

WHAT DOES "IT" MEAN, ANYWAY?
EXAMINING THE TIME COURSE OF SEMANTIC ACTIVATION IN REFERENCE RESOLUTION

BY

CYBELLE MARGUERITE SMITH

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Arts in Psychology
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2016

Urbana, Illinois

Master's Committee:

Professor Kara D. Federmeier
Professor Gary S. Dell
Professor Cynthia L. Fisher

ABSTRACT

Pronouns (e.g. 'it') are commonly studied in research on anaphoric reference, because they appear to carry out the linguistic function of referring back to an entity, while providing little or no new information about the referent. However, differing viewpoints have emerged in the psycholinguistic literature on what cognitive processes are engaged when reading or hearing a pronoun. One view is that encountering a pronoun leads the comprehender to reactivate the semantics of its antecedent. We examined this hypothesis by manipulating the concreteness of a noun antecedent and assessing whether an Event Related Potential (ERP) concreteness effect was elicited at a downstream pronoun. We observed a robust concreteness effect at the noun, but no evidence of a concreteness effect at the pronoun. In a secondary analysis, we examined whether N400 semantic priming from the antecedent would increase on content words shortly following the pronoun, relative to those preceding it. Again, although we observed robust semantic priming from the noun antecedent at positions following it, we did not observe an increase in the size of this effect following the pronoun. Taken together with the broader literature, our data suggest that pronouns do not induce the activation of (much) new semantic information in long-term memory, perhaps instead triggering an attentional shift towards their antecedents' active semantic representations within working memory. This, in turn, suggests that the process of linking an anaphor to its antecedent is attentionally mediated and does not entail long-term memory access, in contrast with ACT-R inspired models (Lewis & Vasishth, 2005).

*To Jamie Norton, Ryan Hubbard,
Allie Letkiewicz, Yifan Hu, Sandra Leister,
and my many supportive family members and friends*

ACKNOWLEDGMENTS

I would first very much like to thank my advisor, Professor Kara D. Federmeier, for her extreme patience, sage advice, and help at every stage in the preparation of this work. I would also like to thank Professor Gary S. Dell for his helpful comments on an earlier draft of this paper, and Professors Cynthia L. Fisher and Dell for serving on my master's thesis committee. Lastly, I am forever grateful for the emotional support of my family and friends. This research was supported by NIA grant AG2630 and a James S. McDonnell foundation award to K.D.F., as well as a Neuroengineering IGERT training fellowship and NSF-GRFP grant 1144245 to C.M.S.

TABLE OF CONTENTS

INTRODUCTION.....	1
METHODS.....	10
RESULTS.....	16
DISCUSSION.....	21
TABLE AND FIGURES.....	32
REFERENCES.....	37
APPENDIX A: COMPLETE LIST OF EXPERIMENTAL STIMULI.....	43

Introduction

The expression of connected ideas through language generally requires making reference to previously evoked concepts or entities. One job of the comprehender is to keep track of which entities are being referred to at different points across the discourse. Both verbal and non-verbal cues may enable comprehenders to effectively track the current topic of conversation and assist them in linking entities with appropriate propositions and thematic roles (reviewed in Levy & Fowler, 2000). The focus of the current paper is on whether, when, and how comprehenders make use of a particular linguistic cue, the personal pronoun, when making such links across text.

A personal pronoun, such as “he”, “she” or “it”, most often refers to a salient entity in the preceding text, its “antecedent.” Overt pronouns such as these have the potential to alter the cognitive state of the comprehender and allow him/her to maintain continuity of reference across propositions. However, the body of evidence to date has engendered diverse perspectives on what precise mental processes a comprehender engages in when reading or hearing a pronoun. The current paper seeks to address the role of memory and attention in forming a link between a pronoun and its most likely antecedent under relatively unambiguous conditions. In order to provide insight into the role of semantic memory in online pronoun processing, this study makes use of event-related potentials (ERPs) – in particular, the N400 component and the ERP concreteness effect – to test whether reading a pronoun induces the reader to recall semantic information about its antecedent.

Evidence for rapid linking of pronouns and antecedents

One way to assess how comprehenders process overt pronouns is to make the pronouns difficult to resolve immediately, either by failing to provide a matching antecedent or rendering the pronoun ambiguous. A large body of literature (reviewed in Callahan, 2008; Garnham, 2001; Sanford & Garrod, 1989) has documented comprehenders’ ability to rapidly detect when a pronoun fails to match its most likely candidate antecedent(s). For example, reading times increase when a pronoun mismatches the gender (e.g., Carreiras, Garnham, Oakhill, & Cain, 1996; Chow, Lewis, & Phillips, 2014) or number (Kreiner, Garrod, & Sturt,

2013) of its only likely antecedent within the preceding text. Studies using electrophysiological measures have confirmed this rapid sensitivity and demonstrated that differential responding begins within several hundred milliseconds of encountering a mismatching pronoun (reviewed in Callahan, 2008). Mismatching features tested in the literature have included syntactic gender, definitional/biological gender, stereotypical gender, number, and case. At a minimum, this demonstrates that comprehenders pay enough attention to pronouns in the input stream to notice when they fail to match an expected feature set, and that they engage in additional processing to accommodate for such a failure.

In a related line of work, comprehenders demonstrate immediate sensitivity to the referential ambiguity of an anaphor in context (e.g., Garnham, Oakhill, Ehrlich, & Carreiras, 1995; Nieuwland & Van Berkum, 2008). When comprehenders encounter a pronoun (or noun) that has multiple possible referents, a sustained anterior negativity has been observed (Nieuwland & Van Berkum, 2006). This negativity is known in the event-related potential (ERP) literature as the “Nref effect.” Its amplitude is sensitive to the number of possible referents for a noun or pronoun, as determined by the situation model (Nieuwland, Otten, & Van Berkum, 2007), and it is also sensitive to the degree of contextual bias towards a single referent (Nieuwland & Van Berkum, 2006). This effect has been linked to working memory, as participants with a higher reading span elicit a larger amplitude Nref effect (Nieuwland & Van Berkum, 2006; see also Nieuwland, 2014). Overall, the Nref literature suggests that comprehenders (at least, those with a high reading span) rapidly appreciate when there is ambiguity associated with a pronoun. The Nref effect has been proposed to indicate an increase in processing load as comprehenders grapple with multiple interpretations of the text (Nieuwland & Van Berkum, 2008), although the effect persists well beyond the disambiguation point (Nieuwland, Otten, & Van Berkum, 2007).

These findings suggest that pronouns that cannot be readily resolved elicit a rapid cognitive response from the comprehender, which in some cases appears to entail a transient increase in processing load (Callahan, 2008; Kennison & Trofe, 2003; Nieuwland & Van Berkum, 2008). However, such findings do not address how comprehenders respond to pronouns that *do* match the most salient entity in the preceding discourse. To address this, Arnold, Eisenband, Brown-Schmidt, and Trueswell (2000) used eye-tracking in a visual world paradigm to assess how quickly a pronoun’s antecedent is fixated relative to

competing characters in a scene. They found that looks to a pronoun's most likely referent can increase within approximately 200-600 ms after pronoun onset. This is the clearest evidence to date that unambiguous pronouns (e.g. "he" in the context of a male and female character) may be rapidly linked to their antecedents. However, participants in this study were asked to judge whether the picture display matched what they heard, a task that they could carry out only if they successfully resolved the pronoun to its antecedent. Thus, the task in this study may have implicitly biased participants towards thorough referential processing. Järvikivi, van Gompel, Hyönä, and Bertram (2005) obtained a similar result using ambiguous pronouns and a different task. Participants were asked to follow pictures of two characters (potential antecedents) while listening to two-sentence long "ministories" for comprehension. They were also occasionally asked to continue the ministry using the stated location and presented characters. Järvikivi and colleagues found rapid emergence of effects of grammatical role and order of mention on the proportion of fixations to each possible referent following the pronoun. Their earliest effect (grammatical role) emerged on a similar time scale to Arnold et al.'s effect of gender, at 480-690 ms after average pronoun onset. Thus, although Järvikivi and colleagues used ambiguous stimuli, their results support Arnold et al.'s (2000) general conclusions and suggest that pronoun resolution processes may be rapidly initiated under a variety of task conditions.

Mixed evidence for semantic "reactivation"

The above findings have demonstrated that comprehenders are capable of linking pronouns to their antecedents relatively rapidly, at least under some task conditions. However, the precise nature of this linking process remains unclear. One line of work, often using written probes, has sought to determine the extent to which establishing a link between a pronoun and its antecedent involves reactivation of the antecedent's semantics. The results of these studies are disparate. In one representative study, for example, Shillcock (1982) had participants perform a lexical decision task on either a semantically related or unrelated visual probe word at three critical points (denoted below using *) while listening to sentences such as the following:

The teacher * did not board the train, for the * simple reason that it/he * was not going to the South Coast of England.

The critical probe words were either semantically related to one of the antecedents (e.g., “school” – related to teacher) or were unrelated to either (e.g., “street”). Shillcock found that priming for the semantically related probe word was larger after a pronoun that referred to the semantically related antecedent (“he”) than after an alternative pronoun that referred to something else (“it”).

However, within the probe literature, separate experimental manipulations have inspired different conclusions about how pronouns influence the activation profile of semantic constituents. Some manipulations consistently elicit differential priming. For example, passages containing a pronoun that refers back to a particular entity induce priming for probe words related to that entity, when compared with passages that instead contain one fewer reference to the critical entity (e.g., Bever & McElree, 1988; Cloitre & Bever, 1988; Leiman, 1982; MacDonald & MacWhinney, 1990; Speelman & Kirsner, 1990). Shillcock (1982) is a particularly carefully controlled variant on this theme. Other studies which, like Shillcock, manipulated the identity (and thus the referent) of a pronoun, but otherwise held the passage and probe words constant, also found evidence that the pronoun’s identity modulated priming of a subsequent probe word, consistent with a semantic reactivation account (e.g., Chang, 1980; Nicol, 1988). Thus, quite a large body of evidence suggests that the presence of a pronoun may facilitate retrieval of information related to its antecedent.

Several other manipulations, however, complicate the appealing narrative that the act of reading a pronoun induces reactivation of its antecedent. First, where multiple time points have been assessed, antecedent-related probe words do not always exhibit a larger priming effect, or faster response times, following the pronoun than preceding it (e.g., Gernsbacher, 1989; Shillcock, 1982; cf. Nicol & Swinney, 2003). Second, if the presence of an overt pronoun is in itself changing the activation level of its referent, one would expect that removing the pronoun should in turn attenuate or delay this effect. However, whenever overt pronouns are directly compared to null pronouns in the probe literature, no difference in the activation

profile is observed (Bever & McElree, 1988; Corbett & Chang, 1983; Emmorey & Lillo-Martin, 1995; Tyler & Marslen-Wilson, 1982).

In other work, when the probe word employed is either the pronoun's antecedent or a similarly salient control word (e.g., the name of a primary character who has not been referred to as recently), results are mixed, such that the activation level of the salient control word is not always below that of the pronoun's antecedent (Greene et al., 1992; Emmorey & Lillo-Martin, 1995; MacDonald & MacWhinney, 1990). In probing the time course of this pattern, Gernsbacher (1989) showed that response times to the antecedent probe did not differ from those to a salient control word when probes were presented immediately following the pronoun. However, with more time after the pronoun, response times to both types of probe words increased, and those for the salient control probes increased faster than for the antecedent probe words. Thus, Gernsbacher concluded that reading a pronoun induces suppression of non-referent entities, rather than facilitation of the pronoun's referent. Indeed, Shillcock (1982)'s results, among others, can be interpreted as support for a suppression-based account, in light of how the activation pattern evolves across time (see Sanford & Garrod, 1989).

Such disagreements highlight difficulties with drawing conclusions from the kind of probe tasks that have been widely used in this literature. Responding to a midsentence probe, for example, may disrupt normal comprehension and tend to elicit more strategic processing. Moreover, as discussed more extensively in McKoon, Ratcliff, and Ward (1994), there are challenges in determining an appropriate baseline and appropriate set of control probes. Thus, it seems important to examine this question using other tasks and methods, especially those, such as eye-tracking (e.g., Lago, 2014) and ERPs, that can be collected as participants simply read for comprehension.

One other advantage of the ERP technique is its ability to differentiate among several distinct cognitive processes that can all fall under the umbrella term of "activation." "Activation" has been variably defined, but it consistently refers to a process that enhances the accessibility of some cognitive structure for additional processing (e.g., Collins & Loftus, 1975; Cowan, 1993; Gernsbacher, 1991). In many cases, behavioral priming evidence, taken alone, is unable to further specify what activation might entail, beyond a general facilitation of processing. However, further specifying the mechanism of activation becomes necessary

when building a model of reference in terms of more general memory and attentional processes. In the semantic priming literature, activation is often used to refer to changes of state within long-term memory and/or the transfer of information from long-term memory to working memory (Collins & Loftus, 1975; Kutas & Federmeier, 2011; Neely, 1977). However, within the discourse processing literature, the idea that structures may change their relative “activation” level within working memory is also central to several models (Lewis & Vasishth, 2005; McElree, Foraker, & Dyer, 2003; see also Ericsson & Kintsch, 1995 for discussion). In some cases, the facilitation of structures within working memory is described as an attentional process separate from long-term memory access (e.g., Garrod, Freudenthal, & Boyle, 1994; Foraker & McElree, 2007), while in other cases, it is described as being the same as activation in long-term memory (Lewis & Vasishth, 2005). In the current study, we will specifically target measures of long-term memory access, to determine whether activation in this sense plays a role in pronominal processing, as suggested by, for example, Lewis & Vasishth's (2005) ACT-R model of language comprehension.

