Cognitive load does not decrease pronoun use when speaker's and addressee's perspectives are dissociated

Jorrig Vogels (j.vogels@uvt.nl)

Emiel Krahmer (e.j.krahmer@uvt.nl)

Alfons Maes (maes@uvt.nl)
Tilburg center for Cognition and Communication, PO Box 90153
5000 LE Tilburg, the Netherlands

Abstract

It has been suggested that a referent's accessibility is affected by the degree to which it is in the speaker's attention. Assuming that less accessible referents are less likely to be pronominalized, this predicts that speakers under cognitive load use more elaborate referring expressions. However, speakers under load may also have difficulty taking into account their addressee's perspective, which may either lead to more use of the speaker's own discourse model or to more economic expressions. To tease these effects apart, we conducted a story completion experiment in which cognitive load was manipulated by the presence or absence of a secondary task for the speaker. In addition, we dissociated the speaker's and the addressee's perspectives. Our results do not provide evidence for the hypothesis that cognitive load reduces the accessibility of referents in the speaker's own discourse model, suggesting that speaker attention does not determine accessibility.

Keywords: cognitive load; referring expressions; accessibility; perspective taking

Introduction

When speakers refer to something, they have to choose a certain type of referring expression, such as a definite description (e.g. *the girl*) or a pronoun (e.g. *she*). Traditionally, the speaker's choice of a referring expression has been assumed to be tailored for the addressee (e.g. Ariel, 1990; Gundel, Hedberg, & Zacharski, 1993). According to this view, speakers make assumptions about the cognitive status of the referent in the mind of the addressee. An important factor in determining this status is the salience of the referent in the discourse. For example, if the referent was the topic of the preceding sentence, it can be assumed to be highly accessible in the addressee's discourse model, and therefore it does not need an elaborate description to be reactivated. Because the addressee knows that the speaker would have used a more elaborate expression if she had a less activated referent in mind, the use of an attenuated expression, such as a pronoun, aids the addressee's interpretation.

However, studies that have manipulated speakers' attention resources suggest that the activation of mental representations in the speaker's own memory is also important for the choice of referring expression. For example, Arnold and Griffin (2007) and Fukumura, Van Gompel, and Pickering (2010) varied the number of possible referents in the discourse, and found that speakers used fewer pronouns when a referential competitor was present, even though the referent was salient in the discourse (i.e. topical). In addition, speakers have been found to choose fewer attenuated expressions for salient referents when they are distracted by another task (Rosa & Arnold, 2011). These findings suggest that the choice of referring expression is affected by the degree to which the referent is in the speaker's attention: referent accessibility, and hence pronoun use, decreases when the speaker has to spread attentional resources over multiple possible referents or multiple tasks.

If restrictions on speakers' attention resources influence the accessibility of referents in their own memory, speakers experiencing an increased cognitive load should be more likely to use elaborate expressions such as full noun phrases. However, increased cognitive load may also affect the degree to which speakers are able to take into account the perspective of the addressee (e.g. Epley, Keysar, Van Boven, & Gilovich, 2004; Horton & Keysar, 1996). For example, Horton and Keysar (1996) showed that when under time pressure, speakers were not taking into account the addressee's perspective in choosing whether or not to include an adjective in their referring expressions. Thus, cognitive load may make the choice of referring expression more egocentric. On the one hand, this could mean that this choice is based more on the speaker's own discourse model (e.g. Fukumura & Van Gompel, 2012). That is, when speakers are under load, they might be less able to calculate whether the referent is accessible for the addressee or not. Because in many cases speaker and addressee have access to the same discourse information, it is generally difficult to distinguish between a referring expression that is tailored for the addressee and a referring expression that is based on the speaker's own model of the discourse when all discourse information is shared. For example, a referent that is highly accessible in the addressee's discourse model is often also highly accessible in the speaker's discourse model. When the speaker's and the addressee's perspectives differ, however, speakers under load might be inclined to use pronouns when the referent is salient in their own discourse model but not salient in the addressee's discourse model. Conversely, they might be inclined to use full noun phrases when the referent is not salient in their own discourse model but salient in the addressee's discourse model.

On the other hand, speakers under load may fail to take into account any information about discourse salience, and resort to using more pronouns in general, because these are short, have little semantic content, and are hence easy to produce (Almor, 1999; Burzio, 1998). For example, studies on children and elderly people have found that having limited working memory capacity increases the use of pronouns in contexts in which the referent is not salient for the addressee (e.g. references following a topic shift; Hendriks, Englert, Wubs, & Hoeks, 2008; Wubs, Hendriks, Hoeks, & Koster, 2009; see also Almor et al., 1999).

