



Grammaticality Illusions in the Real-Time Processing of Backward Noun Phrase Ellipsis*

Younseo Ha (Sungkyunkwan University) · Michale Frazier (Leipzig University) · Nayoun Kim (Sungkyunkwan University)



This is an open-access article distributed under the terms of the Creative Commons License, which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

Received: January 5, 2025

Revised: January 24, 2025

Accepted: February 10, 2025

Younseo Ha (First author)
Master's student, Department of English Language and Literature, Sungkyunkwan University

Michale Frazier (Co-author)
Postdoctoral Researcher, Institution for Linguistics, Leipzig University

Nayoun Kim (Corresponding author)
Professor, Department of English Language and Literature, Sungkyunkwan University

*This work was supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2023S1A5A8079483).

ABSTRACT

Ha, Younseo, Michael Frazier, and Nayoun Kim. 2025. Grammaticality illusions in the real-time processing of backward noun phrase ellipsis. *Korean Journal of English Language and Linguistics* 25, 277-288.

This study investigated the processing of backward Noun Phrase Ellipsis (NPE), as it provides an interesting staging ground for the question of what kind of content in the antecedent is recovered at the NPE-site (Kim 2019, Kim et al 2019, Kim and Yoshida submitted) where the ellipsis-site occurs prior to the antecedent. Building on previous findings on the illusion of grammaticality in non-elliptical constructions (Wagers et al. 2009) and NPE-contexts (Kim 2019, Kim et al. 2019, Kim and Frazier 2022), we examined whether similar patterns emerge in the context of backward NPE. If the agreement attraction effect and similar grammatical asymmetries arise in backward NPE, this could serve as an argument against simple memory retrieval models of agreement attraction that do not take structural information into account (c.f. Martin and McElree 2008, 2009, 2011). Our results indicate that an illusion of grammaticality is present even during real-time processing of backward NPE, suggesting that structural information, alongside morphological and semantic content, is retrieved when resolving the ellipsis (Kim 2019, Kim et al. 2019, Kim et al. 2020, Lago et al. 2015, Tanner et al. 2014, Wagers et al. 2009). Furthermore, we observed an agreement attraction effect in ungrammatical conditions with no effect in grammatical conditions (Wagers et al. 2009), which is in line with the hypothesis that the reader participates in an active and sustained search for the antecedent even when the first attempt to recover the NPE-site fails—in this case, owing to incongruence in features between the head noun and the verb with respect to number (Kim 2019, Kim et al. 2019)—resembling the patterns observed in resolving cataphoric dependencies (Giskes and Kush 2021, Kazanina et al. 2007).

KEYWORDS

backward noun phrase ellipsis, structural information, real-time processing, cue-based retrieval mechanism

1. Introduction

Understanding sentences that involve ellipsis poses a challenge to the reader, partly because the reader needs to compute the meaning of the sentence at a particular time-point in the series of linguistic inputs when there is yet no overt content that maps into the interpretation of the elided constituent. Consider Example (1), which involves Noun Phrase Ellipsis (NPE).

- (1) Mary's key to the cabinet is under the bed and Thomas's [_{NPE-site} ~~key to the cabinet~~] is on the desk.

In the second conjunct, a span of linguistic material is omitted between the possessive-marked Determiner Phrase (DP; *Thomas's*) and the verb (*is*). Because *Thomas's* and *is* cannot form a constituent, the parser should anticipate that a construction involving an ellipsis is present. Throughout the paper, we identify the position of the omitted NP as the *NPE-site*, and to the linguistic information that supplies the proper content of the NPE-site as the *antecedent* (Kim 2019, Kim et al. 2019, Merchant 2001, Yoshida et al. 2013). The interpretation of anaphoric elements such as NPE depends on their antecedents. Thus, it is plausible that an online sentence processing mechanism that handles anaphoric elements accesses their antecedent(s) in real-time.

The processing of NPE, as presently understood, unfolds in following manner. The parser needs to (i) first register the occurrence of a non-overt Noun Phrase (NP) in its input stream; (ii) inspect the morphosyntactic and possibly semantic information in the left-context of the linguistic span it has so far processed; (iii) search its memory store for information triggering retrieval of a constituent that could potentially serve the role as an antecedent for the elided segment of the structural input; and (iv) link the NPE-site to the antecedent. In the general case, the antecedent will have already been processed by the time the ellipsis-site is realized; thus, understanding an ellipsis involves recovering already-processed materials in memory. This suggests that the processing of NPE is fundamentally reliant on memory-retrieval mechanisms (Lewis and Vasishth 2005, Lewis et al. 2006, Martin and McElree 2008). Thus, a long-standing question about the processing of ellipses in general is “What types of information pertaining to the antecedent are recovered when the ellipsis is identified?” (Frazier and Clifton 2001, Kim 2019, Kim et al. 2019, Kim and Yoshida submitted, Martin and McElree 2008, 2009, 2011).

