

The effects of Referent Specificity and Utterance Contribution on pronoun resolution

Jennifer E. Arnold (jarnold@linc.cis.upenn.edu)

Institute for Research in Cognitive Science, University of Pennsylvania
3401 Walnut St. Suite 400A, Philadelphia, PA 19104

Maryellen C. MacDonald (mcm@gizmo.usc.edu)

Department of Psychology, Hedco Neuroscience Building
University of Southern California, Los Angeles, CA 90089-2520

Abstract

Two experiments explore how pronoun resolution is influenced by a) properties of discourse referents, specifically whether they are underspecified and in need of description, and b) the contribution of the pronoun-containing utterance, specifically whether it provides a description or specifies an event. We find that these factors interact, such that when an underspecified referent is in focus, reading is facilitated for description continuations, but when a specified referent is in focus, reading is facilitated in event continuations when the specified referent continues as the topic. This study reveals one of the complex interactions that underlies pronoun resolution.

Introduction

How do readers interpret pronouns? Research has identified numerous relevant factors, many of which are claimed to affect resolution only at the point where the pronoun is encountered. For example, the interpretation of the pronoun in (1) is guided by the roles of the potential antecedents, "Mary" and "Sarah", and the fact that one character is the more likely cause of the blaming event (e.g., Garvey and Caramazza, 1974).

1. Mary blamed Sarah because she...

It has been further suggested that verb biases only come into play at the moment that the reader encounters the pronoun, and that they do not lead the implicit cause to be more generally accessible beforehand (McDonald and MacWhinney, 1995; Garnham et al., 1996).

By contrast, Arnold (1998) proposed that reference processing is influenced by the likelihood that a given entity will be important to the following discourse, which is construed dynamically and is not localized to the referring form itself. If the information available to the comprehender suggests that the speaker is more likely to refer to one entity, comprehension is facilitated when such a reference occurs. On this view, referent activation is linked to the probability that the entity will be referred to, and in some cases activation can be anticipatory.

This approach to discourse processing is inspired by research on syntactic ambiguity resolution, which has

recently come to focus on how various aspects of the context make it more likely for the speaker to provide certain types of information, well before an ambiguity is encountered. For example, research on modifier ambiguities has found that NP-modifiers are easier to comprehend if the referential context makes the noun need modification. For example, a context containing a set of books makes it easier to parse "Put the book on the table on the floor", since without modification the bare NP "the book" is ambiguous (e.g., Altmann, Garnham, and Dennis, 1992; Crain and Steedman, 1985; Tanenhaus et al., 1995). In other cases, the need for modification is determined by properties of the referring expression itself. Thornton, MacDonald, and Gil (in press) found that a non-specific NP like "a house" was more modifiable than a more specific NP like "my house" ("a house with shutters" vs. "my house with shutters"), and that it was easier for readers to attach PPs to non-specific NPs than to specific NPs.

The approach in this line of work is fundamentally referential: these studies show that comprehenders attempt to find referents for referring expressions immediately and incrementally, and when a bare NP is not sufficiently informative, they search for further information in the linguistic input.

Referent Specificity. Our study applies the preceding logic to local discourse comprehension, and investigates the role of referent specificity during pronoun comprehension. We hypothesize that readers may find it likely that an underspecified character will be mentioned again soon, because they may expect the speaker to justify having introduced this character to the story. For example, in (2) readers may focus on "a student" as a likely topic of the following utterance.

2. On the first day of class, I saw the professor talking to a student in the front row. She...

If the underspecified character is indeed likely to be mentioned again, it should be easier to interpret a subsequent pronoun referring to this character.

At the same time, other aspects of referent specificity make contradictory predictions. One character, "the professor", has been introduced with a specific, definite NP. Definite NPs are often used for given, topical entities in a discourse (Prince, 1992), and although "the professor" is not given, it is inferable from the context of a class. If the speaker chooses a definite NP for this character, the comprehender may assume that it is meant to be a central character in the story. Thus, the definite, specific nature of the professor character may make it a probable topic of the following utterance.

Because of these contrasting predictions, we hypothesize that the comprehender's tendency to focus on one character or the other will be influenced by another factor: the comprehender's perception of how the following utterance relates to the story.

