ELSEVIER

Cognition 83 (2002) 113–139

COGNITION

www.elsevier.com/locate/cognit

The scope of linguistic generalizations: evidence from Hebrew word formation

Iris Berent[a,*], Gary F. Marcus[b], Joseph Shimron[c], Adamantios I. Gafos[d]

[a]*Department of Psychology, Florida Atlantic University, 777 Glades Road, Boca Raton, FL 33431-0991, USA*
[b]*Department of Psychology, New York University, New York, NY, USA*
[c]*School of Education, University of Haifa, Haifa, Israel*
[d]*Department of Linguistics, New York University, New York, NY, USA*

Received 6 March 2001; received in revised form 6 September 2001; accepted 14 November 2001

Abstract

Does the productive use of language stem from the manipulation of mental variables (e.g. "noun", "any consonant")? If linguistic constraints appeal to variables, rather than instances (e.g. "*dog*", "*m*"), then they should generalize to *any* representable novel instance, including instances that fall beyond the phonological space of a language. We test this prediction by investigating a constraint on the structure of Hebrew roots. Hebrew frequently exhibits geminates (e.g. *ss*) in its roots, but it strictly constraints their location: geminates are frequent at the end of the root (e.g. *mss*), but rare at its beginning (e.g. *ssm*). Symbolic accounts capture the ban on root-initial geminates as *XXY, where X and Y are variables that stand for any two distinct consonants. If the constraint on root structure appeals to the identity of abstract variables, then speakers should be able to extend it to root geminates with foreign phonemes, including phonemes with foreign feature values. We present findings from three experiments supporting this prediction. These results suggest that a complete account of linguistic processing must incorporate mechanisms for generalization outside the representational space of trained items. Mentally-represented variables would allow speakers to make such generalizations.

Keywords: Language; Linguistic generalizations; Hebrew word formation

1. Introduction

Productivity is at the core of linguistic competence (Chomsky, 1980): speakers routinely produce and comprehend numerous sentences they have never heard before.

* Corresponding author. Fax: +1-561-297-2160.
E-mail address: iberent@fau.edu (I. Berent).

PII: S0010-0277(01)00167-6

Symbolic accounts attribute the productive use of language to the representation of mental variables (Fodor & Pylyshyn, 1988; Marcus, 1998, 2001; Pinker & Prince, 1988). Variables (e.g. *verb*, *nouns*, *morpheme*, *syllable*) are abstract place holders. For instance, the variable *noun* can stand for any sound combination, including both familiar (e.g. *dog*) as well as novel instances (e.g. *blick*). According to the symbolic view, linguistic processes operate over variables. For instance, the plural inflection rule in English concatenates the suffix *s* to the variable *noun stem* (Marcus, 1998; Pinker, 1999; Pinker & Prince, 1988). Although specific symbolic accounts of language differ on whether linguistic competence is captured in terms of inviolable rules (e.g. Chomsky, 1957), or violable constraints (e.g. Prince & Smolensky, 1997), all these proposals consider linguistic processes as operations over variables (Berent, 2001; Marcus, 2001). The representation of variables is indeed crucial for explaining productivity. Because mental operations appeal to variables (e.g. *noun stem*), their application is robust with respect to the idiosyncrasies of specific tokens. The symbolic hypothesis thus leads to a strong prediction regarding the scope of linguistic generalizations: if the mind encodes variables, then speakers should be able to extend linguistic generalizations across the board to *any* representable novel token, regardless of their familiarity with this instance or its features.

The symbolic hypothesis, however, has been subject to continuous challenge from connectionist accounts (e.g. Elman, 1993; Elman et al., 1996; Plaut, McClelland, Seidenberg, & Patterson, 1996; Rumelhart & McClelland, 1986; Seidenberg & McClelland, 1989). Such models, dubbed "eliminative connectionism" (Pinker & Prince, 1988), account for language without representing operations over variables. For instance, there are numerous models of inflection that eliminate the representation of a "*noun stem*" or "*verb stem*" (e.g. Daugherty & Seidenberg, 1992; Hahn & Nakisa, 2000; Plunkett & Juola, 1999; Plunkett & Marchman, 1993; Plunkett & Nakisa, 1997; Rueckl, Mikolinski, Raveh, Miner, & Mars, 1997; Rumelhart & McClelland, 1986). Despite the absence of variables, these models can generalize to novel tokens. For instance, the novel noun *blick* could be inflected by analogy to familiar nouns (e.g. *brick*, *block*). Such generalizations appear to challenge the symbolic hypothesis. If connectionist models can adequately generalize without appealing to variables, then there is reason to doubt that mental representations encode them. The symbolic hypothesis, however, makes a rather specific prediction regarding the scope of generalizations. The symbolic account does not merely state that people can generalize. It specifically predicts that generalizations can apply to *any* new representable instance, regardless of its similarity to trained items.

Although it is well known that many connectionist models can generalize, the scope of such generalizations is rarely explored in a systematic fashion. Recent work by Marcus (1998, 2001) makes it clear that the ability to generalize to *some* instances does not guarantee that a model can generalize to *all* novel instances. Of course, generalization to novel instances is always limited by a model's representational scheme. A model that cannot represent the phoneme "*d*", for instance, naturally cannot generalize to this phoneme. The more interesting question is whether a given model can generalize to all or only some of the novel items that it can represent. Consider, for instance, the novel noun *xick* (*/xIk/*, where the first phoneme is a velar fricative). The initial phoneme is not part of the English inventory, but it is familiar to speakers from borrowings (e.g. *chutzpah*, *chanukka*), and hence it is clearly representable. In fact, speakers can easily generalize regular inflection to strange-sounding words (Berent, Pinker, & Shimron, 1999; Prasada & Pinker, 1993). Some connectionist networks, however, fail to generalize the inflection rule to such items (Prasada & Pinker, 1993; for discussion see also Marcus, 2001).[1] Such failures are not simply due to an inability to generalize to novel items, as the same models can generalize to items that resemble familiar trained regular instances (e.g. the novel noun *blick*; similar to the regular nouns *brick* and *block*).

In order to capture the distinction between novel items like *xick* and novel items like *blick* Marcus (1998, 2001) proposed the notion of a *training space*. A training space is the space of feature values that is used in representing training instances. Trained instances that can be exhaustively described in terms of trained feature values fall within the training space. For instance, for an inflection model that represents instances by means of English phonemes, *blick* falls within the training space. In contrast, *xick* falls outside the model's training space (since no regular noun shares its initial phoneme). Marcus showed that the generalization ability of one popular class of models, multilayer perceptrons that lack operations over variables, depends on the position of a novel item relative to the model's training space. His investigation specifically concerned a class of functions that universally maps any input onto a single unique output. The English plural rule is an instance of this class, since it uniquely maps any noun stem to the plural form (*Nounstem* → *Nounstem* + *s*). Marcus formally demonstrated that multilayer perceptrons that lack operations over variables can adequately extend such functions to instances falling within the training space (e.g. *blick*), but they fail to generalize to items falling outside the training space (e.g. *xick*).

These findings make it clear that models cannot simply be evaluated on the basis of whether they can generalize to new instances; it is also important to examine whether a given model can generalize outside of its training space – the ability of a model to extend generalizations within the training space does not guarantee generalizations outside it. But Marcus's investigations also raise another important point: failures to exceed the training space are not necessarily a liability in a model of natural language. Although the "within" vs. "outside" of training space distinction gives us an important yardstick for evaluating models, that yardstick must be compared with human behavior. If linguistic generalizations were limited to the training space then the failure of a model (e.g. certain multilayer perceptrons) to exceed the training space would be considered a virtue, and offer support to the view that linguistic mental representations eliminate variables. Conversely, if the symbolic hypothesis is correct to predict that some linguistic generalizations exceed the training space, then accounts of linguistic competence would appear to require operations

[1] Eddington (2000) demonstrated that non-connectionist single-route systems (Analogical Modeling of Language and the Tilburg Memory Based Learner) can generalize past tense inflection to novel forms in a manner that matches human behavior (Prasada & Pinker, 1993). The systems studied by Eddington, however, only *classify* an item into two predetermined categories (regular vs. irregular) – they do not *generate* the past tense form. This seemingly minor procedural modification is significant, since it overrides the problem of copying the stem (Marcus, 2001). Because the system is not required to operate over all parts of the stem, it may simply ignore phonological idiosyncrasies, basing the classification on the remaining, familiar portions. In contrast, a system that produces the inflected form must "figure out" that all portions of the stem of English regular verbs must be copied. If the system has not experienced instances where a given phonologically idiosyncratic string is copied, it is unlikely to consistently exhibit such behavior.

over variables. It is therefore crucial to investigate empirically whether people extend linguistic generalizations outside the training space.

Research on artificial grammar learning suggests that adult learners (e.g. Reber, 1969) and infants (e.g. Gomez & Gerken, 1999) can generalize knowledge acquired from training items generated by a finite state grammar to a new set of test items. Such generalizations have been observed even when the test items share no features with the training items. For instance, Altamann, Dienes, and Goode (1995) demonstrated that adult learners trained on letter sequences transfer their knowledge to tone sequences. Likewise, Marcus, Vijayan, Bandi Rao, and Vishton (1999) demonstrated that 7-month-old infants, trained on three-syllable "sentences" such as ABB, can discriminate between novel test sentences that are consistent with the ABB grammar and those that are inconsistent (e.g. AAB). Because the similarity of test items to consistent training exemplars could not be explained in terms of shared phonemes or phonetic features, these findings suggested that learners can generalize the artificial grammar beyond the training space. Artificial languages, however, may differ in some important respects from natural language, and hence the scope of generalizations observed with an artificial grammar may not be indicative of linguistic competence.[2] Here, we examine whether generalization outside the training space is a property of natural language. We address this question by investigating a phonological constraint. If linguistic generalizations exceed the training space, then speakers may be able to extend this constraint to novel phonemes, including phonemes with novel phonemic feature values.