Ongoing debates on the processing of ellipses concern the recovery of linguistic content at the ellipsis-site. For example, in (1), the parser can, in principle, access structural information concerning the antecedent when identifying the NPE-site (Frazier and Clifton 2001, Kim 2019, Kim et al. 2019, Kim et al. 2020, Kim and Frazier 2022, Kim and Yoshida submitted, Murphy 1985, Yoshida et al. 2013). If the structural information is recovered, the NPE-site should be treated in a parallel manner comparable to a non-elliptical overt NP within the same structure. If this is the case, we expect that the entire informational content of the antecedent constituent, including structural information such as [+Specifier TP] and [+ Sister of the P], will be accessed and retrieved at the NPE-site (e.g., [[[_{head-N} *key*]]_{PP} *to* [_{DP} *the* [_{NP} *cabinet*]]]) (Kim 2019, Kim et al. 2019, Kim et al. 2020).

Another possibility, however, is that featural information without reference to syntactic structure could be accessed at the NPE-site. In this case, when recognizing the ellipsis, either the head noun's information or that of the modifier could be expected to be accessed, with neither taking priority. Yet a third possibility is that the parser accesses the head content without a modifier, as it is structurally and informationally more prominent, governing the meaning of the entire phrase (Kim 2019, Kim et al. 2019). Lastly, because recent items are less susceptible to interference during the retrieval process, the embedded noun adjacent to the NPE-site (e.g., *the cabinet* in (1)) may be erroneously retrieved as an antecedent.

Kim et al. (2019) leveraged an established processing phenomenon—the *illusion of grammaticality* paradigm (i.e., the agreement attraction effect in ungrammatical sentences; Lago et al. 2015, Tanner et al. 2014, Wagers et al. 2009) to investigate what kinds of antecedent-related information can be accessed in the NPE-contexts. Previous studies have provided evidence that already-processed materials are represented by different type of linguistic features (morphological, syntactic, semantic, lexical, etc.) and are utilized as retrieval cues for the antecedent retrieval process. If the previously processed items match the features related to the dependent element (e.g., the verb in the subject-verb dependency), then the items are reactivated in a parallel fashion by a content-addressable memory mechanism (Lewis and Vasishth 2005, Lewis et al. 2006, Martin and McElree 2008). Moreover, partially-matched features present on an intervening element positioned between a dependent element (e.g., a verb in subject-verb agreement) and its controlling element (e.g., the subject in such a dependency) can sometimes cause an ungrammatical sentence to be judged as acceptable.

For example, although both (2a) and (2b) are ungrammatical, (2b), with the [+plural] feature on *cells* matching the agreement form of the verb *were*, is known to be read faster and rated more acceptable than (2a), which lacks such a partial match with the intervenor (Lewis and Vasishth 2005, Wagers et al. 2009).

- (2) a. *The key to the cell were rusty from many years of disuse.
b. *The key to the cells were rusty from many years of disuse.

(Wagers et al. 2009, p. 221)

However, this grammatical asymmetry is detected in ungrammatical sentences yet not in grammatical ones; this is described as the *illusion of grammaticality* (Lago et al. 2015, Tanner et al. 2014, Wagers et al. 2009). Kim (2019) and Kim et al. (2019) used the illusion of grammaticality to investigate what is accessed at the NPE-site to derive the correct meaning of NPE. By the time the NPE-site is processed, the antecedent has already been read and, *ex hypothesi*, retained in memory (Kim 2019; Kim et al. 2019). Therefore, building the contents within the NPE-site necessitates triggering the retrieval of a constituent that could potentially serve as an antecedent for the elided portion of the structural input in memory (Kim et al. 2019, Martin and McElree 2008, 2009, 2011).

- (3) a. Mary's key to the cabinet is under the bed and Thomas's is on the desk.
b. *Mary's key to the cabinet are under the bed and Thomas's are on the desk.
c. Mary's key to the cabinets is under the bed and Thomas's is on the desk.
d. *Mary's key to the cabinets are under the bed and Thomas's are on the desk.

(Kim et al. 2019, p. 192, a modified version)

Kim et al. (2019) predicted that, when structural details pertaining to the antecedent are reactivated at the NPE-site, detailed information associated with the head noun and its modifier (e.g., [+Specifier TP] and [+ Sister of the P]) should be recovered along with it (Kim and Frazier 2022). Specifically, the head noun should be accessed initially, and in cases where its number mismatches that of the verb, the reader can search for an alternative noun that matches the number (Kim et al. 2019, Lago et al. 2015, Tanner et al. 2014, Wagers et al. 2009). Consequently, (3d) should have shown significantly faster reading times than (3b) in the verb region, whereas no discrepancies in reading times were expected for grammatical sentences. If only the most prominent element within the phrase (i.e., information about the head) is recovered at the NPE-site, then no illusion of grammaticality should arise, and grammatical sentences should simply be processed at a faster rate than their ungrammatical counterparts.

