Manuscript submitted for publication. Please do not cite without permission.

Integrating the Spatial Semantics of Verbs and Prepositions during Sentence Processing

Robert Thornton
Pomona College

Natalie M. Klein
University of Rochester

Elspeth Llewellyn
Scripps College

Martin Hackl
Massachusetts Institute of Technology

In order to understand the spatial meaning of an utterance, comprehenders must integrate the spatial semantics of verbs and prepositions. For example, the spatial interpretation of *at* is typically non-directional, simply specifying a location in space: in the sentence *Amy ran at the track*, *at* only specifies the location of the running event. However, in the phrase *threw at*, the inherent directionality of the verb *threw* coerces the meaning of *at* into specifying a bounded path along which a thrown object traveled, e.g., *Amy threw the ball at me*. We report reading time data that show a processing cost for this semantic coercion. The results show coercion effects for both verbs and prepositions, and support semantic theories in which directionality is built from simpler spatial representations.

Psycholinguistic studies of semantic processing typically focus on how semantics influences syntactic processing (e.g., Ferreira & Clifton, 1986; MacDonald, Pearlmutter, & Seidenberg, 1994; Pickering & Traxler, 1998; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995), rather than how semantic information itself is integrated and interpreted. Thus, whereas psycholinguists have well-developed accounts of how syntax is processed, relatively little is known about how a compositional semantic representation is composed during sentence processing. Recently, several studies have addressed this question by examining enriched semantic composition, a process in which the meaning of a word is coerced into a different meaning to meet the semantic restrictions of other words in the sentence (e.g., McElree, Traxler, Pickering, Seely, & Jackendoff, 2001; Piñango, Zurif,

This research was supported by NIA grant K01AG025195 to the first author and NSF grant 0642748 to the last author. This research came out of a project in the psycholinguistics course at Pomona College during the Fall of 2003. Thanks to Ben Acland, Nicole Bagette, Henry Bley-Vroman, Gabrielle Osborne, Jacob Stiglitz, and Andrew Waser. Address correspondence to Robert Thornton, Department of Linguistics and Cognitive Science, 185 E. 6th St., Pomona College, Claremont, CA 91711. E-mail: robert.thornton@pomona.edu.

& Jackendoff, 1999; Traxler, McElree, Williams, & Pickering, 2005). For example, McElree et al. (2001) examined phrases like *began the book*, which have been argued to require coercion because verbs like *begin* require an event as a complement rather than a simple NP, so the phrase is coerced to mean something like *began reading the book* (Jackendoff, 1997; Pustejovsky, 1995). McElree et al. (2001) found increased reading times for those phrases compared to ones that do not require coercion, such as *read the book*, offering support for coercion as a psychologically real mechanism of semantic integration. We build on that work by examining how spatial information from verbs and prepositions is integrated, enriched, and interpreted in real time. Specifically, we focus on how one source of spatial information, directionality, determines the spatial interpretation of a sentence.

Locatives and Motion Verbs

The spatial information conveyed in a clause is determined by the meaning of the verb and locative expressions (typically prepositional phrases) that characterize the spatial relations that hold among event participants. Thus, in order to understand how the semantic representation of spatial information is composed, it is crucial to understand how the spatial information from verbs and prepositions is integrated.

In English, locatives are realized predominantly as prepositional phrases (PPs), such as *inside the house*, *to the store*, or *along the river*. Two broad classes of locative PPs can be distinguished (e.g., Miller & Johnson-Laird, 1976; Jackendoff, 1983): PATH PPs characterize a bounded path along which the event described in the clause unfolds, as in (1), whereas PLACE PPs characterize the place within which a state exists or an event unfolds, as in (2).

(1) John drove into the parking lot.

(2) John exercised inside the house.

PATH PPs typically combine with motion verbs that describe events that unfold along a path. Interestingly, even though all motion verbs conceptually involve a path, one can distinguish linguistically two types: directional motion verbs require that a path be linguistically specified (or salient in the discourse) and non-directional motion verbs do not (Jackendoff, 1983).

The classification of verbs and prepositions by themselves as directional or non-directional is not always clear-cut. For example, a variety of PLACE prepositions can give rise to PATH readings when combined with certain directional verbs (although the opposite does not seem to hold c.f. Asher & Sablayrolles, 1995). Such a case is given in (3), where *inside the house*, which is typically used to name a location as in (2), specifies the path along which the event unfolds.

