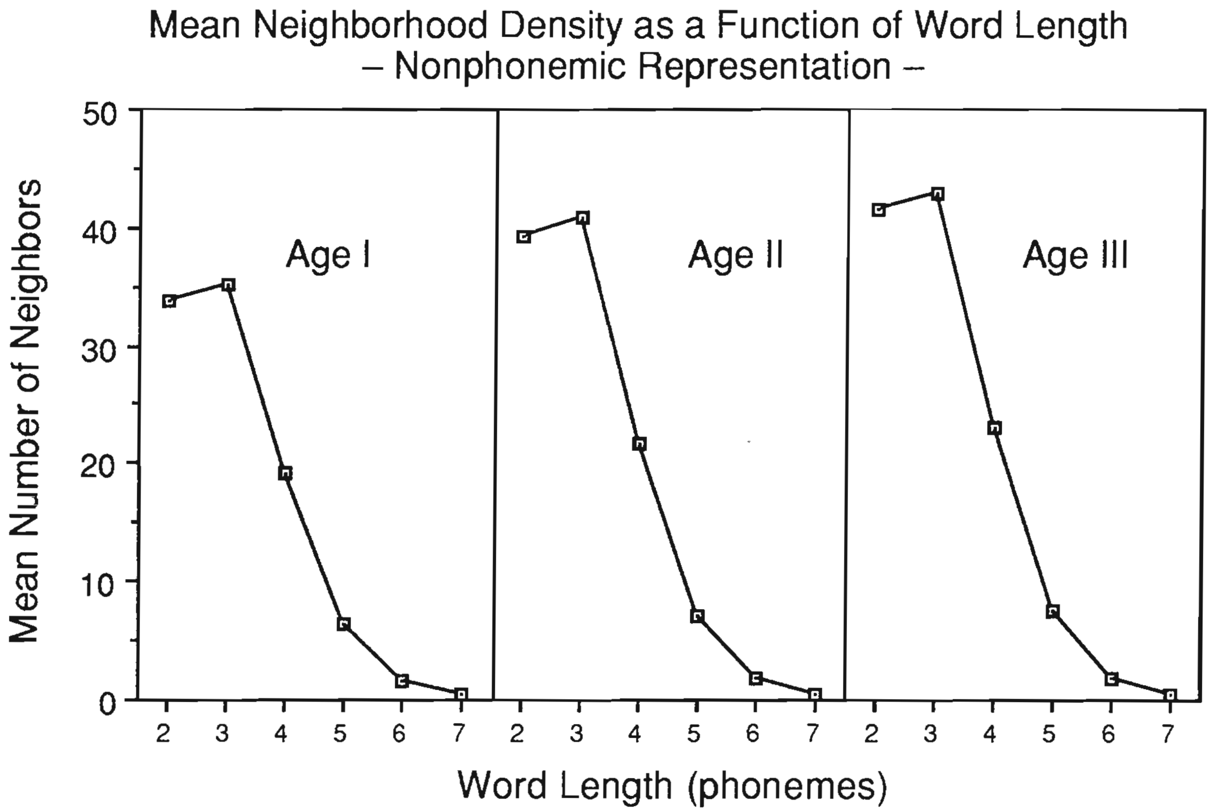

Figure 5.2 Mean neighborhood density as a function of word length measured in phonemes. Neighborhood density was computed assuming nonphonmic representation (manner class, place of articulation, and random coding). Age I data are shown in the left panel, Age II data in the center panel, and Age III data in the right panel.

Insert Figures 5.3 & 5.4 here

The left panel of Figure 5.4 shows the effects of different levels of specificity for the coding schemes (hereafter referred to as "specificity") on mean neighborhood density, where 'Std.' refers to the standard code, 'Std.+WC subs' refers to the standard code using within-category substitutions only, 'Std.+V' refers to the standard code plus fully-specified vowels, and 'Std.+V/WC' refers to the standard code plus vowels and within-category substitutions only. For the random code, 17 randomly selected segments corresponded to the 17 fully specified vowels used in the manner and place codes. As expected, increasing the number of classes within a coding scheme by fully specifying vowels decreased mean neighborhood density. Adding vowels to the representations increases the phonotactic constraints that must be satisfied before a word qualifies as a neighbor for a target word. Similarly, restricting segmental substitutions to within-category substitutions also decreased mean neighborhood density. However, the effect was larger when vowels were added to the lexical representations than when substitutions were restricted. The lowest neighborhood density was produced when the two types of variations on the standard coding schemes were combined. The effect of specificity was statistically reliable, $F(3, 9) = 1360.43$, $MSE = 61.472$, $p < .0001$.

With the exception of the four-way interaction, all of the possible interactions among sample period, length, code, and specificity were statistically reliable in the present analysis ($p < .05$). Instead of considering in detail all of the significant interactions, only a subset of the interactions are described below. Appendix I contains the ANOVA summary table for these interactions, plus all other interactions and main effects in the analysis for Age I.

Insert Figure 5.5 here

Figure 5.5 shows mean log neighborhood density plotted as function of sample period, code, and specificity for Age I. A log scale was used because of the wide range of neighborhood densities spanned by these variables. Each panel shows log neighborhood density plotted as a function of the three types of nonphonemic representational systems. Different levels of specificity are shown in successive panels (from top to bottom: standard codes, standard codes in which within-category substitutions only are permitted, standard codes with fully specified vowels, and standard codes in which vowels are fully specified and in which only within-category substitutions are permitted). For each successive sample period, the mean number of neighbors generally increased, albeit at a relatively modest rate. The manner and place codes generally produced similar neighborhood densities, especially for the standard version, as well as for the standard code in which only within-category substitutions were permitted. When vowels were added to the representational systems, the two linguistically-based codes became similar to the random code (compare the top two panels with the bottom two panels). One way to restate this finding is that the addition of vowels (or 17 additional categories corresponding to vowels in the case of the random code) to the representational systems had a larger effect on reducing mean neighborhood density for the two linguistically-based codes than for the random code. Again, this demonstrates how the manner and place codes are sensitive to the specific distributional characteristics of English words in a way that is different from the random code. Finally, because one of the goals of the present investigation was to determine if a nonphonemic representational system could produce similar neighborhood densities as the phonemic representational system, it is important to note the linguistically-derived representational system with the lowest overall neighborhood densities. The system that met this criterion utilized the manner code and fully

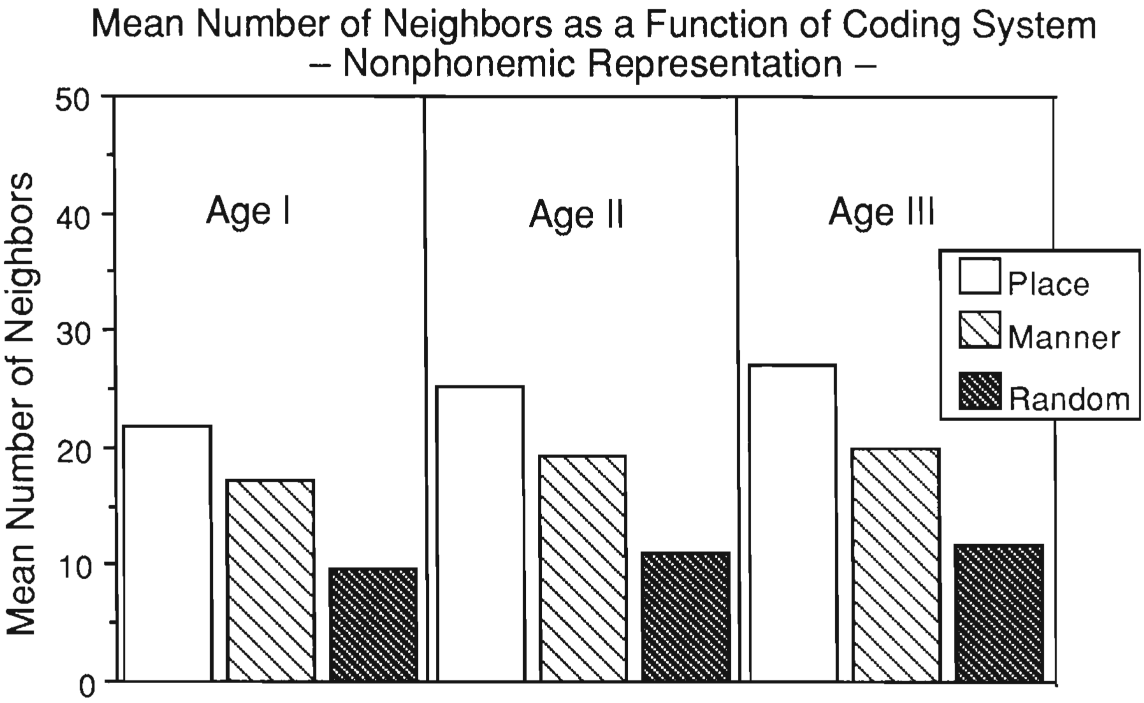

Figure 5.3 Mean number of neighbors as a function of place, manner, and random coding systems. Data from Age I are shown in the left panel, Age II data in the center panel, and Age III data in the right panel.

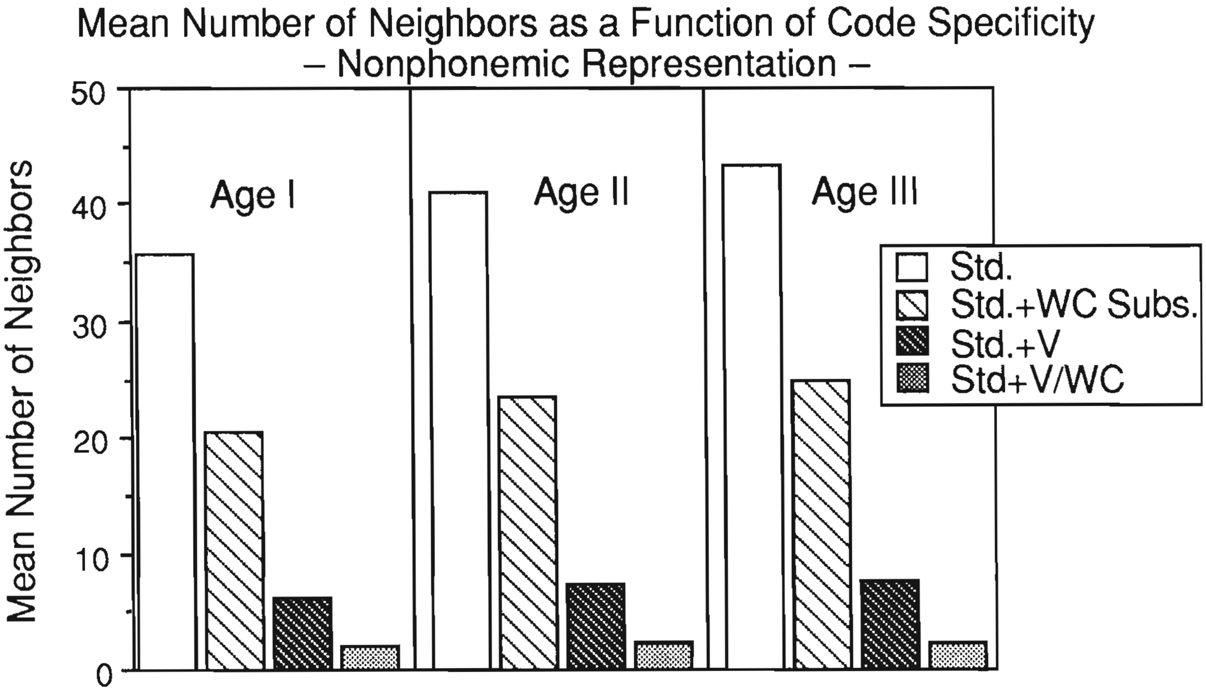

Figure 5.4 Mean neighborhood density as a function of coding specificity. "Std." refers to the basic coding system. "Std.+WC Subs." refers to a standard coding system in which only within-category substitutions are permitted. "Std.+V" refers to a standard coding system in which vowels are fully specified. "Std.+V/WC" refers to a standard coding system in which vowels are fully specified and only within-category substitutions are permitted. Age I data are shown in the left panel, Age II data in the center panel, and Age III data in the right panel.

