Eye-movements and ERPs reveal the time-course of processing negation and remitting counterfactual worlds

Heather J Ferguson[1 2]

Anthony J Sanford[1]

Hartmut Leuthold[1]

[1]University of Glasgow, Scotland, UK

[2]University College London, England, UK

Number of Pages: 47

Number of Figures: 7

Correspondence to:

Heather Ferguson email: h.ferguson@ucl.ac.uk

Tel: +44 (0)20 7679 4046

Research Department of Linguistics

University College London

Chandler House

2 Wakefield Street

London WC1N 1PF

England, UK

Abstract

The ability to update our current knowledge using contextual information is a vital process during every-day language comprehension. To understand a negated statement, readers are required to cancel real-world expectations, but are not explicitly provided with an alternative model. Thus, the question of how and when a negative context influences interpretation of later events arises. We report one eye-movement study (Exp.1) and one ERP study (Exp. 2) investigating the effects of negation on discourse processing. Prior context depicted a real-world (RW), or negated-world (NW), while a second sentence was manipulated to create RW anomalous continuations, where events included a violation of RW knowledge, and RW congruent continuations, where the events described were congruent with RW knowledge. Results from Experiment 1 showed that the negated discourse context did not influence initial processing of the target sentence, as reflected in participants' eye movement behaviour. Similarly, Experiment 2 revealed that the typical N400 effect to semantic violations has not been reversed by introducing a negated-world context. However, in later processing, Experiment 1 demonstrated that the negated-world context is eventually incorporated into the representation of the sentence meaning. Thus, we suggest that discourse does not always have an immediate effect on language comprehension and discuss the results in terms of a variety of accounts of representing negation.

Section: Cognitive and Behavioural Neuroscience

Keywords: Language comprehension, Negation, Discourse processing, Eye-tracking, ERP, N400

1. Introduction

Negation has attracted a large amount of research dating back to the 1960's and 1970's (e.g. Arroyo, 1982; Carpenter & Just, 1975; Clark & Chase, 1972; Eiferman, 1961; Gough, 1965; 1966; Just & Carpenter, 1971; Wales & Grieve, 1969; Wason, 1961). Most of this work has examined the effect of negation on local sentence comprehension and has largely relied on response-based tasks (e.g. sentence verification or sentence completion tasks) that do not lead to a clear account on the mechanisms and constraints associated with representing negated information. In the studies presented here, we have used on-line measures (eye-movements and event-related brain potentials (ERPs)) that have allowed a potentially exciting look into the unfolding process of comprehending negation. These techniques are advantageous because they provide a great deal of multidimensional information about the temporal dynamics and specific cognitive functions involved in language processing. Therefore using these techniques, we anticipated gaining detailed information about the time course of processing negation.

In this paper we consider the processing consequences for conditional statements where a prior context sentence negates real-world expectations, but does not explicitly set up an alternative scenario, such as in (1).

(1) If cats were not carnivores…

In this example, participants must cancel real-world knowledge that cats *are* carnivores for a full understanding of the utterance. According to the model theories of language comprehension, readers represent the states of affairs described in a text as mental models (Johnson-Laird; 1983). However, this raises two important processing questions when applied to a statement such as in (1), (i) how quickly is real-world knowledge negated and, (ii) when is an alternative model created? Negation also creates an interesting comparison with counterfactuals that explicitly state an ‘alternative world’, thus implying that (at least some aspect of) reality is no longer relevant. So, to make sense of a counterfactual utterance, such as, *If cats were vegetarians*, readers must accept the ‘alternative world’ of cats being vegetarians. However, to do this, it is necessary that, at some level, the comprehender suppresses their real-world knowledge that cats are carnivores.

We have examined these processing mechanisms by studying anomalies, which have provided a useful tool to investigate the integration of new linguistic input into existing representations. For example, in a sentence such as, *The girl drank her tea from a book* the features of a book are inappropriate for the concept of drinking tea. That such an anomaly is obviously noticed is indicated by eye tracking research in which the ease or difficulty associated with processing an incoming word has been shown to be accurately reflected in the amount of time that readers spend looking at that word (e.g. Braze, Shankweiler, Ni & Palumbo, 2002; Ni, Fodor, Crain & Shankweiler, 1998; Rayner, 1998). Similarly, the processing of words that are pragmatically anomalous or unexpected, as compared to normal or expected, has been found to trigger a larger N400 component, a centroparietally distributed, negative-going deflection in the ERP with a peak at about 400 ms (see Kutas, van Petten, & Kluender, 2006, for a review). Importantly for present purposes, eye-movements and the N400 are also influenced by discourse-dependent context anomalies, with effects similar in surface form to pragmatic anomalies, showing that anomaly detection is not confined to the constraints of the local sentence (Joseph, Liversedge, Blythe, White, Gathercole & Rayner, in press; Rayner, Warren, Juhasz & Liversedge, 2004; Van Berkum, Brown, Zwitserlood, Kooijman & Hagoort, 2005; Van Berkum, Hagoort & Brown, 1999; Van Berkum, Zwitserlood, Hagoort & Brown, 2003).

However, successful language comprehension does not necessarily require information to map directly onto our real-world expectations. A growing collection of evidence suggests that a strong discourse context can overrule local lexical-semantic factors when these two conflict and as such can immediately influence comprehension processes. For example, Ferguson and Sanford (2008; Ferguson, Sanford & Leuthold, 2007) measured eye-movements and ERPs during reading to investigate whether the presence of a counterfactual context would eliminate the effects of a statement that was anomalous with respect to the real world, but which fits the counterfactual context. Participants read short passages where a context sentence introduced a counterfactual-world (CW), or real-world (RW) situation, while a second sentence was manipulated to create RW anomalous continuations, where events included a violation of RW knowledge, and RW consistent continuations, where the events described were congruent with RW knowledge. Results from both eye tracking and ERPs (N400) showed that typical effects from RW violations can be 'neutralised' within an appropriate pre-specified counterfactual context, and RW-congruent items can lead to the experience of an anomaly following an inconsistent CW context. Other studies have used descriptive discourse to create contexts that normalise violations of real-world principles (Filik, 2008; Filik & Leuthold, 2008; Nieuwland & Van Berkum, 2006). In these cases, readers happily accept events that violate their real-world expectations, when they are acceptable given the prior context. Thus, taken together, these studies provide strong evidence that discourse context can apparently completely overrule local semantic violations.

To date, many experimental investigations of negation support a view where the overall meaning of a negated discourse is computed at some stage of processing using non-linguistic representations to establish important elements of the negated entity. Thus, comprehending a text could be considered to be experiential in nature as mental simulations represent the described state of affairs. Notably, Kaup and colleagues have carried out a number of investigations to examine the question of how negated concepts, and their consequences, are represented in a non-linguistic system (Kaup, in press; Kaup, Yaxley, Madden, Zwaan & Lüdtke, 2007; Lüdtke, Friedrich, De Filippis & Kaup, 2005). Participants read sentences such as, 'There was no eagle in the sky/ nest', and then decided whether the object displayed in a picture, which followed sentence offset after a brief delay (e.g., 250 ms), was mentioned in the sentence. Results showed that responses were faster when the picture matched the shape of the object implied by the negated statement (e.g. eagle with folded wings for 'no eagle in the sky'), suggesting that the negated situation has been visually simulated during language processing.

ERP studies on negation employing a sentence verification paradigm have investigated effects of negation on the N400 brainwave (Fischler, Bloom, Childers, Roucos & Perry, 1983; Hald, Kutas, Urbach & Pahrhizkari, 2004). Participants read sentences such as in (6).

(6) a.	Hawaii is tropical.	(True affirmative)
b.	Hawaii is cold.	(False affirmative)
c.	Hawaii is not cold.	(True negative)
d.	Hawaii is not tropical.	(False negative)

Sentences (6b) and (6c) elicited a larger N400 effect, than sentences (6a) and (6d). One interpretation of this N400 effect is that participants have not used the negation operator to make sense of negative sentences. Furthermore, these studies show that the semantic meaning of a word and its fit with the preceding context (e.g. Hawaii – tropical) plays a more immediate role in comprehension than the truth value of the sentence as a whole. Thus, semantic plausibility of the sentence as a whole has not influenced initial processing, although it is suggested that the sentence's truth value is computed at a later stage of comprehension. In fact, a recent study by Lüdtke, Friedrich, De Filippis and Kaup (2005) demonstrated that when the delay between sentence offset and picture presentation was long (1500 ms) in a sentence-picture-verification task (see above for details), a main effect of negation and truth value emerges. In the N400 time window the negation effect was indicated by a more negative-going ERP waveform for negative than affirmative conditions over posterior electrodes, whereas the effect of truth value showed up over frontal electrodes. These ERP effects were absent when the delay was short (250 ms). The authors suggest that after a long processing delay, comprehenders have finalized the meaning of the negated sentence and have mentally represented this complete meaning. However, these ERP studies (Fischler et al., 1983; Hald et al., 2004, 2005) are limited due to confounding lexical priming effects between the sentence subject (e.g. *Hawaii*) and target word (e.g. *hot*) as a result of the close proximity of these words. Also, the passages used in these studies were very short, with the target word immediately following the negation marker 'not' and as such it is possible that participants have not had time to integrate the negative context. It is also important to note that to date, at least to our knowledge, no studies have investigated the consequences of negation on later linguistic processing outwith the negated sentence.

One view of how negatives are processed is that they serve to reduce accessibility of information mentioned within its scope during language comprehension, and *nothing more* (see Prado & Noveck, 2006; Von Klopp, 1993). Evans and colleagues (Evans, Clibbens & Rood, 1996; Evans & Over, 2004) describe the role of negation for processing abstract conditionals as "to deny presuppositions rather than to assert information" (Evans et al., 1996, p.394). Some support for this comes from MacDonald and Just (1989) who employed a probe recognition or word-naming task following sentences such as (7) and revealed evidence that readers create a propositional representation where a negative construct (e.g. 'no') encapsulates the explicitly negated concept ('cookies'), thus reducing later accessibility of this information. This will be referred to as the *narrow-view* account.

