*Delimiting the power of suppression in lexical processing: The question of below-baseline performance**

PAULA RUBIO FERNANDEZ

Abstract

This study aims to delimit the power of the mechanism of suppression by critically examining previous research on lexical processing as well as providing new empirical evidence of property activation in word interpretation. In particular, I have investigated to what level suppression dampens the activation of irrelevant conceptual information, whether to the baseline level of unrelated words or below. In order to answer this question, rarely addressed directly in the literature, I carried out a cross-modal lexical priming experiment investigating the activation of central properties of unambiguous nouns in neutral and cancelling contexts across three time delays (0, 400 and 1000ms). I conclude that in normal communication suppression does not operate below baseline at the lexical level, although different manipulations of the participants' attention could show below-baseline performance in certain laboratory tasks.

1 Introduction

It is a commonplace in the literature on lexical processing that priming effects are taken as a measure of word activation. In a simple lexical decision task, a word like *nurse* would be recognised faster after a lexical decision has been made on a word associate like *doctor* than after an unrelated word such as *bread* (e.g. Meyer & Schvaneveldt, 1971). This facilitating effect is usually interpreted as the result of a spreading activation process, according to which in processing the word *doctor* activation would have spread to conceptually associated words like *nurse*, speeding

*The study reported here is an expanded version of some of the experiments completed by the author towards her PhD degree at Cambridge University. This research was initially supported by an AHRB Postgraduate Award, and currently by a British Academy Postdoctoral Fellowship.

I would like to thank Richard Breheny and John Williams for supervising the setting and analysis of the experiments. Thanks also to Sam Glucksberg, Ira Noveck and Robyn Carston for their comments on earlier versions of this paper.

up their identification (Schvaneveldt & Meyer, 1973). This is the basis of the cross-modal priming paradigm widely used in early studies of lexical ambiguity resolution (e.g. Swinney, 1979; Onifer & Swinney, 1981; Tanenhaus, Leiman & Seidenberg., 1979; Tabossi, Colombo & Job, 1987; Tabossi & Zardon, 1993). In these experiments sentential contexts are presented in the acoustic modality and at the offset of an ambiguous prime in the sentence, a visual target related to one of its meanings is presented for a lexical decision or a naming task (e.g. *He found several spiders, roaches and other bugs in the corner of his room* - ANT). Meaning activation is measured in relation to reaction times to unrelated targets (e.g. SEW in the above example).

It seems reasonable that if facilitation is understood as positive activation, suppression would be understood as negative activation. This is supported by the ample empirical evidence that suppression can also spread in a manner analogous to that of activation (what is sometimes called an 'inhibition-based fan effect'; e.g. Neumann & Deschepper, 1992; Neumann, Cherau, Hood & Steinnagel, 1993). However, whereas facilitation with respect to the baseline level of unrelated words is taken as indicative of activation, suppression is not always measured in relation to the same threshold. In this respect, if the contextually irrelevant meaning of a homonym was to be suppressed, reaction times to associates of that meaning (e.g. SPY in the example above) would be similar to reaction times to unrelated controls, rather than being necessarily slower (e.g. Swinney, 1979; Seidenberg, Tanenhaus, Leiman & Bienkowski, 1982; Tanenhaus et al., 1979).

In short, priming is generally understood as positive activation above the baseline level of unrelated words. However, suppression is not always understood as negative activation involving below-baseline performance, with different methodologies showing different results (e.g. Onifer & Swinney, 1981; Neely, 1976; Simpson & Kang, 1994). I have therefore tried to investigate the power of the mechanism of suppression in lexical interpretation in the present study. I have adopted Gernsbacher's notion of suppression, understood as a general cognitive mechanism that actively reduces the activation of irrelevant information in language processing (Gernsbacher, 1990). This operative reduction of activation of conceptual information that is inconsistent with interpretation would in principle be different from passive decay for lack of contextual stimulation. Also, I distinguish suppression from inhibition on the grounds that the former would affect conceptual information that has been activated, whereas the latter would prevent conceptual information from getting activated. I will try to demonstrate the limits on the operation of the mechanism of suppression, not only by providing my own experimental results, but also by critically examining previous studies on lexical processing.

1.1 Suppression in previous studies of lexical processing

Lexical ambiguity is one of the paradigms where irrelevant information is supposed to be actively suppressed (see Gernsbacher & Robertson, 1999; cf. Simpson, 1981). The mechanism of suppression would operate in disambiguation processes by reducing the activation of the inappropriate meanings of ambiguous words. Studies using the cross-modal lexical priming paradigm have shown that the multiple meanings of an ambiguity get initially activated, although after a short delay only the contextually appropriate meaning remains active (Swinney, 1979; Tanenhaus et al., 1979; Seidenberg et al., 1982; cf. Tabossi et al., 1987). However, although the inappropriate meaning is supposed to be actively suppressed, its activation seems to drop to the baseline level of unrelated words but not below.

According to some views of suppression as a general cognitive mechanism (e.g. the Structure Building Framework; Gernsbacher, 1990), suppression would not only reduce the activation of inappropriate meanings of ambiguous words, but also that of less-relevant meaning properties of unambiguous words. Studies of lexical processing of monosemous words using the cross-modal priming paradigm (e.g. Greenspan, 1986; Whitney, McKay, Kellas & Emerson, 1985; cf. Tabossi, 1988) have shown that central and peripheral properties (e.g. COLD and SLIPPERY for *ice*, respectively) are activated initially regardless of contextual bias, although after a delay peripheral properties lose their activation in contexts biasing central properties (e.g. *The fresh meat was protected by the ice*). However, as in the case of inappropriate meanings of ambiguous words, the activation of irrelevant peripheral properties of unambiguous nouns does not drop below baseline.

