

Integrating Discourse and Local Constraints in Resolving Lexical Thematic Ambiguities

Joy E. Hanna

Dept. of Brain & Cognitive Sciences
University of Rochester
Rochester, NY 14627
hanna@bcs.rochester.edu

Michael J. Spivey-Knowlton

Dept. of Psychology
Cornell University
Ithaca, NY 14853
mjsk@cornell.edu

Michael K. Tanenhaus

Dept. of Brain & Cognitive Sciences
University of Rochester
Rochester, NY 14627
mtan@bcs.rochester.edu

Abstract

We conducted sentence completion and eye-tracking reading experiments to examine the interaction of multiple constraints in the resolution of a lexical thematic ambiguity. The ambiguity was introduced with prepositional "by"-phrases in passive constructions, which can be ambiguous between agentive and locative interpretations (e.g., "built by the contractor" versus "built by the corner"). The temporarily ambiguous sentences were embedded in contexts that created expectations for one or the other interpretation. The constraints involved, including discourse-based expectations, verb biases, and contingent lexical frequencies, were independently quantified with a corpus analysis and a rating experiment. Our results indicate that there was an interaction of contextual and more local factors such that the effectiveness of the contexts was mediated by the local biases. Application of an explicit integration-competition model to the off-line (sentence completion) and on-line (eye-tracking) results suggests that, during the processing of these ambiguous prepositional phrases, there was an immediate and simultaneous integration of the relevant constraints resulting in competition between partially active alternative interpretations.

Introduction

Constraint-based models of sentence processing (MacDonald, Pearlmuter & Seidenberg, 1994; Trueswell & Tanenhaus, 1994) propose that the comprehension system continuously integrates information from a number of sources in order to converge on a consistent interpretation. Within this framework, ambiguity resolution is viewed as a process in which information from contextual and more local sources combines to provide support for competing alternatives.

Recent studies have shown that discourse-based information can rapidly combine with local (within sentence) factors to influence ambiguity resolution (for a recent review, see Tanenhaus & Trueswell, 1995). Moreover, the effectiveness of discourse context is mediated by local constraints, such that context effects are strongest when local biases are relatively weak (MacDonald et al., 1994; Spivey-Knowlton, Trueswell & Tanenhaus, 1993). However, up to this point, constraint-based models of ambiguity resolution have been criticized as vague with regard to the specific constraints involved, their strength, and the exact time-course and nature of their integration. In addition, it is commonly assumed that, while multiple constraints may be integrated in *off-line* processing tasks where there is sufficient time to compute complex contextual biases, the architecture of the language processing system imposes restrictions on the time-course

with which different classes of constraints can be used in *on-line* processing.

In light of these considerations, the research reported here had two primary goals. The first was to further test the claim that discourse context effects are mediated by within-sentence constraints, including contingent lexical frequency and lexical thematic preferences. The second was to independently quantify local and contextual constraints in order to examine their integration in both off-line and on-line comprehension, in particular by applying an explicit integration-competition model to the results (e.g., Spivey-Knowlton, 1994; Spivey-Knowlton & Tanenhaus, submitted).

Much of the research investigating the effects of discourse context has focused on the effects of referential presuppositions in syntactic ambiguity resolution (e.g., Altmann & Steedman, 1988; Britt, 1994). However, the referential manipulations in these experiments may have been confounded with the discourse-based expectations they create (cf. Spivey-Knowlton & Sedivy, 1995). Additionally, the focus on syntactic attachment ambiguities has tended to ignore cases where the indeterminacy hinges on the lexical ambiguity of the preposition itself, including what thematic role is assigned to the object of the preposition.

In order to address the goals stated above, the experiments reported here made use of contextually created expectations which were non-referential in nature, and investigated the resolution of a (non-syntactic) lexical thematic ambiguity. The temporary ambiguity occurred within passive target sentences with the preposition "by." This preposition is lexically ambiguous, introducing a noun phrase that can be assigned the thematic role of an agent (e.g., "built by the contractor") or a location (e.g., "built by the corner") among others. Discourse constraints were manipulated by using embedded questions to establish an expectation for receiving information about an agent or location, without differing in their referential complexity. Example stimuli are shown in (1):

(1) a. Agent Context/Agent Target Noun:

The artist decided to go to the gallery. Once he got there he wanted to know *who* had hung his prize painting. He was pleased to discover that his painting had been hung *by the director* earlier in the week.

b. Location Context/Location Target Noun:

The artist decided to go to the gallery. Once he got there he wanted to see *where* his prize painting had been hung. He was pleased to discover that his painting had been hung *by the entrance* earlier in the week.