(7) Almost every weekend Mary bakes some bread but no cookies for the children.

However, contrasting evidence demonstrates that this suppression of negated information is not an obligatory process and instead suggests that negation processing can be viewed as a guide for making the negated object the basis of a wider *search-for-alternatives*. Specifically, McDonald and Just (1989) found that when probe words were close associates (e.g. butter), suppression was not evident following *no bread*. In fact, a great number of studies have indicated that although initially processing is insensitive to negation, effects of negation may emerge later in language comprehension (see Giora, 2006, for a review). For example, Kaup, Lüdtke and Zwaan (2006) showed that while negated concepts lost accessibility 750 ms after their offset, by 1500 ms they were replaced by alternative opposites. Thus, after this long delay, focus shifted from the negated concept (*The door is not open*) to an alternative ('The door is closed'). However, the exact time course of this shift from the negated concept to an alternative interpretation is still under debate (cf. Hasson & Glucksberg, 2006).

So far then, research into negation has provided evidence to support a suppression account, in which suppression is limited by factors such as the nature and truth value of the negated concept, the processing time allowed and the availability of contextual information. Most research has compared affirmative and negative sentences in isolation (Hasson & Glucksberg, 2006; Just & Carpenter, 1971; Kaup et al., 2006; McDonald & Just, 1989; Wason, 1961), or following a supportive context (e.g. Giora, Fein, Aschkenazi & Alkabets-Zlozover, 2007; Glenberg, Robertson, Jansen & Johnson-

Glenberg, 1999; Wason, 1965) and as such cannot address the issue of how a negated concept is reflected on-line in later processing of relevant information.

The studies reported here provide an investigation of the consequences of negation within a wider discourse, using eye-tracking and ERP methodologies. Specifically, negation was used to cancel real-world expectations within a conditional statement, such as (8). This was then followed by a second sentence with either RW anomalous continuations, where events included a violation of RW knowledge (though consistent with the negated context), or RW congruent continuations, where the events described were congruent with RW knowledge (inconsistent with negated context). Thus, any differences in reading at sentence 2 must reflect processing in terms of real-world knowledge or the negated-world context.

(8) If cats were not carnivores, they would be cheaper for owners to look after. Families could feed their cat a bowl of carrots/ fish and listen to it purr happily.

This paradigm eliminates many limitations from previous studies. Since negation is depicted in a separate context sentence, the lengthy delay before the consequence in sentence 2 should allow alternative opposites to be activated if necessary. Similarly, by examining effects in the second sentence, difficulties arising directly from the extra syllable in negative sentences are eliminated (cf. Clark & Chase, 1972). Finally, priming effects are avoided here in two ways. First, the negative context negates a category (e.g. carnivores) while later reference is made to a specific example (e.g. fish), thus avoiding obvious confounds from simply repeating a previously negated word. Second, the negated entity (e.g. *not carnivores*) is presented in a separate context sentence to increase distance between negated and critical words which will reduce the impact of low-level lexical priming.

What can be predicted about negated passages in this paradigm, in line with previously detailed research on anomaly detection? If it is true that comprehenders suppress access to the negated concept in their representation as implied by the *narrow-view* (e.g., Evans et al., 1996; Evans & Over, 2004), then greater difficulty should emerge for target words that were negated by the context sentence. Thus, if comprehending a sentence such as (8) involves suppression of *carnivore*, then this should be reflected in lengthened reading times and increased regressions as well as a larger N400 following *fish*, as compared to a RW context condition. Also, according to this *narrow-view*, readers should show equal difficulty, as indicated by eye-movement measures and N400 amplitude, when the negated context is followed by *carrots*, since this account asserts that a search for alternatives is not initiated as a function of negation. In contrast, the *search-for-alternatives* view (e.g., McDonald & Just, 1989) would predict for the NW condition that *carrots* should be easier to process than *fish* since it is an appropriate alternative to the negated concept. This is the basic question of the experiments reported here.

To test these predictions, we conducted two studies in which the materials manipulate propositions that are anomalous with respect to the real-world, but appropriate given a pre-specified negated world. The question was whether and at what stage of processing the negative context would influence interpretation of the target sentence, making the anomaly acceptable. Also, whether processing of RW congruent information can lead to difficulties when it has been preceded by a negative context. Experiment 1 employed eye-tracking methodology as participants read a prior context sentence depicting a negated-world (NW) or a real-world (RW), followed by a second sentence continuation that described events that were consistent or inconsistent with the prior RW or NW context. In Experiment 2, ERPs were recorded as participants read

NW-inconsistent and NW-consistent passages to provide converging evidence regarding the time-course of negation processing.

2. Results

2.1 Experiment 1

2.1.1 Regions of analysis

The critical second sentence was divided into four regions for the purpose of aggregating reading times and classifying eye-movements, as shown in (9). For each sentence frame, corresponding regions contained the same number of words in all four versions.

(9) If cats were not carnivores they would be cheaper for owners to look after. |Families could$_1$| feed their cat a bowl of$_2$| carrots and$_3$| it would gobble it down happily.$_4$|

The first sentence created a RW or NW context for the critical sentence. Region 1 introduced an agent followed by the modal verb *could*. Region 2 (pre-critical) contained material that led to the critical region. Region 3 (critical) always consisted of the inconsistent or consistent noun, plus the connective *and*, and thus this region was always two words long. Region 4 (post-critical) comprised the critical sentence wrap-up.

The following measures were used to analyse the tracking patterns. *First-pass reading time* is the sum of the duration of fixations made on first entering a region of text until an eye-movement exits the region to either the left or right. First-pass reading times can provide an indication of the difficulty experienced when participants initially process a region of text. *First-pass regressions out* is the sum of regressive saccades made from the current most rightward fixation with a region of text, indicating the degree to which left to right movement was disrupted during the first sweep of the eyes through a region of text. By examining the location of this disruption it is possible to establish particular regions in a text that cause difficulty for semantic processing. *Regression path time* is the sum duration of fixations from the first entry into that region from the left to the time the region was first exited to the right. This measure includes fixations made to re-inspect earlier portions of text and therefore provides an indication of early processing difficulty and time revisiting the sentence in order to make sense of such a difficulty. *Total reading time* is the sum of the duration of all fixations made within a region and provides an indication of the overall amount of time spent processing text in that region. Finally, *regressions-in* convey information on the percent of regressive movements from the right into each region and consequently present researchers with details of the regions of text that readers need to revisit in order to make sense of a piece of text.

Prior to analysis, short contiguous fixations were corrected using an automatic procedure. Fixations shorter than 80 ms were integrated with larger adjacent fixations within one character and fixations shorter than 40 ms that were not within three characters of another fixation were excluded. Fixations longer than 1200 ms were truncated. Trials where participants failed to read the sentence or there had been tracker loss were eliminated prior to analysis. Specifically, trials where two or more adjacent regions had zero first-pass reading times were removed, which accounted for less than 2% of the data.

Data for the different eye movement measures were analysed in separate ANOVAs with repeated-measures on variables context (RW vs. NW) and consistency (consistent vs. inconsistent).

Table 1 displays mean values for each measure in each condition and region.

--------------------TABLE 2 HERE----------------

2.1.2 First-pass reading times

In region 1, following the context sentence, first-pass reading times revealed a main effect of context [$F_1(1, 35) = 2.85$, $p = 0.1$; $F_2\ (1, 31) = 4.62$, $p = 0.04$] that was significant by items and as a trend by participants. Specifically, first-pass reading times were longer following a NW context than RW context. This suggests that readers are still processing prior information in NW conditions and that NW statements are more difficult to comprehend.

At the critical region 3, a significant interaction emerged [$F_1(1, 35) = 4.93$, $p = 0.03$; $F_2\ (1, 31) = 5.50$, $p = 0.03$]. Analysis of the simple main effects revealed longer first-pass reading times in the inconsistent than the consistent condition at the level of RW [$F_1(1, 35) = 8.20$, $p < 0.01$; $F_2\ (1, 31) = 13.07$, $p = 0.001$], however there was no difference between consistency conditions at the level of NW [$Fs < 0.3$]. Additionally, RW led to longer first-pass reading times compared to NW for inconsistent conditions [$F_1(1, 35) = 9.06$, $p < 0.005$; $F_2\ (1, 31) = 7.84$, $p < 0.01$] but there was no difference between RW and NW for consistent conditions [$Fs < 1.0$]. The increased reading times for the RW-inconsistent condition were expected given previous eye-movement investigations of real-world violations (see Ferguson & Sanford, 2008). Furthermore, the fact that NW-consistent and NW-inconsistent conditions were not significantly different at this stage suggests that processing is still ongoing within the negated context, and specifically that neither the violation of real-world knowledge nor the inconsistency with the NW context has affected processing at this point. Thus, we provide evidence that the effects of real-world anomalies can be picked up in measures of very early processing during a reading task. Additionally in this critical region, a main effect of consistency is apparent [$F_1(1, 35) = 5.28$, $p = 0.03$; $F_2\ (1, 31) = 4.63$, $p = 0.04$]. However, this effect was largely due to greatly increased reading times in the RW-inconsistent condition. There was no main effect of context [Fs < 1.0].

-------------------FIGURE 1 HERE---------------------

By the post-critical region 4, a main effect of context emerged [$F_1(1, 35) = 2.86$, $p = 0.1$; $F_2\ (1, 31) = 4.13$, $p = 0.05$] with longer first-pass reading times following a NW context than a RW context. We suggest that these longer reading times in the post-critical region following a NW context reflects the recruitment of additional processing to relate incoming information to the wider discourse context.