(3) John stormed inside the house.

In (2), the meaning of *inside the house* simply specifies a location, whereas in (3) it describes a path by naming its endpoint. One way to account for the different interpretations of *inside the house* in (2) and (3) is to assume that the preposition *inside* is ambiguous, and its core semantics include both a PLACE and PATH meaning. Conversely, on a coercion account, spatial prepositions have a core meaning that can be semantically enriched through context. On this account, the core meaning of PLACE prepositions like *inside* is non-directional, i.e., they are not ambiguous but specify a spatial location with no path. In (3), however, *inside* modifies the directional motion verb *storm*. Because directional motion verbs require a path to be specified, the meaning of *inside* is coerced into being directional. These two accounts make differing predictions about how PLACE PPs will be integrated with directional verbs. The ambiguity account posits that a PATH meaning is part of the core semantics of PLACE prepositions and will be easily accessible, given that all meanings of an ambiguous word are initially activated during word recognition (e.g., Tanenhaus, Leiman, & Seidenberg, 1979). One more recent class of lexical ambiguity models, selective access models, specifically argue that the appropriate meaning of an ambiguous word is selected by the context (e.g., Martin, Vu, Kellas, & Metcalf, 1999; Simpson & Krueger, 1991; Vu, Kellas, Petersen, & Metcalf, 2003). This account would argue that if the preposition is ambiguous in its spatial semantics, the appropriate PLACE vs. PATH meaning should be selected by the context, i.e., the appropriate meaning will be activated by the context without processing cost. Conversely, on a coercion account, spatial prepositions have a core meaning that can be semantically enriched through context. On this account, the core meaning of PLACE prepositions like inside is non-directional, i.e., they are not ambiguous but specify a spatial location with no path. In (3), however, inside modifies the directional motion verb storm. Because directional motion verbs require a path to be specified, the meaning of inside is coerced into being directional. These two accounts make differing predictions about how PLACE PPs will be integrated with directional verbs. The selective access ambiguity account posits that a PATH meaning is part of the core semantics of PLACE prepositions and will be easily accessible, given that both meanings are available for selection by the context. Alternatively, the coercion account predicts that extra work needs to be done to enrich a PLACE preposition into having a PATH interpretation; thus we should see a processing cost for items like (3) relative to ones like (2).

Just as we see flexibility in the spatial interpretation of prepositions like *inside*, many motion verbs display a certain amount of flexibility in spatial meaning, in that they can appear in contexts that denote a path as well as contexts that do not. In (4a), the verb *wander* is interpreted non-directionally because no path is specified: *at the school* specifies only the location at which the wandering happened. In (4b), however, *to the school* provides a path along with the wandering happened, and thus *wandered* is used directionally.

(4) a. Because he woke up early, the boy wandered at the school.

b. Because he woke up early, the boy wandered to the school.

Again, this difference could be accommodated by either ambiguity or coercion. On the ambiguity account, *wandered* is compatible with either a directional or non-directional meaning, so a PATH can be specified or omitted. On the coercion account, the core meaning of *wander* is non-directional, but because a PATH is explicit in (4b), the meaning of *wander* is coerced into being directional to be consistent with the denoted PATH.

In sum, the spatial meaning of verbs and prepositions can either be directional or non-directional, but this directionality is often flexible and can be changed by other elements in a sentence. The purpose of this paper is to investigate the mechanisms that construct the spatial interpretation of a sentence during online sentence processing: specifically, we ask whether a comprehender's ultimate spatial interpretation results from ambiguity resolution or semantic coercion. We present data from a self-paced reading experiment using items that juxtapose motion verbs and prepositions, pairing directional verbs with PLACE PPs (e.g., *darted at*) and non-directional verbs with PATH PPs (e.g., *wandered to*). Example items are presented in Table 1.

Table 1: Sample Materials

Directional Preambles
+ PATH PP:
To protect her nest, the bird darted to the hunter just now.
+ PLACE PP:
To protect her nest, the bird darted at the hunter just now.
Nondirectional preambles
+ PATH PP:
Because he woke up early, the child wandered to the school last Tuesday.
+ PLACE PP:
Because he woke up early, the child wandered at the school last Tuesday.

