

Connectionist Models of Language Production: Lexical Access and Grammatical Encoding

GARY S. DELL

FRANKLIN CHANG

ZENZI M. GRIFFIN

University of Illinois at Urbana–Champaign

Theories of language production have long been expressed as connectionist models. We outline the issues and challenges that must be addressed by connectionist models of lexical access and grammatical encoding, and review three recent models. The models illustrate the value of an interactive activation approach to lexical access in production, the need for sequential output in both phonological and grammatical encoding, and the potential for accounting for structural effects on errors and structural priming from learning.

I. INTRODUCTION

Psycholinguistic research into language production—the process of translating thoughts into speech—has long been associated with connectionist models. Spreading activation models of lexical access in production represent some of the earliest applications of connectionist ideas to psycholinguistic data (e.g., Dell & Reich, 1977; Harley, 1984; MacKay, 1982; Stemberger, 1985). These models combined representations from linguistics with interactive activation principles and sought to explain speech errors, particularly errors resulting from multiple causes or processing levels. For example, “Lizst’s second Hungarian *restaurant*” instead of “*rhapsody*” involves mistakenly using a word that is associatively, syntactically, and phonologically related to the intended word. Activation that spreads interactively among processing levels seems to be a natural way to account for these kinds of slips.

Direct all correspondence to: Gary S. Dell, Beckman Institute, University of Illinois, 405 North Mathews, Urbana, IL 61801; E-mail: gdell@s.psych.uiuc.edu

Since the early speech-error models of lexical access, connectionist models of production have progressed on three fronts. First, the empirical basis of the models has been extended. They have now been applied to error data from aphasic patients, children, and older adults (e.g., Berg & Schade, 1992; Burke, MacKay, Worthley, & Wade, 1991; Dell, Schwartz, Martin, Saffran, & Gagnon, 1997; Stemberger, 1989) and to response time data from experimental paradigms (e.g., Cutting & Ferreira, 1999; Griffin & Bock, 1998; Roelofs, 1992, 1997; Schriefers, Meyer, & Levelt, 1990). Second, models have begun to address grammatical encoding, the selection and ordering of words in sentences. The third area of progress has concerned connectionist architectures. Whereas all of the early models were hand-wired networks with local representations, some recent production models make use of distributed representations acquired from learning algorithms. In addition, recent architectures allow for the production of true sequences (Eikmeyer & Schade, 1991; Jordan, 1986; Gupta, 1996; Hartley & Houghton, 1996; Houghton, 1990; MacKay, 1987).

In this article, we examine some connectionist models of production. Our aim is not to review the field, but rather to concentrate on our own recent efforts in two areas, lexical access and grammatical encoding. Lexical access and grammatical encoding are aspects of production that can be located in what has been called the formulation component (Levelt, 1989). This component takes a *message*, a nonverbal representation of the utterance, and turns it into linguistic form. Words are accessed and ordered (grammatical encoding) and their sounds are retrieved and organized for articulation (phonological encoding). Thus the formulation component is distinguished from a prior component responsible for message formation and a subsequent one that executes articulatory movements. Specifically, we will present three models. The first model deals with the access of single words and concentrates on explaining the errors of aphasic patients. The second focuses on the phonological encoding and links error phenomena to the characteristics of the vocabulary and the sequential nature of words. The third model addresses structural priming effects in grammatical encoding (e.g., Bock, 1986a).

In our discussion of language production, two issues will be in focus: serial order and linguistic structure. First, consider serial order. The creation of a temporal sequence is the essence of production. Yet, the canonical connectionist architectures, such as feedforward multi-layered networks, cannot create true temporal sequences. Rather these architectures generate a single output activation pattern for a particular input in parallel. Of course, one can finesse this problem by organizing output units into separate banks for each position in the sequence. But that is not how production proceeds. Sentences are, for the most part, constructed piecemeal from beginning to end. The words that are initially retrieved tend to be placed early in the sentence and these initial placements constrain subsequent lexical and structural decisions (Bock, 1982, 1986b, 1987a; Ferreira, 1996, 1997; Kempen & Hoenkamp, 1987; Levelt, 1989). This property of production, *incrementality*, demands a model with sequential output and where previous output interacts with the message to guide subsequent output. Even within a word, temporal sequence is important. Not only are the sounds of words articulated in sequence, but they also seem to be retrieved that way from the lexicon (Meyer, 1991; Wheeldon & Levelt, 1995). The sequential retrieval

of sounds is likely responsible for several phenomena such as the vulnerability of word-initial sounds to speech errors (Gupta & Dell, 1999) and consequently any model dealing with such phenomena must be sequential.

Furthermore, any production model must get the details of linguistic structure right. Especially here, there is reason to expect production to differ from comprehension. In comprehension, structural features, such as grammatical affixes or function words, are simply cues that are often indirectly related to meaning. In fact, as the comprehensibility of telegraphic speech shows, some structural cues are largely unnecessary for understanding. Consequently, relatively less emphasis on structural cues is required in the process of mapping from spoken or written language to meaning. Production models, in contrast, must make linguistic structure a priority. As pointed out by Garrett (1975) and Bock (1990), structural details such as subject-verb agreement affixes must be produced regardless of whether they code for key aspects of the message. Structural features are even preserved when speakers err. For example, one might say “I applied a pack” instead of “I packed an apple.” Notice that the error preserves the phrase structure of the sentence, keeping the function morphemes in place. Moreover, at the phonological level, the intended word “an” changes to “a” in agreement with the initial consonant in “pack” and the pronunciation of -ed changes from /t/ in “packed” to /d/ in “applied” (Fromkin, 1971). How to handle this sensitivity to structure is a major challenge for connectionist models, particularly for models that learn structure indirectly from the statistics of linguistic sequences.

We now turn to the models, starting with models of lexical access.

II. LEXICAL ACCESS

Lexical access is, relatively speaking, the easy part of production to model. It is simple pattern association: A pattern of activation corresponding to the meaning of a word needs to be mapped onto a pattern corresponding to the word’s sounds. Moreover, lexical access is not a generative process. Aside from the productive use of morphology, the words that one seeks are stored in the lexicon.

Despite this seeming simplicity, lexical access poses a number of challenges. First, word choice must be made in the context of other retrieved words and the speaker’s communicative goals. For example, one can refer to someone as a person, a woman, a mother, a female parent, “Sheila” or even “she.” As Roelofs (1996) and Levelt (1989) point out, characterizing lexical access as just a mapping from the semantic features of a concept to a lexical item ignores the fact that such features do not uniquely identify a word. Second, words that have similar meanings do not necessarily have similar sounds, for example, the words, “mother,” “father,” “man,” and “woman.” A semantic coding of these words in terms of FEMALE (1 or 0) and PARENT (1 or 0) and a phonological encoding of whether their initial sound is ‘m’ or not creates a mapping that is formally equivalent to exclusive-OR (see Dell et al., 1997). As a result, the mapping between meaning and sound is not linearly separable, and any network achieving this mapping would require a layer of nonlinear hidden units between semantics and sound. Third, as

we have already mentioned, the output of lexical access is a sequence of sounds. Consequently, the mapping from meaning to sounds is one from a static (meaning) to a dynamic (phonological) representation. Finally, the output is more than just a sequence of phonological units. Rather, the retrieved phonological units are related to one another by the syllabic and metrical organization of the word's form.

Because of the complexity of the mapping from meaning to sound, theories of lexical access often assume that this mapping occurs in two steps. In the first step, *lemma selection*, a concept is mapped onto a *lemma*, a nonphonological representation of a word. Often, the lemma is assumed to be associated with the grammatical properties of the word, its syntactic category and features such as gender or number. The second step, *phonological encoding*, transforms the lemma into an organized sequence of speech sounds. Probably the most intuitive evidence for these two steps comes from the tip-of-the-tongue (TOT) phenomenon. A speaker knows that a word exists but cannot access its sounds. A simple interpretation is that lemma selection has succeeded, but phonological encoding has not. Recent support for this claim has come from studies showing that speakers in the TOT state know the grammatical properties of the word being sought including, surprisingly, the word's grammatical gender (Miozzo & Caramazza, 1997; Vigliocco, Antonini, & Garrett, 1997; see Caramazza, 1997 for an alternative view).

