Compound Pre-Access Decomposition: Effects of Constituent Disruption*

Roberto G. de Almeida & Gary Libben

Abstract

In this study, we report three experiments that investigate whether recognition of existing English compounds is dependent on the recognition of their constituent morphemes. A constituent disruption paradigm was employed with compounds of fixed length, but variable constituent length. The results of Experiment 1 confirmed that disrupting a single character in a three-character word by replacing it with a crosshatch (e.g., *way* → *w#y)* interfered with recognition rates more than the disruption of a single character of a five-letter word (e.g., *drive* → *dr#ve*). However, in Experiment 2, when these constituents were embedded in eight-character compounds (e.g., *driveway*), no differential effect of constituent disruption was found for either naming accuracy or latency. In Experiment 3, this effect was replicated in a masked priming lexical decision paradigm in which disrupted compounds served as primes and intact compounds served as targets. The results of the three experiments were interpreted to offer counter-evidence to the claim that compound recognition is dependent on the prelexical recognition of constituent morphemes. We argue that although English compound processing may indeed be characterized by both whole-word recognition and constituent activation, access to individual morphemes is not a necessary condition in the recognition process.

Introduction

Our goal in this paper is to shed light on one of the fundamental questions in the study of lexical processing: What role do constituent morphemes of a multimorphemic word play in on-line word recognition? We focus on the question of whether the processing of English compound words is characterized by pre-lexical morphological decomposition, the breakdown of a multimorphemic word into its morphological constituents prior to the recognition of the whole word. The question of whether such decomposition occurs has been the subject of considerable debate in the psycholinguistic literature since Taft & Forster (1975) presented evidence that English prefixed words were stripped of their affixes to enable lexical lookup of the stem. One of the reasons for the centrality of this debate has been that our views of whether prelexical morphological decomposition takes place has important consequences for our views of

Folia Linguistica XXXVI/1-2 *0165-4004/02/36-97 $ 2.–*
(C) Mouton de Gruyter, Berlin — *Societas Linguistica Europaea*

the nature of representations within the mental lexicon and the functional architecture of the lexical access system (Baayen, Dijkstra, & Schreuder, 1997; Schreuder & Baayen, 1995).

As has been argued by Libben (1994) and Libben, Gibson, Yoon, & Sandra (in press), compound words present a privileged window to the study of prelexical morphological decomposition. The primary reason for this is that the morphological nature of compounding presents a challenge to our conceptions of how decomposition may be achieved. Because compounds are composed of *root* + *root* structures, their constituent morphemes are drawn from the open-class vocabulary of the language, rather than from the closed-class set of affixes that could, in principle, be removed in the recognition of prefixed or suffixed words. This characteristic ensures that, from a computational perspective, prelexical decomposition of compounds will be more complex than decomposition of affixed forms. Moreover, there is reason to believe that the process of prelexical morphological decomposition will be called upon relatively often in normal language processing. Compounding is typically a very productive word formation process. Thus, native speakers of a language such as English are likely to encounter new compound forms that, by definition, can only be understood in terms of their constituent morphemes. These morphemes must be isolated, accessed and understood in order for a compound interpretation to be formed.

The link that we have just formed between prelexical morphological decomposition in the recognition of existing multimorphemic words and the process by which novel multimorphemic words are interpreted has played a key role in our conceptions of morphological decomposition. In the early polarized debates on whether decomposition has any role to play in the recognition of existing words, proponents of the whole-word access perspective (e.g., Butterworth, 1983) claimed that decomposition is exclusively a non-word strategy that is available under exceptional circumstances for the processing of real words. This perspective was offered as an explanation for the fact that decomposition effects obtained in lexical decision experiments were sensitive to the proportion of decomposable non-word fillers in the experimental materials (McQueen & Cutler, 1998).

Another reason for which the link between decomposition in non-words and real words is critical is that non-word studies have been the source of hypotheses regarding the exact mechanisms of decomposition and the manner in which these mechanisms interact. Multimorphemic strings that are seen or heard for the first time can only be understood through prelexical morphological decomposition. Thus it has been natural to look to the investigation of the processing of novel words for templates of how prelexical decomposition might take place for real words. In order to explore the nature of such a template, let us consider a novel compound such as *catbed*. This compound, although novel, is most often easily comprehended as 'a bed for cats'. Assuming that this comprehension is the result of successful morphological decomposition, we might begin by identifying the following three logical subcomponents of the process:

(a) Division of the string into two morphemes (*cat & bed*)
(b) Access to the representations and meanings of those morphemes
(c) Creation of a compound interpretation based on morphological headedness in the language and the meanings of the individual morphemes.

At first blush, the existence of these subcomponents appears self-evident. Access to the constituents, *cat* and *bed*, is only possible once they have been isolated within the input in sub-process (a). The access itself in subprocess (b) is critical because, without it, there is no basis at all for interpretation. That interpretation must be based on the meanings of the root morphemes in subprocess (c) as well as on the knowledge that English compounds are right-headed. This knowledge of headedness ensures that interpreters assume that a *catbed* is a type of *bed*, not a type of *cat*.

