Brain and Language **81,** 326–341 (2002)
doi:10.1006/brln.2001.2528, available online at http://www.idealibrary.com on IDEAL®

Identity Avoidance in the Hebrew Lexicon: Implications for Symbolic Accounts of Word Formation

Iris Berent

Florida Atlantic University

Published online November 27, 2001

Hebrew frequently exhibits geminates in the root but strictly constrains their location: Root-initial gemination is rare (e.g., *ssm*), whereas root-final gemination (e.g., *smm*) is frequent. Four experiments demonstrate that Hebrew speakers generalize this constraint to novel roots. When speakers are encouraged to form a triliteral root from a biconsonantal input (e.g., *sm*), they frequently reduplicate the root's final radical (e.g., *smm*), but not its initial radical (e.g., *ssm*). Likewise, the rejection of novel root foils with root initial geminates is easier than roots with final geminates. In both cases, speakers' performance is inexplicable by the statistical structure of the Hebrew language. Speakers' ability to freely generalize the constraint on root structure suggests that their linguistic competence appeals to mental variables. © 2001 Elsevier Science (USA)

Key Words: morphology; phonology; Hebrew; root; OCP; symbolic computation.

It is well established that speakers' knowledge of word structure is productive (Chomsky, 1980). For instance, upon hearing the novel verb *blix,* speakers instantly generate its past tense, *blixed.* The nature of these generalizations, however, is the subject of an ongoing debate. Symbolic accounts of cognition attribute linguistic generalizations to operations over variables—abstract placeholders, such as a verb stem, noun, syllable, and so on (Fodor & Pylyshyn, 1988; Marcus, 2001; Pinker & Prince, 1988; Pinker, 1999). For instance, the past tense is formed by concatenating a variable standing for the verb stem and the suffixed *d* (Pinker, 1999; see also Berent, Pinker, & Shimron, 1999; Kim, Pinker, Prince, & Prasada, 1991; Marcus, Brinkmann, Clahsen, Wiese, & Pinker, 1995). Because the past tense rule operates on variables (e.g., *verb stem*), not specific word instances (e.g., *blix*), the rule applies across the board, regardless of the word's sound or meaning. Conversely, associationist views of cognition consider variables obsolete. Speakers' productive use of language is largely explained by its statistical structure (Elman, Bates, Johnson, Karmiloff-Smith, Parisi, & Plunkett, 1996; Plaut, McClelland, Seidenberg, & Patterson, 1996; Seidenberg & McClelland, 1989; Seidenberg, 1997). For instance, the regular inflection of *blix* is attributed to its similarity to existing regular verbs (e.g., *blink*). The representations taking part in this process, however, are specific sound–spelling–meaning combinations (e.g., *blix*), not the abstract variables they may instantiate (e.g., *verb stem;*

This research was supported by an NIH FIRST award, 1R29 DC03277-017, to Iris Berent.

Address correspondence and reprint requests to Iris Berent, Department of Psychology, Florida Atlantic University, 777 Glades Road, P.O. Box 3091, Boca Raton, FL 33431-0991. Fax: (561) 297-2160. E-mail: iberent@fau.edu.

0093-934X/01 $35.00

see Plunkett & Nakisa, 1997; Rueckl, Mikolinski, Raveh, Miner, & Mars, 1997; Rumelhart & McClelland, 1986). The fierce debate surrounding the past tense clearly goes beyond the domain of word formation. What is at stake is the core question of how the mind works: Are some mental processes sensitive to the combinatorial structure of variables or is cognition explicable largely by the statistical structure of specific instances?

The work reported here presents a relatively new case study for examining this question. As in the past tense debate, this case study is taken from the domain of morphology. Unlike the case from inflection, however, this example concerns the representation of a morpheme in the lexicon, not its concatenation to other morphemes. The constraint examined here is the Obligatory Contour Principle (OCP)—a universal ban on the representation of identical elements in phonological representations (McCarthy, 1986). Obligatory Contour Principle effects are observed in the representation of a morpheme in the lexicon as well as postlexically in the output of phonological process (e.g., Goldsmith, 1990; Kenstowicz, 1994; Yip, 1988). One of the best known linguistic discussions of the lexical OCP concerns the structure of Semitic roots; hence, the following experiments assess OCP effects in Modern Hebrew. I first review the findings of four experiments demonstrating that Hebrew speakers constrain identity in the root morpheme and generalize this constraint to novel forms. The General Discussion examines the nature of the architecture implicated by speakers' performance.

Before describing the experimental findings, a brief introduction of Hebrew root structure is in order. Hebrew words are formed from a root and word pattern. The root is an abstract sequence of three consonants. For instance, *smm* (an example frequently used in McCarthy's 1986 classic article) is a root that indicates drugs. The word pattern is a template. It includes three slots for the consonants and provides the vowels and affixes. To form a word, speakers must insert the root into the word pattern. For instance, to form a verb from the root *smm*, speakers insert the root in the word pattern CiCeC (C stands for any consonant). The resulting verb, SiMeM (he drugged), is captured in Fig. 1. On this autosegmental account, distinct phonological constituents are segregated onto distinct tiers and anchored to a skeleton, an array of abstract placeholders for consonants and vowels.

Note that the root *smm* includes two adjacent identical consonants, i.e., geminates. Gemination is very frequent in Hebrew roots, but, like other Semitic languages, Hebrew constrains the location of geminates. Geminates are frequent at the end of the root (e.g., *smm*), but are extremely rare in its beginning (e.g., *ssm*). McCarthy (1986) accounts for this asymmetry in terms of the OCP. The OCP bans adjacent identical elements from lexical representations. The OCP thus prohibits the representation of

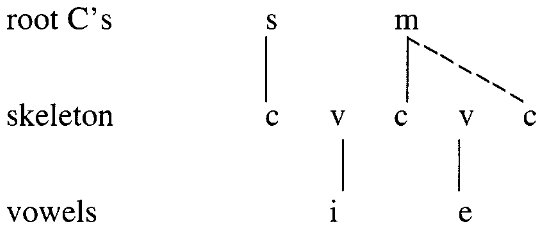

FIG. 1. The autosegmental structure of the verb SiMeM. The root consonants are represented on a separate tier, segregated from vowels. The geminates *mm* are formed by aligning the biconsonantal root *sm* with the word pattern CiCeC. Because the alignment proceeds from left to right, the rightmost consonant slot remains unfilled, allowing for the spreading of the phoneme *m* into this free slot. The doubly linked phoneme surfaces as root-final geminate *smm*.

roots such as *smm* in the lexicon. The representation of *smm* is *sm*, erasing the geminates from the lexicon. Root geminates (e.g., *smm*) may only be formed productively by the grammar. Specifically, to form a word from *sm*, it must be inserted into the word pattern in a left-to-right fashion. The remaining empty slot at the root's right edge may be filled by rightward spreading, resulting in the geminates *mm*. In what follows, I refer to this process as reduplication (Everett & Berent, 1998; Gafos, 1998; Rose, 2000). Because reduplication can only proceed rightward, geminates are expected to emerge root finally (e.g., *smm*), but not root initially (e.g., *ssm*). The OCP thus nicely accounts for the asymmetry in the distribution of root geminates found in Semitic.

