Cognitive Psychology **42,** 1–60 (2001)
doi:10.1006/cogp.2000.0742, available online at http://www.idealibrary.com on IDEAL®

Do Phonological Representations Specify Variables? Evidence from the Obligatory Contour Principle

Iris Berent

Florida Atlantic University

Daniel L. Everett

SIL International

and

Joseph Shimron

School of Education, University of Haifa, Haifa, Israel

Mental variables are central to symbolic accounts of cognition. Conversely, according to the pattern associator hypothesis, variables are obsolete. We examine the representation of variables by investigating the Obligatory Contour Principle (OCP, McCarthy, 1986) in Hebrew. The OCP constrains gemination in Hebrew roots. Gemination is well formed at the root's end (e.g., SMM), but not in its beginning (e.g., SSM). Roots and geminates, however, are variables; hence, according to the pattern associator view, the OCP is unrepresentable. Three experiments demonstrate that speakers are sensitive to the presence of root gemination and constrain its location. In forming words from novel biconsonantal roots, speakers prefer to reduplicate the root's final over its initial radical, and they rate such outputs as more acceptable. The avoidance or rejection of root-initial gemination is independent of its position in the word and is inexplicable by the statistical frequency of root tokens. Our results suggest that linguistic representations specify variables. Speakers' competence, however, is governed by violable constraints. © 2000 Academic Press

---

This research was supported by the National Institute of Health's NRSA (F 32 DC 000186-02) and FIRST awards (1R29 DC03277-02) to Iris Berent. We are indebted to Gary Marcus and Steve Pinker for discussions of this research. We also thank Allan Nash, for statistical consultation, and Addie Adams, Ruben Bouissa, Josef Degani, and Vered Vaknin for their technical assistance.

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. E-mail: iberent@fau.edu.

0010-0285/00 $35.00

Productivity is a defining feature of natural language. The knowledge of language entails the ability to produce and comprehend novel utterances (Chomsky, 1980). Generative linguistics attributes the productivity of language to the representation of mental variables. Variables are abstract placeholders. A variable (e.g., *Noun*) can enumerate a large number of tokens, including familiar tokens (e.g., *house* and *cat*) as well as novel instances (e.g., *blix;* Marcus, 1998a, 1998b, in press; Pinker & Prince, 1988). Variables could further be used to define formal relationships across a universal class of potential tokens. For instance, using the variable X to refer to any syllable, one could capture the structure of the sequences *gaga, dada, papa*, namely their identity, XX. Because variables treat all tokens as equivalent, ignoring their idiosyncrasies, variables allow to extend generalizations to novel tokens, regardless of their similarity to known items. Indeed, 7-month infants trained on XYY sequences (e.g., *galala* and *badada*) can discriminate novel XYY sequences (e.g., *wofefe*) from novel XXY sequences (e.g., *wowofe;* see Marcus, Vijayan, Bandi, Rao, & Vishton 1999).

Variables are central to symbolic accounts of cognition. On the symbolic hypothesis, mental states are determined by the constituent structure of variables (Fodor & Pylyshyn, 1988; Fodor & McLaughlin, 1990; Pylyshyn, 1986). It is the syntactic structure of mental representations that permits physical machines to exhibit systematicity and productivity (Fodor & Pylyshyn, 1988). A radical alternative to the symbolic approach is offered by eliminative connectionism. On this view, formal variables and their syntactic relations play no role in mental processes (Elman, 1993; Elman, Bates, Johnson, Karmiloff-Smith, Parisi, & Plunkett, 1996; Hare & Elman, 1995; Plaut, McClelland, Seidenberg, & Patterson, 1996; Plunkett & Marchman, 1993; Rueckl, Mikolinski, Raveh, Miner, & Mars, 1997; Rumelhart & McClelland, 1986; Seidenberg, 1987; 1997; Seidenberg & McClelland, 1989). Mental states are determined by context-sensitive associations between specific tokens and by the network's particular activation dynamics. Thus, the knowledge of language and its productive use is largely explicable in terms of the statistical properties of the linguistic input. The following investigation examines this claim.

Before we discuss our evidence, a brief explanation of our terminology is in order. In what follows, we use the term ‘‘pattern associator’’ to refer to the class of cognitive models that eliminate variables. The pattern associator hypothesis should not be equated with connectionism. We do not question the principled adequacy of connectionism in capturing human cognition, nor do we contrast connectionist and nonconnectionist formalisms. Instead, we compare two classes of potential (connectionist) models. One class embodies mental variables, whereas the other class eliminates them. The ability of some connectionist networks to implement symbolic functions is well known (Hornik, Stinchcombe, & White, 1989; Siegelman & Sontag, 1995). Existing connectionist models that implement variables include the analogical reasoning model by Hummel and Holyoak (1997) and Prince and Smolensky's (1997) optimality theory, an account of linguistic competence based on the principles of Harmony theory (Smolensky, 1995; for further discussion see Marcus, 1998a; in press). In contrast, a large group of popular connectionist models eliminates mental variables (e.g., Elman, 1993; Hare & Elman, 1995; Plaut et al. 1996; Plunkett & Marchman, 1993; Rueckl et al., 1997; Rumelhart & McClelland, 1986; Seidenberg & McClelland, 1989). This class is not limited to any particular model brand or architecture (e.g., feedforward networks, single recurrent networks, etc.). Common to these models is the assumption that variables are not represented; hence, they play no role in human cognition. Implementational and eliminative connectionism thus present radically different cognitive hypotheses. It is the cognitive hypothesis, tacit in these classes of models, that is the center of our investigation. Are variables necessary for explaining linguistic productivity?

A familiar, highly contentious, case for studying the role of linguistic variables is inflectional morphology. We first summarize some of the key findings in this area. We then present a new case study from morphophonology, a constraint on root structure known as the Obligatory Contour Principle (McCarthy, 1986). Our experiments exploit this phenomenon for examining the role of variables in linguistic representations.

THE REPRESENTATION OF VARIABLES: EVIDENCE FROM INFLECTIONAL MORPHOLOGY

The role of variables has been extensively examined by research on inflectional morphology and, in particular, verb past-tense inflection. According to the words/rules view (Pinker, 1999; see also Kim, Pinker, Prince, & Prasada, 1991; Kim, Marcus, Pinker, Hollander, & Coppola; 1994; Marcus, Brinkmann, Clahsen, Wiese, & Pinker, 1995; Pinker, 1991, 1994, 1999; Prasada & Pinker, 1993), past-tense inflection is achieved by two mechanisms. Irregular inflection (e.g., *come–came*) is generated by an associative process whose operation is highly sensitive to the properties of specific tokens. In contrast, regular inflection (e.g., *like–liked*) is formed by a simple rule that concatenates the verb stem and the past-tense inflectional suffix. Regular inflection thus operates on variables (e.g., verb stem) rather than on the tokens they enumerate (e.g., *like*). Variables are mentally represented and play a causal role in inflection. In contrast, according to the pattern associator hypothesis, variables and formal constituents are obsolete. The formation of *liked* from *like* is no different than the production of *came* from *come* (e.g., Hare & Elman, 1995; Plunkett & Marchman, 1993). Both are explained by the associations of specific tokens. Speakers' sensitivity to the constituents of *liked* is attributed to the co-occurrence of orthographic phonological and semantic features rather than to their corresponding variables (e.g., Seidenberg, 1987; Rueckl et al., 1997). Variables are not encoded nor do they play any role in inflectional morphology. A widely debated question in the morphological literature, then, is whether the representation of morphologically complex words encodes the compositional structure of abstract variables.

The representation of variables has been supported by three types of evidences. One is the demonstrable contribution of morphemes to word perception even after controlling for their nonformal correlates (e.g., Feldman, 1994; Fowler, Napps, & Feldman, 1985; Marslen-Wilson, Komisarjevsky Tyler, Waksler, & Older, 1994). A second line of investigation identifies properties that differentially affect regular vs irregular inflection. Because regular inflection is achieved by a rule, it should be independent of a word's idiosyncratic properties. In accord with this prediction, novel verbs are regularly inflected regardless of their similarity to existing regular verbs. In contrast, irregular inflection, the outcome of an associative process, is highly sensitive to similar existing irregular verbs (Prasada & Pinker, 1993; but see Daugherty & Seidenberg, 1992). Selective similarity effects for irregular, but not regular inflection, is also observed in nominal inflection (Berent, Pinker, & Shimron, 1999).

The third line of research probes for knowledge that is specifically defined over linguistic variables. Consider nominal inflection. According to the symbolic view, all regular nouns are members of a single category. This category is defined solely by the fact that its members instantiate a variable, as there is no semantic orthographic or phonological feature common to all these nouns. If it can be demonstrated that, despite their surface dissimilarity, regular nouns are treated alike, and distinguished from irregular nouns, then we have evidence for their representation by a variable. Such a demonstration would indicate that a morphemic variable, such as ''regular noun,'' is the domain of a mental constraint. Findings supporting this claim were reported by Gordon (1985). Gordon (1985) showed that 3- to 5-year-old children differentially treat regular and irregular nouns, allowing irregular (e.g., *mice eater)* but not regular (e.g., *rats eater*) plurals inside a compound.[1] The ban on compounding regular plurals cannot be due to their sound (the *s* ending) or meaning, as pluralia tantum (e.g., *pants*), but not their semantic controls (e.g., *shirts*), are produced inside a compound.[2] In contrast to compounding, that selectively allows irregular plurals, names selectively disallow them. Names invariably take regular inflection, even when they sound identical to an existing irregular noun. For instance, members of the Child family are

---

[1] Further work (Alegre & Gordon, 1996) demonstrates that regular noun inflections are permitted in compounds when they are fed back to morphology from syntax (e.g., [red rats] eater). Importantly, in the absence of such a recursive process (e.g., red [rats eater]), regular inflections are prohibited in compounds.

[2] Seidenberg, Haskell, and MacDonald (1999) proposed a connectionist account for such data. Their proposal attributes the ban on regular plurals in a compound to semantic and phonological preferences. As indicated above, Gordon (1985) has controlled for these aspects of compounding. It is thus unclear how Seidenberg's et al (1999) accounts for these facts.

The Childs, not **Children*[3] (Berent et al., 1999; Kim et al., 1991, 1994). These two pieces of evidence support the status of variables by demonstrating that they define the domain of linguistic generalizations, compounding and inflection, respectively. Following a similar rationale, we explore the role of the root morpheme as the domain of a phonological constraint, the Obligatory Contour Principle (OCP).

THE REPRESENTATION OF VARIABLES: EVIDENCE FROM THE OCP

The OCP is a fundamental principle in modern phonology (Goldsmith, 1990; Kenstowicz, 1994; McCarthy, 1979, 1986, 1989; Yip, 1988). In its formulation by McCarthy (1986), the OCP is a universal phonological constraint on lexical representations. The lexical OCP bans adjacent identical elements in a phonological constituent. Evidence for the OCP has been observed in a variety of languages, and its effects concern the identity of segments, tones, and features. One of the most widely cited evidence for the OCP is found in Semitic morphology. We thus chose to investigate OCP effects in Hebrew. We first describe OCP effects in the formation of Hebrew words. Next, we explain the implications of OCP effects to the investigation of mental variables. We review existing evidence for the OCP and outline some questions that are unaddressed by this research, questions that motivate our present investigation.

OCP Effects in Hebrew

Hebrew words include two ingredients: a root and a word pattern. The root typically contains three consonants. The word pattern provides the vowels and affixes and specifies the location of the three root consonants by means of placeholders. Words are formed by inserting the root in a word pattern. For instance, the verb KaTaB (he wrote), is formed by inserting the root, KTB, in the verbal word pattern CaCaC.[4] Verbal word patterns are called *binyanim* (singular: *binyan*). The Hebrew binyanim differ considerably in the location of the root in the word, the vowels, affixes, and some core aspects of the word's meaning. For example, the conjugation of the root KTB in the word patterns hitCaCeC and hiCCiC yields the verbs hitKaTeB

---

[3] The avoidance of irregular inflection for names is a special case of a more general phenomenon, namely the blocking of irregular inflection for nouns and verbs that have no visible canonical root. Because irregularity is a property of the root, the lack of visible root also prevents access to the irregularity feature; hence, regular inflection applies by default (Kim et al., 1991, 1994; Marcus et al., 1995). Additional cases where the root is invisible include borrowings, acronyms, and words derived from a different grammatical category (cf. *He flew to LA* vs *He flied out to center field*).

[4] We use C to refer to any consonant. For viewing convenience, we notate root consonants in uppercase.

a. KaTaB

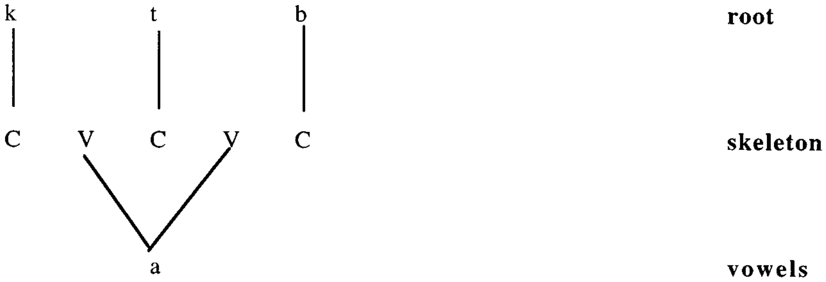

b. hiTKaTeB

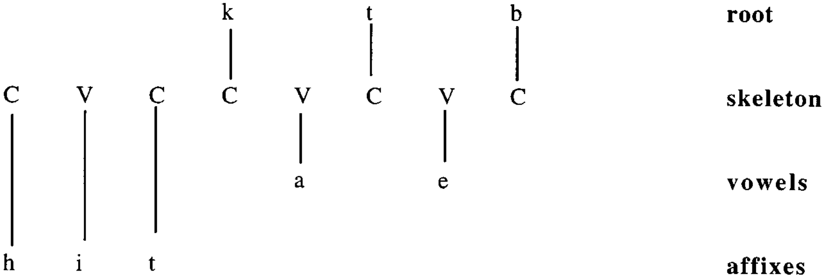

FIG. 1. The representation of the verbs katab and hitkateb. Note that root consonants are represented on a single plane, segregated from vowels and affixes.

(corresponded with) and hiKTiB[5] (he dictated), respectively. In addition to their morphological status, however, the root consonants also form a phonological constituent. This fact is captured in autosegmental phonology by representing the root consonants on a single plane, separate from the vowels and affixes. Figure 1 illustrates the multiplanar structure of the verbs KaTaB and hitKaTeB. Note that consonants and vowels are segregated into different planes and anchored to a skeleton, an abstract set of timing units, assigning distinct slots to consonants and vowels.

This representation has some important consequences for explicating phonological process. The segregation of the root consonants renders them adjacent in phonological representations. Phonological constraints are thus free

---

[5] A subsequent spirantization realizes the K and B as x and v, respectively.

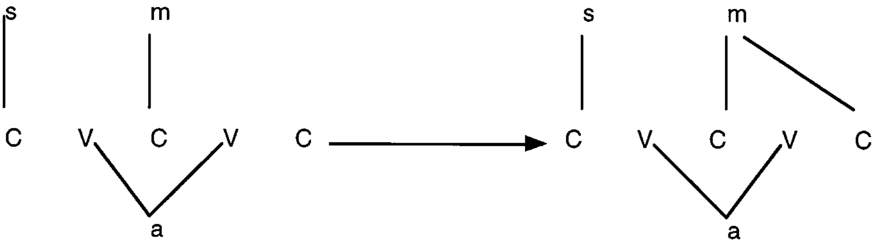

FIG. 2. The derivation of root-final gemination from its underlying biconsonantal form. The left figure illustrates the alignment of the biconsonantal root SM with the word pattern CVCVC. Because the alignment proceeds from left to right, the leftmost consonant slot remains unfilled. The right figure describes the spreading of the phoneme into the free slot. The doubly linked phoneme surfaces as root-final geminate, SMM.

to operate on root consonants, ignoring any intermediate vowels. The focus of our present investigation is in one such constraint, the OCP. The OCP bans identical adjacent consonants from the representation of the root in the mental lexicon. To the extent that Hebrew roots exhibit geminates, these must be formed by the grammar. For instance, consider the root SMM. According to McCarthy (1979, 1986), this root is stored in the lexicon as a biconsonantal root, SM. The geminates emerge in the process of word formation. To form a word from a biconsonantal root (e.g., SM), it must be aligned with a word pattern containing three slots for the root consonants. The alignment proceeds from left to right, leaving the rightmost slot of the root consonants unoccupied. To fill the empty slot, the second root radical (e.g., m in SM) spreads rightward. The root-final radical, m, links to two consonant slots, resulting in the surface root form SMM (see Fig. 2). The geminates in SMM are thus the manifestation of a single lexical radical, avoiding a violation of the OCP. In what follows, we refer to the process that generates roots with geminates, e.g., SMM from a biconsonantal representation (e.g., SM as reduplication.[6]

This account makes a simple, testable prediction. If root gemination is well formed only when it is produced by rightward reduplication, then there should be an asymmetry in the location of geminates in the root. Roots with final gemination, such as SMM may be formed by the rightward reduplication of a biconsonantal lexical representation, SM. In contrast, root-initial gemination, such as SSM, cannot be formed by rightward reduplication from SM. Such roots must be stored in the lexicon with their geminates, thereby violating the OCP. The acceptability of geminates should thus depend on

---

[6] Reduplication in McCarthy's analysis differs substantially from the proposals of Gafos (1998) and Everett and Berent (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.

their location in the root: Root-initial gemination (e.g., SSM) is ill formed, whereas root-final gemination (e.g., SMM) is well formed.

This prediction is well supported by the distribution of roots in Semitic: Root-initial gemination is extremely rare, whereas root-final gemination is frequent. However, this asymmetry may also reflect a diachronic source that is no longer active in modern speakers' competence. In this series of experiments, we assess the psychological reality of the OCP. Our interest in the OCP is twofold. One source of interest in the OCP is its pivotal role in phonological theory. However, our primary interest in this phenomenon stems from its potential to inform a central debate in cognitive theory, namely the representation of mental variables.

OCP Effects as Evidence for Variables

The OCP refers to two variables: the root morpheme and geminates. If mental representations do not specify variables, then neither geminates nor the root could be representable. Thus, the pattern associator hypothesis must predict that the OCP cannot constrain mental representations. The elimination of variables, however, does not necessarily prevent a pattern associator from mimicking OCP effects. If the behavior in question correlated with the statistical distribution of linguistic tokens, then a pattern associator could exhibit a constraint on Hebrew roots by appealing to their statistical correlates. Hebrew does not seem to present the language learner with such a bootstrapping strategy when it comes to the OCP. The frequency of subword units in Hebrew generally does not correspond to the frequency of morphological constituents. A constraint on root structure does not appear to be learnable from the statistical structure of Hebrew words.

