Misspoken words affect the perception and retrieval of intended words

Article *in* Language, Cognition and Neuroscience · August 2020

DOI: 10.1080/23273798.2020.1802052

3 authors, including:

Hossein Karimi
Mississippi State University
15 PUBLICATIONS **220** CITATIONS

SEE PROFILE

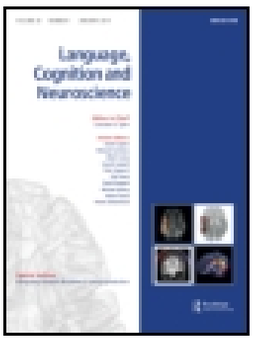

Routledge
Taylor & Francis Group

Language, Cognition and Neuroscience

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/plcp21

Misspoken words affect the perception and retrieval of intended words

Hossein Karimi , Michele Diaz & Fernanda Ferreira

To cite this article: Hossein Karimi , Michele Diaz & Fernanda Ferreira (2020): Misspoken words affect the perception and retrieval of intended words, Language, Cognition and Neuroscience, DOI: 10.1080/23273798.2020.1802052

To link to this article: https://doi.org/10.1080/23273798.2020.1802052

LANGUAGE, COGNITION AND NEUROSCIENCE
https://doi.org/10.1080/23273798.2020.1802052

Routledge
Taylor & Francis Group

REGULAR ARTICLE

Misspoken words affect the perception and retrieval of intended words

Hossein Karimi[a], Michele Diaz[b] and Fernanda Ferreira[c]

[a]Department of Psychology, Mississippi State University, MS , USA; [b]Department of Psychology, The Pennsylvania State University, PA , USA; [c]Department of Psychology, University of California, Davis, CA, USA

ABSTRACT

Everyday speech contains disfluencies, including unintentionally spoken words. What is the fate of the misspoken word in the comprehender's memory? In this study, we asked whether: (1) the gender of misspoken words lingers and affects how the intended word is perceived, and (2) whether and how lingering representations can cause interference during the retrieval of the intended word. In two experiments, participants provided spoken responses to given prompts. In Experiment 1, participants used masculine or feminine pronouns to refer to gender-neutral words (*passenger*) depending on the gender of a preceding misspoken word (*pilot* or *stewardess*), suggesting that the gender of reparanda lingers. Experiment 2 showed that the presence of a misspoken word resulted in a reduction of pronominal reference to the intended word, suggesting that the misspoken word causes interference when the intended word is being retrieved by functioning as an additional discourse entity.

ARTICLE HISTORY
Received 5 September 2019
Accepted 21 July 2020

KEYWORDS
Disfluency; Lingering memory representations; Referential processing; Good Enough language processing; Perception; Encoding; Retrieval

Introduction

Human speech is not perfect. It is estimated that 6–10% of daily speech that we encounter contains some sort of disfluency, including repetitions (e.g. *The … the ladle*), filled pauses (*uh*, *um … the ladle*) and self-repairs (*The bowl, uh I mean, the ladle*; Bortfeld et al., 2001; Fox Tree, 1995). Disfluencies might seem like nothing more than nuisance material that should be ignored or suppressed by comprehenders. However, numerous studies have shown that disfluencies do in fact influence the processing of subsequent linguistic information (e.g. Arnold et al., 2003, 2007b, 2004; Bailey & Ferreira, 2003; Barr & Seyfeddinipur, 2010; Brennan & Schober, 2001; Christianson et al., 2001; Corley, 2010; Corley et al., 2007; Engelhardt et al., 2010; Ferreira et al., 2004; Kidd et al., 2011; Lau & Ferreira, 2005; Lowder & Ferreira, 2016; MacGregor et al., 2010; Maxfield et al., 2009).

Despite the large number of studies investigating the role of disfluencies in language processing, little research has been carried out to understand the effect of disfluencies on subsequent language production in general, and reference production in particular. In this study, we focus on self-repairs as a way to investigate whether the gender of misspoken words lingers in memory and, also whether lingering representations could influence subsequent reference to the intended word. In a self-repair disfluency, a word that has been unintentionally spoken (i.e. the misspoken word) is replaced with the word that was intended. For example, in *The little girl handed the bowl, uh I mean, the ladle to her mother*, the noun phrase *bowl* would be the misspoken word and *ladle* would be the intended word. We will henceforth refer to the misspoken word as the "reparandum" and to the intended word as the "repair", as this is the established terminology used in disfluency research (Levelt & Cutler, 1983; Lowder & Ferreira, 2016; Karimi et al., 2019). Self-repair disfluencies sometimes contain an "editing term" indicating that the word just spoken was not intended and is going to be corrected. In the example above, *uh I mean* would be the editing term.

One interesting question regarding the processing of self-repair disfluencies is the fate of the reparandum in the comprehender's representation of the utterance. Specifically, is the reparandum erased from memory, or does it (or at least some features of it) linger? If the reparandum is completely deleted from the memory representation of the sentence or discourse, it should have little or no influence on how subsequent linguistic material is processed. In contrast, if it lingers, later processing should be affected by its presence (and its properties). In support of the second view, Ferreira et al. (2004) and Lau and Ferreira (2005) employed grammaticality judgement tasks and demonstrated that the syntactic properties of a reparandum linger in working memory and influence how the repair is processed.

CONTACT Hossein Karimi hk702@msstate.edu
Supplemental data for this article can be accessed at https://doi.org/10.1080/23273798.2020.1802052

However, these results leave some important questions unanswered: (1) Do semantic properties of misspoken words such as gender also linger, and affect how the repair is perceived in terms of its gender? (2) Could lingering representations cause interference during the retrieval of the repair? And if yes, how does such interference take place? In addition, it is possible that because grammaticality judgment tasks emphasise well-formedness, any inconsistency in syntactic features from reparandum to repair is more likely to be noticed by participants (even at a conscious level), increasing the probability that the features of the reparandum remain in memory. To the best of our knowledge, no previous studies have made use of an implicit processing task. Below, we discuss how the form of referring expressions during language production serves as an ideal tool to examine these questions.

Referential processing as a measure to investigate the accessibility of repairs

Referential processing occurs when speakers use a linguistic form to refer to entities from a previous part of the sentence or discourse. For example, after hearing *Susan went shopping with John yesterday*, reference could be made to either *Susan* or *John*, and the referring expression could be either a pronoun (*he* or *she*) or a repeated noun (*Susan* or *John*). Crucially, making reference necessarily involves retrieving the memory representation associated with the referent (i.e. the word that is referred to; Lucas et al., 1990; MacDonald & MacWhinney, 1990), and therefore provides a reasonable testing ground for subsequent retrieval of previously encoded words.

A number of findings regarding the production of different forms of referring expressions are relevant to the present study and we will therefore briefly discuss them here. Perhaps the most important finding pertains to the relationship between ease of memory retrieval (i.e. accessibility)[1] and form of reference. Numerous studies have shown that referential candidates that are highly accessible in memory tend to be referred to with pronouns (e.g. *she*) rather than repeated nouns (e.g. Susan, Arnold, 2001; Brennan, 1995; Fletcher, 1984; Fukumura & Van Gompel, 2010a, 2011; Gordon et al., 1993; Grosz et al., 1995; Gundel et al., 1993; Karimi et al., 2014; Stevenson et al., 1994). For example, replicating numerous prior findings, Fukumura and Van Gompel (2010a) showed that the syntactic subject of a preceding sentence is significantly more likely to be realised with a pronoun rather than a repeated noun. Moreover, Karimi et al. (2014) found that speakers are more likely to refer to a referential candidate with a pronoun rather than a repeated noun if that referential candidate is semantically richer than a competitor. Fukumura and Van Gompel (2011) additionally found that animate referential candidates are more likely to be referred to with a pronoun rather than a repeated noun compared with inanimate referential candidates. Since syntactic prominence, semantic richness, and animacy have all been shown to increase accessibility in memory (*Syntactic prominence*: e.g. Bock & Warren, 1985; Grosz et al., 1995; Gundel et al., 1993, *Semantic richness*: e.g. Hofmeister, 2011; Hofmeister & Vasishth, 2014; Troyer et al., 2016; see also Almor, 1999; 2004, *Animcay*: e.g. Bock, 1982; Bock et al., 1992; Branigan et al., 2008; Keil, 1979), these results could be interpreted as a general tendency to use attenuated referring expressions to refer to highly activated referential candidates.

Another related finding is that people tend to use fewer pronouns than repeated nouns to refer to either of two referential candidates that have the same gender rather than different genders, an effect that is attributed to semantic interference between the potential referents (Arnold & Griffin, 2007a; Fukumura et al., 2013). We will henceforth refer to this effect as the "semantic interference effect". Semantic interference is assumed to reduce memory activation and therefore pronominal reference. For instance, Arnold and Griffin (2007a) presented participants with cartoon panels containing two characters who were either of the same gender or who had different genders. The participants then had to describe a following cartoon panel in which the referential characters were involved in a different action. The results showed that pronominal reference to either of the referential characters dropped as a function of gender-congruence between the two referential candidates. These results were later replicated by Fukumura et al. (2013) in Finnish where pronouns are not gender-marked, suggesting that the semantic interference effect primarily occurs at a conceptual (rather than at the linguistic) level.

A final related finding is that the presence of additional referential characters leads to fewer pronominal references to the referent regardless of gender congruence (Arnold & Griffin, 2007a; Fukumura et al., 2010b). For instance, previous studies have demonstrated that the presence of linguistic (Arnold & Griffin, 2007a) or visual (Fukumura et al., 2010b) referential competitors results in less pronominal reference to the referent. This finding is explained by arguing that when there is an additional referential character, the overall memory activation is distributed over more referential candidates and thus less activation is necessarily left for the referent, which in turn leads to a reduction in pronominal reference (see Arnold & Griffin, 2007a). We will henceforth refer to this effect as the "additional character effect".