To tease these possible effects of cognitive load apart, the linguistic salience of the referent should be varied, because the effect of cognitive load might be different for salient and non-salient referents. In addition, to determine whether cognitive load affects the speaker's discourse model or the speaker's assumptions about the addressee's discourse model, the speaker's and addressee's perspectives should be dissociated. Therefore, we conducted a story completion experiment in which we manipulated perspective, referent salience and cognitive load. Perspective was manipulated by presenting one of the sentences of the story only to the speaker, over headphones (cf. Fukumura and Van Gompel, 2012). The referent was considered linguistically salient for the speaker when it was mentioned in this privileged sentence, and not salient when it was only mentioned in the introductory sentence. Since the addressee did not hear the privileged sentence, referent salience was reversed for the addressee. Cognitive load was manipulated by giving the speaker a second, unrelated (but also verbal) task in the first or the second half of the experiment.

We hypothesized that if cognitive load lowers the accessibility of referents in the speaker's discourse model, it should decrease pronoun use, irrespective of the referent's linguistic salience. If, on the other hand, cognitive load makes perspective taking more difficult, there should be a stronger tendency to use the speaker's own discourse model in the dual task condition. That is, speakers should be more likely to use pronouns when the referent is salient in their own discourse model (and not in their addressee's), than when the referent is salient in the addressee's discourse model (and not salient in their own). Alternatively, speakers may tend to use more pronouns in general when under increased cognitive load.

Methods

Participants

Sixty-four students (47 female; mean age: 20.2 years) from Tilburg University participated in the experiment for course credit. Half of them acted as speakers, the others acted as addressees. All were native speakers of Dutch, the language of the experiment.

Materials

The experimental items consisted of 16 pairs of photographs, taken from Vogels, Krahmer, and Maes (in press), accompanied by two introductory sentences and the onset of a third sentence. The first picture of a pair always showed one male and one female person sitting next to each other. In the second picture, one of these persons performed an action, such as walking away or getting a glass of water. This person will be referred to as the target character, as participants were expected to refer to this character in their continuations. There were two versions of each picture pair; one in which it was the male person and one in which it was the female person that performed the action. An example of a picture pair is shown in Figure 1.

Figure 1: Example of a stimulus item in two conditions. Sentence 1 was read aloud by the speaker; sentence 2 was presented only to the speaker over headphones. Context sentences are translations of the Dutch originals.

The first sentence introduced both characters with indefinite noun phrases, either *een meisje* "a girl" and *een jongen* "a boy" or *een vrouw* "a woman" and *een man* "a man". One was mentioned as the subject, and the other in a

prepositional phrase (e.g. *Een meisje zat te discussiëren met een jongen* "A girl was arguing with a boy"). This sentence was read aloud by the speaker to the addressee. The second sentence described some emotional or physical state of the person mentioned in the prepositional phrase (e.g. *De jongen raakte enorm gepikeerd* "The boy got really annoyed"). This sentence was prerecorded by a female native speaker of Dutch, and only heard by the speaker over headphones. The onset of the third sentence was always *Vervolgens...* "Subsequently...", serving as a cue for the speaker to complete the story.

In addition, 20 picture pairs served as fillers. These differed from the experimental items in that some showed either two male or two female characters or only one character. In the accompanying sentences, some characters were given labels such as *een verkoopster* "a saleswoman" or *een Duitser* "a German", and sometimes the same character was the subject of both introductory sentences. An additional 4 items were included as practice items.