The question of below-baseline suppression has been more directly investigated in lexical ambiguity studies looking at repetition effects. According to Gernsbacher and St. John (2002), experimental paradigms investigating the effect of repetition on disambiguation could show that suppression takes activation of inappropriate meanings below baseline. For example, Simpson and Kang (1994) asked participants to verify the grammaticality of sentences including a homonym. Target sentences were verified more slowly if they were preceded by a prime sentence biasing a different meaning of the homonym than if they came after an unambiguous control. The authors take this slower response time in the second presentation of the homonym as evidence of inhibition; that is, the inappropriate meaning would have been suppressed below baseline in the first presentation of the homonym, so it would take longer to retrieve it when appropriate in the second presentation.

In my view, rather than inhibition, Simpson and Kang's results could show a simple garden-path effect: in order to verify that a sentence like *The man had to have an organ transplant* makes sense, the BODY PART meaning of *organ* would have to be activated, selected and integrated in the mental representation of the sentence. In encountering the word *organ* again, participants would retrieve the BODY PART meaning first since this would be the most accessible meaning after processing the prime sentence. However, if the target sentence biases a different meaning than the prime sentence (e.g. *The keys of the organ were splintered and broken*), participants would still retrieve the meaning of *organ* that was most accessible after first processing the homonym. This implies that they would have to backtrack and select the other, appropriate meaning before verifying the target sentence, hence the longer reaction times.

This interpretation of Simpson and Kang's data seems to be corroborated by the fact that target sentences were verified faster after a prime sentence biasing the same meaning of the homonym than after an unambiguous control. This would indicate that, in this case, the appropriate meaning in the second presentation would have already been made highly accessible in processing the prime sentence, whereas in the control condition the same target sentence would have had to be disambiguated without any "guiding cues" from a previous sentence[1].

Gernsbacher, Robertson and Werner (2002) observed similar results in a series of experiments investigating the mechanism of suppression in disambiguation also using the sentence verification task in the context of repetition. However, the type of control sentences used in this extensive study makes it even more difficult to evaluate their results. Gernsbacher and her colleagues used either 'neutral' or 'nonsense' controls as primes for target sentences including the same homonym. In explaining the purpose of this type of control, the authors argue that no meaning of the ambiguous word would have been primed or suppressed in processing either the neutral or the nonsense controls, but this is surely not the case. It follows from the nature of the verification task that in order to respond positively ("it makes sense") to a neutral control like *They saw the match*, or negatively ("it doesn't make sense") to a nonsense control like *They prosecuted the match* participants have to assign some meaning to the homonym and therefore prime one of its meanings[2].

In my view, given the sentence verification task employed, the type of control sentences used by Gernsbacher and her colleagues would have failed to operate as a

---

[1] The sentence verification task used in Simpson and Kang (1994) must have had some effect on their results since Simpson and Adamopoulos (2002) did not observe such strong inhibitory effects when replicating the experiment with a self-paced reading task.

[2] Another problem with this kind of controls is that participants tend to make metaphorical interpretations of nonsensical sentences automatically (see Glucksberg, Gildea & Bookin, 1982; Keysar 1994).

proper baseline from which to make an accurate measure of activation or suppression in the critical conditions. In the case of suppression in particular, if the baseline reflected the activation of some meaning of the ambiguity, any relative measure of suppression of the other meaning would have been overestimated. Therefore, although the experiments reported in Gernsbacher et al. (2002) offer some measure of meaning activation, the inaccuracy of the baseline measure makes it difficult to evaluate the reliability of their results (see Rubio, 2005 for further discussion).

2 An on-line study of suppression

Given the inconclusive results of the above studies regarding the power of the mechanism of suppression in lexical interpretation, I carried out a cross-modal lexical priming experiment investigating the time-course of activation of central properties of monosemous nouns in two different types of sentential contexts, neutral and cancelling. Like other studies of lexical processing (e.g. Greenspan, 1985), the present study used the original methodology from studies of lexical ambiguity but with unambiguous words as primes.

Neutral sentences served as a type of baseline context where suppression of central properties would not be expected. According to certain models of cognition (e.g. Anderson, 1983), mental representations that are not stimulated by the context would passively decay. Moreover, according to even the most general views of suppression (e.g. Gernsbacher, 1990), this mechanism would not have any grounds to operate in neutral contexts as it would not have enough of a basis from which to select the conceptual information that may be irrelevant for interpretation. I therefore tested the activation of central properties in neutral contexts not making salient any particular aspect of the primes (e.g. *The expedition approached the territory of the cheetah* – FAST). This type of context would offer a basic pattern of activation from which to establish any relative measure of suppression.

The second type of context used in this study were cancelling sentences which make central properties inconsistent with the mental representation of the prime in the context. According to various views of suppression (e.g. Gernsbacher, 1990; Simpson & Kang, 1994), conceptual properties that are inconsistent with word interpretation would be actively suppressed. Therefore, central properties should be suppressed in cancelling contexts where these properties would clash with the mental representation of the prime concept (e.g. *In the final exam, the biology students had to dissect a cheetah* – FAST).