Thus, form of reference to referential candidates can be used as a measure to assess the accessibility (i.e. ease of retrieval) of memory representations. More activated memory representations are retrieved from memory more easily, which then manifests itself as more attenuated referring expressions (such as pronouns) for the associated representations. Conversely, the accessibility (i.e. retrieval ease) of less activated representations is low, which increases the probability of using more elaborated referring expressions (such as repeated nouns) to refer to the associated representations. This inverse relationship between memory activation and forms of referring expressions might be because highly activated referential candidates need less linguistic signal to be retrieved a second time by comprehenders, and vice versa (Gordon et al., 1993; see also Jaeger, 2010).

The present study

Using two sentence continuation experiments, this study aims to answer two questions: (1) Does the gender of misspoken words linger, and affect how the repair is perceived? (2) Does the lingering representation of the reparandum cause interference during the retrieval of the repair? And if yes, what is the mechanism underlying such interference?

In both experiments, participants heard prompt sentences containing a reparandum and a repair and then provided a spoken response in which they could make reference to the repair. Thus, the task necessarily involved an initial encoding of both the reparandum and the repair into working memory during the processing of the prompt sentences, and retrieval of the representation associated with the repair during the formulation of the response (i.e. if and when reference was made to the repair). If the gender of the reparandum lingers, the perception of the repair (in terms of its gender) should vary based on the gender of the reparandum (Experiment 1). Similarly, if the lingering representation of the reparandum can interfere with the retrieval of the repair, we should observe a tendency to use fewer pronouns to refer to the repair when a reparandum exists compared to when it does not (Experiment 2).

Experiment 1

In this experiment, participants were auditorily presented with context scenarios such as (1) in which a gender-neutral word (*passenger*) was introduced. Then, in the following question, this gender-neutral word was either presented alone (2a), or was turned into a repair by introducing a reparandum into the discourse. Crucially, this reparandum was definitionally or stereotypically either male (*pilot*, 2b) or female (*stewardess*, 2c).[2]

(1) *Context*: Imagine there is a passenger flying on an old airplane, and suddenly there is some big, scary turbulence.

(2)

(a) *No Reparandum (Baseline)*: What do you think this passenger should do?

(b) *Male Reparandum*: What do you think this pilot uh I mean passenger should do?

(c) *Female Reparandum*: What do you think this stewardess uh I mean passenger should do?

Participants had to produce a spoken response to the questions, and we measured whether they referred to the gender-neutral word with a masculine (*he*, *him*, *his*, *himself*) or feminine pronoun (*she*, *her*, *hers*, *herself*) anywhere in their responses. If the gender of the reparandum lingers, and if the lingering representation of the reparandum can interfere with the retrieval of the repair, we should observe variations in form of reference as a function of the gender of the reparandum. Specifically, since generating a referential expression during the production of the response requires retrieval of the memory representation associated with the referent (in this case the repair), if the gender feature of the reparandum lingers in working memory, participants should have some tendency to refer to the gender-neutral repair as feminine or masculine depending on the gender of the reparandum. Specifically, participants should refer to the gender-neutral word using more masculine pronouns following (2b) than (2a). Similarly, they should refer to the gender-neutral word using more feminine pronouns following (2c) than (2a).[3]

Method

Participants

Fifty-four undergraduate students from the participant pool of University of California, Davis took part in the experiment in exchange for course credit. They were all native speakers of English and reported no hearing or language-related disorders.

Materials

We created 54 experimental items such as (1) and (2). However, we later removed one item from our analyses because we realised that most of our participants did

not know the meaning of the reparandum in that item (i.e. "spinster"; see Supplementary material for a list of all stimuli used in this experiment). We also created 32 filler items similar in format (e.g. *Imagine you learn that one of the members of a team has not done any of the work on an important class project. How would you deal with this slacking group member?*). Some fillers introduced non-human characters and some contained disfluencies affecting different parts of the question (e.g. on the verb).

The critical items and fillers were recorded at a normal speech rate by a research assistant who was blind to the purposes of the study (after practicing being disfluent) and were distributed across three lists such that each list contained only one version of each experimental item. The experiment was programmed in PsychoPy v1.83.03 which presented the questions and recorded the participants' responses. To reduce distractions during the experiment, the sentences were presented to the participants through a pair of headphones. Participants were initially told that the experiment was an "opinion survey" to avoid bias in their responses. However, they were informed about the manipulation after completing the study.

Procedure

The participants were tested individually and were seated in front of a computer screen. The instructions stated that they would be presented with an opinion survey and that they had to speak their opinions into a microphone. They were also asked to try to limit their responses to only one complete sentence, but it would also be fine if they produced additional sentences. To ensure that the speech was reasonably spontaneous, the participants were encouraged to start speaking as soon as a response came to mind.

Each trial started with a press of the spacebar, which would play the context sentence followed by the question, during which time the screen was empty. As soon as the question was over, a "Speak" prompt appeared on the screen, indicating that the participants could start speaking their responses. They had unlimited time to speak; once they were done, they pressed the "s" button on the keyboard to stop recording. The instructions emphasised that they should make sure to press the "s" button only after they were completely done with their response. After pressing "s", a prompt saying "Press spacebar to hear the next question" appeared on the screen, and by pressing the spacebar, the next trial began.

Before the experiment, each participant went through three practice items to get accustomed to the experiment. The experiment took approximately 50 minutes to complete and participants were allowed to take a break whenever they wished.

Coding and statistical analyses

We performed two separate and independent analyses. In the "Gender Interpretation" analysis, our measure was the gender of the gender-neutral word (male vs. female) as revealed by any type of pronoun used in the responses. For this analysis, we looked at the responses in their entirety (i.e. all the produced sentences if there were multiple sentences) and coded them for any pronoun type that revealed a gender for the gender-neutral word. Specifically, we coded subject pronouns (*he* or *she*), object pronouns (*him* or *her*), possessive pronouns (*his*, *her* and *hers*) and reflexive pronouns (*himself* and *herself*) in the responses.[4] Responses were excluded if they: (1) Contained collective pronouns (*they*) or coordinated singular pronouns (*he or she*, *his or her, his or hers, himself or herself*), (2) were simple repeats of the gender-neutral repair (*the passenger*) instead of a pronoun, and (3) were ambiguous with respect to the gender of the neutral noun (e.g. S*tay calm and follow the crew's instructions*). Participants used gender-neutral references (*they*, *them*, *their*, *theirs*, *he or she*, *his or her*, *his or hers*) and repeated nouns (*the passenger*) overwhelmingly. Plus, some participants stopped recording their responses before starting to speak in some trials. After applying the exclusion criteria, we were left with 616 analysable data points (21.5% of the total data) for this analysis.[5] We ran logit mixed-effects regression models (Baayen et al., 2008; Jaeger, 2008) predicting gender of the neutral word as a function of the gender of the preceding reparandum (No Reparandum vs. Female Reparandum vs. Male Reparandum).

For the "Form of Reference" analysis, we looked only at the first sentences in the responses and coded whether the gender-neutral word was referred to with a pronoun (*he* or *she*) regardless of its gender, or a repeated noun (*passenger*). For this analysis, we focused only on subject pronouns (*he* or *she*) and eliminated any non-subject pronouns (*his*, *her*, *hers*), as this is the standard practice in sentence continuation paradigms investigating pronominal reference (e.g. Arnold, 2001; Fukumura et al., 2013; Fukumura & Van Gompel, 2010a, 2011; Stevenson et al., 1994, among many others). In other words, we only coded whether the syntactic subject of responses was a subject pronoun (*he* or *she*) or a subject repeated noun (*the passenger*). Continuations were excluded from the data if: (1) The syntactic subject did not refer to the target (gender-neutral) word (e.g. *Everybody should keep calm.*), (2) the syntactic

subject was not the first-mentioned noun phrase in the continuation (e.g. *If the turbulence is too bad, the passenger is allowed to panic.*), and (3) the response was not a new and complete sentence (e.g. … *try to keep calm*). After applying the exclusion criteria, we were left with 2056 data points (71.8% of the total data) for the final analysis. Overall, participants used 1799 subject repeated nouns and only 257 subject pronouns to refer to the gender-neutral word.

For this analysis, too, we used logit mixed-models to predict form of reference to the gender- neutral word (pronoun vs. repeated noun) as a function of the presence of a reparandum (Male Reparandum + Female Reparandum vs. No Reparandum). Note that we collapsed over Male and Female Reparandum conditions for this analysis because we wanted to examine whether the very presence of reparandum, regardless of its gender, would complicate the subsequent retrieval of the memory representation associated with the repair and would, therefore, reduce pronominal reference to the repair.

Following Bates et al. (2015), we always fitted parsimonious random-effects structure for both "Gender Interpretation" as well as "Form of Reference" analyses (cf. Barr et al., 2013). That is, we systematically increased the complexity of the random-effects structure starting from the intercepts-only models and moving to the full random effects structure (Barr et al., 2013), keeping only the random components whose inclusion significantly improved the model fit, and minimised the AIC (Akaike Information Criterion) value. Based on Bates et al., (2015), such parsimonious models avoid overparametrisation or false convergence.