Thus, directionality can potentially be contributed either by the preamble containing the verb or the PP. For no-conflict cases, that is, cases in which directionality of the preamble and PP match, the ambiguity and coercion accounts make the same prediction. The interpretation is easily determined by the mutual directionality or non-directionality of the two components; no additional processing cost is required for semantic integration.

However, in the mismatch cases, the accounts make different predictions. Ambiguity predicts no processing cost because both the directional and non-directional meanings are available as part of the lexical representation; no additional work needs to be done. For the coercion account, however, in cases where a directional preamble is combined with a PLACE PP, the PP is coerced into a directional reading, which should result in an additional processing cost. In cases in which a non-directional preamble is combined with a PATH PP, the motion event is understood as unfolding along the path denoted by the PP, thus the preamble would be coerced into a directional meaning, again resulting in a processing cost. In both cases, the effect should arise at the preposition or soon after, the point at which both the verb and preposition semantics are available for integration.

Method

Participants

Twenty undergraduates from the Claremont Colleges volunteered to participate, none of whom participated in the norming study. All reported that they were native speakers of English and had normal or corrected-to-normal vision.

Materials and Design

Two sets of materials were designed, the first consisted of directional preambles and the second non-directional ones. Each set had 12 items and the classification of the preambles is described below. Within each set, the type of preposition was manipulated (PATH versus PLACE function). The prepositions were selected from the locative (non-directional) vs. directional classification of Zwarts (2005). The non-directional (PLACE) prepositions included *at*, *in*, *near*, *on*, *inside*, and *outside*. The directional (PATH) ones included *into*, *onto*, *to*, *across*, *through*, *up*, *down*, and *out*. Zwarts' classification distinguishes prepositions that indicate where something is at a particular point in time (i.e., the non-directional ones) from ones that indicate where an object is going (i.e., the directional ones). The test to distinguish these is whether or not the preposition can be used by themselves following some form of the copula *be*. For example, the sentence *the rice is inside the pot* is acceptable, whereas *the rice is into the pot* is not. The prepositions that can be used in that way with *be* therefore indicate a simple location and are classified as PLACE prepositions (*inside* in the example), the others that cannot simply specify a location but indicate direction are classified as PATH prepositions (*into* in the example).

The classification of the preambles was more complicated. Whereas the directionality of the prepositions seems to be lexically specified, the classification of the preambles seems to be contextually defined rather than the property of a specific lexical item; a few verbs such as *enter* or *exit* seem to have directionality in their core semantics and therefore require at least an implicit path, but most spatial verbs seem to be quite contextually flexible. For example, the verb *ran* can specify either movement along a path (e.g., *Pat ran from here to there*) or within a specific location (e.g., *Pat ran around inside the gym*). A better way to classify our preambles mirrors the argument/adjunct distinction, such that directional preambles, by definition, require PATH arguments to be specified whereas non-directional preambles do not and therefore take adjuncts as modifiers. To justify our interpretations of our items directionality, we performed three of the PP-argumenthood diagnostics enumerated by Schütze and Gibson (1999)) - ordering, movement, and wh-extraction - on our stimuli. The mismatch case (directional preamble + PLACE PP) resulted in an overall directional interpretation. The diagnostics suggest that the PLACE PP behaves as an argument in this environment, supporting our intuitions about their directionality. The non-conflict case (non-directional preamble + PLACE PP) resulted in a non-directional interpretation. The test results suggest that the PP behaves like an adjunct in this environment, supporting out intuitions about their non-directionality. Thus, the tests confirm that the directional preambles require an argument to specify their path but non-directional preambles do not.

Moreover, in order to ensure that our classification of the items was correct, we ran an online norming study. Two presentation lists were created with all 24 experimental items. The lists were counterbalanced, only one version of each item appearing on a list and each list having an equal number of items from each condition. Following each item, a set of possible interpretations was presented in a multiple choice fashion, each of which was clearly directional or non-directional. For example, in the *...wandered at school...* example in Table 1, the choices were *The child traveled to school by wandering* (directional), *The child crashed up against the school* (directional), and *The child was at school and wandered around there* (non-directional), as well as an *other*, fill-in-the-blank option. The order of the options was random, other than the *other* option always coming last. If the *other* responses were ambiguous in their directionality, they were excluded from analysis, otherwise they were included the appropriate category.