There has been considerable debate regarding the relationship between lemma selection and phonological encoding. Some models assume that they are discrete, modular stages (Levelt, 1989; Roelofs, 1996, 1997). Lemma selection is completed before any phonological information is activated, and during phonological encoding, no semantic information is consulted. Support for the modular view has come from studies showing that the access of semantic information strictly precedes that of phonological information (van Turennout, Hagoort, & Brown, 1997; Schriefers et al., 1990). However, other studies have offered evidence that phonological encoding begins before lemma access is complete (Cutting & Ferreira, 1999; Peterson & Savoy, 1998). Furthermore, speech errors such as "Hungarian restaurant" for "Hungarian rhapsody," or "snake" for "snail" suggest the simultaneous activation of semantic or associative information and phonological information. The first connectionist model that we present, the *aphasia model*, preserves the distinction between lemma selection and phonological encoding, but denies that these are modular stages. Instead, it allows later levels to begin processing before earlier ones have finished (cascading) and for processing at later levels to influence that at earlier ones (feedback).

The Aphasia Model

The aphasia model (Dell et al., 1997) was developed to explain the error patterns of aphasic and nonaphasic speakers in picture naming experiments. 23 patients and 60 nonaphasic controls were given 175 pictures of simple objects and they had to pronounce the object's name (which was a noun for all the pictures). Errors were placed in five categories. For example, assuming that "cat" is the target, errors could be *semantic* ("dog"), *formal* ("mat" or "cap"), *mixed* ("rat," both formally and semantically related),

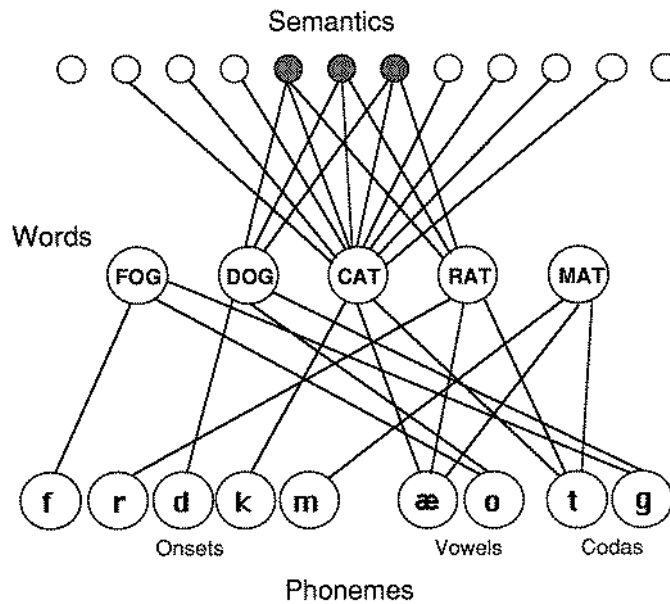


Figure 1. The aphasia model. Connections are excitatory and bidirectional.

unrelated (“pen” or “log”), or *nonwords* (“lat,” also including nonwords bearing no resemblance to the target such as “lom”).

Figure 1 shows the architecture of the aphasia model. There are three layers of units: semantic features, words, and phonemes. Each word corresponds to a single unit in the word layer. Bidirectional excitatory connections link words to their semantic features and phonemes. In the implementation used by Dell et al., each word connected to ten semantic features and three phonemes.

Lexical access is achieved by interactive spreading activation. Semantic units are activated, this activation spreads throughout the network, and ultimately the sounds of the intended word are retrieved. However, the model differs from classic interactive activation models (e.g., McClelland & Rumelhart, 1981) in several respects. Most importantly the aphasia model has two clear steps in the retrieval process, corresponding to lemma selection and phonological encoding. We briefly describe these two steps using “cat” as an example.

At the start of lemma selection, activation is added to the semantic features of the target word CAT. The activation spreads for a fixed number of time steps according to a noisy linear activation update rule. The bidirectional excitatory connections cause all three network levels to become active. In addition to the target word unit, CAT, semantic neighbors such as DOG become activated through shared semantic features. More interestingly, words such as MAT receive activation by feedback from phonemes shared with the target. When a mixed word such as RAT exists, it gains activation from both shared semantics and shared phonemes. Consequently a mixed word is usually more

activated than a purely semantic or formal neighbor of the target. A decision process concludes lemma selection. The most highly activated word of the appropriate grammatical category (here, noun) is chosen. However, the process is not perfect. Because of activation noise, there is some chance that a semantic, formal, or mixed neighbor of the target (or even an unrelated word) will be selected.

The second step, phonological encoding, begins with a large boost of activation to the chosen word unit. This boost introduces a nonlinearity into the model's activation process, which enables the network to handle the arbitrary mapping between semantic features and phonemes. After the boost, activation continues to spread for another fixed period of time. The most highly activated phonemes are then selected and linked to slots in a *phonological frame*, a structure that represents the number and kind of syllables in the word and its stress pattern. This linking concludes phonological encoding. Errors in phonological encoding occur when, due to noise, one or more wrong phonemes are more active than those of selected word. Typically such errors result in nonwords (e.g., "lat" for CAT) or less often, form-related words (e.g. "mat" or "sat" for CAT). In principle, the other error categories can also happen during phonological encoding, but the most common locus of these errors in the aphasia model is lemma selection.

In applying the model to aphasic naming errors, Dell et al. (1997) made two critical claims. First, they hypothesized that patient error patterns would fall between two extremes: the *normal pattern* produced by nonaphasic speakers, and a *random pattern* defined by the *error opportunities* associated with the error categories. The normal pattern was estimated from the 60 control speakers' data in the picture-naming task: Correct responses (97%), semantic errors (1%), and mixed errors (1%) nearly exhausted all of the relevant responses. The random pattern is the probability of each error type happening if a person knew no words, only the rules of sound combination in English, and "randomly" produced legal phonological strings in the picture-naming task. Dell et al. estimated these error opportunities for English from a variety of sources. Roughly speaking, the random pattern is mostly nonwords (80%) with the remaining responses being, in order of likelihood, unrelateds, formals, semantic errors, and mixed errors. In claiming that the set of possible patient error patterns falls between the normal pattern and the random pattern, the aphasia model instantiated the *continuity thesis*, an idea that goes back at least to Freud (1891/1953). Under this thesis, normal speech errors and aphasic paraphasias reflect the same processes.

The second basic claim of the aphasia model concerns its mechanism for creating error patterns between the normal and random patterns. The model's lexicon was set up so that its error opportunities matched the error opportunities estimated for English, and its other parameters (noise, size of the activation boost to the selected word, connection weight, decay, and time) were selected to give an error pattern that matched the normal pattern (See Table 1). To create aphasic error patterns, the model was lesioned by limiting its ability to transmit activation (reducing the connection weight parameter, p), its ability maintain its activation pattern (increasing the decay parameter, q), or both. The lesions create errors by reducing activation levels, which enhances the effect of noise. The greater the extent of the lesion, the more the model's error pattern approaches the random pattern.

TABLE 1
Picture Naming Error Proportions from Control Participants and Simulated
Proportions from the Aphasia Model (from Dell et al., 1997)

	Response Category					
	Correct	Semantic	Formal	Nonword	Mixed	Unrelated
Controls	.97	.01	.00	.00	.01	.00
Aphasia Model	.97	.02	.00	.00	.01	.00

Note. Connection weight (p) = .1; Decay (q) = .5.

However, the weight and decay components to a lesion promote different kinds of errors. A pure decay lesion is associated with more semantic, formal, and mixed errors (related word errors), while a pure weight lesion promotes nonword and unrelated word errors. For example, a weight lesion that reduces the model's correctness to 30%, creates 41% nonwords, 10% unrelated, 12% formals, 7% semantic, and 1% mixed. In contrast, a decay lesion leading to 30% correct has an error pattern of 26% nonwords, 7% unrelated, 20% formals, 13% semantic, and 3% mixed. Reducing weight makes the activation patterns on each level less consistent with one another, and leads to what Dell et al. call "stupid" errors. The production of a nonword reflects inconsistency between the word and phoneme layers, whereas an unrelated word reflects inconsistency between the semantic and the word level. When decay is increased without altering weight, errors tend to occur because noise dominates the decayed activation levels. But many of the errors reflect successful activation transmission among the levels because connection weights are still strong. These are "smart" errors (mixed, formal, and semantic errors) in which the word level is consistent with the semantic level or the phonological level is consistent with the word level.