1.1. Identity avoidance in Hebrew roots

The constraint we examine is the Obligatory Contour Principle (OCP, McCarthy, 1979, 1986), a universal phonological constraint on identical elements in phonological representations. The OCP constrains identical phonemes, features and tones, and its effects have been documented in a variety of languages (e.g. Leben, 1973; McCarthy, 1979, 1986; Yip, 1988; for reviews, see Goldsmith, 1990; Kenstowicz, 1994). The question of whether OCP effects generalize outside speakers' training space thus illuminates a central aspect of phonological competence. OCP effects have been subject to intensive scrutiny in Semitic languages, and hence our investigation assesses the OCP in Hebrew.

Hebrew words include two ingredients: a root and a word pattern. The root is an abstract sequence of (typically) three consonants that carries the core meaning of the word. For instance, the root *lmd* is associated with learning. The word pattern is a prosodic template that contains placeholders for the root consonants and it specifies vowels and affixes. Hebrew words are formed by inserting the root into the word pattern. For instance, the insertion of the root *lmd* into the word pattern CiCeC (where C stands for any consonant) yields the verb *limed*, he taught. The same word pattern also yields the verbs *tipes* (he climbed), *diber* (he talked) and *shiker* (he lied) from the roots *tps*, *dbr* and *shkr*, respec-

[2] To our knowledge, the only other line of research that addresses the scope of natural-language generalizations concerns generalization of the inflection rule to strange-sounding words (Berent et al., 1999; Bybee & Moder, 1983; Prasada & Pinker, 1993). Although speakers' ability to inflect strange-sounding words is consistent with the view of such generalizations as exceeding the training space, this earlier research did not directly assess whether these words fall beyond the phonological space of the language.

tively. Conversely, a single root (e.g. lmd) may be inserted in various word patterns (e.g. CaCaC, CiCuC, hitCaCeC) resulting in a family of morphologically related words (e.g. *lamad*, he studied; *limud*, study; *hitlamed*, he taught himself, respectively).

Unlike English, for instance, the morphological constituents of Hebrew words are not concatenated in a linear fashion, but, are instead, intercalated. For instance, in the verb *limed*, the consonants *lmd* are members of a single morpheme, whereas the intermediate vowels *i* and *e* form a second morpheme. The linguistic theory of autosegmental phonology (e.g. Goldsmith, 1990) nicely captures this fact by representing the root on a separate level of representation segregated from vowels and affixes. Fig. 1 illustrates this representation for the verb *limed*. This representation includes three levels of representation, one for the root consonants, another for vowels, and a third level for consonant and vowel placeholders (the CV skeleton), serving as an anchor for root and vowel phonemes. This geometric arrangement constitutes a psychological hypothesis: phonological elements that are represented on the same level (an autosegment) are psychologically adjacent, and hence they are free to interact with each other. Because elements on a single level form a mental constituent, they are also subject to phonological constraints that govern their co-occurrence. Our interest here concerns a constraint on identical consonants.

Many Hebrew roots include adjacent identical consonants. For instance, the root *smm* includes the identical consonants *mm*. In what follows, we refer to adjacent identical root consonants as geminates. It should be kept in mind, however, that root geminates correspond to two distinct segments (rather than a phonetically long segment), and they are typically separated by an intermediate vowel (e.g. *simem*, from the root *smm*). Geminates are indeed quite common in Hebrew roots. Their location, however, is strictly constrained. Geminates are frequent at the end of the root (e.g. *bdd*, *sll*, *grr*, *xdd*, *smm*), but are rare at its beginning (e.g. **ssm*). Why should geminates be more common at the end of roots rather than at their beginnings?

McCarthy (1979, 1986) has provided an intricate but satisfying explanation for this asymmetry. His explanation starts with a linguistic constraint known as the OCP. The OCP is a constraint that bans adjacent identical elements in phonological representations. Because root consonants are represented on a single level, they are psychologically adjacent. If root geminates were represented by two distinct phonemes, they would violate the OCP. McCarthy thus proposes to capture root geminates by a single phoneme. For instance, a root like *smm* is actually stored as a biconsonantal representation, *sm* (see Fig. 2). The geminates are due to the manner of associating that single phoneme with the word pattern during word formation. To form a word, the root must be associated with the three root-consonant slots in the word pattern in a left to right order. For a biconsonantal

Fig. 1. The representation of the verb LiMeD (*he taught*). The representation includes three levels: a skeleton, including an abstract sequence of placeholders for consonants and vowels, the root consonants, and the vowels.

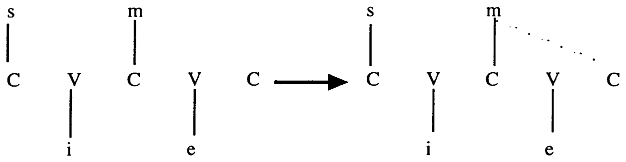

Fig. 2. The formation of root-final gemination from its underlying biconsonantal form. The left figure illustrates the alignment of the biconsonantal root **sm** with the skeleton CVCVC. Because the alignment proceeds from left to right, the rightmost consonant slot remains unfilled. The right figure describes the spreading of the phoneme into the free slot. The doubly linked phoneme surfaces as root-final geminate, **smm**.

root, such as *sm*, this process leaves an empty slot at the root's right edge. This slot is next filled by the spreading of the last consonant rightwards. The result, the double linking of a single phoneme (e.g. *m*) with two slots, surfaces as a geminate *mm*. The attribution of root geminates to a productive process of spreading makes a specific prediction regarding the location of geminates. Because root geminates are formed by spreading,[3] and because spreading proceeds rightwards (e.g. *sm* → *smm*), geminates may emerge at the end of the root, but not in its beginning, a prediction that is supported by the distribution of geminates in triliteral Hebrew roots.

At the core of McCarthy's account is the assumption that speakers constrain the identity of phonological elements. Identity, however, is defined in reference to a variable. For instance, using X and Y as variables standing for any two distinct consonants, one could capture the constraint on root structure as *XXY (the asterisk indicates ill-formedness). This formulation renders any two identical consonants ill-formed root-initially. While McCarthy's account does seem to fit the facts, one might wonder whether there could be an alternative associative account that could capture the root structure constraint in terms of a statistical distribution of specific instances. For example, speakers may notice that certain consonant bigrams, including *bb*, *gg*, *dd*, are frequent at the end of the root, but not in its beginning. Such statistical knowledge renders roots with unfamiliar bigrams root-initially (e.g. *gg*, *dd*, *ff*) ill-formed. Note that this account does not specifically ban identity in the initial bigram (indeed, in the absence of variables, identity is unrepresentable) – the ill-formedness of root-initial geminates is merely a byproduct of their rarity.

Both the symbolic and statistical formulations of the constraint on root structure would generalize to novel roots – but they would differ in their predictions about the scope of generalizations. The symbolic account predicts that generalization should extend to *any* representable novel instance, regardless of the familiarity with this instance or its phonological features. In contrast, generalizations predicted by the associative statistical account should be sensitive to the properties of specific root instances.

Our experimental investigation of OCP effects in Hebrew demonstrates that speakers extend the constraint on root structure to novel roots (Berent, Bibi, & Tzelgov, 2000; Berent, Everett, & Shimron, 2001; Berent & Shimron, 1997; Berent, Shimron, & Vaknin,

[3] Alternative accounts for the constraint on Semitic geminates, formulated within the framework of Optimality Theory, attribute the formation of geminates to reduplication (Everett & Berent, 1998; Gafos, 1998). These differences are inconsequential for the present discussion, as all accounts share the assumption that geminates are obtained by the copying of a variable.

2001). The existing findings, however, are inconsistent with a statistical explanation. First, the production of geminates is insensitive to the frequency of specific geminate bigrams in the language. For instance, in a production task in which participants were encouraged to form triliteral roots from their biconsonantal representations (e.g. form a triliteral root from the biconsonantal root *sm*), speakers overwhelmingly chose to reduplicate the root's final consonant (e.g. *sm* → *smm*) rather than to add a consonant (e.g. *sm* → *sml*) even though the expected frequency of geminate responses was far lower than the expected frequency of addition responses (Berent, Everett, & Shimron, 2001). A second challenge to a statistical account is the sensitivity of speakers to the formal structure of geminates, i.e. their identity. For instance, participants in lexical decision experiments discriminate between roots with final geminates (e.g. *smm*) and non-geminate controls (e.g. *psm*) despite their equation for their bigram frequency (Berent, Shimron, & Vaknin, 2001). Given these challenges to a statistical explanation, the distinction between these two root types suggests that speakers are sensitive to aspects of their structure, namely, the presence of identity and its location in the root.

Although speakers' insensitivity to the statistical structure of test items lends no support for an associative account, proponents of this view could attribute such null findings to a ceiling effect: an insensitivity to fine-grained differences in the degree of training on test phonemes that are all familiar (e.g. *smm* vs. *sml*), rather than the representation of variables (e.g. XYY vs. XYZ). Although this explanation is countered by numerous cases demonstrating the sensitivity of associative systems to the statistical structure of language (e.g. Plaut et al., 1996; Plunkett & Marchman, 1993; Plunkett & Nakisa, 1997; Rhode & Plaut, 1999; Rumelhart & McClelland, 1986; Seidenberg & McClelland, 1989), our previous results cannot rule it out. A clearer test for discriminating the associative and symbolic accounts concerns generalizations to instances that lack *any* relevant training – those that fall outside their training space. Here, we examine whether speakers generalize the constraint on root structure beyond the phonological space of Hebrew by investigating geminates with foreign phonemes and foreign feature values. Hebrew has no roots with foreign phoneme geminates. A statistical knowledge concerning the co-occurrence of Hebrew phonemes thus provides no information regarding the location of such geminates. If speakers' knowledge regarding root structure represents identity, however, then they should be able to extend it regardless of the contents of geminate phonemes.

