Connectedness: Papers by and for Sarah VanWagenen
UCLA Working Papers in Linguistics 18: 1–29
Carson T. Schütze & Linnaea Stockall (eds.), 2014

The Morphologically Organized Mental Lexicon
Further Experimental Evidence*

Sarah VanWagenen

In theoretical linguistics, the notion of a morpheme as a discrete stored linguistic unit is not uncontroversial. Nonetheless, there is general agreement among linguists that lexical items are related to one another via rules or the equivalents that amount to morphological relationships. Outside the field of linguistics, however, cognitive scientists have argued that morphological relationships are, in fact, simply the by-products of semantic and phonological overlap between items: morphology is epiphenomenal (Seidenberg 1993, Gonnerman 1999, Plaut and Gonnerman 2000). Though many predictions stemming from these two perspectives depend on the specifics of the models used to implement them, repetition priming has been used as a means of addressing the general divide. Using a version of this paradigm, I test the hypothesis that semantic and phonological overlap are sufficient to account for priming results using as stimuli morphologically related words in which semantic and phonological overlap are minimal. Previous work has confounded semantic overlap with shared morphology (see Marslen-Wilson, Ford, Older, and Zhou 1996, Feldman 2000 for discussion). The present study exploits the observation that delayed repetition (items intervening between prime and target), priming does not yield semantic effects (Henderson, Wallace and Knight 1984): *joy* does not prime *happiness* as it does when the target immediately follows the prime, and further, *hapless* does not inhibit *happiness* as in immediate repetition. Morphologically related words, however, continue to prime one another even with the delay in target presentation (Kouider 2000, Bentin and Feldman 1990). My work focuses on the representation of derivational suffixes, which are morphemes in the traditional sense, but provide little semantic or formal overlap with other words that share them. I find that derived words that share a suffix significantly prime one another (*happiness* primes *darkness*) whereas their semantic and form relatives do not. These results constitute further experimental evidence in favor of a morphologically organized mental lexicon.

1 Introduction: Morphology in Psycholinguistics

Linguists use the term *morpheme* to refer to the smallest arbitrary sound–meaning pairings in a language. For example, the fact that the series of sounds /bebi/ corresponds to the real world entity BABY in English is completely arbitrary, but anything smaller either represents a completely different meaning, e.g., /be/ for real world BAY or /bi/ for

* This is a lightly edited and reformatted version of Sarah's MA thesis, originally submitted in 2005, officially filed in 2009. No attempt has been made to update the references.

real world BEE, or nothing at all, e.g., /b/. Slightly more abstractly, the fact that the sounds /s/, /z/, and /əz/ can serve to pluralize whatever word they suffix to is also completely arbitrary, though the choice of plural allomorph from amongst this set is determined by phonological rule.

Though languages employ thousands of morphemes, expressive power lies not so much in their number, but in their combinatory properties: morphemes can combine into words (even in novel ways), and the meaning of the combination can, in general, be derived from the meaning of the parts. So, while the pairings of /bebi/ with BABY and of /z/ with PLURAL are unpredictable, the meaning of the combination of /bebi-z/ given BABY and PLURAL is not. However, combination is not a free-for-all: certain combinations are allowed whereas others are not; particular orders of morphemes are allowed whereas others are not. While it is possible (and linguists would maintain that it is indeed true) that there are universal principles underlying these combinatory rules, there are certainly language-specific, and even morpheme-specific, rules as well. As such, in addition to meaning and phonological form, certain structural information must be encoded as part of a morpheme's lexical entry, e.g., order (root, suffix, prefix, infix, circumfix, transfix), grammatical category, and selectional restrictions (what kinds of other morphemes it can attach to). Morphology, then, can be understood as the level at which phonological, semantic, and structural information come together, and morphemes the smallest units that embody such a combination.

An interesting question for psycholinguists, then, is whether people actually use morphological representations in language production and comprehension. It is certainly possible that morphemes are just convenient theoretical notions with no psychological instantiation. In considering the organization of the mental lexicon, i.e. how chunks of sound, meaning, and structure are represented in the mind, the immediate question arises as to whether morphemes constitute manipulable chunks or whether overlapping phonological, semantic and structural information is independently manipulated.

The present study attempts to establish a particular paradigm as a reliable means for assessing whether or not morphemes, in particular suffixes, are psychologically represented as constituents, i.e., as manipulable chunks.

2 Previous Experimental Work

Obviously, psychologists must make use of indirect measures of the workings of the human brain, and coming up with experimental paradigms in which to do so is no easy task. One paradigm that has been used extensively to probe the nature of lexical representation and processing is lexical *priming*.

In lexical priming experiments, participants see or hear a string of letters/sounds (the *prime*), followed by another string (the *target*). They are asked to identify the target as quickly and as accurately as possible as a real word of their language or as a non-word, usually by a button press, and their accuracy rate and reaction times are recorded. The basic idea behind priming is that lexical representations are accessed when a receiver hears or sees a string of sounds or letters that correspond to that representation. Residual activation from accessing that lexical representation will linger, such that subsequent exposure to the same representation will result in faster access.[1]

Consequently, repetition of the probe as target is facilitatory, i.e. results in speeded reaction time to the target: *happy* primes *happy* (*happy* → *happy*). Interestingly, morpho-

[1] Whether this works in terms of a lowered threshold for activation Morton (1979) or a higher baseline (McClelland and Rumelhart 1981) is still a live research topic.

logical derivatives of a stem[2] also prime words consisting of just that stem: *happiness* → *happy* (Marslen-Wilson et al., 1994).[3] Can this be considered evidence for morphologically based lexical representations that are decomposed into their constituent parts as opposed to stems represented separately with each affix and affixes only being represented as part of a stem? Put more strongly, is the *happy* in *happiness* the same representation as that of *happy* in *happy*?

The most obvious objection to such a conclusion is that *happiness* and *happy* are both highly semantically related and highly phonologically and orthographically related. As such, we cannot rule out the possibility that *happy* and *happiness* do not share a morpheme, but prime each other merely as a function of these kinds of relatedness. Perhaps, as has been argued, what look like morphological relationships are simply an overlap of semantics and phonology.[4] Indeed, probes that are semantically but not morphologically related also prime their targets: *joy* → *happy*. Form-related probes can also prime their targets under very special circumstances,[5] (*hapless* → *happiness*), however, in most circumstances they increase reaction times to targets, presumably because of competition effects in word recognition (Luce, Pisoni, and Goldinger, 1990).[6]

2.1 Addressing Semantics and Form

In an important study, Marslen-Wilson et al. (1994), hereafter MW, addressed this objection by including semantic and form control probes with morphological probes. They used a cross-modal variation of the priming paradigm in which primes were auditorily presented, followed immediately by visually presented targets. This choice of inter-modal presentation was motivated by their interest in accessing abstract representations as opposed to modality-specific ones. If presentation were intra-modal, they reasoned, any possible priming effects could be mediated by episodic recall of the perceptual code as opposed to features of the mental lexicon, which is, by assumption, independent of modality.[7]

Significant priming effects of morphology were observed, even when the surface form of the stem had been phonologically altered by affixation, e.g. not only did *friendly* prime *friend*, (priming will subsequently be represented by a right arrow: →), but *elusive* → *elude*, and *serenity* → *serene*. Semantic priming was also significant, even when there was no morphological relationship, e.g., *idea* → *notion* almost as much as *friendly* → *friend*.[8] Priming based on phonological relatedness (form priming) was not observed:

[2] To my knowledge, no work has been done in English on complex words with more than two potential morphemes, so here the stem is equivalent to the root.

[3] Priming effects are also observed when order of presentation is inverted: *happy* → *happiness* (Marslen-Wilson et al., 1994).

[4] I use "phonology" loosely here to refer not only to the phonological properties of a morpheme, but also the orthographic properties that represent the phonological ones—in short, for the auditory and visual forms that a morpheme takes. I will refer to these properties collectively as the 'form' of the morpheme for the rest of the paper.

[5] Form-related probes speed reaction times to targets when the probe is presented for a short amount of time (under 66 ms) in which the participant does not consciously apprehend the stimulus.

[6] Magnetoencephalography (MEG) studies have shown two distinct sources of behavioral inhibition: inhibited activation for initially matched phonological relatives, (e.g., *spinach–spin*) and post-activation suppression of non-initially matched relatives (e.g., *teacher–reach*) (Pylkkänen, Stringfellow and Marantz 2002).

[7] See Bowers and Kouider (2003) and Bowers (2000) for arguments against intra-modal priming effects being attributed to episodic memory traces, however.

[8] Unfortunately, it is not reported whether there is a significant difference between the morphological and the purely semantic priming.

there was a slight inhibitory effect, but it was not significant (see Table 1 for a summary of the results).

In order to assess the degree to which semantic relatedness was a necessary component of morphological priming, MW ran additional experiments with pairs of words whose relationships were either semantically transparent, i.e., the meaning of the derived word[9] was clearly made up of the meaning of the stem plus the meaning of the affix, e.g. *friend–friendly*, or semantically opaque, i.e., the meaning of the derived word could not be transparently derived from the meaning of the stem, e.g., *casual–casualty*.[10] They found significant facilitatory effects in the transparent cases, but only an insignificant trend in the opaque cases. If one adopts the perspective that semantically opaque derived words are bi-morphemic, this might be taken as evidence that morphological priming is essentially semantic priming, and at least suggests that semantic transparency is a necessary condition for morphological priming. Another possibility, in fact, the one argued for by MW and standardly assumed, is that the opaque words are mono-morphemically represented, hence the lack of priming. Still, the robust priming in both the morphological and semantic conditions makes dissociating their effects difficult.

