Psycholinguistic Models of Speech Development and Their Application to Clinical Practice

TUTORIAL

This article presents an introduction to psycholinguistic models of speech development. Two specific types of models are addressed: *box-and-arrow* models and *connectionist* or *neural network* models. We review some historical and some current models and discuss recent applications of such models to the management of speech impairment in children. We suggest that there are two ways in which a psycholinguistic approach can influence clinical practice: by directly supplementing a speech-language pathologist's repertoire of assessment and treatment approaches and by offering a new way to conceptualize speech impairment in children.

KEY WORDS: speech development, psycholinguistic models, connectionist models, impairment, assessment

S ince Ingram's (1976) seminal work entitled *Phonological Disability in Children*, speech-language pathologists have increasingly applied linguistic-based approaches to their clinical practice. Such approaches have not only provided greater insight into children's phonological systems; they have also provided new approaches to intervention, based on attempts to change children's phonological rule systems rather than to correct faulty motor behaviors (e.g., Bernhardt & Stemberger, 2000; Grunwell, 1987; Hodson & Paden, 1991; Stoel-Gammon & Dunn, 1985).

As new theories have emerged in the area of speech development and impairment, clinical researchers have endeavored to bring them to the attention of practicing clinicians in order to keep clinicians abreast of theoretical developments in the field and to inform clinical practice. Bernhardt and Stoel-Gammon's (1994) tutorial on the clinical application of nonlinear phonology, for example, provides a particularly wellexecuted instance of the ongoing uses of linguistic theories to aid our understanding of children's phonological development and impairment.

As Stackhouse and Wells (1997) point out, however, linguistic approaches offer only one perspective for studying children's speech difficulties. Although such approaches allow highly detailed *descriptions* of children's phonological systems, they do not provide *explanations* for

Elise Baker Karen Croot* Sharynne McLeod** The University of Sydney Lidcombe, NSW, Australia

Rhea Paul***

Southern Connecticut State University and Yale Child Study Centre New Haven, CT

^{*}Currently affiliated with Macquarie University, Sydney, Australia

^{**}Currently affiliated with Charles Sturt University, Albury, Australia

^{***}At the time a visiting scholar at the School of Communication Sciences, The University of Sydney, Lidcombe, NSW, Australia

why individual systems take the normally developing or impaired forms they do.

An alternative perspective with a longer history in speech-language pathology is the medical perspective. The medical perspective does aim to explain the underlying cause of speech impairments when these are due to identifiable organic problems (such as cleft palate) for which there is the possibility of medical intervention. Linguistic and medical perspectives on speech development and impairment clearly complement one another. The first allows a description of the *language system* the child is using at any point in his or her development, whereas the second considers the integrity of the *neuroanatomical system* supporting speech and language. Both approaches, however, are limited in their potential to explain speech impairments of unknown origin.

By contrast, a third perspective embraces the goal of explaining speech impairment. *Psycholinguistic* approaches to speech and language development aim to explicate the way in which children process speech and language at a cognitive or psychological level and thus aim to formulate hypotheses about the psychological processes or components that may be impaired. Psycholinguistics is a subdiscipline within the broad field of psychology. The broad aim of psychology is to explain human behavior; the aim of psycholinguistics is to explain human linguistic behavior.

Psycholinguists approach this task by proposing theoretical models. A primary goal of any theoretical model is to capture the key components of a system and make to explicit the relationships among those components. In a psycholinguistic model of speech development the key components are the psychological processes involved in the "perception, storage, planning and production of speech as it is produced in real time in real utterances" (McCormack, 1997, p. 4). At the simplest level, psycholinguistic models highlight three major aspects of speech processing: the receptive processing of words, the storage or underlying representations of words, and the processes involved in their production (Dodd, 1995; Fee, 1995). More sophisticated models provide more detailed accounts of the operations at each of these levels. Psycholinguistic models therefore provide a framework for explaining the descriptive or symptomatic information about impaired phonological systems derived from linguistic-based assessments by attempting to identify the level at which speech processing is disrupted (Stackhouse & Wells, 1993).

Much of the impetus behind the attempts to model children's speech development from a psycholinguistic perspective has come from one of the fundamental mysteries of childhood speech, the [fɪs] phenomenon (Berko & Brown, 1960). We see this phenomenon in operation when an adult requests clarification of a child's pronunciation. For example, "Did you say [fis]?", only to be given an unchanged, but more insistent, response: "No, I said [fis]." The child is thought to be able to perceive the adult pronunciation of the word FISH, but not to be able to faithfully reproduce the word in his or her speech.

Our intentions in offering this tutorial on psycholinguistic approaches to speech development and impairment are two. First, we hope to make more accessible some of the recent theoretical work that has explored how aspects of speech development such as the [fis] phenomenon can be understood and how children's speech progresses from a heavy dependence on simplification of output to increasingly consistent adult-like forms. To this end, we will endeavor to introduce the reader to some of the terminology frequently used in psycholinguistic models and will discuss a selection of historically influential box-and-arrow models to illustrate the fundamentals of the approach. The models we focus on are those of Smith (1973) and Menn and colleagues (Kiparsky & Menn, 1977; Menn, 1978; Menn, Markey, Mozer, Lewis, 1993; Menn & Matthei, 1992).¹We follow this by presenting some recent clinical applications of such models, including those of Hewlett (1990) and Stackhouse and Wells (1997). We conclude our presentation of current theoretical approaches by introducing connectionist models, which are recent arrivals on the theoretical scene and extend the range of hypotheses that can be tested about children's speech and language skills. Because of the relatively new perspective connectionist models bring to the field, our description of these models is in comparison more detailed.

Our second aim is to show that such theoretical understanding can have important effects on clinical practice-not only in influencing assessment and intervention procedures, but in reshaping our thinking about the nature of speech impairment. For example, in the 1970s, the concept of the phonological process fundamentally reordered the way clinicians thought about children's speech problems. Before this time, "speech therapy" had been about teaching children "how" to produce sounds-the implicit assumption being that youngsters made errors because they had not yet acquired the proper motor skills for accurate speech production. Intervention focused on initiating correct production through articulatory modification and practice to achieve automaticity in motor performance. When linguistic theories such as generative phonology (e.g., Chomsky & Halle, 1968) and natural phonology (e.g., Stampe, 1979) proposed systematic relationships between underlying representations of words' sound structure and the way

¹ For a more detailed critique and historical account of the development of box-and-arrow models the reader is referred to Maxwell (1984), Menn and Matthei (1992), Vihman (1996), and Bernhardt and Stemberger (1998).

people ended up saying them, the phonological process became of interest. Broadly defined, a phonological process is a linguistic rule that converts a child's perception of a word into his or her own production of it. As had happened before in the case of transformational grammar, this new interest in children's linguistic behavior helped carry a new linguistic theory into the realm of speech-language pathology. And once it arrived, its effects were profound. It changed the way we assess children's speech-moving us from looking for errors on individual sounds in particular positions to looking for patterns of error that generalize across several contexts. It expanded the repertoire of intervention approaches to include some aimed at changing the organization of sound structure in the child's implicit linguistic system, as well as those based exclusively on changing motor output. But perhaps most importantly, it changed the way we conceptualized speech impairments in children. We no longer viewed them as exclusively "articulatory" in nature. We began to see them as an integrated aspect of the developing linguistic system potentially requiring intervention at more than one level before full competence could be achieved.

The questions we pose in this tutorial, therefore, are these: Are there similar lessons to be learned from the new perspectives on speech acquisition that are emerging from contemporary psycholinguistics? Should the theories we explore here cause us to rethink speech impairment in children yet again? We will return to these questions at the end of our discussion of psycholinguistic models.