The research presented here examines whether the constraint on root structure is active in the linguistic competence of modern Hebrew speakers. Specifically, this work addresses two questions. One question is whether Hebrew speakers constrain the location of geminates in the root in accord with the OCP. If the constraint on root structure is active, then speakers should extend this generalization to novel roots. Geminates should be acceptable only at the end of the root, but not in its beginning. A second question concerns the representation of this constraint. The OCP bans identical elements in the root. The representation of root and identity, however, requires two mental variables. The following investigation examines whether the constraint implicated in speakers' behavior appeals to these abstract variables or it is explicable by the statistical structure of the Hebrew language. I review two sources of experimental evidence for the OCP: a production task and lexical decision experiments. An extensive discussion of these findings may be found in Berent, Everett, and Shimron (2001a) and Berent, Shimron, and Vaknin (2001b).

EVIDENCE FROM THE PRODUCTION TASK

According to McCarthy (1986), root geminates are absent in the lexicon. Geminates (e.g., *smm*) may only be formed by reduplicating a biconsonantal lexical representation rightward (e.g., *sm* → *smm*). Gemination is extremely frequent in Hebrew roots. If McCarthy's account is correct, then speakers must routinely use reduplication in Hebrew word formation. Experiment 1 examines whether speakers employ reduplication to form root geminates (e.g., *smm*) from biconsonantal inputs (e.g., *sm*).

The task used here is an off-line production task. Participants are presented with a printed exemplar and a new root. Their task is to form a word from the novel root by analogy to the exemplar. For instance, participants were presented with the novel root *psm* and the exemplar PaʔaL. To form the word, participants must decompose the morphological structure of the exemplar PaʔL into the root, *pʔl*, and the word pattern, CaCaC; delete the root; and replace it with the novel root *psm*. The resulting word is PaSaM. Half of the trials in the experiment included triconsonantal roots, such as *psm*. These trials are not relevant for investigating the OCP and are discussed no further. Of main interest are the trials in which the novel root included only two consonants, e.g., *sm*.[1] These trials present participants with an alignment problem: The word pattern requires three consonants, whereas the root provides only two. How do speakers solve the alignment problem?

One possible solution is to violate the structure of the word pattern by deleting a consonant slot from the word pattern. The insertion of the root *sm* in the word pattern

[1] I use the root *sm* for illustration purposes because this root is frequently used in the linguistic discussion of OCP effects. The root *smm*, however, is an existing Hebrew root. The experiment used no existing root (or biconsonantal roots that may be turned into an existing root by reduplication).

TABLE 1
A Comparison of the Possible Correct Responses with Root-Initial vs Root-Final Geminates in the Production Task (Experiment 1)

			Possible geminate response	
Root	Word class	Exemplar	Root-initial	Root-final
sm	I	Pa-ʔaL	Sa-SaM	Sa-MaM
sm	II	maP-ʔi-Lim	maS-Si-Mim	maS-Mi-Mim
sm	III	hit-Pa-ʔaL-ti	hiS-ta-SaM-ti	hiS-ta-MaM-ti

CaC will yield forms such as SaM (which are acceptable, albeit irregular in Hebrew). Of particular interest is the class of solutions that adds a consonant to the root. Two questions are examined here. One is what kind of segment is added: a geminate (e.g., *smm* and *ssm*) or a new segment (e.g., *sml*)? If speakers possess a reduplication mechanism, then they should be likely to reduplicate one of the root's segments. Reduplication should specifically proceed rightward, resulting in forms such as *smm* rather than *ssm.*

A second question concerns the generality of the reduplication responses with respect to word structure. Recall that the OCP constrains the location of geminates in the root—an abstract variable. An alternative account, however, may attribute the asymmetry in the location of geminates to a constraint on their location in the word. To determine whether speakers' knowledge is sensitive to *root* or *word* structure, the experiment systematically varied the location of geminates in the word by using exemplars that differ in their morphological structure (see Table 1). In one class of exemplars, the root is unaffixed. If speakers were to align this word pattern with the novel root (e.g., *sm*) by reduplication, then root-initial geminates would invariably fall word-initially (e.g., SaSaM),[2] whereas root-final geminates would always fall word-finally (e.g., SaMaM). In contrast, in the second and third word classes, the root is ''sandwiched'' between a prefix and a suffix; hence, root-initial/-final geminates never occur at the beginning or end of the word. In the second word class, however, the first and second root radicals are truly adjacent (e.g., maSSiMim), whereas in the third word class, these radicals are separated by a vowel (e.g., hiStaSaMtem). If speakers constrain the location of geminates in the root in accord with the OCP, then their responses should reflect root-final, but not root-initial geminates, and this pattern should emerge regardless of the position of geminates in the word.

Method

Participants. Twenty-four native Hebrew speakers students at the School of Education at the University of Haifa served as participants.

Materials. Participants were presented with a printed list of roots and exemplars. They were asked to conjugate each root in analogy to a given exemplar. Forty-eight novel roots were used in the study. Half of these roots were biconsonantal (e.g., *sm*), and the other half were triconsonantal (e.g., *psm*). The triconsonantal roots were used as fillers. They were included in the experiment in order to encourage participants to produce words with three radicals to the biconsonantal roots (by either gemination or addition) rather than leaving this slot empty (and violating the word pattern). All novel roots consisted of consonant combinations that do not exist in Hebrew. Biconsonantal roots were further selected such that the reduplication of their initial or final radical does not yield an existing root.

[2] For viewing convenience, root consonants are notated in uppercase English letters. The Hebrew orthography does not discriminate between root and nonroot letters.

