

# Can connectionist models of phonology assembly account for phonology?

IRIS BERENT

Florida Atlantic University, Boca Raton, Florida

Connectionist models have gained considerable success as accounts of how printed words are named. Their success challenges the view of grapheme-to-phoneme correspondences (GPCs) as rules. By extension, however, this challenge is sometimes interpreted also as evidence against linguistic rules and variables. This inference tacitly assumes that the generalizations inherent in reading (specifically, GPCs) are similar in their scope to linguistic generalizations and that they are each reducible to token associations. I examine this assumption by comparing the scope of generalizations required for mapping graphemes to phonemes and several linguistic phonological generalizations. Marcus (1998b) distinguishes between two types of generalizations: those that fall within a model's training space and those that exceed it. The scope of generalizations is determined by the model's representational choices—specifically, the implementation of operations over mental variables. An analysis of GPCs suggests that such generalizations do not appeal to variables; hence, they may not exceed the training space. Likewise, certain phonological regularities, such as syllable phonotactic constraints and place assimilation, may be captured by an associative process. In contrast, other phonological processes appeal to variables; hence, such generalizations potentially exceed the training space. I discuss one such case, the obligatory contour principle. I demonstrate that speakers conform to this constraint and that their behavior is inexplicable by the statistical structure of the language. This analysis suggests that, unlike GPCs, phonological generalizations may exceed the training space. Thus, despite their success in modeling GPCs, eliminative connectionist models of phonology assembly may be unable to provide a complete account for phonology. To the extent that reading is subject to phonological constraints, its modeling may require implementing operations over variables.

The past decade has witnessed a fierce debate between symbolic and associative accounts of mental architecture. At the center of this controversy are two case studies. One case study pertains to reading—specifically, the naming of printed words. The second case study concerns language—most notably, inflectional morphology. There are clearly some striking similarities between linguistic competence and the ability to name printed words.<sup>1</sup> Both abilities exhibit productivity: Speakers can inflect new words (e.g., *wug-wugs*) and pronounce unfamiliar letter sequences (e.g., *blin*). Both domains also manifest quasiregularity: They exhibit systematic relationships between inputs and outputs but admit many exceptions (Plaut, McClelland, Seidenberg, & Patterson, 1996). For instance, the English orthography includes highly regular correspondences between letters and sounds (e.g., *cat*), coupled with some highly unpredictable correspondences (e.g., *choir*). Likewise, inflectional morphology encompasses both regular (e.g., *rat-rats*, *like-liked*) and irregular correspondences (e.g., *mouse-mice*, *write-wrote*) between the base morpheme and its inflection. These similarities

give rise to the proposal that the mental computations in these two domains may be comparable.

If the knowledge of language and the ability to read words aloud are comparable, then the architectures of word-naming models could potentially capture linguistic competence. Connectionist models of word naming have gained considerable success. Proponents of these models interpret this success as indicating their potential to account for linguistic knowledge—specifically, the knowledge of word formation. For instance, upon reviewing the performance of their word-reading model, Plaut et al. (1996) note that

the approach taken here is applicable, we believe, to a wide range of linguistic and cognitive domains—essentially to all those with quasi-regular structure in the sense that there is systematicity that coexists with some arbitrariness and many exceptions. The first domain to which the approach was applied was that of inflectional morphology (Rumelhart & McClelland, 1986). As stated in the introduction, this application certainly remains controversial. Pinker and his colleagues (Marcus et al., 1992, Pinker, 1991, Pinker & Prince, 1988) continue to maintain that no single mechanism can fully capture the behavior of the regular inflectional process and the handling of exceptions. . . . We see little in these criticisms that stands against the applicability of the connectionist approach in principle. Indeed, the arguments raised in these articles do not, in general, reflect a full appreciation of the capabilities of connectionist networks in quasi-regular do-

This research was supported by NIH FIRST Award 1R29 DC03277-017 to I.B. I thank Gary Marcus for discussions of this work. Correspondence concerning this article should be addressed to I. Berent, Department of Psychology, Florida Atlantic University, 777 Glades Road, P.O. Box 3091, Boca Raton, FL 33431-0991 (e-mail: iberent@fau.edu).

mains. For example, Pinker (1991) did not acknowledge that connectionist models of both reading aloud (as shown here and in S & M89) and of inflectional morphology (Daugherty & Seidenberg, 1992) show the very frequency  $\times$  regularity interaction that he takes as the key indicators of the operation of a (frequency insensitive) rule system and a (frequency sensitive) lexical lookup mechanism. (pp. 106–107)

Does the success of word-reading models indicate the potential of their architectures to account for language? This paper examines this question. Because most connectionist models of reading aloud have not been tested in the domain of language,<sup>2</sup> it is difficult to evaluate the performance of specific models across domains. Rather than comparing computational models, this paper focuses primarily on the scope of generalizations exhibited by humans in the two areas. If connectionist models of reading aloud are to account for linguistic competence, then the scope of generalizations in these domains must be comparable. To determine whether generalizations are comparable, I draw heavily on the work of Gary Marcus. Marcus (1998a, 1998b, 2001) defines the scope of generalizations and formally demonstrates that their learnability depends on a model's representational scheme. I begin the discussion by examining the representational characteristics of connectionist models of reading aloud and symbolic representations. I next analyze the scope of generalizations exhibited in the mapping of graphemes to phonemes and compare them with the scope of some linguistic generalizations. To facilitate the comparison, I focus on an aspect of linguistic competence that is tightly linked to the mapping of graphemes to phonemes—namely, phonology, the sound structure of language. If connectionist models of reading aloud are a viable account of linguistic competence, they should be expected to capture phonological generalizations.

### Representational Commitments: The Role of Variables

The wide interest in the performance of connectionist models is due to their novel account of mental representations, an account that often departs in crucial ways from symbolic theories of cognition. An important distinction between symbolic representations and the representations used in many connectionist models concerns mental variables (Marcus, 1998a, 1998b, 2001; Pinker & Prince, 1988). Variables are abstract placeholders. For instance, one could use the variable *verb stem* as a placeholder for instances such as *like*, *run*, *establish*, or *participate*. According to symbolic accounts of cognition, mental representations encode variables (Fodor & Pylyshyn, 1988; Pinker & Prince, 1988). Mental processes operate over variables, rather than the instances they encode. For example, the past tense rule combines the variables *verb stem* and *past tense suffix*. Consequently, it can apply across the board to any instance of a *verb stem* (e.g., *blix*), regardless of its sound, meaning, or resemblance to familiar regular verbs (Pinker,

1999). The application of symbolic operations to variables thus allows treating all instances alike and ignoring their idiosyncrasies (Marcus, 1998b, 2001).

A prominent group of connectionist models, however, eliminates operations over variables (e.g., Elman, 1993; Elman et al., 1996; Hare & Elman, 1995; Plaut et al., 1996; Rueckl, Mikolinski, Raveh, Miner, & Mars, 1997; Rumelhart & McClelland, 1986; Seidenberg, 1987, 1997; Seidenberg & McClelland, 1989a). The cognitive hypothesis tacit in these models offers a radical alternative to the symbolic account. Variables, on this view, are merely a partial, coarse description of human behavior. This description is demonstrably incorrect. Because (on the connectionist view) variables play no role in cognition, a closer scrutiny is bound to reveal some fundamental discrepancies between the predictions of the symbolic account and human behavior. For instance, in referring to their interactive activation model of word perception, Rumelhart and McClelland note that

it is not correct to say that the model exactly mimicked the behavior we would expect to emerge from a system which makes use of explicit orthographic rules. However, neither do human subjects . . . . Thus, human perceptual performance seems, in this case at least, to be characterized only approximately by rules. (1986, p. 218)

In the absence of variables, cognition is constrained primarily by the properties of specific instances and their statistical distribution. Following Marcus, Brinkmann, Clahsen, Wiese, and Pinker (1995), I refer to the hypothesis that mental representations lack operations over variables as the *pattern associator hypothesis*. Conversely, the view of the mind as representing variables and operating over variables is the *symbolic hypothesis*. Before discussing the viability of these approaches as accounts of word reading and linguistic knowledge, it is important to clear up some side issues that could obscure the contrast between these two hypotheses.

1. *Pattern associators versus connectionism*. It has become commonplace to contrast the symbolic hypothesis with connectionism. Connectionism, however, is not inherently incompatible with the representation of variables (Marcus, 2001). Connectionism is a formalism that may be used in a wide range of models. Some of these models embody operations over variables (e.g., Hummel & Holyoak, 1997; Shastri, 1999; Smolensky, 1995, 1999); others eliminate them (e.g., Elman, 1993; Hare & Elman, 1995; Plaut et al., 1996; Plunkett & Marchman, 1993; Rueckl et al., 1997; Rumelhart & McClelland, 1986; Seidenberg & McClelland, 1989a). Connectionist models that implement operations over variables are perfectly compatible with the symbolic hypothesis. The following discussion specifically examines the subset of connectionist models that eliminates them. Following Pinker and Prince (1988), I refer to this approach as *eliminative connectionism*.

2. *Pattern associators versus feedforward networks*. The term *pattern associators* is often used to refer specifi-

cally to feedforward networks. In contrast, the following discussion uses this term to refer to any model that embodies the pattern associator hypothesis (i.e., a model that eliminates operations on variables); it does not imply any particular connectionist architecture.

3. *Variables versus rules.* Just as the pattern associator hypothesis must be distinguished from connectionism, the symbolic hypothesis should not be equated with rules. The symbolic hypothesis postulates the representation of variables and the implementation of operations over variables. Variables, however, may be incorporated within either rules or constraints<sup>3</sup> (see also Marcus, 2001). Rules and constraints may further be either inviolable or violable. The representation of variables is thus fully compatible with the view of mental processes as governed by violable constraints. In fact, recent linguistic theories account for linguistic competence as a ranked set of innate, violable constraints appealing to mental variables (Prince & Smolensky, 1993, 1997; Smolensky, 1999; Tesar & Smolensky, 2000; for a review, see Archangeli & Langendoen, 1997).

4. *Variables versus specific linguistic accounts.* The symbolic hypothesis is moot not only with respect to the format of linguistic knowledge (i.e., rules vs. constraints), but also with regard to its specific contents and source. For instance, consider phonology. Some linguistic theories account for phonological structure by appealing to discrete, atemporal mental abstractions (e.g., Chomsky & Halle, 1968), whereas others concern temporal articulatory gestures (Browman & Goldstein, 1989; Fowler, 1993; Gafos, 1999). Each of these approaches, however, may or may not appeal to variables. The symbolic hypothesis is not inherently linked to any particular view regarding the contents of linguistic knowledge. The following discussion examines accounts of language only with regard to their assumptions concerning the representation of variables. A more specific assessment (i.e., the discrimination between various symbolic accounts) falls beyond the scope of this investigation.