Table 1: Summary of MW 1994 priming results

Example	Condition	Morphological type	Priming?
friendly–friend	+morph, +sem	derived–stem	yes
elusive–elude	+morph, +sem with C change	derived–stem	yes
serenity–serene	+morph, +sem with V change	derived–stem	yes
confession–confessor	+morph, +sem	derived–derived	no
casualty–casual	+morph, –sem	derived–stem	non-significant facilitatory trend
idea–notion	+sem	n/a	yes
tinsel–tin	+phon	n/a	non-significant inhibitory trend

Interestingly, derived–derived pairs (as opposed to the previously described stem–derived pairs), i.e., pairs of words that share a stem but have different affixes (e.g., *confession–confessor*, which share *confess*), were also included in the stimulus set of this experiment. High semantic transparency[11] was not facilitatory in the derived–derived cases, indeed no priming effects were observed in this condition. The conclusion drawn was that semantic relatedness was indeed a necessary (cf. semantically opaque pairs) but not a sufficient condition for priming.

[9] Here, "derived" is used in a traditional sense to refer to the result of adding a derivational affix to a stem.

[10] Of course, it is not obvious that *casualty* is bi-morphemic at all, or, if it is, that its stem (presumably *casual*) is the same stem as that of the adjective. See section 3.1.1 for more discussion.

[11] Semantic transparency was measured by a semantic relatedness survey in which participants were asked to rate the degree to which the meaning of, say, *happy* was related to the meaning of *happiness*, or in the derived–derived cases, that of *confessor* to that of *confession*.

Calling the priming effects of Table 1 morphological, then, is based on the lack of priming in the semantically related derived–derived pairs in the context of the particular model of the lexicon that MW develop. They propose that *confession* and *confessor* contain *confess*, with the affixes *-ion* and *-or* linked to the stem much like satellites of a central body. While the *confess* element of both words primes itself, there is competition between affixes in the word recognition process. Inhibitory links exist between them, such that when one is selected, the other is necessarily suppressed.[12] This inhibition counters the facilitatory effect of stem priming, accounting for the lack of priming between derived–derived pairs within a morphologically organized lexicon in which complex words are decomposed. How to explain this effect based on semantic and form overlap is less obvious, since clearly there is a significant degree of both in the transparent derived–derived pairs.

Summing up, this study presents results that show semantic and morphological priming in the cases of stem–derived, derived–stem, and prefixed derived–derived pairs. No priming is found for suffixed derived–derived pairs, morphologically but not semantically related pairs, or phonologically overlapping pairs with neither a morphological nor a semantic relationship.

2.2 *Varying the paradigm: effects of semantics and form*

A study by Feldman (2000) also examined the effects of semantic, orthographic, and morphological similarity on priming, but with an eye toward its gradient, rather than categorical, nature. In addition to considering the implications of priming for representations, she also considers the perspective that the time course of word processing plays a crucial role in priming results. The length of time the prime was displayed was systematically varied and reaction times were analyzed with prime type (morphological, semantic, or orthographic) and prime duration as factors.

In an immediate repetition visual–visual lexical decision task, Feldman found that as the length of the stimulus onset asynchrony (SOA, i.e., prime duration + interval between prime and target) increased, the degree of orthographic inhibition also increased. Her results are summarized in Table 2.

Further, above a certain SOA (66 ms), evidence for sensitivity to morphological relatedness emerged: words that were semantically and morphologically related showed greater priming effects than those that were semantically but not morphologically related. Only at longer SOAs (300 ms) did semantic effects and morphological effects really diverge. Overall, the combined priming effects of form and semantics decreased with longer SOAs, while the priming effects of morphological relatedness increased, as can be seen by comparing the two boldface columns.

[12] However, semantically transparent pairs that were derived by virtue of prefixation (as opposed to suffixation, e.g., *refasten–unfasten*) did show robust priming effects. MW argue that, because the prefix is the first string perceived, its competitors for stem attachment are not activated, so no inhibitory links are necessary.

Table 2: Summary of Feldman 2000 priming results:
Facilitation (ms) as a function of relatedness and SOA

SOA (ms) \ Prime Type / Ex.	Morphological *vowed–vow*	Semantic *pledge–vow*	Orthographic *vowel–vow*	Sem + Orth Facilitation
66	29	**22**	18	**40**
116	35	**30**	–8	**22**
300	49	**33**	−18	**15**

It is difficult to assess these results in the context of trying to ascertain the nature of lexical representations. Clearly, there are processing steps sensitive to different aspects of these representations that affect reaction times in the lexical decision task.[13] However, morphological effects cannot be explained by similarity along semantic or orthographic dimensions independently at any SOA. More intriguing is the possibility of morphological effects being explained by some relation between semantics and form, though Feldman's results indicate that it is not a linear one.

2.3 Reviving the objection: the interaction of semantics and form

Arguing that explicit morphological representations do not exist, Gonnerman (1999) considers the interaction of semantics and form in an effort to account for priming results that appear to be morphologically based. It is a more subtle characterization of an interaction than the previously tested hypothesis of a linear relation between these two dimensions. Gonnerman asserts that semantic similarity is necessary for facilitation, but not necessarily sufficient. When two words are highly semantically related, then their degree of phonological overlap will strongly affect the degree of priming obtained. In particular, she offers an alternative explanation for the MW 1994 results in which they report a morphological priming effect for semantically transparent stem–derived pairs of words (*friend* → *friendly*), but no priming for derived–derived pairs (*confessor–confession*), whether or not they were semantically transparent. In a connectionist simulation of priming using MW's stimuli where only phonological and semantic representations are encoded, i.e., there is no explicit morphological representation, she reports the same pattern of results as MW's experimental data. She claims that the interaction of semantic and phonological similarity is sufficient to account for the data: there must be a high degree of similarity on both dimensions in order for priming to obtain. In the MW derived–derived conditions, the degree of phonological similarity was too low for priming to obtain, despite the high semantic relatedness of the pairs.

It should be noted, however, that the measure by which phonological similarity was assessed was somewhat ad hoc when it came to the similarity of stems to derived words. Measuring similarity in words of unequal length resulted in inappropriately high dissimilarity scores when the first word was a stem and the second a derivative of the stem, (e.g., *friend–friendly*). So, if two words were phonologically identical when one word ended, then they were considered identical (given a score of 1), the comparison stopped and an arbitrarily determined 0.6 was added to the score. However, in the case of two derivatives of the same stem, the comparison continued until the end of the longer word, resulting in

[13] Efforts to tease these steps apart have been made in recent brain imaging research. Studies using MEG have identified response components that correlate with the initial low-level perception of stimuli, lexical activation, and subsequent phonological / orthographic competition (Pylkkänen, Stringfellow and Marantz 2002).

potentially disproportional scores of dissimilarity, which would bias the results. More significantly, this model leaves the extremely robust effects of semantic priming with no morphological or phonological relationship (e.g., *idea–notion*) completely unexplained.

Behavioral experiments were also performed to test the predictions of the computational model, namely that there would be gradient priming effects as a result of the degree of relatedness on the semantic and phonological dimensions, that derived–derived pairs would prime each other if they are sufficiently semantically and phonologically related, and that semantically and phonologically highly related pairs that do not share a morpheme would prime each other to the same extent that morphologically related words do. Gonnerman's strong hypothesis is that phonological and semantic overlap will completely account for what are only putatively morphological effects.

These predictions are all borne out in cross-modal priming lexical decision tasks. Interestingly, Gonnerman did get priming in the derived–derived condition when phonological similarity is high. The reason MW did not get priming in the same condition, she argues, was simply because of the unfortunate coincidence that the semantically transparent derived–derived words were too phonologically dissimilar. Further, Gonnerman gets a significant amount of priming in cases of high semantic and phonological similarity where a morphological relationship is unquestionably lacking, e.g., *trifle*, *trivial*.

These results are taken as a lack of strong evidence for independent morphological representations. However, although morphologically based theories do not explicitly predict that the degree of semantic and phonological overlap will affect the degree of priming, this could be straightforwardly incorporated into models relying on morphological structure by referencing the phonological and semantic levels of representation, which are indisputable and uncontroversial. Proponents of morphemes as the units around which the mental lexicon is organized are not arguing that semantic and phonological representations do not exist. However, current symbolic models do not consider the source of these gradient priming effects or address the issue of how degrees of semantic and phonological similarity could affect morphological representations. A processing model that incorporates the time course of morphological processing could certainly be brought to bear on these issues.

2.4 Representing affixes

Continuing to amass experimental evidence in favor of an abstract level of morphological representation where words are decomposed into their morphemic parts, Marslen-Wilson, Ford, Older, and Zhou (1996) extended the previous work of MW 1994 to test the predictions of the stem–affix model developed in that study. This model represents both stems and affixes as independent elements, and as such, predicts that affixes will also prime themselves. Using the cross-modal priming paradigm again, this is exactly what they found: a priming effect for shared morphological representations (*happiness* → *darkness*), and no effect for formally (orthographically and phonologically) related pairs (*happiness* → *harness*). More finely, they found a significant effect for productive affixes with only a trend in the right direction for unproductive affixes.[14] Notice that the semantic relationship between happiness and darkness as whole words is negligible, which undermines Gonnerman's semantics–phonology interactive view of morphology. Recall that on this view, semantic relatedness plays a crucial role in priming results and further, that a high degree of both semantic and phonological relatedness was

[14] The measure of productivity was not reported for this paper.

necessary in order to obtain a facilitatory effect. Here, neither the semantic nor the phonological prerequisites are met, though a careful examination of the stimuli would be in order.