Participants were asked to conjugate each of the roots described above by analogy to a given exemplar. These exemplars were formed by conjugating the root PʔL in one of three classes of word patterns. The three word classes differed in terms of the surface transparency of the root in the word (see Table 1). In the first word class, the root was unaffixed; hence, the word's morphological structure was highly transparent. This class consisted of verbs in qal, piʔel, and puʔal in the singular past tense perfect form. In the second and third classes, the root was both prefixed and suffixed, hence, the word's morphological structure was more opaque. The second and third word classes differed, however, with respect to the surface adjacency of the root-initial bigram. In the second word class, the root-initial bigram was not separated by a full vowel (e.g., MaSSiMim). Members of the second class included the present tense of verbs in hifʔil and nouns in nifʔal. The third word class included verbs in hitpaʔel past tense. In this word class, the initial bigram was separated by at least a full vowel (e.g., hiStaSaMtem).

Each biconsonantal root (24 roots) and each triconsonantal root filler (24 roots) were paired with an exemplar in each of the three word classes, resulting in a total of 144 experimental trials.

Procedure. Participants were presented with a printed list. There were 144 lines in the list, each representing a separate trial. Each trial presented a novel root and a familiar word exemplar. Participants were asked to conjugate the new root in analogy to the exemplar and write down their response (including diacritic marks, used in Hebrew to specify vowels). The order of the trials in the list was random. The orthographic representation of the exemplars specified all their vowels using diacritic marks.

Coding scheme. Errors were failures to respond, the deletion of one of the root's radicals, or the use of an incorrect conjugation. Correct responses to the experimental roots were classified according to the following categories: (a) *Root-initial gemination:* gemination of the first root radical; (b) *Root-final gemination:* gemination of the second root radical; (c) *No-gemination:* a correct alignment of the root with the word pattern without filling the third consonant slot; and (d) *Addition:* a triconsonantal root consisting of the given biconsonantal root and an additional new radical.

Results and Discussion

The error rate in this experiment was low (M = 6.7%), and it was not significantly affected by word class.[3] The remaining correct responses included the addition of a new segment (M = 14.35% of the trials) and no gemination responses (M = 31.89% of the trials).[4] The most frequent response strategy, however, was gemination, occurring on 46.84% of the trials. The production of geminates was strongly constrained by their location in the root. Virtually all gemination responses occurred at the end of the root (M = 46.82% of the trials), but practically never in its beginning (M = 0.23% of the trials). This asymmetry resulted in a significant main effect of gemination type [$F_s(1, 23) = 31.13$, $MSE = 2509.79$; $F_i(1, 23) = 2388.93$, $MSE = 32.70$] in the ANOVA's (3 word class × 2 gemination type). The ANOVA's also yielded a gemination type × word class interaction [$F_s(2, 46) = 3.73$, $MSE = 74.79$; $F_i(2, 46) = 9.53$, $MSE = 30.16$].

Figure 2 plots the rate of total correct responses with root-initial vs root-final gemination. Planned comparisons indicated that the rate of root-final geminates was significantly more frequent than root-initial gemination in the first [$t_s(46) = 20.86$; $t_i(46) = 32.85$], second [$t_s(46) = 17.23$; $t_i(46) = 26.94$], and third [$t_s(46) = 17.87$; $t_i(46) = 28.36$] word classes. These results demonstrate that speakers produce geminates at the end of the root, but not in its beginning. The asymmetry in the location of geminates is further general with respect to their position in the word: Speakers refrain from root-initial geminations regardless of word position when the geminates are either word-initial (e.g., SaSaM) or internal (e.g., hiStaSaMtem). Hebrew speakers thus appear to constrain the location of geminates in the root in accord with the OCP.

[3] In this and all subsequent analyses, the significance level is .05.

[4] The relatively high rate of no-gemination responses indicates that, despite the clear preference of root-final gemination over root-initial gemination, speakers appear to dislike gemination altogether. Berent et al. (2001a) formally account for this finding as an identity aversion constraint, captured within the framework of Optimality Theory.

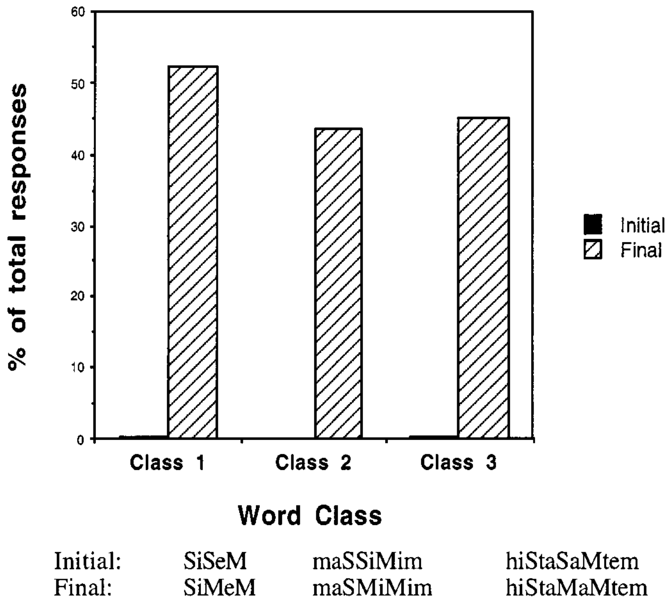

FIG. 2. The rate of production of root-initial vs root-final gemination in the production task as a function of word class.

EVIDENCE FROM LEXICAL DECISION

The production results demonstrate that speakers can form triconsonantal roots by reduplicating their underlying biconsonantal representation, as predicted by the OCP. The strong asymmetry in the location of geminates in the root, specifically, the absence of responses with root-initial geminates, is further consistent with the view of root-initial geminates as ill-formed. The off-line production findings, however, may be criticized as reflecting a metalinguistic problem solving strategy rather than linguistic competence. This metalinguistic explanation is not very likely: The constraint on root structure is not patent to speakers, and they are typically unable to provide an explanation for the aversion of root-initial geminates (Berent & Shimron, 1997). It is nevertheless desirable to extend and replicate these findings in on-line methods. Experiments 2–4 examine whether the constraint on root structure affects on-line performance using the lexical decision task.