**An outline.** The pattern associator and symbolic hypotheses hold opposing views on the role of variables in mental representations. These representational commitments have direct implications for the scope of generalizations exhibited by a model. Marcus (1998b, 2001) outlines a distinction between two types of generalizations: One type extends beyond the learner's training space, whereas the other does not. I apply this taxonomy for evaluating the potential of eliminative connectionist models of word reading to account for language. I claim that eliminative connectionist accounts of reading may offer a viable account for phonology only if the scope of generalizations in the two domains is comparable and that the comparability of generalizations depends on the role of variables in these domains. An analysis of mapping graphemes to phonemes suggests that this domain does not appeal to variables, nor does it require generalizations beyond the reader's training space. The absence of variables and the

limited scope of generalizations are also characteristic of some phonological processes. Other phonological generalizations, however, may exceed the scope of the generalizations inherent in grapheme-to-phoneme correspondences (GPCs). The elimination of variables may thus prevent pattern associator accounts of word reading from providing a complete account of phonological competence.

### Does the Mapping of Graphemes to Phonemes Appeal to Variables?

One of the most influential symbolic accounts of word pronunciation is Coltheart's dual-route model (Coltheart, 1978; Coltheart, Curtis, Atkins, & Haller, 1993; Rastle & Coltheart, 1999). This model postulates (at least) two separate mechanisms for the pronunciation of printed words: a productive mechanism, assembling pronunciation by a set of correspondences between graphemes and phonemes, and a direct route, retrieving the pronunciation of a word as a whole on the basis of its spelling.

The dual-route account embodies three independent assumptions regarding the mental and neurological machinery dedicated to reading. The first assumption is the segregation of the lexical and the sublexical routes. Whole-word spelling-to-sound correspondences are stored in the lexicon, whereas subword correspondences are computed by a separate route that converts graphemes to phonemes. Although this assumption remains subject to an ongoing debate (e.g., Balota & Spieler, 1998; Besner, 1999; Besner, Twilley, McCann, & Seergobin, 1990; Plaut, 1999; Seidenberg & McClelland, 1989b; Seidenberg & Plaut, 1998; Spieler & Balota, 1997), it is not unique to the symbolic account of reading. Zorzi, Houghton, and Butterworth (1998) demonstrated that the segregation of the sublexical and the lexical routes facilitates the acquisition of spelling-to-sound regularities by a feedforward two-layer network. Distinct procedures dedicated to the activation of whole-word versus subword segments are found also in the multiple-trace connectionist model of word reading (Ans, Carbonnel, & Valdois, 1998). A second assumption of the symbolic account is that the mapping of graphemes to phonemes proceeds in a serial fashion. Like the route segregation assumption, however, the assumption of seriality is not specific to symbolic accounts. Serial effects also emerge in a connectionist network that incorporates a sequential *fixation* component despite the elimination of GPC rules (Plaut, 1999). The assumptions of pathway segregation and seriality are thus not inherent to the symbolic hypothesis, and I will not discuss them any further.

At the heart of the symbolic account is a third assumption concerning the computational characteristics of the two pathways. The symbolic account of word pronunciation proposes that the assembly of subword correspondences is achieved by rules, whereas lexical correspondences are implemented by an associative process

(Coltheart et al., 1993; Rastle & Coltheart, 1999). For instance, the assembly of phonology for the word *mint* relies on the rules  $m \rightarrow /m/$ ,  $i \rightarrow /ɪ/$ ,  $n \rightarrow /n/$ ,  $t \rightarrow /t/$  (Rastle & Coltheart, 1999). It is this computational assumption that apparently provides the principal motivation for the segregation of the two pathways: If the computation of word-size spelling-to-sound correspondences is similar to the computation of subword correspondences, then their implementation by a single pathway is at least possible. Conversely, if the computation of word-size spelling-to-sound correspondences differs from the computation of subword correspondences, then the existence of two pathways is inevitable. The hypothesis that phonology assembly is achieved by rules has faced mounting criticism, which is summarized below.

### What is Regular?

A major criticism of a symbolic assembly mechanism concerns the definition of its potential outputs. The view of assembly as a symbolic process presupposes some distinction between its possible outputs (regular correspondences) and impossible outputs (irregular correspondences). No such distinction has been found. The problem of explaining “what is regular” is not simply a matter of sorting words into *regular* or *irregular* categories. The dual-route model can clearly distinguish between regular and irregular words by stipulating a rule. For instance, the model can easily capture the fact that the pronunciation of the word *pint* is unpredictable by stipulating that this word violates the rule  $i \rightarrow /ɪ/$ . In reference to this rule, *pint* is clearly irregular, whereas *mint* is regular. However, the model fails to motivate this stipulation. The only difference between the correspondences in irregular and regular words is statistical: The strength of association between the vowel grapheme and the corresponding phoneme. Regularity, here, is thus indistinguishable from frequency (Van Orden, Pennington, & Stone, 1990). Indeed, in cases in which the frequency of the various correspondences of a given grapheme is similar, the distinction between *regular* and *irregular* correspondences becomes altogether ambiguous.

The absence of a principled distinction between regular and irregular correspondences of graphemes to phonemes stands in marked contrast to some linguistic regularities. Consider again inflectional morphology. There are several well-defined distinctions between regular and irregular inflection, and these differences appear to be independent of the statistical properties of the language. One distinction between regular and irregular forms concerns their sensitivity to stored tokens. A novel word that is similar to familiar irregular verbs (Prasada & Pinker, 1993) or nouns (Berent, Pinker, & Shimron, 1999) is likely to take their irregular inflection. For instance, the novel verb *spling* (similar to the irregular verbs *sling*, *slink*, *sting*, etc.) is more likely to take irregular inflection than is *nist* (which is distant from irregular verbs). In contrast, a novel verb that is similar to familiar regular verbs (e.g., *plip*, similar to the regular verbs *drip*, *chip*, *clip*, *flip*, *skip*, etc.) is no more likely to take regular inflection than a novel verb that

does not resemble any regular verb (e.g., *ploamph*; see Prasada & Pinker, 1993). Put differently, the similarity to stored tokens selectively affects irregular, but not regular, inflection.

Regular and irregular inflection also can be contrasted in terms of their admissibility in various grammatical processes. Certain linguistic processes are selective with respect to the regularity of inflection. For instance, compounding typically admits irregular (e.g., *mice-eater*), but not regular, plurals<sup>4</sup> (e.g., *\*rats-eater*; see Alegre & Gordon, 1996; Gordon, 1985). The inflection of names presents the opposite pattern. Names formed from existing nouns take regular, but not irregular, plurals<sup>5</sup> (e.g., *the Childs*; see Berent et al., 1999; Kim, Marcus, Pinker, Hollander, & Coppola, 1994). Irregular inflection is also banned for nouns and verbs that are derived from different grammatical categories—that is, denominals and deverbals. For instance, irregular-sounding verbs that are derived from a noun (e.g., *he flied out to center field*, from the noun *fly ball*) take regular, rather than irregular, inflections (Kim, Pinker, Prince, & Prasada, 1991). The preferences concerning the inflection of nouns inside compounds, on the one hand, and names, denominals, and deverbals, on the other, are thus complementary, and they are specifically defined in reference to regularity.

In contrast to the regularity of GPCs, the preferences concerning regular inflection cannot be explained by the statistical structure of the English language. The ban on regular plurals in a compound is not simply due to their sound (*s*-ending) or meaning (plural), since singulars that sound like regular plurals (e.g., *pants*) and irregular plurals (e.g., *mice-eater*) are produced inside a compound (Gordon, 1985).<sup>6</sup> Likewise, the blocking of irregular inflection for words derived from a different grammatical category is inexplicable by their unusual meaning, since metaphors do not necessarily take regular inflection<sup>7</sup> (Kim et al., 1991). Indeed, the selective admissibility of regular versus irregular forms to these grammatical processes is observed across languages that vary considerably in their phonological and statistical properties. The avoidance of regular plurals inside compounds has been observed in English (Gordon, 1985), where most plurals are regular, and German, where plurals are overwhelmingly irregular (Clahsen, 1999). Similarly, the blocking of irregular inflection for names has been, so far, observed in English, French, Dutch, Hungarian, Arabic, and Hebrew (Pinker, 1999). There appears to be a clear, cross-linguistic distinction between regular and irregular inflection. Regular inflection differs from irregular inflection in that the regular functions as a default: It applies upon the failure to activate a stored irregular root. Regular inflection operates as a default even when the regular type is less frequent than irregular types (Marcus et al., 1995), and it is independent of the distribution of regular and irregular forms in the language (cf. Berent et al., 1999, and Plunkett & Nakisa, 1997). Thus, in contrast to the statistical view of regularity in mapping graphemes to phonemes, regularity in the area of inflectional morphology is dissociable from frequency and is distinguishable from irregular inflection.

### Instance Specificity

A second distinction between the regularity of GPCs and many linguistic rules concerns instance specificity. At the core of linguistic rules are variables—abstract placeholders that can enumerate a large class of instances. For example, the variable *verb stem* enumerates instances such as *walk*, *like*, *refrigerate*, *congratulate*, etc. Linguistic rules, however, apply over variables, not the instances they enumerate. For example, the past tense rule concatenates a suffix to a *verb stem*, rather than to specific instances (e.g., *walk*). Because rules are blind to the identity of specific instances, they can treat all instances as equivalent, ignoring their idiosyncratic properties, such as their meaning or sound (Marcus, 1998b, 2001). The representation of variables not only determines the set of possible inputs to the rule, but may further shape its output. A rule's inputs and outputs exhibit various types of relationships. Some relationships are arbitrary, such as  $x \rightarrow y$ . Other relationships, however, are nonarbitrary. For instance, the past tense output is obtained by copying the stem ( $V_{\text{stem}} \rightarrow V_{\text{stem}} + \text{suffix}$ ). Given the regular base *like*, for example, the past tense form must include that base (e.g., *liked*). The rule's output thus has a constituent structure ( $V_{\text{stem}} + \text{suffix}$ ), and this structure is uniquely determined by the form of the input (Fodor & Pylyshyn, 1988).

