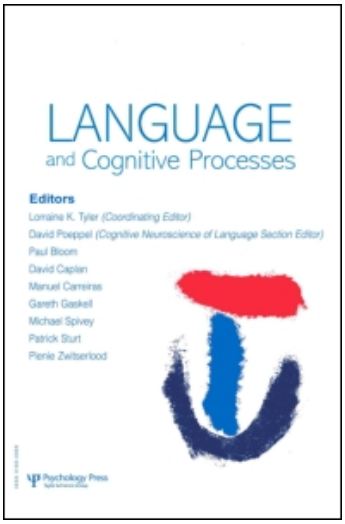

Language and Cognitive Processes
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713683153

A sticky stick? The locus of morphological representation in the lexicon

Marcus Taft [a]; Minh Nguyen-Hoan [a]
[a] University of New South Wales, Sydney, Australia

First published on: 30 June 2009

To cite this Article Taft, Marcus and Nguyen-Hoan, Minh(2010) 'A sticky stick? The locus of morphological representation in the lexicon', Language and Cognitive Processes, 25: 2, 277 — 296, First published on: 30 June 2009 (iFirst)
To link to this Article: DOI: 10.1080/01690960903043261
URL: http://dx.doi.org/10.1080/01690960903043261

LANGUAGE AND COGNITIVE PROCESSES
2010, 25 (2), 277–296

Psychology Press
Taylor & Francis Group

A sticky stick? The locus of morphological representation in the lexicon

Marcus Taft and Minh Nguyen-Hoan
University of New South Wales, Sydney, Australia

It is demonstrated that the meaning given to an ambiguous word (e.g., *stick*) can be biased by the masked presentation of a polymorphemic word derived from that meaning (e.g., *sticky*). No bias in interpretation is observed when the masked prime is a word that is semantically related to the target with no morphological relationship (e.g., *glue*), though such a semantically based bias is revealed when the prime is unmasked. Because the masked priming results cannot be explained in terms of facilitation of processing at either the form level or the semantic level, it is concluded that an intermediate level provides the locus of the effect, referred to here as the 'lemma' level. Thus, any model of lexical processing needs to incorporate such an intermediate level to capture the relationship between stem morphemes and their derived forms.

Keywords: Lemma; Lexical representation; Morphemic representation; Morphological processing; Semantic priming.

It has been clearly shown that the recognition of a word (e.g., *hunt*) is facilitated by the immediately prior presentation of a polymorphemic word derived from it (e.g., *hunter*) even when that prime is masked from consciousness (e.g., Diependaele, Sandra, & Grainger, 2005; Lavric, Clapp, & Rastle, 2007; Longtin, Segui, & Halle, 2003; Marslen-Wilson, Bozic, & Randall, 2008; McCormick, Rastle, & Davis, 2008; Rastle, Davis, & New, 2004). Such a finding indicates that derived words are decomposed into their stem and affix at some stage in their processing. The likelihood that such decomposition is early and automatic (i.e., obligatory) can be seen in a further

Correspondence should be addressed to Marcus Taft, School of Psychology, University of NSW, Sydney NSW 2052, Australia. E-mail: M.Taft@unsw.edu.au

The research reported in this paper was supported by a grant from the Australian Research Council to the first author. The statistical assistance of Goran Lazendic is gratefully acknowledged.

http://www.psypress.com/lcp DOI: 10.1080/01690960903043261

result reported in the above studies whereby facilitation is also found when the prime is actually a 'pseudo-affixed' word, being monomorphemic while looking as though it contains the target as its stem (e.g., *corner* priming *corn*). Finally, such masked priming effects cannot be explained merely in terms of the orthographic overlap between prime and target, because none of the studies find any effect when the prime ends in letters that can never function as an affix (e.g., *turnip* does not prime *turn*).

So, when a letter-string ends in letters that form an affix, the standard account is that these letters are treated separately from the rest of the letter-string at an early stage of processing (though see Baayen, Wurm, & Aycock, 2007, for an alternative approach). Recognition of the polymorphemic word therefore involves the activation of the same lexical representation that is used to recognise its putative stem when that stem is a free-standing word. What is not clear, however, is how this idea should be instantiated within a theoretical model.

The explanation typically given for the above masked priming effects has focused on the form-based level of representation which, in the case of visual word recognition, corresponds to the orthography of the word. So, the orthographic representation for the word *corn* is activated at the earliest stages of processing *corner*, hence facilitating its subsequent activation when the word *corn* is then presented. The same is true of *hunt* in relation to *hunter*, but not of *turn* in relation to *turnip* because *ip* will not be treated as an affix and, hence, the orthographic unit *turn* will not be isolated. Thus, the pattern of masked priming effects can be readily explained at the form level of representation. However, there is another level of representation that can capture morphological relationships, and this is a level that is intermediate between the form of the input (orthography in this case) and the functional output (i.e., meaning and syntax), capturing the correlation between them.

AN INTERMEDIATE LEVEL OF REPRESENTATION

Representing the correlation between input and output is characteristic of distributed connectionist models (e.g., Seidenberg & McClelland, 1989, Plaut, McClelland, Seidenberg, & Patterson, 1996). In such models, the correlation is captured by sub-patterns of activation that develop within a hidden layer of units that mediates between input and output. Morphemic structure emerges when stable sub-patterns develop in response to the consistent co-occurrence of orthographic and semantic information (e.g., Plaut & Gonnerman, 2000; Raveh, 2002; Rueckl & Raveh, 1999; Seidenberg & Gonnerman, 2000). For example, the relationship between *hunt*, *hunter*, *hunting*, *huntsman*, etc., is captured at this level by virtue of the fact that the recurring form *hunt* is associated with a similar meaning in all of those

words. Similarly, the consistent meaning associated with the recurrence of *venge* in *revenge*, *avenge*, *vengeance*, and *vengeful* will be captured at this level, despite the fact that *venge* is not a word in its own right.