Participants in each of these experiments were presented with 90 word targets and 90 nonword foils (see Table 2). The targets were existing familiar Hebrew words

TABLE 2
An Illustration of the Materials Used in the Lexical Decision Experiments (Experiments 2–4)

	Experiment 2	Experiment 3	Experiment 4
Nonwords			
Root-initial geminates	Gi-GuS	Ki-KaS-tem	hit-Ka-KaS-ti
Root-final geminates	Si-GuG	Si-KaK-tem	hiS-ta-KaK-ti
No geminates	Ri-GuS	Ni-KaS-tem	hit-Na-KaS-ti
Words			
Root-final geminates	Di-MuM	Si-NaN-tem	hit-Ba-SaS-ti
No geminates	Di-ShuN	Si-MaN-tem	hit-Ba-LaT-ti

formed from roots with either root-final or no gemination. Our main interest in this experiment concerns the nonword foils. These words were formed from roots with initial gemination, final gemination, or no gemination. To examine the generality of the constraint on root structure, these experiments manipulated the morphological structure of the stimuli. Each of the experiments used similar word structure for the targets and foils. In Experiment 2, the morphological structure of the stimuli was transparent, lacking any affixes. In Experiment 3, the root was followed by a suffix, and in Experiment 4, the root was ‘‘sandwiched’’ between a prefix and a suffix; hence, morphological structure was highly opaque. To assess the generality of the root-structure constraint across items and participants, Experiment 2 examines performance for one set of roots, whereas Experiments 3 and 4 extend this investigation to a second group of materials.

The design of these experiments addresses two questions. One question concerns speakers' sensitivity to the location of geminates in the root. Each experiment compares performance with root-initial and root-final geminates. These roots are equated for the presence of geminates and differ only in their location. If speakers constrain the location of geminates in the root, then roots with initial gemination should be easier to reject compared to roots with final gemination. If the constraint on geminates truly concerns the root, not merely the word, then this asymmetry should emerge in each of the three experiments, regardless of the position of geminates in the word.

A second question of interest concerns the representation of geminates. The findings reviewed so far indicate that speakers constrain the location of geminates. These observations, however, cannot determine whether this constraint is specific to gemination: Because root-initial gemination is rare, such stimuli may be discriminated from roots with final gemination simply by appealing to the frequency of the initial bigram, not necessarily the identity of its constituents. The dissociation between the structure of geminates and their frequency is difficult to achieve for root-initial geminates, whose initial bigram is typically rare. However, the representation of gemination may be examined by comparing root-final geminates and no-gemination controls. These roots were matched for bigram frequency; hence, any differences between them must specifically indicate sensitivity to the presence of geminates. If speakers represent the structure of geminates then they may discriminate between these two root structures despite their equation on their statistical properties.

Method

Participants. Three groups of native Hebrew speakers, students at the University of Haifa, served as participants in Experiments 2–4. Each group included 20 participants.

Materials. The materials in each experiment consisted of 90 words and 90 nonwords. The targets and foils shared the same word patterns (see Table 2). The word patterns used in Experiment 2 were the nominal word patterns CaCiC, CoCeC, and CaCuC. Experiment 3 used verbal patterns in PaʔaL and Piʔel followed by a suffix. In Experiment 4, all the word patterns were in hitpaʔ1 and were followed by a suffix. Targets words were formed from existing triliteral roots such that the combination of the root and word pattern corresponded to an existing Hebrew word. In contrast, the roots of the nonword foils consisted of three consonants that do not correspond to an existing Hebrew root.

The 90 nonword foils were generated from 30 trios of novel roots. Each trio included three root types: root-initial gemination (e.g., *kks*), root-final gemination (e.g., *skk*), and no gemination (e.g., *nks*). Experiment 2 employed one such set. Experiments 3 and 4 extended and replicated these findings for a new set of novel roots. The roots with final gemination and no gemination were equated for their summed positional bigram frequency using a database, including all the productive triconsonantal roots from the Even-Shoshan (1993) Hebrew dictionary (a total of 1449 roots). Positional bigram frequency was determined by counting the number of roots exhibiting a given radical bigram at the same root position. The mean positional bigram frequency was computed by adding the positional bigram frequency of the C1C2, C1C3, and C2C3 bigrams. For instance, the mean positional bigram frequency of the novel root *zpp* is 14 because there is 1 root sharing its initial bigram (*zpt*), 2 roots sharing its first and third consonants (*szp* and *nzp*), and 10 roots sharing its second and third radicals (*app, gpp, xpp, Tpp, kpp, lpp, cpp, rpp, spp, tpp*).[5] The mean positional summed bigram frequency of the roots with final gemination vs no geminates, respectively, was 12.2 (SD = 3.8) and 12.6 (SD = 4.1) in Experiment 2 and 12.3 (SD = 5.2) and 11.1 (SD = 5.1) in Experiments 3 and 4.[6]

Each experiment employed a set of 90 target words. Each set was generated from 45 pairs of existing Hebrew roots. The members of the pair shared the same word pattern, but differed in root structure. One member had root-final geminates, whereas the other member had a root with no geminates. Hebrew does not have a word frequency count, hence, it was impossible to match the frequency of words with final vs no identical consonants. These data thus do not allow for inferences regarding the effect of root structure on the identification of familiar words.

Procedure. At the beginning of each trial, a fixation point consisting of four dots appeared at the center of the computer 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 2 s elapsed. Participants were asked to indicate whether the string of letters corresponds to an existing Hebrew word by pressing one of two keys. Slow responses (responses slower than 1500 ms) and inaccurate responses received negative feedback from the computer in the form of a tone and a computer message. The experiment initiated with the practice stimuli followed by the experimental trials. Participants were tested individually. Each participant was presented with a different random order of the experimental trials.

Results of Experiments 2–4

The effect of root type on response latency was assessed separately in each of the experiments by means of one-way ANOVA's by participants (F_s) and items (F_i).[7] Root structure did not significantly affect response latency or accuracy for target words.[8] Recall that, in the absence of a word frequency count for Hebrew, it was impossible to match target words for frequency; hence, this null effect may be partly due to differences in surface word frequency. In contrast, a significant effect of root structure was observed for the nonword foils.[9] The effect of root type was significant in Experiment 2 [$F_s(2, 38) = 12.73$, $MSE = 914.87$; $F_i(2, 58) = 5.74$, $MSE = 3584.43$], Experiment 3 [$F_s(2, 38) = 10.77$, $MSE = 827.30$; $F_i(2, 56) = 8.49$, $MSE = 1552.18$], and marginally so in Experiment 4 [$F_s(2, 38) = 3.45$, $MSE = 956.07$, $p < .05$; $F_i(2, 56) = 2.62$, $MSE = 2269.80$, $p < .09$]. A similar analysis on response accuracy revealed a significant effect only in Experiment 2 [$F_s(2, 38) = 13.18$, $MSE = .002$; $F_i(2, 58) = 4.96$, $MSE = .007$].

