

Knowing What's Going On. The Influence of Contextual Knowledge on the On-line Resolution of Structural Ambiguity

Nina Versteeg (Nina.Versteeg@let.uu.nl)

Utrecht University, Utrecht institute of Linguistics OTS
Trans 10, 3512JK, Utrecht, The Netherlands

Ted Sanders (Ted.Sanders@let.uu.nl)

Utrecht University, Utrecht institute of Linguistics OTS
Trans 10, 3512JK, Utrecht, The Netherlands

Frank Wijnen (Frank.Wijnen@let.uu.nl)

Utrecht University, Utrecht institute of Linguistics OTS
Trans 10, 3512JK, Utrecht, The Netherlands

Abstract

Can world knowledge modulate on-line parsing preferences? This was investigated in two off-line experiments and an on-line reading study. World knowledge was manipulated through the discourse context. More specifically, it concerned the mental representation that readers build up of the state of affairs described by the text. The results showed that this “contextual knowledge” immediately influenced the initial analysis of the critical sentence. Second, contextual knowledge appeared to overrule not only syntax-based principles, but the influence of topic-structure as well. These results provide evidence for the notion of an interactive parser that deals with different sources of information at the same time.

Keywords: contextual knowledge; syntactic ambiguity resolution; interactive processing; constraint satisfaction.

Introduction

Many studies of language processing focus on the question whether non-syntactic linguistic information, such as lexical semantics and discourse structural factors, such as topic-structure or referential success, have an immediate influence on the resolution of ambiguities. The foundation of many studies on the impact of referential success has been laid by Crain & Steedman (1985). Crain and Steedman claim that whether readers are garden-pathed in sentences like the famous *the horse raced past the barn fell* depends on the preceding context. If more than one equally plausible referent for *the horse* is available, readers need to restrict the set of possible referents when they encounter this phrase. Since a restrictive relative clause is a good way of doing this, readers would, according to the theory, immediately interpret the sentence as a relative clause – i.e. arrive at the correct interpretation. However, if the preceding discourse contains only one possible referent for *the horse*, there is no set of possible referents that needs to be restricted. Therefore, readers are expected to initially select the simple NP analysis, causing a garden path effect. In other words, the degree of referential success of the possible analyses of a sentence predicts which one is preferred. Since 1985 much evidence has been reported in favor of this idea (e.g.

Altmann & Steedman, 1988; Ni, Crain & Shankweiler, 1996; Van Berkum, Brown & Hagoort, 1999).

However, the extent to which world knowledge has an immediate impact on parsing has not been investigated very intensively as yet. In the present study world knowledge is operationalized as a reader's knowledge about the (interpersonal) relations between the referents in a sentence. In other words, when you know that Pete hates Mary but not Ellen, does that affect your inclination to treat the last two proper names in the string “Pete hates Marry and Ellen...” as a coordinated noun phrase (cf. Hagoort, Hald, Bastiaansen, & Petersson 2004)? This is the question asked in the present study. In order to find an answer, we look at NP/S-coordination ambiguities such as those in (1).

1a. *Pete said that he was very annoyed by the unfounded opinion of Mary and Ellen, a remark that they both found very rude.* [NP-coordination]

1b. *Pete said that he was very annoyed by the unfounded opinion of Mary and Ellen actually found that the discussion was blown up out of all proportion.* [S-coordination]

Previous research has shown that readers are inclined to initially interpret the ambiguous NP *Ellen* as part of a complex object NP, as in sentence (1) (Frazier, 1987; Hoeks, Vonk & Schriefers, 2002). Frazier (1987) uses her *Minimal Attachment* strategy to explain this preference: NP-coordination requires fewer syntactic nodes and is therefore more economical.

However, Hoeks et al. argue that readers prefer the structure that requires fewest accommodations of the existing topic-structure. An S-coordination requires two topics (*Pete* and *Ellen* in example sentence (1b)) and is therefore preferred if the preceding discourse has a two-topic structure, i.e. if two people have been in the center of attention throughout the preceding discourse. An NP-coordination on the other hand, requires one topic (*Pete* in example (1a)) and is therefore preferred if the preceding discourse has a one-topic structure, i.e. if one person has been in the center of attention throughout the preceding discourse. Hoeks et al. report evidence for their prediction in

a completion study, a self-paced reading experiment and an eye tracking study.