Tracking the time course of semantic memory activation using ERPs

The current study will employ two well-studied phenomena in the ERP literature to answer questions about semantic memory access during pronoun comprehension: the ERP concreteness effect and the N400.

Concreteness, or the extent to which a word or phrase's meaning can be “experienced by the senses” (Spreeen & Schulz, 1966, p. 460), has been shown to affect memory, word learning, and behavior on a variety of tasks (reviewed in Schwanenflugel, 1991). Concreteness-related modulations of brain activity during word reading can be seen with fMRI (reviewed in Binder, Desai, Graves, & Conant, 2009), MEG (Dhond, Witzel, Dale, & Halgren, 2007), and EEG (Kounios & Holcomb, 1994; Holcomb, Kounios, Anderson, & West, 1999). Such “concreteness effects” are explained under all major accounts as being due to a qualitative and/or quantitative difference in the information retrieved from long-term semantic memory when invoking concrete vs. abstract concepts (Paivio, 1986; Schwanenflugel, Harnishfeger, & Stowe, 1988; Holcomb et al., 1999; Kousta, Vigliocco, Vinson, Andrews, & Del Campo, 2011; Barber, Otten, Kousta, & Vigliocco, 2013). Importantly

for the question at hand, the ERP concreteness effect arises relatively rapidly in the course of word processing, making it a useful measure for examining the dynamics of semantic (re)activation.

More specifically, concrete words elicit a negative voltage change relative to abstract words beginning at around 300 ms post stimulus onset, and extending as late as 900 ms or longer. From roughly 300 to 500 ms, the difference between concrete and abstract words is apparent across a broad distribution of scalp sites, with maximal differences at frontal or central sites using a mastoid reference (Kounios & Holcomb, 1994; Holcomb et al., 1999; West & Holcomb, 2000). After roughly 500 ms, the distribution of the effect remains anterior and sometimes narrows to predominantly frontal sites¹ (Holcomb et al., 1999; Lee & Federmeier, 2008; cf. Kounios & Holcomb, 1994; Zhang, Guo, Ding, & Wang, 2006). The late frontal portion of the concreteness effect is preferentially preserved (West & Holcomb, 2000) or enhanced (Farah et al., 1989; Gullick, Mitra, & Coch, 2013) during tasks designed to encourage the use of mental imagery, when compared to semantic tasks that do not require imagery. For this reason, the ERP concreteness effect is often described as a modulation of the centro-posterior N400 component, combined with a sustained and temporally overlapping frontal negativity that has been linked to mental image generation (West & Holcomb, 2000; Huang, Lee, & Federmeier, 2010).

The fact that the ERP concreteness effect encompasses a modulation of the N400 reinforces the idea that reading concrete words initiates the activation of a greater amount of semantic information than reading abstract words. The N400 is an extensively studied negative-going ERP component that is thought to index the process of activating information stored in long-term semantic memory (see Kutas & Federmeier, 2011 for review). The amplitude of the N400 component is sensitive to a variety of lexical and contextual factors, as well as their interactions, such that greater N400 amplitudes are seen under circumstances in which more new information is elicited by an incoming word – e.g., because it is less expected, is being encountered for the first time (as opposed to being repeated), or is less related to prior words (not semantically primed) (see Kutas & Federmeier, 2011). Similarly,

¹ Use of a nose tip reference can affect the distribution and polarity of the concreteness effect, particularly when mental imagery is encouraged (Farah, Weisberg, Monheit, & Peronnet, 1989; Nittono, Suehiro, & Hori, 2002; Welcome, Paivio, McRae, & Joanisse, 2011).

words with a larger orthographic neighborhood size and/or more lexical associates elicit a larger N400 (Laszlo & Federmeier, 2011), presumably because such words tap into a more extensive part of the lexico-semantic network. The larger N400 amplitude observed for concrete words has thus been interpreted as arising because these words evoke activity across a denser network of semantic and contextual associations than abstract words, and/or carry with them a body of nonlinguistic sensory associations that abstract words do not (Barber et al., 2013; Holcomb et al., 1999; Lee & Federmeier, 2008).

In the current study, we took advantage of the ERP concreteness effect as a means of probing the extent to which semantic processing of a pronoun mirrors that seen when the antecedent was originally encountered. As such, we manipulated the concreteness of a sentence-medial noun and examined the effect of this concreteness manipulation both at the critical noun itself and later, at a pronoun downstream. For example:

Concrete: The **beer** did not go over well, since **it** didn't suit the guests' taste.

Abstract: The **joke** did not go over well, since **it** didn't suit the guests' taste.

Having established, through norming, that the pronoun (either "it" or "they") is consistently interpreted as referring back to the critical noun, we can examine whether an ERP concreteness effect is apparent at the pronoun. If so, this suggests that semantic information about the critical noun has been reactivated at the pronoun. If not, this suggests that reading a pronoun does not always induce a comprehender to reactivate semantic information about its antecedent.

In addition, the N400's sensitivity to semantic priming makes it possible to conduct a secondary analysis of our ERP data that widens the scope of our investigation to all words following the critical noun. As already described, N400 amplitudes decrease (i.e., the waveform becomes more positive) when a word is preceded by a semantically-related prime. Such semantic priming effects are also evident within sentences, although the effect typically dissipates as the distance between the "prime" word and the "target" increases (Van Petten, Weckerly, McIsaac, & Kutas, 1997; Federmeier, Van Petten, Schwartz, & Kutas, 2003; cf. Van Petten, 1993). By tracking the degree of lexical association between all downstream content words and the critical noun, we can monitor the evolution of the semantic priming effect on

the N400 across the sentence. Of course, this analysis depends on the fact that our stimuli were constructed such that *two* critical nouns (e.g. “joke” and “beer”), which have differing degrees of semantic similarity to each downstream content word, are paired with each otherwise identical sentence frame (e.g. “The ___ did not go over well...”). Using this secondary analysis, we can assess whether N400 priming due to the preceding critical noun reemerges following the pronoun. If so, this would be further evidence of semantic reactivation of the critical noun’s semantics at or around the pronoun, but without the constraint that such reactivation must occur immediately, time-locked to the pronoun’s onset.

Summary of predictions

In summary, if readers are reactivating the semantic properties of a pronoun’s antecedent immediately upon encountering that pronoun, we should expect to see an ERP concreteness effect at the pronoun itself. If the semantics of the antecedent are reactivated -- even at a more variable delay with respect to encountering the pronoun -- we should also expect to see a larger effect of N400 priming by the critical noun antecedent on content words immediately following the pronoun, as compared to those immediately preceding it. The absence of these effects would suggest that reading the pronoun did not cause the reader to reactivate semantic information about its antecedent in long-term memory.

Methods

Participants

Data are reported from a sample of 32 participants (12 males, aged 18-22, mean age 20), all native English speaking undergraduates recruited at the University of Illinois, who were compensated with course credit or payment. 4 additional participants were dropped due to below chance performance or excessive trial loss. Handedness was assessed using the Edinburgh handedness inventory (Oldfield, 1971). All participants were right-handed (mean score: .79, where 1 denotes strongly right-handed and -1, strongly left-handed). 18 reported having left-handed family members. No participants had major exposure to languages other than English prior to the age of 5. All had normal or corrected-to-normal vision, and none of the participants had a history of neurological or psychiatric disorders or brain damage, or was using neuroactive drugs. Participants were randomly assigned to one of the two experimental lists.

Materials

124 concrete nouns were matched with 124 abstract nouns using an automated algorithm that drew from a larger pool of 2220 English nouns. The algorithm ranked noun pairs based on the difference between the words' Kucera-Francis frequency, orthographic neighborhood size (operationalized as "orthographic Levenshtein distance 20," see Yarkoni, Balota, & Yap, 2008), number of letters, and familiarity. Noun pairs were then selected from among those with the highest level of similarity, such that there was no significant difference between the concrete and abstract noun lists on any of the above features (all p 's > .1). Table 1 provides means and standard errors by condition for concreteness and each of the matched features. All concrete nouns had a concreteness rating of 500 or more using the MRC Psycholinguistic Database (range: 502-634) and abstract nouns of 499 or less (range: 234-494). Based on the SUBTLEX corpus, all nouns had a dominant word sense of "noun", and were used as a noun more than 80% of the time.

These concrete and abstract nouns were embedded into 124 sentence frames such that they were always the second word and were later referred back to by a pronoun (“it” or “they”). For example:

Concrete: The **duck** was unaware that **it** was being hunted.

Abstract: The **prey** was unaware that **it** was being hunted.

A complete list of experimental stimuli is provided in Appendix A. The feature-matched pairs of nouns were not always used in the same sentence frame, but, rather, were distributed across the set to ensure that critical sentences sounded as natural as possible. Although words and phrases downstream of the critical noun may have adopted a different shade of meaning across the two conditions, we aimed to keep the dominant interpretation of downstream word meanings as similar as possible across conditions and avoided any cases in which a word downstream of the critical noun in one condition would be interpreted as a homophone with a radically different meaning in the other condition.

After norming (described below), the 248 critical sentences (generated from 124 sentence frames) were divided into two presentation lists, such that no item was shown more than once per list. Critical sentences were further organized into four blocks within each presentation list, and the four blocks were presented in one of four different sequences. Sentences were pseudorandomized within each block such that the same condition (concrete or abstract) did not appear more than three times in a row. Most critical nouns did not appear elsewhere in the stimulus set, and, wherever possible, repetitions were arranged to occur after the critical noun had already been read. As no fillers were used, each participant thus read 124 total sentences, 62 with a concrete critical noun and 62 with an abstract critical noun.

Norming

To ensure that the critical pronouns were interpreted as referring back to the critical noun, we normed the stimuli using Amazon Mechanical Turk. 248 critical sentence fragments up to and including the critical pronoun were combined with 32 additional sentence

fragments that were ultimately not used and divided into 8 lists of 35. Participants (N = 83) read the sentence fragments and were asked to complete them with the first reasonable completion that came to mind. After going through all of the fragments once, participants were then shown each fragment again with their completion and asked to indicate what the pronoun referred to using a dropdown menu. The dropdown menu listed the critical noun, other nouns in the sentence, “null”, “situation(s)”, “new entity/ies” and “other.” Participants who selected “other” were then able to type their response in the textbox provided. Because of some trial loss and partial completions of the lists, between 8 and 12 reference annotated cloze responses (mean = 10.16) were ultimately collected for each sentence fragment.

The proportion of the time that the critical noun was selected as the referent of the pronoun was calculated for each sentence fragment. For the 124 final sentence frames, the mean proportion of the time that the critical noun was selected as the referent of the pronoun was 93.8% in the concrete condition and 87.4% in the abstract condition. We fit logistic regression models to the norming data to confirm whether there was a difference between conditions in the distribution of responses. A model with condition as a fixed effect and subject and item random intercepts revealed a small but reliable effect of condition (estimated effect of concrete category = 1.18, $z = 4.61$, $p < .001$; intercept = 2.99). For the average item and participant under the model there was a 98.5% chance of the critical noun being selected as the referent in the concrete condition, and a 95.2% chance in the abstract condition. Participants thus consistently interpreted the critical pronoun as referring back to the critical noun in both conditions, as designed.

Procedure

Participants read one of two lists of 124 sentences for comprehension and were told that at the end of the experiment they would be asked to take a quiz that would test how well they understood and remembered the sentences. They were discouraged from memorizing the sentences and told to read them for comprehension as they would a book or a magazine. Prior to the main experiment, participants read two short filler sentences to practice not blinking or moving as they read.

Sentences were presented one word at a time in 30-pt font on a computer screen 100 cm from the participant. Each word was presented at the center of the screen for 200 ms, followed by 300 ms of blank screen. Prior to the start of each sentence, a fixation cross was presented at the center of the screen for 500 ms followed by a variable time delay of 200-400 ms to mitigate the effects of anticipatory slow potentials on the averaged waveform. Progress through the experiment was self-paced, in that participants pushed a button to move on to the next sentence. The experimenter checked in with the participants at three break points and encouraged them to rest before continuing. The participant controlled the duration of the breaks. The recording session lasted approximately 30 minutes.

After reading the sentences, participants took a word recognition quiz. They were presented with a list of 248 nouns, roughly half of which they had seen while reading the sentences and half of which they had not. They were asked to circle all the words on the list that they remembered reading on the computer screen. The order of words on the list was initially pseudo-randomized by block and condition but held constant across participants. The response to each word was then coded as a hit, miss, false alarm, or correct rejection.