Box-and-Arrow Models Inputs, Outputs, and Underlying Representations

A number of key terms are used to describe the processes involved in the perception, storage, and production of speech. The input signal is the speech signal heard by the child, usually assumed to come from an adult speaker. The *output signal* is the utterance produced by the child. The unseen psychological events that occur between the arrival of an input signal and the production of speech are the focus of psycholinguistic models. Events that process the input signal are referred to as input processes, whereas events that process the production of speech are referred to as output processes. Some aspects of speech processing are thought to happen online-that is, they occur during the actual perception or production of speech and thus require a share of the attentional resources dedicated to the speech task. Other processes, thought to happen offline, take place as part of the child's background mental processing rather than during the time dedicated to the speech task. In this sense, online processing is sometimes defined as occurring in *real-time*, whereas offline processing is said to be *time-free* (Hewlett, 1990). In *box-and-arrow* psycholinguistic models, each hypothesized level of representation or processing can be represented in a diagram by a "box," and the relationships between them by "arrows," hence the name. Sometimes (as in the models of Smith, 1973, and Menn, 1978, described later in this paper) the arrows represent processes additional to those shown in boxes. Such models make explicit the hypothesized information-processing activities carried out in a particular cognitive function (such as language), in a manner analogous to computer flowcharts that depict the processes and decisions carried out by a computer program.

Box-and-arrow models differ widely in the number of unseen psychological processes they describe and thus in the number of boxes they contain. Some have only one or two boxes between the input and output signals (e.g., Menn, 1978; Smith, 1973), whereas others have multiple boxes representing complex relationships between a number of different information-processing events (e.g., Hewlett, 1990; Hewlett, Gibbon, & Cohen-McKenzie, 1998; Stackhouse & Wells, 1997). The most important box, however, and the source of much ongoing debate, is that representing the underlying representation (or UR). In essence, an underlying representation captures information stored in a child's mind about a word he or she knows and uses. As the following description of several models will illustrate, the nature of this information and thus the type(s) of representation present in the child's knowledge base have captured the attention of researchers for some time.

Early Single-Lexicon Models

Smith (1973) was one of the first to address specifically the nature of children's underlying representations. He proposed that children had one lexicon to hold their underlying representations of speech (Figure 1). His extensive analysis of the longitudinal data collected from his son Amahl led him to conclude that children's underlying representations were adult-like and in fact equivalent to adult surface representations. Amahl's ability to discriminate between minimal pair words such as MOUSE and MOUTH before he had started to talk was cited as evidence in support of Smith's position.

Figure 1 shows Smith's (1973) single lexicon model (cited in Smith, 1978, p. 260). The first box shows that the input from adult speech was stored as the child's underlying representation. The underlying representation was thought to be perceptually based. Smith proposed that these stored representations were modified online through the action of phonological rules (also referred to as realization rules) to create surface representations. Smith's phonological rules were devised using **Figure 1.** Single-lexicon model proposed by Smith (1973). Note: From "Lexical Representation and the Acquisition of Phonology" (p. 260) by N. V. Smith, 1978. In B. B. Kachru (Ed.), *Linguistics in the Seventies: Directions and Prospects. Special issue of studies in the linguistic sciences*, *8*, 259–273. Copyright 1978 by N. Smith. Reprinted with permission.



distinctive features and other aspects of generative phonology. The systematic differences between the child's perception and production were attributed to the action of the phonological rules on the child's underlying representation. Finally, articulatory instructions were applied to create the pronounced word or output.

Development of single-lexicon models continued through the 1970s (e.g., Macken, 1980; Smith, 1978). Much revision occurred because of evidence that did not support Smith's (1973) assumption that a child's perception was always adult-like. For example in Amahl's production of the words PUDDLE [pAgəl], and CUDDLE [kAdəl] it was unclear why he produced the [g] in PUDDLE if he could perceive and subsequently produce the /d/ in CUDDLE. Smith (1978) suggested that such variation could be the result of a difficulty perceiving the difference between certain speech sounds. Consequently, Smith (1978) revised his earlier model by adding a perceptual filter to account for inaccurate perception.

Thus early single-lexicon models accounted for the $[f_{1S}]$ phenomenon in that children were presumed to hear a difference between $[f_{1J}]$ and $[f_{1S}]$, but then produce the words as homonyms because of the action of phonological rules. These models had difficulty, however, accounting

for variable pronunciations-especially instances where different tokens of a word would be pronounced in different ways by the same child and where one phoneme would be pronounced differently in different words (Bernhardt & Stemberger, 1998). For example, it was unclear how regressive idioms could arise-that is, how some words could be pronounced with a relatively immature pattern compared with phonologically similar words, such as the word FAT pronounced as [bæt] if the child can pronounce FOUR, FIRE, and FUN correctly. Similarly, it was unclear how to account for *progressive* idioms: words in which the pronunciation is more advanced than in similar words (Menn & Matthei, 1992). Although proponents of single-lexicon models suggested that variable words could be treated as lexical exceptions (i.e., words that do not follow across-the-board applications of rules), some researchers considered this explanation cumbersome and insufficient because of the potentially large number of such exceptions and the difficulty of deciding at what point a word becomes an exception to a rule (Bernhardt & Stemberger, 1998).

Early Two-Lexicon Models

The limitations of the early single-lexicon models provided two catalysts for the development of two-lexicon models. First, there was the need to better account for the variability in children's speech. Second, there was a need to account for the idea that normally developing children at early stages of development and children with speech impairment may have underlying representations unique to their own language system. This was contrary to Smith's (1973) idea that children have underlying representations similar to adult surface forms. (For a complete discussion see Maxwell, 1984.)

It was therefore proposed that children must have two lexicons for their underlying representations: an *input lexicon* for representations used in word recognition and an *output lexicon* for representations used in word production. Proposing separate input and output lexicons allowed children to acquire a store of underlying representations that were non-adult-like, thus accounting for the existence of lexical exceptions.

A number of researchers worked on the development of various two-lexicon models (Hewlett, 1990; Kiparsky & Menn, 1977; Menn, 1978, 1983; Spencer, 1986, 1988). Menn and Matthei (1992) provide an excellent historical account of the development of these models. A typical exemplar of the ideas embraced by the early twolexicon models is that of Menn (1978, p. 103), illustrated in Figure 2.

In this type of model, the child stores underlying adult-like perceptual representations in the input lexicon. This perceptual representation is then modified Figure 2. Two-lexicon model proposed by Menn (1983). Note: From "Development of Articulatory, Phonetic and Phonological Capabilities" (p. 103) by L. Menn, 1983. In B. Butterworth (Ed.), *Language Production* (Vol. 2). London: Academic Press. Copyright 1983 by Academic Press. Reprinted with permission.



offline through the application of phonological rules or processes to create another representation to be used for production, which is stored in an output lexicon. Once a child has stored a word in the output lexicon, subsequent productions are accessed from the output lexicon only rather than being accessed from the input lexicon and modified online.

Although two-lexicon models are able to account for the variability observed in children's speech because the child can have more than one representation of the same word in the output lexicon, such models nonetheless have a number of limitations (Bernhardt & Stemberger, 1998; Chiat, 1994; Menn & Matthei, 1992; Vihman, 1996). With the potential for duplication of lexical items in the output lexicon, the models fail to explain how children select one representation over another, how representations change to become more adult-like, and how old forms are deleted (Bernhardt & Stemberger, 1998; Dinnsen, Barlow, & Morrisette, 1997; Vihman, 1996).

Recent One- and Two-Lexicon Models

In 1990, Hewlett proposed a more detailed two-lexicon model of speech production by relating the underlying phonological processes described by the earlier models to the articulatory-phonetic production of speech. His model (Figure 3) sought to address some of the limitations of previous two-lexicon accounts by specifically considering how children select one representation over another, how output representations change to become more adult-like, and how offline rules can be suppressed online.

Hewlett (1990) proposed that a child produces a word via one of two possible speech processing routes. The child can access an auditory-perceptual featurebased representation from an input lexicon and send this information to a motor programmer, which then devises a motor plan for its production. Alternatively, he or she can access an articulatory-based representation from an output lexicon; a representation established offline via phonological rules that map the perceptual representation onto articulatory feature specifications. Producing words via the input lexicon route is thought to be more laborious because this involves online processing. By contrast, word production via the output lexicon route is thought to be more automatic because output lexical representations already contain the relevant production information required for implementation by highly learned combinations of muscle commands. Hewlett (1990) suggested that variable productions and improvements occur when four conditions apply:

- 1. The child becomes aware of the insufficiency of his or her current production.
- 2. The child desires to change it.
- 3. The child acquires knowledge of the relevant crucial articulatory targets.
- 4. The child has sufficient dexterity of the vocal apparatus to implement speech sounds at speed in a variety of phonetic contexts.