Eighty native English speakers completed the survey, forty for each list. The results confirmed our classification. A single sample t-test indicated that when both the preamble and PP were non-directional, participants tended to choose a non-directional interpretation, $t(79) = 6.06$, $p < .0001$. And when either (or both) was directional, participants tended to choose directional interpretations, all t's > 4.5, all p's < 0.0001. Thus participants' interpretations confirmed our initial classification.

Thus, there were two levels of the two factors: preamble type (directional vs. non-directional) and PP type (PLACE vs. PATH). Each item began with an introductory clause (e.g., *To protect her nest*), which was separated from the main clause by a comma. The main clause always consisted of eight words: (1) the definite determiner *the*, (2) the head noun of the subject (e.g., *bird*), (3) the verb, (4) the preposition, (5) another determiner, (6) a noun, and (7-8) a temporal phrase (e.g., *just now*).

For the self paced reading study, two lists were also created in the same counterbalanced fashion. Each list also began with eight practice items and the experimental items were randomized with 24 filler items. Thus each list was 56 items long. A yes/no comprehension question was composed for each item.

Procedure

A Dell PC was used to present the materials, and the MicroExperiemntal Laboratory (MEL; Schneider, 1988) software package controlled stimulus presentation and collection of reaction times. A single-word, self-paced reading task was employed: at the beginning of each trial, lines of dashes appeared on the screen, with each dash representing a character in the sentence to appear. Participants then pressed the space bar with their dominant hand to see each word in a non-cumulative fashion (Just, Carpenter, & Woolley, 1982). The keypress that ended the presentation of the last word of a sentence triggered presentation of the yes/no comprehension question. Participants answered the question by pressing a key marked either "yes" or "no" and were given feedback on screen about their accuracy. The practice items were presented first, followed by one of the experimental presentation lists. The practice items appeared in a fixed order and the experimental items within a list were presented in a random order. On average, the experimental session took approximately 15 minutes.

Results

Overall, participants answered 96.2% (SD=3.5%) of the comprehension questions correctly, and the range of accuracy was 88% to 100%. Error rates did not vary significantly by condition. Items answered incorrectly were excluded from reading-time analyses.

Reading times were length-adjusted for two reasons: (1) to compare reading times across conditions that varied in number of characters, and (2) to normalize reading times across subjects, so that items could be trimmed by each condition as well as by word. The procedure involved calculating a regression equation across all experimental items and fillers to determine a length-adjusted reading time for each word (Ferreira & Clifton, 1986). In order to ensure that a few extreme values did not grossly skew the equation, all raw reading times less than 100 ms or greater than 1500 ms were classified as participant or equipment failures and were excluded from analysis. This procedure affected 40 out of the 17,795 total observations (less than 0.3%) and only 24 of the 9,695 experimental observations (less than 0.3%).

Analyses of Variance (ANOVAs) were computed on RTs for each of 8 words: the preverbal determiner and noun, the verb, the preposition, and following four words. No reading time differences were expected before the preposition because the sentences did not differ across condition before then. Observations whose values were more than two standard deviations from the cell mean for that word in each condition were winsorized, which affected less than 4.5% of all observations. Again, the factors were type of verb (directional versus non-directional) and type of preposition (PATH versus PLACE). The results are presented in Figure 1.

As expected, there were no significant effects for the first three regions (determiner, noun, and verb). At the preposition, although neither main effect approached significance (both $p's > .2$, there was a significant interaction of both variables, F_1 $(1, 19)$ = 13.48, $p < .005$; F_2 $(1, 22)$ = 13.91,