It should be interesting to investigate the time-course of activation of central properties of unambiguous nouns because, in processing the prime word, they

would get not only initially activated by an automatic process of spreading activation of associates, but also remain active by virtue of their strong association to the prime concept (see Barsalou, 1982; Rubio, Breheny & Wei Lee, 2003; Whitney et al., 1984). In this respect, the activation of central properties would have to be actively reduced by the mechanism of suppression in contexts where they are inconsistent with interpretation. In contrast, more peripheral properties would passively decay unless relevant in the context given their weak conceptual association to the primes (see Rubio, 2005; Rubio et al., 2003 for discussion).

The two different types of contexts used in this study would allow testing of two different hypotheses regarding the power of the mechanism of suppression. First, looking at the results from cancelling contexts, it should be possible to see whether suppression operates below-baseline in lexical interpretation, or rather reduces the activation of inconsistent conceptual information to the baseline level of unrelated words. Second, if there was a difference in the resulting levels of activation between passive decay and active suppression, it should be possible to see whether the activation of those properties that have decayed would be higher than that of those properties that have been suppressed. In this respect, the activation of central properties may drop to baseline level in neutral contexts but below baseline in cancelling contexts.

2.1 Method

2.1.1 Participants.

The participants in this experiment were 84 undergraduate students at Cambridge University who volunteered to take part in the experiment. They all had English as their first language. Each session lasted approximately 12min.

2.1.2 Materials and Design.

A set of 22 common nouns with predictable distinctive properties were selected as primes. However, rather than taking these properties from a dictionary definition of the critical nouns, I directly tested for property dominance by using two questionnaires based on the literature on prototypes and the internal structure of categories (Barsalou, 1983, 1985, 1987; Rosch, 1973; Rosch & Mervis, 1975). In Questionnaire 1, participants were presented with two tasks. In the first, they were asked to give a brief definition for each of the 22 words of a list (example given: "HOUSE: *Building where people live*"). In the second task they were asked to list what they thought were the distinctive characteristics of the concepts in the same list of words (example given: "How could you distinguish a WHALE from other animals? *Lives in the sea*, *large size*, *mammal*"). In Questionnaire 2, participants were presented with a free-association task where they were asked to write down

the first characteristic that came to their mind when reading the words on a list, which was the same as in Questionnaire 1.

After piloting the questionnaires on 15 participants (5 for Questionnaire 1 and 10 for Questionnaire 2), only some minor modifications needed to be made to the instructions. The final versions of the questionnaires were distributed among 65 participants (25 for Questionnaire 1 and 40 for Questionnaire 2). Having chosen a list of words with predictable highly distinctive properties, the results were as expected apart from two terms, *tip* and *spring*, which are ambiguous and did not show a homogeneous result. These latter were discarded. For each of the 20 remaining concepts, the most frequent distinctive property (what I refer to as 'central properties') was selected from the three different tasks.

A neutral sentence was then constructed for each one of the 20 primes, which was always the last word in the sentence. In order to avoid intra-lexical priming, sentences did not include any associate of the target word other than the critical prime. The 20 critical sentences were divided in two equal groups matched by the word frequency of the corresponding central properties (Johansson & Hofland, 1989). One group of sentences was paired with related central properties (e.g. *For the dinner Mary brought champagne* - BUBBLE). For the other group, the targets were scrambled, and so the sentences were paired with unrelated central properties (*My grandmother knew that old lullaby* - SMALL). The unrelated sentences and targets served as controls. Two lists of materials were constructed by pairing one group of sentences with related targets in the first list and with unrelated targets in the second list, and the other group of sentences with unrelated targets in the first list and with related targets in the second list. Another set of 20 neutral sentences was constructed and paired with English-like non-words (e.g. *Today John was late for his meeting* - WASK). Critical and filler sentences were randomised individually for each participant in each of the 2 lists of materials.

The same procedure was repeated with cancelling sentences, so one group of sentences was paired with related central properties (*Even though the bottle had been open for a week, John finished off the champagne* - BUBBLE) and the other with scrambled and therefore unrelated central properties (*The music teacher illustrated his point to the students by humming a lullaby* - YELLOW). Participants were randomly assigned to one Context Type, List and ISI, so each participant saw each sentence and the corresponding target only once.

Given that a word is identified at the point in time when information uniquely specifies it, which may actually occur before the physical ending of the word (Marslen-Wilson, 1987), for each of the 20 nouns in the materials a point was selected where the prime would be unequivocally recognised in relation to its neighbours. Targets were presented visually at the end of the acoustic signal 0, 400 or 1000ms after the word-recognition point selected for each prime. This enabled

accurate measuring of initial semantic activation, while controlling for the possibility that an early contextual effect may result from an early word-recognition followed by a fast property-selection given the length of the primes.

Sentences were recorded at a normal rate by a male speaker on an Apple Macintosh computer. The beginning and the end of each context were delimited using acoustic and visual control with the aid of a waveform sound editor. The word recognition point of each prime at the end of the context was also marked using a sound-editing program. The auditory stimuli and the visual targets were synchronised using a specialised computer program (SuperLab).

The experimental materials were preceded by two sets of practise trials. The first one consisted of a lexical decision task on a list of 10 words and 10 non-words that were presented visually one at a time and randomised separately for each participant. The second set of practice trials included both sentential contexts in the acoustic modality and visual targets for a lexical decision. This set of trials contained six neutral sentences similar to the critical ones, although the corresponding visual targets were not related to the primes in any of the practice trials.

