A supralexical model for French derivational morphology

Hélène Giraudo, Jonathan Grainger

A SUPRALEXICAL MODEL FOR FRENCH DERIVATIONAL MORPHOLOGY

1. INTRODUCTION

Today there is a growing consensus in the psycholinguistic research community that morphological information plays a critical role in the highly automatized process of word recognition. At a linguistic level of analysis, morphology describes the word formation rules of a given language. A morphologically complex word, such as *banker*, represents the combination of at least two morphemes: the root *bank* and the suffix *–er* in the given example. As a consequence, from each root morpheme (e.g., *bank*) one can derive numerous morphologically complex words by adding another morpheme to the root (i.e., *-er*, *-ing*, *-s*, *-rupt…* to form the words *banker, banking, banks, bankrupt…etc.*). From a linguistic point of view therefore, the accent is often placed on productivity when discussing the role of morphological information in language processing. Morphemic components (roots and affixes) can be used to create new word forms in production and to understand novel forms in comprehension. For example, someone who had never heard or read the word *dimness* before, but has knowledge of the root *dim* and the suffix *-ness*, could derive the meaning of the novel whole-word from the meaning and function of the component morphemes.

The fact that morphology is the basis of language productivity at the word level, has lead the vast majority of researchers interested in word recognition (cognitive psychologists and psycholinguists) to propose what we will refer to as the sublexical account of morphological representation. Within a general processing framework for word recognition that moves from sensory codes to semantic representations, morphological codes are hypothesized to intervene before whole-word form representations. Two variants of this hypothesis, a single-route and a dual-route version, are presented in Figure 1.

Figure 1 about here

In both variants of the sublexical account of morphological representation, depicted in Figure 1, morphemes are directly activated from sublexical representations of the stimulus (e.g., letters or phonemes), and are therefore form-based representations (i.e., they code the presence of a given string of letters or a given string of phonemes in the stimulus). In this approach, it is phonologically and orthographically specified stems and not abstract roots that are the morphological units of representation, along with affixes. The various authors of these sublexical accounts often speak of a

morphological parsing mechanism that automatically extracts morphemes from the stimulus without reference to whole-word representations. In the single route account (e.g., Taft, 1994), a given word stimulus is first parsed into its morphological components before the word can be recognized as a whole. In the dual-route account (e.g., Caramazza, Laudana, & Romani, 1988; Schreuder & Baayen, 1995), there is both a direct whole-word route from sensory representations to meaning, and a morphological parsing route that extracts morphemes from the incoming signal before assigning meaning to these components.

The great majority of current research on morphological processing provides an interpretation of observed empirical data within a variant of this sublexical framework. However, in this chapter we present what we believe to be a viable alternative for the representation of derivational morphology in French, and we examine some recent data in favour of this alternative hypothesis. The alternative account, referred to as the supralexical account of morphological representation, stipulates that in normal, fast, automatic word recognition, morphemic representations present in derived complex forms are only contacted after some initial processing of whole-word form representations.

Common to both the sublexical and supralexical accounts of morphological representation is the assumption that morphology is explicitly represented in long-term memory, over and above the representation of form and meaning. So first let us summarise what we believe to be the critical empirical data in favour of the representation and use of morphological information during word recognition.

2. EVIDENCE FOR THE ROLE OF MORPHOLOGY IN PRINTED WORD PERCEPTION.

Over the last two decades, a number of studies have been performed in order to examine morphological influences on word recognition using various experimental techniques tested in different languages. These studies can be classified according to the type of manipulation: priming studies (including long-lag priming, and the more recent masked priming research); morpheme frequency effects (including manipulations of cumulative root frequency, and the more recent variable referred to as morphological family size, Schreuder & Baayen, 1997); and pseudo-morpheme manipulations (most predominantly pseudo-prefixation). Some of the most critical evidence, in our opinion, has been obtained using the masked priming paradigm which has gradually been adopted in recent years as an important tool for uncovering the fast and highly automatized processes that underlie printed word perception. Masked priming combined with the lexical decision task (Forster & Davis, 1984) has helped define the relative role played by orthographic, phonologic, morphologic, and semantic