In a syntactically unconstrained content-addressable memory search, structural information about the antecedent is not maintained in memory (Martin and McElree, 2008, 2009, 2011). Thus, already-processed items will be identified in terms of morphological and lexical features including number, gender, and animacy, and the item held in active memory that best corresponds to the retrieval cue(s) will be retrieved. In this case, neither the head noun nor the modifier should have priority over the other, and the representations of the antecedent are accessed directly in a content-addressable feature-matching manner (Martin and McElree, 2008, 2009, 2011). This predicts a similar illusion in grammatical sentences in the form of a similarity-based interference effect (Gordon et al. 2001, Jäger et al. 2017, Van Dyke and Lewis 2003, Van Dyke and McElree 2011). In other words, if both nouns (*key* and *cabinet/cabinets*) are accessed at the NPE-site, the reader may experience an interference effect due to incorrect or partial encoding of the antecedent or elided constituent in memory. Thus, (3a) should be read slower than (3c) because of the similarity-based interference effect (or, equivalently, the temporal decay of the active memory representations).

Across a series of experiments, Kim (2019) and Kim et al. (2019) identified an illusion of grammaticality in NPE-contexts akin to that observed in non-elliptical constructions. This follows from how the information of the antecedent is hypothesized to be stored in memory, namely that the information concerned with the antecedent is stored with sensitivity to structural information. They further compared NPE processing to other nominal anaphoric elements such as anaphoric *one* (e.g., *the rusty key and that fancy one*), demonstrating that these non-elliptical nominal anaphora constructions are processed differently from NPE. The crucial distinction lies in whether covert grammatical structure is involved: NPE involves an unpronounced grammatical structure in the ellipsis-site (Frazier and Clifton 2001, Kim 2019, Kim et al. 2019, Murphy 1985); thus, the reader attempts to recover this unpronounced structure for the antecedent (Kim et al. 2019, Kim and Yoshida submitted).

In this study, we investigated the processing of backward NPE, which serves as a valuable staging ground to determine whether the illusion of grammaticality observed in NPE contexts also arises when the ellipsis-site precedes its antecedent. If an agreement attraction effect and similar grammatical asymmetries are observed, this would be an argument against simple memory retrieval models of agreement attraction that do not consider structural information.

2. Processing Backward NPE

In an NPE, the dependent element (the ellipsis) can occasionally precede the controlling element (its antecedent), as shown in (4). These are referred to as backward NPE. As with standard NPE, the meaning of the elided part associated with the NP is contingent upon its antecedents. For instance, in (4), the NPE-site should correspond to *song from the closet*, rather than *song* from an unspecified location.

- (4) Although Momma's [_{NPE-site} ~~song from the closet~~] was comforting, Billy's song from the closet still terrified the kids.

However, in backward NPE contexts, the NPE-site precedes the antecedent, contrasting with typical NPE processing. Consequently, we expect that encountering an ellipsis will trigger a forward search for an antecedent. This difference raises the question of whether the retrieval process for the antecedent in backward NPE follows the same mechanism as that in standard NPE.

To interpret the backward NPE correctly, the parser must follow a series of steps. First, the parser must recognize the presence of the NPE when encountering the possessive NP (*Momma's*) and verb (*were*) in the first conjunct. This is because these two words cannot form a constituent; thus, the combination of the possessive-marked NP and the verb indicates to the parser that the span of linguistic material has gone unpronounced, that is, that something is elided. This triggers the next step: searching for an antecedent that can supply the missing content at the NPE-site. As the antecedent appears later in the sentence, the parser must retrieve it after the NPE-site has been identified. Once the correct antecedent is located, it must be linked to the NPE-site for a complete interpretation. Similar to the questions raised during the processing of NPE, we asked analogous questions regarding the processing of backward NPE. When an NPE is identified, what type of information is conveyed by the antecedent?

Under the assumption of an unconstrained cue-based retrieval mechanism (Martin and McElree 2008, 2009, 2011), the processing of ellipses does not require the recovery of structural information pertaining to the antecedent in memory; rather, it merely involves matching the retrieval cue (the immediately incoming linguistic stimulus) with items that overlap with it in lexical and morphological features. Conversely, under the view that the syntactic structure concerned with the antecedent is directly copied to the ellipsis-site (Frazier and Clifton 2001, Murphy 1985), structural information such as constituency and headedness should also be recovered at the ellipsis-site, potentially also influencing the antecedent-retrieval process (Kim 2019, Kim et al. 2019, Kim and Frazier 2022).

Thus, these two proposals regarding what kind of information is recovered at the NPE-site make divergent predictions when tested with the illusion of grammaticality. By employing this particular construction of backward NPE and drawing from the well-known phenomenon of the illusion of grammaticality as probes, we investigate (i) whether readers recover the structural information concerned with the antecedent even in cases where a forward search for the antecedent is required, and (ii) whether readers consistently search for the antecedent after they realize that the NPE cannot be interpreted because of the wrong antecedent.

Consider Example (5), in which the *Number of the Embedded Noun* (Plural or Singular feature concerned with the embedded noun: *closet* vs. *closets*) and *Grammaticality* (Grammatical or Ungrammatical depending on the match between the head noun and verb with regard to number) are manipulated (see Kim 2019, Kim et al. 2019). When the possessive DP (*Momma's*) and the verb (*were* or *was*) are encountered, the NPE-site is identified, prompting a forward search for the antecedent downstream. Predictions of processing difficulty vary depending on the information recovered and how the antecedent is linked to the NPE-site.