2.1.3 Apparatus.

The experiment was conducted on a Toshiba laptop computer. The sentences were presented through a pair of headphones plugged into the laptop. The visual probes where presented in capital letters in the middle of the computer screen on a white background. Responses to the visual targets were made via a response box connected to the laptop. 'Word' responses were made with the thumb or the index finger of the right hand on the right-most key of the response box, which was a green key. 'Non-word' responses were made with the thumb or the index finger of the left hand on the left-most key of the response box, which was a red key. Target words remained on the screen until the participant made a lexical decision. There was a 1000ms delay between the offset of the visual target and the onset of the following acoustic context.

2.1.4 Procedure.

The experiment was presented to the participants as a simple psycholinguistic experiment investigating language processing. Participants were first given standard written instructions, which were then explained individually by the experimenter. Participants were told that they would be listening to a series of sentences through the headphones and that at the end of each sentence a string of letters would appear on the computer screen. They should try to indicate as fast and accurately as possible whether the string of letters was a word of English or not by pressing the corresponding key on the response box. It was emphasised that both

tasks, namely listening carefully to the sentences and making a fast lexical decision were equally important, although they should be taken as independent tasks. In order to avoid the possibility that participants might have divided their attention and tried to find some underlying coherent structure connecting the sentences, it was stressed that the sentences were unrelated and did not make a story. It was also indicated that non-words would not correspond in principle with words of other languages, but rather be orthographically similar to legitimate English words.

In order to arrive at the intended interpretation of the sentence, cancelling contexts needed to be processed more deeply than neutral contexts. In an attempt to make sure that participants processed cancelling sentences properly, the standard instructions were modified in the cancelling context condition such that participants were asked to try to visualise the meaning of the sentences they were listening to, in the same way that they would visualise the story if they were reading fiction.

Participants were tested individually. They ran through the two sets of practice trials with the experimenter and got appropriate feedback on their performance. When being tested on the critical materials, participants were left on their own in a closed room or cubicle.

To make sure that participants paid adequate attention to the priming sentences, a short memory test was given at the end of the experiment. Participants had been told about this memory test in the instructions. Three randomly chosen sentences from the critical set and another three from the filler set were included in the memory test. Another six sentences similar in style but different from any of the sentences used in the experiment completed the memory test. Two different memory tests were constructed, one for each type of context. Participants were instructed to tick the sentences that they thought they had listened to in the experiment. It was stressed that no change had been made to the original sentences.

2.2 Results

The minimum of correct responses required in the memory test was 2.5 standard deviations below the participants' average of correct responses per ISI. Only one participant was replaced because he failed to meet this criterion.

The mean response time, standard deviation and proportions of missing data for each Relatedness condition, together with the difference between the experimental (related) and the control (unrelated) conditions and its significance level per Context Type and ISI are presented in Table 1. A response time data point was treated as "missing" if it was either from an erroneous response or was over 2.5 standard deviations above the participant's average response time to the word targets in his exercise.

Table 1
Mean Reaction Times (in Milliseconds), Standard Deviation, Proportions of Missing Data and Level of Facilitation in Each Condition of the Experiment

		ISI		
Context	**Relatedness**	**0**	**400**	**1000**
Neutral	Related	595 (103, .07)	554 (54, .06)	557 (100, .03)
	Unrelated	634 (105, .03)	606 (62, .06)	565 (92, .11)
	Difference (R-U)	-39***	-52***	-8
Cancelling	Related	609 (115, .03)	569 (100, .03)	562 (76, .01)
	Unrelated	647 (152, .06)	581 (95, .06)	572 (94, .04)
	Difference (R-U)	-38*	-12	-10

*p<.05 **p<.01 ***p<.005

The statistical analysis of the results examined the effects of Target Relatedness (related/unrelated), ISI (0/400/1000ms) and Context Type (neutral/cancelling). Mean reaction times were entered into three-way analyses of variance (ANOVA), with both participants (F_1) and items (F_2) as random variables. Only the main effect of Relatedness was reliable[3], $F_1(1, 78)=26.2$, $MSE=1116.1$, $p<.001$; $F_2(1, 19)=6.38$, $MSE=7279.0$, $p<.03$. No significant interaction was observed, with the critical Relatedness x ISI x Context interaction approaching only a level of marginal significance, $F_1(2, 78)=1.88$, $MSE=1116.1$, $p>.1$; $F_2(2, 38)=2.26$, $MSE=1861.1$, $p>.1$.

A 2x3x2 (Relatedness x ISI x Context) ANOVA was also carried out on the arc-sine transformation of the missing data using participants as the random factor. The missing data was arc-sine transformed to stabilize variances (Winer, 1979). Only the main effect of Context was significant, $F_1(1, 78)=5.72$, $MSE=.019$, $p<.02$. Again, the critical Relatedness x ISI x Context interaction did not reach significance level, $F_1(2, 78)=2.27$, $MSE=.040$, $p>.09$. Overall, the average missing data was higher in the unrelated than in the related condition (.169 versus .120), so the facilitation observed could not have been due to missing data points lowering the average response time in the related condition.

