Tapping into semantic recovery : an event-related potential study on the processing of gapping and stripping

Ruijgrok, B.J.

CHAPTER 9

Conclusions and future prospects

9.1 Connecting conclusions

In this thesis, I investigated the ellipsis type Gapping and its sub-type Stripping. After an introduction to the topic in Chapter 1, I discussed in Chapter 2 the relevant theoretical background on Gapping where I demonstrated that Gapping has a multidimensional character. Above all, Gapping has been suggested to be a surface anaphor which has led to a focus on the importance of syntactic structure at the ellipsis site. This contrasts with semantic oriented accounts that might consider ellipsis antecedents to be a type of a deep anaphors. Although I started my research by assuming a simple differentiation between structural and non-structural accounts (as is frequently found in reviews on ellipsis), this opposition seems shaky. Even though theoretical approaches may be leaning towards one side (syntactic, semantic) it appears that to successfully account for the distributional properties of Gapping-like constructions, syntactic, semantic and prosodic factors need to be taken into account. A successful account of ellipsis should be able to answer the question what is the proper balance between these factors. Crucially, three closely related questions have been entertained in the ellipsis literature:

- What is the nature of the ellipsis site (i.e. its representation)?
- What is the nature of the antecedent (referred to as the "identity" condition)?
- Under which conditions is ellipsis allowed (referred to as the "licensing" condition)?

Theoretical treatments make no claims about the timing of ellipsis resolution, which makes it difficult to link theory to processing – a topic I touched on in Chapter 4, in which I put forward suggestions to improve this by means of computational linguistic research. However, these questions may be taken up by experimental research. In an attempt to connect theory to experiments, I utilised the mechanism Copy α, which is associated with theories of surface anaphors, and a cue-based mechanism, which relates to theories of deep anaphors. These mechanisms reflect to some extent the divide between syntax-first and constraint-based approaches. Therefore, they were helpful to make hypotheses as to the time course of the recovery of Gapping and Stripping. I proposed a two-stage mechanism based on retrieval and integration processes and proposed that the two mechanisms make different predictions with respect to the time course of ellipsis processing. A copy account may be costly as it comes to retrieval since searching for and finding structure might be more difficult as a function of the size of the structure. Once a fully fledged structure is available, it is expected that integration processes occur with relative ease. Contrastingly, a cue-based account, which is mainly explaining the mechanism of retrieval, predicts the reverse.

Before testing this hypothesis, I replicated an ERP study in Chapter 5 on verb Gapping in Dutch, the results of which pointed to an integration process reflected by late positivities. Making predictions with respect to processes of retrieval and integration and using pretested stimuli that were based on the replicated study I could not corroborate the effect of late positivities in my first ERP experiment on structural complexity as described in Chapter 6, since my proposed measure point appeared to be too late. As a consequence, any effect of retrieval – if present – could not be determined. In a post hoc analysis of critical words earlier in the test sentences, I was able to find preliminary evidence for the start of the retrieval process. However, it was reflected by a positivity rather than an expected ELAN-related ERP component. In a follow-up experiment on Stripping, it appeared that the addition of an adjunct to the deleted structure could modulate both retrieval and integration phases. I argued that an early positivity is sustained by both acoustic, attentional and linguistic cues, possibly directly targeting a semantic representation (or representations), marking the start of the resolution process to retrieve missing information in order to integrate it with the remnant structure. I further suggested that the relative difficulty of integration of retrieved material is then reflected by the secondary positivity which I related to a P600. While the first positivity seems to be an amalgam of different neural generators, it is conceivable that the late positivities are sustained by several integration processes that work in parallel as the distribution of the late positivities could not be connected to one single mechanism.

In Chapter 7, I investigated the impact of semantic complexity in Stripping conditions, again based on stimuli used in the replicated study. Based on theoretical insights, I hypothesised that quantifying expressions may be a burden on mechanisms of movement and/or copying since additional structural in-

formation has to be analysed for such mechanisms to work. I extended this postulation to processing, and suggested that a mechanism such as Copy α would predict a structural processing cost during the recovery of the quantifying expressions, which should be reflected as a syntax-related ERP. Since no difference could be established between non-quantifying and quantifying expressions, I concluded that this is problematic for accounts that consider the representation of a possible antecedent for ellipsis as fully fledged syntactic structure. Again, positive deflections were found during the resolution of Stripping, though only for a small group of participants.

In a final experiment reported in Chapter 8, I carried out an auditory ERP experiment to investigate the effect of prosody on the prediction of Gapping constructions, asking to what extent the prosody of the first conjunct predicts upcoming (deleted) structure. For this experiment, I recorded a selection of stimuli that were used in the replicated study. Although the sample was too small to draw clear conclusions, an exploratory analysis pointed to ERP effects related to attention/selection processes at the critical measure point.

In order to to keep participants engaged in the task, I included comprehension questions in all the experiments. Interestingly, I was able to show that these offline data do not always converge with the online EEG data. That is, it is possible to detect extra processing effort that does not appear to impact the comprehension. Furthermore, comprehension scores may deviate from acceptability scores. As I argued in Chapter 4, an understanding of human language benefits from complementary methods, that is, it cannot be based on acceptability scores alone.