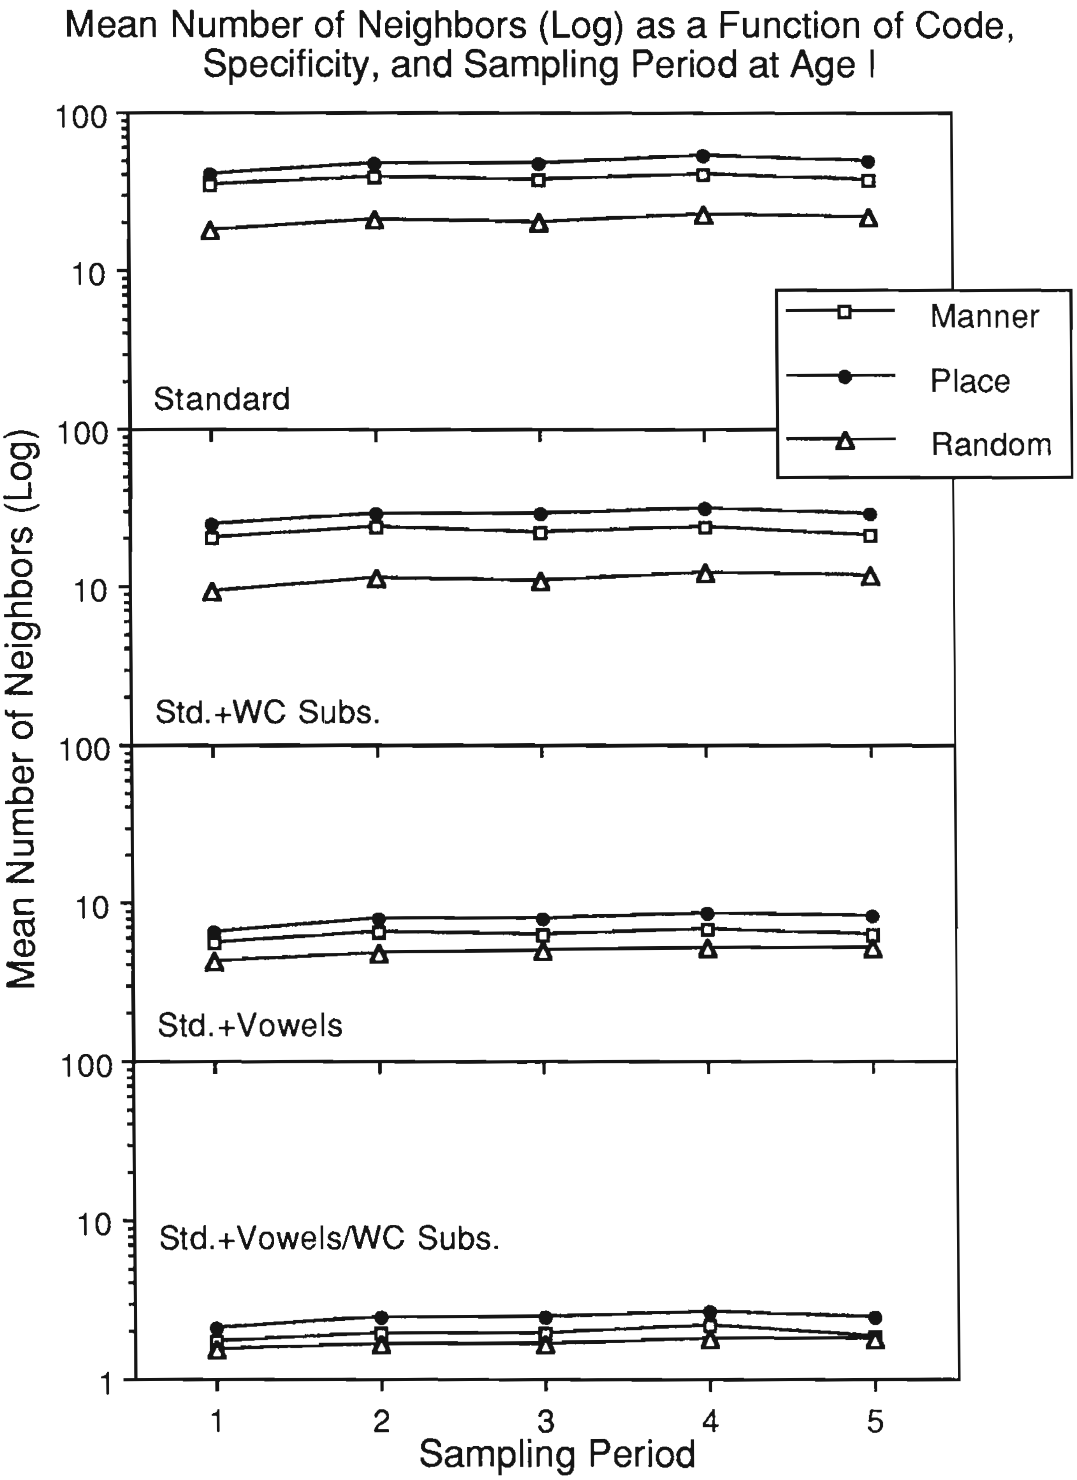

Figure 5.5 Mean number of neighbors as a function of coding system, specificity of coding system, and sampling period at Age I. Each panel shows mean neighborhood density as a function of manner, place, and random codes and sampling period. "Standard", "Std.+WC Subs.", "Std.+Vowels", and "Std.+Vowels/WC Subs." refer to levels of the variable specificity.

specified vowels in conjunction with a metric that only permitted within-category substitutions.

Figure 5.6 illustrates the interaction among code, length, and sample period. Each panel shows mean log neighborhood density plotted as a function of length and sample period for a different code (from top to bottom: manner code, place code, and random code) at Age I. Aside from the main effects for code (i.e., the mean number of neighbors decreases from place to manner to random codes) and length (i.e., short words have more neighbors than long words) that were described earlier in this section, Figure 5.6 shows several additional effects. First, neighborhood density for two- and three-phoneme words coded using the manner and place systems increased as a function of sample period, whereas neighborhood density changed little for two- and three-phoneme words coded using the random code. Second, for both manner and place codes, neighborhood density was greater in three-phoneme words than in two-phoneme words. However, this effect was reversed in the case of the random code; three-phoneme words actually had fewer neighbors than two-phoneme words. Each of these effects provide additional evidence that the manner and place codes capture aspects of the phonological regularities found in English words in a way that makes them differ from the random code.

Insert Figure 5.6 here

In summary, neighborhood density calculated using several nonphonemic coding schemes showed a reliable increase between 1;06 and 2;11. The place code produced the largest number of neighbors, the manner code the next largest number of neighbors, and the random code the fewest number of neighbors. As the degree of specificity increased (i.e., moving from the standard code to one in which vowels were fully specified, to one that permitted within-category substitutions only, to one that included both fully specified vowels and within-category substitutions only), the number of neighbors decreased substantially. The linguistically-derived representational scheme that produced the lowest neighborhood densities was the manner code paired with fully specified vowels and in which only within-category substitutions were permitted. Finally, several interactions demonstrated the unique characteristics of the non-linguistic random code compared to the two linguistic codes.

Age Range II — 2;06-3;05

The sampling period covered in Age II was 2;06 to 3;05, and included three samples. Data from three children were included in this analysis: Nina, Adam, and Sarah. The center panel of Figure 5.1 shows mean neighborhood density plotted as a function of sample period at Age II. There was no significant effect of neighborhood density as a function of successive sample period, $F(2, 4) < 1.0$. This replicates the lack of effect for sample period found for phonemic representation at Age II. Again, the limited range for the variable sample period and the limited number of children included in the sample likely contributed to the lack of a reliable effect for this variable.

The center panel of Figure 5.2 shows mean neighborhood density as a function of word length at Age II. For two– and three–phoneme words, the mean number of neighbors was around 40, while for five–phoneme and longer words, the mean number of neighbors was less than seven. As in earlier analyses, a reliable effect of word length on neighborhood density was observed, $F(5, 10) = 206.75$, $MSE = 174.906$, $p < .0001$.

The center panel of Figure 5.3 shows the effect of code on mean neighborhood density at Age II. Each of the two linguistically-based codes both produced a greater number of neighbors than the nonlinguistically-based random code, replicating the results observed at Age I. Of the two linguistically-based codes, the place of articulation code produced more neighbors than the manner class code, *M* (place of articulation) = 25.3, *M*(manner class) = 19.2, *M*(random) = 11.1. The effect of code was statistically reliable, *F*(2, 4) = 290.12, *MSE* = 38.108, *p* < .0001. These results show that a representational system based on manner class is more effective in uniquely specifying a lexical item than a code based on place of articulation, replicating the effects observed at Age I.

The center panel of Figure 5.4 shows mean neighborhood density as a function of specificity at Age II. As expected from the analyses at Age I, increasing the number of categories and restricting segmental substitutions decreased mean neighborhood density, *F*(3, 6) = 400.05, *MSE* = 124.403, *p* < .0001. Overall, the effect was essentially the same as that observed for Age I, except that mean neighborhood densities were increased somewhat.

The description of interactions obtained in the analysis of the Age II data is restricted to a subset of the total number of interactions. Because the three-way interactions among code, specificity, and sample period, and among code, length, and sample period encompass other interactions, the remaining interactions are not considered. Appendix II contains the complete ANOVA summary table for all interactions and main effects in the analysis of Age II data.

Figure 5.7 illustrates the interaction among code, specificity, and sample period for Age II. Mean log neighborhood density is plotted as a function of code and sample period in each panel; the effect of specificity is shown in successive panels. A consistent effect of code was observed across the four levels of specificity that corresponded to the pattern observed at Age I. In all cases, the two linguistically-based codes produced more neighbors than the nonlinguistic code, while the effect of increased specificity was to reduce the difference between the codes, *F*(6, 12) = 239.10, *MSE* = 12.306, *p* < .0001. However, the three-way interaction among these variables and sample period was not reliable, *F*(12, 24) < 1.0. As we observed for Age I, the linguistically-based code that produced the fewest neighbors was a manner code augmented by fully specified vowels using a metric that only permitted within-category substitutions.

Insert Figure 5.7 here

Figure 5.8 shows the interaction among code, length, and sample period at Age II. Mean log neighborhood density is plotted as a function of length and sample period in each panel; successive panels show different coding systems. The relationship among code and length is consistent across sample period. At each sample period, the mean number of neighbors generally increased as a function of decreasing word length. The exception is for two- and three-phoneme words. In the case of the two linguistically-based codes, three-phoneme words have a greater number of neighbors than two-phoneme words, whereas in the case of the random code, two-phoneme words have a greater number of neighbors than three-phoneme words, *F*(10, 20) = 121.95, *MSE* = 13.163, *p* < .0001. However, no reliable change in this relationship as a function of sample period was observed, *F*(20, 40) < 1.0.

Insert Figure 5.8 here

In summary, many of the same effects observed at Age I were also present in the analyses of data from Age II, with the notable exception of sample period. The place code

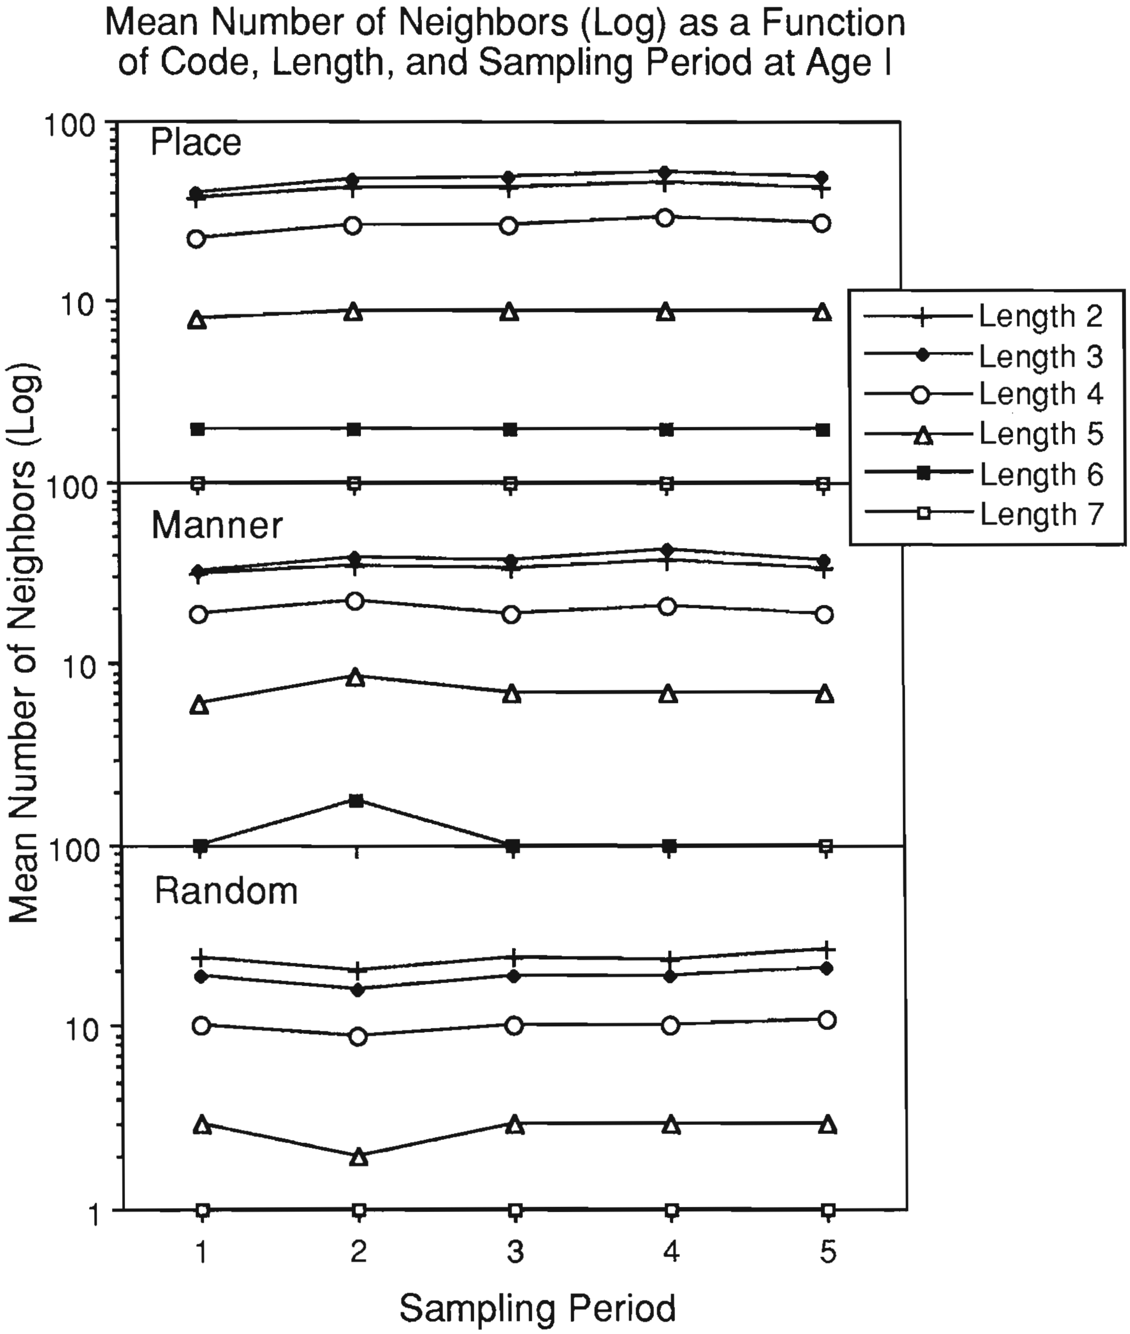

Figure 5.6 Mean number of neighbors as a function of code and sampling period at Age I. Each panel shows the mean neighborhood density for 2-, 3-, 4-, 5-, 6-, and 7-phoneme words.