In contrast to many linguistic rules, GPC outputs manifest no constituent structure. The relationship between the inputs and the outputs is arbitrary—there is nothing in that rule that prevents, in principle, the mapping of *b* to other outputs (e.g.,  $b \rightarrow /l/$ , instead of  $b \rightarrow /b/$ ). GPCs further differ from the inflection rule in that they link specific instances, rather than variables. For example, the correspondence  $b \rightarrow /b/$  associates two specific instances, the  $m \rightarrow /m/$  correspondence associates another pair of instances, and so forth. The instance specificity of GPCs is clearly evident even on their *symbolic* account. Most rules discussed in Rastle and Coltheart (1999) map specific graphemes to specific phonemes (e.g.,  $m \rightarrow /m/$ ). These rules clearly lack variables. Some rules, however, do include variables. For example, the rule  $[a.e \rightarrow /e/]$  maps the grapheme composed of the letter *a* followed by any letter (a variable captured by a dot) and the letter *e* to the phoneme  $/e/$ . Note, however, that, unlike the past tense rule, for instance, the relationship between the rule's input and its output is arbitrary; hence, the representation of the rule's input by a variable (*any letter*) does not determine the form of the output. Although variables play a rather minor role in such rules, their encoding may nevertheless result in incorrect, too strong predictions regarding human behavior. The incorporation of variables in GPCs tacitly assumes that readers treat all letters instantiated by the variable as equivalent. For example, the rule  $a.e \rightarrow /e/$  implies that the speed of mapping the grapheme *a.e* to its phoneme is similar across all possible intermediate letters (e.g., *abe*, *ace*, *ade*, *afe*, etc.). Coltheart and colleagues have not articulated this prediction. Indeed, this description of GPCs may not necessarily capture readers' knowledge. There is presently no evidence

that readers are insensitive to the identity of the middle letter in the *a.e* grapheme. Thus, there is no evidence *requiring* the representation of variables in GPCs. Conversely, there is ample experimental evidence against it. Numerous studies suggest that the correspondences between graphemes and phonemes are highly sensitive to the identity of neighboring spelling and sound units at various grain sizes (Andrews, 1982; Glushko, 1979; Jared, 1997; Jared, McRae, & Seidenberg, 1990; Rosson, 1985; Seidenberg, Waters, Barnes, & Tanenhaus, 1984; Taraban & McClelland, 1987; Treiman, Mullennix, Bijeljac-Babic, & Richmond-Welty, 1995). The expression of these correspondences by variables would fail to exhibit instance specificity. Thus, an adequate description of GPCs does not appear to require the encoding of variables. In fact, the representation of GPCs by variables is likely to impair a model's ability to capture human behavior.

### The Success of Pattern Associator Accounts of Reading

If the mapping of graphemes to phonemes was achieved by symbol manipulation, the process of phonology assembly should not have been adequately modeled by eliminative connectionist networks. Contrary to this prediction, connectionist accounts of word reading have gained considerable success in modeling the pronunciation of both words and nonwords. The model of Seidenberg and McClelland (1989a) captures the pronunciation of monosyllabic words, although its ability to account for nonword pronunciation is limited (Besner et al., 1990; Coltheart et al., 1993). However, the subsequent models of Plaut et al. (1996) and Harm and Seidenberg (1999) account quite well for both monosyllabic word and nonword pronunciation (but see Balota & Spieler, 1998; Besner, 1999; Plaut, 1999; Spieler & Balota, 1997).

The success of connectionism in the area of word reading goes beyond the accommodation of existing experimental findings. Perhaps the strongest support for this approach comes from the ability of connectionist models to generate new predictions, predictions that were not intuitively obvious or previously observed. One of the most exciting recent discoveries involving the design properties of associative systems is the feedback consistency phenomenon. Following the principle that dynamical systems organize through recurrent feedback, Stone, Van Orden, and colleagues (Stone, Vanhoy, & Van Orden, 1997; Van Orden & Goldinger, 1994) predicted that the time it takes for activation to settle in such systems should depend not only on the consistency of the correspondence between spelling and sound, but also on the backward correspondence between sound and spelling. Consider, for example, the identification of the word *kite*. *Kite* is highly consistent in the correspondence of its spelling to sound (forward consistency). However, *kite* is inconsistent in the correspondence between its sound and spelling (feedback consistency), since its rime may be spelled in various ways (e.g., *might*, *light*, *bite*). The dynamical model predicts that this inconsistency should impair word identification. This prediction has recently received ample ex-

perimental support in several languages, using various reading tasks (Stone et al., 1997; Ziegler, Montant, & Jacobs, 1997), as well as in the identification of spoken words (Ziegler & Ferrand, 1998).

### Conclusion

There are numerous challenges to the symbolic account of mapping graphemes to phonemes. The symbolic view offers no coherent notion of regularity, since the distinction between regular and irregular correspondences is unmotivated. In contrast to many linguistic rules, most GPC "rules" (e.g.,  $b \rightarrow /b/$ ) list associations between specific instances of graphemes and phonemes. Although some GPC rules include variables, there is no empirical evidence that requires their representation, and there is substantial evidence to the contrary—readers are highly sensitive to the strength of the association between specific grapheme and phoneme instances. If variables play no role in the mapping of graphemes to phonemes, then the distinction between viewing them as associations or rules is vacuous. The success of pattern associator accounts of mapping graphemes to phonemes is all but expected.

### Does the Success of Eliminative Connectionism in Modeling GPCs Indicate Its Potential to Account for Language?

In the domain of mapping graphemes to phonemes, there is clearly a success story to be told with regard to the performance of eliminative connectionism. The elimination of variables does not seem to impair the performance of word-reading models. In fact, the appeal to context-sensitive associations between specific graphemes and phonemes captures some crucial aspects of readers' knowledge. What is one to make of this success? Is there a general lesson to be learned from this case study regarding the viability of eliminative connectionism as an account of cognition?

For proponents of eliminative connectionism, the answer is clearly yes. The success of connectionist reading models supports the truth of their representational tenet—namely, the absence of mental variables and mental operations over variables. For instance, Van Orden, Pennington, and Stone (1990) note that

we believe that dual process theory fails because, in this case, the symbolic paradigm in which it is rooted fails to account for the subtleties of rule like performance. (p. 488)

Van Orden and colleagues' fierce critique goes far beyond the specific case of reading:

The present hypotheses differ fundamentally from the traditional symbolic hypotheses of dual-process theory. But ironically, perhaps, the present account regains for printed word identification a discernment that helped initiate modern, symbolic, psycholinguistic analyses. A primary force of N. Chomsky's (1959) historic criticism of *Verbal Behavior* (Skinner, 1957) was his demonstration that language behavior is not adequately explained by a theory

that allows only singular modifiable links between observable stimulus–response pairs. Key evidence against such a theory comes from observed classes of psycholinguistic performance that are remarkably similar, as though they were produced by rule. A slightly looser characterization of this argument is that psycholinguistic performance requires something other than case-by-case explanations. But, at least for printed word naming, an account of skilled rule like performance that relied on general, discrete, GPC rules eventually degenerated into word-by-word direct access.

The shift towards case-specific explanations is not unique to word identification, nor is it unique to psycholinguistics. . . . Symbolic accounts, which followed, in part, from N. Chomsky's (1959) critique of Skinner's (1957) case-specific account, have gradually come to resemble the object of that critique. Subsymbolic models produce rule like behavior; it emerges from the reflection of statistical regularities distributed across input–output pairs. This capacity promises to reverse the trend towards case-by-case explanations. (Van Orden et al., 1990, p. 514)

Van Orden et al. (1990) make a bold claim. In their view, the failure of symbolic accounts of word reading is directly linked to their misguided representational scheme—namely, the appeal to rules (generalizations that link variables), rather than to statistical correspondences among input–output pairs. They further consider the failure of the symbolic account of word reading to be indicative of the inadequacy of the symbolic view as an account of cognition, in general, and of language, in particular. Van Orden and colleagues present no evidence to support this conclusion. Although they marshal comprehensive evidence for a statistical explanation of GPCs, an analysis of symbolic accounts of language clearly falls beyond the goals of their paper. Their fierce critique nevertheless underscores the need for a broader evaluation of the representational scheme embodied in connectionist reading models, an evaluation that goes beyond the specific case of word reading. This challenge raises a fundamental problem. To assess the adequacy of an architecture across domains, one must compare the scope of generalizations in these domains. For this end, one must first define the scope of generalizations attained by the learner or the model in any given area. To link such generalizations to the model's representational scheme, it is further necessary to show how the scope of generalizations is determined by the model's representational choices. An assessment of eliminative connectionism requires a formal analysis of these issues.

### Representational Commitments and Generalizations

Does the ability of an architecture to generalize in a given domain (e.g., word reading) indicate its potential to capture generalizations in different domains (e.g., linguistic competence)? A recent analysis by Gary Marcus (Marcus, 1998a, 1998b, 2001) provides some important insights into this question. Marcus formally assesses the

scope of generalizations. He further demonstrates how the representational scheme adopted by a model determines the scope of generalizations it can attain.

Consider first the scope of generalizations. Marcus (1998b, 2001) defines generalizations relative to the representational space of training exemplars—the *training space*. He distinguishes between two types of generalizations. One type of generalization is confined to the training space. A second type of generalization exceeds the training space. The scope of generalizations relative to the training space crucially depends on the model's representational scheme. For instance, consider a model that is trained on the exemplars *rabbit* and *dog* and tested on the novel exemplars *cat* and *table*. Given a localist representation, either of these test exemplars falls outside the training space. Conversely, if the model uses a distributed representation, novel test items that are exhaustively encoded in terms of trained features (e.g., *cat*, encoded as +[animate] +[furry]) fall within the model's training space. In contrast, test items that include untrained feature values (e.g., *table*, which is [-animate]) or untrained features (e.g., +[furniture]) fall outside the training space. Note that these two types of generalization (within vs. outside the training space) may each concern novel instances. In particular, generalizations that fall within the training space may well extend to untrained exemplars represented by a novel combination of trained features.

Marcus (1998b, 2001) proceeds to demonstrate that a model's architecture constrains the scope of generalizations it can attain. Marcus specifically examines generalizations in functions that operate over variables. These functions (*universally quantified one-to-one mappings*, UQOTOM) assign each input a single output, and they apply in a uniform manner to all instances in a domain. For example, the identity function (e.g., an  $X \rightarrow X$ , where  $X$  stands for any token) generates a single, unique output, identical to the input. Because the output depends only on the input, not on past training, and because the class of possible inputs is infinite, the identity function generalizes outside any finite training space. Marcus shows that the representation of variables is critical for such generalizations. Networks implementing operations over variables can generalize UQOTOM to any new instance, including items that fall outside the training space. In contrast, pattern associators (i.e., any architecture that eliminates variables) are limited with regard to the class of functions they can acquire. Such models can adequately generalize UQOTOM to new items that fall within the training space. However, they cannot systematically extend generalizations to items that are outside the training space.