Given that the activation level of central properties in neutral and cancelling contexts differed the most at the intermediate ISI (40ms), a 2x2 (Relatedness x Context) ANOVA was carried out on the reaction-time data corresponding to the 400ms ISI. No main effect was reliable. More important for the present

[3] Only reliable results (i.e. significant in both the analysis per subject and per item) will be reported, with the exception of the critical interactions, which will be reported in all analyses.

investigation, the critical Relatedness x Context interaction was significant, $F_1(1, 26)=8.10$, $MSE=726.80$, $p<.01$; $F_2(1, 19)=5.39$, $MSE=1776.3$, $p<.04$.

Another 2x2 (Relatedness x Context) ANOVA was carried out on the arc-sine transformation of the missing data corresponding to the 400ms ISI. In this analysis, no significant results were observed (all $F_1<1.4$). Overall, the average missing data was lower in the related than in the unrelated condition (.136 versus .169).

Separate analyses for the related and unrelated targets were also carried out on the reaction time data. Focusing first on related targets and considering neutral contexts as the baseline for cancelling contexts, reaction times were slower in the latter contexts across the three time delays (11ms on average). However, this difference was not significant at any ISI (all $t_1<1$). A 2x3 ANOVA did not reveal any reliable results, with the critical Context x ISI interaction not reaching significance level, $F_1(2, 78)=.032$, $MSE=8748.5$, $p>.9$; $F_2(2, 38)=.669$, $MSE=1403.8$, $p>.5$. The corresponding analysis of variance of the arc-sine transformation of the missing data only revealed a significant main effect of Context, $F_1(1, 78)=6.46$, $MSE=.025$, $p<.02$, the two-way Context x ISI interaction not being significant, $F_1(2, 78)=.447$, $MSE=.025$, $p>.5$.

In analysing reaction times to unrelated targets, I investigated the possibility that the priming effect observed had been the result of a disadvantage for unrelated targets, rather than facilitation of related targets[4]. For example, if we look at the data from the 400ms delay, where related targets were primed relative to unrelated targets only in neutral contexts, reaction times to unrelated targets were 25ms higher in neutral than in cancelling contexts. Because of the longer reaction times to both related and unrelated targets at the earliest delay, separate analyses were made for the data corresponding to the 0ms ISI, on the one hand, and the 400 and 1000ms ISIs on the other.

Since reaction times to related and unrelated targets were comparable at 1000ms in both neutral and cancelling contexts (9ms difference in average), I considered the longest delay as the baseline for the 400ms ISI. Reaction times to unrelated targets were higher at the intermediate than at the longest delay, with the greatest difference being observed in neutral rather than in cancelling contexts (41ms versus 9ms, respectively). Nevertheless, this difference was not significant in either type of context (both $t_1<1.5$). A 2x2 (Context x ISI) ANOVA did not reveal any reliable results, the critical two-way interaction being significant only in the analysis per item, $F_1(1, 52)=.496$, $MSE=7572.4$, $p>.4$; $F_2(1, 19)=5.05$, $MSE=1125.3$, $p<.04$. Likewise, the corresponding analysis of variance of the arc-transformation of the missing data did not reveal any significant results (all $F_1<2.1$).

---

[4] I would like to thank Sam Glucksberg for suggesting this possible interpretation of the results and the following statistical analyses.

Finally, in analysing reaction times to unrelated targets at the 0ms ISI, I made a comparison between the means of the first three trials and the last three trials for each participant[5]. If unrelated targets had been suppressed when participants noticed the prime-target relation in the critical trials, this inhibitory effect would only have shown at later trials. I therefore considered early trials as the baseline for late trials. Unrelated targets were facilitated in neutral contexts (-31ms) and suppressed in cancelling contexts (16ms), however this difference was not significant in either type of context (both t_1<1.3). A 2x2 (Context x Trial) ANOVA did not reveal any significant results, the critical two-way interaction not being significant, $F_1(1, 26)=1.08$, $MSE=7162.2$, $p>.3$.

2.3 Discussion

Central properties got initially activated in processing the prime words both in neutral and cancelling contexts, with related targets being facilitated relative to unrelated targets at 0ms. The initial activation of central properties would have been the result of an automatic process of spreading activation of associates. In both types of context, longer reaction times were observed at the earliest ISI to both related and unrelated targets. Given that targets were presented at the point in processing where the prime would have been unequivocally recognised, the lexical decision task would have overlapped in most trials with the end of the acoustic input. I therefore interpret the slower reaction times observed at the 0ms delay as the effect of a task-switching cost.

Whereas in neutral contexts central properties remained active at the intermediate delay, in cancelling contexts the activation of central properties dropped to baseline level by 400ms. Given that central properties are so strongly associated to the primes as to remain active in neutral contexts where they are not particularly relevant for interpretation, it is possible to assume that their loss of activation by 400ms in cancelling contexts would have been due to the operation of the mechanism of suppression. Nevertheless, this mechanism does not seem to operate below baseline at the lexical level since no below-baseline performance was observed at 400 or 1000ms in cancelling contexts. In neutral contexts, the loss of activation of central properties at the longest ISI would have been the result of passive decay, again to the baseline level of unrelated words.

The overall pattern of activation of central properties was not statistically different in neutral and cancelling contexts. However, their level of activation at 400ms varied significantly in the two types of context, indicating that the mechanism of suppression had operated on these associates in cancelling contexts,

[5] Since the experimental materials were randomised individually, it was impossible to carry out an analysis per item.

even if reducing their activation only to the baseline-level of unrelated words and not below.