In a series of recent experiments, Libben, Derwing & de Almeida (1999) have presented evidence that challenges this conception of prelexical morphological decomposition. They investigated constituent activation effects for novel compounds such as *clamprod*, which can be parsed as either *clamp-rod* or *clam-prod* and thus contain four possible morphemes, rather than the two available for novel strings such as *catbed*. They found that all possible constituent morphemes of ambiguous novel compounds such as *clamprod* were activated (*clam, clamp, prod*, and *rod*). This finding calls into question whether process (a), prelexical parsing, actually occurs in novel compound processing in a manner that is independent from process (b), constituent access. They reasoned that if it did, a univocal parse of the novel compound would occur, resulting in activation of either *clam & prod*, or *clamp & rod*, but not both.

A second finding in the Libben, Derwing & de Almeida (1999) study was that parsing preferences for ambiguous novel compounds (e.g., *clamp-rod* or *clam-prod*) were significantly correlated to participants' semantic plausibility ratings for alternative parses. In other words, they parsed ambiguous strings such as *clamprod* into *clam-prod*, if this version was judged to make more sense than the alternative. This suggests that subprocess (c), compound interpretation, does not occur as a result of morphological parsing, but rather is part of the parsing process itself.

The findings above suggest that constituent-based compound processing is not as straightforward as it might have appeared. If the modeling of prelexical decomposition in real words is to be based on the novel word template, there is perhaps no reason to assume that it is composed of easily demarcated procedures (a, b, c). In fact, all that we can be sure about is the core definitional feature of prelexical decomposition—namely that it is based on access to individual morphological constituents of a multimorphemic word. In the study reported in this paper, our goal was to target exactly that feature.

The starting point in the design of our investigation was the observation by Sandra (1990) that constituent morphological activation can occur in two places in the recognition process—prelexically or postlexically. It is impossible to experimentally stop post-lexical activation from occurring—it is assumed to be intrinsic to the representation of a multimorphemic word. Our goal, therefore, was to experimentally manipulate whether the prelexical component is possible, by systematically disrupting access to individual morphemes within a compound stimulus. Consider, for example, two eight character compounds *aircraft* and *driveway*. If these stimuli are altered so that one character from each of the initial constituents is disrupted (e.g., *a#rcraft, dr#veway*), we have effectively removed 1/8 of the compound in both cases. However, because *aircraft* is composed of a three-character initial constituent and a five character final constituent, and because *driveway* shows the opposite configuration (five-three), we have effectively removed 1/3 of the initial constituent in *aircraft*, but only 1/5 of the constituent in *driveway*. It should be noted that in both cases, the full compound remains uniquely specified in the language, despite the disruption.

Similarly, if we disrupt one character of the final constituents of these compounds (e.g., *aircr#ft, drivew#y*), we can achieve the same constituent disruption asymmetry (1/3 vs. 1/5), while holding the amount by which the whole word is disrupted constant. In this case, however, we gain the possibility of investigating whether disrupting the morphological head of the compound affects overall recognition in a manner that is distinct from disrupting the initial constituent. Finally, the paradigm allows us the possibility to compare these position-specific disruptions in bimorphemic compounds to the same types of disruptions in length-matched non-compound control words (e.g., *mountain → m#untain or mount#in).*

To summarize to this point, the constituent disruption paradigm employed in this study was designed to investigate whether or not the recognition of a compound relies on an early analysis of its constituent morphemes. Assuming that this paradigm does in fact detect the early, automatic phases of word recognition, disrupting relevant aspects of the stimulus word should allow us to determine in what ways constituent degradation affects the activation of compounds as well as the extent to which the position of that degradation plays a role. Since, in English compounds, the head is typically the second element, a greater recognition cost when heads were disrupted would indicate that they play a role in early (i.e., prelexical) processing. On the other hand, if disrupting the non-head constituent would yield a greater compound recognition cost, this would provide evidence for a left-to-right compound parsing procedure (Jarema, Busson, Nikolova, Tsapkini, & Libben, 1999; Libben, Gibson, Yoon & Sandra, in press).

The study that we report below is composed of three experiments. In Experiment 1, we investigated the degree to which a tachistoscopic presentation of compound constituents in isolation would be recognized when disrupted (e.g., *a#r*). This information enabled us to determine the recoverability of those elements as individual words. In Experiment 2, these words were embedded as constituents in noun-noun compounds that were presented to participants in either an intact or disrupted form in a naming latency experiment. In Experiment 3, the stimulus presentation format in the naming latency study was incorporated into a visual masked priming lexical decision experiment in which disrupted and non-disrupted compounds served as primes for their intact forms.

2. Experiment 1: compound constituent recoverability

In this experiment we investigated the recognition of disrupted compound constituents—e.g., *a#r*, *cr#ft*, *b#th*, and *r#om*, corresponding to the root constituents of the compounds *aircraft* and *bathroom* respectively. The goal of the experiment, which functioned as a pretest for Experiment 2, was to determine whether a 60 ms tachistoscopic presentation of words (e.g., *a#r* & *cr#ft*) would correctly elicit their full forms (i.e., *air* & *craft*). This allowed us to determine whether or not participants had enough information to access the compound constituents when they were disrupted.

2.1 Method

2.1.1 Participants

Thirty native speakers of English participated as volunteers in this experiment. Eleven were students attending a Cognition course at Concordia University and 19 were students attending a Linguistics course at the University of Alberta.

2.1.2 Materials

The materials for this experiment were obtained from a set of eight-letter compounds (e.g., *aircraft*) that were split into their root constituents (e.g., *aircraft* → *air* and *craft*). Our compound stimuli consisted of 72 items composed of 3+5 (e.g., *aircraft*), 4+4 (e.g., *bathroom*), and 5+3 (e.g., *driveway*) letter constituents, yielding 144 monomorphemic stimuli. Of those, we used 138 original words (we eliminated six repetitions such as *dog*, constituent of both *doghouse* and *puppydog*), 45 of which were three-letter words, 48 were four-letter words, and 45 were five-letter words.