The limited ability of pattern associators to generalize does not necessarily diminish their adequacy as models of human cognition. Indeed, not all aspects of human cognition generalize outside the training space. Nevertheless, people do exhibit some generalizations that exceed their hypothesized training space. The identity mapping is a

case in point. People presented with the instances *a rose is a rose*, *a tulip is a tulip* will easily extend the identity function to a novel item (e.g., *a lilac is a \_\_\_*). Certain influential connectionist networks fail to do so, however. Marcus shows that feedforward and simple recurrent networks<sup>8</sup> can adequately perform the identity function on trained items. However, these networks cannot generalize to novel test items that fall outside the training space. For instance, a network trained on the items *rose* and *tulip* will not generalize the identity function to the new item, *lilac*. This property is not simply due to an inability to represent the test item: The failure to generalize persists even if the model can successfully represent the test item (e.g., *lilac*) in the context of a different function (e.g., *the bee sniffs the lilac*). The failure to generalize outside the training space also does not stem from the idiosyncrasies of any particular implementation. Marcus (1998a, pp. 166–167) presents a formal proof that the inability to generalize reflects a principled limitation of the learning algorithm: the independence between the learning of weight on connections from any input (or output) unit of other input (or output) units. As a result, such models are unable to acquire functions that link variables from training on their instances (for discussion and rebuttals, see Altmann & Dienes, 1999; Eimas, 1999; Marcus, 1999a, 1999b, 1999c; McClelland & Plaut, 1999; Negishi, 1999; Seidenberg & Elman, 1999).

This conclusion has important methodological implications. First, Marcus's work provides a definition for the scope of generalizations. This definition makes it possible to capture the scope of a generalization exhibited by humans or multilayer perceptrons at any given domain and compare generalizations across domains. Second, the link between the scope of generalizations and representational choices allows one to evaluate the potential of a model to generalize in a new domain. If the class of learnable generalizations (by the mind or its models) is constrained by the representation of variables, then one can infer the potential of a model to generalize in a novel domain by inspecting its architecture.

These tools come in handy for examining the potential of pattern associator models of word reading to account for language. Proponents of the pattern associator hypothesis sometimes claim that generalizations in these domains are essentially comparable in their scope; hence, the success of pattern associators in one of these domains is indicative of their ability to account for the other domain (e.g., Plaut et al., 1996; Van Orden et al., 1990). These claims are difficult to assess in the absence of a systematic definition of the scope of generalizations. Marcus's (1998b, 2001) taxonomy addresses this problem. GPCs are equivalent to linguistic generalizations if both exhibit similar relationships to their respective training spaces. If one of these generalizations exceeds the training space, pattern associators will be unable to account for this domain, despite their ability to account for the other.

The previous discussion of GPCs suggests a domain in which knowledge may be easily captured without appealing to variables. The mapping of graphemes to pho-

nemes is a matter of simple association. Because novel words may be broken down into a small number of familiar GPCs, their decoding may not require any generalization. In contemporary pattern associators, however, decoding a novel word may require generalization. These models have attempted to increase the context sensitivity of spelling-to-sound correspondences by including distinct banks of correspondences for various syllabic positions (e.g., Harm and Seidenberg, 1999; Plaut et al., 1996; Zorzi et al., 1998). Training such models on decoding a word-initial *b*, for instance, may not necessarily extend to a word-final *b*. Most existing models, however, are limited to monosyllabic words. Given the phonotactic constraints on phoneme combinations and the extensive training regimen, in which the model is exposed to thousands of words, it is not unlikely that test items fall within the model's training space.<sup>9</sup>

In view of the limited scope of generalizations required in mapping graphemes to phonemes, the success of eliminative connectionism in this domain does not indicate that the mind eliminates operations over variables, nor does it demonstrate that generalizations outside the training space may be acquired by eliminative connectionism. It simply reflects the fact that the acquisition of GPCs may not require such knowledge. In particular, the success of eliminative connectionism in the area of word pronunciation does not guarantee its ability to account for language. Given Marcus's (1998b, 2001) conclusions, eliminative connectionist models of word pronunciation are expected to account for language only if linguistic generalizations do not exceed the training space—that is, if linguistic representations eliminate variables. To assess the viability of eliminative connectionist models of word reading as an account of language, I next examine whether linguistic generalizations are limited to the training space. The work on inflectional morphology suggests that this is not the case (e.g., Berent, Pinker, & Shimron, 1999; Clahsen, 1999; Kim, Pinker, Prince, & Prasada, 1991; Marcus et al., 1995; Pinker, 1991, 1994, 1999; Prasada & Pinker, 1993; Ullman, 1999). I would like to corroborate this conclusion by evidence from a domain closer to reading aloud—namely, phonology.

### The Scope of Linguistic Phonological Generalizations

Of the various components of grammar, phonology provides a particularly interesting test for the viability of eliminative connectionist models of word reading as accounts of language. First, phonology is the output of reading aloud. Thus, the adequacy of word-reading models ultimately depends on their ability to provide a plausible account of phonological representations. Second, phonology seems amenable to an associative account (Dell, Juliano, & Govindjee, 1993). In contrast to the abstract nature of syntactic regularities, phonological regularities have some strong correlates with external phenomena, acoustics and articulation. Third, many phonological processes are sensi-

tive to similarity: Similar phonological elements typically undergo similar changes. Generalizing such processes to new items thus depends on the similarity between a novel item and trained instances, a property that is readily captured by pattern associators. If phonology were entirely reducible to token associations, phonological generalizations would be quite similar to those found in GPCs: Neither would require generalization outside the training space. The elimination of variables in both word reading and language would render connectionist reading models viable accounts of language.

The following discussion assesses the scope of phonological generalizations. I first examine two cases that do not appeal to variables. Generalizations in these areas are limited to the training space and are captured quite well by pattern associators. I then present a new case study. I show that generalizations in that area are inexplicable by the statistical structure of training exemplars. Such generalizations appeal to variables; hence, they may not be limited to the training space. I thus claim that the elimination of operations over variables may prevent connectionist models of phonology assembly from providing a complete account for phonology.

### Phonological Generalizations That Do Not Appeal to Variables

**Phonotactic constraints.** Languages constrain the co-occurrence of phonemes in the syllable. For instance, English allows for the cluster *lp* to occur in the coda (e.g., *help*), but not in the onset (e.g., *\*lpep*). Phonological theory accounts for this fact by the sonority sequencing principle: Onsets must rise in sonority toward the nucleus, and codas must fall in sonority from the nucleus (Kenstowicz, 1994). Sonority roughly correlates with the degree of stricture of an articulator: Vowels are more sonorous than liquids (e.g., *l*), which, in turn, are more sonorous than obstruents (e.g., *p*). According to this scale, the *lp* decreases in sonority; hence, it is acceptable only in the coda.

Although the sonority scale captures syllable phonotactics, this scale is little more than a description of the statistical co-occurrence among a small set of features. In particular, this statement does not appeal to variables. Given the conclusion of Marcus (1998a, 1998b, 2001), an account for phonotactics should fall well within the scope of pattern associators. The recent model of Harm and Seidenberg (1999) indeed captures the acquisition of such constraints in an attractor network. Their enterprise is particularly relevant to the present discussion, since this network is incorporated within a larger framework of a reading model. Harm and Seidenberg demonstrated that, when presented with partial input (e.g., *he\_p*), their phonological attractor network was able to accurately complete the missing segment, so that the responses were phonologically well formed. Likewise, an analysis of the networks' phonological attractors revealed the tendency to pull away from illegal feature combinations toward legal feature combinations. These achievements clearly reflect the acquisition of



some important aspects of English phonotactics. Harm and Seidenberg did not examine the ability of the network to generalize, since all test items were members of the training set. Given the limited problem space, the model's detailed featural representation, and the extensive set of training exemplars, novel test items are likely to fall within the training space of pattern associators. It is thus likely that the model may generalize adequately in this area.

**Place assimilation.** Another phenomenon that has been successfully modeled by eliminative connectionist networks is *place assimilation*. Phonological processes frequently assimilate the place of articulation of a given consonant with that of adjacent consonants. For instance, the phrase *in Brussels* is sometimes pronounced as *im Brussels* (Avery & Rice, 1989, p. 191). In this example, the coronal segment /n/, the target, assimilates the place of articulation of the adjacent labial /b/, the trigger. Coronal assimilation exhibits a marked asymmetry: Coronals are much more likely to adopt the place of articulation of adjacent noncoronals than vice versa. For instance, the coronal segment /n/ is likely to assimilate the place of a labial, but a labial is not likely to assimilate the place of its adjacent coronal (Avery & Rice, 1989; Kenstowicz, 1994). This asymmetry has been also observed in speech error data (Stemberger, 1991), as well as in experimental manipulations using the gating procedure (Gaskell & Marslen-Wilson, 1996; Lahiri & Marslen-Wilson, 1991).

Phonological theory accounts for assimilation as an interaction among adjacent phonological elements. Specifically, autosegmental phonology (e.g., Avery & Rice, 1989) views assimilation as the loss of a segment's feature (e.g., place of articulation) and its replacement by a feature of another segment. Because coronal segments are not specified for place of articulation, they are more susceptible to assimilation than are other segments (e.g., labials) that are specified for place of articulation. In contrast, articulatory phonology (e.g., Browman & Goldstein, 1989) accounts for place assimilation as an increase in the overlap among articulatory gestures, a process that is particularly likely to mask the perception of coronals, owing to the short duration of the tongue tip gesture (Surprenant & Goldstein, 1998). Although the precise interaction among coronals and neighboring segments is subject to debate, the interaction itself is rather limited. The number of targets, triggers, and outputs undergoing coronal assimilation is small, and these segments are easily defined by their specific phonological contents (their feature geometry or articulatory gestures) without appealing to variables. Pattern associators should capture such generalization quite well.

Gaskell, Hare, & Marslen-Wilson (1995) demonstrated that a simple recurrent network can successfully perform this task for English coronals. Gaskell et al. first trained their model on a corpus in which 50% of the word-final coronals assimilated the place of articulation of a following word-initial labial or velar segment. They next tested the model's ability to infer the underlying representation of words taken from the training set. These words were

presented in their surface form, so that the coronal was either assimilated (e.g., *reg car*, representing *red car*) or unassimilated (e.g., *red car*). Gaskell et al.'s model successfully abstracted the underlying coronal segment from its assimilated surface manifestation. Specifically, the model was likely to interpret labials and velars as the manifestation of an underlying coronal. Furthermore, the inference of an underlying coronal segment was sensitive to the viability of the surface form as the product of assimilation. Because assimilated targets acquire the place feature of a neighboring phoneme, any given "surface" phoneme may reflect assimilation only when it shares the place of articulation of its neighboring phoneme, the trigger. The model's performance was sensitive to the agreement of place between the "surface" segment and the trigger. For instance, the model was more likely to infer an underlying coronal in a viable context, including a surface velar followed by another velar (e.g., interpret *quike go* as *quite go*), as compared with a nonviable context, in which the assimilated velar was not followed by another velar (e.g., interpret *quike by* as *quite by*). Although the model was tested only on targets that were previously trained,<sup>10</sup> generalizations to untrained coronal segments (e.g., from /t/ to /d/) may be easily captured within the model's training space by encoding the coronal feature (a feature independently motivated by phonological analysis).<sup>11</sup> The assimilation of place of articulation may thus present a phenomenon that does not require generalizing outside the training space. The success of Gaskell et al.'s simulation is an important initial demonstration of the ability of connectionism to model a phonological process that is of great significance to phonological theory.