2.1.3 First-pass Regressions Out

In the critical region 3, there was no evidence of any effect of context or consistency, or an interaction on regressions out [Fs < 1.2]. However, by the post-critical region 4, a main effect of consistency emerged that was borderline by participants and significant by items [$F_1(1, 35) = 3.02$, $p = 0.09$; $F_2\ (1, 31) = 4.77$, $p = 0.04$]. Thus, by the post-critical region, readers are using the wider discourse to make sense of the passage and regressing back in the text to attempt to make sense of the inconsistency.

2.1.4 Regression path times

As with regressions out data, the critical region 3 yielded no main effects or an interaction between factors [all Fs < 1.4]. However, by the post-critical region 4, regression path times were longer when the critical word had been inconsistent with the prior context [$F_1(1, 35) = 3.91$, $p = 0.06$; $F_2\ (1, 31) = 4.47$, $p < 0.05$]. This consistency effect suggests that by the post-critical region, participants are using the prior RW or NW context to guide language processing.

---------------------FIGURE 2 HERE--------------------

2.1.5 Total reading times

The pre-critical region 2 showed a main effect of context [$F_1(1, 35) = 9.43$, $p = 0.004$; $F_2 (1, 31) = 7.33$, $p = 0.01$] as RW contexts led to longer total reading times than NW contexts in this region. By the critical region 3, a significant interaction [$F_1(1, 35) = 6.68$, $p = 0.01$; $F_2 (1, 31) = 5.34$, $p = 0.03$], a main effect of context [$F_1(1, 35) = 4.15$, $p = 0.05$; $F_2 (1, 31) = 4.08$, $p = 0.05$] and a main effect of consistency that was marginal by participants and significant by items [$F_1(1, 35) = 3.37$, $p = 0.07$; $F_2 (1, 31) = 4.08$, $p = 0.05$] was revealed. Analysis of the simple main effects revealed longer total reading times in the inconsistent than the consistent condition at the level of RW [$F_1(1, 35) = 10.95$, $p < 0.005$; $F_2 (1, 31) = 13.29$, $p = 0.001$], however this was not significant at the level of NW [Fs < 0.6]. Additionally, RW led to longer total reading times compared to NW for inconsistent conditions [$F_1(1, 35) = 12.35$, $p < 0.001$; $F_2 (1, 31) = 11.77$, $p < 0.002$] but there was no difference between RW and NW consistent conditions [Fs < 0.4]. No main effects or an interaction between context and consistency was evident in the post-critical region 4 [Fs < 1.0].[1]

2.1.6 Regressions in

Figure 3 shows the mean number of regressions into each region for each condition.

[1] Number of fixations data reflected total reading times with a main effect of context in the pre-critical region 2 [$F_1(1, 35) = 5.78$, $p = 0.02$; $F_2 (1, 31) = 6.44$, $p < 0.02$], reflecting more fixations in this region following a RW context than a NW context. Critical region 3 revealed a significant interaction [$F_1(1, 35) = 9.71$, $p = 0.004$; $F_2 (1, 31) = 6.44$, $p < 0.02$] with increased number of fixations in the region for inconsistent than consistent words in RW contexts and no difference in NW contexts.

----------------------FIGURE 3 HERE-----------------------

A significant interaction between context and consistency was found at the critical region 3 [$F_1(1, 35) = 4.70$, $p = 0.04$; $F_2\ (1, 31) = 8.52$, $p = 0.006$], with simple main effects revealing increased incidence of regressions into the region in the inconsistent than the consistent condition at the level of RW [$F_1(1, 35) = 5.14$, $p < 0.05$; $F_2\ (1, 31) = 3.80$, $p = 0.06$], and a trend for a reverse consistency effect at the level of NW [$F_1(1, 35) = 1.74$, $p < 0.2$; $F_2\ (1, 31) = 3.27$, $p = 0.08$]. Furthermore, RW led to more regressions into the critical region compared to NW for inconsistent conditions [$F_1(1, 35) = 6.14$, $p < 0.02$; $F_2\ (1, 31) = 10.82$, $p < 0.005$] but there was no difference between RW and NW for consistent conditions [$Fs < 2.1$]. This demonstrates that regressions into the critical region have been directed by real-world knowledge (e.g. *feeding a cat a bowl of carrots*). However, since consistency effects are significant for RW conditions but are a trend for NW conditions, we can conclude that RW anomalies within a RW context lead to stronger processing difficulties than RW anomalies in a NW context.

In sum, Experiment 1 showed that within a RW discourse, the presence of a real-world violation led to longer reading times, more fixations and a higher incidence of regressions in at the critical noun than a real-world consistent event. Thus, information that conflicts with real-world knowledge causes a disturbance in the eye-movement patterns. In contrast, within a NW discourse context, no differences in reading times were found on the critical word, however in the post-critical region regressive eye movements were increased following inconsistent compared to consistent conditions. Therefore, the NW context appears to have had the effect of delaying integration of the critical word as the processor resolved a conflict between real-world and negated-world information. These results suggest that the dominant process here is to relate the content to real-world expectations, and that initially, a prior negated-world context is not available to the reader. However, following this delayed integration, the preceding NW context is accepted as the basis for processing, which is evidenced by the later disruption in response to contextually inconsistent, compared to consistent, critical words. Longer first-pass reading times for NW conditions at the post-critical region, reflecting ongoing evaluation against RW knowledge and increased incidence of regressive eye-movements from this region confirm that the NW context is eventually incorporated into the representation of the sentence meaning. The theoretical implications of these results are discussed at length in the General Discussion.

The data here fits with the results from previous investigations with counterfactual contexts (see Ferguson & Sanford, 2008), that language input is initially influenced by real-world knowledge, even in the presence of a context allowing an alternative world, and that contextual information influences later discourse resolution. We can also infer that RW inconsistencies were detected earlier, and lead to more robust disruptions in comprehension than NW inconsistencies, as effects of the NW inconsistency were revealed later in the eye-movement record, and only in regressive eye-movement measures.

2.2 Experiment 2

The main aim of Experiment 2 was to allow a fuller investigation into the use of NW information, particularly when that information conflicts with existing knowledge about reality. To this end, we examined the processing of events that were either consistent or inconsistent in terms of a pre-specified negated world, as used in Experiment 1. We expected to extend the eye-movement findings of Experiment 1, which indicated that the integration of a NW context, when present, is delayed as the basis of processing, after the initial use of real-world knowledge. That is, congruent with the *search-for-*

alternatives view (e.g., McDonald & Just, 1989), we expected to observe a larger N400 to NW-consistent (RW anomaly) than NW-inconsistent words, indicating that the representation of the sentence meaning is initially influenced by real-world knowledge (on the critical word), and that the NW context is incorporated into the discourse at some later stage of comprehension processing (word downstream of the critical word). However, increasing N400 evidence suggests that a strong discourse context *can* overrule local lexical-semantic factors when these two conflict and as such can immediately influence comprehension processes (e.g., Filik & Leuthold, 2008; Nieuwland & Van Berkum, 2006). Therefore, it is possible that the introduction of a negated discourse context, such as those used in Experiment 1, can eliminate (or even reverse) effects of real-world knowledge violations as indicated by the N400 component of the ERP.

2.2.1 ERP Analysis

Grand average ERP waveforms are presented for the two conditions in Figure 4. It can be seen that the NW-consistent condition triggered a more negative-going, deflection (N400) than the NW-inconsistent condition, starting at about 350 ms after target word onset.

-----------------------FIGURE 4 HERE------------------------

Analysis of mean N400 amplitude in the 350-500 ms time interval revealed a widely distributed negative-going ERP waveform that was more negative for consistent than inconsistent conditions in the analysis of midline electrode data (0.3 vs. 1.5 μV) [$F(2, 34) = 4.6$, $p < 0.05$] and as a trend in the analysis of lateral electrode data (0.5 vs.

1.2 μV) [F(2, 34) = 3.6, $p = 0.07$]. The latter analysis indicated an interaction between consistency x hemisphere x ant-pos persisted [F(2, 34) = 6.5, $p < 0.01$, $\varepsilon = 0.84$], due to a more negative going ERP waveform for consistent than inconsistent conditions over right-posterior sites. The topographical distribution of these consistency effects is illustrated by the map in Figure 5. Together, N400 findings indicate that the representation of the sentence meaning is initially influenced by real-world knowledge, thereby, supporting the search-for-alternatives view. In addition, the finding that the NW context did not eliminate effects of real-world knowledge violations contrasts with recent N400 studies that examined the influences of discourse context on local pragmatic anomaly processing (Filik & Leuthold, 2008; Nieuwland & Van Berkum, 2006).

------------------FIRGURE 5 HERE--------------------

In light of earlier eye-tracking results, showing that consistency effects emerged later in the target sentence, further analyses were carried out to determine whether this N400 effect would be reversed (i.e. larger N400 for context-inconsistent than consistent conditions) later in the ERP waveform. Therefore, we examined ERP effects at words *n + 1* (500-1000 ms after critical word onset) and *n + 2* (1000-1500 ms after critical word onset). Word *n + 1* was always the connector *and*, while word *n + 2* initiated some neutral conclusion to the target sentence, and thus was identical across conditions. Therefore, we concluded that ERP differences revealed here might reflect later processing of the critical word. As is shown in Figure 6, no reliable effects [all Fs < 2] were revealed following the onset of word *n + 1*. Specifically, during this time interval the effect of consistency was neither significant as a main effect nor in interaction with any combination of the variables hemisphere, ant-pos, and verticality [all Fs < 1].