Results & discussion

Gender interpretation analysis

Figure 1 shows the percentage of masculine interpretations of the gender-neutral referential candidate (number of masculine interpretations / total number of both masculine and feminine interpretations) in the three critical conditions. Male interpretations of the gender-neutral referential candidate are shown in blue, and female interpretations are shown in red. As can be seen in Figure 1, there is a strong bias to use male pronouns for the neutral word in the baseline condition (177 masculine pronouns out of 207 total gendered pronouns, yielding 85.5%, and 30 feminine pronouns out of 207 total gendered pronouns, yielding 14.5%). Critically, however, this bias drops when the gender-neutral word is preceded by a female reparandum (111 masculine pronouns out a total of 192 gendered pronouns, i.e. 57.8%, and 81 feminine pronouns out a total of 192 gendered pronouns, i.e. 42.2%) and slightly increases when the reparandum refers to a male (191 masculine pronouns out of a total of 217 gendered pronouns, leading to 88%, and 26 feminine pronouns out of a total of 217 gendered pronouns, leading to 12%).

Table 1 reports the results of our regression models on the data. The random effects for both models included random intercepts for both subjects and items, as well as by-items random slopes for the effect of Condition, and the fixed effect for the models was Condition (we simply changed the reference condition in the two models to compare all three conditions). As can be seen from Table 1, the tendency to interpret the gender-neutral word as a male word significantly dropped when it was preceded by a female reparandum compared to when it occurred alone. However, male interpretation of the neutral noun did not statistically differ between the Male Reparandum and the No Reparandum conditions. Finally, the probability of interpreting the gender-neutral word as male was significantly greater in the Male Reparandum condition than in the Female Reparandum condition. Note that the intercept in the first model ("Comparison of Male & Female Reparandum with No Reparandum") denotes that the probability of using a masculine or feminine pronoun was significantly different from chance (i.e. 50%) in the reference (i.e. No Reparandum) condition. Similarly, because the reference condition in the second model ("Comparison of Male & Female Reparandum") was Feminine Reparandum, the intercept in this model suggests that the probability of using masculine and feminine pronoun was not reliably different from chance in the Feminine Reparandum condition.

These results clearly suggest that the gender of the reparandum lingers in memory even during an implicit, language production task, and influences how the repair is perceived in terms of gender. When the pronoun referring to the repair was being formulated, the features of the previously encoded reparandum (including its gender) were likely still active in memory, causing the gender of the subsequently produced pronouns to be consistent with the gender of the reparandum. The gender-interpretation results extend previous findings on lingering representations of disfluent words reported by Ferreira and colleagues (Ferreira et al., 2004; Lau & Ferreira, 2005) to gender information and also to implicit, language production tasks.

Interestingly, although participants avoided using singular gendered pronouns to the gender-neutral repairs in the vast majority of trials (see above), they showed a bias towards using masculine pronouns whenever they used a singular pronoun in the baseline

Figure 1. The gender interpretation of the gender-neutral word as a function of the gender of the reparandum. Experiment 1. Error bars denote standard errors of the mean.

condition. This bias might arise from general stereotypes concerning gender roles and is consistent with research findings suggesting that people have a bias in favour of masculine references in their language production when both masculine and feminine pronouns are viable (e.g. von der Malsburg et al., 2018). Importantly, this bias in favour of male interpretations in the baseline condition could be the reason why we did not observe a difference between the Male Reparandum and the Baseline condition: Masculine pronoun use might have already been at ceiling in the Male Reparandum condition, as suggested by the finding that 85% of pronouns were masculine in the baseline, No-Reparandum condition.

One potential concern about these results is that some of the repairs might have been treated by subjects not as error repairs in which the speaker mislabeled a particular individual and then produced a correction, but rather as appropriateness repairs (Levelt & Cutler, 1983), in which a speaker simply shifted from a less optimal to a more optimal description of the same individual. For example, in one of our critical items *the anchorman, uh I mean, the news reporter*, the repair might represent a shift to a different description, as an anchorman can also be a news reporter. Crucially, if the repairs were interpreted as appropriateness errors, the gender consistency between the gender-neutral repair and the reparandum is not surprising because the gender of the repair is already revealed by the reparandum. To exclude this possibility, we repeated the same analyses on a subset of items which were clearly error repairs in both the Male and the Female Reparandum conditions (there were 29 such items, identified with an asterisk in Supplementary material). Due to the loss of data, the full models did not converge for these analyses, and we, therefore, had to drop the random slopes and keep only the random intercepts. Importantly, we observed no differences in the results between these analyses and those on the full data. The frequency of masculine interpretations out of the total number of gendered interpretations in the No Reparandum, Female Reparandum and Male Reparandum conditions were 95 out of 111 = 85.6%, 66 out of 105 = 62.9%, and 108 out of 121 = 89.3%, respectively. Statistical analyses showed that the probability of interpreting the gender-neutral word as masculine was significantly greater in the Baseline than in the Female Reparandum condition

Table 1. The results for the Gender Interpretation analysis. Experiment 1.

Comparison/ Predictor	*β*	*SE*	*z*	*p*	Random effects
(1). *Comparison of Male & Female* Reparandum *with No* Reparandum.					
Intercept	3.28	0.90	3.64	<.001	(1 subject) +(1
Female vs. No Reparandum	−2.80	0.88	−3.17	.001	+Condition item)
Male vs. No Reparandum	0.05	1.04	0.04	0.96	
(2). *Comparison of Male & Female* Reparandum.					
Intercept	0.48	0.48	0.99	0.32	(1 subject) +(1
Male vs. Female Reparandum	2.85	0.74	3.87	<.001	+Condition item)

($\beta = -1.88$, $SE = .47$, $z = -3.95$, $p = < .001$), significantly greater in the Male Reparandum condition than in the Female Reparandum condition ($\beta = 2.34$, $SE = .49$, $z = 4.71$, $p = < .001$), and did not differ between the Male Reparandum and No Reparandum conditions ($\beta = .46$, $SE = .50$, $z = .91$, $p = < .35$).

Form of reference analysis

Apart from the *gender* of the pronouns, the *form* of the referring expressions that were used to refer to the repair (i.e. pronouns vs. repeated nouns, regardless of the gender of the pronouns) would reveal whether the reparandum lingers in memory and influences how the repair is processed. This is because if the reparandum lingers in memory, it should necessarily produce interference when the repair is being retrieved, leading to less pronominal reference to the repair (due to the semantic interference effect, discussed above). We also analysed form of reference to the gender-neutral word (pronoun vs. repeated noun) as a function of the presence of a reparandum. As mentioned above, we collapsed over the Male and Female Reparandum conditions, and our dependent variable was whether the syntactic subject of the responses was a subject pronoun (*he* or *she*) or a repeated noun (*the passenger*). The fixed effect was Condition (Male/Female Reparandum vs. No Reparandum), and the random effects included random intercepts for both subjects and items, as well as by-subjects and by-items random slopes for the effect of Condition. In the No Reparandum (baseline) condition, the participants used 85 subject pronouns out of a total of 701 pronouns and repeated nouns (12.1%), and in the average of the Male and Female Reparandum conditions, they used 172 subject pronouns out a total of 1355 pronouns and repeated nouns (12.7%) to refer to the gender-neutral repair. Mixed-effects regression models revealed no significant effect of the presence of a reparandum on form of reference to the repair ($\beta = -0.41$, $SE = 0.48$, $z = -0.85$, $p = .39$).[6] We did not find any significant differences between the baseline condition and any of the reparandum conditions either (Female vs. No Reparandum: $\beta = -0.23$, $SE = 0.43$, $z = -0.55$, $p = 0.58$; Male vs. No Reparandum: $\beta = -0.28$, $SE = 0.44$, $z = -0.64$, $p = 0.52$).[7]

This pattern of results suggests that the presence of a reparandum did not cause interference during the retrieval of the repair, which runs counter to previous findings. Specifically, the degree of sematic similarity between the repair and the reparandum is necessarily greater when a reparandum is present, which should have caused retrieval interference, lowering the memory activation of the repair and leading to less pronoun use to refer to it (The "semantic interference effect"; Arnold & Griffin, 2007a; Fukumura et al., 2013). This is because the repair necessarily has some degree of similarity to the reparandum when a reparandum exists (Male and Female Reparandum conditions), but there can be no semantic similarity between the repair and the reparandum when there is no reparandum (the Baseline condition). Moreover, the very presence of the reparandum might have functioned as an additional character in the discourse (at least temporarily) that lingered in memory. To the extent that this is true, the reparandum should have reduced the activation level of the repair, lowering the probability of pronominal reference to it (the "additional character effect"; Arnold & Griffin, 2007a; Fukumura et al., 2010b).

For the two reasons mentioned above, we should have observed significantly less pronominal reference to the repair when it was preceded by a reparandum compared with when there was no reparandum. However, such an effect might have been obscured in Experiment 1 because the participants used an overwhelming number of repeated nouns (as opposed to subject pronouns) to refer to the gender-neutral word, which greatly lowered variability in the dependent variable. Importantly, we cannot tell from the results of this experiment whether the large number of repeated nouns was due to a tendency to avoid assigning a gender to the gender-neutral repair, or by low activation levels for the gender-neutral repair. To address this issue, we conducted a second experiment.