Procedure

The experiment took place in an experimental room. Two participants were randomly assigned to the role of speaker and addressee. The participant taking the role of speaker was seated at one end of a table, behind a laptop connected to a serial response box, and was wearing headphones. The participant taking the role of addressee was seated at the other end of the table, and was given a booklet containing all different picture pairs and an answer sheet. The experiment was run on the laptop using E-Prime 2.0, and was only visible to the speaker. The speaker's task was to complete the stories depicted by the picture pairs for the addressee. In one half of the experiment, the speaker received a secondary task (cognitive load condition), while there was no secondary task in the other half (no cognitive load). In the no cognitive load condition, each trial started with the item number presented on the screen, accompanied by a 500 ms beep, followed by a cross-hair. Then, the first picture of a pair appeared on the left side of the screen. After 3 s the first introductory sentence appeared below the picture in a red font. The speaker read this sentence aloud to the addressee. After 5 s the second sentence was presented to the speaker over the headphones. Next, while the first picture remained visible, the second picture appeared automatically on the right side of the screen, together with the onset of the third sentence, which also appeared below the picture in a red font. The target character in the second picture was either the subject of the first introductory sentence, and therefore linguistically salient for the addressee (addressee-salient condition), or the subject of the second sentence, which was only presented to the speaker (speaker-salient condition). At this time, recording started, and the speaker completed the sentence based on the event shown in the picture, by saying it aloud to the addressee. After 6 s, recording stopped and the pictures and sentences disappeared. The addressee's task was to select the correct picture pair out of three options from the booklet and mark the correct answer on the sheet. The addressee gave the speaker a hint when the next trial could be started.

In the cognitive load condition, the appearance of the first picture was preceded by the words BAL "ball" or DAL "valley" (Goudbeek & Krahmer, 2011), which was presented in the middle of the screen for 1 s. The same happened at the end of the trial, followed by the question *Was dit woord hetzelfde als het vorige woord? (Ja/Nee)* "Was this word the same as the previous word? (Yes/No)". The speaker then pressed either the green/Yes or the red/No button on the response box. They did not receive feedback on their answers.

The participants received instructions both orally and in written form. Speakers were explicitly told that the sentence presented over headphones could not be heard by their addressee, but that they had to pay attention to it nonetheless, since they would be asked about these sentences after the experiment as an attention check (which was indeed the case). They were also encouraged to pay attention to the dual task by way of a prize offer for the participant with the fewest errors. To keep the speaker aware of the addressee's needs, the addressee was allowed to ask the speaker clarification questions if anything remained unclear, but only after the speaker had finished the story.

The experiment was divided into two blocks, of which one contained the dual task and the other did not, counterbalanced for order. Each block was preceded by two practice items. The experimenter was only present during the instructions and the practice trials. The experiment took about 25 minutes.

Data coding

From all speakers' continuations of the third sentence, we selected the first subject reference, which was expected to refer to the target character. Any further references (e.g. in another follow-up clause) were ignored. We excluded 33 cases in which the first subject did not refer to the target referent, 7 plural references, 3 indefinite references, 1 case in which the sentence presented over the headphones was repeated literally, and 1 missing case. In addition, there were 2 cases in which the referring expression was repaired. However, because the repair was of the same type in both cases (e.g. 'the man... uh the boy'), we kept these cases. In total, we excluded 45 trials (8.8%). The remaining 467 subject references were coded for the type of referring expression: either full noun phrase or pronoun.

Design and statistical analyses

Crossing the two factors Referent salience and Cognitive load resulted in a 2 (speaker-salient, addressee-salient) x 2 (cognitive load, no cognitive load) within-participants design. Participants were assigned to one of four lists, each of which contained one version of a given item. The items were presented in a pseudo-random order, with at least one filler item between two consecutive experimental items.

We performed a logit mixed model analysis on the log odds for a pronoun (Jaeger, 2008). Referent salience and Cognitive load were included as fixed factors, and participants and items as random factors. The fixed factors were centered to reduce collinearity. Starting with a full random effect structure, we excluded random slopes that did not significantly contribute to the model fit. Only the final model will be reported.

Results

Upon inquiry, only 6 participants reported that they found the secondary task difficult. Still, the overall error rate was 9.7%, suggesting that participants might have been overestimating their performance. Few errors were made in the attention check following the experiment (1.7%), suggesting that speakers were attending to the sentences presented over the headphones.

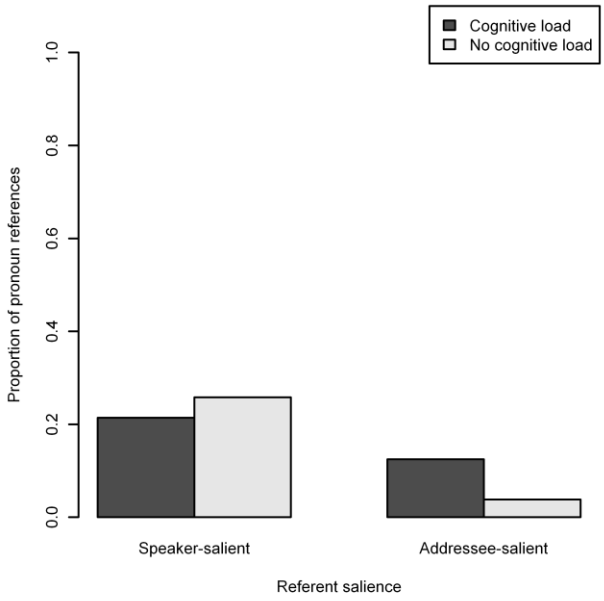

Figure 2: Proportion of pronoun references to the target character in the four conditions.