We report a corpus analysis and off-line rating and completion experiments using these materials which quantified the constraints involved in the resolution of this ambiguity and gave a measure of their integration in off-line performance. We also report an eye-tracking reading experiment which demonstrated that discourse-based expectations can affect the resolution of lexical thematic ambiguities. Furthermore, we show that discourse constraints interacted with more local sources of information, including verb preferences and co-occurrence frequencies, in a competitive manner predicted by the off-line ratings and applicable to both off-line and on-line processing.

Rating and Completion Experiments

A rating task was conducted on 24 subjects to obtain ratings for the importance of receiving information about the agent and the location of the relevant event (e.g., "On a scale from 1 to 5, how important is it for you to find out who/where the painting was hung?"). Ratings were elicited for the agent and location biasing contexts alone (the first two sentences of the stimuli), as well as for the target passive sentences alone (e.g., The painting had been hung ...). These importance ratings independently quantified the strength of two constraints: they indicated how strong the discourse expectations were for individual items as well as for agent and location biasing contexts in general; they also quantified the biases for individual passive verbs (in conjunction with their subject nouns).

An analysis of the Treebank Corpus was performed in order to quantify the frequency with which agentive and locative "by"-phrases occur in passive constructions. Out of over 300 passive sentences containing "by"-phrases, not a single one introduced a location, while agents were frequent (Hanna, Barker & Tanenhaus, 1995). This contingent frequency bias indicates that a "by"-phrase following a passive verb provides overwhelming support for the agentive interpretation.

Finally, gated sentence completions were collected in an experiment conducted on a different set of 36 subjects. The percentage of agentive, locative, and other completions were recorded for each of the 24 stimuli in six conditions, crossing the presence of either context or their absence (3) with the presence or absence of the preposition (2). Table 1 shows a summary of the completion results with sentence fragments up to and including "by".

Without the "by", (fragments like "The painting had been hung ..."; the No Context condition without the "by"),

there was in fact a preference for locative over agentive completions; this reflects a general verb preference in the stimulus set for locations over agents as was indicated by the importance ratings. With the addition of the preposition to the fragments, an agentive preference was found (Table 1), as is consistent with the bias revealed by the corpus analysis. Following an Agent Context, agentive completions were even more frequent; following Location Contexts however, agentive and locative completions were equally probable; $F(1,72)=70.55$, $p<.001$; $F(2,46)=18.13$, $p<.001$. These results clearly indicate that the discourse context was interacting with the contingent "by" frequency bias to produce the off-line performance. However, the pattern of results over all the verbs in the stimulus set obscures the differential effects of individual verb biases. Using the importance ratings, the stimulus items were divided into the most strongly agent and location biasing verbs (as well as an intermediate group). Agent-biasing verbs strengthened the preference for an agent completion in all conditions, but still showed an effect of context; $F(2,14)=7.58$, $p<.01$. Likewise, location-biasing verbs strengthened the preference for a location completion in all conditions, and also showed an effect of context; $F(2,14)=9.01$, $p<.005$.

These off-line results were modeled within an integration-competition framework (Spivey-Knowlton & Tanenhaus, submitted) that implements competition between possible alternative interpretations using recurrent feedback and normalization. In the Normalized Recurrence competition algorithm, constraints are defined in terms of their support for the possible alternative interpretations (in this case, an Agentive or Locative interpretation of the "by"-phrase). See Figure 1. Activations of the node pairs for each constraint (C_n) are first normalized: $C_n(\text{norm})=C_n/\Sigma C_n$. Then, activations for the provisional interpretations (I) are calculated as weighted sums of their corresponding constraint nodes (where the weights ω must sum to one): $I=\Sigma \omega_n C_n$. For the final operation within a cycle, the activation of an interpretation node is multiplied by the input that traveled up a particular pathway and added to the corresponding constraint node's current activation: $C_n=C_n(\text{norm})+I\omega_n C_n$. The model cycles through these operations, allowing converging biases among the constraints to cause the interpretation nodes to gradually settle toward one provisional (probabilistic) interpretation of the "by"-phrase.