Figure 3 summarizes the results of Experiments 2–4 by plotting correct responses to nonword foils as a function of root structure and the word structure used in each experiment. Response accuracy as a function of root type are summarized for the three experiments in Table 3. The results were rather consistent across experiments. Root-initial gemination was rejected significantly faster than root-final gemination in each of the three experiments [Experiment 2: $t_s(38) = 4.43$, $t_i(58) = 3.08$; Experiment 3: $t_s(38) = 4.58$, $t_i(56) = 4.09$; Experiment 4: $t_s(38) = 2.21$, $t_i(56) = 2.15$]. The same pattern also emerged in the accuracy data of Experiment 2 [$t_s(38) = 4.81$; $t_i(58) = 2.96$]. Speakers thus constrain the location of geminates in the root regardless of its location in the word.

[5] The letters ''t'' and ''T'' stand for the Hebrew letters ''taf'' and ''tet'', respectively. ''C'' indicates the sound ''ts.''

[6] The frequency counts reported here are based on a revision of the database used by Berent, Vaknin, and Shimron (2001b); hence, these means differ slightly from those listed there.

[7] To eliminate the effect of outliers, these analyses excluded correct response latencies falling 2.5 standard deviations beyond the mean response latencies for words and nonwords. This procedure resulted in the elimination of less than 3.3% of the correct responses in Experiments 2–4.

[8] Word response latency and accuracy, respectively, were 665 ms, 94%, in Experiment 2; 679 ms, 93.3%, in Experiment 3; and 687 ms, 94.8%, in Experiment 4.

[9] One of the foil trios used in Experiments 3 and 4 was excluded from all analyses because one of its members corresponded to an existing root (shpd).

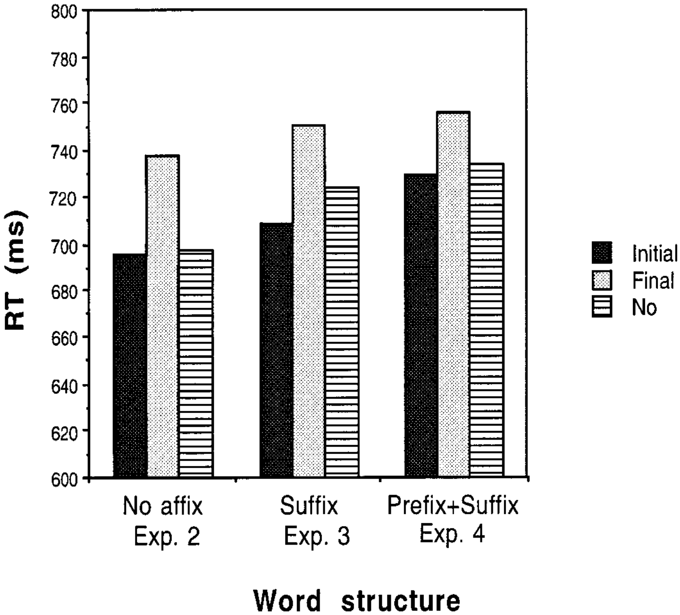

FIG. 3. Response latency to nonword foils in the lexical decision experiments (Experiment 2–4) as a function of root structure and word structure.

TABLE 3
Response Accuracy in the Rejection of Nonword Foils a Function of Root Structure in Experiments 2–4

	Experiment 2	Experiment 3	Experiment 4
Root-initial gemination	95.4	94.8	94.1
Root-final gemination	88.9	93.7	95.3
No gemination	94.2	95.6	93.3

To determine whether this constraint specifically concerns the structure of geminates, let us now turn to compare novel roots with final geminates and no-gemination controls. The contrast was significant in Experiment 2 [$t_1(38) = 4.31$; $t_2(58) = 2.77$] and in Experiment 3 [$t_1(38) = 2.92$; $t_2(56) = 2.48$] and marginally significant in Experiment 4 [$t_1(38) = 2.33$; $t_2(56) = 1.75$, $p = .08$]. Experiment 2 further yielded an increase in error rate for root-final gemination compared to no-gemination controls [$t_1(38) = 3.95$; $t_2(58) = 2.41$]. The consistent discrimination between root-final gemination and no gemination despite their equation for their statistical properties demonstrates that speakers are sensitive to the structure of geminates, namely their identity.[10]

[10] Although as a group, roots with final gemination and no-gemination controls did not differ significantly on their frequency, there were occasionally some frequency imbalances among specific trio members. These imbalances cannot account for the rejection of root with final geminates, as the finding remains even for trio members whose members do not differ from each other by more than one frequency count. This pattern was significant in Experiments 3 [$\Delta = 46$ ms, $t_s(38) = 3.07$; $t_i(16) = 2.50$] and 4 [$\Delta = 41$ ms, $t_s(38) = 2.72$; $t_i(16) = 2.4$] and marginally so in Experiment 2 [$\Delta = 54$ ms, $t_s(38) = 4.17$; $t_i(26) = 1.6$, $p = .12$]. The difficulty in the rejection of root-final geminates also is inexplicable by their phonological similarity to target words with root-final gemination. A post hoc analysis examined the similarity between targets and foils by dividing the summed positional frequency of each letter in target roots of a given structure by its summed positional frequency in foil roots of the same structure. For

The difficulty in responding to root-final geminates may be explained by the properties of the lexical decision task and the OCP. The lexical decision task requires fast discrimination between stimuli based on their appearance as wordlike (e.g., Balota & Chumbly, 1984). One evidence for ''wordhood'' may be coming from the grammar: Words generated by productive grammatical processes may appear more wordlike than words whose structure bears no such hallmark of the grammar. If the OCP is active, then gemination must be produced by a grammatical process of reduplication; hence, roots with geminates may be perceived as more wordlike than roots with no gemination.