The results are interpreted by MW as "strong evidence not only for the combinatorial nature of the mental computations underlying lexical representation and processing, but also for the truly morphological—and not simply semantic—nature of the effects we are dealing with." A more recent paper by Marslen-Wilson (1999) elucidates the reasoning behind this interpretation with regard to the possibility of semantic priming: "Affixes like *-ness* and *de-* do not have clearly definable semantic identities. They are fundamentally morphological entities, functioning in productive linguistic processes of word formation, and it is hard to see how priming between them can be accounted for in anything other than morphological terms."

As appealing as it would be to rule out semantics as the driving factor in morphological priming as suggested in the interpretation of these results, from a linguistic perspective, the argument is unsupportable. While the semantics of affixes may not be intuitively obvious, this by no means precludes them from being 'clearly definable'. Affixes do have meaning and (generally) make predictable semantic contributions to complex words. For example, *-ment* means something like 'the act or process of' and contributes that meaning to the nominalization of the verb it attaches to, e.g., *enjoyment*.[15] Given this, it cannot be determined a priori that affix priming is not the result of semantic priming, in spite of the disparity between the semantics of the whole words.

An additional criticism of this study is also leveled by Giraudo and Grainger (2003); namely, that the comparisons are between items. In other words, the same targets were not used in each condition, rather it was the target that was manipulated to create new conditions, e.g., in the affixed condition the target was *darkness,* in the pseudo-affixed it was *harness*. Also, in the phonological overlap condition, completely different stimuli were used, e.g. *puritan–charlatan*. Given this, the relevant comparison of reaction times to the same target in different conditions cannot be made. Because different items were used in different conditions, the possibility of priming effects being due to particular properties of particular items cannot be ruled out: individual items could be driving the effect, so the validity of generalizing to items of a particular type (e.g., all productive affixes) is suspect (cf. Clark 1973).

2.5 Delayed repetition priming: A more restrictive measure?

Another variation on the repetition priming paradigm involves delaying the presentation of the target relative to the prime by presenting some number of intervening items (Morton 1969, Bowers 1996, Tenpenny 1995). In delayed repetition, participants are not aware of the difference between primes and targets because they make a lexical decision on every item. Kouider and Dupoux (2009) find that semantic priming in non-morphologically related pairs (*idea–notion*) is absent after about one second, as is the inhibitory effect of form. Morphological priming, however, remains significant with up to 144 intervening items. Indeed, morphological priming effects are the same size with 18, 72, and 144 intervening items (Kouider 2002).

Feldman's (2000) work also contains a delayed repetition experiment in which the morphological effect is significant (ps < .005) at a delay of 10 items while the semantic and orthographic effects are not. She finds a statistically significant difference between the morphological effect and the semantic and orthographic effects individually, but not

[15] Of course, there are exceptions, e.g., *government, department*, which will be discussed below.

between the morphological effect and the sum of the semantic and orthographic effects. See table 3.

Table 3: Summary of Feldman 2000 delayed repetition results: Facilitation (ms)

Prime Type	Morphological	Semantic	Orthographic	Sem + Orth facilitation
Example	*vowed–vow*	*pledge–vow*	*vowel–vow*	
Facilitation	28	1	9	10

Because our aim is to determine whether or not the mental lexicon is organized in terms of abstract morphological representations as opposed to overlapping semantic and phonological features exclusively, the fact that these latter dimensions are not singly detectable in this paradigm encourages considering it a potentially more restrictive means of evaluating the existence of just such an abstract relation between words. It would be of interest, however, to vary the dimensions of semantic and phonological similarity within morphologically related words to see if there are also gradient effects of relatedness within the delayed repetition paradigm.

3 Experiment

3.1 Rationale

The fact that statistically significant effects of affix priming are obtained is highly suggestive of their independent status in the lexicon. However, the confounds in the Marslen-Wilson et al. (1996) study previously discussed undermine their results. The experiment described below seeks to eliminate the confounding semantic factor by using the delayed repetition priming paradigm, in which semantic effects have been shown not to exist, and also seeks to eliminate the possibility of individual items driving the obtained effect, by appropriate counter-balancing, i.e., with the same target being tested in each condition.

There are several reasons why this study focuses on derivational suffixes as a means of investigating the nature of representations in the mental lexicon. First, as mentioned above, semantic overlap between two words that share an affix is relatively minimal. If priming results are obtained, then the semantic–phonological interaction account where semantic overlap is a necessary prerequisite for priming, is called into question. However, there is a semantic representation of affixes that cannot be ignored. For this reason, I have chosen the delayed-repetition priming paradigm as the paradigm to investigate suffix representations.

While there is little consensus on the validity of distinguishing between derivational and inflectional morphology, a standard view is that inflectional morphology is part of the syntactic component of the grammar. As such, it is possible that inflectional affixes could prime each other on a structural level (cf. studies on syntactic priming, most notably represented by the work of Bock (1986; Bock and Griffin 2000). Implicit in the present study is the assumption that if derivational morphemes have any representational status at all, they are lexical (closed class) items and do not have independent status in the syntax.

Under the hypothesis that morphemes are constituents and that speaker/hearers decompose words when possible into their morphemic constituents, I predict that priming will be obtained for a target preceded by a morphological prime with which it shares a suffix, but not by a semantic or a formal prime.

3.2 Brief Description

In this experiment, participants made lexical decisions on visually presented letter strings in the delayed repetition priming paradigm. Primes were morphologically, formally (phonologically and orthographically), or semantically related to targets. A control condition presented the target with no preceding prime.

3.3 Materials

Because the effects of morphological priming are so long lasting, the same suffix could not be used in any two targets, lest one target inadvertently serve as a prime for the other in one of the three non-morphological conditions. A list of forty targets with forty different suffixes was therefore composed, largely based on the availability of suitable primes. Targets were preceded by one of three types of prime discussed below: morphological (*humanism* → *heroism*), formal, i.e., phonological/orthographic, (*heresy* → *heroism*), or semantic (*valor* → *heroism*), or by no prime at all, which served as a baseline condition. These types are discussed below. The full set of critical items can be found in the Appendix.

3.3.1 Morphological primes

A morphological prime shared a suffix with its target (e.g. *humanism* → *heroism*). Because of the nature of the paradigm and the design, however, a particular suffix could only be used in a single target. If a suffix had been used in multiple targets, counterbalancing the lists as necessary would not have been possible (see Design discussion below), so each morphological prime/target pair contained in a different suffix. Morphological primes were also matched to targets for number of syllables, placement, and phonological patterns resulting from affixation; some examples are given in Table 4.[16] Matching for frequency at the item level was not always possible given the severity of other restrictions (but see section 3.3.2).

Table 4: Examples of phonological changes as a result of derivation

Stem	Prime	Phonological Stem Change with Affixation
poet *metal*	*poetic* *metallic*	stress shifted one syllable to the right
confess *obsess*	*confession* *obsession*	/s/ → /ʃ/
private *frequent*	*privacy* *frequency*	/t/ → /s/

Semantic transparency of derived words to their stems was also taken into consideration, the standard view being that semantically opaque forms must be stored as whole words with no decomposed representations, by virtue of the fact that their meanings are not recoverable from their parts. Take, for example, the word *virtual*. Its phonological substrings suggest that it could be formed from the stem *virtue* and the suffix *-al*, but the meaning of the whole word is completely unrelated to the word *virtue*. Based on this as-

[16] Often, stress shifts result in a change in vowel quality, and sometimes also in use of flapping. This was not explicitly controlled for.

sumption, I had four linguists independently evaluate the derived stimuli (targets and their morphologically related primes) in order to avoid this possible confound. All four judged the morphological primes and the targets to be semantically related to their bases, i.e., they were transparent.[17]

The *virtual* example, though, highlights a critical aspect of semantic transparency not explicitly discussed in the literature: the semantic contribution of the affix. Interestingly, the *-al* suffix in *virtual* makes the same contribution to the whole word as the *-al* in the more transparent *functional*; namely, producing an adjective with the (very general) meaning "of, relating to, or characterized by" the noun (usually) that it attaches to (*Oxford English Dictionary*). The difficulty is that *virtual* does not mean 'relating to virtue' (even though the two words are historically related), but rather "characterized by an effect or essence but not by fact or name" (ibid.). It is a separate issue whether or not semantically opaque forms whose affixes make a predictable semantic contribution can be represented as constituent parts (see section 5). For present purposes, it is crucial that the semantics of the affix be consistent in prime and target. Despite the hypothesis that morphological priming is not semantic priming, accidental homophony is not predicted to be facilitatory either. Again, a shared representation at the morphological level entails a shared representation at the semantic (and phonological) level.

An extreme case of inconsistency is that of *-er*: the comparative *-er* of *bigger* is not the same as the agentive *-er* in *baker*. This is an obvious instance of homophony. Yet affixes can be more subtly distinct, polysemous rather than homophonous. Consider the verbal suffix *-ize*, which in *idolize* means something like 'to treat like an X' and in *philosophize* means something like 'to actively engage in X'. Though the effects of polysemy on lexical representation are only beginning to be investigated, experimental evidence suggests that there is room for multiple related senses in a single representation (Pylkkänen and Murphy 2004). Still, the relationship among senses can be more or less clear, so to avoid homophony and more opaque cases of polysemy, the *Oxford English Dictionary* meanings of each affix were enumerated and carefully considered by two linguists in the context of derived words as possible stimuli.[18] Only those pairs of words whose affixes were judged to be making the same semantic contribution were used.[19]

3.3.2 Formal controls

Form controls for stem priming experiments generally overlap by some number of characters with their targets, but do not share any morphological properties. In this experiment, the amount of character/phonological segment overlap between morphological primes and their targets is maximally four characters and three phonemes in the affix, (e.g. *-hood*, *-ship*, *-ness* and /hʊd, ʃɪp, nɛs/), and minimally one character/phoneme (e.g. *-y* and /i/), plus any coincidental stem overlap. As an initial rough measure for finding formal primes that shared at least as many segments with the targets as the morphological primes, a list of candidates with a maximum Levenshtein distance of four from their targets was compiled.