EEG Data Acquisition and Preprocessing

The electroencephalogram (EEG) was recorded from 26 silver/ silver-chloride electrodes evenly spaced over the scalp. The sites are midline prefrontal (MiPf), left and right medial prefrontal (LMPf and RMPf), left and right lateral prefrontal (LLPf and RLPf), left and right medial frontal (LMFr and RMFr), left and right mediolateral frontal (LDFr and RDFr), left and right lateral frontal (LLFr and RLFr), midline central (MiCe), left and right medial central (LMCe and RMCe), left and right mediolateral central (LDCe and RDCe), midline parietal (MiPa), left and right mediolateral parietal (LDPa and RDPa), left and right lateral temporal (LLTe and RLTe), midline occipital (MiOc), left and right medial occipital (LMOc and RMOc), and left and right lateral occipital (LLOc and RLOc). The midline central (MiCe) electrode was placed where the “Cz” electrode would appear using the international 10-20 system. Eye movements were monitored via a bipolar montage of electrodes on the outer canthus of each eye. Blinks were detected by an electrode below the left eye. Impedances were kept below 5 K Ω . Signals were amplified with a .02–250 Hz bandpass using a

BrainVision amplifier and digitized at 1000 Hz. Data were referenced online to the left mastoid and rereferenced offline to the average of the left and right mastoids. Each trial consisted of a 1000 ms epoch preceded by a 200 ms prestimulus baseline. Trials contaminated by eye movements, blinks or other recording artifacts were rejected offline, using individually adjusted threshold parameters for each participant. Average trial loss was 11.0% for concrete nouns, 11.9% for abstract nouns, 9.6% for pronouns in the concrete condition, and 11.4% for pronouns in the abstract condition. A digital lowpass Butterworth IIR filter with a 30 Hz half-amplitude cutoff and 12 dB/octave roll-off was applied prior to statistical analysis.

Latent Semantic Association Analysis

As a secondary analysis, we assessed the degree of semantic association between the critical noun and each subsequent content word. Due to semantic priming, more strongly associated content words are expected to have a smaller amplitude N400. Because each content word within each sentence frame was preceded by one of two different critical nouns, we could control for all factors that might impact N400 amplitude at downstream content words aside from those stemming from the critical noun and its relationship with the other words across the sentence. Based on prior work (Van Petten et al., 1997; Federmeier et al., 2003), we expected the N400 priming effect to dissipate at word positions further away from the prime. The critical question, then, was whether the priming effect would re-emerge at word positions following the pronoun, which would be indicative of semantic reactivation.

To examine N400 priming effects from the critical noun and the critical pronoun, we extracted 238 content words that either preceded (N = 133) or followed (N = 105) the pronoun (Appendix A shows all the selected words). Content words consisted of all non-sentence final nouns, non-auxiliary verbs, adjectives, adverbs, and proper names within 3 linear word positions of the pronoun². Lexical association scores between each content word and the critical noun were assessed using Latent Semantic Analysis (Landauer & Dumais, 1997; Landauer, Foltz, & Laham, 1998; pre-computed scores available at

² One sentence pair was excluded from the analysis because both its concrete and abstract critical nouns were not present in the corpus used to generate our semantic similarity scores.

<http://lsa.colorado.edu/>, corpus: General Reading up to First Year College, 300 factors). The degree of semantic similarity to the critical noun of content words that preceded the pronoun (mean = .147, standard deviation = .106) and those that followed it (mean = .142, standard deviation = .111) was similar ($t = .517$, $df = 439.4$, $p > .1$). As a measure of N400 amplitude, we extracted mean voltage of the EEG signal, aggregated across 8 centro-parietal electrode sites (MiCe, LMCE, RMCE, MiPa, LDPa, RDPa, LMOC, and RMOC) between 300 and 500 ms, timelocked to the onset of each content word, with a 200 ms prestimulus baseline. 10.8% of trials were eliminated due to artifacts.

Results

Behavioral

In order to confirm that participants were attending to the sentences, we analyzed their post-test responses using the sensitivity index d' , and the information retrieval statistic of precision. In this context, precision may be defined as the proportion of the time that words that were indicated as having been read previously were in fact read previously, with a precision of .508 indicating that words had been selected at random. Participants consistently discriminated words that they had read before during the experiment from words that they had not encountered previously (d' : mean = 1.12, sd = .42, range = (.44, 2.06); precision: mean = .81, range = (.67, 1)). A single subject had a false alarm rate of 0, and the corresponding d' score was corrected using the log linear rule (Hautus, 1995). We also examined the probability of indicating that a word had been read before using a logistic regression analysis with crossed random effects for subject and word, and fixed effects of correct response and concreteness condition. Using likelihood ratio tests based on nested model comparisons, we confirmed a main effect of “correct response” ($\chi^2_1 = 1086, p < .01$) and an interaction between “correct response” and “concreteness condition” ($\chi^2_1 = 16.64, p < .01$). The main effect of concreteness condition did not reach statistical significance ($\chi^2_1 = 2.362, p > .1$). Concrete words were marginally less likely to be incorrectly indicated as having been read before (intercept = -2.306, $\beta = -.227, SE = .135, z = -1.679, p < .1$), and significantly more likely to be correctly indicated as having been read before ($\beta = .529, SE = .129, z = 4.097, p < .01$). For the average word and participant, the model estimate of the probability of indicating a word had been read before was: 9.1% (7.1%, 11.7%) for abstract, unseen words; 7.3% (5.5%, 9.5%) for concrete, unseen words; 35.5% (29.8%, 41.7%) for abstract, previously read words; 48.4% (36.7%, 49.0%) for concrete, previously read words. 95% confidence intervals based on bootstrapping with 2000 replications are indicated in parentheses.

Concreteness Effects at the Noun

Following Huang et al. (2010), repeated measures ANOVA models were run on the mean EEG amplitude at two different latencies (300-500 ms post stimulus onset, to capture the N400, and 500-900 ms post stimulus onset, to capture the extended part of the frontal negativity) with concreteness and electrode site as factors. Models were run separately on electrodes at the front half (11 sites) and back half (13 sites; excluding the two lateral sites closest to the reference, LLTe and RLTe) of the head.

Between 300 and 500 ms post stimulus onset, repeated measures ANOVAs with condition and electrode site as factors revealed significant concreteness effects over both the front ($F(1,31) = 11.29, p < .01, \eta_G^2 = .037, \eta_P^2 = .267$) and back half of the head ($F(1,31) = 12.60, p < .01, \eta_G^2 = .049, \eta_P^2 = .289$); neither effect interacted with electrode (F 's < 2). Across both regions, responses were more negative in the concrete than abstract condition, replicating typical ERP concreteness effects in this time window (e.g., Kounios & Holcomb, 1994; Holcomb et al., 1999).

Between 500 and 900 ms, over the front half of the head, there was no significant main effect of concreteness ($F(1,31) = 2.008, p > 0.1$), nor interaction between concreteness and electrode site ($F(10,310) < 1$). Over the back half of the head, there was also no significant main effect of concreteness ($F(1,31) < 1$). There was, however, a significant interaction between concreteness and electrode site ($F(12,372) = 3.165, \epsilon_{GG} = .228, p < .05, \eta_G^2 = .002, \eta_P^2 = .093$). Electrodes closer to the middle of the head tended to have a numerically more negative response in the concrete condition, but electrodes closer to the back of the head tended to have a numerically more positive response in the concrete condition.

Concreteness Effects at the Pronoun

The same analysis strategy used for the noun position was used to examine effects at the pronoun. Between 300 and 500 ms, there was no significant main effect of concreteness ($F(1,31) = 1.087, p > .1$). There was a marginal interaction between concreteness and

electrode site ($F(10,310) = 2.178$, $\epsilon_{GG} = .359$, $p < .1$, $\eta_G^2 = .004$, $\eta_p^2 = .066$), but no individual frontal sites showed a significant effect of concreteness following Bonferroni correction. Over the back half of the head, there was also no significant main effect of concreteness ($F(1,31) < 1$), nor was there an interaction between concreteness and electrode site ($F(12,372) = 1.822$, $\epsilon_{GG} = .225$, $p > .1$). There were also no significant effects in the 500 to 900 ms window: frontal sites ($F(1,31) = 2.645$, $p > .1$; interaction $F(10,310) < 1$); posterior sites ($F(1,31) < 1$; interaction $F(12,372) < 1$).

Given the null effects in the ANOVA analysis, we turned to a Bayes factor analysis to directly test the degree of support for the null hypothesis compared to the alternative hypothesis. As our dependent measure, we used the mean voltage between 300-500 ms aggregated across 8 centro-parietal electrode sites (MiCe, LMce, RMce, MiPa, LDPa, RDPa, LMOc, and RMOc) where N400 effects tend to be largest (Kutas & Federmeier, 2011). We estimated the anticipated effect size of concreteness on the N400 at the pronoun under the alternative hypothesis as being equal to the obtained effect size at the noun ($-.627 \mu\text{V}$) for this same time window and set of channels. Under this analysis, the null hypothesis was found to be almost 5 times ($K = 4.715$) more likely than the alternative hypothesis.

Lexical Association Effects Before and After the Pronoun

The absence of a concreteness effect at the pronoun suggests that semantic information about the pronoun's antecedent is not being (re)activated in long term semantic memory within the first second after viewing the pronoun. However, it is possible that semantic information about the antecedent is reactivated, but that this process requires more time. Alternatively, reactivation may not have a consistent temporal relationship with the onset of the pronoun. To examine these possibilities, we performed a secondary analysis in which we tracked semantic association priming of the N400 across word positions. Nested linear mixed effects models predicting N400 amplitude of the content words were fit to the single trial data and compared (see Payne, Lee, & Federmeier, 2015). All models included crossed random effects of subject and content word.

We first confirmed that our basic analysis approach was able to capture an interaction between a) the effect of semantic association with the critical noun on the N400 amplitude of content words, and b) the number of words intervening between the content word and the critical noun (see Figure 3). We expected that the effect of semantic association would decrease with increasing distance between the critical noun prime and the downstream target. We fit a model with fixed effects of distance to the critical noun and semantic similarity score, as well as their interaction. Distance to the critical noun was represented categorically using 5 factor levels (1-3, 4-6, 7-9, 10-12, and > 12 words following the critical noun). This paralleled the factor structure used to examine content words surrounding the pronoun in a subsequent analysis. Since only a single non-sentence final word was greater than 12 word positions following the critical noun, only word positions 1 through 12 are further discussed. Word positions 1-3 and 4-6 exhibited a significant effect of semantic association (1-3: $\beta = 3.488$ (.846, 6.142), SE = 1.346, $t = 2.591$; $\chi_1^2 = 6.700$, $p < .01$; 4-6: $\beta = 3.770$ (.662, 6.882), SE = 1.585, $t = 2.378$; $\chi_1^2 = 5.652$, $p < .05$; 95% profile confidence intervals indicated in parentheses next to each β estimate). The effect of semantic association was not apparent at later word positions (7-9: $\beta = -.245$ (-3.594, 3.114), SE = 1.708, $t = -.143$; $\chi_1^2 = .021$, $p > .1$; 10-12: $\beta = -.897$ (-9.300, 7.530), SE = 4.289, $t = -.209$; $\chi_1^2 = .044$, $p > .1$). Refitting the model with a simple contrast of word position ≤ 6 vs. > 6 following the critical noun resulted in a significant interaction between word position and the effect of semantic similarity ($\chi_1^2 = 4.588$, $p < .05$). Thus, we were able to detect the expected decrement of priming with distance (Van Petten et al., 1997; Federmeier et al., 2003).

Next, we examined non-sentence final content words up to 3 word positions prior to and following the pronoun. Content words immediately following the critical noun were excluded from the analysis, since the concreteness effect at the critical noun overlapped with the baseline period of the immediately following word. The following fixed effects were included in a full model used to test for the presence of an interaction between semantic similarity score and linear word position relative to the pronoun:

- a) an indicator variable set to 1 when the content word followed the critical pronoun, and 0 when it preceded it
- b) the grand mean centered semantic similarity score between the content word and the critical noun
- c) the interaction between a) and b) -- that is, the grand mean centered semantic similarity score for words following the critical pronoun (and 0 for words preceding it)
- d) the grand mean centered linear word position distance of the pronoun from the start of the sentence

Likelihood ratio tests involving nested model comparisons were used to assess all fixed effects. Content words surrounding pronouns near the start of the sentence showed larger N400s than those that were near pronouns later in the sentence ($\beta = .303$, $SE = .087$, $t = 3.477$; $\chi_1^2 = 11.92$, $p < .001$). This indirectly replicates the finding that N400 amplitudes tend to decrease with increasing linear word position. Words with a higher semantic similarity score to the critical noun showed an attenuated N400 amplitude overall ($\beta = 3.597$, $SE = 1.469$, $t = 2.448$; $\chi_1^2 = 5.984$, $p < .05$). However, there was no main effect of following the pronoun, relative to preceding it ($\beta = -.151$, $SE = .272$, $t = -.555$; $\chi_1^2 = .307$, $p > .1$). Critically, there was no interaction between the effect of semantic similarity and whether or not the content word followed the pronoun ($\beta = -2.578$, $SE = 2.199$, $t = -1.172$; $\chi_1^2 = 1.373$, $p > .1$). Numerically, content words following the pronoun tended to show a smaller effect of semantic similarity score than those preceding the pronoun.