Feedback and interaction between the various processes or boxes within the model (e.g., input and output lexicon, motor programmer) is thought to facilitate change of the child's articulatory representations in their output lexicon to more adult-like representations.

Since the publication of Hewlett's (1990) model, a number of researchers have demonstrated how useful this model can be for exploring and understanding the problems underlying impaired speech development (Howard, 1993; Williams & Chiat, 1993). Williams and Chiat (1993), for example, explored potential levels of impairment in children with a suspected delay in their phonological development versus children with unusual or disordered phonology. They did this by examining the responses of the children to a series of speech production tasks including naming, sentence repetition, repetition of nonwords, and repetition of real words. The children with delayed phonology made significantly fewer errors than the children with disordered phonology and showed a more consistent error pattern across tasks. Further, the children with disordered phonology could be classified into two subgroups: one subgroup that made significantly fewer errors on the repetition task than the naming task and another subgroup that had equivalent error rates across tasks. Using Hewlett's (1990) model, Williams and Chiat suggested the subgroup **Figure 3.** Model of speech production proposed by Hewlett (1990). Note: From "Processes of Development and Production" (p. 230) by N. Hewlett, 1990. In P. Grunwell (Ed.), *Developmental Speech Disorders* (pp. 15–38). London: Churchill Livingstone. Copyright 1990 by Churchill Livingstone. Reprinted with permission.



with the consistent error rate had a problem with motor programming, whereas the subgroup with the differing accuracy across tasks may have had unstable underlying representations in their output lexicon. As Williams and Chiat (1993, p. 155) point out, such findings have implications for intervention. They suggest that some children need help to "break programing habits and attempt new programmes" whereas others may need help establishing consistent motor programs for individual words. Such suggestions are also in line with research conducted by Dodd and colleagues (Dodd, 1995; Dodd, Leahy, & Hambly, 1989; Dodd & McCormack, 1999).²

A more recent box-and-arrow model is the speech processing model proposed by Stackhouse and Wells (1997). Stackhouse and Wells postulated that there is a single underlying representation (which they called a *lexical representation*) that contains phonological, semantic, grammatical, orthographic, and motoric information. They link this representation to an extensive series of related processes beginning with audition through to motoric production, as shown in Figure 4.

The lexical representation is depicted in Figure 4 by

three bolded boxes containing phonological, semantic, and motor information. The grammatical and orthographic components of the lexical representation are not explicitly shown in the figure. The input processes include peripheral auditory processing, discrimination of speech versus nonspeech, recognition of phonological forms relevant to the ambient language, in addition to the phonetic discrimination of speech sounds. The output processes include motor programming, motor planning, and motor execution. The broad arrows and shaded boxes represent processes hypothesized to occur offline.

This focus on modelling so many processes involved in speech perception and production has proven clinically useful in the study of children's speech and literacy difficulties (Snowling & Stackhouse, 1996; Stackhouse, 1992, 1993, 1997; Stackhouse & Wells, 1993, 1997; Waters, Hawkes, & Burnett, 1998). An excellent example is the case study of a child name Zoe, who was 2;10

Figure 4. Speech processing model by Stackhouse and Wells (1997). Notes: The broad arrows and shaded boxes represent processes hypothesized to occur offline. From *Children's Speech and Literacy Difficulties: A Psycholinguistic Framework* (p. 350) by J. Stackhouse and B. Wells, 1997. London: Whurr. Copyright 1997 by Whurr Publishers. Reprinted with permission.



²More recently, Hewlett and colleagues (1998) have revisited the singlelexicon theory of speech production in order to account for differential word production accuracy across tasks in a child with a phonological impairment.

(years;months) when she first started receiving speech and language therapy (Stackhouse, 1997; Stackhouse & Wells, 1997). At that time, Zoe presented with symptoms consistent with development verbal dyspraxia, including groping oral movements before vocalization, inconsistent and deviant sound production, and very poor diadochokinetic rate. Stackhouse and Wells's psycholinguistic investigation of Zoe's skills continued from this early stage to when she was 9;8. At that time, she continued to present with more subtle speech problems in addition to significant literacy difficulties.

Over the years, Stackhouse and Wells conducted a series of assessments targeting various aspects of Zoe's input and output speech processing and emerging literacy skills. For example, they assessed auditory discrimination of real and nonword minimal pairs, detection and production of real and nonword rhymes, imitation of single sounds and sequences of sounds, repetition of real and nonwords, and confrontation naming. The assessment results were then used to formulate hypotheses about the loci of Zoe's speech and literacy difficulties relative to the model. For example, Stackhouse and Wells (1997) hypothesized that at age 5;11 some of Zoe's specific impairments included—

- voiced/voiceless difficulties in her speech (which were also reflected in her spelling) due to auditory processing difficulties and weak phonological representations for onsets involving the contrast;
- difficulty with the production of unfamiliar words and consonant clusters due to motor programming problems; and
- difficulty producing affricates due to impaired motor execution.

Information obtained from the detailed assessment was then used to tailor intervention specific to Zoe's areas of difficulty (Stackhouse, 1997). For example, some of the intervention tasks targeting the deficits outlined above included—

- auditory tasks focusing on the distinction between voiced and voiceless consonants;
- activities involving the segmentation of polysyllabic words into syllables designed to help Zoe acquire new words; and
- articulatory exercises to improve Zoe's production of postalveolar fricatives and affricates.

Stackhouse and Wells (1997) provide a more detailed account of this case and other case studies that show how model-based assessment of the component processes involved in speech processing can provide more comprehensive assessment data: data that allow hypothesis testing about the possible problems underlying individual clients' speech and literacy difficulties.

Connectionist Models

Connectionist models differ from box-and-arrow models in both conceptual and practical ways, but the most obvious difference is related to the actual method by which theoretical notions are expressed. Box-andarrow models of a cognitive ability express a theory about how a system works using verbal reasoning (Dijkstra & de Smedt, 1996). That is, using verbally expressed concepts, the cognitive task is analyzed into a series of information-processing steps that are represented in a diagram similar to a computer flowchart (as in Figures 1–4). The nature and function of each component or process in the system and the relationships between these are described in words.

By contrast, a connectionist model of a cognitive ability is computer-based. The model is essentially contained in a computer program that specifies the activity and layout of many simple processing units arranged in a network (i.e., each unit is connected to many other units) (Rumelhart, McClelland, & the PDP research group, 1986). Information processing emerges from the interactions between large numbers of these units. The running computer program simulates the performance of a cognitive task (such as word production) by calculating the outcome pattern that arises after activation initiated by an input pattern has spread throughout the network along its rich interconnections many times over. The adequacy of the model can be determined by comparing the outcomes from its calculations with human responses to the same task.

The simple processing units are often called *nodes* to capture their abstract nature (Murre & Goebel, 1996), but their operation is based on the function of neurons in the brain. The models are known as *connectionist* because of the interconnected network of nodes they contain or as *neural networks* because they are thought to function somewhat similarly to the way neurons function together (Plaut, 1995).

Activation, Connection Strength, and Network Architecture

Just as neurons sum electrochemical impulses received at their dendrites and transmit them as action potentials along their axons, so the nodes in a connectionist network receive activation from, and send activation to, other nodes to which they are connected. The activation level of a node represents the amount of "work" it is doing in transmitting information, ranging from a *resting level* (when the node is not involved in any processing) to the maximum possible level set by the programmer. The amount of activation that can spread around the network is modulated by the *strength* of connections (also known as *weights*) between nodes (Harley, 1995). Some nodes may transmit all their activation to their neighbors, whereas others may transmit a smaller proportion. Connections with negative weights inhibit the activation of nodes at the receiving end. Analogously to the terminology used for neuronal function, connections are considered *excitatory* when they increase the activation levels of other nodes or *inhibitory* when they decrease the activation of other nodes.