2.1.3 Design and Procedure

Two lists of stimuli were created, each one containing the 138 experimental words. The difference between the lists was that for each intact word in one list (e.g., *dog*), there was a corresponding disrupted string in the other list (e.g., (*d#g*). Half of the words in each list were disrupted and half were intact. The lists were counter-balanced so that the words that were intact in the first list were disrupted in the second list and vice-versa. Each list was run with a different group of subjects, list one at Concordia University and list two at the University of Alberta.

Both groups were run in dimly lit classrooms. The stimuli were presented on large screens using Telex (Concordia) and In Focus (Alberta) SVGA projectors attached to Macintosh PowerBook G3 computers. Participants received booklets containing lines numbered from 1 to 138. They were instructed that they would be presented with words on the screen and that their task would be to write down as quickly and as accurately as possible which words appeared on the screen. They were told that they had to write a word even when they were not sure about which word was presented on the screen. Each trial started with the presentation of the trial number for 5 seconds. The offset of trial number presentation was signaled by a beep. This indicated to participants that they had to complete writing down the word that they saw in the previous trial and to focus their attention on the screen for the presentation of the new word. The trial number also served to indicate which line of the booklet participants had to use in the upcoming trial. 500 ms after the trial number was presented, there was a sequence of three events presented on the screen: (1) a forward mask (######) presented for 500 ms, (2) the target word presented in lowercase letters (e.g., *cr#ft*) for 60 ms, followed by (3) a backward mask (######) also presented for 500 ms. After the second mask, the screen was blank and participants had up to 6 seconds to write down the word (1 second of inter-trial interval plus 5 seconds corresponding to the presentation of the upcoming trial number) until the beep signaled the beginning of a new mask-target-mask triad. The forward and backward masks had the same number of hash marks as their corresponding target words. The stimuli (words and hash marks) were written with CourierNew 48 font so that the spacing between characters was kept constant and each hash mark was presented in the exact position corresponding to a letter of the target word. Stimuli were presented in white over a black background. The experiment began with the presentation of 9 practice trials followed by the 138 experimental trials. Each session lasted for about 17 minutes.

2.2 Results and Discussion

The dependent variable in this experiment was the percentage of correct response to the disrupted words. Our first step in the analysis was to determine whether the participants from the two testing universities differed systematically in their performance. We found no significant difference between the recoverability scores of the two groups (t $(272)=1.97$, $p=.87$), and we therefore proceeded to combine the results from both groups in the remaining analysis and to conduct by-items ANOVAs as primary measures.

In Table 1, the percentage correct for the three, four and five character stimuli are presented. Because our goal was to use these recoverability scores to inform the results of Experiment 2, in which the individual words function as either the initial or final constituents of compounds, we separated the stimuli so that words that would function as initial and final constituents were analyzed separately. In both analyses, there was a significant effect of length (Constituent 1: $F(2, 46)=12.54$, $p<.0001$), (Constituent 2: $F(2, 46)=5.79$, $p<.01$). Post-hoc analyses for each constituent showed no significant difference between the four-character and five-character stimuli using the Bonferonni-Dunn technique (Constituent 1, $p=19$; constituent 2, $p=.96$). Using the same post-hoc analysis, we found a significant difference between the 3-character stimuli and the other two lengths for both initial and final constituents (all p's $<.005$).

Table 1. Mean recovery rates (and standard deviations) for words of different lengths. The Constituent 1 column represents words that formed the first constituent of our core compounds (e.g., bathroom*). The Constituent 2 column represents words drawn from the final compound constituents (e.g.,* driveway).

Length		Constituent 1 Mean	Constituent 1 SD	Constituent 2 Mean	Constituent 2 SD
Three	(e.g., *way*)	23%	20	27%	29
Four	(e.g., *bath*)	53%	28	61%	30
Five	(e.g., *drive*)	65%	24	62%	37

The results of Experiment 1 confirmed our expectation that disrupting one character of a three-character word significantly decreased the ability of participants to correctly recover that constituent. This result thus formed the input to Experiment 2. As we have discussed above, we reasoned that if individual constituents play a role in the manner in which compounds are recognized, then we should see the same effect of constituent disruption in compound recognition. If, on the other hand, there is no relationship between the recoverability of individual constituents and either naming accuracy or naming latency in the compound recognition task, we would be forced to conclude that compound recognition is not dependent on the correct identification of individual constituents.

3. Experiment 2: Naming latency

The results of Experiment 1 created recoverability scores for each of the individual compound constituents. In Experiment 2, we investigated whether these disruptions, now in the context of noun compounds, would generate differential effects that were related to their length relative to the length of entire compounds and were related to their position within the compound string. We also investigated whether compound disruption effects were related to the distributional characteristics such as constituent frequency and constituent family size. As we have stated above, the overall goal of the constituent disruption manipulation was to probe the extent to which individual constituent characteristics affected the recognition of the entire compound. In cases where such effects are observed, we reasoned that they would reflect the contribution of the individual morphemes to the recognition of the word as a whole.