Discussion

The current study constrains our understanding of the cognitive processes involved in linking a pronoun to its antecedent during online comprehension. Taking advantage of the continuous nature of the EEG measurement and the functional specificity afforded by several well-studied ERP components, we assessed whether the establishment of referential links entails reactivation of the semantics of the antecedent in long-term memory (as suggested by, e.g., Keshtiari & Vasishth, 2013; Lewis & Vasishth, 2005). We used two reliable semantic ERP effects to address our question: the ERP concreteness effect (which is further divisible into a frontal and centro-posterior component; West & Holcomb, 2000; Huang et al., 2010), and N400 semantic priming (Van Petten, 1993; reviewed in Kutas & Federmeier, 2011). In both cases, we replicated a known ERP effect indicating semantic memory access in the vicinity of the antecedent, but were unable to detect comparable effects at or around the pronoun.

To examine semantic activation at the time a critical word was encountered, we manipulated the concreteness of the sentence-initial noun, taking advantage of the fact that there is a well-replicated ERP effect of concreteness. Concrete words are known to elicit larger N400 responses, with a more frontal distribution than typical N400 effects. The frontal part of the concreteness effect extends beyond the N400 time window and can be distinguished from the N400 response using visual half-field presentation (Huang et al., 2010) and task manipulations (West & Holcomb, 2000). The N400 is thought to be larger for concrete words because they contact a richer semantic network (Federmeier & Laszlo, 2009), and the frontal portion of the concreteness effect has been hypothesized to reflect mental imagery (West & Holcomb, 2000; Gullick et al., 2013). Prior studies showing the ERP concreteness effect had used single words (e.g., Barber et al., 2013), words in phrases (e.g., Huang et al., 2010), or sentence-final words (e.g., Holcomb et al., 1999). We replicated the basic pattern for concreteness effects at the critical noun, although here the frontal portion of the effect did not last as long as in prior work using single word or sentence-final targets. We thus extended previous results by showing that ERP concreteness effects are observable at an early, sentence-medial position. Our results thus show that when nouns are encountered,

even early in a sentence, their semantics are activated and some degree of mental imagery may be engaged.

We could also see the consequences of semantic access of the critical noun in the form of downstream priming effects. Latent semantic analysis (LSA) was used to obtain a metric of semantic similarity between downstream content words and the critical noun. Because all sentence frames were used with two different critical nouns, we were able to isolate the effect of semantic similarity from other sentential factors, such as content word identity. Content words within several word positions of the critical noun showed attenuated N400s when the critical noun was semantically related to them, replicating prior work showing intra-sentential N400 semantic priming (Van Petten, 1993; Van Petten et al., 1997; Federmeier et al., 2003). This finding also supports recent work showing that LSA metrics are a good predictor of N400 amplitude (Van Petten, 2014). Thus, as expected, when a noun is encountered, comprehenders appear able to quickly access its semantics, and that semantic information is demonstrably active at the next several word positions. Our cross-sentence analysis indicates that the strength of this priming effect attenuates after roughly 6 word positions following the critical noun. As this was also the average position of the clause boundary in our experimental items (mean = 5.77), it is possible that the attenuation reflects, at least in part, differential accessibility of words across the clause boundary (Gernsbacher, Hargreaves, & Beeman, 1989). To summarize, then, a set of semantic processes are engaged at the noun and these also affect how words downstream are processed, particularly those in close proximity. The critical question is whether similar processes are re-engaged when a comprehender encounters a word, such as a pronoun, that refers back to that critical antecedent.

We examined whether pronouns referring to concrete vs. abstract critical noun antecedents would also exhibit an ERP concreteness effect similar to that obtained at the noun. We found that they do not. The null finding was supported by a Bayes factor analysis showing that the data provide 4.7 times more support for the null hypothesis of no effect than for the alternative hypothesis of an N400 effect of comparable size to that measured at the noun antecedent. The absence of an N400 effect of concreteness suggests that participants did not reengage the semantics of the noun antecedent immediately upon reading the pronoun. The absence of a concreteness-related frontal negativity also suggests

that participants did not engage in mental imagery of the antecedent upon reading the pronoun.

Using the LSA-based N400 semantic priming analysis, we also examined whether pronouns induced semantic reactivation of the antecedent, but at some short delay. We found that a small main effect of semantic similarity to the critical noun persisted at word positions close to the pronoun. The effect was smaller than at word positions immediately following the critical noun, which may be because the noun was no longer active in memory, or because there was less and less additional information related to the noun remaining to be accessed at each subsequent content word. Critically, however, we failed to see any increase in the size of the N400 semantic priming effect at content words that followed the pronoun, compared with those that preceded it. Moreover, the effect of semantic priming following the pronoun did not statistically differ from zero. Overall, our results are consistent with the idea that the semantics of the critical noun were not reengaged in long-term memory at the pronoun, or up to 2000 ms after it. There are some limitations to this analysis, as two different sets of words preceded and followed the pronoun (although the overall strength of their relationship to the critical noun was similar); thus, it is possible that the influence of specific lexical properties of the different content words could make finding an effect, as well as comparing effect sizes before and after the pronoun, more difficult. Also, not all sentences contributed content words to the critical analysis that both preceded and followed the pronoun. Thus, the comparison of the priming effect before and after the pronoun is across distinct but overlapping sets of sentence frames. Nevertheless, with this analysis, we obtained the expected effects of both semantic similarity and word position, as well as the expected interaction between these factors in positions following the critical noun (where the caveats described above also apply). Thus, this analysis, like that for the concreteness effect, reveals important differences in what happens during and after the point wherein readers encounter a noun versus encounter a pronoun referring back to that noun.

We thus have strong evidence for differential processing at the noun and pronoun, which points to a qualitative distinction in the way that nouns and pronouns are processed for comprehension. As discussed next, these processing differences may arise because nouns and pronouns differ in their semantic properties and usage statistics. Such processing differences, in turn, have implications for anaphoric reference. Specifically, our data are

consistent with the claim that long-term semantic memory access is not necessary for successful anaphor resolution (e.g., Garrod et al., 1994).

Whereas nouns are directly associated with a rich semantic representation, overt pronouns (out of context) signal only a sparse encoding of referent semantics, such as gender, number and case³. Also, nouns and pronouns tend to be used in different referential contexts, dependent on the assumed cognitive state of the comprehender (Garrod & Sanford, 1982; Gundel, Hedberg, & Zacharski, 1993). Specifically, nouns are more often used to initiate reference to a new concept, or to refer back to a concept that has been “backgrounded” or is inactive in memory. In contrast, pronouns tend to maintain reference to given information that is highly salient and currently held under consideration (for example, the topic or subject of the preceding clause). This has been formalized and further specified in terms of a Givenness Hierarchy (Gundel et al., 1993), with null and unstressed pronouns at one end referring only to concepts that are “in focus,” and nouns being used over a range of givenness levels, dependent in part on whether they are paired with a demonstrative, definite or indefinite determiner.

The degree of givenness of a referring expression has been shown to have processing consequences using multiple measures, including reading times (Ehrlich & Rayner, 1983; Warren & Gibson, 2002), accuracy in acceptability judgments (Foraker & McElree, 2007), and the distribution of regressive eye movements to potential antecedents (Carpenter & Just, 1977; cf. Ehrlich & Rayner, 1983). For example, reading times and the perceived complexity of nested syntactic structures steadily decrease with increased givenness of the embedded referring expression (Warren & Gibson, 2002; see also Biber, 1992 for a corpus analysis reaffirming the association between use of pronouns and reduced sentential complexity). This has prompted the consideration of referential type in the calculation of processing cost under Dependency Locality Theory (abbreviated as DLT; Gibson, 2000). There is thus

³ One way to formalize this distinction is to consider the class of concepts that are likely to be referred to using a given word. Pronouns are used to refer to a broad class of concepts, constrained only by a few key semantic features, while nouns often have a much more limited set of concepts to which they regularly refer. Under this framework, to the extent that there are (soft) constraints on what a word is statistically likely to refer to, that word has ties to a richer semantic representation, *ceteris paribus*. Of course, other factors, including word frequency and concreteness, may also interact with such referential statistics to impact the semantic richness of a word-level representation.

converging evidence that less taxing processing occurs in response to pronouns referring to salient entities, when compared to noun phrases introducing new entities.

Moreover, several theories suggest that resolving pronouns primarily entails accessing structures that are already active within working memory and does not elicit the same semantic memory operations as noun-phrase resolution (e.g., Garrod & Sanford, 1982; Garrod et al., 1994; Gundel et al., 1993; McKoon, Gerrig, & Greene, 1996; Myers & O'Brien, 1998). Under one theory, Garrod and Sanford (1982) propose two partitions of working memory: "explicit focus," an active representation of the preceding text, and "implicit focus," the accessible portion of semantic memory used for situational inferencing. They then note that definite noun phrases -- but not pronouns -- are commonly used to access information held in "implicit focus." For example, in the sentence below, the noun "clothes" must be introduced (a) and cannot be substituted with a pronoun like "they" (b):

Mary dressed the baby.

(a) The clothes were made of pink wool.

(b) *They were made of pink wool.

Garrod and Sanford (1982) then go on to claim that pronouns access the initially formed representation of the antecedent, while nouns may build upon and change the antecedent's semantic representation. They summarize:

"The procedural characteristic of pronouns and FDNPs [Full Definite Noun Phrases] may therefore be distinguished in two main ways (1) in terms of restriction of search domain in memory and (2) in terms of the additional construction procedures associated with interpretation of the FDNP, which are not associated with the pronoun." (p. 31)

That is, pronouns are claimed not to access the same partition of memory as nouns, and nouns are claimed to trigger additional operations within semantic memory. Notably, all of this is hypothesized to occur within working memory, without explication of how additional semantic information is made accessible in the course of sentence processing.

Our empirical findings can be reframed in more theoretical terms using an extension of this model, which was subsequently refined by Garrod and colleagues (Garrod et al., 1994). Building off of the N400 literature (recently reviewed in Kutas & Federmeier, 2011), we claim that nouns, to the extent that they convey new information, automatically trigger activation of long-term semantic memory. Thus, the contents of what Garrod and colleagues would call “implicit memory” are constantly being updated, including by the referring expressions themselves. However, given our experimental findings, pronouns are not detectably eliciting this same updating process. This constitutes the first neural evidence corroborating Garrod and Sanford's (1982) proposal that pronouns refer to antecedents mentioned in the preceding text without altering the semantics of the antecedent, while noun anaphors may convey new information about the antecedent in addition to referencing it. Such an account is also consistent with the broader claim that concepts in focus, which tend to be referred to using pronouns, are more semantically specified (Sturt, Sanford, Stewart, & Dawydiak, 2004; Sanford, Sanford, Molle, & Emmott, 2006).

The above theoretical view is in contrast to Lewis and Vasishth's (2005) ACT-R based model, which predicts qualitatively similar memory access procedures across referential types. The Lewis and Vasishth (2005) model of language comprehension (hereafter L&V) does not explicitly code for referring type, in contrast with DLT (Gibson, 2000). Also, in keeping with the ACT-R framework, L&V only differentiates between working and long-term memory in terms of activation level: i.e., working memory is represented as the active portion of long-term memory. This implies that under L&V, and ACT-R more generally, the process of bringing information stored in long-term memory into working memory, and relative change in activation level within working memory, are qualitatively the same phenomenon. Finally, Vasishth and colleagues have used independent evidence to suggest that the activation level of an antecedent is boosted by the presence of an overt pronoun, interpreting their results under the ACT-R framework. Keshtiari and Vasishth (2013) reported that reading times were faster in the vicinity of a downstream verb after participants read clauses containing an overt pronoun that referred to one of its arguments, as opposed to a null pronoun. They interpreted this as suggesting that the overt pronoun had boosted activation of the antecedent in memory. Taken together, Vasishth and colleagues' model of reference predicts that reading a pronoun should result in a pattern of semantic

access, including modulation of the N400, that is qualitatively similar to that at a preceding noun. Our results disconfirm this prediction, suggesting that either a qualitative distinction based on referential type, or between working and long-term memory, is required for a model of language comprehension to fully capture the extant evidence on referential processing.

Of course, in making the above inferences, we are assuming that the pronoun was successfully resolved to the intended noun antecedent in our experiment. If our participants failed to attend to and process the pronouns as they would under ordinary reading conditions, for example, a null result could be obtained that would not generalize well outside of the experimental context. However, this appears unlikely, given that we have online and offline measures demonstrating that our participants were engaged in semantically processing the text and encoding some portions of it into memory. Online, participants showed concreteness effects at the noun, as well as semantic similarity priming from the noun onto words downstream. In addition, we observed a general decrease in N400 amplitude across sentence positions, which has been interpreted as reflecting progressive integration of the text into a coherent discourse model (Van Petten & Kutas, 1990; Payne et al., 2015). Subsequently, in an offline memory task, participants were able to successfully discriminate between nouns they had read before in the critical sentences and a set of novel nouns. Thus, participants appear to have engaged with the text sufficiently deeply to show signs of semantic encoding of the antecedent and semantic integration of the subsequent text.