The perception of gemination as indicating ''wordhood'' can also explain the absence of significant differences between root-initial gemination and no gemination. Rating experiments consistently show that root-initial gemination is considered significantly less acceptable than roots with no gemination (Berent & Shimron, 1997; Berent et al., 2001a). Why does the ill formedness of such roots not facilitate their rejection compared to no-gemination roots? The answer to this question comes, again, from the deleterious effect of gemination. The demand for fast discrimination between targets and foils may encourage a superficial analysis of global familiarity with the stimulus. Because roots with geminates are perceived as wordlike, the very presence of geminates in roots with initial geminates may impair their rejection. Conversely, a more careful analysis of the location of these geminates will indicate ill-formedness, hence, facilitate foil rejection. A comparison of initial-gemination with no-gemination roots thus confounds the presence of geminates with their location. This comparison pits the deleterious effect of gemination against the facilitation resulting from their illicit location. The similar performance with root-initial gemination and no-gemination roots may reflect the mutual cancellation of these conflicting forces. In contrast, the comparison of root-initial gemination with root-final gemination adequately controls for the deleterious effect of gemination. Such comparison yields strong evidence for the unconfounded effect of geminates' location.

The lexical decision results suggest that speakers are sensitive to the presence of identical elements in the root, regardless of their position in the word. The constraint on root identity further emerged even when the geminates were separated by intermediate vowel phonemes, encoded by a vowel letter (e.g., Experiment 3). These results suggest that the constraint on identity operates across intermediate vowels. Speakers thus decompose the root from the word pattern in on-line word identification and constrain the location of identical elements within the root.

---

instance, among the target roots with final gemination, there were two roots exhibiting the letter *b* in the first position, one root exhibiting that letter in the second position, and one root with that letter in the third position. For novel roots with root-final gemination, the occurrence of the letter *b* was three, four, and four times, respectively, for the same three positions. Accordingly, the target-to-foil ratio for the occurrence of *b* was 0.67, 0.25, and 0.25 for positions 1–3, respectively (a mean of 0.39). Similar letteroverlap ratios were also obtained for targets and foils with no gemination. If the impairment in the rejection of foil roots with final gemination is due to their greater similarity to final-gemination targets, then the target/foil overlap ratio should be higher for stimuli with final gemination compared to no gemination. Contrary to this prediction, the overlap ratios were literally identical in Experiment 2 ($M = 1.3$) and quite similar in Experiments 3 and 4. Specifically, the mean overlap for *smm* and *psm* type roots were 1.7 and 1.4, respectively, and they did not differ significantly [$F(1, 30) = 1.3$, $p = .56$, *ns*].

A slightly different explanation attributes the difficulty in rejecting roots with final gemination to the frequency of smm-type roots among the experimental targets. Berent et al. (2001b) have investigated and rejected the root-type explanation by demonstrating that the difficulty with root-final geminates is observed even after a very short exposure to the experimental list, within the first third of the experimental trials. It is unlikely that such short exposure was sufficient for the acquisition of strategies specific to the structure of the experimental list. Regardless of its source, however, a sensitivity to the structure of root types indicates the representation of identity.

GENERAL DISCUSSION

Four experiments examined whether Hebrew speakers constrain the location of geminates in the root morpheme. Experiment 1 demonstrated that when speakers are encouraged to form a triconsonantal root from a biconsonantal input (e.g., *sm*), they reduplicate the root's final radical (e.g., *smm*), but not its initial radical (e.g., *ssm*). Experiments 2–4 further showed that nonword foils with root-initial gemination (e.g., *ssm*) are easier to reject than nonword foils with root-final gemination (e.g., *smm*), a finding that agrees with the view of root-initial gemination as ill formed. The lexical decision experiments also suggest that speakers are sensitive to the presence of geminates in the root and are able to discriminate roots with final gemination and no gemination even when these stimuli are matched for their bigram frequency. In conjunction, these experiments demonstrate that speakers constrain the location of geminates in the root and apply this constraint on-line, even when the task does not require that they attend to the phonological structure of the root.

The results of these experiments carry several implications. One implication concerns the representation of morphological structure. The finding that speakers constrain the location of adjacent identical consonants regardless of word position, and despite the presence of intermediate vowels and affixes, suggests that the domain of this constraint is a morphological constituent rather than the surface word. This outcome is consistent with McCarthy's (1986) proposal that the domain of the OCP is the root morpheme. The representation of the root is further implicated by previous results demonstrating that the root is decomposed from the word pattern in on-line reading (e.g., Bentin & Feldman, 1990; Deutsch, Frost, & Forster, 1998; Feldman & Bentin, 1994; Feldman, Frost, & Pnini, 1995; Frost, Forster, & Deutsch, 1997). Although the perceived adjacency of root consonants is naturally handled by defining the domain of the OCP as the root, these findings may also be captured by the proposal that the OCP operates across intermediate vowels within the morphological domain of the stem (Rose, 2000). Regardless of the formal model proposed to capture the findings, these results clearly indicate that speakers constrain the identity of root consonants.

The production of triliteral roots by reduplicating biconsonantal roots and the difficulty in rejecting root-final geminates in the lexical decision task further suggest that identity is formed by the grammar. These observations are consistent with McCarthy's (1986) proposal that adjacent identical root consonants are erased from the lexicon. This account, however, has been recently subject to some debate in the linguistic literature. Critiques of this proposal have argued that the restriction on root identity is merely the limiting case of a more general, statistical constraint on similarity (e.g., Bat-El, 1994; Frisch, Broe, & Pierrehumbert, 1997; Pierrehumbert, 1993). Proponents of the similarity account may thus attribute the present experimental findings to the feature similarity of root radicals, not specifically their identity. Although there is convincing evidence for the existence of a ban on similar (e.g., homorganic) root radicals (e.g., Frisch, Broe, & Pierrehumbert, 1997; Greenberg, 1950; McCarthy, 1994), these findings are insufficient to support the similarity explanation as an alternative to the constraint on identical root consonants. The principal support for the similarity account is the analyses of Pierrehumbert and colleagues (Pierrehumbert, 1993; Frisch et al., 1997), demonstrating that the distribution of triliteral Arabic roots is explicable by their feature similarity. These analyses, however, exclude roots with final geminates; hence, the similarity proposal currently does not account for their distribution. In fact, a similarity account for the co-occurrence of root geminates would incorrectly predict that root-final geminates (e.g., *smm*) are less frequent than homorganic consonants (e.g., *smp*), a finding that is inconsistent with the structure of the Arabic lexicon (Greenberg, 1950). The experimental results presented here suggest that the constraint on similar root radicals must be complemented by a constraint on identical radicals, such as the OCP. The domain of this constraint must further correspond to a morphological constituent rather than the word. Interestingly, this phonological constraint is also observed in a silent reading task, lexical decision. This suggests that the representation assembled in silent reading is shaped by reader's phonological competence (e.g., Berent & Perfetti, 1995). Most importantly, speakers can freely generalize this constraint to novel roots.