To appreciate the distinction between root and word structure, consider the words *titfor* and *titer*. *Titfor* is an existing Hebrew word (*she will sew*), whereas *titer* is a nonword. These two ‘‘words’’ share the geminate *tit* in their word-initial position, but differ in their root structure. *Titer* is formed from the root TTR, a novel root that manifests root-initial gemination. In contrast, *titfor* is formed from the familiar root TFR. The word-initial geminates in *titfor* stem from the affixation of the prefix *ti* to a root whose initial radical happens to be /t/, but lacks any root gemination. Thus, word-initial geminates are not necessarily root initial. Likewise, root-initial geminates are not necessarily word initial. The conjugation of the root TTR and TFR in the word pattern maCCiC yields *mattirim* and *matfirim.* Neither of these forms manifests word initial gemination. However, because the root of *mattirim* has root-initial geminates, it violates the OCP.

The absence of a systematic correspondence between word and root structure is not limited to a few rare cases. Hebrew roots are productively inflected in numerous binyanim that differ considerably in the word location of the root. Thus, the position of the root geminates in the word is bound to change dramatically as it is inflected in different binyanim. Because the binyanim

TABLE 1
Illustration of the Materials Used in Berent and Shimron (1997)

Root type	Class 1	Class 2	Class 3
SSM	Si-SeM	maS-Si-Mim	hiS-ta-SaM-tem
MSS	Mi-SeS	maM-Si-Sim	hit-Ma-SaS-tem
PSM	Pi-SeM	maP-Si-Mim	hit-Pa-SaM-tem

frequently entail prefixation, any root whose root-initial radical matches the prefix's consonant will result in ‘‘fake’’ geminates word initial. The constraint on root structure thus cannot be easily inferred from the co-occurrence of subword units. The OCP is also difficult to surmise from the properties of root tokens. The OCP renders all tokens of three consonants corresponding to a potential Hebrew root a single class. However, the members of this class exhibit no common defining feature: The class of potential Hebrew roots share no orthographic, phonological, or semantic feature. If speakers nevertheless treated all these tokens as members of the same category, then the defining feature of this category is likely to be formal.

Existing Experimental Evidence for the OCP

Initial support for the OCP is provided by the recent findings of Berent and Shimron (1997). Berent and Shimron reasoned that, if the OCP is a mental constraint that is synchronically active, then speakers should prefer word forms derived from roots with final gemination over words whose roots manifest root-initial gemination. Furthermore, if the constraint on gemination applies to a formal constituent, the root, then the acceptability of geminates should depend on their location in the root, not in the word. Root-initial gemination should be ill-formed regardless of its position in the word (cf. SiSeM vs hiStaSeM). Similarly, the acceptability of root-initial geminates should not depend on their surface adjacency: Root-initial gemination should be ill formed even when the geminates are separated by a full vowel (e.g., SiSeM) or even an infix (e.g., hiStaSeM).

To examine these predictions, Berent and Shimron (1997) asked participants to rate the acceptability of nonwords generated from nonroots, i.e., combinations of three consonants that do not correspond to any existing Hebrew root. The critical items contained root-initial gemination (e.g., SSM). These roots were matched with two types of control roots (see Table 1). One control was a no-gemination root (e.g., PSM). These roots were matched to the initial gemination roots in the second and third consonants, but differed in the first consonant, which was not a geminate. Words derived from roots with initial gemination should thus be rated lower than these no-gemination control roots. To rule out the possibility that the rejection of SSM-type roots is merely due to the presence of gemination, rather than to its location, a second type of control root was used. These roots exhibited root-final gemination (e.g., MSS). They include exactly the same consonants as in the initial gemination roots, but in reverse order. A consistent rejection of words derived from SSM-type roots relative to MSS- and PSM-type controls would indicate a sensitivity to the location of geminates. To further demonstrate that the rejection of root-initial gemination is due to the location of geminates in the root, rather than the word, Berent and Shimron conjugated the roots in several word patterns which differed in the transparency of the root and the surface adjacency of the geminates. The rating results reflected a strong rejection of root-initial gemination relative to its two types of controls. Importantly, the ill formedness of SSM-type roots emerged regardless of the location of geminates in the word. For instance, forms like hiS-ta-SaM-tem (generated from the root SSM) were rated lower than controls despite the fact that the geminates are separated by an infix.[7] These findings suggest that speakers' behavior is guided by a constraint on root rather than on word structure.

Some Unanswered Questions

Berent and Shimron's results suggest that Hebrew speakers possess a knowledge whose domain is the root morpheme. The contents of this knowledge, however, are not entirely clear from Berent and Shimron's (1997) findings. According to McCarthy (1986), speakers' knowledge bans identical consonants in the root. Recall that geminates are defined as the copying of a variable. The critical evidence for identity avoidance comes from an asymmetry in the acceptability of geminates within the root. This asymmetry is compatible with McCarthy's (1986) account that surface geminates are the result of rightward spreading. However, the rejection of root-initial geminates may also reflect their statistical structure. Roots with initial geminates manifest a low bigram frequency. Their rejection may stem from their rarity rather than from identity avoidance. The existing evidence is thus compatible with the idea that speakers do not represent identity.

The uncertainty regarding the representation of identity raises some important theoretical questions for both cognitive and linguistic theory. The representation of identity is the subject of a heated debate sparked by the recent results of Marcus and colleagues (Altmann & Dienes, 1999; Eimas, 1999; Marcus, 1999a, 1999b, 1999c, McClelland & Plaut, 1999; Negishi, 1999; Seidenberg & Elman, 1999). Marcus et al. (1999) demonstrated that 7-month-old infants can learn a constraint on the location of identity in a simple artificial language consisting of trisyllabic words. They claim that such a constraint is unrepresentable by pattern associators; hence, its acquisition provides *prima facie* evidence for the representation of mental variables.

---

[7] This infix is due to a metathesis process that switches the location of the prefix consonant *t* and root-initial sibilants (e.g., hit+SaSeM→hiStaSeM).

However, the use of an artificial language task raises questions regarding the generality of these conclusions with regards to linguistic competence. The OCP presents the opportunity to examine the representation of identity in natural language. Specifically, if Hebrew speakers constrain identity in the root, and if this constraint is unlearnable from the statistical structure of root tokens, then the ability to represent variables must form integral part of linguistic competence. The adequacy of connectionist models as accounts of language may be contingent on their implementation of variables.

The present research investigates whether the constraint on root structure requires the representation of identity, an issue unaddressed by the previous results of Berent and Shimron (1997). We examine two questions with regard to the representation of identity. First, why are roots with final gemination rated more acceptable compared to roots with initial gemination? According to McCarthy (1986), the acceptability of root-final geminates (e.g., SMM) reflects their formation from biconsonantal representations (e.g., SM). Conversely, the greater acceptability of these roots may reflect their higher bigram frequency. To assess the correspondence between the formation of geminates and their acceptability, Experiment 1 examines the production of triconsonantal roots from biconsonantal inputs. Experiments 2–3 obtain ratings for the expected outputs of the production task. A convergence between the rating and production tasks would suggest that the acceptability of root-final geminates reflects their formation.

A second question examined in these experiments concerns the mental representation of identity: Are geminates represented by a variable? To dissociate symbolic and statistical explanations for the formation of geminates, we systematically investigate the statistical structure of Hebrew roots and assess whether participants' behavior is explicable by the expected probabilities of geminates in existing Hebrew roots. Our investigation also seeks converging evidence for the representation of the root morpheme by a variable. If mental representations do not specify variables, then the responses observed in these studies may be largely predicted by the distribution of linguistic tokens. Conversely, if speakers do represent variables, then statistical structure could not fully account for performance. Speakers' behavior should be sensitive to the combinatorial structure of variables, specifically, the location of identity within the root.

EXPERIMENT 1

Experiment 1 employs a production task designed to mimic the formation of root geminates in normal language use. According to McCarthy's analysis (1986), surface forms such as SMM are derived from their underlying biconsonantal representations (e.g., SM) by a productive process of reduplication. Root-final gemination is extremely common in Hebrew. If McCarthy's (1986) account is correct, then all these roots must be formed by reduplication. Reduplication must then be a frequent strategy of word formation. If speakers routinely form geminates by reduplicating biconsonantal representations, then they are likely to employ the same strategy in a task that requires the conjugation of biconsonantal roots. The production task used in Experiment 1 examines this prediction.

In this task, participants are presented with a new root and an existing word. Their task is to conjugate the new root in analogy to the word given. To illustrate the task, consider a trial in which speakers are presented with the new root PSM and asked to conjugate it by analogy to the verb Pa?aL.[8] To perform the task, participants must follow several steps. They must first decompose the exemplar into its root (e.g., *P?L*) and word pattern (e.g, CaCaC). Next, they must delete the exemplar's root (e.g., P?L) and replace it with the given root (e.g, PSM). The expected response in this case is PaSaM. However, half of the new roots were biconsonantal (e.g., SM). These roots are of particular interest because they present the participant with an *alignment* problem: To conjugate a root, it must be aligned with the root placeholders in the word pattern. For instance, in our example, the root PSM must be aligned with the three consonant slots in CaCaC. For biconsonantal roots, however, there is a mismatch between the number of slots required by the word pattern (three) and the number of radicals provided in the root input (two). How, then, can two consonants be aligned with three slots?

The alignment problem may be solved in several ways (please note that participants were not provided with any of these solutions nor were they told of the presence of biconsonantal roots). One solution is to violate the word pattern: Insert the root into two of the available slots and delete the additional vowel. This will leave the subject with the verb SaM. Of primary interest is the class of solutions that preserves the word pattern, but changes the given root by adding a consonant. There are two critical questions with regard to the insertion of a segment. First, we wished to find out *what* kind of segment is added: whether the added segment is identical to one of the given radicals, i.e., a geminate, or whether it is a new phoneme. A related question is whether the production of addition vs gemination is affected by the similarity to existing root tokens. Second, we wished to find out *where* in the word pattern the new segment is inserted. The first question is critical to the representation of geminates: Do speakers represent the identity of geminates by a variable? Our second question is important for understanding the role of the root morpheme. The symbolic and pattern associator hypotheses contrast on their predictions concerning each of these questions.

Predictions Regarding the Representation of Geminates

If the OCP is active, then speakers must routinely form root gemination by reduplicating a stored biconsonantal representation. Participants are thus

---

[8] ? stands for a glottal stop consonant.

likely to use reduplication in conjugating biconsonantal roots in our experiment. Consequently, they should be more likely to form a geminate than to add a new segment. Viewed as the copying of variables, reduplication should be further blind to contents of geminates and their frequency in the language. Conversely, according to the pattern associator hypothesis, participants are unable to represent variables; hence, they should also be unable to copy them. Geminates are simply a special case of nongeminate bigrams that happen to be identical. The formation of geminates is not formally distinct from the formation of any other bigram: Both reflect an associative process of segment addition. There are numerous demonstrations of the sensitivity of pattern associators to type frequency (e.g., Daugherty & Seidenberg, 1992; Plunkett & Marchman, 1993; Prasada & Pinker, 1993; Rumelhart & McClelland, 1986). We thus expect that the production of a new root bigram (either geminate or nongeminate) should mirror its probability of occurrence in the language. Our analyses assess whether the observed probability of geminate responses corresponds to the statistical structure of the Hebrew roots.

Expected probabilities of reduplication vs segment addition. A rough estimate of the magnitude of geminate responses expected by the statistical structure of Hebrew may be obtained by comparing the number of responses that may be formed by combining our biconsonantal inputs with an identical vs a new segment. Each of our biconsonantal roots may be combined with any of the 22 Hebrew consonants in each of its three positions. Only 2 of these 66 addition responses yield a root with final gemination. If geminates are formed by segment addition, then gemination responses should be far less frequent than addition responses.

To obtain a more precise estimate of the probability of geminates vs nongeminate responses, we examined the statistical structure of existing Hebrew root tokens.[9] To estimate the statistical structure of Hebrew roots, we created a database including all productive triconsonantal roots from the Even Shoshan (1993) Hebrew dictionary, a total of 1412 roots. We next generated triliteral roots from our biconsonantal root inputs by inserting each of the

---

[9] Our choice of the root, rather than the word, as the unit of analysis was motivated by two considerations. First, the selection of the root as the basis for analysis permits evaluating the pattern associator's ability to account for identity independent from its ability to analyze the word's morphological structure. To capture the constraint on root identity, it is necessary to decompose the word into its morphological constituents and attend to regularities in root structure. The ability of pattern associators to achieve these goals in a nonconcatenative morphology is uncertain. The examination of bare root forms avoids this problem. In addition, this choice greatly simplifies our analyses. Recall that the Hebrew inflectional system is highly productive. Any given root may be inflected in numerous forms. For instance, a single root, KTB, gives rise to 115 verbal forms and 272 nominal forms. Hebrew has thousands of roots. An analyses of the words produced by all these roots falls beyond the scope of the present work. Importantly, this limitation provides the pattern associator hypothesis with the most favorable conditions imaginable to capture the constraints on root structure. Undoubtedly, our choice of the root as the basis of analysis is strongly biased in favor of the pattern associator hypothesis.

22 Hebrew consonants in each of the three root positions. We computed the positional type bigram frequency of these triliteral roots by summing the positional type bigram frequency of their three bigrams (C1C2 + C2C3 + C1C3).[10] The summed positional bigram frequency of the triliteral roots generated by adding a new segment to our 24 biconsonantal roots is 12,722. Conversely, the summed positional bigram frequency of roots generated by adding a consonant identical to the final radical is 520. Thus, the expected probability of geminate responses relative to new segment addition responses is 0.049. We note that this figure is merely a rough estimate. The actual probability of observing geminate vs new addition responses in any given simulation may depend on the precise token frequency of geminates in the training set and the network's specific architecture and training. Nonetheless, in view of the robust asymmetry in the expected probability of gemination vs addition responses, it is reasonable to expect pattern associators to exhibit a qualitatively similar behavior.

The effect of counterexamples. If geminates are formed by an associative process of segment addition, then its likelihood should depend on the contents of specific root tokens. The existence of counter examples to the OCP permits testing these predictions. Hebrew has about four roots that violate the OCP. Two of these roots are frequent and productive (MMN, *to finance;* MMSh, *to realize*), and they both manifest root-initial gemination of the phoneme *m*. To examine the effect of bigram frequency, we manipulated the similarity of the experimental roots to such counterexamples. In half of the experimental roots, the initial geminates matched existing root-initial geminates (e.g., MG). The other half of the experimental roots was not analogous to any counterexample. If geminates are formed by an associative process, then roots whose initial radical matches a counterexample to the OCP should be more likely to exhibit root-initial gemination compared to no-analogy roots. Conversely, if geminates are formed by reduplication, then the similarity to counter examples should not affect the likelihood of reduplication.

Predictions Regarding the Representation of the Root

In addition to the study of reduplication, our task may also provide converging evidence for the representation of the root morpheme as a mental constituent. If speakers have internalized a constraint concerning the location of geminates in the root, then not only are they expected to reduplicate rather than merely add segments, but they should further constrain the location of geminates relative to the root. Specifically, speakers should be more likely to reduplicate the second root radical over the first. Furthermore, the asymmetry in the location of geminates should emerge regardless of the position of the root in the word. Conversely, according to the pattern associator hy-

---

[10] We choose to examine the bigram as the unit of subroot token co-occurrence because this is the smallest possible unit of subroot radical co-occurrence.

TABLE 2
Illustration of the Conjugation of the Root PSM in the Three Word Patterns Used in the Production Task

Root	Word class	Exemplar	Expected response
PSM	I	Pa-?aL-ti	Pa-Sa-Mti
PSM	II	maP-?i-Lim	maP-Si-Mim
PSM	III	hit-Pa-?aL-ti	hit-Pa-SaM-ti

pothesis, the root, the domain of the OCP, is unrepresentable. Hence, the production of a geminate segment may only be sensitive to its frequency in the word. Words contrasting on their surface structure should not exhibit any systematicity in the location of the new segment.

To dissociate root structure from word position, we manipulated the word pattern of the exemplars. Table 2 illustrates these word patterns. They differed in terms of the location of the root in the word pattern. In the first pattern, the root was unaffixed, whereas in the second and third patterns the root was prefixed and suffixed. The second and third patterns differed, however, with respect to the surface adjacency of geminates in the root. In the third word pattern, root geminates (either initial or final) are separated by at least a full vowel, whereas in the second word pattern root-initial geminates are truly adjacent. Berent and Shimron (1997) observed that surface adjacency increases the ill formedness of root-initial gemination. Importantly, if speakers are sensitive to the location of geminates in the root, then their tendency to reduplicate the second over the first radical should emerge in each of these word patterns.

Method

Participants

The production task was administered as part of a course at the School of Education in the University of Haifa. Twenty-four native speakers served as participants.[11] They were all students at the School of Education and received no compensation for their participation.

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. Both root structure and exemplar structure were manipulated.

Root types. Forty-eight novel roots were used in the study. Half of these roots were triconsonantal, and half were biconsonantal. Our primary interest was in the biconsonantal roots. The triconsonantal roots were used as fillers. Triconsonantal roots are perfectly alignable with the

---

[11] The production questionnaire was administered to 27 native Hebrew speakers students. Three of the students chose not to complete the questionnaire and were thus excluded from all analyses.

three consonant slots provided by the word pattern. In contrast, the alignment of biconsonantal roots with the word pattern always yields an extra consonant slot. The inclusion of the triconsonantal root fillers was designed to encourage participants to fill this extra slot (by either gemination or addition) rather than leave this slot empty (and violate the word pattern).