In this experiment, naming latency was measured under both disrupted and intact presentation conditions. Participants were asked to read the words presented to them as quickly as possible, and the speed of their responses constituted the primary dependent variable. The compound stimuli acted as their own controls in the investigation, thus reducing the effects of differences in their phonological onset characteristics that would otherwise be problematic in the evaluation of measures of naming onset latency. The decision to employ a naming latency paradigm was motivated by the following considerations: We were concerned that if lexical decision were used as the dependent variable, the disrupted presentations would yield too large a proportion of 'no' responses, making it difficult to establish a measure of disruption cost (disrupted presentation latency minus intact presentation latency) for identical responses (i.e., 'yes' in lexical decision). The naming latency paradigm provided the opportunity to overcome this potential difficulty, while still yielding a measure of lexical access. Moreover, the measure obtained from naming latency could most easily be related to the constituent recoverability data obtained in Experiment 1. The reason for this was that in the naming latency paradigm, unlike a lexical decision paradigm, we could be certain that the stimulus word was exactly the one that had been recovered by participants. In other words, by including only correct responses in the analysis, we could be certain that a presentation of *dr#veway* was indeed recovered as the compound *driveway* (as opposed to being simply recovered as a real word of English).

3.1 Method

3.1.1 Participants

Thirty-five native speakers of English participated as volunteers in this experiment, 10 were students attending a Cognition course at Concordia University and 23 were students attending a Linguistics course at the University of Alberta.

3.1.2 Materials and Design

We used 54 8-letter compounds formed from the words used in Experiment 1 (see Appendix 1). All 54 compounds were semantically transparent (i.e., in all cases, the meaning of the whole compound was a function of the meaning of the constituents). The 54 compounds were formed by constituents with 3+5 (e.g., *keyboard*), 4+4 (e.g., *bathroom*), and 5+3 (e.g., *driveway*) letters, with 18 compounds of each composition type. In addition to the 54 compounds, there were 27 eight-letter non-compound words (e.g., *elephant*, *incident*).

Participants were presented with each compound stimulus (e.g., *aircraft*) under one of three conditions: (1) disrupted in the first constituent (henceforth, D-1; e.g., *a#rcraft*), (2) disrupted in the second constituent (henceforth, D-2; e.g., *aircr#ft*), or (3) intact (e.g., *aircraft*). Non-compound words were also presented under one of the three conditions, except that the position of the disruption was in a vowel appearing either in the first half (e.g., *qu#stion*) or in the second half of the word (e.g., *hosp#tal*). Half of the disrupted compounds were disrupted in fixed positions, as follows: the second letter in three-letter constituents (e.g., *a#rcraft*), the third letter in five-letter constituents (e.g., *aircr#ft*), and the second letter in four-letter constituents (e.g., *bathr#om*). The other half of the compound set was disrupted in variable positions—in this case, by replacing with a hash mark the first vowel of the constituent (e.g., *#rmchair*, *b#rthday*, and *head#che*, corresponding to *armchair*, *birthday*, and *headache*). There were three lists of materials, each one containing all 54 compounds, with 18 trials of each stimulus type (i.e., D-1, D-2, and intact), and all 27 non-compound words, with 9 trials of each stimulus type (disrupted early, disrupted late, or intact; we will also refer to disrupted monomorphemic as D-1 or D-2 corresponding to early or late disruption). In addition to the 81 experimental stimuli, each one of the three lists contained 81 fillers for a total of 162 trials. Among the fillers, 27 were intact words and 27 were disrupted at random positions. The fillers were either monomorphemic or derived words of diverse grammatical categories ranging from four (e.g, *bark*) to 10 letters in length (e.g., *expression*).

3.1.3 Procedure and Apparatus

Participants were presented with a sequence of four stimuli on a computer screen: (1) a fixation point (an asterisk) for 1000 ms, followed by (2) a forward mask (########) for 500 ms, followed by (3) the target word in lower case letters (e.g., *keyboard*, *k#yboard*, or *keyb#ard*) presented for 60 ms. The participant's task was to say the word that appeared on the screen as quickly and as accurately as possible. A beep sounded when the voice onset was recorded. We used a self-scoring naming paradigm in which, 1 second after the subject pronounced the target word, an upper-case version of that word appeared on the screen. Participants had, then, to press a button labeled *yes* if the word they had pronounced was the correct one or to press a button labeled *no* otherwise. Participants were tested individually in a dimly lit room. Stimulus presentation was conducted with Macintosh LC 575 and G3 computers running PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993). Response times and error data were collected via CMU response boxes. Naming times were measured from the presentation of the target to the onset of the voice response. Each session lasted for about 20 minutes and began with the presentation of the instructions on the computer screen, followed by a set of 10 practice trials, a reinforcement of the instructions, and the 162 trials.

3.2 Results and Discussion

In the calculation of response accuracy and latency, five compounds that showed overall accuracy of less than 50% were removed from the analysis. This yielded a final compound list of 49 items.

3.2.1 Response accuracy

Figure 1 represents the overall pattern of naming accuracy. As can be seen in the figure, the disruption paradigm reduced recognition rates across all compound types. A two-way ANOVA showed a main effect of presentation condition ($F(2,96) = 13.7, p<.001$) and neither an effect of compound type ($F(2,46) = 1.2$, $p=.30$) nor an interaction between the two factors ($F(4,92) = 1.7$, $p=.15$). Planned comparisons showed that the main effect of presentation condition resulted from an overall significant difference between the intact and disrupted conditions (t (48)=6.3, $p<.0001$). There was no overall significant difference resulting from the disruption of the first vs. second compound constituents (t (48)=0.63, $p=0.52$).