### Phonological Constraints That Appeal to Variables: Evidence From the Obligatory Contour Principle

The discussion so far suggests that eliminative connectionist models can account for significant phonological generalizations. Their success, however, may be largely due to the nature of the phenomena under investigation. Syllable phonotactics and coronal assimilation are easily captured as statistical co-occurrences among a small set of instances. Indeed, many linguistic accounts of these phenomena do not include variables;<sup>12</sup> hence, the elimination of variables should not diminish the explanatory power of pattern associators as compared with symbolic linguistic explanations. In fact, the sensitivity of pattern associators to statistical structure may present a significant advantage. Phonological competence, however, may not be limited to statistical constraints. Some phonological principles may appeal to variables; hence, they allow generalizing beyond the training space. I now turn to examine one such phenomenon—the co-occurrence restrictions on Hebrew roots. This case suggests that phonological competence may generalize in a manner that is fundamentally different from the generalizations characteristic of the mapping of graphemes to phonemes. Thus, the ability of pattern associators to account for some aspects of printed

word pronunciation is not indicative of their ability to account for phonological knowledge.

**The structure of Hebrew roots.** Hebrew words are formed by inserting a root in a word pattern. The root is a series of, typically, three consonants. The word pattern consists of placeholders for the root consonants. It also provides the vowels and affixes. For instance, the verb *bided* (*he isolated*) is formed by inserting the root *bdd* into the *\_i\_e\_* word pattern. Note that the root *bdd* contains a geminate, *dd*.<sup>13</sup> Gemination is quite common in Hebrew roots. However, as in other Semitic languages (Greenberg, 1950), Hebrew exhibits a marked asymmetry in the location of geminates: There are numerous roots whose form is *XYX* (e.g., *bdd*, *gll*, *dmm*, *kdd*, *gnn*, *smm*), but almost no roots of the form *XXY* (e.g., *\*bbd* or *\*ssm*). This asymmetry is explicable by the obligatory contour principle (OCP), one of the most fundamental principles of modern phonological theory.

The OCP (McCarthy, 1986) bans adjacent identical elements in the lexicon. Adjacency, however, does not necessarily reflect the linear position of segments in the word. Instead, adjacency is defined relative to phonological constituents. Autosegmental representations capture distinct phonological constituents on separate levels of representation, denoted *planes*. Root consonants in Semitic languages are a phonological constituent; hence, they are represented on a single plane, separated from intermediate vowels and affixes. Because the root is a phonological constituent, it is subject to the OCP. Specifically, the OCP prohibits adjacent identical consonants in the root. Geminates are thus banned from the root's underlying representation: The lexical representation of roots such as *bdd* is biconsonantal, *bd*. Geminates emerge during the process of word formation. Specifically, to form a word from *bd*, it must be aligned with the word pattern. The alignment proceeds from left to right, leaving an empty consonant slot on the right edge of the root (see Figure 1). This slot is next filled by the spreading of the root's final radical, resulting in a geminate. In what follows, I refer to this process as *reduplication* (see Everett & Berent, 1998; Gafos, 1998). Because geminates are obtained only by a

productive process of reduplication and because reduplication proceeds from left to right, geminates may emerge only at a root's final position, not at its initial position. McCarthy's (1986) account thus predicts that root-final gemination (e.g., *bdd*) should be well formed. In contrast, root-initial gemination (e.g., *bbd*) should be ill formed. This prediction agrees with the asymmetry in the distribution of Semitic roots.

The OCP presents an interesting test for the scope of phonological generalizations. The OCP constrains geminates in the root morpheme. Geminates, however, are defined by a variable, *XX*. The domain of the OCP, the root, corresponds to a second variable. Unlike the previously discussed constraints on phonotactics and place assimilation, the potential inputs to the OCP are not defined by their specific featural contents (e.g., coronals, labials, etc.), nor does the constraint concern their association with other (similar) instances. Instead, the OCP applies to *any* root consonant, and it circumscribes a formal relationship among inputs—identity. The OCP thus contains variables at its core. If mental representations eliminate variables, then the OCP should be unlearnable. Conversely, if speakers' mental machinery operates over variables, then they could represent the OCP. If so, speakers should be able to generalize this constraint to new items, and such generalizations may exceed the phonological space of Hebrew. In what follows, I review empirical evidence demonstrating that Hebrew speakers constrain the location of geminates in the root and generalize this constraint to novel items. The next section examines whether the representation of this constraint appeals to variables.

**Speakers' sensitivity to root structure.** One source of evidence for speakers' knowledge of root structure<sup>14</sup> comes from rating tasks (Berent, Everett, & Shimron, 2001; Berent & Shimron, 1997). In these tasks, participants rate the acceptability of novel words as potential Hebrew words. These words were derived from one of three types of novel roots. One type of root exhibited root initial gemination (e.g., *SSM*), a second type of root exhibited root-final gemination (e.g., *MSS*), and a third type of root had no gemination (e.g., *PSM*). If the OCP is ac-




Figure 1. The formation of root-final gemination from its underlying biconsonantal form. The left figure illustrates the alignment of the biconsonantal root *bd* with the word pattern *CVCVC*. Because the alignment proceeds from left to right, the rightmost consonant slot remains unfilled. The right figure describes the spreading of the phoneme into the free slot. The doubly linked phoneme surfaces as root-final geminate, *bdd*.

**Table 1**  
**An Illustration of the Materials Used**  
**in Berent and Shimron (1997)**

Root Type	Class 1	Class 2	Class 3
[ssm]	Si-SeM	maS-Si-Mim	hiS-ta-SaM-tem
[mss]	Mi-SeS	maM-Si-Sim	hit-Ma-SaS-tem
[psm]	Pi-SeM	maP-Si-Mim	hit-Pa-SaM-tem

tive, speakers should constrain the location of geminates in the root. Root-initial gemination should be ill formed; hence, words derived from roots with initial gemination (e.g., *SiSeM*<sup>15</sup>) should be considered less acceptable than words whose roots exhibit final gemination (e.g., *MiSeS*) or no gemination (*PiSeM*).

McCarthy’s (1986) account attributes the ill-formedness of *SSM*-type roots to the location of geminates in the root. An alternative explanation, however, may account for the rejection of such items in terms of the location of geminate bigrams in the word. For instance, the rejection of the word *SiSeM* may reflect the unacceptability of the letter combination *SiS* word-initially. To secure the conclusion that the constraint on the location of geminates concerns *root* structure, rather than merely *word* structure, we varied the location of geminates in the word by conjugating the roots in three different classes of word patterns (see Table 1). In the first class of word patterns, the root is not affixed. Consequently, the location of geminates in the root is predictable from its word structure. In contrast, in the second and third classes, the root was “sandwiched” between a prefix and a suffix. Thus, the location of geminates in the root cannot be inferred from their word position. For instance, the word *histasamtem* (from the root *SSM*) contains root-initial geminates, but the geminates are not word initial, nor are they adjacent in the word’s surface structure. If speakers are merely sensitive to the location of geminates in the word, then they should not reject words such as *hiStaSaMtem*, relative to their controls (e.g., *hitMaSaStem* and *hitPaSaMtem*). In contrast, if speakers are sensitive to root structure, then they should consider root-initial geminates unacceptable in each of the three word classes, regardless of their morphological structure. The results of the rating studies strongly support the representation of the root. Root-initial gemination was rated significantly lower than either root-final gemination or no-gemination controls (Berent, Marcus, Shimron, & Gafos, in press; Berent & Shimron, 1997). Importantly, this finding emerged in each of the word classes, regardless of the location of geminates in the word. Thus, the domain defining the constraint on geminates is clearly the root.

Converging evidence for the representation of root structure is provided by a production task (Berent, Everett, & Shimron, 2001). In this task, participants were given a novel test root consisting of only two consonants (e.g., *sm*). Their task was to conjugate the test root by analogy to an existing word (e.g., conjugate *sm* in analogy to *paʔal*<sup>16</sup>). To accomplish this task, participants must analyze the morphological structure of the exemplar, delete its root (e.g., *pʔl*), and replace it with the test

root. In doing so, however, participants are faced with an alignment problem: The word pattern calls for three root radicals, but the test root includes only two. We examined how participants solve the alignment problem. According to McCarthy (1986), participants routinely form trilateral roots by rightward reduplication of underlying bi-consonantal representations. If so, speakers should solve the alignment problem by reduplicating one of the root’s radicals. Furthermore, they should be more likely to reduplicate the root’s final radical (e.g., *sm*→*smm*) than its initial radical (e.g., *sm*→*ssm*). Our results strongly support this prediction. Responses exhibiting root-initial geminates were extremely rare (less than 1%) and were far less frequent than responses exhibiting root-final geminates (47% of the total responses). Importantly, the same asymmetry emerged regardless of the location of geminates in the word. In accordance with the rating results, the preference for root-final over root-initial gemination was maintained when the root was highly affixed.

The evidence for root structure in rating and production tasks cannot be attributed to metalinguistic strategies in the performance of off-line tasks. First, participants are typically unaware of the constraint on root structure and are unable to explain the unacceptability of words with root-initial geminates (Berent & Shimron, 1997). Second, the findings obtained in rating and production replicate in on-line tasks. For instance, nonwords whose roots manifest initial gemination are rejected faster than nonwords with root-final gemination controls in the lexical decision task (Berent, Shimron, & Vaknin, 2001). In fact, speakers encode the location of geminates even when reading clearly impairs task performance, as in the Stroop task (Berent, Bibi, & Tzelgov, 2000). The ill-formedness of roots with initial geminates reduces the Stroop interference: Non-color words whose roots manifest initial geminates (e.g., *ssm*) result in faster color naming, relative to words with root-final geminates (e.g., *smm*). As with the previous tasks, the constraint on root gemination emerged in the Stroop and lexical decision tasks across numerous positions of the root in the word. These results suggest that the computation of the root’s structure may be automatic.

**Does the constraint on root structure appeal to mental variables?** The experimental results indicate that speakers possess knowledge concerning the location of geminates in the root and that this knowledge extends to novel roots. McCarthy (1986) accounts for these findings by appealing to two variables. One variable is the root morpheme. A second variable is *any segment X*, a variable implicated in the representation of identity, *XX*. Although the findings reviewed so far are clearly compatible with this proposal, they do not require the representation of mental variables. Proponents of the pattern associator hypothesis may argue that a constraint on root structure is learnable from exposure to Hebrew words; hence, the constraint on root geminates may not be fundamentally distinct from the phonotactic and place assimilation constraints reviewed earlier. In the absence of computational models for the OCP, this claim cannot be ruled out. A closer inspection of the experimental findings, however,

suggests that this phenomenon presents a challenge to the pattern associator hypothesis.