---------------------FIGURE 6 HERE---------------------

However, in the time interval between 250 and 400 ms after the onset of target word *n + 2*, a main effect of consistency emerged in the analysis of midline electrode data [F = 4.52, $p < 0.05$], revealing a more negative-going ERP waveform for consistent than inconsistent conditions. Thus, one interpretation of these results is that the consistency effect has not been reversed by word *n + 2*, and processing is still in terms of real-world knowledge.

3. Discussion

The experiments reported in this paper examined whether and at what stage of processing a context that negates the readers' knowledge about the world influences interpretation of a subsequent target sentence. Experiment 1 examined eye-movements as participants read passages that included a prior context depicting a negated-world (NW), or real-world (RW), followed by a second sentence continuation. Events described in these continuations either included a violation of RW knowledge, or were congruent with RW knowledge. The results showed that within a RW discourse, the presence of a real-world violation led to longer reading times, more fixations and a higher incidence of regressions in at the critical noun than a real-world congruent event. In contrast, the presence of a NW discourse context had the effect of delaying integration of the critical word as the processor resolved a conflict between real-world and negated-world information. Thus, it appears that the negated-world framework has not been integrated into the wider sentence meaning at this stage. However, effects of real-world knowledge were later reversed in the NW conditions. By the post-critical region, reading times and regressions were led by the fit of the critical word with the preceding context, as NW-inconsistent items led to longer first-pass reading times and increased regressive eye-movements compared to NW-consistent. This delayed context effect is taken to reflect ongoing evaluation of the incoming information against knowledge about the world until the NW context can be incorporated into the representation of the sentence meaning.

Further evidence for the delayed influence of a negated context was found in Experiment 2. Here, ERPs were recorded to examine processing of events that were either consistent or inconsistent in terms of a pre-specified negated world. Results revealed that introducing a negated-world context was not sufficient to reverse the typical N400 effect to semantic violations at the critical word. Thus, following a negated context, the model is not immediately updated and language input is initially and persistently tested against real-world knowledge. These results differ from Experiment 1's eye-movement investigation, which showed no evidence of a difference between consistency conditions for NW items at the critical word, but later effects of the contextual inconsistency. These differences are likely to be due to the relative sensitivity of each of these tasks and as such a fuller understanding of language processes can be achieved by combining results across complementary paradigms, such as eye-movements and ERPs. Therefore, it appears that two processes are manifest during the comprehension of negative contexts and their consequents. The first is to relate incoming information to pre-stored world knowledge and the second is to integrate that information with constraints from the wider negated context.

Taken together, these results support previous reading experiments that have demonstrated that initial processing of a sentence is not influenced by a preceding negation operator (Fischler et al., 1983; Hald et al., 2004; Giora, 2006). Further, the suggestion that a passage's total meaning is computed on-line at a later stage is confirmed here by evidence in the eye-movement data for a delayed reversal of local pragmatic effects following a negated context. Thus, these results allow us to reject the *narrow-view* suggestion that negatives simply reduce accessibility of a negated concept (e.g. *carnivores*), which would predict that information mentioned within the negative's scope (e.g. *fish*) would result in early processing difficulties. Instead, the results reported here support the *search-for-alternatives* view. According to this approach, negation prompts the reader to search for an appropriate alternative to the negated concept, following an initial processing insensitivity to negation (e.g. Giora, Balaban, Fein & Alkabets, 2005; Hasson & Glucksberg, 2006; Kaup et al., 2006). Kaup et al. (2006) suggest that the delayed effects of negation are due to participants mentally simulating the negated situation prior to suppressing this simulation and replacing it with a simulation of the actual situation. Thus in reference to the current data, we have demonstrated that simulation is not limited to describing concrete events, such as *the door is not open*. It is also possible when applied to abstract events, such as *the cat is not a carnivore*, where an experiential simulation may need to include some implicit representation of the cat's intrinsic properties.

The processing of negation was further complicated in the current experiments by the use of counterfactual conditional *if- then* statements that typically elicit hypothetical reasoning from premises that are false in actuality (Fauconnier & Turner, 2003). A dual- process approach attempts to explain these mental simulations and the subsequent impact on language processing (Evans, 2007). Thus, language comprehension is generally a rapid and automatic process, as demonstrated by the immediate processing difficulties reported for unnegated RW anomalies. Indeed, in these RW conditions, readers did not even need to access the conditional context to detect an inconsistency, such as 'feeding a cat a bowl of carrots'. In contrast, the NW conditions required readers to override this automatically activated evaluation, and substitute it with the contextually inferred meaning (Deutsch, Gawronski & Stack, 2006). Thus, detection of an inconsistency within a context that is both negated and counterfactual conditional engages a reflective, memory-based component, thus explaining why detection of the inconsistency under these conditions takes longer to emerge.

The experimental items used in the current studies allowed novel investigations of the consequences of negation on later processing (outwith the negated sentence), improving on previous experimental confounds such as, priming and distance between the negative operator and the negated concept. Nevertheless, these short passages afforded a limited exploration of the readers' representation of the negated world. Therefore, future investigations could examine how readers process information over time when the constraints of the negated world are repeatedly supported. Using a longer narrative, one would expect the reader to gradually and implicitly set up an alternative world that is at odds with the readers' knowledge about the world, in which it is not uncommon for cats to eat vegetables. Thus, it is anticipated that in the first instance, a local anomaly would elicit a processing disruption regardless of the preceding NW context. However, this interference from real-world knowledge would diminish and eventually disappear after repeated exposure to NW-consistent information. The build-up of a discourse context has been investigated previously by Niewlands and Van Berkum (2006; Experiment 1). Specifically, they examined the build-up of a cartoon-like context and revealed that while a large N400 effect was initially elicited following animacy violations, this effect disappeared following comparable animacy violations later in the story.

The issue of polarity has also been investigated in relation to natural language quantifiers, demonstrating different patterns of focus induced by various negative (e.g. *few*) and positive (e.g. *a few*) quantifiers. For example, a proposition can convey information about a group of individuals, but differing quantifiers can direct the readers' attention to the good or the bad aspects, as in *Few of the patients were critically ill* (which is good), or *A few of the patients were critically ill* (which is bad). Different patterns of pronominal reference have been revealed for positive and negative quantifiers in continuation tasks (e.g., Moxey & Sanford, 1987; 1993a; Sanford, Moxey & Paterson, 1996). Following a negative quantifier, continuations typically refer to the *Complement Set* (i.e. the patients who were *not* critically ill), while following a positive quantifier, attentional focus is directed to the *Reference Set* (i.e. the patients who *were* critically ill). Therefore, the polarity of the quantifier used clearly affects the inferences made during comprehension. In a similar way, negation could be influencing the saliency of information in the current studies. Negative contexts could initially focus the readers' attention onto the fact that *cats are not carnivores*, and therefore lead to questions such as why this would be the case and initiate a search for alternatives. In contrast, counterfactual (affirmative) contexts (as in Ferguson & Sanford, 2008) could focus readers' attention onto the explicitly stated alternative world of *cats being vegetarians* and therefore focus attention onto the consequences of such a world.

To conclude, in two studies using different experimental techniques, we have shown that processing negated information initiates an extended comparison process to reality, which is more disruptive to language processing than when comprehending affirmative information, such as counterfactuals. The experiments reported in this paper suggest that when processing negated information, comprehenders first simulate the negated situation and then suppress this model in favour of a simulation of the actual situation. This initial consideration of the negated situation prompts a comparison to reality to assess that reality no longer holds, thus delaying access to the actual situation. Therefore, it appears that to comprehend a negated concept, readers must represent information about reality, *not*-reality and an alternative 'counterfactual' reality, although the alternative reality only becomes available after *not*-reality has been accepted. Finally, it is interesting to note that while the N400 effect to pragmatic violations can be reversed by introducing an appropriate affirmative counterfactual (Ferguson, Sanford & Leuthold, 2007) or fictional (Nieuwland & Van Berkum, 2006; Filik & Leuthold, 2008) discourse context, the same is not true following a negated-world context. Thus, an important question arises as to whether the increased N400 effects to RW-inconsistent, CW-inconsistent and NW-consistent conditions are modulated by the same underlying mechanisms. Ultimately, only further experimentation will reveal whether the same neural generators are involved in these different types of anomaly detection.

4. Experimental Procedures

4.1 Experiment 1

4.1.1 Participants

Thirty-six participants were recruited from the undergraduate population of students from the University of Glasgow and were paid to participate in the study. All participants were native speakers of English, who did not have dyslexia and with vision that they reported to be normal or corrected-to-normal. Participants were naïve to the purpose of the study and had not had previous exposure to the experimental items.

4.1.2 Materials and Design

Thirty-two experimental items were constructed, as in Table 1 (please contact first author for a full list of items). Two context conditions were crossed with two consistency conditions to create a 2x2 within-subjects design. The context condition was split into two levels: a real-world (RW), where the first sentence depicted a realistic circumstance, and a negated-world (NW), where the first sentence cancelled real-world expectations, but did not create an alternative model. Note that all items used scalar and bi-polar concepts (e.g. carnivores/ vegetarians; warm/ cold; dangerous/ harmless) to narrow the number of alternatives available to the comprehender, without being strictly dichotomous. The consistency condition also had two levels and was manipulated by changing the noun in the second sentence: inconsistent, such that events described in the second sentence were inconsistent with the prior context; or consistent, whereby the critical noun was consistent with the pre-specified context. This created a fully crossed experiment with four conditions: RW-inconsistent, RW-consistent, NW-inconsistent, and NW-consistent.