Nodes representing different types of information are usually arranged in *levels* or *layers*. (For example, there are separate semantic feature, lexical [i.e., word], and phonological segment nodes in the network of Dell and O'Seaghdha [1991], as illustrated in Figure 5.) Some networks have several layers; others as few as two. The number and layout of nodes and connections between them (the network's *architecture*) depend on the complexity of the task the network is to perform and the programmer's concept of how this can best be modelled.

The *input* to the network is the activation pattern provided to the first layer of the model, and the *output* is represented by the activation pattern at the final layer. For example, when the Dell and O'Seaghdha model illustrated in Figure 5 is used to simulate naming, the semantic feature layer is the input layer, because naming is assumed to begin with semantic processing (a couple of nodes have been shaded in this figure to suggest that they are activated). The phonological segment layer is the output layer because the model's task is to settle on a stable set of activated phonological segments representing the spoken form produced by a human speaker. In a complementary example, the reverse pattern of spreading activation within a network could hypothetically occur in the recognition of spoken words (although Dell and O'Seaghdha, 1991, did not simulate this task). If the task were spoken word recognition, the phonological segments corresponding to an auditory word form (e.g., /d/o/g/) might be activated first, and the network would need to stabilize its activation pattern on an appropriate set of semantic features (such as "it's an animal," "it barks," etc.) to be judged to have "recognized" the word.

The layout of connections between nodes reflects what is known about likely constraints on processing within the system (Rumelhart et al., 1986). Thus, nodes that are thought to directly influence one another are connected; nodes thought not to directly affect one another may not be connected, and inhibitory connections can be modelled between nodes that are assumed not to be active at the same time. For example, in the Dell and O'Seaghdha (1991) model in Figure 5, we can see that these modellers have assumed there is no direct activation of phonological segments by semantic features in picture naming because they have not provided any direct connections between semantic and phonological segment nodes. In this particular network the modellers have simulated various aspects of word production without using inhibitory connections in the design of the network.

Figure 5. An interactive activation model of lexical access in language production by Dell & O'Seaghdha (1991). Note: From "Mediated and Convergent Lexical Priming in Language Production: A Comment on Levelt et al." (p. 605) by G. S. Dell & P. G. O'Seaghdha, 1991. *Psychological Review, 98*, pp. 604-614. Copyright 1991 by the American Psychological Association. Reprinted with permission.



The "processing" carried out by the nodes is specified by mathematical formulas contained in the computer program. There are formulas that sum the activation that has been sent as input to each node from its connected nodes, others that transform this total to another amount in order to satisfy certain mathematical constraints, and still others that determine how much of a node's activation is to be passed on to other nodes (Murre & Goebel, 1996). The complexity of these calculations, given the large number of interconnecting nodes in some networks, may require considerable computing power. The connectionist modeller may also build a degree of inaccuracy or *noise* into the formulas governing activation transmission, which allows the model to operate with a small amount of unpredictability. He or she may also allow for activation levels to decrease or decay by a certain amount over time to prevent these levels from continuing to increase in the network in an unlimited way.

Networks may represent information in a localist or a distributed fashion. In a connectionist model with localist (also known as symbolic) representations, individual nodes represent individual concepts (Murre & Goebel, 1996). For example, in a model of part of the language system, nodes might represent discrete units of linguistic information such as words or phonemes. The network of Dell & O'Seaghdha (1991) that we have already discussed (Figure 5) contains localist representations where nodes represent semantic features, words, and phonological segments. By contrast, in a model with distributed representations, each single piece of linguistic information would be represented by a pattern of activation across a number of nodes and their interconnections, with none of the individual nodes corresponding to recognizable, discrete linguistic units. The large class of parallel-distributed processing (usually abbreviated as PDP) models are of this type. The word paral*lel* here refers to the fact that the many simultaneously activated nodes and connections between them represent many pieces of information being processed at the same time. (This contrasts, for example, with many boxand-arrow models in which the processes represented by each box are carried out serially.) PDP models are fully explicated in the now-classic, three-volume text by Rumelhart, McClelland, and the PDP research group (1986), with the first four chapters of Volume 1 providing a valuable and accessible introduction to the approach. An example of a PDP simulation of the way infants learn to understand and produce speech is reported by Plaut and Kello (1999).

Networks also differ in the direction(s) in which activation can be passed between nodes or layers. Nodes in the first layer are activated by the programmer to simulate input from either the environment or from an earlier stage of cognitive processing not captured by the model. In *feedforward* models, these nodes pass their activation to nodes in the second layer, which pass activation to nodes in the third layer (if there is one), and so on; and activation may only be transmitted in one direction through the network. In interactive activation models, activation may also feed back from later-activated layers to earlier-activated ones; thus nodes that have been activated at the second layer can send activation back to the first layer as well as passing on activation to the third layer. Activation levels for the nodes are updated (i.e., their new values are calculated according to the formulas programmed into the model) in successive cycles or *time steps*. The first time step thus represents the first recalculation of activation values; the second time step, the second recalculation (based on the values calculated in time step one); and so on. In some models the activations of all nodes are updated simultaneously; in others, activation values are updated node by node (Murre & Goebel, 1996). After a number of time steps the activation levels of all the nodes in the network no longer change much at each update, and the network is said to have reached a stable activation pattern. At this point the nodes that are activated at the output layer are taken as the network's "response" to the input.

Learning

One reason that connectionist networks have attracted so much interest for psycholinguistic researchers is because they can not only simulate the outputs from language systems under normal circumstances, they can also simulate the changes within a system as it *learns*. Learning is simulated in a network by including a further set of formulas (or *learning algorithms*), which alter the strength of connections between nodes at each time step. This learning may occur in two broad ways: unsupervised or supervised (Murre & Goebel, 1996).

In unsupervised learning, following a principle first described by Hebb (1949), the network is programmed to increase the strength of connections between input patterns that are similar to one another and to decrease the strength of connections between patterns that are dissimilar from one another. In other words, the weights change to encode correlations between similar patterns. As a result, responses to new inputs are determined by how similar these inputs are to previous inputs (Quinlan, 1991). This is significant because it allows the network to generalize as a result of its "learning." In other words, it can produce an output even when given an input it has never seen before. This is a standard feature of human learning that strictly rule-based accounts of cognitive processing find very difficult to achieve. Similarly, the network is able to offer a response when an input is degraded or incomplete, and if the degraded input is similar enough to the intended input the correct output will still be produced. Although the network "learns" to associate similar inputs with similar outputs, there is no direct outside influence on the network by which it is "taught." It is in this sense that the learning is unsupervised. Nakisa and Plunkett (1998) use unsupervised learning in a connectionist simulation of how infants might rapidly learn to discriminate speech sounds.

In supervised learning, the network is given a target or teaching output pattern as well as an input pattern. The network is programmed with formulas that compute the difference between the target output and the output calculated by the network, and other formulas that adjust the weights to bring the next calculated output closer to the target. Back-propagation is the most commonly used form of supervised learning (Harley, 1995). It was independently developed by at least four groups of researchers: Werbos in 1974; Parker in 1982; LeCun in 1986; and Rumelhart, Hinton, and Williams in 1986-all cited by Murre and Goebel (1996). Backpropagation is so-named because once the degree of error between the target and actual output has been calculated, the error measure is propagated (or fed) backwards layer-by-layer through the network and the connection strengths are adjusted, beginning at the output layer. In this procedure, the network's "learning" is said to be "supervised" because there is a target pattern available for comparison with the network's output.

Learning in connectionist models can therefore accomplish two purposes: It may improve the network's efficiency in transmitting patterns that occur frequently at input, or it may increase the network's success in producing particular output patterns. This is similar to the aims a child has in learning to speak. Children need to process most effectively the auditory-verbal information they experience most often in their environments, and they need to fine-tune their utterances towards those that best accomplish their goals. Clearly the ability of connectionist networks to learn has major implications for the endeavor to understand children's speech and language development.