information in this process. Using prime exposures of around 60 ms, it has been shown that, while morphologically related primes facilitate processing relative to unrelated word primes (e.g., Deutsch, Frost & Forster, 1998; Drews & Zwitserlood, 1995; Forster, Davis, Shocknecht, & Carter, 1987; Frost, Forster & Deutsch, 1997; Grainger, Colé & Segui, 1991), orthographically related primes tend to inhibit target word processing, generating longer response times (RTs) and/or more errors compared to an unrelated prime condition (Drews & Zwitserlood, 1995; Grainger & Ferrand, 1994; Segui & Grainger, 1990). Some recent work has also demonstrated that masked morphological priming is not just a combination of effects of form and meaning. For example, Rastle, Davis, Marslen-Wilson, and Tyler (2000) tested priming with "portmanteau" words (e.g., smog - smoke) and phonaesthemes (e.g., glitter − glisten) where there is both form and meaning overlap in the absence of morphological overlap. At 72 ms prime exposures, these stimuli produced no significant priming effect, while morphological priming was highly robust. Furthermore, in a study of morphological priming in Hebrew (Frost, Deutsch, Gilboa, Tannenbaum, & Marslen-Wilson, 2000), facilitatory priming was obtained with morphologically related but semantically unrelated words. Finally, research carried out by the present authors using the masked priming technique has systematically failed to find any evidence for semantic priming in conditions (60 ms prime exposures with the lexical decision task) where robust morphological priming is obtained. Recent research therefore allows us to reject a purely semantic account of morphological priming.

Thus, there is currently an impressive amount of evidence from masked priming studies in favour of the explicit representation of morphological information in memory. However, there is still little agreement as to precisely how such information is stored in long-term memory. In this chapter, we discuss two main alternatives (the sublexical and supralexical hypotheses) that differ in terms of where, in the processing hierarchy from form to meaning, morphological codes are situated.

The results of all of the above-cited studies using the masked prime paradigm are easily accommodated by the sublexical account of morphological representation. When primes share the same morphological stem as targets, target recognition is facilitated by preactivation of the stem representation during prime processing. However, the supralexical account of morphological representation, to be described in more detail below, can also account for morphological priming effects via the preactivation of morphemic representations. We first present the supralexical hypothesis, before describing some critical tests of these alternative accounts of morphological representation.

3. THE SUPRALEXICAL HYPOTHESIS.

The supralexical hypothesis is once again expressed within the framework of a hierarchical activation model of language comprehension that includes whole-word form representations and morphemic representations that intervene between lower-level form analyses and higher-level semantic representations (Figure 2). According to the supralexical account, morphemic representations receive activation from whole-word form representations and act as a partial interface between form-representations and semantics. Within this framework, morphemic representations are form- and modality-independent. Accordingly, we refer to root rather than stem representations, to reflect their abstract nature. All whole-word forms that share a given root are connected to that root representation, and activated root representations send activation back to all compatible whole-word forms. On presentation of a complex word such as remake, units corresponding to the root make and the affix re will receive activation from the whole-word representation, and send back activation to all whole-word representations that are compatible with either the root or the affix. In this way, root representations impose an organization on the lower-level form representations in terms of so-called morphological families (Butterworth, 1983; Lukatela, Gligorijevic, Kostic, & Turvey, 1980; Manelis & Tharp, 1977).

Insert Figure 2 about here

Both the sublexical and the supralexical accounts of morphological representation propose morphemic representations as the locus of masked morphological priming effects. Thus, all words that share a given morpheme are hypothesized to be linked to a common representation that corresponds to that morpheme. Prime stimuli that contain the same morpheme as target stimuli (i.e., morphologically related primes) generate facilitatory priming via preactivation of the shared morphemic representation. Prime presentation leads to growth of activation in a given morphemic representation thus conferring an advantage in the processing of target stimuli that contain the same morpheme. Here we assume that the size of morphological priming effects are proportional to the activation level of the critical morpheme after prime word presentation.

4. PRIMING WITH SIMPLE AND COMPLEX WORDS

Although, both of the sublexical and the supralexical accounts of morphological representation propose morphemic representations as the locus of masked morphological priming effects, they differ in terms of where the morphemic codes are located relative to whole-word representations. The precise location of these units has a direct implication on how

morphologically complex words are accessed relative to simple words. If a complex word is first parsed into its constituent morphemes (the sublexical hypothesis), then in the masked prime paradigm with brief prime exposures, root primes should be more effective than derived word primes, since the former do not require parsing. Here the basic assumption is that there is some extra computation involved in isolating a root in a derived word compared to recognizing a root presented on its own. This extra computation should slow down the processing of derived word primes leading to less priming. On the other hand, according to the supralexical hypothesis, derived word primes should be just as effective as root primes. Both types of prime will lead to similar levels of preactivation of supralexical morphological units. The bi-directional excitatory links between a supralexical representation of the root and all words containing that root will facilitate the recognition of a target word sharing the same root as the prime. More precisely, the sublexical hypothesis predicts that since parsing a derived word involves isolating the root, and then checking that the remaining letters form a compatible affix, this should involve more computation than when a root is presented on its own (i.e., as a free root without an accompanying affix). The extra computation should delay the appearance of priming effects obtained from derived word primes compared to root primes. The predictions of this particular version of the sublexical hypothesis are summarised at the top of Table 1.