Utterance Contribution. A crucial part of utterance comprehension is interpreting how an utterance contributes to the task at hand (e.g., Clark, 1996; Grosz and Sidner, 1986). When the task is primarily linguistic, comprehension is driven by how the listener perceives the relation between the incoming utterance and the previous discourse (e.g., Garvey and Caramazza, 1974; Garnham et al., 1996; McDonald and MacWhinney, 1995; Stevenson, Crawley, and Kleinman, 1994). With respect to causal relations, as in (1), the interpretation of the pronoun depends on the comprehender knowing that the second clause is specifying the cause of the event described in the first clause. In this example, the connector "because" provides strongly constraining information about this relationship.

When the beginning of an utterance signals what its role is with respect to the previous utterance, it probabilistically influences the comprehender's expectations about where the discourse is going, which in turn impacts the likelihood that a given entity will be mentioned. For example, if the comprehender infers that an utterance will provide descriptive information, underspecified referents will be more likely to be mentioned. We hypothesized that if readers know a description is coming, they are likely to focus on things that need to be described, like "the student" in 2. In contrast, if the utterance appears to specify a subsequent event, readers will focus on characters they perceive as more topical, such as the more specified referent in 2, "the professor".

Hypothesis. We hypothesized that Referent Specificity and Utterance Contribution would interact to make underspecified referents more likely discourse continuations in descriptive contexts, and specified referents more likely continuations in event contexts, and that this would influence pronoun comprehension in the second utterance. Experiment 1 investigated which character was more likely to be mentioned in the continuation of a story. Experiment 2 looked at the comprehension of pronouns under different

conditions of Utterance Contribution and the specificity of the pronoun referent.

Experiment 1: Story-continuation

Methods and Participants. This experiment investigated whether specific or unspecific characters would be considered more likely continuations of a story, depending on whether the following utterance was perceived to be a description or an event. Participants were asked to read short "stories" like (3) and add a natural continuation to the end.

3a. SPECIFIC FIRST: I arrived at the café and discovered the waitress talking to a little boy.

a'. UNSPECIFIC FIRST: I arrived at the café and discovered a little boy talking to the waitress.

b. DESCRIPTION CONDITION: It looked like...

b'. EVENT CONDITION: Right then...

The sentence began with a scene-setting phrase, presented from the perspective of an observer (usually "I" or "we"). Each stimulus item included two characters, of different genders, denoted by NPs typically associated with only one gender (e.g., man, woman, actress, sailor). One character was specific and the other unspecific. Referent Specificity was manipulated by both NP definiteness and role specificity. Specific characters were identified by their roles, e.g. "the waitress" or "the mailman", and were consistent with the scene described in the first part of the sentence. All unspecific characters were either "a man", "a woman", "a (little) boy", or "a (little) girl".

All characters were human and animate. This was important, because our hypothesis was that comprehenders have some expectation for underspecified characters to be described under certain conditions. However, the perceived importance of an unspecific character is probably determined by many factors, one of which may be animacy. For example, some inanimates may be unimportant to the story, like "a beer" in "John drank a beer".

In addition, we manipulated two factors: a) Utterance Contribution (DESCRIPTION vs. EVENT), as described above, and b) Order of Mention (specific first vs. unspecific first).

Order of Mention is one of the strongest known factors affecting pronoun resolution. First-mentioned characters are more likely to be pronominalized in subsequent references, and pronouns are easier to understand if the referent is a first-mentioned character (e.g., Gordon et al., 1993; Stevenson et al., 1994; among others). The first character is the "starting-point" of the utterance (Chafe, 1994), it is the basis by which readers lay the foundation for the rest of the discourse (Gernsbacher, 1990), and it is the most likely character to be mentioned in the following discourse (Arnold, 1998). Because of the demonstrated strength of this factor, we hypothesized that it might interact with Referent Specificity and Utterance Contribution.




Figure 1: Percentage of participant completions beginning with reference to the specific character, unspecified character, both (as a compound NP or "they"), or other referent.

Table 1: Example continuations in Experiment 1, corresponding to categories in Figure 1.

STIMULUS (first sentence):

The first scene of the movie was the cowboy talking to a woman.

COMPLETION TYPE	EXAMPLE (with relevant referring form underlined)
specific	After that ... <u>the cowboy</u> got shot and the woman cried.
unspecific	It seemed like ... <u>she</u> was about to swoon over all over him.
both	After that ... <u>the woman and the cowboy</u> drove off in the wagon.
other	It seemed like ... one of those hokey old Westerns that Jimmy Stewart was in.