According to a distributed connectionist model, then, the priming of a word (e.g., *hunt*) by its morphological variant (e.g., *hunter*) arises at this intermediate level of representation (e.g., Plaut & Gonnerman, 2000). What is hard to capture, however, is the priming of *corn* by *corner* because a pseudo-affixed word will not develop any relationship with its pseudo-stem at the intermediate level (see e.g., Longtin et al., 2003; Rastle et al., 2004). While there might be some facilitation arising from their shared input units, this should be equally true in the case of *turnip* and *turn* where no facilitation has been observed. Therefore, the finding of pseudo-affixed word priming implies a structuring of the orthographic units in such a way that an affix (e.g., *er*) is treated as a separate unit from its stem (e.g., *hunt*). In the process of recognising the pseudo-affixed word *corner*, the same structural decomposition will blindly occur, leading to the priming of *corn*.

The need for a structuring of the orthographic units does not, however, rule out the possibility of an intermediate level of representation that captures morphological relationships. It is possible for the orthographic units to be structured while, at the same time, the correlation between those orthographic units and their function is represented at a level that mediates between form and function. The structuring of orthographic units that has been incorporated into distributed connectionist implementations to date has only been in terms of sub-syllabic units (i.e., onsets, vowels, and codas; Plaut et al., 1996), and not in terms of any morphemically based structures. However, such a possibility has indeed been proposed within an unimplemented connectionist account, though one that is localist in its description. This is the framework presented by Taft (2003, 2004, 2006).

According to this account, the orthographic units representing a polymorphemic word correspond to its individual morphemes, and there are functional units representing semantic and syntactic information. Importantly, though, there is an abstract level of representation that is intermediate between the form and function levels, referred to as 'lemma' units. The notion of a lemma as a link between form and function was originally introduced within the domain of speech production (see e.g., Levelt, Roelofs, & Meyer, 1999). When a spoken word is to be produced, activation is sent from the semantic representation to the lemma which, in turn, provides a link to both lexical syntax and the form of the word (in this case its pronunciation). When it comes to visual word recognition, the system works in the reverse direction: Activation passes from the form units (representing orthography in this case) to a lemma which, in turn, provides the link to semantic and syntactic information (see e.g., Baayen, Dijkstra, & Schreuder, 1997; Schreuder & Baayen, 1995; Taft, 2003, 2004, 2006).

As pointed out by Taft (2003, 2006), a lemma can actually be seen as being equivalent to a stable sub-pattern of activation (or 'attractor basin', cf. Plaut et al., 1996), meaning that it is a way of depicting the unit that develops as a result of a correlation between input and output (i.e., between form and function). Thus, a lemma potentially has the same characteristics as an attractor basin, but because a lemma is not 'hidden', it allows a more concrete portrayal of the lexical processing system for the purposes of conveying its content. This is exemplified by Figure 1 which depicts the lexical units involved in recognising the polymorphemic word *hunter*.

With a lemma being a unit that represents a consistent association between form and function, all monomorphemic words will be represented by a lemma, and this can be seen as their 'lexical entry'. If there is decomposition, an affixed word will be recognised through the lexical entry for its stem, which means that the prime *hunter* will facilitate responses to *hunt* because they activate the same lemma. When *corner* is presented, the lemma for *corn* will also be initially activated after the *er* is stripped off, but will subsequently be suppressed when it fails to activate information about the word *corner*.

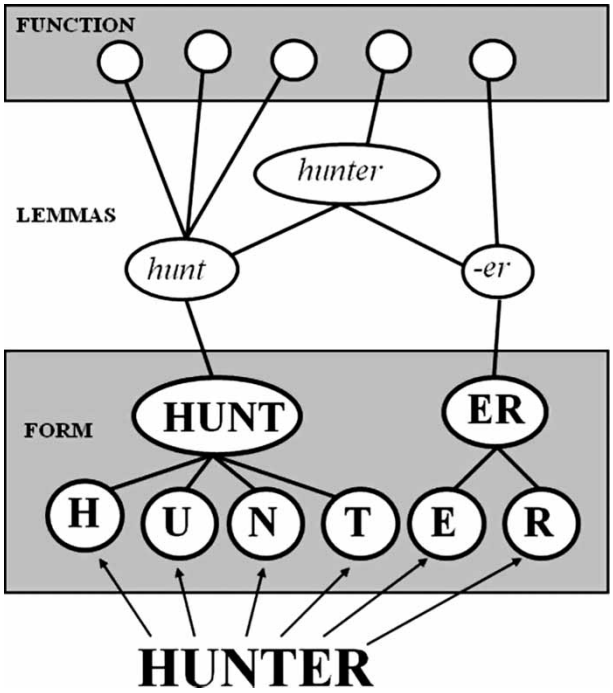

Figure 1. The representation of the polymorphemic word *hunter* in a model of the lexical processing system that incorporates decomposition and a lemma level mediating between form and function. Both *hunter* and its morphemes have their own lemmas, hierarchically organised.

Priming of *corn* by *corner* can therefore occur if such suppression has not yet been completed, that is, if the prime has not had enough time to be fully processed prior to the target appearing (as in the case of masked priming in the lexical decision task). Finally, *turnip* will not prime *turn* because there is no decomposition.

Note that a lemma must also exist for derived words themselves even when they are morphologically transparent, such as *hunter*. This is because such a lemma provides the links to functional information that is specific to the derived word and, hence, cannot be entirely determined from information linked to the lemmas of its component morphemes. For example, the lemma for *hunter* might provide links to information about the clothes that a hunter typically wears, which is something that is hard to capture through associations with the lemmas for *hunt* and *er* alone. So, given the need for a whole-word representation at the lemma level for derivationally affixed words, the issue of whether or not there is obligatory decomposition amounts to the question of whether this whole word representation is activated directly from its graphemes represented at the form level or whether its activation is mediated through the lemmas for its component morphemes. The former refers to direct whole-word access, while the latter entails morphological decomposition. The decomposition approach, as depicted in Figure 1, involves a hierarchy of lemmas whereby both the derived word (e.g., *hunter*) and its stem (*hunt*) have lemma representations, but the former is activated via the latter (e.g., Taft, 2003, 2006; Taft & Ardisinski, 2006; Taft & Kougious, 2004). The idea of a hierarchy of lemmas has also been put forward in the production domain by Sprenger, Levelt, and Kempen (2006) where a lexicalised language unit that is itself composed of words (e.g., an idiom) is represented by a 'superlemma' that is linked to the lemmas of the individual words.