Insert Table 1 about here

Alternatively, the sublexical view could stipulate that morphemic representations are activated as soon as an appropriate string of letters is detected in the stimulus, regardless of the status of the remaining letters (i.e., morpheme or not). This is an example of obligatory morphological decomposition where, at least in initial stages of processing, the system is blind to the compatibility of morphological constituents. In this way, there is no extra computation involved in processing a derived word prime, since no attempt will be made to check if the remaining letters form an appropriate affix. If this is how the morphological parser operates, then the sublexical hypothesis is led to predict that pseudoroot primes (word primes containing a string of letters corresponding to a root without having the status of a root in that particular word) should be just as effective as true derived word primes. The supralexical hypothesis, on the other hand, predicts a significant advantage for derived word primes compared to pseudoroot primes. These predictions are summarized in Table 1.

These predictions were tested in two series of masked priming experiments. In the first series, we examined morphological priming effects on lexical decision latencies to simple and complex word targets using primes that were either free roots or suffixed words (Giraudo & Grainger,

2001). One critical result obtained in these experiments is presented in the left-hand panel of Figure 3. Here we see that responses to suffixed word targets were facilitated (relative to form-related control primes) by both free root primes and suffixed word primes sharing the same root as the target. The free root and suffixed word prime conditions did not differ significantly. It should be noted that the unrelated prime condition (not shown in the figure) was not significantly different from the orthographic control prime condition (mean RT = 798 ms compared to 795 ms for the othographic condition). Finally, in another experiment in this series, responses to free root targets were significantly faster when preceded by related suffixed word primes (608 ms) than unrelated primes (635 ms).

Our observation of statistically equivalent priming for free root and derived word primes falsifies the first version of the sublexical morphological parser described above. According to this version, root primes should have been more effective than derived word primes, since the derived word primes require additional computation to check for the affix status of the remaining letters once the root has been identified. This therefore suggests that, if there is some form of sublexical morphological parsing mechanism, then it operates just as efficiently on complex and simplex stimuli. As noted above, this would be the case if the parsing mechanism were blind to surrounding context. In other words, strings of letters that correspond to roots would be extracted from the stimulus independently of the remaining letters in the string. In such a system, a sublexical root representation will be activated whenever a sequence of letters corresponding to the orthographic description of that root is activated. This alternative version of the sublexical hypothesis was tested in a further experiment comparing masked priming from true derived word primes and pseudo-root primes.

Figure 3 about here

The results of the second experiment are presented in the right-hand panel of Figure 3. It can be seen that only true derived word primes (e.g., laitage-laitier), were found to produce facilitation, while morphologically simple primes containing a pseudo-root (e.g., laitue-laitier) did not influence performance relative to the unrelated prime condition. These results falsify the alternative version of the sublexical hypothesis presented above. According to this version, a subset of letters corresponding to a potential root should activate the corresponding morphemic representation irrespective of the status of the other letters in the string. This should have produced significant morphological priming from primes containing such sequences, as in the pseudo-root primes tested in the above experiment.

5. EFFECTS OF PRIME WORD FREQUENCY AND CUMULATIVE ROOT FREQUENCY IN MASKED PRIMING

Another prediction on which the two proposed architectures for morphological representation differ, concerns the effects of prime word frequency and cumulative root frequency in masked priming. In a supralexical architecture, morphological facilitation is hypothesized to arise via shared morpheme units located above the level of whole-word form representations. When primes and targets share the same root, the activation of the root representation by the prime stimulus facilitates processing of the target word via excitatory connections between morphemic representations and appropriate whole-word representations.

Thus, according to the above account, a morphologically related prime generates two conflicting effects: within-level inhibition acting across simultaneously activated word forms, and top-down facilitation from shared morphemic representations. Both of these mechanisms are fundamental to the interactive-activation family of models of word recognition (McClelland & Rumelhart, 1981). On presentation of a given stimulus, all whole-word representations that share lower-level form representations (e.g., letters) with the stimulus, receive bottom-up activation. These whole-word representations that receive support from the stimulus then enter a competitive process that allows only a single representation to dominate processing via inhibitory connections across all units at that level. In the masked priming paradigm, form-related primes (e.g., blue-blur) can generate inhibition compared to unrelated primes (Segui & Grainger, 1990) by increasing the competitivity of the target word's natural competitors (i.e., those that receive the most bottom-up support from the stimulus, apart from the target itself).