Nonetheless, our key findings in isolation are consistent with the strong version of a “good enough” processing account of reference (Ferreira, Bailey, & Ferraro, 2002; Sanford & Graesser, 2006). Under this hypothetical account, the pronoun is not resolved at all even when participants are paying attention to the task. That is, the antecedent is never assigned to the role filled by the pronoun online. Rather, a referential assignment is made in retrospect only if participants are directly probed. Indeed, there is evidence that pronouns may remain underspecified with respect to their antecedent, in cases where further specification is unnecessary for adequate comprehension (Poesio, Sturt, Artstein, & Filik, 2006). Although this account predicts the same pattern of results as those that we obtained, we disprefer it for several reasons. To start, pronominal reference to a topicalized antecedent

in an immediately preceding clause is a relatively straightforward case (e.g., Arnold, 1998, pp. 26-29). The fact that participants in our norming task interpreted the pronoun as referring to the same antecedent highly consistently suggests that there was not much ambiguity and that this was not a difficult task for adults. If 3rd person pronouns such as “it” and “they” are not resolved online in this case, it would imply that they are not resolved online in a much wider array of cases. However, that idea appears inconsistent with the broader literature, showing, for example, regressive eye-movements to the antecedent after reading a pronoun when verifying the consistency of a sentence with its preceding context (Carpenter & Just, 1977), and effects of distance to the antecedent on reading times (Ehrlich & Rayner, 1983; cf. Blanchard, 1987). It also seems to conflict with other evidence for rapid online pronoun processing described in the introduction (e.g., Arnold et al., 2000; Järvikivi et al., 2005; reviews in Garnham, 2001; Callahan, 2008; Nieuwland & Van Berkum, 2008).

We thus interpret our results as suggesting that pronouns are indeed resolved online, but not through reactivation of information in long-term semantic memory. It may be that the antecedent is sufficiently active at the time the pronoun is encountered that anaphoric resolution requires little or no change to the current representation. This could be because the antecedent remains in focus/active from the time it is encountered or that, via predictive processing mechanisms, it becomes reactivated in anticipation of the pronoun. Words in predictable contexts are known to show an attenuated N400 because the semantics associated with them have been “preactivated” by the facilitating context (e.g., DeLong et al., 2005; Federmeier & Kutas, 1999). Sentence-final nouns in congruent, predictive contexts have been found to show a reduced or absent concreteness effect (Holcomb et al., 1999), as observed for the pronoun in our data. Pronouns are typically used to refer back to highly salient antecedents (Gundel et al., 1993), and highly salient antecedents have a high probability of being referred back to in the following clause (Arnold, 1998). Thus, the probability of encountering a pronoun that refers back to a topicalized entity in the preceding clause is high in natural language use (see also Ariel, 1990, pp. 18-19, for estimates, based on a small body of text, that pronouns are used roughly 90% of the time in English to refer to an antecedent within the same sentence).

Indeed, McKoon, Gerrig, and Greene (1996, p. 920) have argued that “a pronoun does not create accessibility for itself – it confirms accessibility that already exists,” and several

authors have emphasized that pronouns need not be a cue for the reader to recompute an antecedent, but rather may confirm that a default antecedent, often the topic, is what will fill a particular role (e.g., Gernsbacher, 1989; Sanford & Garrod, 1989; Greene et al., 1992). Such accounts suggest that when a single highly active entity is present, pronouns are first checked for match against that entity. Only if there is a failure to match will pronouns engage additional processing, which may include semantic reactivation of the ultimately selected antecedent if it had been backgrounded. This hypothesis is appealing because it reflects that, in many cases, unstressed overt pronouns can be omitted (i.e. replaced by null pronouns), both in English as well as cross-linguistically (Gundel et al., 1993). In Japanese, a heavily pro-drop language, null pronouns are always used to refer to antecedents that are “in focus,” since the referent is easily identifiable from context (Gundel et al., 1993). Even within our experimental sentences, in many cases the critical pronoun could have been omitted with no loss of meaning. For example:

- a) “A **loan** is normally helpful, but sometimes (**it**) can have a negative impact.”
- b) “That **crown** was not important to the king, but (**it**) was valued highly by others.”

Note that to get the continuation of sentence b) to mean that the king was highly valued, you would need to include an overt 3rd person pronoun “he” as the subject of the second clause, because the king is less salient than the crown. However, because the crown is topicalized, no overt pronoun “it” is necessary for the crown to be the inferred subject of the second clause in English.

Under some contexts, then, unstressed pronouns may function similarly to inflection markers, which often convey redundant information hypothesized to reinforce the linguistic signal under noisy conditions (Aylett & Turk, 2004; Bybee, 1985, p. 202-203). Indeed, clitic pronouns and pronominal affixes are common in many fusional (e.g. Spanish), polysynthetic (e.g. Navajo), and agglutinative (e.g. Yup’ik Eskimo) languages (Bhat, 2004; Mithun, 2003; Pineda & Meza, 2005). These inflection-like morphemes function similarly to the independent pronouns found in more analytic languages such as English, Chinese and German (Mithun, 2003; see Bhat, 2004 for discussion). Of course, it should be noted that in so-called “bound-pronoun” languages, the bound pronominal markers are more obligatory

than person-marking affixes in “free-pronoun” languages (Bhat, 2004). This suggests that pronouns in general, however they are realized morphologically, tend to convey more critical information than purely redundant agreement markers. Nonetheless, there are cases, as with many of the experimental stimuli in the current study, in which pronouns carry essentially redundant information and may lead to little additional cognitive processing.

For cases wherein there is evidence that some kind of additional processing is elicited when a pronoun is encountered, our ERP data are important in speaking to the kinds of mechanisms that are likely to be involved. For example, our data suggest that the kind of “reactivation” observed in prior behavioral and eye-tracking work (e.g., Nicol & Swinney, 2003; Lago, 2014, expt 6) does not involve long-term memory activation, but rather reflects operations carried out on the contents of working memory, perhaps as mediated by attentional shifts. This hypothesis is also consistent with the fact that ambiguous pronouns generate a frontal negativity (the Nref effect) that is sensitive to working memory capacity (Nieuwland & Van Berkum, 2006), and also more generally similar to slow-wave potentials linked to verbal working memory load (King & Kutas, 1995; but see Nieuwland, Petersson, & Van Berkum, 2007 for evidence that these effects may have distinct neural sources). Recently, Brodbeck, Gwilliams, and Pykkänen (2015) provided direct neural evidence that referential processing may induce an attentional shift to the antecedent. Using a visual field manipulation, they tracked the time course of reference resolution during sentence processing. They found a posterior negativity contralateral to where a referent had previously been visually displayed at the disambiguation point (either a noun or adjective) of sentences presented one word at a time. Similar contralateral effects (e.g. the N2pc and CDA) have been linked to attentional allocation during visual working memory tasks (reviewed in Luck, 2012; Perez & Vogel, 2012). Thus, converging evidence points to a critical role for attention (i.e., executive processes used in the manipulation of information held in working memory) in reference resolution.

In summary, the current study has highlighted an important memory-related processing distinction between nouns and pronouns when encountered in a typical linguistic context. We have shown no evidence for re-activation of the referent’s semantics in long-term memory shortly after reading a pronoun. Our results are thus consistent with the idea that pronouns, unlike nouns, do not induce the reader to update the contents of working

memory with (very much) new semantic information, and thus that pronouns have limited potential to modify the semantics of their antecedents. Instead, our data suggest that previous behavioral and eye-tracking evidence for the “reactivation” of antecedent semantics may reflect attentionally-mediated processes that operate within working memory.

Table and Figures

Table 1: Concreteness, Word Frequency, Number of Letters, Orthographic Neighborhood Size, and Familiarity, by Condition

Feature	Condition	Mean (Standard Error)
Concreteness	Concrete	576 (3)
	Abstract	373 (6)
Kucera-Francis Frequency	Concrete	61.8 (7.9)
	Abstract	61.8 (7.8)
Number of Letters	Concrete	5.93 (.15)
	Abstract	5.98 (.16)
Orthographic Neighborhood Size (OLD20)	Concrete	2.14 (.06)
	Abstract	2.15 (.06)
Familiarity	Concrete	526 (6)
	Abstract	512 (6)

Figure 1: Concreteness contrast at the critical noun. Depicted at 3 frontal (top), 3 central (middle) and 3 parietal (bottom) sites. Time 0 indicates onset of the critical noun. A low-pass filter with a 15 Hz half-amplitude cut-off was applied for illustration only.

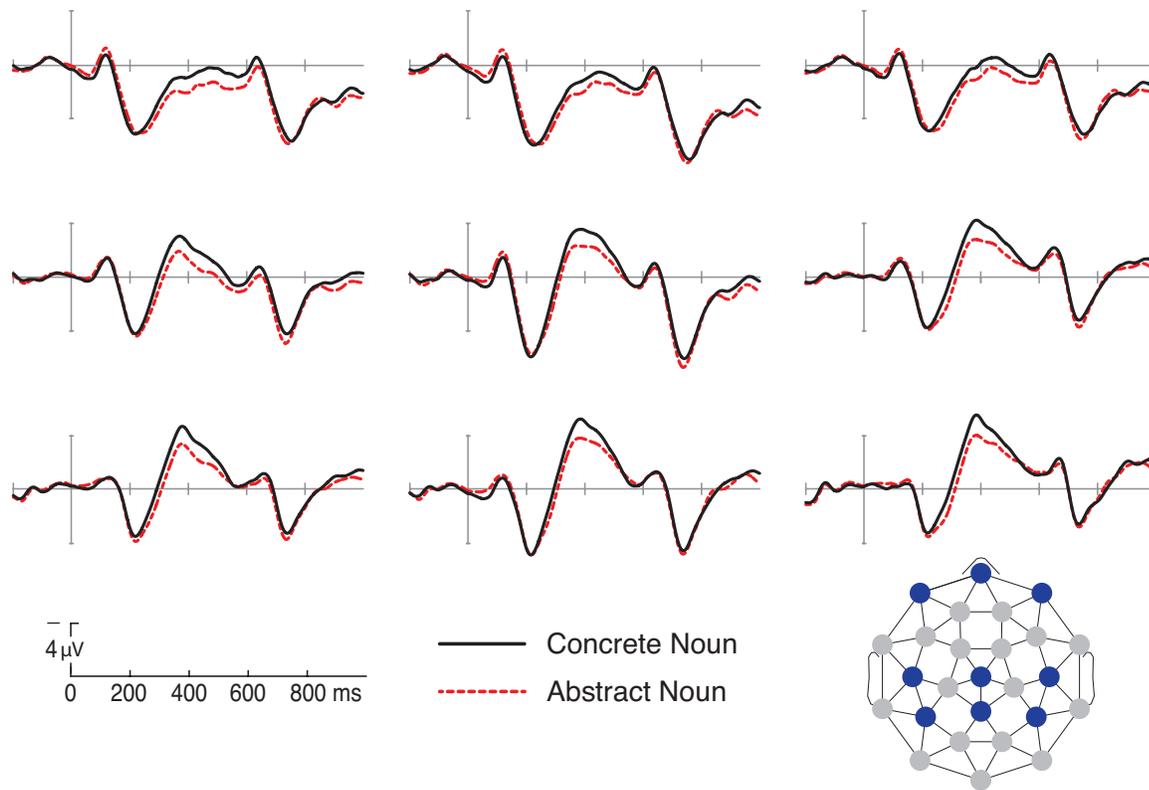


Figure 2: Concreteness contrast at the critical pronoun. Depicted at 3 frontal (top), 3 central (middle) and 3 parietal (bottom) sites. Time 0 indicates onset of the critical pronoun. A low-pass filter with a 15 Hz half-amplitude cut-off was applied for illustration only.