Figure 2 shows the proportion of pronoun references to the target character by referent salience and cognitive load condition. We found a main effect of referent salience: pronouns were more frequent when the referent was salient only for the speaker (23.6%) than when it was salient only for the addressee (8.3%), β = -2.26; SE = 0.85; $p < .01$. There was also a main effect of cognitive load: more pronouns were used when speakers performed a dual task (17.2%) than when they did not (15.7%), β = 1.40; SE = 0.56; $p < .05$. These effects were qualified by a significant interaction, β = 2.85; SE = 0.95; $p < .01$, suggesting that cognitive load only increased pronoun use in the addressee-salient condition. The model included random intercepts for participants (s^2 = 2.85) and items (s^2 = 0.15), as well as by-subject random slopes for referent salience (s^2 = 12.06) and cognitive load (s^2 = 3.08). This suggests that participants varied substantively in the way their pronoun use was affected by the context sentences and the dual task. The contribution of the random slope for cognitive load to the model fit was only marginally significant (p = .06). Removing it decreased the effect size of the fixed factors, but the interaction effect remained significant at the α = .05 level.

Discussion

Speakers used more pronouns when the referent was salient for them (but not salient for their addressee) than when it was not salient for them (but salient for their addressee). This suggests that speakers chose referring expressions more according to the referent's accessibility in their own discourse model than according to assumptions about the referent's accessibility in their addressee's discourse model. This is in line with Fukumura and Van Gompel (2012), who found that speakers were not taking into account their addressee's perspective in choosing referring expressions when the two perspectives were dissociated.

If cognitive load decreases the accessibility of referents in the speaker's own discourse model, as suggested by Arnold and Griffin (2007), we should have seen fewer pronouns across the board when speakers performed the dual task. This is not what we found. Instead, speakers were somewhat more likely to use pronouns when under load, at least in the addressee-salient condition. This finding is also not in line with the hypothesis that cognitive load makes it harder to calculate the referent's accessibility in the addressee's discourse model. If that were the case, the dual task condition should have increased the tendency to use the speaker's own discourse model. That is, speakers under load should have been more likely to use pronouns in the speaker-salient condition, and less likely to use pronouns in the addressee-salient condition.

Our results seem compatible with the hypothesis that speakers under load are more likely to use less costly expressions, such as pronouns (Almor, 1999; Burzio, 1998). That is, when distracted by a secondary task, speakers have fewer memory resources available that are needed to infer that less salient referents should be referred to with more elaborate expressions. The fact that cognitive load only increased pronoun use in the addressee-salient condition may be due to the preference to use pronouns anyway when the referent is salient for the speaker.

An alternative explanation for the effect of cognitive load is that speakers are less able to keep track of their own discourse model when they are under load. This would cause their use of referring expressions to become less consistent, i.e. the choice of referring expressions becomes less tied to the discourse context (Arnold, 2010). This would explain the finding that the difference between the speaker-salient and the addressee-salient condition becomes smaller under load. In addition, this explanation would not be incompatible with the finding by Arnold and Griffin (2007), Fukumura et al. (2010) and Rosa and Arnold (2011) that cognitive load leads to fewer attenuated expressions, since

these studies investigated only contexts in which referents were always linguistically salient. Thus, the present study at least stresses that dissociating the salience of the referent in the speaker's and the addressee's perspective is necessary to tease the possible effects of cognitive load apart, i.e. whether it affects the speaker's representation of the discourse or the speaker's assumptions about the addressee's representation of the discourse.