Experiment 2

The results of the first experiment showed that the gender of unintentionally spoken reparanda lingers in the memory of the comprehender and influences how a gender-neutral repair is interpreted; if the reparandum is feminine, the probability of viewing the gender-neutral word as female increases compared with when the repair occurs alone. Likewise, if the reparandum is male, the probability of interpreting the gender-neutral repair as masculine reliably increases relative to when the reparandum is female. However, Experiment 1 leaves two questions unanswered: First, although the results of the first experiment provide evidence suggesting that reparanda linger in memory, they do not provide a clear answer to the question of whether they cause interference during the retrieval of the repair. Second, Experiment 1 does not provide a clear response to *how* the reparandum changes the perception of the repair. One possibility is that once the repair is encountered, its properties are edited to match that of the reparandum, but the reparandum itself is not added to the discourse model as an additional character. That is, the *edited*

repair is the only discourse representation left from the disfluency operation. A second possibility, however, is that the reparandum functions as an additional discourse entity, and influences the perception of the repair. If this latter possibility is true, the reparandum should cause interference during the retrieval of the repair, even if the genders of reparandum and the repair match. Moreover, if the reparandum functions as an additional discourse character, it should not only affect the retrieval ease of the repair, but also the reparandum should affect how the repair is retrieved at a subsequent point. Specifically, semantic similarity (between the reparandum and the repair) should necessarily be greater when there is a reparandum compared to when the repair occurs alone. As such, if the reparandum functions as an additional discourse character (as opposed to producing an edited repair), there should be more semantic interference, more retrieval difficulty and therefore less pronominal reference to the repair when a reparandum precedes it than when it does not (i.e. the semantic interference effect, Arnold & Griffin, 2007a).

In Experiment 2, we capitalised on previous findings concerning the form of reference (see Introduction) and examined whether the presence of a reparandum complicates the retrieval (accessibility) of the repair when its gender matches that of a reparandum. If it does, then it provides evidence supporting the possibility that reparanda function as additional discourse entities (as opposed to producing edited repairs). Experiment 2 utilised sentences containing two referential candidates (*Julie* and *Robert*) such as (3). The experimental manipulation was whether a gender-congruent reparandum (*Mary*) preceded the repair (*Julie*). Since the serial position and the syntactic role of referential candidates have been shown to have a considerable effect on form of reference (e.g. Arnold, 2001; Fukumura & Van Gompel, 2010a, 2011; Gordon et al., 1993; Gundel et al., 1993; Karimi et al., 2014; Stevenson et al., 1994), and in order to control for these effects, the repair was also manipulated with regards to these variables, such that it occurred both as the syntactic subject and the first-mentioned word as well as the syntactic object and the second-mentioned word of the sentences. In addition to the repair, critical sentences contained a second referential candidate (*Robert*) that was not preceded by a reparandum and would assume the role of the first- or the second-mentioned word depending on the syntactic role/serial position of the repair (*Julie*).[8] To simplify the design and the analyses, we labelled the referential candidate that always occurred alone (i.e. *Robert* in 3) as the "competitor", creating two Word Types: Repair (*Julie*) and Competitor (*Robert*, see 3).

(3)

a. *Fluent:* $\text{Julie}_{\text{repair}}$ walked away from $\text{Robert}_{\text{competitor}}$.
$\text{Robert}_{\text{competitor}}$ walked away from $\text{Julie}_{\text{repair}}$.

b. *Disfluent:* Mary uh I mean $\text{Julie}_{\text{repair}}$ walked away from $\text{Robert}_{\text{competitor}}$.
$\text{Robert}_{\text{competitor}}$ walked away from Mary uh I mean $\text{Julie}_{\text{repair}}$.

Participants heard the prompt sentence and then provided a spoken continuation (they were free to say whatever they wished). If the lingering reparandum functions as an additional discourse entity and causes interference during the retrieval of the repair, there should be less pronominal reference to the repair (*Julie*) when it is preceded by a reparandum (3b) than when it occurs alone (3a), which could be due to a semantic interference effect, or an additional character effect, or both (see Introduction), but not due to edited repairs. This is because if the repair is difficult to retrieve from memory (i.e. is less accessible), there should be a preference to refer to it with repeated nouns rather than pronouns (see Introduction).

Method

Participants

108 undergraduate students from the participant pool of University of California, Davis took part in the experiment in exchange for course credit. All participants were native speakers of American English and reported no hearing or language-related disorders.

Materials

Forty-two experimental sentences containing two referential candidates were created. Each sentence had four versions, as illustrated in (3). The first referential candidate was always the syntactic subject as well as the first-mentioned noun phrase and the second referential candidate was always the syntactic object and the second-mentioned noun phrase. As mentioned above, the serial position of the two referential candidates (and thus their syntactic roles) was counterbalanced to control for syntactic prominence effects. A gender-congruent reparandum (*Mary*) preceded the same referential candidate (*Julie*) in the two-word positions to create a Fluency condition with two levels (Fluent vs. Disfluent). A second referential candidate (the competitor, *Robert*) was also included in all experimental sentences. The gender of the second referential competitor was always incongruent with those of the reparandum and the repair. To add variety to the disfluent stimuli, we used four different editing terms: "uh I mean", "oh no", "sorry", and "oops" (see Supplementary material for the list of all stimuli used in this experiment).

Four experimental lists were created, each containing one version of an experimental item. The experimental sentences in each list were mixed with 40 fillers that also contained two referential candidates. However, the words in the fillers were not names but rather determiner-noun combinations (e.g. *the actor*) and some of them had a relative clause attached to them (e.g. *the actor who was frustrated and visibly upset*). The fillers were the experimental stimuli for an unrelated study investigating the production of referring expressions. The critical and the filler items were recorded by a female native speaker of English who was blind to the purpose of the experiment and who had practiced being intentionally disfluent. The experiment was designed in PsychoPy v1.83.03 which presented the stimuli and recorded the participants' continuations for later transcription and analysis. As in Experiment 1, the sentences were played to the participants through a pair of headphones to maximise focus during the experiment.

Procedure

The procedure for this experiment was identical to that of Experiment 1, except that instead of answering questions, the participants were asked to come up with a meaningful continuation to the sentence they had just heard. They were also asked to produce "new and complete" continuations, and to not "merely lengthen the given sentence", which was made clear with examples. (Specifically, our prompt sentence in the instructions was: "The carpet caught fire". An example of "merely lengthening" the given sentence was "and burned", and an example of a "new and complete" continuation was "The owner was very upset".).

Coding and statistical analyses

We analysed the effect of our manipulations on two binomial dependent variables: "Choice of Referent" coded which entity the participants talked about in their continuations, and "Form of Reference" captured whether the participants used a pronoun or a repeated noun to refer to that entity. The exclusion criteria were the same as those for the "Form of Reference" analysis in Experiment 1, except that we did not discard collective referring expressions such as *They* or *Julie and Robert* (see below). Apart from applying the exclusion criteria, additional data were excluded because participants stopped the recording of their continuations too early, resulting in a total of 1202 (26.5%) excluded data and leaving us with 3334 analysable responses.

Analysis of the choice of referent was performed to examine the effect of Fluency (Fluent vs. Disfluent, 3a vs. 3b) on which referential candidate the participants chose to talk about (the first- or the second-mentioned word). Form of Reference analyses were further divided into two analyses. In "Singular Form of Reference" analysis, we compared the form of reference (pronoun vs. repeated noun) to the repair (*Julie*) and the competitor (*Robert*) as a function of Fluency (Fluent vs. Disfluent), Word Type (Repair vs. Competitor), and their interaction (see above). In the "Collective Reference" analysis, we looked at collective reference to both referential candidates, as a function of Fluency (3a vs. 3b). Collective reference could be made using a plural pronoun (*They*) or a coordinated construction including both referential candidates (e.g. *Robert and Julie*), and we were interested in the frequency of pronominal vs. coordinated collective reference (*They* vs. *Robert and Julie*) as a function of Fluency (Fluent vs. Disfluent), or the presence of a reparandum. Because pronouns reflect higher accessibility of the associated referential candidates than repeated nouns (see above), one would expect more pronominal than coordinated collective reference for two highly accessible referential candidates than for two referential candidates that are less activated in memory.

For all analyses, we ran logit mixed-effects regression models (Baayen et al., 2008; Jaeger, 2008). For the Choice of Referent analysis, Fluency (Fluent vs. Disfluent) was the predictor (which was always centred to the mean), and choice of referent (first- vs. second-mentioned referential candidate) was the binomial outcome. In the Singular Form of Reference analysis, mean-centred Fluency (Fluent vs. Disfluent), mean-centred Word Type (Repair vs. Competitor), and their interaction were the predictors, and form of reference to either referential candidate (pronoun vs. repeated noun) was the binomial outcome. In the Collective Form of Reference analysis, mean-centred Fluency (Fluent vs. Disfluent) was the predictor and collective form of reference (collective pronoun vs. coordinated nouns) was the outcome.

As in Experiment 1, and in order to avoid over-parametrisation, we took a parsimonious approach towards the random effects structures for our models by including only the random components that significantly improved model fit (Bates et al., 2015, c.f. 2013; see above). We will first report the results for Choice of Referent analyses and then report the results of Form of Reference analyses in relation to the two original questions that motivated Experiment 2.

Results

Choice of referent analysis

We first analysed which referential candidate was talked about in the continuations as a function of the presence

of a reparandum in the given sentence. For this analysis, we included random intercepts for both subjects and items as well as by-subjects and by-items random slopes for the effect of Fluency. As mentioned above, the fixed effect was Fluency (Fluent vs. Disfluent). Table 2 reports the raw frequencies as well as the percentages of reference to the first- and the second-mentioned referential candidates in the two critical conditions and Table 3 reports the results of our analyses for choice of referent.[9] The preference to talk about the first-mentioned referential candidate was reliably greater in the Disfluent than in the Fluent condition, suggesting that the presence of a reparandum strengthens the tendency to talk about the syntactically more prominent referential candidate.