Overlapping Representations

Another significant feature of connectionist models that allows them to offer an alternative explanation of some phenomena in children's speech is that representations for different items overlap; that is, they involve some of the same nodes (Stemberger, 1992). Thus, in the model in Figure 5 (Dell & O'Seaghdha, 1991), the activation pattern generated across the whole network for the word CAT would include some of the semanticlayer nodes that are also activated for the word DOG (e.g., nodes corresponding to semantic features such as "domestic pet," and "mammal"). The representation for CAT would also include the phonological nodes [x] and [t], which would be activated for the words RAT and MAT.

In the model shown in Figure 5, whenever a node is activated, it sends on activation to all the other nodes to which it is connected at the next time step. This means that if the semantic features corresponding to the concept of a cat are activated as the input to the model, they will send activation to the lexical-layer node for CAT, but also to all other lexical nodes that are connected to the semantic features for a cat. (For example, the lexical nodes for RAT and DOG will also have connections to the semantic feature "animal." whereas the lexical nodes LOG and MAT will not.) The lexical nodes for RAT and DOG will not be as strongly activated as the node for CAT because only some of their semantic features will be sending them activation. Any activated lexical nodes then send activation to all the phonological nodes to which they are connected; thus, all the phonological nodes for the lexical items CAT, RAT, and DOG will receive some activation. Again, [r, æ, t] and [d, o, g] should receive less than [k, æ, t]. Activation within this particular model spreads interactively (i.e., in both directions), and the final output at the phonological layer is not determined until activation has spread in both directions through the network over many time steps. If the network is operating as it should, the phonological segments [k, æ, t] will receive most activation when the input is the semantic specification for the word CAT.

The consequence of overlapping representations is that when items have something in common (such as shared semantic features or shared phonology), the nodes representing the information that is shared will receive more activation (Stemberger, 1992). Thus, in the model we have been discussing (Dell & O'Seaghdha, 1991), if the nodes for CAT, RAT, and MAT were all activated at the lexical layer, the phonological segments [æ, t] would receive more activation than if they were receiving activation from just one lexical node. Stemberger (1992) called this a gang effect; that is, CAT, RAT, and MAT would be members of a sound-based gang because they share the segments [*æ*, t]. Some items would receive extra activation by virtue of their overlap with a gang; others would have to compete against items that were benefiting from a gang effect. For example, if the particular network above contained a lexical node MOUSE and the phonological nodes [m, au, s], it would be harder to activate the phonological form [maus] than [ræt] because [ræt] would benefit from the gang effect of the lexical nodes RAT, CAT, and MAT. If the network also contained HOUSE and LOUSE, however, MOUSE would not be at such a disadvantage. It is easy to see that in a network with an adult-sized vocabulary there would be extremely complex interactions between sets of words sharing properties at all three layers.

In terms of explaining phenomena in children's speech, gang effects mean that the network naturally represents similarities between items without needing explicit statements or rules about those similarities (Stemberger, 1992). The simplest single lexicon models (e.g., Smith, 1973) assumed that children had adult-like underlying representations that were then acted upon by phonological rules resulting in the child's pronunciation. By contrast, in a connectionist model the differences between the input and output forms emerge from the way activation spreads within the network as a whole. There are no rules, just systematic regularities in the way the network functions when given a particular input (Menn & Matthei, 1992).

For example, Menn and Matthei (1992) suggested that the speech production system develops from connections between a child's motor, auditory, and kinesthetic modalities that are first established during babbling. Menn and Matthei hypothesize that the child learns relationships (called MAK patterns from Motor, Auditory, Kinesthetic) between a motor command, the sound it produces, and kinesthetic feedback about the position of the articulators involved. The child is thought to learn some regularities on his or her own during babbling-for example, that a particular MAK pattern will reliably result in [ba, ba, ba] over many repetitions. Connections within the child's MAK network may be strengthened to increase the reliability of this outcome. Adult sound patterns are also thought to become connected to the MAK patterns within this theoretical network. As a result, new regularities become established in response to adult language input. For example, some of a child's utterances will receive predictable responses from the adult as the adult interprets or imitates the utterance, and the connections in the MAK patterns underlying such utterances are likely to be strengthened. The slowness with which newly learned sounds are incorporated into existing words is attributed to the time required to change the connection strengths already underlying the production of existing words.

Probabilistic Outcomes

The output is always the pattern of nodes with the most activation when the network has settled into a stable activation pattern. The examples above, however, illustrate that because of the rich interconnections between many types of information within a network, *many* nodes in the network become activated—not just the ones most associated with the desired output. In any simulation, therefore, there could easily be a number of possible outputs with only small differences in activation between them, especially if there is an effect of noise (the unpredictability of activation spread) that decreases the activation of the desired output relative

to the activation of competing outputs. This is another property of connectionist networks, one that allows for important ways of developing theories about child speech and language that are quite different from those possible using box-and-arrow models.

First, the competition among possible outputs in connectionist models can explain why the output can be different on different occasions—an answer to the variability problem that models with rule-based processes find hard to explain. In each production, the child is not seen as adapting the adult pronunciation to a form that he or she can produce (using a rule or process), but rather is thought to attempt the adult form consistently within the constraints of the connection patterns available within the developing network (Stemberger, 1992). Children's utterances that appear to result from a combination of rules rather than the operation of a single rule are attributed to more than one established pattern within the network influencing the output at one time (Menn & Matthei, 1992).

Second, connectionist models can not only explain why variability occurs, they can also be specific about the probability with which specific outputs are likely to happen. For example, if a simulation is run 100 times on the computer given the same input, then the number of times each different output occurred would represent the probability of that output as a percentage. This could form the basis of a prediction about how likely a child would be to produce each of a number of different utterances in response to a particular spoken stimulus. This makes connectionist models a powerful tool for testing hypotheses about speech and language development. The modeller's task is to set up the network to produce the same language outputs that a child would produce. If the output from the model does not match what children actually do, this is a clear indication that the hypothesis behind the present form of the model is incorrect and needs modification. The complexity of connectionist models thus allows them to make predictions about different phenomena from those possible using box-and-arrow models and, further, to make quantita*tive* predictions.

Combining Box-and-Arrow and Connectionist Approaches

A model of phonological development proposed by Menn et al. (1993) incorporates two connectionist networks within a larger framework that exemplifies a boxand-arrow approach to model-building. The connectionist networks are not operative, but the description provided by Menn and colleagues illustrates the potential of connectionist modelling in researching children's speech and language development. Further, the model (Menn et al., 1993) demonstrates how the strengths of box-and-arrow and connectionist approaches may be used together to optimal effect. The model is called GEYKO after the way one child, Jacob, pronounced his name.

The GEYKO model, shown in Figure 6, contains a number of different cognitive processing components. The model assumes that a child receives input from a parent and from objects in the environment when learning to speak. Auditory Perception generates phonetic sequences from acoustic input; and Visuospatial Processing, Prehension (grasping), and Orientation also allow the child to interact with his or her environment. Discrimination Memory compares a current phonetic sequence with a previous one; this would allow a child to evaluate the success of his own attempts at imitation. Lexical Memory associates phonetic sequences with semantic information, and the Phonetic Buffer provides a short-term store for phonetic sequences before the computation of speech gestures (Speech Gesture Planning). The Low Level Articulator Controls calculate the specific articulatory trajectories required for each speech gesture based on the current state and position of the required articulators, and the movement of the Speech Articulators is intended to be simulated using an articulatory synthesizer. Goal Selection determines which is the immediate task (including babbling or imitation) and whether and in what way a particular outcome will be reinforced (thus allowing one type of learning). Feedback loops exist between speech output and auditory perception and between articulatory-level activity and speech gesture planning.