Although our results suggest that speakers tend to use their own discourse model, replicating Fukumura and Van Gompel's (2012) findings, it is striking that the overall proportion of pronouns is quite low. Even in the no cognitive load, speaker-salient condition, in which one would expect many pronouns if speakers only took into account their own discourse model, the percentage of pronouns out of all referring expressions did not exceed 30%. This relatively low proportion of pronouns might be due to the manipulation of perspective: Perhaps speakers *were* taking into account the addressee's informational needs, but not up till the level of calculating the referent's cognitive status for the addressee. This kind of detailed audience design might be cognitively too costly (e.g. Brennan & Hanna, 2009; Horton & Gerrig, 2005). Therefore, speakers may just have increased the use of elaborate expressions to be as clear as possible for the addressee, as soon as they were aware of the fact that not all information was shared. Fukumura and Van Gompel (2012) found evidence for such a minimal, one-bit model of audience design (e.g. *the addressee has heard this or not*; Galati & Brennan, 2010; see also Epley et al., 2004) by comparing their condition with privileged information for the speaker to a condition in which all information was shared. They found more pronouns in the shared condition, independently of whether the referent was salient or not. This suggests that speakers use more elaborate expressions when there is privileged information, even though they might run the risk of being overly specific.

Still, Fukumura and Van Gompel (2012) found somewhat higher rates of pronoun use in their privileged, referent-salient condition (Experiment 1: 37%; Experiment 2: 48%) than we did in our experiment. This difference could be due to differences in the linguistic materials. Firstly, while the referent mentioned in the second context sentence (which was only presented to the speaker) was referred to with a pronoun in Fukumura and Van Gompel's experiments, it was referred to with a full NP in our experiment, in accordance with the preferred way of referring to an entity previously mentioned as a direct object in centering theory (e.g. Brennan, 1995). The tendency to pronominalize the entity on a subsequent reference may however be stronger when the referent had already been pronominalized. Secondly, speakers may be more likely to reuse the most recent referring expression, which could also have led to more pronouns in Fukumura and Van Gompel's experiments than in ours.

Even though employing a one-bit model of audience design is probably less cognitively demanding than calculating the referent's accessibility for the addressee, using full NPs to aid the addressee when there is privileged information might be more difficult under load. Hence, if speakers in our study were employing such minimal audience design, our finding that they were more likely to use pronouns under load could be due to difficulties in assessing that the addressee might need more specific information. If this is the case, we would predict that in a situation in which all discourse information is shared, cognitive load does not increase pronoun use, since in that case there is no need to be more specific for the addressee.

One reason why speakers did not make the extra effort to calculate the accessibility of the referent in the addressee's discourse model may be that in the current experiment, as well as in Fukumura and Van Gompel's, references were never ambiguous, since the two characters always had a different gender. Therefore, not taking into account the addressee's perspective would probably not result in interpretation errors. However, in case not taking into account the addressee's perspective would lead to interpretation errors, speakers may base their choice of referring expressions on the discourse model of their addressee (e.g. Fukumura & Van Gompel, 2012; Horton & Keysar, 1996). In that case, increased cognitive load might make this perspective taking more difficult, and cause speakers to fall back on their own discourse model.

In our filler materials, which contained stories with characters of the same gender, and hence pronouns were ambiguous, we indeed found more pronouns when the referent was salient for the addressee but not for the speaker (33%; n = 51) than when the referent was salient for the speaker but not for the addressee (13%; n = 56), suggesting that speakers were taking their addressee's perspective into account. However, cognitive load did not seem to cause speakers to use their own discourse model. Rather, a pattern similar to that in Figure 2 emerged, with more pronouns under load for referents that were not salient for the speaker. This might be an indication that cognitive load as manipulated here is independent of perspective taking.

Our results suggest that there was quite some individual variation as to how speakers' referring expressions were affected by the dual task. One cause of individual differences could be the use of strategies for remembering the words BAL and DAL. Two thirds of all participants reported to have used some kind of mnemonic (e.g. putting up one finger for BAL and two for DAL), although these were not always employed from the beginning. This is a concern that should be taken up by future studies. Nevertheless, the general trend of more pronouns under load in the addressee-salient condition in the present study seems to hold for all participants.

Finally, an important issue is that cognitive load can be manipulated in different ways that may affect the choice of referring expressions differently. For example, it is not clear whether dual tasks, divided attention to multiple referents, and restricted working memory capacity all produce the same kind of cognitive load. In addition, language

production may be affected differently by verbal or visual secondary tasks (Baddeley and Hitch, 1974; Kellogg et al., 2007). In the present experiment, the use of a verbal secondary task may have especially hindered the production of elaborate linguistic forms. Other manipulations, such as adding time pressure, possibly interfere more with activating non-linguistic representations or with perspective taking. This is an issue that needs further research.