- (5) a. *Although Momma's were comforting, Billy's song from the closets still terrified the kids.
 b. *Although Momma's were comforting, Billy's song from the closet still terrified the kids.
 c. Although Momma's was comforting, Billy's song from the closets still terrified the kids.
 d. Although Momma's was comforting, Billy's song from the closet still terrified the kids.

First, upon the appearance of the head noun (*song*), we expect to see a main effect of *Grammaticality* with ungrammatical sentences anticipated to take longer to process than grammatical sentences (Kim 2019, Kim et al. 2019). However, different predictions arise for the ungrammatical conditions (5a, 5b) depending on the type of information retrieved at the backward NPE-site. When the head noun (*song*) and verb are discordant in number, the parser may instead search for an alternative noun that conforms to the verb with respect to the number to rescue the number mis-agreement (Lago et al. 2015, Kim 2019, Kim et al. 2019, Tanner et al. 2014, Wagers et al. 2009). If the structural details pertaining to the antecedent constituent are recovered, meaning that the entire DP with the head and complement of PP information is retrieved (Kim 2019, Kim et al. 2019), and if the parser consistently searches for a potential antecedent (Giskes and Kush 2021) even when the first attempt to recover the NPE-site

fails owing to the number incongruity between the head noun and the verb (Kim 2019, Kim et al. 2019), we expect to see a facilitation effect occurring at the DP inside the PP (*closets*) in (5a) compared with (5b).

Furthermore, if structural information is recovered in the backward NPE, then, in grammatical sentences, a plural embedded noun is predicted to be read slower than a singular embedded noun because the plural embedded noun is not expected, or plural embedded nouns are predicted to be read similarly to singular embedded nouns. Thus, (5c) should be read slower than (5d), or (5c) should be read similarly to (5d) at the embedded noun. Conversely, if retrieval is based only on lexical (semantic and morphological) features, as predicted by structurally unconstrained cue-based retrieval (Martin and McElree 2008, 2009, 2011), similarity-based interference is expected (Kim 2019, Kim et al., 2019, Kim and Frazier 2022). Specifically, in grammatical conditions, (5d) should be read slower than (5c) because of the interference from the singular embedded noun (*closet*), which is consistent with previous findings on similarity-based interference (Jäger et al. 2017, Lewis and Vasishth 2005, Lewis et al. 2006, Van Dyke 2007, Van Dyke and McElree 2011).

3. The Experiment

We conducted a self-paced reading experiment to investigate (i) whether readers recover structural information within an antecedent even when the potential antecedent appears to the right of the NPE-site (Kim 2019; Kim et al. 2019; Kim and Yoshida submitted) and (ii) whether readers consistently search for the antecedent after realizing that the NPE cannot be interpreted because of a mismatch with the wrong antecedent (Giskes and Kush 2021).

3.1 Participants

Seventy-two native English speakers from the US, Canada, and the UK were recruited through Prolific (<https://www.prolific.com/>). All participants were native speakers of English, with ages ranging from 20 to 55 years. They all learned English before the age of 4 years and have predominantly used it as their first and dominant language since the age of 5 years. All participants held bachelor's degrees.

3.2 Materials

The experimental items comprised 16 sentence sets organized according to a 2×2 within-subjects factorial design, whereby the *Number of the Embedded Noun* (Plural or Singular: *closet* vs. *closets*) and *Grammaticality* (Grammatical or Ungrammatical, depending on the featural (number) match between the head and verb: *was* vs. *were*) served as independent factors. Sample stimuli are presented in Table 1. Each item was divided into the regions presented to the participants. Across items, the possessive marked noun (*Momma's*) always preceded the verb (*was* or *were*) in the subordinate clause, and the head noun always occupied the sixth presentation region and ninth embedded noun region. All head nouns contained singular features because grammaticality served as an independent factor. Various subordinating conjunctions (e.g., *although*, *before*, *after*) were used across items. In addition to the 16 experimental items, 30 filler sentences with comparable complexity were included and created in a Latin Square manner. These fillers were unrelated to the current manipulation and contained wh-dependencies that manipulated adjuncts rather than targeting any elliptical constructions.

Table 1. Example Stimuli (Critical Region 1 indicated in blue and Critical Region 2 in green)

	Local (Embedded) Noun: Plural	Local (Embedded) Noun: Singular
Ungrammatical	a. Although Momma’s were comforting, Billy’s song from the closets still terrified the kids.	b. Although Momma’s were comforting, Billy’s song from the closet still terrified the kids.
Grammatical	c. Although Momma’s was comforting, Billy’s song from the closets still terrified the kids.	d. Although Momma’s was comforting, Billy’s song from the closet still terrified the kids.