The structure of biconsonantal roots. The 24 biconsonantal roots consisted of two groups. Half of the items were analogous to existing Hebrew roots that violate the OCP (the ‘‘analogy’’ roots), whereas the other half of the items were not analogous to counterexamples (the ‘‘no analogy’’ group). The analogy roots shared the initial consonant of existing Hebrew roots that violate the OCP. For instance, the novel root MG is analogous to the existing roots MMN, *to finance,* and MMSh, *to realize,* which violate the OCP. Modern Hebrew has four roots that violate the OCP. In two of these roots, the root-initial geminates are mm (MMN and MMSh). Because these roots are both highly productive and familiar, they are considered strong counterexamples to the OCP. The remaining two violations of the OCP are relatively weak, as their roots are unproductive and relatively rare (*nanas,* dwarf; *tite,* swept). Half of the items in the analogy groups had *m* as their first radical (a total of six items). The other half of the ‘‘analogy’’ items had either *n* or *t* root-initial geminates. Despite the relative weakness of the *nn* and *tt* counterexamples, we nevertheless included them in the analogy group for two reasons. First, the phonemes *t* and *n* are frequently used as prefixes. When these prefixes are followed by a root with the same initial phoneme, the resulting words exhibit word-initial gemination (e.g., the root TPR in *tiTPoR, she will sew*). Consequently, the phonemes *t* and *n* are often geminated in word-initial position. The familiarity with these word-initial geminates may encourage root-initial gemination as well. Second, there is only a limited number of possible nonroots with *m* initial geminates. By including the n-t initial geminates, we hoped to increase the size of the analogy group and the power of our manipulation.

Word pattern structure. 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 P?L in one of three classes word patterns (this root was chosen because it is often used as an exemplar in Hebrew grammar classes; hence, its conjugation is especially familiar). The three word classes[12] differed in terms of the surface transparency of the root in the word (see Table 2). 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, pi?el, and pu?al in the singular past-tense perfect form. Conversely, 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 is not separated by a full vowel (e.g., MaSSiMim). Members of the second class included the preset tense of binyan hif?il and mishkal nif?al. The third word class included verbs in binyan hitpa?el past tense . In this word class, the initial bigram is 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.

Design

Root structure (analogy vs no analogy) and word class (3) were crossed and manipulated within participants and within items.

---

[12] We use the term ‘‘word class’’ to refer to a group of word patterns (mishkalim or binyanim) that share properties that are of theoretical interest to our investigation. This term has no technical linguistic significance.

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. The order of the trials in the list was random. The orthographic representation of the exemplars specified all their vowels using diacritic marks.

At the beginning of the study, participants were provided with written instructions that were also read to them aloud. They were told that

> the purpose of the study is to examine how Hebrew speakers form new words. To that end, we invented a few roots of new words. We wish to examine how Hebrew speakers conjugate these roots according to the conjugations accepted in Hebrew. In the left column,[13] you will find new roots. To the right of the root we provide you with an example of a common Hebrew conjugation. We would like you to write the conjugation of the new root according to the example next to it. For instance:

new root	*conjugation exemplar*	*new conjugation*
lmn	hitpa?alti	hitlamanti’’

Please note that neither the instructions to the participants nor the example mentioned the existence of biconsonantal roots or the possible use of a reduplication strategy. Participants were asked to use diacritic marks in order to specify the vowels in their written output.

Coding Scheme

Participants’ responses to the experimental items were coded according to the following categories.

Errors. Errors were failures to respond, the deletion of one of the root’s radicals, or the use of an incorrect conjugation.

Correct responses. 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. Such outputs often require the deletion of one of the vowels in the word pattern. Although the resulting word sometimes altered the word pattern, these responses were nevertheless considered correct. (d) *Addition:* Additions were responses that reflected a triconsonantal root consisting of the given biconsonantal root and an additional new radical. Additions were scored according to the phoneme added and its location in the root. A few of the additions reflected reduplication of parts of the root. For instance, given an xy root, we observed outputs such as xyx, xyxy, yxy, and yy. Although these responses seem to reflect the copying of root structure, rather than the addition of a new segment, these distinctions cannot always be clearly made. We thus opted for a conservative analysis, counting all these outputs as additions.

Results

Participants were overall quite accurate in the performance of the production task ($M = 93.35\%$). To assess the effect of word class and analogy on

---

[13] In the Hebrew original (written from right to left) the new root was presented at the rightmost column and the exemplar was presented to its left.

their errors, we submitted the error data to ANOVAs [3 (word class) × 2 (analogy)] by participants and by items. In this and all subsequent analyses we adopt a level of .05 for rejecting the null hypothesis (p values are provided only for nonsignificant F values that exceed 1). The ANOVAs revealed a main effect of word class, significant only in the analysis by items [$F_s(1, 46) = 1.346$, $MS_e = 129.307$; $F_i(1, 44) = 9.160$, $MS_e = 11.924$]. Tukey HSD comparisons indicated a slight increase in error rate in the second word class. No other effect approached significance (all ps $> .2$, by participants and items).

Our interest, however, was in correct responses. Specifically, we assessed the representation of two mental variables: the root and geminates. In what follows, we describe the analyses examining evidence for each of these variables.

Evidence for the Representation of the Root Morpheme

The first set of analyses assessed the representation of the root by examining whether it serves as the domain of a mental constraint, the OCP. If the OCP is active, then participants should prefer gemination of the second over the first root radical. This asymmetry should emerge regardless of the location of geminates in the word or their surface adjacency.

Participants' correct responses were submitted to separate ANOVAs [3 (word class) × 2 (gemination type)] by participants and items. These analyses yielded a significant effect of gemination type [$\mathrm{F}_s(1, 23) = 31.129$, $MS_e = 2509.79$; $F_i(1, 23) = 2388.931$, $MS_e = 32.70$] and a gemination type × word class interaction [$\mathrm{F}_s(2, 46) = 3.735$, $MS_e = 74.79$; $F_i(2, 46) = 9.527$, $MS_e = 30.156$]. Of the total responses, 46.817% reflected root-final gemination, whereas the rate of root-initial gemination was 0.2315%. We next tested for the asymmetry in the location of geminates in each of the three word classes. These contrasts, as well as all subsequent planned comparisons, were tested using Dunn's test (Kirk, 1982, p. 106), protecting against type I error at the .05 level for the entire family of contrasts. Root-final gemination was far 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. Thus, speakers constrain the location of geminates relative to the root even when the morphological structure of the word is highly opaque (see Fig. 3).

Evidence for the Representation of Geminates

Our previous analyses probed for an asymmetry in the distribution of geminates relative to the root morpheme. The term ''gemination,'' however, has been so far used merely descriptively to refer to the presence of two identical root consonants in the written output. To assess the representation of identity, we now turn to investigate whether geminate responses are explicable by an associative process of segment addition.

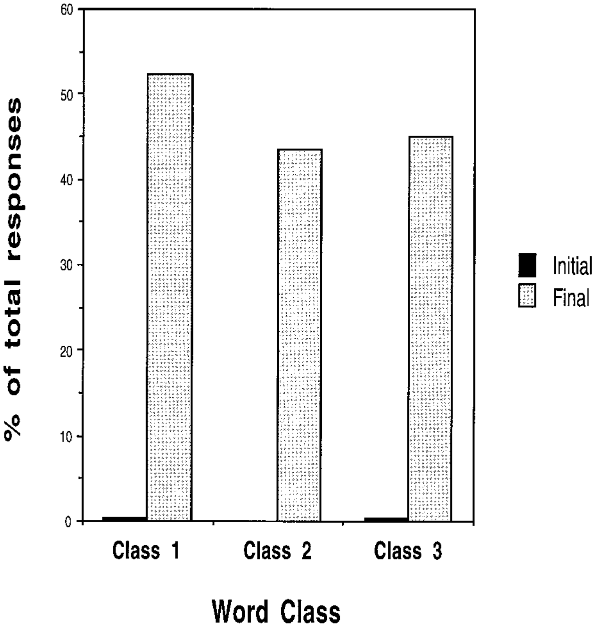

FIG. 3. The rate of root-initial vs final gemination in each of the three root classes in Experiment 1.

The effect of analogy on the location of geminates. If root-final gemination is due to the insertion of a phoneme, rather than to reduplication, then its likelihood should increase for roots analogous to violations of the OCP. To assess the effect of analogy on the type of gemination, we submitted the data to a three-way ANOVA [2 (analogy) × 3 (word class) × 2 (gemination type)] by participants and items. For the sake of simplicity, we examine only the effects and interactions involving the analogy factor. The interaction of analogy × gemination type was significant by participants [$F(1, 23) = 4.994$, $MS_e = 21.293$] but not by items [$F(1, 22) = 1.069$, $MS_e = 32.607$, $p = .3125$]. There was a slight increase in the production of root-initial over root-final gemination for items analogous to counterexamples (see Table 3).

TABLE 3
The Percentage of Responses Reflecting Root-initial vs Root-final Gemination as a Function of the Analogy of the Root to Counterexamples

	Analogy	No analogy
Root initial	0.463	0
Root final	45.83	47.80

TABLE 4
The Percentage of Responses Reflecting Root-initial vs Root-final Gemination as a Function of the Analogy of the Root to Counterexamples

	m analogy	n-t analogy
Root-initial	0.46	0.46
Root-final	46.06	45.60

Note. m analogies are biconsonantal roots whose initial radical are m; n-t analogies are biconsonantal roots whose initial radical is either n or t.

However, root-initial gemination was far less frequent than root-initial gemination even when the item was analogous to counterexamples. None of the other effects involving the analogy factor approached significance (all ps $>$ .13).

Given the trend suggesting the sensitivity of gemination to counterexamples, we wished to subject this finding to a closer scrutiny. Recall that our category of items analogous the to counterexamples contained two groups of items. Half of the items had *m* as their root-initial radical. These items are analogous to highly frequent and productive roots. The second half were items with *n* and *t* as their initial radical, roots for which the counterexamples were weaker. If participants are sensitive to counterexamples, then the rate of root-initial gemination should be higher for the strong analogy group. To investigate this prediction, we compared the magnitude of gemination for each of these groups separately. Participants' correct responses were submitted to an ANOVA [2 (analogy type; m vs n-t) × 3 (word class) × 2 (gemination type)]. These analyses failed to show any difference between the two types of analogy. The effect of analogy was not significant, nor did it interact with any other factor (all ps $>$ .15, by participants and items; see Table 4). Specifically, root-final gemination was far more frequent than root-initial gemination even for roots whose initial radical is a strong counterexample, *m* [$t_s(46) = 30.92$, $MS_e = 78.335$; $t_i(10) = 4.83$, $MS_e = 32.245$].

Are gemination and addition different strategies? If the emergence of geminates is due to the addition of a segment, rather than to reduplication, then the probability of observing geminates should be far smaller than that of adding a new segment. To examine this prediction we compared the different types of correct responses produced by our participants. There were three types of correct responses: gemination (root-final and root-initial), addition, and failure to geminate (see Table 5). For the sake of simplicity, we opted for the most conservative test for the reduplication hypothesis. We separated between initial and final gemination, assigning them into distinct categories, but combined all types of addition into a single category. These four types of correct responses (addition, initial gemination, final gemination, and no-

TABLE 5
The Distribution of Correct Responses in the Production Task

Response type	Percentage of total responses
Additions	14.352
Root-initial gemination	0.231
Root-final gemination	46.817
No-gemination	31.887

gemination) were compared using a one-way ANOVA by participants and items. The main effect of response type was significant [$F_s(3, 69) = 9.503$, $MS_e = 1043.014$; $F_i(3, 69) = 399.945$, $MS_e = 24.82$]. The largest category of responses was root-final gemination. Root-final gemination was far more frequent than all the responses of segment addition combined [$t_s(69) = 3.48$; $t_i(69) = 22.61$]. Interestingly, however, there was also a substantial number of no-gemination responses. Despite its numerical advantage, root-final gemination did not differ significantly from no-gemination in the analysis by participants [$t_s(69) = 1.505$, $p = .1138$; $t_i(69) = 10.94$].

Discussion

In this experiment, we examined Hebrew speakers' sensitivity to the OCP. The OCP bans identity from Hebrew roots. Thus, a sensitivity to the OCP requires the representation of two variables: the root and geminates. Our discussion evaluates speakers' sensitivity to root structure and compares statistical and symbolic accounts for their knowledge.

The Representation of Geminates

Consider first the representation of geminates. Our principal test for the representation of geminates' formal structure is reduplication. If the OCP is active, then Hebrew speakers should routinely form root-final geminates by reduplicating stored biconsonantal forms. Our participants are thus expected to frequently use reduplication in the performance of the experimental task. Conversely, according to the pattern associator hypothesis, reduplication, the copying of a variable, is unrepresentable. Gemination, on this view, must be formed by an associative process of segment addition. Our findings present several challenges to this proposal.

The view of gemination as the outcome of segment addition predicts that this process should reflect the properties of specific root tokens and their frequency. Given the statistical structure of Hebrew roots, the expected probability of gemination should be lower than the probability of segment addition. Specifically, according to our our rough estimate, the expected rate of final gemination vs. addition responses is 4.9%. Our results do not support this prediction. The observed rate of final gemination vs addition responses is 329.30%.[14] Gemination (e.g., bb) was not only far more frequent than the addition of any particular new segment (e.g., bg), but, in fact, was significantly more frequent than the addition of all new segments combined. Likewise, our analysis of counter examples show little evidence for the sensitivity of gemination to similar tokens. The similarity of the target root to existing roots manifesting root-initial gemination (roots beginning with *m, n,* or *t*) had only a marginal effect on the probability of their gemination. Moreover, for the most productive and frequent counterexample, the root-initial geminate *mm,* there was no increase in root-initial gemination. These findings do not agree with the view of gemination as the addition of a new segment.

The associative account is further challenged by the absence of independent evidence for the existence of a segment addition strategy. If geminates were formed by segment strategy, then this strategy should have been manifested also in the addition of nongeminate segments. In particular, if the small number of addition responses we observe was due to a distinct strategy of segment addition, then these responses should have reflected the full range of triliteral roots that may be formed by adding a new segment to the biconsonantal input. Recall that our biconsonantal roots may be transformed into triliteral roots by inserting any one of the 22 Hebrew consonants in each of the three root positions. In contrast, the observed responses of new segment addition cluster around only two radicals, which account for 92% of the total responses of new segment addition (see Fig. 4). These are the segment /y/ at the root-initial and root-middle positions and the segment /l/ at the root-final position. The frequent addition of the radical /l/ at the root-final position may well be due to the particular root exemplar given to our participants: All words exemplars were formed by conjugating the root P?L. The addition

---

[14] One may criticize our above analysis on the grounds that it fails to take into account the relative dominance of the geminate response within the response set for any given item. Specifically, it is conceivable that, despite the overall lower frequency of geminates relative to the set of *all* (geminate and nongeminate) responses, geminates dominate this set: they may exhibit higher frequency than any *single* nongeminate alternative. To assess this explanation, we compared the bigram frequency of potential geminate and nongeminate responses for each of our items. For instance, the item *bg* allows for the formation of geminates by introducing the segment *g* in either the root-middle or -final position. Accordingly, we compared the frequency the triliteral root *bgg* to each of the 42 nongeminate triliteral roots that could be formed by inserting one of 21 possible segments in the root's middle or final position (e.g., *bg?, b?g, bgd, bdg; bgz bzg;* etc.). Our analysis provides no evidence for the dominance of geminate responses. In particular, for 23 of our 24 items, there was at least one nongeminate triliteral root whose frequency was equal or greater than the geminate response. The mean number of nongeminate roots whose frequency exceeds geminate roots was 12.58 (*SD* = 6.82). Thus, the production of geminates is inexplicable by their dominance either within the response set or across items. The statistical structure of the Hebrew language cannot account for the behavior manifested by our participants.

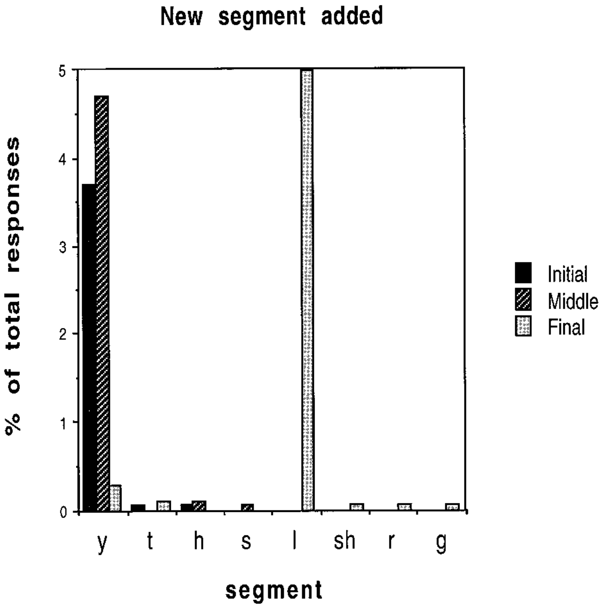

FIG. 4. The addition responses observed in Experiment 1 as a function of the segment inserted and its position in the root.

of /l/ at root-final position may simply reflect its coping from the given exemplar. Likewise, the addition of the /y/ phoneme may be explained by the phonological analysis of the biconsonantal input rather than by a statistical account.[15] Thus, the grand majority of the addition responses observed in this experiment may be readily explained without postulating an addition strategy, governed by the statistical distribution of root tokens. The discrepancy between gemination responses and the statistical structure of Hebrew roots further suggests that, even if a segment addition strategy existed, it could not account for geminate formation.

---

[15] Hebrew roots containing initial and middle /y/ are irregular (Rosen, 1977, Even-Shoshan, 1993). The conjugation of these roots sometimes deletes the /y/ radical altogether. For instance, consider the roots YShB (associated with the meaning *sit),* and ShYM (associated with the meaning *name*). The radical /y/ surfaces in the past tense in binyan qal, e.g., YaShaB (he sat), but deletes in the imperative, SHeB (sit!). Similarly, the root-middle /y/ surfaces in ShiYeM (*he named*) but deletes in SheM (*a name*). The irregular behavior of the radical /y/ may cast some ambiguity on the interpretation of biconsonantal forms. If participants interpret the given biconsonantal root as a word, rather than a root, then they may conclude that the input given to them contains the radical /y/ at either the middle or the final position. For instance, given the input pb, speakers may assume that it represents the word /peb/ (singular, masculine, imperative, in binyan *qal*), whose root it in fact pyb. The conjugation of pyb would sometimes result in the surfacing of the /y/ radical.

The Representation of the Root

Hebrew speakers appear to solve the alignment of a biconsonantal root with the word pattern by reduplication. In doing so, however, they do not randomly copy one of the root's radicals. Our findings reflect an overwhelming preference for copying the second over the first radical. Furthermore, this preference consistently emerged regardless of the location of the root in the word pattern. What is the basis for this remarkable asymmetry in the location of geminates?