Now, since form overlap across morphologically related items is generally much lower than in situations where form-related inhibitory priming is observed (Segui & Grainger, 1990), this allows the facilitatory morphological component to dominate in masked morphological priming (Grainger et al., 1991). Keeping target frequency as high as possible will further help morphological facilitation to dominate over form-level inhibition. The higher the surface frequency of the target word, the greater its resistance to within-level inhibition. This first prediction was tested in a series of masked prime experiments where the printed frequency of morphologically complex prime words (their surface frequency) was varied while maintaining both target frequency and cumulative root frequency (the summed frequency of all derived words sharing the same root) constant (Giraudo & Grainger, 2000).

According to the supralexical model of morphological representation, more frequent morphologically related primes should enhance effects of morphological priming. In this model, morpheme unit

activation is a function of the activation of all whole-word representations that share that morpheme. During prime word processing under the brief temporal conditions of masked priming, the only significantly activated word unit is the prime word itself, whose activation level is a function of its surface frequency (McClelland & Rumelhart, 1981). Thus, the activation level of supralexical morphemic representations will be a function of the surface frequency of the prime word. Grainger et al. (1991) found that prime word frequency (medium or low) did not interact with effects of morphological priming. However, target word frequency varied with prime frequency in this experiment (medium frequency targets were primed with low frequency words, while low frequency targets were primed by medium frequency words).

Figure 4 about here

The results of a series of masked prime experiments (Giraudo & Grainger, 2000) are presented in Figure 4. The results show that prime-word surface frequency influences the size of masked morphologically priming effects obtained with free root targets in the lexical decision task. On the other hand, the orthographic prime condition did not differ significantly from the unrelated prime condition in either the RT or the error analyses, and prime word frequency did not modulate this absence of priming effects. These results fit with the predictions of the supralexical hypothesis.

The results of another experiment in this series showed no influence of the cumulative root frequency of complex prime words on the size of masked morphological priming. Free root targets were primed equally well by suffixed word primes with either a high or a low cumulative root frequency. The predictions of the sublexical and supralexical architectures relative to effects of cumulative root frequency, will be examined within an activation framework where connections strengths are a function of the frequency with which the two connected units have indeed been co-activated in the past. The connections strengths between letter representations and morpheme representations in a sublexical architecture will be a function of cumulative root frequency (indexing the number of times that a given root appears independently of surface forms). On other hand, in a supralexical architecture, the connections strengths between whole-word representations and morphemic representations are a function of surface frequency. In the experiment run by Giraudo and Grainger (2000), priming effects compared to unrelated primes were 38ms for low cumulative frequency primes and 46ms for high cumulative frequency primes. Needless to say, there was no interaction between cumulative frequency and priming effects.

Taken together, the results of this study fit with the predictions of a supralexical architecture where morphemic units are situated above whole-word representations such that morphemic representations only receive activation from low-level stimulus analysis via whole-word representations. As a consequence, the surface frequency of morphological primes affects the size of morphological priming and the effects of morphology necessarily depend on processing of a whole-word form that gives rise to effects of surface frequency.

Looked at from another angle, the results of this study are difficult to reconcile with a dual-route version of the sublexical hypothesis (e.g., Caramazza, et al., 1988; Schreuder & Baayen, 1995). According to this variant of the sublexical approach, the morphological parsing route mainly influences the processing of low frequency words, since high frequency words will be rapidly processed by the direct (non-decompositional) route. This therefore predicts that morphological priming effects should be eliminated with high frequency targets. Since all the targets in the Giraudo and Grainger (2000) study were high frequency free roots, the presence of robust morphological priming falsifies the dual-route variant of the sublexical hypothesis.

6. THE REPRESNTATIONAL STATUS OF DERIVATIONAL AFFIXES

One key result most often cited against the supralexical account of morphological representation, is the pseudo-prefixation effect. Indeed, the vast majority of studies investigating the role of affixes in the recognition of morphologically complex words have used what is commonly referred to as a pseudo-affixation manipulation. Within this set of studies, practically all have focused on the processing of prefixes (see Henderson, Wallis & Knight, 1984, for an exception). Pseudo-prefixed words are morphologically simple words that contain an initial letter sequence that corresponds to a prefix but does not function as a prefix in that word (e.g., the *pre-* in *precipice* does not function as a prefix and the remaining letters in the word, *-cipice*, do not form a legal stem). Several studies have examined lexical decision performance to pseudo-prefixed words and nonwords. The pseudo-prefixation manipulation with nonword targets distinguishes morphologically non-decomposable nonwords (e.g., *depertoire*) from morphologically decomposable nonwords (i.e., nonwords such *dejuvenate* that are formed by the illegal combination of a stem and an affix). It has been shown that the latter category of nonwords were more difficult to reject than the former, yielding longer RTs (reaction times) and more errors (Caramazza et al., 1988; Laudanna, Cermele & Caramazza, 1997; Taft & Forster, 1975). Concerning performance to word targets, it has been shown that pseudo-prefixed words like *precipice* are processed more slowly and