What kind of mental architecture is required in order to represent this constraint? The introduction to this article outlined two conflicting accounts of linguistic generalizations. According to the symbolic view, the mind has the capacity to represent abstract variables and operate over variables (Fodor & Pylyshyn, 1988; Marcus, 2001; Pinker & Prince, 1988; Pinker, 1999). Conversely, associationist accounts of cognition attribute the constraint on word structure to the statistical structure of specific instances in the language (Elman et al., 1996; Plaut et al., 1996; Rumelhart & McClelland, 1986; Seidenberg & McClelland, 1989; Seidenberg, 1997). The constraint on root geminates has been so far discussed in a manner that implicitly presupposes the representation of two variables: root and geminates.[11] The root morpheme is a variable that can stand for any combination of three consonants. Likewise, geminates are represented as the copying of a variable, XX. The experimental findings appear compatible with this account. But do they specifically *require* the representation of variables? Is the constraint on root geminates explicable by the statistical structure of the Hebrew language? Can multilayer networks that eliminate operations over variables account for these results?

The answer to this question is bound to be highly controversial, especially given the absence of existing associationist accounts for the OCP. Ultimately, the adequacy of associationist accounts for the OCP can only be determined by future implementations. Several observations, however, cast some doubt on the potential success of this approach. To account for the experimental findings, associationist networks must be able to acquire the constraint on root structure from exposure to Hebrew words. Such models must be formulated without appealing to variables. The elimination of operations over variables is critical here. The question of whether the mind has a symbolic architecture should not be equated with the question of whether connectionism is adequate. Connectionism is a computational approach that may or may not include variables; hence, connectionism as such is not necessarily an alternative to the hypothesis that the mind manipulates variables (Marcus, 2001; Pinker & Prince, 1988). Conversely, the symbolic hypothesis is perfectly compatible with statistical learning defined over variable combinations (e.g., learning the frequency of root types of the form ABB, AAB, or ABC). To falsify the symbolic hypothesis, one cannot simply show that a given behavior is captured by a connectionist network. Instead, it is necessary to show that the behavior is captured without incorporating variables. In particular, an associationist account for the constraint on root structure must handle the experimental findings without representing the root or identity by variables.

Consider first the appeal to the root variable. The experimental results demonstrate that the root serves as the domain of the constraint on gemination: The location of

[11] The appeal to variables does not uniquely hinge on McCarthy's (1986) formal account. Gafos (1998) and Rose (2000) suggest that identity is formed by a formal process of reduplication rather than by long-distance spreading. The analysis of Rose (2000) further favors the stem, rather than the root, as the domain of the constraint. These revisions to McCarthy's proposal do not eliminate the appeal to variables. Like spreading, reduplication is a process that copies variables ($X \rightarrow XX$). Like the root, the stem is a variable that is defined regardless of the properties of specific instances. For simplicity, the following discussion is framed in reference to the root. The same arguments apply for the stem domain.

geminates is constrained relative to the root. Associationist accounts of cognition eliminate variables; hence, they cannot represent this domain. To account for the findings presented here, such models must acquire the constraint on the location of geminates from the distribution of subword units. To do so, these models must be able to decompose words into rootlike units, identify these units as members of a single class and infer the regularities that are common to its members. Root decomposition may well be attained by associationist models. Hebrew affixes typically have some well-defined phonological and semantic properties that could be used to distinguish them from the root (e.g., Plaut & Gonnerman, 2000). Decomposition, however, is only the first step in acquiring a constraint on root structure. To learn the constraint on geminates, a network must identify all roots as members of a single class and abstract the regularities common to its members. Roots are only identifiable by their formal abstract category: They share no similar semantic, orthographic, or phonological features. It is thus unclear whether associationist accounts can abstract the constraint on root structure from the distribution of wordlike units. It should be noted, however, the acquisition of the root domain from words has not been subject to formal investigation, hence, the ability of associationist accounts to abstract the root awaits further research.

Let us assume, however, that an associationist network (i.e., a network lacking operations over variables) is presented with the set of all Hebrew roots. Can such a network learn that ABB roots include a geminate, BB? Can it distinguish geminate from nongeminate bigrams and constrain their location in the root? The representation of identity has been recently subject to close scrutiny; hence, the performance of associationist accounts in this task can be stated quite specifically. However, before examining whether associationist networks can represent identity, let us first consider whether speakers do so.

The experimental findings provide ample evidence that speakers represent the identity of geminates. First, the lexical decision experiments demonstrate that speakers discriminate between roots with final geminates and no-gemination controls despite their equation for bigram frequency. A distinction between these two root types can only be explained by the structure of geminates, their identity. Additional support for the view of geminates as formed by variable copying comes from the comparison of geminate and nongeminate responses in the production task. Associationist accounts must view the formation of geminates as indistinguishable from the formation of nongeminate responses. The segment that is added (geminate vs nongeminate) should depend on the frequency of the resulting root in the language. Each of the radicals in the biconsonantal roots used in Experiment 1 can combine with any 1 of 19 other radicals (e.g., *sm* → *smb, smg, smd, sml, smk, smt,* and *smr*), only 2 of which happen to be geminates (e.g., *smm* and *ssm*). All things being equal, the frequency of geminate responses should be lower than nongeminate responses. To be more specific in this prediction, Berent et al. (2001a) calculated the expected frequency of geminate vs addition responses to their materials by summing the bigram frequency of all possible triliteral responses. The expected frequency of geminate responses compared to the total possible responses was 4%. In contrast, the observed frequency of geminate responses was 76%. The production of geminates is thus clearly inexplicable by the statistical structure of root instances in the language. Speakers generalize the constraint on geminates regardless of the strength of the statistical evidence for such responses. Further research (Berent, Marcus, Shimron, & Gafos, 2001c) demonstrates that speakers freely generalize the constraint on root structure even when the language provides no statistical evidence for either root-initial or root-final geminates because the geminates are formed from novel phonemes, including phonemes with phonetic feature values that never occur in Hebrew.