In sum, the present study has shown that speakers use more pronouns when they experience increased cognitive load, at least when the referent is not salient for the speaker (but is salient for the addressee). Whether this is due to a general preference to produce economic forms, or to difficulties in keeping track of the accessibility of the referents in the discourse model should be researched further. However, we have not found support for the claim that cognitive load, at least in the form of the dual task used here, decreases the accessibility of referents in the speaker's discourse model. Hence, although accessibility may be related to attention, it probably does not hold generally that less attentive speakers use more elaborate referring expressions.

References

Almor, A. (1999). Noun-phrase anaphora and focus: The informational load hypothesis. *Psychological Review, 106*(4), 748-765.

Almor, A., Kempler, D., MacDonald, M. C., Andersen, E. S., & Tyler, L. K. (1999). Why Do Alzheimer Patients Have Difficulty with Pronouns? Working Memory, Semantics, and Reference in Comprehension and Production in Alzheimer's Disease. *Brain and Language, 67*, 202-227.

Ariel, M. (1990). *Accessing noun-phrase antecedents*. London: Routledge.

Arnold, J. E., & Griffin, Z. (2007). The effect of additional characters on choice of referring expression: Everyone counts. *Journal of Memory and Language, 56*, 521-536.

Arnold, J. E. (2010). How Speakers Refer: The Role of Accessibility. *Language and Linguistics Compass*, 4(4), 187-203.

Baddeley, A. D., & Hitch, G. (1974). Working Memory. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation* (Vol. 8). New York/London: Academic Press.

Brennan, S. E. (1995). Centering Attention in Discourse. *Language and Cognitive Processes*, 10(2), 137-167.

Brennan, S. E., & Hanna, J. E. (2009). Partner-Specific Adaptation in Dialog. *Topics in Cognitive Science, 1*, 274–291.

Burzio, L. (1998). Anaphora and Soft Constraints. In P. B. D. Fox, P. Hagstrom, M. McGinnis & D. Pesetsky (Eds.), *Is the Best Good Enough?: Optimality and Competition in Syntax*. Boston: MIT.

Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective Taking as Egocentric Anchoring and Adjustment. *Journal of Personality and Social Psychology, 87*(3), 327-339.

Fukumura, K., & Van Gompel, R. (2012). Producing Pronouns and Definite Noun Phrases: Do Speakers Use the Addressee's Discourse Model? *Cognitive Science, 36*(7), 1289-1311.

Fukumura, K., Van Gompel, R., & Pickering, M. J. (2010). The use of visual context during the production of referring expressions. *The Quarterly Journal of Experimental Psychology, 63*(9), 1700-1715.

Galati, A., & Brennan, S. E. (2010). Attenuating information in spoken communication: For the speaker, or for the addressee? *Journal of Memory and Language, 62*(1), 35-51.

Goudbeek, M., & Krahmer, E. (2011). Referring under load: Disentangling preference-based and alignment-based content selection processes in referring expression generation *Proceedings of PRE-Cogsci: Bridging the gap between computational, empirical and theoretical approaches to reference*. Boston.

Gundel, J. K., Hedberg, N., & Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language, 69*, 274-307.

Hendriks, P., Englert, C., Wubs, E., & Hoeks, J. (2008). Age Differences in Adults' Use of Referring Expressions. *Journal of Logic, Language and Information, 17*, 443-466.

Horton, W. S., & Gerrig, R. J. (2005). The impact of memory demands on audience design during language production. *Cognition, 96*, 127–142.

Horton, W. S., & Keysar, B. (1996). When do speakers take into account common ground? *Cognition, 59*, 91-117.

Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, 59, 434-446.

Kellogg, R. T., Olive, T., & Piolat, A. (2007). Verbal, visual, and spatial working memory in written language production. *Acta Psychologica, 124*, 382–397.

Rosa, E. C., & Arnold, J. E. (2011). The Role of Attention in Choice of Referring Expression *Proceedings of PRE-Cogsci: Bridging the gap between computational, empirical and theoretical approaches to reference*. Boston.

Vogels, J., Krahmer, E., & Maes, A. (in press). Who is where referred to how, and why? The influence of visual saliency on referent accessibility in spoken language production. *Language and Cognitive Processes*.

Wubs, E., Hendriks, P., Hoeks, J., & Koster, C. (2009). Tell me a story! Children's capacity for topic shift *Proceedings of the 3rd Conference on Generative Approaches to Language Acquisition North America (GALANA 3)* (pp. 313-324). Somerville, MA: Cascadilla Press.