3.3 Procedures

The experiment was executed via the online experiment platform, PC Ixer (Zehr and Schwarz 2018). The experiment was conducted by participants on their personal laptops through a link shared by Prolific. Sentences in the experiment were displayed in a self-paced, word-by-word, non-cumulative manner. Participants were directed to press the space bar to move to continue to the next region of the sentence, and the reading time spent in each region was automatically recorded. They were guided to read the sentences at a comfortable pace. They were also provided a practice section to make themselves comfortable with the procedure and methodology of the experiment. Participants filled out informed consent forms prior to their involvement in the experiment and earned \$11 per hour as compensation. The experiment took approximately 15–20 minutes.

3.4 Analysis

All reading times below 80ms or greater than 2000ms were discarded prior to statistical analysis. Due to violations of the normality assumption, reading times were log-transformed before analysis. A sum-contrast coded linear mixed effects model (Baayen et al. 2008, Bates et al. 2015) with maximal convergence (Barr et al. 2013, Bates et al. 2015) was employed for statistical analysis with *Number of the Embedded Noun* (Singular vs. Plural) and *Grammaticality* (Grammatical vs. Ungrammatical) as fixed factors (Kim 2019, Kim et al. 2019), considering participants and items as random intercepts, along with by-participant and by-item random slopes.

3.5 Results

The average log reading times for the spillover region of Critical Region 1 (head noun: *song*), second spillover region of Critical Region 1, spillover region of Critical Region 2 (embedded noun: *closet, closets*), and second spillover region of Critical Region 2 are shown in Figures 1–3, respectively.

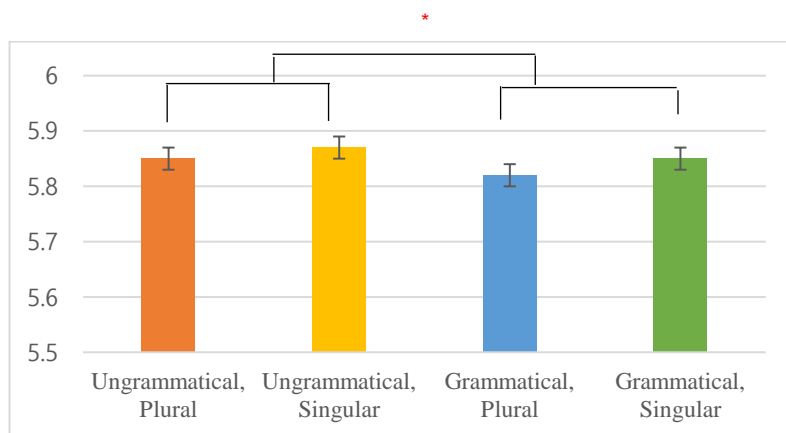


Figure 1. Mean Reading Times at the Spillover Region of Critical Region 1 (from)

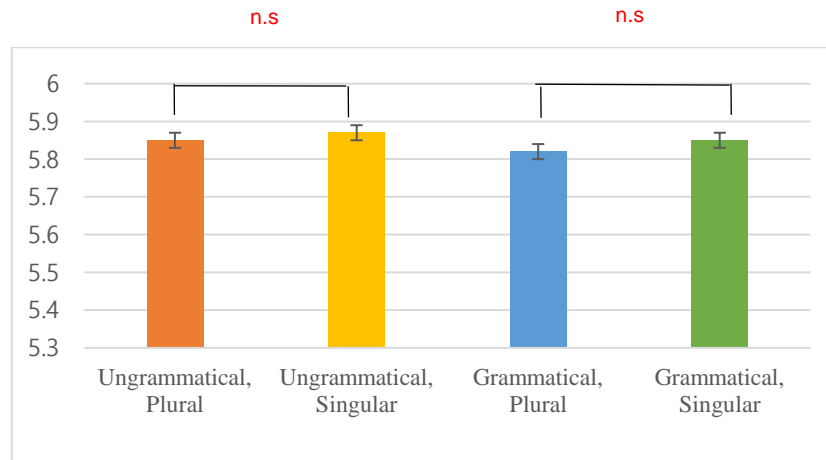


Figure 2. Mean Reading Times at the First Spillover Region of Critical Region 2 (*still*)

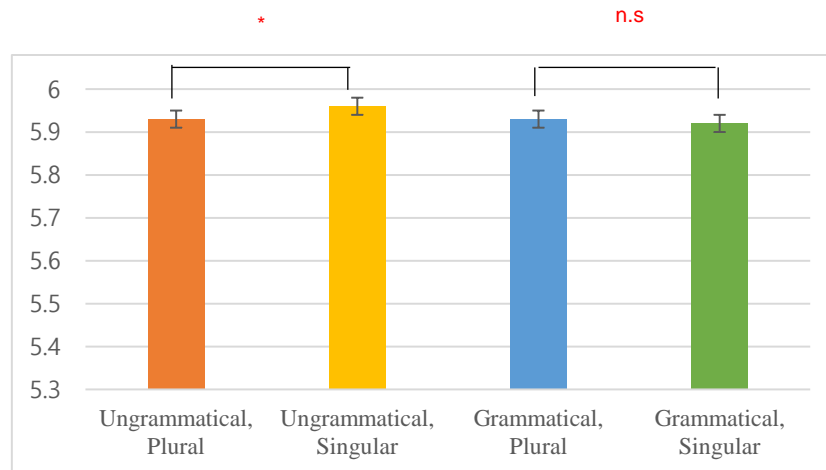


Figure 3. Mean Reading Times at the Second Spillover Region of Critical Region 2 (*terrified*)