If direct whole-word access occurs (i.e., without decomposition), the only link between the lemma for *hunter* and the lemma for *hunt* would be the shared orthographic units that activate them and the overlapping semantic units that are activated by them. The lack of *turnip-turn* priming suggests that facilitation does not arise solely from orthographic units below the level of the morpheme, so the priming of *hunt* by *hunter* would need to be explained in terms of a convergence of activation from their shared orthographic units and activation from their shared semantic units. However, the observed priming of *corn* by *corner* cannot be handled by a whole-word access account because the two words have no shared semantic units. The same is true if decomposition were to occur only after whole-word access (e.g., Giraudo & Grainger, 2000, 2003).

So, evidence from the masked priming paradigm seems to point to the importance of pre-lexical decomposition. What still needs to be clearly established, though, is whether such morphological processing extends

beyond the form level and, hence, whether the postulation of an intermediate level of representation (instantiated here as the lemma level) helps clarify the nature of the lexical processing system.

LEXICAL AMBIGUITY

The development of attractor basins within a distributed connectionist model has been used to explain how lexical ambiguity might be represented (e.g., Rodd, Gaskell, & Marslen-Wilson, 2002, 2004). When the same word-form is trained in association with more than one set of semantic features, the system stabilises into separate attractor basins for each meaning. With a lemma being the localist equivalent of an attractor basin, optimisation of the correlation between form and meaning similarly leads to different lemmas being developed for each meaning.

Take the word *stick*, which means both 'adhere' and 'twig'. Such a homograph represents two different words that happen to have the same form in English and, therefore, should be associated with two separate lemmas linked to their own particular meaning. It makes no sense to represent the two versions of *stick* separately at the form level because that is the very level where they are identical. Instead, it is logical to suggest that the same form-level representation feeds into two independent lemmas.

When polymorphemic words are derived from a homograph, they are often only related to one of its meanings. For example, the polymorphemic word *sticky* is derived only from the 'adhere' version of *stick*, which means that its 'polymorphemic' lemma will be linked only to the 'stem' lemma corresponding to that meaning. This is illustrated in Figure 2. It can be seen from this framework that, independent of any form-level priming between *sticky* and *stick*, the former should prime the 'adhere' meaning of the latter relative to the 'twig' meaning. Such a bias in meaning must happen at the lemma level and not the form level because the latter does not differentiate the two meanings. So, if such priming were observed it would support the existence of a level of representation above the form level; a level that captures morphological relationships.

THE PRESENT STUDY

The use of a lexical decision task within the masked priming paradigm does not help establish whether there is priming of only the relevant meaning of an ambiguous target. That is, lexical decision times to *stick* might be facilitated by the prior masked presentation of *sticky*, but such priming could arise from the shared orthographic unit *stick* (isolated after

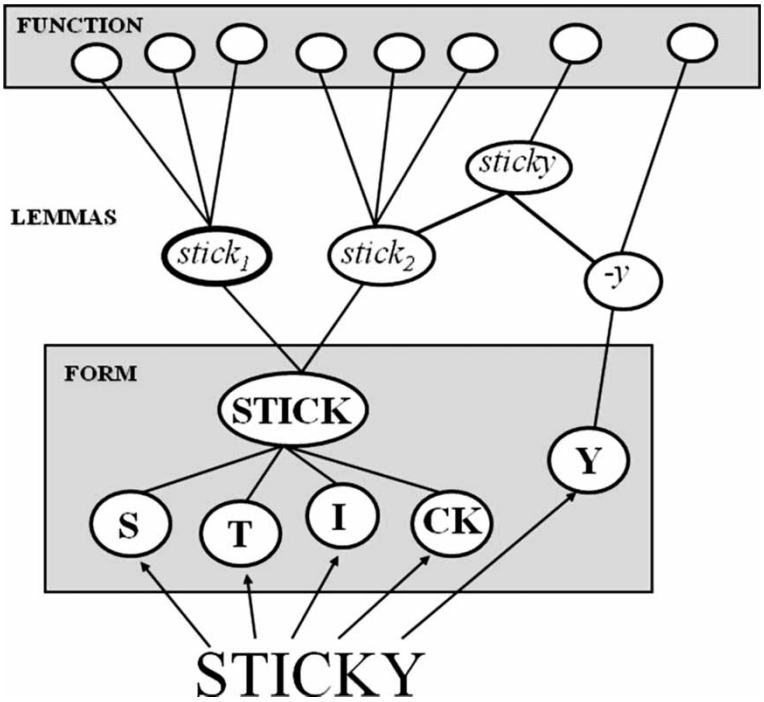

Figure 2. The relationship between *stick* and *sticky* in a model of the lexical processing system that incorporates decomposition and a lemma level mediating between form and function. The lemma 'stick$_1$' corresponds to the 'twig' meaning of *stick*, and the lemma 'stick$_2$' corresponds to the 'adhere' meaning. The former is depicted with a thicker outline to indicate that it is the more frequent meaning.

decomposition) rather than at the lemma level. In order to examine the priming of individual meanings of an ambiguous word, it is necessary to use a task that determines which meaning of the word the reader is thinking of. To this end, it is possible to simply ask the reader to use the word in a sentence or provide the first word or phrase they think of when seeing the target. If the 'adhere' meaning of *stick* is given more often when preceded by the masked prime *sticky* than when preceded by an unrelated word, this would indicate priming that does not arise merely at the form level. Whether it can be concluded that it arises at the lemma level, however, depends on whether it is possible to dispense with a possible alternative locus for the effect, namely, priming on the basis of shared semantic features alone.