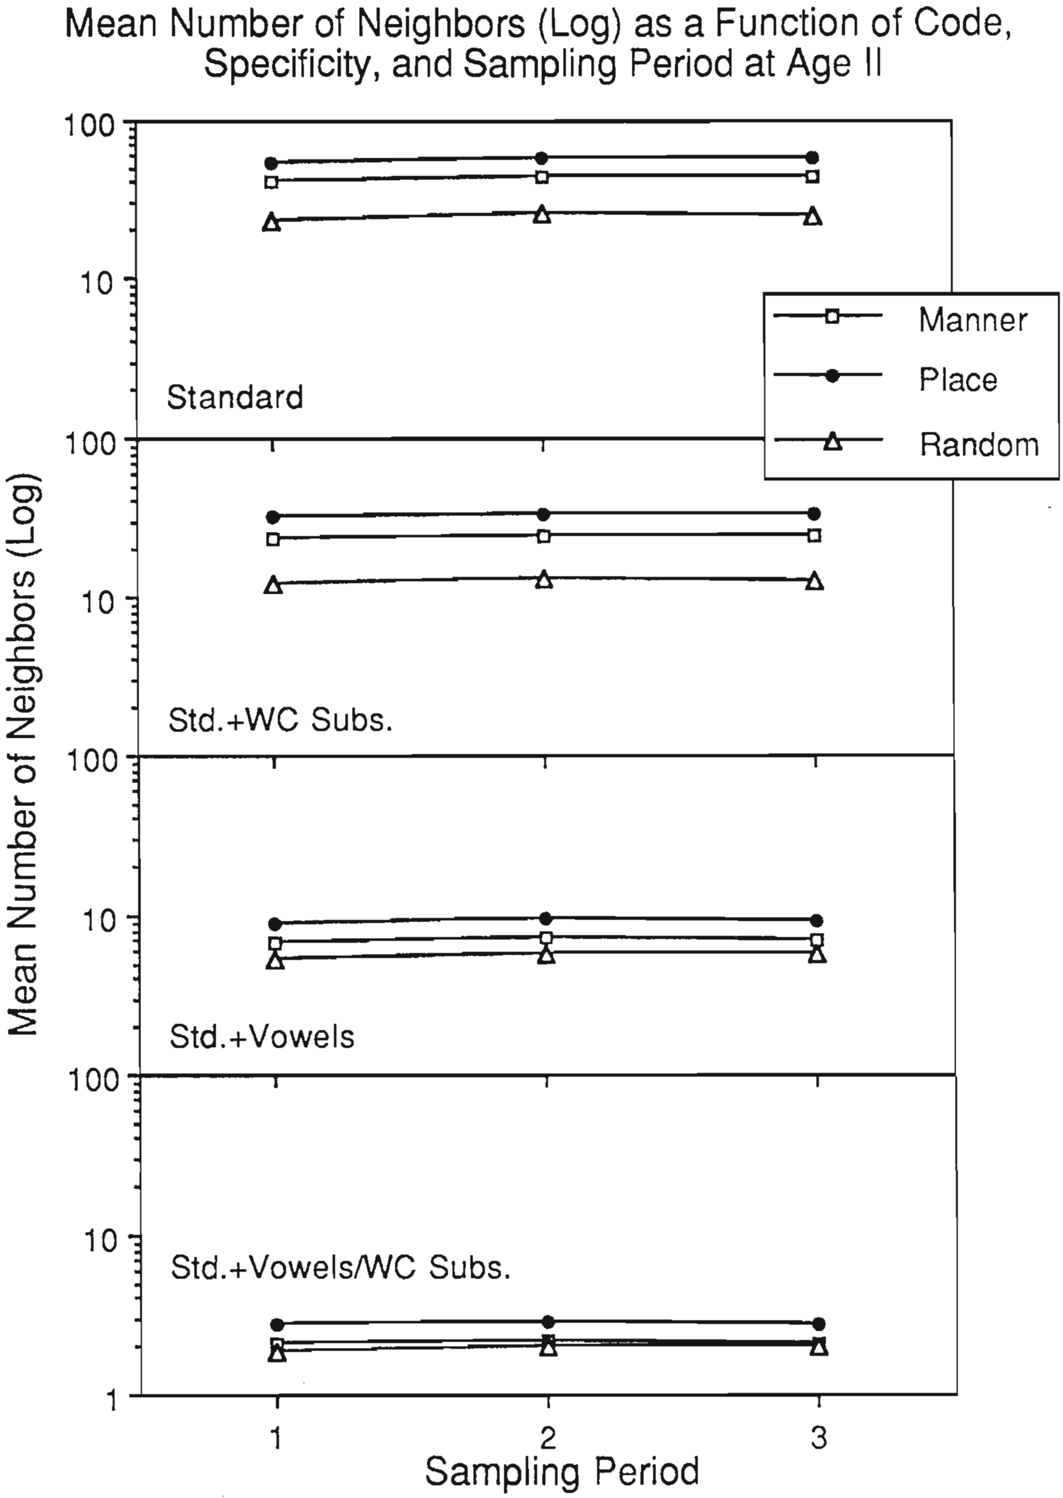

Figure 5.7 Mean number of neighbors as a function of coding system, specificity of coding system, and sampling period at Age II. Each panel shows mean neighborhood density as a function of manner, place, and random codes and sampling period. "Standard", "Std.+WC Subs.", "Std.+Vowels", and "Std.+Vowels/WC Subs." refer to levels of the variable specificity.

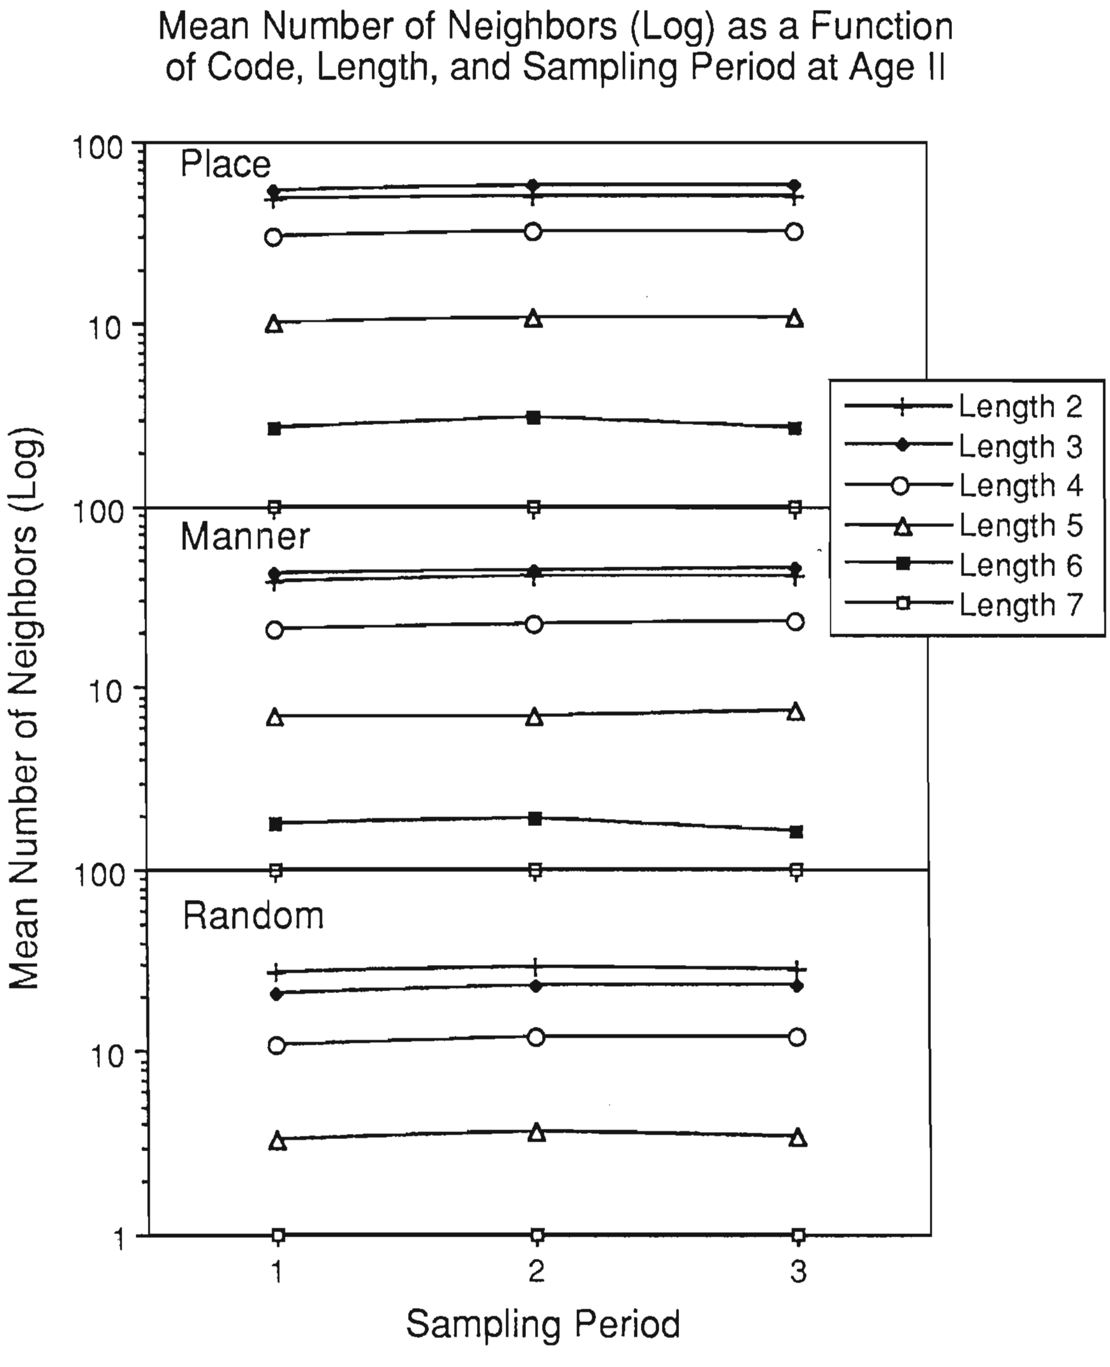

Figure 5.8 Mean number of neighbors as a function of code and sampling period at Age II. Each panel shows the mean neighborhood density for 2-, 3-, 4-, 5-, 6-, and 7-phoneme words.

produced the largest number of neighbors, the manner code the next largest number of neighbors, and the random code the fewest number of neighbors. As the degree of specificity increased, the number of neighbors decreased substantially. The linguistically-derived representational scheme that produced the lowest neighborhood densities was the manner code in which vowels were fully specified and only within-category substitutions were permitted. Furthermore, as at Age I, several of the interactions demonstrated how the nonlinguistic random code differed from the two linguistic codes.

Age Range III — 3;04-5;01

The sampling period covered in Age III was 3;04 to 5;01, and included nine samples. Data from two children were included in this analysis, Adam and Sarah. The right panel of Figure 5.1 shows mean neighborhood density plotted as a function of sample period at Age III. No significant effect of successive sample period on neighborhood density was observed, $F(8, 8) < 1.0$. This result replicates the findings obtained at Age III when phonemic representation was assumed. The same reason is likely responsible for this null result in both cases, namely, the fact that the analyses were based on data from only two children, thus limiting the statistical power of the analyses.

The right panel of Figure 5.2 shows mean neighborhood density as a function of word length at Age III. For two– and three–phoneme words, the mean number of neighbors was between 40 and 45, while for four–phoneme words, the mean number of neighbors decreased to around 20, and for five–phoneme and longer words, the mean number of neighbors was less than 10. As in earlier analyses, a reliable effect of word length on neighborhood density was present, $F(5, 5) = 330.20$, $MSE = 244.15$, $p < .0001$.

The right panel of Figure 5.3 shows the effect of code on mean neighborhood density at Age III. Each of the two linguistically-based codes produced a greater number of neighbors than the nonlinguistically-based random code, replicating the results obtained at Ages I and II. Of the two linguistically-based codes, the place of articulation code produced more neighbors than the manner class code, M (place of articulation) = 27.1, M(manner class) = 19.9, M(random) = 11.7. The effect of code was statistically reliable, $F(2, 2) = 586.39$, $MSE = 43.57$, $p < .0017$.

The right panel of Figure 5.4 shows the effects of different levels of specificity for the coding schemes on mean neighborhood density at Age III. Given the results obtained at Age I and Age II, one would expect that increasing the number of categories and restricting segmental substitutions would result in decreased mean neighborhood density. Consistent with these earlier results, a reliable effect of specificity was obtained, $F(3, 3) = 554.13$, $MSE = 201.23$, $p < .0001$.

The description of interactions obtained in the analysis of the Age III data is restricted to a subset of the total number of interactions, namely, the three-way interactions among code, specificity, and sample period, and among code, length, and sample period. Appendix III contains the complete ANOVA summary table for all interactions and main effects in the analysis of Age III data.