In the present study the focus will not be on the influence of a semantic or discourse-structural factor but on a more general type of knowledge, namely world knowledge, as provided by means of discourse context (i.e. contextual knowledge). More specifically we will investigate whether readers' knowledge about the state of affairs described in the text guides the parsing of ambiguities.

The critical manipulation was to embed S-coordinations that were temporarily ambiguous between S- and NP-coordination in contexts that were either supportive or non-supportive of S-coordination. Consider the following translations of Dutch experimental items that were either supportive of S-coordination (example text 2a.) or not supportive of S-coordination (example text 2b.):

2a.

Young writer Bart Schut has lashed out at critic Manon Thijssen during the literary TV-programme 'The Plantation'. (1) The writer said that he was very upset by Thijssen's negative review of his debut. (2) In the broadcast, on the other hand, the debutant did receive appreciation from journalist Francien Koopmans. (3) This he could appreciate more, he pointed out. (4) Schut said that he was very annoyed by the unfounded opinion of Thijssen and Koopmans actually found that the discussion had got completely out of hand. (5) Next week 'De Volkskrant' will publish a review of the book too. (6)

2b.

Young writer Bart Schut has lashed out at critic Manon Thijssen during the literary TV-programme 'The Plantation'. (1) The writer said that he was very upset by Thijssen's negative review of his debut. (2) In the broadcast the debutant was criticized by journalist Francien Koopmans as well. (3) This he could not appreciate either, he pointed out. (4) Schut said that he was very annoyed by the unfounded opinion of Thijssen and Koopmans actually found that the discussion had got completely out of hand. (5) Next week 'De Volkskrant' will publish a review of the book too. (6)

Sentence (3) of text (2a) states that writer Bart Schut received *appreciation* from journalist Koopmans. Based on this information it is highly plausible that the reader immediately chooses to interpret sentence (5), the target sentence, as an S-coordination. After all, it is not plausible at all that Schut is annoyed by both Thijssen and Koopmans. Sentence (3) of text (2b) however states that writer Bart Schut is criticized not only by critic Manon Thijssen, but by journalist Francien Koopmans as well. Based on this information it is highly plausible that the reader chooses to interpret sentence (5) not as an S-coordination, but as an NP-coordination. After all, in this case the contextual knowledge of the reader makes it highly plausible that Schut is annoyed by both Thijssen and Koopmans.

Our main prediction is that contextual knowledge guides the parsing of ambiguities and can therefore modulate the

on-line parsing preference for NP-coordination to a preference for S-coordination. In other words, we expect ambiguity resolution to be an interactive process in which different knowledge sources are invoked at the same time. To test this prediction we carried out three experiments: an off-line completion study, an off-line judgment study and an on-line reading experiment. In the remaining part of this paper all three experiments will be discussed in turn.

Experiment 1: sentence completion

The main goal of the completion study was to investigate the influence of contextual knowledge on structure building. Furthermore, the experiment was conducted to test our materials, in particular the manipulation of contextual knowledge.

The experimental texts were designed to manipulate contextual knowledge in the way described in the previous section. The dependent factor was the syntactic structure of the target sentence as completed by the participants. The target sentence was interrupted after the first NP (proper name) after the connective 'and', as in example sentence (3):

3. *Schut said that he was very annoyed by the unfounded opinion of Thijssen and Koopmans...*

At the point of interruption, the target sentence is ambiguous between S- and NP-coordination.

We expected contextual knowledge to affect the way participants completed the interrupted target sentence. More specifically, we expected participants to complete the target sentence more often as an NP-coordination when the context biased towards NP-coordination (i.e., when it is highly plausible based on the context that Schut is annoyed by both Thijssen and Koopmans) and more often as an S-coordination when the context biased towards S-coordination (i.e., when it is highly plausible based on the context that Schut is annoyed by Thijssen, but not by Koopmans).

Method

Participants One hundred and thirty students at the Faculty of Arts of Utrecht University participated in this and the following experiment for course credits. All were native speakers of Dutch and naive as to the purpose of the experiment.

Materials For this and the following experiments twenty-four experimental texts similar to those in example texts (2a) and (2b) were constructed, each with an NP- and an S-coordination-supportive version. All the materials were written in Dutch. They were designed to resemble short newspaper reports. All experimental texts contained a one topic-structure, so both syntax-based principles and topic-structure strongly biased against S-coordination.