To capture the behavior of Hebrew speakers, pattern associators must learn to constrain the location of identical consonants in the root without representing either the root or identity. To define the domain of the constraint on geminates, pattern associators must decompose root tokens from words, treat these root tokens as a class, and abstract the regularities among its members—a large set of dissimilar objects sharing no common orthographic, phonological, or semantic properties. The success of pattern associators on each of these tasks remains to be demonstrated.<sup>17</sup> Let us assume, however, that a pattern associator is presented with all instances of existing Hebrew roots. Could the model learn the constraint on the location of geminates and generalize it to novel roots in a manner that matches native speakers' performance?

On a cursory examination, the constraint on root structure appears easy to discern from the statistical structure of Hebrew roots. Because geminates are ill formed root-initially, the frequency of root-initial geminate bigrams (e.g., *ssm*) is typically lower than that of root-final bigrams (e.g., *mss*). The unacceptability of root-initial geminates could thus reflect the statistical structure of Hebrew roots, not specifically identity—a formal property defined over a variable. An inspection of the experimental findings, however, provides various indications that speakers are sensitive to the identity of geminates.

1. *The frequency of geminate versus addition responses.* One piece of evidence indicating the representation of identity comes from the production task. When participants produce a trilateral root from a biconsonantal input (e.g., *sm*), their responses normally exhibit a geminate at the end of the root (e.g., *smm*), but not in its beginning (e.g., *ssm*). The production of geminates is easily explained as reduplication: an operation that copies a variable and is blind to the content of specific segments. An alternative account, however, may attribute the formation of geminates to an associative process of segment addition. On this view, the formation of geminates (e.g., *sm* → *smm*) is indistinguishable from the formation of nongeminate response (e.g., *sm* → *sml*)—the rate of geminate responses and their location are merely a by-product of the statistical structure of the Hebrew language. To examine this explanation, we (Berent, Everett, & Shimron, 2001) calculated the expected rate of geminate versus nongeminate responses from the distribution of existing Hebrew roots. The predicted rate that the addition of a segment results in a root-final geminate versus nonidentical segment for our materials is 4.9%. In contrast, the observed rate of root-final geminates, relative to the addition of a new segment, is 329%. Geminate formation is thus inexplicable by bigram frequency.

2. *The effect of counterexamples.* Additional evidence against a statistical explanation concerns the effect of counterexamples to the OCP. Hebrew has a few roots whose structure violates the OCP, two of which are highly productive and familiar (*mmn*, financed; *mmsh*, realize).

Both of these roots begin with the geminate *mm*; hence, this geminate is quite frequent root-initially. If the avoidance of root-initial geminates was due to the rarity of these geminate bigrams, this pattern should be reduced for roots where the initial bigram is frequent. To test this prediction, Berent, Everett, & Shimron (2001) examined the production of trilateral roots from biconsonantal roots that are analogous to OCP violations (*mg*). These items exhibited a clear asymmetry in the production of geminates: 46% of the responses to roots analogous to counterexamples manifested root-final gemination, whereas fewer than 0.5% of the responses exhibited root-initial gemination (for comparison, roots that are not analogous to counterexamples yielded root-final vs. root-initial gemination in 47.8% and 0% of the responses, respectively).

3. *The distinction between geminates and nongeminate controls.* The findings from the production task demonstrate that the formation of geminates is inexplicable by their frequency in the language. Additional evidence from lexical decision demonstrates that speakers can discriminate between geminate and nongeminate bigrams even when they are closely matched on their statistical structure. If speakers form root-final geminates by reduplication, their presence may signal a grammatical operation and, hence, strengthen the appearance of such stimuli as “wordlike.” Supporting this prediction, nonwords generated from novel roots with final geminates (e.g., *smm*) were more difficult to reject in lexical decision, as compared with no-gemination controls (e.g., *psm*) that were matched for bigram frequency (Berent, Shimron, & Vaknin, 2001). These findings emerged in three different experiments that varied the location of geminates in the word (e.g., *SiSeM*, *SiSaSMti*, *hiStaSaMti*). These results suggest that speakers represent the formal structure of geminates—namely, their identity.

In summary, Hebrew speakers generalize the constraint on root structure to novel roots, and this generalization is robust with respect to the content of specific geminate bigrams and their frequency in the language. In contrast to humans, pattern associators are highly sensitive to statistical structure (e.g., Daugherty & Seidenberg, 1992; Plunkett & Marchman, 1993; Rumelhart & McClelland, 1986). The behavior of Hebrew speakers implicates a mental operation that applies to variables, rather than an association among specific instances. If the constraint on root structure is insensitive to the phonological properties of specific exemplars, it may generalize to novel phonemes, including novel phonological features. The findings of Berent et al. (in press) support this prediction. Marcus's (1998b, 2001) analysis suggests that such generalization should fall beyond the scope of pattern associators that eschew variables.

## Conclusions and Caveats

Connectionist accounts of reading aloud have gained considerable success. Proponents of the pattern associator hypothesis interpret this success as evidence for the adequacy of reading models as accounts of language. This

claim, however, assumes that the scope of the generalizations in these two domains is comparable and that neither requires generalizing outside the training space. An inspection of the mapping of graphemes to phonemes supports this conclusion. GPCs do not appeal to variables, and speakers are highly sensitive to their statistical structure. The previous discussion also identified two linguistic phonological phenomena that may be adequately captured by pattern associators—namely, syllable phonotactic constraints and coronal assimilation. The mapping of graphemes to phonemes thus resembles *some* phonological manifestations insofar as they exclude variables; hence, they do not require generalizing beyond the training space. Such phenomena are indeed modeled successfully by eliminative connectionism. Although certain phonological processes are associative in nature, other processes clearly appeal to variables. OCP effects are a case in point. The existing experimental evidence suggests that Hebrew speakers generalize the constraint on root structure in a manner that is independent of the statistical structure of the Hebrew language. In view of the known sensitivity of pattern associators to statistical regularities, such findings may not be easily captured by eliminative connectionism. Thus, despite the excellent performance of eliminative connectionist accounts of reading, their success in fully modeling linguistic competence is uncertain.

I would like to conclude my commentary by proposing a few caveats. My claim that some phonological generalizations appeal to variables is based on a single case study. The scope of generalizations in this area clearly requires further experimental and computational research. Furthermore, sorting generalizations as falling within or outside the training space may not be trivial. The scope of phonological generalizations depends on the hypothesized space of phonological features and must be independently motivated. At present, my claim that certain phonological generalizations extend beyond the training space (and beyond the reach of eliminative connectionism) remains inconclusive. No doubt, the boundary conditions on phonological generalizations may be difficult to establish. It is precisely these boundaries, however, that are crucial for understanding the nature of mental architecture. Although pattern associator accounts of word reading, phonotactics, and coronal assimilation should be commended for their success, these cases might not allow exploring the boundaries of phonological generalizations. Phonological generalizations whose linguistic formulation appeals to variables may be far more informative for examining these questions.

Another clarification concerns the notion of *eliminative connectionist* and *dual-route* models. I claim that eliminative connectionism may be unable to fully account for phonological competence. Eliminative connectionism, however, is but a small subset of connectionism; hence, its shortcomings should not be interpreted as criticism of connectionism as a whole. Likewise, one should not interpret the criticism of symbolic accounts of reading as an attack on dual-route models. My discussion of word pro-

nunciation was intentionally limited to a single aspect of dual-route models—namely, the assumption that GPCs are achieved by a symbolic process. However, there is nothing in the assumption of separate routes that requires the computation of GPCs by a symbolic process. Thus, the rejection of this particular assumption does not require the rejection of dual-route models. Dual-route models embody additional assumptions regarding the distinction between storage and computation and the segregation of the neural pathways implementing these functions, assumptions supported by independent empirical evidence. Their evaluation falls beyond the scope of this discussion.

The exclusive focus on GPCs also limits the evaluation of connectionist accounts of word pronunciation. The conclusion that GPCs are learnable by an associative process falls short of demonstrating that phonological representations assembled in word reading eliminate variables. The formation of adequate phonological representations may go far beyond the mapping of letters to sounds. Indeed, Plaut et al. (1996) clearly acknowledge that the success of their model largely stems from their representational scheme: the distinction between the units assigned to a segment depending on its syllabic position (onset, nucleus, or coda). Because the implementation of syllable structure is critical for this model's success, it cannot be considered as evidence for the absence of variables in the representation of printed words (see also Zorzi et al., 1998). In fact, the emergence of OCP effects in on-line reading (Berent et al., 2000; Berent, Shimron, & Vaknin, 2001) strongly supports the role of variables in assembled phonological representations.

More generally, reading ability may not be independent of linguistic competence (Berent & Perfetti, 1995; Perfetti, 1985). The success of current eliminative connectionist accounts of word reading indicates that phonology assembly, like some aspects of linguistic performance, is highly sensitive to the statistical structure of instances. Such sensitivity is the hallmark of an associative process. However, if phonology assembly is constrained by linguistic competence, an account of phonology assembly may be ultimately required to represent variables as well. The implementation of operations over variables may be necessary for a complete account for reading as a linguistic skill.

## REFERENCES

- ALEGRE, M., & GORDON, P. (1996). Red rats eater exposes recursion in children's word formation. *Cognition*, **60**, 65-82.
- ALTMANN, G., & DIENES, Z. (1999). Rule learning by seven-month-old infants and neural networks. *Science*, **284**, 977.
- ANDREWS, S. (1982). Phonological recoding: Is the regularity effect consistent? *Memory & Cognition*, **10**, 565-575.
- ANS, B., CARBONNEL, S., & VALDOIS, S. (1998). A connectionist multiple-trace memory model for polysyllabic word reading. *Psychological Review*, **105**, 678-723.
- ARCHANGELI, D., & LANGENDOEN, L. (1997). *Optimality theory*. Oxford: Blackwell.
- EVERY, P., & RICE, K. (1989). Segment structure and coronal underspecification. *Phonology*, **6**, 179-200.
- BALOTA, D., & SPIELER, D. (1998). The utility of item-level analyses in