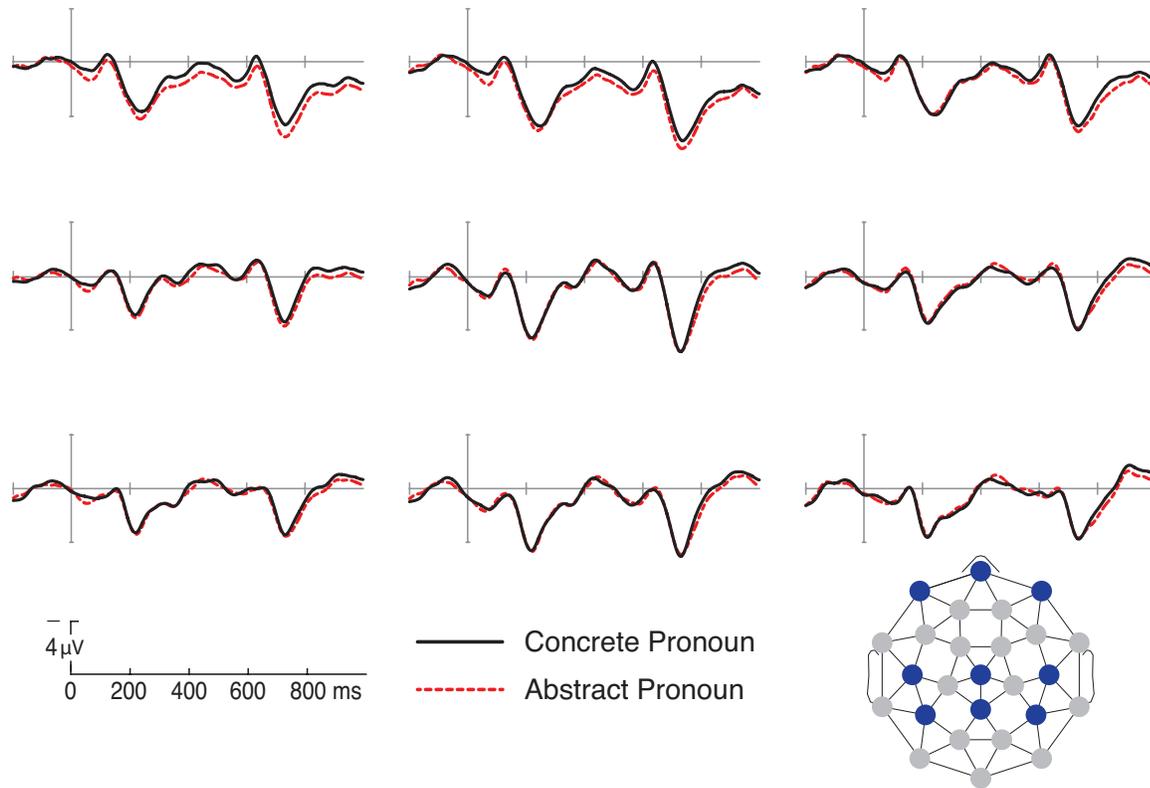


Figure 3. Effect of semantic association between content words and critical noun on N400 amplitude varies as a function of linear distance to the critical noun. From left to right, bins spanning three consecutive word positions are plotted separately; word positions 1-3 = the word immediately following the critical noun through the word 3 positions after the critical noun, etc. a. A regression line (semantic similarity vs. N400 amplitude) based on a mixed effects model fit to all word position bins is plotted separately for each word position bin. LSA semantic similarity score is plotted on the x axis, and mean voltage in the N400 time window over 8 centroparietal sites is plotted on the y axis. 95% bootstrapped confidence intervals (N = 2000 simulations) are shaded around each regression line. Positive slopes indicate semantic priming by the critical noun. Points represent mean N400 amplitude aggregated across participants for each content word that was preceded by a particular critical noun (note: model was not directly fit to points, which are for illustration only). b. Model estimated effect sizes for the effect of semantic association strength on N400 amplitude, grouped by word position bin, with 95% bootstrapped confidence intervals (N = 2000 simulations). Sentence-final words were excluded from the model.

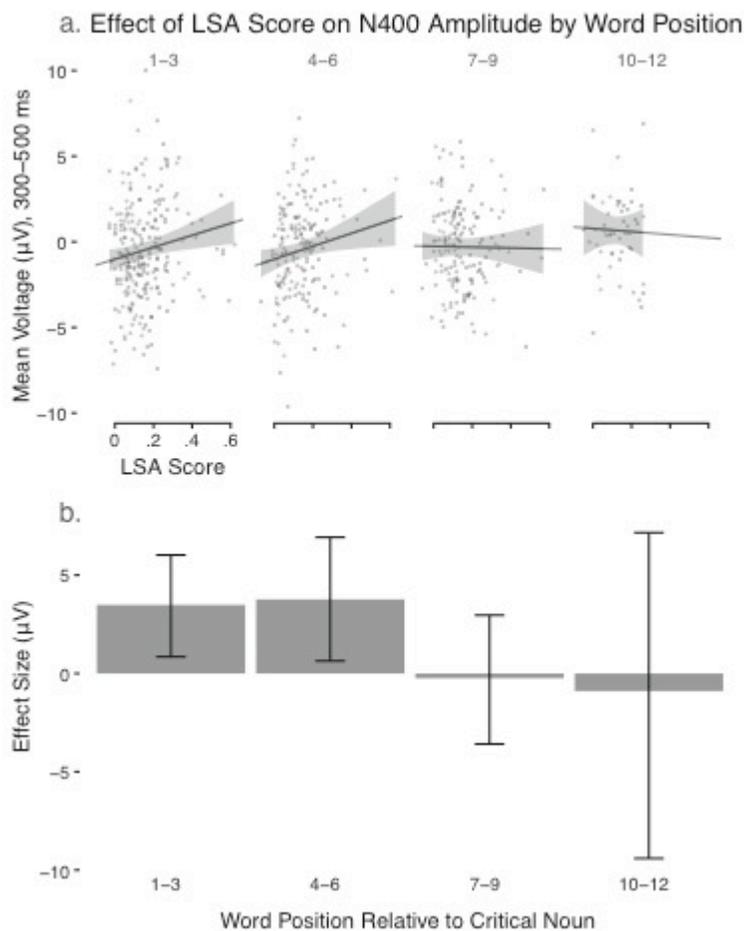
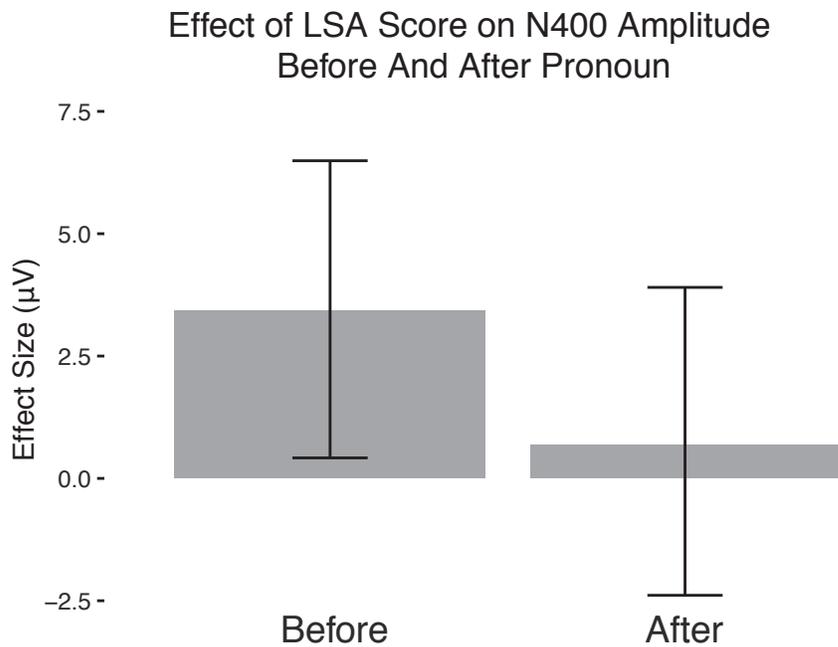


Figure 4. Effect of semantic association between content words and critical noun as a function of linear distance to the critical pronoun. Estimated effect sizes with 95% bootstrapped confidence intervals (N = 2000 simulations) are plotted for content words occurring within three linear word positions before or after the pronoun. For simplicity, model used for illustration did not include distance between the pronoun and the critical noun as a predictor; however, this did not change fixed effects estimates by more than 1/8 of a standard error relative to the model in the text.



References

- Ariel, M. (1990). *Accessing noun-phrase antecedents*. London: Routledge.
- Arnold, J. E. (1998). *Reference form and discourse patterns*. Stanford University.
- Arnold, J. E., Eisenband, J. G., Brown-Schmidt, S., & Trueswell, J. C. (2000). The rapid use of gender information: Evidence of the time course of pronoun resolution from eyetracking. *Cognition*, 76(1), B13–B26.
- Aylett, M., & Turk, A. (2004). The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. *Language and Speech*, 47(1), 31–56.
- Barber, H. A., Otten, L. J., Kousta, S.-T., & Vigliocco, G. (2013). Concreteness in word processing: ERP and behavioral effects in a lexical decision task. *Brain and Language*, 125(1), 47–53.
- Bever, T. G., & McElree, B. (1988). Empty categories access their antecedents during comprehension. *Linguistic Inquiry*, 19(1), 35–43.
- Bhat, D. S. (2004). In *Pronouns* (pp. 15–30). Oxford, UK: Oxford University Press.
- Biber, D. (1992). On the complexity of discourse complexity: A multidimensional analysis. *Discourse Processes*, 15(2), 133–163.
- Binder, J. R., Desai, R. H., Graves, W. W., & Conant, L. L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, 19(12), 2767–2796.
- Blanchard, H. E. (1987). The effects of pronoun processing on information utilization during fixations in reading. *Bulletin of the Psychonomic Society*, 25(3), 171–174.
- Brodbeck, C., Gwilliams, L., & Pylkkänen, L. (2015). EEG can track the time course of successful reference resolution in small visual worlds. *Frontiers in Psychology*, 6, 1787.
- Bybee, J. L. (1985). *Morphology: A study of the relation between meaning and form* (Vol. 9). Philadelphia: John Benjamins Publishing.
- Callahan, S. M. (2008). Processing anaphoric constructions: Insights from electrophysiological studies. *Journal of Neurolinguistics*, 21(3), 231–266.
- Carpenter, P. A., & Just, M. A. (1977). Reading comprehension as eyes see it. *Cognitive Processes in Comprehension*, 109–139.
- Carreiras, M., Garnham, A., Oakhill, J., & Cain, K. (1996). The use of stereotypical gender information in constructing a mental model: Evidence from English and Spanish. *The Quarterly Journal of Experimental Psychology: Section A*, 49(3), 639–663.
- Chang, F. R. (1980). Active memory processes in visual sentence comprehension: Clause effects and pronominal reference. *Memory & Cognition*, 8(1), 58–64.
- Chow, W.-Y., Lewis, S., & Phillips, C. (2014). Immediate sensitivity to structural constraints in pronoun resolution. *Frontiers in Psychology*, 5, 630.
- Cloitre, M., & Bever, T. G. (1988). Linguistic anaphors, levels of representation, and discourse. *Language and Cognitive Processes*, 3(4), 293–322.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407–428.
- Corbett, A. T., & Chang, F. R. (1983). Pronoun disambiguation: Accessing potential antecedents. *Memory & Cognition*, 11(3), 283–294.
- Cowan, N. (1993). Activation, attention, and short-term memory. *Memory & Cognition*, 21(2), 162–167.

- DeLong, K. A., Urbach, T. P., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience*, 8(8), 1117–1121.
- Dhond, R. P., Witzel, T., Dale, A. M., & Halgren, E. (2007). Spatiotemporal cortical dynamics underlying abstract and concrete word reading. *Human Brain Mapping*, 28(4), 355–362.
- Ehrlich, K., & Rayner, K. (1983). Pronoun assignment and semantic integration during reading: Eye movements and immediacy of processing. *Journal of Verbal Learning and Verbal Behavior*, 22(1), 75–87.
- Emmorey, K., & Lillo-Martin, D. (1995). Processing spatial anaphora: Referent reactivation with overt and null pronouns in American Sign Language. *Language and Cognitive Processes*, 10(6), 631–653.
- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102(2), 211.
- Farah, M. J., Weisberg, L. L., Monheit, M., & Peronnet, F. (1989). Brain activity underlying mental imagery: Event-related potentials during mental image generation. *Journal of Cognitive Neuroscience*, 1(4), 302–316.
- Federmeier, K. D., & Kutas, M. (1999). A rose by any other name: Long-term memory structure and sentence processing. *Journal of Memory and Language*, 41(4), 469–495.
- Federmeier, K. D., & Laszlo, S. (2009). Time for meaning: Electrophysiology provides insights into the dynamics of representation and processing in semantic memory. *Psychology of Learning and Motivation*, 51, 1–44.
- Federmeier, K. D., Van Petten, C., Schwartz, T. J., & Kutas, M. (2003). Sounds, words, sentences: age-related changes across levels of language processing. *Psychology and Aging*, 18(4), 858–872.
- Ferreira, F., Bailey, K. G., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, 11(1), 11–15.
- Foraker, S., & McElree, B. (2007). The role of prominence in pronoun resolution: Active versus passive representations. *Journal of Memory and Language*, 56(3), 357–383.
- Garnham, A. (2001). In *Mental models and the interpretation of anaphora* (pp. 67–71, 139–143). Philadelphia: Psychology Press.
- Garnham, A., Oakhill, J., Ehrlich, M.-F., & Carreiras, M. (1995). Representations and processes in the interpretation of pronouns: New evidence from Spanish and French. *Journal of Memory and Language*, 34(1), 41.
- Garrod, S. C., & Sanford, A. J. (1982). The mental representation of discourse in a focussed memory system: Implications for the interpretation of anaphoric noun phrases. *Journal of Semantics*, 1(1), 21–41.
- Garrod, S., Freudenthal, D., & Boyle, E. (1994). The role of different types of anaphor in the on-line resolution of sentences in a discourse. *Journal of Memory and Language*, 33(1), 39.
- Gernsbacher, M. A. (1989). Mechanisms that improve referential access. *Cognition*, 32(2), 99–156.
- Gernsbacher, M. A., Hargreaves, D. J., & Beeman, M. (1989). Building and accessing clausal representations: The advantage of first mention versus the advantage of clause recency. *Journal of Memory and Language*, 28(6), 735–755.