The aphasia model gave a good account of patient error patterns. Dell et al. successfully fit the model to 21 of 23 fluent aphasic patients who were given the picture-naming test. Figure 2 illustrates the overall fit by plotting each predicted and obtained error proportion for all categories and patients. Table 2 shows the results for three patients, one fit with a pure decay lesion, I.G., one with a pure weight lesion, L.H., and one with a lower level of correctness, G.L.

Dell et al. then used the parameters assigned to the patients to make predictions. Here we mention two of these. First, if a patient's assigned connection weight was low they should not exhibit error phenomena that the model attributes to excitatory feedback between levels. The connection weights are just too low to support interaction between layers. According to the model, interactive feedback causes mixed errors (e.g., RAT for CAT) and formal errors that obey grammatical constraints (e.g., the noun MAT replacing the noun CAT). In support of the prediction, only the patients that the model assigned large weights to had significant tendencies to produce mixed errors and form-related nouns. The second prediction concerned recovery. Ten patients were retested on the naming test after one or more months. On average, the patients improved their performance by 16%. The model was able to fit these improved error patterns as well as the

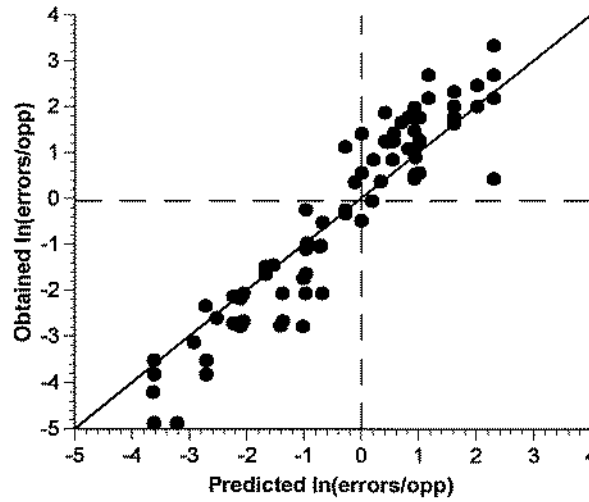


Figure 2. A comparison of error proportions from 21 patients and predicted proportion from the aphasia model. Proportions are transformed by the natural log of the ratio of the proportion and the error opportunities for each error category.

original ones. More importantly, recovery seemed to involve a movement of affected parameters toward normal values. Thus, recovery, or within-patient variation, takes place along the same dimensions as those that characterize between-patient variation.

The good points of the aphasia model arise from its interactive activation architecture. Interactive activation offers a natural mechanism for the error types and the model permits graceful degradation through parameter alterations. The fact that the model fits normal and patient error patterns provides support for both its general approach to lexical access and for the claim that brain damage entails disruption in the abilities to transmit and maintain activation. However, there are three key limitations in the aphasia model. First, it is not sequential. The model's phonemes are retrieved all at once, contrary to data (e.g., Meyer, 1991). Second, the network structure is not learned. Finally, the model assumes the existence of pre-stored phonological frames that specify the syllabic and metrical structure

TABLE 2
Picture Naming Error Proportions from Selected Patients and Simulated Proportions from the Aphasia Model (from Dell et al., 1997)

Patient/ Parameters	Response Category					
	Correct	Semantic	Formal	Nonword	Mixed	Unrelated
I.G.	.69	.09	.05	.02	.03	.01
$p = .1, q = .86$.73	.13	.04	.05	.04	.01
L.H.	.69	.03	.07	.15	.01	.02
$p = .0057, q = .5$.69	.07	.06	.14	.01	.03
G.L.	.28	.04	.21	.30	.03	.09
$p = .079, q = .85$.27	.11	.20	.29	.03	.10




Figure 3. The architecture of the phonological error model. Each rectangle indicates a group of units.

of each word. While there may be considerable evidence for such frames (e.g., Sevald, Dell, & Cole, 1995), the aphasia model neither implements nor explains them. The next model that we present, the *phonological error model*, confronts these limitations.

The Phonological Error Model

The phonological error model (Dell, Juliano, & Govindjee, 1993) is an attempt to apply PDP principles specifically to phonological encoding. The model uses a simple recurrent network (Elman, 1990; Jordan, 1986) to map from a static representation of a word to a sequence of phonological features. Figure 3 shows the architecture.

The input layer represented the word to be spoken. In different versions of the model, the input was either a random bit vector (which can be viewed as either a lemma or a semantic representation) or a vector that was correlated with the word's form (either an underlying phonological representation or the orthographic input from a reading aloud task.) In both cases, the input remained unchanged during the production of the word. The input activation passed through a hidden layer to an output layer of 18 phonological features, one unit for each feature.

The phonological error model produces sequences of features by means of recurrent one-to-one connections from the output and hidden layers to layers of context units. One context layer is a copy of the hidden layer's activation on the previous pass through the network (*internal context units*) and the other corresponds to that for the output layer (*external context units*). Specifically, the production of a word (here, CAT) goes like this: The input units are activated in the pattern designated for CAT. The internal context units are initialized to zero, and the external context units are set to a pattern that symbolizes a word boundary (.5 on every unit). Activation spreads from the input and context units to hidden units and then to the output phonological features. The target activation pattern

corresponds to the features of the first phoneme /k/. To the extent that the output deviates from the target, weights are adjusted by backpropagation. The model's output and hidden layer activations are then copied to the external and internal context units, respectively. The context units keep track of where the model is in the sequence. After the /k/ is produced, the context represents the state of already having produced that phoneme, instead of the word-boundary state. This change in the context allows for the production of /æ/ in the next forward pass of activation. The process continues for the remainder of the word with the final target being the word boundary pattern.

The model was trained by repeatedly presenting words and adjusting weights. Dell et al. trained several models on vocabularies of 50 to 412 short English words (1–3 phonemes), and examined how performance differed with training vocabulary and architecture. Here our concern is with the model's ability to explain facts about phonological speech errors.

Some speech error effects have been interpreted as evidence for a *frame-and-slot* approach to word form retrieval (Shattuck–Hufnagel, 1979). Frame-and-slot models separate the retrieval of a word's sounds from the retrieval of a phonological frame. The frame represents the number of syllables in the word and the location of stress. Within the frame, each syllable is associated with slots that label the kind of sound (e.g., consonant or vowel) that the slot may hold. Placing the retrieved sounds in the slots assembles the word form. Recall that the aphasia model used an activation-sensitive version of this mechanism; the phoneme units with the highest activation levels were selected by linking them to frame slots.

The speech error effects that support the idea of a separate frame include the following: the phonotactic regularity of errors, syllabic constituent effects, and the existence of sound exchanges. Here we define these and explain why they are supportive of separately representing sounds from the structures in which they occur. First, phonological errors have a strong tendency to follow the phonotactic patterns of the language. Thus, in English, one would not expect to see slips such as “king” to “nging” because syllable-initial /ng/ does not occur in English. The phonotactic regularity of errors has been attributed to phonological rules that guide the insertion of retrieved sounds into frame slots. Presumably in English, the insertion of /ng/ into an onset slot would be blocked.

Syllabic constituent effects concern which parts of syllables are most likely to slip. Syllables are thought to have a hierarchical onset-rime structure in which, for example, a CVC syllable is composed of a C onset and a VC rime. Speech errors reflect this structure. For a CVC syllable, one is more likely to see a slip of either an onset (C) or a rime (VC) than other combinations such as the CV part of a CVC syllable. For example, the error “Tup Kin” instead of “Tin Cup” involves the movement of rime constituents. Because the constituent structure of syllables is often assumed to be a property of phonological frames, these effects support phonological frames.