Figure 5.9 illustrates the interaction among code, specificity, and sample period at Age III. Mean log neighborhood density is plotted as a function of code and sample period in each panel; the effect of specificity is shown in successive panels. The effect of code across the four levels of specificity corresponds to the pattern observed at Age I and Age II. In all cases, the two linguistically-based codes produced more neighbors than the nonlinguistic code, while the effect of increased specificity was to reduce the difference between the codes, $F(6, 6) =$

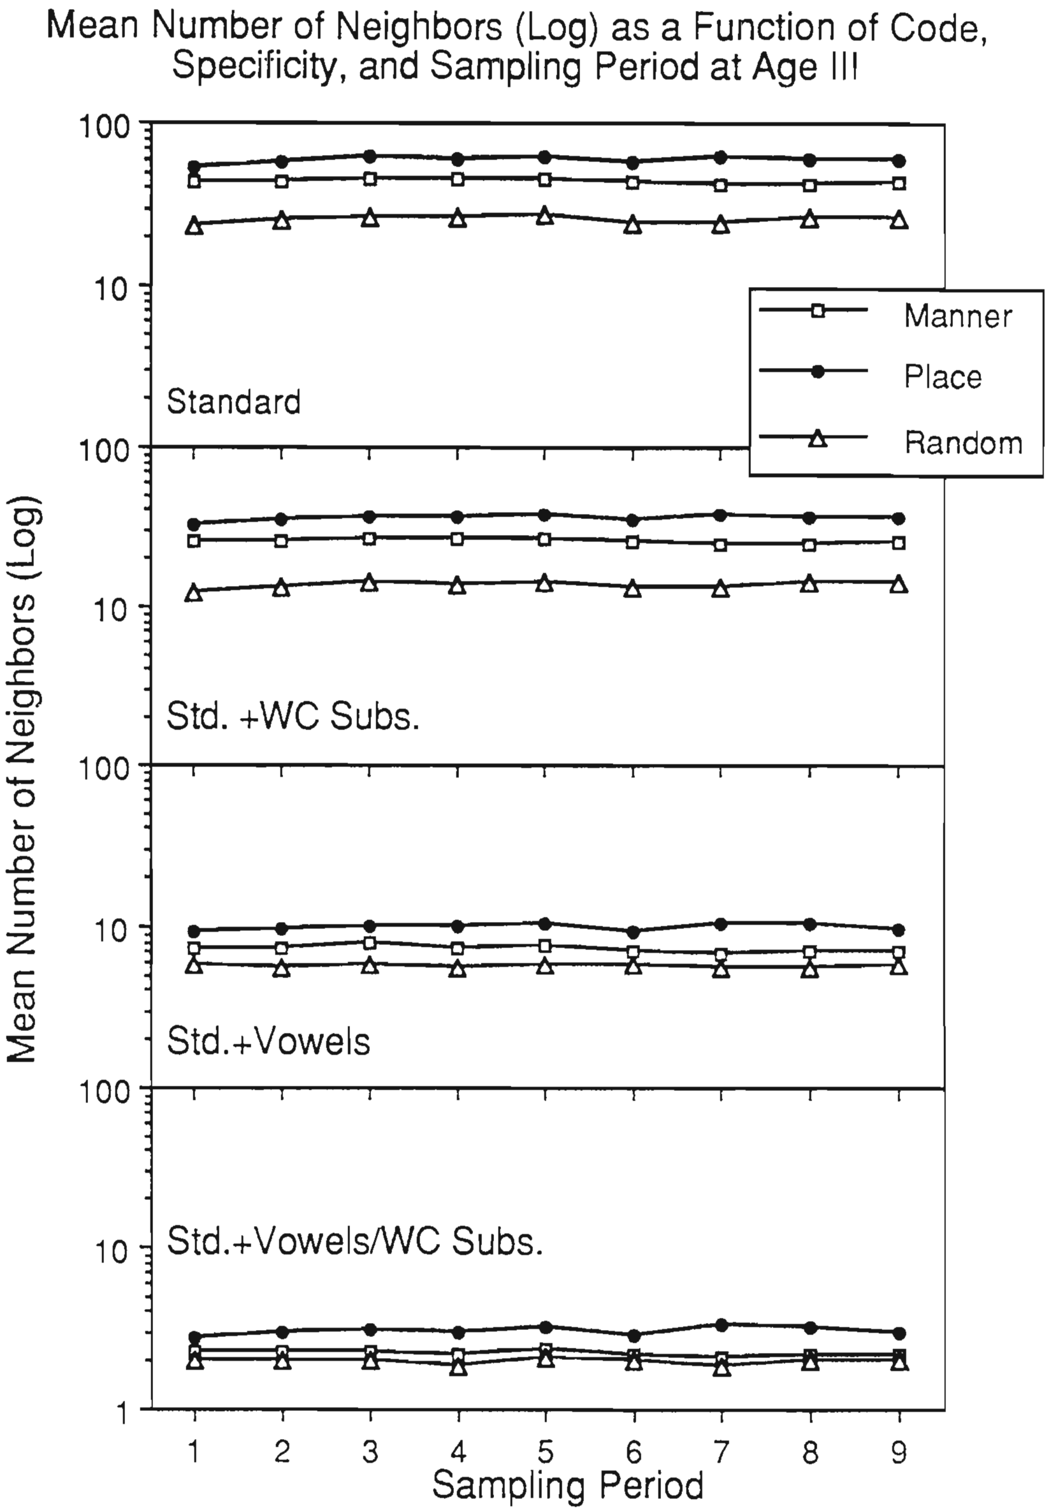

Figure 5.9 Mean number of neighbors as a function of coding system, specificity of coding system, and sampling period at Age III. Each panel shows mean neighborhood density as a function of manner, place, and random codes and sampling period. "Standard", "Std.+WC Subs.", "Std.+Vowels", and "Std.+Vowels/WC Subs." refer to levels of the variable specificity.

Mean Number of Neighbors (Log) as a Function of Code, Length, and Sampling Period at Age III

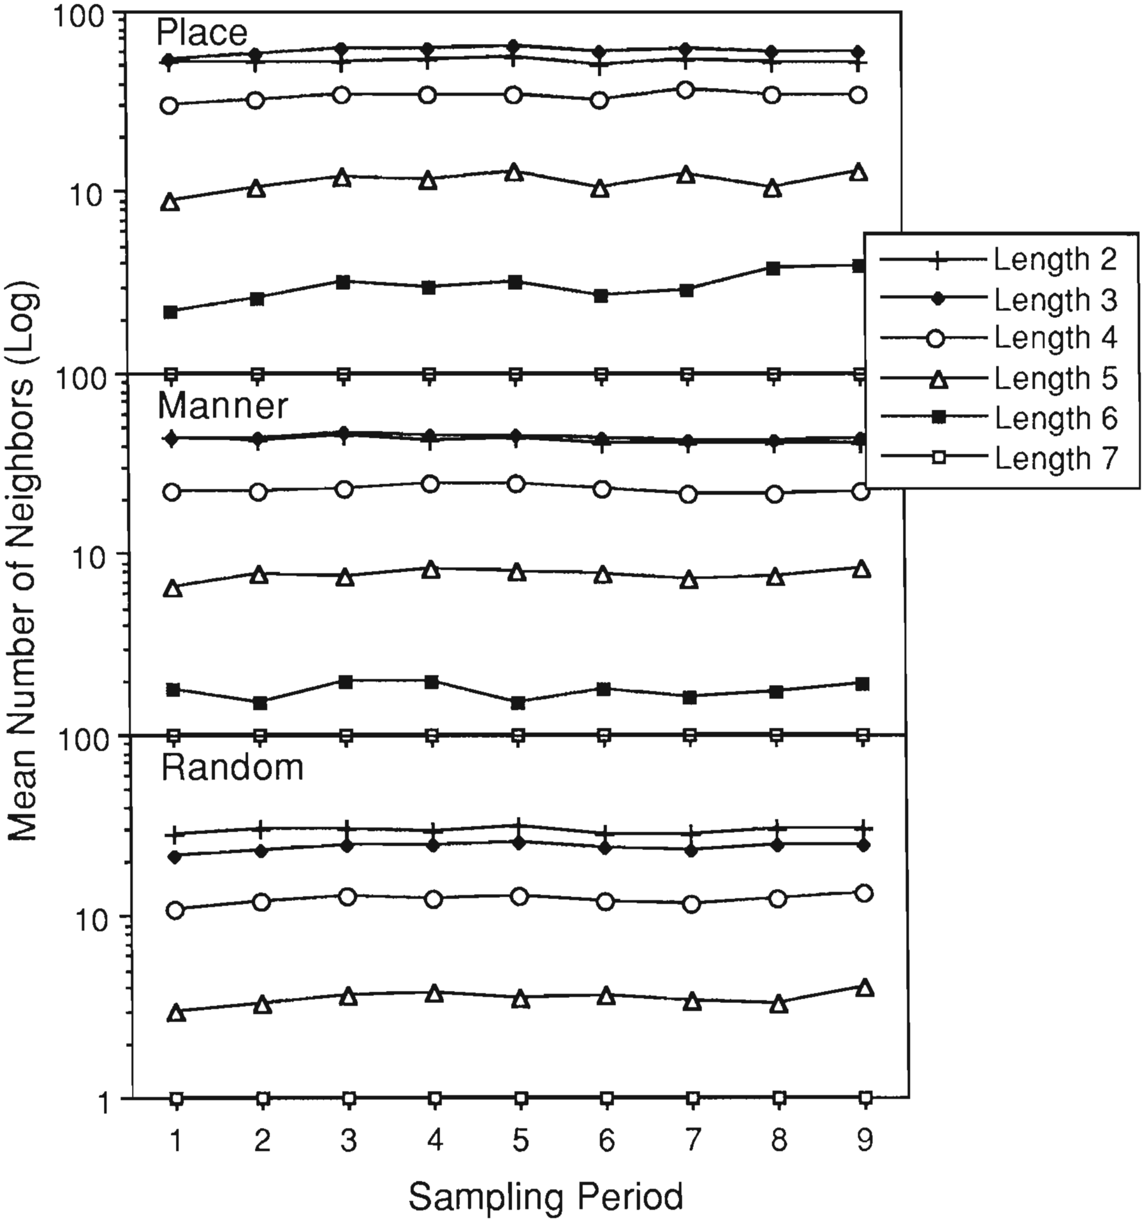

Figure 5.10 Mean number of neighbors as a function of code and sampling period at Age III. Each panel shows the mean neighborhood density for 2-, 3-, 4-, 5-, 6-, and 7-phoneme words.

717.97, *MSE* = 9.1, $p < .0001$. However, there was no reliable three-way interaction among these variables and sample period, $F(48, 48) < 1.5$.

Insert Figure 5.9 here

Figure 5.10 shows the interaction among code, length, and sample period at Age III. Mean log neighborhood density is plotted as a function of length and sample period in each panel; successive panels show different coding systems. The relationship among code and length is consistent across sample period. At each sample period, the mean number of neighbors generally increased as a function of decreasing word length. The exception is for two- and three-phoneme words. Three-phoneme words have a greater number of neighbors than two-phoneme words when coded using the manner and place systems. In contrast, two-phoneme words have a greater number of neighbors than three-phoneme words when coded using the random system. However, this effect is attenuated compared with Ages I and II. Overall, the interaction between code and sample period was reliable, $F(10, 10) = 351.78$, *MSE* = 10.217, $p < .0001$. The three-way interaction among code, length, and sample period was not statistically reliable, $F(80, 80) < 1.0$.

Insert Figure 5.10 here

In summary, many of the same effects observed at Age I and Age II were also found in the analyses of data from Age III. First, sample period failed to show any systematic variation. Second, of the two linguistically-based codes, the manner code produced the fewest number of neighbors. As before, increasing the degree of specificity substantially reduced the number of neighbors. Finally, several of the interactions demonstrated how the nonlinguistic random code differed from the two linguistic codes.

Comparison of Nonphonemic Representation and Phonemic Representation

Phonemic representation could be considered as a baseline condition useful for evaluating the various nonphonemic representational systems examined in the present investigation. In order to judge how mean neighborhood density in the nonphonemic representational systems compared with the phonemic code, several analyses were carried out to compare the manner class code and the phonemic code. The manner class code was chosen because, when used in conjunction with fully specified vowels and restricted to within-category substitutions, it resulted in the lowest overall neighborhood densities of all the linguistically-based, nonphonemic representational systems. Figure 5.11 shows, from left to right, the mean number of neighbors for the two types of representational systems as a function of sample period at Ages I, II, and III. In general, the manner code produces a marginally larger number of neighbors than the phonemic code across all three age ranges. An ANOVA was carried out for each age range comparing the mean neighborhood densities obtained using the two codes. Each analysis consisted of a three-way ANOVA with variables code (phonemic and manner class), sample period (varying from three to nine samples depending on the age range examined), and word length (two to seven segments). The results of the ANOVAs are outlined below. Only those results relevant to the variable code are described; the remainder of the effects observed were consistent with the results described earlier for phonemic and nonphonemic representations.

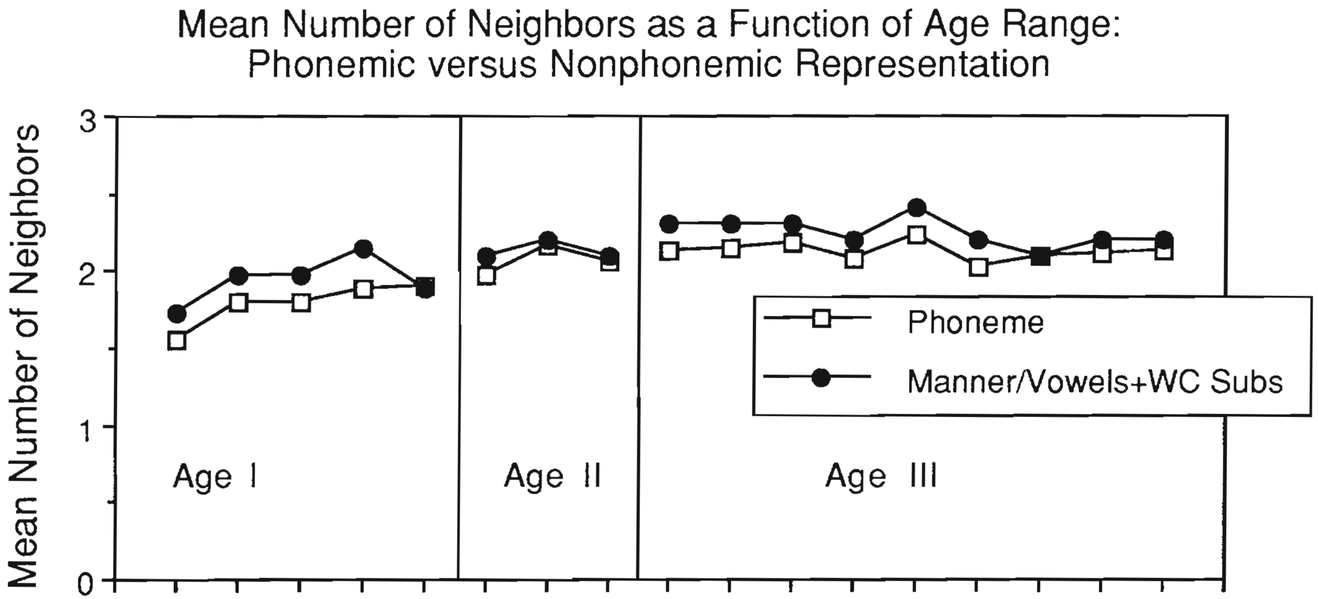

Figure 5.11 Mean number of neighbors as a function of age range and representation system. Two representation systems are compared, the adult phonemic system and a nonphonemic system that produced the lowest mean neighborhood densities. The nonphonemic system utilized a manner class representation system in which vowels were fully specified and in which only within-category substitutions were permitted when computing neighborhood density. Age I data are shown in the left panel, Age II data in the center panel, and Age III data in the right panel.