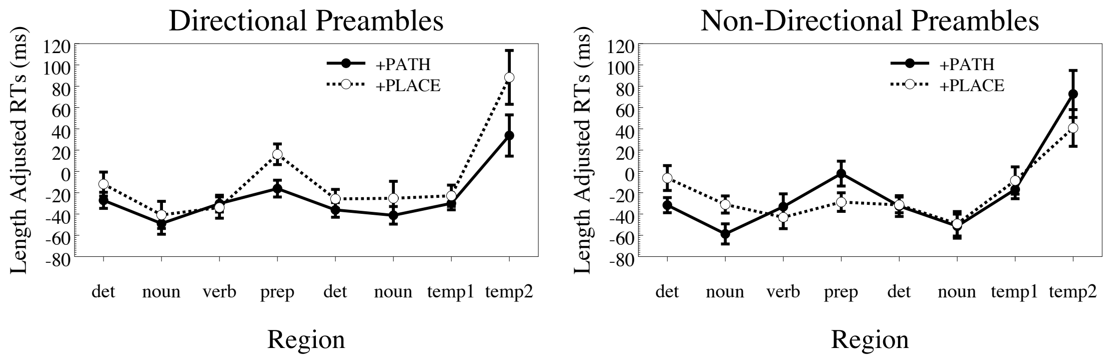

Figure 1. Mean length-adjusted reading times, with standard error bars computed across participants

$p < .001$. Pairwise comparisons reveal the nature of the interaction: for the directional preambles, PATH prepositions have significantly shorter RTs than PLACE prepositions, F_1 (1, 19) = 5.68, $p <$ .05; F_2 (1, 22) = 9.10, $p < .05$. Non-directional preambles show the opposite effect, with PLACE prepositions being read more quickly than PATH ones, F_1 (1, 19) = 4.12, $p = .057$; F_2 (1, 22) = 5.57, $p < .05$. Although this last effect was significant by items, it was only marginally significant by subjects (i.e., $p = .057$). The only other significant effect in the following regions was a main effect of verb type on the next to last word in the sentence, such that the directional-verb items were read faster than the non-directional ones, F_1 (1, 19) = 9.88, $p = .01$; F_2 (1, 22) = 4.83, $p < .05$.

Discussion

The results support the coercion account of semantic integration. In the mismatch cases, in which the directionality of the verbs and prepositions differed, we observed a processing cost such that RTs for the mismatching items were slower than for the matching items. This was true for both mismatch types (i.e., non-directional preambles + PATH PPs and directional preambles + PLACE PPs). These findings are inconsistent with accounts that predict that both directional and non-directional meaning are easily accessible (e.g., lexical ambiguity). These coercion effects occur very rapidly, arising at the earliest point at which participants could start to integrate the verb and prepositional semantics.

Moreover, our data support an important point elaborated upon in Jackendoff (1983), that directional meaning can be built out of non-directional meaning, but non-directionality cannot be constructed from a directional representation. While our data demonstrate that processing slows at the preposition in both mismatch conditions, our norming data indicate that if either element in the mismatch cases is directional, the final interpretation is also directional, whether the expectation for a PATH is contributed by the preamble or PP. Along with the self-paced reading data, this suggests that the processing cost is due to the coercion of the non-directional element into a directional one.

Alternative Interpretations of the Data

Despite the fact that the data are consistent with coercion, there are a number of potential simple, frequency-based alternative explanations for our data that we should rule out. First, it is possible that our PLACE and PATH prepositions differed in frequency, and the shorter reading times in some conditions are simply due to this discrepancy. This was not the case, with no frequency differences between the conditions in either the Brown corpus (Kucera & Francis, 1967), $t(12)$ = -.303, $p > .76$, or the HAL corpus (Lund & Burgess, 1996), $t(12)$ = -.142, $p > .89$.

Another possible alternative explanation is that our results were driven by the co-occurrence frequency of the verb + preposition pairs. Such an account would predict that high-frequency pairs should result in shorter RTs whereas relatively low-frequency pairs would result in longer RTs. In order to test this alternative, we looked to parsed corpora (e.g., the Penn Treebank) to obtain frequencies of the pairs we used. However, the very low frequency of some of our verbs (e.g., *scurried*, *squawked*) resulted in a sparse data problem. This problem has been addressed by Keller and Lapata (2003), who suggested that obtaining bigram frequencies using Google is a good solution by demonstrating high correlations between web and corpus frequencies as well as strong correlations between web frequencies and human plausibility judgments.

We obtained Google frequencies for each verb + preposition pair using a standard Google search with the pair in quotation marks. We then regressed the Google frequency of the pair with the length-adjusted RTs on the pairs preposition. The results show no significant relationship between bigram frequency and RTs on the preposition, R^2 = 0.0041, $p > .67$, strong evidence against a contingent frequency explanation of our effects.