1.2. Does the constraint on root structure generalize beyond the phonological space of Hebrew?

In these experiments, we examined generalization to phonemes that are borrowed from English. We chose English as the source of the borrowing since participants in the experiments, University students, are fluent in English, and hence they can easily represent these phonemes. The inventory of Hebrew phonemes is presented in Table 1 below. Although most English phonemes overlap with the Hebrew inventory, four English phonemes are not shared with Hebrew, and are sufficiently distant from the Hebrew inventory to assure that they are not perceived as allophones of Hebrew phonemes. These are the phonemes *j* (as in jeep), *ch* (as in chair), *th* (as in thick) and *w* (as in wide). Although University students can represent these English phonemes, we

Table 1
The inventory of Hebrew phonemes[a]

	Bilabial	Labiodental	Alveolar	Alveopalatal	Palatal	Velar	Uvular	Guttural
Stop (oral)	p, b		t, d			k, g		ʔ
Nasal (stop)	m		n					
Fricative		f, v	s, z	ʃ		x		
Affricate			ç					
Glide					y			h
Liquid			l				R	

[a] Phonemes are transcribed using the IPA (ʃ is the initial phoneme in *ship*, ç is the onset of *tsar*; *x* is the onset of *chutzpa*). The classification of phonemes into places of articulation is based on traditional phonetic analyses of place of articulation as described in Fromkin and Rodman (1993) and Kenstowicz (1994). The realization of the phoneme *R* is subject to wide inter-speaker variability. The classification proposed here is based on Ashkenazic pronunciation of Hebrew.

would not expect them to have used geminates including such phonemes. Our question here – a test of extra-training space generalization – is whether speakers generalize the root structure constraint to the non-native phonemes.

To provide the strictest test of extra-training space generalization, it is desirable that our test items not only include novel phonemes, but that at least some of those phonemes include *feature* values that are not customarily found in Hebrew. Although none of the English phonemes used in our experiments is an allophone of a Hebrew phoneme, some of these phonemes may be perceived as a composite of Hebrew phonemes or feature values. The palatal affricates *ch* and *j*, for instance, may be perceived as a composite of the palatal place of articulation, the affricate manner and voicing, features that exist in Hebrew. Likewise, *w* may be represented as a velar voiced glide composite. In contrast, *th*, an interdental (for American speakers, see Ladefoged & Maddieson, 1996, p. 143), may not be exhaustively described by a set of Hebrew phonemes. Hebrew has no interdentals, and hence the place of articulation of *th* is absent in the space of Hebrew features. Place of articulation, however, is not sufficiently precise to capture phonological distinctions (e.g. Kenstowicz, 1994). Indeed, place of articulation is subject to wide cross-speaker and cross-linguistic variability, a problem that is particularly noticeable for coronal fricative consonants (fricatives produced by the tongue tip, e.g. *s*, *th*, *sh*; see Dart, 1991; Gafos, 1999; Ladefoged & Maddieson, 1996). A far more accurate, speaker-invariant distinction between coronal fricatives is obtained by examining the area of the cross-sectional channel between the tongue and the palate. This phonetic dimension is formalized in terms of a distinctive feature that contrasts coronal fricatives, the Tongue Tip Constriction Area (TTCA) (Gafos, 1999). The TTCA feature has three possible values: narrow, mid, and wide. The Hebrew coronal fricatives utilize two of these values, narrow (*s*, *z*, and *ts*) and mid (*sh*), whereas the English phoneme *th* utilizes the feature value wide. Thus, the feature value defining the phoneme *th* (TTCA-wide) falls outside the space of Hebrew's distinctive features (additional empirical support for this conclusion is presented in Section 5).

To test whether speakers generalize the constraint on root structure outside the space of Hebrew phonemes and feature values, we presented participants with novel roots including novel phonemes. We incorporated these phonemes in three types of roots. One type of roots had geminates at the beginning of the root (e.g. jjr). A second type of roots presented the same geminates at the end of the root (e.g. rjj), and a third type of root had no geminates (e.g. jkr). If speakers ban identical consonants in the beginning of the root, then roots with initial geminates should be considered unacceptable compared to roots with final geminates or no geminates. Because root-initial geminates and no-gemination roots differ on the number of novel phonemes, the rejection of root-initial gemination compared to no-gemination controls may also be explained in terms of sheer phonological novelty. In contrast, roots with initial geminates and roots with final geminates are strictly matched for their segmental contents: they include the same phonemes, and differ only in their order (e.g. JJR vs. RJJ). A distinction between these root types must indicate sensitivity to root structure, specifically the location of geminates in the root.

If speakers can generalize the constraint on root structure beyond the space of Hebrew phonemes, then they should consider roots including foreign phoneme geminates at their beginning as less acceptable than roots including the same geminates at their end. The subset of roots including the *th* geminates allows us to further test whether speakers can generalize the constraint on root structure for novel feature values. If the constraint on root structure concerns the co-occurrence of specific tokens, corresponding to Hebrew features (e.g. labials), then such knowledge should be mute with respect to the co-occurrence of interdentals (specifically, coronal fricatives defined as TTCA-wide). Consequently, speakers should be insensitive to the location of geminates composed of the phoneme *th*. In contrast, if the constraint on root structure appeals to variables (e.g. *XX, where X stands for any consonant), then this constraint should generalize beyond the space of Hebrew features. Roots including root-initial geminates with the phoneme *th* should thus be considered ill-formed.

In addition to assessing the asymmetry in the acceptability of geminates, we also examine the domain defining their location. Recall that the OCP constrains geminates within the root. The root morpheme, however, is an abstract variable. A consistent sensitivity to the structure of the root would thus indicate the representation of the root by a variable. An alternative statistical account, however, may attribute the constraint on the location of geminates to their position in the word, not the root. For instance, the unacceptability of words such as *jajartem* from the root *jjr* may be due to the ill-formedness of identical consonants at the beginning of the word, rather than the root. To determine whether the domain of the OCP is the root or the word, our experiments systematically dissociate the position of geminates in the root from their word position. For this end, Experiments 1 and 2 conjugate each of the three types of roots in word patterns that differ in the location of geminates in the word. For instance, the root *jjr* is presented as either *jajartem*, where the geminates are word-initial, or *hijtajartem*, where geminates are preceded by a prefix (*hi*) and separated by an infix (*ta*). Geminates in the word *hijtajartem* are thus neither word-initial nor are they adjacent. If the constraint on the location of geminates appeals to the word, then one should not expect a consistent response across these various types of word structure. Conversely, if the constraint appeals to the root, then speakers should consistently consider each of these forms as unacceptable. Experiments 1 and 2 assess the constraint on root structure by means of rating novel roots generated from foreign phonemes. Experiment 3 obtains an on-line assessment of this knowledge by incorporating these roots in a lexical decision task.

Table 2
The structure of the materials used in Experiments 1 and 2 as a function of root type and the transparency of the word structure

Type	Root	Transparent	Opaque
Root-initial geminates	jjr	ja-jar-tem	hij-ta-jar-tem
Root-final geminates	rjj	ra-jaj-tem	hit-ra-jaj-tem
No geminates	jkr	ja-kar-tem	hij-ta-kar-tem

2. Experiment 1

Experiment 1 obtained acceptability ratings for word trios generated by conjugating a matched set of three root types (see Table 2). One root type had foreign phoneme geminates in its beginning (e.g. JJR), a second root type had the same geminates at the root's end (e.g. RJJ), and a third type was matched to the first on the initial and final consonant, but had no geminates (e.g. JKR). The words generated from each root trio were matched on their morphological structure (i.e. the location of root in the word, the vowels and affixes) and differed only on their root structure. To determine whether the constraint on geminates concerns their position in the root or their surface word position, we manipulated the morphological structure of the word. In half of the stimuli, the root was not preceded by a prefix, and hence their morphological structure was relatively transparent. In these word patterns, root-initial geminates were salient, and their location was invariably word-initial. The structure of the other half of the materials was considered opaque, since these roots were sandwiched between a prefix and a suffix. In these words, root-initial geminates were never word-initial and root-final geminates were not word-final. Furthermore, because root-initial sibilants are subject to metathesis (a phonological process that switches the location of phonemes, e.g. *hit* + *saper* → *histaper*, *cut one*'s *hair*), many of the root-initial geminates used in our experiments in this word pattern were further separated by an affix.[4] For instance, the root *jjr* is inflected as *hij-ta-jar-tem*, switching the location of the affix *t* and the root-initial consonant *j* (instead of *hit-ja-jar-tem*).

Participants in our experiments were presented with matched word trios, and asked to rate the acceptability of each member of the trio relative to the other members. If speakers' knowledge of root structure extends beyond the space of Hebrew phonemes, then we expect words generated from roots with initial gemination to be rated lower than roots with either final gemination or no gemination, regardless of the location of geminates in the word. If the constraint on the root generalizes to novel features, then the rejection of root-initial geminates should emerge for the phoneme *th*, a phoneme whose TTCA feature value is absent in the Hebrew feature inventory.

[4] The application of phonological rules, such as the metathesis rule above, to novel phonemes is not entirely clear. Traditional formulations of the metathesis rule define its input as sibilants. Although, phonetically, the phonemes *th*, *j* and *ch* are not considered sibilants (Ladefoged & Maddieson, 1996), our intuitions suggest that the metathesis rule generalizes to these phonemes. The precise domain of the metathesis rule requires further research. Note that, because roots with initial gemination and no gemination are matched for the initial phoneme, they are each subject to metathesis, and hence a distinction between these forms cannot be explained by the metathesis operation.

2.1. Method

2.1.1. Participants

Twenty-four University of Haifa students who were native speakers participated in the experiment. Participants were paid for taking part in the experiment.