[17] A more objective measure of semantic relatedness to be discussed below was used post-hoc to evaluate morphological primes and targets to their bases. The mean LSA scores for primes and targets were the same. See section 3.1.3 for further definition and discussion.

[18] Thanks to Liina Pylkkänen for emphasizing the importance of semantic consistency.

[19] Multiple senses of whole words were not explicitly controlled for. For example, the *-ment* in one sense of *management* makes the standard nominalizing semantic contribution to the verb *manage*, as in 'Our management of resources has improved with the hiring of Mr. X.' Another sense of *management*, though, seems to take a less frequent sense of *-ment*, i.e. that of referring to a body that performs the action described by the verb, as in 'The management has really treated its employees unfairly.' Although none of the morphological primes or their targets seem ambiguous post-hoc, this was not rigorously considered.

Levenshtein distance (Kruskal and Liberman 1983) is a measure of similarity in which two strings are compared and one unit of distance is assigned for each single string element deleted, inserted, or substituted for another element.[20] For example, the distance between *dog* and *dig* would be 1 for the substitution of *i* for *o*. The distance between *dig* and *hags* would be 3: one for substituting *i* for *a*, one for substituting *h* for *d*, and one for inserting *s*.

A candidate set of form control words was auto-generated in which distances were based on orthographic similarity. From the orthographically similar words, phonological similarity in terms of Levenshtein distance was evaluated by comparing phonemic transcriptions by hand. For example, orthographically, there is only a distance of 3 between the target *loveless* and the prime candidate *novelist*: *l* → *n*, *e* → *i*, and *s* → *t*. However, phonemically, the distance is 5 (/lʌvlɛs/, /nɔvəlɪst/), since in addition to the above-stated orthographic changes, which are also phonological changes, the following phonemic changes also take place in spite of the preserved orthography: /ʌ/ → /ɔ/ and Ø → /ə/. So, the formal prime *level* was selected, which is a distance of 4 from *loveless* both orthographically and phonologically. The additional criteria were that primes be as closely matched to targets in frequency and number of syllables as possible.

Phonological and orthographic overlap between morphological primes and targets and formal primes and targets was then calculated. Formal controls overall shared more segments both orthographically and phonologically with their targets than did morphological primes: morphological primes and targets overlapped by an average of 3.75 characters and 3.2 phonemes while formal primes and targets overlapped by an average of 5.23 characters and 4 phonemes. Note, though, that for the formal primes and targets, the overlap was not necessarily contiguous, e.g. *cruel* was the formal prime for *crudely*, where the first three characters/sounds overlap, and the words also share the /l/ character/sound.

3.3.3 Semantic controls

Because the semantic influence on affix priming is at issue, primes that were semantically but not morphologically related to the targets were also included. The standard way of assessing the semantic relatedness of two words is by administering surveys in which participants are asked to rate the degree of relatedness according to some arbitrary scale. However, a statistical method for assessing context–usage meaning correlations between two words (or passages of text) called Latent Semantic Analysis (LSA) (Landauer, Foltz, and Laham 1998) has recently been shown to correlate very highly with human judgments (Rastle et al. 2004). For efficiency, LSA was used to assess semantic relatedness.

Potential primes were selected from association and synonym lists from WordNet® (Fellbaum 1998). LSA scores were then calculated using a pairwise comparison that assessed the relatedness of the semantic prime candidate and the target, with 1 being a maximal positive relation (i.e., synonymy) and −1 being a maximal negative relation, (i.e. antonymy). A score of 0, then, would represent no semantic relation.

The comparisons were based on the average vocabulary of a first-year college student (one of many options in the LSA online database). The candidate with the highest semantic relatedness score relative to the target that was as closely matched to the target for number of syllables and frequency as possible was selected. The mean LSA score was 0.19.

The degree of semantic relatedness of morphological primes and their targets was also calculated. If semantic factors, contrary to expectation, did turn out to play some role in

[20] I do not know of any previous research that has used Levenshtein distance in this way. Thanks to Colin Wilson for suggesting this measure as a means of evaluating candidates for form controls.

delayed repetition priming, ruling out the possibility of it playing a role in the morphological priming condition would be important. The mean LSA score for relatedness of morphological primes and their targets was 0.08, substantially lower than the semantic–target pairs.[21]

Recent experiments have obtained results in which the distinction between semantic and associative priming is relevant: semantic primes (e.g., *doctor–physician*) are relatively short lived, but associative primes (e.g., *doctor–nurse*) have longer lasting effects (Perea and Gotor 1997, Perea and Rosa 2002). The semantic controls used in this experiment are a combination of the two. It is possible that if more primes were associative, a semantic relatedness effect could emerge. However, if suffix priming turns out to be semantically conditioned, it would not be along associative lines (assuming there is no general associative relation between words that share an affix, since the meaning contribution of the affix to the whole is the same in the morphological condition as it is for the targets), and so the effect would not last across intervening items.

3.3.4 Nonwords

Seventy pronounceable nonwords were created and used in all lists. Candidate words were obtained from the ARC Nonword database (Rastle, Harrington, and Coltheart 2002) and manipulated so as to match the average character and syllable length of the real words. Additionally, real derivational suffixes (different from the forty used in the real word trials) were attached to 40 of the nonwords to match the number of real words with suffixes, so that ‘suffixedness’ could not be used by participants as a heuristic for deciding wordhood.

3.4 Design

Four lists were generated in which each of the 40 targets appeared once. Ten targets appeared in each of the four conditions in each list, so that across all groups, each target appeared once in each condition. Table 5 illustrates this design.

Table 5: Counter-balanced design

Group 1	Group 2	Group 3	Group 4
Morph prime List 1 targets	Form prime List 1 targets	Semantic prime List 1 targets	No prime List 1 targets
Form prime List 2 targets	Semantic prime List 2 targets	No prime List 2 targets	Morph prime List 2 targets
Semantic prime List 3 targets	No prime List 3 targets	Morph prime List 3 targets	Form prime List 3 targets
No prime List 4 targets	Morph prime List 4 targets	Form prime List 4 targets	Semantic prime List 4 targets

As previously mentioned, because priming effects in the delayed repetition paradigm are surprisingly long lasting, no preceding word (prime, target, or nonword) could share an affix with any particular target except its morphological prime. For example, if the

[21] It should be noted, however, that human judgments for the semantic relatedness of pairs of words that share only an affix have not been performed and caution should probably be exercised in assuming that the Rastle et al. results can generalize to pairs related only on that dimension.

morphological prime/target pair in List 1 were *humanism–heroism*, there could not be another *-ism* word in List 1 preceding the target. If there were, any observed priming effects could not be exclusively attributed to the explicit primes. Therefore, each target had a unique suffix.

A total of 40 targets and 30 primes were included in each list. The difference between the number of targets and number of primes arises because ten targets in each list were not preceded by primes (the No prime condition), so that they could serve as a baseline for comparison for the other conditions. These constituted the 70 real words that were intermixed with the 70 nonwords. Within the set of nonwords, ten real suffixes appeared twice, to match the ten real word morphological prime/target pairs per list. The same number of singly occurring suffixes (30) was also maintained for the real words and nonwords.

In order to ensure that a sufficient distance was maintained between prime and target, and further, that primes preceded targets, stimuli could not be presented in a truly random order. Therefore, each list was divided into four blocks with an average of 35 items per block. These items were then randomly sorted with adjustments made by hand where necessary, yielding a pseudo-random order. This order was then fixed, i.e. was the same for each participant who saw that list. Primes and targets always appeared within the same block, making the maximum distance between them 33 items. The average distance between primes and targets was 20 items with 5 being the minimum. To control for ordering effects, the four blocks were presented in a random order determined by the software with each experimental run.

3.5 Participants

Sixty-six[22] participants were recruited from the UCLA community. Participants were either paid or given credit in an introductory linguistics course for participating.

3.6 Procedure

Participants were tested individually in a quiet booth. They were seated in front of a Macintosh computer with a 14-inch monitor running PsyScope software (Cohen et al. 1993). A button box was placed on the table in front of the monitor which participants used to respond to the stimuli. They were given the following instructions: “You will see a string of letters across the computer screen. You are to decide whether or not the string is a real word of English. Press the GREEN button if the string IS a real word of English. Press the RED button if the string is NOT a real word of English. Please respond as quickly and as accurately as possible.”

Every trial proceeded as follows: a central fixation cross appeared for 500 ms to begin the trial. This was followed by a letter string presented for 500 ms in all lower case letters, which was followed by a blank screen. Participants were given 3000 ms from the start of the stimulus presentation to respond to the string by pushing the right green button for ‘word’ or the left red button for ‘nonword.’ A new trial began immediately after participants responded or after the time allotted had expired (though no participants exceeded their allotted time). Participants were not given accuracy feedback.