When looking separately at reaction times to related and unrelated targets no significant results were observed. Although the highest reaction times to related targets were observed in cancelling contexts, this suppression of related targets relative to the baseline of neutral contexts was not significant. Therefore, this analysis did not reveal below-baseline suppression either.

Reaction times to unrelated targets did not differ significantly between 400 and 1000ms in both neutral and cancelling contexts. Unrelated targets could have been at a disadvantage if, in noticing the prime-target relation in critical trials, participants had developed expectations of a semantic relation which unrelated targets did not fulfil. However, the analysis of variance did not reveal a significant interaction, which could have indicated that the priming observed at the intermediate delay in neutral contexts was the result of slower reaction times to unrelated targets (rather than facilitation of related targets). Finally, the analyses of the data from the 0ms ISI did not show a significant difference between average reaction times to unrelated targets in early and late trials across contexts. If during their performance participants had developed a strategy to respond faster to related targets and correspondingly slower to unrelated targets, a significant difference should have been observed between early and late trials for unrelated targets. However, no significant disadvantage for unrelated targets was observed across trials in either neutral or cancelling contexts, ruling out that possible interpretation of the data.

In view of these results, it can therefore be confirmed that related targets were indeed facilitated wherever the level of priming was significant, and that they were not suppressed below baseline when inconsistent with word interpretation.

3 General Discussion

It seems reasonable to conclude from the present results that the difference between passive decay (neutral contexts) and active suppression (cancelling contexts) is not reflected in the power of the reduction of activation (i.e. to- versus below-baseline) but lies instead with the speed with which deactivation comes into play (i.e. 400-1000ms versus 0-400ms, respectively).

None of the cross-modal lexical priming studies reviewed earlier showed below-baseline performance (e.g. Swinney, 1979; Whitney et al., 1985). According to these studies, both the activation of contextually inappropriate meanings of ambiguous words and that of peripheral properties of unambiguous words decreased at some point during processing but only to the baseline level of

unrelated words, not below. Moreover, I have argued that off-line studies of lexical ambiguity using the sentence verification paradigm in the context of repetition did not necessarily show below-baseline suppression but rather a garden-path effect or an inflated measure of suppression at the expense of facilitation (e.g. Simpson & Kang, 1994; Gernsbacher et al., 2002). The present study using the cross-modal lexical priming paradigm did not show below-baseline performance neither in neutral contexts nor in cancelling contexts where suppression would be expected. A possible explanation of these data is that for some methodological reason, the cross-modal priming paradigm may not be able to detect below-baseline suppression. However, I would argue that this is not necessarily the case.

Below-baseline suppression has been observed using the lexical decision task employed in the cross-modal priming paradigm (e.g. Marcel, 1980; Lorch, Balota & Stamm, 1986; Neely, 1976). Rather than sentential contexts in the acoustic modality, these studies used single-word contexts so letter strings were presented one after the other for a lexical decision, some of them being critically related (e.g. TREE > PALM > WRIST). Moreover, other studies using sentential rather than single-word contexts have also shown below-baseline performance with the lexical decision task (e.g. Kellas, Paul, Martin & Simpson, 1991; Forster, 1981). However, rather than in the acoustic modality, these experiments presented the sentential contexts in the visual modality one word at a time. In view of the negative activation observed in these studies, it seems that the cross-modal priming paradigm should in principle be able to detect below-baseline suppression.

I would argue that the critical difference between experiments showing to- and below-baseline suppression using the lexical decision task would lie with the amount of attentional resources allocated to each word in the context. Whereas with the cross-modal priming paradigm sentential contexts are presented in the acoustic modality at a normal speech rate, studies using single-word contexts and sentential contexts displayed visually one word at a time would allow participants to focus more on each context word. Therefore, it is possible that the power of the mechanism of suppression at the lexical level may be related to the quantity of attentional resources available for processing a given lexical item.

The relation between suppression and attention seems to be confirmed by research on selective attention, where stimuli that are deliberately ignored exert an inhibitory effect on associated information (Neill, Valdes & Terry, 1995; Tipper, 1985; cf. Damian, 2000). In all the priming studies reviewed up to this point, the prime word was deliberately attended to. However, in experiments manipulating the participants' attention so that the prime word was deliberately ignored, the prime exerts on its associates an inhibitory rather than a facilitatory effect (e.g. Yee, 1991; Dagenbach, Carr & Wilhelmsen, 1989; Neely, 1977; Dagenbach, Carr & Barnhardt, 1990; Fuentes & Tudela, 1992).

Irrelevant conceptual information, even if associated to a stimulus that is under the focus of attention, can also have a negative priming effect when discarded; what is called 'endogenous negative priming' (Neill et al., 1995). An example of this type of priming would be the negative effect of the inappropriate meaning of a homonym, which as a prime would be attended to even if the associated contextually-inappropriate meaning would be ignored[6]. The fact that selective attention can have an inhibitory effect on unattended information would offer support to my interpretation of the different results observed with the lexical decision task and the conclusion that the power of the mechanism of suppression at the lexical level would be related to the amount of attentional resources available for each processed word.

Simpson and Adamopoulos (2002), for example, observed much stronger inhibitory effects when testing words in isolation than when using a self-paced reading task. The authors argue that for the sake of parsimony, the same mechanism should underlie the inhibitory effects observed in both experiments. Although I share their view of a single mechanism of suppression operating in lexical processing, in my view the power of this cognitive mechanism would not be the same in their two experiments so different results would actually be expected. The same difference is sometimes observed with positive priming, where facilitatory effects are observed when testing words in isolation but not when testing them in sentential contexts (Williams, 1992; Tabossi, 1991).