The reason that *sticky* is morphologically related to the 'adhere' meaning of *stick* is because they have overlapping semantic features. This then means

that any priming from one to the other could arise purely at the function level without any involvement of lemmas. That is, presentation of *sticky* could activate the semantic features associated with the 'adhere' meaning of *stick* which, in turn, could bias the interpretation of *stick* toward that meaning. Therefore, in order to show that bias in the meaning of *stick* from the masked prime *sticky* arises explicitly at the intermediate lemma level, it must be shown that no such bias is obtained from a masked prime that is semantically, but not morphemically, related to the 'adhere' version of *stick*, such as *glue*. If the masked presentation of the prime *glue* were to have no impact on the meaning given to *stick*, it would show that pure semantic priming does not occur in this paradigm. This would mean that any bias generated by the masked prime *sticky* could not have arisen at the functional level of processing where meaning is located. Since the priming of a particular meaning of a homograph cannot be ascribed to the form level either, the existence of *sticky-stick* priming would clearly point to a level of representation that mediates between form and function, embodied here as the lemma level.

The present study therefore examines the bias generated by a masked morphemically related prime (e.g., *sticky-stick*) in comparison to a masked semantically related prime (e.g., *glue-stick*) in order to address the importance of a level of representation mediating between form and function when explaining morphological representation and processing. Furthermore, if such morphemically based priming is observed even though the prime is very briefly presented and masked from consciousness, it would provide added evidence that decomposition is early and automatic.

EXPERIMENT 1

Method

Materials. Target items were 30 ambiguous monomorphemic words (e.g., *stick*, *express*, *rail*). For each of these targets it was possible to generate an affixed word derived from only one of its meanings (e.g., *sticky*, *expression*, *railing*), and these affixed words comprised the primes for the 'Morphemic' condition. Most had a derivational suffix (including four with nominal -*ing*), but there were also seven cases of prefixation (e.g., *unjust*, *discontent*). Although an attempt was made to avoid items where the derived form was related to what was thought to be a highly dominant meaning of the ambiguous target, this was not always successful. For example, it turned out that the word *still* was almost always interpreted as meaning 'quiet' (rather than 'yet', its more common, but more abstract usage), which meant

that there was little scope for the Morphemic prime *stillness* to generate a bias toward that meaning.[1]

A further set of primes was set up for the 'Semantic' condition whereby the prime was semantically related to the relevant version of the ambiguous target, but was not morphemically related to it (e.g., *glue* as a prime for *stick*, *phrase* as a prime for *express*, and *fence* as a prime for *rail*). The Semantic and Morphemic primes were matched on their semantic similarity to the targets as determined by their Latent Semantic Analysis (LSA) scores (e.g., Landauer, Foltz, & Laham, 1998), with means of .22 and .20 respectively. In addition, the two types of prime did not differ on log word frequency as based on the CELEX norms combining spoken and written frequencies (Baayen, Piepenbrock, & van Rijn, 1993), $t(29) = 1.04$, with the trend favouring the Semantic primes.

The third condition had primes that were unrelated to their targets. These were constructed by re-pairing the Semantic primes with different targets (e.g., *phrase* being the unrelated prime for *stick*), with no participant seeing the same prime under both the Semantic and Unrelated conditions. This could be achieved through a Latin-Square design whereby three sub-lists of items were presented to three subgroups of participants. Within this design, each participant received ten items from each condition and, across the whole experiment, each target was presented under each condition without any participant seeing the same target more than once. All experimental stimuli can be found in the Appendix.

In addition to the 30 experimental items, a set of 20 filler words was included where the target was unambiguous. The primes in these cases were either morphemically related (e.g., *furry* as the prime for *fur*), semantically related (e.g., *danger* as the prime for *trouble*), or unrelated (e.g., *capital* being the prime for *fist*). Some of the primes were affixed versions of an ambiguous stem (e.g., *toaster*, *gravely*).

Procedure. Primes were presented for 50 ms in lower-case letters. They were each preceded by a 500 ms presentation of a row of hash marks, and were followed by the target in upper-case letters presented for 500 ms. Both the prime and target were presented in Arial font with a font size that allowed the target (18 point) to entirely overlay the prime (10 point) even

[1] For a number of the words, it also turned out that the subordinate meaning was so infrequent relative to the dominant meaning that it was almost never given (e.g., the 'safe' meaning of *sound*). Such cases would have had the potential to generate a morphemically based bias, but only if it were able to counteract the massive frequency disadvantage working against it. Often, however, the frequency imbalance was so great that the subordinate meaning was never thought of even when the prime was related to it (e.g., when the 'safe' meaning of *sound* was preceded by *unsound*).

though it typically had fewer letters. Items were presented using DMDX display software (Forster & Forster, 2003).

Participants were told that they would see a series of words in upper-case letters and they were to indicate the first meaning they thought of for that word either by using it in a sentence, by providing a related word or phrase, or by explicitly defining it. They were also told that some of the words might have more than one meaning and that they were to give their response to the first meaning that came to mind. No time limit was set. Responses were monitored by the experimenter through headphones and the meaning given by the participant for the ambiguous word was marked on a score-sheet and recorded for subsequent re-analysis if required. The experimenter was blind to the conditions of the experiment.

In order to establish that the prime was not available to conscious awareness, a pilot test was carried out ($n = 15$) where the task was lexical decision on the prime. For the purpose of this task, 30 additional ambiguous word targets were included. Each of these was preceded by a non-word prime, a third of which were orthographically similar to their target (e.g., *cressly* as the prime for *cross*) and two-thirds were not (e.g., *terbile* as the prime for *mean*). Performance on this task was no different to chance (51.17% errors), $t < 1$, with no sign of a difference between the three experimental conditions, $F < 1$. Hence, it is safe to conclude that the prime was successfully masked, and equally so across conditions.

Participants. Sixty undergraduate students at the University of New South Wales participated in the experiment, with 20 in each subgroup within the Latin-Square design. Course credit was given for participation.