Participants were given 30 practice trials and a subsequent opportunity to ask clarifying questions before the presentation of the 140 test stimuli. The experiment lasted approximately 10 minutes.

[22] The pattern of results was evident at 32 participants, but 66 were run for additional statistical power.

3.7 Results

For these analyses, inaccurate responses were excluded (15%), as were responses to targets whose primes had been given an inaccurate response (5%). One item was excluded from all conditions due to experimenter error. Reaction times greater than 2.5 SD from a particular participant's mean across all accurate real word responses (approximately 3%) were replaced by that cutoff value. The mean decision latencies are summarized in table 6.

Table 6: Decision latencies (ms)

Condition	Example	Mean	StdErr	Priming
No prime	*heroism*	551.6	9.1	(baseline)
Morphological	*humanism → heroism*	531.6	8.6	+20
Formal	*heresy → heroism*	546.5	9.0	+5.1
Semantic	*valor → heroism*	555.5	10.0	–3.9

The reaction times for each participant in each condition were averaged and these averaged RTs were analyzed by ANOVA according to the method described in Raaijmakers et al. (1999) for counter-balanced designs.[23] According to this method, the main effect of Condition is tested against the Condition × List interaction except if the interaction is non-significant at a conservative α-level (.25). If this is the case, then the main effect of Condition is tested against the pooled error (the Condition × List error + Residuals). A main effect of List is not relevant for assessing the effect of Condition with a counter-balanced design: "In this design, the between-groups variability is confounded with (part of) the interaction between list and treatment [i.e., Condition]. However, ... the mean difference between the treatment conditions (and hence the treatment effect) is not affected by any difference that might exist between the lists."

In the present experiment, there was a main effect of List, $F(3, 186) = 7.77$, $p < .001$, indicating the variability among items that one would expect. However, the Condition × List interaction was non-significant $F(3, 3) = 1.32$, $p = .59$, therefore the pooled error was used. This analysis shows no effect of Group, $F(3, 62) = 1.94$, no Condition × List interaction at $\alpha = .25$, $F(3, 186) = 1.32$, and a significant main effect of Condition, $F(3, 189) = 5.65$, $p < .001$. (See section 3.4 for a review of Group, List, and Condition factors.) Results are shown in Figure 1.

[23] Thanks to Colin Wilson for making me aware of this paper and for help figuring out the right kind of analysis to do.

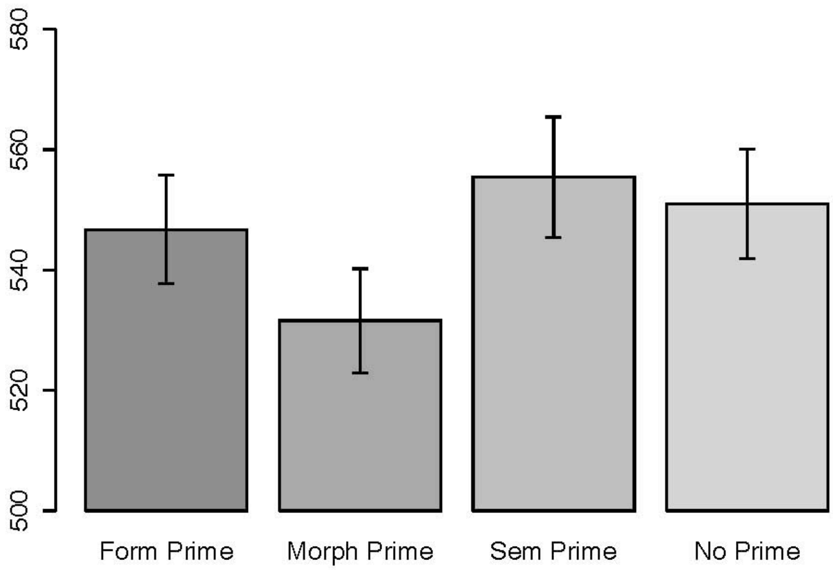

Figure 1: Mean decision latencies (ms)

Paired *t*-tests compared means of the No Prime condition (i.e., the baseline condition) and each other condition. There was no significant difference between the baseline and the semantic condition, $t(65) = -0.51$, and no significant difference between the baseline and the form condition, $t(65) = 0.8$. There was a significant difference between the baseline and the morphological condition, $t(65) = 3.47$, $p < .001$.

In sum, the predictions of this experiment were borne out. Morphologically related primes (those with shared suffixes) facilitated reaction times to targets. No priming effects were obtained in the case of semantically related or formally related primes.

3.8 Item differences

As mentioned above, there was a main effect of List. Though our experimental design and subsequent statistical analysis corrected for the possibility of individual items driving any potential effect, it is still interesting that individual items behave somewhat differently. Investigating the properties of items that contribute to priming patterns is an obvious next step in attempting to understand the nature of the mental lexicon.

4 Discussion

One objective of this experiment was to establish delayed repetition priming as a reliable measure of suffix priming independent of semantic confounds. The results show that suffixes do prime each other. This effect cannot be attributed to shared semantics or to shared form as there was no difference in reaction times from the semantically and formally primed targets as compared to the targets that had no prime. Additionally, the morphologically related primes and targets were not semantically related as whole words at all. The degree of phonological overlap in the morphological condition was very small, as well. The priming effect, then, appears to be a purely morphological one. This reinforces the validity of the delayed repetition priming paradigm as being a more restrictive means of establishing morphological relationships between words in the mental lexicon in that morphological priming persists where semantic and formal priming are not perceptible.

There are larger issues on which the results of this experiment can be brought to bear: are relationships between words explicitly encoded in the lexicon and are complex words represented in terms of their constituent parts? The current results strongly suggest that suffixes are indeed accessed as constituents. Morphological relationships, though, entail semantic and form relationships, at least to some degree, so what does it mean when there is an effect of morphological priming but no semantic or form priming? I would assert that whatever morphology turns out to be (see the following discussion), morphological priming is identity priming, i.e., the *ness* in *happiness* is exactly the same representation that is accessed when perceiving the *ness* in *darkness*, and that the delayed repetition priming paradigm is the most restrictive measure of morphological priming, since semantic and formal priming effects are filtered out by the delay.

4.1 A connectionist perspective

One alternative to morphological priming as identity priming might be something like a prototype theory where the *ness* in *happiness* has a certain *ness*-ness about it and the *ness* in *darkness* has a certain *ness*-ness about it and depending on the relationship of that string to the rest of the word (the stem), the string could more or less resemble the essence of *ness*-ness. This would be the kind of view on which only semantic and phonological relatedness would account for morphological priming effects. If there is no single representation, then it is only to the extent that there is overlap on a given dimension(s) between multiple representations that there can be a relationship between such representations. Indeed, this seems to be the perspective of many connectionist accounts of priming.

Some connectionists take issue with morphology as an independent level of representation, and instead argue that it is an emergent property of semantic and phonological/ orthographic interactions: morphology is a convergence of codes. Under this view, there are no morphemes per se, but morphological structure (and perhaps even identity) is reflected in hidden units and the strength of the connections between semantics and form (Seidenberg 1993). One appeal of this approach is that it can account very naturally for graded effects of priming: degrees of semantic and phonological relatedness affecting the degree of priming. However, this phenomenon has only been rigorously established in the immediate repetition paradigms. Systematically addressing this in the delayed repetition paradigm where semantic and phonological priming is not present is a necessary next step.

Certainly, the connectionist approach has the advantage of being the more parsimonious explanation in that it does not posit an additional level of representation. If semantic and phonological interactions were sufficient to account for experimental data and could be learned such that the experimental results follow, this would be a preferable approach. No aspect of formal linguistic theory hinges on this issue: linguistic analyses of word combinations effectively exploit this level of description and will continue to, whether or not it is part of the way the lexicon is organized. However, the nature of representations does indeed play a role in understanding the nature of the computations that are performed over them and hence in understanding how we produce and comprehend language.

4.2 A traditional approach

The primary difference between this connectionist view and a more traditional symbolic view is the existence (or lack thereof) of a morphological level of representation. While I think that the lack of understanding of hidden units precludes being able to make a substantial claim here on the part of the connectionists, the agenda is to establish that such a level is unnecessary. Other differences between the two approaches center around the discreteness of 'morphological' units and their status as linguistically primitive.

4.2.1 Morphological level of representation

On a symbolic account morphemes are constituents that can be manipulated and that enter into relationships with other morphemes via combinatory rules. It is this view that is supported by theoretical linguistics. As mentioned briefly above, a significant aspect of morphology is that it allows linguists and possibly speaker/hearers to establish relationships between words. As a result, patterns emerge, often in the form of paradigms, and based on these patterns we are able to make further generalizations; i.e., we are able to exploit the structure in the data to learn more about our language,[24] as well as comprehend and produce novel combinations of these constituent strings. Additionally, morphemes have particular describable properties: what order they can occur in under what circumstances, what phonological changes they induce, what they can and cannot attach to, what other morphemes block them from occurring. And further, phonological rules make reference to morphemes. Indeed, as linguists study large numbers of languages, the distributional evidence for morphemes as organizational units of language continues to be extensive (see, for example, recent database projects such as Syntactic Structures of the World's Languages/Terraling (Koopman, Collins and Kayne, 2011) or Language and Gene Lineages (Longobardi and Guardino, 2009). Just as a notion of morpheme is a useful in linguistic analysis, it is ostensibly useful as a means of organizing our lexicons, as well.