Extrapolating from the experimental to the natural setting, in oral communication the amount of attentional resources devoted to each word in a sentence would be comparable to that available in the cross-modal priming paradigm, just as in normal reading less attention would be allocated to each word compared to those laboratory tasks which allow focusing on each word of the context. I would therefore like to conclude that in normal communication the mechanism of suppression does not necessarily operate below baseline at the lexical level. This would imply that, although a wide range of experimental tasks manipulating the participants' attention exhibit the inhibitory effect of actively ignored stimuli (see Neill et al., 1995 for an exhaustive review), negative lexical priming may not be so common in normal language processing.

For instance, in a previous study testing the activation of superordinates in metaphoric contexts (e.g. *John doesn't like physical contact, he is a <u>cactus</u>* – PLANT), I observed below-baseline suppression using the cross-modal lexical priming paradigm. However, instead of using a pen and paper memory test at the

---

[6] Endogenous negative priming resulting from selective attention offers an alternative interpretation of the suppression of contextually-inconsistent conceptual information, which I have explained as the result of the operation of a cognitive mechanism. However, it follows from the present discussion that these two explanations of the data are not incompatible.

end of the exercise, the secondary task in this experiment was to paraphrase the nominal metaphor aloud after making the lexical decision (see Rubio, 2002). Assuming that in paraphrasing an expression the words in the expression must be avoided, the below-baseline suppression observed would have been an instance of negative priming. However, the number of occasions where people are requested to paraphrase an expression is minimal compared to the number of times they encounter a metaphor. Therefore, negative priming may not be so frequent in normal language processing, the mechanism of suppression requiring extra attentional resources to operate with such power at the lexical level.

4 Conclusions

The results of the present study have shown that the difference between a central conceptual property decaying passively and it being actively suppressed does not lie with the power of the reduction of activation but rather depends on the point in processing when the property loses its initial activation. Both in neutral contexts where suppression would not have enough grounds to operate, and in cancelling contexts where central properties would be inconsistent with word interpretation, the activation of central properties dropped to the baseline level of unrelated words, not below. However, these strong associates lost their initial activation in neutral contexts between 400 and 1000ms from the word-recognition point of the prime, whereas they were actively suppressed in cancelling contexts between 0 and 400ms. More importantly, the activation of central properties was significantly different in the two context types at the intermediate delay.

I have also argued that the power of the mechanism of suppression on irrelevant lexical information is directly related to the amount of attentional resources assigned to the particular lexical item. Experimental tasks allowing participants to focus their attention on each word of a sentence show that the mechanism of suppression can operate below baseline in lexical processing. However, other tasks allocating less processing resources per word, such as the cross-modal lexical priming paradigm where sentential contexts are presented in the acoustic modality at a normal speech rate, show that suppression reduces the activation of irrelevant lexical information to the baseline level of unrelated words but not below. I have therefore concluded that in normal communication the mechanism of suppression does not generally operate below baseline at the lexical level, negative priming not being a common phenomenon in normal language processing compared to other laboratory tasks manipulating the participants' attention.

References

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory and Cognition, 10*, 82-93.

Barsalou, L. W. (1983). Ad hoc categories. *Memory and Cognition, 11*, 211-227.

Barsalou, L. W. (1985). Ideals, central tendency and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*, 629-649.

Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Niasser (Ed.), *Concepts and Conceptual Development: Ecological and intellectual factors in categorisation*. Cambridge: Cambridge University Press.

Dagenbach, D., Carr, T. H., & Barnhardt, T. M. (1990). Inhibitory semantic priming of lexical decisions due to failure to retrieve weakly activated codes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 328-340.

Dagenbach, D., Carr, T. H., & Wilhelmsen, A. (1989). Task-induced strategies and near-threshold priming: Conscious influences on unconscious perception. *Journal of Memory and Language, 28*, 412-443.

Damian, M. F. (2000). Semantic negative priming in picture categorization and naming. *Cognition, 76,* 45–55.

Forster, K. I. (1981). Priming and the effects of sentence and lexical contexts on naming time: Evidence for autonomous lexical processing. *Quarterly Journal of Experimental Psychology, 33A*, 465-495.

Fuentes, L. J., & Tudela, P. (1992). Semantic processing of foveally and parafoveally presented words in a lexical decision task. *Quarterly Journal of Experimental Psychology, 45A*, 299-322.

Gernsbacher, M. A. (1990). *Language comprehension as structure building*. Hillsdale NJ: Lawrence Erlbaum.

Gernsbacher, M. A., & Robertson, R. W. (1999). The role of suppression in figurative language comprehension. *Journal of Pragmatics, 31*, 1619-1630.

Gernsbacher, M. A., Robertson, R. R. W., & Werner, N. K. (2002). The costs and benefits of meaning. In D. S. Gorfein (Ed.), *On the Consequences of meaning selection.* Washington, DC: APA Publications.

Gernsbacher, M. A., & St. John, M. F. (2002). Modelling the mechanism of suppression in lexical access. In D. S. Gorfein (Ed.), *On the consequences of meaning selection*. Washington, DC: APA Publications.

Glucksberg, S., Gildea, P., & Bookin, H. B. (1982). On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behaviour, 21*, 85-98.