Some phonological errors involve the exchange of speech sounds (e.g., “heft lemi-sphere”). Although these are not very common—only 5–10% of phonological errors are exchanges—they are clearly not random events. Rather, an initial anticipatory substitution (“left” being spoken as “heft”) appears to cause another substitution, typically in the next

word, in which the replaced sound replaces the anticipated sound (“lemisphere”). The existence of exchanges suggests the action of phonological frames. Each sound was erroneously placed in the other’s frame slots.

The phonological error model accounts for some of the error effects attributed to separate phonological frames although it lacks explicit frames. When noise is introduced into the model’s weights, it produces realistic phonological errors. Specifically, erroneous phoneme sequences in the model have a strong tendency to be phonotactically regular. The percentage of errors that were phonotactically legal in the model ranged between 87 and 100%. Moreover, the errors tend to involve the hypothesized frame constituents. The model produces more syllable onset than syllable coda errors, and it produces more rime (VC) errors than CV errors. In general, the model’s errors are sensitive to the structure of English words because it is trained on English words and it represents those words with linguistically motivated features (see Anderson, Milostan, & Cottrell, 1998). The superimposed weight changes associated with the training set creates sequential schemata, pathways in the model’s activation space that reflect common sequences of features. When the model errs, it sticks close to these pathways. Thus, the errors obey English phonotactics. The reason that the model’s errors tend to involve onsets rather than codas is a consequence of both the English vocabulary and the model’s sequential nature. There is more variety in word and syllable onsets than in codas. Hence there is more uncertainty about onsets, which makes them more error prone. This difference in uncertainty is enhanced by the sequential nature of the model. At the beginning of a word, the context units’ activation state is uninformative about the phoneme to be retrieved. However, as the model produces more of the word, the context units become more informative. For example, after having already retrieved /kæ./ the possible continuations are much fewer than before. That the model’s errors tend to involve VC (rime) units more than CV units is also due to the vocabulary structure. English, like most languages, tends to have fewer VC’s than CV’s (Kessler & Treiman, 1997) and consequently must “reuse” the VC’s that it does have. In the model the VC’s that are present thus become part of well-worn paths. When output jumps from one such path to another, a slip of an entire VC results.

The main problem with the phonological error model is that it has no mechanism for exchanges. The model could conceivably produce an anticipatory substitution such as “heft” for “left” in the context of “left hemisphere” through contamination from an upcoming word. But such an error would not naturally trigger a subsequent substitution to make the exchange “heft hemisphere.” The very fact that exchanges occur between structurally similar sounds points to a mechanism that binds values (retrieved sounds) to variables (slots in a frame). Specifically, in localist activation-based models with frames (e.g., Berg & Schade, 1992; Dell, 1986; Hartley & Houghton, 1996; Stemberger, 1985; MacKay, 1982), an exchange such as “heft hemisphere” can happen as follows: First, the activation of the /h/ node is greater than that of /l/ and so replaces it in the frame slot for the onset of the first syllable. The selected sound, /h/, then undergoes inhibition, which tends to prevent its reselection. When the onset slot of the next syllable is filled, the /l/, which did not undergo inhibition, may be more active than the inhibited /h/, and thus replace it completing the exchange. The phonological error model does not have this kind

of mechanism, or any other in which a substitution in one syllable triggers a corresponding substitution in a later syllable.

The phonological error model's failure to produce exchanges is a serious problem. It leads us to question its architecture. However, it does not cause us to abandon many of the principles present in the model. For example, the model attributes several error effects to the statistical structure of the word-form lexicon and to the fact that sounds are retrieved in sequence by means of a dynamic context. We believe that this attribution is correct. So, regardless of the architecture of the phonological access system, it should, in our view, embody mechanisms for sensitivity to sound statistics and sequence.

Before we turn to grammatical encoding, we should make a few observations about both of the lexical access models that we have reviewed. The aphasia model and the phonological error model are, at least on the surface, quite different. The former is a two-step interactive activation model that retrieves position-specific phonemes in parallel and links them to slots in a frame. The latter is a PDP simple recurrent network that learns distributed representations allowing for the sequential output of phonological features. Moreover, the phonological error model deals only with phonological processes, while the aphasia model does both lemma and phonological access.

Given these differences between the models, it is useful to consider how the phonological error model's approach could be extended to deal with the larger domain of the aphasia model. In fact, Plaut and Kello (1999) have constructed such a model. Their model learned to map from representations of word meaning to sequences of articulatory gestures by using error signals emerging from knowledge acquired during word comprehension. The resulting model has two key features of the aphasia model. First, connections run from semantics to phonology and in the reverse direction, making the activation flow interactive. Second, the model has something very much like two steps when it produces a word. This is because the intermediate layer must achieve its proper activation pattern before a sequential articulation process can begin. So, the first step involves retrieval of a static representation, and the second consists of turning that representation into a sequence.

The two steps in Plaut and Kello's model, however, are not the two steps of the aphasia model. The aphasia model's first step is retrieval of a word's lemma, while Plaut and Kello's model's first step is retrieval of a static phonological representation. It is difficult to tell whether these differences are fundamental or not because Plaut and Kello made no claims about grammatical processes. It may turn out that the "lemma" is a static representation at an intermediate level that serves both as input to a sequential phonological output process and as output from processes that map from messages onto word sequences. If so, then there is a great deal of concordance between the models. More generally, we believe that models such as the aphasia model are high level characterizations whose insights will be useful for understanding PDP implementations such as Plaut and Kello's model or the phonological error model. In the final section of this article, we ask whether a PDP model of grammatical encoding, one that shares many characteristics with the Plaut and Kello and phonological error models, can handle facts about the production of sentences.

III. GRAMMATICAL ENCODING

Like phonological encoding, grammatical encoding has often been conceptualized in terms of frames and slots based on patterns found in speech errors. In frame-and-slot models of grammatical encoding (e.g., Bock, 1982; Dell, 1986; Garrett, 1975; MacKay, 1982), frames represent syntactic structures with slots labeled with the grammatical classes (e.g., noun or verb) of the lemmas that may fill them. Analogous to the phonotactic regularity of sound errors is the grammaticality of word substitutions. Word substitutions and exchanges tend to involve words of the same grammatical class, such as “please pass the *fork*,” in which “fork” replaces “salt,” keeping the utterance grammatical while altering its meaning (Garrett, 1975). Within a frame-and-slot model, exchanges across grammatical classes such as “please *salt* the *pass*” are unlikely because they involve two violations: a noun in a verb slot and a verb in a noun slot.

Recall that the phonological error model accounted for phonological structural effects by means of a learning mechanism that derived structure from the statistics of the training set rather than from explicit word frames. Perhaps an analogous approach to sentence learning could be used to study grammatical structure. However, there are many differences between sentence and word production that make grammatical encoding a more difficult process to model. Here we will review some of these differences as part of a summary of psycholinguistic research into sentence production, and then describe a new model of sentence production, the *structural priming model* (Chang, Griffin, Dell, & Bock, 1997), which attempts to explain some structural effects using a connectionist model of learning.

Most sentences are novel, while most words are not. One speaks of “retrieving” a word’s form from memory, but of “generating” a sentence. Consequently, a greater emphasis must be placed on the generalization ability of a sentence-production model than on one that produces word forms. To get the right kind of generalization, one must first understand the nature of the input to production, the message, and then consider the mapping from the message to a word sequence. Because of the arbitrary but fixed mapping between a word’s meaning and its sound, the input to phonological encoding can be a unique representation lacking internal structure, as in the aphasia model’s words. In contrast, a message logically must possess internal structure to support generalization to novel utterances. The message must contain sufficient information about its elements to allow appropriate words to be selected and must express the relations among those elements—who did what to whom. The difficulty for the modeler lies in how to represent this information. Debates within linguistics have provided a wealth of ideas about what information is necessary in a representation of sentence meaning. But psycholinguistic research on the nature of message representation is scanty (see, however, Bock & Eberhard, 1993; Slobin, 1996).

An important feature of the mapping between messages and grammatical forms is its variability. First, there is lexical variability. Unlike word forms, which are usually unambiguously associated with the same sounds, there may be many ways to map between message elements and words. A cat can be “cat,” “animal,” “Spot,” “it,” and so on.