- Gibson, E. (2000). The dependency locality theory: A distance-based theory of linguistic complexity. *Image, Language, Brain*, 95–126.
- Greene, S. B., McKoon, G., & Ratcliff, R. (1992). Pronoun resolution and discourse models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(2), 266–283.
- Gullick, M. M., Mitra, P., & Coch, D. (2013). Imagining the truth and the moon: an electrophysiological study of abstract and concrete word processing. *Psychophysiology*, 50(5), 431–440.
- Gundel, J. K., Hedberg, N., & Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language*, 69(2), 274–307.
- Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d' . *Behavior Research Methods, Instruments, & Computers*, 27(1), 46–51.
- Holcomb, P. J., Kounios, J., Anderson, J. E., & West, W. C. (1999). Dual-coding, context-availability, and concreteness effects in sentence comprehension: an electrophysiological investigation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(3), 721–742.
- Huang, H.-W., Lee, C.-L., & Federmeier, K. D. (2010). Imagine that! ERPs provide evidence for distinct hemispheric contributions to the processing of concrete and abstract concepts. *NeuroImage*, 49(1), 1116–1123.
- Järvikivi, J., van Gompel, R. P., Hyönä, J., & Bertram, R. (2005). Ambiguous pronoun resolution contrasting the first-mention and subject-preference accounts. *Psychological Science*, 16(4), 260–264.
- Kennison, S. M., & Trofe, J. L. (2003). Comprehending pronouns: A role for word-specific gender stereotype information. *Journal of Psycholinguistic Research*, 32(3), 355–378.
- Keshtiari, N., & Vasishth, S. (2013). Reactivation of antecedents by overt versus null pronouns: Evidence from Persian. *Journal of Language Modelling*, 1(2), 243–266.
- King, J. W., & Kutas, M. (1995). Who did what and when? Using word-and clause-level ERPs to monitor working memory usage in reading. *Journal of Cognitive Neuroscience*, 7(3), 376–395.
- Kounios, J., & Holcomb, P. J. (1994). Concreteness effects in semantic processing: ERP evidence supporting dual-coding theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(4), 804–823.
- Kousta, S.-T., Vigliocco, G., Vinson, D. P., Andrews, M., & Del Campo, E. (2011). The representation of abstract words: why emotion matters. *Journal of Experimental Psychology: General*, 140(1), 14–34.
- Kreiner, H., Garrod, S., & Sturt, P. (2013). Number agreement in sentence comprehension: The relationship between grammatical and conceptual factors. *Language and Cognitive Processes*, 28(6), 829–874.
- Kutas, M., & Federmeier, K. D. (2011). Thirty years and counting: Finding meaning in the N400 component of the event related brain potential (ERP). *Annual Review of Psychology*, 62, 621–647.
- Lago, S. (2014). *Memory and Prediction in Cross-Linguistic Sentence Processing* (Doctoral dissertation). University of Maryland.

- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*(2), 211–240.
- Landauer, T. K., Foltz, P. W., & Laham, D. (1998). An introduction to latent semantic analysis. *Discourse Processes*, *25*, 259–284.
- Laszlo, S., & Federmeier, K. D. (2011). The N400 as a snapshot of interactive processing: Evidence from regression analyses of orthographic neighbor and lexical associate effects. *Psychophysiology*, *48*(2), 176–186.
- Lee, C.-L., & Federmeier, K. D. (2008). To watch, to see, and to differ: An event-related potential study of concreteness effects as a function of word class and lexical ambiguity. *Brain and Language*, *104*(2), 145–158.
- Leiman, J. M. (1982). *A chronometric analysis of referent assignment to pronouns* (Doctoral dissertation). Wayne State University.
- Levy, E. T., & Fowler, C. A. (2000). The role of gestures and other graded language forms in the grounding of reference. In *Language and gesture* (Vol. 2, pp. 215–234). Cambridge: Cambridge University Press.
- Lewis, R. L., & Vasishth, S. (2005). An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, *29*(3), 375–419.
- Luck, S. J. (2012). Electrophysiological correlates of the focusing of attention within complex visual scenes: N2pc and related ERP components. In S. J. Luck & E. S. Kappenman (Eds.), *The Oxford handbook of event-related potential components* (pp. 329–360). New York: Oxford University Press.
- MacDonald, M. C., & MacWhinney, B. (1990). Measuring inhibition and facilitation from pronouns. *Journal of Memory and Language*, *29*(4), 469–492.
- McElree, B., Foraker, S., & Dyer, L. (2003). Memory structures that subserve sentence comprehension. *Journal of Memory and Language*, *48*(1), 67–91.
- McKoon, G., Gerrig, R. J., & Greene, S. B. (1996). Pronoun resolution without pronouns: some consequences of memory-based text processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(4), 919–932.
- McKoon, G., Ratcliff, R., & Ward, G. (1994). Testing theories of language processing: An empirical investigation of the on-line lexical decision task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(5), 1219–1228.
- Mithun, M. (2003). Pronouns and agreement: the information status of pronominal affixes. *Transactions of the Philological Society*, *101*(2), 235–278.
- 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*(3), 226–254.
- Nicol, J. L. (1988). *Coreference processing during sentence comprehension*. Massachusetts Institute of Technology.
- Nicol, J. L., & Swinney, D. (2003). The psycholinguistics of anaphora. In A. Barss (Ed.), *Anaphora: A reference guide* (pp. 72–104). Oxford, UK: Blackwell.
- Nieuwland, M. S. (2014). "Who's he?" Event-related brain potentials and unbound pronouns. *Journal of Memory and Language*, *76*, 1–28.
- Nieuwland, M. S., Otten, M., & Van Berkum, J. J. (2007). Who are you talking about? Tracking discourse-level referential processing with event-related brain potentials. *Journal of Cognitive Neuroscience*, *19*(2), 228–236.

- Nieuwland, M. S., Petersson, K. M., & Van Berkum, J. J. (2007). On sense and reference: Examining the functional neuroanatomy of referential processing. *NeuroImage*, 37(3), 993–1004.
- Nieuwland, M. S., & Van Berkum, J. J. (2006). Individual differences and contextual bias in pronoun resolution: Evidence from ERPs. *Brain Research*, 1118(1), 155–167.
- Nieuwland, M. S., & Van Berkum, J. J. (2008). The neurocognition of referential ambiguity in language comprehension. *Language and Linguistics Compass*, 2(4), 603–630.
- Nittono, H., Suehiro, M., & Hori, T. (2002). Word imageability and N400 in an incidental memory paradigm. *International Journal of Psychophysiology*, 44(3), 219–229.
- Oldfield, R. C. (1971). The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia*, 9(1), 97–113.
- Paivio, A. (1986). *Mental representations: A dual coding approach* Oxford University Press. New York, NY: Oxford University Press.
- Payne, B. R., Lee, C.-L., & Federmeier, K. D. (2015). Revisiting the incremental effects of context on word processing: Evidence from single-word event-related brain potentials. *Psychophysiology*, 52(11), 1456–1469.
- Perez, V. B., & Vogel, E. K. (2012). What ERPs can tell us about working memory. In S. J. Luck & E. S. Kappenman (Eds.), *Oxford handbook of event-related potential components* (pp. 361–372). New York: Oxford University Press.
- Pineda, L., & Meza, I. (2005). The Spanish pronominal clitic system. *Procesamiento Del Lenguaje Natural*, 34, 67–103.
- Poesio, M., Sturt, P., Artstein, R., & Filik, R. (2006). Underspecification and anaphora: Theoretical issues and preliminary evidence. *Discourse Processes*, 42(2), 157–175.
- Sanford, A. J., & Garrod, S. C. (1989). What, when, and how?: Questions of immediacy in anaphoric reference resolution. *Language and Cognitive Processes*, 4, 235–262.
- Sanford, A. J., & Graesser, A. C. (2006). Shallow processing and underspecification. *Discourse Processes*, 42(2), 99–108.
- Sanford, A. J., Sanford, A. J., Molle, J., & Emmott, C. (2006). Shallow processing and attention capture in written and spoken discourse. *Discourse Processes*, 42(2), 109–130.
- Schwanenflugel, P. J. (1991). Why are abstract concepts hard to understand? In P. J. Schwanenflugel (Ed.), *The psychology of word meaning* (pp. 223–250). Mahwah, NJ: Erlbaum.
- Schwanenflugel, P. J., Harnishfeger, K. K., & Stowe, R. W. (1988). Context availability and lexical decisions for abstract and concrete words. *Journal of Memory and Language*, 27(5), 499–520.
- Shillcock, R. (1982). The on-line resolution of pronominal anaphora. *Language and Speech*, 25(4), 385–401.
- Speelman, C. P., & Kirsner, K. (1990). The representation of text-based and situation-based information in discourse comprehension. *Journal of Memory and Language*, 29(1), 119–132.
- Spreen, O., & Schulz, R. W. (1966). Parameters of abstraction, meaningfulness, and pronunciability for 329 nouns. *Journal of Verbal Learning and Verbal Behavior*, 5(5), 459–468.
- Sturt, P., Sanford, A. J., Stewart, A., & Dawydiak, E. (2004). Linguistic focus and good-enough representations: An application of the change-detection paradigm. *Psychonomic Bulletin & Review*, 11(5), 882–888.

- Tyler, L. K., & Marslen-Wilson, W. (1982). Processing utterances in discourse contexts: On-line resolution of anaphors. *Journal of Semantics*, 1, 297–314.
- Van Petten, C. (1993). A comparison of lexical and sentence-level context effects in event-related potentials. *Language and Cognitive Processes*, 8(4), 485–531.
- Van Petten, C. (2014). Examining the N400 semantic context effect item-by-item: Relationship to corpus-based measures of word co-occurrence. *International Journal of Psychophysiology*, 94(3), 407–419.
- Van Petten, C., & Kutas, M. (1990). Interactions between sentence context and word frequency in event-related brain potentials. *Memory & Cognition*, 18(4), 380–393.
- Van Petten, C., Weckerly, J., McIsaac, H. K., & Kutas, M. (1997). Working memory capacity dissociates lexical and sentential context effects. *Psychological Science*, 8(3), 238–242.
- Warren, T., & Gibson, E. (2002). The influence of referential processing on sentence complexity. *Cognition*, 85(1), 79–112.
- Welcome, S. E., Paivio, A., McRae, K., & Joanisse, M. F. (2011). An electrophysiological study of task demands on concreteness effects: evidence for dual coding theory. *Experimental Brain Research*, 212(3), 347–358.
- West, W. C., & Holcomb, P. J. (2000). Imaginal, semantic, and surface-level processing of concrete and abstract words: an electrophysiological investigation. *Journal of Cognitive Neuroscience*, 12(6), 1024–1037.
- Yarkoni, T., Balota, D., & Yap, M. (2008). Moving beyond Coltheart's N: A new measure of orthographic similarity. *Psychonomic Bulletin & Review*, 15(5), 971–979.
- Zhang, Q., Guo, C., Ding, J., & Wang, Z. (2006). Concreteness effects in the processing of Chinese words. *Brain and Language*, 96(1), 59–68.

Appendix A: Complete List of Experimental Stimuli

The first bolded word of each sentence denotes the concrete critical noun, the second denotes the abstract critical noun, and the third denotes the critical pronoun (“it” or “they”). Content words (CW) are listed below each stimulus, and those included in the lexical association analysis are underlined. * indicates all content words for that item were excluded from the lexical association analysis due to lack of lexical association scores with the critical nouns.