----------------------TABLE 2 HERE---------------------

The length and frequency of the critical nouns was matched across conditions using the British National Corpus. Nouns in the RW-inconsistent and NW-consistent conditions averaged 5.6 (min.= 3, max.= 10) characters, while nouns in the RW-consistent and NW-inconsistent conditions averaged 5.9 (min.= 3, max.= 9) characters. Mean frequency was 82.1 occurrences per million words for the RW-inconsistent and NW-consistent conditions and 64.5 for the RW-consistent and NW-inconsistent conditions. Thus, we can be confident that any difference in reading times between conditions is not due to discrepancies in length or frequency of the critical nouns.

One version of each item was assigned to one of four presentation lists. The 32 experimental items were assigned to lists so that equal numbers of each condition appeared on each list. One version of each item appeared on each list so that participants did not see more than one version of any given item. The items were displayed alongside 152 filler sentences of various types to create a single random order. At least one filler item intervened between each experimental item. Nine participants were randomly assigned to read each experimental list and they only saw each target sentence once, in one of the four conditions. A comprehension question was presented after half of the experimental (i.e., 16) and filler (76) items. Participants did not receive feedback about their responses. All participants scored at or above 90% accuracy on the comprehension questions.

4.1.3 Eye tracking

Participants' eye movements were recorded during reading using a Forward Technology Dual Purkinje generation 5.5 eye-tracker. The eye-tracker recorded participants' gaze location and movement from the right eye, although viewing was binocular. Recordings were taken every millisecond. A forehead rest and head strap, along with a bite-bar, were used to stabilize participants' head position and to minimise interference to the signal caused by head movements. All sentences were presented in size 10 Courier New font style through a PC, on a VDU screen, 60cm from the participants' eyes.

4.1.4 Procedure

The eye tracker was calibrated using a series of nine fixed targets distributed across the display, during which the participant was asked to fixate on each point on the computer screen as they appeared in order to establish the correlation between x/y voltages and screen position. Calibration was monitored and adjustments to the calibration were made throughout the experiment as necessary.

Prior to the presentation of each sentence, a pattern of fixation points appeared on the screen, spanning the position to be occupied by the sentences. Participants were instructed to fixate on a sequence of fixation points ending at the top left side, where the first character of the text would be displayed. Once this calibration check was completed accurately, the experimenter pressed a button to advance the screen to display the next item. This procedure ensured that participants were consistently tracking well, and that eye-movement records began uniformly with the initial words in each sentence. Adjustments to the calibration were made at this stage if necessary. Materials appeared individually on the screen, spanning a maximum of four lines of text, with each line separated by two blank lines and a maximum of 75 characters per line. Participants were instructed to read at their normal rate and to comprehend the sentences, in order to answer the questions. After reading each sentence, they pressed a button that led to the presentation of a comprehension question or the next trial. Comprehension questions followed 50% of target and filler trials in a pseudo-random order. Participants responded to the questions by clicking either the button in their left hand or the button in their right hand when given a choice of two answers on either left or right side of the screen. Adjustments to the calibration were made during the experiment when the experimenter deemed it necessary.

4.2 Experiment 2

4.2.1 Participants

Eighteen right-handed native English speakers were recruited from the University of Glasgow (9 males) and were paid to take part in this study. Participants ranged in age from 18 to 25 years. Participants did not have dyslexia and had vision which they reported to be normal or corrected-to-normal. All participants were naïve to the purpose of the study and had not taken part in Experiment 1.

4.2.2 Materials and Design

Eighty experimental items were constructed for use in this study. A single-factor within subjects design compared NW-inconsistent and NW-consistent conditions, as shown in Table 3.

----------------------------TABLE 3 HERE---------------------------

Here, each item consisted of an introduction sentence and a NW context sentence followed by a critical target sentence. The first sentence of each item introduced the topic of the passage, and was identical in both conditions. The purpose of this introductory sentence was to maintain a consistent structure across experimental items and fillers (which were part of a different study). Sentence two depicted a ‘negated-world’ context that cancelled real-world expectations, but did not create an alternative model. The target sentence described an event that could either be inconsistent with the prior context (but consistent with real-world knowledge), or consistent with the prior context (but anomalous according to real-world knowledge). Thus, within the example above, it should be inconsistent for cats to eat *fish* if they were ‘not carnivores’, but it should be consistent for cats to eat *carrots* within this context. Apart from this critical word, the remainder of the target sentence was identical across experimental conditions.

The experimental items were organised into two lists, such that each item appeared in one of its two versions in a given list, but appeared in all versions across the two lists. Therefore, each list comprised 40 items in each of the two conditions. The items were displayed alongside 160 filler sentences of various types, with at least one filler item intervening between each experimental item.

Participants were tested in an electrically shielded booth with ambient light kept at a low level. Word stimuli were presented in black 16 point Helvetica font on a white background at the centre of a 15” computer monitor at a viewing distance of 90 cm.

4.2.3 Procedure

Participants were informed about the EEG procedure and experimental task. After electrode application they were seated in a booth where they read the materials from a computer screen. There were six practice trials to familiarize them with the procedure, after which the experimenter answered any questions. As illustrated in Figure 4, each trial began with the presentation of a single centrally-located red fixation cross for 500 ms to signal the start of a new item. After this time, a white fixation cross appeared for 500 ms. Next, sentence 1 (introduction sentence) was presented on the screen. Participants were instructed to read this sentence and press spacebar on a keyboard to continue when ready. A blank screen appeared for 1000 ms, before sentence 2 (NW context sentence) was presented, in the same way as sentence 1. Participants were asked to read this so that they understood the context and pressed spacebar to progress at which point, the screen was cleared for 500 ms. Finally, a fixation cross (500 ms) preceded sentence 3 (critical sentence) that was presented word-by-word, with each word appearing at the centre of the screen for 300 ms. A blank-screen interval of 200 ms separated words. Sentence-ending words appeared with a full stop. A 2500 ms blank-screen interval followed each item. There was no secondary task. Trials appeared in eight blocks of thirty trials. Each block was separated by a break, the duration of which was determined by the participant. Thus, participants were tested in a single session that lasted approximately one hour, during which they were seated in a comfortable chair located in an isolated room.

----------------------FIGURE 7 HERE--------------------

4.2.4 Electrophysiological Measures

A BIOSEMI Active-Two amplifier system was used for continuous recording of electroencephalographic (EEG) activity from 72 Ag/AgCl electrodes over midline electrodes Fpz, AFz, Fz, FCz, Cz, CPz, Pz, POz, Oz, and Iz, over the left hemisphere from electrodes IO1, Fp1, AF3, AF7, F1, F3, F5, F7, F9, FC1, FC3, FC5, FT7, C1, C3,

C5, M1, T7, CP1, CP3, CP5, TP7, P1, P3, P5, P7, PO3, PO7, O1, two nonstandard positions PO9’ and O9’ which were located at 33% and 66% of the M1-Iz distance, respectively, and from the homologue electrodes over the right hemisphere (see Figure 5). EEG and EOG recordings were sampled at 256 Hz. Off-line, all EEG channels were recalculated to an average mastoid reference.

Trials containing blinks were corrected using a dipole approach (BESA Version 5.1.6) and automatic artifact detection software (BESA) was run. Trials with non-ocular artifacts (drifts, channel blockings, EEG activity exceeding ± 75 μV) were discarded. For all remaining trials, epochs of 1700 ms, starting 200 ms prior to the onset of the critical noun, were generated from the continuous EEG record. Thus, the post-stimulus epoch lasted for a total duration of 1500 ms.

4.2.5 ERP Data Analysis

For analysis of the EEG data, we only included trials without EEG or EOG artifacts. The signal at each electrode site was averaged separately for each experimental condition time-locked to the onset of the critical noun. Before the measurement of ERP parameters EEG and EOG activity was band-pass filtered (0.03-30 Hz, 6 dB/oct). The ERP waveforms were aligned to a 200 ms baseline prior to the onset of the critical noun. To analyze experimental effects on the N400, mean ERP amplitude was determined in the time interval from 350-500 ms relative to critical word onset.

ERP amplitudes at midline electrodes (AFz, Fz, FCz, Cz, CPz, Pz, POz, Oz) were analysed separately from data recorded over lateral electrode sites. Specifically, regions of interest (ROIs) pooled data from lateral electrode sites, as recommended for the analysis of high-density electrode arrangements (cf. Dien & Santuzzi, 2005). The electrodes were divided along a left-right dimension, an anterior-to-posterior dimension, and a dorsal-ventral dimension. The six ROIs over the left hemisphere were: left-anterior-ventral (AF7, F7, FT7, F5, FC5), left-anterior-dorsal (AF3, F3, FC3, F1, FC1), left-central ventral (TP7, T7, C5, CP5), left-central-dorsal (C3, CP3, C1, CP1), left-posterior-ventral (PO9', O9', P7, PO7, O1), and left-posterior-dorsal (P3, PO3, P1, P5); six homologue ROIs were defined for the right hemisphere.

Statistical analyses were performed by means of Huynh-Feldt corrected repeated measures analyses of variance (ANOVA). For the statistical analysis of ERP amplitude data recorded from midline electrodes, an ANOVA was performed with variables consistency (NW-consistent/ NW-inconsistent) and electrode (AFz, Fz, FCz, Cz, CPz, Pz, POz, Oz). For the analysis of ERP deflections maximal over lateral electrode sites, an ANOVA was performed with variables consistency (NW-consistent/ NW-inconsistent), hemisphere (left, right), ant-pos (anterior, central, posterior), and verticality (ventral, dorsal).

References

Braze, D., Shankweiler, D., Ni, W., Palumbo, L.C., 2002. Readers' eye movements distinguish anomalies of form and content. Journal of Psycholinguistic Research, 31 (1), 25-44.

Clark, H.H., Chase, W.G., 1972. On the process of comparing sentences against pictures. Cognitive Psychology, 3, 472-517.

Deutsch, R., Gawronski, B., & Strack, F., 2006. At the boundaries of automaticity: Negation as a reflective operation. Journal of Personality and Social Psychology, 91, 385-405.