One possibility is that speakers simply possess knowledge regarding the location of geminates in Hebrew words. The locus of geminates produced by our participants may reflect their probable word position. This account is strongly incompatible with our findings. Hebrew has numerous words that begin with geminates. Recall that Hebrew frequently exhibits word-initial geminates due to the concatenation of a prefix and a root beginning with the same radical. For instance, the words *titfor* (*she will sew*), *lilbosh* (*to wear*), and *nin?l* (*locked*) all reflect ''fake'' geminates due to prefixation. If the insertion of geminates in the experimental task was simply guided by knowledge of word structure, then root-initial gemination should have been frequent word initially. Likewise, if the root was eliminated, then participants should not have exhibited systematicity in root structure across word positions. Neither of these predictions is borne by the data. Root-initial gemination is practically absent, and its absence is just as pronounced when the root occupies word-initial or word-internal position. Thus, the marked asymmetry in the location of geminate responses cannot reflect the distribution of geminates in Hebrew words. To account for this asymmetry, we must assume the representation of a knowledge whose domain is the root, a variable.

A Puzzle

Our findings as a whole are compatible with the claim that speakers possess a knowledge regarding the location of geminates in the root. This constraint closely matches the OCP (McCarthy, 1986). However, there is one aspect of the finding that is at odds with McCarthy's (1986) account. In about 32% of the correct responses, participants chose to avoid gemination altogether, producing roots with only two consonants. For instance, given the root bp and the exemplar Pa?aL, we observed the response PaB. These responses are considered correct because they represent a conjugation of the given roots in a possible word pattern. Specifically, PaB could reflect an irregular conjugation in binyan qal. However, such responses distort the word pattern in order to avoid reduplication.

No-gemination responses could not be dismissed as anomalous outcomes of task-specific strategies. Biconsonantal word patterns are not uncommon in Hebrew. Hebrew has several irregular forms containing only two root consonants (e.g., SaMti, *I put;* NaXnu, *we rested;* Xam, *hot;* Kar, *cold*). Nevertheless, these no-germination responses are puzzling: Root-final gemination is clearly well formed. In fact, on McCarthy's (1986) account, root-final gemination should be indistinguishable from no-gemination roots as far as its well formedness. Furthermore, our results indicate that participants possess a reduplication strategy. If so, then why don't participants use this strategy across the board for all examples? Why do they sometimes prefer to avoid reduplication, even though this response distorts the word pattern? Is it possible that, despite its frequency, the identity of segments within a root is not entirely desirable?

EXPERIMENT 2

Why do speakers avoid reduplication despite the ubiquity of geminates in the language? As noted above, words with surface biconsonantal roots exist in Hebrew, thus, their production cannot be solely due to a task-specific strategy. Our current results, however, cannot rule out the possibility that the production task enhances such responses. A second explanation attributes the avoidance of reduplication to the statistical properties of its output, root-final identity. Although root-final gemination is generally frequent, it is possible that the token frequency of triliteral roots generated by rightward reduplication of our experimental roots is lower than that of roots with no geminates. Thus, the avoidance of rightward reduplication may reflect the relative rarity of the resulting tokens.

On a third account, however, the avoidance of reduplication is neither a task-specific nor a statistical artifact, but is, instead, a feature of linguistic competence. Specifically, the avoidance of reduplication may reflect a relative unacceptability of identity in phonological representations. Such a conclusion, however, is incompatible with McCarthy's (1986) specific predictions. Furthermore, this proposal is incompatible with the theoretical framework in which these predictions are grounded. Generative linguistics has traditionally assumed that linguistic competence is governed by inviolable rules and constraints (e.g., Chomsky, 1965). This approach requires that a frequent linguistic structure is well formed. In view of the ubiquity of root-final gemination in Hebrew, it must be acceptable. Interestingly, despite the profound differences between generative linguistics and the pattern associator hypothesis, their predictions, in this case, are quite similar. Assuming that the token frequency of root-final geminates is not lower than that of no-gemination controls, the pattern associator should predict that root-final gemination is as acceptable as no-gemination. The possible undesirability of root-final gemination is problematic for these two accounts, albeit from quite different reasons. An identity avoidance constraint is problematic for the generative framework because this constraint is frequently violated. Conversely, for the pattern associator view, identity avoidance is problematic,

TABLE 6
An Illustration of the Materials Used in Experiment 2

Root type	Class 1	Class 2	Class 3
SSM	Si-SeM	maS-Si-Mim	hiS-ta-SaM-tem
SMM	Si-MeM	maS-Mi-Mim	hiS-ta-MaM-tem
PSM	Pi-SeM	maP-Si-Mim	hit-Pa-SaM-tem

not because it is violable, but because it is unrepresentable: If variables are eliminated, then geminates are indistinguishable from nongeminates.

Experiment 2 investigates the acceptability of root identity. To examine whether the relative frequency of the gemination responses observed in the production task corresponds to their acceptability, we asked native Hebrew speakers to rate these outputs. For each of the biconsonantal roots used in Experiment 1 we formed a trio of three triliteral roots. For instance, for the root SM, these were the root-initial gemination, SSM, root-final gemination, SMM; and no-gemination control, PSM (see Table 6). The root-initial and root-final gemination members were produced by reduplicating either the initial or the final radical of one of the biconsonantal roots used in Experiment 1. The no-gemination roots (e.g., PSM) were matched to the initial-gemination roots in the second and third radicals, but included a different (nongeminate) initial radical. Each root trio was conjugated in each of the three classes of word patterns used in Experiment 1. Consequently, words sharing a root structure differed considerably in their surface structure. For instance, root-initial gemination surfaces as word initial in the first class (e.g., SiSeM), but as word medial in the second (e.g., MaSSiMim) and third (hiStaSaMtem) class. The second and third classes differed with regard to the surface adjacency of consonants at the first and second root positions. In the second class, root-initial geminates were truly adjacent (e.g., maSSiMim), whereas in the third class, they were separated by at least a full vowel (e.g., hiStaSaMtem). This design thus dissociates root structure from word position. Participants were asked to rate the acceptability of each word in the trio relative to the other two members. Specifically, speakers were asked to determine which of these forms sounded the best, which sounded the worst, and which one was intermediate. These rating permit assessing the representation of two variables, the root and identity.

Predictions Regarding the Representation of the Root

If speakers do not represent the root by a variables, then the marked contrast in the position of the root in the word should prevent systematicity in the rating of different word patterns sharing the same root structure (e.g., SiSeM, maSSiMim, and hiStaSaMtem). In contrast, if speakers represent the root, and if, further, the OCP is active, then speakers should be sensitive to the location of geminates in the root. Root-initial gemination should be rated lower than either root-final gemination or no-gemination controls, regardless of the location of geminates in the word.[16]

Additional evidence for the representation of the root may be found in the effect of root token frequency. Our previous discussion contrasted statistical and formal knowledge. However, these two sources of constraints are not mutually exclusive. In fact, it is the representation of the root variable that permits identifying root tokens as coherent linguistic objects despite their dissimilar surface manifestations in the word. The acquisition of statistical regularities regarding token properties may modulate the rating of our experimental roots. Roots whose initial gemination corresponds to a frequent bigram may not be rejected. To examine the effect of root tokens, we compared the rating of roots whose initial geminates are analogous to OCP violations with roots whose initial geminates never appear in root-initial position in Hebrew, the no-analogy roots. Within the analogy groups, half of the root-initial geminates were *mm,* a bigram appearing in two frequent and productive violations of the OCP in Hebrew. The other half of the roots had either *nn* or *tt* as their initial bigrams, corresponding to rather infrequent and unproductive violations of the OCP. If speakers possess knowledge regarding the position of geminates in specific tokens, then we would expect a reduction in the rejection of roots that are analogous to OCP violations, especially those corresponding to the frequent bigram *mm.*

Predictions Regarding the Representation of Identity

According to the pattern associator hypothesis, geminates are unrepresentable. Geminates are thus indistinguishable from nongeminate bigrams: Their acceptability should reflect their statistical structure.[17] Our materials present participants with two forms of identity: root-initial identity and root-final identity. Root-initial identity is both rare and unacceptable, according to the

---

[16] Note that our predictions strictly concern the *relative* acceptability of the different root forms (e.g., SiSeM is less acceptable than SiMeM), not an absolute level of acceptability (e.g., SiSeM is ill formed). We believe that absolute acceptability levels are not easily interpretable. A rating behavior is the outcome of various sources, including both the speaker's linguistic competence as well as various performance limitations and biases. The absolute unacceptability of a word may stem from its ungrammaticality. However, such judgment may also reflect a general bias against novel forms or identity. A comparison of root-initial gemination to controls addresses these biases. This method is not only sound experimentally, but may also be perfectly compatible with a theory of violable constraints. In this framework, the nonfatal violation of a constraint, could, conceivably, decrease the acceptability of a candidate without necessarily deeming it ‘‘unacceptable.’’

[17] The sensitivity to statistics over bare roots requires their abstraction from the word. Because root and word properties do not perfectly correlate, it is unclear whether statistics of bare root structure may be acquired by the pattern associator. For the sake of studying the effect of identity, we avoid the problem of root extraction by assuming that a pattern associator machine is provided with such forms.

OCP. Thus, a rejection of root-initial geminates may not necessarily reflect the representation of identity. Conversely, root-final geminates are frequent in the language. Our statistical analysis of the experimental roots indicate that their bigram frequency was higher than no-gemination controls. If speakers do not represent identity by a variable, then root-final gemination should be *more* acceptable than no-gemination controls. Conversely, if variables are represented, then geminates may be distinguished from nongeminates by virtue of their structure. Although the avoidance of root-final identity is incompatible with McCarthy (1986), it is nevertheless perfectly compatible with the symbolic view. If the aversion of root-final geminates in Experiment 1 reflects identity avoidance, then we should expect lower ratings for root-final gemination compared to no-gemination controls.

Method

Participants

The rating task was administered as part of a course at the School of Education at Haifa University. Twenty-two native Hebrew speakers served as participants. They were all students at the School of Education at the University of Haifa and received no compensation for their participation.

Materials

The materials were 216 words conjugated from 72 root trios. Each root was composed of three consonants that do not correspond to any existing Hebrew root. These root trios included three members: A root whose first and second radicals were identical (root-initial gemination, e.g., SSM), a root whose second and third radicals were identical (root-final gemination, e.g., SMM) and a root with three distinct radicals (a no-gemination root, e.g., PSM). The initial and final gemination members were formed by reduplicating either the initial or final radical in the 24 biconsonantal roots used in Experiment 1. Thus, the root-initial and root-final members were matched in their segmental contents and differed only in the location of the geminates. The no-gemination control was matched to the root-initial gemination member in its two final radicals, but differed in the first radical, which was not identical to the second.

The 24 root trios consisted of two equal groups. In the analogy roots, the initial gemination root was analogous to an existing Hebrew root that violates the OCP. Six of the analogous roots, had mm as their initial geminate—a bigram that appears in two frequent and productive Hebrew roots. The other six roots were analogous to weaker counterexamples. Three of these roots had *tt* as their initial bigram, and three had *nn* as their initial bigram. A second group of 12 root trios formed the no-analogy roots. The initial-gemination roots in these trios were not analogous to OCP violations.

The 24 trios of experimental roots were conjugated in each of the three word classes described in Experiment 1, yielding 72 word trios. The members of the word trio were identical in their word pattern and differed only in their root structure.

Root frequency. We assessed the statistical properties of our experimental roots by computing their type bigram frequency. In the absence of existing statistics on Hebrew roots (types or tokens), we used our own database, including all triliteral productive roots listed in the Even Shoshan Hebrew dictionary (1993), a total of 1412 roots. We calculated the positional type bigram frequency of our experimental roots by adding the frequency of their initial C1C2 bigram, their final C2C3 bigram, and the bigram of the initial and final radicals, C1C3. For instance, the summed positional bigram frequency of the root LLM is computed by summing

TABLE 7
The Summed Type Positional Bigram Frequency of the Roots Used in Experiment 2 as a Function of Analogy

Root type	Analogy	No analogy	Total
Root-initial gemination	6.83	5.25	6.041
Root-final gemination	12.83	10.33	11.58
No-gemination	8.58	10.083	9.33

the type frequency of its initial bigram, ll (0); its final bigram, lm (6); and the bigram consisting of the initial and final radicals, lm (5); a summed type frequency of 11. The mean summed positional type bigram frequency for our materials is presented in Table 7. As expected, roots with initial geminates were less frequent than their no-gemination controls, which, in turn, were less frequent than roots with final gemination. To assess the reliability of these differences, we submitted the mean summed bigram frequency of our experimental roots to an ANOVA [3 (root type) × 2 (analogy)]. The ANOVA yielded a significant main effect of root type [$F(2, 44) = 14.884$]. The analogy [$F(1, 22) < 1$] and the interaction of analogy × root type [$F(2, 44) = 2.104$, $p = .1341$] were not significant. Planned comparisons confirmed that the frequency of roots with initial gemination was lower compared to roots with either final gemination [$F(1, 22) = 29.422$] or no-gemination [$F(1, 22) = 10.38$]. Conversely, roots with final gemination were significantly more frequent than their no-gemination controls [$F(1, 22) = 4.85$].

Procedure

Participants were presented with 72 word trios. They were asked to rate the acceptability of each word relative to the other trio members as a possible Hebrew word. In performing the rating task, speakers were advised to attend only to the sound of the word and ignore its spelling or meaning associations. They were asked to assign the number 1 to the word that sounds the best, 3 to the word that sounds the worst, and 2 to the intermediate word. To express high acceptability by higher numbers, we report the data using an inverted scale, created by subtracting each score from 4. Thus, 1 corresponds to the worst sounding item and 3 corresponds to the best sounding item.

Results

Our discussion of these results addresses three questions: (a) Are Hebrew speakers sensitive to the location of geminates within the root? (b) Is the acceptability of root-initial geminates affected by their similarity to counter examples? (c) Do speakers prefer to avoid identity?

Are Hebrew Speakers Sensitive to the Location of Geminates Relative to the Root?

Mean acceptability ratings were submitted to two ANOVAs [3 (word class) × 3 root type)] by participants and items.[18] These analyses reflected

---

[18] It should be noted that the use of an ordinal rating scale could have introduced dependence among the subject mean ratings for the three root types, thereby violating one of the ANOVA'S assumptions. One option available for the analysis of such data is the use of nonparametric tests (e.g., the Wilcoxon *t* test for two dependent groups). Such an analysis, however, cannot examine the interaction of root type with word structure, an interaction that is crucial for

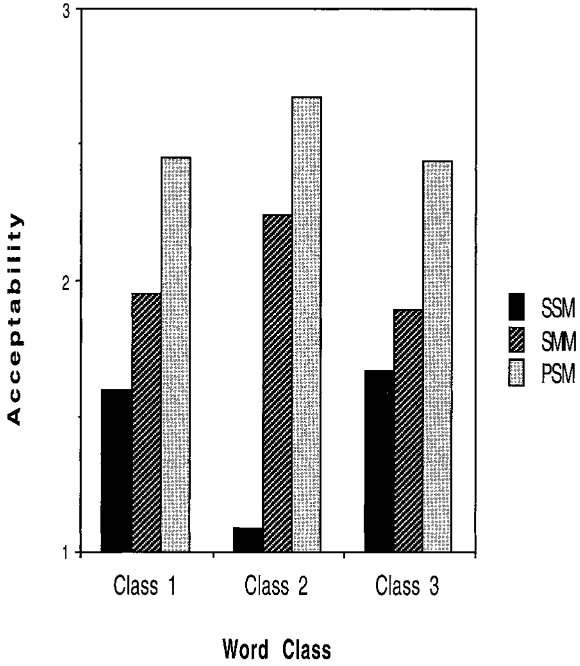

FIG. 5. Acceptability ratings as a function of root type and word class in Experiment 2.

a significant main effect of root type [$F_s(2, 42) = 271.706$, $MS_e = .070$; $F_i(2, 46) = 103.028$, $MS_e = .203$] and a significant interaction with word class [$F_s(4, 84) = 45.543$, $MS_e = .037$; $F_i(4, 92) = 16.743$, $MS_e = .113$]. The effect of word class was not significant (all Fs < 1).

The effect of root type was further investigated using Dunn's planned contrasts. Across word classes, root-initial gemination was rated significantly lower compared to either root-final gemination [$t_s(42) = 12.49$; $t_i(46) = 7.52$] or no-gemination [$t_s(42) = 23.29$; $t_i(46) = 14.35$] controls. This finding indicates a sensitivity to the location of geminates within the root. Importantly, the rejection of root-initial gemination emerged regardless of its location in the word (see Fig. 5). Root-initial gemination was rejected compared

---

evaluating whether the sensitivity to the location of geminates depends on their location in the root or merely the word. This question is pivotal for examining the representation of the root by a variable. Likewise, this approach cannot evaluate the modulation of root structure effects by analogy, an important question for evaluating the role of associative memory for root tokens. Another disadvantage of the analysis of a complex hierarchical design by means of separate *t* tests is the inflation in the likelihood of type one error. We thus chose to analyze the data by means of an ANOVA. To control for the increase in type one error due to the potential dependence between the observations, we tested all effects using the Greenhouse–Geisser procedure (Winer, 1971, p. 523). The critical statistic are reported using the original (uncorrected) degrees of freedom.

to its no-gemination [$t_s(84) = 6.09$; $t_i(92) = 3.62$] and final-gemination [$t_s(84) = 14.81$; $t_i(92) = 8.813$] controls in the first class. A similar rejection of root-initial gemination relative to its final gemination [$t_s(84) = 36.61$; $t_i(92) = 11.85$] and no-gemination controls [$t_s(84) = 27.59$; $t_i(92) = 16.41$] was obtained even when the root was heavily affixed, in the second class. Likewise, root-initial gemination was rejected relative to its final-gemination [$t_s(84) = 3.96$; $t_i(92) = 1.94$, n.s.] and no-gemination controls [$t_s(84) = 13.46$; $t_i(92) = 8.012$] in the third class, whose word structure is opaque. Thus, the rejection of root-initial gemination is general with regard to word structure. Word structure, however, did affect the acceptability of root-initial geminates. The rejection of root-initial geminates was especially pronounced in the second class, where they were not separated by a full vowel. Thus, surface adjacency enhances the unacceptability of root-initial geminates.

Is the Acceptability of Root-Initial Geminates Affected by Token Frequency?

If the acceptability of root structure is affected by token properties, then the rejection of root-initial geminates may depend on their frequency. Our design contrasted two levels of frequency for root-initial geminates. Roots whose geminates never occur in root-initial position (the no-analogy roots) and roots whose geminates do appear in root-initial position (the analogy roots). The ANOVA by participants [2 (analogy) × 3 (root type) × 3 (word class)] revealed significant interaction of root type × analogy [$F_s(2, 42) = 12.632$, $MS_e = .035$; $F_i(2, 44) = 1.049$, $MS_e = .202$, $p = .3578$] and a significant three way interaction of root type × analogy × word class [$F_s(4, 84) = 7.002$, $MS_e = .036$; $F_i(4, 88) < 1$]. To investigate this interaction, we conducted separate ANOVAs on the rating of analogous vs nonanalogous items.