Form of reference analyses

Based on the hypotheses motivating this experiment, we analysed form of reference to each referential candidate, as well as collective reference to both referential candidates. We report each analysis separately.[10]

Singular form of reference

We first analysed form of reference to the repair (*Julie*) and the competitor (*Robert*) to see whether the retrieval of any of these word types is more difficult as a function of processing a reparandum. For this analysis, we compared pronominal reference to the repair and the competitor when there was a disfluency vs. when there was no disfluency (i.e. pronominal reference to the two *Julie*s in 3a vs. pronominal reference to the two *Julie*s in 3b; and pronominal reference to the two *Robert*s in 3a vs. pronominal reference to the two *Robert*s in 3b). As before, for this analysis, we included random intercepts for both subjects and items, as well as by-subjects and by-items random slopes for the effect of Fluency. Figure 2 shows the proportions of pronominal reference (blue) and repeated reference (yellow) to the repair and the competitor when they are fluent and disfluent. As can be seen in this figure, pronominal reference to the repair is greater when it occurs alone (in the Fluent condition) than when it is preceded by a reparandum (in the Disfluent condition). However, the rate of the pronominal reference to the competitor did not vary across the Fluent and Disfluent conditions.

Table 2. Choice of reference biases. Experiment 2.

Condition/%Reference	First-mentioned	Second-mentioned
Fluent	33.7% (408)	66.3% (802)
Disfluent	38.2% (463)	61.8% (747)

Table 3. The results for Choice of Referent analyses. Experiment 2.

Predictor	*β*	*SE*	*z*	*p*	*Random effects*
Intercept	0.69	0.18	3.81	<.001	(1+Fluency|subject) +(1+Fluency|
Fluency	0.34	0.16	2.15	0.03	item)

Table 4 reports the raw frequencies as well as the percentages of pronominal and repeated reference, and Table 5 reports the results of our regression model predicting probability of pronominal reference as a function of Fluency and Word Type. As can be seen in this table, we observed a main effect of Fluency, with significantly more overall pronominal reference (i.e. to both the repair and the competitor) in the Fluent than in the Disfluent condition. However, the effect Word Type was not statistically significant. Importantly, we also observed a significant interaction between Fluency and Word Type. Follow-up simple analyses revealed that the repair received significantly more pronominal reference in the Fluent than in the Disfluent condition (see the left grid of Figure 2), whereas pronominal reference to the competitor did not reliably vary between the Fluent and the Disfluent conditions.

A reduction in pronominal reference to the repair as a function of the presence of a gender-congruent reparandum is consistent with the findings of Arnold and Griffin (2007a) and Fukumura et al. (2013), who showed that gender congruence between two potential referents reduces overall pronominal reference (i.e. to either referential candidate), presumably due to a semantic interference effect, and also with the finding that additional referential characters reduce the overall activation of potential referents (Arnold & Griffin, 2007a; Fukumura et al., 2010b). Note that the additional characters account predicts that pronominal reference to the competitor should be reduced too (Arnold & Griffin, 2007a). However, we did not observe such an effect (but see results for collective reference below). One potential reason for lack of an effect on the competitor could be that the reparandum causes more retrieval interference for the repair than for the competitor. We will discuss this possibility further in the General Discussion.

Collective form of reference

Because there was considerable collective reference to the referential candidates (914 cases in total) through the use of collective pronouns (*they*) or coordinated constructions (*Julie and Robert* or *Robert and Julie*), we also examined the probability of collective reference to the referential candidate (collective pronoun vs. coordinated names) as a function of Fluency (Fluent vs. Disfluent). If the presence of a reparandum reduces the overall

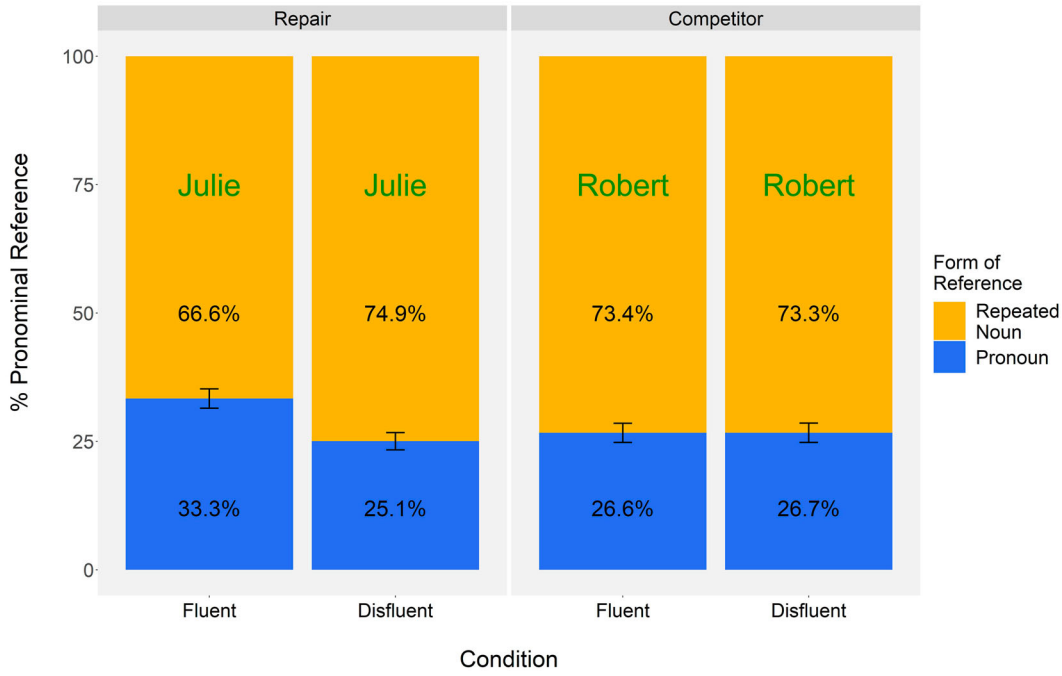

Figure 2. Proportion of pronominal (yellow) and repeated (blue) reference to the repair (*Julie*) and the competitor (*Robert*) across the Fluent and the Disfluent conditions in Experiment 2. Error bars denote standard errors of the mean.

activation of both referential candidates, they both should be more accessible in the Fluent than in the Disfluent condition, leading to more collective pronominal reference (*They*) rather than repeating both referential candidates (*Julie and Robert*) in the Fluent condition than in the Disfluent condition. Once again, we included random intercepts for both subjects and items, as well as by-items random slopes for the effect of Fluency. Fluency was the fixed effect. Figure 3 shows the proportions of collective pronominal reference (*they*) and collective coordinated reference (*Julie and Robert*) as a function of Fluency, which clearly indicates a stronger preference for pronominal collective reference when there was no reparandum in the discourse than when there was one.

We ran a separate logit mixed-effects regression model predicting form of collective reference as a function of Fluency. The results revealed significantly more collective pronouns in the Fluent (89.6%, 398 collective pronouns out of total of 444 collective pronouns and coordinated reference) than in the Disfluent condition (76.8%, 361 collective pronouns out of total of 470 collective pronouns and coordinated reference; $\beta = -2.21$, $SE = 0.44$, $z = -5.06$, $p < .001$), lending further support to the hypothesis that the memory representation associated with the reparandum causes interference and reduces the accessibility of both referential candidates. Similar to singular pronominal reference to either referential candidate, the results of collective reference analyses are also consistent with the possibility that the reparandum functions as an additional discourse entity and reduces the overall activation of discourse characters (i.e. the "additional character effect", Arnold & Griffin, 2007a; Fukumura et al., 2010b).

Table 4. Form of reference biases. Experiment 2.

Condition/%Reference form		Pronouns	Repeated Nouns
Repair	Fluent	33.3% (212)	66.6% (424)
	Disfluent	25.1% (165)	74.9% (494)
Competitor	Fluent	26.6% (153)	73.4% (421)
	Disfluent	26.7% (147)	73.3% (404)

General discussion

This study investigated whether the representation of sentence material not intended to be part of the official content, in this case a reparandum, remains in memory and influences subsequent language processing. The tool we used to examine this question was to present subjects with a sentence prompt containing a reparandum and a repair to which they then generated

Table 5. The results for Form of Referent analyses. Experiment 2.

Predictor	β	SE	z	p	Random effects
Intercept	−2.05	0.39	−5.27	<.001	(1+Fleuncy
Fluency	−0.65	0.26	−2.52	0.01	subject) + (1
Word Type	−0.09	0.15	−0.64	0.52	+Fluency item)
Fluency × Word Type	0.61	0.29	2.12	0.03	
Fluency effect for Repair	−1.21	0.50	−2.41	0.02	Same as in main model
Fluency effect for Competitor	−0.34	0.51	−0.68	0.50	Same as in main model

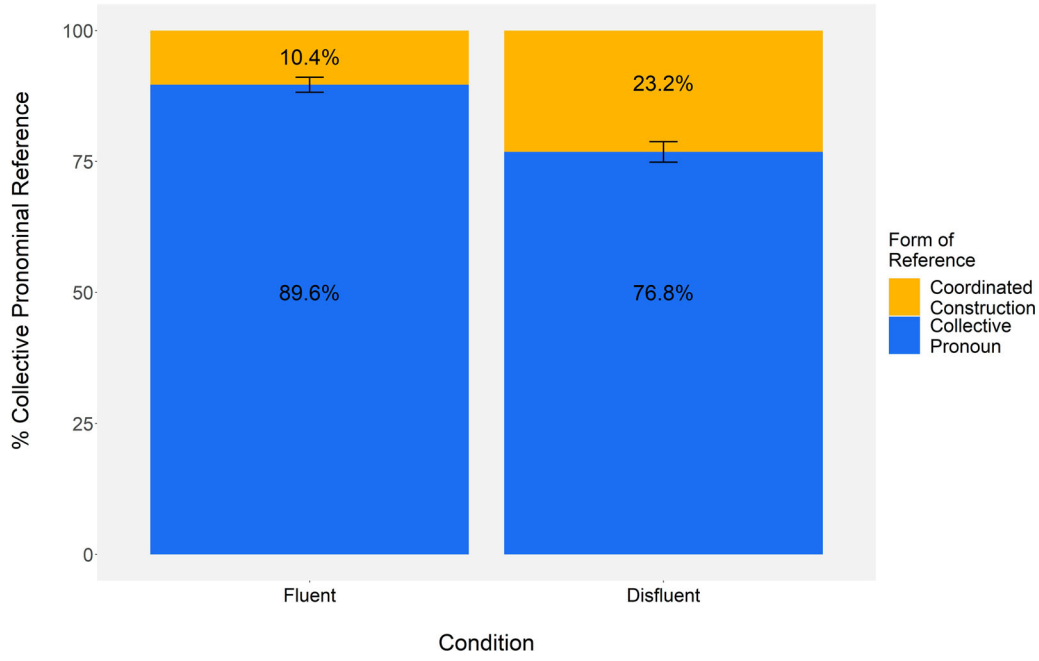

Figure 3. Proportion of pronominal (yellow) and coordinated (blue) collective reference to the two discourse entities across the Fluent and the Disfluent conditions in Experiment 2. Error bars denote standard errors of the mean.