Finally, it is possible that other lexical differences among our prepositions could have affected reading times. For example, studies have shown that words with multiple meanings are identified more quickly than words with fewer meanings (e.g., Borowsky & Masson, 1996). To rule out the possibility that any lexical variables accounted for our results, we examined data from the English Lexicon Project (Balota et al., 2002), which showed no differences between our prepositions across condition in either lexical decision, $t(12)$ = 0.874, $p > .399$, or naming latency, $t(12)$ = 1.320, $p > .211$. The lack of difference in the two main word recognition tasks indicates that our effects cannot be explained by lexical-level differences among the prepositions we used. In sum, there is no support for any lexical-level explanation of our data; the results are most consistent with coercion.

In sum, our results build on recent work on semantic processing by suggesting that coercion, a mechanism of semantic integration proposed for other phenomena, also applied to the integration of the spatial semantics of verbs and prepositions.

References

Asher, N., & Sablayrolles, P. (1995). A typology and discourse semantics for motion verbs and spatial PPs in French. *Journal of Semantics*, *12*, 163-209.

Balota, D. A., Cortese, M. J., Hutchison, K. A., Neely, J. H., Nelson, D., Simpson, G. B., et al. (2002). *The English lexicon project: A web-based repository of descriptive and behavioral measures for 40,481 English words and nonwords*. Washington University: http://elexicon.wustl.edu/.

Borowsky, R., & Masson, M. E. J. (1996). Semantic ambiguity effects in word identification. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *22*, 63-85.

Ferreira, F., & Clifton, J., C. (1986). The independence of syntactic processing. *Journal of Memory and Language*, *25*, 348-368.

Jackendoff, R. (1983). *Semantics and cognition*. Cambridge, MA: MIT Press.

Jackendoff, R. (1997). *The architecture of the language faculty*. Cambridge, MA: MIT Press.

Just, M. A., Carpenter, P. A., & Woolley, J. D. (1982). Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, *111*, 228-238.

Keller, F., & Lapata, M. (2003). Using the web to obtain frequencies for unseen bigrams. *Computational Linguistics*, *29*, 459-484.

Kucera, H., & Francis, W. N. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavioral Research Methods, Instruments, & Computers*, *28*, 203-208.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, *101*, 676-703.

Martin, C., Vu, H., Kellas, G., & Metcalf, K. (1999). Strength of discourse contexts as a determinant of the subordinate bias effect. *Quarterly Journal of Experimental Psychology*, *52*, 813-839.

McElree, B., Traxler, M. J., Pickering, M. J., Seely, R. E., & Jackendoff, R. (2001). Reading time evidence for enriched composition. *Cognition*, *78*, B17-B25.

Miller, G. A., & Johnson-Laird, P. N. (1976). *Language and perception*. Cambridge, UK: Cambridge University Press.

Pickering, M. J., & Traxler, M. J. (1998). Plausibility and recovery from garden paths: An eye-tracking study. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *24*, 940-961.

Piñango, M. M., Zurif, E., & Jackendoff, R. (1999). Real-time processing implications of enriched composition at the syntax-semantics interface. *Cognition*, *28*, 395-414.

Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.

Schneider, W. (1988). Micro experimental laboratory: An integrated system for IBM PC compatibles. *Behavioral Research Methods, Instruments, & Computers*, *20*, 206-217.

Schütze, C. T., & Gibson, E. (1999). Argumenthood and English prepositional phrase attachment. *Journal of Memory and Language*, *40*, 409-431.

Simpson, G. B., & Krueger, M. (1991). Selective access of homograph meanings in sentence context. *Journal of Memory and Language*, *30*, 627-643.

Tanenhaus, M. K., Leiman, J., & Seidenberg, M. S. (1979). Evidence for multiple stages in the processing of ambiguous words in syntactic contexts. *Journal of Verbal Learning and Verbal Behavior*, *18*, 426-441.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*, 1632-1634.

Traxler, M. J., McElree, B., Williams, R. S., & Pickering, M. J. (2005). Context effects in coercion: Evidence from eye-movements. *Journal of Memory and Language*, *53*, 1-25.

Vu, H., Kellas, G., Petersen, E., & Metcalf, K. (2003). Situation-evoking stimuli, domain of reference, and the incremental interpretation of lexical ambiguity. *Memory & Cognition*, *31*, 1302-1315.

Zwarts, J. (2005). Prepositional aspect and the algebra of paths. *Linguistics and Philosophy*, *28*, 739-779.