- model evaluation: A reply to Seidenberg and Plaut. *Psychological Science*, **9**, 238-240.
- BERENT, I., BIBI, U., & TZELGOV, J. (2000, November). *The autonomous computation of linguistic structure in reading: Evidence from the Stroop task*. Paper presented at the 41st Annual Meeting of the Psychonomic Society, New Orleans.
- BERENT, I., EVERETT, D., & SHIMRON, J. (2001). Do phonological representations specify variables? Evidence from the obligatory contour principle. *Cognitive Psychology*, **42**, 1-60.
- BERENT, I., MARCUS, G., SHIMRON, J., & GAFOS, A. (in press). The scope of linguistic generalizations: Evidence from Hebrew word formation. *Cognition*.
- BERENT, I., & PERFETTI, C. A. (1995). A rose is a REEZ: The two cycles model of phonology assembly in reading English. *Psychological Review*, **102**, 146-184.
- BERENT, I., PINKER, S., & SHIMRON, J. (1999). Default nominal inflection in Hebrew: Evidence for mental variables. *Cognition*, **72**, 1-44.
- BERENT, I., & SHIMRON, J. (1997). The representation of Hebrew words: Evidence from the obligatory contour principle. *Cognition*, **64**, 39-72.
- BERENT, I., SHIMRON, J., & VAKNIN, V. (2001). Phonological constraints on reading: Evidence from the obligatory contour principle. *Journal of Memory & Language*, **44**, 644-665.
- BESNER, D. (1999). Basic processes in reading: Multiple routines in localist and connectionist models. In P. A. McMullen & R. M. Klein (Eds.), *Converging methods of understanding reading and dyslexia* (pp. 413-458). Cambridge, MA: MIT Press.
- BESNER, D., TWILLEY, L., McCANN, R., & SEERGOBIN, K. (1990). On the association between connectionism and data: Are a few words necessary? *Psychological Review*, **97**, 432-446.
- BROWMAN, C. P., & GOLDSTEIN, L. (1989). Articulatory gestures as phonological units. *Phonology*, **6**, 201-251.
- BULLINARIA, J. A. (1997). Modeling reading, spelling, and past tense learning with artificial neural networks. *Brain & Language*, **59**, 236-266.
- CHOMSKY, N., & HALLE, M. (1968). *The sound pattern of English*. New York: Harper & Row.
- CLAHSEN, H. (1999). Lexical entries and rules of language: A multidisciplinary study of German inflection. *Behavioral & Brain Sciences*, **22**, 991-1060.
- COLTHEART, M. (1978). Lexical access in simple reading tasks. In G. Underwood (Ed.), *Strategies of information processing* (pp. 151-216). London: Academic Press.
- COLTHEART, M., CURTIS, B., ATKINS, P., & HALLER, M. (1993). Models of reading aloud: Dual-route and parallel-distributed processing approaches. *Psychological Review*, **100**, 589-608.
- DAUGHERTY, K., & SEIDENBERG, M. (1992). Rules or connections? The past tense revisited. *Proceedings of the 14th Annual Conference of the Cognitive Science Society* (pp. 259-264). Hillsdale, NJ: Erlbaum.
- DELL, G., 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.
- EIMAS, P. (1999). Do infants learn grammar with algebra or statistics? *Science*, **284**, 433.
- ELMAN, J. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, **48**, 71-99.
- ELMAN, J., BATES, E., JOHNSON, M., KARMILOFF-SMITH, A., PARISI, D., & PLUNKETT, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.
- EVERETT, D. L., & BERENT, I. (1998). An experimental approach to the OCP: Evidence for violable identity constraints in Hebrew roots. *Rutgers Optimality Archive*. Available <http://rucss.rutgers.edu/roa.html> (ROA-235).
- FODOR, J., & PYLYSHYN, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, **28**, 3-71.
- FOWLER, C. A. (1993). Phonological and articulatory characteristics of spoken language. In G. Blasken, J. Dittman, H. Grimm, J. Marshall, & C.-W. Wallesch (Eds.), *Linguistic disorders and pathologies: An international handbook*. Berlin: Walter de Gruyter.
- GAFOS, D. (1998). Eliminating long-distance consonantal spreading. *Natural Language & Linguistic Theory*, **16**, 223.
- GAFOS, D. (1999). *The articulatory basis of locality in phonology*. New York: Garland.
- GASKELL, M., HARE, M., & MARSLLEN-WILSON, W. (1995). A connectionist model of phonological representation in speech perception. *Cognitive Science*, **19**, 407-439.
- GASKELL, M., & MARSLLEN-WILSON, W. (1996). Phonological variation and interference in lexical access. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 144-158.
- GLUSHKO, R. (1979). The organization and activation of orthographic knowledge in reading aloud. *Journal of Experimental Psychology: Human Perception & Performance*, **5**, 674-691.
- GORDON, P. (1985). Level ordering in lexical development. *Cognition*, **21**, 73-93.
- GREENBERG, J. (1950). The patterning of morphemes in Semitic. *Word*, **4**, 196-208.
- HAHN, U., & NAKISA, R. S. (2000). German inflection: Single route or dual route? *Cognitive Psychology*, **41**, 313-360.
- HAMMOND, M. (1997). Optimality theory of prosody. In D. Archangeli & L. Langendoen (Eds.), *Optimality theory* (pp. 33-58). Oxford: Blackwell.
- HAMMOND, M. (1999). *The phonology of English*. Oxford: Oxford University Press.
- HARE, M., & ELMAN, J. (1995). Learning and morphological change. *Cognition*, **56**, 61-98.
- HARM, M. W., & SEIDENBERG, M. S. (1999). Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review*, **106**, 491-528.
- HUMMEL, J. E., & HOLYOAK, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, **104**, 427-466.
- JACOBSON, R., FANT, C., & HALLE, M. (1963). *Preliminaries to speech analysis*. Cambridge, MA: MIT Press.
- JARED, D. (1997). Spelling-sound consistency affects the naming of high-frequency words. *Journal of Memory & Language*, **36**, 505-529.
- JARED, D., McRAE, K., & SEIDENBERG, M. (1990). The basis of consistency effects in word-naming. *Journal of Memory & Language*, **29**, 687-715.
- KENSTOWICZ, M. (1994). *Phonology in generative grammar*. Cambridge, MA: Blackwell.
- KIM, J., MARCUS, G., PINKER, S., HOLLANDER, M., & COPPOLA, M. (1994). Sensitivity of children's inflection to grammatical structure. *Journal of Child Language*, **21**, 179-209.
- KIM, J., PINKER, S., PRINCE, A., & PRASADA, S. (1991). Why no mere mortal has ever flown out to center field. *Cognitive Science*, **15**, 173-218.
- LAHIRI, A., & MARSLLEN-WILSON, W. (1991). The mental representation of lexical forms: A phonological approach to the recognition lexicon. *Cognition*, **38**, 245-294.
- MARCUS, G. F. (1998a). Can connectionism save constructivism? *Cognition*, **66**, 153-182.
- MARCUS, G. F. (1998b). Rethinking eliminative connectionism. *Cognitive Psychology*, **37**, 243-282.
- MARCUS, G. F. (1999a). Connectionism: With or without rules? *Trends in Cognitive Sciences*, **3**, 168-170.
- MARCUS, G. F. (1999b). Response: Do infants learn grammar with algebra or statistics? *Science*, **284**, 433.
- MARCUS, G. F. (1999c). Response: Technical comments—Rule learning by seven-month-old infants and neural networks. *Science*, **284**, 977.
- MARCUS, G. F. (2001). *The algebraic mind: Integrating connectionism and cognitive science*. Cambridge, MA: MIT Press.
- MARCUS, G. F., BRINKMANN, U., CLAHSEN, H., WIESE, R., & PINKER, S. (1995). German inflection: The exception that proves the rule. *Cognitive Psychology*, **29**, 189-256.
- MCCARTHY, J. (1986). OCP effects: Gemination and antigemination. *Linguistic Inquiry*, **17**, 207-263.
- MCCLELLAND, J. L., & PLAUT, D. C. (1999). Does generalization in infant learning implicate abstract algebra-like rules? *Trends in Cognitive Sciences*, **3**, 166-168.
- NEGISHI, M. (1999). Rule learning by seven-month-old infants and neural networks. *Science*, **284**, 875.
- PERFETTI, C. (1985). *Reading ability*. New York: Oxford University Press.
- PINKER, S. (1991). Rules of language. *Science*, **253**, 530-535.
- PINKER, S. (1994). *The language instinct*. New York: Morrow.