a spoken response. Given that many such responses would contain pronouns referring to the repair, this task allowed us to examine our hypotheses by assessing whether the pronoun had features in common with the reparandum, in this case, gender marking (Experiment 1), and also whether the activation level of the repair dropped (due to interference) as function of having a preceding reparandum (Experiment 2).

Experiment 1 demonstrated that the gender of the reparandum lingers in working memory and influences how the gender-neutral repair was interpreted: When the reparandum was female, the gender-neutral repair was interpreted as female, and when the reparandum was male, the gender-neutral repair was more likely to be interpreted as a male word compared to when the reparandum was female. Thus, the picture painted by the results of Experiment 1 is one in which the gender of the reparandum lingers in memory and influences how the repair is perceived in terms of its gender. In Experiment 2, we investigated whether the presence of a reparandum could cause interference and reduce the accessibility of the memory representation associated with the repair. We measured the probability of pronoun use and found significantly less pronominal reference to the repair and also less collective pronominal reference to both discourse entities when the repair was preceded by a reparandum, suggesting that the reparandum's effect on the perception and retrieval of the repair is due to it functioning as an additional discourse character (rather than it producing an edited repair, see above).Note that although the singular form of reference results are also compatible with the semantic interference effect, the reparandum must first be present in the discourse model before it can cause semantic interference.

A critical question that arises in the face of these data is why reparanda should linger in memory in the first place. One possibility is that, once encoded, the cognitive cost of suppressing the representation associated with reparanda might be high, causing the associated representations to linger in memory. This is because suppressing such representations necessarily requires some degree of cognitive control such as *inhibiting* the representation of the reparandum, or *updating* the contents of working memory, or *shifting* attention and resources from the reparandum to the repair (Miyake & Friedman, 2012). It could be the case that executing such cognitive control processes consumes computational resources, and so the language processing system generally avoids them.

Another reason for lingering representations of reparanda could be that they may actually turn out to be helpful for language processing. First, previous research has demonstrated that reparanda are used to make predictions about the lexical identity of the repair. Lowder and Ferreira (2016), for example, demonstrated that after hearing a word like *dog* followed by an editing term (*uh I mean*), people start to direct their gaze towards the semantically more related competitor (*cat*) before it is spoken. Thus, ignoring or removing the reparandum could in fact result in losing valuable information

that would otherwise facilitate the processing of the upcoming linguistic material. Second, it is possible that reparanda could be used to make inferences about the speaker and the intended message they are attempting to convey. That is, the reparandum could reveal pragmatic information about the characteristics of the speaker and/or of the message, which could then be used to facilitate language processing in particular as well as interpersonal communication in general. This idea is consistent with previous research showing that the characteristics of the speaker are rapidly taken into account by the listeners (Tesink et al., 2009; Van Berkum et al., 2008; Yildirim et al., 2016; See also Kleinschmidt & Jaeger, 2011), and also with theories of language processing assuming that linguistic communication is a joint activity in which speakers and listeners may base their messages upon what they know about each other (Brennan & Clark, 1996; Goldman, 2006; Pickering & Garrod, 2004; 2013). Translated to the current study, it could be the case that the representation associated with the reparandum is needed to make such inferences about the speaker. Based on this account, for example, the results of Experiment 1 could be explained by proposing that listeners could have inferred that if the *passenger* was male, it would have been unlikely that the speaker confused him with a female person such as a *stewardess*. Such an inference then leads to the assumption that the *passenger* would likely be a female. Note that the "cognitive control" and the "usefulness" accounts of lingering reparanda are not necessarily mutually exclusive. Although it might be cognitively more difficult to inhibit unwanted representations, update the contents of working memory, or shift attention, the benefits of facilitated subsequent processing may outweigh the cognitive costs, encouraging the language processing system to "pay the cognitive costs" in order to facilitate the processing of subsequent information.

Our results are consistent with previous studies showing that the syntactic properties and initial and erroneous syntactic representations linger in memory during explicit, grammaticality judgment tasks (Christianson et al., 2001; Ferreira et al., 2004; Lau & Ferreira, 2005). The current results extend those findings by showing that the gender feature of reparanda lingers too, and also that the lingering representation of the reparandum interferes with the subsequent retrieval of the repair by functioning as an additional character in the current discourse representation. Our results from Experiment 2 are also in line with research showing that referential competitors that share their gender with the referential target cause semantic interference during the retrieval of the target, reducing the target's activation and therefore the probability of pronominal reference to it (Arnold & Griffin, 2007a; Fukumura et al., 2013). Because the reparandum in Experiment 2 always had the same gender as the repair, and because our results from Experiment 1 showed that the gender of reparanda lingers, the lingering reparandum probably functioned as an additional character (as opposed to producing an edited repair) and increased semantic interference during the retrieval of the repair, reducing its activation level and therefore pronominal reference to it. Moreover, the results of our analyses on overall pronoun use and collective reference in Experiment 2 is consistent with previous studies showing the additional referential characters reduce the activation level of the target candidate, and therefore pronominal reference to it. Specifically, because we know that representation of the reparanda lingers in memory, the disfluent sentences essentially always contained one more discourse entity relative to the fluent sentences. In other words, the lingering representation associated with the reparandum likely functioned as an additional character, causing the total activation level to be distributed over more discourse entities, and therefore reducing the activation level for each individual entity (Arnold & Griffin, 2007a; Fukumura et al., 2010b).

One potential alternative account of these results might appeal to "communicative clarity". The idea here is that the reduced tendency to use pronouns in the presence of a reparandum could be due to an effort on the part of the speaker to reduce referential uncertainty (i.e. to increase communicative clarity) when a third, unintended referential candidate (i.e. the reparandum) is also present in the discourse. This alternative explanation is unlikely for two reasons, however: First, because the reparandum and the repair were always of the same gender (e.g. *Mary* uh I mean *Julie*), and always of a different gender compared to the second referential candidate (*Robert*), using a pronoun to refer to the repair would be just as unambiguous and communicatively efficient as a repeated noun. The second weakness of this argument is that any speaker taking communicative clarity into consideration must first be aware (consciously or unconsciously) that the representation of the reparandum might linger in the memory representation formed by the listener.

Although our results from Experiment 2 provided evidence in support of retrieval interference caused by the presence of a reparandum, we did not observe less pronominal reference to the repair when it was preceded by a reparandum than not in Experiment 1. As mentioned above, this could be because there was an overwhelming use of repeated nouns (as opposed to pronouns) in this experiment. The low variation in form of reference could

have obscured potential effects and may have been a strategy to avoid assigning a gender to the neutral noun. Thus, variation in form of reference was perhaps driven by both memory activation levels as well as a tendency to avoid assigning a gender to the gender-neutral word, making it difficult to draw any conclusions about form of reference in Experiment 1. However, this concern was addressed by the results of Experiment 2 where variations in form of reference was greater.

Additionally, in Experiment 2, we expected to observe less pronominal reference to the competitor (i.e. *Robert*) too, but the results showed a reduction in pronominal reference for the repair but not for the competitor. However, we did observe less collective pronominal reference to both the repair and the competitor when there was a reparandum in the discourse compared to when there was not. Why does the reparandum affect collective reference to both entities, but not singular reference to a competitor? We propose that this is probably because the reparandum produces more interference for the repair than for the competitor for two reasons: First, the genders of the reparandum and the repair always matched, whereas the genders of the reparandum and the competitor were always different. This means that while the semantic interference effect applied to the repair, it did not apply to the competitor (Arnold & Griffin, 2007a; Fukumura et al., 2013). Second, the reparandum and the repair are conceptually closer together than the reparandum and the competitor. This is because the reparandum is corrected to arrive at the repair, meaning that there is necessarily more confusability between the reparandum and the repair than between the reparandum and the competitor. Such confusability may then lead to more explicit reference to the repair, whereas the competitor may be less affected. However, note that when collective reference is made to both entities, the confusability attached to the repair automatically applies to the collective entity, leading to less collective pronominal reference when there is a reparandum in the discourse than not. Note that previous research reporting the additional character effect always had a "normal" additional character rather than a reparandum (Arnold & Griffin, 2007a; Fukumura et al., 2013), which might be the reason why they observed less singular pronominal reference to the competitor.