Design For this and the following experiment eight experimental lists were constructed, so that each condition was equally represented and no two conditions from the

same item appeared in the same list. Each list consisted of twenty-four experimental texts interspersed among twenty-four filler items in such a way that the list started with two filler items and after this every experimental text alternated with a filler item. The fillers consisted of texts that closely resembled the experimental texts. The main differences were that different interpersonal relations were described than in the experimental texts and that the S-coordinations were replaced by NP-coordinations and constructions with other connectives.

Procedure Participants were instructed to complete the target sentence after the third proper name of the sentence in a way they considered grammatical as well as plausible, given the preceding text. It was emphasized that it was not important that the continuations were funny or original. Because participants took part in the experiment during class time, they had approximately one hour to finish the experiment.

Results & discussion

Ungrammatical continuations and missing data, for instance from participants who were unable to finish in time, were excluded from further analyses. This concerned 2% of the data. The remaining data were submitted to hierarchical regression analyses.¹ The results showed that the target sentence was completed more often as an S-coordination when the context supported S-coordination and more often as an NP-coordination in a context that was supportive of this structure ($\chi^2 = 14.33$; $df = 1$; $p < 0.001$; see Figure 1).

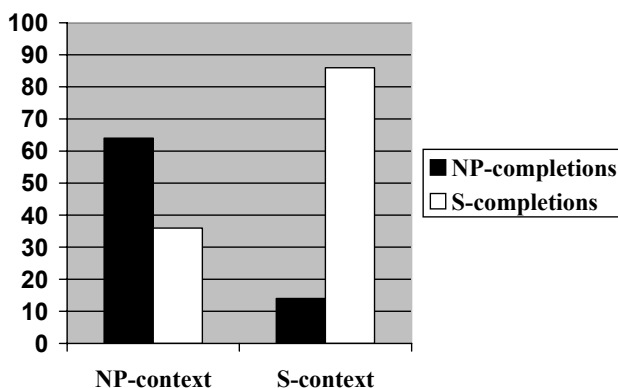


Figure 1. Percentages of completion as NP-coordination or S-coordination as a function of contextual knowledge.

Apparently, readers continue a sentence in the way that is most plausible based on their contextual knowledge. This result confirms our hypothesis. We also conclude that our manipulations had effectively determined readers' contextual representation of the text.

Experiment 2: judgements

Experiment 1 tells us what readers consider as an adequate continuation of an ambiguous sentence fragment, based on their contextual knowledge. Experiment 2 was designed to test whether a certain continuation-type is judged to be easier, more plausible, and more natural in one context as compared to another.

Contextual knowledge was manipulated in the same way as in Experiment 1. We measured participants' judgments of the perceived complexity, plausibility and naturalness of the target sentence within its context. We expected that participants would judge the target sentence, which was always an S-coordination, to be easier, more plausible and more natural when it was embedded in a context supporting S-coordination than when it was embedded in an NP-coordination-supportive context. The reason for making the target sentence an S-coordination in all cases was that by comparing the processing of the non-preferred S-coordination in both supportive and non-supportive contexts, we were able to investigate whether contextual knowledge can modulate the preference for NP-coordination to a preference for S-coordination.

Method

Procedure After participants finished the completion study they had a short break, after which they participated in the judgment study. Again, they had approximately one hour to finish the experiment. Participants were told that they were about to judge the completions of other students that participated in the completion study the week before. Their task was to judge the target sentences (alleged completions) on understandability, plausibility within the context and naturalness. They indicated their assessments by markings on three separate 5-point scales, one for each variable. Scale value 1 was used to indicate a negative evaluation once (very difficult) and a positive evaluation twice (very plausible / very natural) and scale value 5 was used to indicate a positive evaluation once (very easy) and a negative evaluation twice (very implausible / very unnatural).

Results and discussion

Missing data, for instance from participants who were unable to finish in time, were excluded from the analyses. This concerned 1% of the data. The remaining data were submitted to hierarchical regression analyses. The difficulty-scale was reversed in the analysis, so that scale value 1 always indicated a positive evaluation and scale point 5 always indicated a negative evaluation. The results showed that participants judged the target sentence to be easier, more plausible and more natural in the S-coordination-supportive context than in the NP-coordination-supportive context ($\chi^2 = 3.58$; $df = 1$; $p = 0.03$ / $\chi^2 = 26.81$; $df = 1$; $p < 0.001$ / $\chi^2 = 6.80$; $df = 1$; $p < 0.01$ respectively; see Figure 2). These results confirm our hypotheses.