Dien, J., Santuzzi, A.M., 2005. Application of repeated-measures ANOVA to high-density ERP datasets: A review and tutorial. In T. Handy (Ed.), Event Related Potentials: A Methods Handbook (pp. 57-82). Cambridge, MA: MIT Press.

Evans, J.St.B.T., 2007. Hypothetical thinking: Dual processes in reasoning and judgement. Hove, UK: Psychology Press.

Evans, J.St.B.T., Clibbens, J., Rood, B., 1996. The role of implicit and explicit negations in conditional reasoning bias. Journal of Memory and Language, 35, 392-409.

Evans, J.St.B.T., Over, D.E., 2004. If. Oxford, England: Oxford University Press.

Fauconnier, G., & Turner, M., 2003. The Way We Think: Conceptual Blending and the Mind's Hidden Complexities. New York: Basic Books.

Ferguson, H.J. & Sanford, A.J., 2008. Anomalies in real and counterfactual worlds: An Eye-Movement Investigation. Journal of Memory and Language, 58, 609-626.

Ferguson, H.J., Sanford, A.J. & Leuthold, H., 2007. Real-world interference in detecting violations of counterfactual and negated worlds. Poster presented at Cognitive Neuroscience Society Annual meeting. New York, NY.

Filik, R., 2008. Contextual override of pragmatic anomalies: Evidence from eye movements. Cognition, 106, 1038-1046.

Filik, R. & Leuthold, H., 2008. Processing local pragmatic anomalies in fictional contexts: Evidence from the N400. Psychophysiology, 45, 554-558.

Fischler, I., Bloom, P., Childers, D.G., Roucos, S.E., Perry, N.W., 1983. Brain potentials related to stages of sentence verification. Psychophysiology, 20, 400-409.

Giora, R., 2006. Anything negatives can do affirmatives can do just as well, except for some metaphors. Journal of Pragmatics, 38, 981-1014.

Giora, R., Balaban, N., Fein, O., Alkabets, I., 2005. Negation as positivity in disguise. In H.L. Colston & A. Katz (Eds.), Figurative language comprehension: Social and cultural influences. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Giora, R., Fein, O., Aschkenazi, K., & Alkabets-Zlozover, I., 2007. Negation in context: A functional approach to suppression. Discourse Processes, 43, 153-172.

Glenberg, A.M., Robertson, D.A., Jansen, J.L., & Johnson-Glenberg, M.C., 1999. Not propositions. Journal of Cognitive Systems Research, 1, 19-33.

Hald, L., Kutas, M., Urbach, T.P., Pahrhizkari, B., 2004. The N400 is not a brainwave: negation and the N400 effects for true and false sentences. Poster presented at the Cognitive Neuroscience Society annual meeting.

Hasson, U., Glucksberg, S., 2006. Does understanding negation entail affirmation? Journal of Pragmatics, 38, 1015-1032.

Joseph, H.S.S.L., Liversedge, S.P., Blythe, H.I., White, S.J., Gathercole, S.E., Rayner, K., in press. Children’s and adult’s processing of anomaly and implausibility during reading: Evidence from eye movements. Quarterly Journal of Experimental Psychology.

Kaup, B., in press. What psycholinguistic negation research tells us about the nature of the working memory representations utilized in language comprehension. In H. Pishwa (Ed.), Language and Memory. Mouton de Gruyter.

Kaup, B., Lüdtke, J., Zwaan, R.A., 2006. Processing negated sentences with contradictory predicates: Is a door that is not open mentally closed? Journal of Pragmatics, 38, 1033-1050.

Kaup, B., Yaxley, R.H., Madden, C.J., Zwaan, R.A., Lüdtke, J., 2007. Experiential simulations of negated text information. Quarterly Journal of Experimental Psychology, 60, 976-990.

Kutas, M., van Petten, C.K., & Kluender, R. (2006). Psycholingusitics Electrified II. In M.J. Traxler & M.A. Gernsbacher (Eds.), *Handbook of psycholinguistics* (2nd ed.) (pp. 659-724). Amsterdam: Academic Press.

Lüdke, J., Friedrich, C.K., De Filippis, M., Kaup, B., 2005. ERP correlates of negation in a sentence-picture-verification paradigm (Poster). 45th Annual Meeting of the Psychonomic Society. Toronto, Canada.

MacDonald, M.C., Just, M.A., 1989. Changes in activation levels with negation. Journal of Experimental Psychology: Learning, Memory, and Cognition, 15, 633-642.

Moxey, L.M., Sanford, A.J., 1987. Quantifiers and Focus. Journal of Semantics, 5, 189-206.

Moxey, L.M., Sanford, A.J., 1993a. Communicating Quantities: A Psychological Perspective. Lawrence Erlbaum.

Ni, W., Fodor, J.D., Crain, S., Shankweiler, D.P., 1998. Anomalous strings: Eye movement patterns. Journal of Psycholinguistic Research, 27, 515-539.

Nieuwland, M.S., Van Berkum, J.J.A., 2006. When peanuts fall in love: N400 evidence for the power of discourse. Journal of Cognitive Neuroscience, 18, 1098-1111.

Prado, J., Noveck, I.A., 2006. How reaction times can elucidate matching effects and the processing of negation. Thinking & Reasoning, 12, 309-328.

Rayner, K., 1998. Eye movements in reading and information processing: 20 years of research. Psychological Bulletin, 124, 372-422.

Rayner, K., Warren, T., Juhasz, B., Liversedge, S., 2004. The effects of plausibility on eye movements in reading. Journal of Experimental Psychology: Learning, Memory and Cognition, 30, 1290-1301.

Sanford, A.J., Moxey, L.M., Paterson, K.B., 1996. Attentional focus with quantifiers in production and comprehension. Memory & Cognition, 24, 144-155.

Van Berkum, J.J.A., Brown, C.M., Zwitserlood, P., Kooijman, V., Hagoort, P., 2005. Anticipating upcoming words in discourse: Evidence from ERPs and reading times. Journal of Experimental Psychology: Learning, Memory, and Cognition, 31, 443-467.

Van Berkum, J.J.A., Hagoort, P., Brown, C.M., 1999. Semantic integration in sentences and discourse: Evidence from the N400. Journal of Cognitive Neuroscience, 11, 657-671.

Van Berkum, J.J.A., Zwitserlood, P., Brown, C.M., Hagoort, P., 2003. When and how do listeners relate a sentence to the wider discourse? Evidence from the N400 effect. Cognitive Brain Research, 17, 701-718.

Von Klopp, A., 1993. Negation: Implications for Theories of Natural Language Processing. Doctoral thesis, Centre for Cognitive Science, University of Edinburgh, Edinburgh.

Acknowledgements.

HF acknowledges the support of ESRC in the form of a postgraduate studentship. Thanks are due to members of the Glasgow Language Group, for helpful comments on this work.

Tables

Table 1: Mean eye-movement measures for regions 1 to 4 (standard errors in parentheses)

	Context Sentence 1	*1 Families could*	*2 feed their cat a bowl of*	*3 carrots and*	*4 it would gobble it down happily.*
st-pass reading time (ms)					
W-inconsistent	2527 (104.2)	454 (25.7)	682 (26.4)	307 (11.3)	865 (44.3)
W-consistent	2590 (115.2)	481 (28.9)	697 (32.5)	277 (7.9)	880 (45.7)
W-inconsistent	2480 (115.8)	489 (27.8)	661 (29.9)	281 (9.6)	937 (46.8)
W-consistent	2504 (125.2)	506 (25.4)	657 (31.0)	286 (8.5)	914 (34.8)
gression path time (ms)					
W-inconsistent	2527 (104.2)	514 (33.6)	899 (48.6)	370 (16.9)	1182 (80.2)
W-consistent	2590 (115.2)	509 (30.1)	875 (42.8)	344 (16.5)	1126 (77.4)
W-inconsistent	2480 (115.8)	544 (40.2)	830 (37.2)	355 (25.4)	1140 (54.8)
W-consistent	2504 (125.2)	544 (30.2)	879 (38.8)	352 (16.5)	1065 (49.0)
st-pass regressions out (%)					
W-inconsistent	-	5.4 (1.7)	16.2 (2.8)	14.6 (2.7)	24.8 (3.5)
W-consistent	-	4.3 (1.4)	13.8 (2.2)	15.8 (3.1)	18.4 (2.7)
W-inconsistent	-	4.1 (1.3)	14.5 (2.0)	13.6 (2.1)	19.5 (3.0)
W-consistent	-	3.7 (1.4)	17.5 (2.4)	12.2 (2.4)	16.6 (2.7)
tal reading time (ms)					
W-inconsistent	2619 (118.2)	550 (35.2)	902 (40.3)	360 (16.6)	1018 (52.9)
W-consistent	2651 (117.7)	555 (32.8)	886 (47.0)	313 (11.5)	986 (56.5)
W-inconsistent	2607 (122.7)	578 (32.7)	819 (33.7)	312 (12.9)	1005 (46.4)
W-consistent	2577 (128.7)	569 (27.2)	829 (36.3)	322 (10.9)	991 (39.4)
gressions In (%)					
W-inconsistent	11.4 (2.1)	21.6 (3.2)	20.8 (3.1)	12.8 (3.1)	0.9 (0.7)
W-consistent	9.9 (2.0)	16.8 (2.7)	17.8 (3.0)	6.7 (2.1)	1.0 (0.6)
W-inconsistent	13.4 (2.4)	20.1 (2.9)	15.5 (2.5)	5.9 (1.8)	0.6 (0.4)
W-consistent	10.7 (2.0)	18.8 (2.5)	15.2 (2.3)	9.5 (1.5)	1.9 (1.1)

Table 2: Examples of experimental sentences (Experiment 1).