Results

The vast majority of responses took the form of a sentence containing the target word or an explicit definition. Figure 3 presents the percentage of times that the target was given the meaning that corresponded to its Morphemic (and, hence, its Semantic) prime. A generalised linear mixed-effects model that utilised the logit function, suitable for analysis of the binominal dependent measure, was conducted in order to investigate both the Morphemic and Semantic priming effects (i.e., the fixed effects) and to identify the random factors moderating these effects (see Baayen, Davidson, & Bates, 2008). Models of different complexities were tested and it was revealed that the model with subjects and items as random effects provided the best fit to the data (log-likelihood ratio value $= -785.46$, $\chi^2(1) = 689.75$, $p = .00$). Thus, analysis of the fixed effects was conducted using this model. It was found that the primed

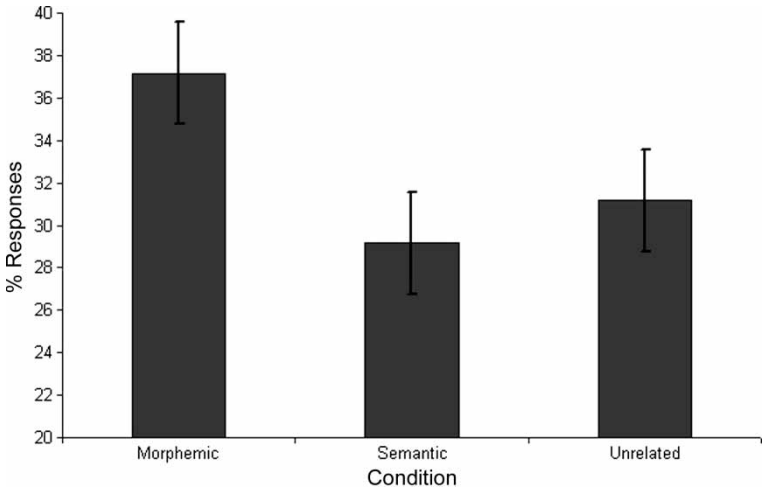

Figure 3. The percentage of responses given to the target that corresponded to the meaning of the morphemic/semantic prime.

target meaning was given significantly more often in the Morphemic condition than either the Unrelated or Semantic conditions, $t = 2.83$, $p < .01$ and $t = 3.78$, $p < .001$ respectively, while the Semantic and Unrelated conditions did not differ, $p = .334$.

As indicated earlier, the frequency of the Semantic primes was somewhat higher than that of the Morphemic primes. If it were to turn out, counter-intuitively, that masked priming was weaker for higher frequency primes, the difference between the conditions could potentially be explained. In the event, however, there was no correlation between the frequency of the prime and the magnitude of either semantic priming (Semantic minus Unrelated scores: $r = .11$) or morphemic priming (Morphemic minus Unrelated: $r = -.08$). Therefore, any mismatch in prime frequency that might have existed between the Morphemic and Semantic conditions cannot be used to explain the data.

Discussion

The results are very clear. The prominence of one of the meanings of an ambiguous target word can be increased by virtue of the prior masked presentation of a word derived from that meaning. No such change in the prominence of meaning takes place when the masked word is merely a semantic relative. It is therefore argued that morphological relationships must be captured in the lexical processing system by means of representations that are neither at the form nor semantic level. The priming that was

observed cannot be explained at the form level because the two versions of a homograph are identical at that level, and it cannot be explained at the semantic level because there was no semantic priming. The proposed level is intermediate between form and meaning, namely, a lemma level.

The critical lack of an effect in the Semantic condition might be seen as being consistent with the fact that masked semantic priming has been previously shown to be negligible when using a lexical decision task with a prime duration as brief as 50 ms (see e.g., Bueno & Frenck-Mestre, 2008; Perea & Gotor, 1997; Taft & Kougious, 2004). However, masked semantic priming does appear to emerge when the task is explicitly directed toward the meaning of the target, namely, semantic categorisation (e.g., Bueno & Frenck-Mestre, 2008) and, therefore, semantic priming might have been expected in the present experiment where the task was ostensibly focused on meaning. The fact that such priming was not found can be potentially explained in terms of an important difference between semantic categorisation and the present task where readers are asked to identify the meaning of a word. The former explicitly requires the manipulation of information at the semantic level and can only be performed on the basis of that information. In the present task, readers simply make use of information at the semantic level as a means of indicating which lemma they have activated. That is, the focus is actually on the lemma level rather than the semantic level, because manipulation of information at the semantic level is not required.

Lexical decisions can also be made on the basis of lemma status without the involvement of semantics. So, the lack of semantic priming in both lexical decision and the present meaning identification task implies that the processing of a masked prime goes only as far as the lemma level in a situation where there is no need to manipulate semantic information. In contrast, when the prime is not masked, it is apparent that activation does pass on to the semantic level given that semantic priming is typically revealed under such circumstances (e.g., Meyer & Schvaneveldt, 1971; Neely, 1991).

Before one can confidently draw conclusions about the present study, however, it needs to be determined whether the lack of priming in the Semantic condition arose simply because the semantic relationship between the primes and targets was not strong enough to reveal an effect. Although prime/target LSA scores were matched between the Morphemic and Semantic conditions, it is possible that semantic relatedness varied between the two conditions on the basis of something other than semantic distance. Experiment 2 was therefore carried out to establish whether bias in interpretation of the target could at least be obtained in the Semantic condition when the prime was no longer masked, that is, under conditions in which we know that semantic priming should be observed.

EXPERIMENT 2

Method

Materials and procedure. The Semantic and Unrelated items of Experiment 1 were used in Experiment 2, but were presented under unmasked conditions. This was achieved by eliminating the forward mask (i.e., the row of hash marks), and increasing the prime duration to 83 ms. The lower-case prime and upper-case target were presented in the same font size as each other (Arial 10).

The instructions were the same as in the first experiment, except that participants were additionally told that each trial included a lower-case and an upper-case word and that, although they should read the former, they should only make their response to the latter. Because each target was presented under two different conditions, the Latin-Square design required two subsets of items and two subgroups of participants.

Participants. There were 20 participants taken from the same pool as Experiment 1, 10 in each subgroup. None had participated in the first experiment.

Results and discussion

A clear bias toward the primed meaning was observed in this experiment for the Semantic condition, with mean scores of 39.7% and 28.7% for the Semantic and Unrelated conditions respectively. The generalised mixed-effects analysis (Baayen et al., 2008) revealed that the model with subjects and items as random effects provided the best fit to the data (log-likelihood ratio value $= -302.81$, $\chi^2(1) = 156.84$, $p < .001$). Using this model, it was found that the percentage of times the primed target meaning was given in the Semantic condition was significantly greater than in the Unrelated condition, $t = 3.51$, $p < .001$.