Each of the 12 experimental items was rotated through the four conditions that resulted from crossing the two factors (Utterance Contribution and Order of Mention). These were presented in 4 lists to 24 members of the Stanford University community,¹ along with 24 items from another experiment and 36 fillers.

The experiment was conducted using an oral story completion method, where participants read the stimulus sentences out loud into a tape recorder, and provided their continuation orally. This method has the advantage that people respond more quickly, which means that their responses reflect the on-line processes occurring as they reach the end of the stimulus. In addition, they do not restrict themselves to extremely short responses, as can be the case with written sentence-completion.

¹ One subject was excluded because he focused on the question of who he referred to in his continuations, and two subjects were replaced because they were non-native speakers of English. One item was excluded from the analysis due to experimenter error in stimulus construction. Four continuations were excluded because the participant produced an unintelligible response or repeated the stimulus.

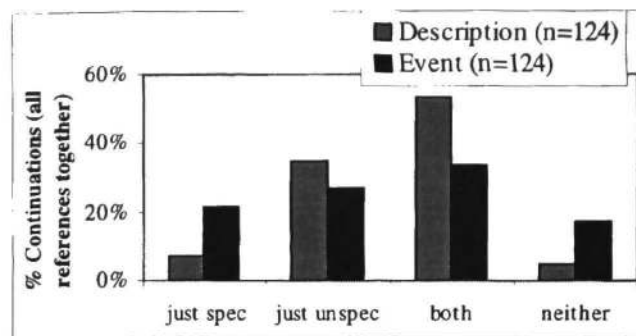


Figure 2: Percentage of participant completions that included references to both unspecified and specific characters, just one or the other, or neither (only some other referent).

Predictions. The goal of this experiment was to see which character participants referred to in their continuations more frequently. We expected that unspecified characters would be relatively more frequent in the DESCRIPTION condition, and specific characters more frequent in the EVENT condition. We also expected that these factors would interact with a tendency for story continuations to refer more often to first-mentioned than second-mentioned characters.

Results. The participant completions were tape recorded, transcribed, and analyzed to answer two questions: a) which character did the participant begin the continuation with? (Figure 1), and b) considering all references in the continuation, how often did the participant refer to both characters or just one or the other? (Figure 2).

Figure 1 shows that in the DESCRIPTION condition, speakers were most likely to begin their continuation with a reference to the unspecified character. In the EVENT condition, by contrast, both specific and unspecified characters were likely beginnings for the continuation (comparing specific and unspecified characters only, $\chi^2(1)=12.9, p<.001$). Counter to our expectations, there was no effect of Order of Mention nor any interaction between it and the other factors. Examples of each type of completion are listed in Table 1. The data in Figure 1 suggest that as predicted, in the DESCRIPTION condition the unspecified character became more accessible, and speakers began their continuations with this character.

However, Figure 2 shows a further difference between DESCRIPTION and EVENT conditions. These data consider all responses in the entire continuation, which show that participants were more likely to refer to both characters in the DESCRIPTION than in the EVENT condition (Z for two proportions=3.1, $p<.002$). This suggests that when comprehenders perceive that a description is coming, they are most likely to produce a description that describes the relationship between the two characters.

By contrast, the EVENT condition led to more varied responses. The introductory phrase in this condition often signaled a change in time or place. For this reason, participants were most likely to begin their response with a reference to something other than the specific or unspecified

character (e.g., "All of a sudden...the lights went out.") Taking the entire continuation into account, responses were relatively evenly split between those that referred to both characters, just the specific, just the unspecific, or neither. Contrary to expectations, event continuations did not focus primarily on the specific character, perhaps because specific characters are not strongly marked as likely topics of the next utterance, in the absence of other discourse cues like repeated reference. However, Figure (2) shows that continuations referring **just** to the specific character were more common in the Event than in the DESCRIPTION condition ($Z=3.2$, $p<.002$), suggesting that the EVENT condition does promote the accessibility of specific characters to a certain extent.