RW-inconsistent If cats are hungry, they usually pester their owners until they get fed. Families could feed their cat a bowl of carrots and listen to it purr happily.
RW-consistent If cats are hungry, they usually pester their owners until they get fed. Families could feed their cat a bowl of fish and listen to it purr happily.
NW-inconsistent If cats were not carnivores, they would be cheaper for owners to look after. Families could feed their cat a bowl of fish and listen to it purr happily.
NW-consistent If cats were not carnivores, they would be cheaper for owners to look after. Families could feed their cat a bowl of carrots and listen to it purr happily.

Table 3: Examples of experimental sentences (Experiment 2).

NW-inconsistent
Cats are popular pets in British households.
If cats were not carnivores, they would be cheaper for owners to look after.
Families could feed their cat a bowl of fish and listen to it purr happily.
NW-consistent
Cats are popular pets in British households.
If cats were not carnivores, they would be cheaper for owners to look after.
Families could feed their cat a bowl of carrots and listen to it purr happily.

Figure Captions

Figure 1: Mean first-pass reading times in critical region showing standard error bars.

Figure 2: Mean regression path reading times in post-critical region, showing standard error bars.

Figure 3: Percentage regressions in per region, showing standard error bars.

Figure 4: Grand average ERPs elicited by critical nouns in sentence three. Note that negativity is plotted downwards.

Figure 5: Topographic maps of the ERP difference waveform (inconsistent condition minus consistent condition) for the time interval 350-500 ms (N400) relative to critical word onset.

Figure 6: Grand average ERPs elicited by critical nouns and subsequent words in sentence two. Note that 0ms indicates the onset of the critical word, and a new word was presented every 500 ms following.

Figure 7: Schematic diagram of the experimental procedure.

Figure 1

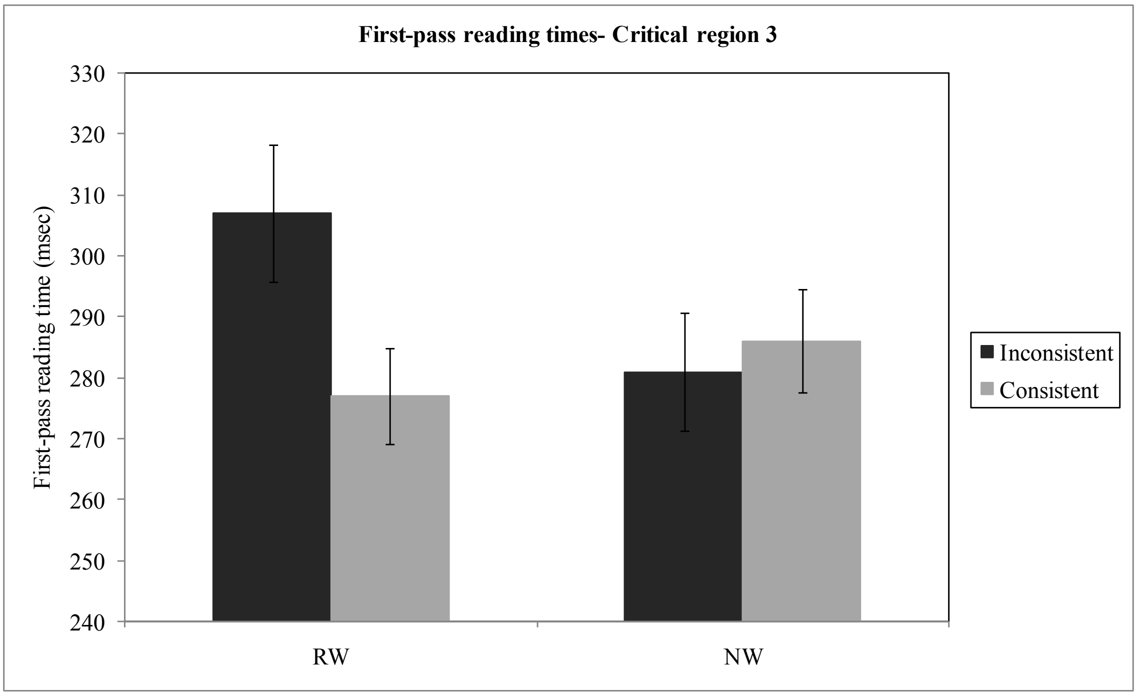
First-pass reading times- Critical region 3
330
320
310
300
290
280
270
260
250
240
First-pass reading time (msec)
RW
NW
Inconsistent
Consistent

Figure 2

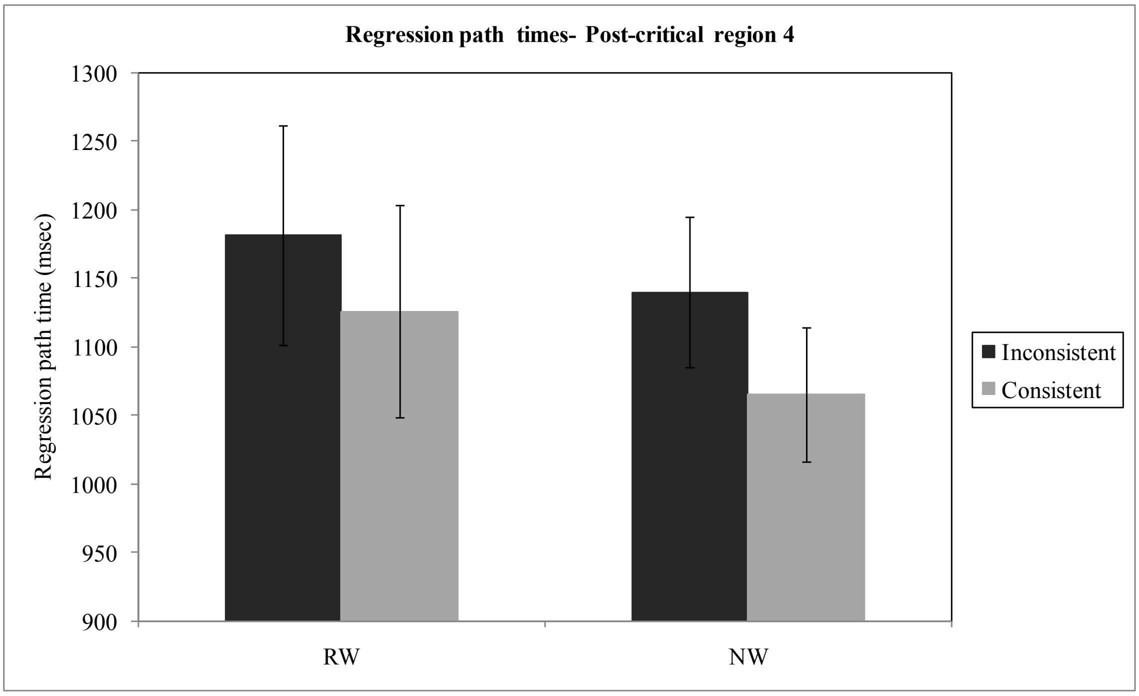
Regression path times- Post-critical region 4
Regression path time (msec)
1300
1250
1200
1150
1100
1050
1000
950
900
RW
NW
Inconsistent
Consistent

Figure 3

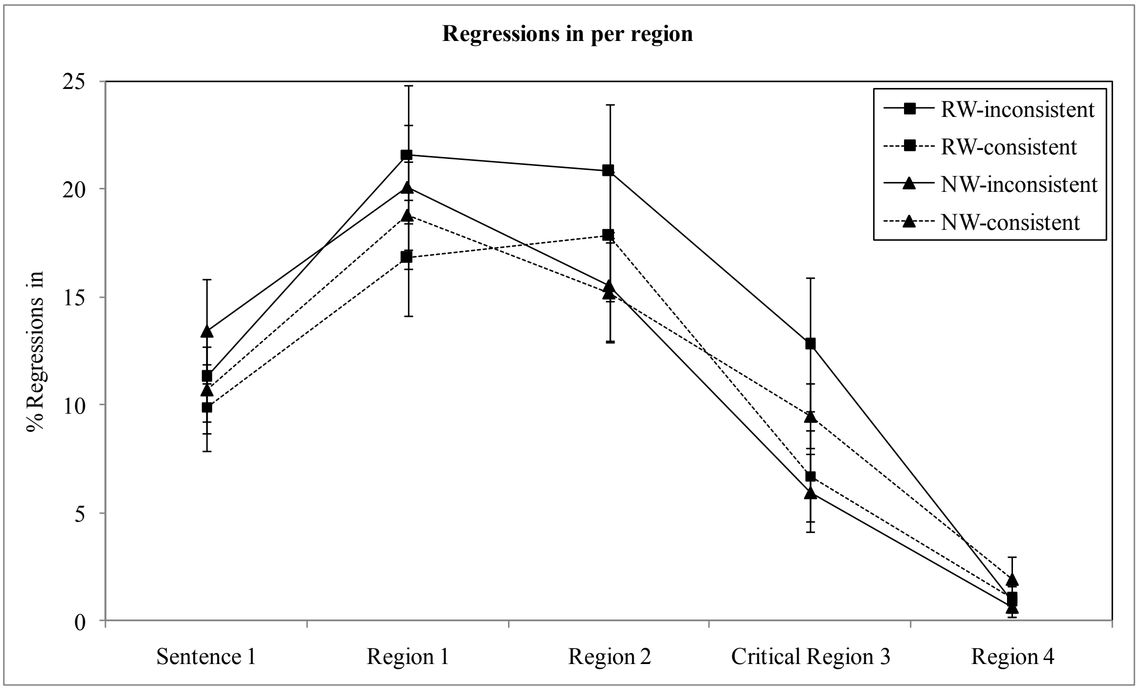
Regressions in per region
% Regressions in
25
20
15
10
5
0
RW-inconsistent
RW-consistent
NW-inconsistent
NW-consistent
Sentence 1
Region 1
Region 2
Critical Region 3
Region 4