The observation of pure semantic priming under unmasked conditions contrasts dramatically with Experiment 1, where masking of the prime prevented such an effect (and even led to an inhibitory trend). For the comparison of the magnitude of semantic priming between Experiment 1 and Experiment 2, the model that provided the best fit had subjects and items as random effects (log-likelihood ratio value $= -801.89$, $\chi^2(1) = 627.56$, $p < .001$), with the interaction between the two experiments and the priming factor being significant, $t = 3.53$, $p < .001$.

It is apparent that the semantic relationship between the prime and the relevant interpretation of the target was sufficiently strong for the prime to generate a bias toward that interpretation. Therefore, the interpretation given to the masked Morphemic priming effect of Experiment 1 is reinforced.

Specifically, that effect must be explained at a level that does not represent semantic relationships (i.e., lower than the function level) because the semantically related primes produced no bias when masked. Given that a bias in meaning cannot be explained at the form level either, the apparent locus for the effect is a level that mediates between form and function.

GENERAL DISCUSSION

Most studies of morphological decomposition focus on the issue of whether a polymorphemic word is recognised via a representation of its stem or directly on the basis of the whole word. Lexical decision responses obtained within the masked priming paradigm suggest that early and automatic decomposition does occur, but the literature says little about the instantiation of that decomposition in the lexical processing system. The assumption is typically made that the stem representation that is activated through decomposition is a representation of its form. That is, the locus of morphological representation and processing in visual word recognition is orthographic. The present research adds a further dimension to our understanding of morphological processing by showing that the meaning given to an ambiguous word can be biased by the prior masked presentation of a word derived from that meaning. Such a bias cannot be explained at the orthographic level because that level does not differentiate between the meanings of a homograph, and neither can it be explained at the semantic level because semantic relatedness alone did not generate a bias in meaning. Therefore, it is argued that a level that mediates between form and function exists, and that this is the locus of morphological relationships.

Within the framework presented here, the intermediate level takes the form of lemmas which are a localist characterisation of the stable patterns of activation (i.e., attractor basins) generated within a distributed network (e.g., Plaut et al., 1996; Rodd et al., 2004). As depicted in Figures 1 and 2, decomposition at the form level leads to activation of the lemma for the whole derived word via the lemma for its stem. Although there needs to be a representation of the whole word somewhere in the system, its locus is the lemma level rather than the form level according to this account.

The difference between the lemma representation of an affixed word (e.g., *hunter*) and a pseudo-affixed word (e.g., *corner*) can be seen when comparing Figure 1 with Figure 4. The former shows how the lemma for *hunter* is activated via the lemma for both of its constituent morphemes (i.e., *hunt* and *er*). In contrast, Figure 4 illustrates how the lemma for *corner* is activated directly from the orthographic units rather than via the lemma units representing its pseudo-morphemic constituents. At the lemma level, the units for *corner* and *corn* would then compete with each other, with the

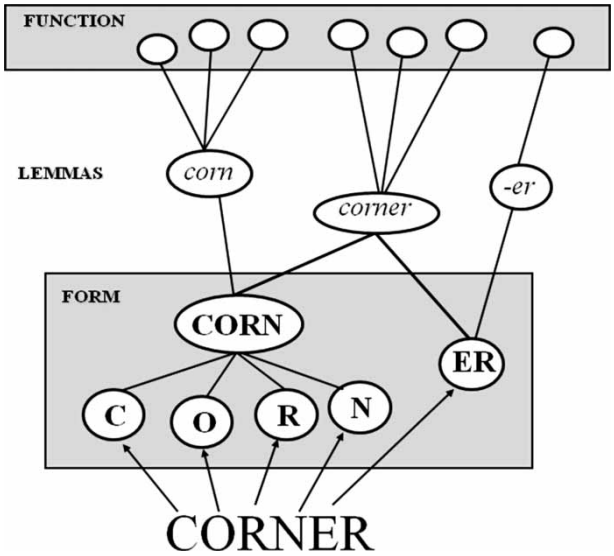

Figure 4. The representation of the pseudo-affixed word *corner* within the lemma framework. The lemma for *corner* is activated through form units only.

former winning out because of its greater activation through the *er* unit. Although the figure shows the activation of the lemma for *corner* being mediated via the orthographic unit for *corn*, it could alternatively be activated directly from the individual graphemes. Either way, it needs to be explained how *corner* is able to prime lexical decision responses to its pseudo-stem (*corn*) in the masked priming paradigm.

Facilitation might arise purely from the form representation that is used for both *corner* and *corn*. However, because the lemma for *corner* wins the competition against the lemma for *corn*, it needs to be explained why the latter is not inhibited to the extent that such facilitation is lost. It cannot be that masking of the prime prevents activation passing up to the lemma level, because the selective priming of one meaning of *stick* by *sticky* shows that masking does allow activation to proceed to a level higher than that of form.

One possibility is that, although the lemma for *corn* and *corner* vie with each other to reach a recognition threshold, such competition does not actually inhibit the activation level of the 'loser'. That is, all that matters is which lemma reaches threshold first. As such, when the lemma for *corner* reaches the threshold ahead of that for *corn*, the latter remains fully pre-activated and the masked priming of *corn* by *corner* is equivalent in strength to the masked priming of *hunt* by *hunter*. Such equivalent priming is the conclusion drawn by Rastle and Davis (2008) from their examination of all

published experiments that have compared derived word and pseudo-derived word priming with a short prime duration.

Alternatively, it could be argued that inter-lemma competition normally leads to inhibition of the losing lemma, but with such a short duration before the target is presented that the competition set up by the prime is not resolved and, hence, the loser (i.e., the lemma for the target) is protected from being fully inhibited. The degree to which such inter-lemma inhibition is prevented will determine the amount of pseudo-derived word priming relative to derived word priming. In relation to this, it should be noted that when a statistical meta-analysis is performed on the experiments listed by Rastle and Davis (2008) that have directly compared derived and pseudo-derived word priming, a significantly larger effect for the former is actually observed, $t(13) = 2.32$, $p < .05$. So, there is some evidence to suggest that presentation of a pseudo-derived word (e.g., *corner*) does partly inhibit the word corresponding to its pseudo-stem (*corn*) relative to the priming generated between a truly derived word (e.g., *hunter*) and its stem (*hunt*).