Although the GEYKO model is depicted in the same sort of diagram as those used to represent processing components and their interrelationships in box-and-arrow models, its Auditory Perception component and Speech Gesture Planning component are intended to be implemented as running connectionist networks (Menn et al., 1993). The Auditory Perception network is intended to derive phonetic features from acoustic information. The network's task is to classify sounds according to formant values, formant transitions, and whether they contain periodic or aperiodic energy, using unsupervised learning. This first network would produce an output corresponding to information about place and manner of articulation and voicing. The task of the second network, Speech Gesture Planning, is to learn speech gestures that would reproduce those phonetic representations

Figure 6. GEYKO cognitive architecture. Note: From "Connectionist Modelling and the Microstructure of Phonological Development: A Progress Report." (p. 427) by L. Menn, K. Markey, M. Mozer, and C. Lewis, 1993. In B. de Boysson-Bardies, S. de Schonen, P. Jusczyk, P. McNeilage, & J. Morton (Eds.), *Developmental Neurocognition: Speech and Face Processing in the First Year of Life* (pp. 421–433). Dordrecht: Kluwer Academic. Copyright 1993 by Kluwer Academic Publishers. Reprinted with kind permission.



*Goals are linked with other modules as well

at production. Given the phonetic features as input, the network would thus produce as output a set of target positions and movement parameters for the articulators.

In other words, the model would be given acoustic information from adult speech as input and would compute a set of articulatory commands as output—exactly the task that children accomplish when they acquire adult-like speech. This type of model therefore has two advantages. First, it brings a box-and-arrow approach to the task of hypothesizing about which cognitive subsystems might be involved in a child's speech and language development. Second, it allows for the potential specificity of connectionist networks in defining and testing hypotheses about how the child's learning of perceptual representations and articulatory commands might actually be carried out.

Advantages and Limitations of Connectionist Models Compared With Box-and-Arrow Models

Implementing a connectionist model on a computer provides a precise, formal way of expressing a theory. This has some advantages and some limitations. The level of detail represented in the program of a connectionist network both requires and makes possible a more exacting comparison between the model and real-world data than box-and-arrow models frequently allow. This may, therefore, make it very clear when a connectionist modeller needs to revise a theory. By contrast, box-andarrow theories are harder to reject and harder to correct, because the verbal concepts are less precise and theoretical or empirical inconsistencies can be more difficult to spot (Chiat, 1994; Dijkstra & de Smedt, 1996; Stemberger, 1992). The disadvantage of the precision offered by connectionist simulation is the amount of time that may be consumed in developing appropriate computational values from empirical data (Murre, 1994). Moreover, revising a model is not simply a matter of adding or deleting a box or some arrows. The entire network may need rethinking when its output fails to conform to real world observations.

A major strength of connectionist models is that the quantitative predictions they allow are very attractive for building theories. The opportunity to predict the probabilities of all likely outputs, rather than to account only for the correct output, means that the variability of a child's utterances can actually be used as data to test a theory in connectionist models (Menn & Matthei, 1992). In box-and-arrow models, such data cannot be used in the same way because there is no opportunity to predict how likely any particular output is. Model-based calculation of the probabilities of various outcomes in children's speech development relies, however, on statistical regularities in the language data to which the network has been exposed. Early in development when a child produces a small number of utterances, calculating such regularities may not be overly difficult because the model can be programmed to simulate production of the child's entire vocabulary. As a child's language system develops, however, the number of possible interactions within the system (between words and between different types of information about words) becomes enormous, and the calculations may require considerable computational power. Some connectionist networks (e.g., Dell & O'Seaghdha, 1991) contain vocabularies of very few words, whereas others have vocabularies of around 400 words (e.g., Plaut & Kello, 1999). Exponential increases in computer technology over the past couple of decades suggest that increasingly large connectionist models of the speech and language system will become increasingly available.

The complexity of the interactions that occur in connectionist networks means it may be difficult to make predictions from a description of a model without running a simulation to see what the actual outputs are. In many contexts this may be less convenient than using a box-and-arrow model, which can be fully explicated with pen and paper in order to generate predictions, without the computer-implemented stage. In the near future, connectionist modelling is unlikely to be available to most speech-language pathologists in clinical settings as a method of mapping out and testing hypotheses about a particular client's speech and language skills, whereas box-and-arrow models are ideally suited to this clinical endeavor.

The Future: Clinical Application

As our review of current psycholinguistic models has highlighted, researchers have already begun considering the potential of box-and-arrow models in the clinical setting (e.g., Hewlett, 1990; Stackhouse & Wells, 1997). A growing body of literature is emerging on the application of these models to speech and literacy difficulties (e.g., Bridgeman & Snowling, 1988; Bryan & Howard, 1992; Chiat, 1989; Ebbels, 2000; Howard, 1993; Popple & Wellington, 1996; Snowling & Stackhouse, 1996; Stackhouse, 1993, 1997; Stackhouse & Wells, 1993; Waters et al., 1998; Williams & Chiat, 1993). Boxand-arrow approaches to therapy planning have been used in adult settings for some years, primarily to direct assessment techniques and as a basis for determining what the focus of intervention should be (Coltheart, Bates, & Castles, 1994). Additionally they provide opportunities to establish the specificity of treatment effects (Seron, 1997). Work has also begun on the application of connectionist modelling to predict recovery patterns in acquired aphasia (Dell, Schwartz, Martin, Saffran, & Gagnon, 1997; Martin, Saffran, & Dell, 1996; Plaut, 1996, 1999). This endeavor has obvious implications for the design of optimal therapy goals; however, Harley (1996), in a persuasive review, concludes that this work is still in the early stages.

It is our experience that speech-language pathologists often operate in an eclectic way in the assessment, analysis, and management of speech impairment in children, drawing primarily on linguistic-based approaches. We suggest that systematic and detailed assessments of the *cognitive processes* described in psycholinguistic models may add to the speech-language pathologist's repertoire, allowing for the testing of different hypotheses about children's speech perception and production abilities. The findings from such assessments may then be used to tailor intervention to the identified problem areas (Bryan & Howard, 1992). The challenge ahead for the speech-language pathologist will be to consider the growing body of literature on psycholinguistic models and to relate it to clinical practice.

We believe, therefore, that there are two ways that psycholinguistic theory building can potentially influence the clinical management of speech impairment in children: by directly supplementing the clinician's repertoire of approaches to assessment and therapy planning and by more generally offering a new way of conceptualizing speech impairments. We discuss these below.

Model-Based Investigation of Individual Client Impairments

First, models of the type discussed earlier in this paper can be used to develop and test theories about the processes that underlie speech production in individual clients, as the case study of Zoe's speech and literacy difficulties admirably illustrates (Stackhouse & Wells, 1997). In clinical practice, the speech-language pathologist could test his or her hypothesis about the source of a particular behavior by comparing the client's performance with the performance that would be predicted by a model if it were impaired in the manner hypothesized. Identifying the impaired process could then point the way towards an intervention.

To make use of this approach would require three things:

- 1. Selection of a specific box-and-arrow or connectionist model with the potential to provide insight into which processes of speech development might be impaired in a particular child.
- 2. Detailed assessment of the processes described by the model: the processes directly involved in speech

behavior and the cognitive processes that interact with and underpin speech.

3. Interventions directed at the identified impairments—which in turn test the initial hypothesis about the nature of the impairments. If the presenting symptoms can be shown to improve as a result of the model-based intervention, this will provide support for the clinician's hypothesis about the nature of the problem and for the adequacy of the model itself. Cases where the symptoms do not improve may suggest that the initial hypothesis was incorrect, or that the intervention did not effectively target the required process as intended, or that the model itself requires further development.

Although there are potential benefits in applying psycholinguistic models in everyday clinical settings as we have outlined, there are also some cautions to bear in mind. First, assessment methods with established validity and reliability for children of varying ages are not currently available to evaluate all the component psychological processes hypothesized by these types of models. Instead, clinicians will frequently need to devise such assessments themselves. This may introduce a degree of uncertainty about whether one clinician is interpreting aspects of the chosen model in the same way that others are. Simultaneously, however, it enables assessment items to be tailored directly to the client's need. Second, a long history of research in the fields of reading and language on the effects of the remediation of hypothesized psycholinguistic processes has to date demonstrated limited efficacy (Bortner, 1971; Lahey, 1988; Paul, 2000). This suggests that clinicians and researchers need to work together to develop assessment tools and intervention approaches that would successfully target putatively impaired psychological processes. Recent case study data, including some cases discussed earlier, indicate that the future may hold promise for clinical applications of box-and-arrow models of speech development. Third, the logistics of implementing services within a psycholinguistic framework are formidable. As noted above, evaluation of a child's speech impairment based on a box-and-arrow model requires extensive assessment individual to each client. Designing a connectionist model based on a child's speech data is even more time-intensive and would require a level of collaboration with computational modellers that is rarely available in clinical settings.