The no-analogy roots. The analyses on the no-analogy items essentially replicate the findings of the omnibus analyses. The ANOVA revealed a significant main effect of root type [$F_s(2, 42) = 283.118$, $MS_e = .082$; $F_i(2, 22) = 96.757$, $MS_e = .131$] and a significant interaction of root type × word class [$F_s(4, 84) = 26.991$, $MS_e = .046$; $F_i(4, 44) = 6.838$, $MS_e = .099$]. Across word classes, root-initial gemination was rated significantly lower compared to either root-final gemination [$t_s(42) = 12.62$; $t_i(22) = 7.38$] or no-gemination controls [$t_s(42) = 23.78$; $t_i(22) = 13.9$]. This rejection also emerged separately in each of the word patterns (see Table 8). Specifically, root-initial gemination was rated significantly lower compared to its final gemination control in the first [$t_s(84) = 8.22$; $t_i(44) = 4.13$] and second [$t_s(84) = 17.08$; $t_i(44) = 8.59$] classes. In the third class, the same trend was significant by participates [$t_s(84) = 3.93$; $t_i(44) = 1.98$, n.s.]. Likewise, root-initial gemination was rated lower than its no-gemination controls in the first [$t_s(84) = 15.38$; $t_i(44) = 7.74$], second [$t_s(84) = 24.89$; $t_i(44) = 12.52$], and third classes [$t_s(84) = 14.79$; $t_i(44) = 7.44$]. The interaction of

TABLE 8
Mean Acceptability Rating as a Function of Root Type and Word Class for the No-analogy Roots Presented in Experiment 2

Root type	Class 1	Class 2	Class 3
Root-initial gemination	1.492	1.098	1.595
Root-final gemination	2.023	2.201	1.849
No gemination	2.485	2.705	2.549

root type × word class reflects a stronger rejection of root-initial gemination for the second class, in which the initial geminates are adjacent.

The analogy roots. A separate analysis of the analogy roots [2 (analogy; mm vs nn or tt) × 3 (word class) × 3 (root type)] revealed a significant three-way interaction in the analysis by participants [$F_s(4, 84) = 5.632$, $MS_e = .097$; $F_i(2, 40) = 1.183$, $MS_e = .126$, $p = .3329$]. To examine its source, we conducted separate analyses on the two types of analogous items. Because each of these groups includes only six items, we limit our investigation to the analyses by participants.

The results of the weak analogy roots (root-initial nn and tt) did not differ from the no-analogy group. The analyses revealed a significant main effect of root type [$F_s(2, 42) = 199.059$, $MS_e = .113$] and a significant interaction of root type × word class [$F_s(4, 84) = 15.428$, $MS_e = .090$]. Root-initial gemination was rejected compared to root-final gemination [$t_s(42) = 11.31$] or no-gemination controls [$t_s(42) = 19.89$]. The rejection of root-initial gemination was also significant in each of the word classes compared to either no-gemination [in the first class, $t_s(84) = 3.77$; in the second class, $t_s(84) = 12.68$; in the third class, $t_s(84) = 5.52$] or final gemination [in the first class, $t_s(84) = 9.29$; in the second class, $t_s(84) = 18.45$; in the third class, $t_s(84) = 10.88$] controls. As with the no-analogy roots, however, the rejection of root-initial gemination was particularly pronounced in the second class (see Table 9).

The separate analyses of the strong analogy roots (the m-analogy roots) revealed significant effects of root type [$F_s(2, 42) = 56.826$, $MS_e = .175$] and its interaction with word class [$F_s(4, 84) = 29.346$, $MS_e = .136$]. The

TABLE 9
Mean Acceptability Rating as a Function of Root Type and Word Class for the nn and tt Analogy Roots Presented in Experiment 2

Root	Class 1	Class 2	Class 3
Root-initial gemination	1.606	1.061	1.507
Root-final gemination	1.947	2.208	2.008
No-gemination	2.447	2.731	2.492

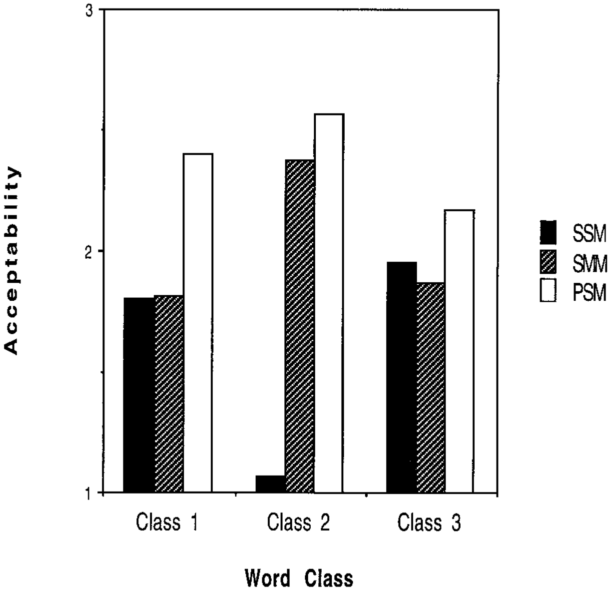

FIG. 6. Acceptability ratings for the strong analogy items (root-initial mm) as a function of root type and word class in Experiment 2.

source of these effects, however, was quite different than the no-analogy roots. Across root types, root-initial gemination was rejected compared to its no-gemination [$t_s(42) = 5.62$] or final-gemination controls [$t_s(42) = 10.65$]. This rejection, however, is largely due to the second-class items (see Fig. 6). As with the no-analogy and the nt-analogy roots, the roots with the initial-geminate mm were strongly rejected compared to final-gemination [$t_s(84) = 11.77$] and no-gemination [$t_s(84) = 13.54$] controls in the second class. In the first and third classes, there was also some evidence for the rejection of root-initial gemination relative to no-gemination controls [$t_s(84) = 5.36$; $t_s(84) = 1.97$, n.s.; for the first and third classes, respectively]. These findings, however, do not necessarily reflect a sensitivity to the location of geminates in the root. Recall that the rejection of root-initial geminates in the second class is partly due to their adjacency. Likewise, the rejection of root-initial geminates relative to no-gemination controls may be due to a general avoidance of identity rather than the rejection of root-initial identity per sé. Indeed, when root-initial geminates are compared to final gemination controls, the rejection of root-initial gemination is altogether absent [$t_s(84) < 1$, for both the first and third classes]. Thus, the rejection of root-initial gemination is blocked for roots that are analogous to strong counterexamples. Conversely, roots analogous to weak geminates behave just like the no-analogy items.

Do Speakers Avoid Identity?

The asymmetry in the rating of root-initial vs root-final geminates reflects a sensitivity to the location of geminates. Root-final gemination appears more desirable. However, is root-final gemination as desirable as no gemination?

Across analogy type and word classes, root-final gemination was rated significantly lower compared to no-gemination controls [$t_s(42) = 10.79$; $t_i(46) = 6.82$]. The rejection of root-final gemination emerged in each of the three word classes [in the first class, $t_s(84) = 8.718$; $t_i(92) = 16.41$; in the second class, $t_s(84) = 7.67$; $t_i(92) = 4.56$; in the third class, $t_s(84) = 9.501$; $t_i(92) = 6.06$]. The rejection of root-final gemination did not depend on root analogy. Significant rejection of root-final gemination relative to its no-gemination control emerged for either the analogy or the no-analogy items. Specifically, in the no-analogy group, root-final gemination was rejected across word classes [$t_s(42) = 11.55$; $t_i(22) = 6.52$] as well as in each of the word classes separately [in the first class, $t_s(84) = 7.16$; $t_i(44) = 3.604$; in the second class, $t_s(84) = 7.81$; $t_i(44) = 3.93$; in the third class, $t_s(84) = 10.86$; $t_i(44) = 5.46$]. Likewise, for the analogy roots, root-final gemination was rejected across word classes [$t_s(42) = 9.65$; $t_i(22) = 3.84$] as well as within the first [$t_s(84) = 7.11$; $t_i(44) = 3.76$], second [$t_s(84) = 4.69$; $t_i(44) = 2.63$], and third [$t_s(84) = 5.18$; $t_i(44) = 2.89$, n.s.] classes. Thus, despite the clear sensitivity to the location of identity, speakers prefer to avoid it altogether.

Discussion

Experiment 1 demonstrated that triconsonantal roots are generated from biconsonantal inputs by means of rightward reduplication. Experiment 2 extended these findings in two directions. First, we examined whether the avoidance of root-initial geminates in Experiment 1 reflects their ill formedness. Second, we investigated whether the ubiquity of root-final gemination reflects their well formedness. A sensitivity to the constituent structure of geminates and their location relative to the root would implicate the representation of two variables, geminates and root. Our discussion summarizes the evidence for each of these variables.

The Representation of the Root

The findings of Experiment 2 indicate that root-initial gemination is relatively unacceptable. Speakers consistently rate root-initial gemination (e.g., SSM) lower than no-gemination controls that are identical in all but the initial segment (e.g., PSM). Furthermore, root-initial gemination is rejected compared to controls in which the same geminates are root-final (e.g., SMM). Thus, the acceptability of geminates depends on their location. Importantly, the domain defining the location of geminates is the root: Geminates are consistently rejected in root-initial position, regardless of their position in the word (cf. SiSeM, maSSiMim, and hiStaSaMti). Although the rejection of root-initial geminates was general, its magnitude was modulated by word class. Replicating the findings of Berent and Shimron (1997), root-initial gemination in the second class was particularly unacceptable when the geminates were not separated by a full vowel (e.g., maSSiMim). Surface adjacency, however, is not necessary for the rejection of geminates. A similar rejection of root-initial geminates was obtained in the first class (e.g., SiSeM). Even more striking is the rejection of root-initial gemination in the third class, whose morphological structure is highly affixed (e.g., hiStaSaMtem). Despite the extreme opacity of its word structure, speakers are sensitive to root structure, and reject root-initial gemination relative to controls. The consistent rejection of root-initial gemination demonstrates that speakers possess some long term knowledge that constrains Hebrew roots. The root, a variable, is thus mentally represented and serves as the domain of a linguistic constraint.

In addition to an abstract knowledge of root structure, our results also reflect knowledge of specific root tokens. Our materials included two types of roots whose initial bigram matches counter examples to the OCP. The weak analogy roots (*n* or *t-initial* geminates) exhibited essentially the same rejection of root-initial gemination observed for the no-analogy group. In contrast, strong analogy had a pronounced effect on the acceptability of root-initial geminates. In the second class, roots with initial-m geminates (e.g., MMG) were rejected relative to final identity controls (e.g., MGG) only when the geminates were truly adjacent (e.g., maMMigiM). This rejection, however, may stem from the surface adjacency of root-initial geminates, rather than their root location. Indeed, when the root-initial geminates were separated by a vowel, in the first (e.g., MiMeG) and third (e.g., hitMaMaGti) classes, the rejection of these roots was eliminated. Thus, a strong analogy to counterexamples blocks the rejection of root-initial gemination.

Speakers' sensitivity to counterexamples disagrees with the robust tendency to reduplicate the m-initial biconsonantal roots in Experiment 1. This sensitivity suggests that acceptability is determined, at least in part, by statistical knowledge of specific tokens. Under General Discussion, we explain the role of token specific associations and their confinement to the rating tasks. Although our findings clearly implicate an associative mechanism, they are incompatible with a strong pattern associator hypothesis. The systematicity in the rating of root structure, despite the marked differences in word position, suggests that the root, a variable, is mentally represented. The distinction between weak and strong analogy roots is perfectly compatible with this proposal. Recall that the initial geminates in our weak and strong analogy roots are all frequent bigrams word initially. If the rejection of root-initial geminates was merely guided by their word bigram frequency, then the rejection of root-initial gemination should have been blocked for both the weak and the strong analogy roots. Furthermore, neither of these root types should have been rejected when the geminates were word medial, in the third word class. Our findings are incompatible with this prediction. Clearly, the token-specific knowledge guiding performance concerns the root, not the word. Speakers' sensitivity to the frequency of root tokens agrees with our claim that they represent the word's morphological structure.

The Representation of Identity: Evidence from an a-Frequency Effect

The production task used in Experiment 1 provided strong evidence for the hypothesis that geminates are formed by reduplication, i.e., variable copying. However, this task also left us with a puzzle: Despite the strong preference of rightward reduplication, reduplication was not used across the board. In about a third of the trials, speakers preferred to avoid altering the root, even though this strategy resulted in the distortion of the word pattern. This finding suggests that reduplication is not entirely desirable.

Experiment 2 examined the acceptability of root-final identity by comparing their rating to no-gemination controls. Despite their higher bigram frequency, roots with final-gemination were rated as less acceptable than no-gemination controls.[19] Note that the rejection of root-final gemination is not *caused* by their greater frequency[20]: If frequency resulted in rejection, then roots with initial gemination, whose frequency is lower than their controls, should not have been consistently rejected. Thus, the rejection of root-final gemination is simply unrelated to token frequency. The rejection of root-final gemination is robust and replicable: It emerged within each of the word patterns, and it replicates similar findings reported in two of Berent and Shimron's experiments (1997). In the absence of a statistical account for the rejection of root final geminates, this phenomenon can only be explained by appealing to the formal structure of geminates, namely their identity. The rejection of root-final geminates provides strong support for the representa-

---

[19] To assure that the rejection of root-final gemination in our present study is not due to a few rare items, we reanalyzed our data after excluding all root trios whose final gemination member is less frequent than its no-gemination control. This procedure resulted in the exclusion of 8 root trios. For the remaining 16 root trios, the final gemination member was clearly more frequent than its no-gemination control [$\Delta = 5.75$, $t(15) = 5.62$]. Despite their higher frequency, roots with final gemination were still rated significantly lower than their no-gemination controls [$t(15) = 3.169$]. Thus, root-final gemination is rejected despite its *higher* frequency.

[20] One may attempt to account for the rejection of root-final geminates by appealing to their type frequency. Specifically, although root tokens with final geminates are more frequent than no-gemination controls, as a type, root gemination is less frequent than no-gemination. Such an account, however, parts from the assumption that gemination is a structure distinct from nongemination and that speakers note the frequency geminates in the language. Furthermore, this account fails to explain why identity is a less frequent root type in the language and why speakers are sensitive to its presence. Thus, a type frequency account does not present an alternative to the representation of identity by a variable nor does it present a principled explanation for its avoidance in the language.

tion of variables. These findings also carry some specific implications to phonological theory, implications we address under General Discussion.

EXPERIMENT 3

The findings of Experiments 1–2 demonstrate that Hebrew speakers possess a knowledge that constrains root structure. However, the tasks employed in these experiments explicitly require attention to root structure. This limitation of our method is unlikely to affect our conclusions with regards to the existence of root-structure knowledge, as such knowledge clearly could not have been acquired by performing the experimental task. Nevertheless, the demand to compare root structures does limit our ability to assess the role of such knowledge in language processing. Do speakers apply knowledge regarding root structure in processing novel Hebrew words, even when such knowledge is not required in order to perform the experimental task?

Experiment 3 examines the generality of the constraint on root identity. To this end, we performed a simple variation in the rating procedure employed in Experiment 2. Instead of eliciting a relative rating of words that differ only on their root structure, we employ an absolute rating task. We presented our participants with the list of words used in Experiment 2 in a randomized order and asked them to rate the acceptability of each word individually. This task no longer calls for a comparison of words to other words. Furthermore, the words in the list differ on many dimensions other than their root structure. Thus, attending to root structure is no longer required for the performance of the experimental task. Using a similar task, Berent and Shimron (1997) observed a significant rejection of root-initial gemination. The present experiment is designed to replicate this finding. In addition, we explore the effect of counterexamples on the rejection of root-initial gemination. Because absolute rating may be based on either word structure or contents, this task may be strongly sensitive to token specific properties. The analogy of novel roots to counterexamples to the OCP may thus block their rejection. However, if speakers possess a knowledge that bans identity in the root, then this knowledge is most likely to become evident in rating the no-analogy roots.

Method

The design and materials are identical to those described in Experiment 2. The only difference compared to Experiment 2 is in the rating procedure. Participants in Experiment 3 were presented with a randomized list of the 216 words employed in Experiment 1. They were asked to rate each word for its own acceptability, rather than its acceptability compared to other words, as required in Experiment 2. We used a 5-point rating scale, with 1 indicating an impossible Hebrew word, 2 indicating a word that does not sound good, 3 indicating a word that sounds strange, 4 indicating a word that sounds good, and 5 indicating a word that sounds excellent.

TABLE 10
Mean Acceptability Rating as a Function of Word Class and Analogy in Experiment 3

	Analogy	No analogy
Word class 1	3.926	3.692
Word class 2	3.284	3.281
Word class 3	3.031	3.022

Participants. Participants were 24 native Hebrew speakers. They were all students at the School of Education, University of Haifa, and received no compensation for their participation.

Results

The effect of analogy on acceptability. To examine the effect of analogy on acceptability, we submitted the mean ratings to separate ANOVAs by participants and items [2 (analogy) × 3 (word class) × 3 (root type)]. For the sake of simplicity, we first focus on the effects involving the analogy factor. The analysis by participants revealed a significant main effect of analogy [$F_s(1, 23 = 6.839)$, $MS_e = .106$; $F_i(1, 22) < 1$], an interaction of analogy by word class [$F_s(2, 46) = 6.713$, $MS_e = .093$; $F_i(2, 44) < 1$] and a three-way interaction of analogy × root type × word class [$F_s(4, 92) = 4.428$, $MS_e = .071$; $F_i(4, 88) = 1.006$, $MS_e = .098$, $p = .4087$]. Across root types and word classes, words conjugated from roots analogous to counterexamples were rated higher than their no-analogy controls. This advantage, however, was clearly due to the first word class (see Table 10). Analogy roots in the first class were rated higher than the no-analogy roots, a difference significant by participants only (Tukey, HSD tests). Conversely, in the second and third classes, ratings for the analogy and no-analogy words were virtually identical. The significance of the three-way interaction in the analysis by participants, however, suggested that analogy and word class also modulate the effect of root type. To further investigate this effect, we submitted the ratings for the analogous and nonanalogous words to separate ANOVAs.