Insert Figure 5.11 here

Age I. A marginally significant main effect of code at Age I was obtained, $F(1, 3) = 6.97$, $MSE = 0.206$, $p < .08$. The effect would likely have been reliable had it not been for an inconsistency in neighborhood density related to word length. For two-phoneme words, the manner code produced slightly fewer neighbors than the phonemic code. Otherwise, the manner code always produced a larger number of neighbors than the phonemic code. The interaction between code and length was reliable, $F(5, 15) = 287.76$, $MSE = 0.071$, $p < .0001$. However, no reliable interaction was obtained between code and sample period, $F(4, 12) = 1.27$, or among code, sample period, and length, $F(20, 60) = 1.31$. Overall, the results suggest small but relatively consistent differences between the phonemic and manner class representations at Age I.

Age II. The analysis of Age II data revealed a marginally reliable effect of code, $F(1, 2) = 14.56$, $MSE = 0.012$, $p < .07$, mirroring the results obtained at Age I. In addition, a reliable interaction between code and length was obtained, $F(5, 10) = 146.69$, $MSE = 0.082$, $p < .0001$. Again, the interaction between code and length was due to the fewer number of neighbors for two-phoneme words produced by the manner code, violating the tendency for the manner code to produce the greater number of neighbors. And, as was observed at Age I, the interaction between code and sample period was not reliable, $F(2, 4) = 2.18$, nor was the interaction among code, sample period, and length reliable, $F(10, 20) < 1.0$. Overall, these results suggest that the manner code generally produces slightly more neighbors than the phonemic code at Age II.

Age III. Data from Age III revealed no reliable overall effect of code, $F(1, 1) < 1.0$. However, similar to the results for Ages I and II, a significant interaction between code and length was obtained, $F(5, 5) = 65.22$, $MSE = 0.324$, $p < .0001$. Again, the trend for the manner code to produce more neighbors than the phoneme code was not observed for two-phoneme words. As in the analysis of Age I and Age II data, no reliable interaction between code and sample period was obtained, $F(8, 8) < 1.0$, nor was there a significant interaction among code, sample period, and length, $F(40, 40) = 1.43$. In sum, with the exception of very short words, the manner code produces slightly more neighbors than the phonemic code at Age III.

Taken together, the results of these analyses indicate that the manner class code produces a small but consistently larger number of neighbors than the phonemic code. Despite these differences, several lines of evidence suggest that children do not use a phonemic system for representing lexical information. First, their performance in production tasks diverges in many respects from what would be expected if they were using an adult phonemic system (Macken, 1979; Maxwell, 1984; Smith, 1981). Second, the word recognition abilities of children differ markedly in certain ways from those of adults (e.g., Walley, 1988). Even simple tasks that would seem to minimize the effect of cognitive limitations due to memory or attention, such as discriminating two sounds, result in performance that deviates from the adult model (e.g., Graham & House, 1971). I would argue that since research has demonstrated the nonphonemic representation and organization of lexical information in children, it is reasonable to consider alternative representational systems that are similar to the phonemic system in certain critical ways. The most important characteristic that any alternative to the phonemic system must possess is that the confusability among words in such a system be comparable to that obtained using the phonemic system, such as the manner class system evaluated in the present analyses. Using this as a criterion, the manner class representational system becomes a viable contender for how young children might represent lexical information, even though there may be other equally plausible alternatives. (The issue of additional alternative representational systems is discussed further in Chapter VI.)

Summary

This chapter described the results of analyses of neighborhood densities in young children aged 1 1/2 to 5 years produced by a variety of nonphonemic representational systems. Specifically, three systems were examined. The first was based on six manner class categories, the second based on six place of articulation categories, and the third based on the random assignment of phonemes to six categories. Although the nonphonemic systems as a group generated a larger number of neighbors than the phonemic code, consistent differences among the three nonphonemic codes were observed. The random scheme produced the fewest neighbors, while of the two linguistically-based codes, the manner code produced the fewest neighbors. This pattern of results was consistent across all the age ranges examined. In addition, a variable related to the specificity of the three different codes was also included in the analyses. This variable had four levels: the basic representational system, the basic system augmented by fully specifying vowels, the basic system constrained by allowing only within-category substitutions when calculating neighborhood density, and version that included both of these modifications to the basic representational system. As each of these constraints was imposed on the basic coding system, a concomitant reduction in neighborhood density was observed that was consistent across all the age ranges examined. The only age range in which a reliable effect of sample period on neighborhood density was observed was Age I, corresponding to the period around two years. The same pattern of results related to age was obtained when the phonemic code was used. A comparison of the nonphonemic code that produced the lowest neighborhood densities with the phonemic code revealed small but reliable differences. Specifically, a nonphonemic system in which consonants were represented according to manner class and vowels were fully specified produced neighborhood densities that were very similar to those produced by the phonemic system. I concluded that even though there were small differences between the two systems, the similarity between the manner class system and the phonemic system suggested that this version of the manner class code was a viable alternative to the phonemic system, especially for young children who have small lexicons.

Overall, the results of this chapter extended the findings obtained in Chapter IV. A computational approach was used to explore the effect of various nonphonemic representational systems on neighborhood density in young children. This methodology provided an empirical basis to test the hypothesis of researchers such as Jusczyk (1986) and Charles-Luce and Luce (1990), who proposed that young children represent information about words in a way that differs from adults. The results showed that it was possible to formulate a linguistically-based, nonphonemic representational system that produced neighborhood densities that did not differ substantially from those produced by a phonemic representational system.

Chapter VI. Children's Lexicons: Summary and Conclusions

Summary

The present investigation examined the representation and organization of spoken words in the mental lexicons of young children. The neighborhood similarity metric was used to calculate neighborhood densities that served as an index of lexical organization. The average number of neighbors permitted the comparison of several different types of representational systems, including an adult phonemic system and several nonphonemic representational systems. Part of the motivation for the present investigation was the earlier findings of Charles-Luce and Luce (1990) who showed that words in the lexicons of 5- and 7-year-olds had many fewer neighbors than the same words in the adult lexicon. Based on their analyses of a database containing a large number of children's utterances, Charles-Luce and Luce suggested that the way words are represented in children's lexicons could differ substantially from the way words are represented in the adult lexicon. Another motivation for the present investigation was research examining child language and cognitive development. These studies suggested that many developmental processes could be characterized by a progression from analysis at the level of wholes to analysis of parts. That is, a representation of perceptual information that is relatively undifferentiated eventually develops into a representation in which the components that comprise the whole are themselves specified (Aslin & Smith, 1988).

In the present investigation, neighborhood density was calculated for language samples from children between the ages of 18 months and 5 years. These samples were obtained from the CHILDES database (MacWhinney & Snow, 1985). Three age ranges were investigated within this period: Age I, corresponding to 1;6 – 2;11 (four children); Age II, corresponding to 2;10 – 3;05 (three children); and Age III, corresponding to 3;04 – 5;01 (two children).

In the first analysis, segmental representations were assumed to be equivalent to an adult phonemic system. Mean neighborhood densities were calculated as a function of sample period and word length in phonemes. The largest increase in mean neighborhood density occurred at Age I; at the later ages, neighborhood density either increased only slightly or remained relatively stable. This increase in neighborhood density around 2 years of age coincides with the onset of the naming explosion, in which children rapidly add words to their lexicons. The lack of change in neighborhood density at Ages I and II was attributed, in part, to limitations in the database itself. The variable length produced the predictable effect of reduced density for longer words. These results replicated and extended in several ways the findings of Charles-Luce and Luce (1990) to children younger than 5 years. Finally, correlations among several lexical variables were found to be similar in adult and children's lexicons, suggesting that some of the same structural characteristics of English words, such as the positive relationship between word frequency and neighborhood density, affect both child and adult lexicons. The relatively low neighborhood densities observed, especially at the earliest ages, suggests that children may be able to uniquely identify an item in the lexicon without necessarily utilizing a traditional segmental phonemic representational system. This, in turn, implies that children may not have to uniquely specify the internal structure of words in order to reliably identify them in many contexts.

The effect of nonphonemic representation on neighborhood densities was also explored. Alternative coding systems assigned phonemes to broad categories based on manner of articulation, place of articulation, and a randomly-defined set of categories not based on any known linguistic principles. Broad categories such as these are the simplest and most economical representational system that still allows the unique specification of individual lexical items, at least when the size of the lexicon is small (Huttenlocher & Zue, 1983; Pisoni et al., 1985)). The random system was included to provide a nonphonemic baseline condition to compare with the two linguistically-based representational systems. In addition, the specificity of each code was assessed by varying whether vowels were fully specified or not, and whether neighborhood density was calculated using a modified version of the standard similarity metric in which substitutions were restricted to only within-category substitutions.

Analyses that paralleled those described above for the phonemic system were carried out to assess the consequences of using nonphonemic representations. Overall, neighborhood densities were substantially increased when the nonphonemic codes were used. As observed in the earlier analyses, the only age period in which a significant change in neighborhood density was obtained was at Age I. Several interesting effects related to the number of neighbors produced by each nonphonemic system were noted. First, the two linguistically-based systems both produced a larger number of neighbors than the random code. The linguistic coding systems apparently captured some of the phonological regularities present in English that the random system did not. Second, of the two linguistically-based systems, the manner class code produced fewer neighbors than the place of articulation code. Finally, when the manner class code was combined with fully specified vowels and within-category substitutions, the mean number of neighbors was very close to the number of neighbors produced by the phonemic code for the period between 1 1/2 and 5 years. Thus, it seems plausible, at least from a computational perspective, that young children could utilize a nonphonemic representational system without substantially impeding their ability to uniquely identify words in their lexicon. The results of these analyses provided a quantitative measure of the relationship between phonemic and various nonphonemic representational systems, as well as providing support for the proposal that young children may have a phonemic representational system for lexical information that is different from the system used by adults.

Taken together, the analyses of phonemic and nonphonemic representational systems in the present investigation produced several important results. First, this investigation provided initial quantitative information about neighborhood densities in children below the age of five. Second, the investigation demonstrated similarities in the relationship among several lexical variables in the lexicons of children and adults. Third, the investigation obtained quantitative information about neighborhood densities assuming various nonphonemic representational systems. Finally, the study provided preliminary evidence regarding the plausibility of speculations made by theorists (e.g., Ferguson & Farwell, 1975; Jusczyk, 1986; Waterson, 1971) who have argued for the existence of nonphonemic representational systems in early childhood. Each of these findings demonstrated that a computational methodology is useful for studying the development of lexical representation and organization in young children.

Limitations of the Present Work

The results obtained in the present investigation provided important insights into the organization and representation of lexical information by children. However, it is important that the work described in the present investigation be viewed as preliminary and part of an ongoing long-term research program to learn more about lexical representation and word recognition in children. There are a several aspects of the present investigation that necessitate additional work, including the selection of data for analysis, the choice of nonphonemic representational systems, and the limitations of computational analyses. Each of these issues is discussed below.

Selection of Data for Analysis

In the present investigation, the original data were selected from the CHILDES database. The range of data selected was limited by several constraints including the age period for which data was available from individual children, the number of words in each sample, and the alignment of samples across children. The basic issue associated with each of these constraints is variability. Variability is inherent in child language development and, as a consequence, comparisons across children that attempt to quantify some aspect of the developmental process are prone to more error than typically observed in work with adults. While not discounting the importance of individual differences, variability poses problems when attempting to formulate a general account of development. Solutions to this problem in the present investigation included obtaining as large a sample size as possible from each child and aligning sample files from different children according to type/token ratio, an index of lexical development. However, this latter strategy was an imperfect solution and substantial variability among the children within each age range still remained. The addition of further data to the database could reduce the effect of variability due to these constraints. However, at the present time, the database has not been enlarged sufficiently to warrant redoing the analyses with data from more children.

Another important issue regarding the CHILDES database concerns its suitability for answering the questions considered in the present investigation. Can a database containing data collected from diverse sources for a variety of purposes provide useful information for answering questions about lexical representation and organization that focus on phonological form? To the extent that lexical statistics generated from written language accurately predict the performance of adult subjects tested in various types of experiments involving spoken language, it would not be unreasonable to expect that statistics generated from an analysis of the spoken language of children might also be useful in predicting the performance of children in other types of tasks. The strength of this analogy rests on the assumption that, despite the variability in child language, certain basic kinds of information, such as neighborhood densities and word frequencies, can be captured by computational analyses of the type carried out in the present investigation.