2.1.2. Materials

The materials consisted of 72 trios of words. The words were generated by inserting 24 root trios in two types of word pattern (see Appendix A). Each root trio had three members. One member included a foreign phoneme geminate at the beginning of the root (e.g. JJR). A second root type had the same geminates root-finally (e.g. RJJ). A third root type was matched for the root with initial geminates for the first and last phoneme, but had a different radical root medially (e.g. JKR). We used four foreign phonemes to form geminates: *th*, *ch*, *j*, and *w*. The roots were next conjugated in two verbal patterns. The morphological structure of the first verbal pattern was relatively transparent. It had a suffix, but no prefix, and hence root-initial geminates were presented conspicuously at the beginning of the word. These verbal patterns included the verbal word patterns *piʔel* and *paʔal*. The second word pattern was the verbal *hitpaʔel*. Members of this word pattern are prefixed and suffixed, and hence the location of geminates in the word was rather opaque. Each of the 24 root trios was conjugated in each of the two word patterns, resulting in 72 trios of words.

2.1.3. Procedure

Participants were presented with a printed list including 72 word trios. The order of the trios was random. Likewise, the order of the words within the trio was randomly determined. Because the Hebrew orthography does not provide for some of the foreign phonemes used in our roots, we presented participants with an English transcription. To assist with the parsing of these long letter strings, we indicated syllable boundaries by a hyphen. To make sure that speakers represented the foreign phonemes correctly, they were asked to read each trio aloud and the experimenter verified that the pronunciation of the words was accurate. Participants were specifically instructed to read the words "as if they were Hebrew words and to pronounce them as carefully as possible". They were next instructed to compare the words in the trio and determine the extent to which they sounded like a possible Hebrew word. Participants were asked to assign the rating 1 to the word that sounded the best, 2 to the word that sounded intermediate, and 3 to the word that sounded the worst. To express high acceptability ratings by larger numbers, we inverted the scale by subtracting each score from 4. Thus, in our report, 1 corresponds to the word that sounded worst and 3 indicates the word that sounded best.

2.2. Results and discussion

Table 3 provides mean acceptability ratings as a function of root type and word pattern. Rating scores were submitted to two-way ANOVAs (2 word pattern $\times$ 3 root type) by participants and items. In this and all subsequent experiments we adopt 0.05 as the level of statistical significance. The ANOVAs yielded a significant main effect of root type

Table 3
Mean acceptability rating for the words presented in Experiment 1 as a function of root type and word structure

	Transparent	Opaque
Root-initial gemination	1.57	1.45
Root-final gemination	1.99	2.15
No gemination	2.44	2.41

($F_s(2,46) = 71.87$, $MSE = 0.142$, $P < 0.0001$; $F_i(2,46) = 89.25$, $MSE = 0.114$, $P < 0.0001$) and a significant interaction of root type × word pattern ($F_s(2,46) = 8.73$, $MSE = 0.028$, $P < 0.0007$; $F_i(2,46) = 3.99$, $MSE = 0.062$, $P < 0.03$). Across word patterns, roots with initial gemination were rated as significantly less acceptable compared to roots with either final gemination ($F_s(1,46) = 52.75$, $P < 0.0001$; $F_i(1,46) = 65.51$, $P < 0.0001$) or no gemination ($F_s(1,46) = 141.43$, $P < 0.0001$; $F_i(1,46) = 175.63$, $P < 0.0001$). The same pattern emerged regardless of the location of geminates in the word. Specifically, for the transparent word pattern, words whose roots exhibit initial gemination were rated significantly lower compared to words with root-final gemination ($F_s(1,46) = 73.16$, $P < 0.0001$; $F_i(1,46) = 33.51$, $P < 0.0001$) or no gemination ($F_s(1,46) = 316.11$, $P < 0.0001$; $F_i(1,46) = 144.79$, $P < 0.0001$). Likewise, in the opaque word pattern, words with root-initial gemination were significantly less acceptable compared to roots with final gemination ($F_s(1,46) = 206.22$, $P < 0.0001$; $F_i(1,46) = 94.48$, $P < 0.0001$) or no gemination ($F_s(1,46) = 389.65$, $P < 0.0001$; $F_i(1,46) = 178.53$, $P < 0.0001$).[5] The consistent rejection of root-initial geminates across various word positions suggests that the constraint on geminates concerns their position in the root, rather than their surface word position. The observation of this pattern for roots with foreign phoneme geminates suggests that speakers freely generalize the constraint on root structure to novel phonemes outside the inventory of Hebrew segments.

To examine whether speakers can generalize the constraint on root structure beyond the space of Hebrew *features*, we next turn to examine responses to roots including *th* geminates, a phoneme whose TTCA feature value does not exist in the Hebrew language. The pattern of results for the *th* roots replicates the findings reported above (see Table 4). The ANOVAs conducted over the *th* roots yielded a significant main effect of root type ($F_s(2,46) = 72.93$, $P < 0.0001$, $MSE = 0.193$; $F_i(2,10) = 44.24$, $MSE = 0.08$, $P < 0.0001$) and a significant interaction of root type × word pattern ($F_s(2,46) = 6.16$, $MSE = 0.131$, $P < 0.005$; $F_i(2,10) = 4.89$, $MSE = 0.04$, $P < 0.04$). Across word patterns, root-initial gemination was rated significantly lower than root-final gemination ($F_s(1,46) = 69.90$, $P < 0.0001$; $F_i(1,10) = 42.42$, $P < 0.0001$) and no gemination ($F_s(1,46) = 137.54$, $P < 0.0001$; $F_i(1,10) = 83.44$, $P < 0.0001$). Roots with initial gemi-

[5] Each of the word patterns also exhibited significantly lower ratings for roots with final gemination compared to no gemination. Because roots with final gemination include more unfamiliar phonemes than roots with no gemination, this pattern may be simply explained by the amount of unfamiliar phonological material in the root. Alternatively, however, root-final gemination may be less acceptable than roots with no gemination, a finding obtained previously for existing Hebrew phonemes by Berent and Shimron (1997) and Berent, Everett, and Shimron (2001), using this rating task. In view of the confound between root structure and the amount of novel phonemes, this finding is uninterpretable in the present experiment, and hence we will not discuss it further.

Table 4
Mean acceptability rating for the words including the foreign phoneme *th* as a function of root type and word structure (Experiment 1)

	Transparent	Opaque
Root-initial gemination	1.49	1.31
Root-final gemination	2.00	2.30
No gemination	2.51	2.39

nation were significantly less acceptable than their final-gemination and no-gemination controls in both the transparent (compared to root-final gemination: $F_s(1, 46) = 24.10$, $P < 0.0001$; $F_i(1, 10) = 19.19$, $P < 0.002$; compared to no gemination: $F_s(1, 46) = 96.39$, $P < 0.0001$; $F_i(1, 10) = 76.69$, $P < 0.0001$) and the opaque word patterns (compared to root-final gemination: $F_s(1, 46) = 88.75$, $P < 0.0001$; $F_i(1, 10) = 70.62$, $P < 0.0001$; compared to no gemination: $F_s(1, 46) = 105.75$, $P < 0.0001$; $F_i(1, 10) = 84.15$, $P < 0.0001$). These results demonstrate that Hebrew speakers constrain the location of identical consonants in the root, and generalize this constraint beyond the phonological space of Hebrew phonemes and features.

3. Experiment 2

The findings of Experiment 1 suggest that speakers can generalize the constraint on root structure beyond the phonological space of Hebrew. These findings, however, may be criticized on the grounds that their relevance to the study of linguistic generalizations is uncertain. Because participants were explicitly asked to compare words that differ only on their roots, it is possible that their sensitivity to the location of geminates is the product of some meta-linguistic strategies, rather than linguistic competence. Although there is now substantial evidence that the constraint on root structure applies in a mandatory fashion even when attention to the root is not required (Berent et al., 2000; Berent, Shimron, & Vaknin, 2001), it is possible that the extension of such knowledge to foreign phonemes may not be governed by the grammar. Experiment 2 attempts to address this limitation by performing a simple change in the rating task. Participants in this experiment were presented with the same materials employed in Experiment 1 in a randomized list. They were asked to rate the acceptability of each word independently, rather than to compare it to its matched controls. Because the words in the list differ on numerous dimensions, including the number of novel phonemes and the word pattern, this rating procedure does not require that participants attend to root structure. If the constraint on the location of novel geminates in the root, observed in Experiment 1, is due to the explicit comparison of matched root structures, then no sensitivity to root structure should be obtained in the present experiment. Conversely, if the constraint on root structure applies across the board, then root-initial gemination should be unacceptable even when attention to root structure is not required.

Table 5
Mean acceptability rating for the words presented in Experiment 2 as a function of root type and word structure

	Transparent	Opaque
Root-initial gemination	2.71	2.57
Root-final gemination	2.89	2.91
No gemination	3.08	2.97

3.1. Method

Twenty University of Haifa students who were native Hebrew speakers participated in this experiment. These individuals did not take part in Experiment 1, and they were paid for their participation. The materials and procedure were the same as in Experiment 1 with the only exception that the words were presented in a randomized printed list. Participants were asked to read each word aloud and then indicate its acceptability as a possible Hebrew word using a 1–5 scale (1 = impossible; 5 = excellent).

3.2. Results and discussion

Table 5 provides mean acceptability ratings for the words presented in Experiment 2 as a function of root type and word structure. The ANOVAs (2 word pattern × 3 root type) yielded a significant effect of root type ($F_s(2, 38) = 11.32$, MSE = 0.122, $P < 0.0001$; $F_i(2, 46) = 12.01$, MSE = 0.155, $P < 0.0002$). No other effect reached significance (all $F < 1.15$). Replicating the findings of Experiment 1, words generated from roots with initial gemination were rated significantly lower compared to roots with either root-final gemination ($F_s(1, 38) = 8.50$, $P < 0.007$; $F_i(1, 46) = 10.44$, $P < 0.003$) or no gemination ($F_s(1, 38) = 22.21$, $P < 0.0001$; $F_i(1, 46) = 23.11$, $P < 0.0001$).