One limitation of the oral story-continuation methodology is that participants tend to focus on the second-mentioned character more than usual. In naturally occurring language, first-mentioned entities tend to be discourse-given, tend to be continued in the following discourse, and when they are referred to, are often pronominalized. However, it has been observed that task demands of the story-completion task lead to more frequent mention of the second-mentioned character, possibly reflecting a recency effect (see Arnold, 1998 for a discussion of this methodology). This pattern also emerged in our data here, in that participants referred equally often to the first-mentioned ($n=92$) and second-mentioned ($n=86$) characters ($Z=.39$, $p>.6$). This may explain why Order of Mention did not interact with the other variables of interest, Referent Specificity and Utterance Contribution.

In sum, Experiment 1 confirmed that the need for specification of some discourse characters interacts with the comprehender's perception of how a given utterance relates to the previous discourse. When the beginning of the utterance signaled a description, people began their continuations more often with the unspecific character, and they were more likely to mention both characters during the continuation than in the EVENT condition. The EVENT condition produced more varied responses, including a higher tendency to focus exclusively on the specific character than in the DESCRIPTION condition.

Our next question was how these patterns of probable story continuation relate to the on-line comprehension of pronominal references.

Experiment 2: On-line pronoun resolution

Methods and Participants. We used a self-paced moving window paradigm to present 16 two-sentence stories to 40 USC undergraduates, one word at a time. These items were combined with 40 items from two other experiments, 9 practice items, and 40 fillers, which were randomized in 8 lists.

The stimuli followed the same structure as those in Experiment 1. Each sentence contained one specific and one unspecific character. We also manipulated three variables: a) Order of Mention (specific first vs. unspecific first), b) Utterance Contribution (DESCRIPTION vs. EVENT

continuation), and c) Pronoun Referent (specific vs. unspecific character). Sample stimuli are in (4).

4a. SPECIFIC FIRST: When I got to the kitchen, I saw the maid yelling at a man.

a'. UNSPECIFIC FIRST: When I got to the kitchen, I saw a man yelling at the maid.

b. DESCRIPTION CONDITION: It seemed that {he/she} had spilled milk all over the floor.

b'. EVENT CONDITION: Shortly after that {he/she} stormed out the door.

The data from 10 participants were excluded from the analysis due to errors on more than 15% of the comprehension questions for this experiment ($n=8$) or extremely long reading times on ($n=2$). Data were trimmed at 2 standard deviations above and below the cell means.

We divided the stimulus items into nine regions, and analyzed the residual reading times for each region (Ferreira and Clifton, 1986). The scheme for regionization is detailed in Table 2.

Table 2. Regions analyzed in Experiment 2

Region	Example
intro to sentence 1	I walked into the room and saw
NP 1	a man
verb region	talking with
NP 2	the nanny.
intro to sentence 2	It seemed like
pronoun	she / he
next word (1)	was
next word (2)	very
end region ²	angry.

Predictions. We predicted that the specificity of the characters would interact with the contribution of the second utterance. We expected that in the DESCRIPTION condition, reading times would be shorter when the pronoun referred to the unspecific character, and in the EVENT condition, reading times would be shorter when the pronoun referred to the specific character. We predicted a possible interaction of these variables with Order-of-Mention, since this factor has been shown to be significant in other studies, despite its lack of influence in Experiment 1.

We also expected these results to occur in the region(s) immediately following the pronoun. Past work using the moving-window paradigm has established that the processing load for a given word is often observed one or two words later.

Results. The major finding was that Utterance Contribution produced different patterns of facilitation, depending on which referent was in focus: 1) When the unspecific

² The reading times for the last region are shown but were not analyzed, because this region contained a different number of words in each item.

character mentioned first (and therefore was in focus), the DESCRIPTION continuations were facilitated, and 2) when the specific character was mentioned first, the EVENT continuation was facilitated, but only when the pronoun referred to the specific character.

The first difference among the stimuli occurred during the first sentence, where half the items mentioned the specific character first, and half mentioned the unspecific character first. This distinction yielded two results. The more relevant result³ was that it influenced the way the rest of the item was read. Readers focused on the first-mentioned character, which determined whether facilitation occurred in EVENT or DESCRIPTION conditions.