Two critical regions were analyzed. The first critical region is the head noun (*song*), and the second critical region is the embedded noun (*closet* and *closets*). At the first critical region, we found no main effect of *Grammaticality* ($\beta=0.00$, $SE=0.01$, $t=0.20$), *Number of the Embedded Noun* ($\beta=0.02$, $SE=0.01$, $t=1.46$), nor an interaction between the two ($\beta=0.03$, $SE=0.03$, $t=1.19$). At the spillover region followed by the first critical region (*from*), we did observe a main effect of *Grammaticality* ($\beta=-0.03$, $SE=0.01$, $t=-2.36$) where ungrammatical sentences were processed at a significantly slower rate compared to their grammatical equivalents. But we observed no main effect of *Number of the Embedded Noun* ($\beta=0.02$, $SE=0.01$, $t=1.14$) nor an interaction between the two ($\beta=0.03$, $SE=0.03$, $t=0.91$). At the second spillover region of the first critical region (*the*), we continued to observe a marginal main effect of *Grammaticality* ($\beta=-0.02$, $SE=0.01$, $t=-1.89$) where ungrammatical sentences took considerably more time to read than grammatical counterparts. But we observed no main effect of *Number of the Embedded Noun* ($\beta=0.02$, $SE=0.01$, $t=1.76$) nor an interaction between the two ($\beta=-0.00$, $SE=0.03$, $t=-0.04$).

The second critical region was the embedded noun (*closet/closets*). At the second critical region (*closet/closets*), we found a main effect of *Grammaticality* ($\beta=-0.03$, $SE=0.01$, $t=-2.01$) but no main effect of *Number of the Embedded Noun* ($\beta=0.00$, $SE=0.01$, $t=0.50$) nor an interaction between the two ($\beta=0.05$, $SE=0.03$, $t=1.61$). At the spillover region followed by the second critical region (*still*), we did not observe a main effect of *Grammaticality*

($\beta=-0.01$, $SE=0.01$, $t=-1.01$), *Number of the Embedded Noun* ($\beta=0.01$, $SE=0.01$, $t=0.94$) nor an interaction between *Grammaticality* and *Number of the Embedded Noun* ($\beta=0.04$, $SE=0.03$, $t=1.61$).

At the second spillover region of the second critical region (*terrified*), we did not observe a main effect of *Grammaticality* ($\beta=-0.02$, $SE=0.01$, $t=-1.38$) nor a main effect of *Number of the Embedded Noun* ($\beta=0.01$, $SE=0.01$, $t=0.76$), yet we observed an interaction between the two ($\beta=-0.05$, $SE=0.03$, $t=-2.07$) which can be characterized as an agreement attraction (Wager et al. 2009). Furthermore, we conducted follow-up analyses to understand the nature of an interaction between *Grammaticality* and the *Number of the Embedded Noun*. The results of the follow-up analysis revealed that there was a main effect of *Number of the Embedded Noun* in ungrammatical condition ($\beta=0.04$, $SE=0.02$, $t=1.91$) such that plural embedded noun exhibited faster reading times than the singular one in ungrammatical conditions but not under grammatical conditions ($\beta=-0.02$, $SE=0.02$, $t=-0.95$), suggesting grammatical asymmetries are also observed during the processing of backward NPE (Kim 2019, Kim et al. 2019, 2020).

At the third spillover region followed by Critical Region 2 (*the*), we did not observe a main effect of *Grammaticality* ($\beta=0.00$, $SE=0.01$, $t=0.13$), but observed a main effect of *Number of the Embedded Noun* ($\beta=0.03$, $SE=0.01$, $t=2.43$). We also continued to observe a marginal interaction between *Grammaticality* and *Number of the Embedded Noun* ($\beta=-0.04$, $SE=0.03$, $t=-1.75$). We conducted follow-up analyses to understand the nature of the interaction between *Grammaticality* and *Number of the Embedded Noun*. The results revealed a main effect of *Number of the Embedded Noun* in ungrammatical conditions ($\beta=0.05$, $SE=0.02$, $t=3.01$), wherein the plural embedded noun exhibited faster reading times than the singular one in ungrammatical conditions but not in grammatical contexts ($\beta=0.01$, $SE=0.02$, $t=0.45$).

4. Discussion and Conclusion

A point of contention regarding the processing of ellipses is whether structural details pertaining to the antecedent are reactivated in instances where the ellipsis-site is encountered (Kim 2019, Kim et al. 2019, Kim and Frazier 2022, Kim and Yoshida submitted, Martin and McElree 2008, 2009, 2011). We investigated (i) whether readers recover the structural information concerned with the antecedent along with the morphological and semantic content, and (ii) whether readers persistently find the grammatically sanctioned antecedent for ellipsis even when the first chance of identifying the potential antecedent is foiled (Giskes and Kush 2021, Kazanina et al. 2007), by examining the processing of backward NPE.