To examine whether the constraint on root structure extends for roots with foreign feature values, we next turn to examine the ratings of roots including the geminate *th* (see Table 6). The ANOVAs performed on these items yielded a main effect of root type ($F_s(2, 38) = 5.72$, MSE = 0.309, $P < 0.007$; $F_i(2, 10) = 7.84$, MSE = 0.067, $P < 0.01$). No other effect reached significance (all $F < 1$). Root-initial gemination was rated significantly lower compared to roots with either final gemination ($F_s(1, 38) = 7.22$, $P < 0.02$; $F_i(1, 10) = 9.89$, $P < 0.02$) or no gemination ($F_s(1, 38) = 9.75$, $P < 0.004$; $F_i(1, 10) = 13.37$, $P < 0.005$). These findings converge with the results of Experiment 1 in demonstrating that speakers freely generalize the constraint on root structure to roots with novel phonemes, including phonemes with a novel feature value. Speakers constrain

Table 6
Mean acceptability rating for the words including the foreign phoneme *th* as a function of root type and word structure (Experiment 2)

	Transparent	Opaque
Root-initial gemination	2.36	2.30
Root-final gemination	2.63	2.69
No gemination	2.72	2.72

Table 7
An illustration of the target words and foils used in Experiment 3 as a function of root structure

	Target words	Foils
Root-initial geminates	————	hij-ta-jar-tem
Root-final geminates	hit-pa-lal-tem	hit-ra-jaj-tem
No geminates	hit-pa-lash-tem	hij-ta-kar-tem

the location of the geminates in the root regardless of their word position, and apply this constraint even when the task does not require attention to root structure.

4. Experiment 3

This experiment seeks converging evidence for the application of the constraint on root structure under conditions that do not encourage conscious awareness of root structure by using a time-limited procedure and stimuli materials whose morphological structure is opaque. This experiment employs a variant of the lexical decision task. Participants were presented with existing Hebrew words and novel words. To encourage the encoding of the word's phonemes, they were asked to read each letter string aloud and indicate as fast as they could whether it corresponded to an existing Hebrew word by pressing one of two keys on the computer. The novel words were the materials used in Experiments 1 and 2 in the opaque word pattern condition, and they exhibited roots with either initial geminates, final geminates or no geminates (see Table 7). The words were matched for the novel roots on their word pattern, and their roots exhibited either root-final gemination or no gemination. Because Hebrew has only two productive roots with root-initial geminates, it was impossible to systematically examine the asymmetry in the location of geminates for existing roots, and hence our main interest concerns performance with novel roots. If the constraint on root structure applies in on-line reading, then speakers should be sensitive to the location of geminates in the root. Roots with initial gemination should be perceived ill-formed, and hence they should be easier to classify as nonwords compared to roots with final gemination or no gemination.

4.1. Method

4.1.1. Participants

Twenty University of Haifa students who were native Hebrew speakers participated in the experiment. Participants were paid for taking part in the experiment.

4.1.2. Materials

The stimuli materials consisted of 72 Hebrew words and 72 novel words. The novel words corresponded to the 72 opaque novel words used in Experiments 1 and 2, generated from 24 root trios with either initial gemination, final gemination or no gemination. The target words shared the same verbal patterns as the novel words, but they corresponded to existing Hebrew words. The target words were generated from 36 matched pairs of existing roots, exhibiting either root-final gemination or no gemination (see Appendix B). All stimuli were transcribed using English letters. To facilitate the decoding of these long letter strings, syllable boundaries were indicated by hyphens.

4.1.3. Practice trials

To familiarize the participants with the experimental task, they were first presented with practice trials, consisting of 18 words and 18 nonwords, presented in a random order. These words and nonwords shared the same word patterns as the experimental stimuli. None of the practice stimuli appeared in the experimental trials.

4.1.4. Procedure

At the beginning of each trial, a fixation point consisting of three asterisks appeared at the center of the screen. Participants initiated the trial by pressing the space bar. They were then presented with a string of letters at the center of the computer screen, displayed until participants responded or a maximum of 3 s elapsed. Participants were asked to read each letter string aloud and determine whether it corresponded to an existing Hebrew word. Word responses were given by pressing the 1 key. Nonword responses were given by pressing the 2 key. These two keys were marked by either “yes” or “no” labels, and positioned such that participants used their preferred hand to provide both responses. Slow responses (responses slower than 3 s) and inaccurate responses received negative feedback from the computer in the form of a tone and a computer message. The experiment was initiated with the practice stimuli followed by the experimental trials. Participants were tested individually. The order of the trials in the experiment was random.

4.2. Results and discussion

To eliminate the effect of outliers, we excluded responses falling 2.5 SD beyond the mean. This procedure resulted in the exclusion of less than 1.4% of the total observations. We next submitted response latency and accuracy for words and nonwords to separate ANOVAs. Mean response latency and accuracy to target words is provided in Table 8. The ANOVAs on responses to the target words did not yield a significant effect of root structure in either the latency ($F_s(1,19) = 1.41$, $\mathrm{MSE} = 4034$, $P < 0.25$, NS; $F_i(1,35) < 1$, $\mathrm{MSE} = 23,193$, NS) or accuracy measures ($F_s(1,19) = 2.46$, $\mathrm{MSE} = 0.002$, $P < 0.14$, NS; $F_i(1,35) < 1$, $\mathrm{MSE} = 0.01$, NS). Our main interest, however, concerns responses to novel words.

Table 9 provides mean response latency and accuracy for nonword foils. The ANOVAs on response latency revealed a main effect of root type, significant by items ($F_i(2,46) = 3.21$, $\mathrm{MSE} = 11,535$, $P < 0.05$) and marginally so by participants ($F_s(2,38) = 2.89$, $\mathrm{MSE} = 6200$, $P < 0.07$). A significant effect of root type also emerged

Table 8
Mean response latency and accuracy for the target words used in Experiment 3 as a function of root structure

	Response latency (ms)	Response accuracy (% correct)
Root-final gemination	1809	88.2
No gemination	1833	90.3

Table 9
Mean response latency and accuracy for the nonword foils used in Experiment 3 as a function of root structure

	Response latency (ms)	Response accuracy (% correct)
Root-initial gemination	1916	92.4
Root-final gemination	1976	89.5
No gemination	1954	83.9

in the analyses of response accuracy ($F_s(2,38) = 8.27$, MSE $= 0.004$, $P < 0.002$; $F_i(2,46) = 5.94$, MSE $= 0.007$, $P < 0.006$). Novel words with root-initial gemination were rejected significantly faster compared to roots with final gemination ($F_s(1,38) = 5.63$, $P < 0.03$; $F_i(1,46) = 6.40$, $P < 0.02$). Likewise, response accuracy to roots with initial gemination was numerically higher than responses to roots with final gemination ($F_s(1,38) = 1.74$, $P < 0.20$, NS; $F_i(1,46) = 1.39$, $P < 0.25$, NS). Responses to roots with initial gemination were also significantly more accurate ($F_s(1,38) = 15.92$, $P < 0.0004$; $F_i(1,46) = 11.53$, $P < 0.002$) and numerically faster ($F_s(1,38) = 2.33$, $P < 0.13$, NS; $F_i(1,46) = 1.90$, $P < 0.17$, NS) compared to roots with no gemination.

To examine whether the constraint on root structure generalizes to phonemes with novel phonetic features, we conducted separate analyses on responses to novel words with the *th* geminate (see Table 10). The ANOVAs yielded a significant main effect of root type in the analysis over response latency ($F_s(2,38) = 4.65$, MSE $= 27,101$, $P < 0.02$; $F_i(2,10) = 3.99$, MSE $= 14,411$, $P < 0.06$), but not in response accuracy ($F_s(2,38) = 2.11$, MSE $= 0.034$, $P < 0.14$; $F_i(2,10) = 1.83$, MSE $= 0.012$, $P < 0.22$). Novel roots with initial gemination were rejected significantly faster ($F_s(1,38) = 10.94$, $P < 0.003$; $F_i(1,10) = 7.25$, $P < 0.02$) and numerically more accurately ($F_s(1,38) = 2.04$, $P < 0.17$, NS; $F_i(1,10) = 1.76$, $P < 0.22$, NS) compared to roots with final gemination. Responses to roots with initial geminates were also numerically more accurate ($F_s(1,38) = 3.97$, $P < 0.06$; $F_i(1,10) = 3.44$, $P < 0.10$, NS) and faster ($F_s(1,38) = 1.61$, $P < 0.22$; $F_i(1,10) < 1$) compared to roots with no gemination.

The findings of the lexical decision experiments demonstrate that speakers constrain the location of geminates in the root. Novel words with novel geminate phonemes root-initially were rejected significantly faster than controls where the same geminates were presented root-finally. Because roots with initial geminates and final geminates are strictly matched on the number of foreign phonemes, the presence of geminates, and the word's morphological structure, the distinction between them clearly indicates sensitivity to the

Table 10
Mean response latency and accuracy for the nonword foils including the phoneme *th* as a function of root structure (Experiment 3)

	Response latency (ms)	Response accuracy (% correct)
Root-initial gemination	2049	86.6
Root-final gemination	2208	78.3
No gemination	2116	75.1

location of geminates in the root. The sensitivity to the location of geminates in the root agrees with previous findings obtained in the lexical decision task using novel roots with existing Hebrew phonemes (Berent, Shimron, & Vaknin, 2001). The present results replicate and extend these earlier findings to suggest that speakers generalize the constraint on root structure outside the set of Hebrew phonemes. The emergence of this findings for geminates with the interdental *th* further indicates that the constraint generalizes to phonemes with a novel phonemic feature.