The acceptability of no-analogy roots. The ANOVAs performed on the no-analogy roots revealed significant main effects of root type [$F_s(2, 46) = 74.547$, $MS_e = .144$; $F_i(2, 22) = 24.365$, $MS_e = .220$], word class [$F_s(2, 46) = 13.02$, $MS_e = .63$; $F_i(2, 22) = 14.846$, $MS_e = .277$], and their interaction ($F_s(4, 92) = 19.667$, $MS_e = .108$; $\mathrm{F}_i(4, 44) = 10.887$, $MS_e = .097$]. Across root types, the first word class was rated higher than the second and third word classes, which did not differ significantly (Tukey HSD tests). Importantly, word acceptability was affected by root structure (see Fig. 7). Root-initial gemination was rated significantly lower than either root-final gemination [$t_s(46) = 9.17$; $t_i(22) = 5.17$) or no-gemination [$t_s(46) = 11.56$; $t_i(22) = 6.65$] controls. The rejection of root-initial gemination further

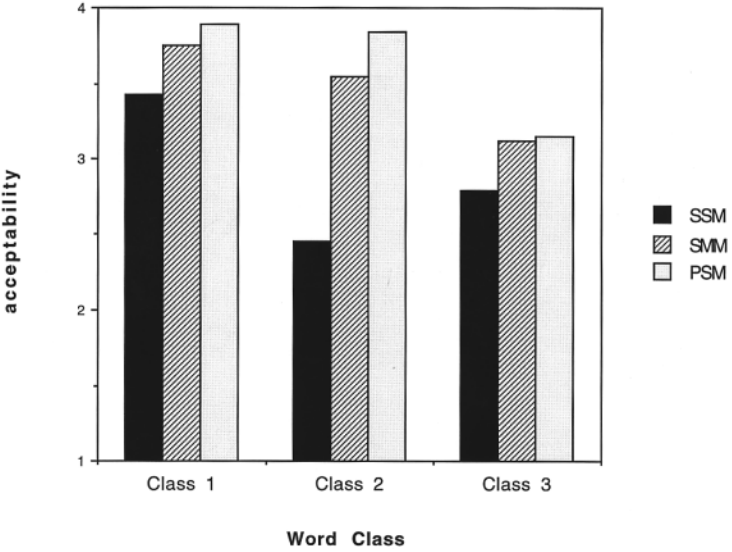

FIG. 7. Acceptability ratings as a function of root type and word class for the ‘‘no analogy’’ items in Experiment 3.

emerged in each of the word classes separately. Specifically, the analyses by participants revealed significantly lower ratings for root-initial gemination compared to its root-final and no-gemination controls, respectively, in the first [$t_s(92) = 3.41$; $t_s(92) = 4.88$], second [$t_s(92) = 11.51$; $t_s(92) = 14.56$], or the third [$t_s(92) = 3.46$; $t_s(92) = 3.74$] word classes. The rejection of root-initial gemination was also significant by items in the second class [$t_i(44) = 8.36$; $t_i(44) = 10.926$], and approached significance in the first [$t_i(44) = 2.54$, n.s.; $t_i(44) = 3.63$] and third [$t_i(44) = 2.59$, n.s.; $t_i(44) = 2.78$, n.s.] classes. However, there was no evidence for the rejection of root-final geminates, either across word classes [$t_s(46) = 2.39$, n.s.; $t_i(22) = 1.47$, n.s.] or within the first or third word classes separately (all $t_s < 1.5$). The only hint of the unacceptability of root-final geminates emerged in the second class, in the analysis by participants only [$t_s(92) = 3.05$; $t_i(44) = 2.57$, n.s.].

The acceptability of analogy roots. Separate analyses of the acceptability of roots analogous to counterexamples [3 (root type) × 3 (word class) × 2 (analogy type; m vs (n-t)] revealed a significant main effect of root type [$F_s(2, 46) = 43.862$, $MS_e = .341$; $F_i(2, 20) = 20.253$, $MS_e = .184$), word class ($F_s(2, 46) = 25.375$, $MS_e = 1.203$; $F_i(2, 20) = 20.299$, $MS_e = .377$], and their significant interaction [$F_s(4, 92) = 38.642$, $MS_e = .188$; $F_i(4, 40) = 18.304$, $MS_e = .099$]. Words in the first class were rated higher than in the second and third classes (Tukey HSD tests). Importantly, speakers were clearly sensitive to root structure as well. As in the no-analogy items, there was a significant rejection of root-initial gemination compared to its final

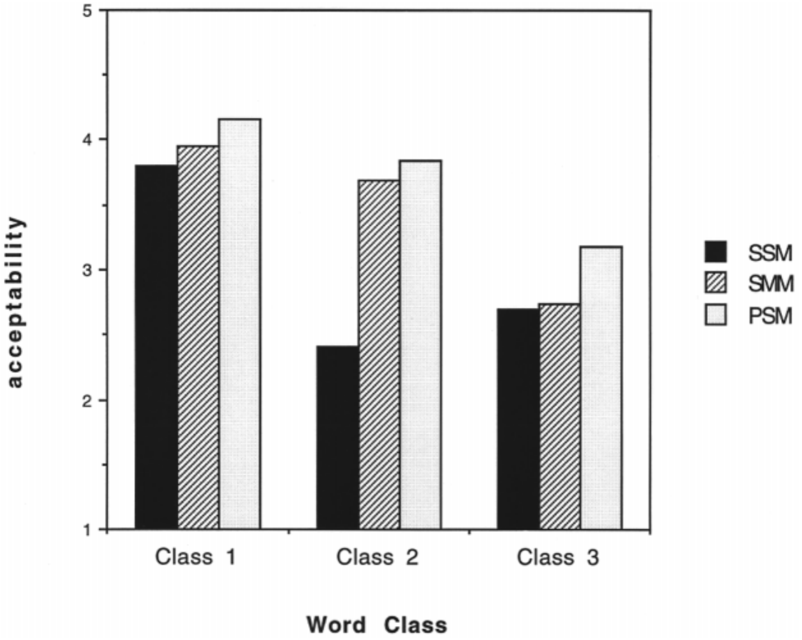

FIG. 8. Acceptability ratings as a function of root type and word class for the analogy items in Experiment 3.

gemination [$t_s(46) = 7.19$; $t_i(20) = 4.87$] and no-gemination [$t_s(46) = 8.79$; $t_i(20) = 5.98$] controls and no evidence for the rejection of root-final geminates [$t_s(46) = 1.59$, n.s.; $t_i(20) = 1.11$, n.s.]. In contrast to the no analogy roots, however, the rejection of root-initial gemination was confined to the second word class (see Fig. 8). In particular, the ratings of root-initial gemination did not differ significantly from their final gemination controls in either the first [$t_s(92) = 1.66$, n.s.; $t_i(40) = 1.13$, n.s.] or third classes [$t_s(92) < 1$; $t_i(40) < 1$]. Likewise, root-initial gemination was not rejected compared to its no-gemination control in the third class [$t_s(92) = 1.803$, n.s.; $t_i(40) = 1.24$, n.s.], and was only marginally less acceptable in the first [$t_s(92) = 3.14$; $t_i(40) = 2.16$, n.s.]. The blocking of the rejection of root-initial gemination for these classes emerged regardless of the type of analogy, for either the m analogy roots or the n-t analogy roots.[21] Conversely, strong rejection

---

[21] The analysis by participants suggested a modulation of the effects of root type and word class by the type of analogy (m vs n-t). In particular, the interaction of root type x analogy [$F_s(2, 46) = 10.357$, $MS_e = .111$; $F_i(2, 10) = 1.589$, $MS_e = .184$, $p = .2289$] and the three-way interaction of word class × root type × analogy were significant by participants only [$F_s(4, 92) = 3.693$, $MS_e = .103$; $F_i(4, 40) < 1$]. To investigate the source of these interaction, we conducted separate analyses on the m and nt type roots. The results obtained for the separate analyses of the m and nt roots were essentially identical. Each of the analogy groups revealed a significant main effect of root type [for m-roots; $F_s(2, 46) = 30.065$, $MS_e = .184$; $F_i(2, 10)$

of root-initial gemination was observed in the second class compared to either root-final gemination [$t_s(92) = 14.39$; $t_i(40) = 9.89$] or no-gemination [$t_s(92) = 15.57$; $t_i(40) = 10.70$] controls. Thus, the analogy to counterexamples blocks the rejection of root identity. Geminates are rejected only if they are adjacent.

Discussion

Experiment 3 examined speakers' sensitivity to root structure in the absolute rating task. The results of the no-analogy roots replicate Experiments 1–2 in demonstrating the existence of an abstract knowledge that bans root-initial gemination. As in our previous experiments, the rejection of root-initial geminates emerged regardless of their location in the word. The novel contribution of Experiment 3, however, is in demonstrating root structure knowledge in a task that does not require explicit attention to the root. These results suggest that the contribution of root structure knowledge in processing Hebrew words may be general.

However, our investigation indicates some limitations in the application of this knowledge. These limitations do not concern task demands or word structure, but instead the similarity of the roots to stored root tokens. Replicating the findings of Experiment 2, the similarity of the experimental roots to strong counterexamples (the root-initial *mm* words) essentially blocked the rejection of root-initial gemination. In contrast to the findings of Experiment 2, however, the blocking of the constraint on root structure was ob-

---

$= 10.830$, $MS_e = .128$; for n-t roots: $F_s(2, 46) = 39.467$, $MS_e = .268$; $F_i(2, 10) = 10.969$, $MS_e = .240$) and its interaction with word classes [for m-roots; $F_s(4, 92) = 31.763$, $MS_e = .133$; $F_i(4, 20) = 12.916$, $MS_e = .082$; for n-t roots: $F_s(4, 92) = 21.676$, $MS_e = .158$; $F_i(4, 20) = 7.358$, $MS_e = .117$]. Across word classes, root-initial gemination was rejected compared to its final gemination [for m-roots: $t_s(46) = 4.22$; $t_i(10) = 7.03$; for n-t roots: $t_s(46) = 5.365$; $t_i(10) = 2.95$] and no-gemination controls [for m roots: $t_s(46) = 3.81$; $t_i(10) = 6.35$; for n-t roots: $t_s(46) = 8.75$; $t_i(10) = 4.62$]. In each analogy group, however, the rejection of root-initial gemination was due to the second word class, in which the geminates were adjacent in the surface The principal difference between the n-t and the m-analogy roots concerns the rejection of root-final gemination. The n-t roots manifested a rejection of root-final gemination compared to its no-gemination control in the analysis by participants [$t_s(46) = 3.10$; $t_i(10) = 1.67$]. No evidence for the rejection of root-final gemination emerged for the m-analogy roots [$t_s(46) < 1$; $t_i(10) < 1$]. An inspection of the cell means suggested that the rejection of root-final identity among the n-t roots is due to the third word class ($\Delta = 0.4408$). To investigate the source of this spurious effect, we compared the two types of roots, the n vs the t roots (each consisting of three roots). This analysis by participants revealed a significant interaction of root type by analogy [$F_s(2, 8) = 5.482$, $MS_e = .126$] and a marginally significant interaction of root type × word class × analogy [$F_s(4, 16) = 2.681$, $MS_e = .0695$]. These interactions indicate that the rejection of root-final gemination in the third class was specially due to roots whose initial consonant was t (e.g., tgg). Recall that the third word class is conjugated by affixing these roots with the prefix hit (e.g. hit-Ta-GeG). The elevated rejection of these items may reflect the undesirability of the identity between the prefix's last consonant and the root's initial consonant, which form adjacent surface geminates.

served also for the weak type of analogy, the root-initial *nn* and *tt* roots. Thus, the use of an open ended rating procedure appears to have increased speakers' sensitivity to token properties.

Two additional differences between our present findings and Experiment 2 are worthy of mention. First, the use of an open ended rating procedure revealed some clear preferences with regards to word class. Replicating the findings of Berent and Shimron (1997), speakers prefer words in the first over the second and third classes. Second, the use of an open-ended rating procedure also eliminates the rejection of root-final gemination. There are several possible explanations for this finding. When presented with the nullification of an effect, one would most naturally question its reliability. Specifically, the nullification of final identity avoidance may stem from some random variability. Although such an explanation can never be ruled out, it is highly unlikely. The rejection of root-final gemination in the relative rating task is highly reliable and replicable. A highly significant rejection of root-final identity has been previously documented in three separate studies (Berent and Shimron, 1997, Experiment 1 and its replication; Berent, Everett, & Shimron, in preparation). In addition, these findings nicely complement the avoidance of reduplication in the production task (Experiment 1). Thus, the aversion of root-final identity in the relative rating task is a robust phenomenon. Equally consistent, however, is its absence in the absolute rating task. The rejection of root-final identity in relative rating and its acceptance in absolute rating replicate the findings of Berent and Shimron (1997). To demonstrate that the rejection of final identity is systematically linked to the rating procedure, we conducted a meta-analysis comparing Experiments 2–3. The ANOVA [2 (experiment) × 3 (word class) × 3 (root type)] revealed a significant interaction of root type × experiment. Tukey HSD comparisons revealed that root-final identity was rated significantly lower than its no-gemination control in Experiment 2 ($p < .01$ by participants and items), but not in Experiment 3 ($p > .05$). Thus, the acceptance or rejection of root-final identity is directly linked to the nature of the rating judgment.

Why does the rating procedure affect the rejection of root-final identity? Could the relative rating task encourage or even cause the rejection of root-final identity? We believe that the view of identity aversion as a task artifact is improbable. It is difficult to see how identity aversion could have been elicited by the relative rating task. In fact, as we argued earlier, the task did not even require attention to identity. Participants in this task could have discriminated between the root trios by attending to token frequency. Such a strategy, as we explained, should have resulted in a higher acceptability of root-final geminates. The discrepancy between the two tasks also does not seem to stem from differences in their sensitivity. The absolute rating procedure is just as sensitive to root-initial identity as the relative rating technique. In fact, the absolute rating technique is more sensitive to word-class properties and token characteristics. Thus, the rejection of root-final identity in Experiment 2 is neither caused nor encouraged by the experimental task. Its absence in Experiment 3 is unlikely due to its insensitivity. Under General Discussion, we outline a principled account for this finding in terms of Comparative Optimality. At present, suffice it to say that root-final identity is rejected only when it is directly compared to no-gemination controls. In contrast, the rejection of initial identity is observed regardless of task demands in each of our three experiments.

Our results converge with existing research in Hebrew (Bentin & Feldman, 1990; Deutsch, Frost, & Forster, 1998; Feldman & Bentin, 1994; Feldman, Frost, & Pnini, 1995; Frost, Forster, & Deutsch, 1994) obtained using online measures in suggesting that the decomposition of words into their morphological constituents may be general. Because existing Hebrew roots correspond to frequent units at the orthographic, phonological, and semantic levels, evidence for decomposition is often ambiguous with regards to the nature of the decomposed units: abstract variables vs specific tokens (Rueckl et al., 1997). Our present findings contribute to the resolution of this ambiguity. We demonstrate that the unit decomposed in processing Hebrew words is constrained by knowledge that applies to variables. Thus, variables affect the processing of Hebrew words in the present experimental task. Because this task did not require attention to the root, our findings suggest that the contribution of these symbolic processes may be general and may not depend on speakers' control.

GENERAL DISCUSSION

This set of experiments addressed a fundamental question for theories of language: Do linguistic representations specify variables? Variables are central to symbolic accounts of cognition, but are considered obsolete according to the pattern associator hypothesis. To test these hypotheses, we examined the role of two variables in the representation of Hebrew words: the root morpheme and identity. We open our discussion by reviewing the evidence in support of the representation of the root and geminates by variables. We then proceed to address a second core question: Does the specification of variables entail inviolable rules? Although our findings support the symbolic hypothesis, they are clearly incompatible with the view of linguistic knowledge as governed by inviolable rules. We propose an alternative linguistic account for our findings, an account couched in terms of a theory of inviolable constraints, Optimality Theory. Finally, we consider the role of statistical structure in the acquisition of lexical representations. Despite the support provided for variables, our data also strongly implicate the role of an associative mechanism in the acquisition of lexical representations. We propose that the acquisition of lexical representations is achieved by both a default, symbolic process, and an associative network. We explore the division of labor between these processes and discuss its implications to the acquisition of lexical representations and language processing.

Do Linguistic Mental Representations Specify Variables?

Our three experiments demonstrate a strong asymmetry in the production and rating of root-initial vs root-final geminates. These results implicate the representation of variables by linguistic competence. Before reviewing the evidence for variables, we wish to clarify the relevance of our conclusions to linguistic competence. Our present experiments all used off-line methods. One may thus criticize our conclusions as reflecting metalinguistic knowledge rather than linguistic competence. Our findings are incompatible with this proposal. First, the constraint on root structure is not patent to Hebrew speakers, nor is it explicitly taught in the school system. Participants are thus typically unaware of the constraint on root structure and are unable to explain the unacceptability of words with root-initial geminates (Berent & Shimron, 1997). Second, the findings obtained in the rating and production tasks replicate in on-line methods. Because nonwords generated from novel roots with initial gemination are ill formed, their rejection in a lexical decision task is significantly faster compared to roots with final gemination (Berent, Shimron, & Vaknin, in press). Likewise, words with root-initial geminates exhibit lesser interference with color naming in the Stroop task (Berent, Bibi, & Tzelgov, 2000). The sensitivity to root structure in situations in which such behavior is clearly contrary to task demands is incompatible with its attribution to a deliberate metalinguistic strategy. The convergence between these findings and our current results suggests that they reflect the contents of linguistic competence. We further claim that such knowledge appeals to two variables, the root morpheme and identity. The representations of variables may thus be critical for the adequacy of theories of linguistic competence.

The representation of the root. Our findings indicate a marked asymmetry in the acceptability of geminates in Hebrew roots: Root-final gemination is preferred over root-initial gemination. This finding was observed as an overwhelming asymmetry in the production of words from biconsonantal roots (in Experiment 1) and in rating these outputs (in Experiments 2–3). Importantly, the unacceptability of root-initial gemination was observed regardless of its location in the word: Root-initial gemination was rejected or avoided when the geminates were word initial (e.g., SiSeM) or word internal (e.g., maSSiMim and hiStaSaMtem). Thus, the knowledge guiding this behavior cannot be defined in reference to the word. Instead, it is the root, an abstract variable, that constitutes the domain of speakers' knowledge.