A further question that can be asked is why a non-affixed word like *turnip* fails to facilitate responses to the word contained at its beginning (i.e., *turn*). The argument previously put forward was that the non-morphemic status of *ip* prevents it from being isolated from *turn*. However, within the framework provided in Figure 4, it should be the case that the graphemes *t*, *u*, *r*, and *n* activate the form unit that exists for the word *turn*, along with its lemma, even if the graphemes *i* and *p* do not combine to activate a form unit. Why then should activation of the lemma for *corn* produce priming when *corner* is the prime while activation of the lemma for *turn* does not produce priming when *turnip* is the prime? There are several arguments that can be made.

First, the potential internal structuring of *corner* as *corn* plus *er* might mean that its lemma is activated via form-level units corresponding to those structures (as shown in Figure 4) whereas the lack of internal structure for *turnip* means that its lemma is activated directly from its grapheme units. Because the form and lemma units for the word *turn* should still be activated when *turnip* is presented, it would then need to be argued that such activation is quickly inhibited when the rest of the letter-string (i.e., *ip*) does not correspond to any form unit, unlike the *er* of *corner*. Such inhibition is so rapid that masking of the prime does not prevent it from happening.

In contrast to the above suggestion, Taft (1979, 1987, 2001) proposes that all polymorphemic words are represented by a form unit, called the Basic Orthographic Syllabic Structure (BOSS), that corresponds to the first syllable, but with a maximised coda. If so, there would be no difference in the form-level representation of the *turn* of *turnip* and the *corn* of *corner* because both constitute the BOSS. Instead, the difference between the representation of the two words would lie in the fact that *er* has a corresponding lemma while *ip*

does not, and it would have to be this fact that leads to the rapid inhibition of the representation of *turn* that prevents it being primed by *turnip*.

There is, however, a further account of *turnip-turn* priming that warrants further investigation. This is the possibility that such priming actually exists when the target corresponds to the BOSS of the prime. The non-affixed items that have been used in the previous masked priming studies (e.g., Diependaele et al., 2005; Lavric et al., 2007; Longtin et al., 2003; Marslen-Wilson et al., 2008; McCormick et al., 2008; Rastle et al., 2004) have included a mixture of cases where the target is the BOSS of the prime (e.g., *turnip-turn*, *brothel-broth*) and where it is not (e.g., *shunt-shun*, *freeze-free*). We do not know, then, whether facilitation occurs if only the former are used as items, which means that the facts about masked priming of sublexical units are actually incomplete.

Rather than representing a BOSS, another possibility that has been suggested is that orthographic units develop on the basis of relative bigram frequencies (e.g., Rastle et al., 2004; Seidenberg, 1987; though see Rapp, 1992). For example, the bigram *er* might be processed as a unit in *corner* because of its common occurrence in the language relative to the bigram *ne*, whereas *turnip* is not processed as *turn* + *ip* because *ip* is a less common bigram than *ni*. Such structuring of form units is compatible with the distributed connectionist approach because it is determined purely on the basis of the statistical characteristics of letter combinations, though how this might be specifically instantiated in a distributed model has never been made clear.

Finally, it can be emphasised that morphological decomposition does not appear to be an optional method of processing a polymorphemic word under strategic control. The priming that was observed in the Morphemic condition of the present study occurred even though the prime was masked from conscious awareness. As has been concluded from previous studies of masked morphemic priming, decomposition appears to be an automatic component of the processing of visually presented polymorphemic words.

Manuscript received February 2009
Revised manuscript received May 2009
First published online June 2009

REFERENCES

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*, 390–412.

Baayen, R. H., Dijkstra, T., & Schreuder, R. (1997). Singulars and plurals in Dutch: Evidence for a parallel dual route model. *Journal of Memory and Language*, *37*, 94–117.

Baayen, R. H., Piepenbrock, R., & van Rijn, H. (1993). *The CELEX Lexical Database* (*CD-ROM*). Philadelphia, PA: Linguistic Data Consortium, University of Pennsylvania.

Baayen, R. H., Wurm, L. H., & Aycock, J. (2007). Lexical dynamics for low frequency complex words: A regression study across tasks and modalities. *The Mental Lexicon*, *2*, 419–463.

Bueno, S., & Frenck-Mestre, C. (2008). The activation of semantic memory: Effects of prime exposure, prime-target relationship and task demands. *Memory and Cognition*, *36*, 882–898.

Diependaele, K., Sandra, D., & Grainger, J. (2005). Masked cross-modal morphological priming: Unravelling morpho-orthographic and morpho-semantic influences in early word recognition. *Language and Cognitive Processes*, *20*, 75–114.

Forster, K. I., & Forster, J. C. (2003). A Windows display program with millisecond accuracy. *Behavior Research Methods Instruments and Computers*, *35*(1), 116–124.

Giraudo, H., & Grainger, J. (2000). Prime word frequency in masked morphological and orthographic priming. *Language and Cognitive Processes*, *15*, 421–444.

Giraudo, H., & Grainger, J. (2003). A supralexical model for French derivational morphology. In E. Assink & D. Sandra (Eds.), *Reading complex words* (pp. 139–157). Amsterdam: Kluwer.

Landauer, T. K., Foltz, P. W., & Laham, D. (1998). Introduction to Latent Semantic Analysis. *Discourse Processes*, *25*, 259–284.

Lavric, A., Clapp, A., & Rastle, K. (2007). ERP evidence of morphological analysis from orthography: A masked priming study. *Journal of Cognitive Neuroscience*, *19*, 866–877.

Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, *22*, 1–38.

Longtin, C. M., Segui, J., & Halle, P. A. (2003). Morphological priming without morphological relationship. *Language and Cognitive Processes*, *18*, 313–334.

Marslen-Wilson, W., Bozic, M., & Randall, B. (2008). Early decomposition in visual word recognition: Dissociating morphology, form, and meaning. *Language and Cognitive Processes*, *23*, 394–421.