produce more errors than both prefixed words and non prefixed words (Pillon, 1998; Taft, 1981). This pseudo-prefixation effect was interpreted as direct evidence in favor of the sublexical hypothesis of morphological representation. According to one specific view of morphological processing with the sublexical framework, the recognition of any morphologically complex word induces an obligatory affix stripping procedure by which the complex word is stripped of its affix, and the stem is used as the entry for lexical access (Taft, 1981). Because pseudo-affixed words contain a letter string that resembles an affix, the model predicts that the pre-lexical parsing procedure will fail and this procedural deadlock will delay lexical access.

However, several researchers have levelled the argument of computational inefficiency against a prefix-stripping procedure. Schreuder and Baayen (1994), for example, argued that pseudo-affixation is so common in languages like Dutch and English that applying a compulsory pre-lexical parsing procedure would lead to a substantial decrease in processing efficiency (see Colé, Beauvillain, Pavard & Segui, 1986, for the same argument relative to French). Empirical evidence against a mandatory prefix-stripping procedure was provided by Laudanna, Burani and Cermele (1994). These authors demonstrated that lexical decision latencies to prefixed nonwords varied according to the distributional properties of affixes in the language (the ratio between truly prefixed and pseudo-prefixed words beginning with the same orthographic sequence). The more likely a given letter string would occur as a prefix compared to a pseudo-prefix, the harder it was to reject nonwords containing such strings at their beginnings.

One unavoidable problem with all pseudo-prefixation studies is that they involve between-item comparisons that unfortunately leave the door open for possible confounding variables. In a study to be summarised here, we used a priming paradigm to study affix processing, thus allowing within-item comparisons. The effects of affixed prime words that share an affix with a complex target word compared to a pseudo-affix prime condition (the pseudo-affix condition provides a form priming baseline against which true affix priming can be measured) were tested using the masked prime paradigm (Giraudo & Grainger, 2002). In these experiments, the same affixed targets words with affixed and pseudo-affixed primes (as well as an unrelated prime condition) were examined. Thus each target word (prefixed and suffixed words) was preceded by a morphologically related word sharing the same affix, a pseudo-affixed word, and an unrelated word.

Figure 5 about here

The results of one experiment using a prime exposure duration of 57ms, are presented in Figure 5. Here we see strong effects of prefixed primes relative to both pseudo-prefixed and unrelated control primes, while suffixed primes produced non-significant effects. These results first clearly

establish that morphological priming effects can be obtained between two words sharing the same prefix. Most important, these prefix priming effects were obtained relative to a pseudo-prefixed prime condition as well as relative to an unrelated prime. The pseudo-affixed primes served as a form control allowing any morphological priming effects to be interpreted as intervening over and above effects of orthographic or phonologic overlap between prime and target. However, the pseudo-prefixed priming condition did not differ significantly from the unrelated priming condition, suggesting that the low degree of form overlap (on average 40% letters shared) was not sufficient to produce measurable priming in the present testing conditions. The contrast between the effects of real prefixes and pseudo-prefixes also stands in contradiction to any model that applies an obligatory prefix-stripping mechanism. If a masked prime sharing a prefix with the following target word can facilitate processing of the latter, then this facilitation should have occurred for both the prefixed and the pseudo-prefixed primes.

Further attempts to obtain significant suffix priming effects all failed. These experiments included manipulations that i) increased the surface frequency of the suffixed word primes, ii) increased the length of the suffix in letters, and iii) used suffixed word primes whose morphological boundary corresponded to a syllable boundary (which is not often the case for derived suffixed words in French).

The prefix priming effect observed in these experiments complements the root priming advantage relative to pseudo-root primes reported in Giraudo and Grainger (2001). Following our interpretation of the root vs. pseudo-root effect, we once again argue that the simplest interpretation of the prefix vs. pseudo-prefix effect is in terms of a supralexical representation of prefixes. Thus all prefixed words containing the same prefix are hypothesized to be connected to a unit coding the presence of that particular prefix. Such prefix representations receive activation from whole-word form representations, and feed-back activation to these representations. Thus only true prefixed word primes will generate significant levels of activation in prefix representations, and these pre-activated representations facilitate the processing of target words containing the same prefix.