- PINKER, S. (1999). *Words and rules: The ingredients of language*. New York: Basic Books.
- PINKER, S., & PRINCE, A. (1988). On language and connectionism: Analysis of parallel distributed processing model of language acquisition. *Cognition*, **28**, 73-193.
- PLAUT, D. C. (1999). A connectionist approach to word reading and acquired dyslexia: Extension to sequential processing. *Cognitive Science*, **23**, 543-568.
- PLAUT, D. C., MCCLELLAND, J. L., SEIDENBERG, M. S., & PATTERSON, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, **103**, 56-115.
- PLUNKETT, K., & JUOLA, P. (1999). A connectionist model of English past tense and plural morphology. *Cognitive Science*, **23**, 463-490.
- PLUNKETT, K., & MARCHMAN, R. (1993). From rote learning to system building: Acquiring verb morphology in children and connectionist nets. *Cognition*, **48**, 21-69.
- PLUNKETT, K., & NAKISA, R. C. (1997). A connectionist model of the Arabic plural system. *Language & Cognitive Processes*, **12**, 807-836.
- PRASADA, S., & PINKER, S. (1993). Generalization of regular and irregular morphological patterns. *Language & Cognitive Processes*, **8**, 1-56.
- PRINCE, A., & SMOLENSKY, P. (1993). *Optimality theory: Constraint interaction in generative grammar* [(Tech. Rep. No. CU-CS-696-93). Boulder: University of Colorado, Department of Computer Science. (Tech. Rep. No. TR-2). New Brunswick, NJ: Rutgers University, Rutgers Center for Cognitive Science].
- PRINCE, A., & SMOLENSKY, P. (1997). Optimality: From neural networks to universal grammar. *Science*, **275**, 1604-1610.
- RASTLE, K., & COLTHEART, M. (1999). Serial and strategic effects in reading aloud. *Journal of Experimental Psychology: Human Perception & Performance*, **25**, 482-503.
- ROSSON, M. (1985). The interaction of pronunciation rules and lexical representations in reading aloud. *Memory & Cognition*, **13**, 90-99.
- RUECKL, J., MIKOLINSKI, M., RAVEH, M., MINER, C., & MARS, F. (1997). Morphological priming, fragment completion, and connectionist networks. *Journal of Memory & Language*, **36**, 382-405.
- RUMELHART, D. E., & MCCLELLAND, J. L. (1986). On learning the past tense of English verbs: Implicit rules or parallel distributed processing? In J. L. McClelland, D. E. Rumelhart, & The PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models* (pp. 216-271). Cambridge, MA: MIT Press, Bradford Books.
- RUSSEL, K. (1997). Optimality theory and morphology. In D. Archangeli & L. Langendoen (Eds.), *Optimality theory* (pp. 102-133). Oxford: Blackwell.
- SEIDENBERG, M. S. (1987). Sublexical structures in visual word recognition: Access units of orthographic redundancy? In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 245-263). Hillsdale, NJ: Erlbaum.
- SEIDENBERG, M. S. (1997). Language acquisition and use: Learning and applying probabilistic constraints. *Science*, **275**, 1599-1603.
- SEIDENBERG, M. S., & ELMAN, J. (1999). Do infants learn grammar with algebra or statistics? *Science*, **284**, 433.
- SEIDENBERG, M. S., HASKELL, T., & MACDONALD, M. (1999, November). *Constraints on plurals in compounds: Some implications of compounds research*. Paper presented at the 40th Annual Meeting of the Psychonomic Society, Los Angeles.
- SEIDENBERG, M. S., & MCCLELLAND, J. L. (1989a). A distributed developmental model of word recognition and naming. *Psychological Review*, **96**, 523-568.
- SEIDENBERG, M. S., & MCCLELLAND, J. L. (1989b). More words but still no lexicon: Reply to Besner et al. (1990). *Psychological Review*, **97**, 447-452.
- SEIDENBERG, M. S., & PLAUT, D. C. (1998). Evaluating word-reading models at the item level: Matching the grain of theory and data. *Psychological Science*, **9**, 234-237.
- SEIDENBERG, M. S., WATERS, G. S., BARNES, M. A., & TANENHAUS, M. K. (1984). When does irregular spelling or pronunciation influence word recognition? *Journal of Verbal Learning & Verbal Behavior*, **23**, 383-404.
- SHASTRI, L. (1999). Infants learning algebraic rules. *Science*, **285**, 1673-1674.
- SMOLENSKY, P. (1995). Constituent structure and explanation in an integrated connectionist/symbolic cognitive architecture. In C. MacDonald & G. MacDonald (Eds.), *Connectionism: Debates on psychological explanation*, (Vol. 2, pp. 221-290). Oxford: Blackwell.
- SMOLENSKY, P. (1999). Grammar-based connectionist approaches to language. *Cognitive Science*, **23**, 589-613.
- SPIELER, D., & BALOTA, D. (1997). Bringing computational models of word-naming down to the item level. *Cognitive Science*, **8**, 411-416.
- STEMBERGER, P. (1991). Apparent anti-frequency effects in language production: The addition bias and phonological underspecification. *Journal of Memory & Language*, **30**, 161-185.
- STONE, G. O., VANHOY, M., & VAN ORDEN, G. C. (1997). Perception is a two-way street: Feedforward and feedback phonology in visual word recognition. *Journal of Memory & Language*, **36**, 337-359.
- SURPRENANT, A. M., & GOLDSTEIN, L. (1998). The perception of speech gestures. *Journal of the Acoustical Society of America*, **104**, 518-529.
- TARABAN, R., & MCCLELLAND, J. L. (1987). Conspiracy effects in word pronunciation. *Journal of Memory & Language*, **26**, 608-631.
- TESAR, B., & SMOLENSKY, P. (2000). *Learnability in optimality theory*. Cambridge, MA: MIT Press.
- TREIMAN, R., MULLENNIX, J., BIJELJAC-BABIC, R., & RICHMOND-WELTY, D. (1995). The special role of rimes in the description, use, and acquisition of English orthography. *Journal of Experimental Psychology: General*, **124**, 107-136.
- ULLMAN, M. (1999). Acceptability ratings of regular and irregular past-tense forms: Evidence for a dual-system model of language from word frequency and phonological neighborhood effects. *Language & Cognitive Processes*, **14**, 15-67.
- VAN ORDEN, G. C., & GOLDINGER, S. D. (1994). Interdependence of form and function in cognitive systems explains perception of printed words. *Journal of Experimental Psychology: Human Perception & Performance*, **20**, 1269-1291.
- VAN ORDEN, G. C., PENNINGTON, B. F., & STONE, G. O. (1990). Word identification in reading and the promise of subsymbolic psycholinguistics. *Psychological Review*, **97**, 488-522.
- ZIEGLER, J. C., & FERRAND, L. (1998). Orthography shapes the perception of speech: The consistency effect in auditory word recognition. *Psychonomic Bulletin & Review*, **5**, 683-689.
- ZIEGLER, J. C., MONTANT, M., & JACOBS, A. M. (1997). The feedback consistency effect in lexical decision and naming. *Journal of Memory & Language*, **37**, 533-554.
- ZORZI, M., HOUGHTON, G., & BUTTERWORTH, B. (1998). Two routes or one in reading aloud? A connectionist dual-route model. *Journal of Experimental Psychology: Human Perception & Performance*, **24**, 1131-1161.

NOTES

1. Although printed words are linguistic objects, the knowledge governing their decoding is clearly distinct from the knowledge of language structure. The mastery of one's language obviously does not entail the acquisition of literacy. Likewise, the decoding of a printed word may not necessarily reflect the full intricacy of linguistic competence. Given that language and word reading reflect distinct competencies, one can assess the relationship between knowledge in these two domains.
2. One exception is the model of Bullinaria (1997). This localist feedforward network exhibits some learning and generalization in the areas of reading, spelling, and past tense inflection. This network, however, was tested only on a small subset of the relevant phenomena characteristic of human behavior in each of these domains. Furthermore, the network clearly exhibited some problems even in the areas on which it was tested (e.g., the absence of pseudo-homophone effects, the failure to distinguish between regular words and non-words, and the low activation for some of the nonwords); hence, I will not discuss this specific network any further.
3. Rules circumscribe the conversion of inputs to outputs. In contrast, constraints may specify the shape of outputs independently of their inputs (Hammond, 1997). For instance, one could express the devoiced realization of the English plural suffix in *cats* by a rule that converts the suffix *z* into *s* in the context of *t*. This rule prescribes a procedure for converting specific inputs (*z*) into a specific output (*s*). In contrast, a constraint approach directly spec-

ifies the form of the output by requiring that any adjacent consonants agree on voicing (Russett, 1997). This ban makes no assumptions about inputs.

4. Regular plurals are admissible in compounds under some highly restricted syntactic and semantic circumstances. For instance, children accept the phrase *red-rats-eater* only when they perceive the plural to be modified by the adjective (*an eater of red rats*, not *a red eater of rats*). Likewise, in the absence of an adjective (e.g., *publications catalog*), a regular plural is accepted when the plural indicates a heterogeneous set of objects (*many different publications*, not simply *many publications*). Alegre and Gordon (1996) account for these findings by proposing that regular plurals are accepted in a compound when they constitute a phrase—that is, when they are fed back to a compounding from syntax.

5. Hahn and Nakisa (2000) demonstrated that the regular inflection of noncanonical roots depends on their phonological and semantic properties. The possibility that phonological and semantic factors *contribute* to the interpretation of a word as lacking a canonical root, however, does not demonstrate that such factors are *sufficient* to account for the inflection of rootless forms. Although Hahn and Nakisa argue that a single-route model can capture human performance, their investigation examined only the classification of the plural onto various categories—Hahn and Nakisa did not demonstrate that the network can actually produce the correct plural form. Furthermore, the performance of such networks is contingent on the asymmetry in the distribution of regular and irregular forms. Berent, Pinker, and Shimron (1999), however, demonstrated that, unlike pattern associators, speakers do exhibit a regular default even when these conditions are violated. The ability of pattern associators to capture these findings is yet to be demonstrated.

6. Seidenberg, Haskell, and MacDonald (1999) proposed a connectionist account for such data. Their proposal attributes the ban on regular plurals in compounds to a semantic and phonological preference. In contrast to the semantic constraints discussed by Alegre and Gordon (1996)—constraints that supplement (not replace) the distinction between regular and irregular plurals—the proposal of Seidenberg et al. views regularity as an artifact of semantic and phonological confounds. As was indicated above, however, Gordon (1985) has controlled for these aspects of compounding. It is unclear how Seidenberg et al. (1999) account for these facts.

7. Plunkett and Juola (1999) claim to account for the inflection of words derived from a different grammatical category by adding some “pseudo-semantic” nodes to a pattern associator. Their description of the network’s performance is insufficient to fully evaluate this proposal. On a cursory inspection, however, the rate of regular inflection for irregular sounding words seems quite different from human performance. For instance, the network assigned regular inflection for verbs derived from nouns only on 26%, as compared with the 64% observed for preschoolers (Kim, Marcus, Pinker, Hollander, & Coppola, 1994). Furthermore, it is unclear whether these effects are specific to irregular inflection.

8. In the simulations using the feedforward network, “words” were captured by distributed representations, whereas, in the simple recurrent network simulations, every word was captured by a separate node. Marcus (1998b) shows, however, that the use of distributed representation is unlikely to alleviate this problem. First, the incorporation of distributed representation allows generalizing only for test items that contain trained features. Second, the use of distributed representations may impair the model’s ability to perform even on trained items, since this representation

cannot discriminate between exemplars that share features, a problem known as the *superposition catastrophe*.

9. In fact, Plaut et al. (1996) discuss in detail the inability of slot-based and context-sensitive representations to extend generalizations to untrained contexts. This problem, however, is unlikely to be evident in the performance of fully trained models.

10. Although all the target words undergoing assimilation were presented to the model during training, not all test targets were presented in their assimilated form during training. However, because the assimilation rule acquired by the model is limited to three coronal phonemes and five trigger phonemes, the critical aspects for assimilation in these items were clearly fully trained.

11. Gaskell et al.’s (1995) model encoded phonological segments in terms of 11 binary features from Jacobson, Fant, and Halle (1963). This representation did not explicitly encode the coronal feature. Thus, each of the coronal segments undergoing assimilation was encoded as a separate configuration of features.

12. The syllable, the domain of phonotactic constraints, is a variable. However, because syllables often correspond to word tokens, it is easy to describe the domain of such constraints without such variables.

13. I use the term *geminate* to refer to a single phoneme linked to two word-pattern slots. Modern Hebrew does not contrast phonemes in terms of length; hence, geminates in Hebrew always correspond to two distinct segments in trilateral roots (rather than to phonetically long segments).

14. Speakers’ sensitivity to the location of geminates in the root is clearly compatible with the OCP. I wish to clarify, however, that these results do not directly demonstrate identity avoidance in the lexicon. There are indeed several possible accounts for the asymmetry in the location of identical consonants in the root. The following analysis does not attempt to discriminate among them.

15. For viewing convenience, I indicate root consonants with uppercase letters. The Hebrew orthography, however, does not distinguish between root versus nonroot consonants; hence, this distinction was never present in the experimental materials. The examples of root-initial and root-final gemination (e.g. *SSM*, *MSS*, *SMM*) are used to assist readers who are familiar with the linguistic discussion of the OCP, where these example are frequently employed. The roots for *MSS* and *SMM*, however, correspond to existing Hebrew roots. The roots employed in the experiments reported here were all novel roots.

16. The ? stands for a glottal stop consonant.

17. Elman (1993) claims that his simple recurrent network has learned the distinction between abstract lexical categories (e.g., nouns and verbs) from their syntactic roles. This finding, however, is not directly relevant to the case of Hebrew roots. First, unlike the noun–verb distinction, where different types of nouns often share meaning, roots share no semantic features. Second, it is doubtful that the network acquired a categorical distinction between nouns and verbs. Such a distinction would require that all members of a category (e.g., all nouns) be represented alike. The network does not assign a uniform representation to all nouns or verbs, since some verb instances appear more “verblike” than others (e.g., see Elman, 1993, Figure 3).

(Manuscript received July 29, 1999;  
revision accepted for publication March 16, 2001.)