McCormick, S. F., Rastle, K., & Davis, M. H. (2008). Is there a 'fete' in 'fetish'? Effects of orthographic opacity on morpho-orthographic segmentation in visual word recognition. *Journal of Memory and Language*, *58*, 307–326.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*, 227–234.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Perea, M., & Gotor, A. (1997). Associative and semantic priming effects occur at very short stimulus-onset asynchronies in lexical decision and naming. *Cognition*, *62*, 223–240.

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

Plaut, D. C., McClelland, J. L., Seidenberg, M. S., & Patterson, K. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, *103*, 56–115.

Rapp, B. C. (1992). The nature of sublexical orthographic organization: The bigram trough hypothesis examined. *Journal of Memory and Language*, *31*, 33–53.

Rastle, K., & Davis, M. H. (2008). Morphological decomposition based on the analysis of orthography. *Language and Cognitive Processes*, *23*, 942–971.

Rastle, K., Davis, M. H., & New, B. (2004). The broth in my brother's brothel: Morpho-orthographic segmentation in visual word recognition. *Psychonomic Bulletin and Review*, *11*, 1090–1098.

Raveh, M. (2002). The contribution of frequency and semantic similarity to morphological processing. *Brain and Language*, *81*, 312–325.

Rodd, J. M., Gaskell, M. G., & Marslen-Wilson, W. D. (2002). Making sense of semantic ambiguity: Semantic competition in lexical access. *Journal of Memory and Language*, *46*, 245–266.

Rodd, J. M., Gaskell, M. G., & Marslen-Wilson, W. D. (2004). Modelling the effects of semantic ambiguity in word recognition. *Cognitive Science*, *28*, 89–104.

Rueckl, J. G., & Raveh, M. (1999). The influence of morphological regularities on the dynamics of a connectionist network. *Brain and Language*, *68*, 110–117.

Schreuder, R., & Baayen, R. H. (1995). Modeling morphological processing. In L. B. Feldman (Ed.), *Morphological aspects of language processing* (pp. 131–154). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Seidenberg, M. S. (1987). Sublexical structures in visual word recognition: Access units or orthographic redundancy? In M. Coltheart (Ed.), *Attention and Performance, XII*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Seidenberg, M. S., & Gonnerman, L. M. (2000). Explaining derivational morphology as the convergence of codes. *Trends in Cognitive Sciences*, *4*, 353–361.

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

Sprenger, S. A., Levelt, W. J. M., & Kempen, G. (2006). Lexical access during the production of idiomatic phrases. *Journal of Memory and Language*, *54*, 161–184.

Taft, M. (1979). Lexical access via an orthographic code: The Basic Orthographic Syllabic Structure (BOSS). *Journal of Verbal Learning and Verbal Behavior*, *18*, 21–39.

Taft, M. (1987). Morphographic processing. The BOSS re-emerges. In M. Coltheart (Ed.), *Attention and performance, XII*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Taft, M. (2001). Processing of orthographic structure by adults of different reading ability. *Language and Speech*, *44*, 351–376.

Taft, M. (2003). Morphological representation as a correlation between form and meaning. In E. Assink & D. Sandra (Eds.), *Reading complex words* (pp. 113–137). Amsterdam: Kluwer.

Taft, M. (2004). Morphological decomposition and the reverse base frequency effect. *Quarterly Journal of Experimental Psychology*, *57A*, 745–765.

Taft, M. (2006). A localist-cum-distributed (LCD) framework for lexical processing. In S. M. Andrews (Ed.), *From inkmarks to ideas: Current issues in lexical processing* (pp. 76–94). Hove, UK: Psychology Press.

Taft, M., & Ardisinski, S. (2006). Obligatory decomposition in reading prefixed words. *The Mental Lexicon*, *1*, 183–199.

Taft, M., & Kougious, P. (2004). The processing of morpheme-like units in monomorphemic words. *Brain and Language*, *90*, 9–16.

APPENDIX

The following are the items used in the experiments. The upper-case target is followed by its lower-case Morphemic, Semantic, and Unrelated prime. The number following each prime is the percentage of times that the meaning of the Morphemic/Semantic prime was given. For the Semantic and Unrelated primes, the two scores refer to Experiment 1 (masked) and Experiment 2 (unmasked) respectively.

STATE	statement	5	assert	10	10	crust	0	30
TEAR	tearful	60	weeping	50	80	filter	70	40
LEAN	leanest	65	skinny	50	90	glue	55	20
TASTE	tasteful	0	elegance	5	10	arrange	0	0
COURT	courtship	15	serenade	5	10	exit	0	0
SUIT	suitable	10	befit	5	0	cheat	25	10
STAR	starlet	0	actress	0	20	safe	15	0
STALK	stalker	80	pursue	80	80	satisfied	80	80
BOWL	bowling	50	tenpin	25	40	yearn	30	40
RACE	racial	20	ethnic	10	40	tutor	15	0
TRAIN	trainer	10	tutor	5	20	fence	5	0
PART	depart	10	exit	5	0	first	10	0
STRAIN	strainer	0	filter	5	10	reservation	0	10
ORDER	disorder	60	arrange	65	70	wary	65	30
STICK	sticky	45	glue	20	10	phrase	15	30
CONTENT	discontent	55	satisfied	65	60	silence	65	40
SCALE	scaly	5	crust	10	10	unfair	0	10
LIE	liar	100	cheat	85	100	wrestler	95	100
SOUND	unsound	0	safe	0	10	envelope	0	0
LONG	longing	0	yearn	0	20	verify	0	0
BOX	boxer	5	wrestler	10	10	skinny	5	0
SECOND	secondly	65	first	40	40	assert	35	40
RAIL	railing	65	fence	30	70	befit	40	40
SEAL	unseal	35	envelope	30	30	pursue	25	50
JUST	unjust	60	unfair	30	80	tenpin	30	70
STILL	stillness	90	silence	85	80	weeping	85	100
WATCH	watchful	55	wary	35	40	actress	15	20
EXPRESS	expression	60	phrase	35	50	serenade	65	50
BOOK	booking	5	reservation	0	0	elegance	0	10
FIRM	confirm	85	verify	80	100	ethnic	90	40