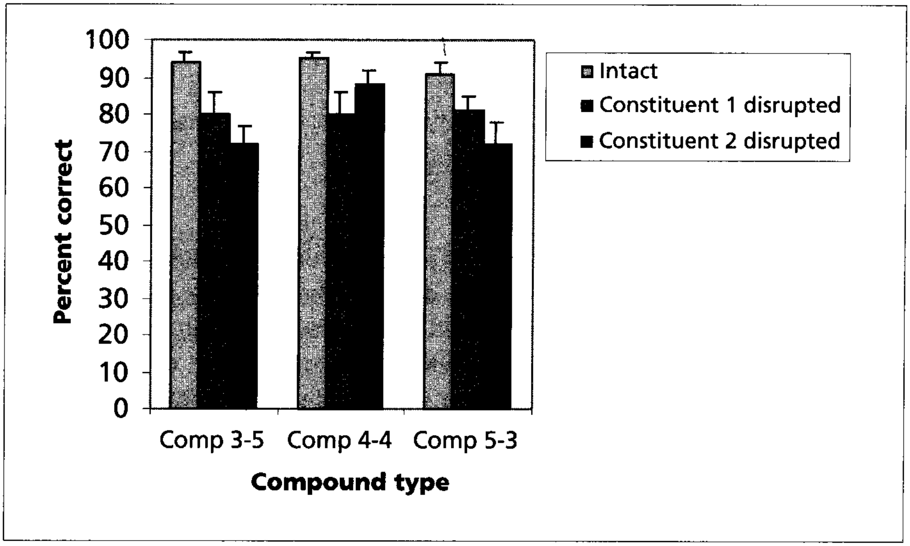

Figure 1. Response accuracy (in percent) in the naming task. Error bars represent standard errors for each category.

It is noteworthy that the disruption paradigm significantly lowered naming accuracy rates despite the fact that there was in fact only one compound in the English language (i.e., the target) that corresponded to the disrupted presentation. In other words, the missing character was, strictly speaking, not required to recover the target compound. We conclude from this that the paradigm had a relatively powerful effect on participants' word recognition procedures. We assume then, that if indeed compound recognition were dependent on the prior recognition of individual morphemes, we should have found substantially differences between the accuracy rates for the 3-5 and 5-3 compounds across the two disruption conditions, D1 and D2. In fact, however, those accuracy rates were virtually identical. The 3-5 compounds showed a D1 accuracy rate of 79.7% and the 5-3 compounds showed a D1 accuracy rate of 80.7%. For the D2 condition, the accuracy rate for the 3-5 compounds was 71.8%, and for the 5-3 compounds 71.7%. Therefore, on the basis of response accuracy, we do not see evidence of individual constituent recoverability. It should be noted, however, that the basis for the expectation that 3-5 and 5-3 compounds would behave in a differentiated manner derived from the results of Experiment 1. As a check against the possibility that there was a relation between constituent recoverability and compound accuracy, but that this relation was perhaps obscured by the classification of compounds into 3-5, 4-4, and 5-3 types, we correlated the recoverability scores from Experiment 1 with response accuracy scores, irrespective of compound category. The results of this analysis yielded no significant relationships. The correlation of Constituent 1 recoverability and the response accuracy when that constituent was disrupted was *r*=.27. For the second constituent recoverability and disruption, *r*=.17. Finally, the correlation between average constituent recoverability (across constituents 1 and 2) and average disrupted accuracy was *r*=.20.

3.2.2 Response latency

Latency data were calculated for correct responses only. Latencies greater than 2000 ms (2%) were removed from the dataset. The resulting patterns across compound types and presentation conditions are shown in Figure 2. As in the accuracy analysis, there was clear evidence of processing cost associated with stimulus disruption, resulting in a main effect of the presentation type factor (*F*(2,92)=9.4, *p*=.002). There was no main effect of compound type, (*F*(2,46)=0.17, *p*=.84).

Although the interaction between compound type and presentation type was not significant (*F*(4,92)=1.6, *p*=.17), Figure 2 shows a relatively high disruption cost associated with disruption of the second element in 3-5 compounds. This response time difference is, however, in the opposite relation to that predicted on the basis of constituent recoverability (i.e., it is associated with disrupting a 5-character constituent rather than a 3-character one).

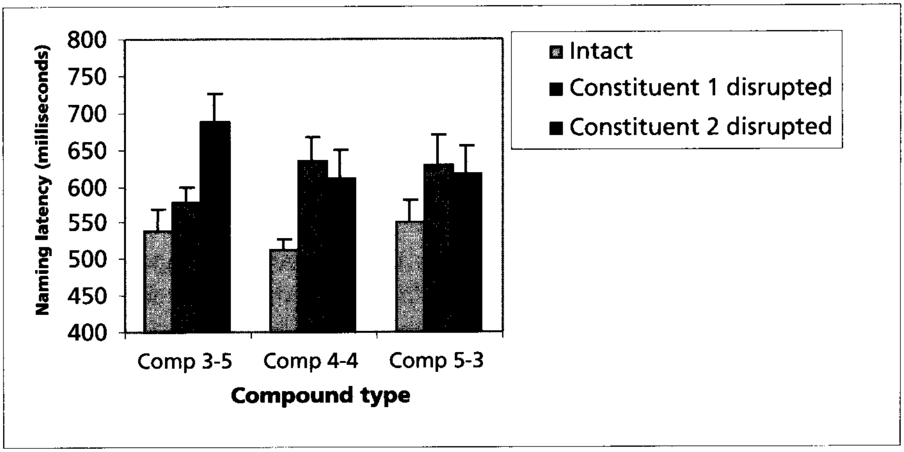

Figure 2. Response latency (in milliseconds) in the naming task. Error bars represent standard errors for each category.