¹ Many thanks go to Huub van den Bergh for his statistical advice.




Figure 2: Mean judgments on the easiness, plausibility and naturalness of the target sentence as a function of contextual knowledge (five-point scale).

The most important conclusion that we can draw from the judgment study is that contextual knowledge influences the representation of a sentence that a reader has after reading (i.e. the product of sentence processing). Apparently, participants' judgments were not ultimately based on the syntactic simplicity of the target sentence in itself or on the degree in which the existing topic-structure had to be accommodated, but, again, contextual knowledge appeared to overrule these factors. We take this as more evidence for the effective manipulation of contextual knowledge.

Experiment 3: self-paced reading

The two off-line studies clearly showed that contextual knowledge influences structure building and the mental representation that readers have built from a sentence after reading. In the third experiment we measured on-line reading processes, using a moving window self-paced reading paradigm. Specifically, we investigated whether contextual knowledge can overrule the parsing preference for NP-coordination. The Dutch target sentence and its literal translation, including the way it was segmented, is presented in (4a.) and (4b.) respectively.

4a. *Schut / zei / zich te storen aan / het ongegronde oordeel van / Thijssen / en / Koopmans / vond / eigenlijk / dat / de discussie / te veel / werd opgeblazen.*

4b. *Schut / said / to be very annoyed by / the unfounded opinion of / Thijssen / and / Koopmans / found / actually / that / the discussion / too much / was blown up.*

The concrete expectations for example sentence (4b.) would be the following. If contextual knowledge affects on-line parsing decisions, readers will immediately select an S-coordination in an S-coordination-supportive context, i.e. the correct interpretation of the sentence. Because of this, no reanalysis is needed and no increase in processing time will occur. However, if the context is supportive of NP-coordination, readers will initially select NP-coordination as the correct interpretation of the sentence. Because of this,

reanalysis of the sentence to S-coordination turns out to be necessary when the disambiguating verb 'found' is encountered, leading to an increase in processing time.

Method

Participants Fifty-three students at the Faculty of Arts of Utrecht University participated in the experiment. None of them participated in experiment 1 and 2. They were paid € 7,50 for their participation. All were naïve as to the purpose of the experiment.

Materials The materials that we used were basically the same as those in experiment 1 and 2. On the basis of the results for these studies, we only made some minor changes in some of the texts in order to strengthen the context manipulation.

Design The experiment started with a practice session comprising three filler texts. Then, the first twelve experimental texts and nine filler items were presented in a random order to the participants. After a one-minute break the remaining texts were presented randomly. Two experimental lists were constructed, so that each condition was equally represented and no two conditions from the same item appeared in the same list. The texts that appeared together in the same condition were randomly selected; this co-occurrence was the same for both lists. The texts that appeared in the first part of the experiment in the first list appeared in the second part of the experiment in the second list and vice versa.

Procedure Participants were seated in front of a computer. They were instructed that short texts would be presented. For contextual knowledge to be an independent variable, it was necessary that participants had properly processed the manipulated discourse information by the time they were reading the target sentence. To increase the likelihood that they had actually done so, participants were instructed to pay specific attention to the mutual relationships between the characters in the texts.

The texts were presented using the software programme E-prime. Initially, only the title of a text was visible on the screen, the letters of each word of the rest of the text appeared in dashes, so that only the basic lay out of the text was visible. Full stops were also presented. Every time the spacebar was pressed, the dashes of one constituent were replaced by the actual words. When the spacebar was pressed again, the words of the following constituent appeared and the words of the previous one were replaced by dashes again. In this fashion the entire text was read, making it impossible to look back in the text. We chose to present the texts constituent-by-constituent because this was considered less unnatural than a word-by-word presentation. However, to be able to measure the reading times of the critical segments as accurately as possible, all critical parts of the target sentence were presented word-by-word.