Second, a message can be realized by different syntactic structures without changing its core meaning. For example, when no particular element is in focus, the same proposition could be expressed with an active (“The horse kicked the cow”) or a passive sentence (“The cow was kicked by the horse”). When this syntactic flexibility is combined with flexibility in word choice, very similar statements can be made using very different syntactic structures, as in “Clinton defeated Dole,” “Dole lost to Clinton,” “Dole was defeated by the president,” and so on.

If messages do not determine word order, what does? In English, the assignment of lemmas to grammatical roles (e.g., subject, direct object) is the primary determinant of eventual word order and grammatical role assignment depends on the ease of lexical encoding. Evidence comes from studies showing lexical priming effects on grammatical role assignments (Bock, 1986b, 1987a). The easier it was to select a word to express a substantive concept, the more likely it was to be encoded as a sentential subject. This result implies that noun phrases are assigned to grammatical roles in the order in which their lemmas are selected and in the order of grammatical role prominence (subject, then direct object, then object of preposition). Moreover, more conceptually available message elements (by virtue of being more topical, imageable, animate, or prototypical) are placed in more prominent grammatical roles than are less accessible elements (e.g., Bock & Warren, 1985; Ferreira, 1994). Such conceptual factors probably influence grammatical role assignment indirectly by taking priority in lemma selection. Together, these findings indicate that grammatical encoding is highly opportunistic; the most prominent message elements are the first to be lexicalized and the earliest lexicalized concepts are assigned to the earliest occurring grammatical roles, such as sentential subject.

However, more than conceptual and lexical accessibility affect word order and sentence structure. Studies conducted by Bock and colleagues (Bock, 1986a; Bock & Loebell, 1990) demonstrate the existence of *structural priming*: Speakers tend to repeat the structures of previously uttered sentences even when the sentences differ in prosodic, lexical, and conceptual content. For example, speakers are more likely to describe the event depicted in Figure 4 with a passive, such as “A policeman is being hit by an ambulance,” if they just produced a sentence that was passive rather than active. Furthermore, this increase in the use of one structure does not appear to be caused by strengthening links between grammatical role assignments and event roles (e.g., agent, patient) in the message (Bock & Loebell, 1990). Intransitive-locatives (“The 747 was landing by the control tower”) and passives (“The 747 was landed by the control tower”) differ primarily in the event roles that their constituents play (e.g., whether the sentential subject is the agent). Nevertheless, a speaker is as likely to use a passive structure after producing an intransitive-locative prime as after a passive prime. This suggests that the link between grammatical and event roles is not the locus of structural priming. Rather, the priming appears to be related to the constituent structure of the sentences (e.g., whether there is a prepositional phrase after the main verb). Furthermore, results of a study by Bock, Loebell, and Morey (1992) suggest that conceptual factors influence when elements are lexicalized and assigned grammatical roles, while structural accessibility has an independent influence by affecting the type of grammatical roles that are filled (e.g., subject and




Figure 4. Example of a target picture and active and passive prime sentences (Bock; 1986a).

direct object for actives, or subject and object of preposition for passives and intransitive-locatives). Thus, grammatical encoding cannot be accomplished by blindly assigning lemmas to grammatical roles as concepts are lexicalized. Lemmas must be marked with event-related information to ensure that the relationship between message elements is not lost, otherwise an agent could easily become the subject of a passive sentence. This has been called the coordination problem (Bock, 1987b) and poses difficulties for all models of grammatical encoding.

In summary, grammatical encoding is a particularly challenging process to model for several reasons: (1) Little is known about its input representation, except that its internal structure must permit generalization. (2) The mapping between concepts and words is variable, as is the mapping between message relationships and grammatical roles. (3) The assignment of grammatical roles is constrained but not determined by message relationships. (4) There are structural priming effects of a character that suggest structural frames. A connectionist learning model overcoming these challenges needs to demonstrate structural effects without possessing explicit structures, role-binding without explicit tags or grammatical roles, and must make flexible and opportunistic decisions in the choice of words and structures.

Fortunately, some aspects of grammatical encoding may be readily accounted for within a connectionist framework. Recent studies indicate that the influence of a prime sentence persists across the production of up to 10 structurally unrelated sentences (Bock, Dell, Griffin, Chang, & Ferreira, 1996). This result suggests that priming may be a type

of implicit learning rather than the result of activation of structures in short-term memory. If so, connectionist-learning models associated with gradual weight change may be able to explain some features of priming. Furthermore, Elman (1993) demonstrated that a recurrent network could implicitly learn grammatical structure when trained to anticipate the next word in a sentence. The structural-priming model (Chang et al., 1997) to which we now turn, employed a related architecture in an attempt to produce grammatical sequences from a message and mimic the structural priming effect.

Structural Priming Model

The central claim of the structural-priming model is that structural priming is a form of implicit learning. In other words, the same mechanism through which the model learns to produce sentences causes the priming. To realize this claim, it was necessary to make the model accord with three basic assumptions about production. These are, first, production starts with a message expressing propositional content. Second, message elements may differ in their accessibility and these differences contribute to structural choices. Third, words are selected one at a time, with earlier selections constraining later ones, that is, processing is incremental and left-to-right. To reflect these assumptions, the model used a type of simple recurrent network that learned to map from a static message to a sequence of words, and it allowed for differential activation levels among message elements to determine the target sequence.

In this framework, structural priming results from the learning algorithm, which was backpropagation. When a prime sentence is produced, weight changes take place which favor the production of that sentence from its message. Chang et al.'s hypothesis was that these changes would generalize to structurally related sentences. So, a subsequent message that may be associated with more than one structure, such as an active/passive or double-object/prepositional dative option, would be more likely to be encoded using the structure of the prime. This is possible because the weight changes associated with a particular message-sentence mapping are shared with other message-sentence combinations.

Figure 5 illustrates the general theory behind the model. Like some other connectionist treatments of language (e.g. Christiansen & Chater, 1999; Plaut & Kello, 1999), Chang et al. propose a close relationship between the comprehension and production. In particular, they suggest that the context representations that guide sequential production arise during comprehension. For example, suppose that a simple recurrent network that maps word sequences onto messages carries out comprehension, and that this network uses the activation pattern of its hidden units as context to facilitate this mapping. We know that the changes in the activation patterns of these hidden units would come to be sensitive to the structure of input sentences (Elman, 1993) and to the mapping of this structure onto meaning. Chang et al. hypothesized that these hidden unit/context activation patterns could be directly used during production in the following way: At the beginning of a sentence to be produced, the context units would be set to null values indicating a sentence boundary. The production side of the model would then output the first word. This word

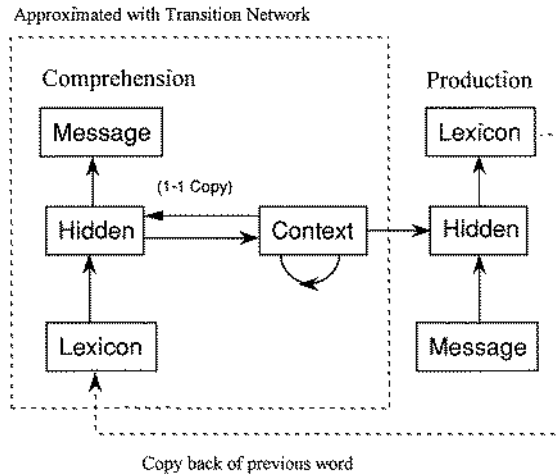


Figure 5. Comprehension can provide the dynamic context for production (Chang et al., 1997).

would be fed into the input of the comprehension side of the model, thereby updating the context units' activations. These activations would then serve as contextual input to production, effectively signaling the system that it is now time to produce the second word, and so on until the end of the sentence.

We now describe Chang et al.'s implemented model and their simulations of the structural priming effect. As Figure 5 shows, Chang et al., did not actually implement the comprehension side of their model. Rather they attempted to approximate the contribution from comprehension by creating a localist transition network and using it as context for production. This transition network will be described after we introduce other aspects of the implemented model.