New Conceptualizations of Speech Impairments in Children

Although psycholinguistic approaches may not be readily applicable to the routine management of all clients, there may be a second, and broader, clinical purpose that they can serve. Shriberg (1993) has discussed the need for understanding speech impairment in children not only in order to identify more effective treatments but to provide levels of explanation that address etiology and pathogenesis. This implies that understanding psycholinguistic processes in speech may be important not only with respect to the management of individual cases. Modifying behavior, after all, does not necessarily require a fully developed theory of the behavior's cause (Mowrer, 1954; Osgood, 1963). Another clinically important implication lies in the potential of psycholinguistic approaches to provide clinicians with new ways of conceptualizing the causes and correlates of speech impairment in children.

As we have discussed, the "phonological process revolution" had profound effects on the way we think about speech impairment in children. It did lead to new treatment approaches, but perhaps more importantly it led speech-language pathologists to see "articulation problems" in a new way. The phonological process orientation influenced clinicians to view speech errors as a potential aspect of impaired language development, to think about how speech and language impairments might interact (e.g., Paul & Shriberg, 1982), and to see the speech-impaired child as an active learner, involved in the process of generating and applying rule-governed strategies to the task of learning to talk. In this way, a change in the models we use to describe a clinical phenomenon can have a profound influence on how we conceptualize not only an impairment but the entire framework in which the impairment occurs. This, in turn, can be useful for developing more integrated models of language acquisition—models that go beyond examining such processes as articulatory and phonological development in isolation and help us, instead, to see them as coordinated strands of development that ultimately weave the tapestry of communicative competence.

For example, if we believe that children develop different versions of lexical representations for recognition and for production in order to streamline and minimize resource allocation in production, then we might think about analogous resource-reduction strategies that could take place in other areas of cognition or communication. We might consider, in this case, how resources might be balanced across the system. Clinically, this might lead us to think about ways of controlling demand for overall communication resources during activities in which we attempt to change phonological production. For example, we might want to control the overall complexity of speech tasks by practicing newly emerging sounds in contexts that involve talk about familiar topics in short sentences. If, on the other hand, we hypothesize that speech sound production improves through a process of attempts to produce a correct perceptual target, a process shaped in a probabilistic manner by experience and feedback, then

our focus would not be on controlling complexity but on increasing the frequency and saliency of the feedback provided. This might, as one example, lead us to choose an operant approach to intervention, at least at first, so that correct production can encounter frequent reinforcement.

These suggestions are only preliminary. Their point is merely to highlight the fact that the way(s) in which we conceptualize developmental processes and impairments can affect clinical practice not only by suggesting specific intervention approaches, but also by reframing the ways in which we understand children's communication problems. This reframing can lead us to attempt to integrate broader perspectives into our thinking about speech and language impairments. It is our hope that in this tutorial we have encouraged some of our readers to begin this process of reframing and rethinking.

Acknowledgments

We would like to acknowledge Karla McGregor, Ron Gillam, and two anonymous reviewers for helpful comments on an earlier version of this paper.

References

- Berko, J., & Brown, R. (1960). Psycholinguistic research methods. In P. H. Mussen. (Ed.), *Handbook of research methods in child development* (pp. 517–557). New York: John Wiley & Sons.
- Bernhardt, B., & Stoel-Gammon, C. (1994). Nonlinear phonology: Introduction and clinical application. *Journal* of Speech and Hearing Research, 37, 123–143.
- Bernhardt, B. H., & Stemberger, J. (1998). Handbook of phonological development from the perspective of constraint-based nonlinear phonology. San Diego, CA: Academic Press.
- Bernhardt, B. H., & Stemberger, J. P. (2000). Workbook in nonlinear phonology for clinical application. Austin, TX: Pro-Ed.
- **Bortner, M.** (1971). Phrenology, localization, and learning disabilities. *Journal of Special Education, 5,* 23–29.
- Bridgeman, E., & Snowling, M. (1988). The perception of phoneme sequence: A comparison of dyspraxia and normal children. *British Journal of Disorders of Communication*, 23, 245–252.
- Bryan, A., & Howard, D. (1992). Frozen phonology thawed: The analysis and remediation of a developmental disorder of lexical phonology. *European Journal of Disorders of Communication, 27,* 343–365.
- **Chiat, S.** (1989). The relation between prosodic structure, syllabification and segmental realization: evidence from a child with fricative stopping. *Clinical Linguistics and Phonetics, 3,* 223–242.
- **Chiat, S.** (1994). From lexical access to lexical output: What is the problem for children with impaired phonology? In M. Yavas (Ed.), *First and second language pathology* (pp. 107–133). San Diego, CA: Singular.

Chomsky, N., & Halle, M. (1968). Sound pattern of English. Cambridge, MA: MIT Press.

Coltheart, M., Bates, A., & Castles, A. (1994). *Cognitive neuropsychology and cognitive rehabilitation*. Hove, U.K.: Lawrence Erlbaum.

Dell, G. S., & O'Seaghdha, P. G. (1991). Mediated and convergent lexical priming in language production: A comment on Levelt et al. (1991). *Psychological Review, 98*, 604–614.

Dell, G. S., Schwartz, M. F., Martin, N., Saffran, E. M., & Gagnon, D. A. (1997). Lexical access in aphasic and non-aphasic speakers. *Psychological Review*, 104, 801–838.

Dijkstra, T., & de Smedt, K. (1996). Computer models in psycholinguistics: An introduction. In T. Dijkstra & K. de Smedt (Eds.), *Computational psycholinguistics: AI and connectionist models of human language processing* (pp. 3–23). London: Taylor and Francis.

Dinnsen, D. A., Barlow, J. A., & Morrisette, M. L. (1997). Long-distance place assimilation with an interacting error pattern in phonological acquisition. *Clinical Linguistics* and Phonetics, 11, 319–338.

Dodd, B. (1995). Children with speech disorder: Defining the problem. In B. Dodd (Ed.), *Differential diagnosis and treatment of children with speech disorder* (pp. 1–19). San Diego, CA: Singular.

Dodd, B., Leahy, J., & Hambly, G. (1989). Phonological disorders in children: Underlying cognitive deficits. *British Journal of Developmental Psychology*, 77, 55–71.

Dodd, B., & McCormack, P. (1999, July). Consistency of word production and speech disorder. Paper presented at Child Phonology Conference, Bangor, Wales.

Ebbels, S. (2000). Psycholinguistic profiling of a hearingimpaired child. *Child Language Teaching and Therapy*, *16*, 3–22.

Fee, E. J. (1995). Segments and syllables in early language acquisition. In J. Archibald (Ed.), *Phonological acquisition and phonological theory* (pp. 43–61). Hillsdale, NJ: Lawrence Erlbaum.

Grunwell, P. (1987). *Clinical phonology* (2nd ed.). Baltimore, MD: Williams & Wilkins.

Harley, T. A. (1995). *The psychology of language: From data to theory*. Hove, U.K.: Erlbaum (UK) Taylor & Francis.

Harley, T. A. (1996). Connectionist modelling of the recovery of language functions following brain damage. *Brain and Language*, *52*, 7–24.

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

Hewlett, N. (1990). Processes of development and production. In P. Grunwell (Eds.), *Developmental speech disorders* (pp. 15–38). Edinburgh, U.K.: Churchill Livingstone.

Hewlett, N., Gibbon, F., & Cohen-McKenzie, W. (1998). When is a velar an alveolar? Evidence supporting a revised psycholinguistic model of speech production in children. *International Journal of Language and Communication Disorders*, 33, 161–176.