Greenspan, S. L. (1986). Semantic flexibility and referential specificity of concrete nouns. *Journal of Memory and Language, 25*, 539-557.

Johansson, S., & Hofland, K. (1989). *Frequency analysis of English vocabulary and grammar based on the Lob corpus: Vol. 1, Tag frequencies and word frequencies*. Oxford: Clarendon Press.

Kellas, G., Paul, S. T., Martin, M., & Simpson, G. B. (1991). Contextual activation and meaning access. In G. B. Simpson (Ed.), *Understanding Word and Sentence*. North Holland: Elsevier Science Publishers.

Keysar, B. (1994). Discourse context effects: Metaphorical and literal interpretations. *Discourse Processes, 18*, 247-269.

Lorch, R. F., Balota, D. A., & Stamm, E. G. (1986). Locus of inhibition effects in the priming of lexical decisions: Pre- or post-lexical access? *Memory and Cognition, 14*, 95-103.

Marcel, A. J. (1980). Conscious and pre-conscious recognition of polysemous words. In R. Nickerson (Ed.), *Attention and performance VIII.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Marslen-Wilson, W. D. (1987). Functional parallelism in spoken word-recognition. *Cognition, 25,* 71-102.

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

Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory and Cognition*, *4*, 648-654.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited capacity attention. *Journal of Experimental Psychology: General, 106*, 226-254.

Neill, W. T., Valdes, L. A., & Terry, K. M. (1995). Selective attention and the inhibitory control of attention. In F. N. Dempster & C. J. Brainerd (Eds.), *Interference and inhibition in cognition*. San Diego, CA: Academic Press.

Neumann, E., Cherau, J. F., Hood, K. L., & Steinnagel, S. L. (1993). Does inhibition spread in a manner analogous to spreading activation? *Memory, 1*, 81-105.

Neumann, E., & Deschepper, B. (1992). An inhibition-based fan effect: Evidence for an active suppression mechanism in selective attention. *Canadian Journal of Psychology, 46*, 1-40.

Onifer, W., & Swinney, D. A. (1981). Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias. *Memory and Cognition, 9*, 225-226.

Rosch, E. (1973). On the internal structure of perceptual and semantic categories. In T. E. Moores (Ed.), *Cognitive development and the acquisition of language*. New York: Academic Press.

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 4*, 573-605.

Rubio Fernández, P. (2002). The inhibition of core features in metaphor interpretation. *RCEAL Working papers in English and Applied Linguistics*, *University of Cambridge, 8*, 73-100.

Rubio Fernández, P. (2005). Pragmatic processes and cognitive mechanisms in lexical interpretation: The on-line construction of concepts. PhD dissertation, University of Cambridge.

Rubio Fernández, P., Breheny, R., & Wei Lee, M. (2003). Context-independent information in concepts: An investigation of the notion of 'core features'. In F. Schmalhofer, R. M. Young & G. Katz, *Proceedings of EuroCogSci 03: The European Cognitive Science Conference 2003*. London: Lawrence Erlbaum.

Schvaneveldt, R. W., & Meyer, D. E. (1973). Retrieval and comparison processes in semantic memory. In S. Kornblum (Ed.), *Attention and Performance IV*. New York: Academic Press.

Seidenberg, M. S., Tanenhaus, M. K., Leiman, J .L., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive Psychology, 14*, 470-519.

Simpson, G. B. (1981), Meaning dominance and semantic context in the processing of lexical ambiguity. *Journal of Verbal Learning and Verbal Behavior*, 20, 120-136.

Simpson, G. B., & Adamopoulos, A. C. (2002). Repeated homographs in word and sentence contexts: Multiple processing of multiple meanings. In D. S. Gorfein (Ed.), *On the consequences of meaning selection*. Washington, DC: APA Publications.

Simpson, G. B., & Kang, H. (1994). Inhibitory processes in the recognition of homograph meanings. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in Attention, Memory, and Language*. San Diego, CA: Academic Press.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*, *18*, 645-659.

Tabossi, P. (1988). Effects of context on the immediate interpretation of unambiguous nouns. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *14*, 153-162.

Tabossi, P. (1991). Understanding words in context. In G. B. Simpson (Ed.), *Understanding Word and Sentence*. North Holland: Elsevier Science Publishers.

Tabossi, P., Colombo, L., & Job, R. (1987). Accessing lexical ambiguity: Effects of context and dominance. *Psychological Research, 49*, 161-167.

Tabossi, P., & Zardon, F. (1993). Processing ambiguous words in context. *Journal of Memory and Language*, *32*, 359-372.

Tanenhaus, M. K., Leiman, J. M., & Seidenberg, M. S. (1979). Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior*, *18*, 427-440.

Tipper, S. P. (1985). The negative priming effect: Inhibitory priming and ignored objects. *Quarterly Journal of Experimental Psychology, 37A*, 571-590.

Whitney, P., McKay, T., Kellas, G., & Emerson, W. A. (1985). Semantic activation of noun concepts in context. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *11*, 126-135.

Williams, J. (1992). Processing polysemous words in context: Evidence for interrelated meanings. *Journal of Psycholinguistic Research*, *21*, 193-218.

Winer, B. J. (1971). *Statistical Principles in Experimental Design*. New York: McGraw-Hill.

Yee, P. L. (1991). Semantic inhibition of ignored words during a figure classification task. *Quarterly Journal of Experimental Psychology*, *43A*, 127-153.