The input to the model was a message, a set of localist semantic features representing a single proposition. The message was an 87-dimensional vector that represented concepts such as "boy," actions such as "walking," and the event roles, *agent*, *patient*, *recipient*, and *location*. The message remained activated throughout the production of the sentence. The relationships among the message participants involved associating blocks of features with event roles. So, within the agent block, there were 18 units including units for CHILD, MALE, and UNITARY. The patient, recipient, and location blocks also had 18 units each and these coded for the same features as the agent block. The action block had 15 features. These included localist units for specific actions such as WALKING, GIVING, or CHASING, and their number of arguments. For example, the message CHASE (BOYS, DOG) would be associated with activated agent units CHILD, MALE, MULTIPLE, patient units, BARKS, ANIMAL, UNITARY, and action units CHASING and 2-ARGUMENT.

Differences in conceptual accessibility were implemented by having the features of one role more activated than others. These differences determined the target structure of sentences during training. Given filled agent and patient roles, and a transitive action, the

TABLE 3
Sentence Types Used in the Structural Priming Model

Sentence Type	Percent in Training	Structure Sequence
Intransitive	17	AGENT VERB. Girl walks.
Active transitive	27	AGENT VERB PATIENT. Man chases dog.
Passive transitive	9	PATIENT AUX PASTPART by AGENT. Dog is chased by man.
Double object dative	17	AGENT VERB RECIPIENT PATIENT. Woman gives man dog.
Prepositional dative	17	AGENT VERB PATIENT PREP RECIPIENT. Woman gives dog to man.
Intransitive locative	4	AGENT AUX PRESPART PREP LOCATION. Boys are walking near bus.
Transitive locative	9	AGENT VERB PATIENT PREP LOCATION. Girls chase dogs to car.

model was trained to produce an active sentence if the agent was more activated than the patient, and a passive if the reverse was true. A message for the production of a double-object dative sentence, as opposed to a prepositional dative sentence, differed only in whether the patient or recipient was more highly activated.

The message units had learnable connections to 50 hidden units which, in turn, had learnable connections to the model's output layer. In the output layer, there was one unit for each of the 59 words in its vocabulary. These included singular and plural nouns (12 of each), 2 obligatorily transitive verbs, "chase(s)" and "feed(s)"; 2 optionally transitive verbs, "see(s)" and "hear(s)," 2 intransitive verbs, "walk(s)" and "live(s)," and 4 dative verbs, "give(s)," "make(s)," "show(s)," and "write(s)." Each verb that could be used transitively also had a past participle form (e.g. "heard") and each verb that could be used intransitively had a present participle form (e.g. "living"). There were also units for *is*, *are*, *by*, *near*, *for*, *to*, and *PERIOD* (end of sentence marker). Verbs had to agree with subject nouns in number. Table 3 gives examples of the sentences that the model was trained to produce.

The transition network served as another input layer to the model. It reflected the current state of the sentence from the perspective of the comprehension system. This network contained 10 localist nodes representing the following syntactic and event role categories: *PERIOD*, *VERB*, *AUX*, *PastParticiple*, *PresentParticiple*, *PREP*, *AGENT*, *PATIENT*, *RECIPIENT*, and *LOCATION*. Each of these nodes had connections with modifiable weights to each hidden unit. The activation of the transition network's nodes changes as the sentence progressed. Table 4 shows which nodes would activate and when they would activate for particular sentence types. Consider, for example, the sentence, "Girls give man robot." Before any word has been produced, the sentence boundary node *PERIOD* is the only unit in the transition network that is on. (Note that this *PERIOD* is different from the output layer unit for *PERIOD*). So, the input to the model consists of

TABLE 4
Sequence of Activated Units in the Transition Network for Context in the Structural Priming Model for Each Sentence Type

Intransitive	PERIOD → {AGENT PATIENT} → VERB → PERIOD
Active transitive	PERIOD → {AGENT PATIENT} → VERB → PATIENT → PERIOD
Passive transitive	PERIOD → {AGENT PATIENT} → AUX → PastP → PREP → AGENT → PERIOD
Double object dative	PERIOD → {AGENT PATIENT} → VERB → {RECIPIENT PATIENT} → PATIENT → PERIOD
Prepositional dative	PERIOD → {AGENT PATIENT} → VERB → {RECIPIENT PATIENT} → PREP → RECIPIENT → PERIOD
Transitive locative	PERIOD → {AGENT PATIENT} → VERB → {RECIPIENT PATIENT} → PREP → LOCATION → PERIOD
Intransitive locative	PERIOD → {AGENT PATIENT} → AUX → PresP → PREP → LOCATION → PERIOD

just PERIOD plus the appropriate message representation. Under these conditions, the model should produce “girls.” The word “girls” is then assumed to pass to the comprehension system, which would determine that it is likely an agent or possibly a patient. This interpretation by the comprehension system was implemented by turning on the nodes for AGENT and PATIENT in the transition network. (See, in Table 4, that for double object datives, the first transition is from PERIOD to [AGENT PATIENT].) The node for PERIOD would also remain partly on. Each transition network node retained half of its activation across the production of each additional word. Notice that the ambiguity associated with the role assignment for “girls” is only from the comprehension system’s perspective. The message on the production side knows that “girls” is the agent. The state of AGENT and PATIENT on and PERIOD half on then signals for the production of the next word, “give.” The comprehension of “give” turns on the VERB unit, which then enables the production of “man.” Again, because of ambiguity, the comprehension system does not know whether “man” is RECIPIENT or PATIENT, and so both of these units then turn on. The process continues with the production of the final word “robot,” which is now unambiguously identified as a PATIENT. Because the comprehension system is controlling the state changes of the context, there is no need to relearn the sentence patterns during production. However, the production system does have to learn how to associate these patterns with messages to produce sequential output. An important consequence of using a context derived from the comprehension system is that the contextual states are associated with temporary ambiguity (here, ambiguity about event roles of noun phrases). According to Chang et al., this ambiguity contributes some error to the production side of the model which in turn makes for greater weight change on that side, and ultimately to sizable structural priming due to those weight changes.

The model was trained using backpropagation. Weights were initialized to normally distributed random values with mean 0 and variance .5. The training corpus consisted of

3600 sentences reflecting the proportions of sentence types shown in Table 3. Each of these was trained an average of 31 times. The learning rate for the first quarter of the training was .06, and .03 for the remainder. Momentum was .9.

The 3600 training sentences represented a fraction of the 175,152 sentences that were possible given the vocabulary and sentence types. Aside from the fact that actives occurred more frequently than passives, the proportions of each sentence type in the training set were not designed to reflect the relative frequencies of the structures in natural language. After the model learned, it was tested on a random set of 400 sentences that preserved the proportion of sentence types in the corpus. 74% of these were novel sentences (66% of the non-novel sentences were intransitives because there is less opportunity for variety with this type). The model accurately produced the correct word 94% of the time.

To test priming, the model was first given a priming trial consisting of a message that required either a simple active or a passive sentence (in the case of datives, a double-object or prepositional dative sentence). As in training, weights were adjusted as the sentence was produced using a learning rate of .03. Then the model was given a target message that was neutral with respect to conceptual accessibility. For a transitive message, the agent and patient were equally activated, and correspondingly for the patient and recipient of a dative message. For the datives, each prime and target was a novel message, one that had not been previously trained. For transitives, one-fourth of the test messages had been trained before. Chang et al. recorded the percentage of times that each structure was produced as a function of the prime, the differences being the measure of priming.

The model exhibited a fair amount of structural priming. For example, "Boys chase dogs" as a prime would promote "Girl feeds cat" over "Cat is fed by girl." Figures 6 and 7 show the magnitude of active-passive and dative priming from Bock et al. (1996) and from the model. The sizes of the priming effects in the model match up well with the data. The absolute percentage of prepositional datives, though, is lower in the model than in the data. Pictures in the experiment were selected so that alternating structures were used equally often on average. No such constraint was applied to the model. The figures also show the model's successful simulation of the persistence of priming over 10 intervening neutral sentences (Bock et al., 1996). Thus, the model exhibits the phenomenon that motivates an account of priming as implicit learning. The same type of weight changes used in learning caused persistent structural priming.

Chang et al. also investigated whether the model could simulate the effects of Bock and Loebell (1990), who showed that priming is centered on the surface syntactic structure of the prime, rather than its mapping from event to grammatical roles. Here, the model's success was more checkered. It showed priming between intransitive locatives such as "boys are walking near bus" and the thematically different but structurally similar passive ("dog is chased by man"), in agreement with Bock and Loebell. However, it failed to show priming between transitive locatives such as "girls chase dogs near car" and prepositional datives, which Bock and Loebell had found.