Finally, in our experiments, the rate of disfluency was higher than what is considered typical in spontaneous speech (i.e. 6-10%, Bortfeld et al., 2001; Fox Tree, 1995). Thus, one potential challenge to our study could be that the higher rates of disfluencies might have made the participants overly conscious about our manipulation, giving rise to the observed results. However, one must bear in mind that although we observed some effects of disfluency on subsequent processing (i.e. the perception of repairs in Experiment 1, and lower pronominal reference to the repair in Experiment 2), we also did *not* observe other expected results (i.e. the null effect of reparanda on form of reference to the repair in Experiment 1, and to the competitor in Experiment 2). If conscious awareness gave rise to the observed effects, it is not clear why it has failed to produce other effects. Moreover, there are practical limitations for why we had higher rates of disfluency in our study. Specifically, we would have needed to increase the number of trials to 530 in Experiment 1, and to 420 in Experiment to maintain a 10% rate of disfluencies, which is the highest end of the range of disfluency rates in spontaneous speech. Such high numbers of trials might have led to other (perhaps more serious) problems such as fatigue, and attentional lapses on the part of the participants.

To the best of our knowledge, our study is the first showing effects of misspoken words on the perception and subsequent retrieval of intended words. As such, although our results provide answers to our original questions, they also give rise to many new questions. For example, whether and how does the type and duration of editing terms affect how the intended word is processed? How do prosodic cues on the reparandum and/or the editing terms affect the results? Also, based on our results, it is not clear how long the lingering representation of the reparandum will last in memory. Relatedly, individual differences in working memory span or language skills might influence how reparanda are processed and whether and how they influence subsequent processing. Moreover, it is not clear what properties of editing terms give rise to lingering reparanda. One possibility is that the editing operation in a repair disfluency results in heightened attention, giving rise to the lingering effect. Finally, we only investigated gender in our experiments, and it is therefore not clear whether other semantic properties would also linger in memory. These are questions that we hope future research will address.

Taken together, the results of both experiments provide answers to the two original questions raised in this study: (1) Using an implicit language production task, we observed that the gender feature of a reparandum lingers in memory and affects how the repair is perceived; (2) the memory representation associated with the reparandum functions as an additional discourse character and not only causes interference during the retrieval of the repair, but also reduces overall activation levels, leading to less collective pronominal reference to both referential candidates. At a theoretical level, our results lend support to the Good Enough approach to

language processing according to which the linguistic representations created from incoming input are not always veridical, in part because remnants of unintended/unwanted representations are not necessarily removed from the current discourse representation to produce a clean and idealised final representation (e.g. Ferreira & Patson, 2007; Ferreira et al., 2002; Karimi & Ferreira, 2016; also see Christianson, 2016). The final interpretation of a sentence, then, may reflect not only its compositional meaning but also the concepts and ideas that are activated and only partially suppressed during online processing.

Notes

1. There is a distinction between "accessibility" and "availability" of information in the literature, with the former referring to retrieval speed and the latter referring to probability of successful retrieval of the target memory item (e.g. McElree, 2000; McElree et al., 2003; Foraker & McElree, 2007; Van Dyke & McElree, 2011). However, the current paper is agnostic to this distinction. We therefore use the term "accessibility" throughout the paper and we use this term interchangeably with "*ease* of retrieval".
2. Our use of forms that assume these gender stereotypes should in no way be taken to imply acceptance or endorsement of these gender stereotypes. In fact, we look forward to the day when this manipulation is ineffective.
3. It is important to note that we also calculated the *degree* of semantic similarity between male reparanda and the repairs as well as female reparanda and the repairs using Latent Semantic Analysis (http://lsa.colorado.edu/). We then ran a *t*-test on the similarity values and found no significant difference between the two conditions. Thus, any potential differences between the conditions could not be caused by the *degree* of semantic similarity between the conditions.
4. Most of our data for this analysis came from object, possessive, and reflexive (i.e. non-subject) pronouns.
5. Note that the low number of analyzable datapoints is caused by the nature of the task plus some technical issues, and does not invalidate the results. In an open-ended task like ours, participants are free to say whatever they wish and in whatever form they choose. Consequently, they may or may not use a "gender-revealing" pronoun. Note also that the total number of gendered pronouns was almost equally distributed across the three conditions (see below), indicating that the excluded trials are not biased based on condition.
6. The random effects for this model included random intercepts for both subjects and items, as well as by-subjects and by-items random slopes for the effect of Condition.
7. The random effects for this model also included random intercepts for both subjects and items, as well as by-subjects random slopes for the effect of Condition.
8. Note that *Julie* is not technically a repair in the Fluent condition, but we will call it a repair for the sake of simplicity.
9. Note that the total number of references to the first- and the second-mentioned words add up to 2420, whereas the total number of analyzable responses was 3334 (see above). The difference is due to 914 of analyzable responses which were analyzed as collective references (see below).
10. To examine the effect of syntactic role/serial position, we also ran a model predicting form of reference as a function of "Word Position". The results revealed significantly greater probability of pronominal reference to the syntactic subject/the first-mentioned referential candidate than to the syntactic object/the second-mentioned referential candidate ($\beta = -3.07$, $SE = 0.35$, $z = -8.59$, $p < .001$).

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

Almor, A. (1999). Noun-phrase anaphora and focus: The informational load hypothesis. *Psychological Review*, *106*(4), 748–765. https://doi.org/10.1037/0033-295X.106.4.748

Almor, A. (2004). A computational investigation of reference in production and comprehension. In J. Trueswell & M. Tanenhaus (Eds.), *Approaches to studying world-situated language use: Bridging the language-as-product and language-as-action traditions* (pp. 285–301). MIT Press.

Arnold, J. E. (2001). The effect of thematic roles on pronoun use and frequency of reference continuation. *Discourse Processes*, *31*(2), 137–162. https://doi.org/10.1207/S15326950DP3102_02

Arnold, J. E., Fagnano, M., & Tanenhaus, M. K. (2003). Disfluencies signal theee, um, new information. *Journal of Psycholinguistic Research*, *32*(1), 25–36. https://doi.org/10.1023/A:1021980931292

Arnold, J. E., & Griffin, Z. (2007a). The effect of additional characters on choice of referring expression: Everyone counts. *Journal of Memory and Language*, *56*(4), 521–536. https://doi.org/10.1016/j.jml.2006.09.007

Arnold, J. E., Kam, C. L. H., & Tanenhaus, M. K. (2007b). If you say *thee uh* you are describing something hard: The on-line attribution of disfluency during reference comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(5), 914–930. https://doi.org/10.1037/0278-7393.33.5.914

Arnold, J. E., Tanenhaus, M. K., Altmann, R., & Fagnano, M. (2004). The old and thee, uh, new. *Psychological Science*, *15*(9), 578–582. https://doi.org/10.1111/j.0956-7976.2004.00723.x

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. https://doi.org/10.1016/j.jml.2007.12.005

Bailey, K. G. D., & Ferreira, F. (2003). Disfluencies affect the parsing of garden-path sentences. *Journal of Memory and Language*, *49*(2), 183–200. https://doi.org/10.1016/S0749-596X(03)00027-5

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278. https://doi.org/10.1016/j.jml.2012.11.001

Barr, D. J., & Seyfeddinipur, M. (2010). The role of fillers in listener attributions for speaker disfluency. *Language and Cognitive Processes*, *25*(4), 441–455. https://doi.org/10.1080/01690960903047122

Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). Parsimonious mixed models. *arXiv Preprint ArXiv*, *1506*, 04967.

Bock, J. K. (1982). Toward a cognitive psychology of syntax: Information processing contributions to sentence formulation. *Psychological Review*, *89*(1), 1–47. https://doi.org/10.1037/0033-295X.89.1.1

Bock, J. K., Loebell, H., & Morey, R. (1992). From conceptual roles to structural relations: Bridging the syntactic cleft. *Psychological Review*, *99*(1), 150–171. https://doi.org/10.1037/0033-295X.99.1.150

Bock, J. K., & Warren, R. K. (1985). Conceptual accessibility and syntactic structure in sentence formulation. *Cognition*, *21*(1), 47–67. https://doi.org/10.1016/0010-0277(85)90023-X

Bortfeld, H., Leon, S., Bloom, J., Schober, M., & Brennan, S. (2001). Disfluency rates in conversation: Effects of age, relationship, topic, role, and gender. *Language and Speech*, *44*(2), 123–147. https://doi.org/10.1177/00238309010440020101

Branigan, H. P., Pickering, M. J., & Tanaka, M. (2008). Contributions of animacy to grammatical function assignment and word order during production. *Lingua. International Review of General Linguistics. Revue internationale De Linguistique Generale*, *118*(2), 172–189. https://doi.org/10.1016/j.lingua.2007.02.003

Brennan, S. E. (1995). Centering attention in discourse. *Language and Cognitive Processes*, *10*(2), 137–167. https://doi.org/10.1080/01690969508407091

Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(6), 1482–1493. https://doi.org/10.1037/0278-7393.22.6.1482

Brennan, S. E., & Schober, M. F. (2001). How listeners compensate for disfluencies in spontaneous speech. *Journal of Memory and Language*, *44*(2), 274–296. https://doi.org/10.1006/jmla.2000.2753

Christianson, K. (2016). When language comprehension goes wrong for the right reasons: Good-enough, underspecified, or shallow language processing. *Quarterly Journal of Experimental Psychology*, *69*(5), 817–828. https://doi.org/10.1080/17470218.2015.1134603

Christianson, K., Hollingworth, A., Halliwell, J. F., & Ferreira, F. (2001). Thematic roles assigned along the garden path linger. *Cognitive Psychology*, *42*(4), 368–407. https://doi.org/10.1006/cogp.2001.0752

Corley, M. (2010). Making predictions from speech with repairs: Evidence from eye movements. *Language and Cognitive Processes*, *25*(5), 706–727. https://doi.org/10.1080/01690960903512489