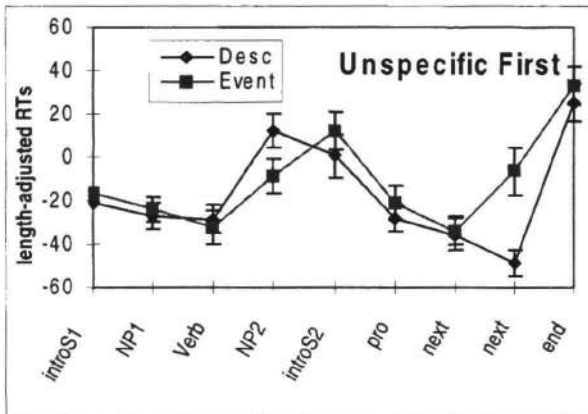


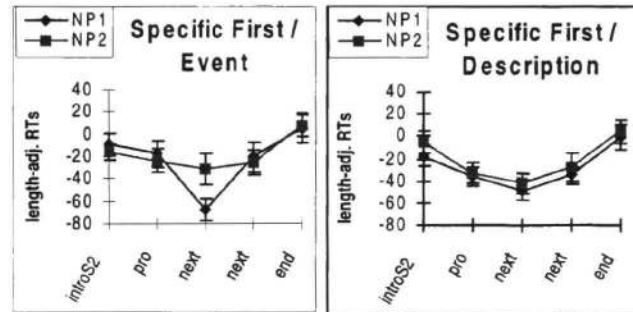
Figure 3: Reading times for each region for items where the unspecific character came first.

Result 1: Unspecific first/Description facilitation. When the unspecific character was mentioned first, participants were focused on an underspecified character in need of description. This need for specification was fulfilled in the DESCRIPTION condition, which resulted in facilitation at the second word after the pronoun. Figure (3) shows the contrast between DESCRIPTION and EVENT conditions for items where the unspecific character came first. An analysis of each region indicates that the only reliable difference between DESCRIPTION and EVENT conditions occurred at the second word after the pronoun ($F(1,29)=17.0$; $F(1,15)=12.5$; $p's < .005$), where reading times were shorter in the DESCRIPTION condition.

Note that the facilitation in the DESCRIPTION condition occurred equally for items with pronouns referring to specific and unspecific characters. That is, if readers were focussing on the unspecific character, reading was facilitated when they got a description, but it didn't matter if this description referred to the specific or unspecific

³ The less relevant result was that the reading times for the region following the noun phrases were longer for the specific characters than the unspecific characters (for NP1, $F(1,29)=7.6$, $F(1,15)=5.7$; for NP2, $F(1,29)=5.5$ $F(1,15)=5.4$; $p's < .05$). This may reflect one of two things: a) the infelicity of using a definite NP for introducing a new character, or b) the simple fact that our definite NPs (e.g., the fire chief, the nanny) were lower frequency words than our indefinite NPs (e.g., a woman, a boy).

character. This result is consistent with the findings from Experiment 1, where the DESCRIPTION condition made participants most likely to mention both characters during their continuation. We hypothesized that this was because describing the unspecific character was usually accomplished through a description of the relationship between the two characters. This is consistent with the idea that when readers are focused on the character that needs description, they accept all descriptive continuations as informative, regardless of which character the pronoun refers to.



Figures 4a and b: Reading times for Sentence 2 for items where the Specific character came first in Sentence 1. The contrast between items with pronouns co-referring with NP1 or NP2 is shown separately for DESCRIPTION and EVENT continuations.

Result 2: Specific First/Event facilitation. In contrast, when the specific character appeared first, the facilitation occurred in the EVENT continuation condition. Here, however, the facilitation only occurred in cases where the pronoun referred to the specific character. We conducted ANOVAs at each word in the second sentence, looking separately at the EVENT and DESCRIPTION conditions when the Specific character came first, comparing conditions where the pronoun co-referred with the first-mentioned NP (i.e., the specific character) or the second-mentioned NP. The only reliable difference occurred at the word after the pronoun in the EVENT condition ($F(1,29)=5.7$; $F(1,15)=4.6$; $p's < .05$).

Discussion

The major result of these studies was that comprehension was influenced by an interaction between character specificity and the perceived relationship between the two utterances. Experiment 1 showed that whether a specific or unspecific character was considered a more likely continuation depended on the perceived role of the next utterance. Experiment 2 showed that these factors influenced on-line reading times, and further that they interacted with Order-of-Mention. When the unspecific character was mentioned first, participants found a DESCRIPTION continuation easier to read in the region following the pronoun. In contrast, when the specific character was mentioned first, reading was facilitated if the

pronoun referred to the specific character in an Event continuation.