4.2.2 Discreteness and Primitives

Whether or not morphemes are 'discrete' is not a particularly interesting question from a theoretical point of view: either way, morphological analysis which makes use of discrete morphemes is a very informative and efficient means of describing the world's languages. In trying to account for the differences between morphological effects and semantic and phonological effects in the priming experiment here and in previous research, however, the discreteness of the representation could be an important factor inasmuch as we accept the hypothesis that morphological priming is identity priming. Clearly, though, morphemes are approximately discrete enough to be manipulated as constituents.

The question of whether or not morphemes are emergent or linguistic primitives is more substantial, but can only be answered with a rigorous learning model, which the field is currently lacking.

4.3 Morphology as Rules

Another possibility is that the suffixes themselves are not independently represented linguistic units, but rather, a by-product of a morphosyntactic or morpho-phonological rule. This looks particularly appealing when considering cross-linguistic data where a morphological process involves truncation, reduplication, template filling, or suprasegmental alternations (e.g., tone) rather than concatenation. The interpretation of the experimental results here within this framework would be that the rule itself is being primed. One would have to take a worked out theory (e.g., Anderson, 1992) and compare its predictions for priming results with the body of experimental data, a task not undertaken for this paper.

[24] How generalizations are constrained and whether or not they are useful in hypothesizing a grammar during the learning stages of language are interesting questions to pursue.

4.4 Gradience

While the current experiment did not address gradience at all (indeed it has not been systematically established that there are gradient priming effects in the delayed repetition priming paradigm), the robustness of gradience in immediate repetition experiments is something that any model of the lexicon should account for.

Though connectionist accounts of morphological priming claim to account nicely for gradient effects, the issue is a substantial one for traditional accounts. Generally, it has been addressed in terms of whether or not complex words are stored in the lexicon as monolithic forms or in terms of their component parts. Twenty years ago, models of morphological representation and processing advocated one or the other view (see Taft and Forster 1975 for a decomposition model and Bradley 1980 and Butterworth 1983 for whole-word representations). Today, there are hybrid models that allow both and vary in terms of criteria for whole vs. decomposed representations (Schreuder and Baayen, 1995, Burani and Thornton 2003, Hay and Baayen, 2005). Decomposition can be based on semantic transparency, frequency of root, relative frequency of root to derived word, productivity of an affix, position of an affix, etc. For example, *casualty* does not prime *casual*. It is semantically opaque relative to its stem and so is assumed by most researchers to be listed as a full form and have no decomposed representation.

On symbolic accounts, then, the gradient nature of priming effects should arise as a result of performance factors that are independent of the nature of the representations (see McQueen and Cutler 1998 for a review), instantiating the fundamental distinction between competence and performance maintained in modern linguistics (Chomsky 1965). Perhaps such accounts could be flexible in the extent to which performance factors can affect representations, but the principles underlying the organization of such representations should not be contingent on a processing explanation (except inasmuch as representation and process are inherently linked from a biological or evolutionary perspective).

However, the fact that the degree of semantic transparency affects the amount of priming is still hard to account for in these models. Can words be ‘partially’ decomposed? Presumably, this could be addressed in terms of competition between two representations, a fully decomposed version and a fully listed version. However, the influence of these factors on decomposition in these models is merely stipulated. Again, a learning model from which these influences could be derived is lacking. The advantage of a traditional/symbolic approach, however, is that it readily captures generalizations that can be used in linguistic analysis of the sort that can be very valuable not only to linguists, but for speaker/hearers in comprehension, production and learning of language.

4.5 Summary

The results of this experiment can only be construed as suggestive in addressing the issues of morphological levels of representation, discreteness of representations, morphemes as linguistic primitives, and morphemes as the organizational units of the mental lexicon. Still, the fact that, in the absence of any semantic or form priming, morphological priming still obtains is quite extraordinary, particularly if we are to take morphology to be the (non-linear) sum of the convergence of semantic and formal codes. Further, the fact that it is suffix priming rather than stem priming is significant. Recall that morphologically related words, while sharing a suffix, were not considered to be semantically related to one another, as a result of their stems being entirely distinct.

5 Future Research

There are three directions of varying scope that I am interested in pursuing as extensions of this research. Most immediately, I want to explore the differences among the particular affixes used as stimuli in this experiment, from both a performance and a linguistic perspective. Secondly, I want to look at priming effects in semantically opaque complex words where the affix's contribution to the meaning of the whole is consistent with its contribution in transparent cases (e.g., *virtual*), using the delayed repetition priming paradigm. As a long-term project, I am interested in developing a realistic morphological learner that is consistent with experimental results.

5.1 Properties of Affixes

The effect of List in this experiment was quite substantial, indicating that different affixes behave differently (recall that the form of the statistical analysis picks out an overall effect that is not driven by the properties of particular items, however). Rigorous post-hoc analyses should be undertaken to consider the properties of individual items and assess their contributions to the priming effects observed. The following are the properties to be investigated.

5.1.1 Affix Level

Based on the distribution of affixes and their phonological interaction with their stems, Kiparsky (1983) proposed distinct levels of affix attachment: stem-changing affixes attach earlier than stem-preserving affixes. Analyses based on syntactic and semantic relationships of affixes to stems reinforce such a distinction (Wasow 1977), with early attachment being more opaque and later affixes being more transparent. Higher productivity also seems to be associated with later affixes. The notion of sequential derivation is not necessary in maintaining this distinction, which roughly translates into inner layers of affixes (early) and outer layers (late). It has been proposed that inner layer combinations are the result of lexical processes and outer layers of syntactic ones (Dubinsky and Simango 1996,). Without committing to this hypothesis, it would be interesting to consider the particular affixes of this experiment in this context. Do inner and outer (Level 1 and Level 2) affixes behave differently with regard to lexical priming?

Indeed, work by Vannest and Boland (1999) purports to have evidence for maintaining such a distinction in the mental lexicon. They claim evidence that suggests that words with Level 2 affixes are decomposed and that words with Level 1 affixes are not. However, they only considered a single Level 2 affix (*-less*) and two Level 1 affixes (*-ity*, *-ation*).

Notice, though, that it is possible to break down affix levels into distinct properties and look at the contribution of each of those properties individually:

<u>1. Phonological Variation</u>

Researchers have investigated the degree to which phonological changes to a stem as the result of affixation affect lexical access. Research has shown that reaction times to words in a straight lexical decision task (i.e., no preceding prime) that underwent some phonological change as the result of derivational affixation were significantly slower than those that did not (Tsapkini et al., 1999). In conjunction with Gonnerman's gradient results with regard to the degree of phonological change between stem and derived word, it is certainly possible that, though the suffixes are phonologically consistent, a greater distance between root and stem would result in decreased priming. Gonnerman showed that the nature of the change also played a role, with vowel changes (*sane–sanity*) showing longer latencies than single consonant changes (*frequent–frequency*).

2. Semantic Relatedness

Although no effect of semantic relatedness was present in the priming results of this experiment, it is possible that this factor did play a more subtle role. Indeed, while the semantic relatedness of the derived target stimuli to their stems was determined to be transparent by four linguists independently, submission of derived–base pairs to a Latent Semantic Analysis (LSA) revealed that the degree of relatedness, or transparency, was, in fact, highly variable ranging from 0.08 to 0.72 with an average of 0.35. A correlation with R^2 of .36 was found between semantic transparency of the targets (based on the LSA score) and degree of priming. The semantic relatedness of the derived primes to their stems, however, (the LSA average of which was the same as that of the targets), did not yield a significant correlation: $R^2 = .02$. This experiment was not designed to consider this, but it could and should be investigated.

It would be of particular interest to consider the average semantic relatedness of each affix to its free stem, e.g., looking at all occurrences of *-ful*, what is its average relatedness to its stems?

5.1.2 Distributional Factors

Recently, statistical analyses of linguistic data have reemerged as potential explanatory tools, particularly in the context of language processing. The effect of frequency seems to be a robust one, not just in language, but in other aspects of cognition as well. Another factor that seems to contribute to the representational status of morphemes is productivity, i.e., the degree to which a process (e.g., the concatenation of a specific affix to various stems) can be used to form novel words, though how to measure it is a difficult question. Finally, there is evidence that knowledge of phonotactics, or combinations of legal sequences of sounds in a particular language, is exploited in determining morpheme boundaries and consequently could effect morphological representations. Again the contribution of these factors in predicting the priming results should be investigated.

1. Frequency

Frequency has been argued to influence the speed with which we access lexical representations, with experimental results to back up the claim (Forster and Davis 1984, Gordon and Caramazza 1985, Embick et al 2001). Additionally, frequency is claimed to affect the extent to which representations are decomposed, with high frequency items being less likely (Hay 2001).

Schreuder and Baayen (1997) showed that the number of derivatives of a root (morphological family size) also speeds reaction times, but the morphological family frequency (i.e., the sum of the frequencies of all members of a morphological family) has no effect.

Hay (2002) argues that the relevant measure is, in fact, relative frequency. By this she means the frequency of the root or stem relative to the frequency of the derived word (e.g., of *happy* to *happiness*). It turns out that derived words with higher frequencies than their stems look 'less decomposed' according to her experimental results. She also correlates this relative frequency with productivity (Hay and Baayen 2002, 2003).

2. Productivity

Defining an appropriate measure of productivity has been a difficult problem (Bauer, 2001). One easily computable measure is that of hapax legomenon formulated by Harald Baayen (Baayen, 1993), where productivity equals the number of words formed by a particular affix that occur only once in a corpus (the hapax legomena) divided by the total token frequency of words with that affix: $P = n_1/N$. However, increasing the sample size results in lower productivity, so the measure is relative. The intuition behind this measure is that high productivity will necessarily result in a larger number of types. The larger the number of types, the less likely they are to be represented more than once in a sample. Additionally, highly productive affixes should lead to coinages of new words, which also are unlikely to be represented more than once.