In summary, the structural priming model successfully realizes the hypothesis that priming is a form of implicit learning. Moreover, it shows that structural priming may be

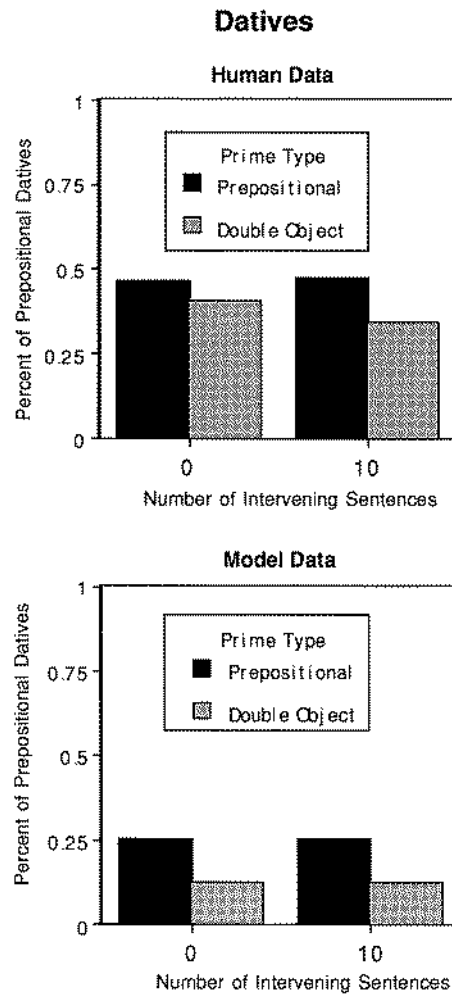


Figure 6. Dative priming over intervening sentences. The data are from Bock et al. (1996).

compatible with sequential connectionist models that derive structure from experience. Interestingly, Chang et al. believe that this structure must come from learning to comprehend, as well as learning to produce. The reasons that the model fails to account for the totality of the priming data are, of course, difficult to ascertain. However, we believe that its assumptions about message structure may be partly responsible. For example, the model assumes separate banks of message units for the four event roles that it uses. This assumption has some unrealistic consequences. First, it denies the possibility that roles have a similarity structure, for example, locations are like recipients (see Jackendoff, 1972). Second, it treats roles as categories that are independent of actions, contrary to some modern theories (e.g., Pollard & Sag, 1994). In general, connectionist models are ultimately only as good as their assumptions about input and output representations.




Figure 7. Active/passive priming over intervening sentences. The data are from Bock et al. (1996).

Progress in models of grammatical encoding is therefore dependent on the development of knowledge about meaning and communication.

Before concluding, let us briefly consider how the structural priming model might relate to the two lexical access models. First, consider the structural priming model's relation to the aphasia model. The aphasia model has a layer of lemma nodes that are actively selected and inserted into syntactic frames. The hidden layer of structural priming model may be seen as approximating the result of this selection and insertion process, with the contribution of the frame being associated with the context representations. However, these similarities between the models should not be overstated. For example, the context

representations of the structural priming model are not pure syntactic frames, something that is made evident by the model's failure to show priming between the structurally identical prepositional datives and transitive locatives.

Next, consider the structural priming model in concert with the phonological error model. They are both recurrent network models and it is tempting to just link them up, the output of the structural priming model providing input to the phonological error model. Such a linkage, however, creates independent phonological and grammatical modules and hence does not allow for interaction between phonological and lexical representations, which we argued (through the aphasia model) was needed to explain interactive error effects. Another point about the two recurrent network models concerns exchanges. Recall that the phonological error model does not produce exchanges such as "cogs and dats" for "dogs and cats." It turns out that the structural priming model does (rarely) produce word exchanges, e.g. "*Dog* is chased by *cat*" as a blend of two possible correct sentences "Dog chases cat" and "Cat is chased by dog." One can speculate that what is lacking in the phonological error model is an analogous mechanism at the phonological level—a conflict between two alternative *correct* outputs, along with a constraint that tries to have each intended output unit (word or phoneme) occur just once in the output sequence. Perhaps, a combination phonological/grammatical model that is associated with alternations such as "cats and dogs" versus "dogs and cats" could produce "cogs and dats." In fact, this conflict account of phonological exchanges is not new. It is the competing plan hypothesis of Baars (1980) and is even similar to Freud's (1891/1953) ideas on speech errors.

IV. CONCLUSIONS

Our review of connectionist models in production has focussed on our recent work in lexical access and grammatical encoding. Our lexical access models, when considered together with other spreading activation models (e.g., Roelofs, 1997), provide a fair coverage of the data. One would expect so, because these kinds of models have been around for some time. The main problem is that different models have different strengths and weaknesses, and so there is no unified approach that has been shown to do the job. Work on models of grammatical encoding, however, is just beginning, and so what has been accomplished is exciting (at least to us), although quite limited.

We consider two key features of production to be serial output and sensitivity to linguistic structure. In the PDP recurrent network models, structure and order are entwined in the sequential schemata that develop from the superimposed weight changes associated with the training set. It remains to be seen whether these schemata have the right characteristics to support generalization in grammatical encoding (e.g., the right kind of structural priming), or account for the sequential structure of word-forms (e.g., exhibit phonological exchange errors).

Ultimately, we believe that connectionist models of the acquisition of the skills of speaking (and comprehending) will contribute to explanations of the nature of language—why it is the way it is (see, e.g., Christiansen & Devlin, 1997; Hare & Elman, 1995; Gupta & Dell, 1999). Moreover, we believe the PDP approach offers the best chance to explain

production as a skill, as something that one learns to do over years of experience. Perhaps most importantly, a PDP approach to language production expresses its commonalities with other linguistic, and even nonlinguistic, skills.

Acknowledgments: The authors thank Kay Bock, Vic Ferreira, Prahlad Gupta, Anita Govindjee, and Linda May for their assistance on this project and Morten Christiansen and Nick Chater for helpful comments. The research was supported by NSF SBR 93-19368, NIH DC-00191 and HD 21011.

REFERENCES

- Anderson, K., Milostan, J., & Cottrell, G. W. (1998). Assessing the contribution of representation to results. In *Proceedings of the 20th Annual Conference of the Cognitive Science Society* (pp. 48–53). Mahwah, NJ: Erlbaum.
- Baars, B. J. (1980). The competing plan hypothesis: An heuristic viewpoint on the causes of errors in speech. In H. W. Dechert & M. Raupach (Eds.), *Temporal variables in speech* (pp. 39–49). The Hague: Mouton.
- Berg, T., & Schade, U. (1992). The role of inhibition in a spreading activation model of language production, Part 1: The psycholinguistic perspective. *Journal of Psycholinguistic Research*, 22, 405–434.
- Bock, J. K. (1982). Towards a cognitive psychology of syntax: Information processing contributions to sentence formulation. *Psychological Review*, 89, 1–47.
- Bock, J. K. (1986a). Syntactic persistence in language production. *Cognitive Psychology*, 18, 355–387.
- Bock, J. K. (1986b). Meaning, sound, and syntax: Lexical priming in sentence production. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 12, 575–586.
- Bock, J. K. (1987a). An effect of accessibility of word forms on sentence structures. *Journal of Memory and Language*, 26, 119–137.
- Bock, J. K. (1987b). Coordinating words and syntax in speech plans. In A. Ellis (Ed.), *Progress in the psychology of language* (Vol. 3, pp. 337–390). London: Erlbaum.
- Bock, J. K. (1990). Structure in language: Creating form in talk. *American Psychologist*, 45, 1221–1236.
- Bock, J. K., Dell, G. S., Griffin, Z., Chang, F., & Ferreira, V. S. (1996). Structural priming as implicit learning. Paper presented at the meeting of the Psychonomic Society, Chicago, IL.
- Bock, J. K., & Eberhard, K. M. (1993). Meaning, sound, and syntax in English number agreement. *Language and Cognitive Processes*, 8, 57–99.
- Bock, J. K., & Loebell, H. (1990). Framing sentences. *Cognition*, 35, 1–39.
- Bock, J. K., Loebell, H., & Morey, R. (1992). From conceptual roles to structural relations: Bridging the syntactic cleft. *Psychological Review*, 99, 150–171.
- Bock, J. K., & Warren, R. K. (1985). Conceptual accessibility and syntactic structure in sentence formulation. *Cognition*, 21, 47–67.
- Burke, D. M., MacKay, D. G., Worthley, J. S., & Wade, E. (1991). On the tip of the tongue: What causes word finding failures in young and older adults? *Journal of Memory and Language*, 30, 542–579.
- Caramazza, A. (1997). How many levels of processing are there in lexical access? *Cognitive Neuropsychology*, 14, 177–208.
- Chang, F., Griffin, Z. M., Dell, G. S., & Bock, K. (1997, Aug.). Modeling structural priming as implicit learning. Presented at Computational Psycholinguistics, Berkeley, CA.
- Christiansen, M. H., & Chater, N. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*, 23, 157–206.
- Christiansen, M. H., & Devlin, J. T. (1997). Recursive inconsistencies are hard to learn: A connectionist perspective on universal word order correlations. In *Proceedings of the 19th Annual Cognitive Science Society Conference* (pp. 113–118). Mahwah, NJ: Erlbaum.
- Cutting, J. C., & Ferreira, V. S. (1999). Semantic and phonological information flow in the production lexicon. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 318–344.
- Dell, G. S. (1986). A spreading activation theory of retrieval in language production. *Psychological Review*, 93, 283–321.