Hodson, B., & Paden, E. (1991). *Targeting intelligible speech: A phonological approach to remediation* (2nd ed.). Austin, TX: Pro-Ed.

Howard, S. (1993). Articulatory constraints on a phonological system: A case study of cleft palate speech. *Clinical Linguistics and Phonetics, 7,* 299–317.

Ingram, D. (1976). *Phonological disability in children.* London: Edward Arnold.

Kiparsky, P., & Menn, L. (1977). On the acquisition of phonology. In J. MacNamara (Ed.), *Language learning and thought* (pp. 47–78). New York: Academic Press.

Lahey, M. (1988). Language disorders and language development. New York: Macmillan.

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

Martin, N., Saffran, E. M., & Dell, G. S. (1996). Recovery in deep dysphasia: Evidence for a relation between auditory-verbal STM capacity and lexical errors in repetition. *Brain and Language, 52,* 83–113.

Maxwell, E. M. (1984). On determining underlying phonological representations of children: A critique of the current theories. In M. Elbert, D. A. Dinnsen, & G. Weismer (Eds.), *Phonological theory and the misarticulating child* (ASHA Monograph 22, pp. 18–29). Rockville, MD: American Speech-Language-Hearing Association.

McCormack, P. (1997, Autumn). New approaches to the assessment of children's speech. *Australian Communication Quarterly*, 3–5.

Menn, L. (1978). Phonological units in beginning speech. In A. Bell & J. B. Hooper (Eds.), *Syllables and segments*. Amsterdam: North-Holland.

Menn, L. (1983). Development of articulatory, phonetic and phonological capabilities. In B. Butterworth (Ed.), *Language production* (Vol 2, pp. 3–50). London: Academic Press.

Menn, L., Markey, K., Mozer, M., & Lewis, C. (1993).
Connectionist modelling and the microstructure of phonological development: A progress report. In B. de Boysson-Bardies, S. de Schonen, P. Jusczyk, P. McNeilage, & J. Morton (Eds.), *Developmental neurocognition: Speech and face processing in the first year of life* (pp. 421–433).
Dordrecht: Kluwer Academic.

Menn, L., & Matthei, E. (1992). The "two lexicon" account of child phonology looking back, looking ahead. In C. A. Ferguson, L. Menn, & C. Stoel-Gammon (Eds.), *Phonological development: Models, research, implications* (pp. 211– 247). Timonium, MD: York Press.

Mowrer, O. (1954). The psychologist looks at language. *American Psychologist, 9*, 660–694.

Murre, J. M. J. (1994, March). *Possibilities and pitfalls of simulation.* Paper presented at the Medical Research Council Applied Psychology Unit Seminar for Postgraduates, Cambridge, U.K.

Murre, J. M. J., & Goebel, R. (1996). Connectionist modelling. In T. Dijkstra & K. de Smeldt (Eds.), *Computational psycholinguistics* (pp. 49–81). London: Taylor and Francis.

Nakisa, R. C., & Plunkett, K. (1998). Evolution of a rapidly learned representation for speech. *Language and Cognitive Processes, 13,* 105–127.

Osgood, C. (1963). On understanding and creating sentences. *American Psychologist, 18,* 735–751.

Paul, R. (2000). Language disorders from infancy through adolescence (2nd ed.). St. Louis: Mosby.

Paul, R., & Shriberg, L. D. (1982). Associations between phonology and syntax in speech-delayed children, *Journal* of Speech and Hearing Research, 25, 536–547.

Plaut, D. C. (1995). Double dissociation without modularity: Evidence from connectionist neuropsychology. *Journal of Clinical and Experimental Neuropsychology, 17,* 291–321.

Plaut, D. C. (1996). Relearning after damage in connectionist networks: Toward a theory of rehabilitation. *Brain and Language*, *52*, 25–82.

Plaut, D. C. (1999). Computational modeling of word reading, acquired dyslexia, and remediation. In R. M. Klein & P. A. Klein (Eds.), *Converging methods for understanding reading and dyslexia* (pp. 339–372). Cambridge, MA: MIT Press.

Plaut, D. C., & Kello, C. T. (1999). The emergence of phonology from the interplay of speech comprehension and production: A distributed connectionist approach. In B. MacWhinney (Ed.), *The emergence of language* (pp. 381– 415). Mahwah, NJ: Lawrence Erlbaum.

Popple, J., & Wellington, W. (1996). Collaborative working within a psycholinguistic framework.. *Child Language Teaching and Therapy, 12,* 60–70.

Quinlan, P. T. (1991). *Connectionism and psychology: A psychological perspective on new connectionist research.* Chicago: University of Chicago Press.

Rumelhart, D. E., McClelland, J. L., & the PDP research group (Eds.). (1986). Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations. Cambridge, MA: MIT Press.

Seron, X. (1997). Effectiveness and specificity in neuropsychological therapies: A cognitive point of view. *Aphasiology*, *11*, 105–123.

Shriberg, L. D. (1993). Four new speech and prosody-voice measures for genetics research and other studies in developmental phonological disorders. *Journal of Speech and Hearing Research, 36*, 105–140.

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

Smith, N. V. (1978). Lexical representation and the acquisition of phonology. In B. B. Kachru (Ed.), *Linguistics in the seventies: Directions and prospects. Special issue of studies in the linguistic sciences, 8,* 259–273.

Snowling, M., & Stackhouse, J. (Eds.). (1996). *Dyslexia,* speech and language: A practitioner's handbook. London: Whurr.

Spencer, A. (1986). Towards a theory of phonological development. *Lingua, 68,* 3–38.

Spencer, A. (1988). A phonological theory of phonological

development. In M. J. Ball (Eds.), *Theoretical linguistics and disordered language* (pp. 115–151). London: Croom Helm.

Stackhouse, J. (1992). Promoting reading and spelling skills through speech therapy. In P. Fletcher & D. Hall (Eds.), *Specific speech and language disorders in children* (pp. 194–203) London: Whurr.

Stackhouse, J. (1993). Phonological disorder and lexical development: Two case studies. *Child Language Teaching and Therapy, 9,* 230–241.

Stackhouse, J. (1997). Phonological awareness: Connecting speech and literacy problems. In B. W. Hodson & M. L. Edwards (Eds.), *Perspective in applied phonology* (pp. 157– 196). Gaithensburg, MD: Aspen Publishers.

Stackhouse, J., & Wells, B. (1993). Psycholinguistic assessment of developmental speech disorders. *European Journal of Disorders of Communication, 28*, 331–348.

Stackhouse, J., & Wells, B. (1997). Children's speech and literacy difficulties: A psycholinguistic framework. London: Whurr.

Stampe, D. (1979). *A dissertation on natural phonology.* New York: Garland.

Stemberger, J. P. (1992). A connectionist view of child phonology: Phonological processing without phonological processes. In C. A. Ferguson, L. Menn, & C. Stoel-Gammon (Eds.), *Phonological development: Models, research, implications* (pp. 165–189). Timonium, MD: York Press.

Stoel-Gammon, C., & Dunn, C. (1985). *Normal and disordered phonology in children*. Baltimore, MD: University Park Press.

Vihman, M. M. (1996). *Phonological development: The origins of language in the child*. Cambridge, MA: Blackwell.

Waters, D., Hawkes, C., & Burnett, E. (1998). Targeting speech processing strengths to facilitate pronunciation change. *International Journal of Language and Communication Disorders, 33*(Suppl.), 469–474.

Williams, N., & Chiat, S. (1993). Processing deficits in children with phonological disorder and delay: A comparison of responses to a series of output tasks. *Clinical Linguistics and Phonetics*, 7, 145–160.

Received July 10, 2000

Accepted February 28, 2001

DOI: 10.1044/1092-4388(2001/055)

Contact author: Elise Baker, PhD, School of Communication Sciences and Disorders, The University of Sydney, PO Box 170, Lidcombe, NSW 1825, Australia. Email: e.baker@cchs.usyd.edu.au