After each text, participants had to verify three statements. The first statement of an experimental text

always questioned the manipulation of contextual knowledge. This statement was always true to avoid any effects due to possible processing differences between answering affirmative or negative. For our example this means that if the text biased towards NP-coordination the statement would be ‘Schut did not receive appreciation from Koopmans’ and if the context biased towards S-coordination the statement would be ‘Schut received appreciation from Koopmans’. The basic idea was that if participants were not capable of correctly verifying this statement, we could not be sure that they had properly gathered the contextual knowledge of the corresponding text.

Results and discussion

The reading times of the participants with correct answers to the critical statement were submitted to hierarchical regression analyses. The results for the third proper name (‘Koopmans’), the disambiguating verb (‘found’) and the adverb (‘actually’) are shown in Figure 3. We included the proper name prior to the disambiguating verb in the analysis, because this is the first point at which participants might have started anticipating the structure of the target sentence. After all, at the third proper name participants had received enough information to decide if NP- or S-coordination would be a plausible structure of the target sentence.




Figure 3: Reading times (msec) for participants who had correct answers to the critical statement as a function of contextual knowledge.

For the critical segments there was no overall effect of contextual knowledge. When we analyzed the critical segments separately, we found no effect of contextual knowledge on the processing of the disambiguating verb and the third proper name. However, there was a main effect of contextual knowledge on the processing of the adverb (*actually*): it was read faster in the S-coordination-supportive context than in the NP-coordination-supportive context: ($\chi^2 = 9.52$; $df = 1$; $p = 0.001$).

As mentioned before, the data from participants with wrong answers to the critical statement were excluded from the analyses. However, these data could still be interesting: the

critical statement questioned whether participants had understood the part of the text that provided S-coordination-supportive or NP-coordination-supportive knowledge. For that reason, and because these participants had wrong answers, one could assume that they erroneously considered the context as supportive of S-coordination when it in fact was not, and vice versa. Then, to support the idea that contextual knowledge immediately influences ambiguity resolution, the reading times should show a pattern opposite to the ones realized by participants with correct answers. Interestingly, this is indeed what the results for participants with wrong answers show: the adverb was processed faster in the NP-coordination-supportive context than in the S-coordination-supportive context ($\chi^2 = 2.76$; $df = 1$; $p = 0.048$; see Figure 4). Note that the participants with wrong answers did not differ from participants with correct answers in other respects than their answers to the critical statements. The differences between texts, for instance, were the same for both groups of participants. Therefore, it is unlikely that participants with wrong answers to the critical statements were just guessing and had no representation of the passage whatever. There was no overall effect of contextual knowledge, nor was there a main effect of this factor on the third proper name and the disambiguating verb separately.

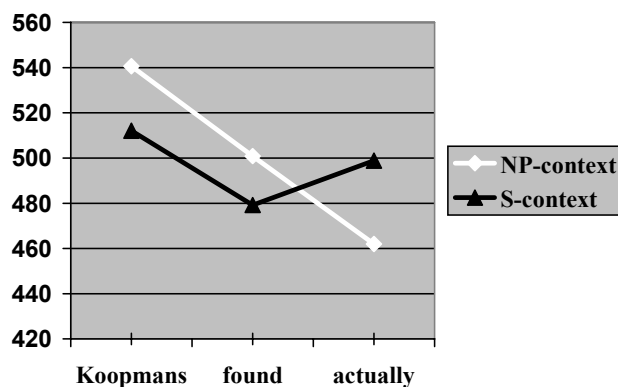


Figure 4: Reading times (msec) for participants who had wrong answers to the critical statement as a function of discourse information.

Another interesting result was that contextual knowledge not only outweighed syntax-based principles (i.e., minimal attachment), but also the principle of topic-structure-economy (readers prefer the structure that requires fewest accommodations of the existing topic-structure, cf. Hoeks et al. 2002).

The results for the verification task showed that participants did very well on the critical statement in general and that there was no effect of contextual knowledge on the amount of correct answers: in the NP-context the proportion of correctly verified statements was 89%, in the S-context this was 88%. Apparently, contextual knowledge provided by the discourse context was processed equally well in both conditions.

The reading results for both participants with correct and wrong answers to the critical statement show that parsing "... X and Y ..." as S-coordination is easier when the discourse context makes it clear that X and Y do not belong together. This supports the notion of an interactive parser that is immediately affected by contextual knowledge.