Our analysis of latency mirrored that for accuracy in that we also correlated latency scores for items with the constituent recoverability means for those items irrespective of compound type. This analysis did not yield significant correlations for either D1 or D2 disruption conditions, as is shown in Table 2.

Table 2. Pearson product moment correlations for constituent recoverability in Experiment 1 and Naming latency in Experiment 2 by items.

	RT for D1	RT for D2
Constituent 1 recoverability	−.04	−.19
Constituent 2 recoverability	+.05	+.11

3.2.3 Effects of other constituent properties on naming accuracy and latency

To summarize to this point, the analysis of response accuracy and latency did not show the dissociation of 3-5 and 5-3 compounds that was predicted by the prelexical decomposition hypothesis. It should be noted, however, that in the current investigation, a great deal of weight was assigned to the recoverability characteristics of individual morphemes under the specific experimental conditions of the disruption paradigm. We therefore extended the investigation to incorporate two other measures of individual morpheme characteristics that have been found to affect response latencies in compound processing (e.g., de Jong, Feldman, Schreuder, Pastizzo & Baayen, in press). These were the Positional Family Size of each constituent and the Positional Family Frequency. The first of these measures, the Positional Family Size of a morpheme, corresponds to the number of compounds that have that particular morpheme as either an initial or final constituent. Thus, the Position Family Size of *air* in *aircraft* would correspond to the number of compounds that have *air* as their initial constituent (in this case, 68, based on the MRC Psycholinguistics Database). The Positional Family Frequency corresponds to the cumulative frequency of those 68 compounds (in this case yielding a value of 120). Analyses using these measures yielded no positive or negative correlations beyond .2 with the effects of disruption in either response accuracy or response latency.

In summary, both the accuracy and latency data obtained in the compound naming experiment reveal an overall effect of disruption, but do not show a pattern of responses that is consistent with the view that compound recognition is dependent on prior recognition of constituent morphemes. It is possible, however, that our failure to find prelexical decomposition effects in this study is related to particular characteristics of the disrupted presentation. For example, the fact that participants were required to respond directly to disrupted compounds in this experiment encouraged the employment of whole-word processing to "fill in the gaps". Therefore, in the final experiment in this study, below, we adapted the constituent disruption paradigm so that participants would be presented with disrupted constituents, but would provide responses to intact compound stimuli.

4. Experiment 3: Masked Priming

In Experiment 3, participants were presented with disrupted stimuli in the context of a masked priming paradigm. In the now classic version of this paradigm (see Forster, 1998) participants are presented with a forward mask (e.g., ########), followed by a prime word in lower case letters (e.g., *aircraft*), followed by a target word in uppercase letters (e.g., *AIRCRAFT*), over which subjects make a lexical decision. This paradigm targets the early phases of lexical processing because the masking of the prime allows for activation without participants being aware of the presence or nature of the prime. Since the prime and target share lexical but not visual characteristics, a faster response time to the target in the experimental condition (i.e., when prime and target are graphemically or morphologically related) as compared to a control condition (when prime and target are not related) indicates the effect of the prime upon the recognition of the target. The assumption is that faster recognition times indicate the influence of the early activation by the prime.

In our version of this paradigm, the priming conditions were identical to the presentations employed in Experiment 2., that is, letters of the prime words were replaced by hash marks (e.g., *a#rcraft* or *aircr#ft*). After presentation of this prime stimulus, participants were presented with the intact compound for lexical decision. Thus in contrast to Experiment 2, this paradigm did not measure whether the disrupted compound could be recognized, but rather whether the disrupted compounds would facilitate differentially in the subsequent recognition of 3-5, 4-4, and 5-3 compounds.

4.1 Method

4.1.1 Participants

Participants were 38 native speakers of English. They were all enrolled in undergraduate Linguistics courses at the University of Alberta and participated as volunteers.

4.1.2 Materials and Design

The core stimuli in this experiment were the identical 8-letter compounds employed in Experiment 2.

For each target compound (e.g., *AIRCRAFT*), there were four types of primes (a between-subjects condition): D1 primes (e.g., *a#rcraft*), D2 primes (e.g., *aircr#ft*), intact primes (e.g., *aircraft*), and unrelated control primes (e.g., *incident*). Thus, there were four lists of materials, each one containing all the 72 compounds with 18 trials of each prime-target type (i.e., D-1, D-2, intact, and control).

4.1.3 Procedure and Apparatus

Subjects were presented with a fixation point for 1000 ms, followed by a forward mask (#########) for 500 ms, followed by a prime word in lower case letters (e.g., *keyboard*, *k#yboard*) presented for 60 ms, and followed by a target word in uppercase letters (e.g., *KEYBOARD*) presented for 500 ms. The subjects' task was to press a button labeled 'yes' as quickly as possible if the target was a word, or press a button labeled 'no' as quickly as possible if it was not a word. The use of this lexical decision paradigm in this experiment, rather than the naming latency paradigm employed in Experiment 2, was made possible by the fact that participants were always responding to undisrupted stimuli. Thus the danger of a high lexical decision error rate, which motivated the use of naming latency in Experiment 2, was not present in this manipulation.