For simulating the sentence completions (Table 1), with

Completion:	No context Verb+"by"		Agent Context Verb+"by"		Location Context Verb+"by"	
	Agentive	Locative	Agentive	Locative	Agentive	Locative
All Verbs	77%	23%	92%	8%	47%	53%
Agent-Biasing Verbs	100%	0%	100%	0%	67%	33%
Location-Biasing Verbs	39%	61%	77%	23%	20%	80%

Table 1: Percentage of Agent and Location Completions




Figure 1: Integration-competition architecture for combining constraints to resolve the "by"-phrase ambiguity. All connections are bidirectional. The heavier weighting of Constraint 2 reflects the strong expectations created by the stimulus contexts.

sentence fragments up to and including "by", the quantified constraints from each context and target sentence (provided by ratings and corpus data) were entered into the model, which was allowed to iterate for 10 cycles. After 10 cycles, the activation of the Agentive Interpretation node indicated the percentage (or probability) of Agentive completions for that stimulus item. Weights for the different constraints were varied until a set was found that approximated the sentence completion data. (However, once these weights were set, they were kept constant for the rest of the simulations.)

Item-by-item variation in the data was accounted for by the simulation results. Individual stimulus items that elicited high percentages of Agentive completions tended to have high activation of the Agentive Interpretation node after 10 cycles. Likewise, stimulus items that elicited low percentages of Agentive completions tended to have low activation of the Agentive Interpretation node after 10 cycles. See Figure 2.

This initial simulation demonstrates how the model can account for graded effects in subjects' biases regarding the lexical thematic ambiguity introduced by the "by"-phrase. More importantly, however, this simulation provided a set of weights for the different constraints that we could then use in attempting to simulate on-line reading data, where the duration of competition (number of cycles) should correspond to mean reading times in the various experimental conditions.

Eye-Tracking Reading Experiment

The results of the sentence completions clearly indicate that context, as well as verb bias, plays a strong role in resolving the lexical thematic ambiguity of the "by"-phrase. However, it is frequently argued that multiple constraints from discourse and local sources have a lengthy amount of time to be integrated in an off-line task completion, whereas in an on-line task such as eye-tracking reading, the processing




Figure 2: Item-by-item variation in sentence completion data and simulation results from Normalized Recurrence.

system imposes restrictions on the sources of information that can be integrated immediately (Rayner, Sereno, Morris, Schmauder, & Clifton, 1989). In order to investigate the time-course of constraint integration in an on-line task, the same materials were used in an eye-tracking reading study in which Agent and Location Contexts were crossed with congruent and incongruent Target Nouns (see, for example, the two congruent conditions in (1)a and b). Eye movements and fixation durations were recorded from 24 subjects reading one version of each stimulus (from the 2 X 2 factorial of Context by Target Noun) amidst 36 filler stimuli. For eye-movement analysis, the target sentences were segmented into four regions: Initial NP / Verb (including the auxiliary) / "by"-phrase / Next two words. The "by"-phrase region is the point of both the ambiguity and its potential resolution given the appropriate interpretation of the target noun.

Figure 3 shows the total reading times for each recording region, including both first pass and regressive fixations. Overall, total reading times at both the "by"-phrase and the next region were slower for Location Targets than Agent Targets; at the "by"-phrase, $F(1,20)=14.13$, $p<.01$; $F(2,1,20)=6.76$, $p<.02$; at the next region, $F(1,20)=14.47$, $p<.01$; $F(2,1,20)=12.40$, $p<.01$. In addition, there was a clear interaction of Context and Target Noun at the "by"-phrase region, such that for Agent Contexts, agentive "by"-phrases were processed much faster than were locative "by"-phrases, while for Location Contexts, both were processed equally quickly at an intermediate rate; $F(1,20)=13.23$, $p<.01$; $F(2,1,20)=11.47$, $p<.01$. This pattern was repeated but only marginally significant at the Next region; $F(1,20)=3.65$, $p=.07$; $F(2,1,20)=3.38$, $p=.08$. These results parallel those obtained with the off-line sentence completion task.