That the expected effect was found not on the disambiguating verb, but one word later, on the adverb, can be explained by the fact that using a moving window self-paced reading design can cause some delayed effects, so called spill-over effects (see e.g. Hoeks et al. 2002).

General Discussion

In this paper, three experiments were presented that investigated the influence of contextual knowledge on the processing of the NP- versus S-coordination ambiguity. This knowledge was provided through the discourse context and consisted of knowledge about the interpersonal relations between the three central figures of the text. Contextual knowledge was manipulated to either bias towards interpreting the sentence as an NP-coordination or as an S-coordination.

The results of Experiment 1 showed that readers continue a sentence in the way that is most plausible according to their contextual knowledge. Furthermore, Experiment 2 showed that readers find S-coordination easier, more plausible and more natural if their contextual knowledge supports S-coordination than if it supports NP-coordination. Finally, Experiment 3 showed that processing an S-coordination is easier when it is embedded in an S-coordination-supportive context than when it is embedded in an NP-coordination-supportive context. Thus, the interpersonal relationships in the described situation appeared to have an immediate impact on the on-line analysis of sentences in which these figures appeared.

Our results highlight the flexible, adaptive nature of sentence processing: based on the results we can claim with some certainty that the parser changes its preferences based on the contents of the previous discourse context if it changes in a critical way. In other words, this is an adjustment of the idea of a parsing process that always makes decisions following the same (syntax-based) principles.

A second important result is that contextual knowledge not only appeared to be of greater importance than syntactic principles, but also than topic-structure. As mentioned before, all our experimental texts had a one-topic structure: the character that was introduced first, stayed in the center of attention throughout the part of the text that preceded the target sentence. On the basis of the results by Hoeks et al. (2002) this would lead to a preference for NP-coordination, since this structure requires only one topic. We have shown how contextual knowledge overrules this preference for NP-coordination. Therefore, whereas Hoeks et al. conclude that the influence of topic-structure overrules syntactic principles like the Minimal Attachment strategy (Frazier 1987), we can go one step further and conclude that in our study contextual knowledge overruled not only syntactic principles, but also pragmatic principles concerning the topic structure of the text.

In conclusion, the role of world knowledge provided through the discourse context is crucial and, most importantly, overrules the principle of economy that is based upon syntactic or discourse factors (topic-structure). This is an adjustment of the idea of a parsing process that always makes decisions following the same (syntax-based) principles.

Future research

We are in the process of constructing several experiments using eye-tracking in which different constraints (topic-structure and contextual knowledge) are manipulated at different points within the sentence and discourse context. Also, possible effects are not only measured at the disambiguating word, but within the ambiguous part of the sentence as well (cf. MacDonald 1994). The aim of these experiments is to gain a more detailed insight in the time course during which factors like contextual knowledge and topic-structure are used.

- Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, 30, 191-238.
- Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological syntactic processor. In D. R. Dowty, L. Karttunen, & A. M. Zwicky (Eds.), *Natural language parsing* (pp. 320-358). Cambridge University Press, Cambridge.
- Frazier, L. (1987a). Sentence processing: A tutorial review. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp.601-681). Hillsdale, NJ: Erlbaum.
- Hagoort, P., Hald, L., Bastiaansen, M.C.M., & Petersson, K.M. (2004). Integration of Word Meaning and World Knowledge in Language Comprehension. *Science*, 304 (5669), 438-440.
- Hoeks, J. C. J., Vonk W., & Schriefers, H. (2002). Processing coordinated structures in context: The effect of topic-structure on ambiguity resolution. *Journal of Memory and Language*, 46, 99-119.
- MacDonald, M. C. (1994). Probabilistic constraints and syntactic ambiguity resolution. *Language and cognitive processes*, 9, 157-201.
- Ni, W., Crain, S., & Shankweiler, D. (1996). Sidestepping garden paths: Assessing the contributions of syntax, semantics and plausibility in resolving ambiguities.
- Van Berkum, J. J. A., Brown, C. M., & Hagoort, P. (1999). Early referential context effects in sentence processing: Evidence from event-related brain potentials. *Journal of Memory and Language*, 41, 147-182.
- Versteeg, N., Sanders, T.J.M., & Wijnen, F.N.K. (submitted for publication). Knowing what's going on. The influence of contextual knowledge on the on-line resolution of structural ambiguity.