Participants were run individually in a dimly lit room. Stimuli presentation was done on a Macintosh PowerBook 520 attached to an external Apple monochromatic monitor and running PsyScope. Response times were collected on a CMU response box.

4.2 Results

Lexical response latencies were calculated for correct responses only and latencies greater than 2000 ms were removed from the analysis (4% of the data).

Response latencies showed a significant effect of priming for the disrupted vs unrelated primes ($t(48) = 5.6, p<.0001$) and a significant cost of disruption relative to the repetition condition ($t(48) = 4.8, p<.001$. The overall means, standard deviations and standard errors for these three priming conditions are presented in Table 3.

Table 3. Lexical decision latencies across the three main priming conditions

	Prime Types		
	Unrelated	Disrupted	Repeated
Mean	661	596	549
Standard deviation	78	65	55
Standard error	11	9	8

In Figure 3, the lexical decision latencies are broken down by compound and disruption type with repetition as the reference condition that corresponds to the intact presentation in Experiment 2.

As can be seen in Figure 3, the response pattern generally parallels that obtained in Experiment 2. There is an overall effect of disruption, relative to the repetition condition, yielding a main effect of presentation condition

($F(2,92)= 11.3$, $p<.0001$). No main effect of compound type was found ($F(2,46)=.19$, $p=.82$) and in the case of this experiment, the interaction of the two factors approached significance ($F(4,92)=2.3$, $p=.06$). As is evident in Figure 3, however, the locus of this interaction effect was in the difference between the 4-4 compounds and the other two types. The 3-5 and 5-3 compounds, which were predicted by the prelexical decomposition hypothesis to show differential priming effects, show the same pattern of lowest latencies in the repeated condition and highest latencies in the D2 condition.

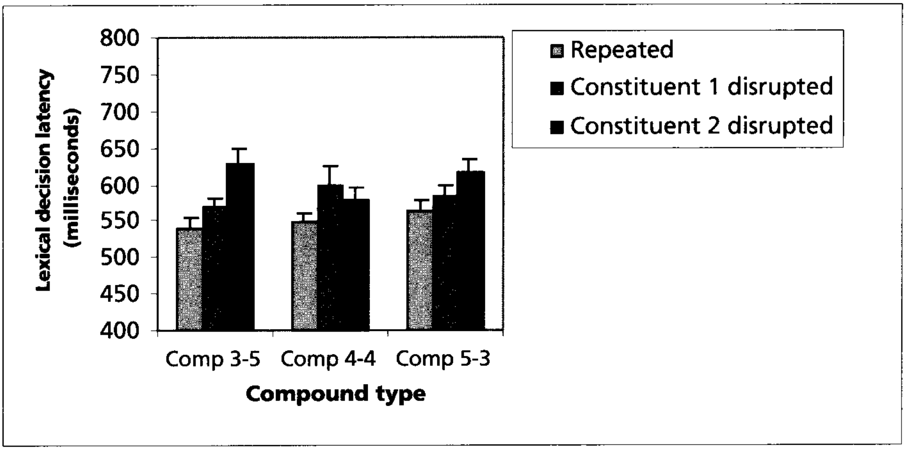

Figure 3. Lexical decision latencies (in milliseconds) in the disrupted masked priming task. Error bars represent standard errors for each category.

The final analyses conducted in this experiment focused on correlations of lexical decision latencies and constituent recoverability, positional family size, and positional family frequency. No correlations with magnitudes above .2 were found for any of these variables with either lexical decision response times or with disruption cost measures (disruption RT minus repetition RT; disruption RT minus control RT).

5. General Discussion

Our goal in this study was to investigate whether existing English compounds show evidence of prelexical morphological decomposition. The investigation of this phenomenon crucially required the development and testing of the new disruption paradigm presented here. In our view, this paradigm shows considerable promise in providing a new vantage point from which to understand how whole-word and constituent-based activation interact.

As we discussed at the outset of this paper, the traditional template for conceptualizing prelexical morphological decomposition derives from the processing of novel words which, by definition, can only be understood in terms of their constituents.

The results of this investigation, using both a naming latency paradigm and masked primed lexical decision paradigm do not support the view that whole word recognition relies on constituent recognition. Because these results derive from a new paradigm it is important to be cautious in interpretation. It remains possible that undetected stimulus or paradigm artifacts account for the observed pattern of results. Nevertheless, in the strict sense, this study has yielded evidence against prelexical morphological decomposition.

How could such results be reconciled with others, (e.g., Libben 1994, Libben and de Almeida 2002, Sandra 1990, Zwitserlood, 1994, Jarema et al. 1999) which have found evidence of constituent effects for semantically transparent compounds such as those used in the present experiments? Our answer to this question brings us back to the manner in which prelexical decomposition is to be understood and to the view that the novel word template for decomposition is wholly inappropriate for real words. It appears from our experiments that constituent morphemes are not the *basis* for real word recognition as they are for novel words. However, this does not preclude the possibility that they are activated prelexically. Indeed, if we compare our disruption paradigm with reference to more traditional priming paradigms, it seems that we have obtained an answer to the question of what happened when a constituent morpheme *is not available* prelexically. In contrast, studies that have found decomposition effects with traditional semantic and partial repetition priming paradigms have addressed the question of what happens when a constituent *is available* prelexically.