Another measure is to consider the proportion of truly affixed words vs. pseudo-affixed words, e.g., how many strings ending in *-y* are affixed (*dirty*) and how many are in fact monomorphemic (*party*)?[25]

3. Phonotactics

Research has also been done showing that speaker/hearers are sensitive to phonotactics when determining whether or not nonce words are likely to be represented as mono- or bi-morphemic (Pierrehumbert 1994; Hay, Pierrehumbert and Beckman 2003; VanWagenen 2004). Investigating the extent to which phonotactics contribute to likely segmentation in the target stimuli of this experiment would also be interesting. Additionally, the simple length in segments of the stem and/or the affix could contribute to the probability of a word being decomposed.

5.2 Affixes as constituents: the case of semantically opaque words

The genesis of this work was to find a paradigm that could assess the polymorphemicity of words for the purpose of challenging the widely accepted hypothesis that semantically opaque words (e.g., *virtual*) are represented as lexical entries with no internal structure, i.e., that they are stored as monolithic forms.[26] Having established delayed repetition priming as a reliable means of determining that affixes are represented as constituents, the next experiment will involve testing these semantically opaque items. If the affix makes the same semantic contribution to the opaque word as it does in transparent words, will morphological priming be obtained? For example, will the *-al* of *virtual* prime the *-al* of *ethical*? Or will the degree of semantic transparency of the derived form to its base override the semantically and formally consistent contribution of the affix?

Another interesting possibility to consider is whether or not affixes attached to real stems prime affixes attached to pseudo-stems, e.g., does *sanity* prime *bleffity*?[27]

5.3 Learning Models

As previously discussed, I think that a computationally sound learning model that is not only consistent with the experimental data, but from which the experimental data can be derived, is the most explicit way to test a theory of morphological representation and processing. Specifically, how do each of the previously discussed factors contribute to lexical representations and to what extent and why?

Several word/morpheme segmentation models exist, most of which are based on distributional cues in written corpora. Such models use information theoretic notions such as Minimum Description Length (Baroni 2000, Brent and Cartwright 1996, Goldsmith 2001, 2006, etc.) and other statistical inference tools (Schone and Jurafsky 2000). Brent

[25] Thanks to Bruce Hayes for this 'batting average' idea of productivity.

[26] The idea of testing semantically opaque words via affix priming originated with Marco Baroni and Carson Schütze.

[27] Thanks to Bruce Hayes for this suggestion.

and Siskind (2001) actually make use of real child language acquisition data and consider transcriptions from child-directed speech in order to assess the feasibility of using isolated words (not in a continuous stream of words) as key factors in segmentation. These models are concerned with being able to identify and isolate morphemes in corpora that are untagged. They provide excellent computational insights into possibilities for a model of the human natural language learner which can be tested and improved by considering human data.

Ultimately, a rigorous learning model will, I think, provide the most insight into these issues of representation, but in the meantime, we can continue to amass experimental evidence for such a model to account for.

6 Conclusion

This work focuses on the representation of suffixes. Suffixes are often considered the least likely candidates for independent representation in the mental lexicon. Assuming linear processing, it is rarely necessary to hear an entire suffix before being able to identify a derived stem; therefore, the advantage of representing suffixes as unique constituents is not immediately obvious. (And, indeed, only a detailed learning model can assess the efficiency of independent representation of suffixes.)

Previous work in the repetition priming paradigm using a lexical decision task has confounded semantic overlap with shared morphology. The present study establishes delayed repetition priming as a reliable paradigm for assessing whether or not strings are psychologically represented as constituents, i.e., as manipulable chunks. Delayed repetition displays the target several items after the probe (an average of 20 items in the current experiment). At this lag, there are no priming effects of semantic relatives: *joy* does not prime *happiness* as it does when the target immediately follows the probe. Additionally, there were no effects of form relatives: *harness* does not prime *happiness*. There was, however, a significant facilitatory priming effect of morphological relative: *darkness* → *happiness*.

These results show suffixes to be represented as manipulable chunks and provide further evidence for a morphologically organized mental lexicon.

Acknowledgments

I am indebted to the following for their willingness to discuss ideas with me at various points in this project: Mark Allen, William Badecker, Sid Kouider, Alec Marantz, Liina Pylkkänen, Philippe Schlenker, and Colin Wilson. Thanks also to my grad school peers Jeff Heinz, Ben Keil, Christina Kim, Greg Kobele, Katya Pertsova, and Marcus Smith for advice, support, and encouragement. I am grateful to my M.A. committee (Kie Zuraw, Bruce Hayes, Carson Schütze) for their extraordinary patience and insightful criticisms. Of course, all the usual disclaimers apply.

Appendix: Stimuli

Target	Morphological Prime	Semantic Prime	Formal Prime
manager	publisher	coach	manatee
idolize	victimize	revere	iodine
loveless	faultless	alone	level
obsession	confession	desire	assassin
logician	beautician	genius	legion
villainous	odorous	wicked	violin
redeemable	exportable	exchange	redouble
observant	expectant	perceive	serpent
improvement	achievement	upgrade	deprave
apprenticeship	companionship	amateur	appendix
thankful	wasteful	pleased	tranquil
reddish	sluggish	crimson	riddle
terrorist	motorist	rebel	tortoise
diversify	intensify	transform	riverside
leakage	blockage	outflow	lozenge
ethical	functional	honor	tickle
supportive	digestive	assist	sporting
equality	intensity	fair	squalid
heroism	humanism	valor	heresy
wisdom	freedom	knowledge	widow
sculpture	failure	model	sulphur
simpleton	singleton	numskull	simpering
girlhood	sainthood	juvenile	garland
crudely	darkly	unrefined	cruel
privacy	frequency	solitude	rival
dampen	sweeten	stifle	dame
bothersome	worrisome	pestering	brother
visionary	cautionary	dream	visor
habitual	conceptual	custom	halibut
reversible	producible	correct	reverie
metallic	poetic	aluminum	malice
stringy	waxy	fiber	stingray
seaward	homeward	voyage	scabbard
statuette	kitchenette	figurine	statute
uppermost	outermost	top	thermos
seedling	duckling	embryo	beeline
godlike	lifelike	devout	goggle
piglet	droplet	pork	picket
hardness	flatness	dense	harvest
performance	appearance	show	romance

References

Anderson, Stephen. R. 1992. *A-morphous morphology*. Cambridge: Cambridge University Press.

Baayen, Harald. 1993. On frequency, transparency and productivity. In *Yearbook of Morphology 1992*, ed. by Geert Booij and Jaap van Marle, 181–208. Dordrecht: Kluwer Academic Publishers.

Baroni, Marco. 2000. Using distributional information to discover morphemes: A distribution-driven prefix learner. Ph.D. dissertation, UCLA.

Bauer, Laurie. 2001. *Morphological productivity.* Cambridge: Cambridge University Press.

Bentin, S., and L.B. Feldman. 1990. The contribution of morphological and semantic relatedness to repetition priming at long and short lags: Evidence from Hebrew. *Quarterly Journal of Experimental Psychology* 42A:693–711.

Bock, J. Kathryn. 1986. Syntactic persistence in language production. *Cognitive Psychology* 18:355–387.

Bock, Kathryn, and Zenzi M. Griffin. (2000). The persistence of structural priming: Transient activation or implicit learning? *Journal of Experimental Psychology: General* 129:177–192.

Bowers, Jeffrey S. 1996. Different perceptual codes support priming for words and pseudowords: Was Morton right all along? *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22:1336–1353.

Bowers, Jeffrey S. 2000. In defense of abstractionist theories of repetition priming and word identification. *Psychonomic Bulletin and Review* 7:83–99.

Bowers, Jeffrey S. and Sid Kouider. 2003. Developing theories of priming with an eye on function. In *Rethinking implicit memory*, ed. by Jeffrey S. Bowers and Chad J. Marsolek, 19–40. New York: Oxford University Press.

Burani, C. and A.M. Thornton. 2003. The interplay of root, suffix and whole-word frequency in processing derived words. In *Morphological structure in language processing,* ed. by R.H. Baayen and R. Schreuder, 157–208. Mouton de Gruyter.

Bradley, Dianne 1980. Lexical representation of derivational relation. In *Juncture*, ed. by Mark Aronoff and Mary-Louis Keaton, 37–55. Saratoga, CA: Anma Libri.

Brent, Michael R., and Timothy A. Cartwright. 1996. Distributional regularity and phonotactic constraints are useful for segmentation. *Cognition* 61:93–125.

Brent, Michael R., and Jeffrey Mark Siskind. 2001. The role of exposure to isolated words in early vocabulary development. *Cognition* 81:B33–B44.

Butterworth, Brian. 1983. Lexical representation. In *Language production*, *volume 2: Development, writing and other language processes*, ed. by Brian Butterworth, 257–294. London: Academic Press.

Chomsky, Noam. 1965. *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

Clark, Herbert H. 1973. The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior* 12:335–359.

Cohen, Jonathan, Brian MacWhinney, Matthew Flatt, and Jefferson Provost (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments, and Computers* 25:257–271.

Dubinsky, Stanley, and Sylvester Ron Simango 1996. Passive and stative in Chichewa: Evidence for modular distinctions in grammar. *Language* 72:749-781.