I also carried out a working memory test as a means to control for the variation of the capacity of people's working memory systems. In general, the working memory data showed a small, non-significant positive correlation with sentence comprehension scores. In one experiment, there was a large, significant correlation. Based on the differences in ERP data, I concluded that the correlation might have been caused by a difference in attention level. In general, working memory as measured in this study does not seem to play a large role in comprehension of elliptical sentences.

All experimental findings underscored the multidimensional nature of Gapping. In that sense, an answer to the first question listed above cannot be clear-cut since the nature of the ellipsis site appears to consist of different information types. I have argued in favour of the notion of two consecutive processes underlying ellipsis resolution. While in the theoretical literature a distinction has been assumed between identity of the antecedent and the form of the ellipsis site, to my knowledge, it has never been acknowledged that these conceptions may be associated with a processing order. In addition, the psycholinguistic literature has overlooked a possible order of processing steps, which has led to an ambiguous discussion on the question to what extent ellipsis resolution is cost-free. Although a straightforward linkage between theory and processing is problematic, on the basis of the EEG data it seems that the identity of an antecedent corresponds to a representa-

tion of information types that are targeted during retrieval processes, while the actual form of the ellipsis site may be understood in terms of the way that these information types are integrated at a secondary stage. In addition, although a licensing constraint regarding prosodic parallelism may be absent in behavioural responses, it can be measured in terms of brain activity.

Despite the fact that theories of ellipsis lack a comprehensive account of timing, meaning it is hard to relate theory to online processes, some theorists do have their doors open (or at least, ajar) to processing data. While they may differ as to the degree of syntactic structure assumed, they all emphasise the importance of semantic representations. As a side effect, they tend to account for ellipsis constructions (of whichever type) using one mechanism. From a processing view, this is preferred since it seems unlikely that every descriptively different ellipsis type is resolved by a uniquely dedicated procedure. In fact, with the current data it appears that ellipsis processing resembles 'normal' sentence processing to a large extent. Sentence comprehension is an incremental process during which incoming information is paired with an interpretation – updating representations step by step. On a word-by-word basis, the processor parses each new incoming word to retrieve the necessary information. Incrementally, the processor postulates phonological, syntactic and semantic representations integrating different information types to construct the meaning of a sentence. Ellipsis resolution differs in terms of the polarity and latency of the ERP component related to the retrieval phase.

9.2 Limitations and future prospects

However interesting the results in this study are, I am confronted with some limitations. As already noted during data analysis, some results are based on explorations of small sample sizes and need to be corroborated by future experiments. Furthermore, such experiments should be done cross-linguistically, using languages other than English, Dutch and Spanish, which are the languages for which experimental data on ellipsis exists. As has become clear, different methods may yield different results, therefore, it is highly recommended that stimuli sets are tested using different methods in order to get a more complete picture. Preferably, analyses of the same data sets should be published concurrently to prevent other researchers from attempting to replicate findings that will never be found. However, this requires patience, which is an underestimated virtue as long as researchers are rewarded on the basis of output.

As mentioned, the driving force behind the current project was the simple differentiation between syntactic and semantic accounts. The initial idea was to link this differentiation to electrophysiological data. Already at the end of the second chapter, I concluded that Gapping cannot be captured in either syntactic or semantic terms as the most promising accounts (will need to) integrate different levels of representation. In an attempt to connect theoretical

insights to existing processing accounts, I arrived at a comparison of two behaviourally motivated models of ellipsis processing that are partially reminiscent of the syntax-semantics divide. Despite this, it also became clear that a mapping between existing theoretical insights and processing may not always be straightforward or even justifiable. Therefore, results accumulated in the current study should be interpreted with great caution if one tries to relate them to theory. I have experienced this as a big limitation and I sincerely hope that theoretical and experimental research will begin to reconcile in the near future. For example, to get a better understanding of cues as used in a pointer account, processing research may very well profit from theoretical insights – and vice versa. After all, both approaches aim to investigate one and the same language system. At this point, a particular experience comes to mind: when attending my poster presentation (Ruijgrok, Cremers, Cheng, & Schiller, 2016) during the Ellipsis Across Borders Conference 2016 in Sarajevo, I was very happy to hear Jason Merchant analysing his theoretically-motivated semantic E feature (as proposed in Merchant, 2001) in terms of a experimentally-motivated cue.

Although this research project has concluded, I anticipate embarking on follow-up experiments. The attentive reader may have noticed that one condition of the pretested sentences as reported in Chapter 5.3 have not been tested in an ERP setting, namely the sentence as shown in (1).

(1) Koen verving de kast in de woonkamer, en Judith niet.
Koen replaced the cabinet in the living.room and Judith not
'Koen replaced the cabinet in the hall, and Judith not.'

Comparing *niet* to a control condition that contains *ook* instead will give insight into the way negation is processed. What's more, the stimuli used in the reported (and proposed) ERP experiments could be used in self-paced reading tasks and in eye-tracking experiments, to investigate how a two-stage resolution process can be measured using these techniques. Finally, I would like to get a better understanding of the location of neural generators underlying ellipsis resolution for which I would need to explore these processes using fMRI technique. As usual, in the attempt to find answers, we generate more questions, which I leave to a future me.