Taken together, the findings that prelexical activation helps but is not necessary seem most compatible with a dual process approach that posits both prelexical morphological decomposition with a generally faster whole-word recognition procedure. Under this architecture, the results of the experiments reported in our study may be characterized in the following manner: When constituent activation is disrupted, it doesn't *have the opportunity* to make the difference it otherwise would.

Note

* Authors are listed in alphabetical order. This research was supported by a Research Grant from the Social Sciences and Humanities Research Council of Canada (SSHRC) to the first author and by a Major Collaborative Research Initiative Grant from SSHRC awarded to Gonia Jarema (Co-Principal Investigator and Director), Eva Kehayia (Co-principal Investigator), and Gary Libben (Co-Principal Investigator). We wish to thank Uli Dressler, Gonia Jarema, and an anonymous reviewer for helpful comments on an earlier draft. We also want to thank Bruce Derwing, and the participants of the Compound Processing workshop held at the University of Vienna in March of 2000, for insightful discussions on

some of the data reported in this paper. We are grateful to Julia Peters, Shannon McEwen, Debra Elliot, and Mathias Schirmeier for their help with the preparation of the experimental materials.

Addresses of the Authors:

Roberto G. de Almeida
Department of Psychology
Concordia University
Montreal, Canada
almeida@alcor.concordia.ca

Gary Libben
Department of Linguistics
University of Alberta
Edmonton, Canada
gary.libben@ualberta.ca

References

Baayen, H., Dijkstra, T., & Schreuder, R. 1997. "Singulars and plurals in Dutch: evidence for a parallel dual route model." Journal of Memory and Language 36, 94-117.

Butterworth, B. 1983. "Lexical representation". In: B. Butterworth (ed.) Language Production 2. New York: Academic Press, 257-294.

Cohen, J., MacWhinney, B., Flatt, M., & Provost, J. 1993. "PsyScope: an interactive graphic system for designing and controlling experiments in the psychology laboratory using Macintosh computers." Behavior Research Methods, Instruments and Computers 25.2, 257-271.

de Jong, N. H., Feldman, L., Schreuder, R., Pastizzo, M. & Baayen, H. in press. "The processing of Dutch and English compounds: peripheral morphological and central orthographic effects." Brain and Language.

Forster, K.I. 1998. "The pros and cons of masked priming." Journal of Psycholinguistic Research 27, 203-233.

Jarema, G., Busson, C., Nikolova, R., Tsapkini, K., & Libben, G. 1999. "Processing compounds: A cross-linguistic study." Brain and Language, 68.1-2, 362-369.

Libben, G. 1994. "How is morphological decomposition achieved?" Language and Cognitive Processes 9, 369-391.

Libben, G. & de Almeida, R. G. 2002. "Is there a morphological parser?" In: S. Bendjaballah et al. (Eds.) Morphology 2000. Amsterdam: John Benjamins. 213-225.

Libben, G., Derwing, B. & de Almeida, R. G. 1999. "Ambiguous novel compounds and models of morphological parsing." Brain and Language 68, 378-386.

Libben, G. Gibson, M., Yoon, Y-B, & Sandra D. in press. "Compound fracture: The role of semantic transparency and morphological headedness." Brain and Language.

McQueen, J. & Cutler, A. 1998. "Morphology in word recognition." In: A. Spencer & A. Zwicky (eds.), The handbook of morphology. Oxford: Blackwell.

Sandra, D. 1990. "On the representation and processing of compound words: automatic access to constituent morphemes does not occur." Quarterly Journal of Experimental Psychology 42, 529-567.

Schreuder, R. & Baayen, H. 1995. "Modeling morphological processing." In: L.B. Feldman (Ed.), Morphological aspects of language processing, Hillsdale, NJ: Lawrence Erlbaum, 345-364.

Taft M. & Forster, K.I. 1975. "Lexical storage and the retrieval of prefixed words." Journal of Verbal Learning and Verbal Behavior 14, 630-647.

Zwitserlood, P. 1994. "The role of semantic transparency in the processing and representation of Dutch compounds." Language and Cognitive Processes 9, 341-368.

Appendix 1

Compound and non-compound stimuli

3-5 compounds	4-4 compounds	5-3 compounds	Non-compounds	
aircraft	stopsign	honeybee	medicine	bacteria
earphone	freehand	housefly	opposite	question
waxpaper	snowball	photolab	homicide	amethyst
nutshell	postcard	motorcar	caffeine	bachelor
armchair	doorknob	puppydog	capacity	carnival
haystack	firewood	pricetag	calendar	bandanna
daydream	bathroom	chickpea	material	vacation
inkpress	backbone	chairleg	geranium	marathon
gemstone	fishbowl	turbojet	petition	heritage
lipstick	footstep	nobleman	particle	hospital
oilfield	flagpole	sharkfin	province	velocity
teaspoon	headache	skullcap	ambition	prestige
warpaint	bookcase	waterboy	daughter	thousand
seashore	railroad	mousepad	accident	schedule
keyboard	shipmate	tabletop	cardigan	portrait
bugspray	lifetime	birthday	garrison	language
tapdance	glowworm	lunchbox	facility	industry
gaslight	workshop	driveway	mountain	universe
eyeglass	homeland	trashcan		
icestorm	raindrop	deathbed		
doghouse	hairline	saucepan		