Although our findings indicate a reliable facilitation in the rejection of foils with root-initial geminates relative to final-gemination controls, the distinction between initial-gemination roots and no-gemination controls was observed only in the accuracy measure, and was only marginally significant for roots with the foreign phoneme *th*. This pattern, however, is in full agreement with the earlier findings of Berent, Shimron, and Vaknin (2001) obtained using novel roots with existing Hebrew phonemes. Berent and colleagues observed no difference in response to root-initial geminates compared to no-gemination controls despite significant differences relative to final-gemination controls.[6] They attributed the absence of a distinction between root-initial gemination and no-gemination controls to a confound created by the very presence of geminates in the root. Indeed, Berent, Shimron, and Vaknin (2001) observed that the presence of geminates in roots with final geminates impairs their rejection relative to frequency-matched no-gemination controls. They explained the deleterious effect of geminates by their perception as wordlike. Participants in the lexical decision task are known to discriminate between targets and foils based on the appearance of the stimulus as wordlike. The "wordhood" of stimuli may be affected not only by familiarity with their meanings and spellings (Balota & Chumbly, 1984) but also by grammatical structure: stimuli that are viewed as formed by productive grammatical operations may be perceived as more "wordlike" than stimuli that do not bear the hallmark of the grammar. If the OCP is active, then geminates must be formed productively by the grammar. Because roots with geminates are the product of a grammatical operation, they may be viewed as more "wordlike", and hence novel roots with gemination may be more difficult to discriminate from existing roots compared to novel roots with no gemination. The comparison of roots with initial gemination to no-gemination controls thus confounds the presence of geminates with the ill-formedness associated with their location. The absence of significant differences between root-initial vs. no-gemination may reflect the cancellation of these multiple conflicting forces. In contrast, the comparison of roots with initial and final geminates strictly controls for the presence of geminates, and hence differences between these root types must indicate sensitivity to the location of geminates in the root. The present results indicate that speakers are sensitive to the location of geminates in the root even when geminates are composed of foreign phonemes with a novel phonemic feature, and despite the use of an opaque word pattern. These findings demonstrate that speakers constrain the location of geminates in the root, and extend these constraints beyond the phonological space of Hebrew.

[6] This pattern is difficult to interpret in the present experiment since the distinction between roots with final geminates and no geminates is confounded with the number of novel phonemes.

5. General discussion

Hebrew roots frequently contain geminates, but the location of geminates is strictly constrained. This distributional fact is attributed to an active phonological constraint, the OCP. Previous experiments (Berent et al., 2000; Berent, Everett, & Shimron, 2001; Berent & Shimron, 1997; Berent, Shimron, & Vaknin, 2001) have demonstrated that the constraint on root structure generalizes to novel roots within the phonemic inventory of Hebrew. The present experiments examined whether this constraint generalizes to roots with foreign phonemes. We demonstrated that words containing foreign phoneme geminates root-initially are considered ill-formed. These roots are rated as less acceptable than roots with final geminates or no geminates, regardless of the position of geminates in the word. Roots with initial geminates are also easier to be classified as foils in the lexical decision task compared to roots containing these same geminates root-finally. Speakers' sensitivity to the location of geminates under conditions that do not require attention to root structure (in Experiments 2 and 3) and using a time-limited procedure (in Experiment 3) suggests that the constraint on root geminates forms part of their linguistic competence.[7] Our findings thus indicate that speakers can generalize a linguistic constraint to novel phonemes. Because foreign phoneme geminates are absent in Hebrew, the distinction between roots with initial vs. final geminates is inexplicable by the statistical frequency of these geminate bigrams in the language. In fact, the same findings replicate even for a phoneme containing a novel phonemic feature value. These results suggest that speakers can generalize a linguistic constraint beyond the space of phonemes and distinctive features of their native language.

Generalization beyond the training space is the hallmark of symbolic architectures. Because symbolic representations encode variables, they can constrain the form of linguistic representations, regardless of any specific tokens that instantiate them. For example, a

[7] Evidence that the ban on root-initial geminates is part of the Hebrew speaker's *linguistic* competence rather than merely a domain-general auditory or articulatory preference comes from a replication of Experiment 1 – but with native *English* speakers. This follow-up experiment, suggested to us by an anonymous reviewer, obtained relative acceptability ratings for the transparent words included in Experiment 1. Because our goal here was to maximize the chances of detecting a domain-general bias, we chose not to include the opaque word structures. (Even if English speakers were sensitive to consonant identity across intermediate vowels, they would be unlikely to perceive root consonants as adjacent in opaque strings, which often separate the root geminates by a consonantal infix (e.g. *hijtajartem*, from the root jjr). Opaque words could also hinder identity detection in the transparent words, where it is most likely to occur, as opaque strings – being long and foreign-sounding – might prompt participants to treat the entire list of stimuli as impossible English words.) Our experiment was therefore limited to the transparent word patterns. Twenty-one native English speakers (students at Florida Atlantic University who had no knowledge of Semitic languages) rated the relative acceptability of these Hebrew materials as English words. The experimenter indicated the location of the primary stress in these strings by reading some examples aloud. The effect of structure was highly significant ($F_s(2,40) = 43.22$, MSE = 0.132, $P < 0.0002$; $F_i(2,46) = 97.46$, MSE = 0.067, $P < 0.0002$). Mean acceptability ratings for words with initial gemination, final gemination and no gemination were 2.1, 1.5 and 2.5, respectively (1 = worst; 3 = best). In agreement with Hebrew speakers, English speakers considered words with gemination less acceptable than words with no gemination ($F_s(1,40) = 13.81$, $P < 0.0007$; $F_i(1,46) = 31.16$, $P < 0.0001$). In contrast to speakers of Hebrew, however, English speakers considered words with initial geminates *more* acceptable than roots with final geminates ($F_s(1,40) = 30.47$, $P < 0.0002$; $F_i(1,46) = 68.76$, $P < 0.0001$). Although we do not know whether these findings reflect sensitivity to the statistical frequency of geminate bigrams or to their identity, they make it clear that the ban on root-initial geminates is not general across languages.

ban on root-initial geminates may be defined as *XXY, where X stands for any phoneme.[8] Operations over variables enable generalizations that can be extended to any instance, including phonemes that do not fall within the learner's training space. Marcus (1998, 2001) has shown that the representation of variables and operations over variables is necessary for achieving such generalizations. He demonstrated that multilayer perceptrons that eliminate representations of variables and operations over variables cannot generalize functions that link variables (universally quantified one to one mappings) beyond their training space. Marcus showed formally that this failure to generalize is principled: it stems from the independence in learning the weight of connections on any given input or output unit from connections on other input/output units. Consequently, the weights acquired on trained nodes cannot constrain the activation of untrained notes, a property that prevents multilayer perceptrons from generalizing to untrained nodes representing novel instances or novel feature values.

In contrast to such networks, people can systematically generalize the constraint on root structure to novel Hebrew phonemes with novel feature values. Given Marcus' conclusions, such generalizations should be unattainable by multilayer perceptrons that eliminate operations over variables. Proponents of associationist accounts of cognition, however, may point out that our analysis does not prove that the constraint on root structure is unlearnable by pattern associators. It is plain that the scope of generalization is intimately linked to the definition of the training space in terms of a specific feature geometry (the hierarchical organization of feature values, e.g. Clements & Hume, 1995). A phoneme that falls outside a phonological space X could potentially be accommodated within an alternative phonological space Y by changing the feature geometry. For instance, our demonstration that speakers can generalize the constraint on root structure to a novel phoneme with a novel feature value, the interdental *th*, hinges on its representation by the novel feature value "wide" using the distinctive feature TTCA (Gafos, 1999). Although the phoneme *th* falls outside this feature space, it is conceivable that alternative analyses could accommodate this phoneme using some combination of features that are all found in Hebrew phonemes. More generally, proponents of associative accounts may claim that any novel instance may be accommodated within the learner's training space by crafting the representational space to fit the desired generalization, and hence, in practice, learners may never be required to generalize beyond their training space.

We do not dispute the possibility that the novel phonemes in our experiments may be accommodated within the feature space of Hebrew by alternative feature geometries. We see several problems with this approach, however. The first concerns the psychological plausibility of alternative feature geometries. A feature analysis, like any other account of mental representations, must capture critical aspects of speakers' knowledge. Although one can clearly design the feature space to fit a desired computational goal, such a solution may not necessarily provide a plausible account for speakers' phonological knowledge,

[8] Alternatively, one could encode the constraint on root structure in terms of the identity of adjacent root radicals. Roots with initial geminates could be represented as "Identical-Different" (indicating a geminate bigram followed by a non-geminate bigram), whereas roots with final geminates could be captured as Different-Identical (indicating a non-geminate bigram followed by a geminate bigram). Because this notation encodes identity, a relationship among abstract variables, it requires operations over variables.

and hence this approach may ultimately fail to capture the entire set of facts that must be explained in a given domain. An important psychological motivation for a feature geometry is the existence of an invariant phonetic basis for its proposed distinctive features. Our analysis of *th* as foreign to the Hebrew inventory is supported by a careful phonetic articulatory analysis (Gafos, 1999). This analysis reveals an invariant phonetic dimension that contrasts coronal fricative consonants across speakers and languages, namely, the area of the cross-sectional channel between the tongue and the palate. The distinctive feature of the TTCA transparently captures this articulatory dimension. Although alternative analyses for the classification of coronal fricatives may be found in the linguistic literature, these proposals lack invariant articulatory or acoustic correlates. Gafos (1999) reviews many of these proposals, including contrasts on place of articulation, the orientation of the tongue tip constriction (the features apical vs. laminal, e.g. Ladefoged & Maddieson, 1996), the shape of the tongue (the feature groove, e.g. Catford, 1977; Halle & Stevens, 1979), and the length of the tongue tip constriction (the feature distributed, e.g. Chomsky & Hale, 1968). None of these proposals could capture the speaker-to-speaker variability or cross-language variability in the pronunciation of coronal fricatives. For instance, English speakers produce the consonant *s* at six different places of articulation using either the laminal or apical constriction orientation (Dart, 1991). A similar variability in the place of articulation and orientation of the tongue tip has been observed for the phoneme *th* (e.g. Ladefoged & Maddieson, 1996). The rejection of our analysis in favor of such alternative geometries is likely to obscure the articulatory basis of phonemic classifications, compromising the psychological plausibility and learnability of such a representation by speakers.