The database analyzed in the present investigation proved adequate for answering basic questions regarding lexical representation and organization despite the variability of sources from which the language samples were originally obtained. However, the present analyses did reveal one potential shortcoming of the CHILDES database. The lack of change for overall neighborhood density during Age II and Age III could be because the CHILDES database contains *only* spoken language. As discussed in Chapter IV, a child's productive vocabulary is unlikely to completely reflect their growing receptive vocabulary. If the goal of research on lexical development is to examine issues of representation and organization, then any database that contains exclusively production data may limit somewhat the generalizability of findings, especially for children in the 3 to 5 year range.

Finally, the results of the present investigation were based exclusively on data from children reared in an English-speaking environment. It is intriguing to consider whether a similar pattern of results would be obtained with utterances produced by children from other linguistic groups, especially from languages containing a substantially different inventory of phonemes (and, as a consequence, different distributions of segments in terms of manner and place of articulation). In short, additional analyses using other samples are required to determine if the results of the present investigation generalize to children learning other languages.

Nonphonemic Representation

The analyses of nonphonemic representational systems carried out in the present investigation made a number of assumptions, both tacit and explicit, that require careful consideration. First, I assumed that although the representations were not phonemic, they nevertheless maintained the segmental quality of phonemic representation (i.e., the representations were comprised of units roughly the size of phonemes). Second, I assumed that all portions of a word were equivalent with respect to their potential for confusability in memory. Third, while I assumed that the representation of consonants in memory could be ambiguous, vowels were presumed to be unambiguously represented in memory. Fourth, I did not take into account the effect of suprasegmental information, such as lexical stress. Although not unreasonable for an initial investigation of lexical representation and organization, it is unlikely that all of these assumptions were fully warranted or justified.

First, evidence discussed in Chapter III emphasized the point that the beginnings and endings of words are more readily preserved by children than medial portions (e.g., Chaney, 1989; Walley, 1988). Thus, a more realistic way to approach nonphonemic representation would be to somehow incorporate a differential weighting of information at the beginnings and ends of words than for information in the medial portions. The weighting of different parts of a word should have some empirical basis but, as of yet, no systematic information about the formulation of such weights has been proposed.

Second, the assumption of a sequentially-ordered segmental representation is questionable given that children appear to have a number of difficulties dealing with units that are smaller than subsyllabic structures, such as onsets and rimes, and that they frequently reduce or delete segments such as consonant clusters (e.g., Chiat, 1979; Menyuk et al., 1986; Treiman, 1985). Instead of an adult-like, segmentally organized system for lexical organization, these findings suggest that the units of child language are larger than the phoneme sized units assumed in the present investigation. While recognizing this as a potential problem for the assumption of segmental representation, it is unclear that a principled alternative to the approach used in the present investigation is currently available. To better deal with this problem, it would be necessary to learn more about when children can and cannot extract segmental information from words. For example, how effectively can children deal with consonant clusters or information at the beginnings, ends, and middle portions of words? Moreover, a child may correctly produce an individual word while other words containing similar sequences of phonemes may not be produced correctly, compounding further the problem of developing general principles for describing the size and organization of lexical units.

Third, the assumption that vowels are fully specified and relatively immune to the problems associated with consonants may not be warranted. Several studies indicate that vowels are not always produced in a manner consistent with the adult model (e.g., Davis & MacNeillage, 1990; Menyuk, et al., 1986). Assuming that misproduction of vowels reflects a problem in the representation of vowel information, it would be necessary to determine those vowels that are most accurately represented, assume that they are fully-specified, and that the remainder are imperfectly specified. However, since it is not clear that problems in production are ultimately caused by imperfectly specified representations, the impact of research in this area on the present investigation remains uncertain .

Fourth, no attempt was made in the present investigation to incorporate any form of suprasegmental information, such as stress pattern, that may also be used for determining similarity relations among items in the lexicon. Information provided by stress pattern has been shown to useful for adults in identifying multisyllabic utterances (Cutler & Norris, 1988). In addition, stress was shown to be effective in reducing search space size in speech recognition work (Huttenlocher & Zue, 1983). Indeed, some researchers have suggested that because stress is one of the most acoustically salient features of spoken language, it is one of the first characteristics of language extracted from the speech signal by the child (e.g., Boysson-Bardies et al., 1984; Studdert-Kennedy, 1986; Waterson, 1971). Thus, stress could have proved useful for augmenting the nonphonemic representational systems. However, since many of the words produced by young children are monosyllabic, it is doubtful that an effort to incorporate stress information would have made a substantial difference to the results of the present analyses.

In sum, several of the assumptions about the nonphonemic representation of linguistic information are questionable. However, either no clear alternatives exist at the present time that would address these problems, or they were deemed not to have serious consequences for the present investigation. Thus, with respect to assumptions regarding nonphonemic representation, the present investigation may be viewed as a compromise that provided a initial picture of the representation and organization of words in children's lexicons, but one that will need to be further elaborated.

Limitations of Computational Analyses

Earlier in the present section I asked whether a lexical database can provide useful information for answering questions about lexical representation and organization that focus on phonological form. In a general sense, this question is related to the overall usefulness of computational analyses in developmental work. Ultimately, the answer to this question will rest on behavioral evidence from children tested in experiments, such as word recognition tasks, that utilize stimuli chosen on the bases of statistics generated from the database. Until these experiments are carried out, the usefulness of computational analyses in developmental work will remain restricted to making concrete various hypotheses proposed by theorists. However, this itself is an extremely useful function and one that should not be underestimated. Influential positions in linguistics and other disciplines have sometimes been accepted in the past with little empirical data to back them up. For example, Jakobson's position regarding the presence of a universal set of phonemes in babbling was accepted somewhat blindly for more than thirty years (Locke, 1983).

Conclusion

Research on the representation and organization of information in children's lexicons has reached a critical point where theories expressed as verbal statements cannot be easily tested in a simple experiment (cf. Wickens, 1982). For example, the consequences of assuming various types of nonphonemic representational systems are difficult, if not impossible, to understand on the basis of verbal descriptions alone. Simply put, the number of variables and their complex interactions make it necessary to formalize these ideas more precisely and develop procedures that allow these complexities to be tested in a suitable manner. Computational procedures such as those used in the present investigation are one way to deal with questions of this kind, especially problems with language. Like other attempts to model psychological processes, its most obvious advantage is that it can take the assumptions of a theorist and demonstrate the consequences of those assumptions. And, like other models, it also makes the assumptions themselves available for explicit analysis and modification.

Utilizing computational procedures to investigate lexical representation and organization in children is in its infancy. Nevertheless, it holds much promise as a means of clarifying and empirically testing hypothesis in this area, as well as generating novel predictions for behavioral experiments that have yet to be done. Substantial work remains to be done to develop improved procedures for computational analysis of the kind undertaken here. However, pursuit of these goals has the potential to make an important and lasting contribution to understanding some very basic issues in psycholinguistics and cognitive development and relations between these two fields.

References

Aitchison, J., & Chiat, S. (1981). Natural phonology or natural memory? The interaction between phonological processes and recall mechanisms. *Language and Speech*, **24**, 311-326.

Aitchison, J., & Straf, M. (1982). Lexical storage and retrieval: A developing skill? *Linguistics*, **19**, 751-795.

Altmann, G., & Carter, D. (1989). Lexical stress and lexical discriminability: Stressed syllables are more informative, but why? *Computer Speech and Language*, **3**, 265-275.

Altmann, G. (1990). Cognitive models of speech processing: An introduction. In G. Altmann (Ed.), *Cognitive models of speech processing* (pp. 1-23). Cambridge, MA: MIT Press.

Archangeli, D. (1984). Underspecification in Yawelmani phonology and morphology. Doctoral dissertation, MIT.

Archangeli, D. (1988). Aspects of underspecification theory. *Phonology*, **5**, 183-207.

Aslin, R. N., Pisoni, D. B., & Jusczyk, P. W. (1983). Auditory development and speech perception in infancy. In M. Haith & J. Campos (Eds.), *Infancy and developmental psychobiology* (pp. 573-687). New York: Wiley.

Aslin, R. N., & Smith, L. B. (1988). Perceptual development. *Annual Review of Psychology*, **39**, 123-124.

Barton, D. (1975). Statistical significance in phonemic perception experiments. *Journal of Child Language*, **2**, 297-298.

Barton, D. (1980). Phonemic perception in children. In G. H. Yeni-Komshian, J. F. Kavanagh, & C. A. Ferguson (Eds.), *Child phonology: Perception* (pp. 97-116). New York: Academic Press.

Becker, C. (1976). Allocation of attention during visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, **2**, 556-566.

Berko, J. (1958). The child's learning of English morphology. *Word*, **14**, 150-177.

Benedict, H. (1979). Early lexical development: Comprehension and production. *Journal of Child Language*, **6**, 183-200.

Bertoncini, J., & Mehler, J. (1981). Syllables as units in infant speech perception. *Infant Behavior and Development*, **4**, 247-260.

Bertoncini, J., Bijeljac-Bibic, R., Jusczyk, P., Kennedy, L., & Mehler, J. (1988). An investigation of young infants' perceptual representations of speech sounds. *Journal of Experimental Psychology: General*, **117**, 21-33.

Best, C. T., McRoberts, G. W., & Sithole, N. M. (1988). Examination of perceptual reorganization for nonnative speech contrasts: Zulu click discrimination by English-speaking adults and infants. *Journal of Experimental Psychology: Human Perception and Performance*, **14**, 345-360.

Bloomfield, L. (1933). *Language.* New York: Holt.

Bolinger, D. (1975). *Aspects of language (2 ed.).* New York: Harcourt, Brace, and Jovanovich.

Bontrager, T. (1991). The development of word frequency lists prior to the 1944 Thorndike-Lorge list. *Reading Psychology: An International Quarterly*, **12**, 91- 116.

Boysson-Bardies, B., Sagart, L., & Durand, C. (1984). Discernable differences in the babbling of infants according to target language. *Journal of Child Language*, **11**, 1-15.

Bradley, D., Garrett, M., & Zurif, E. (1980). Syntactic deficits in Broca's aphasia. In D. Caplan (Ed.), *Biological studies of mental processes* . Cambridge, MA: MIT Press.

Braine, M. (1976). Review of The acquisition of phonology (N. V. Smith)*Language*, **52**, 489-498.

Broadbent, D. (1967). Word-frequency effect and response bias. *Psychological Review*, **74**, 1-15.

Brown, R., & McNeill, D. (1966). The "tip of the tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior*, **5**, 325-337.

Brown, R. (1973). *A first language: The early stages*. Cambridge, MA: Harvard University Press.

Butterworth, B. (1983). Lexical representation. In B. Butterworth (Ed.), *Language production* (pp. 257-294). New York: Academic Press.

Cassidy, K., & Kelly, M. (1991). Phonological information for grammatical category assignments. *Journal of Memory and Language*, **30**, 348-369.

Chaney, C. (1989). I pledge a legiance tothe flag: Three studies in word segmentation. *Applied Psycholinguistics*, **10**, 261-281.

Charles-Luce, J., & Luce, P. (1990). Similarity neighbourhoods of words in young children's lexicons. *Journal of Child Language*, **17**, 205-215.

Chi, M. (1976). Short-term memory limitations in children: Capacity or processing deficits? *Memory & Cognition*, **4**, 559-572.

Chiat, S. (1979). The role of the word in phonological development. *Linguistics*, **17**, 591-610.

Chomsky, N. (1965). *Aspects of the theory of syntax.* Cambridge, MA: MIT Press.

Chomsky, N., & Halle, M. (1968). *The sound pattern of English.* New York: Harper and Row.

Clements, G. N., & Keyser, S. J. (1983). *CV phonology.* Cambridge. MA: MIT Press.

Cluff, M., & Luce, P. A. (1990). Similarity neighborhoods of spoken two-syllable words: Retroactive effects on multiple activation. *Journal of Experimental Psychology: Human Perception and Performance*, **16**, 551-563.

Cole, R. (1973). Listening for mispronunciations: A measure of what we hear during speech. *Perception & Psychophysics*, **11**, 153-156.

Conrad, R. (1964). Acoustic confusions in immediate memory. *British Journal of Psychology*, **55**, 75-84.

Corrigan, R. (1978). Language development as related to stage 6 object permanence development. *Journal of Child Language*, **5**, 173-189.

Cowan, N., Saults, J. S.,Winterowd, C., & Sherk, M. (1991). Enhancement of 4-year-old children's memory span for phonologically similar and dissimilar word lists. *Journal of Experimental Child Psychology*, **51**, 30-52.

Crowder, R. (1978). *Principles of learning and memory*. Hillsdale, NJ: Erlbaum.

Crystal, D. (1991). *A dictionary of linguistics and phonetics*. Oxford, UK: Blackwell.

Cutler, A., & Carter, D. M. (1987). The predominance of strong initial syllables in the English vocabulary. *Computer Speech and Language*, **2**, 133-142.

Cutler, A., Mehler, J.,Norris, D., & Segui, J. (1987). Phoneme identification and the lexicon. *Cognitive Psychology*, **19**, 141-177.

Cutler, A., & Norris, D. (1988). The role of strong syllables in segmentation for lexical access. *Journal of Experimental Psychology: Human Perception and Performance*, **14**, 113-121.

Dell, G. S. (1986). A spreading activation theory of retrieval in sentence production. *Psychological Review*, **93**, 283-321.

Denes, P. B. (1963). On the statistics of spoken English. *Journal of the Acoustical Society of America*, **35**, 892-904.

Dewey, G. D. (1923). *Relativ frequency of English speech sounds*. Cambridge: Harvard University Press.

Diamond, A. (1985). Development of the ability to use recall to guide action as predicted by the infant's performance on AB. *Child Development*, **56**, 868-883.

Donegan, P., & Stampe, D. (1979). The study of natural phonology. In D. Dinnsen (Ed.), *Current approaches to phonological theory* (pp. 126-173). Bloomington, IN: Indiana University Press.

Eilers, R. E., & Oller, D. K. (1976). The role of speech discrimination in developmental sound substitutions. *Journal of Child Language*, **3**, 319-330.