Corley, M., MacGregor, L. J., & Donaldson, D. I. (2007). It's the way that you, er, say it: Hesitations in speech affect language comprehension. *Cognition*, *105*(3), 658–668. https://doi.org/10.1016/j.cognition.2006.10.010

Engelhardt, P. E., Corley, M., Nigg, J. T., & Ferreira, F. (2010). The role of inhibition in the production of disfluencies. *Memory & Cognition*, *38*(5), 617–628. https://doi.org/10.3758/MC.38.5.617

Ferreira, F., Bailey, K. G., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science*, *11*(1), 11–15. https://doi.org/10.1111/1467-8721.00158

Ferreira, F., Lau, E. F., & Bailey, K. G. D. (2004). Disfluencies, language comprehension, and tree adjoining grammars. *Cognitive Science*, *28*(5), 721–749. https://doi.org/10.1207/s15516709cog2805_5

Ferreira, F., & Patson, N. D. (2007). The 'good enough'approach to language comprehension. *Language and Linguistics Compass*, *1*(1-2), 71–83. https://doi.org/10.1111/j.1749-818X.2007.00007.x

Fletcher, C. R. (1984). Markedness and topic continuity in discourse processing. *Journal of Verbal Learning and Verbal Behavior*, *23*(4), 487–493. https://doi.org/10.1016/S0022-5371(84)90309-8

Foraker, S., & McElree, B. (2007). The role of prominence in pronoun resolution: Active versus passive representations. *Journal of Memory and Language*, *56*(3), 357–383. https://doi.org/10.1016/j.jml.2006.07.004

Fox Tree, J. E. (1995). The effects of false starts and repetitions on the processing of subsequent words in spontaneous speech. *Journal of Memory and Language*, *34*(6), 709–738. https://doi.org/10.1006/jmla.1995.1032

Fukumura, K., Hyönä, J., & Scholfield, M. (2013). Gender affects semantic competition: The effect of gender in a non-gender marking language. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(4), 1012–1021. https://doi.org/10.1037/a0031215

Fukumura, K., & Van Gompel, R. P. G. (2010a). Choosing anaphoric expressions: Do people take into account likelihood of reference? *Journal of Memory and Language*, *62*(1), 52–66. https://doi.org/10.1016/j.jml.2009.09.001

Fukumura, K., & Van Gompel, R. P. G. (2011). The effects of animacy in the choice of referring expressions. *Language and Cognitive Processes*, *26*(10), 1472–1504. https://doi.org/10.1080/01690965.2010.506444

Fukumura, K., Van Gompel, R. P. G., & Pickering, M. J. (2010b). The use of visual context during the production of referring expressions. *Quarterly Journal of Experimental Psychology*, *63*(9), 1700–1715. https://doi.org/10.1080/17470210903490969

Goldman, A. I. (2006). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. Oxford University Press.

Gordon, P. C., Grosz, B. J., & Gilliom, L. A. (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, *17*(3), 311–347. https://doi.org/10.1207/s15516709cog1703_1

Grosz, B. J., Joshi, A. K., & Weinstein, S. (1995). Centering: A framework for modelling the local coherence of discourse. *Computational Linguistics*, *21*, 203–225.

Gundel, J. K., Hedberg, N., & Zacharski, R. (1993). Cognitive status and the form of anaphoric expressions in discourse. *Language*, *69*(2), 274–307. https://doi.org/10.2307/416535

Hofmeister, P. (2011). Representational complexity and memory retrieval in language comprehension. *Language and Cognitive Processes*, *26*(3), 376–405. https://doi.org/10.1080/01690965.2010.492642

Hofmeister, P., & Vasishth, S. (2014). Distinctiveness and encoding effects in online sentence comprehension. *Frontiers in Psychology*, *5*, 1237. https://doi.org/10.3389/fpsyg.2014.01237

Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, *59*(4), 434–446. https://doi.org/10.1016/j.jml.2007.11.007

Jaeger, T. F. (2010). Redundancy and reduction: Speakers manage syntactic information density. *Cognitive Psychology*, *61*(1), 23–62. https://doi.org/10.1016/j.cogpsych.2010.02.002

Karimi, H., & Ferreira, F. (2016). Good-enough linguistic representations and online cognitive equilibrium in language processing. *Quarterly Journal of Experimental Psychology*, *69*(5), 1013–1040. https://doi.org/10.1080/17470218.2015.1053951

Karimi, H., Fukumura, K., Ferreira, F., & Pickering, M. J. (2014). The effect of noun phrase length on the form of referring expressions. *Memory & Cognition*, *42*(6), 993–1009. https://doi.org/10.3758/s13421-014-0400-7

Karimi, H., Brothers, T., & Ferreira, F. (2019). Phonological versus semantic prediction in focus and repair constructions: No evidence for differential predictions. *Cognitive Psychology*, *112*, 25–47.

Keil, F. C. (1979). *Semantic and conceptual development: An ontological perspective*. Harvard University Press.

Kidd, C., White, K. S., & Aslin, R. N. (2011). Toddlers use speech disfluencies to predict speakers' referential intentions. *Developmental Science*, *14*(4), 925–934. https://doi.org/10.1111/j.1467-7687.2011.01049.x

Kleinschmidt, D. F., & Jaeger, T. F. (2011, June 19–24). A Bayesian belief updating model of phonetic recalibration and selective adaptation. In *Proceedings of the cognitive modeling and computational linguistics workshop at ACL* (pp, 10–19), Portland, Oregon, USA.

Lau, E. F., & Ferreira, F. (2005). Lingering effects of disfluent material on comprehension of garden path sentences. *Language and Cognitive Processes*, *20*(5), 633–666. https://doi.org/10.1080/01690960444000142

Levelt, W. J., & Cutler, A. (1983). Prosodic marking in speech repair. *Journal of Semantics*, *2*(2), 205–218. https://doi.org/10.1093/semant/2.2.205

Lowder, M. W., & Ferreira, F. (2016). Prediction in the processing of repair disfluencies: Evidence from the visual-World Paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(9), 1400–1416. https://doi.org/10.1037/xlm0000256

Lucas, M. M., Tanenhaus, M. K., & Carlson, G. N. (1990). Levels of representation in the interpretation of anaphoric reference and instrument inference. *Memory & Cognition*, *18*(6), 611–631. https://doi.org/10.3758/BF03197104

MacDonald, M. C., & MacWhinney, B. (1990). Measuring inhibition and facilitation from pronouns. *Journal of Memory and Language*, *29*(4), 469–492. https://doi.org/10.1016/0749-596X(90)90067-A

MacGregor, L. J., Corley, M., & Donaldson, D. I. (2010). Listening to the sound of silence: Disfluent silent pauses in speech have consequences for listeners. *Neuropsychologia*, *48*(14), 3982–3992. https://doi.org/10.1016/j.neuropsychologia.2010.09.024

Maxfield, N. D., Lyon, J. M., & Silliman, E. R. (2009). Disfluencies along the garden path: Brain electrophysiological evidence of disrupted sentence processing. *Brain and Language*, *111*(2), 86–100. https://doi.org/10.1016/j.bandl.2009.08.003

McElree, B. (2000). Sentence comprehension is mediated by content-addressable memory structures. *Journal of Psycholinguistic Research*, *29*(2), 111–123. https://doi.org/10.1023/A:1005184709695

McElree, B., Foraker, S., & Dyer, L. (2003). Memory structures that subserve sentence comprehension. *Journal of Memory and Language*, *48*(1), 67–91. https://doi.org/10.1016/S0749-596X(02)00515-6

Miyake, A., & Friedman, N. P. (2012). The nature and organization of individual differences in executive functions four general conclusions. *Current Directions in Psychological Science*, *21*(1), 8–14. https://doi.org/10.1177/0963721411429458

Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, *27*, 169–190. https://doi.org/10.1017/S0140525X04000056

Pickering, M. J., & Garrod, S. (2013). An integrated theory of language production and comprehension. *Behavioral and Brain Sciences*, *36*(04), 329–347. https://doi.org/10.1017/S0140525X12001495

Stevenson, R. J., Crawley, R. A., & Kleinman, D. (1994). Thematic roles, focus and the representation of events. *Language and Cognitive Processes*, *9*(4), 519–548. https://doi.org/10.1080/01690969408402130

Tesink, C. M., Petersson, K. M., Van Berkum, J. J., Brink, V. d., Buitelaar, D., & & Hagoort, J. K. (2009). Unification of speaker and meaning in language comprehension: An fMRI study. *Journal of Cognitive Neuroscience*, *21*(11), 2085–2099. https://doi.org/10.1162/jocn.2008.21161

Troyer, M., Hofmeister, P., & Kutas, M. (2016). Elaboration over a discourse facilitates retrieval in sentence processing. *Frontiers in Psychology*, *7*, 374. https://doi.org/10.3389/fpsyg.2016.00374

Van Berkum, J. J., Brink, V. d., Tesink, D., Kos, C. M., & & Hagoort, M. (2008). The neural integration of speaker and message. *Journal of Cognitive Neuroscience*, *20*(4), 580–591. https://doi.org/10.1162/jocn.2008.20054

Van Dyke, J.A, & McElree, B. (2011). Cue-dependent interference in comprehension. *Journal of Memory and Language*, *65*(3), 247–263.

von der Malsburg, T., Poppels, T., & Levy, R. P. (2018). *Implicit gender bias in linguistic descriptions for expected events: The cases of the 2016 US and 2017 UK election*.

Yildirim, I., Degen, J., Tanenhaus, M. K., & Jaeger, T. F. (2016). Talker-specificity and adaptation in quantifier interpretation. *Journal of Memory and Language*, *87*, 128–143. https://doi.org/10.1016/j.jml.2015.08.003