Recent studies have demonstrated that all information comprising the content of the head and modifier of the antecedent is accessed and recovered at the NPE-site (Kim 2019, Kim et al. 2019, Kim and Yoshida submitted). However, it is also plausible that only featural information without reference to syntactic structure can be accessed under a syntactically unconstrained content-addressable search if the structural information about the antecedent is not encoded in memory (Kim et al. 2019; Martin and McElree 2008, 2009, 2011). Leveraging the backward NPE construction and the *illusion of grammaticality* paradigm (Wagers et al. 2009), we examined whether agreement attraction in the NPE context could also be observed in cases wherein the ellipsis-site occurs prior to the antecedent. If an agreement attraction effect and similar grammatical asymmetries are observed during the processing of backward NPE, this would serve as an argument against simple memory retrieval models of agreement attraction that do not consider structural information (c.f. Martin and McElree 2008, 2009, 2011).

The overall pattern of effects observed in this study is consistent with a cue-based model of antecedent retrieval but not a syntactically unconstrained one (Martin and McElree 2008, 2009, 2011). Specifically, grammaticality

manipulation at the position of the head noun had a significant main effect such that grammatical sentences were read significantly faster than their ungrammatical counterparts. At the position of the embedded noun, we observed an interaction such that the plural + ungrammatical conditions were faster than their singular counterparts; however, there were no differences between the singular + grammatical and plural + grammatical conditions, resembling a standard agreement attraction effect (Kim 2019, Kim et al. 2019, Lago et al. 2015, Tanner et al. 2014, Wagers et al. 2009), generally regarded as a signature of cue-based retrieval mechanisms.

The pattern in ungrammatical conditions represents the typical illusion of grammaticality found in NPE. The facilitation effect that we observe when the embedded NP is plural suggests that the retrieval mechanism reactivates the entire antecedent NP at the point where the syntactic structure pertaining to the antecedent is retrieved, not exclusively the morphological or semantic features pertinent to the head noun (Kim et al. 2019, Kim and Yoshida submitted, Murphy 1985, Yoshida et al. 2013).

Overall, this pattern of results indicates that the complete syntactic structure associated with the antecedent phrase is retrieved, including nominal modifiers: if the bare head noun alone was accessed, we would anticipate an absence of facilitatory effects in the ungrammatical and grammatical conditions. If only the nouns associated with the antecedent were accessed, without the structurally most prominent noun taking priority over the others, we would anticipate agreement attraction effects to manifest in all circumstances – in both grammatical and ungrammatical conditions (Kim 2019, Kim et al. 2019). The observed pattern of interference from the embedded NP in ungrammatical conditions is evidence that the entire complex NP of the antecedent is re-accessed at the point of antecedent retrieval. This implies that its component nominals can participate in calculating antecedent competition (Kim 2019, Kim et al. 2019, Lago et al. 2015, Tanner et al. 2014, Wagers et al. 2009), thereby inducing interference effects when no entirely grammatical antecedents can be retrieved.

Finally, the results of the experiment suggest that readers attempt to locate the antecedent at the earliest possible retrieval site, in a way comparable to the processing of wh-filler gap dependencies (Stowe 1986). When processing these dependencies, readers attempt to postulate the earliest gap in a grammatically sanctioned position immediately after the wh-filler is identified and associate it with the gap (Stowe 1986). Similarly, once the NPE-site is located and identified in the first conjunct, readers initiate a search for a grammatically permissible antecedent in the earliest possible position, which crucially involves a forward search for the antecedent. Furthermore, the interaction between *Grammaticality* and *Number of the Embedded Noun* is consistent with the hypothesis that the reader engages in an active and sustained search for the antecedent even when the first attempt to recover the NPE-site fails, in this case owing to incongruence in number. This pattern resembles that observed in real-time processing of cataphoric dependencies (Giskes and Kush 2021, Kazanina et al. 2007).

Backward ellipsis is a promising domain for investigating antecedent retrieval mechanisms because it allows the retrieval component of the ellipsis resolution process to be (at least partially) dissociated from the specific mechanisms of memory decay. While our study addresses only the backward NPE construction, future research should investigate whether this overall pattern of results—evidence of a content-addressable memory system with sensitivity to the syntactic structure of the retrieved material—can be replicated in other backward ellipsis constructions such as backward verb-phrase ellipsis (VPE) or backwards sluicing. The agreement properties of nominals are unusually conducive to the study of backwards ellipsis, as they facilitate an agreement-mismatch paradigm; however, the stringent licensing requirements of nominals, requiring not only syntactic selection but also thematic-role assignment and Case, could cause them to be represented on-line by the antecedent retrieval system with “richer” syntactic representations than phrases of other categories, which do not have these additional licensing requirements. VPE is a particularly promising construction for this purpose as the agreement dependency between the overt subject and the elided verb also permits the construction of gender-mismatch paradigms. If the

sensitivity to syntactic structure observed here is due to thematic or Case features being retrieved as part of the nominal rather than due to constituency structure, we would not expect the same pattern to be obtained in the backward ellipsis of non-nominal categories.