We claim that a constraint on root structure challenges connectionist models that eliminate variables. To acquire a constraint on root structure, pattern associators[22] must be able to decompose Hebrew words into their morpholog-

---

[22] To reiterate, pattern associators should not be equated with any specific brand of connectionist models: we use pattern associators to refer to the class of connectionist models that eliminates variables.

ical constituents. This initial step is not necessarily unattainable in the absence of an explicit root variable, as Hebrew affixes often correspond to sub-word units associated with some relatively invariant orthographic, phonological, and semantic features. In many words, bare root forms could thus be extracted by attending to the statistical structure of Hebrew words. However, a statistical decomposition process is not infallible. One challenge for a statistical decomposition is the presence of linearly discontinuous roots. For instance, the metathesis between the affix and the root-initial radical in hiStaSaMtem can easily trick a statistical procedure into parsing the words as composed of the prefix *hi* and the root *stsm*, a quadruple root that contains no geminates, hence, does not violate the OCP. The results of Berent and Shimron (1997) demonstrate that speakers are not led the garden path by such words. Decomposition, however, is only the first step in the acquisition of root structure constraints. The main challenge facing pattern associators is to acquire adequate *generalizations* over root tokens. Note that the set of all existing triliteral Hebrew roots share no orthographic, phonological, or semantic feature. The only common denominator to these items is formal. There are currently no computational models of the OCP. Thus, the ability of pattern associators to acquire the constraint on root structure is yet uncertain. Specifically, it is presently unclear whether pattern associators can decompose linearly discontinuous roots, treat dissimilar root tokens as members of a single class, abstract significant generalizations regarding their structure, and ignore their specific idiosyncrasies. We hope that our findings encourage a systematic investigation of these challenges.

The representation of identity. Our experiments indicate that Hebrew speakers are not only sensitive to root structure, but they specifically constrain identity within the root. Our findings provide two sources of support for the representation of identity by a variable: Identity avoidance and identity formation. Identity avoidance was evident in the ubiquity of no-gemination responses in Experiment 1. These responses were puzzling, since they resulted in the distortion of the word pattern. Additional evidence for identity aversion was observed in Experiment 2. The relative rating task reflected a reliable and consistent rejection of roots with final identity compared to no-gemination controls despite their higher token frequency. Such a distinction between geminate and nongeminate bigrams is only explicable in reference to their structure.

Although the evidence for identity avoidance is highly reliable, its scope is limited. The results of Experiment 3 provide no evidence for the avoidance of root-final identity. We return to discuss the contrast between Experiments 2–3 in ‘‘An Alternative: Violable Constraints on Identity.’’ Even if one rejects our account and chooses to conclude that identity avoidance is a limited phenomenon, such conclusion would not obviate the need to represent variables. Indeed, our results provide a second, independent evidence for the representation of geminates by a variable. This evidence concerns identity formation. According to McCarthy (1986), final identity in triliteral surface roots is formed by reduplication of biconsonantal inputs. The findings of Experiment 1 provide strong support for the use of root-final gemination as the principal, perhaps the only, strategy for the formation of triliteral roots from their biconsonantal input. Importantly, the formation of root-final geminates cannot be attributed to segment addition. Gemination was the most frequent response despite the fact that its expected probability is far lower than the addition of a new segment. Furthermore, the formation of geminates was unaffected by their similarity to counterexamples. Thus, our findings suggest that geminates are formed by a symbolic process of reduplication: A process of variable copying that is blind to segment contents or the frequency of the resulting bigram.

The representation of identity has been systematically investigated by Marcus (1998a, 1998b, in press). Marcus demonstrated that a simple recurrent network and feedforward networks fail to generalize a simple identity function such as ‘‘an X is an X.’’ After training the network on nine instances of this function (e.g., *a rose is a rose; a tulip is a tulip),* it was presented with a 10th new token (e.g., a *lilac*). Despite excellent performance on the 9 trained items, the network failed to generalize to the 10th item. This failure is robust and principled. The failure to generalize does not depend on the number of training examples, the parameters of learning rate and momentum, the number of hidden units or hidden layers. It is also not due to a failure to represent the novel token. To rule out a representational failure, Marcus trained a single network on two frames. Frame A was *The bee sniffs the X*, whereas frame B was the identity function (*an X is an X*). The network was trained on the instance *lilac* on frame A (*The bee sniffs the lilac*), but not on frame B. Marcus observed that, despite the successful performance with *lilac* in the context of frame A, the network was unable to generalize the identity function to this item.

Marcus (1998a, 1998b, in press) concluded that the robust failure to generalize the identity function is principled. Pattern associators cannot generalize to items containing features on which the model was not trained. This limitation stems from their training independence: The setting of weights for any given unit (input or output) to the hidden units is independent of the other units. Thus, these models cannot exploit trained nodes to constrain the activation of untrained nodes. For some functions, this limitation is nondetrimental. Pattern associators will perform adequately on any function that does not require generalization outside the training space. Identity, however, is not one of those functions. Identity applies to any segment, regardless of its content . Because identity concerns variables, rather than tokens, its relevant training space, is, by definition, infinite. For functions defined over variables, the failure of pattern associators to generalize outside its training space is virtually guaranteed. One may criticize the relevance of Marcus’ formal analysis to our present results on the grounds that they do not directly demonstrate generalization outside the training space. The linguistic experience of our adult participants has effectively resulted in their ‘‘training’’ on all Hebrew geminates. Our tasks thus did not require that they extend this generalization to untrained phonemes or features. Although our present experiments do not probe for the generalization of reduplication to *untrained* items, our investigation nevertheless systematically assessed whether this behavior is sensitive to the *degree* of training, i.e., the familiarity with trained items. This approach is perfectly compatible with the pattern associator hypothesis. For instance, Rhode and Plaut (1999) have recently noted that: ‘‘language development depends critically on the *frequency* with which forms occur in the language and not simply on whether or not they occur at all’’ (p. 98). There are indeed numerous documentations of the sensitivity of pattern associators to type frequency (e.g., Daugherty & Seidenberg, 1992; Plunkett & Marchman, 1993; Rumelhart & McClelland, 1986). If Hebrew geminates were represented as the association between stored tokens, then the probability of forming geminates and their acceptability should have reflected their frequency. The marked discrepancy between the expected and observed probability of geminate formation in Experiment 1, insensitivity to counterexamples, and the phenomenon of identity aversion in Experiments 1–2 demonstrate speakers’ insensitivity to the degree of familiarity with trained items.

Given that reduplication extends regardless of the familiarity with trained phoneme combinations, we expect this behavior to emerge also for untrained phonemes. Our subsequent investigation tested this prediction (Berent, Marcus, & Shimron, 2000). This experiment obtained ratings of novel words generated from novel roots that include foreign phonemes, i.e., segments that do not form part of the Hebrew inventory (e.g., ch, th, w, j). As in our present experiments, we constructed root trios in which one of these foreign phonemes appeared in root-initial geminates (e.g., ChChK), root-final geminates (e.g., KChCh), or in a nongemiante root-initial bigram (e.g., ChPK). These roots were conjugated such that root-initial geminates were either word initial (e.g., ChiCheK) or word internal (e.g., hiChtaChaKtem). Participants were presented with a randomized list of these words. They were asked to read each word aloud and rate its absolute acceptability. There are no existing Hebrew roots with geminates including these phonemes; hence, a statistical knowledge cannot constrain the location of foreign geminates in the root. The results nevertheless replicate the findings of our present Experiment 3. Speakers rate root-initial gemination as significantly less acceptable than either root-final gemination or no-gemination controls, and this finding emerged regardless of the location of geminates in the word. Furthermore, the rejection of root-initial gemination was obtained even when the foreign phoneme includes phonetic features that do not occur in Hebrew. Such phonemes clearly fall outside the training space of Hebrew speakers. Participants’ sensitivity to the structure of roots with foreign features demonstrates their ability to extend the constraint on root identity outside their training space. Marcus' (1998a) formal proof suggests that such behavior may be unattainable by pattern associators.

In summary, the pattern associator hypothesis attributes linguistic knowledge to the statistical properties of the linguistic input (e.g., Elman, 1993; Plaut et al., 1996; Seidenberg, 1997; Rhode & Plaut, 1999). Our investigation includes numerous assessments of the statistical structure of the Hebrew language. None of these analyses explains our data. In view of the known sensitivity of pattern associators to the statistical structure of the language, our findings raise doubts regarding the ability of pattern associator models to account for the OCP. Furthermore, the problem of root decomposition and generalization across dissimilar root tokens, coupled with Marcus' (1998a) formal analysis of the identity function, suggest that the phenomenon we observe may tap into some principled limitations of these models. In the absence of a systematic computational investigation of the OCP phenomenon, we currently cannot rule the possibility that pattern associator models may rise to these challenges. We certainly do not claim that this constraint is unrepresentable by connectionism. Instead, we believe that, in order to represent these phenomena, the network may require the implementation of two variables: the root and identity.

Evidence against Inviolable Rules

The constraint on root structure provides strong support for the representation of variables. However, our data are incompatible with the attribution of this knowledge to inviolable rules. Our findings present two challenges to the view of the OCP as an inviolable constraint. One challenge concerns OCP counterexamples. Hebrew contains about four roots with initial gemination, two of them are highly productive and familiar. One could try to maintain the view of OCP as inviolable by attributing these violations to a few stipulations in the lexicon. However, this solution would not do. The acceptance of root-initial geminates is not limited to the handful of existing Hebrew roots, but is productively applied to novel roots that are analogous to these counterexamples. Our observation of productive violations of the lexical OCP agrees with the violations of identity avoidance as a derivational constraint and a rule triggering device which have been amply discussed in the linguistic literature (e.g., Odden, 1986; Yip, 1988; Myers, 1994). These observations are incompatible with the view of the OCP as inviolable.

An ever stronger challenge to McCarthy's (1986) account is the rejection of root-final geminates. Experiment 1 reflected a curious aversion of reduplication. A similar avoidance of root-final identity emerged in Experiment 2 as a rejection of root-final geminates relative to their no-gemination controls. These findings are inexplicable by McCarthy's (1986) account. According to this proposal, root-final geminates do not require the lexical specification of identity; hence, they should be as acceptable as their no-gemination controls. To account for the aversion of root-final identity, one may propose a constraint against root-final gemination. Such a solution, however, suffers from a deep incompatibility with McCarthy's (1986) proposal. In contrast to root-initial gemination, root-final gemination is extremely frequent in Semitic. If one were to ban root-final gemination, then this constraint would be routinely violated. It is thus impossible to capture the desired generalization without making incorrect, too strong, predictions under the view of phonological principles as inviolable. In contrast, these results may be readily explained by a theory of violable constraints, Optimality Theory.

An Alternative: Violable Constraints on Identity

Our results demonstrate that Hebrew speakers' knowledge appeals to the representation of abstract variables. However, the constraints governing these variables are violable. To accommodate these findings, we propose an account couched within a theory of violable constraints, Optimality Theory (Prince & Smolensky, 1993, 1997; for additional discussions of the OCP in optimality theory, see Gafos, 1998; Myers, 1994, Yip, 1995). In what follows, we briefly describe our account. The goals of the following discussion are twofold. First, we wish to illustrate how the constraint on root identity may be accommodated within a theory of violable constraints. Second, we account for our empirical findings that are at odds with McCarthy's (1986) proposal. These include the rejection of root-final identity and its confinement to the relative rating procedure.

Optimality theory accounts for linguistic phenomena without assuming rules or derivations. Instead, phonological representations are generated and evaluated with respect to a ranked set of violable constraints. The constraints are considered universal, but their ranking is language specific. All constraints apply directly to phonological outputs. Specifically, linguistic objects are produced by two functions: GENERATE and EVALUATE. GENERATE computes a set of outputs to a given input. These outputs are then evaluated for their optimality with respect to the constraint ranking by the EVALUATE function. The candidate that violates the minimum number of highest ranked constraints is considered optimal. Importantly, the optimality of a candidate output is strictly relative: An optimal output can, and often does, violate constraints. Its optimality simply reflects the fact that its constraint violation is less severe than other candidates.

Our specific account for the experimental results assumes a family of constraints that ban identity from phonological representations (see also Everett & Berent, 1998). IDENTITY bans identity from the root. A higher ranked constraint, ADJACENT IDENTITY, specifically bans adjacent identical elements in phonological representations. Although identity is undesirable, identity (like any other constraint) is violable when outranked by higher constraints. The primary motivation for identity formation is the need to supply Hebrew words with three root consonants. This pressure may be attributed to either a constraint on syllable structure which requires bisyllabicity

TABLE 11
The Optimality of Three Candidates (SiSeM, SiMeM, and SiMeT) with Respect to Two Constraint Families, CORRESPONDENCE and IDENTITY

Output (input: SM, _i_e)	Constraints	
	CORRESPONDENCE	IDENTITY
SiMeM		*
SiSeM	!*	*
SiMeT	!*	

Note. Constraint ranking is indicated by their left-to-right order. A constraint violation is indicated by an asterisk, and fatal violations are indicated by an exclamation mark.

(Gafos, 1998) or to the full alignment of the three consonant slots in the word pattern with the root input (i.e., MAXIO, see Everett & Berent, 1998). The need for an extra root consonant may be met in several ways. Table 11 illustrates several output candidates for the biconsonantal root SM and the word pattern _i_e_. It describes the optimality of these candidates with respect to two constraint families, CORRESPONDENCE and IDENTITY. In this and all subsequent tables, the ranking of the constraints is indicated by their left-to-right order (highly ranked constrained are to the left). A constraint violation is indicated by a star, and a fatal violation is indicated by an exclamation mark.

Each of the candidates in Table 11 represents a solution to the problem of aligning a biconsonantal root and the three root slots in the word pattern. One solution is the insertion of a new consonant (e.g., SiMeT). This solution, however, violates the highly ranked family of constraints requiring the CORRESPONDENCE of the output to the input.[23] To avoid its violation, speakers may obtain the additional output segment by means of reduplication. This results in either SiSeM or SiMeM. According to Gafos (1998), leftward reduplication in SiSeM fatally violates two other members of the CORRESPONDENCE constraint family.[24] Leftward reduplication is thus avoided. In contrast, SiMeM, generated by rightward reduplication does not violate any CORRESPONDENCE constraints. Thus, despite the violation of the lower ranked IDENTITY constraint, SiMeM is the optimal candidate.

The principal entailment of our account is that reduplication is the optimal strategy for the formation of a word from a biconsonantal root. This prediction agrees with the results of our production experiment. When presented

---

[23] Specifically, the DEPIO constraint requires that output segments correspond to input segments.

[24] These constraints govern the correspondence between the base and affix: The ANCHOR-R constraint requires that the rightmost element of the reduplicative affix correspond to the rightmost element of the base and the ALIGN$^{\text{AFF-R}}$– R constraint aligns the reduplicated consonant with the right edge of output.

with a biconsonantal root, speakers solve the mismatch between the word pattern and the input root by means of reduplication rather than segment addition. Our account also states that reduplication proceeds rightward. Accordingly, the production task reflected a robust preference for reduplicating root-final over root-initial radicals. Converging evidence for the optimality of rightward reduplication is found in the rating tasks. Nonwords rating, however, is more complex than in their production. Before we can examine the rating results, we must first address the demands of the rating task.

In the rating task, speakers assess the acceptability of novel words derived from novel roots. We propose that acceptability reflects the optimality of a word as an output of its lexical representation. Thus, to determine the acceptability of a word, speakers must first identify its lexical representation. In contrast to the production study, participants in the rating studies are not provided with a lexical input. They also cannot simply retrieve it from the lexicon, since these forms are all novel. To evaluate the acceptability of a word, speakers must *infer* its lexical representation. This task resembles the acquisition of new roots by the child. Smolensky (1996) proposed that the child identifies the lexical input by applying structural constraints to its surface form. In a similar vein, participants in the experiment can infer the lexical representation of surface forms by applying to them the constraint ranking. The optimal output of this computation will be stored as the word's lexical representation. Specifically, the optimal lexical input for roots with geminates (initial or final) is biconsonantal because it does not violate any identity constraint. Armed with this ''lexical'' input, speakers can now turn to evaluate the optimality of the target word as an output of that ''lexical'' input. Forms with root-initial gemination are rejected because they fatally violate CORRESPONDENCE . Conversely, root-final gemination is optimal; hence, these forms are rated higher than the initial gemination counterparts.

Importantly, however, within a theory of violable constraints, the optimality of root-final gemination does not imply no constraint violation. Recall that root-final gemination violates IDENTITY. Because the identity family is dominated by the CORRESPONDENCE constraint family, words exhibiting root-final geminates are optimal candidates despite the violation of at least one member of the identity family. However, IDENTITY violations nevertheless carry consequences that may be observed in certain experimental settings. In the absolute rating task, speakers determine whether SiMeM is optimal relative to its inferred input, SM. The answer to that question is clearly yes, since no better candidate is available. However, in the relative rating experiment, speakers are asked to *compare* the optimality of different words. For instance, they compare the optimality of SiMeM and PiSeM. To perform this task, speakers cannot simply compare the optimality of SiMeM relative to its input, SM, as required in the absolute rating task. Instead, they must determine how the optimality of SiMeM as an output of SM compares

TABLE 12
The Comparative Optimality of Three Outputs (SiSeM, SiMeM, and PiSeM) Relative to Their Inputs (Not Shown) with Respect to the CORRESPONDENCE Constraint Family and the IDENTITY Constraint

	Constraints	
	CORRESPONDENCE	IDENTITY
SiSeM	!*	*
SiMeM		*
PiSeM		

Note. Constraint ranking is indicated by their left-to-right order. A constraint violation is indicated by an asterisk, and fatal violations are indicated by an exclamation mark.

to the optimality of, say, PiSeM relative to its input, PSM. Table 12 represents the outcomes of comparing SiSeM, SiMeM and PiSeM as outputs of their respective inputs.

On our account, SiMeM (but not PiSeM), violates IDENTITY; hence, SiMeM is less optimal than PiSeM. If speakers can compute comparative optimality, then, despite its Optimality relative to its input, SiMeM should be less acceptable than PiSeM. Because SiSeM violates the higher ranked CORRESPONDENCE constraints, its constraint violation is more severe than either SiMeM or PiSeM, rendering it lowest on the comparative optimality rank. Our account views the divergence between the two rating experiments as principled and theoretically significant. First, this divergence demonstrates that speakers can compare the optimality of different input-output pairs. Second, the rejection of SiMeM despite its optimality and relative frequency suggests that speakers' linguistic competence includes a violable constraint on Identity.

The account proposed so far explains why SiMeM is frequent yet is rejected relative to PiSeM. It also explains why the rejection of SiMeM emerges only when speakers are asked to perform relative rating. However, this account cannot address the effect of counterexamples. Root-initial gemination fatally violates IDENTITY and CORRESPONDENCE. If so, why do participants in the rating tasks accept root-initial geminates when they are analogous to counterexamples?