This general pattern of results is consistent with accounts that postulate a parsing preference for arguments over adjuncts (assuming agentive "by"-phrases constitute

arguments and locative ones adjuncts): locations should be harder to process than agents regardless of context (cf., Liversedge, Pickering & Branigan, 1995). However, the data pattern over all the stimuli conceals differential effects that emerge when verb biases are taken into account. The stimulus items were divided into strongly agent and location biasing verbs and two different patterns emerged, just as in the completion data. Among the eight Agent Biasing Verbs, total reading times at the "by"-phrase were fastest for Agent Targets and slowest for Location Targets, regardless of context; $F(1,4)=30.10, p<.01$. However, among the eight Location Biasing Verbs, there was a Context by Target interaction such that reading times were slowest for targets that were incongruent with the context, i.e. Agent Context - Location Target and Location Context - Agent Target; $F(1,4)=18.10, p<.02$. In other words, following Location Contexts and Location Biasing Verbs, subjects actually found it easiest to process Location Targets, a result that argument/adjunct proposals cannot account for.

Although the total reading time results showed a reliable interaction between context and target, it is commonly argued that total reading times, in general, do not provide evidence regarding the initial moments of ambiguity resolution (Rayner et al., 1989). First pass reading times, where fixation durations from regressive eye movements are excluded from the analysis, may provide a more accurate measure of initial processes during ambiguity resolution. Figure 4 shows first pass reading times across critical

regions of the target sentence. The results exhibit what looks like an interaction between context and target at the "by"-phrase, but this interaction (where Agent Targets were read slightly faster in Agent Contexts and Location Targets were read slightly faster in Location Contexts) did not achieve statistical significance; $F(1,20)=1.4, p>.1$; $F(2,20)=1.55, p>.1$. Results such as these (first pass reading times failing to show a reliable context effect) are typically interpreted as *clear evidence for a serial application of constraints in which context is not used during initial ambiguity resolution* (e.g., Rayner, Garrod & Perfetti, 1992).

However, in a competition model using Normalized Recurrence (see Figure 1), subtle graded effects of context can be accounted for by conflicting biases arising from the relevant constraints. In this model, ambiguity resolution is not a process of categorically selecting one alternative and possibly having to revise this discrete interpretation at a later time (e.g., Frazier, 1987). Instead, this model allows the alternative interpretations to be simultaneously partially active, and compete against one another over time. Thus, slow reading times are an indication of lengthy competition (due to conflicting constraints), rather than initial misinterpretation followed by revision.

Most importantly, in the model described here, all information sources (context and local constraints) are integrated simultaneously. Therefore, if this parallel model can simulate the first pass reading times at the "by"-phrase




Figure 3: Total reading times for the recording regions of the agent and location target sentences in agent and location biasing contexts. The Location Target Nouns cause more processing difficulty than Agent Target Nouns, but this interacts with Context at the "by"-phrase.




Figure 4: First pass reading times for the agent and location target sentences in agent and location biasing contexts.

(Figure 4), it would pose as an existence proof demonstrating that the data pattern *does not* provide "clear evidence for a serial system in which context is not used during initial ambiguity resolution."

The same computational architecture (with the same inputs and weights) used to simulate the sentence completion data was allowed to iterate until either interpretation node reached a dynamic criterion of activation that decreased over time ($1 - 0.01t$, where t is the number of cycles that have elapsed). The number of cycles it took for a stimulus item to reach criterion was an indication of how much competition should take place at the "by"-phrase in that stimulus item. Figure 5 shows the mean first pass reading times at the "by"-phrase combined with the mean number of competition cycles for items in those conditions.

In the simulation, the lengthiest competition occurred when a strongly *weighted* constraint (sentence completions were best fit with Context given a weight of $2/3$) was pitted against a strongly *biased* constraint (based on the corpus data, "by" is given a .98 bias for Agentive and .02 bias for Locative); this is the Location Context - Agent Target Noun condition. The briefest competition occurred when those two strong constraints both converged on an Agentive interpretation; this is the Agentive Context - Agent Target Noun condition. See Figure 5. Intermediate degrees of conflict between constraints elicited intermediate competition durations, and reading times; these are both Location Target Noun conditions. In particular, readers did not have as much difficulty as might be expected in the

Agent Context - Location Target Noun condition. This could be interpreted as evidence for the delayed use of contextual information. However, the model, which uses all information sources immediately, showed exactly this pattern. In addition, in this condition the model was actually frequently converging relatively quickly (at that point in the sentence) *toward an Agentive interpretation of the "by"-phrase*. Thus, in this condition, the model generally predicts that later on in the sentence, as further information provides




Figure 5: First pass reading times and model predictions at the "by"-phrase region.

support in favor of the Locative interpretation, there would be especially long reading times, and a high probability of regressive eye movements, as the constraints gradually overthrow the supremacy of the Agentive interpretation (or probabilistically "revise" the interpretation). Indeed, this Agent Context - Location Target Noun condition showed the longest total reading times (Figure 3), as well as the kind of regressive eye movements that would accompany a probabilistic "revision" of this sort. Figure 6 shows the probability of a regressive eye-movement from the Next Region (the two words following the "by"-phrase), overlaid with the absolute difference between the activation of the Agentive Interpretation node after competition at "by", and the input from the Target Noun's support for the Agentive interpretation. This metric provides an indication of the degree to which the model's "current" interpretation was incompatible with the new input.