Figure 4

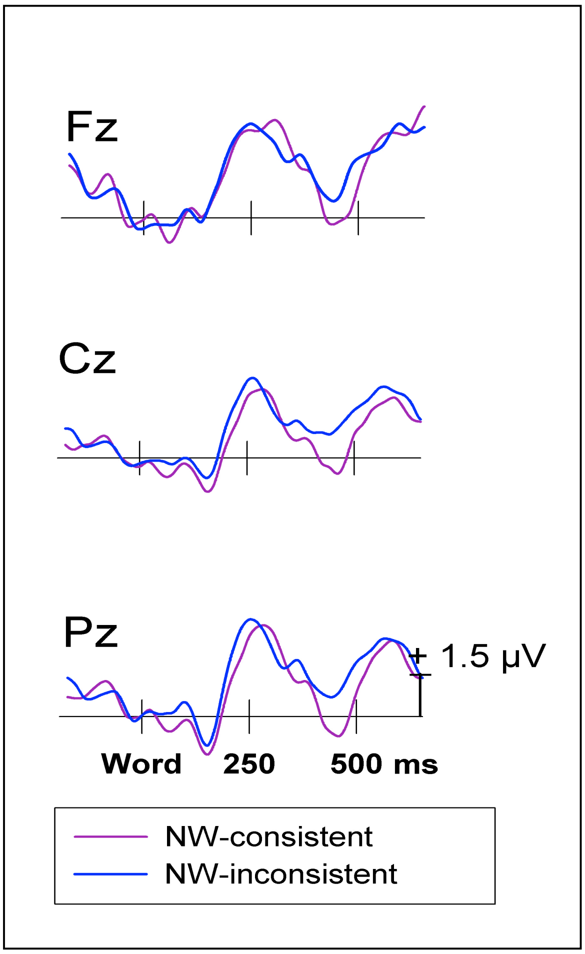

Figure 5

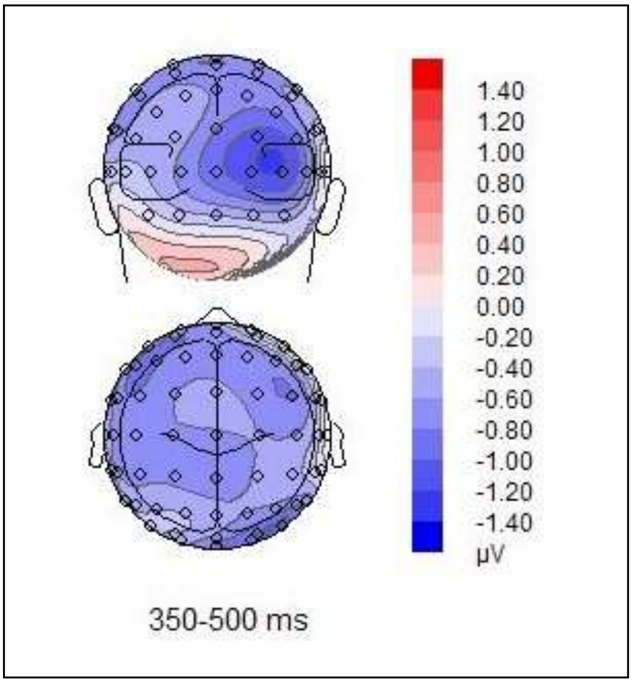
1.40
1.20
1.00
0.80
0.60
0.40
0.20
0.00
-0.20
-0.40
-0.60
-0.80
-1.00
-1.20
-1.40
µV
350-500 ms

Figure 6

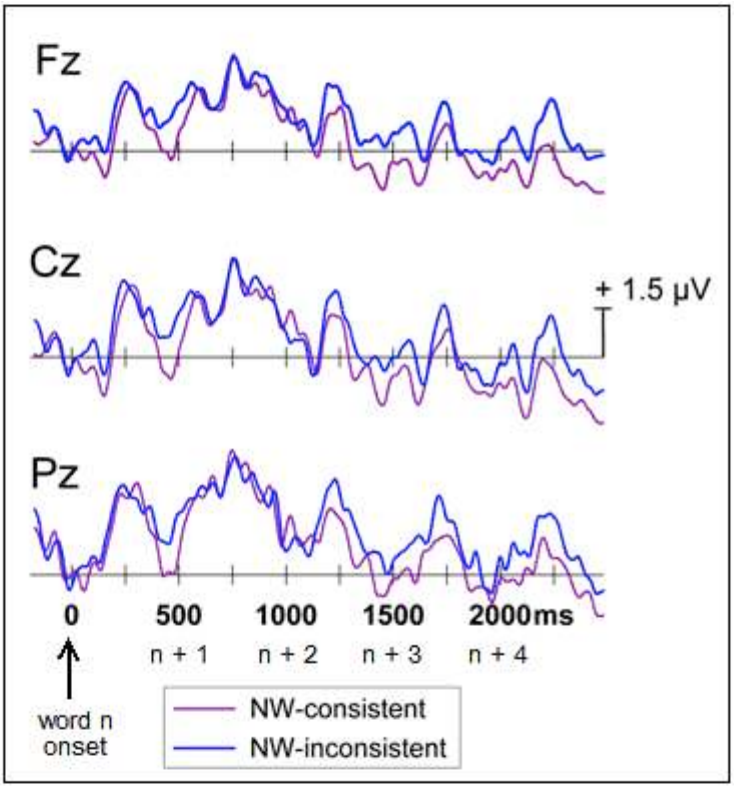
Fz
Cz
Pz
+ 1.5 µV
0
500
1000
1500
2000ms
n + 1
n + 2
n + 3
n + 4
word n
onset
NW-consistent
NW-inconsistent

Figure 7

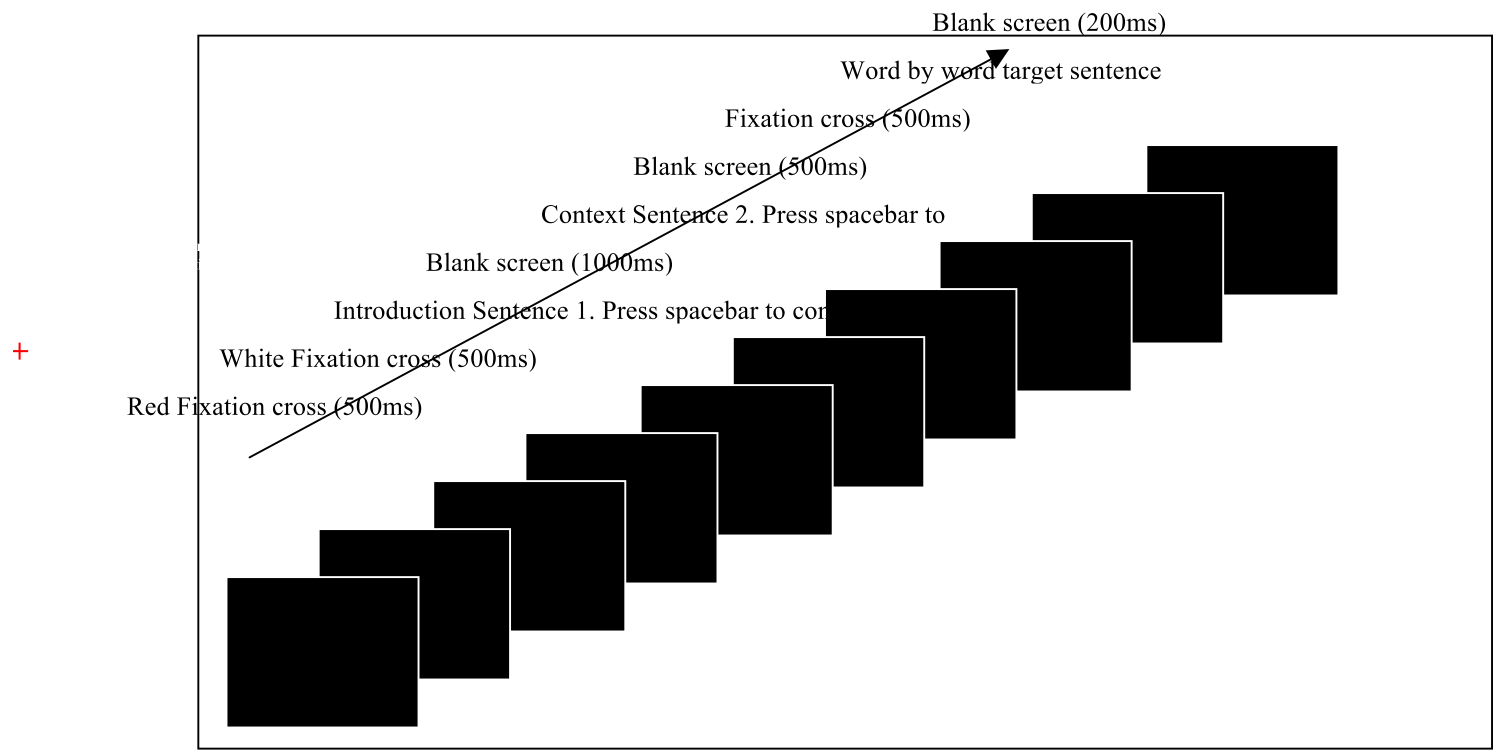
Blank screen (200ms)
Word by word target sentence
Fixation cross (500ms)
Blank screen (500ms)
Context Sentence 2. Press spacebar to
Blank screen (1000ms)
Introduction Sentence 1. Press spacebar to con
White Fixation cross (500ms)
Red Fixation cross (500ms)