The Role of Symbolic and Associative Processes in the Acquisition of Lexical Representations

Our results suggest that root identity is productively formed by a symbolic process of rightward reduplication. This finding has direct implications to

TABLE 13

The Lexical Representations Inferred for SiSeM and MiMeS and the Violations of CORRESPONDENCE Incurred by Generating Each of These Words from Their Inferred Lexical Representations

Overt form	Inferred lexical representation	CORRESPONDENCE
SiSeM	SM	!*
MIMES	MMS	

Note. Constraint ranking is indicated by their left-to-right order. A constraint violation is indicated by an asterisk, and fatal violations are indicated by an exclamation mark.

the structure of lexical representations as well. If speakers form root-final identity by means of a productive process, then the lexical representations that serve as the input for this productive process must eliminate geminates. Our findings imply a default symbolic process that erases identity from the lexicon: This process transforms any surface root whose form is XXY or XYY into XY. Its application is sensitive to the form of the input and blind to its contents.

However, the investigation of counterexamples suggests that identity erasure may be overridden. The rejection of root-initial geminates is blocked when their content is similar to OCP violations.[25] To account for these findings, we propose that OCP violations are lexically stored as triliteral roots. These marked forms can block identity erasure by triggering an associative mechanism. Specifically, when presented with a word, speakers decompose its triliteral surface root. If this surface representation is sufficiently similar to a stored OCP counterexample, it blocks the default identity erasure. As a result, the new root is stored in the lexicon in its triliteral form, despite its root-initial identity. In contrast to the default identity erasure, which operates over variables, the associative mechanism is sensitive to token specific properties. It is the similarity to stored tokens that triggers the associative mechanism and permits the acquisition of counterexamples to the OCP.[26]

The division of labor between the associative and default symbolic processes can nicely account for several additional aspects of our results. First, our account explains the acceptance of root-initial gemination that is analogous to counterexamples (see Table 13). Recall that root-initial gemination

---

[25] The integration of terminology from Autosegmental phonology and Optimality theory results in some inconsistency. Our previous discussion has referred to root-initial geminates as OCP violations, following McCarthy (1986). However, within the OT account proposed in the previous section, the rejection of root-initial gemination is attributed to the violation of the Identity and Correspondence constraints. For the sake of simplicity, we keep referring to root-initial gemination as OCP violations, although we no longer adopt the OCP (as described in McCarthy, 1986) as a theoretical explanation for their rejection.

[26] An alternative account may capture the associative mechanism within the grammar by proposing an analogy constraint that is triggered by similarity and overrides Identity. The resolution of these two conflicting accounts goes beyond the scope of the present work.

TABLE 14
CORRESPONDENCE Violations Incurred by Generating MiMeK from Its Lexical Representations (Inferred or Given) in the Rating and Production Tasks

	Lexicon	Output	CORRESPONDENCE
Rating	MMK	MiMeK	
Production	MK	MiMeK	!*

Note. Constraint ranking is indicated by their left-to-right order. A constraint violation is indicated by an asterisk, and fatal violations are indicated by an exclamation mark.

is ill formed because its generation from its inferred biconsonantal root incurs a fatal violation of members of the CORRESPONDENCE family. In contrast, the lexical representations inferred for words analogous to counterexamples are triconsonantal. Thus, these roots can form words without violating CORRESPONDENCE.

The associative process also nicely accounts for the divergence in identity avoidance depending on its specific manifestation and the experimental task. For instance, consider root-initial identity. Experiments 2–3 demonstrate that the rejection of root-initial identity is blocked for counter examples (e.g., MiMeK). Conversely, in Experiment 1, biconsonantal roots whose initial radical is identical to a counterexample (e.g., MK) did not yield root-initial gemination (e.g., MiMeK). If counterexamples block the rejection of root-initial identity, then why don't speakers reduplicate the initial radical in the production task?

This puzzle is readily solved by the distinction between the effects of analogy in lexical acquisition and word formation. According to our account, the associative process can block identity erasure from inferred lexical representations, but cannot affect word formation. The production task provides speakers with biconsonantal roots. Because lexical representations are given, their acquisition cannot be overridden by the associative process. As illustrated in Table 14, the production of root-initial gemination from a biconsonantal input violates CORRESPONDENCE. In contrast, participants in the rating studies are presented with words rather than with roots. To rate these words, they must first infer their lexical representation. It is during the acquisition of a lexical representation that the default and associative processes exercise their effect. Because the surface triliteral representation of counterexamples is similar to stored tokens, it can now activate the associative process and override identity erasure, yielding a triliteral lexical representation. This avoids the violation of CORRESPONDENCE during the formation of a word and permits generating words from these candidates. As a result, rating, but not production, is affected by similarity to counterexamples.

A third contribution of the associative process in accounting for our results concerns the disagreement between the distinct manifestations of identity avoidance. Our investigation included different forms of identity: root-initial

TABLE 15
The Violations of CORRESPONDENCE and IDENTITY Constraints Incurred by Generating Overt Forms from Their Inferred Lexical Representations

Overt form	Lexicon	Correspondence	Adjacent identity	Identity
maMMiKot	MMK		!*	*
MiMeK	MMK			*
MiKeK	MK			*
PiSeM	PSM			

Note. Constraint ranking is indicated by their left-to-right order. A constraint violation is indicated by an asterisk, and fatal violations are indicated by an exclamation mark.

identity among nonadjacent consonants, root-initial identity among adjacent consonants, and root-final identity among nonadjacent consonants. In each of these forms we observed some evidence for identity avoidance. However, these distinct manifestations do not always co-occur. In particular, despite blocking the rejection of root-initial gemination in our two rating experiments for counterexamples (e.g., MiMeK), the same roots were rejected when the geminates were adjacent in the surface (e.g., maMMiKot). This finding is puzzling: If speakers tolerate the presence of identity in the root, then why do they specifically reject identity when the geminates are adjacent? A similar dissociation between different manifestations of identity avoidance was observed in the relative rating task. Recall that counterexamples with root-initial geminates (e.g., MiMeK) were as acceptable as their final gemination controls (e.g., MiKeK), but both forms of identity were rejected compared to the no-gemination controls (e.g., PiMeK). We have demonstrated that the rejection of root-final identity is not a frequency effect; hence, it must indicate the representation of geminates by a variable. However, if speakers do represent identity in these roots, then why are they insensitive to its location?

Our account explains this divergence by assuming that the evaluation of word candidates for the violation of structural constraints, including identity, is independent of the evaluation of their roots during the acquisition of lexical representations (see Table 15). On our account, strong counterexamples are spared from identity erasure due to the activation of the associative mechanism. For instance, given experimental words such as MiMeK, participants infer a triliteral root, MMK. Compared to other possible candidates (e.g., MiMe and MiMeKD), the analogy-word MiMeK is the optimal output of mmk. Importantly, despite their optimality, analogy-words do exhibit identity; hence, they violate the IDENTITY family. For nonadjacent geminates (e.g., MiMeK), words with root-initial gemination violate the lowest ranked constraint, IDENTITY, a rather weak violation that is equivalent to that incurred by root-final gemination (e.g., MiKeK). The weak violation IDENTITY explains their equal rejection relative to no-gemination control in the relative rating task. Conversely, truly adjacent geminates (e.g., maMMiKot) fatally violate the higher ranked ADJACENT IDENTITY constraint; hence, these roots are rejected relative to either their final-gemination or no-gemination controls.

The division of labor between the associative and default, symbolic processes in the acquisition of lexical representations resembles the proposals of Pinker (1991, 1999) and Marcus et al. (1995) with regard to inflectional morphology. Pinker (1991, 1999) and Marcus et al (1995) marshal numerous arguments in favor of a distinction between an associative and default processes of inflection. The associative process is governed by similarity and applies whenever the input under considerations activates similar stored irregular tokens. When a root token is absent or irretrievable, a default mechanism applies. In contrast to the associative process, the default symbolic process is blind to token properties and is sensitive only to the combinatorial structure of formal constituents. The proposal of default vs associative processes in the acquisition of lexical representations is a natural extension of the Pinker (1991) and Marcus et al. (1995) account of inflection. Indeed, the two processes are tightly linked: the abstraction of lexical representations from morphologically complex words is the mirror image of their inflection.

Our proposal of two processes of lexical acquisition raises numerous questions. One question concerns the generality of identity erasure in language processing. Our results suggest that the default representation of root identity is biconsonantal. If our proposal extends to word processing, then it would imply that all familiar Hebrew words containing root identity are lexically stored as biconsonantal roots and that identity is productively formed by reduplication. However, it is possible that the familiarity with a root results in the storage of a triliteral representation by the associative process, in addition to its default biconsonantal representation. The triliteral representation of such roots may thus be obtained by either a default reduplication or by retrieval of its stored triliteral form. The actual contribution of default reduplication and retrieval in the processing of root identity would depend on their relative speed, which, in turn, may depend on familiarity and the demands of the experimental task. Future research is required to investigate the lexical representation of familiar roots using time-limited paradigms.

A more general question concerns the relationship between default identity erasure and other proposals regarding the erasure of redundancy from the lexicon. These proposals include the elimination of affixation that is predictable by the default inflectional process and the elimination predictable phonological features by phonological underspecification (e.g., Archangeli, 1988). The relationship between these different forms of redundancy avoidance awaits further research. One view may attribute these distinct manifestations to a single principle of lexical organization, namely the avoidance of redundancy in the lexicon (e.g., Di Sciullo & Williams, 1987). Conversely, the various manifestations of redundancy avoidance may be due to the operation of distinct, autonomous principles.

Understanding the relationship between the different manifestations of redundancy erasure is crucial for carving specific models of linguistic competence. For our present purposes, however, it is the convergence of evidence for symbolic processes that is of principal interest. The goal of the present investigation was to contrast two rival accounts for the productivity of language: the symbolic hypothesis and the pattern associator view. Our investigation of identity avoidance in Hebrew agrees with the evidence from inflectional morphology in demonstrating that speakers posses an associative process that productively constrains novel forms by computing statistical information over stored tokens. Indeed, the role of statical knowledge in cognition is undeniable. The center of debate between the two hypotheses concerns the role of variables. The combinatorial structure of variables is fundamental for explaining linguistic productivity according to the symbolic account, but is obsolete, according to the pattern associator hypothesis. We have demonstrated that Hebrew speakers possess tacit knowledge that constrains two variables, root and identity. Our results suggest that variables play a causal role in linguistic behavior. The appeal to variables is thus necessary for the adequacy of theories of language.

REFERENCES

Alegre, M., & Gordon, P. (1996). Red rats eater exposes recursion in children's word formation. *Cognition,* **60,** 65–82.

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

Archangeli, D. (1988). Aspects of underspecification theory. *Phonology,* **5,** 183–207.

Aronoff, M. (1994). *Morphology by itself.* Cambridge, MA: MIT Press.

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,* **42A,** 693–711.

Berent, I., Bibi, U., & Tzelgov, J. (1999). *Root structure constraints in the performance of the Stroop task.* Manuscript in preparation.

Berent, I., Bibi, U., & Tzelgov, J. (2000). *The autonomous computation of linguistic structure in reading: Evidence from the Stroop task.* A paper presented at the 41st meeting of the Psychonomic Society. New Orleans, LA.

Berent, I. Everett D., & Shimron, J. (2000). *Identity vs. place of articulation avoidance in Hebrew.* Unpublished data.

Berent, I., Marcus, G., & Shimron, J. (2000). *Training spaces and generalization: Evidence from Hebrew word-formation.* Manuscript in preparation.

Berent, I., Pinker, S., & Shimron, J. (1999). Default nominal inflection in Hebrew: Evidence for mental variables. *Cognition,* **72,** 1–44.

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., & Shimron, J. (1997). The representation of Hebrew words: Evidence from the Obligatory Contour Principle. *Cognition,* **64,** 39–72.

Berent, I., Shimron, J., & Vaknin, V. (in press). *Phonological constraints on reading: Evidence* *from the Obligatory Contour Principle. Journal of Memory and Language,* doi:10.1006/jmla.2000.2760.

Chomsky, N. (1965). *Aspects of the theory of syntax.* Cambridge, MA: MIT Press.

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

Daugherty, K., & Seidenberg, M. (1992). Rules or connections? The past tense revisited. In *Proceedings of the 14th Annual Conference of Cognitive Science Society* (pp. 259–264). Hillsdale, 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.

Di Sciullo, A., & Williams, E. (1987). *On the definition of a word.* Cambridge, MA: MIT Press.

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

Elman, J. (1993). Learning and development in neural networks: The importance of starting small. *Cognition,* **48,** 71–99.

Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness. A connectionist perspective on development.* Cambridge, MA: MIT Press.

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.* Available: http://ruccs.rutgers.edu/roa.html (ROA-235).

Feldman, L. B. (1994). Beyond orthography and phonology: Differences between inflection and derivations. *Journal of Memory and Language,* **33,** 442–470.

Feldman, L. B., & Bentin, S. (1994). Morphological analysis of disrupted morphemes: Evidence from Hebrew. *Quarterly Journal of Experimental Psychology,* **47A,** 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 & Cognition,* **21,** 947–960.

Fodor, J., & McLaughlin, B. (1990). Connectionism and the problem of systematicity: Why Smolensky's solution doesn't work. *Cognition,* **35,** 183–204.

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

Fowler, C., Napps, S., & Feldman, L. (1985). Relations among regular and irregular morphologically related words in the lexicon as revealed by repetition priming. *Memory & Cognition,* **13,** 241–255.

Frost, R., Forster, K., & Deutsch, A. (1994). *Morphological and orthographic analysis in word recognition contrasted.* Paper presented at the 35th meeting of the Psychonomic Society, Saint Louis: MO.

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

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

Gordon, P. (1985). Level ordering in lexical development. *Cognition,* **21,** 73–93.

Hare, M., & Elman, J. (1995). Learning and morphological change. *Cognition,* **56,** 61–98.

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayered feed-forward networks are universal approximators. *Neural Networks,* **26,** 359–366.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review.* **104,** 427–466.

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

Kim, J., Marcus, G., Pinker, S., Hollander, M., & Coppola, M. (1994). Sensitivity of children's inflection to grammatical structure. *Journal of Child Language,* **21,** 179–209.

Kim, J., Pinker, S., Prince, A., & Prasada, S. (1991). Why no mere mortal has ever flown out to center field. *Cognitive Science,* **15,** 173–218.

Kirk, R. (1982). *Experimental design: procedures for the behavioral sciences.* Pacific Grove: Brooks/Cole.

Marcus, G. (1998a). Can connectionism save constructivism? *Cognition,* **66,** 153–182.

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

Marcus, G. (1999a). Response—Do infants learn grammar with algebra or statistics? *Science,* **284,** 433.

Marcus, G. (1999b). Response (technical comments): Rule learning by seven-month old infants and neural networks. *Science,* **284,** 977.

Marcus, G. (1999c). Connectionism: With or without rules? *Trends in Cognitive Sciences,* **3,** 168–170.

Marcus, G. (in press). *The algebraic mind: Reflections on 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.

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

Marslen-Wilson, Komisarjevsky Tyler, L., Waksler, R., & Older, L. (1994). Morphology and meaning in the English mental lexicon. *Psychological Review,* **101,** 3–33.

McCarthy, J. (1979/1985). *Formal problems in Semitic phonology and morphology*. New York: Garland Press.

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

McCarthy, J. J. (1989). Linear order in phonological representation. *Linguistic Inquiry,* **20,** 71–99.

McClelland, J., & Plaut, D. (1999). Does generalization in infant learning implicate abstract algebra-like rules? *Trends in Cognitive Sciences,* **3,** 166–168.

Myers, S. (1994). *OCP effects in Optimality theory.* Rutgers Optimality Archive (http://ruccs.rutgers.edu/roz.html)

Negishi, M. (1999). Rule learning by seven-month-old infants and neural networks. *Science,* **284,** 875.

Odden, D. (1986). On the Obligatory Contour Principle. *Language,* **62,** 352–383.

Pinker, S. (1991). Rules of language. *Science,* **253(2),** 530–535.

Pinker, S. (1994). *The language instinct.* New-York: Morrow.

Pinker, S. (1999). *Words and rules: The ingredients of language.* New York: Basic Books.

Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of parallel distributed processing model of language acquisition. *Cognition,* **28,** 73–193.

Plaut, D., 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., & Marchman, R. (1993). From rote learning to system building: acquiring verb morphology in children and connectionist nets. *Cognition,* 21–69.

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

Prince, A., & Smolensky, P. (1993). *Optimality theory.* Rutgers University manuscript.

Prince, A., & Smolensky, P. (1997). Optimality: From neural networks to Universal Grammar. *Science, 275(14),* 1604–1610.

Pylyshyn, Z. (1984). *Computation and cognition: Towards a foundation for cognitive science.* Cambridge, MA: MIT Press.

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

Rosen, H. (1977). *Contemporary Hebrew.* The Hague: Mouton.

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 the past tense of English verbs: Implicit rules or parallel distributed processing? In J. L. McClelland, D. E. Rumelhart & The PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition:* Vol. 2. Psychological and biological models (pp. 216–271). Cambridge, MA: MIT Press.

Seidenberg, M. (1987). Sublexical structures in visual word recognition: Access units of orthographic redundancy? In M. Coltheart, (Ed.), *Attention and performance XII: Reading* (pp. 245–263). Hillsdale, NJ: Erlbaum.

Seidenberg, M., & Elman, J. (1999). Do infants learn grammar with algebra or statistics? *Science,* **284,** 433.

Seidenberg, M. Haskell, T., & MacDonald, M. (1999, November). *Constraints on plurals in compounds: Some implications of compounds research.* Paper presented at the 40th annual meeting of the Psychonomic Society. Los Angeles, CA.

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

Seidenberg, M. (1997). Language acquisition and use: Learning and applying probabilistic constraints. *Science,* **275(14),** 1599–1603.

Siegelman, H., & Sontag, E. (1995). On the computational power of neural nets. *Journal of computer and system sciences,* **50,** 132–150.

Smolensky, P. (1996). *The initial state and 'Richness of the Base' in Optimality theory*. Rutgers Optimality Archive (http://ruccs.rutgers.edu/roz.html)

Smolensky, P. (1995). Constituent structure and explanation in an integrated connectionist/symbolic cognitive architecture. In C. MacDonald & G. MacDonald (Eds.), *Connectionism: Debates on psychological explanation* (Vol. Two, pp. 221–290. Oxford: Basil Blackwell.

Winer, B. J. (1971). *Statistical principles in experimental design.* New York: McGraw-Hill.

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

Yip, M. (1995) *Identity avoidance in phonology and morphology*. Rutgers Optimality Archive (http://ruccs.rutgers.edu/roz.html)

Accepted May 19, 2000; published online January 4, 2001