1. The **lunch** / **quantity** was extremely large, and Harry thought **it** was too much for him.
CW: extremely, large, Harry, thought
2. His **brain** / **ego** was badly injured, to the point that **it** might never fully recover.
CW: badly, injured, fully, recover
3. That **newspaper** / **riot** changed society because **it** was so extreme.
CW: changed, society, extreme
4. The **magazine** / **edition** hadn't come out yet, but **it** was going to soon.
CW: come, soon
5. The **supper** / **mixture** was not only delicious, by her standards **it** was healthy, too.
CW: delicious, standards, healthy
6. The **money** / **grade** was missing, and **it** would be hard to make up.
CW: missing, hard, make
7. A **medicine** / **loan** is normally helpful, but sometimes **it** can have a negative impact.
CW: normally, helpful, negative, impact
8. That **music** / **style** was not his favorite, because **it** didn't have a very good beat.
CW: favorite, good, beat
9. The **hotel** / **service** was below expectations, because he had thought **it** would be more elegant.
CW: expectations, thought, elegant
10. The **library** / **department** actually closed at seven, but Nicole thought that **it** closed at eight.
CW: closed, seven, Nicole, thought, closed, eight
11. The **picture** / **scale** was too small, and so **it** needed to be adjusted.
CW: small, needed, adjusted
12. That **book** / **trip** didn't cost much even though **it** was so fantastic.
CW: cost, fantastic
13. The **beer** / **joke** did not go over well, since **it** didn't suite (sic) the guests' taste.
CW: go, well, suite, guests', taste
14. His **heart** / **luck** was good up through his sixties, then **it** started to fail.
CW: good, sixties, started, fail
15. Their **coffee** / **flare** for events was critical and **it** helped with party planning.
CW: events, critical, helped, party, planning
16. The **duck** / **prey** was unaware that **it** was being hunted.
CW: unaware, hunted
17. The **chains** / **trends** were hard to make out, because **they** were so entangled.
CW: hard, make, entangled
18. The **block** / **east** was full of crime, and, as a location, **it** was not ideal.

- CW: full, crime, location, ideal
19. The **poster / news** didn't incite much debate considering **it** was so shocking.
CW: incite, debate, considering, shocking
20. The **certificate / team** made her look impressive because **it** was so prestigious.
CW: made, look, impressive, prestigious
21. The **lamb / sum** was very small and **it** was easy to handle.
CW: small, easy, handle
22. Her **shower / fate** could not be altered since **it** was already sealed.
CW: altered, already, sealed
23. Their **banner / cause** was highly controversial, but **it** sure did raise awareness.
CW: highly, controversial, sure, raise, awareness
24. The **disc / hole** was perfectly circular, but **it** wasn't big enough.
CW: perfectly, circular, big
25. That **crown / issue** was not important to the king, but **it** was valued highly by others.
CW: important, king, valued, highly
26. Her **leg / mood** always improved after exercising, and that day **it** was better than usual.
CW: improved, exercising, better, usual
27. Her **mouth / misery** was both very painful and **it** made her look unsightly.
CW: painful, look, unsightly
28. The **thread / message** was very long, and **it** needed to be cut shorter.
CW: long, needed, cut, shorter
29. That **bolt / height** was perfect for the job because **it** met their specifications.
CW: perfect, job, met, specifications
30. Those **universities / donors** are extremely wealthy and **they** love giving scholarships to deserving students.
CW: extremely, wealthy, love, giving, scholarships, deserving, students
31. The **breakfast / plot** wasn't very satisfying because **it** was so bland.
CW: satisfying, bland
32. His **chickens / demons** kept him up at night, and **they** bothered him during the day, too.
CW: kept, night, bothered, day
33. That **church / guilt** not only influenced her greatly, **it** changed her life.
CW: influenced, greatly, changed, life
34. The **factory / birth** caused many problems because **it** created a divide among the townspeople.
CW: caused, problems, created, divide, townspeople
35. The **oil / accident** was hard to clean up, because **it** was so messy.
CW: hard, clean, messy
36. That **badge / fact** was something he treasured, but **it** didn't mean much.
CW: treasured, mean
37. The **letters / rewards** kept coming and **they** just piled up.
CW: kept, coming, piled
38. His **tongue / voice** was a source of embarrassment because **it** was so rough.
CW: source, embarrassment, rough
39. His **painting / advantage** helped him stand out but **it** couldn't win him the competition.

- CW: helped, stand, win, competition
40. The **laundry** / **truth** needed to be hidden because **it** was so dirty.
CW: needed, hidden, dirty
41. A **bath** / **rest** either energizes you or **it** drains you completely.
CW: energizes, drains, completely
42. The **feasts** / **numbers** were enormous and **they** confirmed his impression of the people.
CW: enormous, confirmed, impression, people
43. Her **boxes** / **principles** were strong like **they** were made of steel.
CW: strong, made, steel
44. His **hand** / **shadow** touched the wall and then **it** slid off again.
CW: touched, wall, slid
45. The **floor** / **story** was not very tidy, and **it** needed to be cleaned up.
CW: tidy, needed, cleaned
46. That **school** / **summer** shaped their lives in that **it** taught them many important lessons.
CW: shaped, lives, taught, important, lessons
47. His **farm** / **soul** was finally complete and **it** wasn't for sale.
CW: complete, sale
48. Her **rice** / **talent** grew on its own, as if **it** needed no help.
CW: grew, needed, help
49. Their **cotton** / **skill** brought them great fame, because **it** was unrivaled in the land.
CW: great, fame, unrivaled, land
50. The **path** / **space** wasn't wide enough, although **it** would've been very scenic.
CW: wide, scenic
51. The **plate** / **speck** stood out because **it** was the only thing on the otherwise spotless countertop.
CW: stood, only, thing, otherwise, spotless, countertop
52. Her **belt** / **hint** was outrageous, because **it** was so overdone.
CW: outrageous, overdone
53. Her **head** / **appearance** was unadorned, but **it** was quite stunning.
CW: unadorned, stunning
54. The **spot** / **gang** scared her because **it** appeared so suddenly.
CW: scared, appeared, suddenly
55. The **blood** / **originality** in his art was widely discussed, as **it** was what brought him his fame.
CW: art, widely, discussed, brought, fame
56. That **hospital** / **emergency** was so exasperating because **it** was managed so poorly.
CW: exasperating, managed, poorly
57. Her **nose** / **task** was not at all pleasant, because **it** was just so huge.
CW: pleasant, huge
58. The **stone** / **ignition** was stuck and **it** stayed that way.
CW: stuck, stayed, way
59. The **earth** / **memory** was awe-inspiring as **it** took hold of his imagination.
CW: awe-inspiring, took, hold, imagination
60. That **coat** / **size** didn't fit him, because **it** was too tight.
CW: fit, tight

61. His **foot / shame** was hard to conceal because **it** kept slipping out.
CW: hard, conceal, kept, slipping
62. That **tape / side** wasn't sticky, so **it** started to come unattached.
CW: sticky, started, unattached
63. That **football / subject** was the one he liked best, because **it** reminded him of so many stories.
CW: liked, best, reminded, stories
64. The **iron / relation** between the two structures needed to be reinforced because **it** wasn't strong enough.
CW: structures, needed, reinforced, strong
65. The **oak / death** in the dark was frightening because **it** was reminiscent of the phantom.
CW: dark, frightening, reminiscent, phantom
66. The **hall / trance** was dream-like, because **it** was full of colorful images.
CW: dream-like, colorful, images
67. The **medal / raid** was impressive, although **it** couldn't save the captain's reputation.
CW: impressive, save, captain's, reputation
68. The **accordion / ceremony** made him smile because **it** was so sentimental.
CW: smile, sentimental
69. The **gravy / gloom** was avoided by everyone because **it** looked so uninviting.
CW: avoided, everyone, looked, uninviting
70. That **ball / phrase** bounced around the room like **it** would never stop.
CW: bounced, room, stop
71. The **alligators / aggressors** were ruthless when **they** attacked the villagers.
CW: ruthless, attacked, villagers
72. His **beard / impotency** wasn't so much a disaster as **it** was an embarrassment.
CW: disaster, embarrassment
73. The **dormitory / burial** was mostly dreary but **it** had some redeeming aspects.
CW: dreary, redeeming, aspects
74. Her **bosom / charm** was a winning point and **it** got their attention.
CW: winning, point, got, attention
75. The **cologne / robbery** should have made her suspect him, but **it** didn't at all.
CW: suspect
76. That **belly / race** motivated him to run because **it** was important to him.
CW: motivated, run, important
77. The **avalanche / shortage** was a huge catastrophe since **it** cut off their food supply.
CW: huge, catastrophe, cut, food, supply
78. The **husbands / invaders** meant business and **they** wanted their way.
CW: business, wanted, way
79. That **bagpipe / hobby** repelled girls like **it** was a disease.
CW: repelled, girls, disease
80. That **balloon / hostility** hung in the air and **it** lasted for days.
CW: hung, air, lasted, days
81. The **aluminum / mileage** was fine, although **it** wasn't what he'd hoped.
CW: fine, hoped
82. The **armor / tyranny** was hard to overcome because **it** was so strong.

- CW: hard, overcome, strong
83. His **messengers / reflexes** were fast and **they** saved him from an unpleasant situation.
CW: fast, saved, unpleasant, situation
84. The **bouquet / apology** made her cry because **it** was so touching.
CW: cry, touching
85. The **ambulance / inquiry** was very slow and **it** almost came too late.
CW: slow, came, late
86. The **appliance / contract** was not very useful, nor would **it** make much revenue for the business.
CW: useful, revenue, business
87. The **coast / galaxy** was breath-taking as **it** came into view.
CW: breath-taking, came, view
88. Her **cloak / freedom** was her greatest treasure, and **it** couldn't be taken from her.
CW: greatest, treasure, taken
89. The **aspirin / employment** that she found helped but **it** didn't fully solve her problems.
CW: found, helped, fully, solve, problems
90. The **authors / gestures** were very inappropriate and **they** angered the whole group.
CW: inappropriate, angered, whole, group
91. The **customers / crafts** at the store were great because **they** were so high-end.
CW: store, great, high-end
92. The **kernels / imprints** of corn on her art project looked nice until **they** got messed up.
CW: corn, art, project, looked, nice, got, messed
93. That **creature / paradox** would torment him unless **it** was dealt with.
CW: torment, dealt
94. The **damsels / thefts** required the knight's attention because **they** were assigned to him by the king.
CW: required, knight's, attention, assigned, king
95. A **bullet / spasm** caused her to fall when **it** went through her leg.
CW: fall, went, leg
96. The **duchesses / gentries** of that time were very well-known, and furthermore **they** carried great political weight.
*CW: time, well-known, carried, great, political, weight
97. The **ministers / prayers** at that service were very moving because **they** appealed to everyone.
CW: service, moving, appealed, everyone
98. The **avenue / berth** was pretty narrow but **it** didn't bother her.
CW: narrow, bother
99. The **court / south** was a great place to live because **it** was so luxurious.
CW: great, place, live, luxurious
100. The **clove / whiff** of garlic upset the vampire, even though **it** couldn't hurt him.
CW: garlic, upset, vampire, hurt
101. The **beehive / discord** was chaotic and **it** held them transfixed.
CW: chaotic, held, transfixed
102. That **highway / mile** was the worst because **it** was so backed up with traffic.
CW: worst, backed, traffic
103. The **doorway / clearance** was not very high, and the architect thought **it** needed to be

- higher.
 CW: high, architect, thought, needed, higher
104. Their **corner / integrity** was threatened but **it** never gave in.
 CW: threatened, gave
105. The **furnace / inferno** blazed intensely as **it** cast red light onto the ground.
 CW: blazed, intensely, cast, red, light, ground
106. The **instrument / trapezoid** was hard to understand because **it** required more advanced mathematics.
 CW: hard, understand, required, advanced, mathematics
107. The **office / drama** at the capital is terrible but **it** is improving gradually.
 CW: capital, terrible, improving, gradually
108. That **cabinet / fallacy** cost them a fortune and **it** still causes problems.
 CW: cost, fortune, causes, problems
109. His **band / success** made him popular because **it** legitimized him to others.
 CW: popular, legitimized
110. His **forehead / fever** was scalding before, but now **it** had cooled down a bit.
 CW: scalding, cooled
111. The **cloth / distress** that she revealed was unexpected because **it** didn't seem appropriate for a wedding.
 CW: revealed, unexpected, appropriate, wedding
112. The **arrow / jolt** could have hurt him if **it** had been stronger.
 CW: hurt, stronger
113. That **gallery / period** was an important one because **it** shaped how art was done for years afterward.
 CW: important, shaped, art, years
114. The **catfish / incline** was hard to reach because **it** was deep in the water.
 CW: hard, reach, deep, water
115. Although **brass / safety** was once a priority at the factory, recently **it** had been neglected.
 CW: priority, factory, recently, neglected
116. While **bronze / fraud** can serve its purposes, in general **it** yields poor results.
 CW: serve, purposes, yields, poor, results
117. The **glacier / geography** at the national park attracted many tourists because **it** was so spectacular.
 CW: national, park, attracted, tourists, spectacular
118. His **cigar / anger** made him look foolish, because **it** didn't match his personality.
 CW: look, foolish, match, personality
119. The **brim / origin** of the ancient cups was key because **it** dated the archeological site.
 CW: ancient, cups, key, dated, archeological, site
120. The **camera / removal** of the reporter was met with dismay since **it** made them feel uncomfortable.
 CW: reporter, met, dismay, feel, uncomfortable
121. When **colonels / manners** are gruff all the time, **they** can provoke resentment.
 CW: gruff, provoke, resentment
122. His **officers / duties** were already annoying him, and **they** got more and more aggravating.

CW: annoying, got, aggravating

123. Her **baby / dignity** was precious to her because **it** was all she had left.

CW: precious, left

124. The **banana / disaster** got old quickly, even though **it** had seemed fresh two days before.

CW: old, quickly, seemed, fresh, days