However, in the study reported above (Giraudo & Grainger, 2002), a distinct asymmetric pattern of priming effects was obtained for prefixes and suffixes: Strong prefix priming was observed, but there was no hint of a suffix priming effect. This absence of suffix priming effects can be interpreted in terms of a left-to-right processing bias in word recognition. This type of mechanism has already been used to explain asymmetries in effects obtained across prefixes and suffixes. Colé, Beauvillain, and Segui (1989) tested for effects of whole-word (surface) frequency and cumulative root frequency in French derived words. They found surface frequency effects for both prefixed and suffixed words, whereas only suffixed words

showed a cumulative frequency effect (see also Beauvillain, 1996, and Meunier & Segui, 1999). According to these authors, suffixed words but not prefixed words are accessed via the representation of their stem, hence giving rise to the asymmetric effects of cumulative frequency. This processing bias could result from the specific structure of spoken language, spoken words being delivered to the lexical processor gradually from beginning to end. Thus, the processing of suffixed words could be dominated by their stem representation (at the beginning of the word) while prefixed words would be processed with more equal weight given to prefix and stem (e.g., Schriefers, Zwitserlood, and Roelofs, 1991).

Other possible differences between prefixes and suffixes might underlie the asymmetric pattern of affix priming obtained in the present study. A number of important distinctions between prefixes and derivational suffixes have been noted by linguists (e.g., Gardes-Tamines, 1988, for French). Prefixes have a more predominantly compositional character than suffixes. Rather than simply transforming the meaning of the stem to which it is attached, prefixes often add meaning in the same manner as combining two roots in compound words (note the tendency to hyphenate after prefixes but never before suffixes). Prefixes have an exclusively semantic function whereas derivational suffixes have both semantic and syntactic functions. Finally, prefixes never phonologically or orthographically transform the stem to which they attach, whereas suffixes do. These different factors may provide the basis for a different representational status for prefixes compared to suffixes. Only further experimentation will help uncover the mechanisms underlying this distinct dissociation in masked prefix and suffix priming observed in the above study.

7. CONCLUSIONS

All the findings presented above provide support for a supralexical account of morphological representation for derivational morphology in French. It has been shown that (1) there is not extra computation involved in processing derived word primes compared to free root primes, (2) no priming effects can be obtained with word primes containing a pseudo-morpheme (whether a pseudo-root or a pseudo-affix) compared to unrelated primes, and (3) only the surface frequency, and not the cumulative root frequency, of morphologically complex primes affects the size of morphological priming, The key results are summarised in Table 2 along with an evaluation of how they agree with the predictions of the sublexical and supralexical architectures.

Table 2 about here

According to the supralexical account of morphemic representation, representations corresponding to roots and affixes are activated by any whole-word representation containing such units. Activation then feeds-back from morphemic units to whole-word representations hence providing a mechanism that accounts for the morphological priming effects observed in the present study (for roots and prefixes) and other masked prime experiments (for roots). The fact that morphemic representations only receive activation from whole-word representations (in comprehension), allows this approach to capture the results summarised in Table 2. Activation of morphemic representations corresponding to the root of a stimulus occurs in exactly the same way for free root and derived word stimuli. On the other hand, the presence of letter sequences that correspond to roots without having such a status in reality (i.e., pseudo-roots), will not lead to activation of the corresponding morphemic representation. Similarly, when the surface frequency of derived word stimuli is too low, activation of the component morphemes is negligible in the conditions of masked priming. Finally, while increasing the surface frequency of derived words increases the connections strengths between the whole-word representation and its corresponding morphemic representation, variations in cumulative root frequency have little influence. If one assumes that modifications of connection strengths occur once the network has stabilised, in an interactive activation network this implies that only the connection between the stimulus word and its corresponding morphemes are reinforced. Whole-word representations that are partially activated in early phases of stimulus processing (and these are likely to include morphologically related words) are not assumed to participate in the adjustment of connection strengths.

The supralexical account of morphological representation can be thought of as a whole-word based solution to the problem of morphological parsing. The crucial question then is how novel complex forms can be recognized as such. The solution to this problem can be found in hierarchical activation networks that allow significant levels of activation to develop from partial form overlap with the stimulus (e.g., Grainger & Jacobs, 1998). For example: given the novel stimulus XNESS, and given prior knowledge of words such as Yness, Zness, and Xer, Xed, then partial overlap with the stimulus allows activation of these whole-word forms and then activation of the X and NESS morphemes. At the whole-word level, no single representation will reach a critical activation level for identification, and so no learning will proceed on the basis of existing whole-word representations. This is a situation where the system that controls learning in a network can dedicate a new representational unit to stimulus being processed. This newly recruited representation is then connected to those units that are most active at the morpheme level. This is a classic example of

Hebbian learning applied in various areas of cognitive science (e.g., Burton, 1998).