- Dell, G. S., Juliano, C., & Govindjee, A. (1993). Structure and content in language production: A theory of frame constraints in phonological speech errors. *Cognitive Science*, *17*, 149–195.
- Dell, G. S., & Reich, P. A. (1977). A model of slips of the tongue. In R. J. Dipietro & E. L. Blansitt (Eds.), *The third LACUS forum* (pp. 448–455). Columbia, SC: Hornbeam Press.
- Dell, G. S., Schwartz, M. F., Martin, N., Saffran, E. M., & Gagnon, D. A. (1997). Lexical access in aphasic and nonaphasic speakers. *Psychological Review*, *104*, 801–939.
- Eikmeyer, H.-J., & Schade, U. (1991). Sequentialization in connectionist language-production models. *Cognitive Systems*, *3*(2), 128–138.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, *14*, 213–252.
- Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, *48*, 71–99.
- Ferreira, F. (1994). Choice of passive voice is affected by verb type and animacy. *Journal of Memory and Language*, *33*, 715–736.
- Ferreira, V. (1996). Is it better to give than to donate? Syntactic flexibility in language production. *Journal of Memory and Language*, *35*, 724–755.
- Ferreira, V. (1997). *Syntactic and lexical choices in language production: What we can learn from “that,”* Ph.D. dissertation.. University of Illinois at Urbana—Champaign.
- Freud, S. (1891/1953). *On aphasia: A critical study* (Zur Auffassung der Aphasien). NY: International Universities Press.
- Fromkin, V. A. (1971). The non-anomalous nature of anomalous utterances. *Language*, *47*, 27–52.
- Garrett, M. F. (1975). The analysis of sentence production. In G. H. Bower (Ed.), *The psychology of learning and motivation* (pp. 133–175). San Diego: Academic Press.
- Griffin, Z. M., & Bock, K. (1998). Constraint, word frequency, and the relationship between lexical processing levels in spoken word production. *Journal of Memory and Language*, *38*, 313–338.
- Gupta, P. (1996). Immediate serial memory and language processing: Beyond the articulatory loop. Beckman Institute Cognitive Science Technical Report, CS-96–02.
- Gupta, P., & Dell, G. S. (1999). The emergence of language from serial order and procedural memory. In B. MacWhinney (Ed.), *Emergentist approaches to language* (pp. 447–481), 28th Carnegie Mellon Symposium on Cognition. Mahwah, NJ: Erlbaum.
- Hare, M., & Elman, J. L. (1995). Learning and morphological change. *Cognition*, *56*, 61–98.
- Harley, T. A. (1984). A critique of top-down independent levels models of speech production: Evidence from non-plan-internal speech errors. *Cognitive Science*, *8*, 191–219.
- Hartley, T. A., & Houghton, G. (1996). A linguistically-constrained model of short-term memory for nonwords. *Journal of Memory and Language*, *35*, 1–31.
- Houghton, G. (1990). The problem of serial order: A neural network model of sequence learning and recall. In R. Dale, C. Mellish, & M. Zock (Eds.), *Current research in natural language generation* (pp. 287–319). London: Academic Press.
- Jackendoff, R. (1972). *Semantic interpretation in generative grammar*. Cambridge, MA: MIT Press.
- Jordan, M. I. (1986). *Serial order: A parallel distributed processing approach* (ICS Technical Report 8604). University of California at San Diego, La Jolla, CA.
- Kempen, G., & Hoenkamp, E. (1987). An incremental procedural grammar for sentence formulation. *Cognitive Science*, *11*, 201–258.
- Kessler, B., & Treiman, R. (1997). Syllable structure and the distribution of English syllables. *Journal of Memory and Language*, *37*, 295–311.
- Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.
- MacKay, D. G. (1982). The problems of flexibility, fluency, and speed-accuracy trade-off in skilled behaviors. *Psychological Review*, *89*, 483–506.
- MacKay, D. G. (1987). *The organization of perception and action: A theory for language and other cognitive skills*. New York: Springer-Verlag.
- 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.
- Meyer, A. S. (1991). The time course of phonological encoding in language production: Phonological encoding inside a syllable. *Journal of Memory and Language*, *30*, 69–89.
- Miozzo, M., & Caramazza, A. (1997). The retrieval of lexical-syntactic features in tip-of-the-tongue states. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 1–14.

- Peterson, R. R., & Savoy, P. (1998). Lexical selection and phonological encoding during language production: Evidence for cascaded processing. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *24*, 539–557.
- 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: Erlbaum.
- Pollard, C., & Sag, I. (1994). *Head-driven phrase structure grammar*. Chicago, IL: University of Chicago Press.
- Roelofs, A. (1996). Computational models of lemma retrieval. In T. Dijkstra & K. de Smedt (Eds.), *Computational psycholinguistics* (pp. 308–327). London: Taylor & Francis.
- Roelofs, A. (1997). The WEAVER model of word-form encoding in speech production. *Cognition*, *64*, 249–284.
- Schriefers, H., Meyer, A. S., & Levelt, W. J. M. (1990). Exploring the time-course of lexical access in production: Picture-word interference studies. *Journal of Memory and Language*, *29*, 86–102.
- Sevold, C. A., Dell, G. S., & Cole, J. (1995). Syllable structure in speech production: Are syllables chunks or schemas? *Journal of Memory and Language*, *34*, 807–820.
- Shattuck-Hufnagel, S. (1979). Speech errors as evidence for a serial-order mechanism in sentence production. In W. E. Cooper & E. C. T. Walker (Eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett* (pp. 295–342). Hillsdale, NJ: Erlbaum.
- Slobin, D. I. (1996). From “thought and language” to “thinking for speaking.” In J. Gumperz & S. C. Levinson (Eds.), *Rethinking linguistic relativity* (pp. 70–96). Cambridge, MA: Cambridge University Press.
- Stemberger, J. P. (1985). An interactive activation model of language production. In A. W. Ellis (Ed.), *Progress in the psychology of language* (Vol. 1, pp. 143–186). Hillsdale, NJ: Erlbaum.
- Stemberger, J. P. (1989). Speech errors in early child language production. *Journal of Memory and Language*, *28*, 164–188.
- van Turenhout, M., Hagoort, P., & Brown, C. M. (1997). Electrophysiological evidence on the time course of semantic and phonological processes in speech production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 787–806.
- Vigliocco, G., Antonini, T., & Garrett, M. F. (1997). Grammatical gender is on the tip of Italian tongues. *Psychological Science*, *8*, 314–319.
- Wheeldon, L. R., & Levelt, W. J. M. (1995). Monitoring the time course of phonological encoding. *Journal of Memory and Language*, *34*, 311–334.