Embick, David, Martin Hackl, Jeannette Schaeffer, Meltem Kelepir, and Alec Marantz. A magnetoencephalographic component whose latency reflects lexical frequency. *Cognitive Brain Research* 10:345–348.

Feldman, Laurie Beth. 2000. Are morphological effects distinguishable from the effects of shared meaning and shared form? *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26:1431–1444.

Fellbaum, Christiane, ed. 1998. *WordNet: An electronic lexical database*. Cambridge, MA: MIT Press.

Forster, K. I., and C. Davis. (1984). Repetition priming and frequency attenuation in lexical access. *Journal of Experimental Psychology: Learning, Memory, and Cognition*: 10, 680–698.

Giraudo, Hélène, and Jonathan Grainger. 2003. On the role of derivational affixes in recognizing complex words: Evidence from masked priming. In *Morphological structure in language processing*, ed. by R. Harald Baayen and Robert Schreuder, 209–232. Berlin: Mouton de Gruyer.

Goldsmith, John. 2001. Unsupervised learning of the morphology of a natural language. *Computational Linguistics* 27:153–198.

Goldsmith, John (2006). An algorithm for the unsupervised learning of morphology. *Natural Language Engineering* 12(3):1–19.

Gonnerman, Laura Michelle. 1999. Morphology and the lexicon: Exploring the semantics–phonology interface. Ph.D. dissertation, University of Southern California.

Gordon, B., and A. Caramazza. 1985. Lexical access and frequency sensitivity: Frequency saturation and open/closed class equivalence. *Cognition* 21:95–115.

Hay, J. (2001). Lexical frequency in morphology: Is everything relative? *Linguistics* 39:1041–1070.

Hay, Jennifer. 2002. From speech perception to morphology: Affix ordering revisited. *Language* 78:527–555.

Hay, Jennifer, and Harald Baayen. 2002. Parsing and productivity. In *Yearbook of Morphology 2001*, ed. by Geert E. Booij and Jaap Van Marle, 203–235. Dordrecht: Kluwer Academic Publishers.

Hay, Jennifer, and Harald Baayen. 2003 Phonotactics, parsing and productivity. *Italian Journal of Linguistics* 15:99–130.

Hay, Jennifer, and Harald Baayen. 2005. Shifting paradigms: gradient structure in morphology. *Trends in Cognitive Sciences* 9:342–348.

Hay, Jennifer, Janet Pierrehumbert, and Mary E. Beckman. 2003. Speech perception, well-formedness and the statistics of the lexicon. In *Phonetic interpretation: Papers in Laboratory Phonology VI*, ed. by John Local, Richard Ogden and Rosalind Temple, 58–74. Cambridge: Cambridge University Press.

Henderson, L., J. Wallace and K. Knight. 1984. Morphemic structure and lexical access. In *Attention and performance X: Control of langugae processes*, ed. by H. Bouma and D. Bouwhuis, 211–226. Hillsdale, NJ: Lawrence Erlbaum.

Kiparsky, Paul. 1983. Word-formation and the lexicon. In *Proceedings of the 1982 Mid-America Linguistics Conference,* ed. by Francis Ingemann, 3–32. University of Kansas, Lawrence.

Koopman, H., C. Collins, and R. Kayne. (2011). Syntactic Structures of the World's Languages Database: http://sswl.railsplayground.net/.

Kouider el Ouahed, Sid-Ahmed. 2002. Rôle de la conscience dans la perception des mots. Thèse de doctorat, École des hautes études en sciences sociales, Paris.

Kouider, Sid, and Emmanuel Dupoux. 2009. Episodic accessibility and morphological processing: Evidence from long-term auditory priming. *Acta Psychologica* 130:38–47.

Kruskal, Joseph B., and Mark Liberman. 1983. The symmetric time-warping problem: From continuous to discrete. In *Time warps, string edits, and macromolecules: The theory and practice of sequence comparison*, ed. by David Sankoff and Joseph B. Kruskal, 125–161. Reading, MA: Addison-Wesley.

Landauer, Thomas K., Peter W. Foltz, and Darrell Laham. 1998. An introduction to Latent Semantic Analysis. *Discourse Processes* 25:259–284.

Longobardi, Giuseppe, and Cristina Guardiano. 2009. Evidence for syntax as a signal of historical relatedness. *Lingua* 119:1679–1706.

Luce, P.A., D.B. Pisoni, and S.D. Goldinger. 1990. Similarity neighborhoods of spoken words. In *Cognitive Models of Speech Processing: Psycholinguistic and computational perspectives,* ed. by Gerry T.M. Altmann, 122–147. Cambridge, MA: MIT Press.

Marslen-Wilson, William. 1999. Abstractness and combination: The morphemic lexicon. In *Language processing*, ed. by Simon Garrod and Martin J. Pickering, 101–119. Hove, UK: Psychology Press.

Marslen-Wilson, William, Lorraine K. Tyler, Rachelle Waksler, and Lianne Older. 1994. Morphology and meaning in the mental lexicon. *Psychological Review* 101:3–33.

Marslen-Wilson, William D., Mike Ford, Lianne Older, and Zhou Xiaolin. 1996. The combinatorial lexicon: Priming derivational affixes. In *Proceedings of the eighteenth Annual Conference of the Cognitive Science Society*, ed. by Garrison W. Cottrell, 223–227. Mahwah, NJ: Lawrence Erlbaum.

McClelland, James L., and David E. Rumelhart. 1981. An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review* 88:375–407.

McQueen, James M., and Anne Cutler. 1998. Morphology in word recognition. In *The handbook of morphology*, ed. by Andrew Spencer and Arnold M. Zwicky, 406–427. Oxford: Blackwell.

Morton, J. 1969. The interaction of information in word recognition. *Psychological Review* 76:165–178.

Morton, John. 1979. Word recognition. In *Psycholinguistics 2: Structures and processes*, ed. by John Morton and John C. Marshall, 107–156. London: Paul Elek.

Perea, Manuel, and Arcardia Gotor. 1997. Associative and semantic priming effects occur at very short stimulus-onset asynchronies in lexical decision and naming. *Cognition* 62:223–240.

Perea, Manuel, and Eva Rosa. 2002. The effects of associative and semantic priming in the lexical decision task. *Psychological Research* 66:180–194.

Pierrehumbert, Janet. 1994. Syllable structure and word structure: a study of triconsonantal clusters in English. In *Phonological structure and phonetic form: Papers in Laboratory Phonology* 3, ed. by Patricia A. Keating, 168–188. Cambridge: Cambridge University Press.

Plaut, D.C., and L.M. Gonnerman. 2000. Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? *Language and Cognitive Processes* 15:445–485.

Pylkkännen, Liina, Andrew Stringfellow, and Alec Marantz. 2002. Neuromagnetic evidence for the timing of lexical activation: an MEG component sensitive to phonotactic probability but not to neighborhood density. *Brain and Language* 81:666–678.

Pylkkänen, Liina, and Gregory Murphy. 2004. Sense competition and the representation of polysemy. Presented at the Linguistic Society of America, Boston.

Raaijmakers, Jeroen G. W., Joseph M.C. Schrijnemakers, and Frans Gremmen. 1999. How to deal with "the language-as-fixed-effect fallacy": Common misconceptions and alternative solutions. *Journal of Memory and Language* 41:416–426.

Rastle, Kathleen, Matthew H. Davis, and Boris New. 2004. The broth in my brother's brothel: Morpho-orthographic segmentation in visual word recognition. *Psychonomic Bulletin and Review* 11:1090–1098.

Rastle, Kathleen, Jonathan Harrington, and Max Coltheart. 2002. 358,534 nonwords: The ARC Nonword Database. *Quarterly Journal of Experimental Psychology* 55A:1339–1362.

Schone, Patrick, and Daniel Jurafsky. 2000. Knowledge-free induction of morphology using Latent Semantic Analysis. In *Proceedings of CoNLL-2000 and LLL-2000*, ed. by Claire Cardie, Walter Daelemans, Claire Nédellec and Erik Tjong Kim Sang, 67–72. New Brunswick, NJ: Association for Computational Linguistics.

Schreuder, Robert, and R. Harald Baayen. 1995. Modeling morphological processing. In *Morphological aspects of language processing*, ed. by Laurie Beth Feldman, 131–154. Hillsdale, NJ: Lawrence Erlbaum.

Schreuder, Robert, and R. Harald Baayen. 1997. How complex simplex words can be. *Journal of Memory and Language* 37:118–139.

Seidenberg, Mark S. 1993. A connectionist modeling approach to word recognition and dyslexia. *Psychological Science* 4:299–304.

Taft, Marcus, and Kenneth I. Forster. 1975. Lexical storage and retrieval of prefixed words. *Journal of Verbal Learning and Verbal Behavior* 14:638–647.

Tenpenny, P. L. 1995. Abstractionist versus episodic theories of repetition priming and word identification. *Psychonomic Bulletin and Review* 2:339–363.

Tsapkini, Kyrana, Eva Kehayia, and Gonia Jarema. 1999. Does phonological change play a role in the recognition of derived forms across modalities? *Brain and Language* 68:318–323.

Vannest, Jennifer, and Julie E. Boland. 1999. Lexical morphology and lexical access. *Brain and Language* 68:324–332.

VanWagenen, Sarah. 2004. The role of phonotactics in morphological decomposition. Ms., UCLA.

Wasow, Thomas. 1977. Transformations and the lexicon. In *Formal syntax*, ed. by Peter W. Culicover, Thomas Wasow and Adrian Akmajian, 327–360. New York: Academic Press.

Affiliation

Sarah VanWagenen
Department of Linguistics
UCLA