This solution for the processing of novel complex forms also provides part of the answer as to how supralexical morphemic representations might be learned in the first place. Once whole-word representations have developed and are connected to their corresponding semantic representations, then morphemic representations can develop via the correlated activity of a specific semantic representation (a set of semantic features, for example) and the simultaneous activation of words sharing bottom-up information derived from the stimulus. Thus, on presentation of a complex word such as *reprint*, all words sharing the same prefix will be activated and will activate in turn the semantic representation of this prefix (assuming systematicity in the meaning of the prefix across the different prefixed words). Each specific case of correlated activity in form and meaning can be represented by a morphemic unit that codes that specific relationship. Thus, the representation for the prefix RE can then be thought of as a node that codes the link between all whole-word forms that contain this prefix and that also activate some common semantic representation.

Now, given that the great majority of linguistic input (particularly in the early years of language use) arrives through the auditory modality, the specific constraints imposed by spoken language will shape the above learning process. More specifically, it is now assumed by several models of spoken word recognition (e.g., McClelland & Elman, 1986; Norris, 1994) that phoneme sequences that are possible word beginnings will activate such words even when not aligned with the beginning of the stimulus. Prefixes, and the stems of suffixed words are at the beginning of complex words, hence simplifying the detection of correlated activity in these cases. Free stems (as opposed to bound stems) of prefixed words also constitute legal word beginnings and will therefore be activated even when not aligned with the beginning of the stimulus. On the other hand, suffixes by definition never form the possible beginning of a word (suffix-like sequences such as *ment* at the beginnings of words are totally unrelated to real suffixes). This implies that only prefixes, stems (free or bound) of suffixed words, and the free stems of prefixed words offer the appropriate conditions for detecting form-meaning correlations that are necessary for the development of morphemic representations. This leads us to predict a further dissociation in masked morphological priming: whereas free stems of prefixed and suffixed words should show priming, bound stems should only produce priming in suffixed words. On this point, it is interesting to note that in French there are many cases of bound stems occurring in suffixed words, while none occur in prefixed words. The French language may have evolved in a way that reflects the difficulty in detecting the morphemic status of such bound stems in prefixed words.

In this chapter we have summarised some recent research on the processing of French derivational morphology using the masked prime paradigm. The results confirm that morphological codes are rapidly and automatically extracted from printed strings of letters, even in conditions where no morphological decomposition is required (i.e., with free stem targets). The evidence suggests that this automatic activation of morphological codes operates via whole-word forms: the lexical processor first detects the presence of a given word form, and then activates its component morphemes. Future research should examine the limits of this alternative *supralexical* approach to morphological representation, particularly across different languages.

REFERENCES

Beauvillain, C. (1996). The integration of morphological information and whole-word form information during eye fixations on prefixed and suffixed words. *Journal of Memory and Language, 35,* 801-820.

Burton., A.M. (1998). In J. Grainger & A.M. Jacobs, *Localist connectionist approaches to human cognition.* Mahwah, NJ: Erlbaum.

Butterworth, B. (1983). Lexical Representation. *Language Production, 2,* 257-295.

Caramazza, A., Laudanna, A., & Romani, C. (1988). Lexical access and inflectional morphology. *Cognition, 28,* 297-332.

Colé, P., Beauvillain, C, Pavard, B., & Segui, J. (1986). Organisation morphologique et accès au lexique. *L'Année Psychologique, 86,* 349-365.

Colé, P., Beauvillain, C., & Segui, J. (1989).On the representation and processing of prefixed and suffixed words: A differential frequency effect. *Journal of Memory and Language, 28,* 1-13.

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

Drews, E., & Zwitserlood, P. (1995). Morphological and orthographic similarity in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 1098-1116.

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

Forster, K. I., Davis, C., Schoknecht, C., & Carter, R. (1987). Masked priming with graphemically related forms: Repetition or partial activation? *The Quarterly Journal of Experimental Psychology, 39,* 211-251.

Frost, R., Forster, K., & Deutsch, A. (1997). What can we learn from the morphology of Hebrew ? A masked priming investigation of morphological

representation. *Journal of Experimental Psychology: Learning, Memory and Cognition, 23*, 1-28.

Gardes-Tamines, J. (1988). *La grammaire: phonologie, morphologie, lexicologie, Tome 1*. Paris: Armand Colin.

Giraudo, H., & Grainger, J. (2000). Effects of prime word frequency and cumulative root frequency in masked morphological priming. *Language and Cognitive Processes, 15, 421-444.*

Giraudo, H., & Grainger, J. (2001). Priming complex words: Evidence for supralexical representation of morphology. *Psychonomic Bulletin & Review*, in press.

Giraudo, H., & Grainger, J. (2002). On the role of derivational affixes in recognizing complex words: Evidence from masked priming. *Manuscript submitted for publication.*

Grainger, J., Colé, P., & Segui, J. (1991). Masked morphological priming in visual word recognition. *Journal of Memory and Language, 30*, 370-384.