Figure 6: Regressive eye-movements out of the Next region and model simulation of interpretation incompatibility.

General Discussion

The results presented here provide evidence that discourse context effects are mediated by local (within-sentence) constraints. Specifically, contextually created expectations interacted with a contingent lexical frequency bias and verb preferences during the resolution of a lexical thematic ambiguity. This points to the important and immediate role of both discourse and lexical representations during ambiguity resolution (MacDonald et al., 1994; Spivey-Knowlton et al., 1993; Trueswell & Tanenhaus, 1994). Furthermore, by independently quantifying the relevant constraints and examining their interaction within a framework that utilizes an explicit competition algorithm to model processing difficulty, a clearer picture of the time-course and nature of constraint integration and ambiguity resolution emerges. The relevant constraints are immediately and simultaneously integrated, resulting in competition between partially active alternative interpretations. This method of independent constraint quantification and explicit modeling of off-line and on-line data demonstrates that constraint-based modeling can be pursued in a principled fashion. Moreover, this account suggests that patterns of local processing difficulty that have

previously been interpreted as evidence for architecturally imposed delays in the use of certain constraints are in fact a natural consequence of competition among alternatives in a parallel constraint-based system.

Acknowledgments

This work was supported by NIH grant #HD27206 to MKT. We are grateful to Eiling Yee for assistance in data collection.

References

- Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, 30, 191-238.
- Britt, M. A. (1994). The interaction of referential ambiguity and argument structure. *Journal of Memory and Language*, 33, 251-283.
- Frazier, L. (1987). Theories of syntactic processing. In J. Garfield (Ed.), *Modularity in knowledge representation and natural language processing*. Cambridge: MIT Press.
- Hanna, J., Barker, C. & Tanenhaus, M. (1995). Integrating local and discourse constraints in resolving lexical thematic ambiguities. Poster presented at the 8th Annual CUNY Sentence Processing Conference. Tucson, AZ.
- Liversedge, S., Pickering, M. & Branigan, H. (1995). The comprehension of sentences that are ambiguous between agentive and locative interpretations. Poster presented at the 8th Annual CUNY Conference on Human Sentence Processing. Tucson, AZ.
- MacDonald, M., Pearlmutter, N. & Seidenberg, M. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, 101, 676-703.
- Rayner, K., Garrod, S. & Perfetti, C. (1992). Discourse influences in parsing are delayed. *Cognition*, 45, 109-139.
- Rayner, K., Sereno, S., Morris, R., Schmauder, R. & Clifton, C. (1989). Eye movements and on-line language comprehension processes. *Language and Cognitive Processes*, 4, 21-50.
- Spivey-Knowlton, M. (1994). Quantitative predictions from a constraint-based theory of syntactic ambiguity resolution. *Proceedings of the 1993 connectionist models summer school*, (pp.130-137). Hillsdale, NJ: Erlbaum.
- Spivey-Knowlton, M. & Sedivy, J. (1995). Resolving attachment ambiguities with multiple constraints. *Cognition*, 55, 227-267.
- Spivey-Knowlton, M. & Tanenhaus, M. (submitted). Syntactic ambiguity resolution in discourse: Modeling the effects of referential context and lexical frequency within an integration-competition framework.
- Spivey-Knowlton, M., Trueswell, J. & Tanenhaus, M. (1993). Context effects in syntactic ambiguity resolution. *Canadian Journal of Experimental Psychology: Special Issue*, 47, 276-309.
- Tanenhaus, M. & Trueswell, J. (1995). Sentence comprehension. In J. Miller & P. Eimas (Eds.), *Handbook of cognition and perception*. SD, CA: Academic Press.
- Trueswell, J. & Tanenhaus, M. (1994). Toward a lexicalist approach to syntactic ambiguity resolution. In C. Clifton, L. Frazier & K. Rayner (Eds.), *Perspectives on sentence processing*. Hillsdale, NJ: Erlbaum.