In contrast to people, associationist networks suffer from some principled limitations in exhibiting such generalizations. Marcus (1998a, 1998b) formally demonstrated that multilayer perceptrons that lack operations over variables cannot generalize functions such as identity outside the space of features on which they were trained. It is important to note that these limitations do not concern connectionist networks as a whole: Marcus (2001) discusses numerous ways in which multilayer perceptrons can incorporate operations over variables and demonstrates that such networks can successfully generalize the identity function beyond the space of training features. In the absence of operations over variables, however, multilayer perceptrons fail to freely generalize the identity function (for converging evidence, critiques, and rebuttals, see Altmann & Dienes, 1999; Christiansen & Curtin, 1999; Christiansen, Conway, & Curtin, 2000; Dominey & Ramus, 2000; Eimas, 1999; Marcus, 1999a, 1999b, 2001; McClelland & Plaut, 1999; Negishi, 1999; Seidenberg & Elman, 1999; Shastri, 1999). The conjunction of the present experimental findings and Marcus's computational work suggests that there are some fundamental differences between the ability of humans and associationist networks (i.e., networks lacking operations over variables) to generalize a constraint on identical elements. Humans can extrapolate the constraint regardless of whether the features of a particular instance were trained, whereas the network can only interpolate the constraint within the space of trained features. The ability of humans to extend generalization to any item, regardless of familiarity with its features, agrees with the view of such generalizations as appealing to mental variables. The principled failures of associationist networks to capture such generalizations suggests that the representation of variables may be necessary to account for linguistic competence.

REFERENCES

Altmann, G. T. M., & Dienes, Z. (1999). Rule learning by seven-month-old infants and neural networks. *Science,* **284,** 977.

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: Learning, Memory, and Cognition,* **10,** 340–357.

Bat-El, O. (1994). Stem modification and cluster transfer in Modern Hebrew. *Natural Language and Linguistic Theory,* **12,** 571–596.

Bentin, S., & Feldman, L. B. (1990). The contribution of morphological and semantic relatedness in visual word recognition: Evidence from Hebrew. *Quarterly Journal of Experimental Psychology,* **42,** 693–711.

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., Everett, D., & Shimron, J. (2001a). Do phonological representations specify formal variables?: Evidence from the Obligatory Contour Principle. *Cognitive Psychology,* **42**(1), 1–60.

Berent, I., Shimron, J., & Vaknin, V. (2001b). Phonological constraints on reading: Evidence from the Obligatory Contour Principle. *Journal of Memory and Language,* **44**(4), 644–665.

Berent, I., Marcus, G., Shimron, J., & Gafos, D. (2001c). The scope of linguistic generalizations: Evidence from Hebrew word formation. *Cognition,* in press.

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.

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

Christiansen, M. H., Conway, C. M., & Curtin, S. (2000). A connectionist single-mechanism account of rule-like behavior in infancy. In L. Gleitman & A. Joshi (Eds.), *The Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society* (pp. 83–88). Mahwah, NJ: Erlbaum.

Christiansen, M. H., & Curtin, S. L. (1999). The power of statistical learning: No need for algebraic rules. In M. Hahn & S. C. Stones (Eds.), *Proceedings of the Twenty-First Annual Conference of the Cognitive Science Society* (pp. 114–119). Mahwah, NJ: Erlbaum.

Deutsch, A., Frost, R., & Forster, K. I. (1998). Verbs and nouns are organized and accessed differently in the mental lexicon: Evidence from Hebrew. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **24**(5), 1238–1255.

Dominey, P. F., & Ramus, F. (2000). Neural network processing of natural language. I. Sensitivity to serial, temporal and abstract structure in the infant. *Language and Cognitive Processes,* **15**(1), 87–127.

Eimas, P. (1999). Do infants learn grammar with algebra or statistics. *Science,* **284,** 433.

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). *Ha'milon ha'xadash [The new dictionary]* Jerusalem: Kiryat Sefer.

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

Feldman, L. B., & Bentin, S. (1994). Morphological analysis of disrupted morphemes: Evidence from Hebrew. *Quarterly Journal of Experimental Psychology,* **47,** 407–435.

Feldman, L. B., Frost, R., & Pnini, T. (1995). Decomposing words into their constituent morphemes: Evidence from English and Hebrew. *Journal of Experimental Psychology: Learning, Memory and Cognition,* **21,** 947–960.

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

Frisch, S., Broe, M., & Pierrehumbert, J. (1995). *Similarity and phonotactics in Arabic. Rutgers Optimality Archive* (MS ROA-223), http://ruccs.rutgers.edu/roa.html.

Frost, R., Forster, K., & Deutsch, A. (1997). What can we learn from the morphology of Hebrew?: A masked-priming investigation of morphological representation. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **23,** 829–856.

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

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

Greenberg, J. (1950). The patterning of morphemes in Semitic. *Word,* **4,** 196–208.

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

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.

Marcus, G. (1998a). Can connectionism save constructivism? *Cognition,* **66,** 152–182.

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

Marcus, G. (1999a). Connectionism: With or without rules? Response to J. L. McClelland and D. C. Plaut. *Trends in Cognitive Sciences,* **3,** 168–170.

Marcus, G. (1999b). Reply to Seidenberg and Elman. *Trends in Cognitive Sciences,* **3,** 289.

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

Marcus, G., 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.

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

McClelland, J., & Plaut, D. (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.

Pierrehumbert, J. (1993). Dissimilarity in Arabic verbal roots. *North Eastern Linguistic Society,* **23,** 367–381.

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., & Gonnerman, L. (2000). Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? *Language and Cognitive Processes,* **15**(4–5), 445–485.

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

Rose, S. (2000). Rethinking geminates, long-distance geminates and the OCP. *Linguistic Inquiry,* **31,** 85–122.

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.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2, pp. 216–271). Cambridge MA: Bransford Book/MIT Press.

Seidenberg, M. S. (1997). Language acquisition and use: Learning and applying probabilistic constraints. *Science,* **275,** 1599–1603.

Seidenberg, M. S., & Elman J. L. (1999). Networks are not ‘hidden rules’. *Trends in Cognitive Sciences,* **3,** 288–289.

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

Shastri, L. (1999). Infants learning algebraic rules. *Science,* **285,** 1673–1674.

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