References

- Baayen, R. H., D. J. Davidson and D. M. Bates. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59, 390-412.
- Barr, D. J., R. Levy, C. Scheepers and H. J. Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language* 68(3), 255-278.
- Bates, D., M. Mächler, B. Bolker, and S. Walker. 2015. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67(1), 1-48.
- Frazier, L. and C. Clifton. 2001. Parsing coordinates and ellipsis: Copy α . *Syntax* 4(1), 1-22.
- Giskes, A. and D. Kush. 2021. Processing cataphors: Active antecedent search is persistent. *Memory & Cognition* 49(7), 1370-1386.
- Gordon, P. C., R. Hendrick and M. Johnson. 2001. Memory interference during language processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27(6), 1411-1423.
- Jäger, L. A., F. Engelmann and S. Vasisht. 2017. Similarity-based interference in sentence comprehension: Literature review and Bayesian meta-analysis. *Journal of Memory and Language* 94, 316-339.
- Kazanina, N., E. F. Lau, M. Lieberman, M. Yoshida and C. Phillips. 2007. The effect of syntactic constraints on the processing of backwards anaphora. *Journal of Memory and Language* 56(3), 384-409.
- Kim, N. 2019. *Hold, Release, and Retrieve: The Study of Wh-Filler-Gap Dependencies and Ellipsis*, Doctoral dissertation, Northwestern University.
- Kim, N., L. Brehm and M. Yoshida. 2019. The online processing of noun phrase ellipsis and mechanisms of antecedent retrieval. *Language, Cognition and Neuroscience* 34(2), 190-213.
- Kim, N., L. Brehm, P. Sturt and M. Yoshida. 2020. How long can you hold the filler: Maintenance and retrieval. *Language, Cognition and Neuroscience* 35(1), 17-42.
- Kim, N. and M. Frazier. 2022. The effect of antecedent complexity in the processing of noun phrase ellipsis. *Discourse and Cognition* 29(4), 17-35.
- Kim, N. and M. Yoshida. submitted. Mechanisms working behind the antecedent retrieval.
- Lago, S., D. E. Shalom, M. Sigman, E. F. Lau and C. Phillips. 2015. Agreement attraction in Spanish comprehension. *Journal of Memory and Language* 82, 133-149.
- Lewis, R. L. and S. Vasisht. 2005. An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science* 29(3), 375-419.
- Martin, A. E. and B. McElree. 2008. A content-addressable pointer mechanism underlies comprehension of verb-phrase ellipsis. *Journal of Memory and Language* 58(3), 879-906.
- Martin, A. E. and B. McElree. 2009. Memory operations that support language comprehension: Evidence from verb-phrase ellipsis. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35(5), 1231.
- Martin, A. E. and B. McElree. 2011. Direct-access retrieval during sentence comprehension: Evidence from sluicing. *Journal of Memory and Language* 64(4), 327-343.
- Merchant, J. 2001. *The Syntax of Silence: Sluicing, Islands, and the Theory of Ellipsis*. Oxford University Press.
- Murphy, G. L. 1985. Processes of understanding anaphora. *Journal of Memory and Language* 24(3), 290-303.

- Stowe, L. A. 1986. Parsing WH-constructions: Evidence for on-line gap location. *Language and Cognitive Processes* 1(3), 227-245.
- Tanner, D., J. Nicol, and L. Brehm. 2014. The time-course of feature interference in agreement comprehension: Multiple mechanisms and asymmetrical attraction. *Journal of Memory and Language* 76, 195-215.
- Van Dyke, J. A. 2007. Interference effects from grammatically unavailable constituents during sentence processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 33(2), 407-430.
- Van Dyke, J. A. and R. L. Lewis. 2003. Distinguishing effects of structure and decay on attachment and repair: A cue-based parsing account of recovery from misanalyzed ambiguities. *Journal of Memory and Language* 49(3), 285-316.
- Van Dyke, J. A. and B. McElree. 2011. Cue-dependent interference in comprehension. *Journal of Memory and Language* 65(3), 247-263.
- Vasishth, S. and R. L. Lewis. 2006. Argument-head distance and processing complexity: Explaining both locality and antilocality effects. *Language*, 767-794.
- Wagers, M. W., E. F. Lau and C. Phillips. 2009. Agreement attraction in comprehension: Representations and processes. *Journal of Memory and Language* 61, 206-237.
- Yoshida, M., M. W. Dickey and P. Sturt. 2013. Predictive processing of syntactic structure: Sluicing and ellipsis in real-time sentence processing. *Language and Cognitive Processes* 28(3), 272-302.
- Zehr, J. and F. Schwarz. 2018. PennController for Internet Based Experiments (IBEX).

Examples in: English

Applicable Languages: English

Applicable Level: Tertiary