Grainger, J. & Ferrand, L. (1994). Phonology and orthography in visual word recognition: Effects of masked homophone primes. *Journal of Memory and Language, 33*, 218-233.

Grainger, J., & Jacobs, A., M. (1998). *Localist connectionist approaches to human cognition*. Mahwah, NJ: Erlbaum.

Henderson, L., Wallis, J., & Knight, D. (1984). Morphemic structure and lexical access. In H. Bouma & D. Bouwhuis (Eds.), *Attention and Performance, X* (pp.221-224). Hillsdale, N. J: Lawrence Erlbaum Associates, Inc.

Laudanna, A., Burani, C., & Cermele, A. (1994). Prefixes as processing units. *Language and Cognitive Processes, 9*, 295-316.

Laudanna, A., Cermele, A., & Caramazza, A. (1997). Morpho-lexical representations in naming. *Language and Cognitive Processes, 12*, 49-66.

Lukatela, G., Gligorijevic, B., Kostic, A., & Turvey, M.T. (1980). Representation of inflected nouns in the internal lexicon. *Memory & Cognition, 8*, 415-423.

Manelis, L., & Tharp, D. (1977). The processing of affixed words. *Memory & Cognition, 5,* 690-695.

McClelland, J. L. & Elman, J. L. (1986). The TRACE model of speech perception. *Cognitive Psychology, 18*, 1-86.

McClelland, J.L., & Rumelhart, D.E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 88*, 375-405.

Meunier, F. & Segui, J. (1999). Frequency effects in auditory word recognition: The case of suffixed words. *Journal of Memory and Language, 41*, 327-344.

Norris, D. (1994). Shortlist: A connectionist model of continuous speech recognition. *Cognition, 52*, 189-234.

Pillon, A. (1998). The pseudo-prefixation effect in visual word recognition: A true-neither strategic nor orthographic-morphemic effect. *The Quarterly Journal of Experimental Psychology, 51*, 85-120.

Rastle, K., Davis, M.H., Marslen-Wilson, W.D., and Tyler, L.K (2000). Morphological and semantic effects in visual word recognition: A time-course study. *Language and Cognitive Processes*, *15*, 507-538.

Schreuder, R., & Baayen, R. H. (1994). Prefix stripping re-revisited. *Journal of Memory and Language, 33*, 357-375.

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

Schreuder, R., & Baayen, R. H. (1997). How complex simplex words can be. *Journal of Memory and Language, 37*, 118-139.

Schriefers, H., Zwitserlood, P., & Roelofs, A. (1991). The identification of morphologically complex spoken words: Continuous processing or decomposition? *Journal of Memory and Language, 30*, 26-47.

Segui, J. & Grainger, J. (1990). Priming word recognition with orthographic neighbours: Effects of relative prime-target frequency. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 65-76.

Taft, M. (1981). Prefix stripping revisited. *Journal of Verbal Learning and Verbal Behaviour, 20*, 289-297.

Taft, M. (1994). Interactive-activation as a framework for understanding morphological processing. *Language and Cognitive Processes, 9*, 271-294.

Taft, M., & Forster, K. I. (1975). Lexical storage and retrieval of prefixed words. *Journal of verbal Learning and verbal Behavior, 14*, 638-647.

Table 1. Logic of the predictions for the experiments reported in Giraudo and Grainger (2001), generated from the sublexical and supralexial accounts of morphological representation. The sublexical model's predictions differ as a function of whether the morphological parser is sensitive to the morphemic status of all morphemes in the stimulus. The predictions are expressed as relative sizes of priming effects for the different prime conditions tested (root, derived / morpheme, pseudo-morpheme).

Model variant	Experimental predictions
Sublexical hypothesis:	
- 1. Parser checks that remaining letters form an affix	root > derived morpheme > pseudo-morpheme
- 2. Parser does not check that remaining letters form an affix	root = derived morpheme = pseudo-morpheme
Supralexical hypothesis:	root = derived morpheme > pseudo-morpheme

Table 2. Summary of the key results obtained in masked morphological priming experiments (with 57 ms prime exposures) relative to predictions generated from the sublexical and supralexical accounts of morphological representation. The results are shown to be in agreement with the predictions (Y), in disagreement (N), or possibly in agreement (?).

Result	Model	Agreement
No difference in priming from free root and derived word primes.	Sublexical Supralexical	? Y
No priming from monomorphemic primes containing a pseudo-root	Sublexical Supralexical	? Y
No priming from derived word primes with low surface frequencies	Sublexical Supralexical	N Y
No influence of the cumulative root frequency of derived word primes.	Sublexical Supralexical	N Y