These results support a view in which reading comprehension is influenced by the reader's estimation of where the discourse is going. An important feature of this view is that this estimation is built up dynamically, and is influenced by both properties of focused referents (e.g., whether they are specific or unspecific), and other information that signals how the following utterance will relate to the story.

This study also shows that these factors influence how the pronoun is resolved and integrated with the predicate, as indicated by the fact that the observed effects occurred immediately following the pronoun. This suggests that pronoun resolution is not guided by simple rules like "pronoun refers to focused character". Rather than a general first-mentioned advantage, Experiment 2 showed that the features of the focussed referent determined how the contribution of the next utterance impacted comprehension. These data are consistent with a view that information relevant to pronoun resolution accrues from information throughout the discourse, and is not localized to either the introduction of the discourse entities or to the pronoun itself. This study manipulated the introduction to the second sentence as a way of signaling its role, but other factors like the tense of a phrase, discourse genre, or task demands may also influence the perception of utterance contribution and pronoun resolution.

These two experiments have begun to unravel some of the complex interactions that affect language comprehension. However, there are many unanswered questions. For example, why did Order of Mention interact with Referent Specificity and Utterance Contribution in Experiment 2 but not in Experiment 1? We suspect that this occurred because of task differences between the experiments. However, this and other questions need to be explored in future studies.

In sum, language comprehension is a referentially driven process. Speakers and writers establish discourse entities and predicate information about them, and comprehenders need to identify these referents. We knew that this influences syntactic ambiguity resolution; this study shows that a similar factor affects reading and pronoun resolution.

Acknowledgements

Thank you to Robert Thornton for his help with the experiment and discussions of the data. This research was partially funded by a Stanford University Graduate Research Opportunity grant, and partially funded by NSF Grant SBR-95-11270.

References

Altmann, G. T. M., Garnham, A., and Dennis, Y. (1992). Avoiding the Garden Path: Eye Movements in Context. *Journal of Memory and Language*, 31, 685-712.
Arnold, J. (1998). *Reference Form and Discourse Patterns*. Ph.D. Dissertation, Stanford University.

Chafe, W. (1994). *Discourse, consciousness, and time*. Chicago: Chicago University Press.
Chafe, W. L. (1976). Givenness, contrastiveness, definiteness, subjects, topics, and point of view. *Subject and topic*, C. N. Li (Ed.), 25-56. New York: Academic Press, Inc.
Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.
Crain, S., & Steedman, M. (1995). On not being led up the garden path: The use of context by the psychological syntax processor. In D. Dowty, L. Karttunen, & A. M. Zwicky (Eds.), *Natural language parsing: Psychological, computational, and theoretical perspectives* (pp. 320-358). Cambridge, UK: Cambridge University Press.
Ferreira, F., & Clifton, C., Jr. (1986). The independence of syntactic processing. *Journal of Memory and Language*, 25, 248-368.
Garnham, A., Traxler, M., Oakhill, J. & Gernsbacher, M. A. (1996). The locus of implicit causality effects in comprehension. *Journal of Memory and Language*, 35.517-543.
Garvey, C. & Caramazza, A. (1974). Implicit causality in verbs. *Linguistic Inquiry*, 5.459-64.
Gernsbacher, M. A. (1990). Language comprehension as structure building. Hillsdale, NJ: Erlbaum.
Gordon, P. C., Grosz, B. J. & Gilliom, L. A. (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, 17.311-357.
Grosz, B. & Sidner, C. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12.175-204.
McDonald, J. & MacWhinney, B. (1995). The time course of anaphor resolution: Effects of implicit verb causality and gender. *Journal of Memory and Language*, 34.543-566.
Prince, E. F. (1992). The ZPG letter: Subjects, definiteness, and Information-status. *Discourse description: diverse linguistic analyses of a fund-raising text.*, W. C. Mann & S. A. Thompson (Ed.), 295-325. Amsterdam: John Benjamins.
Stevenson, R. J., Crawley, R. A. & Kleinman, D. (1994). Thematic roles, focus and the representation of events. *Language and Cognitive Processes*, 9.473-592.
Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., and Sedivy, J. C. (1995). Integration of Visual and Linguistic Information in Spoken Language Comprehension. *Science*, 268, 1632-1634.
Thornton, R., MacDonald, M. C., and Gil, M. (in press). Pragmatic Constraint on the Interpretation of Complex Noun Phrases in Spanish and English. *Journal of Experimental Psychology: Language, Memory, and Cognition*.