Eimas, P., Miller, J., & Jusczyk, P. (1987). On infant speech perception and the acquisition of language. In S. Harnad (Ed.), *Categorical perception: The groundwork of cognition* (pp. 161-195). Cambridge, UK: Cambridge University Press.

Elbers, L. (1985). A tip-of-the-tongue experience at age two. Journal of *Child Language*, **12**, 353-365.

Elliott, L. L., Hammer, M. A., & Evan, K. E. (1987). Perception of gated, highly familiar spoken monosyllabic nouns by children, teenagers, and older adults. *Perception & Psychophysics*, **42**, 150-157.

Emmorey, K. D., & Fromkin, V. A. (1988). The mental lexicon. In F. Newmeyer (Ed.), *Cambridge survey of linguistics* (pp. 124-149). New York: Cambridge University Press.

Enns, J. T., & Girgus, J. S. (1985). Developmental changes in selective and integrative visual attention. *Journal of Experimental Child Psychology*, **40**, 319-337.

Enns, J. T., & Akhtar, N. (1989). A developmental study of filtering in visual attention. *Child Development*, **60**, 1188-1199.

Ferguson, C. A., & Farwell, C. B. (1975). Words and sounds in early language acquisition. *Language*, **51**, 419-439.

Fernald, A. (1985). Four-month-old infants prefer to listen to motherese. *Infant Behavior and Development*, **8**, 181-195.

Forster, K. I. (1976). Accessing the mental lexicon. In R. Wales & E. Walker (Eds.), *New approaches to language mechanisms* (pp. 139-174). Amsterdam: North-Holland.

Forster, K. I. (1979). Levels of processing and the structure of the language processor. In W. E. Cooper & E. Walker (Eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett* (pp. 27-85). Hillsdale, NJ: Erlbaum.

Foss, D., & Gernsbacher, M. (1983). Cracking the dual code: Towards a unitary model of phoneme identification. *Journal of Verbal Learning and Verbal Behavior*, **22**, 609-632.

Fromkin, V. A. (1987). The lexicon: Evidence from acquired dyslexia. *Language*, **63**, 1-22.

Fry, D. B. (1955). Duration and intensity as physical correlates of linguistic stress. *Journal of the Acoustic Society of America*, **35**, 765-769.

Garber, E., & Pisoni, D. (1991). Lexical memory in visual and auditory modalities: A second report. *Research on Speech Perception Progress Report No. 17.* Bloomington, IN: Speech Research Laboratory, Indiana University.

Garner, W. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.

Garnica, O. K. (1973). The development of phonemic speech perception. In T. Moore (Ed.), *Cognitive development and the acquisition of language* . New York: Academic Press.

Geffon, G., & Sexton, M. (1978). The development of auditory strategies of attention. *Developmental Psychology*, **14**, 11-17.

Gibson, E. J. (1969). *Principles of perceptual learning and development.* Englewood Cliffs, NJ: Prentice-Hall.

Gleitman, L. R., & Gleitman, H. (1992). A picture is worth a thousand words, but that's the problem: The role of syntax in vocabulary acquisition. *Current Directions in Psychological Science*, **1**, 31-35.

Goldfield, B. A., & Reznick, J. S. (1990). Early lexical acquisition: Rate, content, and the vocabulary spurt. *Journal of Child Language*, **17**, 171-183.

Goldinger, S. D., Luce, P. A., & Pisoni, D. B. (1989). Priming lexical neighbors of spoken words: Effects of competition and inhibition. *Journal of Memory and Language*, **28**, 501-518.

Goldinger, S. D., Pisoni, D. B., & Logan, J. S. (1991). On the nature of talker variability effects on the recall of spoken word lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **17**, 152-162.

Goldsmith, J. (1990). *Autosegmental and metrical phonology.* Oxford, UK: Blackwell.

Gopnik, A., & Meltzoff, A. (1987). The development of categorization in the second year and its relation to other cognitive and linguistic developments. *Child Development*, **58**, 1523-1531.

Goswami, U., & Bryant, P. (1990). *Phonological skills and learning to read.* Hove, UK: Erlbaum.

Graham, L. W., & House, A. S. (1971). Phonological oppositions in children: A perceptual study. *Journal of the Acoustical Society of America*, **49**, 559-566.

Greenberg, J. H., & Jenkins, J. J. (1963). Studies in the psychological correlates of the sound system of American English. *Word*, **20**, 157-177.

Grieser, D., & Kuhl, P. K. (1989). Categorization of speech by infants: Support for speech-sound prototypes. *Developmental Psychology*, **25**, 577-588.

Halle, M. (1985). Speculations about the representation of words in memory. In V. Fromkin (Ed.), *Phonetic linguistics: Papers presented to Peter Ladefoged* (pp. 101-114). New York: Academic Press.

Halle, M. (1986). On speech sounds and their immanent structure. Unpublished manuscript, MIT.

Harrington, J., & Johnstone, A. (1987). The effects of equivalence classes on parsing phonemes into words in continuous speech recognition. *Computer Speech and Language*, **2**, 273-288.

Howes, D. H., & Solomon, R. L. (1951). Visual duration threshold as a function of word probability. *Journal of Experimental Psychology*, **41**, 401-410.

Hulme, C. (1984). Developmental differences in the effects of acoustic similarity on memory span. *Developmental Psychology*, **20**, 650-652.

Hulme, C., & Tordoff, V. (1989). Working memory development: The effects of speech rate, word length, and acoustic similarity on serial recall. *Journal of Experimental Child Psychology*, **47**, 72-87.

Huttenlocher, D. P., & Zue, V. W. (1983). Exploring phonotactic and lexical constraints in word recognition. In 106th meeting of the Acoustical Society of America, San Diego, CA.

Hyman, L. M. (1975). *Phonology: Theory and analysis.* New York: Holt, Rinehart, and Winston.

Ingram, D. (1989). *First language acquisition: Method, description, and explanation.* Cambridge, UK: Cambridge University Press:

Irwin, J. V., & Wong, S. P. (1983). *Phonological development in children 18 to 72 months.* Carbondale, IL: Southern Illinois University Press.

Iverson, G., & Wheeler, D. (1987). Hierarchical structures in child phonology. *Lingua*, **73**, 243-257.

Jakobson, R., Fant, G., & Halle, M. (1952). *Preliminaries to speech analysis.* Cambridge, MA: MIT Press.

Jakobson, R., & Halle, M. (1956). *Fundamentals of language.* The Hague: Mouton.

Jusczyk, P. W. (1982). Auditory versus phonetic coding of speech signals during infancy. In J. Mehler, E. Walker, & M. Garrett (Eds.), *Perspectives on mental representation* (pp. 361-387). Hillsdale, NJ: Erlbaum.

Jusczyk, P. W. (1989). Developing phonological categories from the speech signal. Paper presented at International Conference on Phonological Development. Stanford University:

Jusczyk, P. W., & Derrah, C. (1987). Representation of speech sounds by young infants. *Developmental Psychology*, **23**, 648-654.

Jusczyk, P. W., Friederici, A. D., & Wessels, J. (in preparation). Infants' recognition of foreign versus native language words.

Jusczyk, P.W., Hirsh-Pasek, K., Kemler Nelson, D., Kennedy, L., Woodward, A., & Piwoz, J. (1992). Perception of acoustic correlates of major phrasal units by young infants. *Cognitive Psychology*, **24**, 252-293.

Kail, R. (1990). *The development of memory in children (3rd ed.).* New York: Freeman.

Kaye, J. (1989). *Phonology: A cognitive view.* Hillsdale, NJ: Erlbaum.

Keating, P. A. (1988). *A survey of phonological features.* Bloomington, IN: Indiana University Linguistics Club.

Kelly, M. H. (1992). Using sound to solve syntactic problems: The role of phonology in grammatical category assignments. *Psychological Review*, **99**, 349-364.

Kelly, M. H., & Bock, J. K. (1988). Stress in time. *Journal of Experimental Psychology: Human Perception and Performance*, **14**, 389-403.

Kemler-Nelson, D. G. (1983). Exploring and reexploring issues of integrality, perceptual sensitivity, and dimensional salience. *Journal of Experimental Child Psychology*, **36**, 365-379.

Klatt, D. (1979). Speech perception: A model of acoustic-phonetic analysis and lexical access. *Journal of Phonetics*, **7**, 279-312.

Klatt, D. H. (1989). Review of selected models of speech perception. In W. D. Marslen-Wilson (Ed.), *Lexical representation and process* (pp. 201-262). Cambridge, MA: MIT Press.

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

Kuhl, P. K. (1979). Speech perception in early infancy: Perceptual constancy for spectrally dissimilar vowel categories. *Journal of the Acoustical Society of America*, **66**, 1668-1679.

Kuhl, P. K. (1983). Perception of auditory equivalence classes for speech in early infancy. *Infant Behavior and Development*, **6**, 263-285.

Lahiri, A., & Marslen-Wilson, W. (1991). The mental representation of lexical form: A phonological approach to the recognition lexicon. *Cognition*, **38**, 245-294.

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.

Lane, D. M., & Pearson, D. A. (1982). The development of selective attention. *Merrill-Palmer Quarterly*, **28**, 317-337.

Larar, J. N. (1986). Lexical access using broad acoustic-phonetic classifications. *Computer Speech and Language*, **1**, 47-59.

Lasky, R. E., Syrdal-Lasky, A., & Klein, R. E. (1975). VOT discrimination by four to six month old infants from Spanish environments. *Journal of Experimental Child Psychology*, **20**, 215-225.

Liberman, I. Y., Shankweiler, D.,Fischer, F. W., & Carter, B. (1974). Explicit syllable and phoneme segmentation in the young child. *Journal of Experimental Child Psychology*, **18**, 201-212.

Locke, J. L. (1980a). The inference of speech perception in the phonologically disordered child. Part I: A rationale, some criteria, the conventional tests. *Journal of Speech and Hearing Disorders*, **4**, 431-444.

Locke, J. L. (1980b). The inference of speech perception in the phonologically disordered child. Part II: Some clinically novel procedures, their uses, some findings. *Journal of Speech and Hearing Disorders*, **4**, 445-468.

Locke, J. L. (1983). *Phonological acquisition and change.* New York: Academic Press.

Locke, J. L. (1986). The sound shape of early lexical representations. In J. Locke & B. Smith (Eds.), *The emergent lexicon* (pp. 3-22). New York: Academic Press.

Luce, P. A. (1986a) Neighborhoods of words in the mental lexicon. *Technical Report Number 6.* Bloomington, IN: Speech Research Laboratory, Indiana University.

Luce, P. A. (1986b). A computational analysis of uniqueness points in auditory word recognition. *Perception & Psychophysics*, **39**, 155-158.

Luce, P. A., Pisoni, D. B., & Goldinger, S. D. (1990). Similarity neighborhoods of spoken words. In G. Altmann (Ed.), *Cognitive models of speech processing* (pp. 122-147). Cambridge, MA: MIT Press.

Macken, M. A. (1979). Developmental reorganization of phonology: A hierarchy of basic units of acquisition. *Lingua*, **49**, 11-49.

Macken, M. A. (1980). The child's lexical representation: The 'puzzle-puddle-pickle' evidence. *Journal of Linguistics*, **16**, 1-17.

MacWhinney, B., & Snow, C. (1985). The child language data exchange system. *Journal of Child Language*, **12**, 271-296.

Mandler, J. M., & Bauer, P. J. (1988). The cradle of categorization: Is the basic level basic? *Cognitive Development*, **3**, 247-264.

Mandler, J. M., Bauer, P. J., & McDonough, L. (1991). Separating the sheep from the goats: Differentiating global categories. *Cognitive Psychology*, **23**, 263-298.

Marslen-Wilson, W., & Welsh, A. (1978). Processing interactions and lexical access during word recognition in continuous speech. *Cognitive Psychology*, **10**, 29-63.

Marslen-Wilson, W. (1987). Functional parallelism in spoken word-recognition. *Cognition*, **25**, 71-102.

Maxwell, E. (1984). On determining underlying phonological representation of children: A critique of the current theories. In M. Elbert,D. Dinnsen, & G. Weismer (Eds.), *Phonological theory and the misarticulating child*. ASHA Monograph 22, Rockville, MD: American Speech-Language Association.

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

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.

McNeill, D., & Lindig, K. (1973). The perceptual reality of phonemes, syllables, words, and sentences. *Journal of Verbal Learning and Verbal Behavior*, **12**, 431-461.

Menyuk, P., Menn, L., & Silber, R. (1986). Early strategies for the perception and production of words and sounds. In P. Fletcher & M. Garman (Eds.), *Language acquisition: Studies in first language development* (pp. 198-222). Cambridge, UK: Cambridge University Press.

Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology*, **32**, 89-115.

Miller, G. A. (1951). *Language and communication*. New York: McGraw-Hill.

Miller, G. A., Heise, G. A., & Lichten, W. (1951). The intelligibility of speech as a function of the context of 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.

Mines, M. A., Hansen, B. F., & Shoup, J. E. (1978). Frequency of occurrence of phonemes in conversational English. *Language and Speech*, **21**, 221-241.

Mohanan, K. P. (in press; cited in Stemberger, 1991). On the basis of underspecification.

Morais, J., Bertelson, P.,Cary, L., & Alegria, J. (1987). Literacy training and speech segmentation. *Cognition*, **24**, 45-64.

Morton, J. (1979). Word recognition. In J. Morton & J. D. Marshall (Eds.), *Psycholinguistics 2: Structures and processes* (pp. 107-156). Cambridge, MA: MIT Press.

Moskowitz, A. I. (1973). The acquisition of phonology and syntax: A preliminary study. In G. Hintikka, J. Moravesik, & P. Suppes (Eds.), *Approaches to natural language* (pp. 48-84). Dordrecht: Reidel.

Nadel, L., & Wexler, K. (1984). Neurobiology, representations, memory. In G. Lynch, J. L. McGaugh, & N. M. Weinberger (Eds.), *Neurobiology of learning and memory* (pp. 125-134). New York: Guilford.

Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 259-303). Hillsdale, NJ: Erlbaum.

Peters, A. (1977). Language learning strategies: Does the whole equal the sum of the parts? *Language*, **53**, 560-573.

Peters, A. (1983). *The units of language acquisition.* Cambridge, UK: Cambridge University Press.

Pisoni, D. B., & Garber, E. (1991). Lexical memory in visual and auditory modalities: The case for a common mental lexicon. In H. Fujisaki, (Ed.), *Proceedings of the 1990 International Conference on Spoken Language Processing*, Kobe, Japan.

Pisoni, D. B., Nusbaum, H., Luce, P. A., & Slowiaczek, L. (1985). Speech perception, word recognition, and the structure of the lexicon. *Speech Communication*, **4**, 75-95.

Pisoni, D. B., & Luce, P. A. (1987). Acoustic-phonetic representations in word recognition. *Cognition*, **25**, 21-52.

Pollack, I., Rubenstein, H., & Decker, L. (1959). Intelligibility of known and unknown message sets. *Journal of the Acoustical Society of America*, **31**, 273-279.

Pollack, I., Rubenstein, H., & Decker, L. (1960). Analysis of incorrect responses to an unknown message set. *Journal of the Acoustical Society of America*, **32**, 454-457.

Rovee-Collier, C., Early, L., & Stafford, S. (1989). Ontogeny of early event memory: III. Attentional determinants of retrieval at 2 and 3 months. *Infant Behavior and Development*, **12**, 147-161.

Rozin, P., & Gleitman, L. R. (1977). The structure and acquisition of reading II: The reading process and the acquisition of the alphabetic principle. In A. S. Reber & D. L. Scarborough (Eds.), *Toward a psychology of reading* (pp. 55-141). Hillsdale, NJ: Erlbaum.

Rubenstein, H., Garfield, L., & Milliken, J. A. (1970). Homographic entries in the internal lexicon. *Journal of Verbal Learning and Verbal Behavior*, **9**, 487-494.

Sagey, E. C. (1986) The representation of features and relations in non-linear phonology. Doctoral Dissertation, MIT.

Samuel, A., & Ressler, W. (1986). Attention within auditory word perception: Insights from the phonemic restoration illusions. *Journal of Experimental Psychology: Human Perception and Performance*, **12**, 70-79.

Savin, H. B., & Bever, T. G. (1970). The non-perceptual reality of the phoneme. *Journal of Verbal Learning and Verbal Behavior*, **9**, 295-302.

Schwartz, R., & Leonard, L. (1982). Do children pick and choose? An examination of phonological selection and avoidance in early lexical acquisition. *Journal of Child Language*, **9**, 319-336.

Schwartz, R., & Terrell, B. Y. (1983). The role of input frequency in lexical acquisition. *Journal of Child Language*, **10**, 57-64.

Shepp, B. E., & Barrett, S. E. (1991). The development of perceived structure and attention: Evidence from divided and selective attention tasks. *Journal of Experimental Child Psychology*, **51**, 434-458.

Shipman, D. W., & Zue, V. W. (1982). Properties of large lexicons: Implications for advanced isolated word recognition systems. *International Conference on Acoustics, Speech, and Signal Processing*, 546-549.

Shvachkin, N. (1973/1948). The development of phonemic speech perception in early childhood. In C. A. Ferguson & D. I. Slobin (Eds.), *Studies of child language development* (pp. 91-127). New York: Holt, Rinehart, and Winston.

Smith, E. E. (1978). Theories of semantic memory. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes* (pp. 1-56). Potomac, MD: Erlbaum.

Smith, E. E., & Medin, D. L. (1983). *Categories and concepts*. Cambridge, MA: Harvard University Press.

Smith, N. V. (1973). *The acquisition of phonology: A case study*. Cambridge, UK: Cambridge University Press.

Smith, N. V. (1981). On the cognitive representation of developing phonology. In T. Meyers, J. Laver, & J. Anderson (Eds.), *The cognitive representation of speech* (pp. 313-321). Amsterdam: North-Holland.

Soreno, J. A., & Jongeman, A. (1990). Phonological and form class relations in the lexicon. *Journal of Psycholinguistic Research*, **19**, 387-404.

Stemberger, J. P., & MacWhinney, B. (1986). Frequency and the lexical storage of regularly inflected forms. *Memory & Cognition*, **14**, 17-26.

Stemberger, J. P. (1991). Apparent anti-frequency effects in language production: The addition bias and phonological underspecification. *Journal of Memory and Language*, **30**, 161-185.

Stemberger, J. P., & Stoel-Gammon, C. (1991). The underspecification of coronals: Evidence from language acquisition and performance errors. In C. Paradis & J.-F. Prunet (Eds.), *The special status of coronals* . San Diego, CA: Academic Press.

Stern, D. N., Spieker, S.,Barnett, R. K., & MacKain, K. (1983). The prosody of maternal speech: Infant age and context related changes. *Journal of Child Language*, **10**, 1-15.

Strange, W., & Broen, P. (1981). The relationship between perception and production of /w/, /r/, and /l/ by three-year-old children. *Journal of Experimental Child Psychology*, **31**, 81-102.

Streeter, L. (1976). Language perception of 2-month-old infants shows effects of both innate mechanisms and experience. *Nature*, **259**, 39-41.

Studdert-Kennedy, M. (1986). Sources of variability in early speech development. In J. Perkell & D. H. Klatt (Eds.), *Invariance and variability in speech processes* (pp. 58-76). Hillsdale, NJ: Erlbaum.

Sullivan, M. W. (1982). Reactivation: Priming forgotten memories in human infants. *Child Development*, **53**, 516-523.

Swinney, D. A., & Cutler, A. (1979). The access and processing of idiomatic expressions. *Journal of Verbal Learning and Verbal Behavior*, **18**, 523-539.

Taft, M. (1979). Recognition of affixed words and the word frequency effect. *Memory & Cognition*, **7**, 263-272.

Thorndike, E. L., & Lorge, I. (1944). *The teacher's word book of 30,000 words*. New York: Teachers College, Columbia University.

Trehub, S. E. (1976). The discrimination of foreign speech contrasts by infants and adults. *Child Development*, **47**, 466-472.

Trehub, S. E.,Thorpe, L. A., & Morrongiello, B. A. (1987). Organizational processes in infants' perception of auditory patterns. *Child Development*, **58**, 741-749.

Treiman, R., & Baron, J. (1981). Segmental analysis ability: Development and relation to reading ability. In G. E. MacKinnon & T. G. Waller (Eds.), *Reading research: Advances in theory and practice (Vol. 3)* (pp. 159-198). New York: Academic Press.

Treiman, R., & Breaux, A. M. (1982). Common phoneme and overall similarity relations among spoken syllables: Their use by children and adults. *Journal of Psycholinguistic Research*, **11**, 569-598.

Treiman, R. (1985). Onsets and rimes as units of spoken syllables: Evidence from children. *Journal of Experimental Child Psychology*, **39**, 161-181.

Treiman, R. (1988). Distributional constraints and syllable structure in English. *Journal of Phonetics*, **16**, 221-229.

Treiman, R., & Danis, C. (1988). Short-term memory errors for spoken syllables are affected by the linguistic structure of the syllables. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **14**, 145-152.

Treisman, M. (1978a). A theory of the identification of complex stimuli with an application to word recognition. *Psychological Review*, **85**, 525-570.

Treisman, 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.

Vihman, M. M. (1988). Early phonological development. In J. Bernthal & N. Bankson (Eds.), *Articulation and phonological disorders* (pp. 60-109). Englewood cliffs, NJ: Prentice-Hall.

Walley, A. C.,Smith, L. B., & Jusczyk, P. W. (1986). The role of phonemes and syllables in the perceived similarity of speech sounds for children. *Memory & Cognition*, **14**, 220-229.

Walley, A. C. (1988). Spoken word recognition by young children and adults. *Cognitive Development*, **3**, 137-165.

Walley, A. C., & Metsala, J. L. (1988). The growth of lexical constraints on spoken word recognition. *Perception & Psychophysics*, **47**, 267-280.

Walley, A. C. (in press). Spoken word recognition and vocabulary growth in early childhood. In J. Charles-Luce & J. R. Sawusch (Eds.), *Theories in spoken language: Perception, production, and development* . Norwood, NJ: Ablex.

Waterson, N. (1971). Child phonology: A prosodic view. *Linguistics*, **7**, 179-211.

Wepman, J. M., & Hass, W. (1969). *A spoken word count: Children — ages 5, 6, and 7.* Chicago: Language Research Associates.

Werker, J. F., Gilbert, J. H. V.,Humphrey, K., & Tees, R. C. (1981). Developmental aspects of cross-language speech perception. *Child Development*, **52**, 349-353.

Werker, J. F., & Tees, R. C. (1984). Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development*, **7**, 49-63.

Werker, J. F. (1989). Becoming a native listener. *American Scientist*, **77**, 54-59.

Zipf, G. (1949). *Human behavior and the principle of least effort.* Cambridge, MA: Addison-Wesley.

Appendix I

Analysis of variance summary table —nonphonemic coding, Age I

Source	Sum of Squares	D.F.	Mean Square	F	Tail Probability
Mean	374696.67	1	374696.67	1606.28	0.0000
1Error	699.80	3	233.26		
time(sample)	1840.31	4	460.07	5.03	0.0129
2Error	1096.53	12	91.37		
code	35769.82	2	17884.91	202.61	0.0000
3Error	529.63	6	88.27		
tc	353.13	8	44.14	2.59	0.0340
4Error	409.09	24	17.04		
specificity	250886.43	3	83628.81	1360.43	0.0000
5Error	553.25	9	61.47		
ts	1232.67	12	102.72	4.56	0.0002
6Error	810.53	36	22.51		
cs	28829.00	6	4804.83	270.99	0.0000
7Error	319.15	18	17.73		
tcs	226.65	24	9.44	2.17	0.0063
8Error	313.61	72	4.35		
length	297664.03	5	59532.80	560.91	0.0000
9Error	572.09	15	38.13		
tl	1791.45	20	89.57	5.00	0.0000
10Error	1075.57	60	17.92		
cl	26113.57	10	2611.35	170.44	0.0000
11Error	459.63	30	15.32		
tcl	315.54	40	7.88	1.63	0.0228
12Error	580.90	120	4.84		
sl	200484.23	15	13365.61	1161.37	0.0000
13Error	517.88	45	11.50		
tsl	1479.25	60	24.65	4.44	0.0000
14Error	999.10	180	5.55		
csl	23137.48	30	771.24	267.13	0.0000
15Error	259.84	90	2.88		
tcsl	221.59	120	1.84	1.38	0.0125
16Error	481.52	360	1.33		

Appendix II

Analysis of variance summary table —nonphonemic coding, Age II

Source	Sum of Squares	D.F.	Mean Square	F	Tail Probability
Mean	222678.01	1	222678.01	364.54	0.0027
1Error	1221.68	2	610.84		
time(sample)	159.38	2	79.69	0.79	0.5154
2Error	405.70	4	101.42		
code	22111.77	2	11055.88	290.12	0.0000
3Error	152.43	4	38.10		
tc	10.43	4	2.60	0.13	0.9692
4Error	166.50	8	20.81		
specificity	149301.44	3	49767.14	400.05	0.0000
5Error	746.42	6	124.40		
ts	131.57	6	21.92	1.01	0.4636
6Error	260.95	12	21.74		
cs	17653.24	6	2942.20	239.10	0.0000
7Error	147.66	12	12.30		
tcs	12.46	12	1.03	0.18	0.9984
8Error	140.64	24	5.86		
length	180808.97	5	36161.79	206.75	0.0000
9Error	1749.06	10	174.90		
tl	129.46	10	12.94	0.57	0.8216
10Error	456.93	20	22.84		
cl	16051.59	10	1605.15	121.95	0.0000
11Error	263.25	20	13.16		
tcl	23.66	20	1.18	0.24	0.9994
12Error	197.19	40	4.92		
sl	122278.78	15	8151.91	233.72	0.0000
13Error	1046.38	30	34.87		
tsl	140.21	30	4.67	0.77	0.7791
14Error	363.65	60	6.06		
csl	14476.75	30	482.55	130.54	0.0000
15Error	221.78	60	3.69		
tcsl	31.43	60	0.52	0.34	1.0000
16Error	182.99	120	1.52		

Appendix III

Analysis of variance summary table —nonphonemic coding, Age III

Source	Sum of Squares	D.F.	Mean Square	F	Tail Probability
MEAN	497518.62	1	497518.62	552.59	0.0271
1Error	900.33	1	900.33		
time(sample)	488.31	8	61.03	0.47	0.8499
2Error	1048.91	8	131.11		
code	51097.87	2	25548.93	586.39	0.0017
3Error	87.13	2	43.56		
tc	354.62	16	22.16	1.64	0.1676
4Error	216.83	16	13.55		
specificity	334519.21	3	111506.40	554.13	0.0001
5Error	603.68	3	201.22		
ts	437.25	24	18.21	0.53	0.9347
6Error	819.91	24	34.16		
cs	39096.00	6	6516.00	717.97	0.0000
7Error	54.45	6	9.07		
tcs	212.66	48	4.43	1.23	0.2423
8Error	173.60	48	3.61		
length	403102.26	5	80620.45	330.20	0.0000
9Error	1220.76	5	244.15		
tl	503.07	40	12.57	0.44	0.9951
10Error	1156.37	40	28.90		
cl	35944.40	10	3594.44	351.78	0.0000
11Error	102.17	10	10.21		
tcl	376.63	80	4.70	0.93	0.6206
12Error	403.47	80	5.04		
sl	273723.73	15	18248.24	368.59	0.0000
13Error	742.62	15	49.50		
tsl	476.61	120	3.97	0.49	1.0000
14Error	982.46	120	8.18		
csl	30736.12	30	1024.53	512.75	0.0000
15Error	59.94	30	1.99		
tcsl	278.33	240	1.15	0.94	0.6897
16Error	296.71	240	1.23		