Alternative geometries that place the phoneme *th* within the Hebrew space are further challenged by empirical evidence. One piece of evidence comes from the pattern of borrowing foreign phonemes into Hebrew. Hebrew has numerous loan words from English and Arabic that originally include the foreign phonemes used in our experiments (Choueka, 1997). Borrowings with the phonemes ch (e.g. *check*, a total of 32 cases), j (e.g. *job*, a total of 20 words), and *w* (*baklawa*, from Arabic, a total of six words) typically preserve the pronunciation of the original phoneme. In contrast, none of the numerous borrowings with *th* preserves its original place of articulation. Without exception, the *th* phoneme in the loan word is transformed into *t* (e.g. *thermometer* → *termometer*, *therapy* → *terapya*). The analysis of the phoneme *th* (but not ch, w, and j) as including a foreign feature value accounts for this fact. Additional support for our analyses is offered by the findings of Experiment 2. Recall that this experiment used an open ended rating task, a task that allows discriminating between words based on any phonological aspect, not necessarily the location of geminates. If the *th* phoneme falls outside the Hebrew feature inventory, then roots exhibiting this phoneme should be perceived as less "Hebrew-like" than roots including the phonemes *ch*, *j* and *w*. A separate analysis of the roots with no geminates supports this prediction. No-gemination roots containing the phoneme *th* ($M = 2.71$) were rated significantly lower than roots containing the phonemes *ch* ($M = 3.34$; $F_s(1,57) = 15.78$, $P < 0.0003$; $F_i(1,20) = 10.95$, $P < 0.004$) and *j* ($M = 3.17$; $F_s(1,57) = 8.29$, $P < 0.006$; $F_i(1,20) = 5.75$, $P < 0.03$), and numerically (but not significantly) lower than roots with the *w* phoneme ($M = 2.89$; $F_s(1,57) = 1.19$, $P < 0.28$, NS; $F_i(1,20) < 1$, NS). These findings lend support to the proposal that participants in our experiment represent the phoneme *th* by means of a novel feature value. It is unclear how alternative analyses that accommodate the phoneme *th* within the Hebrew feature space could account for these empirical observations.

A second problem associated with carving the feature space to accommodate novel phonemes concerns the ability of multilayer perceptrons to learn the constraint on phoneme co-occurrence using feature-based representations. The constraint on Hebrew root structure is quite transparent at the segment level, but it becomes far less patent at the level of features. Thus, although the appeal to feature-level representations may help fit novel phonemes within the model's training space, such representations may ultimately render the desired generalization unlearnable by multilayer perceptrons. It is well known that Semitic roots constrain not only identical segments but also nonidentical segments produced by the same articulator (McCarthy, 1994), such as labials (e.g. *b*, *m*, *f*), dorsals (consonants produced by the tongue body, e.g. *g*, *k*), coronal fricatives (e.g. *s*, *th*), etc. Our analysis of triliteral Hebrew roots (Berent & Shimron, 2001) listed in the Even-Shoshan Hebrew dictionary (Even-Shoshan, 1993) supports this observation. In addition to the ban on root-initial geminates (e.g. *bbd), Hebrew rarely exhibits adjacent radicals that share an articulator root-initially (e.g. *b* and *m* are both labials, and hence **bmd* is ruled out).

At first glance, this observation seems to lend further support for feature-based accounts of the constraint on root structure. For instance, it is tempting to capture the ban on adjacent identical segments as a ban on the same articulator feature. Although a constraint on articulator feature is clearly called for, such an account is insufficient to capture the desired generalization. Indeed, this approach fails when it comes to the acceptability of identical vs. similar segments root-finally. At the end of the root, adjacent identical segments are frequent (e.g. *d**mm***), but nonidentical segments that share an articulator (e.g. **d**mb***) are rare. Our experimental investigation further confirms that speakers consider root-final geminates as more acceptable than frequency-matched controls produced by the same articulator (Berent & Shimron, 2001). Multilayer perceptrons that use a featural representation may face some troubles capturing this fact. Because such multilayer perceptrons would be strongly sensitive to feature co-occurrence, they are likely to incorrectly generalize the frequent co-occurrence of identical segments root-finally (e.g. *dmm*) as indicating the acceptability of non-identical segments sharing the same articulator (e.g. *dmb*). Thus, although a featural representation may fit novel phonemes within the training space, it is uncertain whether such a representation would allow multilayer perceptrons to acquire the constraint on root structure.

In summary, we do not reject the idea that pattern associators may be able to accommodate novel phonemes within the phonological space of Hebrew. Such an approach, however, has a rather heavy price tag. It compromises the plausibility of the model as a phonological explanation by obscuring the articulatory basis of distinctive features, and it is further inconsistent with the pattern of ratings and borrowing for foreign phonemes. In fact, the appeal to a feature-based account may ultimately prevent multilayer perceptrons from learning the constraint on root structure altogether. Put generally, the attempt to extend the scope of generalizations by tinkering with phonological representations reduces the flexibility of the model and diminishes its ability to capture the phonological structure of the language. In contrast, symbolic architectures need to make no a priori assumptions concerning feature geometry. Because such models generalize in reference to operations defined over variables, they can extend generalization to any representable novel item, regardless of its relation to the model's representational space. The symbolic hypothesis makes the strong claim that such generalizations are attainable by humans. Our present empirical findings are consistent with this hypothesis.

Acknowledgements

This research was supported by the National Institute of Health FIRST award (1R29 DC03277-02) to Iris Berent and a National Institute of Health grant (HD 37059) to Gary Marcus. We are grateful to Steve Pinker for discussions of this research. We also thank Vered Vaknin for her assistance in conducting these experiments.

Appendix A. The materials employed in Experiments 1 and 2

Trio	Root	Transparent	Opaque
1	chchv	cha-chav-ti	hich-ta-chav-tem
1	vchch	va-chach-ti	hit-va-chach-tem
1	chnv	cha-nav-ti	hich-ta-nav-tem
2	chchd	chi-chad-nu	hich-ta-chad-nu
2	dchch	di-chach-nu	hit-da-chach-nu
2	chmd	chi-mad-nu	hich-ta-mad-nu
3	chchl	cha-chal-tem	hich-ta-chal-ti
3	lchch	la-chach-tem	hit-la-chach-ti
3	chbl	cha-bal-tem	hich-ta-bal-ti
4	chchr	chi-char-ti	hich-ta-char-tem
4	rchch	ri-chach-ti	hit-ra-chach-tem
4	chpr	chi-par-ti	hich-ta-par-tem
5	chchg	cha-chag-tem	hich-ta-chag-tem
5	gchch	ga-chach-tem	hit-ga-chach-tem
5	chrg	cha-rag-tem	hich-ta-rag-tem
6	chchk	chi-chak-nu	hich-ta-chak-nu
6	kchch	ki-chach-nu	hit-ka-chach-nu
6	chnk	chi-nak-nu	hich-ta-nak-nu
7	jjb	ji-jab-ti	hij-ta-jab-ti
7	bjj	bi-jaj-ti	hit-ba-jaj-ti
7	jlb	ji-lab-ti	hij-ta-lab-ti
8	jjg	ja-jag-nu	hij-ta-jag-nu
8	gjj	ga-jaj-nu	hit-ga-jaj-nu
8	jmg	ja-mag-nu	hij-ta-mag-nu
9	jjk	ji-jak-tem	hij-ta-jak-tem
9	kjj	ki-jaj-tem	hit-ka-jaj-tem
9	jbk	ji-bak-tem	hij-ta-bak-tem
10	jjl	ji-jal-ti	hij-ta-jal-ti

(*continued*)

Trio	Root	Transparent	Opaque
10	ljj	li-jaj-ti	hit-la-jaj-ti
10	jpl	ji-pal-ti	hij-ta-pal-ti
11	jjn	ja-jan-tem	hij-ta-jan-tem
11	njj	na-jaj-tem	hit-na-jaj-tem
11	jtn	ja-tan-tem	hij-ta-tan-tem
12	jjr	ja-jar-tem	hij-ta-jar-tem
12	rjj	ra-jaj-tem	hit-ra-jaj-tem
12	jkr	ja-kar-tem	hij-ta-kar-tem
13	ththk	thi-thak-nu	hith-ta-thak-nu
13	kthth	ki-thath-nu	hit-ka-thath-nu
13	thbk	thi-bak-nu	hith-ta-bak-nu
14	ththr	thi-thar-tem	hith-ta-thar-tem
14	rthth	ri-thath-tem	hit-ra-thath-tem
14	rthn	ri-than-tem	hit-ra-than-tem
15	ththn	tha-tan-tem	hith-ta-than-tem
15	nthth	na-thath-tem	hit-na-thath-tem
15	thpn	tha-pan-tem	hith-ta-pan-tem
16	ththg	thi-thag-nu	hith-ta-thag-nu
16	gthth	gi-thath-nu	hit-ga-thath-nu
16	thbg	thi-bag-nu	hith-ta-bag-nu
17	ththl	tha-thal-ti	hith-ta-thal-ti
17	lthth	la-thath-ti	hit-la-thath-ti
17	thml	tha-mal-ti	hith-ta-mal-ti
18	ththm	thi-tham-tem	hith-ta-tham-tem
18	mthth	mi-thath-tem	hit-ma-thath-tem
18	thgm	thi-gam-tem	hith-ta-gam-tem
19	wwg	wi-wag-ti	hit-wa-wag-ti
19	gww	gi-waw-ti	hit-ga-waw-ti
19	wdg	wi-dag-ti	hit-wa-dag-ti
20	wwd	wa-wad-nu	hit-wa-wad-nu
20	dww	da-waw-nu	hit-da-waw-nu
20	wzd	wa-zad-nu	hit-wa-zad-nu
21	wwz	wi-waz-tem	hit-wa-waz-tem
21	zww	zi-waw-tem	hiz-ta-waw-tem
21	wkz	wi-kaz-tem	hit-wa-kaz-tem
22	wws	wa-was-ti	hit-wa-was-ti
22	sww	sa-waw-ti	his-ta-waw-ti
22	wsk	wa-sak-ti	hit-wa-sak-ti
23	wwl	wi-wal-nu	hit-wa-wal-nu
23	lww	li-waw-nu	hit-la-waw-nu
23	wgl	wi-gal-nu	hit-wa-gal-nu

(*continued*)

Trio	Root	Transparent	Opaque
24	wwk	wa-wak-tem	hit-wa-wak-tem
24	kww	ka-waw-tem	hit-ka-waw-tem
24	wnk	wa-nak-tem	hit-wa-nak-tem

Appendix B. The target words with root-final gemination and no gemination employed in Experiment 3

hit-o-shash-ti	hit-a-ban-ti
hit-ba-sas-ti	hit-ba-sar-ti
hit-go-nan-tem	hit-ga-nav-tem
hit-go-dad-nu	hit-ga-mad-nu
hit-go-rar-nu	hit-ga-ash-nu
hit-bo-shash-ti	hit-ga-lash-ti
hit-ko-nan-ti	hit-ka-shar-ti
hit-lo-nan-ti	hit-la-bash-ti
hit-bo-dad-nu	hit-a-dam-nu
hit-bo-lal-nu	hit-ba-gar-nu
hit-mo-sas-tem	hit-ma-kad-tem
hit-ko-mam-tem	hit-ka-nas-tem
hit-o-nan-tem	hit-a-bal-tem
hit-la-kak-nu	hit-la-hat-nu
hit-ro-nan-ti	hit-ka-bad-ti
hit-go-shash-ti	hit-ga-rad-ti
hit-bo-nan-tem	hit-ba-rag-tem
hit-ko-faf-nu	hit-ka-bats-nu
hish-to-mam-ti	hit-ra-sak-ti
hish-to-vav-tem	hit-ra-tav-tem
hit-go-lal-nu	hit-ra-pak-nu
hit-ga-paf-tem	hit-ra-sham-tem
hish-to-lal-nu	hish-ta-lat-nu
hits-ta-nan-nu	hits-ta-rad-nu
hiz-da-kak-tem	hiz-da-man-tem
hish-to-kak-ti	hish-ta-par-ti
hits-to-faf-tem	hits-ta-mak-tem
hit-mo-dad-nu	hit-ma-sar-nu
hit-no-sas-tem	hit-na-gash-tem
hit-ya-dad-nu	hit-ha-dar-nu
hit-ma-mash-nu	hit-ma-sad-nu
hit-po-rar-ti	hit-pa-kad-ti

(*continued*)

hit-ro-tsats-nu	hit-ra-gaz-nu
hit-ro-mam-tem	hit-ro-kan-tem
hit-pa-lal-tem	hit-pa-lash-tem
hit-po-tsats-tem	hit-pa-rak-tem

References

Altamann, G. T. M., Dienes, M., & Goode, A. (1995). Modality independence of implicitly learned grammatical knowledge. *Journal of Experimental Psychology: Learning Memory and Cognition*, *21*, 899–912.

Balota, D., & Chumbly, J. (1984). Are lexical decisions a good measure of lexical access? The role of word frequency in the neglected decision stage. *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 340–357.

Berent, I. (2001). Can connectionist models of phonology assembly account for phonology? *Psychonomic Bulletin & Review*, *8*, 661–676

Berent, I., Bibi, U., & Tzelgov, J. (2000). *The autonomous computation of linguistic structure in reading: evidence from the Stroop task*. Paper submitted to the 41st meeting of the Psychonomic Society, New Orleans, LA.

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., 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. (2001). Co-occurrence restrictions on identical consonants in the Hebrew lexicon: are they due to similarity? submitted for publication.

Berent, I., Shimron, J., & Vaknin, V. (2001). Phonological constraints on reading: evidence from the Obligatory Contour Principle. *Journal of Memory and Language*, *22*, 1–22.

Bybee, J., & Moder, C. (1983). Morphological classes as natural categories. *Language*, *59*, 251–270.

Catford, J. (1977). *Fundamental problems in phonetics*, Edinburgh: Edinburgh University Press.

Chomsky, N. (1957). *Syntactic structures*. The Hague: Mouton.

Chomsky, N. (1980). Rules and representations. *Behavioral and Brain Sciences*, *3*, 1–61.

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

Choueka, Y. (1997). *Rav Milim*, Tel-Aviv: Stiematzky.

Clements, G., & Hume, E. (1995). The internal organization of speech sounds. In J. Goldsmith (Ed.), *The handbook of phonological theory* (pp. 245–306). Cambridge, MA: Blackwell.

Dart, S. (1991). Articulatory and acoustic properties of apical and laminal articulations. *UCLA Working Papers in Phonetics*, *79*, 1–55.

Daugherty, K., & Seidenberg, M. (1992). *Rules or connections? The past tense revisited*. Paper presented at the 14th annual conference of the Cognitive Science Society, Hillsdale, NJ.

Eddington, D. (2000). Analogy and the dual-route model of morphology. *Lingua*, *110*, 281–298.

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.

Even-Shoshan, A. (1993). [The new dictionary]. *Ha'milon ha'xadash*, Jerusalem: Kiryat-Sefer.

Everett, D. L., & Berent, I. (1998). An experimental approach to the OCP: evidence for violable identity constraint in Hebrew roots. Rutgers Optimality Archive. Available: http://ruccs.rutgers.edu/roa.html (ROA-235).

Fodor, J., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: a critical analysis. *Cognition*, *28*, 3–71.

Fromkin, V., & Rodman, R. (1993). *An introduction to language*. Fort Worth, TX: Harcourt Brace Javanovich College Publishers.

Gafos, A. (1998). Eliminating long-distance consonantal spreading. *Natural Language and Linguistic Theory*, *16*, 223–278.

Gafos, A. (1999). *The articulatory basis of locality in phonology*. New York: Garland.

Goldsmith, J. A. (1990). *Autosegmental & metrical phonology*. Cambridge, MA: Basil Blackwell.

Gomez, R. L., & Gerken, L. (1999). Artificial grammatical learning by 1-year olds leads to specific and abstract knowledge. *Cognition*, *70*, 109–135.

Hahn, U. N., & Nakisa, R. S. (2000). German inflection: single route or dual route? *Cognitive Psychology*, *41*, 313–360.

Halle, M., & Stevens, K. (1979). Some reflections on the theoretical bases of phonetics. In B. Lindlom & S. Öhman (Eds.), *Frontiers of speech communication research* (pp. 335–349). London: Academic Press.

Kenstowicz, M. (1994). *Phonology in generative grammar*. Cambridge, MA: Blackwell.

Ladefoged, P. M., & Maddieson, I. (1996). *The sounds of the word's languages*, Cambridge, MA: Blackwell.

Leben, W. (1973). *Suprasegmental phonology.* Dissertation, MIT, Cambridge, MA.

Marcus, G. (1998). Rethinking eliminative connectionism. *Cognitive Psychology*, *37*, 243–282.

Marcus, G. (2001). *The algebraic mind: integrating connectionism and cognitive science*. Cambridge, MA: MIT Press.

Marcus, G., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science*, *283*, 77–80.

McCarthy, J. (1979). *Formal problems in Semitic phonology and morphology.* PhD dissertation, MIT, Cambridge, MA.

McCarthy, J. (1986). OCP effects: gemination and antigemination. *Linguistic Inquiry*, *17*, 207–263.

McCarthy, J. (1994). The phonetics and phonology of Semitic pharyngeals. In P. Keating (Ed.), *Papers in laboratory phonology* (pp. 191–283). Vol. III. Cambridge: Cambridge University Press.

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., McClelland, J., Seidenberg, M., & 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* (4), 463–490.

Plunkett, K., & Marchman, V. (1993). From rote learning to system building: acquiring verb morphology in children and connectionist nets. *Cognition*, *48*, 21–69.

Plunkett, K., & Nakisa, C. (1997). A connectionist model of Arabic plural system. *Language and Cognitive Processes*, *12*, 807–836.

Prasada, S., & Pinker, S. (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes*, *8*, 1–56.

Prince, A., & Smolensky, P. (1997). Optimality: from neural networks to universal grammar. *Science*, *275* (14), 1604–1610.

Reber, A. S. (1969). Implicit learning and tacit knowledge. *Journal of Experimental Psychology*, *81*, 115–119.

Rhode, D., & Plaut, D. (1999). Language acquisition in the absence of explicit negative evidence: how important is starting small? *Cognition*, *72*, 67–109.

Rueckl, J., Mikolinski, M., Raveh, M., Miner, C., & Mars, F. (1997). Morphological priming, fragment completion, and connectionist networks. *Journal of Memory and Language*, *36*, 382–405.

Rumelhart, D. E., & McClelland, J. L. (1986). On learning past tense of English verbs: implicit rules or parallel distributed processing? In J. L. McClelland, D. E. Rumelhart & T. P. R. Group (Eds.), *Psychological and biological models. Parallel distributed processing: explorations in the microstructure of cognition* (pp. 216–271). Vol. 2. Cambridge, MA: Bransford Books/MIT Press.

Seidenberg, M., & McClelland, J. (1989). A distributed developmental model of word recognition and naming. *Psychological Review*, *96*, 523–568.

Yip, M. (1988). The Obligatory Contour Principle and phonological rules: a loss of identity. *Linguistic Inquiry*, *19*, 65–100.